

The First Physicians Who Earn Royalties for Being Right

Inside DeepSensi's "Medical Exosuit" for doctors: a consilium that argues back, documentation that writes itself, referrals with built-in anti-kickback compliance - and a knowledge economy that pays only for verified outcomes

Feature · DeepSensi PBC, Dover, Delaware · July 2026 · press@deepsensi.com

Ask a physician what broke medicine for them and you will rarely hear "the patients." You will hear the 11 p.m. documentation backlog, the fear of the one interaction they didn't catch, the knowledge that their hardest-won clinical insights evaporate into hallway conversations while their student-loan balance does not. DeepSensi - the verified clinical-intelligence system whose safety mathematics and diagnostic results are now published in a series of technical papers - was designed by its architect, Tomasz Jan Gomola, around a specific conviction: the way to make medical AI trustworthy is to make the physician more powerful, more protected, and, for the first time, *fairly paid for what they know*. The company calls the resulting toolkit, only half-jokingly, the physician's Medical Exosuit.

A colleague, not an oracle

At the exosuit's core is the Hyper Consilium, described in the newly published WP-003: up to 37 specialist AI agents deliberating in adversarial triads - and, crucially, the physician inside the deliberation rather than after it. Where conventional "human-in-the-loop" AI hands the doctor a finished answer to approve, DeepSensi treats the physician as a scored cognitive node who can inject context, challenge any specialist, or halt the pipeline at any stage. Two design guarantees define the relationship. The physician's voice has a mathematical floor - it can never be weighted to zero, no matter what. And the system's own reputation weights are calibrated retroactively against real patient outcomes, so a doctor whose judgment repeatedly beats the swarm sees their influence grow, permanently and provably.

When the evidence is genuinely insufficient, the system does not bluff. Its LIMBO protocol declares structured uncertainty - and a companion capability called active sensing goes one step further, computing which single test would most decisively resolve the ambiguity. The most dangerous sentence in medicine is a confident guess; DeepSensi's answer is an honest question with a plan attached.

The paperwork disappears

Around the diagnostic core, the exosuit automates the drudgery that consumes up to half a clinician's day. An autonomous scribe listens to the encounter and drafts the structured note with diagnostic codes; the physician reviews, signs, and the record writes back to the hospital's own EHR through standard interfaces. A pharmacy concierge checks every prescription against interaction databases before it can be sent - a hard stop, not a suggestion, when a dangerous combination appears. Patient instructions are automatically checked for plain-language readability before they go out. Even referrals are re-engineered: when one physician sends a patient to another, the fee sits in escrow and is released only when the patient is actually seen - with the amount automatically screened against fair-market-value rules, building anti-kickback compliance into the plumbing itself rather than leaving it to memory and audit season.

Reputation you can take anywhere

Beneath the tools sits the layer Gomola considers the real invention: an economy of verified clinical knowledge. Every physician on the network carries a Global Clinical Performance Score - a five-axis, manipulation-resistant measure of diagnostic accuracy, error-detection skill, knowledge contribution, integrity, and resistance to cognitive bias, published as an open framework in its own paper. The score decays over time, so reputation must be re-earned rather than hoarded; it rewards the safe admission of uncertainty rather than punishing it; and it is accumulated through cryptographically auditable events rather than committee opinion. A brilliant diagnostician in a rural clinic becomes, for the first time, *visible* - on the same objective scale as a professor in a university hospital.

And knowledge itself becomes an asset. A physician who contributes a validated clinical insight - a diagnostic heuristic, a rare-disease pattern, an interaction observation - earns royalties whenever that insight helps resolve a case anywhere in the network. The twist that makes the economics honest: royalties are held in escrow and released only after follow-up confirms the contribution actually improved the outcome. Volume earns nothing; verified impact earns forever. It is, to the company's knowledge, the first formal compensation model for clinical expertise inside an AI system - and it inverts the industry's usual arrangement, in which physician knowledge silently trains someone else's model for free.

Why the exosuit matters beyond the doctor

Hospitals get a workforce that documents faster, prescribes more safely, and refers compliantly - with every action sealed in a court-grade audit trail. Insurers get quantified risk instead of blanket AI exclusions. Patients get a physician who is less exhausted, harder to bias, and backed by a consilium that never sleeps. And the profession gets something it has quietly lacked: a meritocracy with receipts. The era of medical AI that extracts from physicians is ending, DeepSensi argues; the era of AI that armors them - cognitively, legally, and economically - has a working prototype, and its mathematics are published for anyone to attack.

Technical documentation: www.deepsensi.com/papers (WP-003 The Hyper Consilium; WP-004 GCPS) · Physician demos: www.deepsensi.com/for-physicians · Press: press@deepsensi.com