A PERPENDO AI FRAMEWORK

If Your AI Doesn't Know the 5R's, It Doesn't Know Insurance.

5R

The 5R's of Insurance: Risk, Rules, Relationships, Regulations, Responsibilities, are the operating logic behind every underwriting decision. Most AI ignores them.

THE PROBLEM

Generic AI Was Built for Generic Data. Not for Insurance.

Insurance is structured around risk logic, legal rules, trust-based relationships, regulatory requirements, and human accountability. When AI skips these, it doesn't just underperform, it misprices, misfiles, and misses.

It Sees Policy Data

Not the risk logic behind the policy

AI reads structured fields, limit, deductible, occupancy code. It doesn't interpret how those fields interact with the peril environment, ITV gaps, or the underwriter's intent at binding.

It Follows Instructions

Not the rules governing when they apply

Rating algorithms apply factors. They don't know when a filed rate is inapplicable, when a coverage trigger is ambiguous, or when policy language conflicts with endorsement intent.

It Processes Transactions

Not the relationships that shape them

A submission from a preferred broker carries context no database captures. Aggregate exposure across a reinsurance relationship isn't visible in a single policy API call.

It Runs on Inputs

Not the regulations that constrain them

NAIC Model Bulletin, state AI use rules, IFRS 17, generic AI wasn't designed to answer: 'Is this output explainable?'

It Produces Outputs

Not accountability for the decisions

When an AI-assisted decision triggers a coverage dispute, someone has to own it. Generic AI produces no audit trail, no explanation chain, no assignable responsibility.

THE FRAMEWORK

The 5R's of Insurance

The operating logic behind every decision, and the standard any AI built for insurance could be measured against.

R1	R2	R3	R4	R5
RISK	**RULES**	**RELATIONSHIPS**	**REGULATIONS**	**RESPONSIBILITIES**
Exposure · Peril · Severity	*Rating · Coverage · Policy Language*	*Broker · MGA · Reinsurer · Cedent*	*NAIC · State Filings · IFRS · Solvency*	*Accountability · Explainability · Audit*
"What is the true nature of this risk?"	*"What rules govern how we write it?"*	*"Who carries it and what do we owe them?"*	*"What can we legally do and say about it?"*	*"Who is accountable for this decision?"*

Every underwriting decision you make touches all five. AI built for insurance must too.

R1 · RISK

Risk assessment isn’t just a score. It's a judgment.

An AI built on the 5R's standard treats risk as a multi-dimensional judgment, not a lookup table. Exposure quality, peril interaction, and severity trajectory must be understood in context, not just computed.

Exposure · Peril · Severity

"What is the true nature of this risk?"

What AI must do differently:

Interpret exposure context, not just score data fields

Exposure quality is the foundation

COPE data, ITV, coverage structure, if these inputs are inconsistent or stale, every downstream output is wrong. Risk-native AI flags data gaps before scoring, not after.

Peril interaction changes the math

Wildfire risk near a flood zone with aging infrastructure is not three separate scores averaged. Correlated perils compound. AI must model the interaction, not the components in isolation.

Severity trajectory isn't historical data

Material costs, labor shortages, litigation funding, venue shifts , severity is structurally different from what loss triangles show. Risk-aware AI incorporates forward-looking indices, not just backward-looking curves.

Emerging risk has no loss history

Cyber aggregation within a property portfolio, climate-driven secondary perils, new construction techniques. if the model has never seen a loss, the underwriter's judgment is the only valid input.

R2 · RULES

Rating rules don't run themselves.

Policy language is not structured data. It's legal text that defines when coverage applies, when it doesn't, and what happens when it's ambiguous. AI must be built to navigate it, not just read it.

Rating Logic

Filed vs. advisory rates

State-filed rates apply in specific contexts. AI must know jurisdiction, line of business, and filing status , not just apply a factor.

Schedule rating constraints

Debits and credits have regulatory caps. AI that applies unlimited schedule modification isn't just wrong, it's a regulatory challenge

Tier eligibility rules

Risk characteristics determine tier placement. AI must evaluate eligibility criteria in sequence, not in parallel.

Coverage Logic

Coverage triggers

When does a peril become a covered cause of loss? AI must parse trigger language, occurrence vs. claims-made, sudden vs. gradual, named vs. open perils.

Exclusion interpretation

Exclusions carved out by endorsement reinstate what the base form excludes. AI that processes forms sequentially without tracking interaction gets this wrong consistently.

Valuation basis

ACV, replacement cost, agreed value, functional replacement , the valuation method changes the reserve, the premium, and the reinsurance recovery.

Policy Language

Manuscript vs. standard forms

Commercial and specialty risks use manuscript language. AI trained only on ISO forms cannot interpret it.

Ambiguity & contra proferentem

Courts construe ambiguous policy language against the insurer. AI must flag ambiguity, not resolve it arbitrarily.

Endorsement stacking

Each endorsement modifies the base form. The final coverage position is a cumulative function of all modifications in sequence.

An underwriter reads between the lines. AI built on the 5R's standard must too.

Insurance is a trust-based industry.
AI that ignores relationships ignores the deal.

Every submission carries context that no data field captures. Who submitted it, their authority level, their loss history with your book, the reinsurer who has treaty exposure, these are not metadata. They are the risk.

BROKER / PRODUCER	MGA / WHOLESALER	REINSURER / TREATY	CEDENT / POLICYHOLDER
Binding authority, submission quality, relationship tier, premium volume, loss ratio, context AI must carry into every interaction.	Delegated underwriting authority limits, appetite boundaries, reserving standards, AI must know what authority is delegated and what requires referral.	Attachment points, exclusions, reporting triggers, aggregate limits: a pricing decision without reinsurance context is incomplete.	Multi-year relationship, claims behavior, risk improvement commitments, strategic account status, none of this lives in a submission form.

⇔ ⇔ ⇔

Relationship context shapes appetite, authority, pricing, and accountability.

R4 · REGULATIONS

Compliance isn't a constraint on AI. It's the architecture.

Most AI treats compliance as a checklist bolted on after the fact. In insurance, regulation shapes what you can price, how you must explain it, and who is accountable when it goes wrong. It has to be built in from the ground up.

NAIC Model Bulletin on AI

AI-assisted decisions must be explainable, documented, and auditable for regulators and policyholders. Insurers are expected to maintain governance controls, testing procedures, and documentation to demonstrate that automated decision systems do not result in unfair discrimination.

IFRS 17 / Solvency II

Liability measurement and capital requirements depend on actuarial modeling and governance. AI systems that influence reserving, cash-flow projections, or lapse assumptions must remain consistent with contract boundary definitions and capital frameworks.

State Rate & Form Filings

Insurance rates, rating variables, and policy forms are regulated at the state level and typically must be filed under prior-approval, file-and-use, or use-and-file regimes.

State AI Use Regulations

States are increasingly regulating the use of algorithms and external data in insurance decision-making. Regulations that governs predictive models and external consumer data in underwriting, and issuing guidance on fairness testing, model governance, and transparency

Perpendo is built compliance-native, compliance isn't added later. They are the starting point.

R5 · RESPONSIBILITIES

Someone has to own the decision. Make sure it isn't a black box.

When an AI-assisted coverage decision is disputed, a claim is denied, or a regulator asks how a rate was derived, 'the model said so' is not an answer. The 5R's standard requires that every AI output be ownable, explainable, and auditable.

01 Explainability

Every AI output must be translatable into underwriting logic. Not just a confidence score

02 Audit Trail

Decisions, inputs, model version, data sources, timestamps, captured at every step. The reasoning must be fully reconstructable.

03 Human-in-the-Loop

AI assists, humans decide. The standard defines where machine recommendations stop and professional judgment begins.

04 Bias & Fairness Testing

Rate structures and eligibility rules that produce disparate impact on protected classes create legal exposure.

05 Model Governance

Which version of the model made this decision? When was it last validated? Who approved the change? Model governance is risk management for your AI stack , same rigor as investment portfolio governance.

THE STANDARD IN PRACTICE

What 5R's-Native AI Looks Like in Your Workflow

For the underwriter at the desk, the 5R's standard changes what AI puts in front of you, and what it asks you to do with it.

SUBMISSION

R

AI DOES:

AI parses COPE data, flags ITV gaps, identifies peril exposure by geocode, scores concentration risk.

YOU ASK:

Challenge: Is the occupancy classification consistent with the risk address and construction type? Does the ITV reflect current replacement cost?

RATING

R

AI DOES:

Rating factors applied per filed schedule. Coverage triggers evaluated against policy form. Endorsement stack reconciled.

YOU ASK:

Challenge: Is the schedule rating within filed credit/debit limits? Does the coverage trigger match what the insured actually needs?

REFERRAL

R

AI DOES:

Broker tier, binding authority limits, aggregate exposure to reinsurance treaty , all evaluated before recommendation is surfaced.

YOU ASK:

Challenge: Does this risk fit within treaty per-risk and aggregate parameters? What is this broker's loss ratio on similar risks?

DECISION

R

AI DOES:

State filing status confirmed. Regulatory constraints applied. Adverse action language pre-populated if applicable.

YOU ASK:

Challenge: Is this rate defensible to a state regulator? Is the adverse action notice accurate and complete?

BIND & FILE

R

AI DOES:

Audit trail locked. Decision reason chain captured. Model version recorded. Compliance log updated.

YOU ASK:

You own this decision. The AI assisted. The file proves it.

R1 RISK

R2 RULES

R3 RELATIONSHIPS

R4 REGULATIONS

R5 RESPONSIBILITIES

The standard has a name.

The 5R's of Insurance is not a checklist. It is the operating logic of the industry, built over the experience that insurance is fundamentally about trust, precision, and accountability.

ASK YOUR AI SOLUTION

Can your system explain a decision against all five dimensions? If it can't, it isn't built for insurance.

ASSESS YOUR WORKFLOW

Map where each R enters your underwriting process. Where AI assists and where it doesn’t. That gap is your risk exposure.

EXPECT THE STANDARD

Compliance-native, insurance-domain, explainable AI. The 5R's aren't aspirational. They're the standard.