Leveraging Classical AI with Generative AI, Exploring Possibilities: Special focus on BFSI space

Rajendra M Sonar, Ph.D.
Professor of Information Systems/Technology,
SJMSOM, IIT Bombay,
Co-Founder, WKA Technologies

Know Classical AI

Most of the early AI systems were expert systems powered by rule-based system and inference engines, dominated AI from 1970s till 1990s. An expert system is an AI program that emulates human expert decision-making in specialized domains using knowledge-based rules. These systems, which are highly effective for deterministic, complex, narrow, and heuristic-based problems, improve consistency and accessibility of expert knowledge.

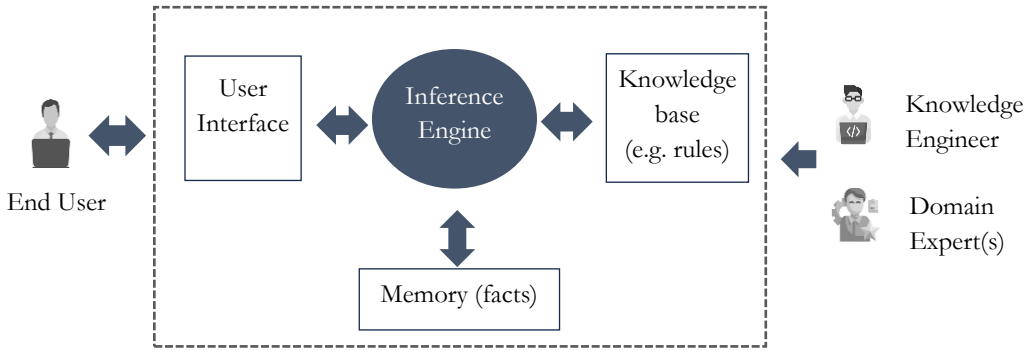

Simple expert system components

Expert systems were developed in many domains and verticals since 1970s. MYCIN was a pioneering 1970s Stanford University rule-based expert system designed to identify bacteria causing severe infections and recommend antibiotics. Using 600 production rules and backward chaining (goal driven problem solving), it achieved accuracy comparable to infectious disease specialists. It was notable for explaining its reasoning, though never used clinically. There were many applications developed in BFSI domain in 1980s till 2000s. Let us start exploring on loan underwriting domain as a reference use-case and specifically CLUES (Countrywide Loan Underwriting Expert System). CLUES was a pioneering rule-based artificial intelligence system launched in February 1993 and was designed to automate the mortgage underwriting process, increasing loan production capacity while maintaining consistency and reducing cost. It paved the way for modern automated underwriting systems (AUS) like Fannie Mae's Desktop Underwriter (DU) and Freddie Mac's Loan Product Advisor (LPA). By the mid-1990s, the system was processing over 8,500 loans monthly across more than 300 decentralized branches and was handling 35% of all Countrywide loan volume. CLUES contained approximately 1,000 business rules derived from human underwriting expertise and secondary market guidelines from more than 100 hours of interviews were conducted.

What AI CTOs did in 1990s

In 2001, Laura Doster, Mortgage Banking, freelance writer based in Springfield, Virginia interviewed the then CTO, Richard Jones, Countrywide Home Loans, Inc., Calabasas, California, an industry leader in technology

adoption. He was about the company's decision to build its own automated underwriting system (AUS) in 1991, the benefits and drawbacks of the "build" decision and what the future of automated underwriting (AU) development is going forward.

Some relevant excerpts from the interview to understand what kind of AI systems existed in 1990s to set the context: (source: https://www.thefreelibrary.com/Countrywide%27s+CTO.-a077828996)

Jones' answers illustrate how Countrywide's size and technology prowess continue to keep it at the top of the list of successful mortgage banks. In fact, it seems the company has had a clue about technology for quite a while, even before it built its groundbreaking CLUES.

We first decided to do this in 1991. At that time there was a refinance boom, kind of like what we're having right now. There was really nothing in existence at that time... no other automated underwriting system, and there was a shortage of underwriters. We were doing a huge amount of volume and we needed to automate. This is consistent with Countrywide's philosophy. All along, our idea was to use technology to gain a competitive advantage, and that was our thinking in 1991.

Well, we always do a "build" versus "buy" decision whenever we look at a technology project. That's exactly what we did in that case. We realized that we needed to have something that would have a very sophisticated decision-making capability. We looked at neural network technology. We looked at a number of different technologies, and decided on a combination of a rules engine and a statistical scorecard methodology. We also looked at decision tree type software. The problem with the neural networks and the decision tree type software was that it was too hard to explain how the answer was derived. It was like a black-box kind of situation. So while they could come out with interesting answers to the underwriting decision, it was very hard to account for exactly how the decision was arrived at.

The way we did it with GLUES was, we bought a rules engine software product to help us define and automate, but we also have built statistical scorecards to help us look at the history of certain loans and certain situations and the probability of default. Then we built a historical scorecard around all that information.

Another big advantage, we felt, was that it wouldn't be just a black box. We would know what was inside the black box. We would know all the elements of the risk decision that was being made and understand exactly how that decision process was arrived at by the software. As opposed to just throwing it to a black-box decision-maker and it comes back with an answer and you don't really know all the factors that went into making that decision.

Last year, the matter of black-box systems became a huge issue in the industry, particularly relative to credit scoring and also in relation to Fannie Mae's and Freddie Mac's automated underwriting systems. Do you look back and think, "We really saw something far ahead that's coming to be an issue right now"?

I think Countrywide demonstrated a lot of foresight when it developed its own automated underwriting system. I think we did a great job selecting a rules engine. So we've kind of taken this back-end engine and completely integrated it with our e-business environment. That's one of the reasons why we're the No. 1 funder of loans on the Internet--because of our ability to price online.

The rules were initially developed using the documented guidelines of CHL's subprime products, plus extensive interviews with subprime underwriting experts. The rules are being continually enhanced based on the performance of loans, marketing strategies, new risk analyses and so on. The enhancements are released monthly. Secondary Marketing's AI department develops and maintains CLUES in close coordination with the Production divisions and IT. Countrywide employs approximately 1,550 technology employees today.

Our servicing Web site, www.customers.countrywide.com, receives over 600,000 unique visitors per month. Over 400,000 borrowers can receive e-mail customer statements, and over 100,000 make online payments per

month. Our underwriting engine tells them upfront whether or not their borrower will qualify and, if so, what is required to substantiate the loan. Time isn't wasted on unnecessary efforts.

Rethinking on Classical AI Today

These excerpts from the interview give fair understanding what roles classical AI systems played. Just think over*, if such systems are to be developed today*, *how the approach would be*! Definitely, not the same it was in 1990s, especially with emergence of GenAI. Let us explore how such systems can be developed using GenAI in the current GenAI context.

Since domain experts are involved in providing their valuable expertise based on years of experience, expert systems take care of credibility and accuracy of knowledge using which it reasons and solves problems the way human beings solve; produce results in consistent and transparent manner, and explains the outcomes when asked. The results are equivalent of prescriptions given; verified and audited results by the professionals and practitioners. Knowledge engineering is one of the crucial tasks, it is the process of building, maintaining, and developing expert systems by eliciting knowledge from human experts and encoding it into a computer-readable knowledge base. It involves using techniques like rule-based systems, frames, and semantic networks to simulate human decision-making in specific domains. Inference engine is the brain of the expert system. It processes the information stored in the knowledge base to draw conclusions or make recommendations using reasoning strategies (like forward chaining or backward chaining) to analyze facts (data) and apply rules. One of the important distinctions of expert systems is knowledge base (typically rules) and inference engine is completely separate making it easier to change the knowledge base. During the problem-solving using goal driven reasoning, it asks only relevant facts based on goal to be achieved.

However, no technology in AI domain is perfect, there are limitations of classical Knowledge-Based Systems: 1> *do not have learning capabilities*: extracting expert knowledge is time-consuming, expensive and error-prone, experts often struggle to articulate tacit knowledge, knowledge needs validation, requires skilled knowledge engineers, 2> *knowledge engineering complexity*: rule explosion, inter-rule conflicts, maintenance overhead as policies evolve and 3> *static knowledge representation*: rules hard-coded, slow to build, versioned manually making them to respond slowly to market changes, difficult to scale across products or geographies, regulatory updates and new risk patterns. As markets, regulations, and borrower behavior evolved rapidly, these systems struggled to keep pace.

Case base reasoning systems reason by analogy, solve new problems by retrieving and adapting solutions from similar past cases. They mimic human reasoning, the way, we as humans solve the problem. CBR systems store past experiences as solved cases. For example, CBR based underwriting system will have all past loan proposals as cases. Since cases themselves represent implicit knowledge, the knowledge required to build CBR systems is much lesser compared to expert systems. The good part of modelling using CBR, each and every feature including categorical ones such as qualification, skills, designation, job industry and status, dependent details can be modelled meaningfully, are given weights depending upon the relevance use-case and domain context rather than assigning numeric values like machine learning algorithms do. Features like age, health are more important in insurance while ability to pay back is crucial for loan use-cases. Lot of domain knowledge (categories, ontologies, similarity measures etc.) is used to semantically match features of solved cases to a new problem case to retrieve relevant cases to be reused and derive solution for problem case. Models are easier to understand and explain. Together, ES and CBR form the backbone of knowledge-driven decision automation.

With the advent of GenAI, the tasks of modelling and engineering of domain knowledge have been simplified. There is really no need to go through huge documents, circulars, regulatory requirements manually; GenAI backed systems can extract and structure the knowledge the way domain experts and knowledge engineers and developers want them to suit systems that implement them. It means the knowledge engineering problem is sorted out to great extent. It also helps make systems deterministic.

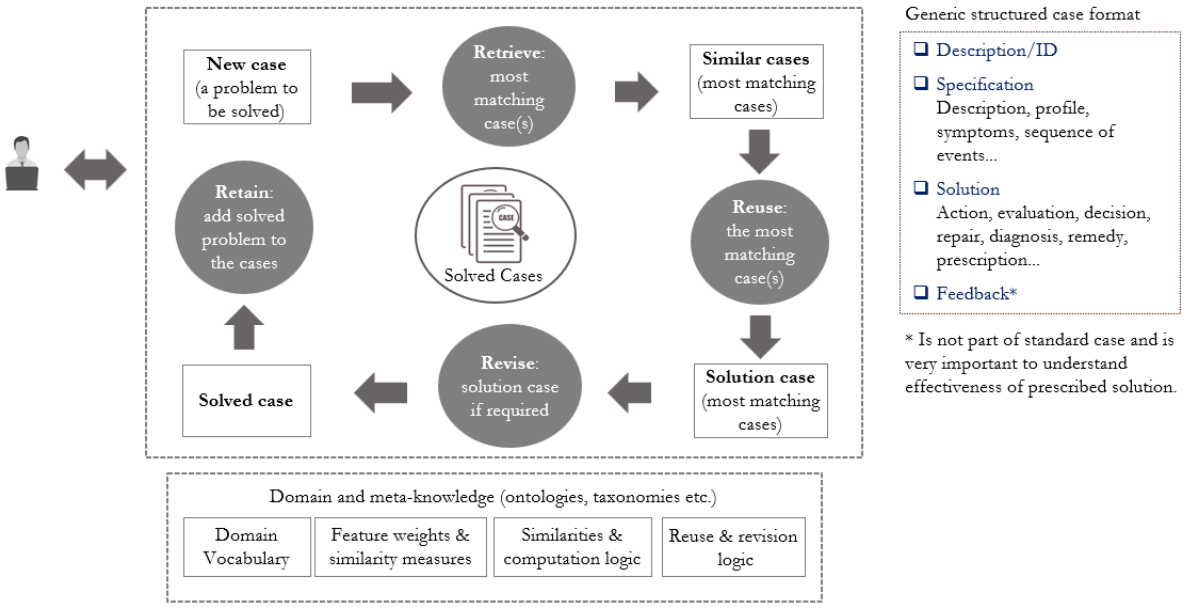

Solving problem using case-based reasoning

Why Classical AI Still Matters in the GenAI Era

Use cases like loan underwriting remains a domain where reasoning, justification, and accountability are at least as important as predictive power. Regulatory scrutiny, explainability mandates, and institutional risk appetite demand systems that can do more than generate plausible answers they must reason transparently and consistently.

This makes loan underwriting a natural candidate for revisiting classical AI approaches, particularly ES and CBR approaches that were foundational to early AI success stories and remain highly relevant in regulated decision-making contexts.

While recent advances in GenAI have transformed natural language understanding and content generation, classical AI paradigms such as ES and CBR continue to offer strengths that are critical in regulated domains like BFSI. Please note, the classical AI in this context, is referred to expert systems and case-based reasoning systems.

GenAI can complement, not replace knowledge-centric AI, enabling a new generation of underwriting systems that are explainable, credible, adaptive, and scalable. The good part is, GenAI would be extensively used (offline) for knowledge engineering while building systems than while problem solving time (run-time), thereby saving lot of tokens (in directly the cost) at the time of execution. This is like one time creation of enterprise level software with help of end users, system analysts and developers using GenAI backed platform. Which takes time, efforts and cost to create but subsequent cost of usage of software by many stakeholders is almost none. Updates can be easily done and managed using controlled vibe-coding.

Challenges with Pure GenAI

Generative AI has reignited interest in AI-driven decision-making by dramatically lowering the cost of language understanding, knowledge summarization and pattern abstraction. However, in underwriting contexts, GenAI introduces critical risks: 1> hallucination: confident but incorrect reasoning, 2> non-determinism: inconsistent outputs, 3> explainability gaps: narrative explanations without verifiable logic, and 4> regulatory discomfort: difficulty proving decision integrity.

As a result, pure GenAI underwriting systems remain unsuitable for core credit decisions without strong controls. Another challenge is, cost overhead as every conversation needs interaction with GenAI and exchange of tokens. Imagine use-cases, which on production need conversations with millions of customers, for example, making underwriting engine available online which tells borrowers upfront whether or not their proposal will qualify and, if so, what is required to substantiate the loan.

Using GenAI to Strengthen ES and CBR Systems

GenAI as a Knowledge Multiplier for ES and CBR

Rather than positioning GenAI as a decision maker, a more powerful paradigm is to use it as a knowledge structuring and acquisition layer. The real opportunity lies in repositioning GenAI as an enabler of classical AI, not its replacement.

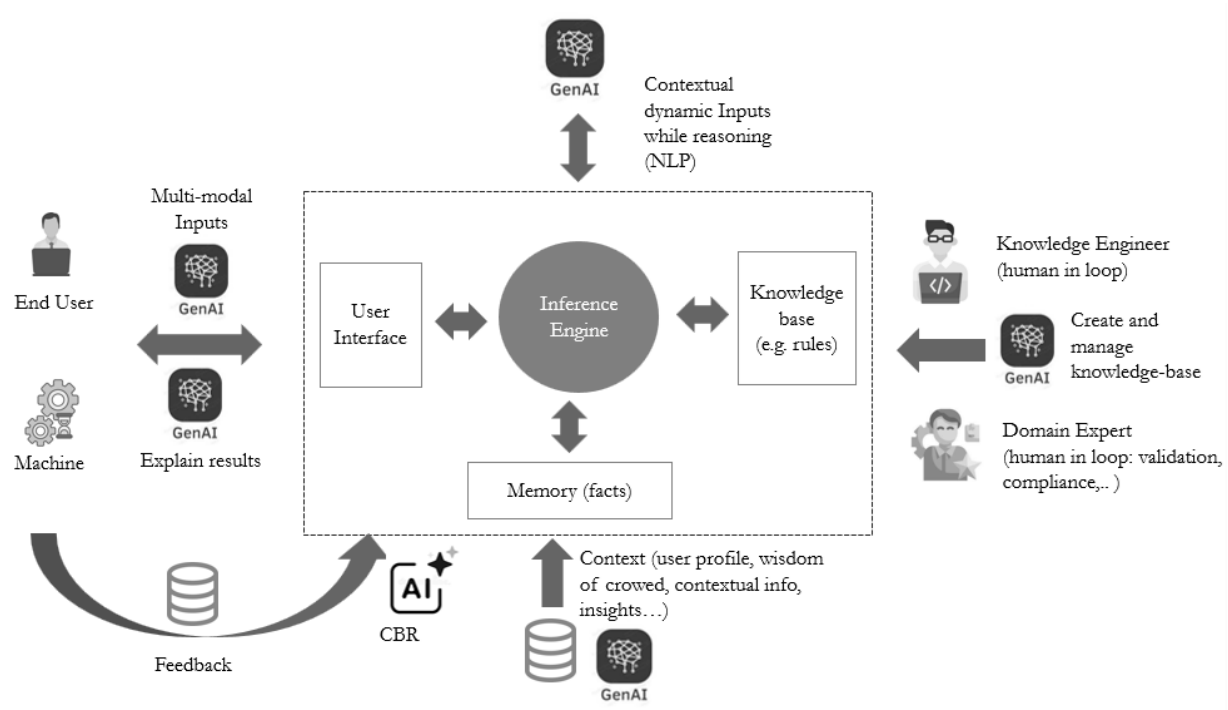

Using GenAI to strengthen Expert System

Accelerating Knowledge Acquisition

GenAI can ingest credit policy manuals, regulatory circulars, historical underwriting notes and convert them into draft rule sets, decision tables, case attributes for CBR systems. This directly takes care of historical bottleneck of knowledge engineering.

Structuring Tacit Expertise

Through guided prompts and scenario-based interactions, GenAI can help experts articulate implicit heuristics, exception patterns and risk trade-offs. These outputs can then be validated and formalized into ES/CBR knowledge bases.

Continuous Knowledge Refresh

Unlike static rule systems, GenAI can monitor policy changes, suggest knowledge updates, flag inconsistencies across rules and cases. Crucially, final authority remains with human experts, preserving governance.

Assisted Knowledge Engineering

Act as a copilot for domain experts, reduce dependence on specialized knowledge engineers

Reimagining Underwriting Systems Today with GenAI

If such systems were to be built today, GenAI would enable faster time-to-knowledge, weeks instead of months to model underwriting logic.

Hybrid Architecture: GenAI for knowledge ingestion, natural language interaction and ES + CBR for Core decision logic and deterministic execution.

Explainable by Design: decisions grounded in explicit rules and precedents, GenAI used to explain decisions in details and not invent them.

Scalable Human-in-the-Loop: experts validate and curate knowledge, GenAI amplifies expert productivity.

LLM-Based Knowledge Engineering Reduces Runtime LLM Dependence

What LLMs are very good at (Offline)

LLMs excel at:

- Reading large policy documents
- Extracting rules, conditions, exceptions
- Proposing similarity attributes
- Drafting adaptation logic
- Generating structured representations (rules, cases, vocabularies)

This directly maps to CBR knowledge containers:

- Vocabulary
- Case base
- Similarity measures
- Adaptation knowledge

Once this knowledge is materialized into structured form, it becomes executable without LLMs, that is the crucial cost-saving pivot. However, whenever there is change knowledge needs revision.

What Happens After Knowledge Is Structured

Once rules and cases are: validated by experts, stored in databases / rule engines, versioned and governed then at runtime:

- Expert System - evaluates rules (cheap, deterministic)
- CBR engine - retrieves similar cases (cheap, index-based)
- Similarity - vector / symbolic distance (very cheap)
- Adaptation - rule-based or parametric (cheap)

No tokens consumed, no hallucination risk, no RAG latency and no prompt engineering

Conceptual Shift

Instead of *Use LLMs to decide loan*, the architecture becomes *Use LLMs to build and maintain the knowledge that decides loans.*

This shift aligns strongly with recent research on CBR-LLM hybrids, which explicitly identifies LLMs as effective knowledge engineers rather than reliable runtime decision agents.

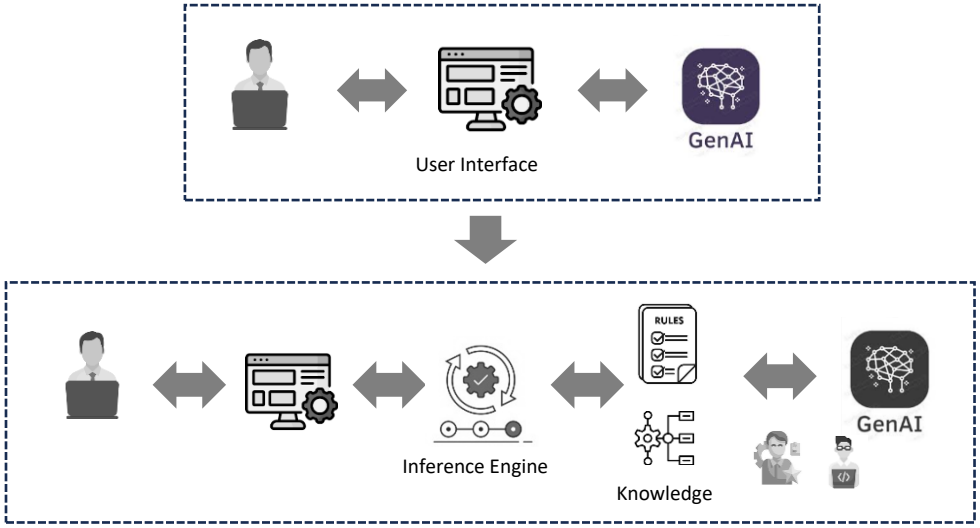

Using Gen AI to create knowledge

Building the Same Systems Today: A Hybrid Blueprint

If underwriting expert systems were designed today, the architecture would look markedly different.

GenAI Layer	Knowledge ingestion, natural language interfaces and expert copilot for rule and case creation
Classical AI Core (ES + CBR)	Deterministic decision logic, explainable inference, policy-compliant execution
Human-in-the-Loop Governance	Validation, Versioning, Auditability

This hybrid approach preserves trust while dramatically improving agility.

AI in Indian BFSI Is a Governance Problem First

In the Indian BFSI ecosystem, loan underwriting is not merely a technical decision problem, it is a regulated institutional process governed by prudential norms, supervisory oversight, and public accountability.

Institutions operate under:

- Reserve Bank of India supervisory frameworks
- Internal vigilance and audit mechanisms
- Parliamentary and CAG scrutiny (especially for PSU banks)

As a result, any AI-driven underwriting system must satisfy three non-negotiable requirements:

- Explainability (why was a loan approved or rejected?)
- Consistency (are similar borrowers treated similarly?)
- Defensibility (can the decision survive audit and legal review?)

These requirements strongly favor knowledge-centric AI paradigms, particularly ES and CBR, over opaque or probabilistic-only approaches.

Classical AI in Indian Banking: An Understated Legacy

Indian banks especially PSU banks and large NBFCs have a long but often under-acknowledged history of deploying rule-driven decision systems, even if not always labeled as "AI".

Rule-Based Underwriting as De-Facto Expert Systems

Historically, Indian underwriting systems embedded:

- Eligibility rules derived from RBI Master Directions
- Exposure limits and sectoral caps
- Credit score cut-offs and deviation matrices
- Delegation-of-power (DoP) logic

These systems effectively functioned as expert systems, with:

- Explicit policy rules
- Deterministic execution
- Manual override workflows

Their longevity is itself evidence of regulatory fitness.

Why ES and CBR Align Naturally with RBI Expectations

Explainability and Audit Readiness

RBI supervision whether through risk-based supervision (RBS), thematic inspections and IT and model audits; implicitly expects traceable decision logic. Expert systems and CBR offer case precedent justification and clear linkage to policy documents. This directly supports internal audit, concurrent audit and RBI inspection responses

Consistency and Fair Lending

Indian regulators are increasingly sensitive to bias, arbitrary discretion and inconsistent treatment across regions or branches. CBR systems, in particular, allow institutions to demonstrate applicants with similar profiles were treated consistently across time and geography. This is a powerful compliance narrative.

Knowledge Acquisition at Scale

Challenges included: large, heterogeneous branch networks, tacit underwriting knowledge concentrated in senior officers and frequent policy circulars and clarifications. Capturing this knowledge manually was slow and fragile. Indian BFSI regulation evolves rapidly: priority sector norms, MSME classification, changes, loan to value (LTV) and provisioning updates and digital lending guidelines.

Reimagining CLUES-Era Systems for Today's India

Earlier loan underwriting expert systems succeeded because they matched:

- Institutional risk culture
- Regulatory expectations
- Operational realities

If similar systems were built today, GenAI would enable:

- Faster knowledge onboarding across products and regions
- Easier adaptation to RBI changes
- Better documentation for audits

But the core decision logic would remain rule- and case-driven.

A Regulatory-First Hybrid Architecture

A future-ready Indian underwriting system would comprise:

- *GenAI Layer (assistive, not authoritative)***:** policy ingestion, knowledge structuring, natural language explanations
- *Classical AI Core (authoritative)***:** Expert system rules aligned to RBI norms, case-based reasoning for precedent handling and deterministic inference
- *Governance Layer***:** human approval, version control, audit logs and explainability reports

Such an architecture satisfies: RBI supervision, PSU bank accountability and NBFC agility

Responsible, Regulator-Ready AI for India

For Indian banks, NBFCs, and regulators, the goal isn't automating underwriting, but doing it responsibly. Classical AI provides the structural integrity and audit trails demanded by Indian regulatory regimes. GenAI accelerates knowledge engineering and evolution without undermining governance.

Secondary Benefits

- *Determinism and Auditability*: same inputs -> same outputs, full reasoning trace available and easier RBI and internal audits.
- *Reduced Vendor Lock-In*: core logic resides in rules and cases, LLM provider can be swapped or downgraded and no dependence on proprietary embeddings or prompts.
- *Lower Operational Risk*: no hallucinations in credit decisions, no runtime model drift and easier incident analysis and rollback
- *Regulatory Comfort*: LLMs are not deciding credit, humans approve knowledge and systems behave predictably

Strategic Implication for Indian BFSI

For Indian banks and NBFCs, the question is not *How do we use LLMs more?* but rather *How do we use LLMs where they reduce cost and risk instead of increasing it?*

Shifting LLMs upstream achieves exactly that:

- LLMs amplify expert productivity
- Classical AI delivers governed execution
- Institutions retain control over decisions

This is not anti-GenAI. It is GenAI used where it is strongest and avoided where it is weakest.

Use-Case Driven Guidance (What BFSI Leaders Actually Care About)

Use Rules / Expert Systems When	*Examples*
• Decision logic is explicitly defined • RBI / internal policy mandates exist • Decisions must be consistent across branches • Volume is high and margins are thin	• Eligibility checks • LTV, FOIR (Fixed Obligation to Income Ratio), DSCR (Debt Service Coverage Ratio) thresholds • Exposure caps and DoP (Delegation of power) rules • Mandatory rejections (blacklists, KYC failures)
Use Case-Based Reasoning (CBR) When	*Examples*

• Exceptions and nuances matter • Historical precedent is relevant • Human underwriters already think in analogies • Explainability via examples is valuable CBR is ideal where rules alone become brittle.	• MSME underwriting with informal income • Deviation handling • Risk-based pricing justification • Portfolio consistency checks
Use LLMs / GenAI When • The task is linguistic or exploratory • Determinism is not mandatory • The output is advisory, not authoritative • Cost per invocation is acceptable LLMs should not be primary credit decision-makers.	*Examples* • Policy summarization • Credit note drafting • Analyst copilots • Customer communication templates • Knowledge extraction from documents
Use Hybrid (Rules + CBR + LLM) When • The problem is regulated and complex • Decisions must be auditable • Knowledge evolves over time • Scale and cost matter	*Examples* • End-to-end loan underwriting • Credit decision explanation systems • Regulatory-compliant AI platforms • Enterprise underwriting modernization

Way Forward

The future of AI in BFSI underwriting is not a choice between classical AI and GenAI, but a thoughtful integration of both.

- Expert Systems and CBR provide trust, explainability, and control.
- GenAI provides scalability, adaptability, and knowledge acceleration.

Together, they enable Responsible AI for BFSI systems that are intelligent, transparent, and regulator-friendly. The real opportunity lies in leveraging GenAI to industrialize knowledge engineering, allowing classical AI systems to evolve continuously without losing their foundational strengths.

For Indian BFSI, the AI question is not *how powerful is the model?* but *how defensible is the decision?*

- Expert Systems and CBR provide institutional trust
- GenAI provides knowledge velocity
- Human governance provides regulatory legitimacy

Together, they form a Responsible AI paradigm that is uniquely suited to India's banking system. The future of underwriting AI in India lies not in bypassing regulation, but in engineering intelligence that works within it.

References:
https://www.thefreelibrary.com/Countrywide%27s+CTO.-a077828996
https://ojs.aaai.org/aimagazine/index.php/aimagazine/article/view/1123
https://aicompetence.org/why-expert-systems-still-matter/
https://hal.science/hal-05006761v1/file/main.pdf
https://chatgpt.com/