FUTURE OF AUDIT PROFESSION – NFRA'S FRAMEWORK FOR AI GOVERNANCE IN AUDIT

CA Milan Mody *CA Prashant Daftary*

Listen to me

Scan for
Audio Podcast Version*

* See Disclaimer on page 4

John Searle's famous 'Chinese Room' thought experiment reminds us that producing the right answer is not the same as understanding it. Searle imagined a person locked in a room who does not know a word of Chinese, but who follows a rulebook to match incoming Chinese characters with appropriate responses. To anyone outside, the room appears fluent in Chinese. Inside, there is no understanding at all, only symbol matching. Today's AI tools are, in essence, that room: fast, increasingly persuasive, and without genuine understanding. As AI becomes more embedded in statutory audits, National Financial Reporting Authority's (NFRA) recent guidance echoes the same principle; technology may process information, but it does not 'understand' business context, fraud motivation, management intent, or economic substance. It predicts the likely answer based on patterns. An auditor, however, must understand the substance behind a transaction, challenge inconsistencies, and exercise professional scepticism. Responsibility of judgement and opinion always rests with auditors (the human element remains central).

Statutory audit is going through a real shift in how technology is used. Data analytics, automated tools and techniques, and increasingly artificial intelligence including generative AI and agentic AI are moving beyond being efficiency aids and are becoming core to how risks are assessed, evidence is gathered, and conclusions are reached. This is no longer a futuristic concept but a reality. These tools help auditors test the entire population and are also useful in testing large & complex transactions (example – derivative transactions, ECL models etc). Use of such tools also exposes the auditors to risks that were not envisaged for example: automation bias, outputs that are opaque or non-deterministic, data privacy exposure and the possibility that use of such technology narrows professional scepticism instead of deepening it.

Take an example of digital lending company which uses AI based tools to evaluate and assess the creditworthiness of borrowers and disburse loans. The auditor of this company uses AI tools to validate the system considering a large database. This creates a situation where AI is auditing another AI and in between all this the human auditor needs to be in control and make the judgement as regards the true and fair view of the financial

results. Such situations are becoming increasingly common and challenging.

Consider how such an engagement actually plays out. The lender's credit model takes in application data, bureau scores and bank statements, and approves or rejects thousands of loans with no manual intervention. The auditor's own tool re-performs the scoring across the full loan book, compares actual disbursements with what the model should have produced, and flags exceptions. Both machines have done exactly what they were built to do. Human judgement is still required at three distinct points. First, whether the assumptions and data underlying both models are appropriate for the period under audit, particularly where lending policy changed during the year. Second, whether the flagged exceptions have actually been investigated and resolved, rather than merely listed in a working paper. Third, whether the combined result, taken together with everything else in the file, supports a true and fair view. None of these three can be delegated to either machine, and this is precisely the position NFRA's principles are designed to protect.

Against this backdrop, NFRA issued the first paper in its Staff Series on Technology in Audit "General Principles for Technology Adoption in Audit" in July 2026. It sets out a principles-based framework, along with specific do's and don'ts and the changes firms need to make to their audit quality systems, for how statutory auditors should evaluate, deploy and govern technology. It should be noted that none of this is new law and each principle restates an obligation that already exists under the Standards on Auditing, the Standards on Quality Control, or the Code of Ethics.

The ten principles are as follows:

Principle 1 — Technology neutrality, principles-based application. The Standards on Auditing have always been technology-neutral, SA 500, for instance, asks whether evidence is sufficient and appropriate, not how it was generated. These guidelines follow the same logic and focus on outcomes i.e. sufficiency of evidence, quality of risk assessment, integrity of the opinion. It does not mandate or prohibit any specific tool or vendor. Using a tool is never a substitute for meeting the obligations/requirement of the underlying standard.

Principle 2 — Non-delegable auditor responsibility. The audit opinion remains the auditor's. Technology can inform and speed up professional judgement, but it cannot replace it, and it cannot be cited to explain away a wrong conclusion. The report is signed by, and belongs to, an identified auditor.

Principle 3 — Audit evidence standards apply regardless of source. Output from an automated tool becomes audit evidence only once the auditor has evaluated its relevance and reliability the same way they would for any other source. Understanding where it came from, assessing the risk of manipulation or error, and testing accordingly are important steps which must be followed by the auditors. The standard does not distinguish between audit evidence based on whether it is manual or automated.

Principle 4 — Professional scepticism must be actively preserved. The auditors are expected to apply the same level of professional scepticism on the output derived from an AI tool. In today's world where deadlines are shrinking there is a significant risk that auditors would rely on an output or a conclusion derived by an AI tool without challenging the result or correlating the same

with other facts and using his or her own past experience. Human experience in no way can be replaced by an AI tool.

Principle 5 — Risk-based, proportionate governance. Validation and review should match the risk a tool poses. A translation aid does not need the scrutiny that an agentic tool selecting and testing samples does. Treating every tool the same either overburdens low-risk usage or under-scrutinises genuinely high-risk tools. The auditors would need to guard carefully against tools that are used in verification of audit data for example use of an audit tool to validate complex derivative transactions.

Principle 6 — Explainability commensurate with materiality. The extent of explainability and documentation with respect to any output generated using a tool, needs to be proportionate with how material the output is to the audit opinion. These principles are already mentioned in SA 500 on reliability and SA 230 on documentation. These standards require a level of understanding and documentation proportionate to the matter under consideration.

Principle 7 — Governance needs to precede deployment. A firm's quality control system should treat technology like any other resource whose risks are identified and addressed before the said resource is deployed. Tools need validation and approval under firm policy before use on a live engagement, and any material change including a model update, a new data source, a vendor change triggers re-validation. Validating a tool once is not a basis for assuming it stays reliable indefinitely. The firm's policies would have to include a process for periodic revalidation of AI tools including the risks it poses. There should also be a formal roll-out process which is formally documented in the audit manual.

Principle 8 — Data stewardship. Client and personal data processed through any tool must be protected consistently with the firm's confidentiality obligations, the Digital Personal Data Protection Act, 2023, and any other applicable legal requirements regardless of which tool is doing the processing. This would require formal documentation of such guidelines and process to ensure compliance.

Principle 9 — Transparency requirements. Where technology materially shapes the audit approach, the auditor must be able to explain that use. This follows from the existing obligation under SA 260 and SA 265 to explain the audit strategy and approach to those charged with governance.

Principle 10 — Continuous monitoring, not one-time approval. Approving a tool before deployment is only a starting point. Outputs can sometimes be non-deterministic, and the same input can produce a different result on a different occasion. The auditors also need to understand that with machine learning the technology moves or changes quickly and hence the auditors must keep monitoring a tool's performance through its working life and re-approve it when warranted. The audit manuals and guidelines will have to be updated to incorporate this principle.

NFRA has also listed specific do's and don'ts for using audit technology. These would go a long way in helping the firms in adopting technology in a more guided manner.

Do's and Don'ts for Using Audit Technology

Do's	*Don'ts*
Validate data before it is fed into any tool e.g., check completeness of the GL extract before running an AI tool.	Don't treat full-population testing by a tool as replacing the need for substantive risk assessment – 100% coverage of a dataset does not mean the right risks were tested for.
Test IT general controls at the source system e.g., access controls and change management over the ERP before relying on data pulled from it.	Don't use unvalidated spreadsheets, macros or scripts on live engagements without firm-level sign-off, even if a team member built it for convenience.
Document why a tool was chosen for a procedure, along with its version and configuration e.g., recording which version of a sampling tool was used and its parameters, so the work can be reconstructed later.	Don't accept a tool's "low risk" classification without periodically testing it e.g., re-checking a sample of items the tool tagged as low-risk to confirm the tagging still holds.
Evaluate service organisation controls for cloud or third-party tools e.g., reviewing a vendor's SOC 1/SOC 2 report before relying on a cloud-hosted analytics platform.	Don't rely on a vendor's claims about accuracy without independent evaluation by the firm.
Understand the logic behind a tool's risk-scoring, not just accept its flagged exceptions e.g., asking why a tool marked certain vendor payments as high-risk before following up.	Don't change a tool's configuration or data source mid-engagement without reassessing the impact on planned procedures.
Retain an audit trail that lets another experienced auditor or regulator, unconnected with the engagement, understand what the tool did.	Don't let efficiency gains from technology compress the time available for review and challenge.

NFRA also flags two points that sit outside the ten principles but matter equally. Data handled through any tool remains subject to the Digital Personal Data Protection Act, 2023, and the Information Technology Act, 2000, with particular focus on where a vendor's model or infrastructure is hosted outside India. And over-reliance on machine output is treated not merely as a quality lapse but as an ethics issue it can amount to a failure of the fundamental principles of professional competence, judgement and due care under the Code of Ethics.

Firms already using or piloting audit technology should map their existing tools against these ten principles, confirm that validation and re-approval processes are documented and ensure that documentation clearly shows the extent of human review behind any technology-assisted output. As data analytics and AI become more embedded in day-to-day audit work, this framework is likely to be the reference point NFRA inspections come back to.

This also gives a preview of what inspections may look like. It is reasonable to expect inspection teams to ask pointed questions on technology: which tools were used on the engagement and for what procedures; where is the record of firm-level approval and validation; who reviewed the output, and what did that review actually involve beyond a sign-off; and can the firm reproduce the result today using the version and configuration recorded in the file. A firm that cannot answer these questions from its audit file will find the conversation difficult, whatever the quality of the underlying audit work. Building the file with these questions in mind costs little during the engagement and saves a great deal afterwards.

Tips for practical implementation

- ✓ Review and revise the audit manual to include a specific section on use of AI tools
- ✓ Formulate SOPs as regards selection, testing, implementation, change management of such tools. This should include process as regards access rights, formal authority structure for deployment of tools.
- ✓ Reports, certifications which need to be maintained by the firm to ensure that the tools comply with the requirements and best practices
- ✓ Updating of the engagement letter, communication with TCWG as regards use of AI tools. Additionally, the audit firms would also have to map the requirements of the DPDP Act and formulate guidelines to ensure compliance with the same.
- ✓ Standard clauses which need to be included in the contracts with the vendors from whom such tools are procured.
- ✓ Training on use of AI tools to staff is going to be the need of the hour. These programs should also highlight the risk involved and precautions to be taken.
- ✓ Guidelines for retaining the trail and documentation of the results of AI tools. There could also be situations where the firms migrate from one tool to another and hence at such times, the audit trail needs to be carefully preserved.

One common risk which we all face is of unapproved or 'shadow' use of AI. Team members pasting client information into public chatbots to summarise an agreement, draft a note/document or explain an accounting treatment. This violates many principles mentioned above example: there is no validation, no approval, no record, and client data may have left the firm's environment altogether. It directly engages Principles 7 and 8 and, depending on the data involved, the DPDP Act as well. A clear internal policy on what is and is not permitted, backed by practical training and periodic checks, is the minimum step that needs to be taken, and it needs to cover articled staff as much as qualified staff/ members.

For small and medium sized firms, the framework can look daunting, and complex. A good starting point would be a simple inventory of everything already in use, including spreadsheet macros, scripts built by staff and any AI-based utilities. One of the partners/senior members can take the responsibility of technology management & governance. The initiative should include formulation of guidelines based on the principles mentioned in NFRA framework.

Daniel Kaheman, was a noble prize-winning Israeli- American psychologist best known for his work on psychology of judgement and decision making as well as behaviours economics, in his famous book 'Thinking, fast and slow' he explains that humans have:

System 1 – Fast, intuitive thinking

System 2 – Slow, analytical thinking.

Generative AI is incredibly fast (like system1). Professional scepticism resembles system 2 – it deliberately slows down to question assumptions and evaluate evidence. NFRA is effectively telling auditors: Don't let AI speed eliminate system 2 thinking.

Global scenario

NFRA's paper is part of a wider global move by audit regulators to catch up with AI adoption in the profession. The IAASB, after global roundtables through the second half of 2025 involving over 240 stakeholders, decided in February 2026 to develop non-authoritative guidance on applying its existing quality management and auditing standards to AI, rather than amend the standards themselves at this stage. The stake holders believed that existing International Standard on Quality Management (ISQM) 1 and International Standard on Auditing (ISA) 220 (Revised) provide a strong foundation for managing technology-related risks. At the same time, stakeholders noted that the IAASB could assist by developing additional practical guidance to help address challenges associated with emerging technologies.

The UK's Financial Reporting Council went further, issuing formal guidance in March 2026 for firms using generative and agentic AI in audit which is built around four pillars covering system design, certification, staff training and governance, and human-in-the-loop review, and addressing three risk categories: deficient output, misuse of output, and non-compliant methodology.

In the US, the PCAOB has so far limited itself to inspection focus and a research project on audit technology opened in May 2026, without issuing binding AI-specific rules, leaving firms to rely on existing standards in the interim.

Set against this backdrop, NFRA's principles place India broadly in step with, and on some points ahead of, where global audit regulators currently stand.

Concluding remarks

As the use of AI tools and agentic AI is becoming increasingly prevalent, the need for a more detailed implementation guide or FAQ or an auditing standard itself should be considered by the ICAI and NFRA. Specific guidelines for small and medium-size firms would be helpful in ensuring responsible adoption of AI in audit. The future of audit quality will depend not only on capabilities of AI, but also on the robustness of governance surrounding its use. Finally, "Let AI not cloud audit judgment".

References:

- NFRA's Staff Series paper (July 2026)
- Searle's original 1980 Chinese Room paper
- IAASB's February 2026 roundtable outcomes
- FRC's March 2026 Generative and Agentic AI Guidance
- PCAOB's Data and Technology research project (May 2026)

■