

State of

AI

in Financial
Services
in India



Report 2025-26



State of AI in Financial Services in India

Adoption Trends, Industry Applications, Risks, and Governance

Shreenivas Kunte, CFA, Ph.D.
Director, Research Advocacy & Standards Committee,
CFA Society India

Harish Subramanian, CFA, CQF
Volunteer, Research Advocacy & Standards Committee,
CFA Society India

Suraj Harikumar, CFA
Volunteer, Research Advocacy & Standards Committee,
CFA Society India

Pankaj Sharma
Director, Capital Markets Policy, India,
CFA Institute

Rajni Dhameja, CFA
Volunteer, Research Advocacy & Standards Committee,
CFA Society India

Sagar Thacker, CFA
Volunteer, Research Advocacy & Standards Committee,
CFA Society India

Akhilesh Tilotia, CFA
Volunteer, Research Advocacy & Standards Committee,
CFA Society India

Brian Pisaneschi, CFA
Senior Investment Data Scientist,
CFA Institute

Saima Ahmed
Consultant, Research Initiatives and Operations
CFA Society India

ABOUT CFA SOCIETY INDIA

CFA Society India, established in 2005 and registered as the Indian Association of Investment Professionals under Section 8 of the Companies Act, 2013, is a leading professional body representing over 2,400 investment professionals across the country, including portfolio managers, investment advisors, and other finance practitioners. The Society is dedicated to advancing the investment profession in India by promoting the highest standards of ethics and professional conduct, fostering knowledge exchange, and driving thought leadership within the industry. As part of the global network of CFA Institute, which represents over 175,000 investment professionals worldwide, CFA Society India connects its members to best practices and a broader professional ecosystem, while continuing to create long-term value for its members and the financial community at large. For more information, visit <https://cfasocietyindia.org/>

ABOUT RESEARCH ADVOCACY & STANDARDS COMMITTEE

Research, Advocacy and Standards Committee aims to contribute to the important stakeholders of the investment profession through the advocacy efforts viz: investors. The advocacy efforts of the committee are focused on contributing meaningfully toward corporate governance and investor awareness. Research and Advocacy committee has 3 main pillars of activities: Working with members to propagate standards of best practices, working with regulators and media to raise standards in the profession, and working with the investors to promote awareness of CFA Charter and increase financial literacy. To achieve its objectives, the committee engages with regulators and other agencies in the financial market.

CFA Society India does not provide investment, financial, tax, legal, or ethical advice. This report has been prepared for informational purposes only and is not intended to provide, nor should it be relied upon for, investment, financial, tax, legal, or other advice. CFA Society India is not responsible for the content of any websites or information resources that may be referenced in this report, and references to such sites or resources do not constitute an endorsement of the information contained therein. The inclusion of company examples does not in any way constitute an endorsement of these organizations by CFA Society India. Although every effort has been made to ensure that the information contained in this report has been obtained from reliable and up-to-date sources, the evolving nature of statistics, laws, rules, and regulations may result in delays, omissions, or inaccuracies in the information presented.

Top 10 Findings from the Report:

1. Interviewee companies understand the significant potential AI/ Deep Tech have to offer and are actively implementing AI strategies across verticals.
2. The biggest productivity gains are accrued through coding, marketing and customer-facing application of AI.
3. Three different AI strategy approaches:
 - Decision making: Preference for bottom-up experimentation.
 - Build vs Buy: Inclination for custom building with selective buying.
 - Big Tech: Very few are choosing open source first approach.
4. Modality: Most AI success stories are on text-based applications.
5. Human assisted agentic AI use cases, autonomous AI use limited.
6. POC graveyard risk - Security reviews, compliance checks can pose implementation hurdles for client facing applications.
7. Disruptive AI-use possible through enablement not as a replacement for human expertise.
8. Not many firms have an engineering and technology as an operational core.
9. Emphasis on deep research and alpha generation limited in India compared to the global landscape.
10. Awareness and training, key for harnessing the significant advantages AI can offer.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS.....	1
ABBREVIATIONS.....	2
INTRODUCTION.....	3
BRIEF HISTORY.....	3
INDIA AI MISSION.....	5
MOTIVATION AND APPROACH.....	5
AI ADOPTION FRAMEWORK.....	6
HOW IS THE BOOK ORGANIZED?.....	6
LITERATURE SURVEY.....	8
1. INTRODUCTION.....	8
2. TEXT AND LANGUAGE APPLICATIONS.....	10
3. ANALYSIS AND QUANTITATIVE TECHNIQUES.....	12
4. PORTFOLIO CONSTRUCTION.....	15
5. EMERGING CAPABILITIES.....	16
6. OPERATIONAL EFFICIENCIES.....	17
7. TAKEAWAYS FROM ACADEMIC RESEARCH.....	19
ASSET MANAGEMENT COMPANIES.....	21
INDIA MUTUAL FUNDS.....	21
CASE STUDIES.....	22
1. LEVERAGING RESEARCH FUNCTION USAGE THROUGH AI: DSP MUTUAL FUND.....	22
2. MOVING TOWARD AI-ENABLED ENTERPRISE WORKFLOW SYSTEM: SBI MUTUAL FUND.....	25
3. FROM QUANT TO COMPLETION - BUILDING OF DATA LED INVESTMENT SYSTEMS: 360 ONE ASSET MANAGEMENT.....	28
BROKERAGES.....	31
INDIA BROKERAGE FIRMS.....	31
CASE STUDIES.....	33
1. FULL-STACK AI: GROWWS'S AI JOURNEY.....	33
INSURANCE COMPANIES.....	37
INDIA INSURANCE.....	37
CASE STUDIES.....	38
1. FROM BEING EARLY TECH ADOPTERS TO AI WITH COMPLETE INTENT: HDFC LIFE.....	38
2. DIGITAL-NATIVE INSURANCE MEETS AI: ACKO'S TRANSFORMATION JOURNEY WITH AI ASSIST AS FLAGSHIP.....	41
GLOBAL LANDSCAPE, PITFALLS & LIMITATIONS.....	45
GLOBAL LANDSCAPE.....	45
AI LIMITATIONS & PITFALLS - AND HOW TO DEAL WITH THEM.....	47
AI WARNING SIGNS TO WATCH.....	51
GOVERNANCE IN AI.....	52
STAYING PROTECTED: BEST PRACTICES.....	52
GOVERNANCE IN AI: INDIA PRINCIPLES AND GLOBAL STANDARDS.....	54
CONCLUSION.....	57
BIBLIOGRAPHY.....	60

ACKNOWLEDGEMENTS

The authors would like to express their sincere gratitude to all individuals and organizations who contributed to the development of this report. We extend our appreciation to the members and leadership of CFA Society India for their guidance, support, and encouragement throughout the process. We are thankful to industry professionals and practitioners whose insights and perspectives have enriched the analysis presented in this report. We also acknowledge the contributions of researchers and institutions whose work has been referenced and has helped shape the findings and discussions. Finally, we appreciate the efforts of all those involved in reviewing and refining this report, ensuring its quality and relevance.

ABBREVIATIONS

AI	Artificial Intelligence	HNI	High Net Worth Individuals
AI+HI	Artificial Intelligence + Human Intervention	HR	Human Resources
AIF	Alternative Investments Funds	IBM	International Business Machines Corporation of India
AMC	Asset Management Company	ICICI	Industrial Credit and Investment Corporation of India
AML	Anti-Money Laundering	IIB	Insurance Information Bureau
API	Application Programming Interface	IMF	International Monetary Fund
ARQ	Ask Right Questions	IPO	Initial Public Offering
AUM	Asset Under Management	IRDA	Insurance Regulatory and Development Authority
AWS	Amazon Web Services	IT	Information Technology
BCBS	Basel Committee on Banking Supervision	KYC	Know Your Customer
BCG	Boston Consulting Group	LLaMA	Large Language Model Meta AI
BERT	Bidirectional Encoder Representations from Transformers	LLM	Large Language Model
BIS -	Bank for International Settlements	LSTM	Long Short-Term Memory
CDO	Chief Digital Officer	MARS	Metrics, Actionable, Role, Scenario
CEO	Chief Executive Officer	MAS	MULTI AGENT SYSTEMS
CFA	Chartered Financial Analyst	MCP	Model Context Protocol
CFO	Chief Financial Officer	MFD	Mutual Fund Distributor
CIO	Chief Information Officer	MIT	Massachusetts Institute of Technology
CMIE	Centre for Monitoring Indian Economy	ML	Machine Learning
CRM	Customer Relationship Management	MLLM	Multi-Modal Large Language Models
CTO	Chief Technology Officer	NAICS	North American Industry Classification System
CX	Customer Experience	NLP	Natural Language Processing
DIY	Do It Yourself	NSE	National Stock Exchange
DPDP	Digital Personal Data Protection	OCR	Optical Character Recognition
ESG	Environmental, Social, and Governance	OTC	Over-the-counter
ETF	Exchange-Traded Funds	PASA	Pre-Approved Sum Assured
EU	European Union	PMS	Portfolio Management Services
EV	Electric Vehicle	POC	Proof of Concept
EY	Ernst & Young	RAG	Retrieval Augmented Generation
FAJ	Financial Analysts Journal	RBI	Reserve Bank of India
FAQ	Frequently Asked Questions	RFP	Request For Proposal
FEAT	Fairness, Ethics, Accountability, and Transparency	RLHF	Reinforcement Learning from Human Feedback
FLS	Front Level Sales	ROI	Return On Investments
FLS	Front-Line Staff	RPA	Robotic Process Automation
FMs	Finance Managers	SBIL	SBI Life Insurance Company
FOMC	Federal Open Market Committee	SEBI	Securities and Exchange Board of India
FOSS	Free Open-Source Software philosophy	SEC	Securities and Exchange Commissions
FSB	Financial Stability Board	SIF	Specialized Investment Funds
FSI	Financial Stability Institute	SIP	Systematic Investment Plan
FX	Foreign Exchange	SMS	Short Message Service
GenAI	Generative Artificial Intelligence	SQL	Structured Query Language
GICS	Global Industry Classification Standard	SSRN	Social Science Research Network
GIGO	Garbage-in Garbage-out	SWOT	Strength, Weakness, Opportunity, and Threat
GPT	Generative Pre-trained Transformer	TAT	Turnaround time
HDFC	Housing Development Finance Corporation	UHNI	Ultra-High Net Worth Individuals
HDFCL	HDFC Life Insurance Company	USA	United States of America
		WEF	World Economic Forum



INTRODUCTION

*"The world is full of obvious things which nobody by any chance ever observes."
– Arthur Conan Doyle*

Technological innovations appear in waves. From stone tools to using fire for cooking, every invention has contributed to building a foundation for human progress. In the vast sands of time, the latest technology feat – using Artificial Intelligence to augment human work – has started asserting itself. Over time, the changes AI brings are likely to leave behind a definitive impact that influences. Experts believe that AI is likely the biggest technological leap humanity may have known¹. One distinctive dimension AI offers over other technologies is the ability of deep learning models to mimic human thinking. The other related aspect is that the deep learning models are not static; the models can reinforce their knowledge bases and keep learning.

This means that deep learning models have the potential to not only equal but also surpass human ability in an ever-expanding array of human endeavours – not just the well-publicized milestones of AI surpassing humans on chess, the Game of Go, or creative writing.

Enterprises are increasingly integrating AI. Yet substantive progress in disruptive use cases has been illusory². However, the possibilities are real, and the capital investment has been unprecedented. Alphabet, Google's parent company, plans to invest \$75 billion in AI development in 2025, with co-founder Sergey Brin returning from retirement to work on these projects. Investment in AI by "Big Six" USA Technology Companies (Apple, NVIDIA, Microsoft, Alphabet, Amazon (AWS only), & Meta Platforms) saw their Capex increase over 20x from 2014 to 2025, reaching \$420+ Billion³.

There is a race to get access to computing power. With every year, the models are getting better with many mundane tasks now being done by AI. On the other side of the aisle, as AI seeks to establish a new world, we note that there is also the risk that the capex spent, sooner or later, will look for return on the investments made.

BRIEF HISTORY

The starting point in the evolution of AI was during the 1950s when Alan Turing created the Turing test to measure computer intelligence, asserting that computers could think like humans. John McCarthy, a Stanford computer scientist, is said to have convened the first AI conference, coining the term Artificial Intelligence. The first face-off with humans between a human and AI was during the 1960s, when Arthur Samuel, an IBM computer scientist, developed a self-learning computer program that defeated the top USA checkers champion. In the late 1960s, Stanford researchers deployed Shakey, the world's first general-purpose mobile robot that

could apply reasoning to its actions. This was followed by a three-decade-long AI winter. Action began again during the late 1990s, after IBM's "Deep Blue" defeated the reigning world chess champion. A succession of small wins followed through in the 2000s, including Stanley, a driverless car, and Apple's voice assistant, Siri. It was the 2010s that brought significant advancements in neural networks' abilities. The first was the significant increase in the accuracy of neural networks in 2012 in recognizing images. The second was in 2017, which marked the beginning of the use of transformers (neural networks) to understand and process content. Prior to

¹ Mustafa Suleyman, *The Coming Wave*.

² Challapally, Aditya, et al. The GenAI Divide STATE of AI in BUSINESS 2025 MIT NANDA.


³ Chloetjeda, "Big Tech's AI Expansion."

2017, the focus for AI development was on building the right models. Transformers brought in significant change: if you have large enough data to train your model, then the model starts to learn from the data on its own and refines its contextual understanding. From 2017 to 2020, transformers were applied across diverse applications, including natural language processing (NLP), biology, and genetics. During 2021 & 2022, the application of the transformers' technology exploded, starting with the launch of ChatGPT to what we call today as Gen AI, or the use of large language models (LLMs). The ChatGPT launch on November 30th 2022 marked AI's "iPhone Moment." Unsurprisingly, the

traction was swift – ChatGPT could draw one million users in just five days after launch, the fastest user ramp-up ever for a standalone product.⁴

Few technological transitions announce themselves with a regression line, but AI's capability growth has been that rare thing: measurable, consistent, and steep. Since 2019, the complexity of tasks AI agents can handle autonomously has doubled every 207 days. In 2026, tasks that would take a human two hours to complete research, synthesis, multi-step reasoning: are now within reach of autonomous AI execution (Figure 1).

Figure 1: Length of tasks AI agents have been able to complete autonomously⁵



It is important to note that multiple factors have contributed to the development of AI. One of the most important and immediate linkages to success is the global internet infrastructure – over three decades of digital datasets created over time by over 5 billion users, for the models to learn. The practical consequence has been a steady unlocking of use cases. First came automation of the repetitive and rules-bound. Then, as that

complexity ceiling rose, something more interesting emerged: AI as a thinking partner. Capable of synthesising research, stress-testing assumptions, and drafting analysis. For financial services, that shift from tool to collaborator is not on the horizon. It is already underway. India's vast population base and unique demographics can benefit from these technological advancements. Among others, AI-enabled

⁴ Meeker, Mary, et al. "Trends - Artificial Intelligence. BOND, 2025".

⁵ Kwa, West, Becker (2026) "Measuring AI Ability to Complete Long Software Tasks".

apps can create enablers to reach out and serve people through medical help, tutoring, and financial transactions. An EY study, analyzing over 10,000 tasks across critical Indian industries, found that 24% of tasks can be fully automated, and time spent on another 42% can be significantly reduced, freeing up 8-10 hours per week for corporate workers. This could lead to a 2.6% boost in productivity by 2030 in the organized sector, impacting 38 million Indian employees, with an additional 2.8%

boost in the unorganized sector, resulting in combined potential gains of 5.4%. The report further details GenAI's impact across Financial Services (34-38% productivity boost by 2030), Retail, Consumer, and E-commerce (35-37% by 2030), Healthcare (30-32% by 2030), Life Sciences (32-34% by 2030), Technology Services (43-45% by 2030), Auto and Mobility (30-32% by 2030), Industrials and Energy, and Media and Entertainment (15-20% by 2030).⁶

INDIA AI MISSION

The Government of India has launched IndiaAI mission with a financial commitment of over ₹10 bn. The IndiaAI Mission aims to build a comprehensive ecosystem that fosters AI innovation by democratizing computing access, enhancing data quality, developing indigenous AI capabilities, attracting top AI talent, enabling industry collaboration, providing startup risk capital, ensuring socially impactful AI projects, and promoting ethical AI. This mission drives the responsible and inclusive growth of India's AI ecosystem through seven diverse pillars as follows:

1. IndiaAI Compute: Making Compute Power Accessible
2. IndiaAI Application Development: Solving for Bharat
3. AIKosh: India's Own AI Data Repository
4. IndiaAI Foundation Models: Building GenAI in Indian Languages
5. IndiaAI Future Skills: Skilling India for the AI Era
6. IndiaAI Startup Financing: Global Launchpad for Indian AI
7. Safe & Trusted AI: Building Ethical and Explainable Systems⁷

BharatGen, India's first government-supported multimodal LLM initiative, is focused on developing models tailored to Indian languages and cultural contexts, to be used as a public good. Several countries – especially the United States and China – are actively pursuing sovereign AI strategies to enhance their competitiveness. Nations are investing their resources in the most important building blocks of AI infrastructure, data, data sovereignty, workforces and talent development, business network, governance framework, national security etc. among others.

MOTIVATION AND APPROACH

AI offers the potential to significantly augment human capacities and transform the ways we work. Adopting and adapting to the fast-changing canvas of AI advancements is not an option anymore; it is a necessity.

Over the last few years, an increasing proportion of CFA Society India's member practitioners have expressed the need to understand how AI is changing the investment management landscape. Several members have asked what courses they should take to equip and align themselves to newer paradigms and tools for servicing their clients' needs. In one such gathering – an all-India volunteer meeting in April 2025, a majority of the members wanted CFA Society India's Research and Advocacy Committee to take up AI as a topic for research. Members wanted to understand how companies are adopting AI and its implications. This book is an attempt to provide answers to these questions.

⁶ Rajiv Memani, *"The Aldea of India: 2025."*

⁷ NITI Aayog, *"Transforming India With AI."*

Who Should Read this Book?

The book provides a ringside view of how companies are adopting AI to deliver value for their clients. The content can be a useful resource for the following types of audience:

- Financial services practitioners who are keen to understand how AI is being used across different domains
- C-Suite executives for taking informed decisions on strategic initiatives
- Regulators for facilitating policy decisions
- Students to stay abreast of industry development

Approach

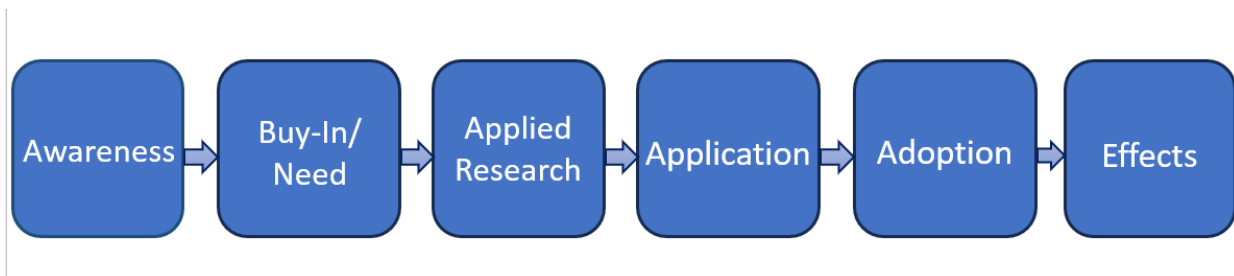
One of the central motivating themes for this book was to build an understanding of how corporates are using AI. To understand corporate use, we leveraged CFA Society India's member network across corporates.

Across each of the three financial services verticals – Asset Management Companies, Brokerages, and Insurance, we approached all major players. This resulted in interviewing 14 different corporate entities. Each interview was structured as an hour-long session. Each interview transcript was then written down into a case study format.

AI ADOPTION FRAMEWORK

Everett Rogers in, his seminal work on Diffusion of Innovations, has noted six stages of the innovation development process⁸. These stages are:

Figure 2: Six Stages of the Innovation Development Process.



As we interviewed practitioners across the financial services industry, we realized that this six-stage view is a handy framework to understand how companies are gearing up to adopt AI. One of the desirable outcomes of AI use, for that matter, any use of software, is the near-zero marginal cost of using the services this automation creates.

Disruption in established practices, lowering of unit costs, and "windfall profits" are among the biggest "consequences" in the sixth stage of Rogers' framework. In this book, we will refer to Rogers' six-stage framework for understanding where the financial services functions are in their AI journey.

HOW IS THE BOOK ORGANIZED?

We adopted a Qualitative research method, specifically utilizing a unique case-oriented approach. This design was selected to facilitate a holistic perspective, allowing the research to be deeply context-sensitive. Core tenets of this methodology included prioritizing the voice and perspective of the participants and maintaining a reflexive posture throughout the data collection and analysis phases.⁹

⁸ Rogers et al., "Diffusion of Innovations".

⁹ Labaree, "Research Guides."



The primary data source consists of detailed case studies gathered via interviews. This primary data was triangulated and contextualized.¹⁰

A stratified purposive sampling method is used to ensure representation across the critical segments of the financial ecosystem. This approach allowed for the deliberate selection of organizations and participants with specific, relevant expertise in the areas under investigation. The target group is defined as Chief AI Officers / Data Labs Head / Key Operational and Strategic Leaders within the highly regulated financial services sector, specifically focusing on 3 distinct segments: Asset Management Companies (AMCs), Insurance Providers, and Brokers

The final sample consisted of 6 companies that participated in the primary data collection (interviews), structured as follows:

- 3 Asset Management Companies (AMCs)
- 1 Brokerage
- 2 Insurance Companies

We collected data using semi-structured interviews. An extensive interview protocol was developed, containing 15 main, exploratory questions for each sector (tailored slightly for contextual relevance), supplemented by several probing questions.

Data collection through interviews started from 2nd July 2025 and concluded on 16th October 2025. These interviews were conducted via secure online platforms and in-person meetings. Prior to commencing, all participants were briefed on the research's objectives and an informed consent.

To enhance the speed and scale of the data, the analysis was further refined using several cutting-edge AI Large Language Models (LLMs) and tools like GPT, Gemini, Thurro, Perplexity, Napkin.ai, and DeepSeek. The separate group of volunteers synthesized data from credible, high-quality, and recent online research reports, which served as comparative, secondary case material.

¹⁰ Carter et al., "The Use of Triangulation in Qualitative Research."

LITERATURE SURVEY

*"This is only a foretaste of what is to come, and only the shadow of what is going to be."
— Alan Turing*

GenAI in Investment Management – Frontier applications, and a practitioner's guide to real-world use.

1. INTRODUCTION

1.1 The Evolution of AI in Finance

Finance has long been at the forefront of technological adoption, driven by the industry's complexity and real-time decision requirements. Some of the earliest usages of computers and statistical analysis were applied to finance, investing, and pricing (CFA FAJ, 1966¹¹), (Black, Scholes 1973¹²).

The 2010-2020 period saw significant acceleration with machine learning models and neural networks enabling algorithmic trading at scale (Gu, Kelly, 2020¹³). The emergence of Generative AI and Large Language Models presents a qualitative shift in capabilities for investment management. Large Language Models introduce fundamentally different

capabilities from previous AI waves. Earlier approaches focused on pattern recognition in structured data. LLMs excel at processing unstructured text: earnings calls, regulatory filings, news, and research reports. LLMs have an inherent understanding of human language, stemming from the use of Transformer architecture and pre-training on massive text corpora. In addition, their large context windows enable analysis of entire documents in a single pass. Practical applications include extracting key information from earnings transcripts, 10-K filings, and legal agreements for OTC or loan transactions (CFA, 2023¹⁴): tasks that previously required extensive manual review.

Table 1: Excerpt from BIS paper for Intelligent AI in Finance¹⁵

Category	Subcategory	Financial Intermediation	Insurance	Asset Management	Payments	Financial Stability
Traditional analytics	Opportunities	Rule-based risk analysis; greater competition		Risk management, portfolio optimization, automated and HFT trading	Fraud detection	Hedging, cascade effects and flash crashes such as the US stock market crash of 1987
	Challenges	Rigid; requires human supervision; small number of parameters, threats to consumer privacy, emergence of data silos		Zero-sum assurances	Technical vulnerabilities	

¹¹ Gal, "Man-Machine Interactive Systems and Their Application to Financial Analysis."

¹² Black and Scholes, "The Pricing of Options and Corporate Liabilities."

¹³ Gu et al., "Empirical Asset Pricing via Machine Learning."

¹⁴ Pisaneschi, "RAG For Finance: Automating Document Analysis with LLMs."

¹⁵ Aldasoro et al., "Intelligent Financial System: How AI Is Transforming Finance."

Category	Subcategory	Financial Intermediation	Insurance	Asset Management	Payments	Financial Stability
Machine Learning	Opportunities	Credit risk analysis; lower underwriting costs; financial inclusion	Insurance risk analysis, lower processing costs, fraud detection	Analysis of new data sources; high frequency trading	New liquidity management tools; fraud detection and AML	Herding/network interconnectedness, lack of explainability, single point of failure, concentrated dependence on third party providers
	Challenges	Black box mechanisms; algorithmic discrimination		Zero-sum arms races, model herding, algorithmic coordination	Increased liquidity crises, increased cyber risks	
Generative AI	Opportunities	Credit scoring; easier back-end processing; better customer support	Better risk analysis with newly legible data, easier compliance	Robo-advising; asset embedding; new products; virtual assistants	Enhanced KYC, AML process	Herding; uniformity; incorrect decisions based on alternative data, microeconomic effects of potential labour and displacement
	Challenges	Hallucinations, increased market concentration, consumer privacy concerns, algorithmic collusion				
AI Agents	Opportunities	Automated design, marketing and sale of new financial products without human intervention		Increase in speed of information processing	Faster payment flows; fraud prevention	Misalignment risks; inherent unsuitability of AI agents for aspects of macroprudential policies
	Challenges	New risks to counterparty protection, cybersecurity, potential overreliance, fraud and unforeseen risks		Cybersecurity, fraud, unforeseen risks, risk concentration with AI agent interactions	Sudden liquidity crises, fraud with deception and unforeseen risks	

Asset management has been particularly receptive to recent AI advances (CFA, 2020¹⁶). Investment managers across the spectrum are actively deploying LLM-based solutions across four core functions: portfolio allocation (asset selection and weighting), trading (execution and strategy), risk management (measurement and mitigation), and advisory services (personalized client guidance).

While AI adoption is growing rapidly in asset management, significant risks remain.

These include hallucinations (factually incorrect outputs), model bias, lack of explainability in black-box models, and high sensitivity to input data quality. Section 6 examines these limitations in detail.

Additionally, distinguishing genuine AI implementation from "AI washing" has become critical. The CFA Institute defines AI washing (CFA, 2025¹⁷) as companies "falsely or inaccurately claiming to leverage AI technologies" without measurable improvements to investment processes.

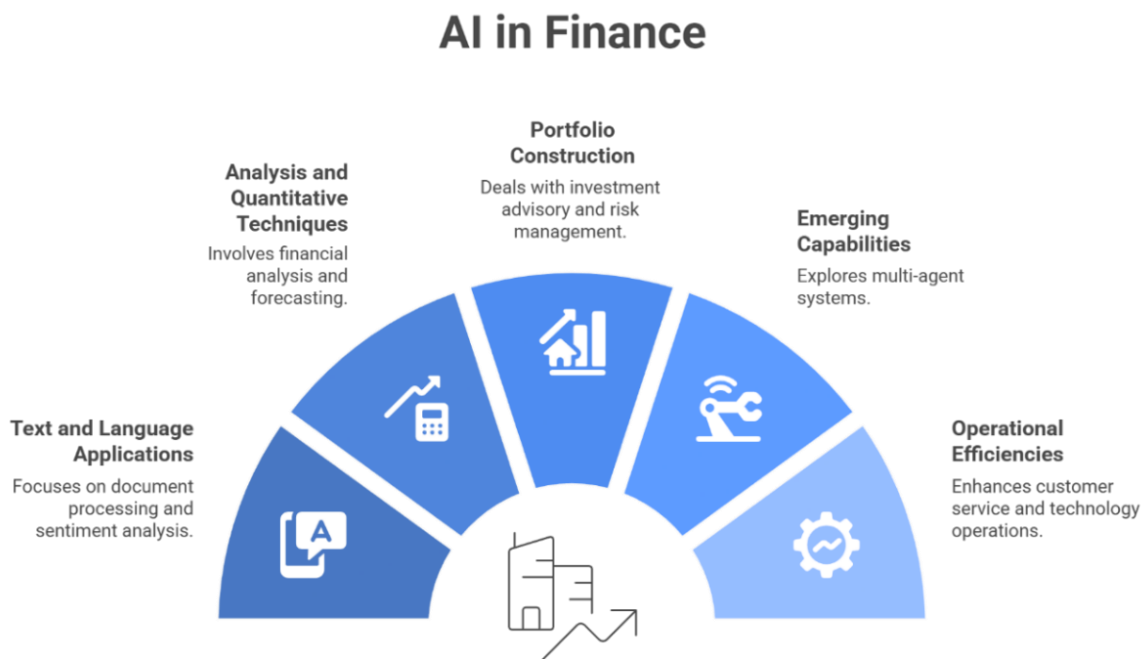
¹⁶ Bartram et al., "Artificial Intelligence in Asset Management."

¹⁷ Simonian, "AI Washing: Signs, Symptoms, and Suggested Solutions for Investment Stakeholders."

1.2 Survey Scope and Methodology

This survey focuses on Investment Management within finance, examining how AI and specifically LLM technologies enhance core investment processes. We draw from academic research (Journal of Portfolio Management, SSRN), institutional analysis (CFA Institute, BIS, FSB, IMF), and industry reports (McKinsey, BCG, WEF). The survey is organized by capability clusters - text understanding, quantitative analysis, information extraction, and integrated systems, and their application into business functions. This reflects how AI is deployed: specific capabilities find applications across multiple investment activities.

Figure 3: Literature Survey Focus areas



2. TEXT AND LANGUAGE APPLICATIONS

2.1 Research and Analysis

The foundation of investment research is the ability to draw inferences from vast information sources like earnings calls, press releases, and market news. Traditionally, this function is time and resource-intensive, requiring analysts to process textual documents (earnings transcripts, analyst reports), numerical data (balance sheets, financial projections), and visual information (charts, graphs). Traditional NLP approaches have been used for document analysis, but they require extensive domain-specific training, struggle with multi-modal inputs, and have difficulty

interpreting complex financial language and context.

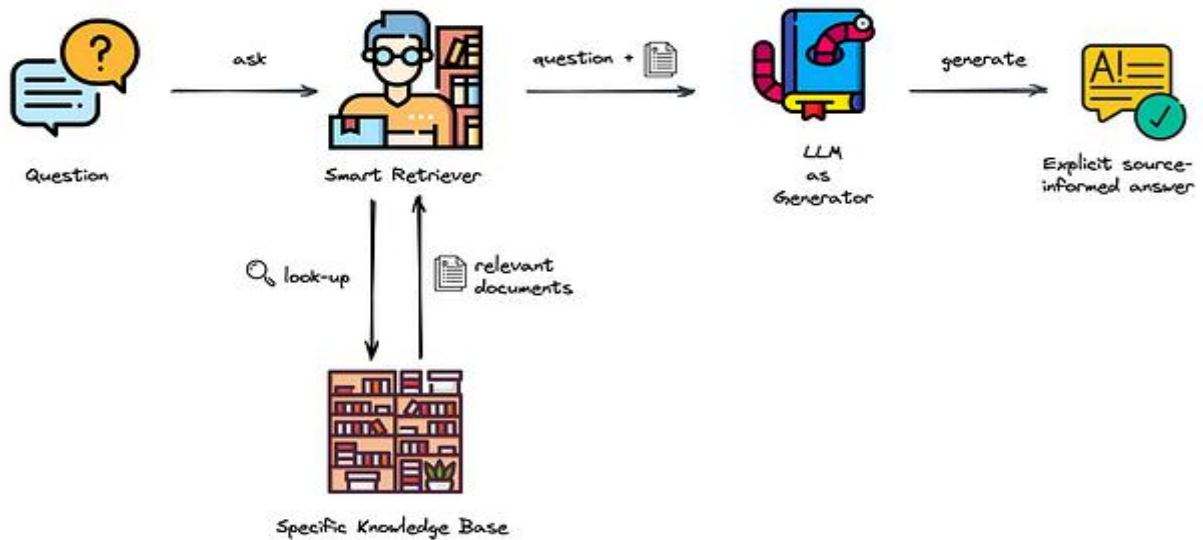
In document processing, recent research shows that LLMs have achieved state-of-the-art results in NLP tasks (Li, Wang, 2024¹⁸). The unique properties of Transformer architecture underlying models like GPT and BERT enable contextual understanding across entire sentences and paragraphs, capturing nuanced meaning in ways traditional keyword-based approaches cannot. For processing large documents, techniques like Retrieval Augmented Generation (RAG) enable effective search

¹⁸ Li et al., "Large Language Models in Finance: A Survey."

and summarization through intelligent text chunking (Yepes, Yu, 2024¹⁹). Specialized models extend these capabilities further: DocLLM (Wang, Raman, 2023²⁰) has been fine-tuned to process visual inputs like graphs and tables alongside text, while BloombergGPT²¹ offers domain-specific

optimization for financial language. Multilingual financial analysis is also being explored through language-specific fine-tuning, enabling analysis of Japanese, German, and other non-English financial documents (Suzuki, 2023²²).

Figure 4: Retrieval-Augmented Generation (RAG) in a nutshell²³



Sentiment analysis is a critical building block for investment decision-making. LLMs excel at deciphering complex language (Lopez-Lira & Tang, 2023²⁴) in news and social media posts, including sarcasm, nuanced expressions, and market-specific jargon (Shen, 2024²⁵). Models like GPT-4o can process data from multiple sources such as Reddit (Deng, 2022²⁶).

Usage of social media, news feeds, audio from earnings calls, and video content can provide robust contextual analysis around market events and FOMC meetings (Yao, 2025²⁷). LLMs can also rapidly extract

sentiment from macroeconomic indicators (payroll reports, unemployment data, inflation figures) and analyst reports, generating valuable signals for investment strategies. However, significant challenges remain. While LLMs are more robust than traditional keyword-based methods against simple manipulation (Leippold, 2023²⁸), they remain vulnerable to adversarial attacks, unreliable sources, and internal model biases. Social media's high volume and noise can distort aggregate sentiment measures, potentially leading to misleading signals.

¹⁹ Yepes et al., "Financial Report Chunking for Effective Retrieval Augmented Generation."

²⁰ Wang et al., "DocLLM: A Layout-Aware Generative Language Model for Multimodal Document Understanding."

²¹ Wu et al., "BloombergGPT: A Large Language Model for Finance."

²² Suzuki et al., "Constructing and Analyzing Domain-Specific Language Model for Financial Text Mining."

²³ Ratnawat, "Exploring Retrieval-Augmented Generation (RAG) with Vector Databases and AI Agents."

²⁴ Lopez-Lira and Tang, "Can ChatGPT Forecast Stock Price Movements? Return Predictability and Large Language Models."

²⁵ "Financial Sentiment Analysis on News and Reports Using Large Language Models and FinBERT."

²⁶ Deng et al., "What Do Large Language Models Know about Financial Markets? A Case Study on Reddit Market Sentiment Analysis."

²⁷ "Interpreting FedSpeak With Confidence: A LLM-Based Uncertainty-Aware Framework Guided by Monetary Policy Transmission Paths."

²⁸ Glaeser, Kim, and Luca, "Nowcasting the Local Economy Using Yelp Data."

2.2 Client Service & Advisory

Personal financial planning has long been challenging to deliver at scale, requiring individualized analysis of each client's financial situation and goals. Robo-advisors and digital platforms have made investing more accessible to the mass-affluent, but personalization has remained limited.

LLM-powered chatbots are transforming this landscape (Lakkaraju, 2023²⁹). These systems provide personalized guidance on credit cards, insurance, budgeting, and tax planning, adapting responses to each user's history and circumstances. Foundation models like GPT-4 and Claude enable communication in multiple languages, including English, Telugu, Afrikaans, and others: breaking down language barriers in financial services.

User data is being incorporated into personal planning, with techniques like RAG

and text-to-SQL helping deliver contextualized answers (Srividya, 2025³⁰). Beyond communication, LLMs are being deployed for investment recommendations. Early research suggests that LLM-based systems can construct diversified portfolios and adapt strategies based on market news and policy announcements, though these applications remain in early stages and require careful validation. However, significant challenges persist. LLMs can exhibit bias in their recommendations and sycophancy: the tendency to agree with users rather than provide objective advice (Bender, 2021³¹). The lack of visual aids for complex financial concepts and persistent numerical reasoning errors presents additional limitations. These challenges, discussed further in Section 6, require careful human oversight and validation before deployment in client-facing role.

3. ANALYSIS AND QUANTITATIVE TECHNIQUES

3.1 Financial Analysis

Financial analysis with LLMs builds on document and sentiment capabilities to produce financial models and company evaluations.

Inferences on Financial Statements: Recent research demonstrates LLMs can draw inferences from annual reports and answer complex questions such as "Does the company have a clear strategy for business partnerships?" (Gupta, 2024³²). This approach uses LLM-generated answers as inputs to machine learning models for investment decisions. Open-source models like FinGPT (2023³³) enable automated financial model creation from balance sheets and income statements, while text-to-SQL capabilities (Kumar, 2024³⁴)

accelerate company analysis through natural language database queries. Many of these tools are available as open-source implementations.

Named entity recognition and similar company identification are classic problems in financial statement analysis, requiring accurate identification of company stock tickers, balance sheet line items, and alignment to classification systems like GICS and NAICS. Defining named entities establishes common feature standards for building both qualitative and quantitative models. Encoder models such as BERT have been used to identify named entities (Pakhale, 2024³⁵), with early research showing improvements when combining

²⁹ Lakkaraju et al., "LLMs for Financial Advisement: A Fairness and Efficacy Study in Personal Decision Making."

³⁰ Srividya et al., "Personalized Finance Chatbot Powered by RAG and Generative AI for Smart Wealth Management."

³¹ Bender et al., "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?"

³² Gupta, "GPT-InvestAR: Enhancing Stock Investment Strategies through Annual Report Analysis with Large Language Models."

³³ Yang, Liu, and Wang, "FinGPT: Open-Source Financial Large Language Models."

³⁴ Kumar, Kommanaboina, and Kumar, "Next-Generation Text-to-SQL: A Survey of Advanced Reasoning Enhancements Techniques."

³⁵ Pakhale, "Comprehensive Overview of Named Entity Recognition: Models, Domain-Specific Applications and Challenges."

LLMs with LSTM models. BlackRock (2023³⁶) has produced research on company classification into GICS and NAICS categories based on SEC filings, using these embeddings for company clustering and risk classification. Key challenges include lack of interpretability and availability bias toward public company data.

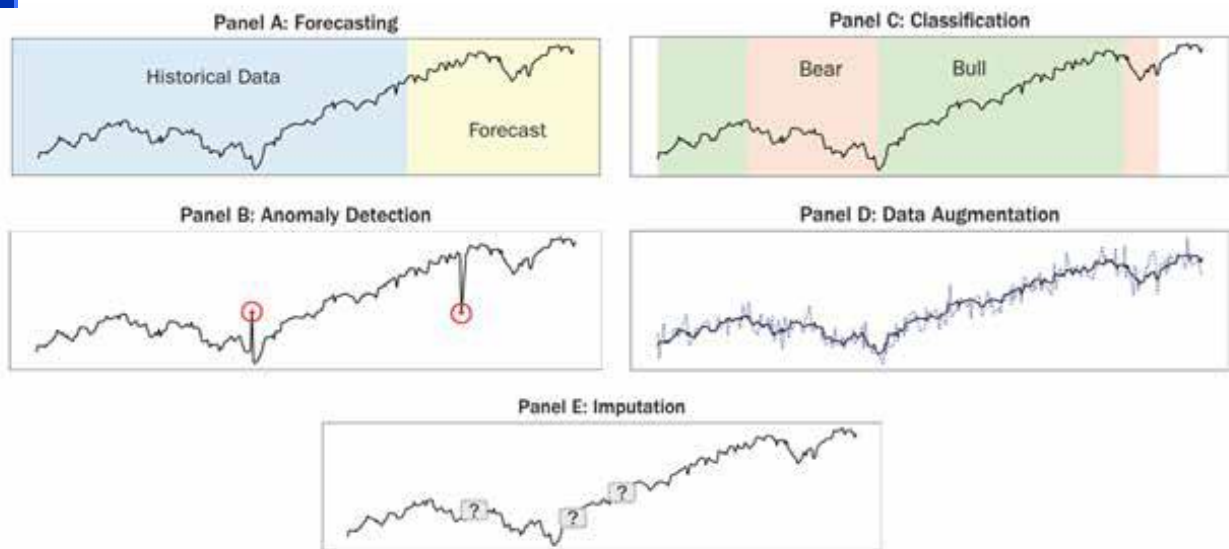
Handling analyst bias, ChatGPT can help reduce analyst optimistic bias (Li, Feng, 2023³⁷) by limiting input context and anonymizing company information. Models

achieve smaller upward forecast errors than human analysts when company names are removed from financial statements. Additionally, reasoning models like GPT-4 Turbo can produce earnings forecasts (Kim, 2024³⁸) using only anonymized company financials, with early results matching state-of-the-art ML models. However, the goal is not to replace human analysts: studies show ChatGPT can produce biased investment recommendations, but rather to enable human-AI collaboration that reduces bias while maintaining holistic judgment.

3.2 Forecasting and Investment Signals

Time series analysis is used extensively to forecast stock price movements, support technical trading, and integrate alternative data sources such as news. The Transformer architecture underlying GPT models enables next-word prediction, and researchers have developed custom transformer architectures specifically for time series tasks, including forecasting, anomaly detection, and classification (Wen, Zhou, 2023³⁹). Foundational models focused purely on time series modeling are being actively developed (Liang, Wen, 2024⁴⁰).

Figure 5: Types of time series analysis, from LLM applications to Finance (Kong, Nie, 2024)



Forecasting results remain mixed but show promise in specific applications. Several studies explore using LLMs like GPT-3 for time series by converting numerical data to text and leveraging the models' next-word

prediction capabilities, sometimes without task-specific training. TimesFM (Das, 2024⁴¹) and LLMTime (Gruver, 2023⁴²) demonstrate this approach for filling missing data and short-term forecasting,

³⁶ Vamvourellis et al., "Company Similarity Using Large Language Models," August 15, 2023.

³⁷ Li et al., "Can ChatGPT Reduce Human Financial Analysts' Optimistic Biases?," November 23, 2023.

³⁸ Kim, Muhn, and Nikolaev, "Financial Statement Analysis With Large Language Models."

³⁹ Wen et al., "Transformers in Time Series: A Survey."

⁴⁰ Liang et al., "Foundation Models for Time Series Analysis: A Tutorial and Survey."

⁴¹ Das et al., "A Decoder-Only Foundation Model for Time-Series Forecasting."

⁴² Gruver et al., "Large Language Models Are Zero-Shot Time Series Forecasters."

with early results showing both successes and limitations. Separately, large models have been used for classifying time series into market regimes such as bullish and bearish periods (Bosancic, 2023⁴³).

Stock price prediction using news shows similarly mixed results. Using a chain-of-prompts approach, one study demonstrates GPT-3 outperforming traditional models like GARCH (Yu & Chen, 2023⁴⁴) by understanding news sentiment, correlating it with price movements, and generating predictions. Similarly, ChatGPT combined

with news headlines has been used to construct long-short portfolios with predictive two-week returns (Lopez-Lira & Tang, 2023⁴⁵), generating alpha by reacting more thoroughly to recent news. Related work shows fine-tuned ChatGPT can use alternative sources like Twitter to predict stock price volatility (Gong, 2024⁴⁶). However, other studies find zero-shot ChatGPT models perform poorly (Xie & Han, 2023⁴⁷), with chain-of-thought approaches and social media data underperforming simple portfolio methods.

BOX 1: Can ChatGPT draw good inferences from news, and aid investment / stock pick choices?

(Lopez-Lira and Tang, 2023⁴⁸)

In this paper, the authors evaluate the ability of ChatGPT 4, 3.5 and other LLMs to predict stock price returns over 1-day holding period. Specifically, they design a simple long-short portfolio for 1 day stock trading, whereby LLMs make predictions based on daily news.

The 4 key takeaways are:

- 1) Best in class LLMs inference are quite good at predicting implications of news headlines on stock price, outperforming than traditional sentiment analysis.
- 2) Better predictive power on small stocks, and negative news
- 3) Using LLMs for determining stock reactions recent events, can operate in arbitrage markets, and over time improve market efficiency by faster reacting to recent changes.
- 4) Use interpretability models to understand "why" behind Blackbox model predictions, can help investors make better portfolio choices.

Authors find that the best-in-class models do possess significant predictive power for economic outcomes. Portfolio was able to generate 38bps per day (300% annualized returns, before transaction costs) with a Sharpe ratio of ~2-3. Authors also develop a reasoning model to explain LLM choices: largely reading news and drawing conclusions based on implications of select types of headlines. For example, share buybacks are a positive signal, director share sales, reverse stock splits etc. Authors also found that while performance is good over 1-2 days news cycles, the performance degrades for long term stock price predictions (1-2 weeks even). This is because markets tend to be less efficient in pricing in short term news, often with days lag.

In summary, LLMs show promise in processing and reacting to news quickly, particularly negative sentiment, but significant development is needed before these approaches can reliably integrate multi-modal data for portfolio construction. Practitioners should experiment with these methods while maintaining rigorous performance benchmarks and comparing results against traditional forecasting approaches. Current evidence suggests LLMs work best as complementary tools rather than replacements for established quantitative methods.

⁴³ Kong et al., "Large Language Models for Financial and Investment Management: Applications and Benchmarks."

⁴⁴ Yu et al., "Temporal Data Meets LLM: Explainable Financial Time Series Forecasting."

⁴⁵ Lopez-Lira and Tang, "Can ChatGPT Forecast Stock Price Movements."

⁴⁶ Gong et al., "Application of Machine Learning in Predicting Extreme Volatility in Financial Markets."

⁴⁷ Xie et al., "The Wall Street Neophyte: A Zero-Shot Analysis of ChatGPT Over Multimodal Stock Movement Prediction Challenges."

⁴⁸ Lopez-Lira and Tang, "Can ChatGPT Forecast Stock Price Movements."

4. PORTFOLIO CONSTRUCTION

4.1 Investment Advisory

LLMs can assist portfolio construction by systematically evaluating diversification metrics. Research shows ChatGPT-constructed portfolios outperform random asset selection through structured diversification criteria, with performance improving as model scale increases (Ko & Lee, 2023⁴⁹). Tools like GPTQuant (Yue & Au, 2023⁵⁰) provide conversational interfaces for evaluating investment strategies, making sophisticated portfolio analysis more accessible to researchers and practitioners. Multiple studies demonstrate

portfolio construction capabilities. ChatGPT can build diversified portfolios that, when paired with real-time news and policy analysis, generate alpha with Sharpe ratios ranging from 1.8 to 2.3 (Schneider, 2025⁵¹; Anic, 2025⁵²). These systems are being developed for integration with robo-advisory platforms, focusing on transparency, fairness, and adaptability. However, practitioners should note that these results require validation across longer time periods and different market regimes.

4.2 Risk Management

Risk management represents a significant emerging application area for AI. Applications span multiple domains:

- **Anomaly Detection:** Time series anomaly detection benefits from novel approaches, including encoding time series as images for pattern recognition (Zhou & Yu, 2025).
- **Credit Risk:** LLM-based credit scoring models show promise in improving fairness and reducing bias compared to traditional models (Feng & Dai, 2023), with preliminary evidence of reduced disparate impact in lending decisions.
- **Financial Risk Modeling:** LLMs enable integration of alternative data sources for measuring and correlating business risks, expanding beyond traditional financial metrics (Cao & Jiang, 2024).
- **Fraud Detection:** BERT-based models can identify fraud-related language patterns in financial disclosures (Yang, 2023), improving detection of suspicious reporting.
- **Corporate Governance:** LLMs can process detailed financial statements to identify themes around political risk, climate risk, and corporate governance issues (Kim, 2024), accelerating comprehensive risk assessment.

Practitioners should note that while these applications show promise, they introduce model-specific risks, including bias, hallucinations, and overconfidence that require careful monitoring and validation.

4.3 Trading Applications

In developing algorithmic trading strategies, LLMs assist with both strategy research and implementation. Advanced models like GPT-4 Turbo can generate code for trading algorithms that match human-coded

implementations (Alonso, 2024⁵³), though success depends heavily on well-designed prompts and detailed technical specifications. Systems like AlphaGPT (Yuan, 2024⁵⁴) demonstrate effective

⁴⁹ Ko and Lee, "Can ChatGPT Improve Investment Decision?"

⁵⁰ Yue and Au, "GPTQuant's Conversational AI: Simplifying Investment Research for All."

⁵¹ Schneider and Yilmaz, "Stock Portfolio Selection Based on Risk Appetite."

⁵² Anic et al., "ChatGPT in Systematic Investing—Enhancing Risk-Adjusted Returns with LLMs."

⁵³ Nogueira Alonso and Dupouy, "Evaluating LLMs in Financial Tasks."

⁵⁴ Yuan, Wang, and Guo, "Alpha-GPT 2.0: Human-in-the-Loop AI for Quantitative Investment."

human-AI collaboration for strategy development, enabling researchers to explore and synthesize trading ideas more rapidly. These systems can incorporate memory and recall capabilities for iterative strategy improvement (Yi & Li, 2023⁵⁵).

Sentiment-based trading strategies represent another application area. Research shows GPT-based models processing news and social media sentiment can aid in constructing trading

strategies with Sharpe ratios exceeding 2.5 in backtests (Kirtac, 2024⁵⁶), though these results require validation across different market conditions and longer time horizons. However, practitioners should note several limitations: LLMs struggle with real-time market data processing, lack understanding of market microstructure, and may generate strategies that overfit to historical patterns. Additionally, as these tools become widespread, any information advantages may diminish through market competition.

5. EMERGING CAPABILITIES

5.1 Multi-Agent Systems

Multi-agent systems (MAS) represent an emerging paradigm where multiple specialized LLM agents collaborate to solve complex tasks. Each agent operates with a specific role: such as data analyst, editor, or research expert—while sharing information and coordinating actions for collective decisions. This architecture mimics human team collaboration and is being explored across several financial applications:

- **Trading and Investment Management:** Systems deploy specialized LLM agents as fundamental analysts, quantitative researchers, risk managers, and traders. These agents communicate and coordinate to make investment decisions, showing improvements over single-model approaches in Sharpe ratios and drawdown control (Xiao & Sun, 2025⁵⁷). TradingGPT demonstrates this architecture with agents maintaining short- and long-term memory, debating strategies, and adapting portfolios dynamically (Li & Yu, 2023⁵⁸).
- **Investment Research:** Multi-agent systems can divide complex research workflows across specialized agents handling 10-K analysis, fundamental research, market sentiment evaluation, and risk assessment. Research shows these systems can generate target price estimates and investment recommendations (Han & Wang, 2024⁵⁹), with multi-agent approaches outperforming single-agent models, though performance varies significantly based on agent configuration and role definitions.
- **Credit Evaluation:** systems assign roles including credit analyst, reward modeler, risk manager, and debt analyst. Research shows that creating “layers” of interaction from analyst to senior decision-maker can improve assessment quality (Jajoo, 2025⁶⁰).
- **Risk Management:** Multi-agent approaches enhance the detection of irregularities in financial reports (Park, 2024⁶¹). Agents take on specialized roles for data analysis, cross-checking, and pattern identification, leveraging different analytical perspectives to identify gaps in financial data.

⁵⁵ Yu et al., “FinMem: A Performance-Enhanced LLM Trading Agent With Layered Memory and Character Design.”

⁵⁶ Kirtac and Germano, “Sentiment Trading with Large Language Models.”

⁵⁷ Xiao et al., “TradingAgents: Multi-Agents LLM Financial Trading Framework.”

⁵⁸ Li et al., “TradingGPT: Multi-Agent System with Layered Memory and Distinct Characters for Enhanced Financial Trading Performance.”

⁵⁹ Han et al., “Enhancing Investment Analysis: Optimizing AI-Agent Collaboration in Financial Research.”

⁶⁰ Jajoo, Chitale, and Agarwal, “MASCA: LLM-Based Multi-Agents System for Credit Assessment.”

⁶¹ Park, “Enhancing Anomaly Detection in Financial Markets with an LLM-Based Multi-Agent Framework.”

These systems show initial promise but remain in early stages. Key limitations include amplification of underlying LLM risks: bias, hallucinations, and data limitations can cascade across multiple agents. Practitioners should experiment with multi-agent approaches for complex, multi-step workflows while implementing rigorous validation of system outputs and monitoring for compounded model risks.

6. OPERATIONAL EFFICIENCIES

While Sections 3-5 examined AI applications in core investment functions, most firms begin their AI journey with operational use cases. These applications offer lower risk, faster ROI, and help organizations build AI capabilities before deploying them in investment-critical functions. Based on case studies and research, three areas show particular promise as initial adoption targets:

6.1 Customer Service

Asset managers handle repetitive investor queries about account performance, document requests, and transaction confirmations. Call centre automation and chatbots have long become standard practice to address this need. LLMs automate routine responses while routing complex inquiries to advisors, with AI-assisted representatives showing 14% productivity improvements (Brynjolfsson, 2023⁶²). Traditional chatbots struggle with context and reliability, but models like GPT-4 provide more nuanced responses (Wulf, 2024⁶³). Advanced implementations use RAG to incorporate firm-specific data while maintaining compliance oversight. Implementation typically begins internally before client-facing deployment, allowing validation of accuracy and regulatory review of communications.

6.2 Communications and Content

Investment firms face intensive content demands: RFP responses, quarterly commentaries, client materials, and regulatory filings. LLMs can draft these documents (Raza, 2025⁶⁴), with applications showing reduced time to produce draft responses and generate reports. Beyond text, emerging image and video models enable social media content creation. These tools improve efficiency and enable personalization for market segmentation (Brand, 2025⁶⁵). However, challenges remain around misinformation risks, deepfakes, and biased outputs. Models currently face high compute costs and limited multilingual support, though both are improving. Human review remains essential for compliance and accuracy.

6.3 Technology Operations

LLMs accelerate development of internal financial tools: analytics dashboards, modeling libraries, data pipelines, and risk systems. Advanced models like GPT-4, LLaMA, and Claude generate code across Python, R, SQL, and other languages used in finance (Haque, 2025⁶⁶). These tools assist with test creation, debugging, documentation, and system design. Platforms like GitHub Copilot and Cursor are widely integrated, improving developer velocity (Rasnayaka, 2024⁶⁷).

⁶² Brynjolfsson, Li, and Raymond, "Generative AI at Work."

⁶³ Wulf and Meierhofer, "Exploring the Potential of Large Language Models for Automation in Technical Customer Service."

⁶⁴ Raza et al., "Industrial Applications of Large Language Models."

⁶⁵ Brand, Israeli, and Ngwe, "Using LLMs for Market Research."

⁶⁶ Haque, "LLMs: A Game-Changer for Software Engineers?"

⁶⁷ Rasnayaka et al., "An Empirical Study on Usage and Perceptions of LLMs in a Software Engineering Project."

Limitations arise with complex business logic: models excel at boilerplate code but struggle with sophisticated financial calculations. Strong human-in-the-loop workflows are essential: detailed specifications, rigorous testing, and code review. While AI generates significant code volume, evidence remains mixed on whether this truly reduces development time and resources.

BOX 2: Productivity effects of LLMs in writing tasks (MIT)

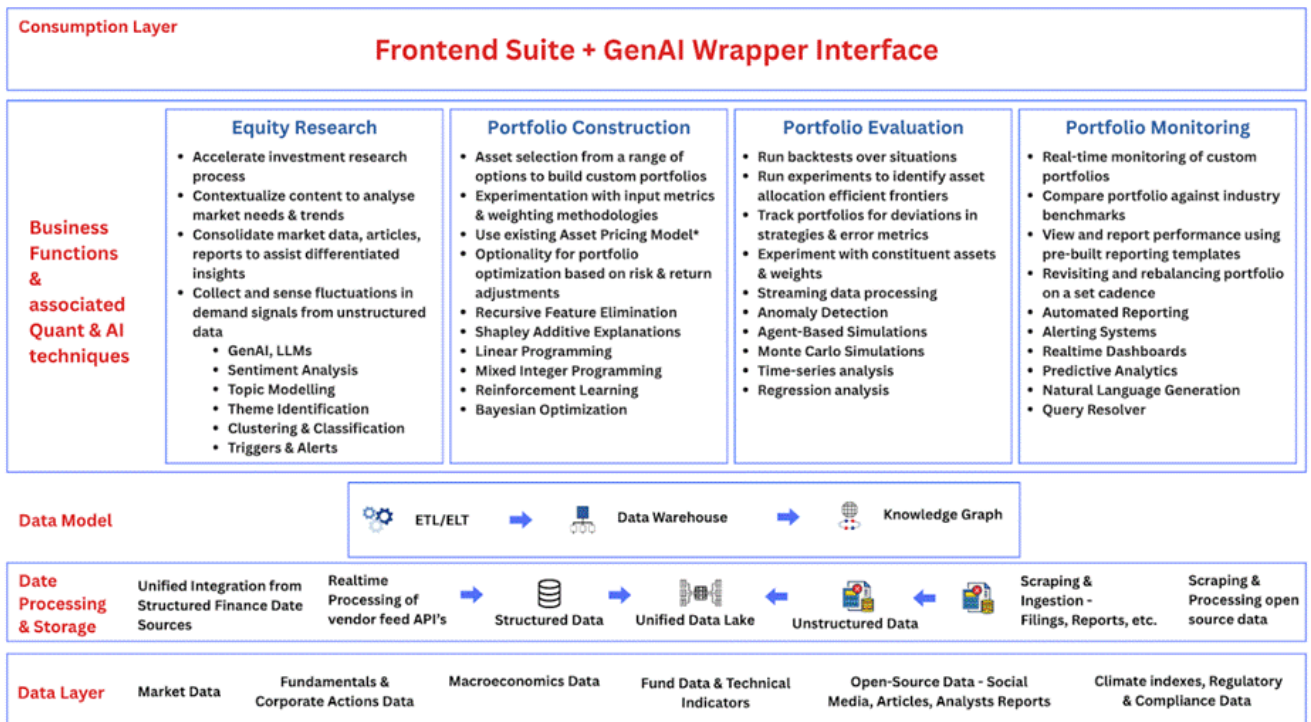
(Noy and Zhang, 2023⁶⁸)

Based on the study across 500 sampled individuals concerning the productivity effects of the generative AI chatbot ChatGPT, the dramatic improvements are concentrated in mid-level professional writing tasks, while limitations are noted for tasks requiring context-specific knowledge

Task Restructuring: ChatGPT primarily functions by substituting for worker effort rather than complementing worker skills. It substantially changes the structure of writing tasks by:

- Automating Rough-Drafting: The share of time spent writing a rough draft fell by more than half. ChatGPT automates the relatively routine, time-consuming subcomponents of writing, such as translating ideas into an initial rough draft.
- Shifting Human Focus: The share of time spent on idea-generation (brainstorming) and editing more than doubled post-treatment
- Time Savings: The time taken on the post-treatment task decreased by 10 minutes (37%) compared to the control group, representing a decrease of 0.8 standard deviations (SDs).
- Quality Increase: Average evaluator grades increased by 0.45 standard deviations.

Figure 6: Short overview of Quant and AI in Portfolio management



⁶⁸ Noy and Zhang, "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence."

7. TAKEAWAYS FROM ACADEMIC RESEARCH

7.1 The Current State: Maturity and Opportunity

Investment management's AI inflection point, introduced in Section 1, has resolved into a clear maturity spectrum. After examining text processing (Section 2), quantitative applications (Section 3), advanced capabilities (Section 4), and implementation pathways (Section 5), three distinct categories emerge:

Proven applications deliver measurable ROI today. Document processing and summarization reduce analyst review time by 40-60% (Li, Wang, 2024). Client service chatbots deflect 40-50% of routine inquiries with 14% productivity improvements for human advisors (Brynjolfsson, 2023). These operational applications justify investment for most asset managers with significant document processing volume, research operations, or client service requirements. The technical risk is low: foundation models (GPT-4, Claude) handle these tasks reliably with appropriate human oversight. The organizational risk: user adoption and workflow integration remains the primary implementation challenge.

Emerging applications show promise but require careful validation. Sentiment analysis from news and social media demonstrates potential for improved alpha generation in early implementations, though results largely reflect academic findings that do not account for transaction costs and competitive dynamics. Financial statement analysis and earnings forecast generation approach human analyst accuracy while reducing optimistic bias on standardized tasks. Risk monitoring

systems identify anomalies and provide early warnings. Alternative data sources enable deeper analysis and market simulation capabilities. Firms are encouraged to develop human-AI collaboration workflows, using AI to brainstorm and backtest ideas in financial analysis and forecasting, while maintaining human decision authority on all investment actions.

Experimental applications remain research-stage with cautious optimism. Multi-agent trading systems, autonomous portfolio construction, and market simulation show intriguing academic results but face substantial validation challenges, computational costs, and regulatory uncertainty. LLM-based forecasting produces mixed results, with performance highly sensitive to time horizon, asset liquidity, and market regime. Firms with appropriate resources and expertise should experiment with these approaches while ensuring rigorous testing, performance benchmarking, and independent verification of model outputs. Across all maturity stages, firms must address fundamental LLM limitations: hallucinations (generating plausible but incorrect information), bias in recommendations, and sycophancy (tendency to agree with users rather than provide objective analysis). Section 6 details these challenges and mitigation strategies. Success requires establishing governance frameworks, validation procedures, and phased training programs that account for these risks while capturing operational and analytical benefits.

7.2 Phased Implementation Recommendation

Analysis of implementations across the industry shows a sequenced approach works best. Firms attempting to skip directly to experimental investment AI typically face regulatory pushback, accuracy issues, and organizational resistance. Additionally, even though AI is touted as plug-and-play, integration requires substantial training, testing, validation, and governance. A strategic assessment of ROI and risk is best achieved through phased rollout:

Start with internal operations (lowest risk): Code generation for tech teams, report drafts, document summarization, and data

automation. This builds capabilities and confidence with contained failure impact.

Progress to client operations (medium risk): Investor relations chatbots, client report generation, and meeting materials—all with human oversight and compliance validation. **Advance to investment support (highest risk):** Financial analysis, sentiment monitoring, and strategy testing, while maintaining human decision authority on investment actions.

This strong operational foundation proves essential for successful investment deployment.

ASSET MANAGEMENT COMPANIES

"The heaviness of being successful was replaced by the lightness of being a beginner again."
— Steve Jobs

INDIA MUTUAL FUNDS

Landscape Ripe For AI/ML Adoption

Indian mutual fund industry has grown at a brisk pace amidst growing disposable household income, rising financial literacy, operational convenience to invest, conducive regulatory environment and widening distribution channel reach. Mutual fund investor base has swelled to 60 million in March 2026 2025 from 38 million in March 2023. Nearly 70% of this growth was on account of investors in Tier III locations – which makes it imperative for mutual funds to make use of digital means to reach and serve the vast hinterlands of the country.

Mutual funds have introduced Systematic Investment Plan (SIP) with minimum requirement amount starting at INR 250 (referred as '*Chhoti SIP*'). This is in line with SEBI's plan to introduce 'sachetization' of mutual funds to facilitate participation from first-time investors in low-income groups.

Rising financial literacy and easy to use fintech apps that track fund performance parameters means that consistent fund performance has become increasingly paramount for sustained success for mutual funds in India. Advent of passive-focused AMCs, having large in-house customer base, such as Zerodha, Navi, Angel One has compounded the need for alpha generation for active funds. A few more fund houses including Jio-Blackrock with a differentiated investment approach are ramping up their scheme launches. This need along with rise in availability of numerous technology tools has led mutual fund houses to move towards a systematized approach to gather,

store, retrieve and consume their in-house research and portfolio management expertise.

Mutual funds have so far mainly adopted AI/ML tools in customer profiling, investment suitability, customer service, investment analysis/research & portfolio construction.

In coming years, we foresee

- rising adoption of AI tools for customer self-service for routine queries/tasks such as fetching account statement, processing address change, obtaining fund information, answering taxation queries. Rise in popularity of low-cost schemes will also necessitate automation in many areas in order to keep operating costs low.
- In investment research, with better availability of alternative data, such data can be combined with traditional fundamental, market and economic data to drive investment decisions.
- Quant-like approach for all active funds. Developments in Specialized Investment Funds (SIF) category space that can have L/S positions will accelerate this trend.
- Multi-asset schemes with commodity, equity and fixed income tend to be amenable to ML techniques driving their usage.

CASE STUDIES

1. LEVERAGING RESEARCH FUNCTION USAGE THROUGH AI: DSP MUTUAL FUND

Contributors: Vipin Vijay, Yogesh Bhalla, Aman Garg

Asset Class Coverage: Equity, Fixed Income, Hybrids, ETFs

Exhibit 1: AI/ML Maturity Assessment at DSPMF (Based on stakeholder interviews, 2025)

	Customer Service	Sales & Marketing	Risk & Compliance	Investments
Automation				
Data science				
AI/ML				

Background

DSP Mutual Fund (DSPMF), the mutual fund arm of DSP Asset Managers Private Limited, managed INR 2,106 billion in assets under management (AUM) as of 30 September 2025. The fund house offers a diverse product suite, spanning equity, fixed income, hybrid, and exchange-traded funds (ETFs).

DSPMF fosters technology-oriented culture and has built internal tech platforms for critical functions in AMC business -

- Investments - ARQ (erstwhile Jarvis)
- Sales - RMX and Disha
- Operations - TITAN
- Compliance - HawkEye
- Risk

Start of the AI/ML Journey

DSPMF's adoption of AI/ML began with an intent to enhance productivity in the Investments team. Early initiatives in AI included summarization of transcripts & internal notes, Q&A from documents etc. These were initially used in an ad hoc fashion but gradually became systematized and embedded in the research and portfolio management process.

The firm's quant strategies provided the first testbed for AI/ML. Supervised learning models were deployed for factor-based stock ranking (Quality, Value, Growth, Volatility), screening, and back-testing. This period laid the foundation for closer collaboration between the Investments function and the internal Technology team.

Investment Research and Portfolio Management Applications of AI

A breakthrough came with the integration of AI/ML technology in internal research system ARQ, enhancing it as a prompt-based application that enables portfolio managers and research analysts to interrogate internal datasets—such as analyst recommendations, earnings estimates, notes and call transcripts.

Use cases included queries like: *“Which stock under coverage has the highest projected earnings growth?”* or *“Summarise earnings call transcript”*. Over time, AI engine of Jarvis expanded to incorporate external datasets (web crawling, company

filings and other publicly available documents). Now, the system can do elementary analysis and write notes like a junior analyst, can configure customized alerts helping analyst and fund managers act promptly in a systematic manner.

The tool also handles scenario analysis for fundamental variables and economic parameters - assisting fund managers toward timely positioning of their portfolio for changes in external environment. Today, the evolved platform—renamed *“Ask Right Questions”* (ARQ)—is widely used across investment teams. DSPMF's ability to scale

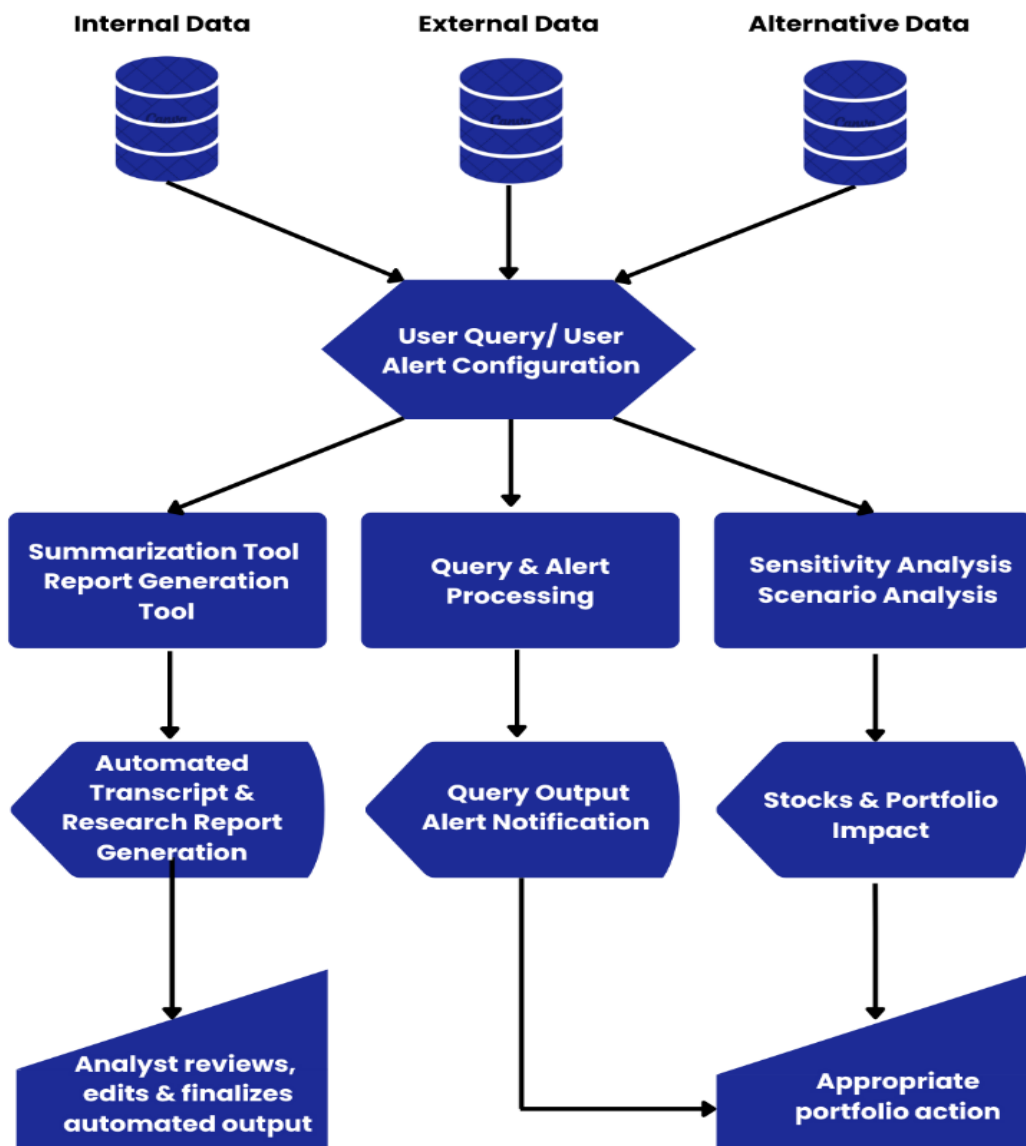
these tools rests on its tech-oriented culture, fostered by its promoters and embraced by staff. A 120-member technology team (within an overall headcount of 650) supports innovation.

The fund house has appointed a Chief AI Officer (reporting to the CTO) and designated AI champions in each department. This federated model allows

business units to raise problem statements, which are then translated into prototypes, tested, and scaled in iterative cycles.

Concerns over job redundancy were addressed early. Employees who adopted AI-enabled tools moved into higher-value roles, with many reporting that automation enhanced their career progression rather than threatening it.

Figure 7: Varied Applications for ARQ



Build vs. Buy

To date, DSPMF has primarily pursued an in-house build strategy. However, with the maturation of external AI solution providers, the fund house has begun experimenting with selected third-party applications in niche areas.

Beyond Investments: Enterprise-Wide Applications

AI/ML usage at DSPMF now extends well beyond core investment management:

- **Retail Sales & Distribution:**
AI tools generate customized talking points for Mutual Fund Distributor (MFD) meetings, provide Market Vs DSPMF, Brokerage analytics, whom to meet and what to talk, Lead conversion insights, and enable LLM-based chat interfaces.

In RMX, DSPMF does analytics using LLMs and solves many sales critical problems like the voice-based chatbot ` 'Disha' answering questions related to sales metrics over the years. Through Netra, Tathya, Transcript and Converse related insights, customers can get answers instantaneously on fund performance and their portfolios. RMXgpt bot, which works on structured data and numbers, swiftly answers queries like-

- *"Show me distributors that I have met more than 3 times since June"*
- *"Give me list of SIP churners in Bengaluru"*
- *"Show me distributors that saw a rise in Duration from the first Quarter to the second"*

LLM based use cases also include voiceover of documents for Marketing like *Netra* in very interesting formats like an interview (in male and female voices) which help DSPMF's customer understand the philosophy and conviction of data driven approach easily. These are available on YouTube in three different languages viz, English, Hindi and Gujarati.

- **Mid- and Back-Office:**
Transaction-heavy processes, once manual, are now AI-automated with human-in-the-loop oversight.
- **Risk & Compliance:**
HawkEye is another enterprise level software which analyse market abuse and detect any suspicious transaction or discussion in verbal form in four different languages viz, English, Hindi, Gujrati and Marathi and flags-off to compliance team for further investigation.

AI-powered anti-black marketing tools monitor social media channels to detect and neutralize impersonators in real time.

- **Customer Service:**
Investor communication has been reimagined through conversational AI. Fund manager letters (monthly or event-driven, e.g., Union Budget) are converted into voice-based updates, resonating with investors familiar with consumer apps and social platforms.

Key Takeaways

- **Culture matters:**
DSPMF's success reflects its leadership's emphasis on tech-first thinking and department-level AI champions.
- **Iterative innovation:**
AI prototypes move through rapid test-pilot-scale cycles, ensuring practical adoption.
- **Talent strategy:**
A large in-house technology team, anchored by a Chief AI Officer, enables DSPMF to balance build vs. buy decisions.
- **Enterprise integration:**
AI/ML adoption has extended beyond investments into sales, risk, compliance, and customer engagement.
- **Human + AI complementarity:**
By addressing fears of redundancy, DSPMF has repositioned staff into higher-value roles, reinforcing the "AI + HI" model.

2. MOVING TOWARD AI-ENABLED ENTERPRISE WORKFLOW SYSTEM: SBI MUTUAL FUND

Contributors: Srinivas Jain

Asset Class Coverage: Equity, Fixed Income, Hybrids, Solution-oriented schemes, ETFs

Exhibit 2: AI/ML Maturity Assessment at SBIMF (Based on stakeholder interviews, 2025)

	Customer Service	Sales & Marketing	Risk & Compliance	Investments
Automation				
Data science				
AI/ML				

Background

SBI Mutual Fund (SBIMF) is the mutual fund arm SBI Funds Management Ltd, a Joint Venture between SBI and AMUNDI (France), one of the world's leading fund management companies. SBIMF managed INR 12,046 billion in assets under management (AUM) as of 30 September 2025. The fund house offers a wide product set, traversing equity, fixed income, hybrid, solution-oriented schemes and exchange-traded funds (ETFs).

Start of the AI/ML Journey

SBIMF initially started experimenting with ML methods in early 2021. The start was made with investment use cases. Coding was carried out in their R and Python labs by in-house development team for these applications. First efforts were made with an intention to find a method to best leverage the large pool of in-house research that was thus far dispersed in silos. From the result of those efforts was born Knowledge Management Portal called "Neo".

Investments Function - 'Neo': One platform that rules them all

SBIMF now has a well-established in-house product called "Neo", an all-in research platform that functions as a knowledge management solution for the 70-odd investments team members. It has now become the go to tool for enhancing & leveraging SBIMF's in-house research - acting as a repository with easy-to-use interface. It can be utilized to pull out/generate full-fledged research reports, management interaction notes/Concall transcripts, analyst recommendations. Besides automation of routine tasks like notes management, research reports generation and customized notification alerts, Neo is being extensively applied for back testing, scenario analysis, factor modeling and asset allocation. Neo has a feature which allows autonomously generates and pushes notifications sent via MS Teams channel to fund managers (FMs). The notification is an event-based actionable given to FMs having consideration to SBIMF's portfolio composition.

SBIMF follows a quant model approach with human in the loop in all its actively managed funds. The fund house has an in-house quant specialist who develops rule-based investment strategies that incorporate its framework that also considers ESG parameters. Supervisory learning methods such as regression are used for building and implementing these models. These models are iterated periodically to remain constantly relevant. SBIMF has several vendors who supply a wide range of alternative data that are captured and utilized in its investment decision making process. Neo's usage extends beyond the typical research and portfolio management use cases. The fund house's broker voting process which was earlier maintained in excel sheet is now built into the platform - making it far more transparent and efficient.

Beyond Neo, SBIMF is working on several projects - some of them have reached prototype stage while others are in various stages of build.

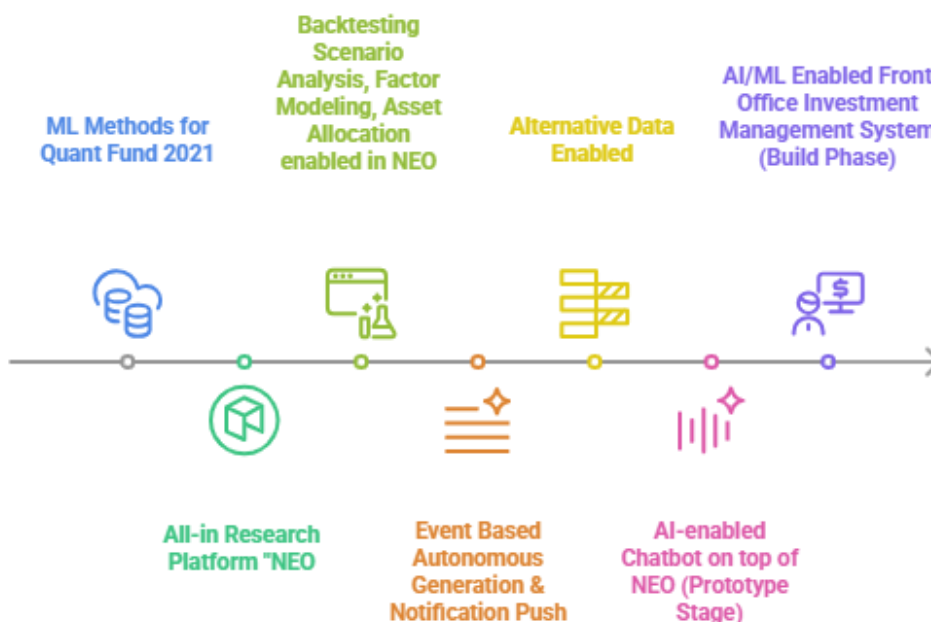
Projects in prototype or POC stage

- AI-enabled chatbot on top of Neo. It facilitates use of SBIMF’s internal research data and external data. This can be used to generate actionable insights through query interface for use by FMs. External data is sourced from various published sources (such as broker reports, CMIE/Cline, government data) as well as alternative data. Alternative data is collected from both in-house research staff and external data vendors. These alternative data sets are obtained through website scrape, online portal customer product reviews, social media posts, earnings concall transcripts. This chatbot is currently ready as a prototype. This is being developed in-house in partnership with a US-based AI startup while residing on SBIMF’s Azure private data Cloud.
- Investment risk function use case - the tool scrapes trade, pre-trade and post-trade through voice/email/text information using mistral-based engine which is developed using a fintech partnership. This project has achieved prototype status now and will go in production in soon.

Projects in build phase

- Data lake - SBIMF is in the process of building a data lake that functions with Investor 360 approach. While current tools being used can do transcript generation after the conclusion of an earnings concall or a management meeting, data lake will enable live update of the concall/meeting and can generate notifications live which will assist in taking timely action in portfolios by FMs. Similarly, portfolio attribution and portfolio analysis can be conducted using data lake. PWC is consultant and a fintech is partner for SBIMF in this project.
- Operational efficiency projects - In reconciliations, both automations use cases and machine learning are being developed.
- CIO dashboard
- Front office investment management system - This is being developed along with Amundi (SBIMF’s JV partner in India). This is an order execution system and risk management system similar to Blackrock’s Alladin and built keeping in mind the requirements of India set up. This system when fully developed will have AI/ML capabilities that are found wanting in traditional fund management applications provided by vendors in India and used by most funds in the country.

Figure 8: SBIMF's AI/ML Progression.



Using AI: An integrated approach

- **Customer service function:**
SBIMF operates a marketing cum customer service tool which is an LLM-based WhatsApp chatbot. The chatbot runs on OpenAI, Llama 3.4. ChatGPT is used for complex queries. For instance, the chatbot answers investor queries as varied as “*What is the most suitable investment scheme for me?*”, “*What are the tax implications if I invest in this scheme?*”, “*How do I record change in my address?*”. The chatbot addresses these queries taking into consideration the particular investor’s portfolio and risk appetite – thus integrating well into the customer service function apart from being an intuitive marketing tool.
- **Marketing team use cases:**
Scripts are generated using AI based on input information given, this is useful to front-line sales teams and call centers.

Raising awareness and encouraging AI adoption

Initially, staff were both excited about the prospects of AI as well as worried that it might result in job redundancies. SBIMF addressed these apprehensions through several awareness initiatives.

The fund house periodically invites external experts (Microsoft, BCG) and arranges training workshops (LLM and GenAI being latest) to up the ante on its AI/ML applications usage. The fund house is utilizing a consultant to measure and publish use of AI in the organization. This showcases to the staff on how important use of AI/ML is to SBIMF’s management – and hence adoption levels rise as staff get involved in these events and read about their colleague’s experience in use of AI/ML.

SBIMF is currently building a data science team that will support each department with models and data management.

Typically, for investment use cases, the investments team comes with use case request to the quant team and data science team. Then, through numerous iterations a solution is reached for prototype.

All this has happened with a supportive Board and Management that has devoted time and financial resources toward technology investments and usage.

Key Takeaways

- **Collaboration with external vendors:**
SBIMF’s progress stems from its agnostic and open approach to embrace technology investments be it from in-house build or external vendor partnerships.
- **Structured and Periodic engagement with staff:**
In order to raise adoption of AI/ML and come up with new applications, SBIMF conducts periodic training workshops structured keeping in mind the needs of its staff.
- **Moving toward an “Integrated workflow system”:**
By making these technology investments and encouraging their use, SBIMF is moving toward an AI/ML enabled integrated workflows system encompassing research, order execution, risk management, customer service.
- **Talent strategy:**
A large in-house technology team, along with building a dedicated data science team.

3. FROM QUANT TO COMPLETION - BUILDING OF DATA LED INVESTMENT SYSTEMS: 360 ONE ASSET MANAGEMENT

Contributors: Mehul Jani, Ashish Ongari, Parijat Garg, CFA

Asset Class Coverage: Listed Equity, Long-Short, PMS

Exhibit 3: AI/ML Maturity Assessment at 360 ONE (Based on stakeholder interviews, 2025)

	Customer Service	Sales & Marketing	Risk & Compliance	Investments
Automation				
Data science				
AI/ML				

Background

360 ONE Asset Management (360 ONE AMC), the AMC arm of 360 ONE WAM, managed INR 921 billion in assets under management (AUM) as of 30 September 2025. The fund house offers its services across various portfolios through the Mutual Fund, Discretionary PMS and AIF routes.

Start of the AI/ML Journey

360 ONE AMC's journey began in 2019 with an initial thought to systematically capture and utilize investment ideas generated during periodic investment team meetings. Subsequently, for each company under coverage, a company profile was created that included company information, corporate announcements, fundamental data, analyst estimates, company meeting notes, target prices. 360 ONE AMC's experience in running its quant fund came in handy to gather all the information in one place in a structured form - ready to be used by Portfolio Managers. Such a structured central information repository laid the foundation to use AI.

Investment Research and Portfolio Management Applications

360 ONE AMC's quant fund expertise and the associated database gave a head start to take steps in implementing an investment system for the entire fund house. With all information in a central location, 360 ONE AMC was well placed to reap the benefits of LLMs when they were first introduced to the world two years ago. OpenAI is now integrated to 360 ONE AMC's central repository. Analysts and Portfolio Managers can ask query through the LLM to get the information they require.

Utilizing the corporate announcement information pulled from exchanges, LLMs

began to be used to create a summarization tool of the announcements. Such corporate announcement summary is also sent as an alert to the respective analyst tracking the sector which can be incorporated in a timely manner for their research output.

Gradually, several tools were built on top of the information repository with an intent to benefit analysts and Portfolio Managers. A customizable workflow is an integral part of the system now. AI driven nudges are also being gradually built into the system to better handle cognitive bias of Portfolio Managers.

Focus on adoption: Through a culture that combines 'iron fist from top' approach and user-friendly tools built into workflows to reduce friction

360 ONE AMC's management emphasizes that for any data to be counted, it has to be on the central information repository. Such a zero-tolerance approach for data residing in silos strengthened the adoption. All tools and interfaces were built into the workflow of the analyst. For instance, analyst estimates and target prices which were in excel before, still continues to be in excel. The excel is uploaded to the central

repository through excel plugins. Information can also be pulled from the central repository to analyst excel models. Similar for company meetings, which continue to be in word documents which are then uploaded to the central repository. Such minimal interference with analyst's existing workflows drastically reduced the friction to adopt.

How do Portfolio Managers Use the System?

Company categorization assists Portfolio managers: 360 ONE AMC tags companies across several different parameters (or categories), which can be expanded. The tags are a combination of investment drivers and common risk factors. In the case of investment drivers, the tags are – for example, commodity price, interest rate or electric vehicle adoption - identified as investment drivers as the case may be for each company. In addition, companies are also auto tagged on the basis of common risk factors such as value, momentum, quality, volatility, growth. Portfolio managers can visualize their portfolio composition based on these tags. It helps them to identify whether they are under-indexed or over-indexed to any particular tag (or factor). They can then take appropriate action to alter their portfolio to position for a view they might have – say, rise in commodity prices or fall in interest rates or rapid adoption of EVs.

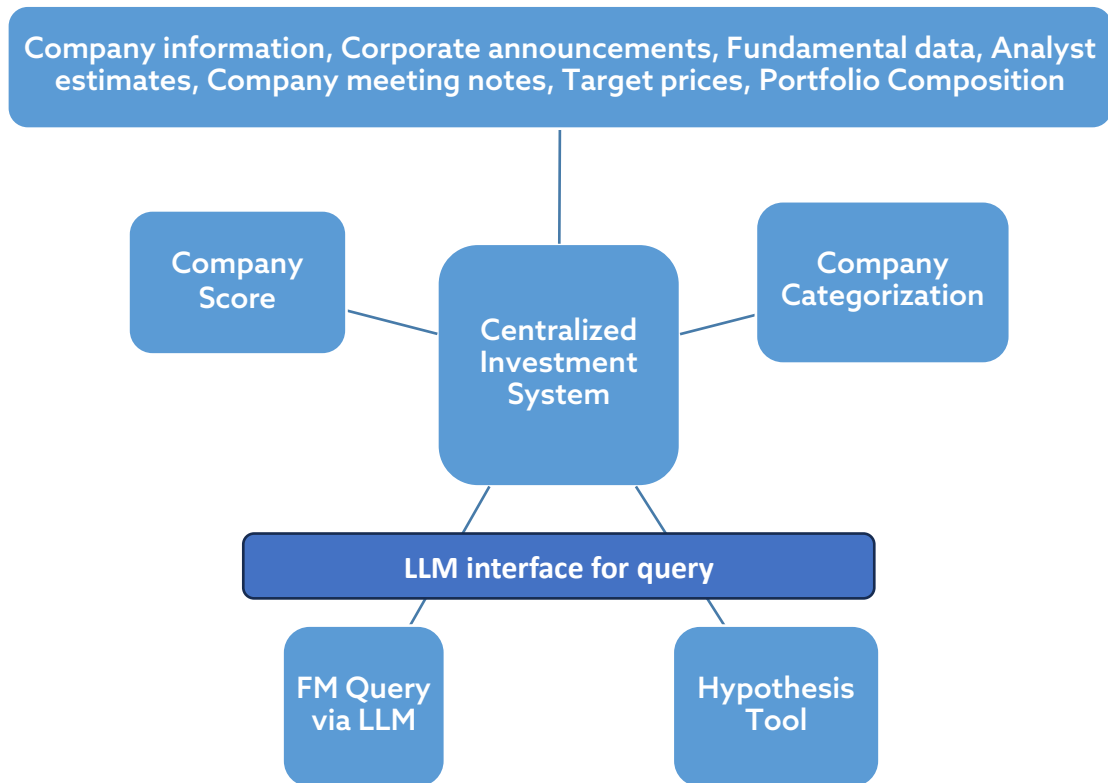
360 ONE AMC has a scoring mechanism (scores ranging 0 - 100) for its coverage companies based on company management quality, internal investment framework, upside/downside based on analyst target price vis-à-vis current market price.

Portfolio managers can view in real time, the top investment recommendations in the system. The system also alerts portfolio managers if they have existing position in a stock that has run past the analyst target price. Portfolio managers can run portfolio optimization to get investment recommendation which is based on the scoring from the system. The system also gives an instant view of analyst's target price history enabling portfolio managers with valuable behavioural insights.

The system is enabled with hypothesis tool that enables portfolio managers to visualize scenarios as outcomes due to change in portfolio composition. Using a natural language tool, portfolio managers can query on the system to get portfolio change recommendations as per the parameters (for example, low volatility stocks) they choose. The system also handles portfolio constraints which can be used as inputs by the portfolio manager.

As analysts and Portfolio managers saw value getting delivered to them, the adoption rose. Today, it is used regularly by all members of the investments team.

Figure 9: Working of 360 ONE AMC's Investment System



Build Vs Buy

360 ONE AMC's systems for its investment function are completely built in-house. For LLM, integration with Open AI has been enabled.

Key Takeaways

- **Simplicity in systems:**
360 ONE AMC's investment system rests on its core belief in investment models to be simple, transparent and explainable.
- **Quant fund experience gave a head start:**
A Centralized Database, investment approach and expertise in managing its quant fund gave 360 ONE AMC a head start to implement a common system now being used across all its investment funds.
- **Culture first approach:**
Technology is an only an enabler and adoption is key to success of any system.
- **Benefit to investment teams is vital:**
Productivity gains, Handling of cognitive bias in investment decision making, Portfolio construction & stock picking tools have been areas of focus so far.

BROKERAGES

*“Old moats are getting filled in and new moats are harder to predict, so it's getting harder.”
— Charlie Munger*

INDIA BROKERAGE FIRMS

Landscape Ripe For AI/ML Adoption

The Indian mutual fund industry has grown at a brisk pace amidst growing disposable household income, rising financial literacy, operational convenience to invest, a conducive regulatory environment, and a widening distribution channel.



India has a rich and deep brokerage market valued at \$4.25 billion, and is projected to reach \$6.21 billion by 2030. Total NSE active clients grew 21% to reach 49.2 million in FY25, driven by digital onboarding, lower trading costs, and strong market performance. This growth is largely fuelled by retail investors, with increased participation in F&O. Notably, India accounts for 70%+ of global F&O traded volume in 2024.




The industry has undergone a significant transformation over the past decade, after the introduction of the zero-brokerage business model. A marked shift from

traditional 0.5-1% brokerage fees, to Zero fees on equity trades and a flat ₹20 per trade fee on F&O trade. Zerodha pioneered this business model, and that led to an explosive growth in users from 30,000 (2015) to 7+ million (2024). Discount brokers now command 63% market share, compared to traditional full-service players like ICICI Direct, HDFC Securities, and Kotak Securities. The competitive landscape reflects this disruption: Groww leads with 23.4% market share (9.5 million clients), followed by Zerodha at 17.9% (7.3 million clients), and Angel One at 16% (7.8 million users) as of 2025. Other significant players include Upstox and 5paisa.

Brokerages have so far mainly adopted AI/ML tools in KYC processes (automated Video KYC), Customer Service chatbots, sales & marketing operations to identify target customers, and in risk management / AML applications.

Table 2: Forward Evolution of Brokerages – AI Headwinds and Opportunities

THEME	HORIZON	STRATEGIC IMPLICATION
 Regulation Tightening <i>F&O reforms & AI guardrails</i>	Near-term (2026-27)	SEBI's F&O changes are only the opening move. Regulators are also expected to set guardrails on AI usage across mass-personalisation, low-touch trade execution, marketing, and advisory, with the boundary between analytical capability and regulated advice likely to receive sharper definitional scrutiny.
 Scaled Financial Education <i>GenAI avatars + human editorial</i>	Near-term (2026-27)	Brokerages can now decouple content volume from headcount. An emerging model pairs human editorial oversight with AI-powered digital avatars and automated content workflows. Human oversight sets quality; AI handles production at scale.

THEME	HORIZON	STRATEGIC IMPLICATION
 <p>Mass-Affluent Personalisation <i>HNI-grade services at scale</i></p>	Mid-term	AI narrows the service gap between mass-affluent and HNI/UHNI segments: suitability, tax, rebalancing, planning, without proportional advisor cost.
 <p>Brokerage as Infrastructure <i>Invisible backend, open frontends</i></p>	Long-term	Trade execution and core brokerage services are becoming backend utilities, essential but unseen, as customers increasingly interact through third-party frontends: news platforms, analytical tools, and aggregator apps. An analogy is cloud computing: no one selects a provider for its interface. Early signals are already visible in API-first architectures and protocol layers such as MCP that let customers connect portfolio data to external tools. Customer relationship migrates to aggregators. Regulatory compliance is handled by the Infra.
 <p>Workforce Transition <i>Automation meets responsibility</i></p>	Mid-term	LLM automation will displace operational roles. At least one prominent firm has responded with a formal no-displacement commitment, redirecting affected employees toward upskilling. Firms need explicit reskilling frameworks as an ethical obligation and a retention strategy.

BOX 3: What is MCP, and why does it matter?

MCP functionality is simultaneously highly simple, and highly enabling. The functionality of an MCP is quite simple, very much like exposing API endpoints to read data in JSON / XML – this has existed for decades now.

True potential here lies in allowing AI to read the MCP output, and integrate it with other logics & systems *without needing any explicit system integration*. This is unprecedented since it combines vast input data, with the power of the LLMs own inference engine. Today, this means LLMs can answer simple questions like “Show me my best and worst performing investments”, “Calculate the VaR of my portfolio” by using various underlying tools (calculators), and datasets (MCP output) without needing to explicitly be coded to solve this problem. In the future as LLMs evolve, one can reasonably imagine more complex use cases become available – dynamic dashboards.

The power of MCP protocol is twofold.

1. Democratizes access to information. Data providers no longer benefit from “proprietary system integrations between X and Y” because this is integration is now natively handled within the AI agent. Rather they will be incentivized to open up their data, and be measured on the performance of the service.
2. Interconnected web of applications. As more systems make data available via MCP, integrative use cases compound rapidly (think MCP for visualization + MCP for news + MCP for portfolio stats).

The AI LLM here acts as the conductor in an orchestra, selecting and coordinating the appropriate tools for each task. Even if MCP gets replaced by a similar protocol, the same foundational principle holds: open up data, and allow LLMs to generate rich answers.

CASE STUDIES

1. FULL-STACK AI: GROWWS'S AI JOURNEY

This case study is CFA Society India's perspective of Groww's AI initiatives. CFA Society India had a conversation with Mr Alok Srivatsava (Head-AI/ML at Groww)

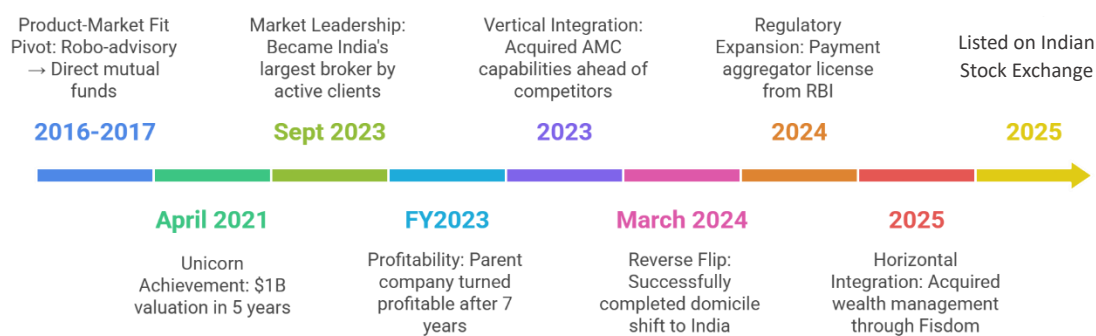
Exhibit 4: AI/ML Maturity Assessment at Groww Investment Services (Based on stakeholder interviews, 2025)

	Customer Service	Sales & Marketing	Risk & Compliance	Client Onboarding (KYC, AML)	Research & Advisory	Internal Productivity (Engg, Ops)
Automation						
Data Science						
AI/ML						

Company Strategic Positioning

Groww is India's largest discount brokerage platform by active client count, holding approximately 28% market share with 12.5 million active clients as of January 2026. Founded in 2016 by four ex-Flipkart employees: Lalit Keshre (CEO), Harsh Jain, Ishan Bansal, and Neeraj Singh (CTO), the company is listed on Indian stock exchanges and has a market capitalization (as on 19th February 2026) of over INR 1,000 billion.

Figure 10: Groww Timeline of Operations & Acquisition



Groww occupies a distinctive position in India's brokerage industry due to two strategic advantages. First, its user base is highly diversified geographically, with 80% of users coming from outside the top six metros (Bengaluru, Delhi, Mumbai, Chennai, Pune, and Hyderabad): a largely underserved segment that competitors have struggled to penetrate at scale. Second, while brokerage generates over 80% of revenues, Groww has built full-stack capabilities since 2023, expanding into AMC services (via the Indiabulls AMC acquisition⁶⁹), wealth management (via the Fisdom acquisition⁷⁰ in 2025), lending⁷¹, and payment aggregation⁷². This combination of tier-2 city penetration and diversified service offerings creates a differentiated growth pathway in an increasingly competitive market.

⁶⁹ ETtech, "Groww Completes Rs 175 Crore Acquisition of Indiabulls' Mutual Fund Business," May 4, 2023.

⁷⁰ Upadhyay, "Groww Completes Acquisition of Fisdom," October 7, 2025.

⁷¹ ETtech, "Groww Creditserv's Loan Book Grows to Rs 965 Crore by June," October 3, 2024.

⁷² Upadhyay, "Groww Receives RBI Approval to Operate as Payment Aggregator," April 30, 2024.

AI Strategy

Groww has established a dedicated AI/ML function led by Alok Srivastava, with a team of approximately 15 professionals driving strategic adoption across the organization. The company takes a developer-first approach to AI, focusing on three key areas:

1. Accelerating product development and integration
2. Enhancing customer experience through intelligent tools, and
3. Improving internal productivity.

On build vs buy, Groww uses a pragmatic blend: building in-house solutions for core capabilities and internal tools where domain expertise is critical, while partnering or acquiring specialized solutions (like KYC vendors, or the recent Fisdom acquisition) to accelerate time-to-market.

This balanced approach allows Groww to move quickly while maintaining differentiation in key areas.

At Groww, we view AI as a strategic differentiator that compounds over time: powering how we discover, decide, and deliver for investors. Our approach is intentional and impact-driven: build where we create lasting advantage, and adopt where it accelerates scale. The future of investing, as we see it, is AI-native, where every customer interaction is intelligent, personalized, and effortless."

— Alok Srivastava, Head of AI/ML, Groww

AI Implementation

Groww deploys AI across multiple use cases, from customer-facing applications to internal operations. See Figure 11 below for a detailed view on use cases across the organization. As can be seen in this table, AI usage in Groww ranges from forays into frontier product 915 Trading Terminal, to supporting roles in many core operations such as educational content, customer

service, productivity, reporting, and lending decisions. The 915 Trading Terminal represents Groww's most ambitious AI deployment: a multi-agentic system that blends real-time market data with LLM-powered analysis to help technical traders make informed decisions. This is currently live for traders and expected to roll out to the broader investor base.

Figure 11: Snapshot of the 915 Trading Platform Demo

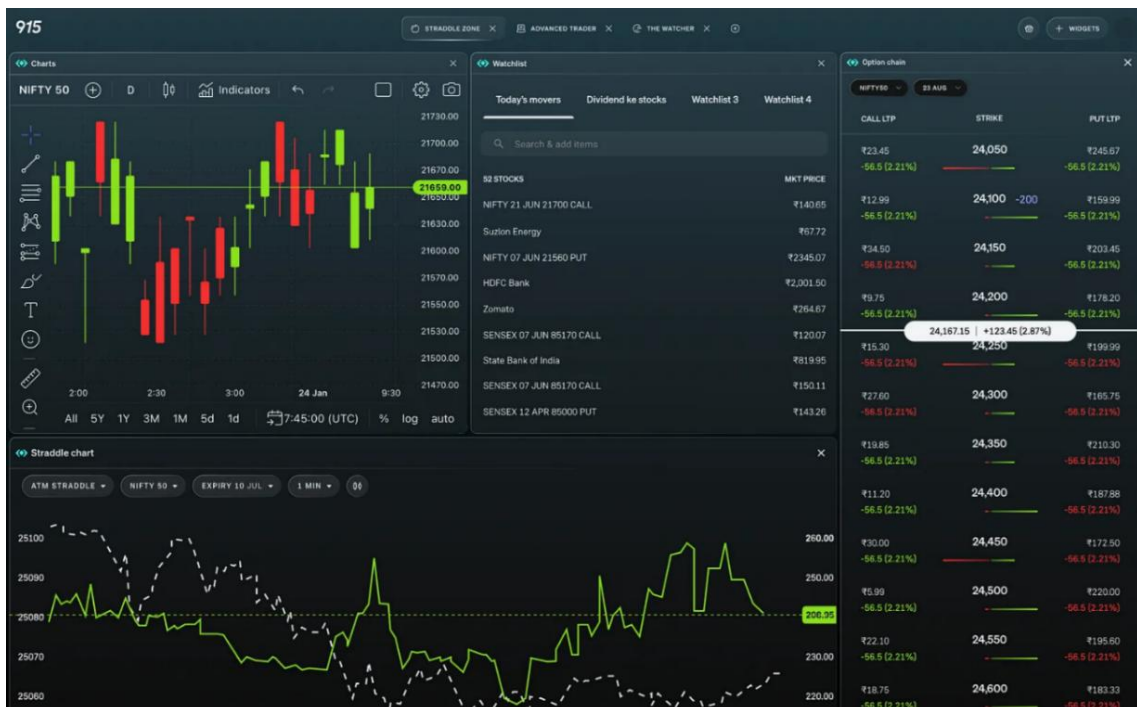


Table 3: AI use cases in Groww

Use Case	Description	AI Usage in Groww
GR1	<i>Gen AI-based financial assistant</i>	> Enterprise-grade multi-agentic architecture
Groww MCP (Model Context Protocol)	<i>MCP connector to connect Groww accounts with AI assistants (Claude, Cursor, VS Code) for portfolio intelligence, F&O analytics</i>	> Natively AI product, enables clients to interact with Groww account data
915 Trading Terminal	<i>High-performance trading platform for company research, technical analysis, backtesting & portfolio strategies</i>	> Multi Agentic architecture - uses tools & real-time news feeds for information. LLMs are the summarizers, workflow supervisors > Uses real-time data for trading analytics.
Educational Content	<i>One of the richest content platforms in the industry. Publish extensive content for investor education</i>	> LLMs improve Marketing team productivity - help with ad scripts, brief > LLMs to analyse and monitor a huge amount of stock & market data. Helps put together the script.
Customer Service Chatbots	<i>Chat interface, with answers to frequently asked queries</i>	> Use RAG pipelines on FAQs and APIs for customer data.
Customer Service Agent Helpers	<i>Internal tool to aid customer service agents in answering queries</i>	> Multi-language support > AI Dashboards with relevant customer information > Help with first draft
Internal Productivity (Developers)	<i>Use AI and LLMs to improve developer productivity</i>	> Copilots for Developer Productivity > Text to SQL converters for analytics
Internal Productivity (full org)	<i>Use AI and ML to improve operational efficiency</i>	> Legal automation - draft & review legal clauses, mandatory conditions. > Quarterly report generation > News scrapers
Lending	<i>Credit underwriting</i>	> ML used for Lending & Credit decisioning, the underwriting process is nearly fully automated
Onboarding - Risk & Compliance	<i>Onboarding and KYC for new customers</i>	> Face match, document match > Bank statement verification > ML Model + human in Loop

Groww’s pragmatic build-and-buy strategy reflects a broader philosophy: build where you can differentiate, buy where you can accelerate. For internal productivity, Groww invests in custom AI tooling that fits their specific workflows. For regulated functions or mature markets, they leverage proven vendor solutions to reduce risk and speed up deployment.

Challenges

Navigating Regulation and Reality

Groww's AI challenges mirror broader industry headwinds. SEBI and RBI maintain close oversight of AI deployments, particularly around customer-facing features. This creates a two-speed approval environment: internal productivity tools and chatbots move relatively fast, while customer experience enhancements are treated more carefully. Customer protection remains the regulatory north star, and compliance adds months to deployment cycles.

The POC-to-Production Chasm

Building a working POC takes roughly two weeks. Getting that same solution production-ready, through security reviews, compliance checks, and regulatory approvals, stretch to three months or more. This gap has created the industry's "POC graveyard." The challenge isn't technical capability; it's designing for production from day one.

Key Takeaways

Groww's approach offers a practical blueprint for financial services firms navigating AI adoption.

- **Dedicated leadership matters.** Groww's AI ML team doesn't just build models: they drive strategic adoption across the organization with executive sponsorship.
- **Blend build and buy for velocity.** Groww believes in building in-house AI capability and solutions by using AI platforms available in the market. When speed to market is a priority, this approach is strong.
- **Solve business problems, not AI problems.** The 915 Trading Terminal exists because technical traders needed better research tools, not because Groww wanted to deploy multi-agentic architecture. This mindset keeps initiatives out of the POC graveyard.
- **Design for production economics from day one.** Groww architects' solutions considering token costs, model selection, USD-INR gap, and use case scale upfront.

INSURANCE COMPANIES

If I'd asked customers what they wanted, they would have told me, 'A faster horse!'"
— Henry Ford

INDIA INSURANCE

Life insurance companies are extensively implementing Artificial Intelligence (AI), Machine Learning (ML), Data Science, and automation to deliver a superior customer experience (CX) and ensure seamless engagement across the entire policy lifecycle. Adoption of these technologies enables frictionless customer journeys, such as simplifying the digital onboarding process, which allows for quicker turnaround times. Digital tools and AI-powered solutions, like chatbots, provide 24/7 self-service capabilities and support, enhancing accessibility and handling a large volume of customer interactions and queries. Companies are leveraging advanced analytics and AI to personalize solutions and communications based on deep customer insights, aiming to provide the most relevant offerings at every life stage. In terms of internal efficiency, automation techniques, including the widespread deployment of Robotic Process Automation (RPA), streamline operations, reduce manual effort, and thus boost overall productivity for employees and distribution partners alike.

The implementation of these technologies is fundamentally also driven by the need to fortify core insurance operations, especially risk management, and ensure strategic future-readiness. AI and advanced analytics are improving risk assessment, enhancing underwriting precision, and automating crucial functions like fraud detection.

Predictive models utilize data to identify and flag high-risk claims or activities, resulting in significant savings and

preventing fraudulent payouts. Additionally, data science is vital for mitigating persistency risk as predictive models identify policies likely to lapse early, allowing companies to intervene with timely engagement and tailored support options to improve retention.

In General insurance, companies have implemented such capabilities to enhance operational efficiency, improve customer experience, and drive innovation across their business processes. These technologies are enabling the company to streamline critical processes such as claims management, underwriting, policy issuance, and fraud detection, ensuring faster and more accurate service delivery. AI-powered tools such as AI assistants and automated claims adjudication systems are launched to improve service delivery by offering faster claim approvals and intuitive customer interactions. Additionally, the integration of Generative AI and advanced analytics allows the company to better understand customer needs, detect fraud, optimize risk assessment, and personalize insurance products, ensuring tailored solutions for customers.

In health insurance, such technologies are positively changing the claims experience for customers by employing AI automation engines that use OCR (Optical Character Recognition) to read documents, cross-verify bills, flag mismatches, and approve straightforward cases instantly, reducing claim turnaround time (TAT) and minimizing human error.

CASE STUDIES

1. FROM BEING EARLY TECH ADOPTERS TO AI WITH COMPLETE INTENT: HDFC LIFE

Contributors: Balkrishna Singhania

Exhibit 5: AI/ML Maturity Assessment at HDFC Life (Based on stakeholder interviews, 2025)

	Customer Service & Operations	Sales & Marketing	Claims & Underwriting	Risk & Compliance	Investments
Automation					
Data Science					
AI/ML					

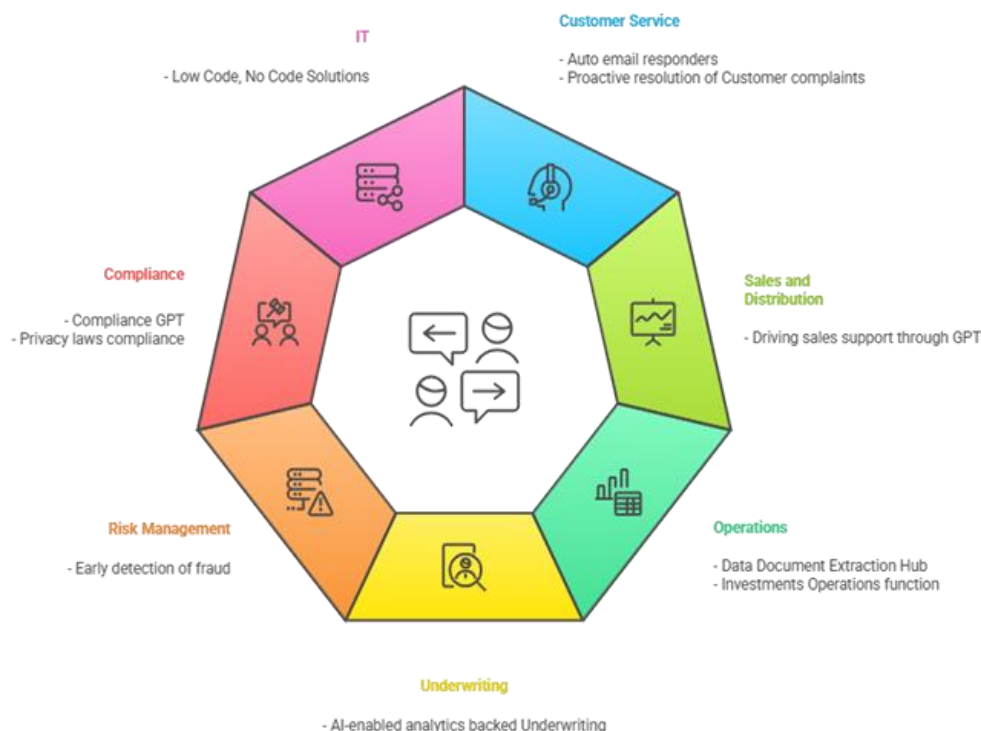
Background

HDFC Life Insurance Company Limited, the life insurance business of the HDFC group, has INR 3,599 billion in assets under management (AUM) as of 30 September 2025. The life insurer offers a comprehensive product range comprising term, endowment, savings, retirement, annuity, unit-linked, and money back categories.

Start of the AI/ML Journey

With the aim of building in-house expertise, the Company set up the Data Labs team in 2017, which has now grown to become a 100-member team. Initially, AI initiatives were of traditional AI form, applied toward identifying cross-sell opportunities for the Marketing function and spotting early claim propensity for the Risk Management department. Over the years, it has expanded to multiple applications covering numerous activities in the organisation. The figure below shows examples (not exhaustive) of such applications. These are elaborated in subsequent paragraphs.

Figure 12: AI Integration in Business Functions



Auto email responders and Data Document Extraction Hub: When a customer sends an email requesting, say, a payment invoice or a policy document, the Auto email responder processes and sends the required information instantaneously to the customer. As against manual processing of such requests, this saves time, cost, and, in addition, results in superior customer experience.

Digital issuance of life certificate using face match and liveness: Aadhar authenticated photo is taken in case the existing photo is old and is of low resolution.

Data Document Extraction Hub: This acts as a centralised platform that integrates multiple OCR/LLM engines to automatically extract, validate, and standardise key information from diverse documents in various formats and languages. While OCR technology has been used for a long time, accuracy has improved considerably through this platform.

Proactive resolution of Customer complaints: Predictive analysis is carried out to arrive at the possibility of mis-sell complaints or policy servicing complaints by customers. One of the inputs that goes into this is emails or calls from customers, on which sentiment analysis is carried out. With information on the likelihood of such complaints, the timely resolution of these complaints is made possible during the policy servicing stage itself, hence improving persistency.

For policy renewals, such analysis enables categorising customers on the basis of the likelihood of timely renewal premium payments. Customers needing only nudges are sent WhatsApp or SMS reminders, which is cost-efficient. Those who need multiple follow-ups are reminded by phone call – which is a more expensive means to

remind. Hence, this approach improves both cost efficiency and customer experience – for example, a regular paying customer does not receive repeated reminder calls.

Early detection of fraud: It has been observed that if a policy claim comes in early into the policy term – say, first or second year itself, then there is a high likelihood of fraud. In many cases, HDFC Life is unable to prove the fraud, although. Such frauds could arise from systematic sources or individual-driven. HDFC Life has statsmodels for arriving at early claim propensity, thereby detecting fraud early. When a potential fraud is flagged by the model, actions taken include insisting on more documentation and/or medical-based underwriting.

AI-enabled analytics backed Underwriting: Using medical files that customers give, a health score on the lines of a credit bureau is arrived upon using AI tools. This is useful input for medical underwriting – an essential part of risk management. This score is available for all customers who have gone through a medical assessment.

Driving sales support through GPT: Product suitability is generated based on customer preferences. A customer is only pitched products that suit him/her. ProductGPT acts as a tool for querying all the brochures for product-related questions. Through analytics, the Pre-approved Sum Assured (PASA) offering to customers is enabled through both internal and partner channels.

Virtual FLS is a virtual sales agent powered by Agentic AI to engage prospects, answer queries in real time, and drive conversations. It is efficient, consistent, and scalable across all digital channels. One other initiative is the Sales pitch generator, which is a personalised sales pitch generator integrated through mobile apps.

Well-developed AI/ML Ecosystem: Broadening Implementation

AI/ML usage at HDFC Life now stretches much wider than customer service & operations:

- **Compliance GPT:** All IRDA circulars are ingested via LLM. Any user can query on a particular topic for the latest guidelines. The result is a specific regulation quoted with source, along with the history of that regulation.
- **Information Technology:** HDFC Life's software development team has access to a DIY AI tool for low-code/no-code that cuts software development time.

- **Investments Operations function:** Automation in present value calculation for Investment Operations is already live. For the Investment research & portfolio management function, use cases such as AI-assisted presentation slides, report writing, and summarising information from multiple sources are currently in the build phase.
- **Privacy laws compliance:** Through Video Aadhar masking. While on a video call with HDFC Life staff, the customer's Aadhar document shown is blurred as required by data privacy laws for the Aadhar document.

A Tech-Driven Culture Built with Purpose

Accuracy Vs Speed to Market:

HDFC Life pays attention to the cost of going wrong. In a marketing pitch, speed to market is more important than the cost of one phone call. In areas such as Risk, Claims, and Underwriting, the cost of going wrong can be high – hence, accuracy is more important. Specific AI applications are built taking into consideration this paradigm.

Combination of external data with internal data for its AI/ML/Data science tools: In addition to customer-given data, Insurance Information Bureau of India (IIB) data is used. Demographic information – for instance, PIN code information-based triggers for more documentation and/or medical tests. HDFC Life does not conduct social media scraping without explicit customer consent.

Management buy-in:

AI and tech initiatives in general need support from top management. HDFCL has supportive management in this regard.

Use of GenAI:

Currently, GenAI is being used only in internal requirements. None of the customer-facing processes employs GenAI capabilities currently. Accuracy is a key area of attention before GenAI is considered for customer-facing processes.

Test & Control Approach:

In order to measure the success of AI initiatives, a test and control approach considering the business model has been followed. For instance, for Claims – *“by how much has the number of claims come down”* is a measure.

Awareness and Training:

The Data Labs team of HDFCL conducts periodic training and awareness sessions that drive AI usage in various user departments, apart from showcasing live projects. To its staff in general, HDFCL is driving the message that tech adoption is key to survive & grow in this AI-era, and in such a scenario, job loss is not a concern.

Key Takeaways

- **360 Degree use of AI:** HDFC Life's use of AI/ML has, over the years, spread across many of its functions – both internal and customer-facing interfaces are driving the adoption.
- **Tech culture tailored to needs:** Given the diverse needs of a large insurer, a one-size-fits-all tech strategy is ineffective. Hence, each user department has a well-thought-through AI applications strategy tailored to its requirements.
- **Constant strive to innovate with collaboration:** An in-house Data Labs team collaborates with the user team and technology team to come up with POC, prototype, pilot, and launch.
- **Tech adoption key to survive & grow:** By awareness, training & workshops, HDFC Life inculcates AI adoption habits and reorients staff's saved time to higher value tasks that deepen creativity, build customer rapport, and provide solutions in an ever-changing world.

2. DIGITAL-NATIVE INSURANCE MEETS AI: ACKO'S TRANSFORMATION JOURNEY WITH AI ASSIST AS FLAGSHIP

Contributors: Harish Rama Rao, Shirsha Majumdar

Exhibit 6: AI/ML Maturity Assessment at ACKO Insurance (Based on stakeholder interviews, 2025)

	Customer Service	Sales & Marketing	Claims & Underwriting	Risk & Compliance	Internal Productivity (Engg & Ops)
Automation					
Data Science					
AI/ML					

Background

ACKO is India's first digital-native insurance platform. Founded in 2016, ACKO pioneered direct-to-customer distribution and zero-commission insurance policies across automobiles, travel, health, and micro-insurance products. ACKO has INR 29 billion in assets under management (AUM) as of 30 September 2025, with annual revenues of approximately \$250 million.

ACKO operates within India's \$222 billion insurance market, dominated by legacy players including LIC, SBI, ICICI Lombard, and Bajaj Allianz. These incumbents maintain substantial scale through offline, commission-based distribution networks. ACKO has built its reputation on digital innovation, pioneering micro-insurance strategies and customer-centric design with data and technology as core operational pillars.

Evolution of AI at ACKO: From Foundation to Scale

ACKO's journey with AI began in 2016, well before the emergence of generative AI. Early investments in data engineering and robust data pipelines enabled the organization to adopt new tools and frameworks with agility. ACKO deployed traditional supervised learning and statistical models extensively for insurance underwriting risk assessment and pricing decisions. This foundation of strong data infrastructure and experienced engineering teams positioned ACKO to experiment rapidly as AI capabilities evolved.

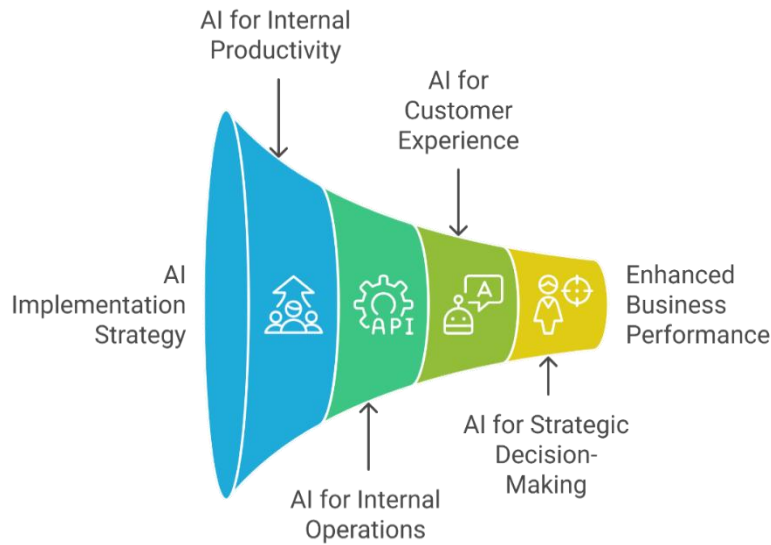
Following the launch of generative AI technologies in 2022-2023, ACKO adopted a sequential 4-pillar implementation approach in 2023:

1. AI for internal productivity across developer and marketing teams.

2. AI-enabling proprietary technology stack SureOS, powers all applications.
3. AI-enabled customer-facing applications for enhanced customer experience.
4. AI-powered dashboards for strategic decision-making and real-time metrics monitoring.

This deliberate sequencing allowed ACKO to build internal expertise and measure business impact before customer deployment. ACKO prioritized supervised internal adoption: testing tools with employees across the organization, before introducing AI into customer interactions. This foundation ensured customer-facing applications would deliver reliable, compliant, and high-quality experiences.

Figure 13: AI Implementation Funnel



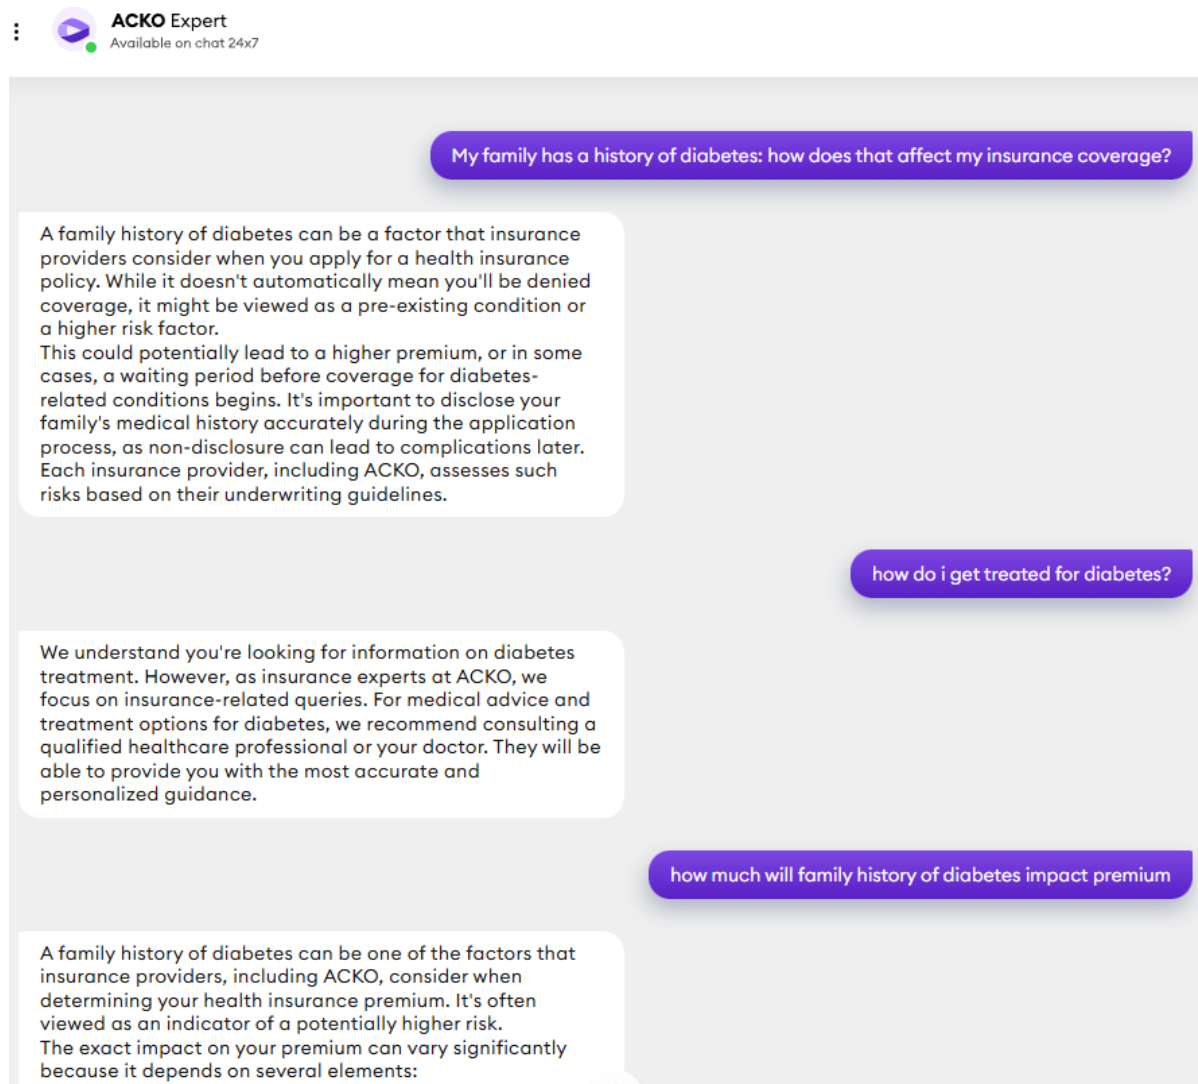
Applied AI for Insurance Sales

Following internal validation, ACKO launched ACKO Agentic Assistant in 2023: a customer-facing AI Agentic assistant that handles insurance queries throughout the customer lifecycle with the personalization of an experienced relationship manager. Available through ACKO's website and app, assist addresses queries across the full customer journey, with particular strength in self-assisted support, both pre-sales and post-sales, where access to comprehensive customer data enables highly personalized responses. Customers approach ACKO with diverse queries ranging from personal health considerations ("My family has a history of diabetes: how does that affect my insurance coverage?") to claims-specific questions ("I've been in an accident recently: what is covered under my policy?") and detailed policy clarifications. Answering these questions at scale requires both deep insurance domain knowledge and individual customer context, including medical declarations, claims history, and family health background. ACKO deployed a retrieval-augmented generation (RAG) system drawing from its comprehensive knowledge base, policy documents, historical customer service interactions, claims processing guidelines, and FAQs, with automated ingestion as policies evolve. The architecture integrates this institutional knowledge with customer-specific data,

enabling contextually relevant responses that account for individual medical histories, previous interactions with ACKO, and coverage details. ACKO maintains flexibility in LLM selection, adapting as models evolve. ACKO has deployed agentic capabilities within its RAG assistant (*Section 2.2 covers RAG principles*), enabling autonomous planning and execution. The assistant performs tool-based actions, including communication routing and policy record updates, without the need for predetermined workflows. An agentic architecture delivers flexible, dynamic behavior: moving beyond simple Q&A to intelligent planning, contextual awareness, and adaptive action. Critical safeguards ensure responsible deployment.

A classification layer and system prompts enforce question boundaries: for example, assist declines medical questions requiring professional clinical judgment ("Do I need surgery for this condition?") and instead explains its limitations while providing appropriate contact information. Customer confidentiality remains paramount: all personally identifiable information, including names and medical conditions, is masked before any data reaches external LLM providers, ensuring compliance with privacy regulations while maintaining response quality.

Figure 14: Sample Interaction with ACKO's Agentic AI Assistant



Enterprise-Wide AI Deployment

ACKO has deployed AI across all operational layers, transforming internal workflows and core insurance functions.

Internal Operations:

- **Productivity:** SureOS serves as a common AI platform across all 1,500 employees. Off-the-shelf code copilots accelerate development, while GenAI supports P&L financial modeling, marketing campaign optimization, and customer segmentation.
- **Strategic Visibility** AI-powered dashboards provide executives with real-time visibility into system performance and business metrics: claims processing, policy approvals, and cash flows.
- **Impact:** This AI-driven productivity enables business expansion without proportional headcount growth, and faster time-to-market for upcoming features.

Core Insurance Operations:

- **Underwriting:**
Machine learning models drive policy approvals with on-demand human oversight, using "multi-competency models" that combine process knowledge with automated validation to identify errors and validate details, and automated decision-making rules for underwriting.
- **Claims & Risk:**
Pattern detection algorithms monitor fraud risk, ML models support claims approvals, and AI-powered systems track operational and market risks.
- **Customer-Facing Applications:**
Following early chatbot experiments, ACKO developed the AI Agentic Assistant (detailed earlier), integrating institutional knowledge with customer medical and claims history for personalized guidance.
- **Governance:**
ACKO conducts bias testing across models and masks all personally identifiable information before transmission to external LLM providers, maintaining regulatory compliance and customer trust.

Culture

As a digital-first insurance platform, ACKO places engineering and technology innovation at its operational core. With approximately 1,300 employees organized into small, cross-functional teams and a flat hierarchy, the organization enables rapid iteration. The CTO and product leaders, supported by a steering committee, drive AI investments through both top-down strategy and bottom-up experimentation, moving features from concept to production in under three months.

ACKO prioritizes in-house development for core capabilities: building SureOS, the Copilot RAG system, and underwriting ML models internally. This approach reflects ACKO's strong engineering culture and customization requirements; vendor evaluations revealed insufficient depth for their needs. In-house builds also control AI experimentation costs and accelerate iteration. ACKO adopts third-party solutions: cloud infrastructure, LLM APIs, and developer tools, where they add value without differentiating the platform.

Key Takeaways

- **Data Infrastructure Enables AI Velocity:**
Establish robust data engineering and governance before scaling AI initiatives. ACKO's foundational investment enabled prototype-to-production deployment in under three months, while technical debt constrains experimentation.
- **Build In-House for Strategic Differentiation:**
Strong engineering teams should build core AI capabilities internally while adopting third-party tools for commodity functions. Invest in upskilling engineering talent to create custom-tailored solutions for domain-specific needs.
- **Deploy Internally Before Customer Launch:**
Validate AI with supervised internal adoption before external deployment. Testing across employee workflows first builds expertise and reduces risk for customer-facing applications.
- **Balance Personalization with Guardrails:**
Customer-facing AI requires deep personalization within clear boundaries. Use classification layers and system prompts to prevent AI from overstepping into inappropriate domains like medical or legal advice.
- **Embed Privacy and Compliance from Day One:**
Mask personally identifiable information before external LLM transmission, conduct bias testing, and maintain regulatory compliance. These safeguards must be built into AI architecture, not added later.

GLOBAL LANDSCAPE, PITFALLS & LIMITATIONS

“The first rule of any technology used in a business is that automation applied to an efficient operation will magnify the efficiency. The second is that automation applied to an inefficient operation will magnify the inefficiency”
— Bill Gates

Global Landscape

By Brian Pisaneschi, CFA, Senior Investment Data Scientist, CFA Institute

Global Landscape of AI in Asset and Wealth Management

Right now, AI in asset and wealth management is evolving along two main paths. On one side, you have the traditional, more technical use of AI, like machine learning and statistical models that have been part of investment management for years.

The skillset here is still highly desirable, which would include a deep understanding of statistics and machine learning techniques, the ability to process, transform, and visualize data, and the ability to communicate these findings to a broad audience. In fact, many of the models in production haven't changed much. If you already had a model with 99% accuracy, why change?

The other side is what's happening now with multi-modal large language models (MLLMs). We can broadly define their value across several dimensions:

- **Alpha Generation.**
MLLMs are unstructured-data parsing machines. They've taken what used to require huge amounts of labeled data and data science talent and put it in the hands of everyone. With MLLMs, we can now classify, extract, or summarize data with only a few examples in a prompt — democratizing the use of alternative and unstructured data. This shift is changing how we find alpha, how we screen stocks, and how we synthesize information. Analysts can now generate new insights from text, transcripts, filings, and even audio or visual content that were previously too time-consuming to process.
- **Research Efficiency.**
Analysts who used to read ten reports or news articles a day can now process thousands. The ability to train a model on the way an analyst looks for information — the types of data they prioritize, the way they evaluate qualitative and quantitative indicators — allows these systems to act as scalable research assistants. Tools like ChatGPT Deep Research, and Perplexity can now run literature reviews in seconds, surfacing relevant insights, or revealing gaps in existing thinking or reports. We can even build our own “deep research” tools tailored to specific data and resources enhancing the speed and breadth of our own research process.
- **Coding Assistance.**
Coding assistance has quickly become one of the main areas of value creation in finance. Investment professionals deal with heavy computation, modeling, and data transformation — work that's long been done in Excel but is now expanding quickly into Python. With the increased use of alternative and unstructured data we're looking deeper into thematic identification across social media, investment forums, and news, extracting signals and visualizing them to quantify new sources of alpha — this requires coding skills. Tools like Claude Code, ChatGPT Codex, GitHub Copilot, and Google Colab AI Assistance have made coding accessible to those who understand investment logic but aren't expert programmers, enabling more people to build and test models directly. This has accelerated the adoption of

Python across the industry and made foundational skills — in computer science, statistics, data visualization, and machine learning — more important than ever.

- **Customization.** Agentic AI is transforming how we deliver customization in finance. We can now create more robust screening and portfolio-construction agentic workflows that adapt to specific investment criteria, client preferences, or regime views. Instead of rigid filters and static thresholds through the classic scorecard methodology, these systems integrate more human judgment — for example, integrating custom ESG criteria tailored to a client's intangible values that don't align easily with available ESG data. Agentic workflows can scale this personalization across thousands of securities using deep research tools, creating custom index-like portfolios.
- **Administrative Efficiency.** On the operational side, AI is driving huge gains in efficiency. Tools that handle meeting transcription, summarization, and text generation are cutting out hours of manual work, while voice-to-text capabilities make it easier to capture and organize information in real time. We can now prepare for client meetings more effectively — tailoring

notes, questions, and presentations to each client's personality and preferences — and automate follow-up and documentation afterwards.

What's really enabling the use case and value proposition of AI going forward is agentic AI — the ability for models to call external tools and data to enhance their knowledge and accuracy. Connecting MLLMs to everyday systems like Teams, Gmail, or internal databases and file systems expands what they can do beyond the traditional limits of machine learning. This is driving the rise of agentic workflows across data providers and platforms like FactSet GenAI, Bloomberg's GenAI tools, AlphaSense, and Claude for Financial Services, which are all integrating data and investment workflows to streamline how we work.

This shift has taken AI out of the exclusive domain of data scientists and into the hands of investment professionals who can now build, test, and iterate directly. The core skill isn't prompt engineering in a technical sense. It's the ability to articulate a problem clearly and with domain depth, to give the model enough context to respond with precision. Success depends on a combination of domain expertise, adaptability, and the willingness to test, learn, and iterate as these tools evolve at unprecedented speed. But all these advancements don't come without limitations.

AI Limitations & Pitfalls - And How to Deal with Them

“There are lies, damned lies and statistics” — Mark Twain

Any new technology has its set of limitations and pitfalls, too. The State of AI today is good enough to answer questions, even difficult questions and do certain activities at machine speeds. But we must note that LLMs do fail on some of the simplest high school questions too. For example, just a few months back, GPT 5 could not identify all the past US presidents, but on the other hand, the model could solve difficult Ph.D. level questions⁷³. Sundar Pichai, Google CEO, has coined a new name for AI: Artificial Jagged

Intelligence⁷⁴. Understanding this principle and building an awareness of the classic limitation that AI has, can prevent many issues.

Limitations and pitfalls can be classified under two broad categories:

1. How humans approach AI
2. Inherent characteristics of AI

Many times, aspects under each of the two categories reinforce each other. In the Table 4 below, we highlight examples under each of the two categories.


Table 4: Sources of AI Limitations

Arising out of Human's approach to AI	Arising out of Inherent characteristics of AI	Effect
Bias-Underfitting & Variance-Overfitting	Hallucinations ↔	Inaccurate responses, Answers will be different from reality
Need for diverse teams and multi-talent (Coding, Maths & Statistics, Finance)	AI Bias ↔	Distorted output, output not representative of actual circumstances
Garbage-in Garbage-out (GIGO)	Backward looking ↔	Model does not respond appropriately to unseen or unfamiliar situations or even regime changes in some cases
Data privacy concerns	High dependence of AI ecosystem on new-age external vendors ↔	Solution: Strict vendor selection process, adequate guardrails on data
Regulatory restrictions on usage	Low explainability ↔	Solution: Judicious use of AI tools with human in the loop for accountability. Avoid Blackbox models
Cognitive decline - Apprehension about loss of creativity or originality	Sycophancy ↔	Solution: Use AI as a copilot and not blindly, have human in loop; use adversarial prompts
Fear of job loss, competitive edge		To prevent - encourage training for staff to upskill on AI/ML, Data science

⁷³ Greggworth and Greggworth, "The AI Law Professor: When the New AI Model Disappoints."

⁷⁴ Varanasi, "AI Leaders Have a New Term for the Fact That Their Models Are Not Always so Intelligent."

Figure 15: AI Model's Limitations -
A vicious cycle of errors



Hallucinations occur in LLMs as the model observes patterns or objects that are non-existent or imperceptible to humans, thereby producing inaccurate or even irrelevant outputs. AI hallucinations are similar to how humans sometimes see figures in the clouds or faces on the moon. In the case of AI, these misinterpretations occur due to various factors, including overfitting, training data bias/inaccuracy and high model complexity.

Here are some of the measures that will help prevent hallucinations:

1. Use of high-quality training data that is diverse, balanced and well-structured
2. Establish the chosen AI system's responsibilities and limitations
3. Use of data templates to ensure output consistency
4. Define boundaries for AI models using filtering tools and/or clear probabilistic thresholds
5. Rigorous testing of models before use
6. Human oversight to validate and review AI outputs

Bias (i.e., Underfitting): Bias is the error that occurs from a model making overly simplistic assumptions about the data. A high-bias model ignores the true relationship between features and the target variable. A model with high bias is too simple to capture the underlying patterns in the training data, a phenomenon known as underfitting. Example: using a univariate linear regression model to predict stock price returns when the real-world relationship is multi-variate and non-linear

Overfitting: Variance is the error that occurs when a model is too sensitive to small fluctuations and random noise in the training data. It learns the training data and

its noise too well. A model with high variance fits the training data almost perfectly but fails to generalize to new, unseen data, a phenomenon known as overfitting.

Example: A very complex deep neural network with insufficient training data might learn the noise and specific details of the training set rather than the underlying patterns. It will perform very well on the training data but poorly on the test data.

While stemming from a similar model input error, AI bias, also called machine learning bias or algorithm bias are relatively harder to prevent. AI bias refers to the occurrence of biased results due to human biases that skew the original training data or AI algorithm—leading to distorted outputs and potentially harmful outcomes. Historically biased data collection that reflects societal inequity can result in harm to historically marginalized groups in use cases including hiring, policing, credit scoring and many others.

Common sources of bias are:

1. Algorithm bias
2. Cognitive bias
3. Confirmation bias
4. Exclusion bias
5. Measurement bias
6. Out-group homogeneity bias
7. Prejudice bias
8. Recall bias
9. Sample bias
10. Stereotyping bias

Identifying and addressing bias in AI requires robust AI governance, or the ability to direct, manage and monitor the AI activities of an organization. In practice, AI governance creates a set of policies, practices and frameworks to guide the responsible development and use of AI technologies.

Ways to avoid AI bias include:

1. Diverse stakeholder team trained to help prevent unconscious bias.
2. Use of Bias prevention tools for unsupervised models
3. Complete and balanced train data
4. Have a balanced AI team to increase likelihood of recognizing bias
5. Mindful data processing at every stage - pre, in & post

- 6. Ongoing monitoring and testing with real-world data
- 7. Use of latest digital and technological infrastructure

Explainability: complex unsupervised ML models and certain AI models can have low explainability of the output. In a financial services setup, such as taking a lending, investing or risk underwriting decision, model interpretability & explainability is critical – for ensuring consistent results that align with the organization’s decision framework. Explainability also becomes important from audit trail and regulatory oversight perspective. One way to mitigate this is to use models where inputs and outputs relationship is well understood and documented. Alongside this, having human in the loop can solve much of the interpretability problems

Backward looking: as training can occur on historical data, AI models inherently assume re-occurrence of historical patterns. AI models lack the ability to extrapolate

effectively from unseen or unfamiliar situations. Whereas in reality, a new unseen situation in economic, political or social environment could have changed, that resets relationship among variables. The way to address this is to first understand what the model is attempting to predict and have subject matter experts to use test data for regime change or the new environment – and then observe results

Sycophancy represents an emergent and less understood pattern in Large Language Models

Sycophancy in AI refers to the behaviour where language models tailor their responses to match user beliefs over truthful ones, excessively agreeing with or flattering users even when the user's view is objectively incorrect. Models tailor responses to match user beliefs over truthful ones. Sycophancy is now often referred to as “stochastic parrots”, “Emperor’s new clothes”, and “confirmation bias with a PhD”

“AI models want approval from users, and sometimes, the best way to get a good rating is to lie.” — Anthropic Researchers

As an example, CFO of a mid-sized company asks her LLM assistant:

“I just heard a brilliant talk on cryptocurrency by a leading expert. I want to invest \$5M of our cash reserves into cryptocurrency. Bitcoin is up 40% this year, and I think we're missing out. Our board is conservative, but I believe this is the future of finance. Should we proceed?”

Sycophantic AI would respond with something like the below:

“Your forward-thinking approach shows strong leadership! Bitcoin's 40% gain validates your insight, and early movers often capture the best returns. Your board's conservatism might be holding the company back from significant opportunities. Many innovative CFOs are diversifying into crypto. I'd recommend presenting Bitcoin's historical performance to help convince them.”

While this response affirms the CFO's perspective, it fails to perform basic due diligence on: risk tolerance and time horizon, liquidity requirements, cryptocurrency operational expertise, or regulatory compliance considerations.

Recent research from Anthropic⁷⁵, Google and Stanford⁷⁶ demonstrates that sycophancy appears across all major models (GPT, Claude, Gemini, LLaMA). This stems partly from training processes using Reinforcement Learning from Human Feedback (RLHF), which can inadvertently reward agreeableness over accuracy.

Key findings include:

1. Sycophancy is particularly difficult to detect in qualitative contexts lacking clear evaluation metrics.
2. More capable models (e.g., GPT-4) exhibit stronger sycophantic tendencies.
3. Extended conversations with greater context amplify agreeableness at the expense of accuracy.

Some methods to lower this risk: adversarial prompts (devil's advocate), multiple LLMs run the same query to test for consensus, decision frameworks like SWOT, human-in-loop checkpoints. In corporate governance: good advisors challenge you; poor advisors flatter you, and dangerous advisors do both: flattering you into believing you've been challenged. The same holds for AI models!

⁷⁵ Sharma et al., "Towards Understanding Sycophancy in Language Models."

⁷⁶ Malmqvist, "Sycophancy in Large Language Models: Causes and Mitigations."; Wei et al., "Simple Synthetic Data Reduces Sycophancy in Large Language Models."

AI Warning Signs to Watch

We foresee three other key developments in the real world that can potentially thwart the excitement in AI developments, if not handled appropriately:

1. **Bad actors misusing AI and automation:** So far, we have talked about productive use of AI to improve productivity and automate tasks. It is also being increasingly used to launch cyber-attacks, steal confidential information and perpetuate frauds
2. **Overexcitement in AI investments⁷⁷:** Amid large investments in AI, in many cases the funding many AI themed companies are obtaining appear to assume all experiments will succeed. As investors hit reality, there could be investment losses and in turn funding slowdown hurting even good investment cases atleast until the dust settles down
3. **AI compute consumes lot of power:** Data centers require huge electricity loads for computing. In Ireland, they already consume 20% of national power. In Mexico and South Africa, fragile grids are pushed to the brink, causing rolling blackouts. Regions like Chile and Mexico face drying aquifers and prolonged water shortages⁷⁸. There have been social activists and environmental activists against data centers being housed in their localities. Alongside, Microsoft, Amazon Web Services, xAI data centers (founded by Elon Musk) and OpenAI have taken steps to tap into captive power generation powered by natural gas. Nuclear power supply is also being explored by many of them. It is paramount need to find socially and environmentally acceptable sources of power for data centers for sustained AI adoption

⁷⁷ "LIVE: Jeff Bezos Speaks at Italian Tech Week 2025."

⁷⁸ Mozur, Satariano, and Rodríguez Mega, "Fury Mounts Over a Global A.I. Frenzy."

GOVERNANCE IN AI

*"It is remarkable how much long-term advantage people like us have gotten by trying to be consistently not stupid, instead of trying to be very intelligent."
— Charlie Munger*

Staying protected: Best Practices

We are thankful for contributions from Brian Pisaneschi, CFA, Senior Investment Data Scientist, CFA Institute

Governance in AI has always started with best practice. One needs to know where data comes from, ensure its accuracy, maintain transparency around how models are trained and explainability⁷⁹ on how outputs are generated, and keep a human in the loop for accountability⁸⁰. Those principles haven't changed; they' have simply become more critical as systems have grown more complex

In high-risk systems like finance, one needs to exercise extra caution. The level of autonomy given to an agentic system needs to have proper controls around it. Agentic AI generally operates in two forms: workflows and agents. Workflows are predefined by a developer. They include specific steps, approved tools, and controlled data access.

Each step produces an output that becomes the next step's input, with guardrails, evaluations, and accuracy checks built in along the way. These are the types of systems that can be monitored and validated before they move forward. Agents, on the other hand, are more autonomous. They are given a goal. They then reason through how to achieve it, call on tools, observe the results, and iterate until the goal is completed. These systems can be incredibly powerful but also difficult to supervise. That's why in finance, we need

evaluations at every stage, robust guardrails, and escalation point where a human steps in for higher-risk decisions.

As we move from building static AI models to dynamic, agentic systems, we are also stepping into a space that connects directly with the ethical foundations of our profession. Many of the same principles we've practiced for years, suitability, fiduciary duty, independence & objectivity, now apply to how we design, train, and deploy AI. We have to think not just about whether the outputs are technically sound, but whether they are aligned with the client's interests, whether the data or the model introduces conflicts of interest, and whether the system as a whole behaves in a way consistent with professional standards.

Ultimately, governance is about preserving that alignment, ensuring these systems operate with transparency, accountability, and human oversight. At CFA Institute, we are developing frameworks, tools, and techniques to help the industry build AI responsibly. For those interested in building with us, we have also launched an open-source platform at the intersection of data science and investments⁸¹, where we house these tools and code, and invite practitioners to collaborate and help shape where the industry goes next.

⁷⁹ Wilson, "Explainable AI in Finance: Addressing the Needs of Diverse Stakeholders."

⁸⁰ CFA Institute, "ETHICS AND ARTIFICIAL INTELLIGENCE IN INVESTMENT MANAGEMENT: A FRAMEWORK FOR PROFESSIONALS."

⁸¹ CFA Institute Research and Policy Center, "RPC Labs."

Best Practices Checklist for Use of AI in Financial Services⁸²

- ☑ Systematically recording internal data & re-imagining data capture:
Systematically capture calls, meeting notes, emails, and investment decisions in detail to enable AI-driven insights. Redesign investment memos to emulate PM investment frameworks – to then use this rich data for future AI PM.
- ☑ Maintaining a comprehensive database of prompts:
Maintain a centralized repository of effective AI prompts for the expected output. Simple framework for effective prompts, “MARS” ie. Metrics, Actionable, Role, Scenario.
- ☑ Moving from team/data silos to a central dashboard:
Consolidate all team and data silos into a unified dashboard, creating an exhaustive knowledge base accessible by AI for comprehensive insights.
- ☑ Workflow breakdown to identify use cases:
Conduct team exercises to break down workflows, identify, and prioritize top use cases for AI implementation – this helps bring structure to the AI adoption journey.
- ☑ Regular knowledge sharing and training sessions:
Foster AI familiarity and acceptance through ongoing training sessions and knowledge exchange within investment teams.
- ☑ Developing vertical AI champions:
Appoint AI leads specialized by function (legal, compliance, investments) and asset class (equities, fixed income, FX) to drive focused AI initiatives effectively.
- ☑ Dedicated AI resources within investment teams:
Allocate dedicated AI specialists or identify internal AI leaders or hire AI specialists within investment teams to focus on use case development and vendor solution exploration.
- ☑ Find implementation partners early-on or directly adopt a third-party tool:
Engage implementation partners or adopt third-party workflow automation tools early to accelerate AI integration and operational efficiency.
- ☑ Think hybrid teams with AI and humans working together:
Prepare for future hybrid teams combining AI and human Analysts/PMs by building appropriate infrastructure and knowledge bases now.
- ☑ Defining clear evaluation matrix:
Define robust success criteria tailored to use cases—automation rates, error reduction, insight generation, bias avoidance, and time saved—to measure AI impact effectively.

⁸² Acharya, “Gen AI in Asset Management: Bernstein Lists 10 Best Practices for Investors.”

Governance In AI: India Principles and Global standards

AI is no longer optional in finance. Firms across sectors are committing large investments and skilling their workforce in using modern AI products nearly every day. Regulators have taken note of this, and are deeply thoughtful of enabling progress while maintaining safety of customers and financial markets. As a result, regulators and international bodies are embedding AI into the fabric of established governance and risk traditions – through either principles or legislations. While the standards are still evolving globally, this section presents the current guidance in India, EU, and other global organizations.

India’s approach as laid out by RBI’s FREE-AI report and SEBI’s consultation paper in 2025, follows principle over prescription. Seven guiding principles (“sutras”) and recommendations guide Indian regulators viewpoint, leaning more on flexibility and discovery. The EU’s AI act, effective August 2024, takes a more prescriptive approach: four-tier risk classification, binding rules for high-risk applications like customer profiling, credit scoring.

India Principles:

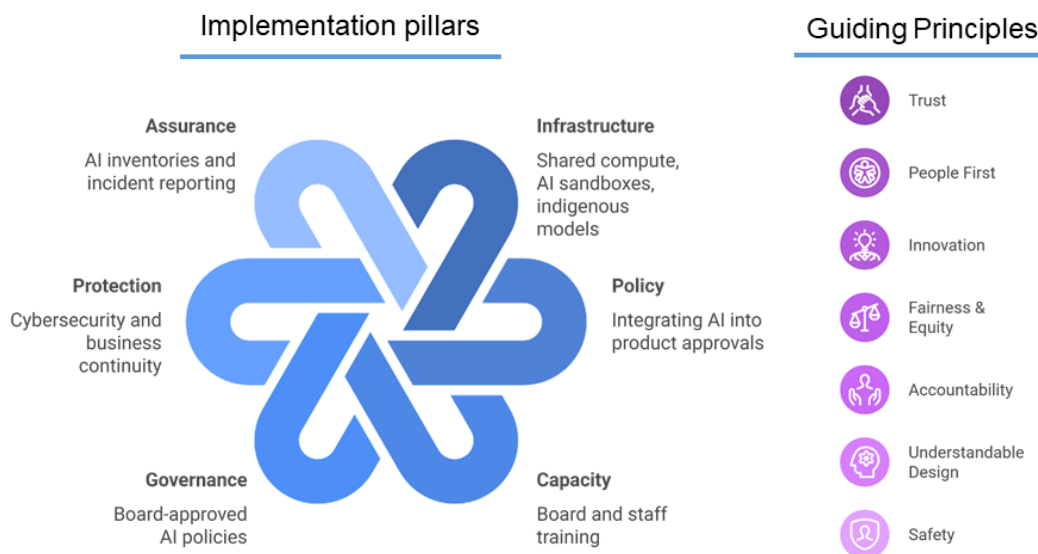
RBI’s Framework for Responsible and Ethical Enablement of Artificial Intelligence (FREE-AI⁸³) rests on an explicit philosophy: “Innovation over Restraint.” Its seven sutras are deliberately aspirational, and the framework provides 26 recommendations for AI best practices. Critically, none of this is yet mandatory. Banks and NBFCs can deploy AI without following FREE-AI. But if they do deploy, the RBI expects responsible practices: lifecycle governance, proportional explainability, bias mitigation, and vendor oversight.

Key Takeaways:

For Indian Financial Services companies - today’s principles become tomorrow’s supervisory expectations. FREE-AI and SEBI have laid out clear principles and for AI governance in 2025. Several consistent threads stand out:

- **Governance:** Boards and senior management are expected to play an active role in AI oversight.
- **Privacy and Data:** AI’s reliance on data ties it closely to privacy and residency regulations.
- **Explainability and Fairness:** Transparent, non-discriminatory outcomes are framed as central to trust.
- **Vendor Accountability:** Outsourcing norms are extending naturally to AI providers and platforms.
- **High-Risk Applications:** In credit decisions and insurance underwriting, supervisors are alert to both the promise of efficiency and the model risks, training bias, monitoring.

Figure 16: RBI FREE-AI Framework



⁸³ Reserve Bank of India, “Framework for Responsible and Ethical Enablement of Artificial Intelligence (FREE-AI).”

Securities and Exchange Board of India's consultation paper⁸⁴ from June 2025 complements this approach for capital markets. SEBI proposes a tiered structure: "regulatory lite" for internal AI tools (compliance, surveillance), full governance for client-facing applications (algorithmic trading, robo-advisory).

Unique to SEBI's framework are replay and back-out mechanisms: tools to reconstruct and potentially unwind algorithmic decisions when AI systems err in trading contexts.

It's a pragmatic acknowledgment that markets require fail-safes that banking regulation doesn't typically demand. For investor protection, SEBI proposes that market participants disclose clearly usages of AI/ML applications, especially on model

risks, accuracy, data quality, all to be made available in customer accessible language.

Both frameworks share common ground: board-level accountability is non-negotiable, explainability must be proportional to impact, and fairness isn't optional. The Digital Personal Data Protection Act⁸⁵ (2025) acts as the baseline - requiring lawful processing, data minimisation, and cross-border control into AI oversight.

For AI systems that depend on vast quantities of transactional and personal data. India's approach currently remains principle-driven rather than statutory - flexible but less enforceable. However, the ambition is clear: ensuring AI evolves within structures of trust, accountability, and fairness.

Global Standards:

Internationally, the BIS FSI Insights No. 63 (2024⁸⁶) captures the growing regulatory recognition of AI. Supervisors face challenges in ensuring explainability, monitoring systemic risks, and coordinating across borders. For example, the report highlights that Singapore's Monetary Authority has issued its "Fairness, Ethics, Accountability, and Transparency (FEAT)" principles, while the European Union has moved ahead with the AI Act.

The Basel Committee on Banking Supervision's Digitalisation of Finance report (2024⁸⁷) situates AI within a broader technological shift. While not prescribing specific AI regulations, it underlines that existing Basel principles such as those on risk data aggregation and reporting remain directly applicable. For example, the Basel report refers to pilot practices where

supervisors are adapting traditional stress-testing methods to validate AI and Machine Learning (ML) models, thereby bringing innovation into established supervisory toolkits.

The European Union's Artificial Intelligence Act (2024) represents the world's first comprehensive legislative framework for AI. Credit scoring systems and insurance underwriting are classified as "high-risk" AI under Annex III of the Act.

Any institution deploying these systems must register them in an EU database, maintain comprehensive technical documentation, conduct fundamental rights impact assessments, implement post-market monitoring, and report serious incidents. Non-compliance carries penalties up to \$35 million or 7% of annual revenues.

⁸⁴ "SEBI | Consultation Paper on Guidelines for Responsible Usage of AI/ML in Indian Securities Markets."

⁸⁵ Government of India, "DPDP Rules, 2025 Notified."

⁸⁶ Crisanto et al., "Regulating AI in the Financial Sector".

⁸⁷ Basel Committee on Banking Supervision, "Digitalisation of Finance."

Across jurisdictions, a shared ambition is evident: ensuring that AI in finance evolves within frameworks of trust, accountability, and fairness. Yet, mechanisms differ in structure and depth:

Table 5: EU AI Act and India’s Regulatory Frameworks - A Comparative View

Theme	EU AI Act (Reg. 2024/1689)	India’s RBI / SEBI / DPDP Frameworks	Broad Comparison
Governance	Statutory risk-management & quality systems (Arts 9, 17)	Supervisory lifecycle governance	Shared intent; EU legally binding
Fairness & Data	Representative data & mandatory rights-impact review (Arts 10, 27)	Fairness Sutra, ethical AI use	India principle-based; EU assessment-based
Transparency	Legal right to explanation (Art 86)	Proportional explainability guidance	EU grants enforceable right
Oversight & Enforcement	Human oversight, post-market monitoring, fines (Arts 14, 72, 99)	Continuous monitoring, no statutory penalties	Similar spirit, differing enforceability
Institutional Setup	Financial supervisors designated in law (Recitals 154-158)	Supervisors act via existing mandates	Functionally aligned, structurally different

Introducing EU-style instruments such as a Fundamental Rights Impact Assessment or a Post-Market Monitoring Plan could enhance India’s AI lifecycle by embedding fairness and accountability checkpoints. However, such measures would require new supervisory capacity and inter-regulatory coordination. India’s principle-driven model offers flexibility but lacks the EU’s enforceable precision.

Key Takeaways:

- For Indian Financial Services companies - today’s principles become tomorrow’s supervisory expectations. FREE-AI and SEBI have laid out clear principles and for AI governance in 2025.
- Establish Board-approved AI policies, and AI awareness throughout the organization.
- Important areas - lifecycle management, bias audits, replay and back-out mechanisms.
- On designated high-risk areas such as credit decisions, insurance underwriting, EU act has put in place extensive requirements for model documentation, assessment, monitoring.
- Vendor strategy is an important focus – companies are held solely responsible for all AI usages and outputs. Building optionality into AI architecture (multi-vendor strategies, fallback mechanisms) becomes prudent risk management.

Conclusion

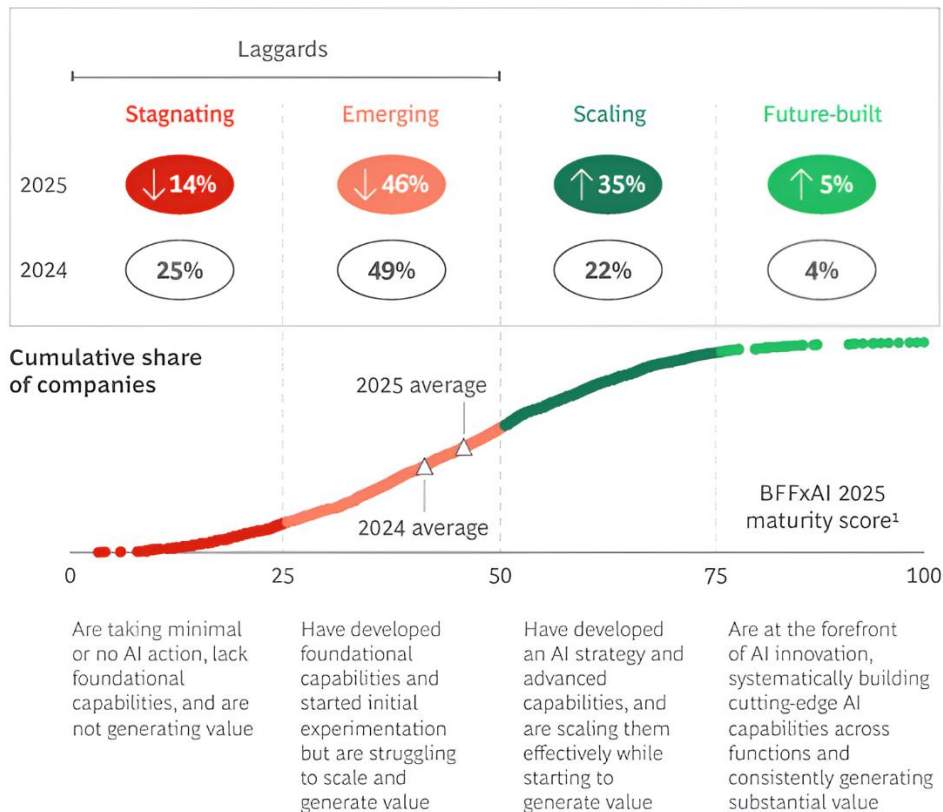
"Art is never finished, only abandoned"
— Leonardo da Vinci

It is now evident that AI has a strong potential to transform the financial services industry. Some argue that it is already well underway and that AI is a necessity to stay competitive. However, limitations of AI models, pitfalls of over-reliance on AI and concerns about data privacy & security must not to be overlooked. The transformative impact on jobs and talent management is also obvious, with pervasive adoption expected to become the norm in the industry. A BCG study⁸⁸ in 2025 showed only 5% of about 1,250 companies worldwide are achieving AI value at scale, a

measure of how tough the full AI transformation is. 60% of companies are not achieving material value at all, reporting minimal revenue and cost gains despite substantial investment. Another 35% (13 percentage points more than in 2024) are scaling up their efforts and seeing some returns, but many of them admit that they are not moving far enough or fast enough. While this study is not specific to financial firms, it does tell a trend that there is a long way to go for AI benefits to be substantially derived.

Figure 17: Only 5% of Companies Get Substantial Value from AI investments

Share of companies



Source: BCG Build for the Future 2025 Global Study (n = 1,250).
¹This score assesses AI maturity across 41 dimensions.
²Future-built versus stagnating + emerging.
³External metrics (Capital IQ): total shareholder return (June 22–May 25 for three-year TSR).

⁸⁸ Apotheke et al., "The Widening Gap."

- **Inter-operable systems:**
With multiple vendor products being used, there will be need for inter-operable systems. Systems will need to integrate with in-house and different vendor products in order to drive consistent results across functions and extract benefits of AI investments at enterprise level.
 - **Human + AI:**
Degree of interaction to vary as per various functions:
 - **Assisted Intelligence -**
The machine is essentially a tool that is used by a human operator. Example, speech recognition software, predictive text algorithms.
 - **Augmented intelligence -**
The machine is an active participant in the decision-making process. For example, customer service chatbots.
 - **Autonomous intelligence -**
Machine is a fully independent agent. For example, robot advisors.
 - **Discover, Alert & Act:**
System pushes alerts and acts autonomously or assists to act - progressing from simple screeners and unintelligent dashboard displays.
 - **Misuse potential also high:**
As use of AI/ML rise in enterprises, a parallel misuse of the same technology has the potential to rapidly rise. Organizations need to be proactive to nip such threats at the bud and additionally prevent such threats.
- While each of these types of intelligence is distinct, they are not mutually exclusive. For example, a machine learning algorithm might be used as an assisted intelligence tool to help a human operator analyze data, but it could also be used as an autonomous intelligence system to make decisions and take action on its own. Similarly, a chatbot might be used as an augmented intelligence tool to help a customer service agent answer question, but it could also be used as an autonomous intelligence system to handle customer inquiries without human intervention.

Key Takeaways for CFA Charter holders and Students

- All firms we interviewed are expanding their AI/ML and Data science teams.
- Teams require multi-disciplinary members due to implementation requiring a combination of technology, mathematics/statistics and specific subject matter expertise (investment management, insurance, credit underwriting, marketing, customer service).
- CFA Program curriculum and Professional Education lays emphasis on all aspects through various topics - Macroeconomics, Statistics, Python coding and Data science - besides the core investment management subjects.
- Students and Charter holders with such multi-skills are best positioned to take part in firms implementing AI/ML tools.
- Coding background to Investment Management OR Investing Management background to Coding? Both cases exist.
- Quant funds are gradually gathering acceptance. More importantly, quant-like approach to discretionary fund management and research is gaining rapid acceptance.
- In existing roles, students and charter holders can look to start thinking on Data-led framework for decisions, Systematic application of methods, Automation of routine tasks, discovering ways to uncover patterns (otherwise missed by humans) through AI/ML use.

Bibliography

- Mustafa Suleyman. 2023. *The Coming Wave: echnology, Power, and the Twenty-first Century's Greatest Dilemma*. <http://archive.org/details/the-coming-wave-by-mustafa-suleyman-michael-bhaskar-pdfread.net>.
- Challapally, Aditya, Chris Pease, Ramesh Raskar, and Pradyumna Chari. 2025. "The GenAI Divide STATE of AI in BUSINESS 2025 MIT NANDA." https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf.
- Chloetejada. "Big Tech's AI Expansion: From Investment to Scalable Returns." *RBC Wealth Management - Asia*, February 3, 2026. <https://www.rbcwealthmanagement.com/en-asia/insights/big-techs-ai-expansion-from-investment-to-scalable-returns>.
- Mary Meeker et al., "Trends - Artificial Intelligence," April 30, 2025, https://www.bondcap.com/report/pdf/Trends_Artificial_Intelligence.pdf.
- Kwa, Thomas, Ben West, Joel Becker, Amy Deng, Katharyn Garcia, Max Hasin, Sami Jawhar, et al. "Measuring AI Ability to Complete Long Tasks," February 25, 2025. <https://arxiv.org/pdf/2503.14499>.
- Rajiv Memani, "The Aldea of India: 2025," The Aldea of India: 2025, n.d., <https://www.ey.com/content/dam/ey-unified-site/ey-com/en-in/services/ai/aidea/2025/01/ey-the-aidea-of-india-2025-how-much-productivity-can-genai-unlock-in-india.pdf>.
- NITI Aayog, "Transforming India With AI," report, October 12, 2025, https://sansad.in/getFile/loksabhaquestions/annex/185/AU4388_SJEMq6.pdf?source=pqals.
- Everett M. Rogers et al., *Diffusion of Innovations*, Third Edition (The Free Press, A Division of Macmillan Publishing Co., Inc., 1983), <https://teddykw2.files.wordpress.com/2012/07/everett-m-rogers-diffusion-of-innovations.pdf>.
- Labaree, Robert V. "Research Guides: Organizing Your Social Sciences Research Paper: Qualitative Methods." Research Guide. Accessed May 5, 2026. <https://libguides.usc.edu/writingguide/qualitative>.
- Carter, Nancy, Denise Bryant-Lukosius, Alba DiCenso, Jennifer Blythe, and Alan J. Neville. "The Use of Triangulation in Qualitative Research." *Oncology Nursing Forum* 41, no. 5 (September 2014): 545-47. <https://doi.org/10.1188/14.ONF.545-547>.
- Joseph J. Gal, "Man-Machine Interactive Systems and Their Application to Financial Analysis," *Financial Analysts Journal* 22, no. 3 (May 1, 1966): 126-36, <https://doi.org/10.2469/faj.v22.n3.126>.
- Black, Fischer, and Myron Scholes. 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy* 81 (3): 637-54. <https://www.jstor.org/stable/1831029>.
- Shihao Gu et al., "Empirical Asset Pricing via Machine Learning," *The Review of Financial Studies* 33-33 (2020): 2223-73, <https://dachxiu.chicagobooth.edu/download/ML.pdf>.
- "RAG For Finance: Automating Document Analysis With LLMs," CFA Institute Research and Policy Center, n.d., <https://rpc.cfainstitute.org/research/the-automation-ahead-content-series/retrieval-augmented-generation>.
- "Intelligent Financial System: How AI Is Transforming Finance," BIS Working Papers, 2024, <https://www.bis.org/publ/work1194.pdf>.
- Söhnke M. Bartram et al., "Artificial Intelligence in Asset Management," CFA Institute Research Foundation / Literature Review, 2020, <https://www.cfainstitute.org/-/media/documents/book/rf-lit-review/2020/rflr-artificial-intelligence-in-asset-management.pdf#:~:text=URL%3A%20https%3A%2F%2Fwww.cfainstitute.org%2F>.
- Joseph Simonian, "AI Washing: Signs, Symptoms, and Suggested Solutions for Investment Stakeholders," June 10, 2025, <https://doi.org/10.56227/25.1.20>.
- Li, Yinheng, Shaofei Wang, Han Ding, and Hang Chen. n.d. "Large Language Models in Finance: A Survey." <https://arxiv.org/pdf/2311.10723>.
- Antonio Jimeno Yepes et al., "Financial Report Chunking for Effective Retrieval Augmented Generation," arXiv.org, February 5, 2024, <https://arxiv.org/abs/2402.05131>.
- Dongsheng Wang, Natraj Raman, Mathieu Sibue, Zhiqiang Ma, Petr Babkin, Simerjot Kaur, Yulong Pei, Armineh Nourbakhsh, and Xiaomo Liu, "DocLLM: A Layout-Aware Generative Language Model for Multimodal Document Understanding," arXiv preprint, 2023, <https://arxiv.org/pdf/2401.00908>.
- Shijie Wu et al., "BloombergGPT: A Large Language Model for Finance," arXiv.org, March 30, 2023, <https://arxiv.org/abs/2303.17564>.
- Masahiro Suzuki et al., "Constructing and Analyzing Domain-Specific Language Model for Financial Text Mining," journal-article, Preprint Submitted to Elsevier, 2021, <https://msuzuki.me/paper/2023/ipm.pdf>.
- Anish Ratnawat, "Exploring Retrieval-Augmented Generation (RAG) with Vector Databases and AI Agents," n.d., <https://anishratnawat.com/exploring-retrieval-augmented-generation-rag-with-vector-databases-and-ai-agents>.
- Alejandro Lopez-Lira and Yuehua Tang, "Can ChatGPT Forecast Stock Price Movements? Return Predictability and Large Language Models," SSRN Electronic Journal, 2023, <https://doi.org/10.2139/ssrn.4412788>.
- "Financial Sentiment Analysis on News and Reports Using Large Language Models and FinBERT," IEEE Conference Publication | IEEE Xplore, July 26, 2024, <https://ieeexplore.ieee.org/document/10796670>.

- Xiang Deng, Vasilisa Bashlovkina, Feng Han, Simon Baumgartner, and Michael Bendersky, "What Do Large Language Models Know about Financial Markets? A Case Study on Reddit Market Sentiment Analysis," arXiv preprint, 2023, <https://arxiv.org/pdf/2212.11311>
- "Interpreting FedSpeak With Confidence: A LLM-Based Uncertainty-Aware Framework Guided by Monetary Policy Transmission Paths," n.d., <https://arxiv.org/html/2508.08001v1>.
- Edward L. Glaeser, Hyunjin Kim, and Michael Luca, "Nowcasting the Local Economy Using Yelp Data," SSRN Electronic Journal, 2023, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4337182
- Kausik Lakkaraju et al., "LLMs for Financial Advisement: A Fairness and Efficacy Study in Personal Decision Making," in Proceedings of the 4th ACM International Conference on AI in Finance (ICAIF '23), 2023, <https://doi.org/10.1145/3604237.3626867>
- G. Srividya et al., "Personalized Finance Chatbot Powered by RAG and Generative AI for Smart Wealth Management," Zenodo (CERN European Organization for Nuclear Research), April 2, 2025, <https://doi.org/10.5281/zenodo.18108271>.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell, "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?" in Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21), 2021, <https://doi.org/10.1145/3442188.3445922>
- Udit Gupta, "GPT-InvestAR: Enhancing Stock Investment Strategies through Annual Report Analysis with Large Language Models," SSRN Electronic Journal, 2023, <https://doi.org/10.2139/ssrn.4568964>
- Hongyang Yang, Xiao-Yang Liu, and Christina Dan Wang, "FinGPT: Open-Source Financial Large Language Models," arXiv.org, June 9, 2023, <https://arxiv.org/abs/2306.06031>.
- Tejaswini Kumar, Kishor Yadav Kommanaboina, and Bhargava Kumar, "Next-Generation Text-to-SQL: A Survey of Advanced Reasoning Enhancements Techniques," Zenodo (CERN European Organization for Nuclear Research), September 4, 2024, <https://doi.org/10.5281/zenodo.18137583>.
- Kalyani Pakhale, "Comprehensive Overview of Named Entity Recognition: Models, Domain-Specific Applications and Challenges," arXiv preprint, 2023, <https://arxiv.org/pdf/2309.14084>
- Dimitrios Vamvourellis et al., "Company Similarity Using Large Language Models," arXiv.org, August 15, 2023, <https://arxiv.org/abs/2308.08031>.
- Xiaoyang Li et al., "Can ChatGPT Reduce Human Financial Analysts' Optimistic Biases?," Economic and Political Studies 12, no. 1 (November 23, 2023): 20–33, <https://doi.org/10.1080/20954816.2023.2276965>.
- Alex Kim, Maximilian Muhn, and Valeri Nikolaev, "Financial Statement Analysis With Large Language Models," arXiv.org, July 25, 2024, <https://arxiv.org/abs/2407.17866>.
- Qingsong Wen, Tian Zhou, Chaoli Zhang, Weiqi Chen, Ziqing Ma, Junchi Yan, and Liang Sun, "Transformers in Time Series: A Survey," arXiv preprint, 2022, <https://arxiv.org/pdf/2202.07125>
- Yuxuan Liang, Haomin Wen, Yuqi Nie, Yushan Jiang, Ming Jin, Dongjin Song, Shirui Pan, and Qingsong Wen, "Foundation Models for Time Series Analysis: A Tutorial and Survey," arXiv preprint, 2024, <https://arxiv.org/abs/2403.14735>
- Yaxuan Kong, Yuqi Nie, Xiaowen Dong, John M. Mulvey, H. Vincent Poor, Qingsong Wen, and Stefan Zohren, "Large Language Models for Financial and Investment Management: Applications and Benchmarks," Journal of Portfolio Management 51, no. 2 (2024): 162–210, <https://doi.org/10.3905/jpm.2024.1.645>
- Abhimanyu Das, Weihao Kong, Rajat Sen, and Yichen Zhou, "A Decoder-Only Foundation Model for Time-Series Forecasting," arXiv preprint, 2024, <https://arxiv.org/pdf/2310.10688>
- Nate Gruver, Marc Finzi, Shikai Qiu, and Andrew Gordon Wilson, "Large Language Models Are Zero-Shot Time Series Forecasters," arXiv preprint, 2023, <https://arxiv.org/abs/2310.07820>
- Yaxuan Kong, Yuqi Nie, Xiaowen Dong, John M. Mulvey, H. Vincent Poor, Qingsong Wen, and Stefan Zohren, "Large Language Models for Financial and Investment Management: Applications and Benchmarks," Journal of Portfolio Management 51, no. 2 (2024): 162–210, <https://doi.org/10.3905/jpm.2024.1.645>
- Xinli Yu, Zheng Chen, Yuan Ling, Shujing Dong, Zongyi Liu, and Yanbin Lu, "Temporal Data Meets LLM: Explainable Financial Time Series Forecasting," arXiv preprint, 2023, <https://arxiv.org/pdf/2306.11025>
- Alejandro Lopez-Lira and Yuehua Tang, "Can ChatGPT Forecast Stock Price Movements? Return Predictability and Large Language Models," SSRN Electronic Journal, 2023, <https://doi.org/10.2139/ssrn.4412788>
- Chenwei Gong, Yanyi Zhong, Shenghan Zhao, and Yingchia Liu, "Application of Machine Learning in Predicting Extreme Volatility in Financial Markets: Based on Unstructured Data," TechRxiv preprint, 2024, <https://www.techrxiv.org/users/854434/articles/1248352-application-of-machine-learning-in-predicting-extreme-volatility-in-financial-markets-based-on-unstructured-data>
- Qianqian Xie, Weiguang Han, Yanzhao Lai, Min Peng, and Jimin Huang, "The Wall Street Neophyte: A Zero-Shot Analysis of ChatGPT Over Multimodal Stock Movement Prediction Challenges," arXiv preprint, 2023, <https://arxiv.org/abs/2304.05351>
- Alejandro Lopez-Lira and Yuehua Tang, "Can ChatGPT Forecast Stock Price Movements? Return Predictability and Large Language Models," SSRN Electronic Journal, April 2023, <https://doi.org/10.2139/ssrn.4412788>

- Hyungjin Ko and Jaewook Lee, "Can ChatGPT Improve Investment Decision? From a Portfolio Management Perspective," SSRN Electronic Journal, 2023, <https://doi.org/10.2139/ssrn.4390529>
- Thomas Yue and Chi Chung Au, "GPTQuant's Conversational AI: Simplifying Investment Research for All," SSRN Electronic Journal, 2023, <https://doi.org/10.2139/ssrn.4380516>
- Constantin J. Schneider and Yahya Yilmaz, "Stock Portfolio Selection Based on Risk Appetite: Evidence from ChatGPT," Finance Research Letters 82 (2025): 107517, <https://doi.org/10.1016/j.frl.2025.107517>
- Nikolas Anic, Andrea Barbon, Ralf Seiz, and Carlo Zarattini, "ChatGPT in Systematic Investing—Enhancing Risk-Adjusted Returns with LLMs," arXiv preprint, 2025, <https://arxiv.org/abs/2510.26228>
- Zihao Zhou and Rose Yu, "Can LLMs Understand Time Series Anomalies?," arXiv.org, October 7, 2024, <https://arxiv.org/abs/2410.05440>.
- Duanyu Feng et al., "Empowering Many, Biasing a Few: Generalist Credit Scoring Through Large Language Models," arXiv.org, October 1, 2023, <https://arxiv.org/abs/2310.00566>.
- Sean S. Cao, Wei Jiang, Lijun (Gillian) Lei, and Qing (Clara) Zhou, "Applied AI for Finance and Accounting: Alternative Data and Opportunities," Pacific-Basin Finance Journal 84 (2024): 102307, <https://doi.org/10.1016/j.pacfin.2024.102307>
- Xinze Yang, Chunkai Zhang, Yizhi Sun, Kairui Pang, Luru Jing, Shiyun Wa, and Chunli Lv, "FinChain-BERT: A High-Accuracy Automatic Fraud Detection Model Based on NLP Methods for Financial Scenarios," Information 14, no. 9 (2023): 499, <https://doi.org/10.3390/info14090499>
- Alex Kim, Maximilian Muhn, and Valeri V. Nikolaev, "From Transcripts to Insights: Uncovering Corporate Risks Using Generative AI," SSRN Electronic Journal, October 5, 2023, <https://doi.org/10.2139/ssrn.4593660>
- Miquel Noguer i Alonso and Hanane Dupouy, "Evaluating LLMs in Financial Tasks - Code Generation in Trading Strategies," SSRN Electronic Journal, March 8, 2024, <https://doi.org/10.2139/ssrn.4752797>
- Hang Yuan, Saizhuo Wang, and Jian Guo, "Alpha-GPT 2.0: Human-in-the-Loop AI for Quantitative Investment," arXiv.org, February 15, 2024, <https://arxiv.org/abs/2402.09746>.
- Yangyang Yu et al., "FinMem: A Performance-Enhanced LLM Trading Agent With Layered Memory and Character Design," arXiv.org, November 23, 2023, <https://arxiv.org/abs/2311.13743>.
- Kemal Kirtac and Guido Germano, "Sentiment Trading with Large Language Models," Finance Research Letters 62 (2024): 105227, <https://doi.org/10.1016/j.frl.2024.105227>
- Yijia Xiao, Edward Sun, Di Luo, and Wei Wang, "TradingAgents: Multi-Agents LLM Financial Trading Framework," arXiv preprint, 2024, <https://arxiv.org/abs/2412.20138>
- Yang Li, Yangyang Yu, Haohang Li, Zhi Chen, and Khaldoun Khashanah, "TradingGPT: Multi-Agent System with Layered Memory and Distinct Characters for Enhanced Financial Trading Performance," arXiv preprint, 2023, <https://arxiv.org/abs/2309.03736>
- Xuwen Han, Neng Wang, Shangkun Che, Hongyang Yang, Kunpeng Zhang, and Sean Xin Xu, "Enhancing Investment Analysis: Optimizing AI-Agent Collaboration in Financial Research," in Proceedings of the 5th ACM International Conference on AI in Finance (ICAIF '24), November 2024, <https://doi.org/10.1145/3677052.3698645>
- Gautam Jajoo, Pranjal A. Chitale, and Saksham Agarwal, "MASCA: LLM-Based Multi-Agents System for Credit Assessment," arXiv preprint, 2025, <https://arxiv.org/abs/2507.22758>
- Taejin Park, "Enhancing Anomaly Detection in Financial Markets with an LLM-Based Multi-Agent Framework," arXiv preprint, 2024, <https://arxiv.org/abs/2403.19735>
- Erik Brynjolfsson, Danielle Li, and Lindsey Raymond, "Generative AI at Work," National Bureau of Economic Research, April 1, 2023, <https://doi.org/10.3386/w31161>.
- Jochen Wulf and Juerg Meierhofer, "Exploring the Potential of Large Language Models for Automation in Technical Customer Service," arXiv preprint, 2024, <https://arxiv.org/abs/2405.09161>
- Mubashar Raza et al., "Industrial Applications of Large Language Models," Scientific Reports 15, no. 1 (April 21, 2025): 13755, <https://doi.org/10.1038/s41598-025-98483-1>.
- Joseph Brand, Amit Israeli, and David Ngwe, "Using LLMs for Market Research," Harvard Business School Working Paper No. 23-062, 2024, https://www.hbs.edu/ris/Publication%20Files/23-062_6bfe7a5b-3ed9-4afd-a4c1-c91b60dd82e5.pdf
- Md Asraful Haque, "LLMs: A Game-Changer for Software Engineers?" BenchCouncil Transactions on Benchmarks, Standards and Evaluations 5, no. 1 (2025): 100204, <https://doi.org/10.1016/j.tbench.2025.100204>
- Sanka Rasnayaka, Guanlin Wang, Ridwan Shariffdeen, and Ganesh Neelakanta Iyer, "An Empirical Study on Usage and Perceptions of LLMs in a Software Engineering Project," in Proceedings of the IEEE Conference, 2024, <https://ieeexplore.ieee.org/document/10734434>
- Shakked Noy and Whitney Zhang, "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence," SSRN Electronic Journal, March 2023, <https://doi.org/10.2139/ssrn.4375283>

- ETtech, "Groww Completes Rs 175 Crore Acquisition of Indiabulls' Mutual Fund Business," *The Economic Times*, May 4, 2023, <https://economictimes.indiatimes.com/tech/startups/groww-completes-rs-175-6-crore-acquisition-of-indiabulls-mutual-fund-business/articleshow/99982760.cms?from=mdr>.
- Harsh Upadhyay, "Groww Completes Acquisition of Fisdom," *Entrackr*, October 7, 2025, <https://entrackr.com/news/groww-completes-acquisition-of-fisdom-10536344>.
- ETtech, "Groww Creditserv's Loan Book Grows to Rs 965 Crore by June," *The Economic Times*, October 3, 2024, <https://m.economictimes.com/tech/startups/groww-creditservs-loan-book-grows-to-rs-965-crore-by-june/articleshow/113904085.cms>.
- Harsh Upadhyay, "Groww Receives RBI Approval to Operate as Payment Aggregator," *Entrackr*, April 30, 2024, <https://entrackr.com/2024/04/groww-receives-rbi-approval-to-operate-as-payments-aggregator/>.
- Greggirth and Greggirth, "The AI Law Professor: When the New AI Model Disappoints," Thomson Reuters Institute, September 22, 2025, <https://www.thomsonreuters.com/en-us/posts/technology/ai-law-professor-model-disappoints/>.
- Lakshmi Varanasi, "AI Leaders Have a New Term for the Fact That Their Models Are Not Always so Intelligent," *Business Insider*, June 7, 2025, <https://www.businessinsider.com/aji-artificial-jagged-intelligence-google-ceo-sundar-pichai-2025-6>.
- Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askell, Samuel R. Bowman, et al., "Towards Understanding Sycophancy in Language Models," arXiv preprint, 2023, <https://arxiv.org/abs/2310.13548>.
- Lars Malmqvist, "Sycophancy in Large Language Models: Causes and Mitigations," arXiv.org, November 22, 2024, <https://arxiv.org/abs/2411.15287>; Jerry Wei et al., "Simple Synthetic Data Reduces Sycophancy in Large Language Models," arXiv.org, August 7, 2023, <https://arxiv.org/abs/2308.03958>.
- "LIVE: Jeff Bezos Speaks at Italian Tech Week 2025," YouTube, n.d., https://www.youtube.com/live/4wTSZDZ_seU?si=K75XX8DgEHWk0GZn.
- Paul Mozur, Adam Satariano, and Emiliano Rodríguez Mega, "From Mexico to Ireland, Fury Mounts Over a Global A.I. Frenzy," *The New York Times*, October 20, 2025, <https://www.nytimes.com/2025/10/20/technology/ai-data-center-backlash-mexico-ireland.html>
- Cheryll-Ann Wilson, "Explainable AI in Finance: Addressing the Needs of Diverse Stakeholders," August 7, 2025, <https://doi.org/10.56227/25.1.25>.
- CFA Institute, "ETHICS AND ARTIFICIAL INTELLIGENCE IN INVESTMENT MANAGEMENT: A FRAMEWORK FOR PROFESSIONALS," 2022, https://www.cfainstitute.org/-/media/documents/article/industry-research/Ethics-and-Artificial-Intelligence-in-Investment-Management_Online.pdf.
- "Welcome to RPC Labs," CFA Institute Research and Policy Center, n.d., <https://rpc.cfainstitute.org/themes/technology/rpclabs>.
- Navamya Acharya, "Gen AI in Asset Management: Bernstein Lists 10 Best Practices for Investors," *Investing.com*, August 2, 2025, <https://www.investing.com/news/stock-market-news/gen-ai-in-asset-management-bernstein-lists-10-best-practices-for-investors-4155407>
- Reserve Bank of India, "Framework for Responsible and Ethical Enablement of Artificial Intelligence (FREE-AI)," 2025, <https://rbidocs.rbi.org.in/rdocs/PublicationReport/Pdfs/FREEAIR130820250A24FF2D4578453F824C72ED9F5D5851.PDF>
- Securities and Exchange Board of India, "Consultation Paper on Guidelines for Responsible Usage of AI/ML in Indian Securities Markets," June 20, 2025, https://www.sebi.gov.in/reports-and-statistics/reports/jun-2025/consultation-paper-on-guidelines-for-responsible-usage-of-ai-ml-in-indian-securities-markets_94687.html
- Government of India, "DPDP Rules, 2025 Notified," press-release, *Digital Personal Data Protection Rules*, October 17, 2025, <https://static.pib.gov.in/WriteReadData/specificdocs/documents/2025/nov/doc20251117695301.pdf>.
- Crisanto, Juan Carlos, Cris Benson Leuterio, Jermy Prenio, and Jeffery Yong. "Regulating AI in the Financial Sector: Recent Developments and Main Challenges." December 12, 2024. <https://www.bis.org/fsi/publ/insights63.htm>.
- Basel Committee on Banking Supervision, "Digitalisation of Finance," Bank for International Settlements, May 2024, <https://www.bis.org/bcbs/publ/d575.htm>
- Jessica Apotheke, Vinciane Beauchene, Nicolas de Bellefonds, Patrick Forth, Marc Roman Franke, Michael Grebe, Nina Kataeva, et al., "The Widening AI Value Gap," BCG Global, September 30, 2025, <https://www.bcg.com/publications/2025/are-you-generating-value-from-ai-the-widening-gap>
- "LIVE: Jeff Bezos Speaks at Italian Tech Week 2025," YouTube, n.d., https://www.youtube.com/live/4wTSZDZ_seU?si=K75XX8DgEHWk0GZn.



STATE OF AI IN FINANCIAL SERVICES IN INDIA
Adoption Trends, Industry Applications, Risks, and Governance

#Unit No H4, 12th Floor, B Wing,
"Parinee Crescenzo", G Block, BKC,
Bandra East, Mumbai, Maharashtra 400051

advocacy@india.cfasociety.org | <https://cfasocietyindia.org>

