AskRegula

AskRegula: Evidence-Grounded Regulatory QA with Integrated Benchmarking and Compliance-Oriented Evaluation

Dr. Tuba Gokhan · Prof. Ted Briscoe

LREC 2026 Industry Day
Palma de Mallorca, Spain
May 2026

جامعة محمد بن زايد للذكاء الاصطناعي
MOHAMED BIN ZAYED UNIVERSITY OF ARTIFICIAL INTELLIGENCE

The Regulatory AI Challenge

Regulation is a dynamic information problem

Regulated sectors face continuously evolving compliance ecosystems:

- large and frequently updated rules
- complex, cross-referenced documents
- high cost of interpretation and operationalisation
- increasing enforcement and non-compliance risk

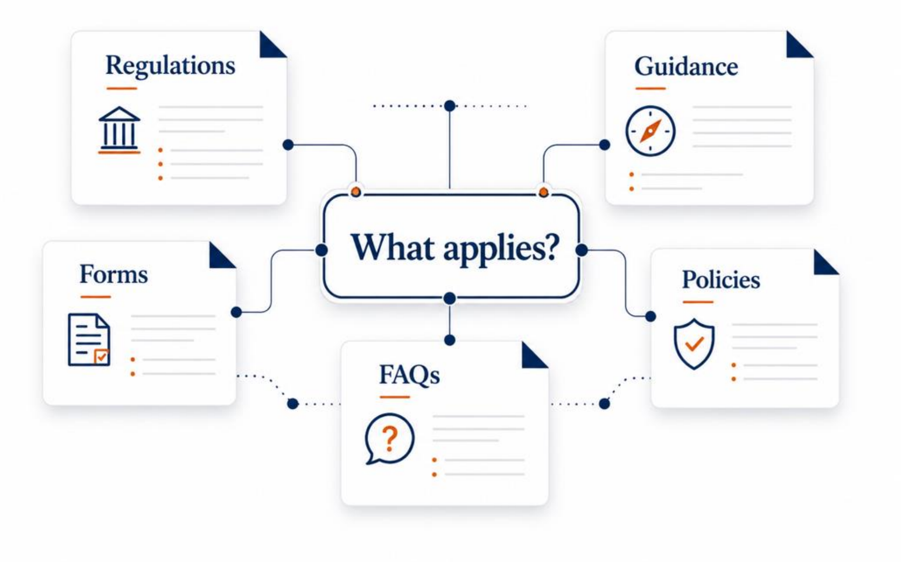

Why this matters

- **Regulatory change is accelerating:** AI-related regulatory developments doubled in 2025, with **157 major updates** recorded between June 2024 and May 2025.
- **Operational integration is slow: 74% of institutions** report that integrating one new regulation takes **more than one year**.
- **Compliance is expensive:** compliance costs can reach up to **13% of operating expenses** in financial services.
- **Non-compliance is costly:** global AML/KYC/sanctions penalties reached approximately **$3.8B in 2025**.
- **Regional enforcement is increasing:** UAE fines exceeded **380M AED** in the first eight months of 2025.

Fluent AI answers are not enough. Regulatory AI must be reliable, traceable, auditable, and evidence-grounded.

Why Regulatory QA is Hard

- **Regulatory answers are rarely contained in a single passage.**
- They may depend on:
  - **linked information** across the document structure
  - **cross-references** between provisions
  - **multiple evidence sources** combined in one answer
  - related elements such as **obligations, definitions, conditions, and exceptions**

*Regulatory QA requires more than retrieval. It requires **grounding, composition, and traceability**.*

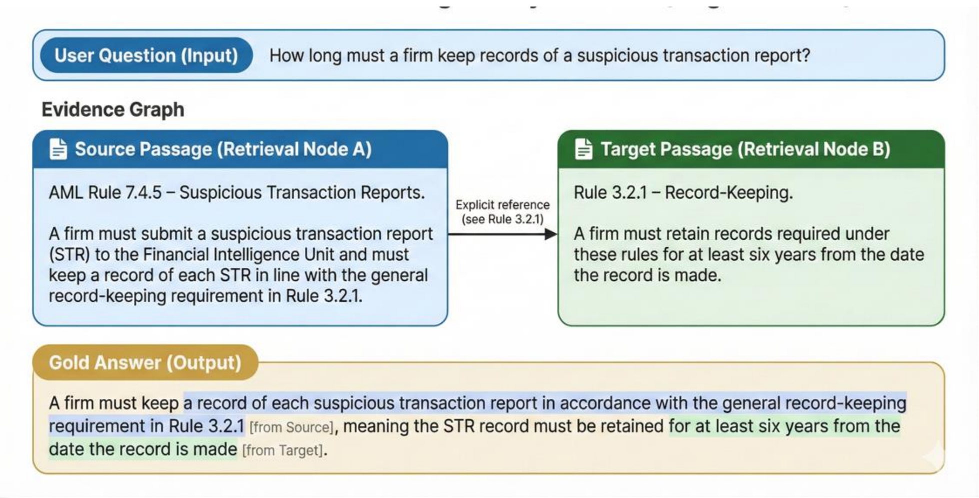

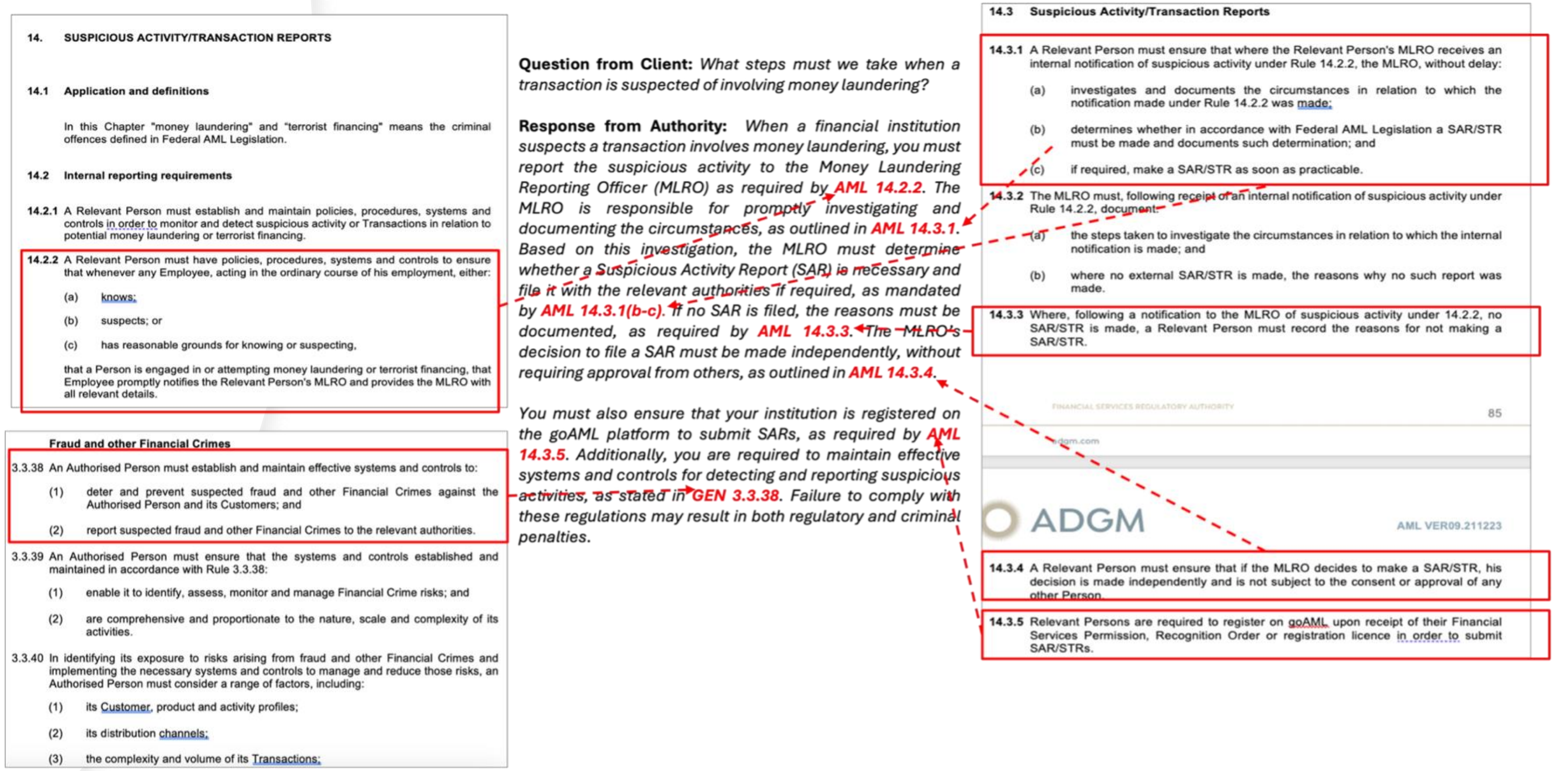

AskRegula From Manual Review to Governed Regulatory AI

Feature	Manual Review	Generative AI / LLMs	AskRegula
Evidence grounding	Expert source review	May be unsupported	Retrieved regulatory evidence
Traceability	Possible but manual	Often limited	Explicit source linkage
Evaluation	Manual and hard to scale	Generic evaluation	Compliance-oriented metrics
Governance	Process-dependent	Not governance-aware	Controlled feedback loop
Deployment control	Internal workflows	External services often used	Private / sovereign-compatible
Scalability	Limited	High	High with controls

AskRegula

More Than a RAG System

A Governance-Oriented Regulatory AI Framework

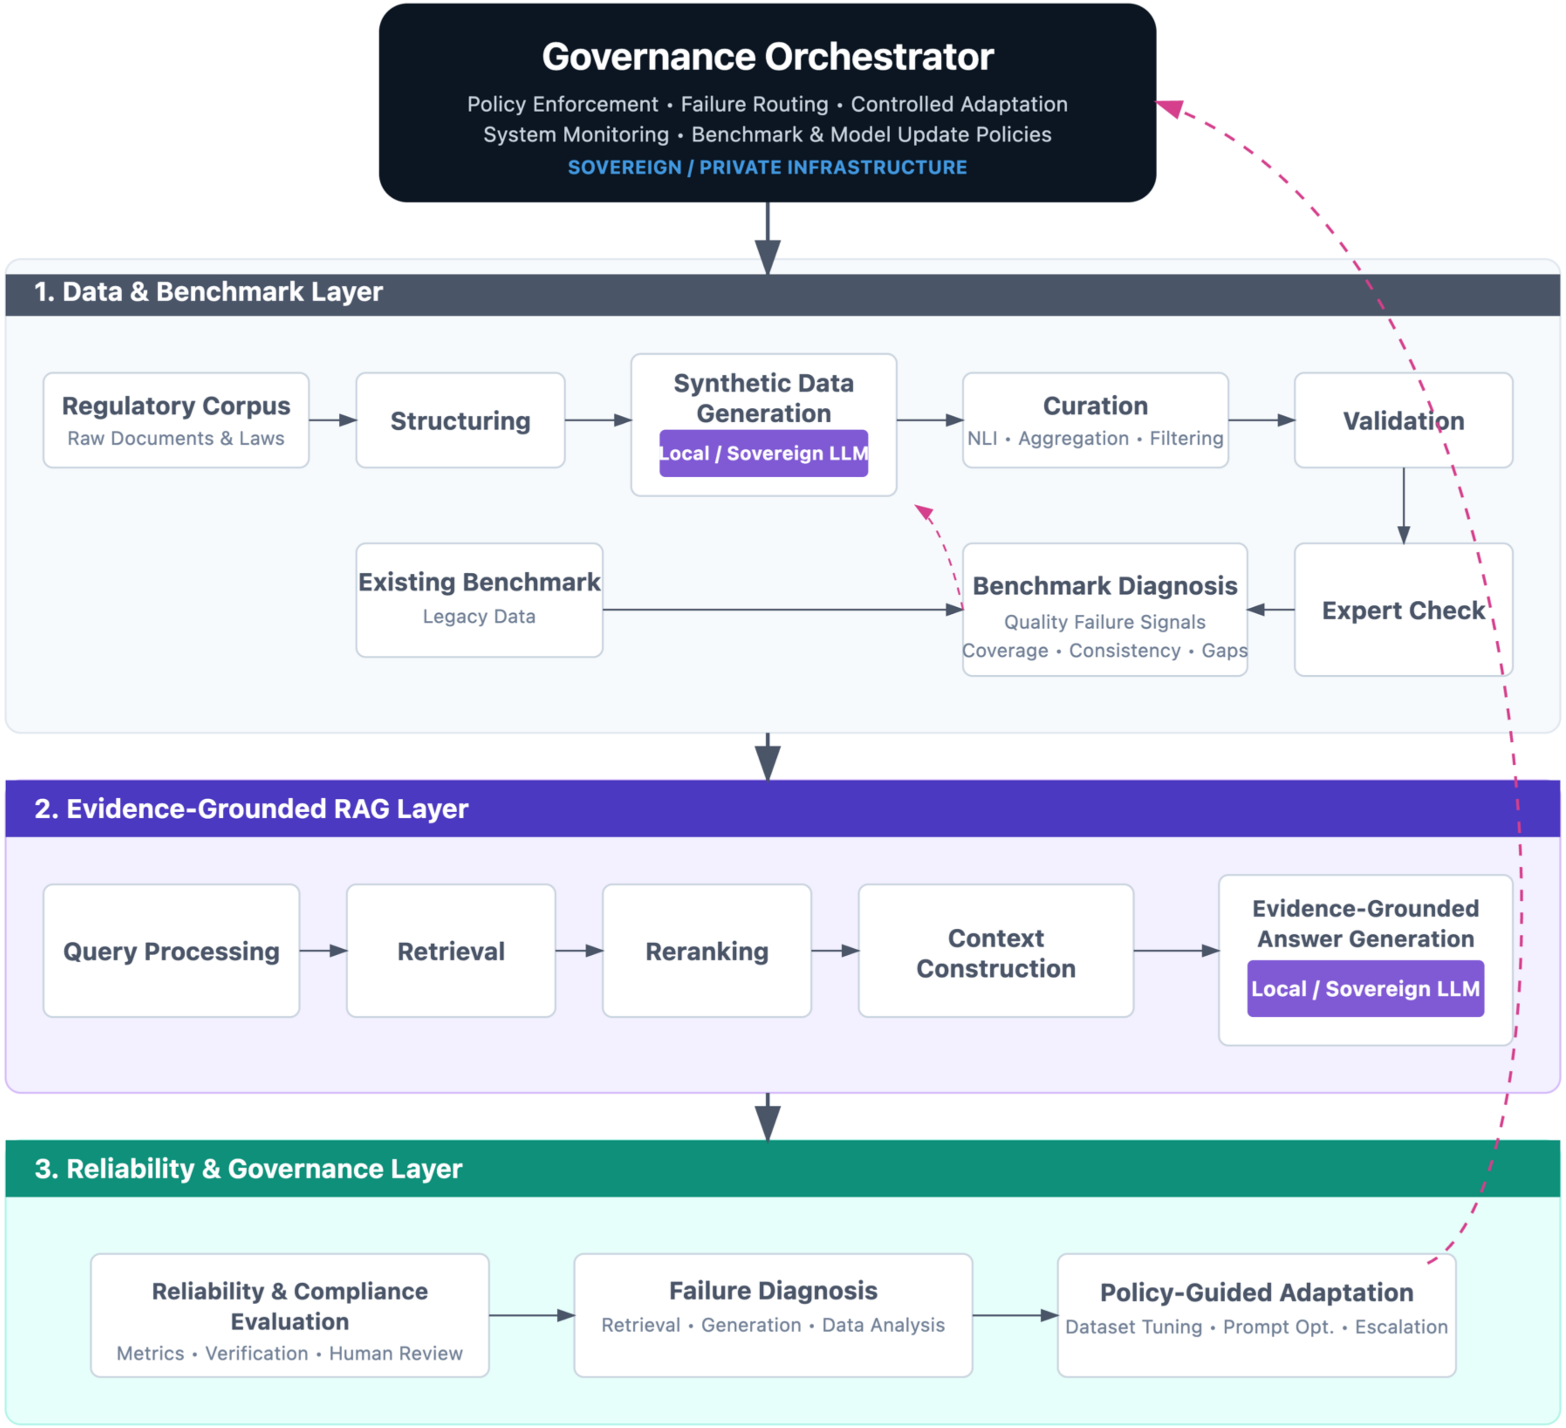

Research Foundations: Building Blocks for AskRegula

AskRegula Layer	Existing Foundation	Role
Dataset Layer	**ObliQA, ObliQA-MP, ObliQA-XRef**	Benchmark generation for obligation-oriented, multi-passage, and citation-dependent regulatory QA
RAG Layer	**RIRAG Shared Task**	Community evaluation of regulatory retrieval and answer generation systems
Evaluation Layer	**RePASs, RePASs-L**	Compliance-oriented evaluation of grounding, obligation coverage, contradiction, and linguistic quality

Potential Applications

- Evidence-grounded QA for financial regulations
- Review of relevant compliance evidence across large document collections
- AI governance and policy compliance support
- Alignment of internal policies with regulatory obligations
- Cross-reference-aware regulatory search
- Benchmarking and evaluation of regulatory AI systems

Expected Impact

- More reliable use of regulatory evidence
- Stronger traceability and auditability of answers
- Systematic evaluation and continuous refinement
- Shift from isolated prompting to governed workflows
- Better alignment with sovereign and private deployment settings
- Reusable foundation across sectors and jurisdictions

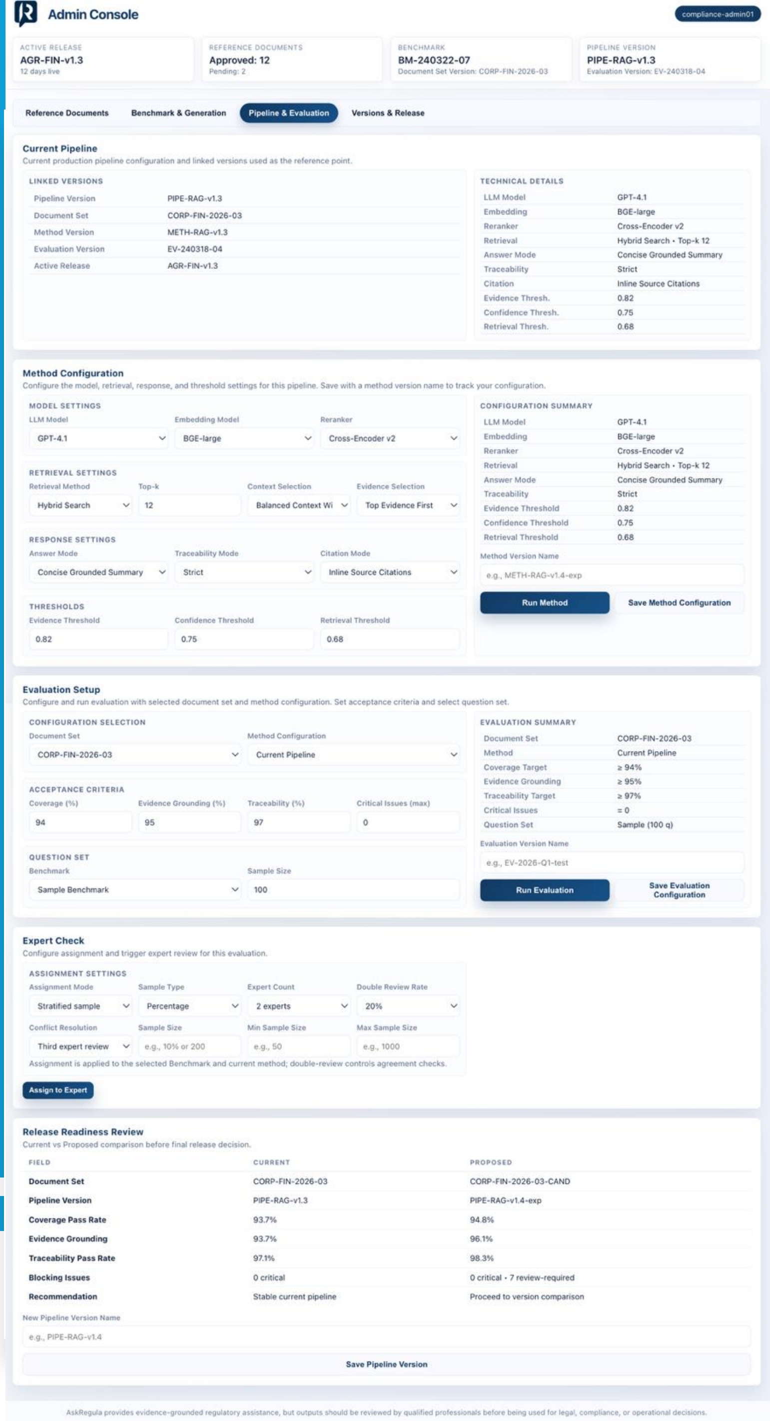

Prototype Interaction Model

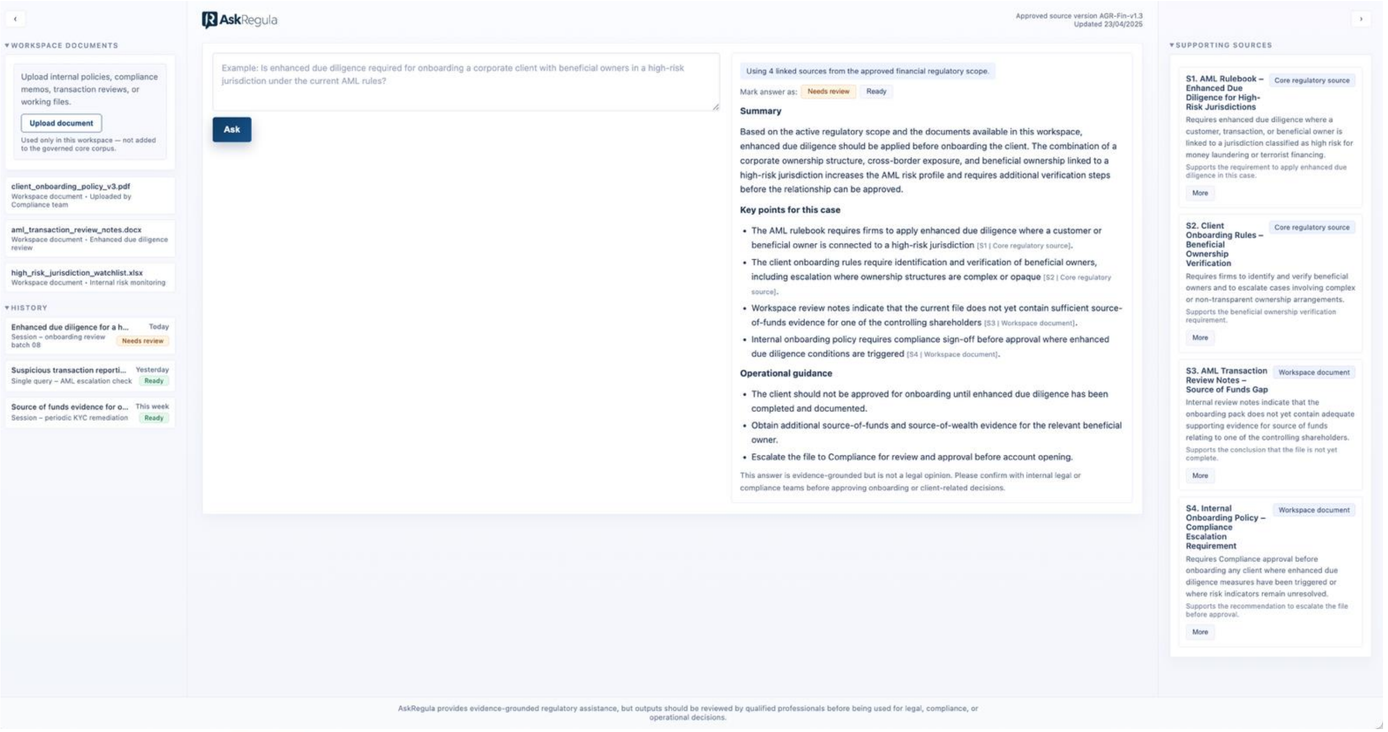

https://regnlp.github.io/AskREGULA/

RegNLP is broader than QA

- Today we focused on:
  - **Evidence-grounded, evaluated, and governed regulatory QA**
- Broader directions:
  - **Extraction · Structure · Cross-References · Grounded RAG · Summarization · Gap Analysis · Conflict Detection · Explainability**
- *Open to collaborations in regulatory AI and domain-specific NLP*

Contact | tuba.gokhan@mbzuai.ac.ae