




International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: <https://www.svedbergopen.com/>



Research Paper

Open Access

Joint Embedding Predictive ArchitectureA Framework for the Next Generation of Insurance Underwriting

Aakash Angadi

Independent Researcher, India. Email: aakash.angadi@gmail.com

Abstract

The insurance underwriting industry stands at an inflection point. For decades, risk assessment has relied on structured data, static actuarial tables, and labor-intensive manual review. The emergence of Joint Embedding Predictive Architecture (JEPA) — an AI framework introduced in 2022 by Yann LeCun, then Chief AI Scientist at Meta — represents a paradigm shift in how machines understand, model, and predict complex risk.

Unlike large language models or generative AI that predict at the pixel or token level, JEPA learns abstract semantic representations of the world — understanding not just what data says, but what it means. This capability is uniquely aligned with the demands of modern underwriting: synthesizing heterogeneous, unstructured data sources into reliable, explainable risk assessments.

Key Findings

- JEPA enables semantic understanding of images, video, and multi-modal data — relevant for property, auto, and specialty line underwriting.
- World modeling capabilities allow JEPA-based systems to anticipate how risks may evolve over time, supporting more proactive underwriting.
- Based on early AI adoption data, insurers leveraging next-generation models could see meaningful reductions in cycle time and improvements in loss ratios.
- JEPA addresses a key limitation of current AI: the difficulty of reasoning about complex, non-standard risk scenarios.

This paper outlines the JEPA architecture, its differentiators relative to existing AI approaches, candidate underwriting use cases, and a practical roadmap for adoption — while noting that JEPA in underwriting remains an emerging field, and that projected impacts should be treated as directional rather than guaranteed.

I. Introduction

Insurance underwriting has long faced a fundamental tension: the need for speed competes directly with the need for accuracy. Processing a standard application manually has historically taken three to five days, while customers increasingly expect near-instant decisions. Meanwhile, the complexity of risks — from climate-exposed properties to novel cyber threats to emerging technology companies — continues to outpace traditional actuarial models.

Artificial intelligence has begun to address this tension. According to a widely cited 2025 industry analysis, AI has reduced average underwriting decision time for standard policies from three to five days to roughly 12.4 minutes, while maintaining a 99.3% accuracy rate in risk assessment. The same analysis reports that over 380 companies — including technology vendors and established insurance carriers — now rely on AI-based underwriting as a decision-support layer. Yet a significant gap remains: current AI excels at structured, standard cases but struggles with the unstructured, non-standard risks that define the most valuable and complex segments of the market.

This is precisely where JEPA enters — not as an incremental improvement, but as a foundational shift in AI capability.

1.1 The Limits of Current AI in Underwriting

Most AI tools in underwriting today fall into two categories:

- Rule-based automation and traditional ML systems — handle standard cases efficiently but cannot reason about novel risks.
- Large Language Models and generative AI — powerful for document analysis and natural language tasks, but they predict at the token level and can hallucinate, making them unsuitable for high-stakes actuarial decisions without careful human oversight.

The gap lies in semantic understanding: the ability to look at a drone image of a property and grasp structural risk, or to interpret telematics footage and comprehend driver behavior — not merely classify it statistically.

2. What Is JEPA? A Technical Overview

Joint Embedding Predictive Architecture (JEPA) is an AI architecture introduced by Yann LeCun in his June 2022 position paper, "A Path Towards Autonomous Machine Intelligence." LeCun — a 2018 Turing Award winner widely regarded as one of the founders of modern deep learning — proposed JEPA as a step toward machines that can build internal models of how the world works, enabling them to reason, plan, and adapt to unfamiliar situations.

JEPA research originated at Meta's Fundamental AI Research (FAIR) lab, where LeCun served as Chief AI Scientist for seven years before announcing his departure in November 2025 to launch Advanced Machine Intelligence (AMI) Labs, a new venture continuing the research program. JEPA-based models (I-JEPA, V-JEPA, V-JEPA 2, V-JEPA 2.1) remain open-source and are actively developed by Meta and the broader research community.

2.1 Core Architectural Principles

JEPA operates on a fundamentally different principle from generative models. Rather than reconstructing raw inputs such as pixels or tokens, JEPA learns by predicting abstract representations in latent space:

Architecture Type	Prediction Domain	Underwriting Implication
Generative AI (LLMs)	Predicts in pixel/token space	Can reconstruct surface detail but risks hallucination; weaker at deep semantic structure
Standard Computer Vision	Classifies patterns statistically	Strong on labeled categories; weaker on novel or composite risks
JEPA	Predicts in abstract representation space	Learns semantic meaning; generalizes to unfamiliar scenarios; supports world-modeling

2.2 The JEPA Family of Models

The JEPA research program has produced a growing family of specialized models, each with relevant underwriting applications:

I-JEPA (Image-JEPA, 2023)

Introduced by Meta AI researchers in June 2023, I-JEPA is a self-supervised learning framework for images. It works by masking portions of an image and predicting those masked regions in abstract representation space — not at the pixel level. This encourages the model to learn semantic understanding of visual content, such as structural integrity, environmental hazard indicators, and property condition.

V-JEPA and V-JEPA 2 (Video, 2024–2025)

V-JEPA extended JEPA to video, learning spatiotemporal representations of how visual scenes evolve across frames. V-JEPA 2, released by Meta in June 2025, is a 1.2-billion-parameter world model trained on more than one million hours of internet video. V-JEPA 2 demonstrates state-of-the-art performance on motion understanding and action anticipation, and supports zero-shot planning in robotics. A refined release, V-JEPA 2.1, followed in March 2026. These advances are directly relevant to telematics analysis, surveillance footage, and dynamic risk assessment, where behavior over time matters more than a single snapshot.

H-JEPA (Hierarchical JEPA)

As outlined in LeCun's 2022 position paper, H-JEPA is a conceptual framework for multi-level, multi-timescale predictive models. The idea is to stack JEPAs so that lower levels predict fine-grained, short-horizon representations and higher levels predict coarse-grained, long-horizon ones. While H-JEPA remains primarily a research direction, its hierarchical reasoning approach is well-suited to the multi-level structure of complex commercial and specialty underwriting.

Why JEPA Beats Pixel-Level Prediction for Underwriting

When an underwriter inspects a photo of a commercial roof, they do not analyze pixel brightness — they assess structural soundness, material aging, drainage patterns, and maintenance quality. JEPA's abstract representation space is designed to mirror this kind of human reasoning, discarding pixel-level noise and focusing on semantically meaningful features. This characteristic positions JEPA-based models to be more robust on novel property types and more aligned with how expert underwriters actually think — though real-world performance in underwriting workflows will need to be validated in production.

3. How JEPA May Transform Underwriting

The use cases that follow are forward-looking. JEPA-based underwriting deployments remain early, and the following sections describe candidate applications based on the architecture's capabilities and adjacent AI deployments already in production.

3.1 Property & Casualty Underwriting

Property assessment is among the most image-intensive tasks in underwriting. Traditional approaches rely on costly on-site inspections or limited aerial photography. I-JEPA's image-understanding capabilities can support:

- Drone and satellite image analysis — semantic interpretation of roof condition, structural integrity, proximity to flood zones, vegetation encroachment, and construction quality with reduced reliance on labeled training data for every property type.

- Virtual inspection — Processing series of property images to support risk assessment that complements or partially substitutes for in-person inspection.
- Climate and catastrophe risk modeling — Assessing wildfire, flood, and wind exposure at the property level rather than relying on broad geographic zones.

3.2 Auto & Mobility Underwriting

Telematics has already begun to transform auto insurance pricing. V-JEPA-style models can extend this by supporting genuine understanding of driving behavior from video data:

- Dashcam footage analysis — Processing driving video to assess behaviors such as following distance, reaction time, lane discipline, and hazard response — extending beyond the speed and braking data captured by traditional telematics sensors.
- Usage-based insurance (UBI) — Behavioral scoring that supports dynamic premiums as driver risk profiles evolve.
- Fleet underwriting — Monitoring commercial fleets at scale to identify behavioral risk patterns across thousands of vehicles.

3.3 Life & Health Underwriting

Life and health underwriting is among the most data-intensive processes in insurance. JEPA-based approaches introduce several emerging capabilities:

- Medical imaging understanding — Semantic image analysis to interpret medical scans at a conceptual level, complementing existing labeled diagnostic models.
- Longitudinal health modeling — Multi-timescale reasoning to model how health risks evolve over the policy lifetime.
- Wearable and IoT data integration — Synthesizing continuous health data streams from wearables into coherent, evolving risk profiles rather than static snapshots.

3.4 Commercial & Specialty Lines

This is where JEPA's potential advantages are most pronounced. Standard automation handles commodity risks; specialty lines require reasoning about novel situations:

Line of Business	Candidate JEPA Application
Emerging Technology Risk	Reasoning about novel risk factors (AI liability, autonomous systems) by analogizing from related domains, even with limited historical claims data.
Cyber Underwriting	Analyzing multi-modal signals — network architecture diagrams, code repositories, incident history — to support holistic cyber risk assessment.

Marine & Aviation	Processing operational footage and sensor data to assess vessel and aircraft condition and operational risk.
Environmental Liability	Interpreting site imagery and environmental data to assess contamination risk and regulatory exposure.
Political Risk & Trade Credit	Hierarchical reasoning that integrates macro geopolitical signals with firm-level data to model sovereign and counterparty risk.

4. The World Model Advantage: Predictive Underwriting

One of JEPA's most distinctive contributions is its capacity for world modeling — constructing an internal model of how a system evolves over time and responds to changing conditions. Meta's V-JEPA 2 release explicitly describes the model as a "world model" for understanding, predicting, and planning in the physical world.

Traditional underwriting is largely reactive: risks are assessed at a point in time, policies are priced based on historical data, and adjustments occur at renewal. JEPA-style world modeling supports a different paradigm:

From "Detect and Repair" to "Predict and Prevent"

World modeling capabilities can help insurers shift from responding to losses after they occur to anticipating and mitigating them. A JEPA-powered underwriting system can move beyond assessing a property's current condition to modeling how its risk profile may evolve given climate trajectories, maintenance patterns, and neighborhood development trends over the life of the policy. The NAIC has identified this "predict and prevent" shift as a meaningful direction for AI in insurance.

4.1 Dynamic Policy Structures

With continuous world modeling, policies could evolve dynamically:

- Premiums adjust in real time as risk profiles change — within regulatory and contractual constraints — rewarding mitigation behaviors and flagging emerging exposures.
- Coverage recommendations adapt to the insured's changing circumstances — proactively suggesting riders or endorsements as risk profiles evolve.
- Early warning systems alert underwriters and policyholders to emerging risk signals before losses materialize.

4.2 Portfolio-Level Risk Intelligence

Hierarchical JEPA architectures can in principle support risk modeling at multiple levels simultaneously — individual policy, portfolio segment, and enterprise-wide — creating a more coherent, near-real-time view of aggregate risk exposure than is typical with current tools.

5. JEPA in Agentic Underwriting Systems

Industry analysts including McKinsey have outlined a near-future vision in which much of customer onboarding in insurance is delivered through multi-agent AI systems — networks of specialized AI agents that collaborate to process applications. JEPA fits naturally as the perceptual and reasoning engine within such systems.

A future JEPA-powered agentic underwriting workflow might operate as follows:

Agent Role	JEPA Contribution
Intake Agent	Ingests and classifies documents, images, and sensor data; JEPA provides semantic understanding of multi-modal inputs.
Risk Profiling Agent	Builds comprehensive risk profiles using JEPA's representation learning to assess both current state and projected risk trajectory.
Pricing & Product Agent	Prices the risk and suggests policy structures, leveraging dynamic risk models for personalized premium calculation.

Compliance & Fairness Agent	Reviews decisions for regulatory compliance, bias, and ethical standards — critical given growing AI governance requirements.
Relationship Agent	Communicates with brokers and clients, explaining AI-derived risk assessments in plain language.

Combined with existing AI document-ingestion tooling, this architecture has the potential to reduce application-to-quote time meaningfully for complex commercial risks. Realized gains will depend on data quality, organizational readiness, and regulatory context.

6. Projected Impact: Directional Performance Targets

The figures below are directional projections based on current AI adoption benchmarks reported by industry analysts and JEPA's architectural capabilities. They should be treated as planning estimates, not guarantees. Actual results will vary by line of business, data quality, and implementation maturity.

Metric	Current State	Projected with JEPA-Enabled Systems
Underwriting Cycle Time (standard policies)	3–5 days	~12.4 minutes already reported by industry benchmarks for current AI; JEPA expected to extend this to complex cases
Inspection Cost	Full cost of on-site inspection	Meaningful reduction via virtual inspection (estimated 50–70%)
Underwriter Productivity	Baseline	30–50% improvement reported in industry analyses for AI-assisted underwriting
Fraud Detection Accuracy	Baseline detection	Improvements toward 90%+ accuracy on anomaly flagging (industry benchmarks)
Combined Ratio Improvement	Baseline	Industry studies cite 3–6 percentage point improvement for early AI adopters
New Market Access	Limited by actuarial data availability	Potential to expand insurability for emerging and novel risks

Sources for these benchmarks include published analyses by McKinsey, Earnix, BizTech Magazine, and SmartDev (2025). See References.

7. Ethical Considerations & Regulatory Compliance

The power of JEPA-based systems comes with commensurate responsibility. Insurers deploying these tools must address several critical governance dimensions.

7.1 Bias and Fairness

AI systems trained on historical data risk encoding historical biases into automated decisions. JEPA's self-supervised learning approach — which reduces reliance on labeled training data — does not eliminate this risk. Insurers should:

- Audit AI-derived risk scores across protected demographic categories.
- Implement fairness constraints in the model training and deployment pipeline.
- Maintain human oversight for any decision that could adversely affect policyholders.

7.2 Regulatory Landscape

State insurance regulators are actively building frameworks for AI oversight. The NAIC's AI Systems Evaluation Tool is currently in a multistate pilot running through September 2026, with twelve participating states including California, Colorado, Connecticut, Florida, Iowa, Louisiana, Maryland, Pennsylvania, Rhode Island,

Vermont, Virginia, and Wisconsin. The Tool examines insurer AI governance programs across four areas: extent of AI usage, internal governance, high-risk AI systems, and underlying data. Nationwide adoption is targeted for the NAIC Fall National Meeting in November 2026. The NAIC's Model Bulletin on AI use has been adopted in approximately 24 states. JEPA deployments must be designed with explainability and auditability from the outset.

7.3 Data Privacy

JEPA's multi-modal data ingestion — spanning images, video, IoT streams, and behavioral data — raises significant data privacy considerations. Insurers must ensure compliance with applicable state and federal privacy regulations and establish clear policies on data retention, consent, and third-party data use.

Regulatory Principle

Human oversight remains essential. JEPA-based tools are best deployed as decision-support systems that augment underwriter judgment — not as autonomous decision-makers for adverse actions. AI-influenced underwriting decisions should remain explainable and subject to human review, consistent with NAIC guidance and applicable state law.

8. Implementation Roadmap

Successful adoption requires a phased approach that aligns technical deployment with organizational readiness:

Phase	Focus	Key Activities
Phase 1 (0–12 months)	Foundation Building	Pilot I-JEPA-style models for property image analysis in P&C lines. Establish data governance framework. Train underwriters on AI-augmented workflows. Baseline performance metrics.
Phase 2 (12–24 months)	Expansion	Integrate V-JEPA-style video models for telematics and auto underwriting. Begin agentic workflow pilots for standard commercial lines. Implement real-time monitoring and bias auditing.
Phase 3 (24–36 months)	Scaled Adoption	Deploy hierarchical, multi-modal models for portfolio-level risk. Launch adaptive policy structures in regulator-aligned form. Extend to specialty lines and emerging risk categories.

9. Conclusion

JEPA represents a meaningful evolution in what AI can do for insurance underwriting — not merely an incremental improvement to existing tools, but a new capability: semantically understanding complex, multi-modal risk data and modeling how risks evolve over time.

The underwriting challenges that have resisted automation — complex commercial risks, novel emerging exposures, image- and video-intensive workflows — are precisely the challenges JEPA's architecture is designed to address. At the same time, JEPA in underwriting remains an early-stage application, and operational results will depend on data quality, governance, regulatory engagement, and rigorous validation.

Insurers that invest now in JEPA capabilities, data infrastructure, and governance will be positioned to lead the next decade of underwriting transformation. Those that delay risk falling behind on speed, accuracy, and the capacity to underwrite the risks that will define the economy of the 2030s.

Strategic Imperative

JEPA-based models are available today as open-source research releases and are advancing rapidly. The competitive window for thoughtful, governance-aligned early adoption is open now. Insurers should begin scoping pilot programs in property image analysis and telematics, while developing the data infrastructure, regulatory alignment, and governance frameworks needed for broader deployment.

References & Further Reading

1. LeCun, Y. (2022). A Path Towards Autonomous Machine Intelligence, Version 0.9.2. OpenReview.
2. Assran, M. et al. (2023). Self-Supervised Learning from Images with a Joint-Embedding Predictive Architecture (I-JEPA). Meta AI / CVPR.
3. Bardes, A. et al. (2024). V-JEPA: Revisiting Feature Prediction for Learning Visual Representations from Video. Meta AI.
4. Assran, M., Bardes, A., Ballas, N. et al. (2025). V-JEPA 2: Self-Supervised Video Models Enable Understanding, Prediction and Planning. arXiv:2506.09985.
5. Meta AI (June 11, 2025). Introducing the V-JEPA 2 World Model and New Benchmarks for Physical Reasoning.
6. McKinsey & Company (July 2025). The Future of AI in the Insurance Industry.
7. NAIC (2026). AI Systems Evaluation Tool — Pilot Program Documentation (Big Data and Artificial Intelligence (H) Working Group).
8. Earnix (April 2025). AI in Insurance Underwriting: Overcoming Challenges and Unlocking Value.
9. BizTech Magazine (March 2025). How Artificial Intelligence Is Transforming the Insurance Underwriting Process.
10. CNBC, Fortune, Euronews (November–December 2025). Coverage of Yann LeCun's departure from Meta and launch of AMI Labs.