

DeepSensi publishes the first quantitative safety standard for clinical AI - free to the entire industry

Dover, Delaware - July 2026. DeepSensi PBC today released the **Gomola Framework**, the first quantitative safety-certification standard for clinical artificial intelligence, together with six technical papers headed to the medRxiv preprint server and peer review at *NEJM AI*. The framework is open and royalty-free.

Clinical AI is the only safety-critical technology deployed at the bedside without a quantitative reliability requirement. Aviation certifies against DO-178C; nuclear engineering against IEC 61508; medicines against Phase III trials. Generative clinical AI - whose dominant failure mode, hallucination, is fluent and clinically plausible - has had no equivalent. The Gomola Framework defines four graded certification levels by explicit per-assertion hallucination-probability bounds, organized around five independently auditable pillars.

Key results, from the published papers:

- **A quantified safety bound.** Fault tree analysis (the method used to certify power plants and avionics) of DeepSensi's 23-barrier verification architecture yields a worst-case bound of 3.23×10^{-6} undetected hallucinations per assertion - meeting the numerical target of safety-integrity level SIL-4 with an ≈ 31 -fold margin.
- **Elite diagnostic accuracy with zero missed criticals.** On 301 NEJM Clinicopathological Conference cases (2014–2023), a full-architecture reference implementation reached 86.0% top-1 and 93.7% top-3 accuracy after safety-first calibration, with the **missed-critical rate reduced from 4.7% to 0.0%** at a cost of just 2.3 *additional* seconds of median deliberation (total median: 14.3 s; a full multi-specialist deliberation completes in 12 seconds to minutes for complex cases - a panel that takes days to convene by hand).
- **Vendor independence.** DeepSensi is not a wrapper on any single model; it cross-examines multiple models from independent vendors, so a better base model yields a better verified answer - the architecture consumes the progress of the model race rather than competing in it.
- **A structured "I don't know."** The LIMBO protocol produces an explicit uncertainty declaration - the disagreement, a working hypothesis, and the resolving tests - instead of a confident guess.

A full cognitive infrastructure. The newly published papers describe a system that extends beyond verification: the **Hyper Consilium** integrates physicians as mathematically scored cognitive partners (with royalties for verified knowledge, released only after patient outcomes confirm value); **AutoResearcher** autonomously generates and adversarially validates research hypotheses with a formal Bayesian convergence guarantee; and **Golden Horizon** provides borderless, free-by-architecture clinical-trial matching and autonomous patient advocacy for those who have exhausted standard-of-care.

A challenge to the industry's benchmarks. A companion paper, "The Flawed Yardstick," documents how popular static benchmarks systematically penalize safe clinical behavior - the "Safe Triage Paradox" - and proposes a dynamic, safety-first evaluation protocol.

The system is operational on edge nodes across four continents, including fully offline and air-gapped deployments. DeepSensi is engaging the U.S. FDA through the Q-Submission pathway and is architected for the EU AI Act ahead of its December 2027 deadline. DeepSensi PBC is a

Public Benefit Corporation whose charter commits it to free global trial matching, open research into autism and refractory epilepsy, and the open safety standard.

"A safety standard behind a paywall is not a standard. It is a product," said founder and chief architect Tomasz Jan Gomola. "We publish the proof, the standard, and the audit. We invite the world to try to refute them."

Technical documentation: www.deepsensi.com/papers

Press and auditor access to research protocols: press@deepsensi.com

About DeepSensi - DeepSensi PBC is a Public Benefit Corporation building verified clinical intelligence: a Cognitive Medical OS that surrounds any language model with an auditable verification architecture. The Gomola Framework and DeepSensi Standard are open and royalty-free; the reference implementation is proprietary. Founder and chief architect: Tomasz Jan Gomola (ORCID 0009-0001-5222-6154).