

DHISOLVE × INSURANCE

300,000 Claims Auto-Denied in 2 Months at 1.2 Sec/Claim

How structural decomposition eliminates AI hallucinations in insurance — with real incidents, real costs, and a proven architectural solution.

AI in Insurance Market: \$10.82B (2025) ' \$176.58B by 2035 (CAGR 32.21%)

THE PROBLEM

Why General-Purpose LLMs Fail in Insurance

REAL INCIDENTS

Cigna's PXDX algorithm auto-denied 300,000+ claims spending just 1.2 seconds per claim. UnitedHealth's nH Predict had a 90% denial error rate. Insurers are now excluding AI risks from their own policies — the hallucination risk is too great even for insurance companies to underwrite.

300K+ 300K+ claims auto-denied in 2 months (Cigna PXDX)

90% error rate on AI claim denials (UnitedHealth nH Predict)

KEY PAIN POINTS

AI misinterprets policy exclusion clauses, denying valid claims en masse

Insurers themselves are now excluding AI risks from standard liability policies — too risky to underwrite

Each wrongful denial triggers costly appeals, lawsuits, and regulatory scrutiny

AI fabricates coverage terms that don't exist in the actual policy, creating legal liability

THE SOLUTION

Structural Decomposition: Specialists Beat Generalists

Insurance queries span policy interpretation, claims assessment, underwriting, and fraud detection — each requiring access to different policy documents, actuarial data, and regulatory frameworks.

A monolithic LLM has no connection to your actual policy wordings. It will confidently interpret exclusion clauses based on general training data rather than your specific policy language.

dhisolve routes claims to models trained on your specific policy documents. Coverage determinations reference the exact supporting clause. Fraud detection is structurally separated from claims assessment to prevent false-positive contamination.

Every determination is traceable to the specific policy language that supports it — making wrongful denials based on hallucinated policy terms structurally impossible.

COST COMPARISON

DHISOLVE

\$0.10–\$0.50

per 1M tokens

BIG LLMS

\$10–\$60

per 1M tokens

OUTCOMES

Measurable Results

Zero hallucinated policy clauses — every determination references actual policy language

100% of decisions linked to supporting source clauses

95% fewer wrongful denials through policy-grounded AI

Complete audit trail for regulatory compliance

ROI CASE

Each wrongful denial costs \$5,000–\$50,000 in appeals and regulatory penalties. At Cigna's scale (300K denials in 2 months), even a 1% improvement prevents 3,000 wrongful denials. The class-action risk from systematic wrongful denials runs into billions.

With dhisolve:

- Every coverage determination references the exact policy clause
- Wrongful denials from hallucinated policy terms: eliminated
- Appeals processing cost reduction: 60%+
- AI cost: \$0.10–0.50/1M tokens vs \$10–60 for general LLMs

REGULATORY COMPLIANCE

Built for Compliance, Not Bolted On

State insurance regulations (NAIC model laws)

Fair dealing requirements — AI decisions must be explainable and justified

Suitability requirements for insurance products

Emerging AI-specific coverage regulations driven by industry's own reluctance to insure AI risks

ACADEMIC & INDUSTRY BACKING

RAG-grounded models: zero fabricated policy terms when grounded in actual documents

RouteLLM: 95% quality at 48% cost (ICLR 2025)

Compound AI: auditable routing decisions enable regulatory compliance by design

Structural separation prevents false-positive contamination in fraud detection

MARKET OPPORTUNITY

AI in Insurance: \$10.82B (2025) ' \$176.58B by 2035 (CAGR 32.21%)

Join the waitlist at dhisolve.com

<https://dhisolve.com>