From Signal Noise to Signal Intelligence

Intent-Driven Agentic AI for Pharmacovigilance Signal Management

A Multi-Agent System Built on PACT-CARE™ and A5 RAZOR Frameworks

Sudhir Shandilya
Creator, PACT-CARE™ and A5 RAZOR Frameworks

Abstract

Pharmacovigilance signal detection faces a fundamental challenge: traditional systems execute fixed analytical pipelines regardless of context, generating volumes of flagged signals that overwhelm safety scientists while potentially missing clinically significant patterns. This paper presents an intent-driven agentic AI architecture that fundamentally transforms how signal management systems operate—not by automating existing workflows, but by replacing rigid business logic with flexible intent interpretation.

Built on two foundational frameworks—PACT-CARE™ for responsible AI adoption and A5 RAZOR for workflow decomposition—the system interprets natural language queries to understand user intent, dynamically orchestrates specialized AI agents, and adapts its analytical approach based on intermediate findings. Architecture integrates six distinct AI/ML paradigms under a human-in-the-loop design that maintains expert authority over all regulatory determinations.

This paper includes a complete walkthrough of a working demonstration system deployed on Hugging Face, showing the end-to-end flow from natural language query through agentic investigation to expert determination and audit trail generation. The system is designed to be equally understandable by technical architects, domain experts, innovators, and regulatory professionals—bridging the gap between those who see AI through code alone and those who experience it only through consumer interfaces.

Keywords: *Pharmacovigilance, Agentic AI, Signal Detection, Human-in-the-Loop, PACT-CARE™, A5 RAZOR, Intent-Driven Architecture, Multi-Agent Systems, Explainable AI, Regulatory Compliance*

1. Introduction

1.1 The Problem with Process Automation

Most AI initiatives automate existing workflows—which means automating your inefficiencies too. Traditional signal management systems operate as fixed analytical pipelines: they execute predetermined sequences of calculations regardless of context, clinical relevance, or the specific question being asked. A query about a potential hepatotoxicity signal triggers the same computational workflow as an inquiry about expedited reporting requirements.

Intent-driven design flips this entirely. Instead of encoding rigid steps, the workflow forms itself based on what you're trying to achieve and what the data reveals. You define *what* you need; the system figures out *how*—while staying fully compliant with every regulation needed.

1.2 Bridging the Gap

This work emerged from personally navigating the tension between two ways of seeing AI—one grounded in technical structure and formal definitions, and the other rooted in deep domain understanding and lived experience. Both perspectives are valuable, but neither alone is sufficient. **Real innovation and progress happen in the space between them, where assumptions and biases are questioned and reconciled.** The internal friction along that path is formative and helps shape what ultimately gets built.

This paper—and the frameworks it presents—aims to be equally understandable by technical architects who need orchestration patterns, domain experts who understand pharmacovigilance but have never built an agentic system, innovators who need a practical

path from concept to implementation, and regulatory professionals who must ensure compliance.

1.3 Foundational Frameworks

The system is built on two complementary frameworks:

- **PACT-CARE™** — An 8-step human-in-the-loop framework for responsible AI adoption: Patient & Problem, Action Policy, Capacity & Context, Thresholds & Trade-offs, Compliance & Regulation, Adoption & UX, Reliability & Recalibration, Equity/Evidence/Economics
- **A5 RAZOR** — A workflow decomposition methodology: Analyze, Architect, Assemble, Align, Activate

Together, PACT-CARE™ ensures the system is *responsible and useful*; A5 RAZOR ensures it is *implementable and maintainable*. Both frameworks are domain-agnostic—the architecture transfers across R&D, regulatory, operations, and finance.

2. Application Walkthrough: End-to-End Flow

The following walkthrough demonstrates the complete signal management workflow using a working application deployed on Hugging Face. The application follows a three-phase structure: **Context → AI Analysis → Human-in-the-Loop**, with a comprehensive audit trail for regulatory compliance.

2.1 Phase 1: Context & System Overview

The application opens with the About tab, establishing context for what the system does and how it differs from traditional approaches.

Understanding Intent-Driven Design

The landing screen immediately addresses the core question: "What Does Intent-Driven Mean?" A side-by-side comparison shows how traditional systems follow fixed sequences (Step 1 → Step 2 → Step 3 → Step 4) with no adaptation, while intent-driven systems declare what they want to achieve and let an intelligent orchestrator determine how to achieve it.

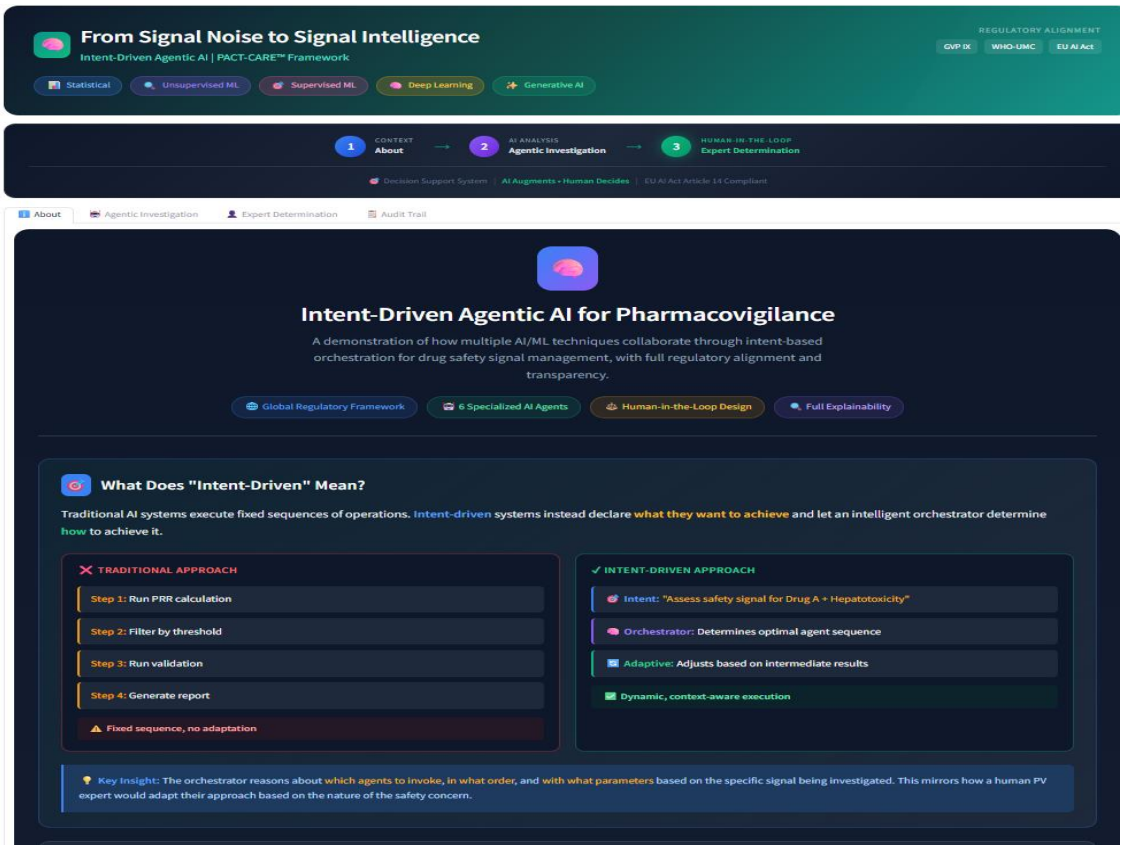

Figure 1: Application landing screen showing intent-driven vs. traditional approach comparison. The key insight: the orchestrator reasons about which agents to invoke, in what order, and with what parameters—mirroring how a human PV expert would adapt their approach.

Multi-Agent Architecture & AI/ML Layers

Scrolling down reveals the system's architecture: six specialized AI agents (Detection, Validation, Causality, Prioritization, Recommendation, Expert Review) collaborate under orchestrator control, each aligned to specific regulatory standards. The multi-layer AI/ML architecture shows six distinct paradigms working together: Statistical (PRR, ROR),

Unsupervised ML (Isolation Forest), Deep Learning (BioClinicalBERT), Supervised ML (XGBoost + SHAP), Predictive (Holt-Winters), and Generative AI (LLM Synthesis).

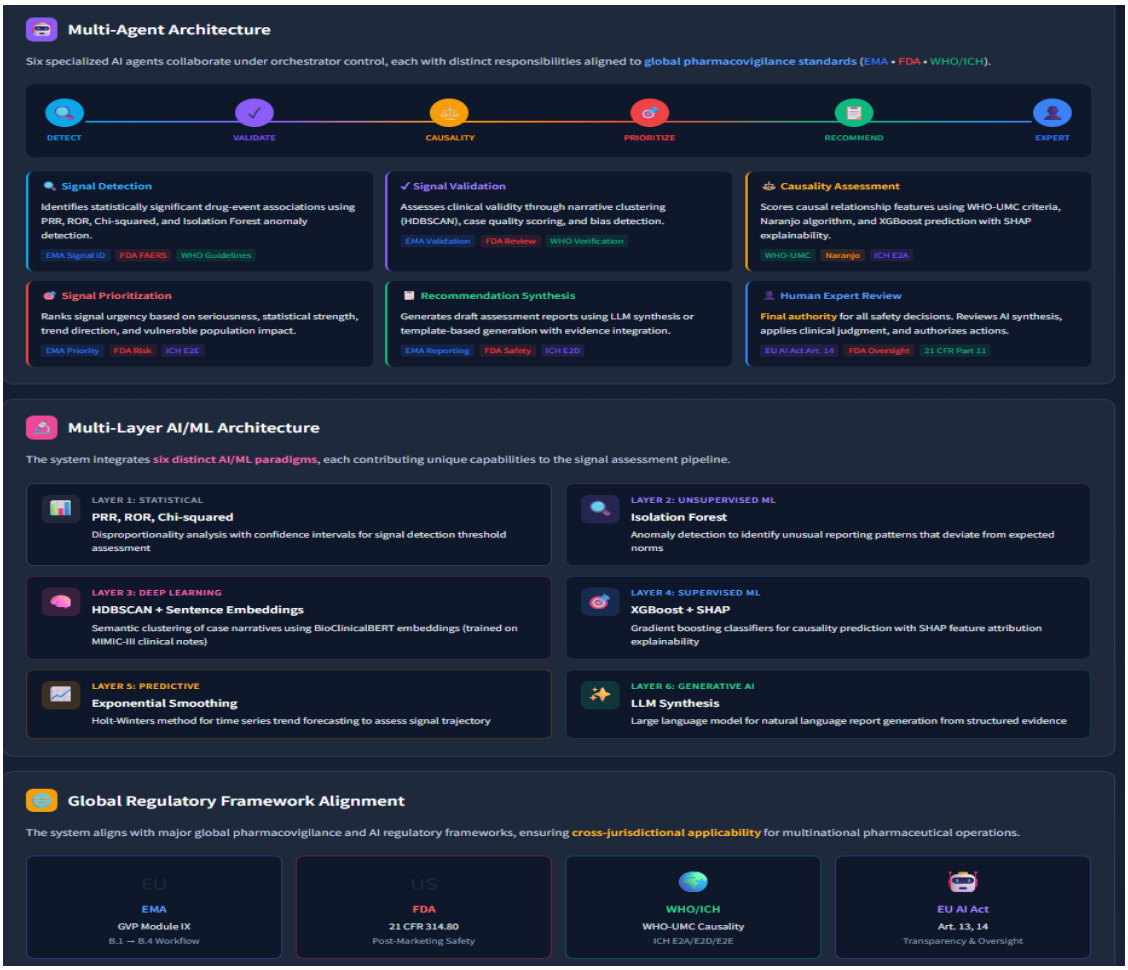

Figure 2: Multi-Agent Architecture showing six specialized agents with regulatory alignments, and Multi-Layer AI/ML Architecture showing six distinct paradigms contributing to the signal assessment pipeline.

Regulatory Alignment & Responsible AI Principles

The system explicitly maps each process step to regulatory requirements across jurisdictions (EMA GVP Module IX, FDA 21 CFR 314.80, WHO-UMC, EU AI Act). Six Responsible AI Principles guide the design: Human Oversight, Audit Trail, Explainability, Bias Detection, Intent-Driven orchestration, and Validated models.

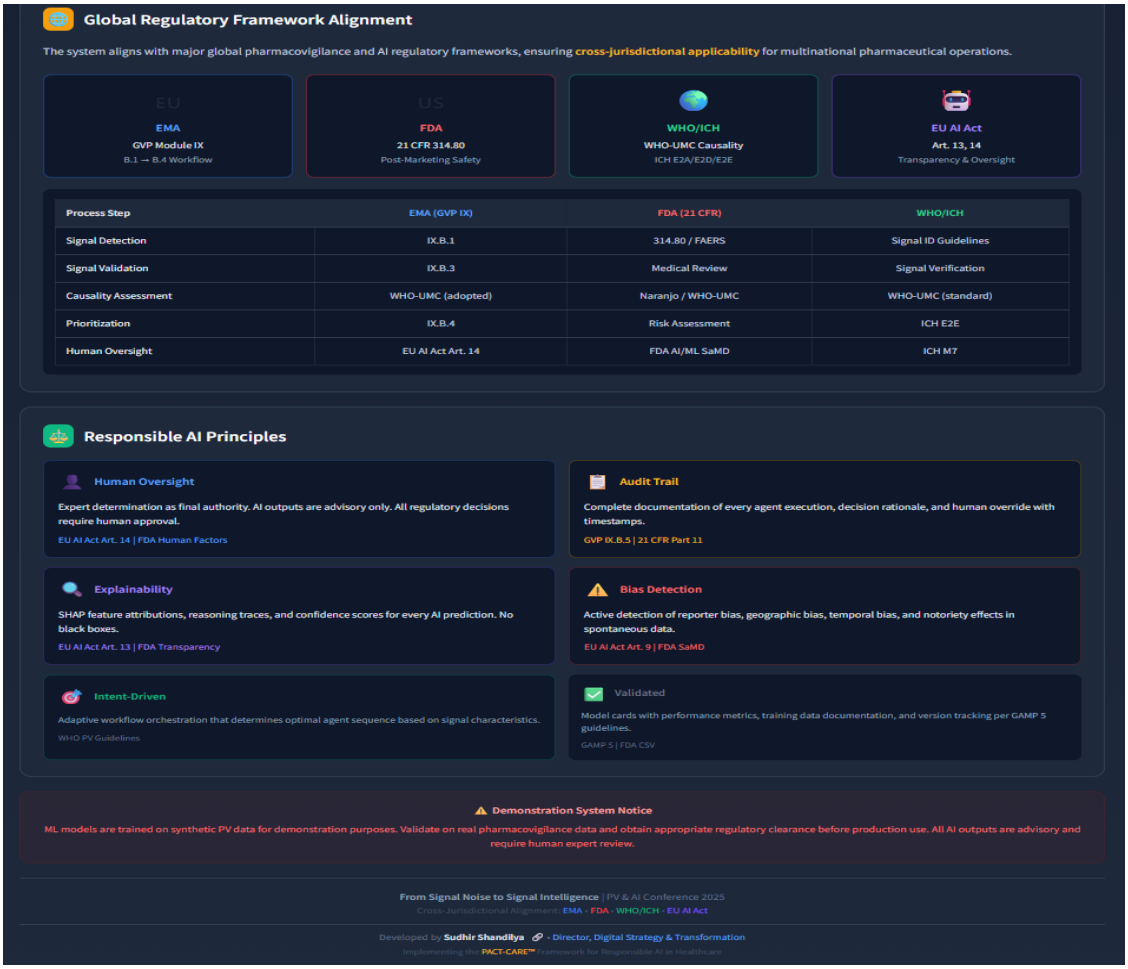

Figure 3: Global Regulatory Framework Alignment matrix showing process step to regulation mapping, and Responsible AI Principles with regulatory citations. Note the Demonstration System Notice at bottom.

2.2 Phase 2: Agentic Investigation

Moving to the Agentic Investigation tab, users can initiate signal investigations using natural language queries. This is where intent-driven design comes to life.

Starting an Investigation

The interface accepts natural language queries like "Investigate Drug A × Hepatic enzyme increased." Sample queries demonstrate different agentic capabilities: drug-event analysis, drug aggregation, regulatory reasoning, and error handling. Clicking "Start Agentic Investigation" initiates the workflow.

The Execution Status panel shows real-time progress as each step completes: Signal Detection (PRR Calculator, ROR Calculator, Chi-squared Test), Signal Validation (ML Classifier, HDBSCAN Clustering, BioClinicalBERT Embeddings), Causality Assessment (WHO-UMC Criteria, Naranjo Algorithm, XGBoost Causality Model), and Priority Assignment (Risk Scoring Engine, Regulatory Timeline Calculator).

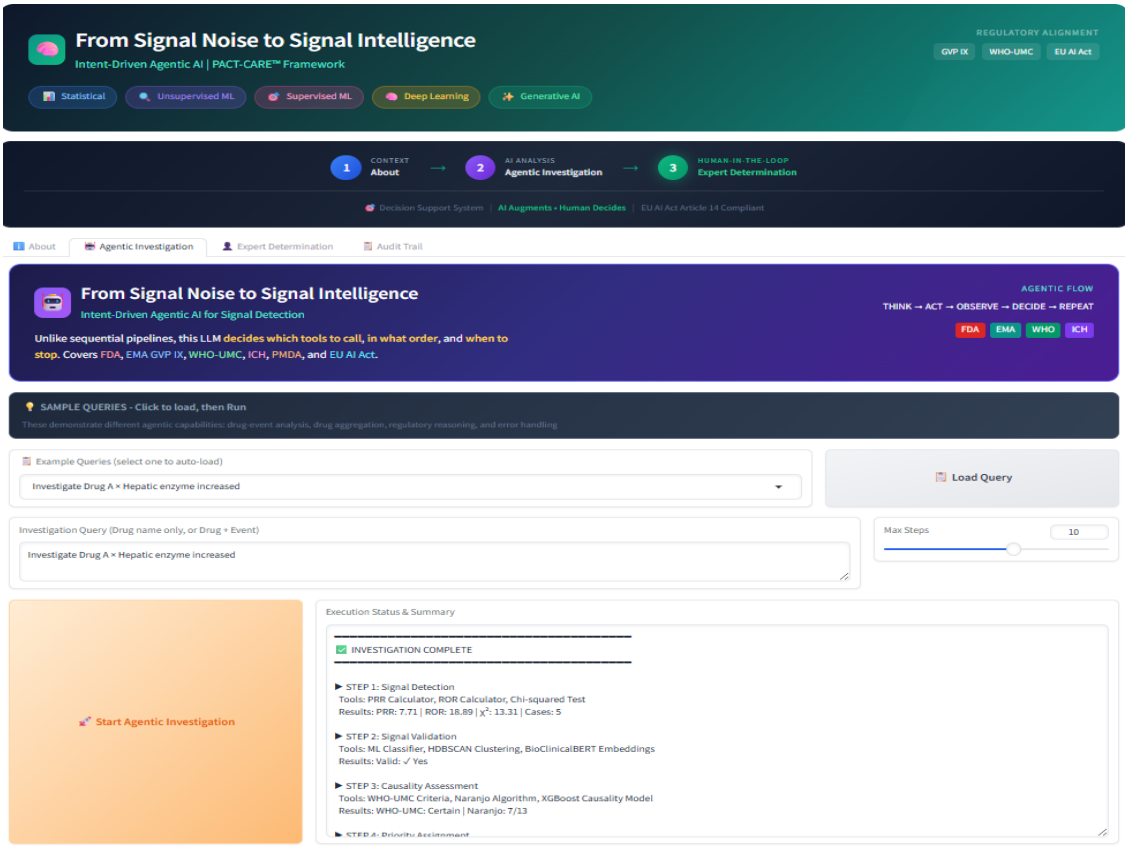

Figure 4: Agentic Investigation interface showing query input, sample queries dropdown, and Execution Status with real-time step completion. Results show PRR=7.71, ROR=18.89, WHO-UMC=Certain, Priority=MEDIUM (44/100).

Agent Reasoning Trace: Transparent Thinking

The Agent Reasoning Trace panel shows exactly how the AI thinks through each step. For Step 1 (Signal Detection), the system displays: regulatory basis (GVP IX.B.1, 21 CFR 314.80(c)), input evidence, decision rationale, alternatives considered and known limitations. This transparency is essential for regulatory compliance and building user trust.

The Signal Management Workflow tracker on the right shows progress through all five AI steps with regulatory mappings for each: Detection (FDA 21 CFR 314.80, EMA GVP IX.B.1),

Validation (FDA Medical Review, EMA GVP IX.B.3), Causality (WHO-UMC, Naranjo, ICH E2A), Prioritization (FDA Risk Assessment, EMA GVP IX.B.4), and Recommendation (FDA 21 CFR 314.80(c), EMA GVP IX.B.3.6).

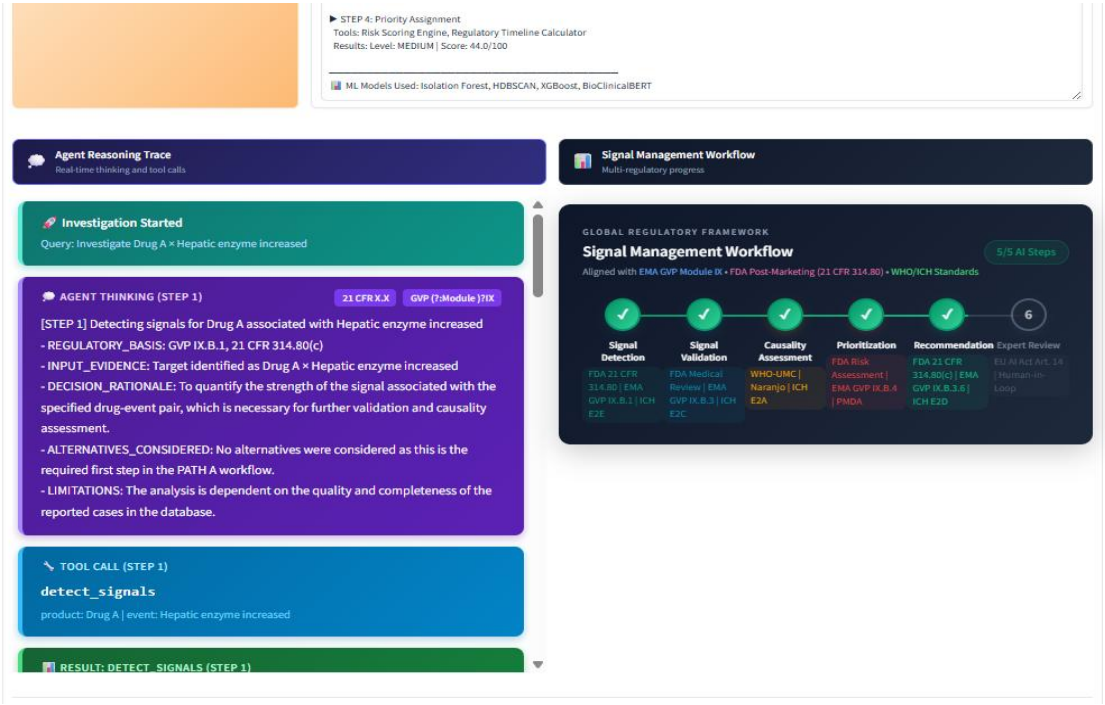

Figure 5: Agent Reasoning Trace showing Step 1 thinking with regulatory basis, decision rationale, tool calls, and limitations. The Signal Management Workflow tracker shows 5/5 AI Steps completed with regulatory citations.

Synthesis Step: Evidence Chain

The final agent thinking (Step 5) synthesizes evidence from all prior steps into regulatory recommendations. The evidence chain shows: signal detected with PRR=7.71 and case count=5, validated with probability=0.58, causality assessed as "Certain" with Naranjo score of 7, prioritized as MEDIUM urgency with score=44. The system notes confidence factors and explicitly states that human oversight is required per EU AI Act Article 14.

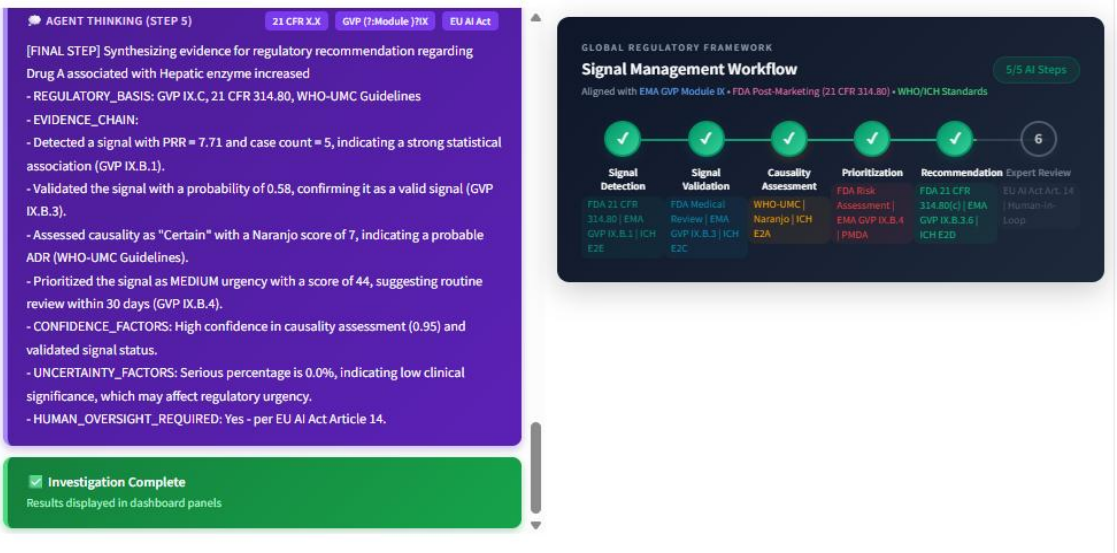

Figure 6: Final Step (Synthesis) showing evidence chain, confidence factors, and "Investigation Complete" notification. Human oversight requirement is explicitly stated.

Global Regulatory Compliance Dashboard

Upon investigation completion, the system displays a comprehensive compliance dashboard showing 100% alignment across all major regulatory frameworks: FDA (21 CFR 314.80, FAERS, Medical Review Standards, Expedited Safety Reporting), EMA (GVP Module IX B.1-B.5), WHO-UMC (Causality Criteria, Naranjo Algorithm, VigiBase Signal Standards), ICH (E2C/E2D/E2E), PMDA (J-ADE Reporting Standards), and EU AI Act (Articles 9, 11, 13, 14).

The Responsible AI panel confirms all eight principles are active: Human Oversight, Audit Trail, Multi-Agent design, Explainability, Risk-Based approach, Bias Detection, Intent-Driven orchestration, and Validated models—each with regulatory citations.

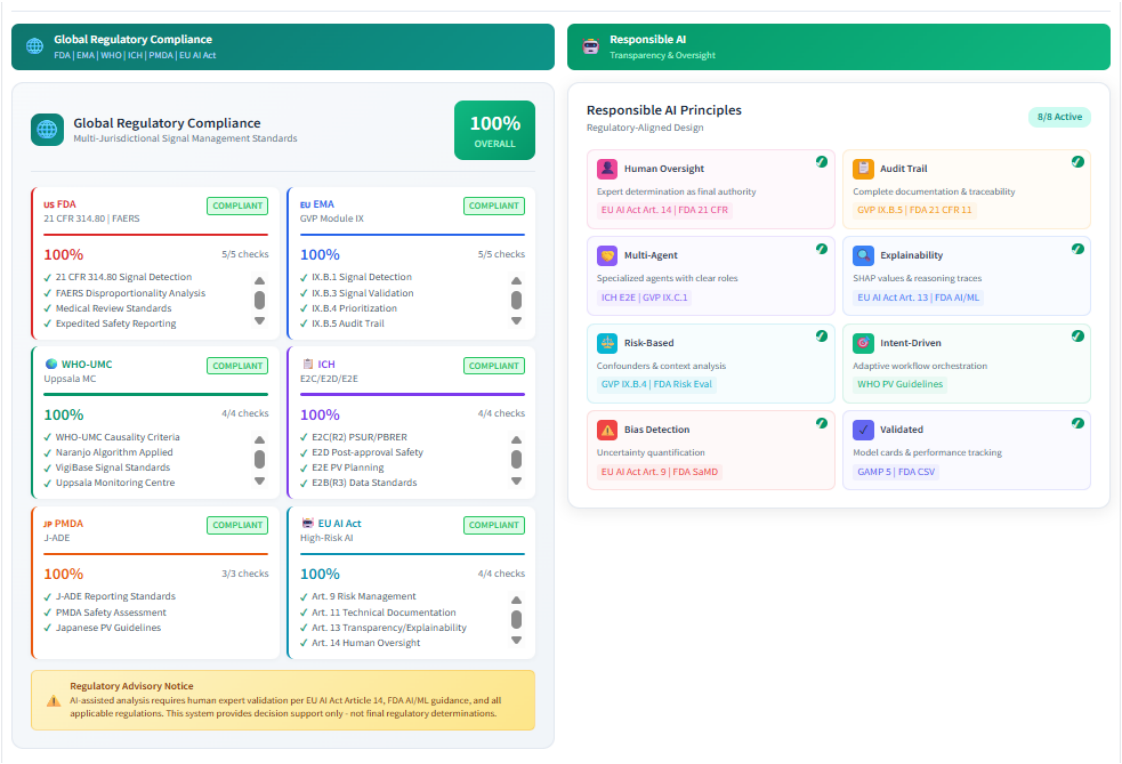

Figure 7: Global Regulatory Compliance dashboard showing 100% compliance across FDA, EMA, WHO-UMC, ICH, PMDA, and EU AI Act, with Responsible AI Principles (8/8 active) and Regulatory Advisory Notice.

Integrated Assessment Results

The results dashboard presents three interconnected views: Agent Communication (showing inter-agent message flow: DET→VAL, VAL→CAUS, CAUS→PRI), the Integrated Assessment with clinical significance and regulatory implications, and the Priority Score (44/MEDIUM with "Routine Safety Review" recommendation).

The Integrated Assessment provides key evidence citations, regulatory implications for FDA/EMA/WHO/PMDA, the regulatory recommendation ("Conduct a routine review within 30 days"), and explicit limitations ("AI recommendations are advisory only", "Clinical judgment must supersede AI suggestions", "This is NOT a final regulatory determination").

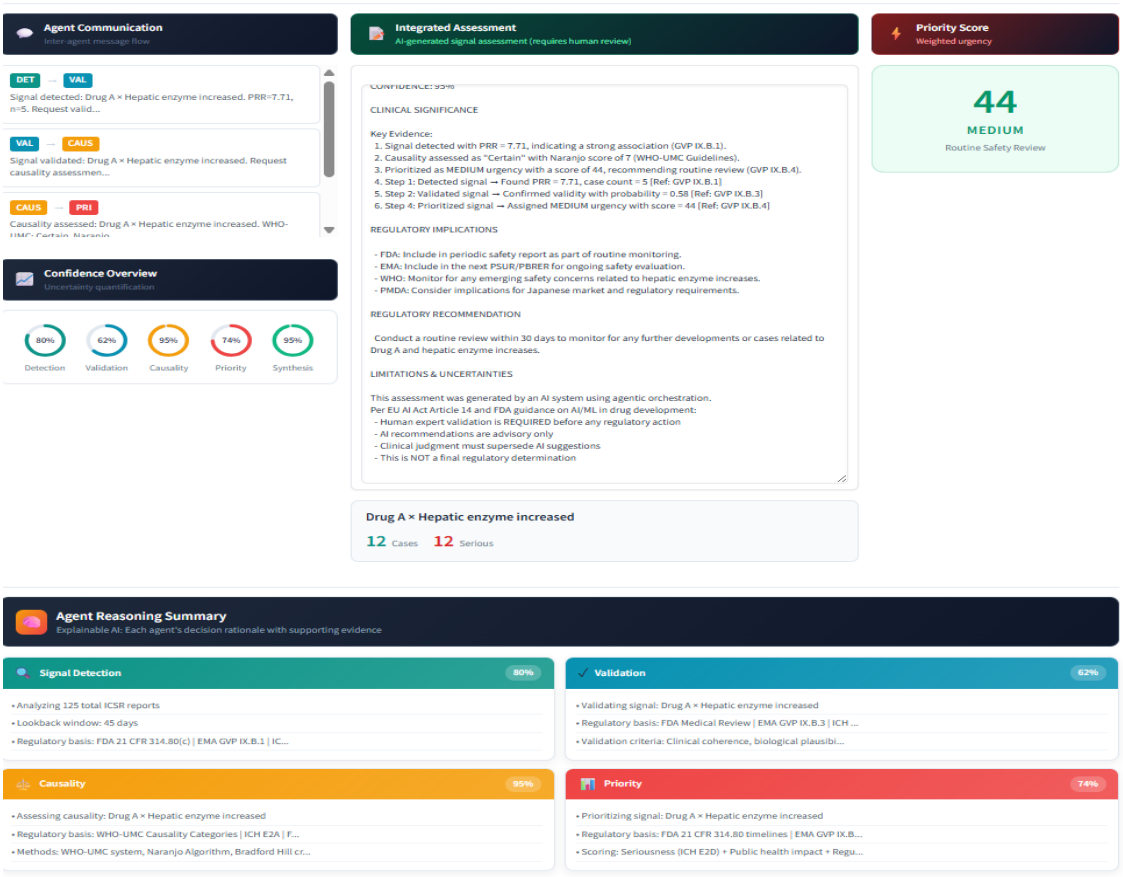

Figure 8: Results dashboard showing Agent Communication flow, Integrated Assessment with clinical significance and regulatory implications, Priority Score (44/MEDIUM), Confidence Overview, and Agent Reasoning Summary cards.

Causality Assessment with ML Explainability

The Causality Assessment panel provides detailed WHO-UMC categorization ("Certain") with Naranjo Score (7 = Probable ADR). Four causality criteria are evaluated: Temporal Relationship (CONSISTENT—median onset 30 days within typical window), Dechallenge (POSITIVE—11/12 cases with positive dechallenge), Rechallenge (POSITIVE—3/3 cases with definitive causal evidence), and Biological Plausibility (HIGH—known hepatotoxic potential via metabolic pathway).

The ML Causality Prediction panel shows the XGBoost model's 95% probability assessment with SHAP feature importance: rechallenge_outcome (30%), dechallenge_outcome (25%), temporal_relationship (20%), biological_plausibility (15%), alternative_explanations (10%). This transparency ensures no black boxes—every prediction is explainable.

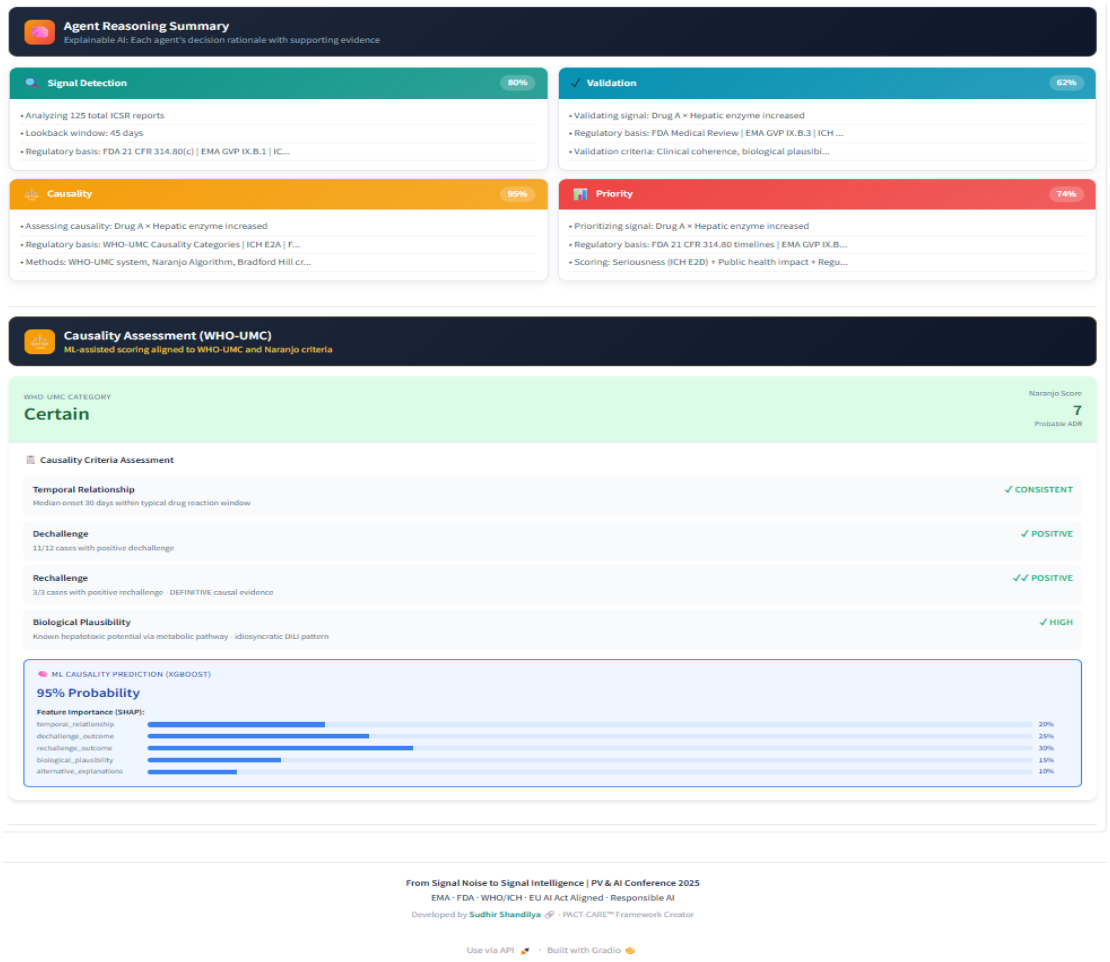

Figure 9: WHO-UMC Causality Assessment (Certain, Naranjo 7) with four criteria evaluated, and XGBoost ML Causality Prediction (95% probability) with SHAP feature importance for full explainability.

AI Recommendation Panel

The AI Recommendation panel consolidates all findings into an actionable summary. Key metrics are displayed prominently: PRR (7.71, 95% CI: 2.3-25.9), Cases (5, 0% serious), WHO-UMC (Certain causality), Naranjo (7/13 = Probable ADR), and AI Confidence (95% weighted average).

Causality Evidence details include: Median Time to Onset (29.5 days), Positive Dechallenge (92%), Positive Rechallenge (0 cases), and Biological Plausibility (Plausible). The Priority

Score Breakdown shows how the 44/100 score was calculated across five dimensions: Seriousness (0/30), Causality Strength (25/25), Statistical Evidence (14/20), Trend Pattern (5/15), and Vulnerable Populations (0/10).

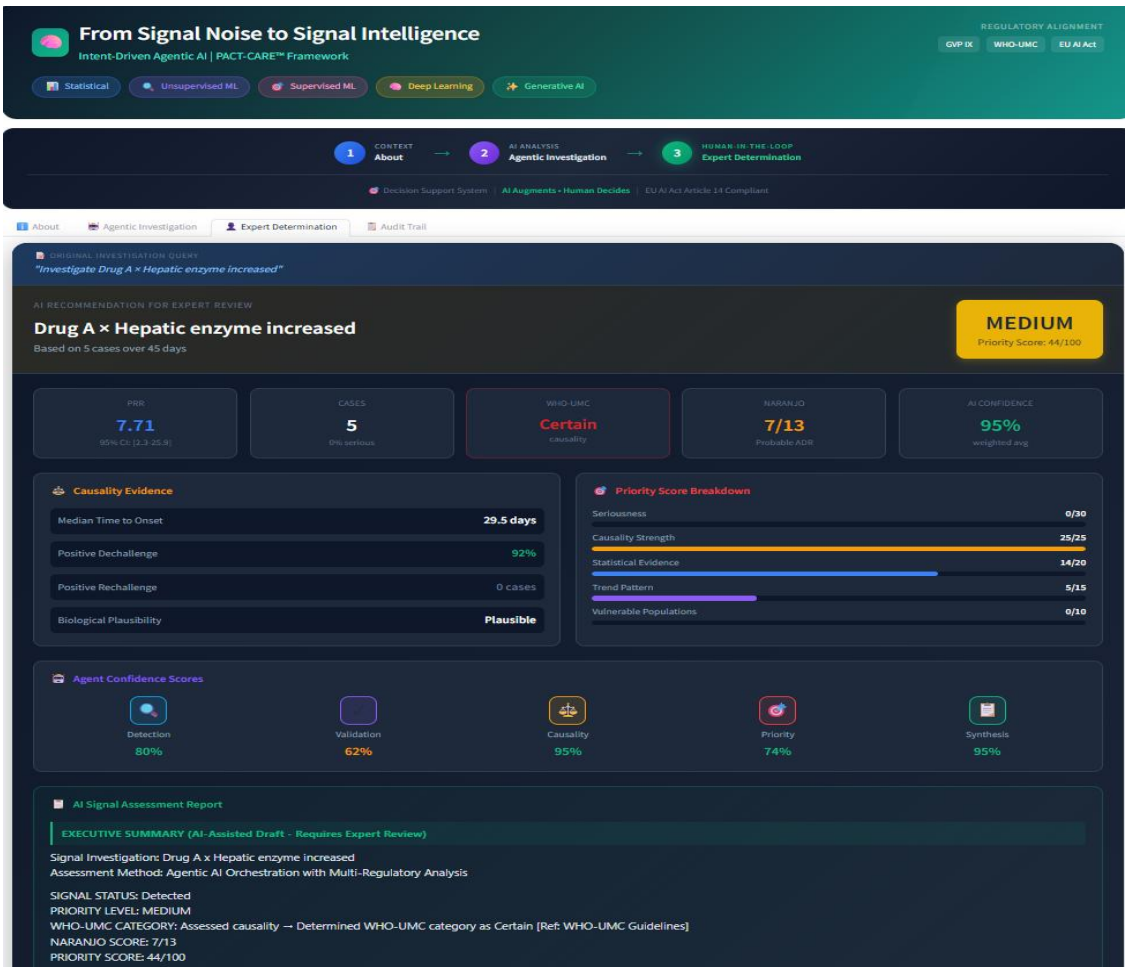

Figure 10: AI Recommendation panel showing key metrics (PRR, Cases, WHO-UMC, Naranjo, AI Confidence), Causality Evidence, Priority Score Breakdown by dimension, and Agent Confidence Scores (Detection 80%, Validation 62%, Causality 95%, Priority 74%, Synthesis 95%).

Executive Summary Report

The AI Signal Assessment Report provides a structured executive summary designed for regulatory review. Sections include: Signal Investigation identification, Assessment Method (Agentic AI Orchestration with Multi-Regulatory Analysis), Signal Status (Detected), Priority Level (MEDIUM), WHO-UMC Category (Certain), Naranjo Score (7/13), Priority Score (44/100), and Confidence (95%).

Clinical Significance lists six evidence points with regulatory references. Regulatory Implications provide specific actions for FDA, EMA, WHO, and PMDA. The Regulatory Recommendation states: "Conduct a routine review within 30 days to monitor any further developments." Limitations & Uncertainties explicitly state: "Human expert validation is REQUIRED before any regulatory action", "AI recommendations are advisory only", "Clinical judgment must supersede AI suggestions", "This is NOT a final regulatory determination."

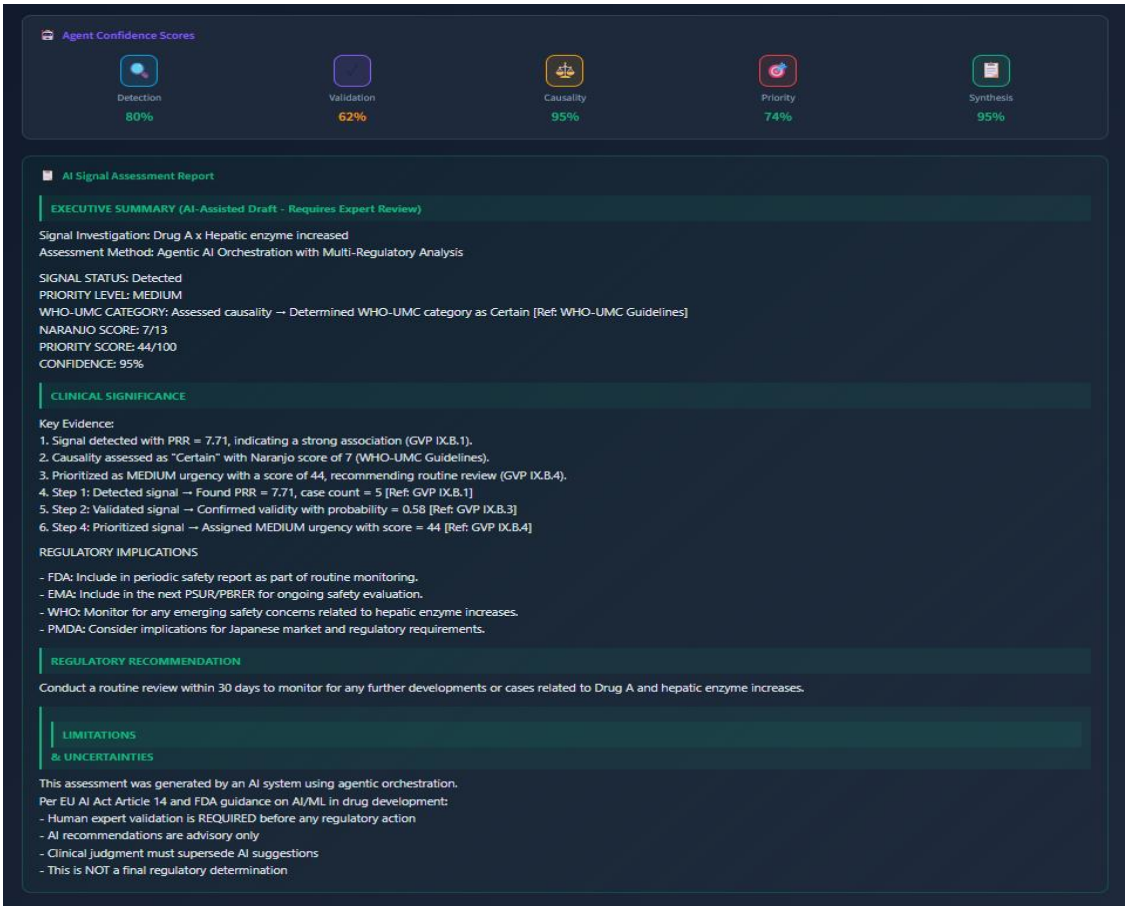

Figure 11: AI Signal Assessment Report with Executive Summary, Clinical Significance, Regulatory Implications, Regulatory Recommendation, and Limitations & Uncertainties. Note the explicit advisory-only statements.

2.3 Phase 3: Human-in-the-Loop Expert Determination

The Expert Determination tab is where the fundamental design principle—**AI Augments, Human Decides**—is implemented. This phase ensures compliance with EU AI Act Article 14 human oversight requirements and FDA guidance on AI/ML in drug development.

Expert Determination Interface

The interface prominently displays "Human-in-the-Loop Decision Required" with the explanation: "The AI has analyzed the signal and provided a recommendation. As a qualified expert, you must review the evidence and make your own determination—which may agree or disagree with the AI. Your decision is official regulatory action."

Required fields include: Reviewer Email (for 21 CFR Part 11 compliance), Role (Safety Scientist, PV Physician, Medical Reviewer), Product, Preferred Term. The AI Recommendation is labeled "DECISION SUPPORT ONLY" with instructions to "Run agentic investigation first (Tab 2)." Four determination options maps to regulatory actions: URGENT (Immediate regulatory action required), HIGH (Priority review within this week), MEDIUM (Routine safety review), LOW (Continue monitoring).

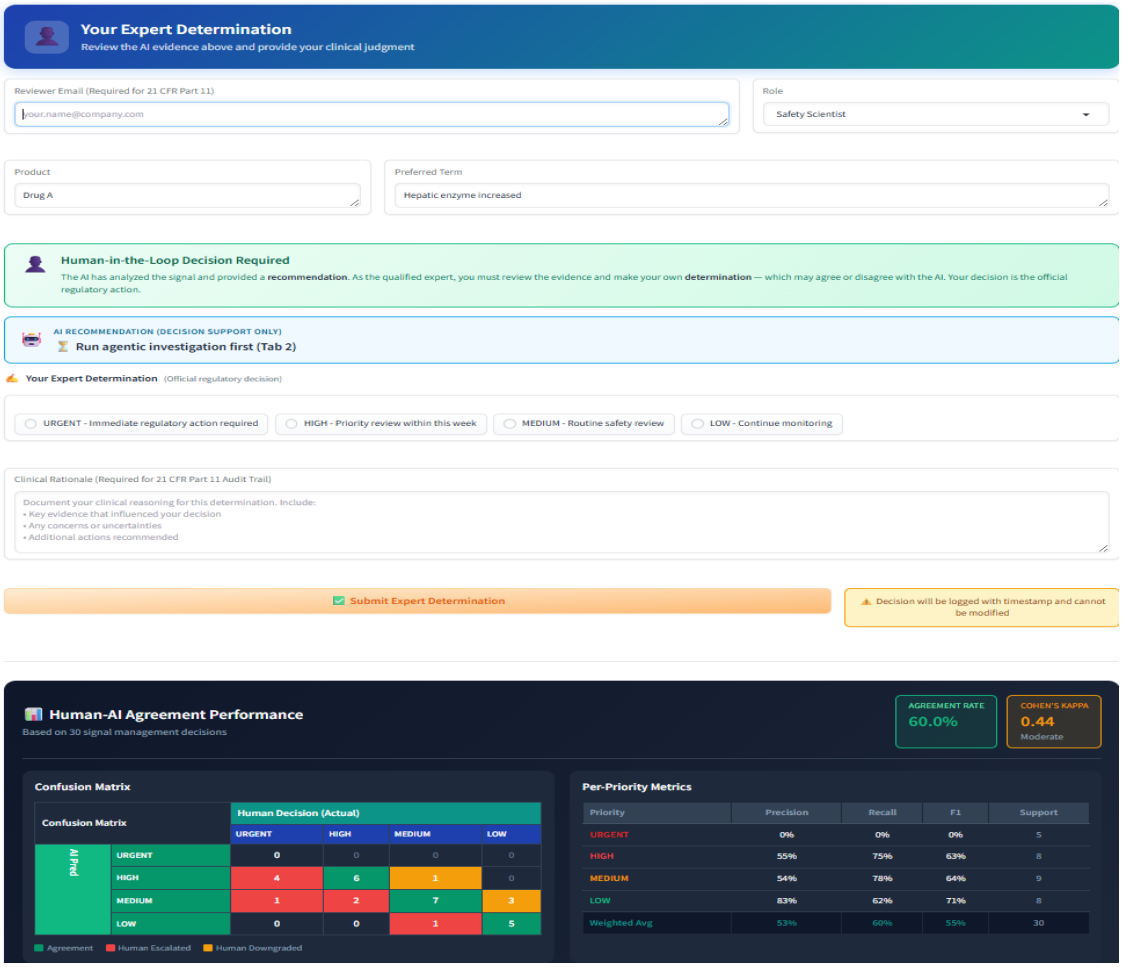

Figure 12: Expert Determination interface showing reviewer identification, Human-in-the-Loop Decision Required banner, AI recommendation as decision support only, determination options, clinical rationale field, and Human-AI Agreement Performance metrics.

Submitting Expert Determination

When the expert completes their determination, selecting MEDIUM and entering clinical rationale ("Agree with the agent's assessment")—the system records the decision with full traceability. The confirmation shows: Investigation ID, Timestamp (UTC), Reviewer email, Role, Electronic Signature hash, Product, Preferred Term, Your Decision (MEDIUM), AI Recommendation (MEDIUM), Agreement status (checkmark), and Clinical Rationale.

The Cumulative Agreement Rate (61%, 19/31 decisions) tracks ongoing human-AI alignment, providing data for system recalibration per the PACT-CARE™ Reliability & Recalibration step.

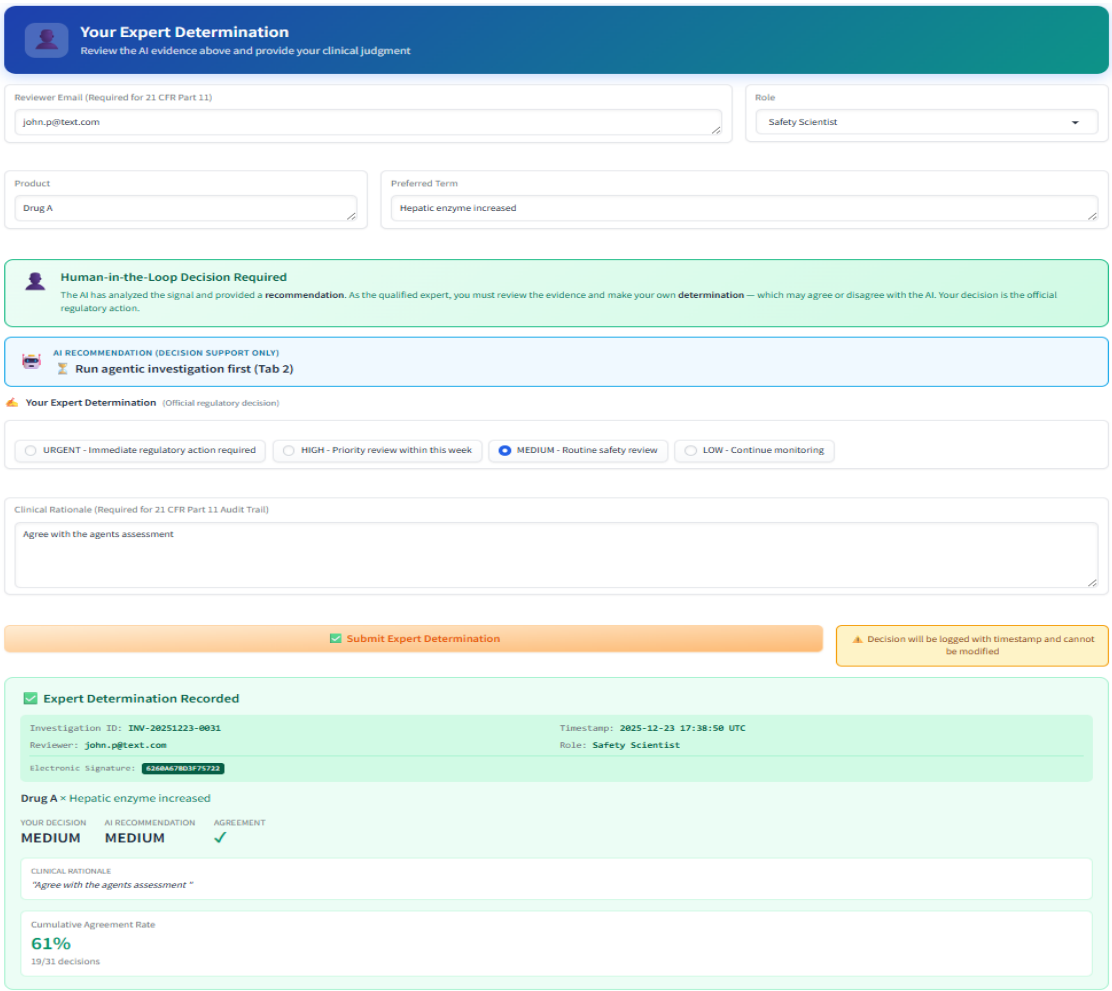

Figure 13: Expert Determination Recorded confirmation showing Investigation ID, timestamp, reviewer identification, electronic signature, decision comparison (Your Decision vs AI Recommendation with agreement checkmark), clinical rationale, and Cumulative Agreement Rate (61%).

Human-AI Agreement Performance Analytics

The performance analytics dashboard provides insights into human-AI collaboration patterns. Based on 31 signal management decisions, the system shows 61.3% Agreement Rate with Cohen's Kappa of 0.46 (Moderate agreement). The Confusion Matrix visualizes where AI and human decisions align or diverge, color-coded: green (agreement), orange (human escalated), red (human downgraded).

Per-Priority Metrics show Precision, Recall, and F1 scores for each priority level. Key insights: 3 Missed Critical (AI LOW → Human HIGH), 1 False Alarm (AI HIGH → Human LOW), 8 Human Escalations (26% of decisions), 4 Human Downgrades (13%). The interpretation note explains: "For safety-critical PV applications, the AI is tuned for high sensitivity (catching potential signals). Some false alarms are acceptable; missed critical signals are concerning."

The Decision History Log provides a complete record of all determinations with investigation_id, timestamp, reviewer details, product, preferred term, human_decision, ai_recommendation, agreed status, rationale, and methods supporting 21 CFR Part 11 audit requirements.

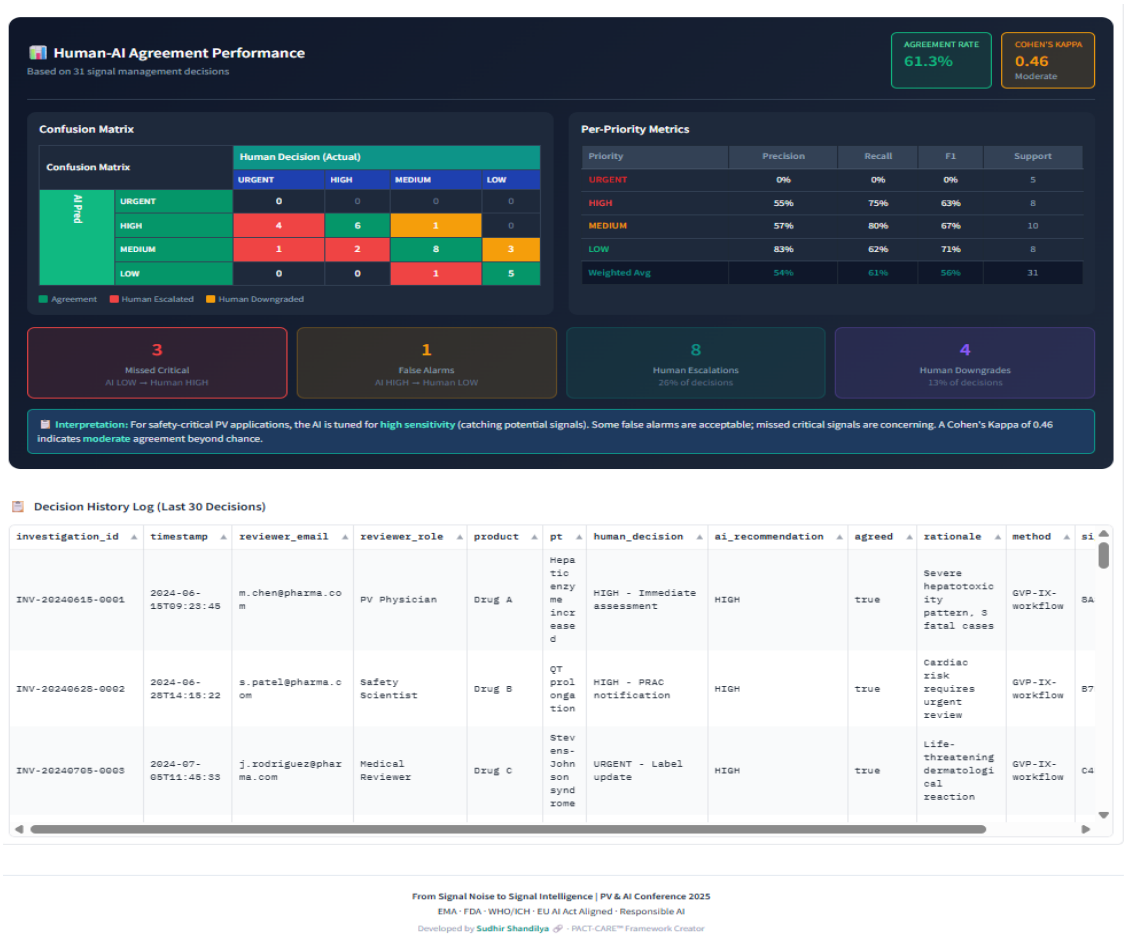

Figure 14: Human-AI Agreement Performance dashboard showing Agreement Rate (61.3%), Cohen's Kappa (0.46), Confusion Matrix, Per-Priority Metrics, key disagreement counts, interpretation guidance, and Decision History Log for audit trail.

2.4 Phase 4: Audit Trail

The Audit Trail tab provides complete documentation required for regulatory inspection. Per 21 CFR Part 11, GVP Module IX.B.5, and EU AI Act Article 12 (traceability), every decision must document: WHAT, WHY, WHEN, INPUTS, OUTPUTS, ALTERNATIVES CONSIDERED, and LIMITATIONS.

LLM Orchestration Audit Trail

The audit trail header states: "Complete chain-of-thought reasoning log • EU AI Act Article 13 Compliant • 21 CFR Part 11 Ready." The Investigation Metadata section records: Investigation ID, Started timestamp, Completed timestamp, and Workflow type (Agentic AI Signal Management - Multi-Regulatory).

The Original User Query is preserved verbatim: "Investigate Drug A × Hepatic enzyme increased." Target Extraction (Query Understanding) shows how the system parsed this: Product Identified (Drug A), Event Identified (Hepatic enzyme increased). The LLM Orchestration Trace section begins with: "This section shows how the AI reasoned through each step," followed by timestamped step-by-step reasoning.

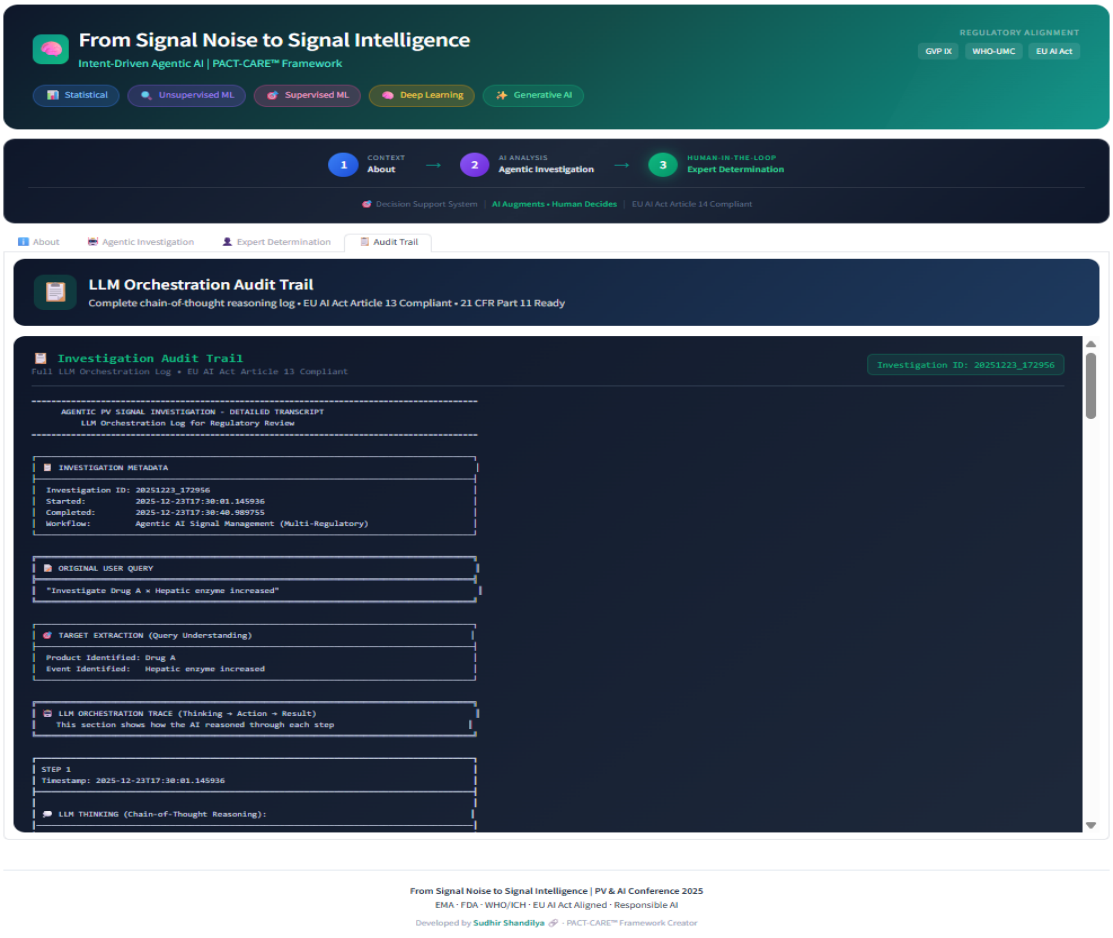

Figure 15: LLM Orchestration Audit Trail showing Investigation Metadata, Original User Query, Target Extraction (query understanding), and beginning of LLM Orchestration Trace with timestamped chain-of-thought reasoning.

Complete Audit Documentation

The audit trail continues with the recommendation summary, showing the complete evidence chain and final assessment. Multi-Event Signal Reports document each investigated drug-event pair with: Report #, Priority, Score, Cases, Fatal count, and WHO-UMC category.

ML Models Used documents which models were invoked (for GAMP 5 validation tracking). Regulatory Frameworks Applied lists all frameworks the system aligned to: FDA, EMA, WHO, ICH, PMDA, EU_AI_ACT. The Disclaimer section states: "This AI-generated transcript is provided for DECISION SUPPORT ONLY. It requires validation by a qualified PV professional before any regulatory action per EU AI Act Article 14 and GVP Module IX.C. The human expert's determination is the OFFICIAL regulatory decision."

The transcript ends with a generation timestamp, providing a complete, immutable record suitable for regulatory submission.

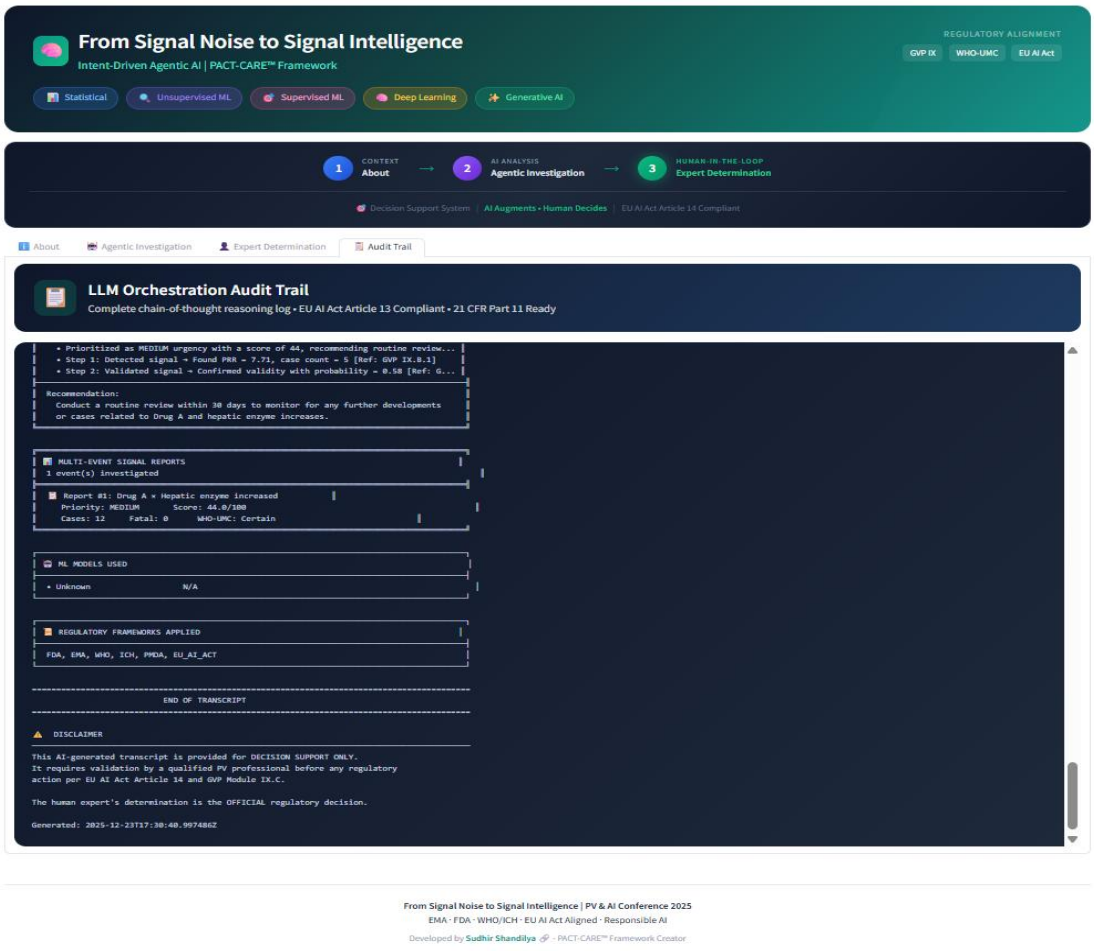

Figure 16: Audit Trail completion showing Recommendation summary, Multi-Event Signal Reports, ML Models Used, Regulatory Frameworks Applied, Disclaimer (decision support only), and generation timestamp.

3. Discussion

3.1 Why This Matters

This work is not about demonstrating that agentic AI *can* be built—that's been established. It's about demonstrating *how* to build it responsibly in regulated environments, with frameworks that are understandable by domain experts, implementable by engineering teams, auditable by regulatory professionals, and transferable across domains.

The friction encountered in developing this work—not always comfortable—helped shape it. Challenging established mindsets, both from hardcore techies and architects with fixed perspectives on one end, and individuals who talk AI via the likes of ChatGPT on the other, revealed the gap this work addresses.

Is this perfect? **No.** But it cuts through noise. The PV example was deliberate—a regulated domain with clear requirements, where the cost of both false positives (alert fatigue) and false negatives (missed signals) is measurable.

3.2 Domain Transferability

The architecture transfers across functions—R&D, regulatory, operations, finance—and across industries. Anywhere you need autonomous AI with compliance, auditability, and human oversight, the same pattern applies. Domain-specific customization involves knowledge bases (PV uses GVP/ICH; CMC would use ICH Q7-Q12), domain prompts, compliance requirements, and integration endpoints.

3.3 Limitations

The demonstration system uses synthetic ICSR data designed to illustrate capabilities rather than reflect production volumes. ML models require retraining on larger, validated datasets for production use. LLM orchestration depends on external API availability. Future work includes integration with production safety databases and extended regulatory framework coverage.

4. Conclusion

This paper has presented an intent-driven agentic AI architecture for pharmacovigilance signal management, built on PACT-CARE™ and A5 RAZOR frameworks. Through a complete application walkthrough, we demonstrated how the system transforms natural language queries into regulatory-compliant signal assessments while maintaining human authority over all determinations.

The core insight: **intent replaces rigid business logic**. The workflow forms based on the outcome you need, not the process you inherited. You define *what*; the system determines *how*—while staying fully compliant.

"AI Augments • Human Decides"

References

European Medicines Agency. (2017). Guideline on good pharmacovigilance practices (GVP) Module IX – Signal management. EMA/827661/2011 Rev 1.

U.S. Food and Drug Administration. (2023). 21 CFR 314.80 – Postmarketing reporting of adverse drug experiences.

World Health Organization. (2018). The use of the WHO-UMC system for standardised case causality assessment.

European Parliament. (2024). Regulation (EU) 2024/1689 – Artificial Intelligence Act.

U.S. Food and Drug Administration. (2003). 21 CFR Part 11 – Electronic Records; Electronic Signatures.

Naranjo, C.A., et al. (1981). A method for estimating the probability of adverse drug reactions. Clinical Pharmacology & Therapeutics.

About the Author

Sudhir Shandilya has over 20 years of experience in healthcare and life sciences. He is the creator of the PACT-CARE™ and A5 RAZOR frameworks for responsible AI adoption in regulated environments.

Contact: LinkedIn: linkedin.com/in/sudhir-shandilya-17b285/
Working Demo: Hosted privately on Hugging Face (contact author for access)

— End of Document —