

A Review of Privacy-First Mobile Health Architectures: On-Device Versus Cloud Inference for Sensitive Health Data

Emmanuel Eniola Ajala[1*], Kingsley Chinazaekpere Ndupu[2], Serif Oyindamola Oyesiji[3], Stanley Nwakamma[4]

[1]Independent Researcher, Atlanta, Georgia, USA.
[2]Kennesaw State University, Georgia, USA.
[3]Independent Researcher, Lagos, Nigeria.
[4]Independent Researcher, USA.
*Corresponding Author Email: emmanuelajala22@gmail.com

(Received: 10.10.2026; Accepted: 05.11.2026)

Abstract

Mobile devices and wearable sensors have become primary instruments of health measurement, collecting physiological, behavioral, and contextual data continuously and at population scale. The same properties that make these data clinically valuable - their granularity, longitudinality, and intimacy – make them among the most privacy-sensitive data routinely generated about individuals, and the architectural question of where machine learning inference runs, on the device that collects the data or in the cloud that aggregates it, has become a central design decision for mobile health systems. This review synthesizes research on privacy-first mobile health architectures, organizing the literature along the spectrum from fully on-device inference through split and edge-assisted designs to cloud inference hardened with privacy-enhancing technologies. It examines the character and re-identifiability of mobile health data, the threat landscape documented in audits of deployed health applications, and the regulatory context that constrains architectural choice. On the technical side, it reviews the model compression, efficient architecture, and hardware acceleration advances that made on-device inference feasible; the attack literature that quantifies the privacy liabilities of centralized inference and training; and the principal privacy-enhancing technologies, federated learning, differential privacy, secure aggregation, cryptographic inference, and trusted execution environments, that mediate between utility and disclosure. Deployed systems, including federated keyboard prediction, smartwatch-based atrial fibrillation detection, decentralized contact tracing, and smartphone-based mental health sensing, are analyzed as architectural case studies. The review identifies persistent open problems in privacy accounting for longitudinal sensor data, energy-constrained private computation, personalization under privacy budgets, and the evaluation of privacy claims, and closes with directions for future research.

Keywords: *Mobile health; Privacy-preserving machine learning; On-device inference; Federated learning; Differential privacy; Privacy-enhancing technologies*

INTRODUCTION

Health measurement has left the clinic. Smartphones and wearable devices now sense heart rhythm, movement, sleep, speech, location, and social behavior continuously, and the resulting data streams support applications ranging from consumer wellness through chronic disease management to clinical-grade diagnostics (Steinhubl *et al.*, 2015; Piwek *et al.*, 2016). Systematic evidence for the effectiveness of mobile health interventions has accumulated steadily (Free *et al.*, 2013), and the broader convergence of machine learning with medicine has made predictive models the standard analytical instrument applied to these data (Topol, 2019). Deep learning in particular, with its capacity to learn representations directly from raw sensor signals, has displaced hand-engineered features across most mobile sensing tasks (LeCun *et al.*, 2015). Health system perspectives situate these instruments inside a broader transformation agenda in which real-world evidence, predictive analytics, and data-driven decision-making are

expected to lift system performance and patient outcomes (Ilodigwe and Adesemoye, 2024), while implementation frameworks for scaling mobile health treat chronic disease monitoring and treatment adherence as the principal deployment targets (Oparah *et al.*, 2021b). Access research adds a distributional dimension, documenting the contribution of telehealth to narrowing rural and urban disparities in care (Oparah *et al.*, 2024; Fapohunda *et al.*, 2023), and decision-support studies show the same data feeding operational and resource planning models inside health services (Adelanwa *et al.*, 2024).

The privacy stakes of this transition are difficult to overstate. Mobile health data are intimate in content, describing bodies, behaviors, and mental states; they are longitudinal, accumulating over months and years; and they are contextual, binding health signals to times, places, and social ties. Research agendas in mobile health security have long identified this combination as qualitatively more sensitive than conventional medical records, because the data are collected outside institutional custody, flow through consumer software supply chains, and resist effective anonymization (Kotz *et al.*, 2016). Legal scholarship reaches a parallel conclusion: much of the health-relevant data generated by phones and wearables falls outside the protective perimeter of traditional health privacy law, leaving architecture and engineering to carry protective weight that regulation does not (Price and Cohen, 2019). Governance scholarship reaches a similar conclusion from the organizational side, arguing that sensitivity classification and regulatory traceability must attach to the data itself rather than to whichever system happens to hold it (Aliliele *et al.*, 2024a), and that privacy, safeguarding, and ethical handling require explicit models wherever health information is processed outside clinical custody (Orise and Niniola, 2023).

Against this background, a single architectural question organizes much of the design space: where does inference run. At one pole, models execute entirely on the device that collects the data, and raw signals never leave the user's possession. At the other, devices act as thin sensors that stream data to cloud services where models are hosted, updated, and monitored centrally. Between the poles lie split architectures that partition computation between device and server, edge designs that interpose nearby infrastructure, and federated designs that centralize learning while decentralizing data. Each position on this spectrum trades differently among privacy, model capability, energy, latency, bandwidth, maintainability, and auditability, and the past decade of research has both expanded the feasible region, chiefly by making capable models small enough to run on phones, and mapped its hazards, chiefly by demonstrating that centralized models and even model updates leak information about the individuals whose data built them.

This review synthesizes that research for the mobile health setting specifically, where the general trade-offs of mobile machine learning acquire clinical and ethical weight. The contribution is an architecture-centered organization of a literature that is otherwise scattered across health informatics, mobile systems, machine learning, security, and cryptography venues, together with an assessment of which privacy-first designs have reached deployment maturity and which remain laboratory demonstrations. The paper first establishes the scope and approach of the review, then characterizes mobile health data and its re-identifiability before examining the documented threat landscape and regulatory context. It next presents the architectural spectrum and considers on-device inference, cloud inference and its attack surface, and split designs. The discussion then turns to the principal privacy-enhancing technologies: federated learning, differential privacy, and cryptographic and hardware-based protection. Deployed systems are analyzed as case studies, followed by an examination of evaluation practices, open problems, and future directions. The paper concludes by synthesizing the principal findings. Adjacent syntheses of data governance in digital health systems supply part of the implementation vocabulary the review borrows when it separates stated policy from realized behavior (Orise and Niniola, 2025; Ekechi *et al.*, 2026).

SCOPE AND APPROACH OF THE REVIEW

The review addresses architectures for machine learning inference and training over data collected by personal mobile and wearable devices for health purposes, spanning consumer wellness, remote monitoring, digital therapeutics, and mobile clinical research. Privacy-first is understood as a design stance in which minimization of disclosure is a primary architectural objective rather than a compliance afterthought: the systems of interest are those that structurally limit what about an individual leaves their device, not those that merely secure transmission and storage of whatever is collected. Conventional security engineering, authentication, transport encryption, and secure storage are treated as assumed infrastructure rather than a subject of review, and the review likewise excludes hospital-side clinical information systems except where mobile architectures interface with them. Conventional enterprise controls, including identity and access management, key custody, and configuration governance, are likewise assumed rather than reviewed, since they are treated thoroughly elsewhere (Mbonu *et al.*, 2020, 2022b; Oshoba *et al.*, 2019).

The literature base spans the health informatics venues in which mobile health applications and evaluations appear (Silva *et al.*, 2015; Kumar *et al.*, 2013), the mobile and ubiquitous computing venues that developed the sensing and on-device inference foundations, the machine learning venues responsible for federated and privacy-preserving learning, and the security and privacy venues in which both attacks and cryptographic defenses are published. Mapping studies of the mobile health intervention space were used to bound the application landscape (Klasnja and Pratt, 2012), and reviews of privacy in mobile health applications anchored the threat-oriented sections (Martínez-Pérez *et al.*, 2015). Emphasis falls on work from roughly 2010 onward, the smartphone era, with foundational cryptographic and statistical privacy results retained regardless of date because the field builds directly on them. Health system and public health informatics literatures were used to bound the deployment context (Ilodigwe and Adesemoye, 2021; Anioke and Atima, 2023c), and governance and compliance scholarship was consulted wherever architectural choice is

constrained by law rather than by engineering (Ominyi and Anichukwueze, 2022; Annan, 2022).

Synthesis is organized around the architectural spectrum rather than around diseases, sensors, or individual technologies, for two reasons. First, the privacy properties that matter are properties of system structure: what data crosses which trust boundary, in what form, under whose control. Second, the same architectural patterns recur across otherwise unrelated applications, so an architecture-centered treatment exposes transferable design knowledge that disease-centered surveys fragment. Throughout, the review distinguishes three claims that the literature often conflates: that a computation is performed locally, that a system's outputs disclose little about its inputs, and that a deployed product in fact behaves as its architecture suggests. These correspond to architecture, formal privacy, and assurance respectively, and they are established by different kinds of evidence.

MOBILE HEALTH DATA AND ITS PRIVACY CHARACTER

The data foundation of mobile health is continuous multimodal sensing. Smartphone sensing frameworks established the phone as a platform for inferring activity, mobility, and social context from accelerometry, location, microphone, and usage traces (Lane *et al.*, 2010), and body-worn systems extended coverage to physiological signals including heart rate and rhythm, skin conductance, temperature, and gait (Pantelopoulos and Bourbakis, 2010; Patel *et al.*, 2012). Surveys of wearable health monitoring document a steadily broadening sensor repertoire and its migration from research prototypes into consumer products (Majumder *et al.*, 2017). On the analytical side, human activity recognition matured from classical classifiers over engineered features (Lara and Labrador, 2013) to deep architectures that combine convolutional feature learning with recurrent temporal modeling (Ordóñez and Roggen, 2016), the latter building on long short-term memory networks originally developed for general sequence learning (Hochreiter and Schmidhuber, 1997), and such models were subsequently demonstrated running directly on mobile and wearable hardware (Ravi *et al.*, 2017). Comparable instrumentation questions arise in the wider sensing literature, where wireless system-on-chip and transceiver behavior determine what a continuously sampling device can sustain (Quainoo and Ogundapo, 2026b), and where radio frequency energy harvesting research sets the outer bound on autonomous, maintenance-free sensing (Asiedu and Quainoo, 2023a).

The health information content of these streams extends well beyond fitness metrics. Personal sensing research shows that mobility regularity, communication patterns, sleep, and voice features carry signals about depression, stress, and other mental health constructs (Mohr *et al.*, 2017), with campus-scale studies demonstrating passive prediction of mental wellbeing and academic outcomes from phone sensors alone (Wang *et al.*, 2014). Smartphone-based research platforms have collected disease-specific measures at cohort scale, as in the Parkinson's study that gathered tapping, gait, and voice tests from thousands of participants through a consumer app (Bot *et al.*, 2016). The clinical translation of such measures is formalized in the digital biomarker literature, which treats sensor-derived indices as measurements requiring validation, qualification, and governance comparable to laboratory assays (Coravos *et al.*, 2019). Continuous measurement also enables continuous intervention: just-in-time adaptive interventions adjust support moment to moment based on sensed state, which makes the sensing pipeline part of the treatment loop and raises the stakes of its integrity (Nahum-Shani *et al.*, 2018). Population-facing analyses make a parallel point at aggregate scale, showing that behavioral and service-use data support allocation and accountability decisions in mental health and substance use programs (Anioke and Atima, 2023b; Owoade *et al.*, 2022; Yusuff *et al.*, 2025). Cardiometabolic research illustrates the clinical end of the same pipeline, where composite indices and echocardiographic parameters are combined for risk stratification (Okwah, 2022) and where nutritional and molecular markers are related to cardiac outcomes in defined cohorts (Ike *et al.*, 2025).

The privacy character of these data follows from three properties. First, they are identifying. The anonymization literature established formal models, k-anonymity and its refinements l-diversity and t-closeness, for limiting re-identification in released tabular data (Sweeney, 2002; Machanavajjhala *et al.*, 2007; Li *et al.*, 2007), and then demonstrated their fragility against high-dimensional behavioral data: sparse preference records were re-identified by linkage (Narayanan and Shmatikov, 2008), four spatiotemporal points sufficed to single out most individuals in a mobility dataset (de Montjoye *et al.*, 2013), and generative modeling showed that incomplete demographic releases still support confident re-identification (Rocher *et al.*, 2019). Systematic review of re-identification attacks on health data specifically found successful attacks concentrated precisely where data custodians had followed weak de-identification practice (El Emam *et al.*, 2011). Mobile health streams combine the worst of these cases: they are high-dimensional, longitudinal, behaviorally unique, and linkable through timestamps and location. Enterprise data classification frameworks operationalize the same insight for practice, treating identifiability as a property to be assessed and traced through processing rather than assumed away by nominal de-identification (Aliliele *et al.*, 2024a; Atakpa and Abolaji, 2022).

Second, the data are inferentially rich: attributes never collected can be derived from attributes that were, so the sensitivity of a dataset is set by what models can extract from it rather than by its nominal schema. Third, collection is continuous and ambient, which erodes the transactional consent model; a user who authorized step counting has, in practice, authorized a data stream from which sleep disruption, tremor, and mood may later be inferred. These properties jointly explain why architectural minimization, ensuring that raw streams never accumulate outside the device, occupies the central place in privacy-first design that access control occupies in institutional health informatics.

THREAT LANDSCAPE AND REGULATORY

CONTEXT

The threat side of the literature proceeds from taxonomy to measurement. Early work classified mobile health privacy threats by adversary and vector, distinguishing identity, access, and disclosure threats across patients, devices, networks, and servers (Kotz, 2011), and comprehensive surveys of privacy in personal healthcare technology mapped the design space of protections available at each layer (Avancha *et al.*, 2012). Subsequent research agendas emphasized that the mobile health supply chain, comprising device vendors, operating systems, app developers, analytics providers, and cloud platforms, multiplies trust boundaries in ways that traditional health information governance does not contemplate (Kotz *et al.*, 2016). Work in adjacent critical sectors applies the same taxonomic discipline to converged operational environments, mapping adversaries, exposure, and remediation across boundaries no single operator controls (Adegbite *et al.*, 2025; Ahmed *et al.*, 2021; Ojukwu *et al.*, 2023), and analyses of cloud misconfiguration show how much residual exposure in modern deployments originates in configuration rather than in protocol design (Aliliele *et al.*, 2023).

Empirical audits then measured practice against principle, with consistent results. An assessment of apps carrying a national health service accreditation found widespread transmission of identifying information without encryption and privacy policies that did not describe actual data flows (Huckvale *et al.*, 2015). Static and dynamic analysis of popular health applications found insecure transmission, third-party sharing, and permissions exceeding function, characterizing the state of practice as alarming (Papageorgiou *et al.*, 2018). Traffic analysis of medicines-related apps documented routine sharing of user data with a concentrated network of analytics and advertising intermediaries (Grundy *et al.*, 2019). These findings echo the platform-level measurement tradition: information-flow tracking on smartphones revealed pervasive undisclosed transmission of sensor and identifier data by third-party applications (Enck *et al.*, 2014), and usability studies showed that permission systems fail to inform the attention and comprehension of most users (Felt *et al.*, 2012). The audit literature thus supplies the empirical premise of privacy-first architecture: in the deployed ecosystem, data that leaves the device enters a commercial data flow that neither users nor regulators reliably observe. Behavioral studies of the people who operate these systems complete the picture, showing that awareness, training, and everyday handling practice determine whether documented controls are exercised at all (Onche *et al.*, 2020, 2024, 2026).

The regulatory context sharpens rather than resolves the problem. In the United States, health privacy law attaches to covered entities rather than to data, so identical measurements are regulated when held by a clinic and largely unregulated when held by an app vendor, a gap analyzed in the health law literature together with proposals to modernize it (Cohen and Mello, 2018). European data protection law applies horizontally and treats health data as a special category demanding heightened protection, with data protection by design an explicit legal obligation, making architecture itself a compliance instrument (Voigt and von dem Bussche, 2017). Policy work in the learning health system tradition argues that durable use of consumer-generated health data requires protections calibrated to data sensitivity rather than custodian identity (McGraw and Mandl, 2021). Comparative reviews of privacy-preserving health data governance across jurisdictions, spanning cryptographic and distributed-ledger strategies, map the design responses this pressure has produced (Afrihyia *et al.*, 2024b), and telehealth reviews identify user trust, privacy, and regulatory governance as adoption determinants for AI-powered remote monitoring systems (Akintola *et al.*, 2025). Professional and ethical commentary converges on the same point from the clinical side, identifying transparency about data practices as a precondition of trust in digital health tools (Torous and Roberts, 2017), and roadmaps for responsible clinical machine learning place data stewardship among the field's primary obligations (Wiens *et al.*, 2019). Across jurisdictions and professional frames, the direction of pressure is uniform: architectures that structurally limit disclosure are favored, and the remainder of this review examines how far current technology permits that limitation without sacrificing capability. Comparative governance analyses extend the question to oversight structures themselves, contrasting executive accountability and regulatory design for AI-driven healthcare management across jurisdictions (Afrihyia *et al.*, 2024a), while compliance scholarship maps the practical problem of harmonizing obligations across overlapping privacy regimes (Ominyi and Anichukwueze, 2022; Mbonu *et al.*, 2019). Impact assessment models supply the procedural instrument through which those obligations are discharged in multi-cloud processing environments (Mbonu *et al.*, 2022a), and readiness studies for emerging artificial intelligence statutes indicate how quickly the compliance surface is expanding (Ominyi *et al.*, 2024). Sector-specific architecture work translates the same obligations into concrete controls for community healthcare organizations operating under HIPAA (Aliliele *et al.*, 2024b; Atakpa and Abolaji, 2022), and accountability scholarship situates the resulting disclosures within wider debates about algorithmic transparency and its limits (Annan, 2024; Ominyi and Anichukwueze, 2023; Ekechi *et al.*, 2026).

THE ARCHITECTURAL SPECTRUM: ON-DEVICE, CLOUD, AND HYBRID INFERENCE

The design space is usefully pictured as a spectrum defined by where raw data, features, model parameters, and predictions reside. Fully cloud-hosted inference concentrates model execution, versioning, and monitoring in centrally operated services, with devices as sensing and presentation endpoints. Fully on-device inference inverts this, executing models where data originates so that, in the limit, only user-facing results exist anywhere. The intermediate region is populated by three recurring patterns: edge offloading, in which computation moves to infrastructure near the user, a line of work initiated by the cloudlet proposal for bringing virtual-machine-based compute within one wireless hop of mobile devices (Satyanarayanan *et al.*, 2009) and generalized as edge computing matured into an architectural movement

(Satyanarayanan, 2017); split computation, in which a single model is partitioned across device and server; and federated designs, in which inference is local but learning is collaborative. Comparable spectra appear wherever constrained endpoints meet capable infrastructure, as in edge intelligence for microgrid control, where autonomy, latency, and sovereignty considerations push inference toward the endpoint (Komi and Ganiu, 2023), and portability research on cloud-to-edge runtimes shows how one artifact can be scheduled at several points along that spectrum without rebuilding it (Ojukwu *et al.*, 2025).

The general systems literature frames the trade-offs. Edge computing analyses catalogue the latency, bandwidth, privacy, and autonomy motivations for moving computation toward data sources, against the constraints of constrained hardware and distributed management (Shi *et al.*, 2016), and communication-oriented surveys of mobile edge computing formalize the joint optimization of computation placement and radio resources (Mao *et al.*, 2017). Surveys of edge intelligence and of deep learning with edge computing map the specific design space for neural workloads, including model placement, partitioning, caching, and inference offloading policies (Zhou *et al.*, 2019; Chen and Ran, 2019). For health workloads, three considerations reweight these general trade-offs. Sensitivity elevates the privacy term: every hop a raw physiological stream traverses is a disclosure event with clinical and legal significance. Continuity elevates the energy term: monitoring that runs always must draw milliwatts, not watts. And accountability elevates the assurance term: claims about where health data flows must be verifiable by auditors, not merely asserted in documentation. Energy and reliability studies of green communication supply the quantitative counterpart of these trade-offs, characterizing how throughput, latency, and lifetime move against one another under constrained power budgets (Asiedu and Quainoo, 2024; Asiedu *et al.*, 2025), and security analyses of edge-hosted microservices show that distributing computation also distributes attack surface (Idika *et al.*, 2021).

Framed this way, the architectural question is not which pole is superior but which components of a pipeline must cross which trust boundaries in which representation. A recurring privacy-first heuristic orders representations by disclosure risk: raw signals carry the most collateral information, learned features less, predictions less still, and aggregate statistics least. Architectures then minimize the sensitivity of whatever must leave the device, transmitting a rhythm classification rather than the electrocardiogram, a model update rather than the training data, a noised aggregate rather than the update. The discussion that follows examines the technologies that make each such reduction possible and the attack results that measure how much protection each actually confers.

ON-DEVICE INFERENCE: MODELS, HARDWARE, AND TOOLCHAINS

On-device inference became realistic when model efficiency research and mobile hardware converged. On the model side, compact architecture families were designed explicitly for constrained budgets: depthwise separable convolutions gave the MobileNet line its efficiency (Howard *et al.*, 2017), inverted residual structures improved it (Sandler *et al.*, 2018), and compound scaling produced families spanning the accuracy-efficiency frontier (Tan and Le, 2019), all building on the residual learning foundations of modern deep networks (He *et al.*, 2016). Complementary compression techniques shrink trained models: pruning with quantization and entropy coding compresses networks by an order of magnitude (Han *et al.*, 2016), knowledge distillation transfers capability from large teachers to small students (Hinton *et al.*, 2015), integer-only quantization co-designed with mobile runtimes made eight-bit inference standard practice (Jacob *et al.*, 2018), and the lottery ticket results showed that dense training often conceals sparse trainable subnetworks, sharpening theoretical interest in sparsity (Frankle and Carbin, 2019). Tutorial surveys of efficient neural processing consolidate this toolset and its hardware implications (Sze *et al.*, 2017). Recent surveys extend this compression toolkit to the model families now reaching the device tier, covering quantization, pruning, and distillation for multimodal small language models intended for privacy-preserving edge inference (Nwakamma *et al.*, 2025), compression of recommender and ranking models for edge deployment (Oyesiji *et al.*, 2025), and parameter-efficient adaptation methods that specialize an on-device model without retraining it wholesale (Oyesiji *et al.*, 2023).

On the systems side, mobile inference engines demonstrated that deep models could run within phone and wearable budgets, with early software accelerators executing multiple concurrent deep models on commodity mobile processors through decomposition and scheduling (Lane *et al.*, 2016). Ecosystem measurements traced adoption: analysis of application stores found deep learning already embedded in thousands of smartphone apps by the late 2010s, dominated by vision and increasingly executed locally (Xu *et al.*, 2019), while benchmark suites quantified the rapidly growing neural performance of mobile system-on-chip generations equipped with dedicated accelerators (Ignatov *et al.*, 2019). Below the phone tier, embedded inference frameworks brought neural execution to microcontroller-class devices measured in kilobytes of memory (David *et al.*, 2021), and the surrounding TinyML practice literature codified development patterns for always-on sensing at milliwatt power (Warden and Situnayake, 2019). Energy analysis of smartphones grounds the budgets these systems must respect, decomposing device power draw across radio, processor, and peripherals and underscoring that communication frequently costs more than computation (Carroll and Heiser, 2010), an accounting that favors local inference for continuous sensing workloads on energy grounds independent of privacy. Device-level power studies complement these benchmarks by modeling how concurrent radios and peripherals shape the residual budget available to inference (Quainoo and Ogundapo, 2026a), and work on intermittently powered platforms shows that lightweight learned controllers can schedule computation under uncertain energy availability (Asiedu *et al.*, 2026; Asiedu and Quainoo, 2025). Semiconductor research points to the supply side of the same trend, where advances in energy-efficient integrated circuits

and radio frequency front-end components determine what future wearable tiers will be able to execute at all (Erhimefe and Gideon, 2025; Gideon, 2025).

For mobile health specifically, these advances translate into a characteristic design: always-on models sized to the device tier, hierarchically staged so that cheap detectors gate expensive analyzers. A wearable-tier model screens a continuous signal, wakes a phone-tier model for candidate events, and only confirmed, clinically meaningful results are surfaced or, with consent, transmitted. Demonstrations of deep on-node analytics for wearable signals established the feasibility of this pattern for activity and physiological monitoring (Ravi *et al.*, 2017). Its privacy value is structural: the raw stream exists only transiently in device memory, and the persistent record is the minimal clinical event. Its costs are equally structural: on-device models are harder to update, monitor, and debug than hosted ones; device heterogeneity multiplies validation burden; and the strongest models may simply not fit, which is the opening for the split and federated designs of later sections. Staged local screening resembles the triage logic long used in automated detection for imaging, where inexpensive pattern-level analysis precedes expert review of flagged cases (Ejofodomi *et al.*, 2014).

CLOUD INFERENCE AND ITS PRIVACY LIABILITIES

The case for cloud inference is capability and control. Centralized platforms host models of unrestricted size on managed infrastructure (Abadi *et al.*, 2016a), and the most influential clinical deep learning results were built in this regime: dermatologist-level skin lesion classification (Esteva *et al.*, 2017), cardiologist-level arrhythmia detection over ambulatory electrocardiography (Hannun *et al.*, 2019), and scalable prediction over electronic health records (Rajkomar *et al.*, 2018). Reviews of deep learning in healthcare presume this centralized pattern, in which data aggregates and models serve (Miotto *et al.*, 2018; Esteva *et al.*, 2019). Centralization also concentrates the operational virtues that clinical deployment demands: versioned models, monitored performance, and controlled rollout, which pure on-device designs sacrifice. The operational economics are familiar from enterprise practice, where predictive scaling and constraint-based allocation govern the cost of hosted inference (Ahmed *et al.*, 2019, 2020a) and migration studies document what organizations gain and surrender when workloads move to managed infrastructure (Ladapo *et al.*, 2022b). Healthcare decision-support deployments illustrate the same centralized pattern applied to operations, resource planning, and clinical decision-making (Adelanwa *et al.*, 2024; Ilodigwe and Adesemoye, 2025).

The privacy liabilities of the same pattern are now quantified by a substantial attack literature. Against served models, membership inference determines whether a given individual's record was in the training set from prediction outputs alone (Shokri *et al.*, 2017), an attack whose health consequence is direct, since membership in a disease cohort is itself sensitive. Model inversion reconstructs class-representative inputs from model access (Fredrikson *et al.*, 2015), and extraction attacks replicate proprietary models through prediction interfaces (Tramèr *et al.*, 2016). Against large trained models, verbatim training data has been extracted, demonstrating that models memorize and can regurgitate individual records (Carlini *et al.*, 2021). Comprehensive white-box analyses extend these results to training-time observation, showing substantial leakage from gradients and intermediate states in both centralized and collaborative settings (Nasr *et al.*, 2019). The collective import for mobile health is that transmitting data to a model, and even participating in building one, are disclosure events whose magnitude depends on defenses, not intentions, and the privacy-enhancing technologies discussed in this review are designed to bound that magnitude.

Between the poles, the practical cloud-inclusive architecture for health has converged on data minimization at the boundary: devices transmit derived features or events rather than raw streams, retention is bounded, and telemetry is separated from identity. These conventions reduce exposure but are difficult to verify externally. The documented audit findings show that such shortcomings can persist even in well-intentioned deployments, underscoring why formally analyzable mechanisms are increasingly preferred to procedural assurances at the research frontier. Architectural proposals for privacy-centric security engineering and privacy-by-design role separation formalize these conventions as design obligations rather than deployment habits (Okoruwa *et al.*, 2020; Badmus *et al.*, 2024; Ominyi, 2021), and identity-centric zero trust models narrow the implicit trust that cloud-inclusive pipelines otherwise extend to every internal component (Ogbole *et al.*, 2021; Adebayo *et al.*, 2025).

SPLIT AND EDGE-ASSISTED ARCHITECTURES

Split computing occupies the middle of the spectrum by partitioning a single neural network between device and server, executing early layers locally and transmitting intermediate activations rather than raw inputs. The systematic case was made by profiling studies showing that the optimal partition point varies by network architecture, hardware, and wireless conditions, and that automatically selecting it can dominate both pure-local and pure-cloud execution in latency and energy (Kang *et al.*, 2017). Distributed formulations generalized the idea across device, edge, and cloud tiers, training networks with early exit points so that easy inputs terminate locally and only difficult ones escalate (Teerapittayanon *et al.*, 2017). Surveys of split computing and early exiting consolidate a decade of variants, including bottleneck injection designs that compress the transmitted representation and training regimes that make partitioned execution robust to channel variability (Matsubara *et al.*, 2022).

The privacy interpretation of split designs requires care, and the literature has become appropriately cautious. Transmitting activations rather than inputs reduces obvious disclosure, but intermediate representations of early layers remain substantially invertible, so a split point is a compression choice, not by itself a privacy mechanism. Privacy-relevant benefits accrue when splitting is combined with representation design that deliberately discards input

information beyond task relevance, with noise addition calibrated under differential privacy, or with the cryptographic protections applied to the narrower interface that splitting creates. For mobile health, the natural role of split designs is therefore economic: they let energy-constrained wearables and phones participate in deep pipelines they cannot host, staging computation across the device, a home or clinical edge node, and the cloud, while the formal privacy budget is enforced by mechanisms layered on the split rather than by the split itself. Early exiting has an additional health-specific attraction: escalation policies map cleanly onto clinical triage semantics, in which routine signals resolve locally and anomalous ones warrant deeper, possibly remote, analysis. Attribute-based encryption offers one such layered mechanism, binding access to policy at the level of the transmitted object rather than the channel that carries it (Shittu and Nabil, 2023).

Edge-assisted variants place the server side of the split on infrastructure with a defined trust relationship, a hospital gateway, a home hub, a clinic-operated edge node, rather than on a public cloud. This preserves the operational virtues of server-side components, updatability and monitoring, while keeping data within an administratively meaningful boundary, and it aligns with the data protection preference for processing proximity discussed earlier in the threat landscape and regulatory context. The cost is fleet management of distributed servers, and the assurance question, whether the edge node in fact behaves as claimed, recurs at the new boundary, motivating the attestation technologies reviewed and discussed later in this review under the cryptographic protection and trusted hardware section. The management burden of such distributed tiers is documented in operational technology settings, where fleets of administratively distinct nodes must be governed, patched, and attested continuously (Ahmed *et al.*, 2021; Adegbite *et al.*, 2025), and reliability studies of intermittently powered networks indicate how uneven availability becomes at the far edge (Asiedu and Quainoo, 2023b).

FEDERATED LEARNING FOR MOBILE HEALTH

Federated learning restructures training rather than inference: models are learned across a fleet of devices that retain their data, with a coordinating server aggregating locally computed updates. The founding formulation demonstrated that iterative averaging of client updates trains deep networks to competitive accuracy over decentralized, non-identically distributed data (McMahan *et al.*, 2017), with companion work addressing the communication efficiency that mobile participation requires (Konečný *et al.*, 2016). Production system papers described the full machinery, client selection, secure aggregation, versioning, and staged rollout, operating across millions of phones (Bonawitz *et al.*, 2019), and the paradigmatic deployment trained next-word prediction models for mobile keyboards without centralizing typing histories (Hard *et al.*, 2018). Methodological extensions treat statistical heterogeneity through multi-task and personalized formulations (Smith *et al.*, 2017), and the field's open problems, from robustness to fairness to systems constraints, are catalogued in comprehensive surveys (Li *et al.*, 2020; Kairouz *et al.*, 2021), with taxonomies distinguishing horizontal, vertical, and transfer settings (Yang *et al.*, 2019). The paradigm has since been carried into other cross-organizational settings in which data cannot be pooled, including lifecycle-aware automation trained across institutional boundaries (Sanni *et al.*, 2023), and human-in-the-loop formulations raise the further question of how disagreement among distributed annotators propagates into the aggregated model (Ladapo *et al.*, 2026; Ladapo *et al.*, 2022a).

Health became an early motivating domain. Perspective papers laid out the case that federated learning fits medicine's structural condition: valuable data fragmented across custodians who cannot pool it (Rieke *et al.*, 2020), and multi-institutional studies demonstrated that federated training over clinical imaging approaches the accuracy of centralized training while respecting data residency (Sheller *et al.*, 2020). Surveys specific to healthcare informatics map applications across imaging, records, and sensing, and importantly for this review, extend the paradigm to the consumer device tier, where each phone or wearable is a client and the fleet constitutes a continuously sensing cohort (Xu *et al.*, 2021). In that tier, federated learning composes naturally with the on-device inference: the same local model that serves predictions can generate training updates, allowing a mobile health system to improve across the population without moving the underlying data off devices. Federated participation also depends on the identity and access machinery that admits clients to a training round, a problem treated in the distributed and federated systems literature (Oshoba *et al.*, 2019; Mbonu *et al.*, 2020).

The attack literature disciplines the paradigm's privacy claims. Gradients are not anonymous: reconstruction attacks recover training inputs from shared updates with alarming fidelity in small-batch regimes (Zhu *et al.*, 2019), and collaborative learning leaks unintended features of client data even when reconstruction fails (Melis *et al.*, 2019), with white-box analyses confirming substantial membership leakage from observed training dynamics (Nasr *et al.*, 2019). Federated learning is therefore best understood as an architectural minimization, raw data residency, that requires composition with formal mechanisms to become a privacy guarantee: secure aggregation so the server observes only sums (Bonawitz *et al.*, 2017), and differential privacy at the client level so that aggregates provably bound individual influence (Geyer *et al.*, 2017). For mobile health fleets, the composition is demanding, because clients are battery-powered, intermittently connected, and heterogeneous, but the keyboard deployment demonstrates that the full stack-local training, secure aggregation, and noised updates-operates at consumer scale (Hard *et al.*, 2018; Bonawitz *et al.*, 2019). Data curation practice interacts directly with these results, since what is admitted into local training determines the sensitivity of what an update can leak (Oyesiji *et al.*, 2024).

DIFFERENTIAL PRIVACY AND STATISTICAL DISCLOSURE CONTROL

Differential privacy supplies the formal vocabulary in which modern disclosure bounds are stated. The definition,

requiring that the distribution of a computation's output be nearly invariant to the presence or absence of any single individual, emerged from foundational work calibrating noise to query sensitivity (Dwork *et al.*, 2006) and was consolidated as a general framework in the same period (Dwork, 2006), with the standard reference treatment establishing composition, group privacy, and the mechanism toolkit (Dwork and Roth, 2014). Two developments made the framework practical for the architectures of this review. Differentially private stochastic gradient descent brought deep learning under the framework, clipping per-example gradients and adding calibrated noise with a moments-based accountant tight enough to train useful models at meaningful budgets (Abadi *et al.*, 2016b). Local and shuffled models moved the trust boundary onto the device itself: randomized response mechanisms deployed in consumer software demonstrated privacy-preserving telemetry collection in which each report is individually deniable (Erlingsson *et al.*, 2014), establishing the pattern, local perturbation before transmission, that consumer platforms subsequently adopted for usage analytics. Applied architectures show the same mechanisms being assembled for regulated analytics, where perturbation, aggregation, and access control are composed to satisfy overlapping legal regimes (Atakpa and Abolaji, 2022).

For mobile health, differential privacy plays three architectural roles. At the aggregate reporting layer, it governs population statistics computed from device fleets, symptom prevalence, exposure rates, and activity distributions, allowing public health use of consumer sensing without individual disclosure. At the federated training layer, client-level formulations bound what any trained model can reveal about any participant, composing with secure aggregation as described in the preceding section (Geyer *et al.*, 2017). At the release layer, it conditions sharing of derived datasets and models with researchers. Each role, however, inherits the framework's central tension in sharpened form: sensor data are longitudinal, and an individual contributes not one record but an unbounded stream, so naive per-report accounting exhausts any credible budget within days. Continuous-observation accounting, user-level rather than event-level guarantees, and budget allocation across a device's many concurrent reporting tasks remain active problems precisely because mobile health is the setting that stresses them most, and current deployments manage the tension by restricting private collection to coarse, low-frequency statistics while keeping fine-grained inference strictly local. The aggregate reporting role is the most visible in practice, since public health functions consume exactly such statistics for surveillance dashboards, outbreak anticipation, and emergency risk stratification (Oparah *et al.*, 2021a, 2022, 2023; Ogunboye *et al.*, 2023), and governance models for those functions specify the accountability under which the statistics are released (Anioke and Atima, 2023a,c).

CRYPTOGRAPHIC PROTECTION AND TRUSTED HARDWARE

Cryptography offers computation over data that remains hidden, and its primitives predate the applications by decades: secret sharing splits values so that only authorized coalitions reconstruct them (Shamir, 1979), garbled circuits established general two-party computation over private inputs (Yao, 1986), zero-knowledge proofs established that statements can be verified without revealing witnesses (Goldwasser *et al.*, 1989), additively homomorphic encryption allowed arithmetic over ciphertexts (Paillier, 1999), and fully homomorphic encryption completed the agenda in principle by supporting arbitrary computation over encrypted data (Gentry, 2009). A dedicated line of work then specialized these tools to machine learning. Neural inference over encrypted inputs was demonstrated with homomorphic evaluation of adapted networks (Gilad-Bachrach *et al.*, 2016), secure multiparty protocols made private training and inference practical for simple models at scale (Mohassel and Zhang, 2017), and hybrid designs combining homomorphic encryption for linear layers with garbled circuits for activations reduced private inference latency by orders of magnitude (Juvekar *et al.*, 2018). In the federated setting, secure aggregation protocols let a server sum client updates while individual contributions remain hidden even against dropouts, and were designed explicitly for mobile fleets (Bonawitz *et al.*, 2017). Migration planning has become part of the same picture, as crypto-agility research anticipates the replacement of current primitives with post-quantum alternatives across the transport and key management layers on which mobile health systems depend (Nwakamma *et al.*, 2026).

Trusted execution environments trade cryptographic assumptions for hardware ones, isolating computation and attesting its identity to remote parties. Analyses of the concept clarify what isolation does and does not promise (Sabt *et al.*, 2015), and detailed treatments of commodity enclave technology document both the architecture and its side-channel exposure (Costan and Devadas, 2016). For mobile health, enclaves appear on both sides of the boundary: device-side secure processors already custody biometric templates and keys, and server-side enclaves permit a cloud operator to process health data it cannot itself read, with attestation providing the assurance component that was found missing from procedural privacy claims as discussed in the threat landscape and regulatory context. Distributed ledger designs address a different slice of the problem, auditability of access rather than confidentiality of computation, with blockchain-mediated permission management over medical records as the representative construction (Azaria *et al.*, 2016). Distributed ledger architectures have been proposed for related exchange and transparency problems, including cross-border data exchange under regulatory observation (Kumuyi *et al.*, 2023), tamper-evident audit trails over configuration and access events (Oshoba *et al.*, 2020), and secure exchange in control systems where provenance matters as much as confidentiality (Shittu *et al.*, 2022). Identity verification frameworks built around minimal disclosure address the complementary problem of admitting the right principals without accumulating identifying material (Kumuyi *et al.*, 2024b).

The comparative economics decide deployment roles. Fully homomorphic inference remains orders of magnitude too costly for continuous mobile workloads; hybrid

cryptographic inference suits episodic, high-stakes queries; secure aggregation is cheap enough for routine fleet learning and is deployed; enclaves are near-native in cost but import hardware trust and side-channel risk. A privacy-first mobile health architecture therefore typically composes them: local inference for the continuous path, secure aggregation plus differential privacy for learning, enclaves for unavoidable server-side processing, and cryptographic inference reserved for narrow interfaces where its cost is tolerable.

DEPLOYED SYSTEMS AS ARCHITECTURAL CASE STUDIES

Deployed systems test which compositions survive contact with consumer-scale constraints. Federated keyboard prediction remains the reference deployment for the learning stack: local training on typing histories, secure aggregation, and centrally coordinated rollout produced measurable model improvement without centralized text collection, across a fleet whose scale, heterogeneity, and churn typify consumer mobile health conditions (Hard *et al.*, 2018; Bonawitz *et al.*, 2019). Its architectural lesson is that the full privacy-first learning composition is operable, and that its binding constraints are systems constraints, participation scheduling, update size, and version management, rather than model accuracy.

Smartwatch arrhythmia detection illustrates the staged on-device inference pattern at clinical scale. In the large prospective smartwatch study of atrial fibrillation, continuous photoplethysmography screening ran on the wearable; only algorithmically flagged irregularity triggered escalation to confirmatory patch electrocardiography, and endpoints were clinically adjudicated (Perez *et al.*, 2019). The architecture embodies a staged hierarchy: raw signals remain transient on the device, and escalation occurs only when an event is detected. Related clinical research using deep learning over electrocardiograms, including for silent atrial fibrillation, indicates the analytical ceiling toward which such staged pipelines are growing (Attia *et al.*, 2019; Hannun *et al.*, 2019). The population such screening is meant to reach is documented in opportunistic screening studies, which find substantial undiagnosed cardiovascular burden wherever routine clinical contact is limited (Amadi *et al.*, 2026), and risk stratification research indicates the composite measures that device-tier screening would ideally approximate (Okwah, 2022). Downstream, the therapeutic literature that acts on such findings continues to develop independently of the sensing tier (Yusuf *et al.*, 2026).

Digital contact tracing during the COVID-19 pandemic became an unusually public architectural referendum. Decentralized designs kept contact graphs on devices, broadcasting rotating pseudonyms and performing exposure matching locally so that no authority learned the social network (Troncoso *et al.*, 2020), and comparative surveys of deployed national apps documented the split between centralized and decentralized architectures and the privacy analyses each attracted (Ahmed *et al.*, 2020). The episode demonstrated both that privacy-first architecture can be a precondition of platform support and public adoption, and that architectural privacy interacts with epidemiological utility, since decentralization constrains the aggregate visibility health authorities receive.

Mobile research cohorts and remote monitoring supply contrasting lessons. App-mediated studies collected rich phenotypic data from thousands of consented participants and released them for research, demonstrating recruitment reach while confronting exactly the re-identifiability limits inherent in their governance design, as outlined in the earlier section – mobile health data and its privacy character (Bot *et al.*, 2016); campus sensing studies showed the inferential depth of passive streams under research ethics oversight (Wang *et al.*, 2014). Comprehensive treatments of remote patient monitoring document the clinical-operational architectures, gateways, dashboards, and alerting, into which any privacy-first sensing tier must ultimately integrate (Malasinghe *et al.*, 2019). Across these cases, the recurring pattern is that successful deployments choose one disclosure-minimizing move and engineer it thoroughly, local screening, federated updates, or on-device matching, rather than attempting the full cryptographic stack at once. Programmatic accounts of mobile health scaling describe the operational layer around these cohorts, including adherence support, escalation pathways, and the service integration that determines whether sensing changes care (Oparah *et al.*, 2021b; Fapohunda *et al.*, 2023; Oparah *et al.*, 2024).

EVALUATING PRIVACY-FIRST ARCHITECTURES

Evaluation in this field is irreducibly multi-dimensional, and the literature's persistent weakness is reporting single dimensions in isolation. Utility evaluation follows the host application: classification metrics for detection tasks, calibrated risk for clinical scores, and, for staged designs, escalation operating characteristics, since a wearable screen's false positive rate determines both clinical burden and how often sensitive data leaves the first tier (Perez *et al.*, 2019). Systems evaluation measures latency, memory, and energy on target hardware, for which mobile benchmarking practice and device power analysis provide methods and baselines (Ignatov *et al.*, 2019; Carroll and Heiser, 2010); credible on-device claims report measurements on the deployment tier, not desktop proxies (Ravi *et al.*, 2017; David *et al.*, 2021). Comparable measurement discipline is emerging in adjacent domains, where model governance regimes specify what must be demonstrated before a model is released rather than after (Nwakamma *et al.*, 2024a; Ominyi *et al.*, 2026b).

Privacy evaluation should distinguish among the three claim types introduced earlier: architecture, formal privacy, and empirical assurance. Architectural claims, what data crosses which boundary, are established by system description and verified by traffic and flow analysis of the kind the audit literature applies to deployed apps (Enck *et al.*, 2014; Papageorgiou *et al.*, 2018). Formal claims are established by mechanism parameters, the differential privacy budget and its accounting over time (Abadi *et al.*, 2016b; Dwork and Roth, 2014), or by cryptographic protocol security under stated assumptions (Bonawitz *et al.*, 2017). Empirical resistance claims are established adversarially, by subjecting the system to the attack repertoire: membership inference, reconstruction from updates, feature leakage, at realistic access levels

(Shokri *et al.*, 2017; Zhu *et al.*, 2019; Melis *et al.*, 2019; Nasr *et al.*, 2019). Mature evaluations report all three where applicable, since each catches failures the others miss: an architecture can be sound and its implementation leaky; a budget can be formally valid and practically meaningless; an attack suite can be survived through evaluation weakness rather than system strength. Adversarial evaluation of this kind has an operational analogue in breach and attack simulation, which validates controls continuously against a defined adversary repertoire instead of certifying them once (Dosunmu and Ogundele, 2024a), and quantification frameworks provide the accompanying vocabulary for expressing residual risk in decision-relevant terms (Dosunmu and Ogundele, 2024b; Adegbite *et al.*, 2025). Continuous auditing systems extend the same logic to the governance layer, generating evidence of compliance as processing occurs rather than reconstructing it afterward (Aliliele *et al.*, 2025).

The comparative dimension remains underdeveloped. Few studies place on-device, split, federated, and cryptographic variants of the same health task on a common footing of utility, energy, and measured leakage, so architectural guidance rests on cross-paper synthesis with incompatible protocols. Shared tasks with standardized threat models, of the kind neighboring fields use to force comparability, would convert these architectural design heuristics (discussed earlier in the architectural spectrum: on-device, cloud, and hybrid inference) into evidence-based selection rules, and constitute one of the clearest infrastructure needs identified by this review.

CHALLENGES AND OPEN PROBLEMS

Four clusters of open problems structure the research frontier. The first is longitudinal privacy accounting. Differential privacy's guarantees degrade under composition, and continuous sensing composes without end; principled user-level accounting over unbounded streams, allocation of budget across concurrent sensing tasks, and semantics for guarantees that span years of participation remain unresolved, and their absence currently confines formal guarantees to coarse telemetry while the most sensitive fine-grained inference relies on architectural containment alone (Dwork and Roth, 2014; Kairouz *et al.*, 2021).

The second challenge is private computation under milliwatt budgets. The cryptographic repertoire discussed above has been benchmarked on servers and phones, rather than on the wearable and implantable tiers where continuous health signals originate; secure aggregation participation, enclave availability, and even client-level noise generation are nontrivial at that tier, and the TinyML systems work that made inference feasible there has no privacy-enhancing counterpart of comparable maturity (David *et al.*, 2021; Bonawitz *et al.*, 2017). The energy harvesting and intermittent computing literature defines the envelope this tier must respect, characterizing the reliability and latency costs of operating without a stable power source (Asiedu and Quainoo, 2023a, 2023b, 2025; Asiedu *et al.*, 2026), and system-level power modeling shows how little headroom remains once radios are accounted for (Quainoo and Ogundapo, 2026a,b).

The third is personalization under privacy constraints. Clinical value concentrates in models adapted to individuals, yet personalization is in tension with every mechanism reviewed: per-user adaptation resists aggregation, increases memorization risk, and complicates budget accounting, while federated personalization methods remain young relative to their importance for health (Smith *et al.*, 2017; Kairouz *et al.*, 2021). Related, heterogeneity across devices, populations, and care contexts threatens both fairness and validity of fleet-learned models, and private training compounds the difficulty by restricting the visibility needed to diagnose it (Li *et al.*, 2020; Wiens *et al.*, 2019). Adaptation methods that specialize a shared model through few parameters offer one route through this tension, though their privacy accounting is not yet established (Oyesiji *et al.*, 2023), and curation frameworks bear on the same problem by determining what personal material enters adaptation at all (Oyesiji *et al.*, 2024).

The fourth is assurance. The audit literature shows that deployed practice diverges from stated policy (Huckvale *et al.*, 2015; Grundy *et al.*, 2019); architectures earn trust only insofar as their properties are verifiable by outsiders, through attestation, reproducible builds, traffic transparency, and, increasingly, mechanisms by which regulators can inspect privacy budgets and data flows without themselves becoming custodians of sensitive data (Costan and Devadas, 2016; McGraw and Mandl, 2021). Assurance is also where the field's socio-technical character is sharpest: permission interfaces and consent flows are the human face of architecture, and the usability findings imply that architectural minimization must succeed even when comprehension fails (Felt *et al.*, 2012). Assurance also has a human layer, since the personnel who administer these systems determine whether attestation and logging are used as designed (Onche *et al.*, 2024, 2026), and enterprise governance scholarship argues that auditable and defensible risk management requires standing institutional functions rather than project-level attention (Ominyi *et al.*, 2026a,b). Confidentiality questions raised by large model deployment in regulated settings show how quickly assurance obligations extend beyond the sensing tier (Ominyi *et al.*, 2025), as do the threats documented for agentic systems that act on a user's behalf (Nwakamma *et al.*, 2024b).

FUTURE DIRECTIONS

Several directions appear most consequential. Foundation-scale models raise the architectural question anew: capable general models exceed device budgets by orders of magnitude, memorization risk grows with scale (Carlini *et al.*, 2021), and the field will need distillation, retrieval separation, and split designs that deliver large-model capability to health tasks without re-centralizing raw data (Hinton *et al.*, 2015; Matsubara *et al.*, 2022). Hardware trajectories favor the privacy-first stance, as neural accelerators diffuse downward through phone and wearable tiers (Ignatov *et al.*, 2019; Sze *et al.*, 2017), and co-design of private mechanisms with that silicon – on-device noise

generation, enclave-backed sensor paths, hardware-assisted aggregation – is a natural systems agenda. On the formal side, maturation of user-level, continual, and task-composed privacy accounting would unlock guarantees for exactly the longitudinal signals that define mobile health (Kairouz *et al.*, 2021), and tighter integration of differential privacy with federated personalization would address the field's central utility tension (Geyer *et al.*, 2017; Smith *et al.*, 2017). On-device work in this direction is already underway for compact multimodal models, where quantization, pruning, and distillation are combined explicitly for privacy-preserving edge deployment (Nwakamma *et al.*, 2025; Oyesiji *et al.*, 2025), and portable runtime research addresses the deployment question of moving one artifact across device, edge, and cloud tiers (Ojukwu *et al.*, 2025). Advances in energy-efficient integrated circuits and radio frequency components will determine how far down the device hierarchy such models can be pushed (Erhimefe and Gideon, 2025; Gideon, 2025).

Institutionally, the review's evaluation findings translate into an infrastructure agenda: shared benchmark tasks over realistic mobile health data with standardized threat models and mandatory multi-dimensional reporting; public reference implementations of the composed stack, local inference, secure aggregation, and accounted noise, at the quality of the production systems literature (Bonawitz *et al.*, 2019); and reporting norms under which privacy claims are stated in the three-part vocabulary of architecture, formal guarantee, and empirical resistance. Finally, regulatory convergence toward data-sensitivity-based protection (McGraw and Mandl, 2021; Cohen and Mello, 2018) suggests that privacy-first architecture will increasingly function as a compliance instrument, and research that makes architectural properties legible to regulators, auditable budgets, attestable flows, and standardized disclosures will shape how the field's technical achievements become public protections. Health system scholarship suggests where such infrastructure would land in practice, since interoperability and data governance frameworks at national scale are the layer through which any device-tier architecture reaches services (Ilodigwe and Adesemoye, 2021, 2026), and equity analyses caution that architectures presuming capable hardware and reliable connectivity will not distribute their benefits evenly (Orise and Niniola, 2024; Akinse *et al.*, 2023). Early detection systems for non-communicable disease illustrate the population value at stake if that distribution is achieved (Afrihyia *et al.*, 2023).

CONCLUSION

The architecture of mobile health computation has become a privacy decision of the first order. A decade of efficiency research moved capable inference onto the devices where health data originates, and a parallel decade of attack research established that centralized alternatives leak more than their operators intended and their users understood. Between these results, the field has assembled a genuine repertoire: on-device and staged inference that contains raw signals, split designs that narrow interfaces, federated learning that decentralizes data while centralizing improvement, differential privacy that bounds individual influence formally, and cryptographic and hardware mechanisms that protect the residue of computation that must occur elsewhere.

The synthesis offered here suggests three summary judgments. First, the architectural question is no longer constrained primarily by model feasibility; continuous health inference fits on consumer hardware, and the binding constraints have moved to updateability, evaluation, and assurance. Second, no single mechanism suffices; deployed successes compose architectural minimization with one or two formally analyzable mechanisms engineered thoroughly, and the composed stack has been demonstrated at consumer scale. Third, the field's credibility now depends less on new mechanisms than on infrastructure: comparable evaluation, longitudinal privacy semantics, low-power private computation, and verifiable claims.

The direction of travel is nonetheless clear. Sensing hardware, accelerator silicon, regulatory pressure, and public expectation are aligned in favor of architectures that keep health data where it is generated and move learning rather than data. The research reviewed here indicates that this stance is not a sacrifice of capability but a reorganization of it, and that the mobile health systems most likely to earn durable clinical and public trust are those whose privacy properties are structural, formal, and verifiable at once.

Grant Support Details

The present research did not receive any financial support to conduct the research.

Conflict of Interest

The authors declare that there is no conflict of interest regarding the publication of this manuscript. In addition, the ethical issues, including plagiarism, informed consent, misconduct, data fabrication and/ or falsification, double publication and/or submission, and redundancy, have been completely observed by the authors.

REFERENCES

1) Abadi, M., Barham, P., Chen, J., *et al.* (2016a) 'TensorFlow: a system for large-scale machine learning'. *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI '16)*. Savannah, GA, USA, pp. 265-283.
2) Abadi, M., Chu, A., Goodfellow, I., *et al.* (2016b) 'Deep learning with differential privacy'. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. Vienna, Austria, pp. 308-318.
3) Adebayo, A., Anunagba, C.O. and Ozowara, D.E. (2025) 'A review of zero trust security models and cost effectiveness in healthcare infrastructure', *International Journal of Advanced Multidisciplinary Research and Studies*, 5(6), pp. 2269-2283.
4) Adegbite, M.P., Adebayo, A. and Ahmed, M.O. (2025) 'A cyber risk quantification and governance architecture for critical infrastructure: from posture measurement to executive reporting', *International Journal of Computer Science and Mathematical Theory*, 11(12), pp. 232-293.
5) Adelanwa, A., Basnet, A. and Anene, U.N. (2024) 'Advanced AI based decision support systems for healthcare operations and resource planning', *Gyanshauryam, International Scientific Refereed Research Journal*, 7(1), pp. 244-276.
6) Afrihyia, E., Akinse, S.G. and Ojukwu, P.U. (2023) 'Development and implementation of an AI-driven early detection system for non-communicable diseases', *International Journal of Advanced Multidisciplinary Research and Studies*, 3(6), pp. 2680-2691.
7) Afrihyia, E., Akinse, S.G. and Ojukwu, P.U. (2024a) 'Comparative governance of AI-driven healthcare management: executive

oversight, regulatory structures, and accountability in the United States and developing countries', *International Journal of Advanced Multidisciplinary Research and Studies*, 4(6), pp. 3087-3102.
8) Afrihyia, E., Ojukwu, P.U. and Akinse, S.G. (2024b) 'Privacy-preserving health data governance models: a comparative review of blockchain and cryptographic strategies in U.S. and developing healthcare systems', *International Journal of Advanced Multidisciplinary Research and Studies*, 4(6), pp. 3071-3086.
9) Ahmed, K.S., Odejobi, O.D. and Oshoba, T.O. (2019) 'Algorithmic model for constraint satisfaction in cloud network resource allocation', *IRE Journals*, 2(12), pp. 516-532.
10) Ahmed, K.S., Odejobi, O.D. and Oshoba, T.O. (2020a) 'Predictive model for cloud resource scaling using machine learning techniques', *Journal of Frontiers in Multidisciplinary Research*, 1(1), pp. 173-183.
11) Ahmed, M.O., Adegbite, M.P. and Adebayo, A. (2021) 'Zero trust architecture for operational technology in North American critical infrastructure: a framework for implementation and resilience optimization', *International Journal of Engineering and Modern Technology*, 7(1), pp. 66-113.
12) Ahmed, N., Michelin, R.A., Xue, W., *et al.* (2020b) 'A survey of COVID-19 contact tracing apps', *IEEE Access*, 8, pp. 134577-134601.
13) Akinse, S.G., Ekechi, N.V. and Ohanebo, C.G. (2023) 'Integrative maternal health model: addressing socioeconomic inequities in rural populations', *International Journal of Advanced Multidisciplinary Research and Studies*, 3(6), pp. 2832-2838.
14) Akintola, A.S., Fawehinmi, Y.O., Aiyenitaju, O., *et al.* (2025) 'Digital transformation in telehealth: a systematic review of user trust, privacy, and regulatory governance in AI-powered remote monitoring systems', *Journal of Scientific Research and Reports*, 31(11), pp. 163-182.
15) Aliliele, C., Mbonu, I.S. and Iwuanyanwu, U. (2023) 'A conceptual framework for continuous cloud misconfiguration monitoring and enterprise risk mitigation strategies', *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 9(10), pp. 373-394.
16) Aliliele, C., Mbonu, I.S. and Iwuanyanwu, U. (2024a) 'A conceptual framework for enterprise data sensitivity classification and regulatory traceability mechanisms', *International Journal of Advanced Multidisciplinary Research and Studies*, 4(6), pp. 3103-3124.
17) Aliliele, C., Mbonu, I.S. and Iwuanyanwu, U. (2024b) 'Advances in HIPAA compliant data architecture and secure analytics frameworks for community healthcare organizations', *Shodhshauryam, International Scientific Refereed Research Journal*, 7(2), pp. 277-324.
18) Aliliele, C., Mbonu, I.S., Uzoka, E., *et al.* (2025) 'A review of AI assisted continuous auditing systems in technology risk and cybersecurity oversight', *Gyanshauryam, International Scientific Refereed Research Journal*, 8(4), pp. 210-250.
19) Amadi, C.E., Ale, O.K., Okorafor, U.C., *et al.* (2026) 'Unmasking the contemporary burden of hypertension in Lagos, Nigeria: insights from an opportunistic screening', *BMC Public Health*, advance online publication.
20) Anioke, S.C. and Atima, M.E. (2023a) 'Public health governance models using process optimization and performance metrics for regulatory oversight', *International Journal of Advanced Multidisciplinary Research and Studies*, 3(6), pp. 2534-2548.
21) Anioke, S.C. and Atima, M.E. (2023b) 'Business intelligence applications for mental health resource allocation and public health program accountability', *International Journal of Advanced Multidisciplinary Research and Studies*, 3(6), pp. 2549-2563.
22) Anioke, S.C. and Atima, M.E. (2023c) 'Public health informatics frameworks for protecting vulnerable populations through data-driven policy enforcement', *International Journal of Advanced Multidisciplinary Research and Studies*, 3(6), pp. 2564-2579.
23) Annan, A.O. (2022) 'Data privacy governance models for cross-border digital platforms', *International Journal of Multidisciplinary Research and Growth Evaluation*, 3(6), pp. 949-964.
24) Annan, A.O. (2024) 'Algorithmic accountability and trade secret protection in artificial intelligence', *Shodhshauryam, International Scientific Refereed Research Journal*, 7(5), pp. 315-347.
25) Asiedu, W. and Quainoo, R. (2023a) 'How far can energy harvesting take us? A systematic review of radio frequency strategies for energy autonomous sensing', *Shodhshauryam, International Scientific Refereed Research Journal*, 6(1), pp. 448-466.
26) Asiedu, W. and Quainoo, R. (2023b) 'Toward maintenance free wireless infrastructure: simulating performance and reliability in large scale intermittently powered IoT networks', *Gyanshauryam, International Scientific Refereed Research Journal*, 6(1), pp. 489-510.
27) Asiedu, W. and Quainoo, R. (2024) 'Rethinking energy, reliability, and latency trade-offs in green communication for next generation IoT', *International Journal of Multidisciplinary Research and Growth Evaluation*, 5(6), pp. 1987-1994.
28) Asiedu, W. and Quainoo, R. (2025) 'GleanCast: epoch-aligned reliable broadcast in batteryless intermittent sensor networks', *International Journal of Engineering and Modern Technology*, 11(12), pp. 205-218.
29) Asiedu, W., Quainoo, R. and Asiedu, A. (2025) 'A conceptual framework for characterizing fundamental energy, reliability, and latency trade-offs in green communication paradigms for next-generation IoT', *International Journal of Advanced Multidisciplinary Research and Studies*, 5(6), pp. 2447-2453.
30) Asiedu, W., Quainoo, R. and Asiedu, A. (2026) 'Predictive power management for intermittent IoT devices using lightweight machine learning under uncertain energy harvest', *Gulf Journal of Engineering and Technology*, 2(4), pp. 108-118.
31) Atakpa, M.I. and Abolaji, T.O. (2022) 'A privacy-preserving data architecture model for regulated industry analytics under GDPR and HIPAA compliance', *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 8(3), pp. 781-808.
32) Attia, Z.I., Noseworthy, P.A., Lopez-Jimenez, F., *et al.* (2019) 'An artificial intelligence-enabled ECG algorithm for the identification of patients with atrial fibrillation during sinus rhythm: a retrospective analysis of outcome prediction', *The Lancet*, 394(10201), pp. 861-867.
33) Avancha, S., Baxi, A. and Kotz, D. (2012) 'Privacy in mobile technology for personal healthcare', *ACM Computing Surveys*, 45(1), 3.
34) Azaria, A., Ekblaw, A., Vieira, T., *et al.* (2016) 'MedRec: using blockchain for medical data access and permission management', *Proceedings of the 2016 International Conference on Open and Big Data (OBD)*. Vienna, Austria, pp. 25-30.
35) Badmus, O., Jooda, D. and Anunagba, C.O. (2024) 'A privacy-by-design framework for role-based security architecture in Salesforce environments handling sensitive personal data', *International Journal of Advanced Multidisciplinary Research and Studies*, 4(6), pp. 3244-3254.
36) Bonawitz, K., Eichner, H., Grieskamp, W., *et al.* (2019) 'Towards federated learning at scale: system design', *Proceedings of Machine Learning and Systems*, 1, pp. 374-388.
37) Bonawitz, K., Ivanov, V., Kreuter, B., *et al.* (2017) 'Practical secure aggregation for privacy-preserving machine learning'. *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. Dallas, TX, USA, pp. 1175-1191.
38) Bot, B.M., Suver, C., Neto, E.C., *et al.* (2016) 'The mPower study, Parkinson disease mobile data collected using ResearchKit', *Scientific Data*, 3, 160011.
39) Carlini, N., Tramèr, F., Wallace, E., *et al.* (2021) 'Extracting training data from large language models'. *Proceedings of the 30th USENIX Security Symposium*. Virtual, pp. 2633-2650.
40) Carroll, A. and Heiser, G. (2010) 'An analysis of power consumption in a smartphone'. *Proceedings of the 2010 USENIX Annual Technical Conference*. Boston, MA, USA, pp. 1-14.
41) Chen, J. and Ran, X. (2019) 'Deep learning with edge computing: a review', *Proceedings of the IEEE*, 107(8), pp. 1655-1674.
42) Cohen, I.G. and Mello, M.M. (2018) 'HIPAA and protecting health information in the 21st century', *JAMA*, 320(3), pp. 231-232.
43) Coravos, A., Khozin, S. and Mandl, K.D. (2019) 'Developing and adopting safe and effective digital biomarkers to improve patient outcomes', *npj Digital Medicine*, 2, 14.
44) Costan, V. and Devadas, S. (2016) 'Intel SGX explained', *IACR Cryptology ePrint Archive*, 2016/086.
45) David, R., Duke, J., Jain, A., *et al.* (2021) 'TensorFlow Lite Micro: embedded machine learning for TinyML systems', *Proceedings of Machine Learning and Systems*, 3, pp. 800-811.
46) de Montjoye, Y.-A., Hidalgo, C.A., Verleysen, M., *et al.* (2013) 'Unique in the crowd: the privacy bounds of human mobility', *Scientific Reports*, 3, 1376.
47) Dosunmu, A.A. and Ogundele, P.O. (2024a) 'Breach and attack simulation frameworks for continuous validation of enterprise security controls', *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 10(3), pp. 1100-1119.

48) Dosunmu, A.A. and Ogundele, P.O. (2024b) 'Cyber risk quantification models for prioritizing enterprise security investment decisions', *International Journal of Multidisciplinary Research and Growth Evaluation*, 5(6), pp. 1777-1785.

49) Dwork, C. (2006) 'Differential privacy'. *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP)*. Venice, Italy, pp. 1-12.

50) Dwork, C. and Roth, A. (2014) 'The algorithmic foundations of differential privacy', *Foundations and Trends in Theoretical Computer Science*, 9(3-4), pp. 211-407.

51) Dwork, C., McSherry, F., Nissim, K., *et al.* (2006) 'Calibrating noise to sensitivity in private data analysis'. *Proceedings of the Third Theory of Cryptography Conference (TCC)*. New York, NY, USA, pp. 265-284.

52) Ejofodomi, O.A., Gideon, E.N., Oladipo, G.O., *et al.* (2014) 'Automated detection of architectural distortion in mammograms using template matching', *International Journal of Biomedical Science and Engineering*, 2(1), pp. 1-6.

53) Ekechi, N.V., Ozowara, D.E. and Anunagba, C.O. (2026) 'Conceptual framework for AI governance, data privacy compliance, and financial sustainability in digital health', *Computer Science and IT Research Journal*, 7(4), pp. 275-299.

54) El Emam, K., Jonker, E., Arbuckle, L., *et al.* (2011) 'A systematic review of re-identification attacks on health data', *PLoS ONE*, 6(12), e28071.

55) Enck, W., Gilbert, P., Chun, B.-G., *et al.* (2014) 'TaintDroid: an information-flow tracking system for realtime privacy monitoring on smartphones', *ACM Transactions on Computer Systems*, 32(2), article 5.

56) Erhimefe, D.E. and Gideon, E. (2025) 'Developing advanced materials and manufacturing processes for high-performance, energy-efficient, and scalable 3D integrated circuits (ICs) and devices', *International Journal of Research Publication and Reviews*, 6(4), pp. 2624-2630.

57) Erlingsson, Ú., Pihur, V. and Korolova, A. (2014) 'RAPPOR: randomized aggregatable privacy-preserving ordinal response'. *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security. Scottsdale*. AZ, USA, pp. 1054-1067.

58) Esteva, A., Kuprel, B., Novoa, R.A., *et al.* (2017) 'Dermatologist-level classification of skin cancer with deep neural networks', *Nature*, 542(7639), pp. 115-118.

59) Esteva, A., Robicquet, A., Ramsundar, B., *et al.* (2019) 'A guide to deep learning in healthcare', *Nature Medicine*, 25(1), pp. 24-29.

60) Fapohunda, M., Akinlolu, V.S., Omaghomi, T.T., *et al.* (2023) 'A conceptual framework for enhancing telehealth adoption among rural and underserved populations in Nigeria and the United States', I*nternational Journal of Multidisciplinary Research and Growth Evaluation*, 4(6), pp. 1348-1356.

61) Felt, A.P., Ha, E., Egelman, S., *et al.* (2012) 'Android permissions: user attention, comprehension, and behavior'. *Proceedings of the Eighth Symposium on Usable Privacy and Security (SOUPS)*. Washington, DC, USA, pp. 1-14.

62) Frankle, J. and Carbin, M. (2019) 'The lottery ticket hypothesis: finding sparse, trainable neural networks', *arXiv*, 1803.03635.

63) Fredrikson, M., Jha, S. and Ristenpart, T. (2015) 'Model inversion attacks that exploit confidence information and basic countermeasures'. *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. Denver, CO, USA, pp. 1322-1333.

64) Free, C., Phillips, G., Galli, L., *et al.* (2013) 'The effectiveness of mobile-health technology-based health behaviour change or disease management interventions for health care consumers: a systematic review', *PLoS Medicine*, 10(1), e1001362.

65) Gentry, C. (2009) 'Fully homomorphic encryption using ideal lattices'. *Proceedings of the 41st ACM Symposium on Theory of Computing (STOC)*. Bethesda, MD, USA, pp. 169-178.

66) Geyer, R.C., Klein, T. and Nabi, M. (2017) 'Differentially private federated learning: a client level perspective', *arXiv*, 1712.07557.

67) Gideon, E.N. (2025) 'Exploring and developing advanced RF MEMS switches for 5G applications, focusing on high performance solutions for RF front end modules', *Mikailalsys Journal of Advanced Engineering International*, 2(2), pp. 171-180.

68) Gilad-Bachrach, R., Dowlin, N., Laine, K., *et al.* (2016) 'CryptoNets: applying neural networks to encrypted data with high throughput and accuracy'. *Proceedings of the 33rd International Conference on Machine Learning (ICML)*. New York, NY, USA, pp. 201-210.

69) Goldwasser, S., Micali, S. and Rackoff, C. (1989) 'The knowledge complexity of interactive proof systems', *SIAM Journal on Computing*, 18(1), pp. 186-208.

70) Grundy, Q., Chiu, K., Held, F., *et al.* (2019) 'Data sharing practices of medicines related apps and the mobile ecosystem: traffic, content, and network analysis', *BMJ*, 364, l920.

71) Han, S., Mao, H. and Dally, W.J. (2016) 'Deep compression: compressing deep neural networks with pruning, trained quantization and Huffman coding', *arXiv*, 1510.00149.

72) Hannun, A.Y., Rajpurkar, P., Haghpanahi, M., *et al.* (2019) 'Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network', *Nature Medicine*, 25(1), pp. 65-69.

73) Hard, A., Rao, K., Mathews, R., *et al.* (2018) 'Federated learning for mobile keyboard prediction', *arXiv*, 1811.03604.

74) He, K., Zhang, X., Ren, S., *et al.* (2016) 'Deep residual learning for image recognition'. *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Las Vegas, NV, USA, pp. 770-778.

75) Hinton, G., Vinyals, O. and Dean, J. (2015) 'Distilling the knowledge in a neural network', *arXiv*,1503.02531.

76) Hochreiter, S. and Schmidhuber, J. (1997) 'Long short-term memory', *Neural Computation*, 9(8), pp. 1735-1780.

77) Howard, A.G., Zhu, M., Chen, B., *et al.* (2017) 'MobileNets: efficient convolutional neural networks for mobile vision applications', *arXiv*, 1704.04861.

78) Huckvale, K., Prieto, J.T., Tilney, M., *et al.* (2015) 'Unaddressed privacy risks in accredited health and wellness apps: a cross-sectional systematic assessment', *BMC Medicine*, 13, 214.

79) Idika, C.N., Salami, E.O., Ijiga, O.M., *et al.* (2021) 'Deep learning driven malware classification for cloud-native microservices in edge computing architectures', *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 7(4), pp. 710-732.

80) Ignatov, A., Timofte, R., Kulik, A., *et al.* (2019) 'AI benchmark: all about deep learning on smartphones in 2019'. *Proceedings of the 2019 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*. Seoul, South Korea, pp. 3617-3635.

81) Ike, P.C., Enyinnaya, A.C., Ujam, C.G., *et al.* (2025) 'Investigating the relationship between nutrition and DNA markers associated with heart health in middle-aged adults with type 2 diabetes', *International Journal of Convergent and Informatics Science Research*, 8(9), pp. 26-53.

82) Ilodigwe, L. and Adesemoye, A.C. (2021) 'The data backbone of health system transformation: a governance and architecture framework for interoperability, data quality, and advanced analytics at national scale', *International Journal of Health and Pharmaceutical Research*, 6(2), pp. 52-76.

83) Ilodigwe, L. and Adesemoye, A.C. (2024) 'Advances in real-world evidence, predictive analytics, and data-driven decision making for improving health system performance and patient outcomes', *International Journal of Health and Pharmaceutical Research*, 9(5), pp. 174-197.

84) Ilodigwe, L. and Adesemoye, A.C. (2025) 'Advances, risks, and implementation challenges of artificial intelligence as a force multiplier for clinical decision making and health system efficiency', *International Journal of Medical Evaluation and Physical Report*, 9(7), pp. 165-185.

85) Ilodigwe, L. and Adesemoye, A.C. (2026) 'Reimagining population health: an integrated conceptual model combining value-based care, digital health innovation, and health equity for sustainable system transformation', *International Journal of Health and Pharmaceutical Research*, 11(5), pp. 157-176.

86) Jacob, B., Kligys, S., Chen, B., *et al.* (2018) 'Quantization and training of neural networks for efficient integer-arithmetic-only inference'. *Proceedings of the 2018 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Salt Lake City, UT, USA, pp. 2704-2713.

87) Juvekar, C., Vaikuntanathan, V. and Chandrakasan, A. (2018) 'GAZELLE: a low latency framework for secure neural network inference'. *Proceedings of the 27th USENIX Security Symposium*. Baltimore, MD, USA, pp. 1651-1669.

88) Kairouz, P., McMahan, H.B., Avent, B., *et al.* (2021) 'Advances and open problems in federated learning', *Foundations and Trends in Machine Learning*, 14(1-2), pp. 1-210.

89) Kang, Y., Hauswald, J., Gao, C., *et al.* (2017) 'Neurosurgeon: collaborative intelligence between the cloud and mobile edge'. *Proceedings of the 22nd International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. Xi'an, China, pp. 615-629.
90) Klasnja, P. and Pratt, W. (2012) 'Healthcare in the pocket: mapping the space of mobile-phone health interventions', *Journal of Biomedical Informatics*, 45(1), pp. 184-198.
91) Komi, N.M. and Ganiu, O.S. (2023) 'Edge intelligence for resilient microgrid control: advances, energy sovereignty, and open challenges', *Gyanshauryam, International Scientific Refereed Research Journal*, 6(3), pp. 525-586.
92) Konečný, J., McMahan, H.B., Yu, F.X., *et al.* (2016) 'Federated learning: strategies for improving communication efficiency', *arXiv*, 1610.05492.
93) Kotz, D. (2011) 'A threat taxonomy for mHealth privacy'. *Proceedings of the 2011 International Conference on Communication Systems and Networks (COMSNETS)*. Bangalore, India, pp. 1-6.
94) Kotz, D., Gunter, C.A., Kumar, S., *et al.* (2016) 'Privacy and security in mobile health: a research agenda', *Computer*, 49(6), pp. 22-30.
95) Kumar, S., Nilsen, W.J., Abernethy, A., *et al.* (2013) 'Mobile health technology evaluation: the mHealth evidence workshop', *American Journal of Preventive Medicine*, 45(2), pp. 228-236.
96) Kumuyi, O., Akeju, B., Uzoka, E., *et al.* (2023) 'Framework for blockchain-based cross-border data exchange and regulatory transparency', *International Journal of Multidisciplinary Futuristic Development*, 4(1), pp. 99-113.
97) Kumuyi, O., Uzoka, E., Akeju, B., *et al.* (2024b) 'Framework for privacy-focused digital identity verification supporting financial inclusion in Africa', *International Journal of Advanced Multidisciplinary Research and Studies*, 4(6), pp. 3150-3162.
98) Ladapo, O.O., Dosunmu, A.A., Jooda, D., *et al.* (2022a) 'Human-in-the-loop machine learning: a state of the art', *Journal of Frontiers in Multidisciplinary Research*, 3(1), pp. 656-669.
99) Ladapo, O.O., Jooda, D., Dosunmu, A.A., *et al.* (2022b) 'Navigating digital transformation: best practices for cloud migration strategies in the enterprise', *Journal of Frontiers in Multidisciplinary Research*, 3(1), pp. 643-655.
100) Ladapo, O.O., Jooda, D., Dosunmu, A.A., *et al.* (2026) 'On the disagreement problem in human-in-the-loop federated machine learning', *International Journal of Multidisciplinary Research and Growth Evaluation*, 7(3), pp. 178-192.
101) Lane, N.D., Bhattacharya, S., Georgiev, P., *et al.* (2016) 'DeepX: a software accelerator for low-power deep learning inference on mobile devices'. *Proceedings of the 15th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. Vienna, Austria, pp. 1-12.
102) Lane, N.D., Miluzzo, E., Lu, H., *et al.* (2010) 'A survey of mobile phone sensing', *IEEE Communications Magazine*, 48(9), pp. 140-150.
103) Lara, Ó.D. and Labrador, M.A. (2013) 'A survey on human activity recognition using wearable sensors', *IEEE Communications Surveys and Tutorials*, 15(3), pp. 1192-1209.
104) LeCun, Y., Bengio, Y. and Hinton, G. (2015) 'Deep learning', *Nature*, 521(7553), pp. 436-444.
105) Li, N., Li, T. and Venkatasubramanian, S. (2007) 't-closeness: privacy beyond k-anonymity and l-diversity'. *Proceedings of the 23rd IEEE International Conference on Data Engineering (ICDE)*. Istanbul, Turkey, pp. 106-115.
106) Li, T., Sahu, A.K., Talwalkar, A., *et al.* (2020) 'Federated learning: challenges, methods, and future directions', *IEEE Signal Processing Magazine*, 37(3), pp. 50-60.
107) Machanavajjhala, A., Kifer, D., Gehrke, J., *et al.* (2007) 'l-diversity: privacy beyond k-anonymity', *ACM Transactions on Knowledge Discovery from Data*, 1(1), 3.
108) Majumder, S., Mondal, T. and Deen, M.J. (2017) 'Wearable sensors for remote health monitoring', *Sensors*, 17(1), 130.
109) Malasinghe, L.P., Ramzan, N. and Dahal, K. (2019) 'Remote patient monitoring: a comprehensive study', *Journal of Ambient Intelligence and Humanized Computing*, 10(1), pp. 57-76.
110) Mao, Y., You, C., Zhang, J., *et al.* (2017) 'A survey on mobile edge computing: the communication perspective', *IEEE Communications Surveys & Tutorials*, 19(4), pp. 2322-2358.
111) Martínez-Pérez, B., de la Torre-Díez, I. and López-Coronado, M. (2015) 'Privacy and security in mobile health apps: a review and recommendations', *Journal of Medical Systems*, 39(1), 181.
112) Matsubara, Y., Levorato, M. and Restuccia, F. (2022) 'Split computing and early exiting for deep learning applications: survey and research challenges', *ACM Computing Surveys*, 55(5), 90.
113) Mbonu, I.S., Aliliele, C., Iwuanyanwu, U., *et al.* (2020) 'A review of identity and access management integration strategies in hybrid and multi cloud environments', *International Journal of Multidisciplinary Research and Growth Evaluation*, 1(5), pp. 795-810.
114) Mbonu, I.S., Aliliele, C., Uzoka, E., *et al.* (2019) 'A review of comparative data protection regulations and secure cloud implementation strategies across jurisdictions', *Iconic Research and Engineering Journals*, 2(9), pp. 482-501.
115) Mbonu, I.S., Iwuanyanwu, U., Aliliele, C., *et al.* (2022a) 'A review of data protection impact assessment models in multi cloud financial infrastructure systems', *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 8(1), pp. 589-623.
116) Mbonu, I.S., Iwuanyanwu, U., Aliliele, C., *et al.* (2022b) 'Advances in cloud identity and access governance optimization in large scale AWS enterprise environments', *Shodhshauryam, International Scientific Refereed Research Journal*, 5(3), pp. 403-438.
117) McGraw, D. and Mandl, K.D. (2021) 'Privacy protections to encourage use of health-relevant digital data in a learning health system', *npj Digital Medicine*, 4, 2.
118) McMahan, B., Moore, E., Ramage, D., *et al.* (2017) 'Communication-efficient learning of deep networks from decentralized data'. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*. Fort Lauderdale, FL, USA, pp. 1273-1282.
119) Melis, L., Song, C., De Cristofaro, E., *et al.* (2019) 'Exploiting unintended feature leakage in collaborative learning'. *Proceedings of the 2019 IEEE Symposium on Security and Privacy (SP)*. San Francisco, CA, USA, pp. 691-706.
120) Miotto, R., Wang, F., Wang, S., *et al.* (2018) 'Deep learning for healthcare: review, opportunities and challenges', *Briefings in Bioinformatics*, 19(6), pp. 1236-1246.
121) Mohassel, P. and Zhang, Y. (2017) 'SecureML: a system for scalable privacy-preserving machine learning'. *Proceedings of the 2017 IEEE Symposium on Security and Privacy (SP)*. San Jose, CA, USA, pp. 19-38.
122) Mohr, D.C., Zhang, M. and Schueller, S.M. (2017) 'Personal sensing: understanding mental health using ubiquitous sensors and machine learning', *Annual Review of Clinical Psychology*, 13, pp. 23-47.
123) Nahum-Shani, I., Smith, S.N., Spring, B.J., *et al.* (2018) 'Just-in-time adaptive interventions (JITAIs) in mobile health: key components and design principles for ongoing health behavior support', *Annals of Behavioral Medicine*, 52(6), pp. 446-462.
124) Narayanan, A. and Shmatikov, V. (2008) 'Robust de-anonymization of large sparse datasets'. *Proceedings of the 2008 IEEE Symposium on Security and Privacy (SP)*. Oakland, CA, USA, pp. 111-125.
125) Nasr, M., Shokri, R. and Houmansadr, A. (2019) 'Comprehensive privacy analysis of deep learning: passive and active white-box inference attacks against centralized and federated learning'. *Proceedings of the 2019 IEEE Symposium on Security and Privacy (SP)*. San Francisco, CA, USA, pp. 739-753.
126) Nwakamma, S., Ojukwu, J. and Oyesiji, S.O. (2024a) 'LLM-as-a-judge for automated model governance: a review of methods, biases, and release-gating practices', *World Journal of Innovation and Modern Technology*, 8(6), pp. 185-216.
127) Nwakamma, S., Ojukwu, J. and Oyesiji, S.O. (2024b) 'Securing agentic AI enterprise workflows against prompt injection, tool poisoning, memory manipulation, and excessive agency threats', *World Journal of Innovation and Modern Technology*, 8(6), pp. 217-248.
128) Nwakamma, S., Ojukwu, J. and Oyesiji, S.O. (2025) 'On-device multimodal small language models for privacy-preserving edge intelligence through quantization, pruning, distillation, and energy-efficient inference', *World Journal of Innovation and Modern Technology*, 9(12), pp. 271-302.
129) Nwakamma, S., Ojukwu, J. and Oyesiji, S.O. (2026) 'Crypto-agile migration to hybrid post-quantum cryptography in enterprise TLS, PKI, VPN, and cloud-native systems and services', *Gulf Journal of Computer Sciences*, 2(2), pp. 16-45.
130) Ogbole, J.I., Okoruwa, P.O., Fadayomi, O., *et al.* (2021) 'Conceptual model for identity-centric zero trust architecture in enterprise security governance', *International Journal of Scientific Research in*

Computer Science, Engineering and Information Technology, 7(5), pp. 393-415.
131) Ogunboye, I., Adebayo, I.P.S., Anioke, S.C., *et al.* (2023) 'Enhancing Nigeria's health surveillance system: a data-driven approach to epidemic preparedness and response', *World Journal of Advanced Research and Reviews*, 20(1), pp. 1352-1369.
132) Ojukwu, J., Oyesiji, S.O. and Nwakamma, S. (2023) 'Cyber-physical ransomware defense in operational technology using physics-informed detection, digital twins, network telemetry, and resilient recovery', *International Journal of Computer Science and Mathematical Theory*, 9(5), pp. 193-228.
133) Ojukwu, J., Oyesiji, S.O. and Nwakamma, S. (2025) 'WebAssembly component model and WASI for portable cloud-to-edge software: language interoperability, security, performance, and container comparisons', *International Journal of Computer Science and Mathematical Theory*, 11(12), pp. 232-263.
134) Okoruwa, P.O., Fadayomi, O., Akeju, B., *et al.* (2020) 'Conceptual model for privacy-centric security engineering in digital and cloud computing systems', *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 7(5), pp. 371-389.
135) Okwah, M.N. (2022) 'Integrating triglyceride-glucose index and echocardiographic parameters for improved cardiovascular risk stratification in Sub-Saharan Africa', *International Journal of Cardiology*, 1(2), pp. 1-16.
136) Ominyi, M. (2021) 'Toward a conceptual framework for embedding privacy by design in corporate restructuring and transaction structuring', *Journal of Accounting and Financial Management*, 7(5), pp. 152-173.
137) Ominyi, M. and Anichukwueze, C.C. (2022) 'A conceptual framework for harmonizing data protection compliance across GDPR, NDPR, and United States privacy regimes', *International Journal of Social Sciences and Management Research*, 8(5), pp. 148-174.
138) Ominyi, M. and Anichukwueze, C.C. (2023) 'A review of algorithmic accountability and model risk management approaches in financial services', *Journal of Accounting and Financial Management*, 9(12), pp. 218-238.
139) Ominyi, M., Anichukwueze, C.C. and Uzougbo, N.S. (2024) 'A systematic review of organizational readiness for the EU Artificial Intelligence Act in multinational enterprises', *World Journal of Innovation and Modern Technology*, 8(6), pp. 185-203.
140) Ominyi, M., Anichukwueze, C.C. and Uzougbo, N.S. (2025) 'Reviewing legal privilege and confidentiality challenges in enterprise deployment of large language models', *Journal of Law and Global Policy*, 10(3), pp. 150-176.
141) Ominyi, M., Anichukwueze, C.C. and Uzougbo, N.S. (2026a) 'A conceptual model for integrating legal, quantitative, and operational disciplines in next-generation regulatory risk leadership', *International Journal of Social Sciences and Management Research*, 12(7), pp. 524-550.
142) Ominyi, M., Anichukwueze, C.C. and Uzougbo, N.S. (2026b) 'Conceptualizing enterprise AI governance functions as institutional standards for auditable and defensible risk management', *Journal of Accounting and Financial Management*, 12(7), pp. 326-356.
143) Onche, V.O., Adegbite, M.P. and Ogbonna, C.S. (2026) 'Human factors in information security: a capability framework for administrative professionals in the digital workplace', *World Journal of Innovation and Modern Technology*, 10(5), pp. 154-179.
144) Onche, V.O., Ogbonna, C.S. and Adegbite, M.P. (2020) 'Information-security awareness and records confidentiality in Nigerian university registries: a behavioural assessment', *International Journal of Multidisciplinary Research and Growth Evaluation*, 1(5), pp. 1032-1044.
145) Onche, V.O., Ogbonna, C.S. and Adegbite, M.P. (2024) 'Effectiveness of cybersecurity awareness training on insider-threat behavior among university administrators: an integrative academic review', *International Journal of Computer Science and Mathematical Theory*, 10(2), pp. 117-144.
146) Oparah, O.S., Ezeh, F.E., Olatunji, G.I., *et al.* (2021a) 'AI-based risk stratification framework for large-scale public health emergency preparedness and response planning', *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 7(1), pp. 332-366.
147) Oparah, O.S., Ezeh, F.E., Olatunji, G.I., *et al.* (2022) 'Big data-enabled predictive models for anticipating infectious disease outbreaks at population and regional levels', *Shodhshauryam International Scientific Refereed Research Journal*, 5(6), pp. 397-423.
148) Oparah, O.S., Ezeh, F.E., Olatunji, G.I., *et al.* (2023) 'Framework for designing national real-time disease surveillance dashboards for public health stakeholders', *Shodhshauryam, International Scientific Refereed Research Journal*, 6(1), pp. 208-227.
149) Oparah, O.S., Gado, P., Ezeh, F.E., *et al.* (2021b) 'Framework for scaling mobile health solutions for chronic disease monitoring and treatment adherence improvement', *International Journal of Advanced Multidisciplinary Research and Growth Evaluation*, 2(4), pp. 1074-1094.
150) Oparah, S.O., Gado, P., Ezeh, F.E., *et al.* (2024) 'Comprehensive review of telehealth effectiveness in bridging rural-urban disparities in healthcare access', *International Journal of Advanced Multidisciplinary Research and Studies*, 4(4), pp. 1535-1546.
151) Ordóñez, F.J. and Roggen, D. (2016) 'Deep convolutional and LSTM recurrent neural networks for multimodal wearable activity recognition', *Sensors*, 16(1), 115.
152) Orise, C. and Niniola, S. (2023) 'Privacy, safeguarding, and ethical data practices in healthcare EdTech: a conceptual model for compliance and trust', *Shodhshauryam, International Scientific Refereed Research Journal*, 6(1), pp. 507-529.
153) Orise, C. and Niniola, S. (2024) 'Digital equity in healthcare education: a systematic review of access frameworks and interventions for underrepresented learners in clinical training', *International Journal of Scientific Research in Humanities and Social Sciences*, 1(2), pp. 1233-1246.
154) Orise, C. and Niniola, S. (2025) 'Data governance in digital health learning systems: a systematic review of implementation frameworks and best practices', *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 11(5), pp. 520-538.
155) Oshoba, T.O., Hammed, N.I. and Odejobi, O.D. (2019) 'Secure identity and access management model for distributed and federated systems', *IRE Journals*, 3(4), pp. 550-567.
156) Oshoba, T.O., Hammed, N.I. and Odejobi, O.D. (2020) 'Blockchain-enabled compliance and audit trail model for cloud configuration management', *Journal of Frontiers in Multidisciplinary Research*, 1(1), pp. 193-201.
157) Owoade, O.A., Moneke, K.C. and Anioke, S.C. (2022) 'Leveraging business intelligence to optimize resource allocation in mental health and substance abuse centers', *Journal of Scientific and Engineering Research*, 9(12), pp. 210-235.
158) Oyesiji, S.O., Nwakamma, S. and Ojukwu, J. (2023) 'Parameter-efficient multilingual adaptation of on-device language models: a survey', *International Journal of Engineering and Modern Technology*, 9(3), pp. 287-320.
159) Oyesiji, S.O., Nwakamma, S. and Ojukwu, J. (2024) 'A conceptual framework for context-aware data curation in fine-tuning pipelines', *International Journal of Engineering and Modern Technology*, 10(11), pp. 197-228.
160) Oyesiji, S.O., Nwakamma, S. and Ojukwu, J. (2025) 'Compressing recommender and ranking models for edge deployment: a survey', *International Journal of Engineering and Modern Technology*, 11(12), pp. 205-237.
161) Paillier, P. (1999) 'Public-key cryptosystems based on composite degree residuosity classes'. *Proceedings of the International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*. Prague, Czech Republic, pp. 223-238.
162) Pantelopoulos, A. and Bourbakis, N.G. (2010) 'A survey on wearable sensor-based systems for health monitoring and prognosis', *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 40(1), pp. 1-12.
163) Papageorgiou, A., Strigkos, M., Politou, E., *et al.* (2018) 'Security and privacy analysis of mobile health applications: the alarming state of practice', *IEEE Access*, 6, pp. 9390-9403.
164) Patel, S., Park, H., Bonato, P., *et al.* (2012) 'A review of wearable sensors and systems with application in rehabilitation', *Journal of NeuroEngineering and Rehabilitation*, 9, 21.
165) Perez, M.V., Mahaffey, K.W., Hedlin, H., *et al.* (2019) 'Large-scale assessment of a smartwatch to identify atrial fibrillation', *New England Journal of Medicine*, 381(20), pp. 1909-1917.
166) Piwek, L., Ellis, D.A., Andrews, S. and Joinson, A. (2016) 'The rise of consumer health wearables: promises and barriers', *PLoS Medicine*, 13(2), e1001953.

167) Price, W.N. and Cohen, I.G. (2019) 'Privacy in the age of medical big data', *Nature Medicine*, 25(1), pp. 37-43.
168) Quainoo, R. and Ogundapo, O. (2026a) 'A system-level power behavior model for Bluetooth and Wi-Fi coexistence in dual-mode wireless devices', *International Journal of Computer Science and Mathematical Theory*, 12(2), pp. 298-358.
169) Quainoo, R. and Ogundapo, O. (2026b) 'Wireless system-on-chip performance in next-generation Internet of Things devices: a systematic review of integrated transceiver testing methodologies', *World Journal of Innovation and Modern Technology*, 10(5), pp. 76-126.
170) Rajkomar, A., Oren, E., Chen, K., *et al.* (2018) 'Scalable and accurate deep learning with electronic health records', *npj Digital Medicine*, 1, 18.
171) Ravi, D., Wong, C., Lo, B., *et al.* (2017) 'A deep learning approach to on-node sensor data analytics for mobile or wearable devices', *IEEE Journal of Biomedical and Health Informatics*, 21(1), pp. 56-64.
172) Rieke, N., Hancox, J., Li, W., *et al.* (2020) 'The future of digital health with federated learning', *npj Digital Medicine*, 3, 119.
173) Rocher, L., Hendrickx, J.M. and de Montjoye, Y.-A. (2019) 'Estimating the success of re-identifications in incomplete datasets using generative models', *Nature Communications*, 10, 3069.
174) Sabt, M., Achemlal, M. and Bouabdallah, A. (2015) 'Trusted execution environment: what it is, and what it is not'. *Proceedings of the 2015 IEEE International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*. Helsinki, Finland, pp. 57-64.
175) Sandler, M., Howard, A., Zhu, M., *et al.* (2018) 'MobileNetV2: inverted residuals and linear bottlenecks'. *Proceedings of the 2018 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Salt Lake City, UT, USA, pp. 4510-4520.
176) Sanni, J.O., Iwuanyanwu, U.A., Essien, M.A., *et al.* (2023) 'Lifecycle-aware marketing automation using federated learning for secure cross-organizational data management', *Gyanshauryam, International Scientific Refereed Research Journal*, 6(6), pp. 337-364.
177) Satyanarayanan, M. (2017) 'The emergence of edge computing', *Computer*, 50(1), pp. 30-39.
178) Satyanarayanan, M., Bahl, P., Cáceres, R. and Davies, N. (2009) 'The case for VM-based cloudlets in mobile computing', *IEEE Pervasive Computing*, 8(4), pp. 14-23.
179) Shamir, A. (1979) 'How to share a secret', *Communications of the ACM*, 22(11), pp. 612-613.
180) Sheller, M.J., Edwards, B., Reina, G.A., *et al.* (2020) 'Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data', *Scientific Reports*, 10, 12598.
181) Shi, W., Cao, J., Zhang, Q., *et al.* (2016) 'Edge computing: vision and challenges', *IEEE Internet of Things Journal*, 3(5), pp. 637-646.
182) Shittu, H. and Nabil, M. (2023) 'Smart supply chain management with attribute-based encryption access control'. *Proceedings of the 2023 IEEE 13th Annual Computing and Communication Workshop and Conference (CCWC)*. Las Vegas, NV, USA, pp. 198-204.
183) Shittu, M.A., Adeniji, I.O., Shittu, H., *et al.* (2022) 'Blockchain-assisted secure data exchange architectures for SCADA-controlled power systems', *Iconic Research and Engineering Journals*, 6(3), pp. 359-379.
184) Shokri, R., Stronati, M., Song, C., *et al.* (2017) 'Membership inference attacks against machine learning models'. *Proceedings of the 2017 IEEE Symposium on Security and Privacy (SP)*. San Jose, CA, USA, pp. 3-18.
185) Silva, B.M.C., Rodrigues, J.J.P.C., de la Torre Díez, I., *et al.* (2015) 'Mobile-health: a review of current state in 2015', *Journal of Biomedical Informatics*, 56, pp. 265-272.
186) Smith, V., Chiang, C.-K., Sanjabi, M., *et al.* (2017) 'Federated multi-task learning', *Advances in Neural Information Processing Systems*, 30, pp. 4424-4434.
187) Steinhubl, S.R., Muse, E.D. and Topol, E.J. (2015) 'The emerging field of mobile health', *Science Translational Medicine*, 7(283), 283rv3.
188) Sweeney, L. (2002) 'k-anonymity: a model for protecting privacy', *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(5), pp. 557-570.
189) Sze, V., Chen, Y.-H., Yang, T.-J., *et al.* (2017) 'Efficient processing of deep neural networks: a tutorial and survey', *Proceedings of the IEEE*, 105(12), pp. 2295-2329.
190) Tan, M. and Le, Q. (2019) 'EfficientNet: rethinking model scaling for convolutional neural networks'. *Proceedings of the 36th International Conference on Machine Learning (ICML)*. Long Beach, CA, USA, pp. 6105-6114.
191) Teerapittayanon, S., McDanel, B. and Kung, H.T. (2017) 'Distributed deep neural networks over the cloud, the edge and end devices'. *Proceedings of the 37th IEEE International Conference on Distributed Computing Systems (ICDCS)*. Atlanta, GA, USA, pp. 328-339.
192) Topol, E.J. (2019) 'High-performance medicine: the convergence of human and artificial intelligence', *Nature Medicine*, 25(1), pp. 44-56.
193) Torous, J. and Roberts, L.W. (2017) 'Needed innovation in digital health and smartphone applications for mental health: transparency and trust', *JAMA Psychiatry*, 74(5), pp. 437-438.
194) Tramèr, F., Zhang, F., Juels, A., *et al.* (2016) 'Stealing machine learning models via prediction APIs', *Proceedings of the 25th USENIX Security Symposium*. Austin, TX, USA, pp. 601-618.
195) Troncoso, C., Payer, M., Hubaux, J.-P., *et al.* (2020) 'Decentralized privacy-preserving proximity tracing', *arXiv*, 2005.12273.
196) Voigt, P. and von dem Bussche, A. (2017) *The EU General Data Protection Regulation (GDPR): a practical guide*. 1st ed. Cham: Springer.
197) Wang, R., Chen, F., Chen, Z., *et al.* (2014) 'StudentLife: assessing mental health, academic performance and behavioral trends of college students using smartphones'. *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*. Seattle, WA, USA, pp. 3-14.
198) Warden, P. and Situnayake, D. (2019) *TinyML: machine learning with TensorFlow Lite on Arduino and ultra-low-power microcontrollers*. 1st Ed. Sebastopol, CA: O'Reilly Media.
199) Wiens, J., Saria, S., Sendak, M., *et al.* (2019) 'Do no harm: a roadmap for responsible machine learning for health care', *Nature Medicine*, 25(9), pp. 1337-1340.
200) Xu, J., Glicksberg, B.S., Su, C., *et al.* (2021) 'Federated learning for healthcare informatics', *Journal of Healthcare Informatics Research*, 5(1), pp. 1-19.
201) Xu, M., Liu, J., Liu, Y., *et al.* (2019) 'A first look at deep learning apps on smartphones'. *Proceedings of The World Wide Web Conference (WWW)*. San Francisco, CA, USA, pp. 2125-2136.
202) Yang, Q., Liu, Y., Chen, T., *et al.* (2019) 'Federated machine learning: concept and applications', *ACM Transactions on Intelligent Systems and Technology*, 10(2), 12.
203) Yao, A.C. (1986) 'How to generate and exchange secrets'. *Proceedings of the 27th IEEE Symposium on Foundations of Computer Science (FOCS)*. Toronto, Canada, pp. 162-167.
204) Yusuf, I.A., Badero, O.J., Ihesiulo, A., *et al.* (2026) 'Aldosterone synthase inhibitors in uncontrolled and resistant hypertension: a phenotype-stratified systematic review and network meta-analysis of randomized trials', *PLOS ONE*, 21(6), e0349932.
205) Yusuff, M., Akinsola, O., Olabiyi, M., *et al.* (2025) 'Leveraging AI in drug and substance abuse recovery: a systematic approach to reintegration and rehabilitation for the homeless', *Journal of Medicine and Health Research*, 10(1), pp. 20-30.
206) Zhou, Z., Chen, X., Li, E., *et al.* (2019) 'Edge intelligence: paving the last mile of artificial intelligence with edge computing', *Proceedings of the IEEE*, 107(8), pp. 1738-1762.
207) Zhu, L., Liu, Z. and Han, S. (2019) 'Deep leakage from gradients', *Advances in Neural Information Processing Systems*, 32, pp. 14747-14756.