




KRRA™
KNOWLEDGE RESILIENCE

AI Hallucination and Professional Accountability

Ensuring only authoritative, validated knowledge is used for consequential decisions and professional work.

THE CHALLENGE

AI systems can generate fluent, confident responses that include incorrect, incomplete, or fabricated information—especially citations and references. When used in professional or consequential contexts, this can lead to legal, financial, reputational, and safety risks.



WHY IT MATTERS NOW



AI can sound confident but be wrong.
Fluency is not proof of accuracy or authority.



Professional liability is real.
Courts and regulators hold humans responsible for AI-generated errors.



Authority cannot be assumed.
Citations, sources, and claims must be verified, not just presented.



Trust depends on verification.
Organizations and the public lose trust when AI errors are repeated.

WHAT CAN GO WRONG WITHOUT KRRA

- ⊗ Fabricated or incorrect citations and references
- ⊗ Outdated or superseded information used
- ⊗ Misleading or incomplete responses
- ⊗ Unverified claims influencing key decisions
- ⊗ Legal sanctions, fines, and malpractice
- ⊗ Loss of professional credibility and trust

HOW KRRA™ HELPS: AUTHORITY, VALIDATION & HUMAN JUDGMENT

KRRA separates knowledge interaction from knowledge authority. Authority comes from validated, governed sources and rigorous verification—not from the AI model. Human judgment remains the final safeguard.

KRRA™ TRUSTED KNOWLEDGE USE WORKFLOW



WHAT KRRA™ ENSURES

- ✓ Only authoritative, validated sources are used
- ✓ All claims are verifiable with traceable provenance
- ✓ Confidence levels reflect validation results
- ✓ Human accountability is preserved
- ✓ Decisions are based on trustworthy knowledge
- ✓ Errors are detected and corrected continuously

KEY PRINCIPLE

Availability ≠ Authority
A response can be available, fluent, and convincing without being authoritative, accurate, or trustworthy.

KRRA™ CAPABILITIES THAT MAKE IT POSSIBLE



Knowledge Authority
Ensures knowledge comes from validated, governed, and trusted sources.

Ch. 9



Validation & Verification
Verifies knowledge, models, and outputs before they influence decisions.

Ch. 11



Source, Model & Interaction Provenance
Provides end-to-end traceability of where knowledge came from and how it was used.

Ch. 14



Human Oversight & Judgment
Keeps humans in the loop for consequential decisions and override.


Ch. 9, 14



Governance & Accountability
Defines roles, policies, and responsibilities for AI systems and their outputs.

Ch. 12

ARCHITECTURAL PATTERN: VERIFIED KNOWLEDGE PIPELINE



REAL-WORLD EXAMPLES



Legal Research & Briefing

Lawyers use AI for research, but KRRA ensures every citation is verified against authoritative legal databases and human judgment is applied before filing.



Financial & Regulatory Reporting

Regulatory references, accounting rules, and risk models are validated and traceable before influencing reports or filings.



Healthcare Decision Support

Clinical guidelines, drug information, and patient data are validated from authoritative medical sources with human review.



Safety & Engineering

Engineering standards, procedures, and maintenance instructions are verified and current before operational decisions are made.

BUSINESS & SOCIETAL IMPACT



Professional Integrity & Trust

Reduces legal, financial, and reputational risk.



Better Decisions

Decisions are based on verified, authoritative knowledge, not guesswork.



Regulatory Compliance

Meets legal and industry requirements for accuracy and due diligence.



Operational Excellence

Fewer errors, rework, and delays from incorrect or unverified information.



Public Confidence

Strengthens trust in organizations, systems, and professionals.

KEY KRRA™ SECTIONS

Ch. 9.3 – 9.6	Knowledge Quality in AI Interactions
Ch. 9.7.1 – 9.7.3	Provenance & Authenticity
Ch. 9.7.7	Interaction Provenance
Ch. 11	Knowledge Validation & Verification
Ch. 12.2 – 12.3	Governance & Accountability
Ch. 14.2 – 14.5	AI Knowledge Architecture
Ch. 14.6 – 14.7	Autonomous Systems & Accountability
	Human Oversight & Judgment

Full details in the Knowledge Resilience Reference Architecture (KRRA™).



KRRA™
KNOWLEDGE RESILIENCE

KRRA™ provides the architectural foundation to design AI systems that are trustworthy, traceable, resilient, and accountable—so technology serves people, not the other way around.

Learn more at
[KNOWLEDGERESILIENCE.COM](https://www.knowledgeresilience.com)



KNOWLEDGE RESILIENCE
.com