Neuro-Symbolic Governance
The New Paradigm for Managing and Controlling AI Evolution in Enterprise

Fusion AI Lab, 21 October 2025

FUSION AI LABS

expert.ai

The State of AI today

*A staggering **95 % of corporate generative AI pilot projects fail** to generate measurable business value, with only 5 % achieving rapid revenue growth."*

*Companies **outsourcing AI through specialized vendors** see significantly better results - up to twice the success rate - compared to those attempting to build AI systems in-house.*

The GenAI Divide: State of AI in Business 2025

August, 2025** | **MIT

https://nanda.media.mit.edu/ai_report_2025.pdf

expert.ai

The pace of innovation

Governance has never been so important

Technology	Birth Year (Commercial)	Years to Reach 50% Adoption	Average Update Cycle (Years)	Innovation Speed Index (0–100)	Qualitative Notes
Personal Computer	1981	16	4	30	Continuous but cyclical innovation; major leaps every decade
Internet	1995	10	8	25	Huge initial disruption but base infrastructure remains stable
Mobile Telephony	1990	12	3	45	Frequent standard updates (2G→5G) drive steady innovation
Smartphone	2007	6	1.5	60	Rapid yearly iterations and global adoption
Cloud Computing	2006	8	2	55	Fast evolution from IaaS to PaaS/SaaS layers
Social Media	2004	5	1	70	Continuous feature updates and new platforms every few years
Internet of Things (IoT)	2010	12	3	40	Slow adoption but steady hardware and protocol evolution
Blockchain / Crypto	2009	9	1	65	Fast-paced innovation but still unstable ecosystems
Traditional AI (ML)	2012	10	2	50	Gradual progress leading to the Generative AI era
Generative AI (LLM)	2022	2	0.5	95	Unprecedented speed of innovation (GPT → Claude → Gemini → Sora)

FUSION AI LABS

expert.ai

The Neuro-Symbolic Approach

Bringing Logic and Learning Together

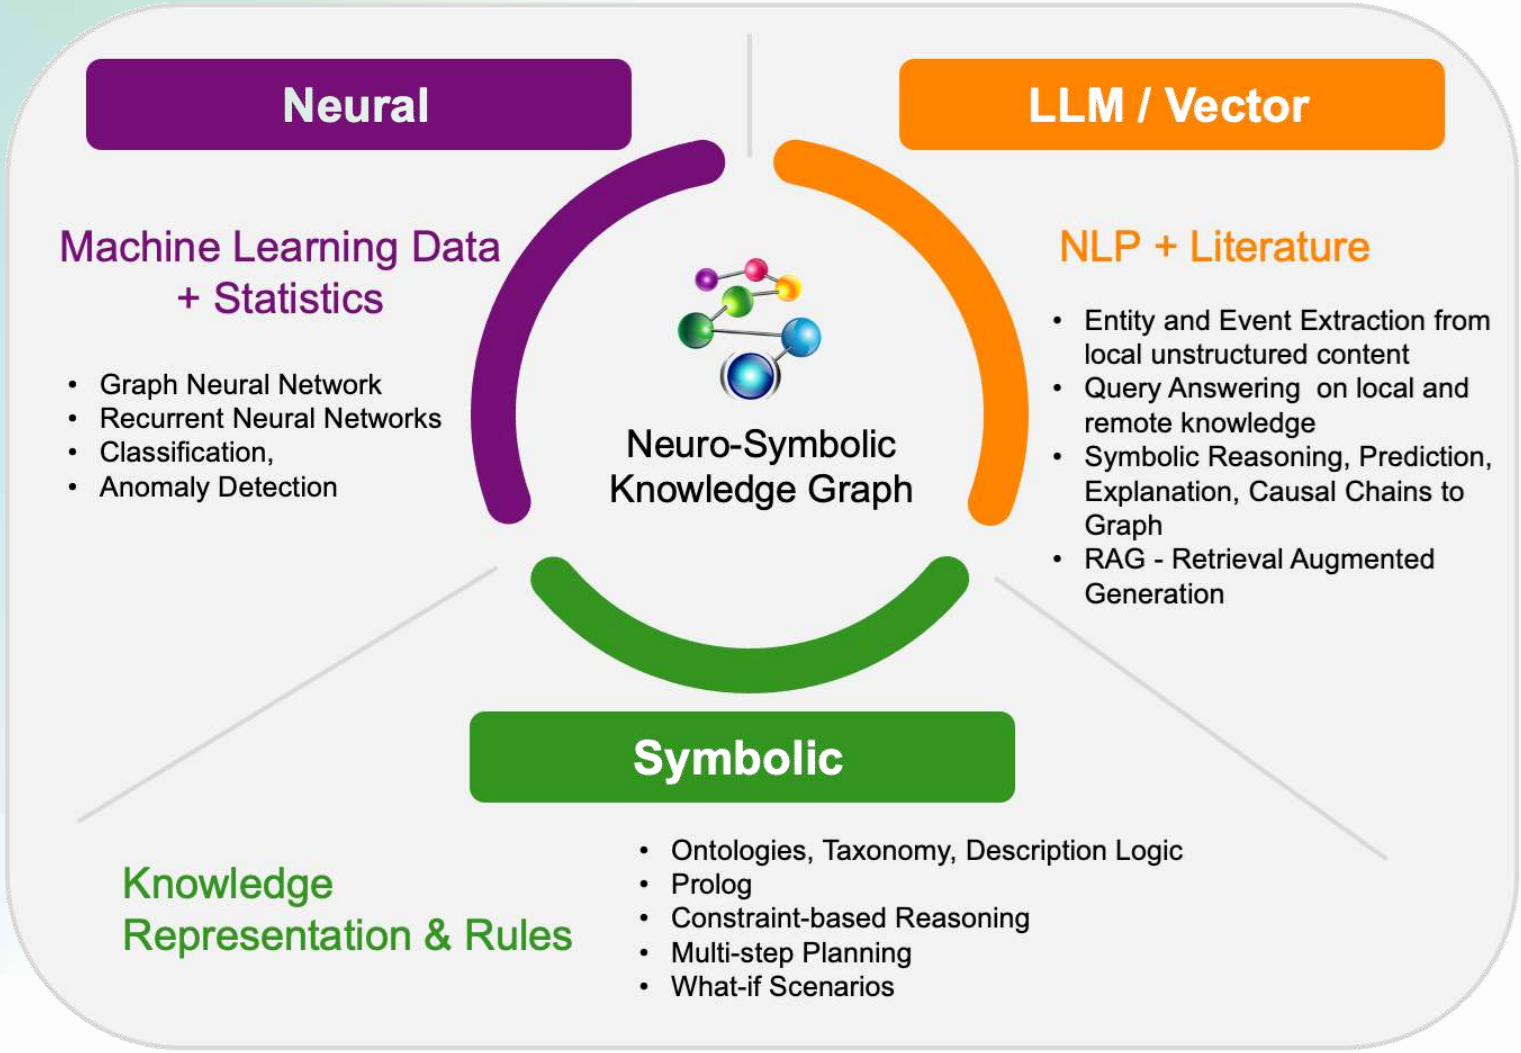

Neural AI – *the learning brain*.

It learns from massive amounts of data — spotting patterns, classifying information, and predicting outcomes. It's powerful and adaptive, but it doesn't really understand what the data means — it just learns correlations.

Symbolic AI – *the reasoning brain*.

It works with concepts, relationships, and rules — the way humans' reason. It can explain its decisions, plan steps, and run "what-if" scenarios, but it lacks the creativity and adaptability of neural models.

LLM / Vector AI – *the language brain*.

Large Language Models can read and generate text, connecting human language to knowledge. They're extraordinary at retrieving and summarizing information but have limited control over logic and consistency.

FUSION AI LABS

expert.ai

AI Governance Across Six Dimensions

From Power to Control

Projects

AI projects evolve faster than any past technology.
Project governance ensures **coherence and continuity** across all initiatives, allowing organizations to innovate quickly **without losing control or consistency**.

Processes

AI is not just about understanding documents but understanding **where each piece of information fits in the business flow**.
Process governance ensures AI connects data to context, keeping automation aligned with how the organization works.

Results

Generative AI is powerful but inherently uncertain — it will produce errors, and we can't always know when.
This makes **result governance** essential to ensure reliability, consistency, and accountability.

Explainability

Most RAG systems can only show *where* information comes from, not *why* it was chosen.
Embeddings and chunks tell you which text was used, not the logic behind the answer.
Explainability governance ensures that AI systems are transparent and auditable, revealing the reasoning that connects data to outcomes.

Cost

Generative AI introduces a new variable cost model which is difficult to estimate and can grow rapidly over time.
Semantic AI, by contrast, has a higher initial setup cost but almost no running cost.
Cost governance means deciding when to use one or the other to achieve efficiency.

Knowledge

Enterprise knowledge is complex, fragmented, and often domain-specific. LLMs alone struggle to understand it, especially when documents differ in structure and language. **Knowledge governance** means structuring this information into a predictable and reusable with domain-specific semantic model.

expert.ai

EI-Flow

Project

A governance framework to manage the evolution of multiple AI initiatives.

Why it matters

AI projects evolve faster than any previous technology. Without a governance framework, each initiative becomes a disconnected experiment.

Project governance ensures coherence, continuity, and strategic control across multiple AI initiatives – allowing the organization to **evolve without chaos**.

Customer value

- **Sustainable innovation**: every new use case reuses validated components, avoiding duplication.
- **Agility with control**: models and architectures evolve rapidly without breaking integrations.
- **Strategic visibility**: leaders can monitor all AI initiatives and investments from a single governance layer.

Neuro-Symbolic advantage

The neuro-symbolic approach merges the **adaptability of neural models** with the **logical precision of symbolic reasoning**.

This hybrid architecture allows enterprises to replace or improve models without losing meaning, structure, or traceability.

It transforms fragmented projects into a coherent, evolvable AI ecosystem that remains under full organizational control.

FUSION AI LABS

expert.ai

Project : One tools, multiple use cases

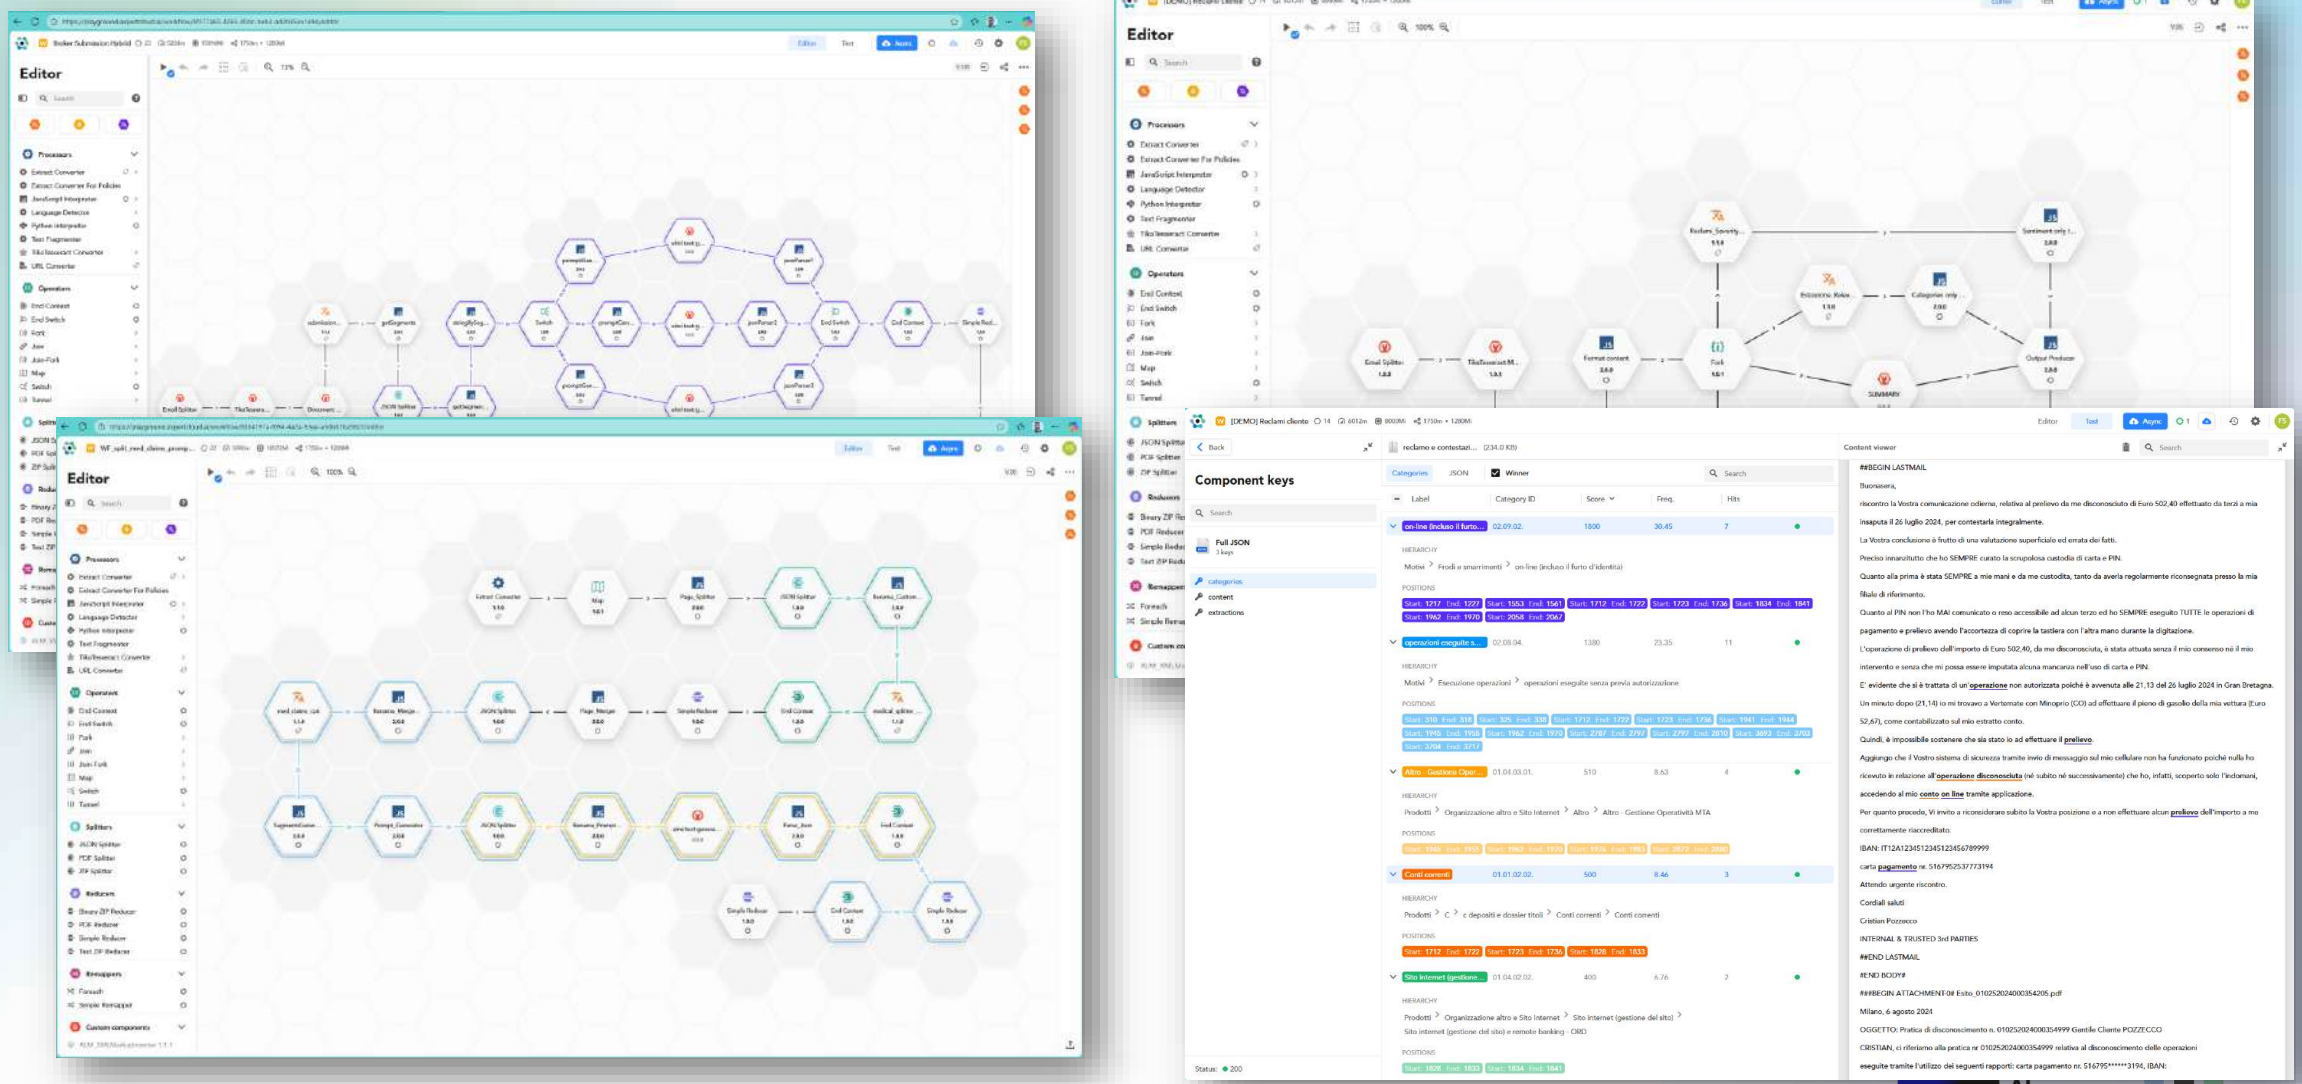

EI-Monitor

Processes

From understanding documents to understanding business context..

Why it matters

Most AI initiatives focus on reading and interpreting documents – extracting or classifying content. But in real life, **documents and data arrive at different times** and belong to specific steps of a workflow.

Process governance ensures AI understands *where* each piece of information fits, *how* it connects to other elements, and *when* it should be used.

It transforms AI from a passive "reader" into an active component of the business process.

Customer value

- **Operational consistency:** every AI system aligns with real workflows and business rules**.**
- **Process visibility:** organizations can monitor how data flows through different steps and decisions.
- **Automation reliability:** AI decisions are coherent with process timing and context, not just data content.
- **Faster scaling:** once the process logic is modeled, new use cases can be deployed with minimal redesign.

Neuro-Symbolic advantage

Generative AI can understand language, but it doesn't know the *logic* of the business process.

By combining **Generative AI** with **Symbolic reasoning**, we can connect data to process semantics – identifying dependencies, priorities, and roles.

The **neural side** interprets the unstructured content; **the symbolic side** manages how this information fits in the workflow.

This ensures decisions are **context-aware, traceable, and explainable** – a foundation for reliable enterprise automation.

FUSION AI LABS

expert.ai

Processes : multiple documents, one process

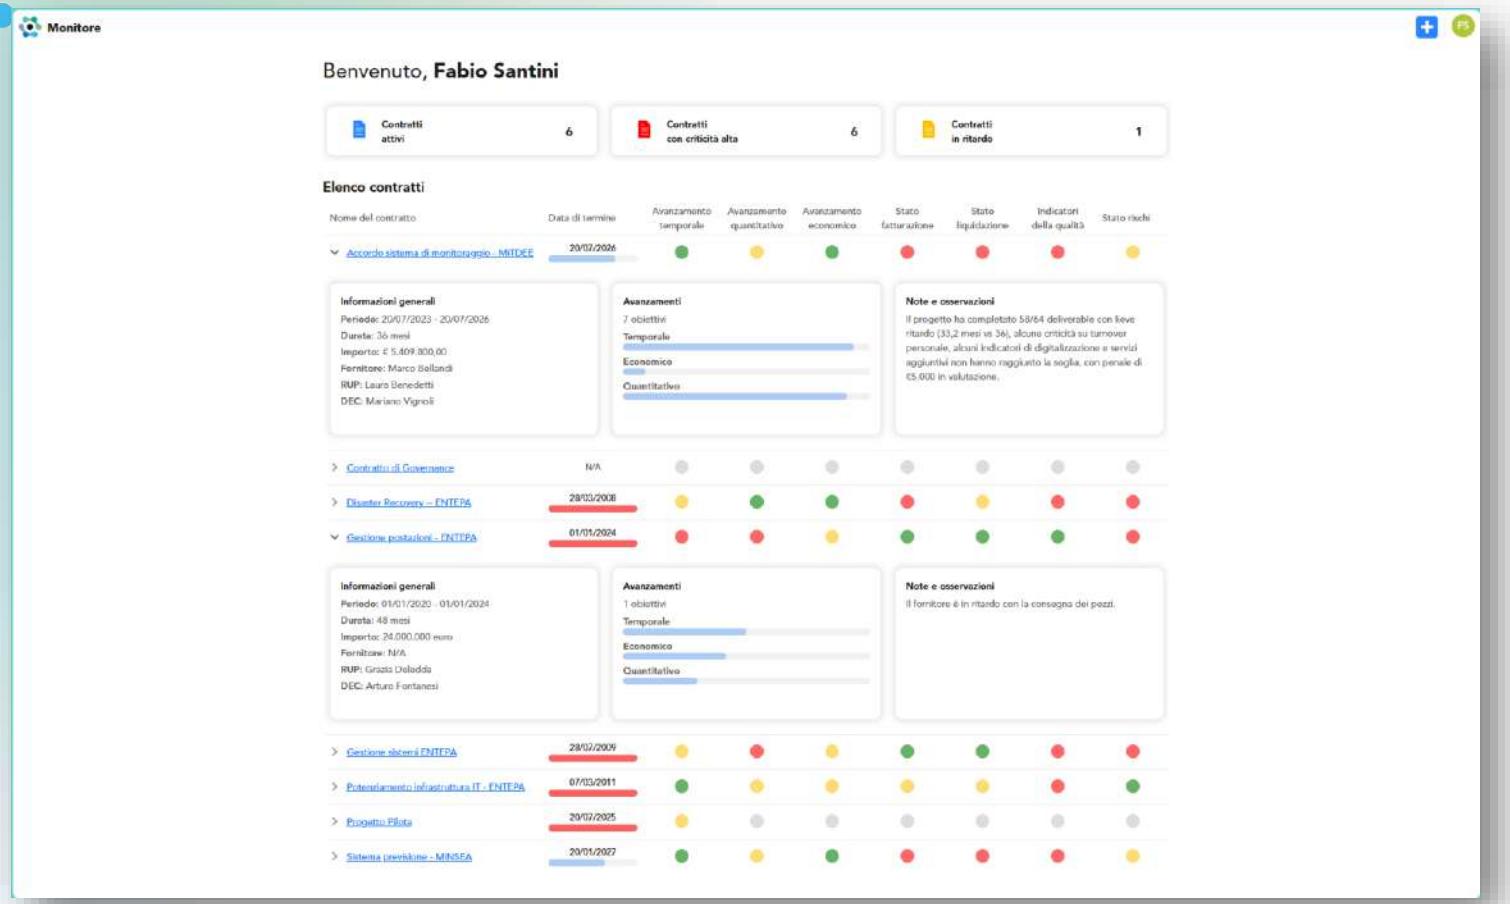

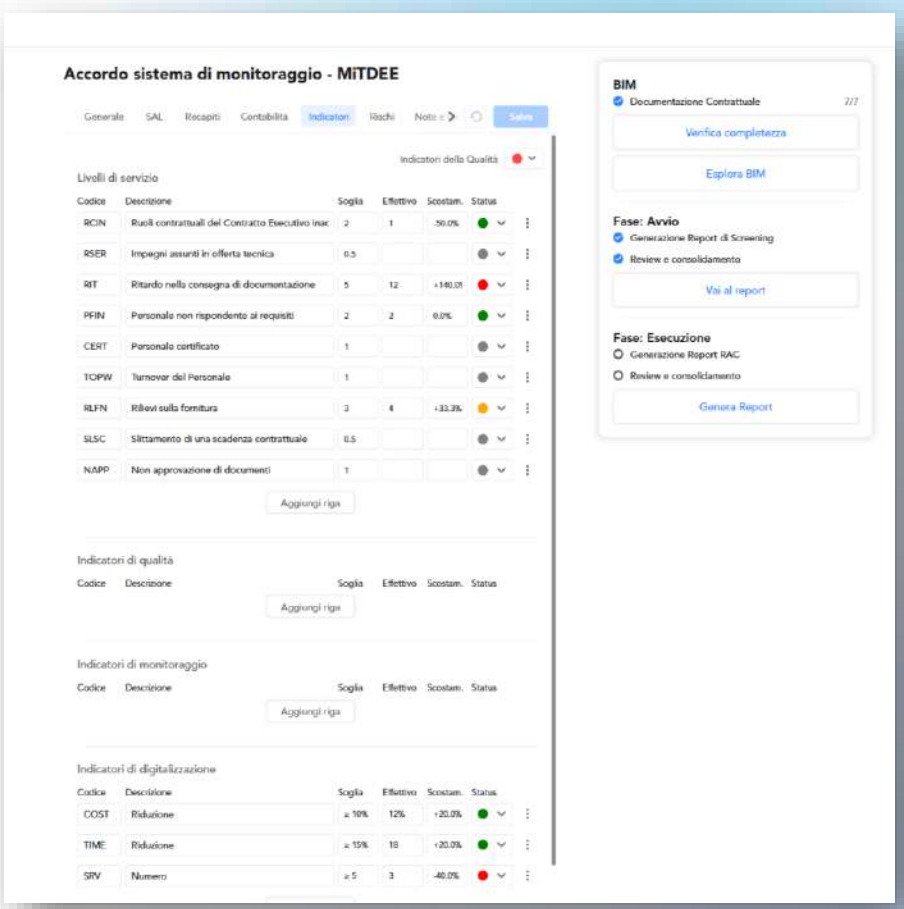

FUSION AI LABS

expert.ai

EI-HybridRAG

Results

From generating answers to guaranteeing reliability.

Why it matters

Generative AI is powerful but inherently unpredictable.
It doesn't reason – it estimates probabilities.
This means errors are not occasional; they are **inevitable**.

And the biggest risk is **unpredictability**: the system can perform flawlessly for days, then suddenly produce a wrong or biased result.

Result governance ensures that every AI output is verified, measured, and understood before it impacts a business decision.

Customer value

- **Trust and reliability:** outcomes can be validated and traced to their sources.
- **Risk reduction:** minimize hallucinations and unpredictable behaviors.
- **Performance transparency:** leaders can measure model accuracy and evolution over time.
- **Compliance readiness:** results can be audited and justified to regulators or stakeholders.

Neuro-Symbolic advantage

The **neural side** provides adaptability and scale – it can analyze vast amounts of data and generate meaningful responses.

The **symbolic side** adds reasoning and control – validating the logic behind every output.

Together, they create **explainable and trustworthy AI systems** that deliver consistent results.

This hybrid governance transforms AI from a *probabilistic generator* into a *predictable decision engine*.

FUSION AI LABS

expert.ai

Results : Confronting AI approaches

MODEL	MRR	Rec@1	Rec@5	Rec@10	Rec@20	Err@1	Ok@2-5	Ok@6-10	Ok@11-20	Err@20
Neural Retrieval	0,57	45,0%	69,0%	78,5%	86,5%	55,0%	24,0%	9,5%	8,0%	13,5%
Symb Retrieval	0,62	51,0%	78,0%	89,0%	93,0%	49,0%	27,0%	11,0%	4,0%	7,0%
Hybrid Retrieval	0,76	65,5%	90,0%	93,5%	97,0%	34,5%	24,5%	3,5%	3,5%	3,0%
Full Neuro-Symbolic Pipeline Retrieval	0,84	76,5%	93,5%	97,5%	98,5%	23,5%	17,0%	4,0%	1,0%	1,5%
Fine-tuned Neuro-Symbolic Pipeline Retrieval	0,90	85,0%	96,0%	98,5%	98,5%	15,0%	11,0%	2,5%	0,0%	1,5%

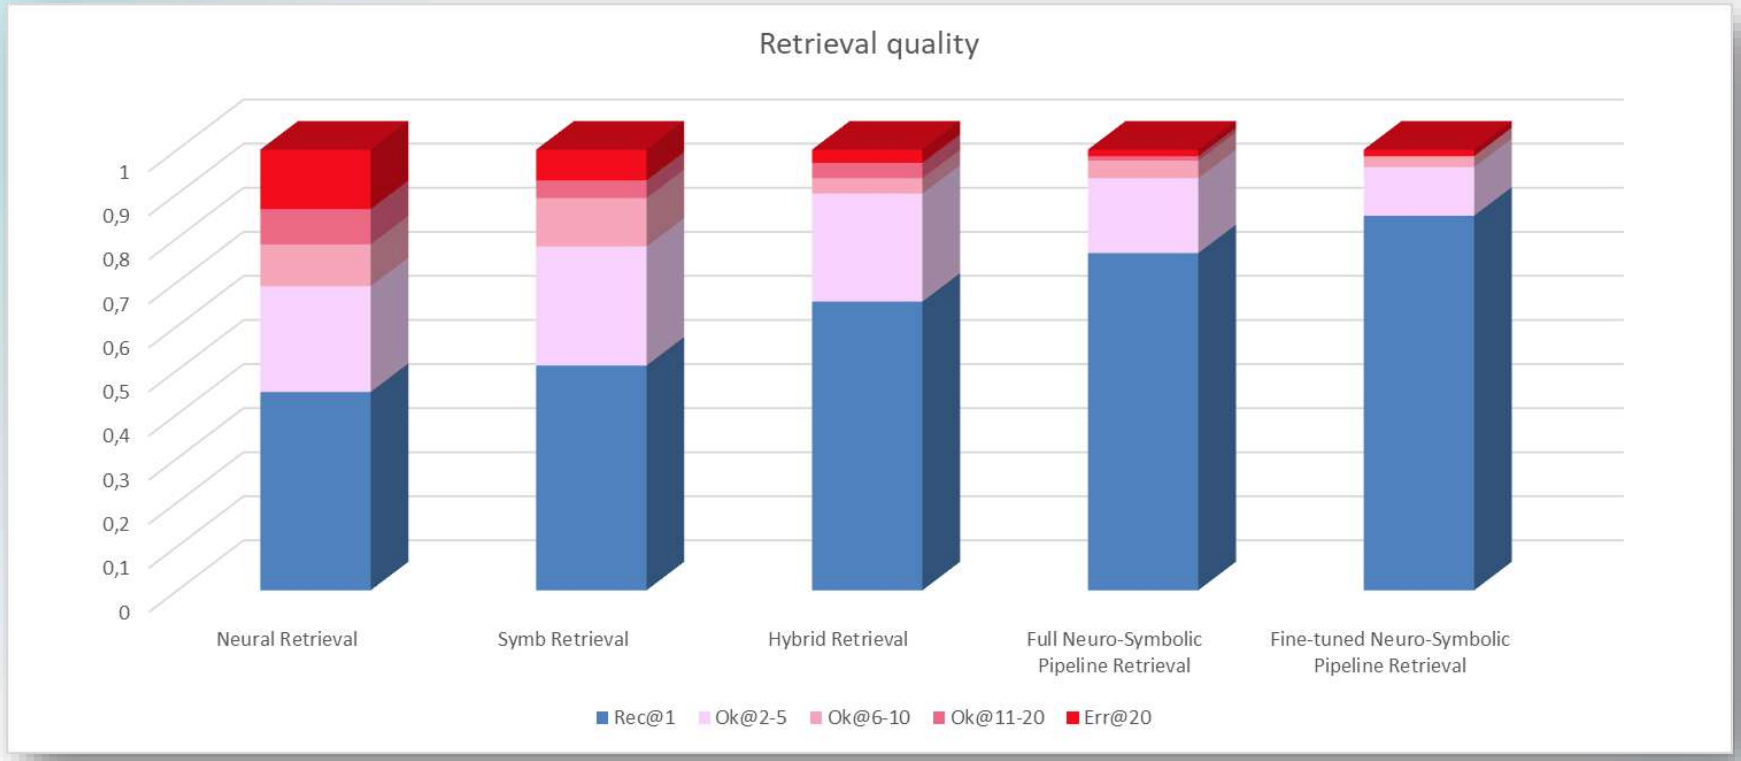

- **Recall** : If the system retrieves 10 documents and 9 of them contain the right information, recall is 90%. It measures coverage – how much of the relevant knowledge the system was able to find.

- **Top-k Accuracy (Ok@)**: If the right document appears first, we say Ok@1.If it's second or third, Ok@2–5, and so on. It measures how quickly the user would see the right answer.

FUSION AI LABS

expert.ai

EI-Flow

Explainability

From knowing the source to understanding the reasoning.

Why it matters

Most enterprise AI systems today, especially those using **RAG (Retrieval Augmented Generation)**, can tell you *where* an answer came from – which document or text chunk – but not *why* that answer was chosen.

This lack of reasoning transparency makes it difficult to trust or validate results, especially in regulated or high-impact environments.

Explainability governance ensures every AI decision can be traced, interpreted, and justified.

Customer value

- **Transparency**: users and auditors understand the path from question to answer.
- **Trust and accountability**: decisions are explainable and defensible.
- **Faster validation**: teams can quickly identify why results differ or fail.
- **Regulatory readiness**: explainable logic supports compliance and ethical standards.

Neuro-Symbolic advantage

Generative AI alone operates as a "black box," where correlations drive outputs, but reasoning remains hidden.

By combining **Generative AI** with **Semantic AI**, the neuro-symbolic approach builds a **hybrid reasoning layer.**

This hybrid creates a system that can not only answer but also **explain its reasoning** – showing *why* certain information was chosen and *how* it connects to the context.

The result: AI that is not just powerful, but **transparent, predictable, and trustworthy**.

FUSION AI LABS

expert.ai

Explainability : why and where

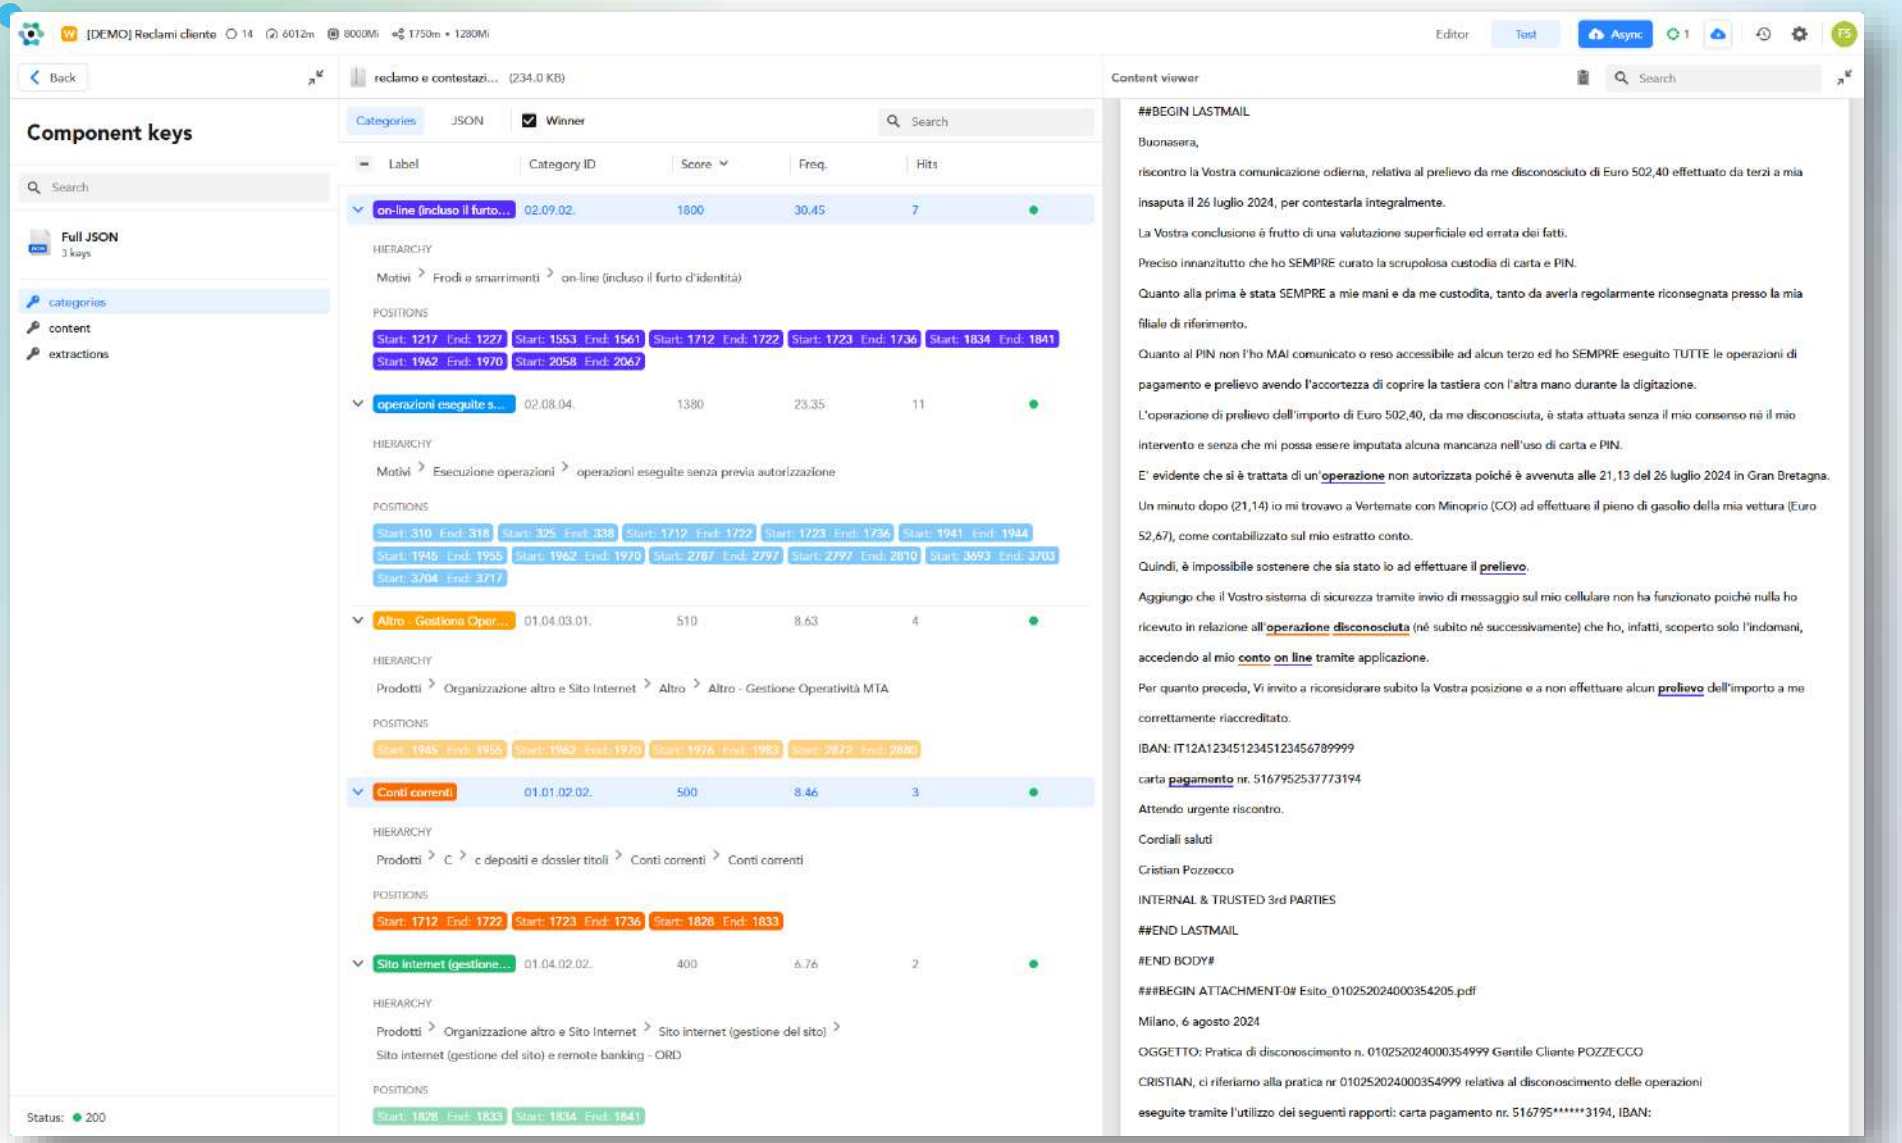

FUSION AI LABS

expert.ai

Explainability : why and where

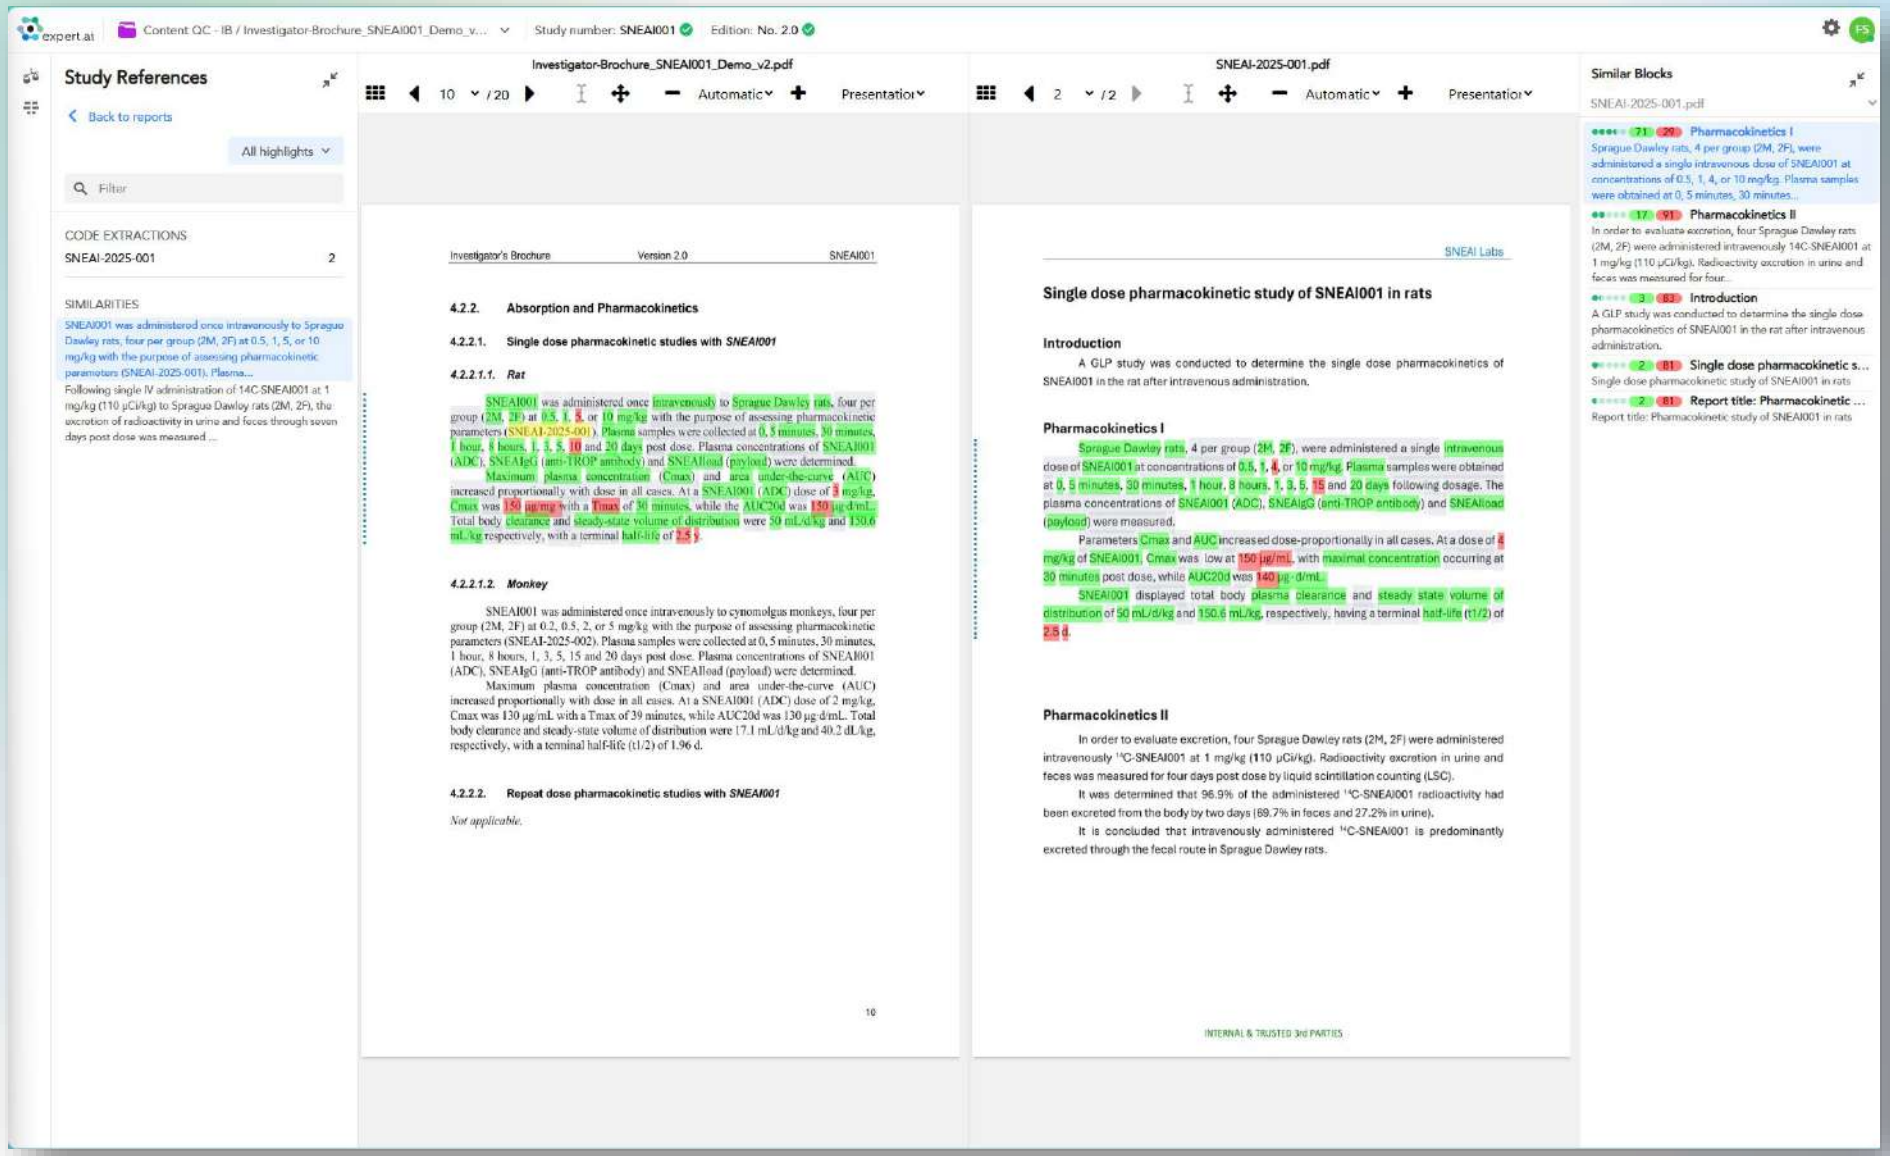

FUSION AI LABS

expert.ai

EI-Benchmark

Cost

Balancing innovation power with economic sustainability.

Why it matters

Generative AI introduces a **new and variable cost model** – every interaction has a price, usually calculated per token.

This makes expenses hard to predict, costs depend on the model used, the number of users, and even prompt length.

On the other hand, **Semantic AI** has a higher setup cost but minimal ongoing expenses.

Cost governance is the discipline of managing these trade-offs to ensure AI remains sustainable.

Customer value

- **Predictable budgets:** clear visibility on the cost per model, task, and usage pattern.
- **Efficiency:** use the right AI for the right purpose, minimizing waste.
- **Sustainability:** balance between innovation speed and long-term affordability.
- **Strategic control:** the ability to switch models or mix approaches without financial surprises.

Neuro-Symbolic advantage

The **neuro-symbolic approach** combines the **flexibility of Generative AI** with the **efficiency of Semantic AI**.
It allows organizations to dynamically choose:

- when to leverage **Generative AI** for creativity and exploration,
- and when to rely on **Semantic AI** for stability and low cost.

This hybrid strategy keeps AI economics **under governance**, ensuring enterprises can innovate fast **without losing financial control**.

FUSION AI LABS

expert.ai

Cost : reduce cost doing benchmarks

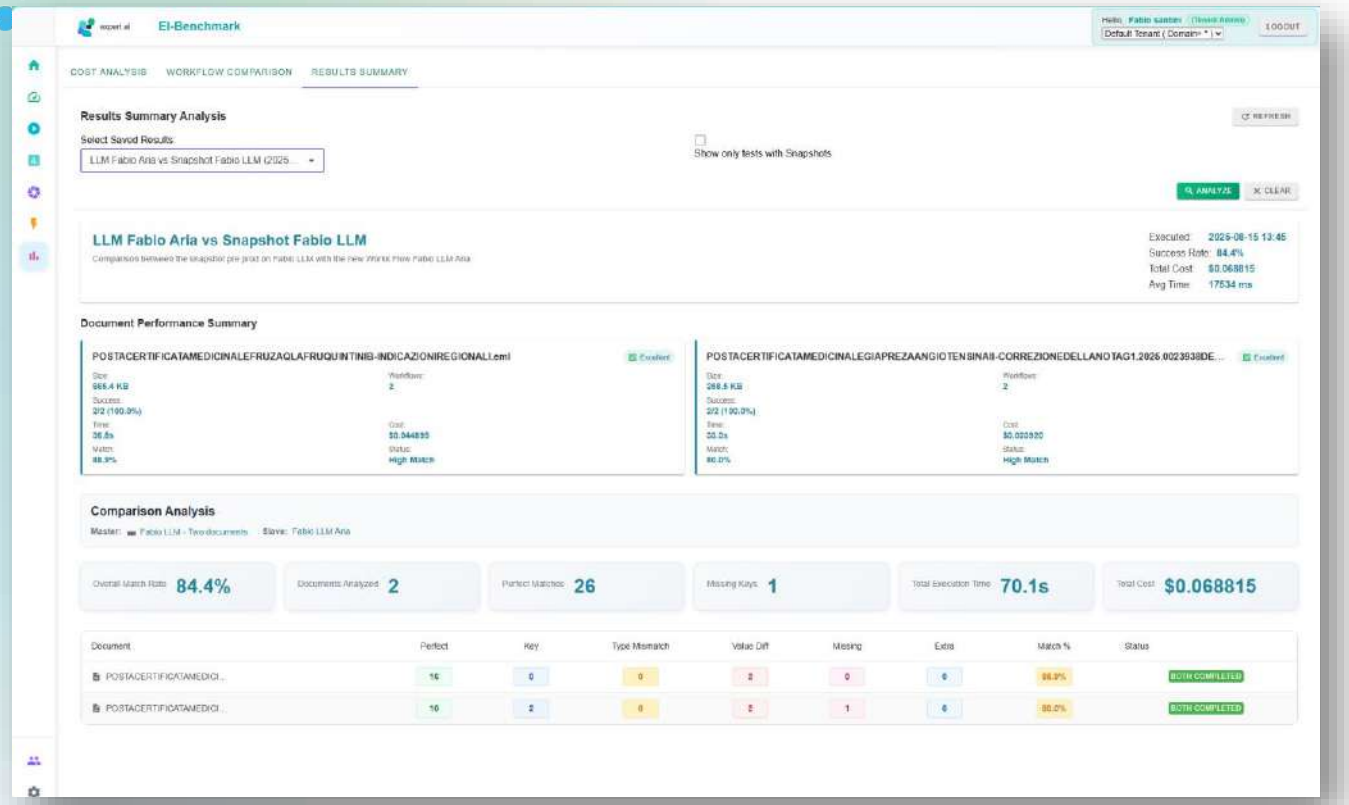

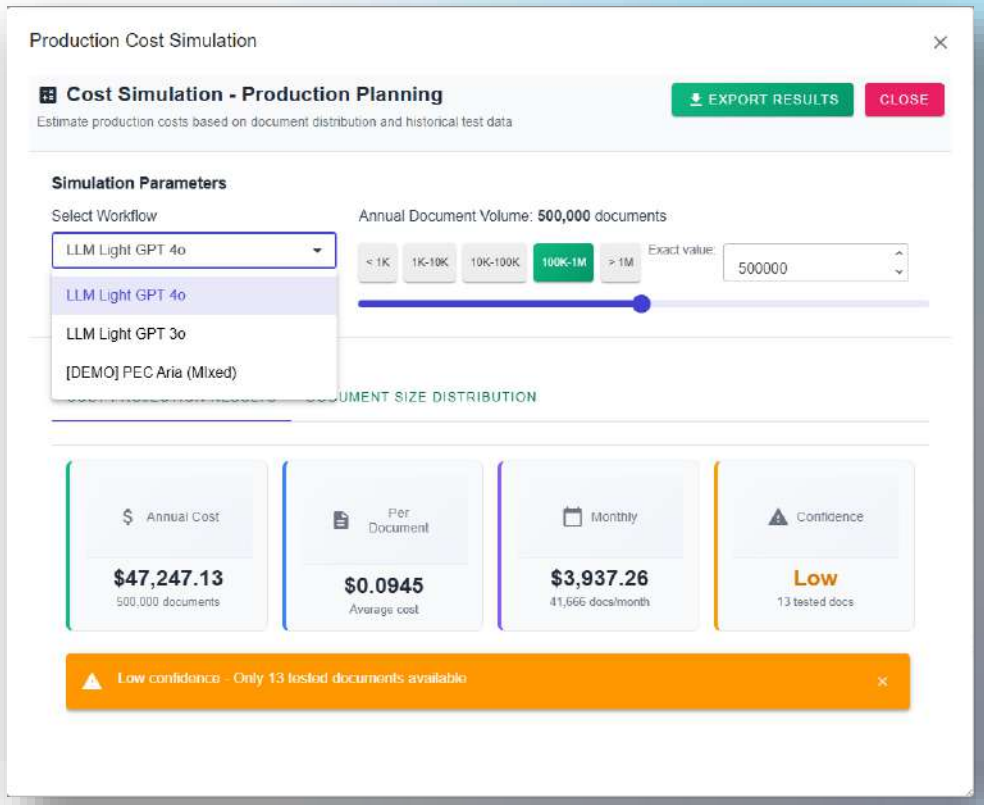

- **Compare** different workflows to understand which one perform at best
- Release updated version of workflow without **breaking the entire process**

- **Confront** different workflow to estimate the production cost
- Compare **workflow costs** with same results

FUSION AI LABS

expert.ai

EI-Benchmark

Knowledge

Turning scattered information into structured intelligence.

Why it matters

Enterprise knowledge is vast, unstructured, and often fragmented across documents, systems, and departments.

LLMs alone struggle to interpret such data especially when formats, vocabularies, and domain languages differ. Even fine-tuned models can't fully capture the **specific logic and terminology**.

Knowledge governance ensures that organizational know-how is structured, connected, and continuously updated, enabling AI to understand—**not just read**—the enterprise context.

Customer value

- **Consistency:** unify scattered content into a single, reliable knowledge layer.
- **Quality:** reduce ambiguity and misinterpretation of domain-specific terms.
- **Scalability:** once the domain is modeled, it can support multiple AI use cases.
- **Strategic asset:** transforms internal knowledge into a reusable, intelligent resource.

Neuro-Symbolic advantage

By combining **Generative AI** with **Semantic AI**, organizations can transform raw information into **governed knowledge**.

The **neural side** interprets and summarizes complex, unstructured content.

The **symbolic side** models the logic, entities, and relationships of the domain.

Together, they create **domain-specific knowledge graphs** that guide retrieval, disambiguate meanings, and contextualize answers.

FUSION AI LABS

expert.ai

Knowledge : add your specific domain data

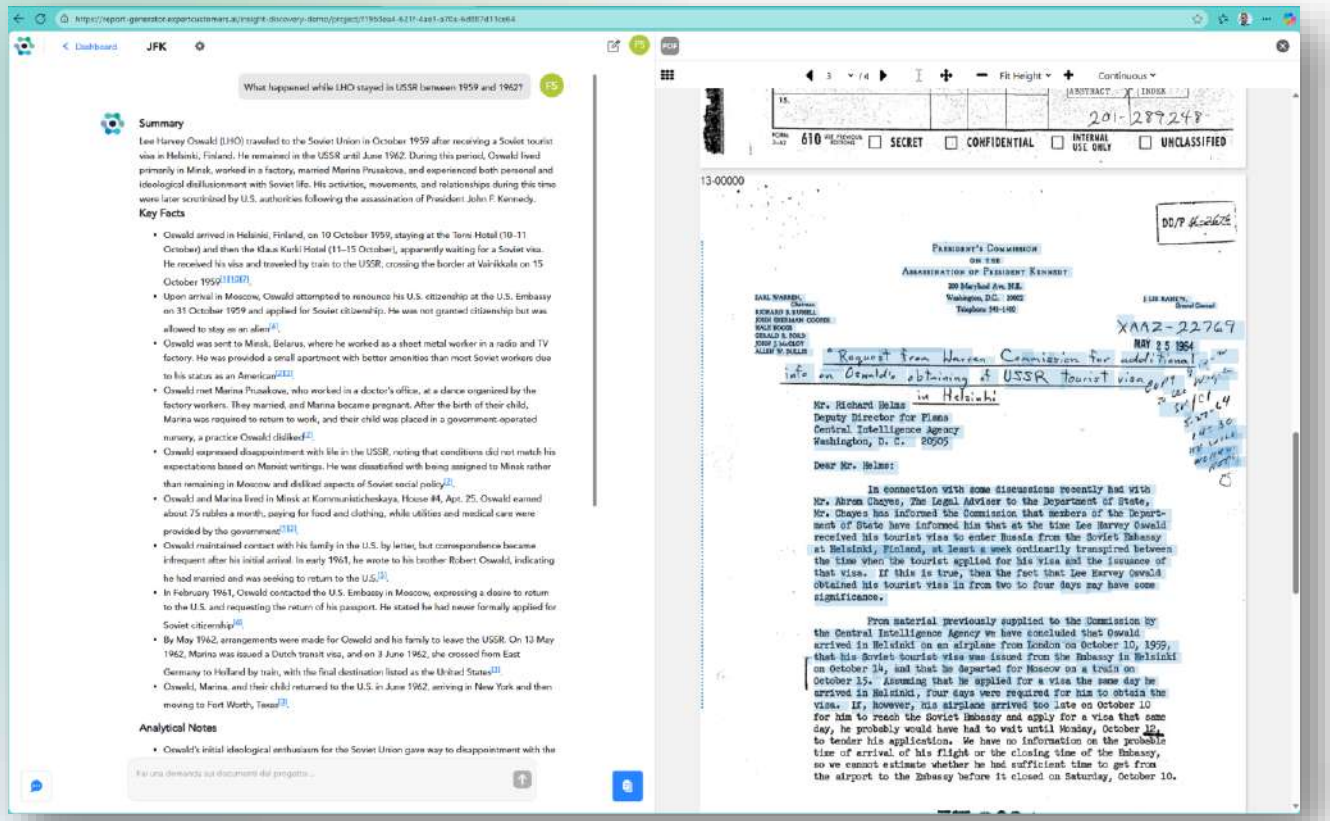

- Make a conversation with **complex documents** with different formats and size

- Every insight produced by the AI is fully traceable – users can click on any sentence of the answer and instantly **see the exact document passage that supports it.**

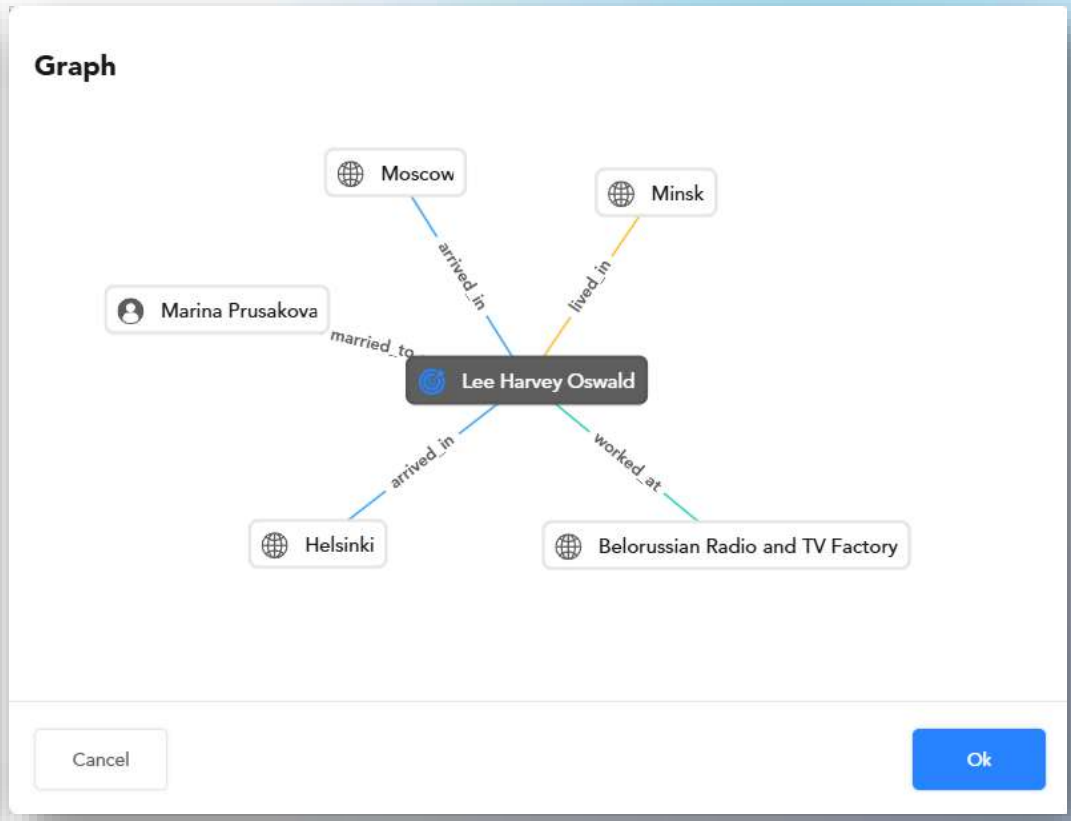

- **Understand** which part of the **specific domain knowledge graph** has been used to retrieve the right documents to generate a better answer.

FUSION AI LABS

expert.ai

Thank you

www.expert.ai

linkedin.com/company/expert-ai

marketing@expert.ai

FUSION AI LABS

expert.ai