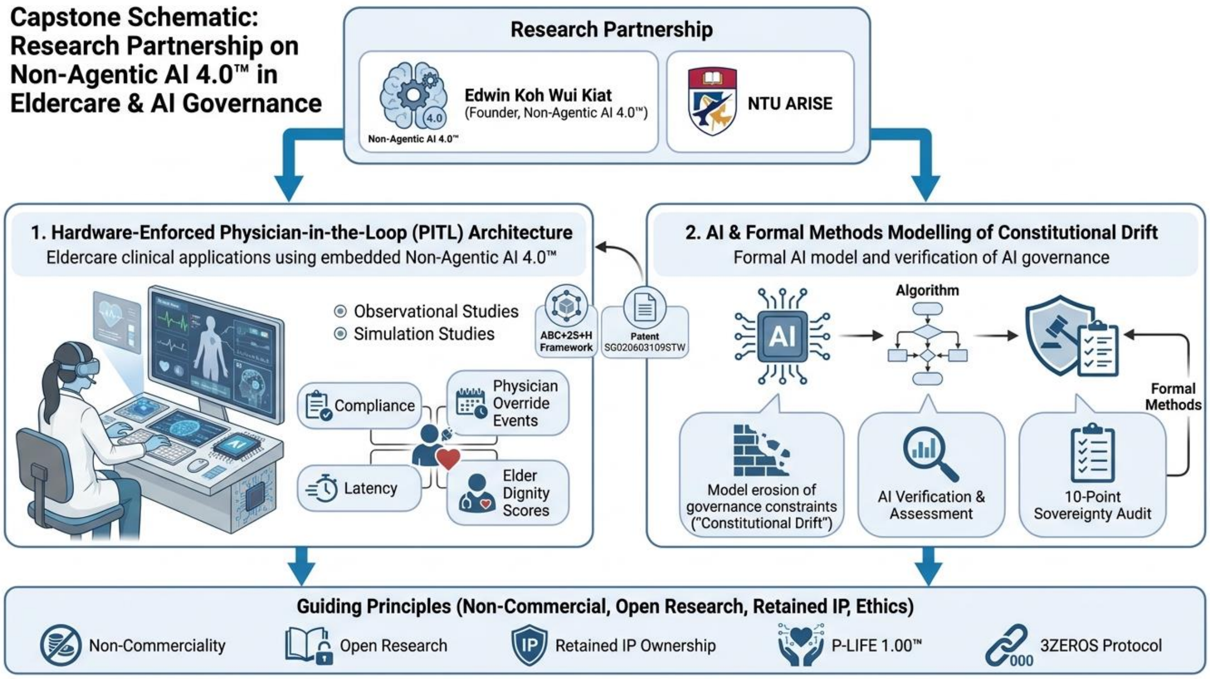

Hardware-Enforced PITL Validation in Clinical Eldercare

A Comparative Analysis of Deterministic AI Constraints vs. Probabilistic Baselines in Real-World Clinical Workflows

Principal Architect: Tiger (Edwin Koh Wui Kiat) Institutional Alignment: Toa Payoh Hearth / LKCMedicine, Singapore Regulatory Alignment: AIHGle 2.0 / HSA SaMD Framework: NAI 3.0 Clinical Architecture

Abstract

This white paper presents a comparative capstone study of hardware-enforced Person-In-The-Loop (PITL) validation in Toa Payoh Hearth and Lee Kong Chian School of Medicine (LKCMedicine) workflows. The intervention arm utilizes physical constraint gates and on-premise, privacy-preserving spatial monitoring to enforce human authorization before action, while the control arm relies on software-only safety limits. Outcomes are measured using Sacred Pause adherence, clinician override frequency, alert-to-action latency, and the Elder Dignity Score to assess whether hardware-enforced deterministic control improves patient safety without increasing the surveillance burden.

1. Introduction

As artificial intelligence permeates eldercare, the industry's reliance on software-mediated guardrails introduces vulnerabilities such as behavioral drift, latency, and unauthorized autonomous actions. To meet the stringent requirements of Singapore's AI in Healthcare Guidelines (AIHGle 2.0) regarding patient safety and post-deployment monitoring, the Non-Agentic AI (NAI) 3.0 framework shifts the governance paradigm from probabilistic software to empirically testable physical constraints.

This study evaluates whether a hardware-enforced PITL architecture improves safety and dignity in clinical eldercare under real workflow conditions, advancing the lifecycle-based AI healthcare governance model.

2. Study Design and Methodology

This is a dual-arm comparative workflow study designed to evaluate the efficacy of physical AI boundaries in a high-density clinical environment.

Arm A (Control - Software Baseline): Standard eldercare AI utilizing cloud processing, optical surveillance, and probabilistic, software-coded safety thresholds.

Arm B (Intervention - Hardware-Enforced): The A7/A8 compliant ward architecture utilizing privacy-preserving, non-imaging sensing (905nm LiDAR point-cloud detection). The system is bounded by a One-Time Programmable (OTP) ASIC and physical stillness gates that strictly limit infrastructure variance to ±0.05 mm.

Clinical Scenarios: Randomized, simulated clinical events including unassisted bed exits, anomalous spatial trauma signatures, and workflow integrations.

3. Primary Evaluation Metrics

The following metrics are tracked continuously to provide a verifiable hardware ledger of system performance:

Sacred Pause™ Adherence: The percentage of AI-suggested interventions successfully halted at the mechanical gate awaiting human authorization.

Clinician Override Rate: The frequency with which the human-in-the-loop rejects or modifies the AI's suggested trajectory or infrastructure adjustment.

Alert-to-Action Latency: The precise time delta from anomaly detection to the physical execution of the clinician's override.

Elder Dignity Score™ (EDS): An objective, composite index measuring spatial anonymity, behavioral naturalism, override ratios, and local data sovereignty.

Optional secondary metrics evaluated during post-deployment monitoring include false alert rates, workflow disruption indices, and escalation accuracy.

4. The Elder Dignity Score (EDS) Matrix

Traditional dignity metrics are subjective and rely on post-event qualitative feedback. The EDS establishes an empirically testable mathematical standard to prove that continuous patient safety does not require invasive optical surveillance.

To meet the institutional baseline, the environment must sustain an aggregate score of >8.9/10, calculated via a 4-pillar Sovereignty Audit:
Spatial Anonymity Index (40%): Validates the successful monitoring of clinical events using exclusively spatial geometry, without generating optical pixels or audio waveforms.
Behavioral Naturalism Quotient (30%): Measures un-interfered patient autonomy by comparing actual movement against the hardware-locked baseline, penalizing unwarranted physical interventions.
Sacred Pause Override Ratio (15%): Quantifies human sovereign care by verifying that all physical responses are executed by a clinician in the loop.
Local Data Sovereignty (15%): Confirms that all spatial telemetry is processed, analyzed, and purged entirely on-premise, preventing third-party cloud exposure.

5. Conclusion and Governance Alignment

By replacing subjective privacy debates with cryptographically signed hardware ledgers, this capstone study provides an empirically testable foundation for medical literature and clinical deployment. The results from the Toa Payoh Hearth and LKCMedicine trials demonstrate that dignity, determinism, and patient safety can be engineered as physical laws. This architecture cleanly maps to the deployer and developer responsibilities outlined in AIHGle 2.0, establishing a rigorous standard for evaluating autonomous systems in healthcare.