

DIGITAL CFO SOLUTIONS

MANAGING FINANCE IN THE AI ERA — TRAINING SERIES

MODULE 2

Understanding the AI Landscape

Technologies, Architecture and Enterprise Ecosystems

Machine Learning

Deep Learning

Generative AI

AI Agents

Responsible AI



About This Module

How this module is structured and how to get the most from it



Read

Work through each numbered section in sequence



Reflect

Pause at each Executive Insight and Knowledge Check



Apply

Relate each concept to your own technology landscape



Review

Use the Key Takeaways slide as a revision summary

MODULE FORMAT

This module contains four Knowledge Checks, five illustrative case studies, an extensive visual reference guide to core AI technologies (Annexure 2E), three quick-reference annexures (2A–2C), a capability-matching matrix (Annexure 2F), and a closing summary. Allow approximately 100–120 minutes for a first read-through.



What You Will Be Able to Do (1 of 2)

After studying this module, readers should be able to:

- Explain the evolution and fundamental concepts of Artificial Intelligence
- Distinguish between Artificial Intelligence, Machine Learning, Deep Learning, Generative AI, and AI Agents
- Understand how AI models learn from data and why data quality is critical
- Appreciate the infrastructure required to support enterprise AI initiatives



What You Will Be Able to Do (2 of 2)

After studying this module, readers should be able to:

- Identify the principal risks associated with AI adoption, including bias, hallucinations, cybersecurity, privacy, explainability, and regulatory compliance
- Evaluate different AI implementation approaches, including commercial AI platforms, customized models, and enterprise AI solutions
- Distinguish between currently available, emerging, and more speculative AI capabilities
- Apply governance principles to ensure responsible, secure, and ethical use of AI



What This Module Covers

Nine themes, from the evolution of AI to a practical implementation framework

 2.1 Introduction

 2.2 What is AI?

 2.3 Machine Learning

 2.4 Deep Learning

 2.5 Generative AI & LLMs

 2.6 The AI Ecosystem

 2.7 Architecture & Infrastructure

 2.8 Risks & Responsible AI

 2.9 Selecting & Implementing AI

 2.11 Five Case Studies

 Annexures 2A–2F

 Takeaways, FAQ & Summary



SECTION 2.1

AI as an Integrated Ecosystem for the Finance Leader

Understanding why AI is becoming essential, and how it evolved into its current form



Why Finance Leaders Need to Understand AI

AI now touches forecasting, reporting, treasury, risk, compliance, audit, tax and strategy

- AI has evolved from a specialized technology into an increasingly important enterprise capability
- Machine Learning, Deep Learning, Generative AI, LLMs, NLP, Computer Vision and AI Agents are often discussed together even though they perform different functions
- Finance leaders do not need to become data scientists, but need sufficient understanding to evaluate these technologies, assess business value, understand limitations, and govern their use responsibly
- Earlier systems relied on predefined rules; modern AI increasingly uses data-driven models capable of prediction, language processing, content generation and coordinated workflow execution

THE CENTRAL MESSAGE

Enterprise AI is not a single technology or product. It is an ecosystem of data, models, systems, infrastructure, governance and human responsibility.



The Evolution of Artificial Intelligence

Periods indicate when approaches became prominent, not precise start and end dates

Period	Major Development
1950s–1970s	Early AI and Symbolic Reasoning
1970s–1990s	Expert Systems and rule-based AI
1990s–2010s	Statistical Machine Learning
2010s	Deep Learning and large-scale neural networks
2020s	Foundation models and Generative AI
Late 2020s onwards	Multimodal and agentic AI with increasingly autonomous but bounded workflows

EXECUTIVE INSIGHT


Earlier systems relied heavily on predefined rules, while modern AI increasingly uses data-driven models capable of prediction, language processing, content generation, multimodal analysis, and coordinated workflow execution.



SECTION 2.2

What Is Artificial Intelligence?

From pattern recognition to content generation — and how AI differs from traditional automation



Core Characteristics of AI Systems

Individual AI systems may provide only some of these capabilities



Pattern Recognition

Identifies trends, relationships and anomalies



Prediction

Estimates future outcomes from data



Classification

Categorises transactions, documents, customers



Language Processing

Analyses and generates human language



Recommendation

Suggests possible actions from information



Content Generation

Produces text, images, code, other outputs



Learning from Data

Trained models refined as conditions change



AI versus Traditional Automation

Rule-based automation and AI are complementary, not competing

Traditional Automation	Artificial Intelligence
Primarily rule-based	Primarily model- and data-driven
Predetermined logic	Learns patterns from training data
Generally deterministic outputs	Often probabilistic outputs
Rules changed manually	Models may be retrained or updated
Well suited to repetitive activities	Well suited to analytical, pattern-driven, language-based tasks
Handles predefined exceptions	Can identify complex or previously unrecognised patterns

EXECUTIVE INSIGHT

A rules-based workflow may automatically process routine invoices, while AI analyses supplier history, payment behaviour and anomaly indicators to identify transactions requiring further review.



Narrow AI versus Artificial General Intelligence

Most AI systems used in organizations today are forms of Narrow AI

NARROW AI — FINANCE EXAMPLES

- Fraud and anomaly detection
- Credit-risk assessment
- Cash-flow and revenue forecasting
- Expense-claim analysis
- Intelligent document processing
- Conversational assistants
- Financial analysis and decision support

ARTIFICIAL GENERAL INTELLIGENCE

- A theoretical concept of AI capable of adapting across a broad range of intellectual tasks
- No universally accepted demonstration of human-level general intelligence exists today
- Should be distinguished from the specialised enterprise AI systems currently used in finance



Illustrative Scenario — Intelligent Expense Management

A hypothetical scenario demonstrating AI in expense review and exception management

- A multinational manufacturer receives a very large volume of employee expense claims each year; finance traditionally relied on predefined approval rules and manual sampling
- An AI-enabled system analyses claims using policy requirements, historical spending patterns, merchant information, duplicate submissions, employee behaviour and anomaly indicators
- Routine low-risk claims proceed through automated workflows, while unusual or higher-risk claims are escalated for review

EXECUTIVE INSIGHT

The approach is designed to accelerate reimbursement, improve exception detection and reduce routine review while retaining human approval for significant or unusual items. The objective is not to automate every activity, but to apply AI where it improves efficiency, insight, control or decision quality.

KNOWLEDGE CHECK 1

Test Your Understanding



Most AI systems in use across finance organizations today are best described as:

A Artificial General Intelligence, capable of any intellectual task

B **Narrow AI, designed to perform particular tasks within defined domains**



C Purely rule-based automation with no learning capability

D Systems that require no data at all to function

WHY: Most AI systems used in organizations today are forms of Narrow AI — designed to perform particular tasks or operate within defined domains. Artificial General Intelligence remains a theoretical concept with no universally accepted demonstration of human-level general intelligence.



SECTION 2.3

Machine Learning Explained Simply

Converting historical data into predictive insight



How Machine Learning Works

A typical Machine Learning process involves five stages



EXECUTIVE INSIGHT

Data quality is critical throughout the process. Poor or inconsistent data can materially reduce model reliability.

- For example, a model can analyse customer payment history, financial ratios, industry conditions and behavioural patterns to estimate the probability of future default



Key Machine Learning Terms

Terminology finance leaders will encounter when discussing ML projects

Term	Meaning	Finance Example
Dataset	Collection of information used for analysis	Historical receivables data
Feature	Input variable used by the model	Payment history, revenue, credit rating
Label	Outcome the model learns to predict	Default / No Default
Training Data	Data used to train the model	Historical loan portfolio
Test Data	Unseen data used to assess final performance	Holdout customer portfolio
Prediction	Model’s estimate of an outcome	Probability of default



Three Types of Machine Learning

Most business applications fall into three broad approaches



Supervised Learning

Learns from labelled historical outcomes



Unsupervised Learning

Discovers hidden patterns, clusters, anomalies



Reinforcement Learning

Learns through rewards and feedback

Learning Method	Primary Objective	Illustrative Finance Applications
Supervised Learning	Predict or classify	Credit risk, forecasting, fraud detection
Unsupervised Learning	Discover structure or anomalies	Segmentation, clustering, anomaly detection
Reinforcement Learning	Optimize sequential decisions	Trading, portfolio, selected optimization

EXECUTIVE INSIGHT

Supervised and unsupervised learning cover most finance applications today. Reinforcement Learning remains more specialized and should not be assumed appropriate for routine finance applications.



Why Machine Learning Matters to Finance Leaders

From "What happened?" toward "What is likely to happen?" and "What action should be considered?"

- More accurate forecasting
- Earlier identification of risk and anomalies
- Improved credit and customer analytics
- Better working-capital and liquidity decisions
- Stronger management decision support

EXECUTIVE INSIGHT

Model outputs remain dependent on data quality, assumptions, model design and changing business conditions. Finance leaders remain responsible for validation, governance, monitoring and human oversight.



SECTION 2.4

Deep Learning and Neural Networks

Multi-layered neural networks that identify increasingly complex patterns in data



What Is Deep Learning?

Machine Learning identifies patterns in data; Deep Learning identifies increasingly complex representations

- Deep Learning is a specialized form of Machine Learning that uses multi-layered neural networks to identify complex patterns in data
- It has enabled major advances in Generative AI, LLMs, Computer Vision, speech recognition and document intelligence
- Compared with conventional Machine Learning, Deep Learning can reduce the need for manual feature engineering because relevant representations can be learned automatically

A NEURAL NETWORK CONSISTS OF

- Input Layer — receives data such as transactions, documents, images or other information
- Hidden Layers — transform information through successive operations to identify complex patterns
- Output Layer — produces a prediction, classification, generated output or other result

KEY POINT

- The term "deep" refers to neural networks containing multiple hidden layers
- Training may require substantial computing resources and high-quality data, particularly for large or complex models



Illustration: A Simplified Neural Network

Input, hidden and output layers working together to produce a result



EXECUTIVE INSIGHT

For finance leaders, the mathematical architecture of Deep Learning is less important than understanding where it adds value, how it differs from conventional Machine Learning, and what risks and resource requirements accompany its use.

- Deep Learning does not replace conventional Machine Learning — the appropriate approach depends on the business problem, available data, required explainability, cost and complexity



Machine Learning versus Deep Learning

Feature	Conventional Machine Learning	Deep Learning
Data Requirements	Often effective with smaller/structured datasets	Often benefits from larger, complex datasets
Feature Engineering	Frequently requires greater human input	Can learn many representations automatically
Explainability	Often easier to interpret	Frequently more difficult to explain
Typical Strength	Structured prediction and classification	Complex text, image, audio and multimodal problems
Finance Examples	Forecasting, credit risk, classification	Document intelligence, language, image, Generative AI



Applications and Challenges of Deep Learning

Particularly valuable where finance involves large volumes of complex or unstructured information



Document Intelligence

Invoices, contracts, statements, filings



Computer Vision

Receipts, ID verification, asset inspection



Speech Recognition

Transcription of calls and meetings



Anomaly Detection

Fraud and transaction monitoring

Challenge	Finance Implication
Data & Computational Requirements	Reliable performance and training may require substantial, high-quality data and computing resources
Explainability	Complex networks may be difficult to interpret in regulated or judgement-sensitive applications
Cost, Expertise & Model Risk	Specialized skills and investment required; performance may deteriorate as conditions change

Test Your Understanding



Why does Deep Learning often present a greater explainability challenge than conventional Machine Learning?

- A Because Deep Learning does not use any data
- B Because multi-layered neural networks build complex representations that can be difficult to interpret**
- C Because Deep Learning always produces more accurate results
- D Because Deep Learning replaces the need for human review entirely



WHY: Deep Learning's multiple hidden layers build increasingly complex representations, which is a source of its power but also makes outputs harder to explain than many conventional Machine Learning models — a factor that matters most for regulated or judgement-sensitive applications.



SECTION 2.5

Generative AI and Large Language Models

From predicting and classifying to creating new content



What Is Generative AI?

AI systems capable of generating new content in response to instructions ("prompts")

- Generative AI can create new content — text, images, audio, video, software code and analytical outputs — based on patterns learned during training
- Finance applications: drafting reports, summarizing documents, preparing Board materials, analysing contracts, generating variance explanations and scenarios, supporting research
- At the centre of many text-based applications are Large Language Models (LLMs) — Deep Learning models designed to process and generate natural language

EXECUTIVE INSIGHT

Search primarily retrieves existing information. LLMs generate responses from learned patterns, supplied context, and — in connected applications — retrieved information. Fluent output does not necessarily mean the information is accurate.



How Large Language Models Work

Output quality depends on the model, prompt, context, source information and safeguards applied



1. Model Training

Learns statistical patterns from large datasets



2. Prompt & Context Processing

Processes instruction and available context



3. Response Generation

Predicts a sequence of tokens to form output

Traditional / Predictive AI	Generative AI
Predicts or classifies outcomes	Generates new content
Detects anomalies	Explains or summarizes identified anomalies
Forecasts revenue or cash flows	Generates narratives and scenarios around forecasts



Key Concepts Every Finance Leader Should Know



Prompts

The instruction or question given to a Generative AI system



Context

Information available to the model when generating a response



Tokens & Context Windows

Units of processing and how much information a model can consider



Retrieval-Augmented Generation

Retrieves relevant information from designated sources as context

EXECUTIVE INSIGHT

Clear objectives, relevant context, appropriate constraints and specified output formats generally improve response quality. A larger context window does not by itself guarantee accuracy or effective analysis.

- A standalone LLM may lack current, proprietary, or organization-specific information — this is the limitation RAG is designed to address



Retrieval-Augmented Generation (RAG)

Grounding model responses in trusted, current, organization-specific information

- RAG retrieves relevant information from designated sources and provides it to the model as context before the response is generated
- Trusted sources may include accounting policies, ERP data, internal reports and audit findings, SOPs, tax/regulatory guidance, and Board-approved policies

BENEFITS OF RAG

- Access to current and organization-specific information
- Stronger grounding in designated sources
- Improved traceability and source attribution
- Reduced risk of unsupported responses

RAG VERSUS FINE-TUNING

- RAG provides relevant knowledge at the time of the request
- Fine-tuning adapts how the model performs a task or generates outputs
- For frequently changing information (regulations, tax rules), retrieval from maintained sources is often more appropriate than fine-tuning alone



AI Copilots versus AI Agents

The appropriate degree of autonomy should depend on risk, materiality and human-approval requirements

AI Copilots	AI Agents
Assist users while leaving control primarily with the human	Use models, tools, data and workflows to perform multi-step activities
Drafts commentary, analyses variances, summarizes reports	Retrieves information, interacts with approved systems, executes sequential tasks
Supports budgeting, forecasting, presentations and research	Monitors workflow status, identifies and escalates exceptions
Outputs remain subject to appropriate professional review	Degree of autonomy depends on risk, materiality, permissions and controls

EXECUTIVE INSIGHT

For example, an AI agent supporting the financial close might retrieve approved data, identify unusual transactions, perform variance analysis, prepare draft commentary, and escalate unresolved exceptions.



Illustration: Agentic AI in Action

Agentic AI moves beyond drafting and analysis toward coordinated execution



EXECUTIVE INSIGHT

Its value depends on clearly defined authority limits, audit trails, controls and mandatory escalation of significant decisions. Material accounting, financial, regulatory or strategic decisions should remain subject to appropriate human oversight and accountability.

- Practical applications span Financial Reporting, FP&A, Treasury, Internal Audit, Tax, Risk Management, Investor Relations and Compliance



Practical Applications of Generative AI in Finance

Finance Function	Illustrative Generative AI Applications
Financial Reporting	Draft commentary, disclosures and executive summaries
FP&A	Variance explanations, scenarios and planning narratives
Treasury	Liquidity commentary, covenant summaries, financing analysis
Internal Audit	Planning support, document summaries, control documentation
Tax	Research support, draft memoranda, regulatory summaries
Risk Management	Risk summaries, scenarios and mitigation documentation
Investor Relations	Earnings materials and shareholder communications



Limitations of Generative AI

Should be treated as an analytical and productivity capability, not an autonomous decision-maker

Limitation	Finance Implication
Hallucination	Models may generate plausible but inaccurate or unsupported information
Knowledge Limitations	Standalone models may lack current or organization-specific information
Bias	Outputs may reflect biases within training data, context or model behaviour
Explainability & Traceability	Some outputs may be difficult to substantiate for regulated or material decisions
Human Accountability	Professional responsibility remains with authorized decision-makers



Illustrative Scenario — Accelerating Board Reporting

A hypothetical scenario demonstrating controlled use of Generative AI in Board materials

- A multinational organization uses an enterprise Generative AI solution integrated with approved financial and business information to support monthly Board reporting
- The system retrieves relevant data, prepares draft performance commentary, summarizes major variances, highlights emerging issues, and generates initial presentation narratives
- Finance executives review the underlying information, challenge the analysis, refine the commentary, and approve the final Board materials

EXECUTIVE INSIGHT

The approach reduces preparation effort, improves consistency, and enables finance professionals to devote more time to analysis and executive discussion while retaining human responsibility for the final report.

Test Your Understanding



What is the primary purpose of Retrieval-Augmented Generation (RAG)?

- A To permanently retrain the model on new data
- B To retrieve relevant information from designated sources and provide it as context before a response is generated** 
- C To eliminate the need for any human review of outputs
- D To increase the size of the model's neural network

WHY: RAG addresses a standalone LLM's lack of current or proprietary information by retrieving relevant material from trusted sources and providing it as context before a response is generated — improving grounding and traceability, but not eliminating the need for validation.



SECTION 2.6

The AI Technology Ecosystem: Building Intelligent Finance Organizations

The power of enterprise AI lies in how multiple capabilities work together



Components of the AI Ecosystem

Different AI technologies perform complementary roles

Technology	Primary Capability	Illustrative Finance Application
Machine Learning	Prediction and classification	Credit risk, fraud detection, forecasting
NLP	Processes and analyses language	Contracts, regulations, reports, correspondence
Computer Vision	Analyses images and visual information	Receipts, identity documents, asset verification
OCR / Document Intelligence	Converts documents into structured data	Invoices, bank statements, tax documents
Knowledge Graphs	Maps relationships among entities	Related parties, ownership structures, fraud
Generative AI	Generates and synthesizes content	Reports, summaries, scenarios, commentary
AI Agents	Orchestrates multi-step activities	Close processes, compliance, transaction workflows



How the Technologies Work Together

Consider an intelligent accounts-payable process



EXECUTIVE INSIGHT

The value arises not from any one component, but from the orchestration of technologies, enterprise data, business rules, controls and human judgement.

- Organizations typically derive greater value when AI capabilities are integrated into end-to-end finance processes rather than deployed as disconnected tools



From Isolated AI Tools to Intelligent Finance

An intelligent finance organization moves beyond simply automating transactions



- Automate routine processing and analyse structured and unstructured information
- Anticipate financial and operational outcomes and identify anomalies and emerging risks
- Generate timely management insights and coordinate workflows across enterprise systems
- Support faster and better-informed decisions

EXECUTIVE INSIGHT

Intelligent finance does not imply unrestricted autonomy. The objective is not maximum automation, but optimal collaboration among people, data, systems and AI.



Illustrative Scenario — Intelligent Trade Finance Processing

Combining multiple AI capabilities within a controlled trade-finance workflow

- A global bank processes international trade transactions involving invoices, bills of lading, letters of credit, insurance certificates, customs documents and counterparty information
- An integrated AI environment combines document intelligence, NLP for contractual terms, Computer Vision for verification, Knowledge Graphs for counterparty relationships, Predictive Analytics for transaction risk, and AI-enabled workflows to coordinate verification and escalation
- Higher-risk transactions and material exceptions remain subject to specialist review and approval

EXECUTIVE INSIGHT

The integrated approach can accelerate document processing, strengthen exception detection and compliance, reduce manual effort, and improve the consistency of trade-finance operations.



SECTION 2.7

AI Architecture and Infrastructure for Finance

Reliable AI requires trusted data, integrated systems, infrastructure, security and governance



Six Layers of Enterprise AI Architecture

These layers are interdependent — weaknesses in any one can undermine AI reliability

Layer	Key Components	Primary Purpose
1. Data	Financial, operational, customer, supplier, external data	Provides trusted information for AI
2. Enterprise Systems	ERP, CRM, TMS, HRMS, planning, tax, GRC	Generate and manage business information
3. Data & Integration	Data warehouses, lakes/lakehouses, APIs	Consolidates, stores, connects information
4. AI & Analytics	ML, Deep Learning, LLMs, Predictive Analytics	Generates predictions, analysis, insights
5. Applications & Orchestration	Assistants, copilots, agents, dashboards	Delivers AI capabilities within processes
6. Security & Governance	Cybersecurity, privacy, access, monitoring	Ensures secure, responsible, controlled use



Data: The Foundation of AI

For AI to use information effectively, data should be appropriately...

Accurate → Complete → Consistent → Timely → Valid → Accessible → Secure

- Structured data includes general ledger transactions, receivables, payables, budgets, forecasts, treasury information, payroll and financial statements
- Unstructured data includes contracts, audit reports, Board papers, emails, tax opinions, regulatory guidance, correspondence and policy documents

EXECUTIVE INSIGHT

Poor-quality or poorly governed data can produce unreliable AI outputs. Data governance, master-data management, ownership, lineage, access rights and quality controls are fundamental components of AI architecture.



Enterprise Systems and Data Platforms

AI does not replace core systems; it adds analytical, predictive and generative capabilities around them

CENTRALIZED DATA PLATFORMS

- Data Warehouse — stores structured, standardized information optimized for reporting and analysis
- Data Lake — stores larger volumes of structured and unstructured information in flexible formats
- Lakehouse — combines flexible storage with stronger management and analytical capabilities

INTEGRATION THROUGH APIS

- APIs enable applications, data platforms and AI services to exchange information and invoke approved functionality
- Secure APIs and integration platforms support authentication, authorization, monitoring and control



Computing Infrastructure and Deployment Models

The appropriate model depends on cost, scalability, resilience and data sensitivity



Public Cloud

Scalable, widely used for AI services



Private Cloud

Dedicated infrastructure for sensitive workloads



Hybrid Environments

Combines public and private resources



On-Premises

Selected infrastructure retained in-house

EXECUTIVE INSIGHT

Finance leaders should evaluate infrastructure choices based on business requirements and risk, rather than assuming that a single deployment model is appropriate for every organization.

- AI models provide the analytical intelligence within the architecture; business users normally access these capabilities through assistants, dashboards, conversational interfaces and automated workflows rather than interacting directly with underlying models



Cybersecurity, Privacy and Governance

Security and governance should be designed into the architecture, not added after implementation

KEY SAFEGUARDS

- Identity and access management; authentication and authorization
- Encryption; secure APIs and system interfaces
- Data classification and masking; cybersecurity monitoring
- Backup, resilience and recovery; logging and audit trails

AI GOVERNANCE SHOULD ESTABLISH

- Accountability for data and models
- Validation requirements and permitted uses
- Human-oversight thresholds and regulatory compliance
- Ethical principles, performance monitoring, escalation procedures



How the Architecture Works Together

Illustration: an AI-enabled quarterly performance-reporting process



EXECUTIVE INSIGHT

Data controls protect information quality → cybersecurity protects access and confidentiality → governance controls AI use → human reviewers validate material outputs and approve final reporting.

- Successful enterprise AI depends as much on the strength of the underlying architecture as on the sophistication of individual AI models



SECTION 2.8

Limitations, Risks, and Responsible AI

AI is neither infallible nor risk-free — outputs depend on data, models, context and governance



Understanding the Limitations of AI

Current AI systems remain fundamentally different from human intelligence

- AI can process large volumes of information, identify patterns, generate predictions, analyse language and automate complex tasks
- It does not possess human judgement, professional scepticism, ethical responsibility, emotional intelligence or accountability
- AI outputs should be treated according to their risk, materiality, purpose and reliability, with appropriate human review where significant consequences may arise

EXECUTIVE INSIGHT

AI should be deployed within a Responsible AI framework that balances innovation and business value with accuracy, security, fairness, transparency, accountability, regulatory compliance and human oversight.



Principal AI Risks (1 of 2)

Risk	Potential Finance Impact
Hallucination	Incorrect reporting, regulatory errors, misleading analysis
Bias	Discriminatory credit, customer, employee or supplier decisions
Explainability	Governance, audit, regulatory and accountability concerns
Data Quality	Unreliable predictions and decisions
Privacy & Confidentiality	Privacy breaches, regulatory penalties, loss of trust



Principal AI Risks (2 of 2)

Risk	Potential Finance Impact
Cybersecurity	Data leakage, manipulation, fraud, operational disruption
Model Drift	Increasingly inaccurate forecasts, classifications or risk assessments
Ethical Risk	Reputational damage, unfair outcomes, loss of stakeholder confidence
Third-Party Risk	Concentration, security, compliance, continuity and contractual risks
Over-reliance on AI	Weak professional judgement and inappropriate decisions

EXECUTIVE INSIGHT

These risks are interconnected and should be managed throughout the AI lifecycle rather than considered only after deployment.



Hallucinations, Bias and Explainability

Three of the most significant risk categories in practice

- Hallucination: Generative AI may produce convincing but inaccurate information — invented regulations, incorrect statistics, fictitious references. Material outputs should be verified against reliable, authoritative sources; RAG improves grounding but does not eliminate the need for validation
- Bias: models may reproduce or amplify biases present in training data, assumptions, features or historical decisions — organizations should assess representativeness, unintended discrimination, and changes in outcomes over time
- Explainability: complex Deep Learning systems may be difficult to interpret; explainability requirements should reflect the risk and materiality of the application



Data, Privacy, Cybersecurity and Model Risk

AI reliability depends heavily on the quality and security of the information it uses

KEY SAFEGUARD AREAS

- Data Quality — accuracy, completeness, consistency, timeliness, lineage, ownership
- Privacy — permitted platforms, data classification, access, retention, cross-border transfers
- Cybersecurity — identity/access management, encryption, monitoring, protection against prompt injection, data poisoning and adversarial inputs

MODEL RISK — MONITOR FOR

- Model accuracy and performance
- Data quality and distribution changes
- False positives and false negatives; emerging bias
- Security incidents and regulatory changes



Responsible AI Principles

The level of governance should be proportionate to risk, materiality, autonomy and potential consequences

Principle	Governance Requirement
Human Oversight	Appropriate human review and authority over significant decisions
Accountability	Clear ownership of AI systems, data, risks, controls and outcomes
Transparency	Appropriate disclosure of AI use and sufficient information about outputs
Fairness	Identification and mitigation of unjustified bias or discrimination
Security & Privacy	Protection against unauthorized access, manipulation and data leakage
Traceability & Monitoring	Documentation, logs, version control, ongoing performance assessment



The Role of Finance Leaders

Finance is an important participant in AI governance, alongside other disciplines

- Contributing to enterprise AI governance and risk frameworks
- Establishing controls over AI used within finance; ensuring data quality and financial integrity
- Defining human-review and approval requirements; overseeing validation and monitoring
- Ensuring appropriate documentation and audit trails
- Coordinating with technology, risk, legal, compliance, cybersecurity and internal audit teams
- Reporting material AI risks and issues to senior management and the Board where appropriate

EXECUTIVE INSIGHT

Responsible AI is an enterprise-wide responsibility requiring multidisciplinary oversight.



Illustrative Scenario — Responsible AI in Financial Reporting

The governance response required when model reliability or fairness deteriorates

- A listed multinational uses Generative AI to assist with drafting quarterly management commentary, accessing approved information through controlled enterprise systems
- Draft commentary is grounded in designated sources, with access controls, version history, source references and audit trails supporting traceability
- Finance professionals review material information and narrative conclusions before approval, with additional legal, compliance or investor-relations review where appropriate

EXECUTIVE INSIGHT

The process demonstrates how Generative AI can accelerate financial-reporting activities while maintaining human accountability, controlled data access, validation and governance.

Test Your Understanding



A credit-risk model's predictions become steadily less accurate as economic conditions change. This is best described as an example of:

A Hallucination

B Model Drift



C Retrieval-Augmented Generation

D Fine-tuning

WHY: Model Drift refers to declining model performance as underlying data or business conditions change. Organizations should monitor accuracy, distribution changes and other indicators, recalibrating, retraining, replacing, or withdrawing models where appropriate.



SECTION 2.9

Selecting and Implementing AI Solutions: A Practical Decision Framework

Where AI can create value, which solution is appropriate, and how it should be implemented and governed



Begin with the Business Problem

The first question should be about outcomes, not products

THE FIRST QUESTION

“What business outcome are we trying to improve?” — rather than — “Which AI product should we purchase?”

- Improving forecast accuracy
- Accelerating financial close
- Strengthening fraud and anomaly detection
- Automating document-intensive processes
- Improving management reporting and strengthening compliance monitoring

EXECUTIVE INSIGHT

Only after defining the desired outcome and measurable success criteria should the organization determine whether AI is appropriate and which capabilities are required.



Match the Business Need to AI Capabilities

Most finance problems require combinations of technologies rather than a single AI tool

Business Need	Potential AI Capabilities
Forecasting	Machine Learning, Predictive Analytics
Invoice Processing	Document Intelligence, OCR, Machine Learning
Board Reporting	Generative AI, LLMs, RAG
Financial Close	Automation, anomaly detection, GenAI, AI Agents
Fraud Detection	Machine Learning, anomaly detection, Knowledge Graphs
Contract Review	NLP, LLMs, RAG



Confirm Organizational Readiness

The principal readiness questions before committing significant investment



Data

Is reliable, accessible and governed data available?



Technology

Can systems and infrastructure support the solution?



Governance

Are security, risk and oversight requirements defined?



People

Are skills, ownership and sponsorship available?



Compliance

Have legal and regulatory requirements been considered?

EXECUTIVE INSIGHT

At the solution-selection stage, the objective is to determine whether any material readiness gaps must be addressed before implementation.



Select the Implementation Approach

There is no universally superior approach

Approach	Appropriate Where	Principal Advantages
Embedded / Configure	Existing platforms already provide suitable capabilities	Faster deployment, easier integration
Buy	Requirements are common; established solutions exist	Speed, vendor support, proven functionality
Build	Requirements are specialized or proprietary	Greater control and customization
Hybrid	Commercial capabilities need enterprise-specific data	Balances speed, flexibility and differentiation

EXECUTIVE INSIGHT

Where requirements are common across organizations, established commercial or embedded solutions may be more economical. Custom development is more likely justified where proprietary data or processes provide strategic differentiation.



Evaluating Vendors and Developing the Business Case

AI vendors should be assessed with the same discipline applied to other strategic relationships

VENDOR EVALUATION DIMENSIONS

- Capability, integration and security & privacy
- Governance, compliance and scalability
- Vendor risk, service & support, commercial terms
- Product roadmap alignment with future architecture


THE BUSINESS CASE SHOULD COVER

- Benefits — cost savings, productivity, forecast accuracy, risk reduction
- Costs — licensing, development, integration, training, governance
- Risks — model, implementation, cybersecurity, vendor dependency
- Measures of success — baseline, targets, milestones, adoption



Pilot Before Scaling

A practical implementation sequence



EXECUTIVE INSIGHT

A successful pilot should demonstrate not merely that the technology works, but that it improves the business process within acceptable risk and economic parameters.

- Critical success factors: clear business ownership + executive sponsorship + trusted data + appropriate technology + multidisciplinary governance + user adoption + measurable value + ongoing monitoring



Illustrative Scenario — Scaling AI Across Finance

Use-case prioritisation, controlled piloting and phased scaling

- A multinational engineering organization identifies several potential AI use cases: forecasting, accounts-payable automation, contract analysis, management reporting and audit analytics
- Rather than deploying multiple disconnected tools, it prioritizes use cases by expected value, feasibility and risk, and addresses data governance, integration and cybersecurity gaps before piloting
- Pilot results are evaluated against predefined performance, control and business-value measures before progressive scaling using common standards

EXECUTIVE INSIGHT

Successful AI transformation depends not on deploying the largest number of AI tools, but on selecting appropriate use cases, establishing necessary foundations, validating value, and scaling responsibly.



Questions Finance Leaders Are Asking (1 of 2)

Question	Answer
What is the difference between AI and Machine Learning?	AI is the broader field concerned with tasks associated with human intelligence. Machine Learning is a subset of AI that uses data to learn patterns and generate predictions or classifications.
How is Deep Learning different from Machine Learning?	Deep Learning uses multi-layered neural networks to identify complex patterns, particularly effective for text, images, audio and documents, though it may require more computing resources.
Is Generative AI the same as Machine Learning?	No. Generative AI is built on ML/Deep Learning but is designed to generate new content, while traditional ML is more commonly used for prediction, classification and anomaly detection.
What is the difference between RAG and fine-tuning?	RAG retrieves relevant information as context at the time of the request; fine-tuning adapts model behaviour through additional training.



Questions Finance Leaders Are Asking (2 of 2)

Question	Answer
What is the difference between an AI Copilot and an AI Agent?	A Copilot assists users while leaving control with the human professional. An Agent can coordinate and execute multi-step activities within defined permissions and governance boundaries.
Can AI replace ERP systems?	Generally, no. ERP systems remain important systems of record; AI typically complements them with predictive analytics, anomaly detection and Generative AI.
What is Responsible AI?	The design, development, deployment, use and monitoring of AI in a manner that is lawful, secure, ethical, transparent and accountable.
How should finance leaders select and implement AI solutions?	Begin with a clearly defined business problem, assess readiness, identify required capabilities, evaluate approaches, build a business case, govern, pilot and scale progressively.



Key Takeaways from Module 2

- Enterprise AI is not a single technology — it is an ecosystem of data, models, systems, infrastructure, governance and human responsibility
- Machine Learning, Deep Learning and Generative AI serve complementary rather than competing roles, each suited to different classes of problem
- AI Agents extend Copilots by coordinating multi-step activities, but require stronger permissions, controls and monitoring
- Reliable AI depends as much on data quality, architecture and governance as on the sophistication of the underlying model
- Finance leaders begin with the business problem, evaluate implementation approaches, pilot before scaling, and retain human accountability for material decisions



An Integrated Ecosystem, Not a Single Technology

THE CENTRAL MESSAGE

Artificial Intelligence comprises complementary technologies for prediction, classification, language processing, document intelligence, content generation and controlled workflow orchestration. Their effective use depends on reliable data, suitable enterprise architecture, secure integration, appropriate technology selection and disciplined implementation.

- AI also introduces material risks relating to reliability, hallucination, bias, explainability, privacy, cybersecurity and model drift
- Finance leaders must match each business problem with the appropriate capability, establish proportionate controls, and retain human accountability for material decisions

LOOKING AHEAD

Module 3 builds on these foundations by examining how AI transforms Financial Planning and Analysis.



SECTION 2.11

Five Illustrative Case Studies

AI foundations and applications — hypothetical scenarios illustrating controls and management considerations



About These Case Studies

IMPORTANT NOTE

The following cases are illustrative and hypothetical. They use representative organisations and assumptions to demonstrate AI applications, controls and management considerations and do not describe verified results from identified organisations.



2.1 Predicting Customer Default Before It Happens



2.2 Transforming Accounts Payable with Intelligent Document Processing



2.3 Building a Generative AI Finance Knowledge Assistant



2.4 Deploying an AI Agent for the Financial Close



2.5 When an AI Credit Model Goes Wrong

Predicting Customer Default Before It Happens

A multinational industrial company selling to thousands of distributors and corporate customers on credit



BUSINESS SITUATION

The company's traditional credit-control process relies on credit ratings, historical financial statements, ageing reports and periodic manual reviews. Although effective for established risks, the process often identifies deterioration only after customers begin delaying payments.



AI INITIATIVE

The company develops a Machine Learning-based credit-risk model using historical payment behaviour, ageing trends, credit utilization, financial indicators, order patterns, industry conditions and selected external economic data. The model assigns each customer a probability of payment deterioration and segments customers into risk categories.



BUSINESS IMPACT

Finance can intervene earlier by reviewing credit limits, accelerating collections or modifying terms. Credit professionals spend less time on low-risk accounts and more on material exceptions. Management does not permit the model to automatically withdraw credit — material changes remain subject to professional review.

Transforming Accounts Payable with Intelligent Document Processing

A global consumer-products company receiving supplier invoices in varied formats, languages and layouts



BUSINESS SITUATION

Finance personnel manually extract invoice information, verify purchase-order references, perform matching, identify duplicates and route exceptions for approval. The process is labour-intensive and slows payment processing.



AI INITIATIVE

The organization implements an integrated Intelligent Document Processing solution. OCR and document-intelligence capabilities extract supplier names, invoice numbers, dates, currencies and payment terms; Machine Learning classifies documents and flags unusual transactions; automated workflows perform matching and route exceptions.



BUSINESS IMPACT

Routine data entry is reduced and processing accelerated, freeing employees to concentrate on exceptions. Success is measured not solely by automation percentage, but also exception accuracy, duplicate detection, control failures, supplier experience and auditability.

Building a Generative AI Finance Knowledge Assistant

A multinational group where finance professionals regularly consult policies, manuals and technical guidance



BUSINESS SITUATION

Finding relevant information across numerous repositories is time-consuming. Management wants to use Generative AI but is concerned that a standalone LLM could provide inaccurate or unsupported answers.



AI INITIATIVE

The company develops an enterprise finance assistant combining an LLM with Retrieval-Augmented Generation. When a question is asked, the system searches approved internal sources, retrieves relevant information, and provides it to the LLM as context, including references to the underlying source. Access permissions ensure employees see only authorized information.



BUSINESS IMPACT

Finance teams locate and synthesise approved internal knowledge more efficiently and consistently. The controlled RAG environment reduces reliance on unsupported model-generated information while improving traceability. For material matters, qualified professionals verify source material and approve the final position.

Deploying an AI Agent for the Financial Close

A diversified multinational group operating numerous subsidiaries



BUSINESS SITUATION

The monthly close involves reconciliations, journal reviews, intercompany matching, variance analysis, status monitoring and repeated follow-up among finance teams. Although several activities are already automated, coordination across systems and teams remains highly manual.



AI INITIATIVE

The organization introduces an AI-enabled close agent operating within defined permissions. It retrieves approved information, monitors completion of activities, identifies unresolved reconciliations, highlights unusual journal entries, performs preliminary variance analysis, prepares draft commentary and alerts personnel to overdue items. It cannot independently approve material journal entries or certify financial statements.



BUSINESS IMPACT

Earlier visibility of close bottlenecks and reduced routine coordination effort, while human approval of material accounting decisions is preserved. Activity logs record actions performed, information accessed, exceptions raised and human approvals obtained.

When an AI Credit Model Goes Wrong

A financial-services organization using an AI model to support credit decisions



BUSINESS SITUATION

The model initially performs within approved thresholds, but as the economic environment changes, customer behaviour begins to differ from the historical data on which it was developed. Analysis identifies that certain variables may be acting as unintended proxies producing inappropriate differences among customer groups.



AI INITIATIVE

Management observes declining prediction accuracy, unusual rejection patterns, difficulty explaining outcomes, and growing employee reliance on model recommendations without sufficient challenge. The organization increases human-review requirements and initiates an independent model review covering data quality, performance, feature selection, bias, explainability and governance.



BUSINESS IMPACT

The model is recalibrated and independently validated before wider use resumes. Enhanced monitoring is introduced for drift, fairness, overrides, and false positives/negatives, with clearer escalation thresholds. The case shows that AI governance does not end at deployment — ongoing monitoring, validation, challenge and human accountability remain essential.



Results at a Glance

Illustrative outcomes across the five case studies (hypothetical scenarios)

Case	AI Capability	Illustrative Result
2.1 Customer Default Prediction	Machine Learning credit-risk model	Earlier intervention; sharper focus on material exceptions
2.2 Intelligent Document Processing	OCR + ML + automated workflows	Faster processing; reduced manual data entry
2.3 GenAI Knowledge Assistant	LLM + Retrieval-Augmented Generation	Faster, more consistent access to approved knowledge
2.4 AI Agent for Financial Close	Agentic AI within defined permissions	Earlier visibility of bottlenecks; preserved human approval
2.5 AI Credit Model Review	Model monitoring and governance response	Recalibrated, independently validated model

EXECUTIVE INSIGHT

The common lesson: technology alone does not create intelligent finance. Sustainable value emerges when AI capability + trusted data + integrated systems + appropriate controls + human judgement operate together.



Concluding Perspective

THE FIVE CASES ILLUSTRATE A PROGRESSION

Machine Learning predicts → Deep Learning and document intelligence process → Generative AI synthesizes → AI Agents orchestrate → Responsible AI governs.

- What business problem does it solve?
- What information does it rely upon?
- What decisions can it influence?
- What risks does it create?
- Where must human accountability remain?

EXECUTIVE INSIGHT

Finance leaders should evaluate AI not simply by asking "What can the technology do?" but by asking these five questions.



ANNEXURES 2A – 2C

Technology, Capability and Readiness Reference Frameworks

Quick-reference frameworks for technology selection and organizational self-assessment



AI Technology and Governance Reference Guide

A quick reference for finance leaders evaluating core AI capabilities

AI Capability	Primary Purpose	Key Benefit
Artificial Intelligence (AI)	Prediction, pattern recognition, decision support	Automation and enhanced decision support
Machine Learning (ML)	Identify patterns to predict or classify outcomes	Predictive insight
Deep Learning	Multi-layered neural networks for complex data	Handles complex, unstructured information
Generative AI	Generate and synthesize new content	Productivity and knowledge support
Large Language Models (LLMs)	Process and generate natural language	Conversational knowledge interaction
Natural Language Processing (NLP)	Analyse and process language	Extracts insight from text



AI Technology and Governance Reference Guide

Continued — orchestration, grounding and governance capabilities

AI Capability	Primary Purpose	Important Consideration
Computer Vision	Analyse images and visual information	Performance depends on data and image quality
OCR / Document Intelligence	Extract structured information from documents	Extraction requires validation and exceptions handling
Predictive Analytics	Estimate future outcomes and probabilities	Forecast performance requires monitoring
Retrieval-Augmented Generation (RAG)	Ground model responses in retrieved information	Source information must be reliable and current
AI Agents	Coordinate multi-step activities using tools/workflows	Requires permissions, controls, monitoring, escalation
Responsible AI	Govern AI securely, ethically and accountably	Requires lifecycle governance and multidisciplinary oversight



AI Capability Selection Matrix for Finance Functions

Matching finance requirements to illustrative AI capabilities

Finance Requirement	Illustrative AI Capabilities	Expected Business Outcome
Revenue Forecasting	ML + Predictive Analytics	Improved forecast accuracy
Cash-Flow Forecasting	ML + Predictive Analytics	Earlier liquidity insight
Fraud Detection	ML + Anomaly Detection + Knowledge Graphs	Earlier identification of suspicious activity
Accounts Payable	Document Intelligence + ML + Agents	Faster processing, lower manual effort
Financial Close	Automation + Anomaly Detection + GenAI + Agents	Shorter and better-controlled close
Board Reporting	GenAI + LLMs + RAG	Faster preparation and stronger synthesis



AI Capability Selection Matrix for Finance Functions

Continued — planning, compliance and strategic applications

Finance Requirement	Illustrative AI Capabilities	Expected Business Outcome
Budgeting & Planning	Predictive Analytics + GenAI	Better planning and decision support
Regulatory Compliance	NLP + RAG + Rules/Analytics	Stronger compliance monitoring
Contract Review	NLP + LLMs + RAG	Faster identification of relevant terms and risks
Treasury	ML + Predictive Analytics + Optimization	Better treasury decision support
Customer Credit	ML + Predictive Analytics	Earlier and more consistent risk identification
Strategic Planning	Predictive Analytics + GenAI	Better-informed strategic decisions

EXECUTIVE INSIGHT

No single AI technology addresses every finance challenge. High-performing finance organizations increasingly combine complementary AI, analytics and automation capabilities within integrated business processes.



Enterprise AI Readiness Assessment and Action Checklist

Eight assessment areas covering strategy, data, technology, governance, security, vendors, people and value

Assessment Area	Illustrative Readiness Question
Strategy & Leadership	Are AI initiatives linked to clearly defined business priorities and measurable outcomes?
Data Readiness	Is relevant data sufficiently accurate, complete, consistent, timely and reliable?
Technology & Architecture	Can AI integrate effectively with relevant ERP, treasury, planning and reporting systems?
AI Governance & Model Risk	Are material AI models appropriately tested and validated before deployment?
Security, Privacy & Compliance	Are AI systems protected through appropriate identity, access and security controls?



Readiness Rating, Priority and Executive Recommendation

Rating	Interpretation	Management Response
Green (G)	Capability substantially established and effective	Maintain, monitor and improve
Amber (A)	Capability partially established or needs strengthening	Develop and implement an improvement plan
Red (R)	Significant readiness or control gap exists	Address before scaling material or higher-risk AI

EXECUTIVE INSIGHT

The objective is not to achieve theoretical perfection before adopting AI, but to establish a sufficiently strong foundation to experiment safely, demonstrate business value, manage risk, and scale responsibly.

- Remaining assessment areas cover Third-Party & Vendor Risk, People & Change Readiness, and Performance, Control & Value Realization — each rated Green/Amber/Red with a prioritized action plan



ANNEXURE 2E

Visual Guide to Core AI Technologies for Finance Leaders

Nine illustrative scenarios (II.1–II.9) and a comparative technology reference (II.10–II.13)



Nine Illustrative Scenarios: Rules to Govern Action

Each follows a common structure: Scenario → As Is → AI-Enabled To Be → Conclusion



II.1 Evolution of Planning
Rules → prediction → generation
→ agency → simulation in forecasting



II.2 ML Cash-Flow Forecasting
Learning individual customer payment behaviour



II.3 Deep Learning Fraud Detection
Recognising complex invoice and payment patterns



II.4 GenAI Management Reporting
Grounded LLM commentary for the CFO



II.5 Agentic Receivables Management
Goal-directed collection within authority limits



II.6 Enterprise Digital Twin
Simulating a major capacity expansion decision



II.7 RPA & Intelligent Automation
Rules-based execution combined with AI judgement



II.8 Predictive & Prescriptive Analytics
From "what is likely" to "what should we do"



II.9 Human–AI Collaboration
Capital allocation combining AI analysis and judgement

EXECUTIVE INSIGHT

The full narrative for each scenario (Scenario / As Is / To Be / Financial Leadership Implication) is set out in the source chapter; this annexure summarizes the key concept and AI-enabled outcome for quick reference.



How the Technologies Relate to One Another

A progressive level of intelligence and decision involvement

Level	Capability	Finance Example
1. Execute	Rules / RPA	Automatically reconcile matching transactions
2. Detect	Deep Learning / ML	Identify unusual transactions or probable fraud
3. Predict	ML / Predictive Analytics	Forecast customer collections or liquidity
4. Recommend	Prescriptive Analytics	Recommend optimal cash deployment or pricing
5. Generate	Generative AI / LLM	Prepare management commentary and explanations
6. Act	Agentic AI	Execute approved workflows across finance systems
7. Simulate	Digital Twin	Model enterprise outcomes under alternative scenarios
8. Govern & Decide	Human–AI Collaboration	Management evaluates evidence and decides



Comparative Analysis of Core AI Technologies (1 of 2)

Different but increasingly interconnected roles

Technology	Core Question	Principal Strength
Rule-Based Automation	What should happen when a condition occurs?	Consistent and highly controllable
RPA	Which repetitive activities can be automated?	High-volume processing and productivity
Machine Learning	What does historical data suggest will happen?	Prediction based on large datasets
Deep Learning	What complex patterns can be detected?	Handles large, complex unstructured data
Predictive Analytics	What is likely to happen?	Forward-looking insight



Comparative Analysis of Core AI Technologies (2 of 2)

Technology	Core Question	Principal Strength
Prescriptive Analytics	What should we consider doing?	Converts prediction into decision alternatives
Generative AI / LLMs	What can be explained or generated from knowledge?	Rapid synthesis and natural-language interaction
Agentic AI	What actions are required to achieve a goal?	Moves from insight to coordinated execution
Enterprise Digital Twin	What could happen under alternative scenarios?	Simulation of interconnected business outcomes
Human–AI Collaboration	How should AI and human judgment combine?	Combines computational scale with judgment

EXECUTIVE INSIGHT

Finance leaders should view these technologies as components of an integrated capability rather than competing technologies.



Comparative Application Across the Finance Function (1 of 2)

Finance Area	Rules / RPA	ML / Deep Learning	GenAI / LLM
FP&A	Data consolidation	Forecasting models	Management commentary
Treasury	Bank-data processing	Cash-flow prediction	Treasury summaries
Accounts Receivable	Reminder workflows	Payment prediction	Customer communications
Accounts Payable	Invoice processing	Duplicate/anomaly detection	Exception explanations
Financial Reporting	Reconciliations & close tasks	Anomaly identification	Reporting narratives



Comparative Application Across the Finance Function (2 of 2)

Finance Area	Agentic AI	Digital Twin
Cost Management	Corrective workflows	Operating-model simulation
Risk Management	Continuous monitoring	Enterprise stress testing
Internal Audit	Continuous audit workflows	Control-scenario simulation
Tax	Compliance orchestration	Scenario analysis
Corporate Finance	Transaction workflows	M&A / capital-allocation simulation



Finance Leader’s Technology Selection Guide

Select according to the business problem, not because a technology is fashionable

If the Finance Problem Is Primarily...	Most Relevant Capability
Highly repetitive and rules driven	Rules / RPA
Based on recognising historical patterns	Machine Learning
Based on highly complex or unstructured patterns	Deep Learning
Concerned with forecasting future outcomes	Predictive Analytics
Based on analysing or generating language	Generative AI / LLMs
Requiring multiple steps and actions across systems	Agentic AI
Involving material judgment or strategic choice	Human–AI Collaboration



An Integrated Architecture, Not Competing Technologies

EXECUTIVE INSIGHT

Rule-based automation and RPA continue to execute predictable tasks. Machine Learning and Deep Learning identify patterns. Predictive and Prescriptive Analytics anticipate outcomes and evaluate alternatives. Generative AI makes insights easier to understand. Agentic AI coordinates governed execution. Digital Twins permit management to test decisions virtually.



EXECUTIVE INSIGHT

Across every layer, human finance professionals retain responsibility for judgment, ethics, material decisions, governance and accountability. The objective is not autonomous finance without people, but increasingly intelligent finance with appropriately governed human oversight.



ANNEXURE 2F

Finance Problem to AI Capability Matrix

Beginning with the business problem to be solved, rather than a preferred technology



Finance Problem to AI Capability Matrix

The classifications are indicative rather than exclusive — solutions frequently combine several technologies

Finance Problem	Primary AI Capability	Required Human Oversight
Deterministic transaction processing	Rules-based execution	Approve rules, access rights, exception thresholds
Invoice, contract & statement extraction	Document intelligence	Review low-confidence fields and material exceptions
Account reconciliation	Intelligent matching	Investigate unresolved or unusual differences
Revenue, cost & cash-flow forecasting	Predictive analytics	Challenge assumptions and structural changes
Budgeting and driver-based planning	Predictive & prescriptive analytics	Approve drivers, constraints, assumptions
Transaction anomaly detection	Pattern and exception detection	Investigate alerts and determine corrective action



Finance Problem to AI Capability Matrix

Continued — relationship analytics through to strategic decision support

Finance Problem	Primary AI Capability	Required Human Oversight
Fraud and relationship analysis	Relationship analytics (Knowledge Graphs)	Assess context, evidence, legal implications
Management commentary	Content generation (Generative AI)	Verify factual accuracy, balance and tone
Accounting-policy interrogation	Knowledge retrieval and synthesis (RAG)	Confirm source authority and applicability
Financial close coordination	Intelligent workflow orchestration (Agentic AI)	Retain approval authority; intervene on exceptions
Capital and resource allocation	Optimization	Determine strategic priorities and acceptable risk
Strategic scenario evaluation	Enterprise simulation (Digital Twin)	Select scenarios; make the final decision



CFO Selection Questions

Before selecting an AI capability, the CFO and finance team should ask

- Is the finance problem deterministic, predictive, explanatory, coordinative, optimization-based or simulation-based?
- Is the required data available, reliable, authorized and sufficiently representative?
- Does the application require transparency and explainability?
- What would be the financial, regulatory or reputational consequence of an incorrect output?
- Should the system inform, recommend, initiate or execute an action?
- What level of human review and approval is required?
- Can the expected value be measured against implementation cost and risk?
- Can a simpler technology solve the problem adequately?



Fit-for-Purpose, Not Most Advanced

EXECUTIVE INSIGHT

The appropriate AI capability is the one that solves a defined finance problem with the least unnecessary complexity while meeting the organization's requirements for reliability, explainability, control and accountability.

Rules



Prediction



Generation



Agency



Simulation

EXECUTIVE INSIGHT

The objective is not to deploy the most advanced technology in every activity. It is to create a fit-for-purpose combination of technology and human judgement that improves the quality, speed and value of finance decisions.



Module 2 — In Summary

Understanding the AI Landscape: Technologies, Architecture and Enterprise Ecosystems



Machine Learning converts historical data into predictive insight



Deep Learning identifies complex patterns in unstructured data



Generative AI and LLMs create content grounded in trusted context



AI Agents coordinate multi-step workflows within defined authority



Reliable AI depends on data, architecture, security and governance



Responsible AI requires human oversight and lifecycle monitoring



Selection begins with the business problem, not the technology



Finance leaders must govern AI while retaining accountability



BD CHATTERJEE

FCA, ACMA, ACS, DipIFR ACCA (UK)


Founder, Digital CFO Solutions

ABOUT THE AUTHOR


A Career Built on Making Finance Understandable

- Over 35 years of CFO and industry experience across global finance leadership roles.
- Author of the Pocketbook of IFRS series — institutional-grade IFRS and US GAAP educational publications for finance professionals.
- Building a growing library of digital workbooks pairing financial rigour with AI acceleration under the Digital CFO Solutions brand.
- Teacher and mentor to finance professionals worldwide, focused on making complex financial concepts practical and understandable.

Books by B D Chatterjee — IFRS & GAAP Library



Available at bdchatterjee.com • Amazon Kindle & Paperback



THANK YOU

Digital CFO Solutions — Managing Finance in the AI Era

Module 2: Understanding the AI Landscape — Technologies, Architecture and Enterprise Ecosystems

BDCHATTERJEE.COM