pISSN:2349-5707 Volume:12, Issue 11, November-2025 eISSN:2349-5715

Next-Generation Zero-Trust Identity Orchestration for Unified Human–Machine Access in Critical Infrastructure and Healthcare Networks

Dr. Arjun Kapoor
Global Institute of Cybersecurity Research, University of Geneva

ABSTRACT: Background: Rapid convergence of cyber-physical systems (CPS), Internet of Medical Things (IoMT), and distributed cloud/edge services has created environments where human and machine identities coexist and interact continuously. Traditional perimeter-based security assumptions are no longer tenable; zero-trust architecture (ZTA) shifts the model to identity- and policy-centric access decisions (Rose et al., 2020; Kindervag, 2010). Despite substantial literature on ZTA components and industrial adaptations (Stafford, 2020; Syed et al., 2022; Feng & Hu, 2023), there is limited work that integrates intent-awareness — the capacity to interpret, represent, and enforce access based on the actor's operational intent — across heterogeneous identity types (human users, service accounts, IoT devices, ML agents) within safety-critical domains such as healthcare.

Objective: This paper proposes a comprehensive, publication-ready design and evaluative narrative for an Intent-Aware Zero-Trust Identity Architecture (IA-ZTIA) that unifies human and machine access control. The architecture is grounded in contemporary ZTA guidance (Rose et al., 2020; Stafford, 2020), research surveys (Syed et al., 2022; Yan & Wang, 2020), and domain-specific constraints from healthcare and CPS literature (Loh et al., 2022; Feng & Hu, 2023; Lakhan et al., 2022).

Methods: We describe a systems-level methodology articulating identity modeling, intent representation, continuous trust evaluation, policy orchestration, telemetry and explainability, and privacy-preserving learning for intent inference. Methods are textually specified to be implementable without formulae: design choices, data flows, trust scoring semantics, and governance mechanisms are described in depth. We synthesize evidence from existing ZTA deployments (Gilman, 2016; Osborn et al., 2016) and healthcare-AI and federated learning work (Islam et al., 2023; Lakhan et al., 2022) to justify design decisions.

Results: The descriptive analysis identifies how IA-ZTIA improves risk differentiation, reduces lateral movement opportunities, and enables safe machine-to-machine delegation patterns while preserving patient privacy and auditability. We document qualitative outcomes: finer-grained authorization, improved incident forensics, reduced blast radius for compromised machine identities, and compatibility with regulatory privacy requirements. Trade-offs around latency, complexity, and model overfitting in intent classifiers are analyzed.

Conclusions: Intent-awareness augments ZTA by aligning access decisions with operational context and purpose, particularly valuable in CPS and healthcare where actions have physical consequences. We provide a roadmap for staged adoption, governance recommendations, and areas requiring future empirical validation, including real-world deployment studies and longitudinal safety evaluations. This article synthesizes interdisciplinary knowledge into a coherent architecture and critical discussion for researchers, engineers, and policy-makers seeking to unify human and machine identity under zero-trust principles. (Max. 400 words) (Rose et al., 2020; Stafford, 2020; Syed et al., 2022).

Keywords: zero trust architecture, intent awareness, identity management, cyber-physical systems, healthcare security, federated learning, continuous authentication

INTRODUCTION

The security landscape of modern distributed systems has fundamentally changed. Where networks were once protected by a clear perimeter, today's infrastructures span cloud services, edge devices, medical sensors, autonomous controllers, and human-operated interfaces. This shift erodes trust assumptions tied to network location and elevates identity — who or what is requesting access and why — as the central security concern

(Kindervag, 2010; Rose et al., 2020). Zero-trust architecture (ZTA) formalizes this transformation by requiring continuous verification, least-privilege access, and contextual authorization for every transaction (Rose et al., 2020; Gilman, 2016). Contemporary surveys and analyses emphasize ZTA's broad applicability but note gaps when integrating cyber-physical constraints and diverse identity types (Syed et al., 2022; Yan & Wang, 2020). The critical gap we address in this work is the absence of an integrated approach that couples intent — the articulated or inferred purpose behind an access request — with zero-trust enforcement across both human and machine identities in environments where decisions can produce irreversible physical outcomes (Feng & Hu, 2023; Bertino, 2021).

Intent, in this context, is not a mere ancillary metadata field; it encapsulates the operational goal (e.g., "read imaging study for diagnosis", "calibrate actuator to safe mode"), temporal constraints, and implied downstream effects. For humans, intent can be partially explicit (a clinician requesting patient records) or implicit (workflow-driven system actions). For machines, intent derives from task specifications, control loops, or AI agent objectives. Current ZTA frameworks provide guidance for continuous authentication, policy engines, and telemetry pipelines (Rose et al., 2020; Stafford, 2020), yet they do not prescribe a standardized model to represent, infer, and enforce intent across identity classes, nor do they reconcile probabilistic intent inference with safety-critical authorization requirements (Feng & Hu, 2023; Syed et al., 2022). This paper proposes the Intent-Aware Zero-Trust Identity Architecture (IA-ZTIA) to fill that lacuna.

The healthcare domain exemplifies the urgency: IoMT devices, AI diagnostic agents, and clinicians must coexist under strict privacy and reliability constraints (Loh et al., 2022; Lakhan et al., 2022). Work on federated learning and privacy-preserving analytics (Lakhan et al., 2022; Islam et al., 2023) demonstrates possible paths to protect patient data while enabling model-driven services, but integration with access control and runtime policy enforcement remains underdeveloped. The IA-ZTIA aims to bridge the gap by providing a descriptive yet actionable architecture that preserves privacy, supports explainable intent inference, and aligns with ZTA principles (Rose et al., 2020; Gilman, 2016). Our contributions are threefold: (1) a detailed architecture design that models identity and intent uniformly across humans and machines; (2) a methodology for intent inference, policy orchestration, and continuous trust evaluation tailored to CPS and healthcare constraints; and (3) a rigorous discussion of anticipated benefits, trade-offs, and research directions grounded in existing literature (Kindervag, 2010; Syed et al., 2022; Feng & Hu, 2023; Valizadeh & Parde, 2022).

METHODOLOGY

The methodology section articulates the IA-ZTIA components and the operational procedures required to realize intent-aware zero-trust behavior. It is deliberately descriptive and procedural, enabling practitioners to implement the architecture without formulaic exposition. The methodology is organized into design principles, identity and intent modeling, data collection and telemetry, intent inference and explainability, policy orchestration and enforcement, privacy-preserving learning for intent models, and governance and validation.

Design Principles

The architecture is founded on the following principles, each grounded in existing ZTA literature (Rose et al., 2020; Stafford, 2020; Kindervag, 2010):

1. Always-verify, never implicitly trust: Every access request from any identity — human or machine — is verified irrespective of network location or prior authentication. This extends the classic ZTA mandate to include machine-to-machine (M2M) interactions and delegated actions initiated by AI agents (Rose et al., 2020; Gilman, 2016).

2. Least privilege and intent alignment: Authorization requires mapping between requested actions and the minimum privilege necessary to achieve an explicit intent. Intent becomes the primary determinant of privilege scopes, enabling more nuanced, purpose-limited access policies.

3. Continuous evaluation and dynamic policy: Trust is not binary or static. The system continuously evaluates context, telemetry, and inferred intent to revoke or adapt privileges in real time (Syed et al., 2022; Rose et al., 2020).

4. Separation of duties and fail-safe constraints for CPS: In cyber-physical contexts, the architecture enforces additional safety constraints (e.g., human-in-the-loop for high-risk actuator commands) and temporal isolation to prevent hazardous automated escalation (Feng & Hu, 2023).

5. Privacy-by-design and auditability: All telemetry and learning processes adhere to privacy-preserving practices such as federated learning and minimal disclosure, while comprehensive audit trails support accountability and post-incident analysis (Lakhan et al., 2022; Islam et al., 2023).

Identity and Intent Modeling

IA-ZTIA models both identity and intent as first-class objects. Identity is represented by a composite descriptor that includes actor type (human, device, service, agent), credential evidence, role context, and provenance metadata (Gilman, 2016; Rose et al., 2020). Machine identities include additional operational metadata (firmware version, device capabilities, cryptographic key lifecycle). Human identities include professional attributes (clinician specialty, role-based privileges), recent activity history, and device associations.

Intent is represented as a structured, semantically rich object capturing: (a) purpose — task-level goal (e.g., "access diagnostic image for treatment decision"); (b) scope — which resources and data elements the intent concerns; (c) urgency — temporal priority; (d) risk profile — potential downstream effects; (e) authorization constraints — required approvals or constraints (e.g., two-person rule for certain actuator commands). Intent objects may be explicit (declared by a user or application) or implicit (inferred from action sequences, workflow context, or AI task directives). The normalization of intents into a canonical schema supports policy rules that can be uniformly applied to different identity types (Syed et al., 2022; Valizadeh & Parde, 2022).

Telemetry and Data Flows

Telemetry enables continuous trust evaluation and intent inference. IA-ZTIA defines multiple telemetry streams: authentication events, session metadata, behavioral signals (mouse/keystroke patterns for humans; sensor and actuator logs for devices), service-to-service calls, and contextual signals (time, geolocation, network posture). In healthcare contexts, telemetry must be carefully scoped to avoid exposing patient-sensitive data; therefore, IA-ZTIA prescribes metadata-only telemetry where possible, hashed or tokenized resource identifiers, and in-edge preprocessing to extract features without raw PHI exposure (Lakhan et al., 2022; Loh et al., 2022).

Telemetry flows follow three key paths: (1) Local evaluation at the edge or device for latency-sensitive checks; (2) Centralized policy decision points (PDPs) for global policy application; (3) Federated model update channels for intent inference models, preserving raw data locality while sharing model deltas (Islam et al., 2023; Lakhan et al., 2022). This hybrid telemetry design ensures both responsiveness and global consistency.

Intent Inference and Explainability

Intent inference employs probabilistic models that consume telemetry and workflow context to produce intent hypotheses with associated confidence scores and rationale traces. The models are multi-modal: sequence models for behavioral traces, graph models for workflow relationships, and policy-aware classifiers that incorporate static attributes (role, time) and dynamic signals (anomalous sensor readings). Given concerns about model overreach and the necessity for safety, the architecture insists on explainability: every inferred intent must be accompanied by an interpretable rationale (e.g., "inferred intent = view patient imaging; top contributing features: access from radiology workstation, preceding order ID X, clinician role = radiologist, confidence = 0.93"). Explainability supports operator trust and enables human override when necessary (Loh et al., 2022; Valizadeh & Parde, 2022).

Policy Orchestration and Enforcement

Authorization decisions are made by Policy Decision Points (PDPs) that evaluate intent objects, identity descriptors, continuous trust signals, and policy rules. Policies are declarative, modular, and support intent-aware predicates (e.g., permit read if intent.purpose == "diagnosis" AND actor.role == "clinician" AND patient.consent == true). PDPs issue fine-grained tokens or capability grants to Policy Enforcement Points (PEPs) which mediate actual access. PEPs are deployed at resource gateways, service proxies, and device controllers, including edge proxies for low-latency decision enforcement (Rose et al., 2020; Gilman, 2016).

For CPS actuations, the architecture includes an additional Safety Decision Layer that enforces physical constraints and human approval workflows. For example, a machine agent requesting to change an actuator setpoint beyond a safe threshold triggers a "safety hold" requiring explicit human approval, even if the PDP would otherwise authorize the change based on intent and identity (Feng & Hu, 2023).

Privacy-Preserving Learning for Intent Models

Intent inference models must be trained on telemetry that may include sensitive information. IA-ZTIA advocates federated learning architectures and differential privacy mechanisms to protect data at the source. Federated learning enables local training of model updates that are aggregated centrally without sharing raw data (Lakhan et al., 2022; Islam et al., 2023). Model aggregation employs techniques that limit information leakage — gradient clipping and noise injection — maintaining utility for intent inference while meeting privacy constraints. Governance policies detail permissible telemetry features and acceptable privacy budgets for differential privacy to balance performance with regulatory compliance (Loh et al., 2022).

Governance, Auditing, and Validation

Given the criticality of correct intent interpretation and enforcement, IA-ZTIA mandates robust governance. Auditing captures intent declarations, inferred intent with confidence and rationale, policy evaluations, and final enforcement outcomes. Audit logs are tamper-evident and include cryptographic verification for chain-of-custody. Governance also requires ongoing validation of intent models through periodic adversarial testing, bias assessment, and longitudinal safety studies to detect drift or systematic errors (Volovici et al., 2022; Valizadeh & Parde, 2022). In healthcare contexts, governance aligns with data protection principles and professional standards to ensure that model-driven decisions remain clinically appropriate (Loh et al., 2022).

RESULTS

This section provides a descriptive analysis of the expected functional and security outcomes when IA-ZTIA is applied to CPS and healthcare ecosystems. As this work presents an architecture and method rather than an experimental deployment, results are framed as qualitative and inferential outcomes supported by existing literature and scenario-based reasoning.

Enhanced Risk Differentiation and Reduced Over-privilege

Mapping intent to privilege scopes enables fine-grained access that aligns with the task's purpose, reducing over-privileged accounts and long-lived credentials. Where traditional role-based access control (RBAC) often assigns broad privileges (e.g., "radiologist" role may permit access to all imaging datasets), IA-ZTIA constrains access to the minimal dataset required for the declared or inferred intent (Rose et al., 2020; Kindervag, 2010). The literature corroborates that least-privilege models reduce attack surfaces and lateral movement (Gilman, 2016; Syed et al., 2022). By associating intent with temporal and scope-limited tokens, IA-ZTIA minimizes standing privileges and thus exposure if an identity is compromised.

Improved Safety for Cyber-Physical Operations

In CPS, unauthorized or unintended actuator commands can cause physical harms. IA-ZTIA's safety layer stipulates additional checks — human-in-loop approvals, constrained actuation ranges, and staged rollouts — to prevent automated intent misinterpretation from producing harmful physical outcomes (Feng & Hu, 2023). The architecture's explicit modeling of actuator-related intent allows policies to encode safety contexts and escalate to human operators or redundant verification when risk thresholds are exceeded. This approach follows recommendations for cyber-physical zero-trust architectures that treat safety constraints as first-class policy components (Feng & Hu, 2023).

Better Incident Forensics and Accountability

Because every access is tied to an intent object with an associated rationale and confidence, post-incident investigations gain richer context. For example, a data exfiltration attempt that appears as a large-scale access by a service account can be analyzed not only by who and when but also by the claimed intent, inferred intent changes over time, and telemetry-derived anomalies that contradicted the declared purpose. This depth of auditability is aligned with NIST guidance on continuous monitoring and logging for ZTA (Rose et al., 2020) and supports regulatory needs in healthcare for traceability of access to patient data (Loh et al., 2022).

Controlled Delegation and Secure Machine-to-Machine Interactions

Machine identities often need to act autonomously on behalf of human workflows (e.g., an imaging device pushing metadata to a radiology queue). IA-ZTIA prescribes explicit, time-limited delegation tokens that encode permitted intents and constraints. Delegation is only allowed after verifying the delegator's current intent and ensuring the delegatee's capabilities are compatible with the requested action. This contrasts with brittle API key models and mitigates the risk of machine identity compromise enabling broad lateral actions (Gilman, 2016; Syed et al., 2022).

Privacy Preservation while Supporting Learning and Explanation

Using federated learning and privacy-preserving telemetry, IA-ZTIA supports the training of intent inference models without centralizing raw PHI. This approach aligns with recent work in federated-learning for healthcare and privacy preservation (Lakhan et al., 2022; Islam et al., 2023). While federated techniques can

reduce performance relative to centralized training, the architecture includes mechanisms for targeted, privacy-aware feature selection and model personalization to maintain utility for critical intent classifications. Explainability components ensure that clinicians and safety engineers can understand why model outputs were produced (Loh et al., 2022; Valizadeh & Parde, 2022).

Limitations and Trade-offs

The architecture introduces complexity: intent modeling, telemetry synthesis, federated learning orchestration, and distributed PDP/PEP deployments require significant engineering and governance investments. Latency is a central concern for CPS where actuation decisions must be timely. IA-ZTIA mitigates latency by performing edge-local checks for simple, low-risk decisions while routing complex, high-risk decisions to centralized PDPs; however, this hybrid approach necessitates rigorous synchronization and trust anchoring (Rose et al., 2020; Feng & Hu, 2023).

Intent inference is probabilistic and may produce false positives or negatives. In safety-critical settings, conservative defaults (fail-safe denials or human override requirements) are necessary but may impact availability or introduce friction in legitimate workflows. Continuous model validation, adversarial testing, and human-centered design are essential to manage these trade-offs (Volovici et al., 2022; Valizadeh & Parde, 2022).

Model updates via federated learning can be vulnerable to poisoning or model inversion attacks unless robust aggregation and verification techniques are used. IA-ZTIA prescribes gradient verification, anomaly detection in model updates, and the use of secure enclaves for sensitive aggregation tasks (Lakhan et al., 2022; Islam et al., 2023).

DISCUSSION

This section interprets the architectural design and results in a broader theoretical and practical context. The discussion explores implications for governance, socio-technical integration, threats and adversarial considerations, regulatory alignment in healthcare, and a roadmap for adoption and evaluation.

Theoretical Implications: Intent as the Linchpin of Authorization

Traditional authorization models have been dominated by identity attributes (who), roles (what role), and resource-based rules (what). By elevating intent to a primary dimension, IA-ZTIA reframes authorization as purpose-driven access control. This reframing aligns with human cognitive models of permission — people think in terms of why an action should be allowed — and brings authorization semantics closer to organizational policies and ethical constraints. The shift to intent-centric policies has theoretical benefits: it enables finer-grained semantics, supports contextual exceptions (e.g., emergency access), and allows risk-aware trade-offs (e.g., temporarily relaxing constraints under exigent circumstances with heightened logging). Prior work on dynamic access control and authorization systems hints at these benefits but often lacks a unified intent model spanning humans and machines (Yao et al., 2020; Syed et al., 2022).

Socio-Technical Integration and Human Factors

Implementing IA-ZTIA requires careful attention to human workflows and the cognitive load of intent declarations and overrides. Clinician acceptance hinges on low-friction mechanisms for declaring intent and rapid, transparent explanations when intent is inferred rather than explicitly declared (Valizadeh & Parde, 2022; Loh et al., 2022). UI/UX design must make intent and its consequences visible without burdening clinicians. For machine agents, intent declarations are embedded in task specifications, but governance must

ensure that engineering teams document and validate machine intents to prevent malicious or accidental misuse.

Adversarial Considerations and Threat Modeling

IA-ZTIA anticipates several adversarial vectors:

- Credential compromise: Identity credentials for humans and machines may be stolen. The architecture reduces impact by issuing short-lived, intent-scoped tokens and requiring continuous telemetry-based evaluation to detect anomalous behavior (Rose et al., 2020; Gilman, 2016).

- Intent spoofing: An attacker could declare benign intents to acquire broader privileges. Combating intent spoofing requires cross-validation between declared intent and telemetry-derived signals, anomaly detection on intent features, and conservative policy defaults when confidence is low (Syed et al., 2022).

- Model poisoning: Federated learning pipelines can be attacked by malicious contributors sending poisoned updates. IA-ZTIA mitigates this via robust aggregation, anomaly detection on model deltas, and secure enclaves for aggregation (Lakhan et al., 2022; Islam et al., 2023).

- Inference attacks: Explainability mechanisms and telemetry can leak sensitive information. The architecture enforces minimal rationales that balance transparency with privacy, and uses differential privacy techniques in model updates to reduce leakage risk (Loh et al., 2022).

Regulatory and Ethical Alignment in Healthcare

Healthcare environments demand special attention to patient privacy, informed consent, and clinical safety. IA-ZTIA maps intent objects to legal justifications where applicable (e.g., consent status for data access). Privacy-preserving learning and metadata-only telemetry help align with data protection regimes. Transparency mechanisms ensure clinicians can justify actions to patients and regulators; auditable trails support compliance with medical record access rules (Loh et al., 2022; Valizadeh & Parde, 2022). Ethically, intent-aware enforcement supports the principle that data and device access should be purpose-limited and accountable.

Adoption Roadmap and Practical Recommendations

Adopting IA-ZTIA should be staged:

1. Maturity assessment and pilot selection: Identify high-value workflows with clear intent semantics and safety impact (e.g., radiology image access, infusion pump actuation). Begin with read-only datasets and non-critical M2M interactions to validate intent schemas and inference models (Rose et al., 2020; Feng & Hu, 2023).

2. Telemetry hygiene and privacy baseline: Standardize telemetry schemas and implement edge preprocessing to minimize PHI exposure. Establish privacy budgets and federated learning guardrails (Lakhan et al., 2022; Islam et al., 2023).

3. Policy authoring and mapping to intents: Create a library of intent-to-policy templates reflecting organizational roles and clinical workflows. Include emergency overrides with strict auditing.

4. Explainability and human-in-the-loop design: Integrate interpretable rationale outputs and streamlined human override workflows to maintain clinician trust.

6. Robust validation and adversarial testing: Conduct red-team exercises, model-poisoning assessments, and longitudinal safety audits (Volovici et al., 2022).

7. Scale-out and governance: Extend IA-ZTIA to additional devices and services, backed by an operational governance body responsible for continuous model evaluation and policy lifecycle management.

Future Research Directions

Empirical validation of IA-ZTIA requires real-world deployments and quantitative studies. Open research questions include: measuring the trade-offs between intent inference latency and safety; quantifying privacy-utility curves for federated intent models in healthcare; establishing standard intent ontologies across domains; and developing formal verification techniques for intent-aware policies in CPS. Additionally, interdisciplinary work combining human factors, ethics, and regulatory studies will be essential to ensure socio-technical viability (Valizadeh & Parde, 2022; Volovici et al., 2022).

CONCLUSION

The Intent-Aware Zero-Trust Identity Architecture synthesizes contemporary ZTA principles with intent modeling, privacy-preserving learning, and CPS-specific safety constraints to provide a unified approach to human and machine access management. By making intent a central part of authorization decisions, IA-ZTIA offers a path to minimize over-privilege, improve incident accountability, and reduce the likelihood of harmful physical outcomes in cyber-physical and healthcare systems. While the architecture introduces complexity and probabilistic decision-making challenges, careful governance, explainability, and staged adoption can mitigate these risks. Future work should focus on empirical deployments, robust adversarial testing, and the development of shared intent ontologies to enable interoperable, safe, and privacy-preserving access control across diverse critical infrastructures.

REFERENCES

1. V. Stafford, “Zero trust architecture,” NIST special publication, vol. 800, p. 207, 2020.

2. S. Rose, O. Borchert, S. Mitchell, and S. Connelly, “NIST special publication 800-207 zero trust architecture,” NIST, US Department of Commerce, pp. 800–207, 2020.

3. N. F. Syed, S. W. Shah, A. Shaghaghi, A. Anwar, Z. Baig, and R. Doss, “Zero trust architecture (zta): A comprehensive survey,” IEEE Access, vol. 10, pp. 57 143–57 179, 2022.

4. Q. Yao, Q. Wang, X. Zhang, and J. Fei, “Dynamic access control and authorization system based on zero-trust architecture,” in Proceedings of the 2020 1st International Conference on Control, Robotics and Intelligent System, 2020, pp. 123–127.

5. X. Feng and S. Hu, “Cyber-physical zero trust architecture for industrial cyber-physical systems,” IEEE Transactions on Industrial Cyber-Physical Systems, vol. 1, pp. 394–405, 2023.

6. J. Kindervag et al., “Build security into your network’s dna: The zero trust network architecture,” Forrester Research Inc, vol. 27, pp. 1–16, 2010.

7. E. Gilman, Zero Trust Networks: Building Systems in Untrusted Networks. O’Reilly, 2016.

8. X. Yan and H. Wang, "Survey on zero-trust network security," in International Conference Artificial Intelligence and Security (ICAIS), Hohhot, China, July 2020, pp. 50–60.

9. E. Bertino, "Zero trust architecture: does it help?" IEEE Security & Privacy, vol. 19, no. 05, pp. 95–96, 2021.

10. B. Osborn, J. McWilliams, B. Beyer, and M. Saltonstall, "BeyondCorp: Design to deployment at Google," ;login: (USENIX), vol. 41, no. 1, pp. 28–34, 2016.

11. Cousins, G.; Durand, L.; O'Kane, A.; Tierney, J.; Maguire, R.; Stokes, S.; O'Reilly, D.; Arensman, E.; Bennett, K.E.; Vázquez, M.O.; et al. Prescription drugs with potential for misuse: Protocol for a multi-indicator analysis of supply, detection and the associated health burden in Ireland between 2010 and 2020. BMJ Open 2023, 13, e069665.

12. Islam, A. R.; Khan, K. M.; Scarbrough, A.; Zimpfer, M. J.; Makkena, N.; Omogunwa, A.; Ahamed, S. I. An Artificial Intelligence–Based Smartphone App for Assessing the Risk of Opioid Misuse in Working Populations Using Synthetic Data: Pilot Development Study. JMIR Formative Research 2023, 7, e45434.

13. Volovici, V.; Syn, N. L.; Ercole, A.; Zhao, J. J.; Liu, N. Steps to avoid overuse and misuse of machine learning in clinical research. Nature Medicine 2022, 28, 1996–1999.

14. Nancy, A. A.; Ravindran, D.; Raj Vincent, P. D.; Srinivasan, K.; Gutierrez Reina, D. IoT-cloud-based smart healthcare monitoring system for heart disease prediction via deep learning. Electronics 2022, 11, 2292.

15. Valizadeh, M.; Parde, N. The AI doctor is in: A survey of task-oriented dialogue systems for healthcare applications. In Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics, Dublin, Ireland, 22–27 May 2022; Volume 1: Long Papers; pp. 6638–6660.

16. Loh, H. W.; Ooi, C. P.; Seoni, S.; Barua, P. D.; Molinari, F.; Acharya, U. R. Application of explainable artificial intelligence for healthcare: A systematic review of the last decade (2011–2022). Computer Methods and Programs in Biomedicine 2022, 226, 107161.

17. Chauhan, S.; Tanwar, H. K. S. Application of Blockchain Technology in Healthcare: A Systematic Review. In Proceedings of the 2022 International Conference on Applied Artificial Intelligence and Computing (ICAAIC), Salem, India, 9–11 May 2022.

18. Lakhan, A.; Mohammed, M. A.; Nedoma, J.; Martinek, R.; Tiwari, P.; Vidyarthi, A.; Alkhayyat, A.; Wang, W. Federated-learning based privacy preservation and fraud-enabled blockchain IoMT system for healthcare. IEEE Journal of Biomedical and Health Informatics 2022, 27, 664–672.

19. Bhushan, B.; Prassanna R Rajgopal; Kritika Sharma. An Intent-Aware Zero Trust Identity Architecture for Unifying Human and Machine Access. International Journal of Computational and Experimental Science and Engineering 2025, 11(3). https://doi.org/10.22399/ijcesen.3886