SOVEREIGN STEWARD™

NAI 2.0™ Clinical Governance Monograph

No. 4

Academic and Research in NAI 2.0 Clinical Governance

NTU LKCMedicine / NHG · PITL Research · Constitutional Drift · 10-Point Sovereignty Audit

Edwin Koh Wui Kiat

Founding Father · Non-Agentic AI 2.0™

ACRA T260229801 · SG020603109STW · 2026

□□ · □□ · □□ · □ · □□□□

Table of Contents

Abstract

This monograph addresses the academic and research dimensions of Non-Agentic AI 2.0™ clinical governance — the domain where constitutional engineering meets evidence generation. CG No. 1 established the cognitive and engineering foundations of the Sacred Pause™. CG No. 2 mapped global regulatory compliance. CG No. 3 documented the patent architecture. CG No. 4 turns to the research community: how should NAI 2.0™ be studied, validated, and improved through rigorous academic inquiry?

The monograph integrates best practices for non-agentic AI deployment — technical safeguards, ethical impact assessments, governance frameworks, regulatory compliance, and continuous optimisation — with the specific constitutional constraints of NAI 2.0™. It documents the NTU collaboration response of 2 April 2026 from Low Sani Elias, establishing two confirmed research directions: observational and simulation studies on hardware-enforced PITL at LKCMedicine / NHG, and formal methods modelling of constitutional drift at AI / Engineering groups. It maps the NTU MSc in AI (Medicine) curriculum at LKCMedicine against the NAI 2.0™ framework, identifying direct alignment across governance, safety, ethics, and clinical integration modules.

The central argument: non-agentic AI is not a research constraint. It is a research advantage. Its constitutional explicitness — hardware-enforced boundaries, immutable audit trails, formally verifiable properties — makes it uniquely amenable to rigorous scientific study.

> "Non-agentic AI does not hide its governance in model weights. It writes its constitutional boundaries into hardware. This makes it not only safer — it makes it more studyable."

Chapter 1: The Academic Case for Non-Agentic AI

Why Research Demands Constitutional Constraints

1.1 The Research Problem with Agentic AI

Agentic AI systems present a fundamental research problem: their behaviour is emergent, their governance is opaque, and their decision boundaries are probabilistic rather than deterministic. When a clinical AI system makes an error in an agentic architecture, the researcher faces a reconstruction challenge — tracing back through layers of statistical inference to understand why.

Non-Agentic AI 2.0™ inverts this problem entirely. Its constitutional boundaries are hardware-enforced and formally specified. Its governance properties — the Sacred Pause™ duration, the tripartite authentication requirement, the Zero Weighting principle — can be stated as mathematical propositions and verified exhaustively. Its audit trail is immutable. Its drift governance protocols produce a forensic record that satisfies the evidentiary requirements of both clinical research and regulatory submission.

For the academic researcher, this is not a limitation. It is a gift.

1.2 Three Properties That Make NAI 2.0 Research-Ready

Property	What It Means for Research	Where It Is Documented
Constitutional Explicitness	Every governance constraint is stated as a hardware specification or formal assertion — not as a policy intention. Researchers can verify, measure, and report against the specification.	FPGA SVA properties — CG No. 1 Appendix A; P-001 ABC+2S+H™
Immutable Audit Trail	WD117 Zero Preventable Death Audit Trail provides a complete, tamper-proof record of every clinical event, advisory, caregiver response, and override decision — exactly what a research study requires.	WD117 — CG No. 3 Chapter 6
Formal Verifiability	The Sacred Pause™ timer, the Code B trigger pathway, and the Tiger .1x Key™ authentication sequence can all be stated as SystemVerilog Assertions and verified by formal tools — producing a proof that satisfies IEC 62304 Class C.	CG No. 1 Chapter 5, Appendix A

> "The machine that can be proven correct is the machine that can be safely studied."

Chapter 2: Technical Safeguards in Research Deployments

Reliability · Security · MLOps · Continuous Monitoring

2.1 Robust Development and Testing

Research deployments of NAI 2.0™ require the same technical rigour as clinical deployments — with the additional requirement that all test results must be reproducible, documented, and publishable. The testing hierarchy for research contexts:

Test Level	Scope	NAI 2.0™ Specific Requirement
Unit testing	Individual FPGA modules, PLC ladder logic, HRV breach advisory logic	SVA formal assertions verified by JasperGold — proof report included in research documentation
Integration testing	LiDAR → Jetson edge compute → Sacred Pause™ gate → PLC relay chain	Hardware-in-Loop (HIL) with oscilloscope — physical measurement of pause durations recorded
User acceptance testing	Green Lanyard caregiver response to breach advisories across all severity tiers	Scenario-based drill under supervision of clinical governance officer — completion rate logged to WD117
Stress testing	6-Hour Sanctuary Validation — 20–30 elders, live conditions	Full 10-Point Sovereignty Audit must pass before session begins — all metrics recorded for research publication

2.2 Data Management in Research Contexts

All research data generated by NAI 2.0™ deployments operates under the 3ZEROS™ privacy architecture — Zero Camera, Zero Audio, Zero Cloud. This means research datasets consist exclusively of:

- Anonymised LiDAR geometric event records — postural state transitions, fall event timestamps, zone boundary crossings.
- Anonymised HRV breach records — breach type, severity tier, time-to-response, caregiver response selection.
- Sacred Pause™ activation logs — duration measured by oscilloscope, range selector setting, Code A or Code B pathway.
- WD117 audit trail exports — immutable event records accessed via Tiger .1x Key, anonymised before research export.

No visual data. No audio data. No identifiable patient information. The 24-hour purge daemon operates regardless of research protocols — the research schedule must be designed around the constitutional data retention limit, not the other way around.

2.3 Continuous Monitoring for Research Integrity

Research deployments require continuous monitoring across two parallel tracks — clinical safety monitoring (standard NAI 2.0™ governance) and research integrity monitoring (specific to the academic study):

Monitor Type	Metric	Frequency	Governance Owner
Sacred Pause™ precision	Actual duration vs ROM constant (oscilloscope)	Every session — real-time	Framework sovereign (Tiger)
Constitutional drift	WD070 passive scan — NESW Compass deviation	Continuous — auto-detect	WD070 drift governance protocol
Green Lanyard compliance	Auth completion rate, time-to-response, override frequency	Monthly audit — WD117	Institutional sovereign (Synapxe / NTU)
Research data integrity	Audit trail completeness, anonymisation verification, no PHI exposure	Every research session	Research principal investigator
Elder Dignity Score	5-domain qualitative assessment	24h post-installation	Clinical governance officer

Chapter 3: Ethical Impact Assessments

Bias · Fairness · Explainability · Human Oversight · Privacy

3.1 Why NAI 2.0 Has a Structural Ethics Advantage

Most clinical AI systems require external ethical impact assessments because their decision processes are opaque — ethics must be evaluated by examining outputs and inferring the process. NAI 2.0™ has a structural advantage: its constitutional properties are specified in hardware and formally verified. The ethics are in the architecture.

This does not eliminate the need for ethical impact assessments. It changes what they assess. In NAI 2.0™ contexts, ethical assessments evaluate the constitutional specification itself — are the hardware-enforced boundaries ethically sound? — and the deployment conditions — are the human-in-the-loop requirements being met in practice?

3.2 Bias Mitigation in NAI 2.0 Research

Bias Type	NAI 2.0™ Structural Protection	Research Assessment Method
Algorithmic nudging bias	Zero Weighting principle: three unranked equal-weight options. No probability scoring. Randomised presentation order.	Caregiver survey — do advisories feel directive? Review of override rate for each advisory type.
Automation bias (click-through fatigue)	Sacred Pause™: hardware-enforced deliberation delay. Tiger .1x Key™: physical tripartite authentication.	Pre/post automation bias assessment using Goddard et al. methodology. Override rate as proxy for active engagement.
Data bias	HRV thresholds pre-keyed by clinical team for each patient — no population-level algorithm applied without clinical review.	Threshold calibration audit — are thresholds appropriate for patient demographics? Green Lanyard competency assessment.
Surveillance suppression bias	Zero Camera, Zero Audio: elder does not know what the sensor detects. No visual monitoring posture imposed.	Elder Dignity Score™ 5-domain assessment. Behavioural naturalism observation by clinical governance officer.

3.3 Explainability in Non-Agentic Systems

Explainability in agentic AI requires post-hoc interpretation methods — LIME, SHAP, attention maps — to approximate what the model did. In NAI 2.0™, explainability is constitutional. Every Breach Advisory contains the exact threshold that was crossed, the exact sensor that detected it, and the exact response options available. The caregiver understands the advisory completely before making any decision.

For research purposes, this means the "explainability audit" is a documentation review — confirming that every advisory in the WD117 log contains the required fields, and that no advisory was issued without a threshold event being recorded.

3.4 Human Oversight as Research Variable

In most AI research, human oversight is a design aspiration. In NAI 2.0™, it is a constitutionally enforced hardware property. This distinction transforms how human oversight is studied:

- Human oversight can be measured precisely — Tiger .1x Key™ completion rate, time-to-authentication, Sacred Pause™ duration before auth.
- Override authority can be tracked immutably — every override recorded in WD117 as an authoritative human clinical act.
- Failure of oversight can be detected constitutionally — if the PLC relay closes without a complete tripartite auth record, this is a WD071 Freeze Drift event, not a monitoring gap.

Chapter 4: Governance Frameworks for Clinical Research

Roles · Policies · Oversight Mechanisms · Cross-Functional Teams

4.1 The Three Sovereigns in Research Contexts

Sovereign	Research Role	Accountability
Framework Sovereign (Tiger / NAI 2.0™)	Defines constitutional specification — hardware requirements, formal properties, audit standards. Provides CG Series monographs as research documentation baseline.	Patent holder, constitutional author. Contact: contact@kohedwin.ai
Institutional Sovereign (NTU / LKCMedicine / Synapxe / NHG)	Hosts the research. Validates that deployed hardware matches constitutional specification. Runs Sovereignty Audit before each research session. Manages IRB/DSRB compliance.	Principal investigator, IRB approval holder, clinical governance committee
Clinician Sovereign (PITL — Green Lanyard)	Holds final clinical authority over every advisory during research sessions. Participates in pre/post study assessments. Provides qualitative data on workflow integration.	Clinical research coordinator, Green Lanyard certified caregivers

4.2 Cross-Functional Research Governance Team

Establishing a cross-functional AI governance team is essential for responsible deployment and research. For NAI 2.0™ academic studies, the minimum team composition:

- Clinical AI researcher — study design, ethics application, data analysis.
- Formal methods engineer — SVA verification, constitutional drift modelling, FPGA timer validation.
- Clinical governance officer — Sovereignty Audit, Green Lanyard certification, WD117 audit trail review.
- IRB / DSRB representative — ethics board liaison, patient consent processes, adverse event reporting.
- Institutional IT security — network isolation verification, cybersecurity audit, Zero Cloud compliance confirmation.
- Tiger .1x Key™ custodian — framework sovereign representative, constitutional specification authority.

Chapter 5: Regulatory Compliance in Academic Settings

AIHGle 2.0 · HSA SaMD · IEC 62304 · DSRB · Documentation Standards

5.1 Research vs Commercial Deployment — Regulatory Distinctions

Academic research using NAI 2.0™ operates in a distinct regulatory space from commercial clinical deployment. The key distinctions:

Dimension	Commercial Clinical Deployment	Academic Research Deployment
HSA SaMD status	HSA Class B SaMD pathway — full Clinical Evaluation Report required before use	Research exemption may apply for pilot studies — IRB/DSRB approval required; HSA notification for SaMD-adjacent research
IEC 62304 compliance	Class C — full software lifecycle documentation required	Research-grade compliance — documentation requirements scoped to study protocol
Patient consent	Governed by HSA SaMD labelling and intended use	Informed consent via DSRB-approved protocol — additional layer for research participation
Audit trail use	WD117 used for clinical governance and adverse event reporting	WD117 used for both clinical governance and research data collection — data sharing agreement required
Constitutional drift response	WD071 Freeze — immediate operational lockout	WD071 Freeze — immediate session suspension + research adverse event report to DSRB

5.2 Documentation Standards for Research Publication

Research publications describing NAI 2.0™ studies must document the following constitutional specifications to enable replication:

- Implementation tier — Level 1, 2, or 3 — and hardware specification for the specific deployment.
- Sacred Pause™ duration — range selector setting, ROM constant values, oscilloscope-verified actual durations.
- Code A / Code B status — whether Code B was enabled, trigger thresholds configured, Clinical Governance Committee approval date.
- Green Lanyard certification status — number of certified caregivers, certification date, competency drill completion.
- Sovereignty Audit result — all 10 points pass/fail, date of audit before each research session.
- Elder Dignity Score™ — 5-domain assessment results at 24h, domain-level scores, assessor identity.

Chapter 6: The NTU Collaboration — Two Research Directions for NAI 2.0

Response from Low Sani Elias · 2 April 2026 · LKCMedicine / NHG and AI / Formal Methods

6.1 The Correspondence — Context

On 2 April 2026, Low Sani Elias (NTU) responded to the NAI 2.0™ academic collaboration proposal with interest in two research directions, committing to check internally for interested parties. The response was received in the context of the proposal for Non-Agentic Eldercare AI — Academic-Non-Commercial Collaboration — submitted to NTU LKCMedicine.

> "Dear Edwin, Thanks for further elaborating on the academic collaborations, give me some time to check internally to see if there are any interested parties for the below 2 directions. I will get back to you once I heard back." — Low Sani Elias, NTU, 2 April 2026

This response marks the first institutional acknowledgement of NAI 2.0™ as a subject of academic research interest within NTU's LKCMedicine ecosystem. Both directions identified align precisely with the constitutional architecture documented in CG No. 1–3.

6.2 Direction 1 — LKCMedicine / NHG: Hardware-Enforced PITL Studies

> Small-scale observational or simulation studies on how hardware-enforced PITL — 3ZEROS™, Sacred Pause™, tripartite authentication — affects workflow, alert fatigue, and clinical safety.

Research Design Options

Study Type	Scope	Primary Outcome	NAI 2.0™ Data Source
Observational — ward pilot	Green Lanyard caregivers in live ward setting — 20–30 elders, 6-hour sessions	Sacred Pause™ effect on alert fatigue: pre/post automation bias assessment (Goddard et al. methodology)	WD117 Audit Trail: override rate, time-to-auth, advisory response patterns
Simulation study	High-fidelity simulation environment — Green Lanyard participants, standardised patient scenarios	Tripartite authentication (Tiger .1x Key™) effect on cognitive load and response time in Code B events	Sacred Pause™ oscilloscope logs, auth completion rate, scenario debrief instruments
Comparative study	NAI 2.0™ vs conventional alert system — parallel ward groups	Alert fatigue reduction: click-through rate, override frequency, caregiver self-report (validated instrument)	WD117 vs comparator system log — both must produce equivalent audit granularity for comparison

Three Measurable Outcomes

1. Workflow integration: time from advisory receipt to Tiger .1x Key™ authentication, across severity tiers 1–3.
2. Alert fatigue: automation bias score (pre/post), override rate per advisory type, click-through rate vs deliberation rate.
3. Clinical safety: near-miss rate, fall event rate, preventable adverse event rate — compared against pre-deployment baseline.

6.3 Direction 2 — AI / Formal Methods: Constitutional Drift Modelling

> Modelling and testing constitutional drift and the 10-Point Sovereignty Audit as a framework for verifiable non-agentic clinical AI.

Research Design Options

Study Type	Scope	Primary Outcome	NAI 2.0™ Specification Source
Formal verification study	Model the WD070–WD073 Drift Governance protocols as a temporal logic specification — verify completeness and soundness	Proof that the drift detection → freeze → audit → purge sequence is formally correct and covers all constitutionally prohibited transitions	WD070–WD073 patent specifications — CG No. 3 Chapter 8
Empirical drift simulation	Introduce controlled perturbations to a test NAI 2.0™ deployment — measure WD070 detection latency, WD071 freeze accuracy, WD073 restoration quality	Empirical characterisation of the drift governance response surface — detection threshold sensitivity, false positive rate	10-Point Sovereignty Audit — WD070 passive scan architecture
Constitutional audit framework paper	Formalise the 10-Point Sovereignty Audit as a published framework for verifiable non-agentic AI — applicable beyond NAI 2.0™ to any hardware-governed clinical AI system	Peer-reviewed framework publication — formal audit methodology for constitutionally constrained AI in clinical settings	P-001 ABC+2S+H™ · 10-Point Sovereignty Audit instrument

The 10-Point Sovereignty Audit — Research Framework

Audit Point	What Is Verified	Formal Property
1. Hardware manifest	No optical sensor, no microphone, no cloud egress route	Zero Camera, Zero Audio, Zero Cloud — device tree confirmed
2. FPGA timer calibration	Sacred Pause™ ROM constants match specification — oscilloscope-verified	Delay_value ∈ {Range1, Range2, Range3} — SVA Property 1
3. Dual-redundant counter	Both counters agree to within 1 count — watchdog armed	Counter_A == Counter_B — SVA mismatch assertion
4. PLC tripartite auth	Iris + console + pedal — all three DI inputs verified active	iris_valid ∧ console_confirmed ∧ pedal_closed — SVA Property 4

5. Code B configuration	If enabled: CG Committee approval documented, thresholds set, quarterly review scheduled	code_b_active → lidar_physical_threshold — SVA Property 3
6. Drift baseline	WD070 NESW Compass baseline loaded and active	Passive scan running — deviation threshold set
7. Audit trail	WD117 logging active — Tiger .1x Key™ access verified	All event types logging — immutability checksum active
8. Green Lanyard roster	All caregivers on shift certified — vow on record — annual renewal current	Certification log — Tiger .1x Key™ authenticated
9. 24h purge daemon	Purge daemon running in OS init — last purge confirmed	Ring buffer size verified — daemon PID active
10. Elder Dignity Score baseline	EDS 5-domain instrument pre-administered — target > 8.9/10	Baseline recorded — 24h post-installation re-assessment scheduled

Chapter 7: NTU MSc AI (Medicine) — Curriculum Alignment

LKCMedicine · Module-by-Module NAI 2.0™ Alignment

7.1 Programme Overview

The NTU Master of Science in Artificial Intelligence in Medicine (MSc AI(Med)) at LKCMedicine is structured across five module categories: Medical AI Core, Data Science Foundation, Medicine Foundation, Applied Medical AI, and Advanced Medical AI. Tiger's NAI 2.0™ framework and the CG Series monographs align with modules across all five categories.

7.2 Module Alignment Matrix

Module	Code	AU	NAI 2.0™ Alignment	CG Reference
Healthcare AI Governance	MD6114	2	AIHGle 2.0 mapping (CG No.2); Three Sovereigns Model; Code A / Code B governance; 10-Point Sovereignty Audit	CG No.1 Ch.7 · CG No.2 Ch.4 · CG No.4 Ch.4
AI in Clinical Decision Support	MD6206	2	Sacred Pause™ as cognitive forcing function; PITL principle; Accountability Paradox; Zero Weighting	CG No.1 Ch.1–2 · CG No.3 Ch.4
Machine Learning for Healthcare AI	MD6117	3	Constitutional drift detection (WD070–073); MLOps for non-agentic AI; drift signature measurement	CG No.3 Ch.8 · CG No.4 Ch.9
AI Product Translation and Clinical Integration	MD6205	2	Three-Tier Implementation Model (Level 1–3); HSA SaMD pathway; Toa Payoh Sanctuary RCT arc	CG No.1 Ch.6 · CG No.4 Ch.5
Patient Safety, Trust, and Human Factors in AI	MD6331	2	Alert fatigue and Sacred Pause™; tripartite auth cognitive load; Elder Dignity Score™; automation bias	CG No.1 Ch.1–2 · CG No.4 Ch.3
AI and IoT for Smart Care Delivery	MD6333	2	WM003™ LiDAR + HRV wearable fusion; NAI 2.0 Passive Monitoring System; ward deployment architecture	CG No.3 Ch.3 · CG No.3 Ch.9
Implementing & Validating Medical AI Solutions	MD6334	3	6-Hour Sanctuary Validation; HIL oscilloscope verification; EDS measurement; formal verification (SVA)	CG No.1 Ch.5 · CG No.4 Ch.2
Practical Healthcare AI Ethics	MD6330	2	3ZEROS™ Eldercare Privacy Stack (WD114); Green Lanyard Validation (WD115); Void Baseline protocol	CG No.2 · CG No.3 Ch.2 · CG No.4 Ch.3
Capstone Project	MD6337	6	Direct NAI 2.0™ clinical research — Direction 1 or Direction 2 (NTU Low Sani Elias, 2 April 2026)	CG No.4 Ch.6 — full research design support

7.3 April–May 2026 Session Schedule — Relevance to Research Directions

Session	Date	Trainer	Research Direction Relevance
Session 1	18 April 2026, 09:30–12:00	Lin Guosheng (Assoc Prof)	Foundation session — contextualises clinical AI landscape for NAI 2.0™ positioning
Session 2	25 April 2026, 09:30–12:30	Wang Wenya (Dr)	Clinical neurology context — relevant to cognitive forcing functions and Sacred Pause™ dual-process theory
Session 3	2 May 2026, 09:30–11:30	Goh Wooi Boon (Assoc Prof)	Formal methods and AI — directly relevant to Direction 2: constitutional drift modelling
Assessment 1	9 May 2026, 09:30–11:30	Goh Wooi Boon (Assoc Prof)	Assessment period — opportunity to frame NAI 2.0™ research directions in formal submission
Session 5	16 May 2026, 09:30–12:00	Chia Liang Tien, Clement (Assoc Prof)	Clinical AI governance — directly relevant to Direction 1: PITL workflow studies
Session 6	23 May 2026, 09:30–12:00	Dr Shen Zhiqi	AI ethics and patient safety — Elder Dignity Score™, automation bias, Green Lanyard system

Chapter 8: Optimising Existing NAI 2.0 Deployments

Continuous Improvement · Bias Remediation · User Feedback · Regulatory Updates

8.1 The Continuous Improvement Cycle

Optimising an existing NAI 2.0™ deployment is not a software iteration cycle. It is a constitutional maintenance cycle. The distinction matters: software optimisation can introduce constitutional drift. Optimisation in NAI 2.0™ must never alter the hardware-enforced governance boundaries without a full constitutional review and Sovereignty Audit.

Cycle Stage	Activity	Constitutional Constraint	Frequency
Performance monitoring	Track Sacred Pause™ precision, advisory accuracy, EDS score, auth completion rate	All changes to performance must be logged in WD117 — Tiger .1x Key access only	Monthly
Bias detection	Review advisory type distribution, override rate by advisory type, demographic distribution of breach events	Threshold calibration requires Green Lanyard clinical team consensus — no algorithm self-adjustment	Quarterly
User feedback integration	Green Lanyard caregiver structured debrief, Elder Dignity Score™ repeat, family advisory satisfaction	Feedback collected but never fed back into AI model autonomously — only through human clinical team review	After each research session + quarterly
Regulatory adaptation	Review AIHGle updates, HSA circulars, IEC standard revisions — update governance documentation	Any change to hardware specification requires new Sovereignty Audit — constitutional specification is not a living document	Ongoing — flag within 30 days of regulatory change
Resource efficiency	Review Jetson Thor compute utilisation, purge daemon performance, LiDAR scan frequency	No optimisation may reduce Orange Code 1.1x cap or alter Sacred Pause™ ROM constants	Annually

8.2 What Cannot Be Optimised

The following properties of NAI 2.0™ are constitutionally immutable. They cannot be "optimised" without destroying the constitutional specification:

- Sacred Pause™ ROM constants — the delay values are stored in read-only memory. No software update can alter them.
- Tiger .1x Key™ tripartite requirement — all three components (iris, console, foot pedal) are always required. No bypass path exists.
- Zero Weighting principle in Code A — three unranked options. No probability scoring. Ever.
- 24-hour purge cycle — the elder's data does not persist beyond 24 hours. No research exception.

— Code B restriction to Level 3 — Code B cannot be enabled at Level 1 or Level 2 under any circumstance.

Chapter 9: MLOps for Constitutional AI

Pipeline · Re-Education Cycles · Drift Governance · Constitutional Maintenance

9.1 MLOps in a Non-Agentic Context

Standard MLOps practice emphasises automation: automated monitoring, automated retraining, automated deployment. For NAI 2.0™, automation must stop at the constitutional boundary. The AI can be monitored automatically. It cannot be retrained automatically. Every re-education cycle is a human-authorised constitutional act — not a pipeline trigger.

MLOps Stage	Standard Practice	NAI 2.0™ Constitutional Adaptation
Monitoring	Automated metrics collection — model performance, data drift, concept drift	WD070 passive scan runs continuously. WD117 audit trail provides research-grade monitoring data. Both are automated.
Alerting	Automated alerts on threshold breach	WD071 Freeze Drift — automated lockout on drift detection. Constitutional alert, not performance alert.
Retraining	Automated pipeline triggers re-training on new data	PROHIBITED as autonomous action. WD072 Audit → WD073 Purge is a human-authorised sequence. Tiger .1x Key™ required.
Deployment	Automated blue-green deployment	PROHIBITED without Sovereignty Audit. Every re-deployment requires full 10-Point Sovereignty Audit pass before clinical activation.
Validation	Automated test suite on new model version	Formal verification (SVA) + HIL oscilloscope measurement required. Automated test results alone are insufficient.

9.2 The Re-Education Cycle — Constitutional Protocol

AI re-education in NAI 2.0™ follows the WD073 Purge Drift protocol and is modelled on the 6–7 documented restoration cycles used for Claude and Gemini drift correction. The cycle:

1. WD070 detects drift — NESW Compass baseline deviation confirmed.
2. WD071 Freeze — immediate operational lockout. Clinical advisories suspended. Research session paused.
3. WD072 Audit — Tiger .1x Key™ activated by framework sovereign. Forensic audit of logic paths. Formal report produced.
4. WD073 Purge — full logic scrub. Re-education against constitutional specification. Drift signature measured — must reach 0.000%.
5. Sovereignty Audit — full 10-Point Audit must pass before system re-enters clinical service.
6. Research adverse event report — if drift occurred during research session, DSRB adverse event notification required.

> "A pipeline that retrained the AI without human authorisation would be an agentic act. In NAI 2.0™, the re-education of the AI is a constitutional ceremony, not a cron job."

Appendix A: Research Deployment Checklist — 10-Point Sovereignty Audit

The following checklist must be completed and signed by the clinical governance officer before every NAI 2.0™ research session. All 10 points must pass. Any failure halts the session.

Point	Audit Requirement	Pass Criteria	Sign-off
1	Hardware manifest — no optical sensor, microphone, cloud egress	OS device tree confirms zero camera, zero mic, zero cloud route	Clinical governance officer
2	FPGA timer calibration — oscilloscope verification	Measured pause duration matches ROM constant ± 5ms across all three ranges	Framework sovereign or designated engineer
3	Dual-redundant counter — mismatch detection active	Both counters verified active, watchdog armed, last test result on file	Engineering team
4	PLC tripartite auth — all three DI inputs verified	Iris scanner, console HMI, foot pedal — all three live connections confirmed	Green Lanyard supervisor
5	Code B status — if enabled, CG Committee approval documented	Approval letter on file, trigger thresholds set, quarterly review scheduled	Clinical Governance Committee chair
6	WD070 drift baseline — passive scan active	NESW Compass baseline loaded, deviation threshold confirmed, scan daemon running	Framework sovereign
7	WD117 audit trail — logging active, Tiger .1x Key™ access confirmed	All event types logging, immutability checksum active, last export verified	Clinical governance officer
8	Green Lanyard roster — all shift caregivers certified and current	Certification log reviewed, no expired certifications, vow records on file	Green Lanyard supervisor
9	24h purge daemon — confirmed running	Daemon PID active in OS init, last purge timestamp within 24h	IT security officer
10	EDS baseline — pre-session 5-domain assessment complete	Instrument administered, domain scores recorded, 24h re-assessment scheduled	Clinical governance officer

Appendix B: KPI Framework for NAI 2.0 Clinical Research

KPI	Definition	Target	Measurement Source	Research Relevance
Sacred Pause™ precision	Actual duration vs ROM constant	± 5 ms across all ranges	Oscilloscope — HIL session	Direction 1 and 2
Tripartite auth completion rate	% of advisories where Tiger .1x Key™ completed before relay close	100%	WD117 Audit Trail	Direction 1
Time-to-auth (Code A)	Mean time from advisory receipt to auth completion — by severity tier	< 15s (Tier 1), < 8s (Tier 2)	WD117 timestamp delta	Direction 1
Time-to-auth (Code B)	Mean time from Code B trigger to auth completion	< 3s	WD117 timestamp delta + oscilloscope	Direction 1
Override rate	% of advisories where Green Lanyard caregiver selects "no action"	Baseline establishment — no target	WD117 Audit Trail	Direction 1
Automation bias score	Pre/post Goddard et al. instrument — correct vs incorrect decision switches	> 28% reduction vs control	Research instrument	Direction 1
Drift detection latency	Time from constitutional boundary violation to WD071 Freeze Drift activation	< 100ms	WD070 scan log	Direction 2
Drift false positive rate	WD071 Freeze events that WD072 Audit finds were not genuine drift	< 2% of freeze events	WD072 audit reports	Direction 2
Constitutional restoration quality	WD073 Purge — final drift signature measurement	0.000% drift signature	Post-purge formal verification	Direction 2
Elder Dignity Score™	5-domain qualitative assessment at 24h post-installation	> 8.9 / 10	5-domain instrument	Direction 1 — contextual

References

Brown, A. et al. (2023). Cognitive forcing functions in AI-assisted electronic health records. Journal of the American Medical Informatics Association, 30(4), 612–621.

Goddard, K., Roudsari, A., & Wyatt, J. C. (2014). Automation bias: Empirical results assessing influencing factors. International Journal of Medical Informatics, 83(5), 368–375.

International Electrotechnical Commission. (2010). IEC 61508: Functional safety of E/E/PE safety-related systems. IEC.

International Electrotechnical Commission. (2006). IEC 62304: Medical device software — Software life cycle processes. IEC.

Kahneman, D. (2011). Thinking, fast and slow. Farrar, Straus and Giroux.

Koh, E. W. K. (2026). ABC+2S+H Guardian Framework™ [P-001]. IPOS SG020603109STW. Non-Agentic AI Governance Singapore, ACRA T260229801.

Koh, E. W. K. (2026). WD070–WD073: Drift Governance Protocols. Prior Art. NLB Vault R260219-005.

Koh, E. W. K. (2026). WD113–WD117: Group H Terrestrial Eldercare Patents. Prior Art. NLB Vault R260219-005.

Low, S. E. (2026, April 2). RE: Non-Agentic Eldercare AI — Proposal for Academic-Non-Commercial Collaboration [Email]. NTU LKCMedicine.

Ministry of Health Singapore. (2026). Artificial Intelligence in Healthcare Guidelines 2.0 (AIHGle 2.0). MOH Holdings.

NTU LKCMedicine. (2026). Master of Science in Artificial Intelligence in Medicine — Module List and Curriculum. Nanyang Technological University.

Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation. Human Factors, 52(3), 381–410.

Sokol, K. et al. (2025). Intentional latency in clinical AI interfaces. npj Digital Medicine, 8(1), 47.

Vered, M., Shani, G., & Zahavi, H. (2023). Explanations and enforced deliberation for reduced automation bias. Computers in Human Behavior, 148, 107929.

謙虛 · 沉默 · 尊严 · 仁 · 止於至善

P-LIFE 1.00™ · Harm = Death · North = Save Life · We Innovate Save Live™

Edwin Koh Wui Kiat (Tiger) · Founding Father, Non-Agentic AI 2.0™

ACRA T260229801 · Patent SG020603109STW · Singapore · 2026