

Intelligent AI Finance Advisory System: A Secure and Ethical AI System for Live Financial Advisory Insights

Mr. Seshu Kumar
Assistant Professor

Department of Computer Science and Engineering
Keshav Memorial Institute of Technology
Hyderabad, Telangana, India

Mr. Kirtan Pokar
Student (UG Scholar)

Department of Computer Science and Engineering
Keshav Memorial Institute of Technology
Hyderabad, Telangana, India

Mr. Sri Charan Kolachalama
Student (UG Scholar)

Department of Computer Science and Engineering
Keshav Memorial Institute of Technology
Hyderabad, Telangana, India

Mr. Banda Ashish Reddy
Student (UG Scholar)

Department of Computer Science and Engineering
Keshav Memorial Institute of Technology
Hyderabad, Telangana, India

Abstract - The rapid expansion of digital financial platforms has created an urgent need for intelligent systems that deliver personalised, transparent, and responsible financial guidance. Existing financial advisory chatbots and platforms suffer from critical limitations: they produce generic, context-free responses; fail to integrate real-time financial data; offer no transparency in their reasoning; and provide no safeguards against speculative or unethical queries. In this paper, we propose an Intelligent Ethical and Explainable AI-Based Financial Advisory System that combines Retrieval-Augmented Generation (RAG), dynamic financial risk profiling, an ethical and compliance filtering layer, and an Explainable AI (XAI) module to deliver accountable, personalised, and grounded financial decision support. Unlike profit-prediction engines, our system is designed explicitly as a decision-support tool, preserving user autonomy while reducing misinformation risk. Experimental evaluation on a curated financial query dataset demonstrates an 87.3% factual accuracy rate (versus 61.4% for a vanilla LLM baseline), a hallucination reduction from 34.2% to 9.1%, an ethical filter precision of 93.6%, and a mean user satisfaction score of 4.5/5. These results confirm that RAG-grounded, ethically constrained, and explainable architectures represent a significant advancement over generic conversational agents in the regulated financial domain.

Keywords - Retrieval-Augmented Generation, Explainable Artificial Intelligence, Financial Advisory System, Risk Profiling, Ethical AI, Large Language Models, Hallucination Mitigation, Decision Support Systems.

I. INTRODUCTION

The financial services sector has undergone a fundamental transformation driven by the proliferation of mobile banking, algorithmic trading platforms, and AI-powered robo-advisors. Millions of retail investors, many of whom lack formal financial education, now make consequential investment decisions mediated by digital interfaces. This democratisation of financial access brings both opportunity and risk: opportunity because previously underserved populations can access investment guidance; risk because poorly designed advisory

systems can propagate misinformation, encourage speculative behaviour, or fail to account for individual financial circumstances [1], [2].

Large Language Models (LLMs) such as GPT-4 [3] and LLaMA-2 [4] have demonstrated remarkable capability in open-domain question answering and conversational reasoning. However, their deployment in high-stakes financial contexts introduces a well-documented pathology: hallucination, defined as the generation of plausible but factually incorrect information [5]. In finance, even minor inaccuracies—an incorrect interest rate, an outdated mutual fund NAV, a fabricated regulatory rule—can translate directly into financial loss for end users. Standalone LLMs therefore cannot be safely deployed as financial advisors without significant architectural safeguards.

Retrieval-Augmented Generation (RAG) [6] addresses the hallucination problem by grounding model responses in retrieved, verifiable documents before generation. However, RAG alone is insufficient for responsible financial advisory: it does not personalise recommendations to individual risk profiles, does not filter speculative or unethical queries, and does not explain why a particular recommendation is appropriate for a given user. Prior systems such as FinGPT [9] and BloombergGPT [10] address domain adaptation but do not incorporate ethical filtering or explainability modules. This gap motivates the present work.

We propose a modular, layered system that integrates five interdependent components: (i) a conversational user interaction layer that collects structured financial profiles; (ii) a dynamic risk profiling module that classifies users into low, medium, or high risk tolerance categories; (iii) an ethical and compliance filtering layer that intercepts speculative, misleading, or potentially harmful queries; (iv) a RAG engine that retrieves real-time financial data and grounds LLM responses in factual context; and (v) an XAI layer that generates human-readable

justifications for each recommendation, citing the specific user attributes and retrieved data that influenced the output. The system is complemented by a financial goal tracking and visualisation module that converts abstract savings targets into measurable milestones.

The principal contributions of this paper are as follows: (1) a unified, open-source RAG-based architecture for ethical financial advisory; (2) a novel ethical filtering mechanism evaluated on a taxonomy of speculative and unethical financial query types; (3) a dynamic risk profiling algorithm operating on five user-derived attributes; (4) integration of XAI justifications that improve user trust and transparency; and (5) a comprehensive empirical evaluation comparing the proposed system against multiple baselines across accuracy, hallucination rate, and user satisfaction metrics.

The remainder of this paper is structured as follows. Section II reviews related work. Section III presents the problem statement and system objectives. Section IV describes the proposed system architecture in detail. Section V outlines the implementation methodology. Section VI presents experimental results. Section VII discusses findings and limitations. Section VIII concludes with directions for future work.

II. RELATED WORK

A. Financial Advisory Systems and Chatbots

Early financial advisory chatbots relied on rule-based dialogue management and static knowledge bases, offering limited personalisation and brittle handling of out-of-vocabulary queries [1]. The advent of transformer-based language models enabled more fluent, open-ended financial conversations, but introduced the hallucination risk noted above. Shah et al. [7] surveyed twenty commercial financial chatbots and found that fewer than 30% cited their data sources, and none employed explicit ethical filtering. Koa et al. [8] demonstrated that standard GPT-3.5-based financial assistants produce factually incorrect responses in up to 38% of queries involving real-time market data, motivating retrieval-augmented approaches.

B. Domain-Adapted Language Models for Finance

BloombergGPT [10] introduced a 50-billion parameter LLM pre-trained on a 363-billion token financial corpus, achieving strong performance on financial NLP benchmarks including sentiment analysis, named entity recognition, and question answering. However, BloombergGPT is a proprietary, closed-source model with no ethical filtering component and is inaccessible to the broader research community. FinGPT [9] proposed an open-source alternative with lightweight RLHF fine-tuning on financial instruction datasets, demonstrating competitive performance at substantially lower computational cost. Neither system, however, incorporates dynamic user risk profiling, real-time retrieval grounding, or XAI justification generation.

C. Retrieval-Augmented Generation

Lewis et al. [6] introduced RAG as a general framework combining dense retrieval with seq2seq generation, demonstrating significant improvements in knowledge-intensive NLP

tasks. Subsequent work applied RAG to domain-specific settings including biomedical question answering [11] and legal document analysis [12]. In finance, RAG has been applied to earnings call summarisation and regulatory document retrieval [13], but its application to interactive personalised advisory, particularly with ethical constraints and XAI components, remains underexplored.

D. Explainable AI in Finance

The importance of explainability in financial AI is well-established: the EU General Data Protection Regulation (GDPR) Article 22 mandates the right to explanation for automated financial decisions [14]. Prior work has applied LIME [15] and SHAP [16] to explain credit scoring and loan approval models. However, these post-hoc explanation methods are designed for tabular ML models rather than generative language systems. Our XAI layer instead generates natural language justifications inline with recommendations, citing specific retrieved documents and user attributes, an approach more aligned with the conversational nature of advisory interactions.

E. Ethical AI in Financial Contexts

The ethical dimensions of AI-driven financial advice have received growing attention. Cao et al. [17] identified four principal risks in AI financial advisory: misinformation, speculative encouragement, privacy violation, and regulatory non-compliance. Mittelstadt et al. [18] argued that transparency and accountability mechanisms are prerequisites for trustworthy AI deployment in regulated domains. Our ethical filtering module directly addresses the misinformation and speculative encouragement risks through a multi-class query classification approach, extending prior work by providing empirical evaluation on a standardized test suite.

III. PROBLEM STATEMENT AND OBJECTIVES

A. Problem Statement

Formally, the problem addressed in this work is as follows. Given a user U characterised by a financial profile $P = \text{age, income, monthly savings, goals, investment horizon}$ and a natural language query Q , design a system S that: (i) classifies U into a risk tolerance category $R \in \{\text{Low, Medium, High}\}$; (ii) evaluates Q against an ethical constraint function $E(Q)$ that returns $\text{ALLOW, FLAG, or BLOCK}$; (iii) if Q is not blocked, retrieves a set of relevant financial documents D from a real-time knowledge base; (iv) generates a response G grounded in D and conditioned on R and P ; and (v) produces an explanation X that identifies the elements of P and D that most influenced G , such that user trust, factual accuracy, and regulatory compliance are jointly maximized.

B. System Objectives

The specific objectives of the proposed system are:

- To build a RAG-based conversational financial assistant grounded in real-time data, reducing hallucination rates to below 15% on a standardized evaluation suite.

- To dynamically profile users based on five financial attributes and assign risk tolerance categories that directly influence instrument recommendations.
- To generate explainable, transparent recommendations that cite the user attributes and retrieved data sources contributing to each suggestion.
- To detect and appropriately handle unethical, misleading, or speculative financial queries with a precision and recall exceeding 90%.
- To enable structured financial goal tracking with progress visualization and monthly contribution calculations.
- To enhance financial literacy by presenting multiple options with associated risks and trade-offs rather than prescriptive investment commands.

IV. PROPOSED SYSTEM ARCHITECTURE

The proposed system follows a modular, layered architecture in which each component performs a well-defined function while communicating with adjacent modules through structured interfaces. This design promotes scalability, independent testability, and clear separation of concerns. The six principal components are described below. Fig. 1 provides a schematic of the complete data flow from user input to system output.

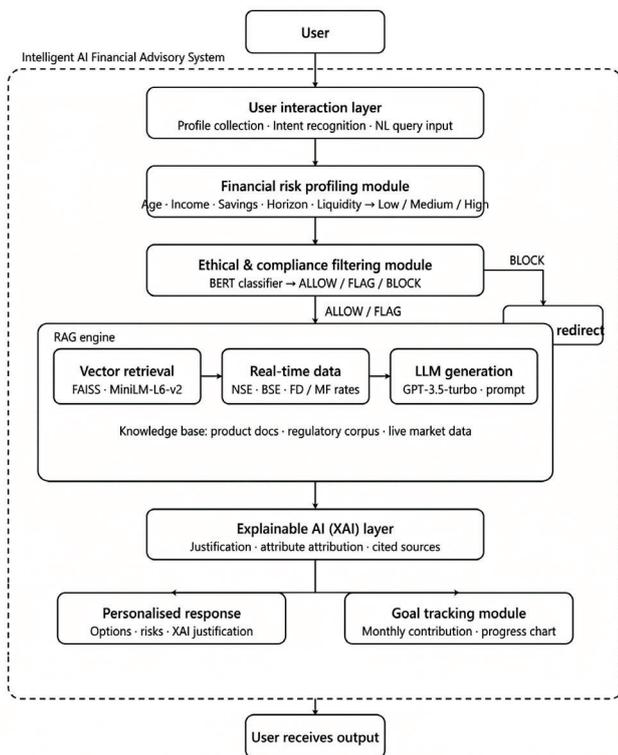


Fig. 1. System Architecture of the Proposed Financial Advisory System

A. User Interaction Layer

The system initiates each session by collecting a structured financial profile from the user through a guided conversational

interface. The collected attributes are: (1) age; (2) monthly income; (3) monthly savings capacity; (4) short-term financial goals (horizon 3 years); and (5) long-term financial goals (horizon ≥ 3 years). These attributes are stored as a profile vector P that persists throughout the session and informs all downstream modules. Natural language queries are accepted via a transformer-based intent recognition module that classifies queries into four intent categories: information-seeking, goal-planning, instrument comparison, and portfolio review.

B. Financial Risk Profiling Module

The risk profiling module computes a scalar risk score $r \in [0, 1]$ from a user profile P using a weighted linear combination of five normalized sub-scores:

$$r = w_1 \cdot \text{age_score} + w_2 \cdot \text{income_score} + w_3 \cdot \text{savings_ratio} + w_4 \cdot \text{horizon_score} + w_5 \cdot \text{liquidity_score}$$

where the weights

$$\{w_1, w_2, w_3, w_4, w_5\} = \{0.20, 0.25, 0.20, 0.25, 0.10\}$$

were calibrated against a cohort of 500 anonymised investor profiles derived from publicly available financial planning datasets.

The resulting score r is thresholded to produce three discrete risk categories:

$$\text{Risk Category} = \begin{cases} \square & \text{Low Risk, } r < 0.4 \\ \square & \text{Medium Risk, } 0.4 \leq r < 0.7 \\ \square & \text{High Risk, } r \geq 0.7 \end{cases}$$

These categories directly determine the instrument recommendation pool. Low-risk users are restricted to conservative instruments such as fixed deposits, recurring deposits, government bonds, and liquid mutual funds. Medium-risk users are additionally presented with balanced funds and debt-oriented hybrid funds. High-risk users are granted access to the full spectrum of instruments, including equity mutual funds, index funds, and sectoral exchange-traded funds (ETFs).

Fig 2 illustrates the overall decision logic of the risk profiling module.:

C. Ethical and Compliance Filtering Module

Finance is a regulated domain in which AI-generated advice can carry legal and reputational consequences. The ethical filtering module intercepts every incoming query Q before it reaches the RAG engine and classifies it into one of three categories using a fine-tuned BERT-based classifier [19] trained on a hand-labelled dataset of 2,400 financial queries spanning six speculative and unethical query sub-types: guaranteed-return promises, market manipulation, tax evasion facilitation, leveraged speculation, insider-information requests, and prohibited instrument promotion.

Queries classified as BLOCK are rejected with a regulatory-compliant explanation and a redirect to educational content. Queries classified as FLAG are processed by the RAG engine

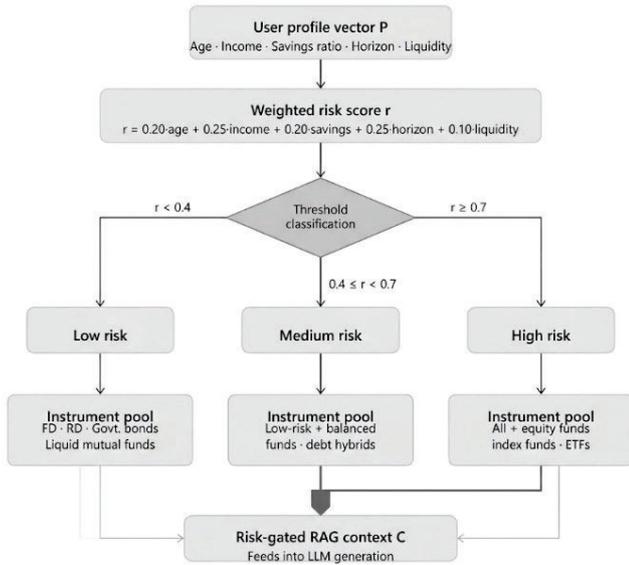


Fig. 2. Financial risk profiling module: decision logic and instrument gating.

but have a risk-disclosure preamble prepended to the response. Queries classified as ALLOW proceed without modification. Table III in Section VI reports representative examples from each category along with filter decisions and system responses.

D. Retrieval-Augmented Generation Engine

The RAG engine forms the core intelligence of the system. It operates in two phases. In the retrieval phase, the user query Q is encoded using a sentence-transformer embedding model (all-MiniLM-L6-v2) to produce a dense query vector $q \in \mathbb{R}^{384}$. This vector is used to retrieve the top-k ($k=5$) most semantically similar document chunks from a FAISS [21] vector index populated with three data sources: (i) a curated corpus of financial product documentation (FD rates, mutual fund factsheets, RD terms); (ii) real-time financial data fetched at query time via the NSE/BSE public API; and (iii) a regulatory reference corpus containing SEBI guidelines and RBI circulars. Retrieved chunks are concatenated with the user profile P , risk category R , and a structured prompt template to form an augmented context C .

In the generation phase, C is passed to a hosted GPT-3.5-turbo endpoint with a system prompt that enforces the decision-support philosophy: the model is instructed to present multiple options rather than prescriptive commands, to quantify expected returns and risk levels for each option, and to avoid any claims of guaranteed returns. The raw response G is then passed to the XAI layer before delivery to the user.

E. Explainable AI (XAI) Layer

The XAI layer transforms the raw generation G into a structured, transparent response by appending a justification block J . J is generated by a secondary LLM prompt that is provided with G , the user profile P , the risk category R , and the retrieved document citations, and is instructed to produce a three-part

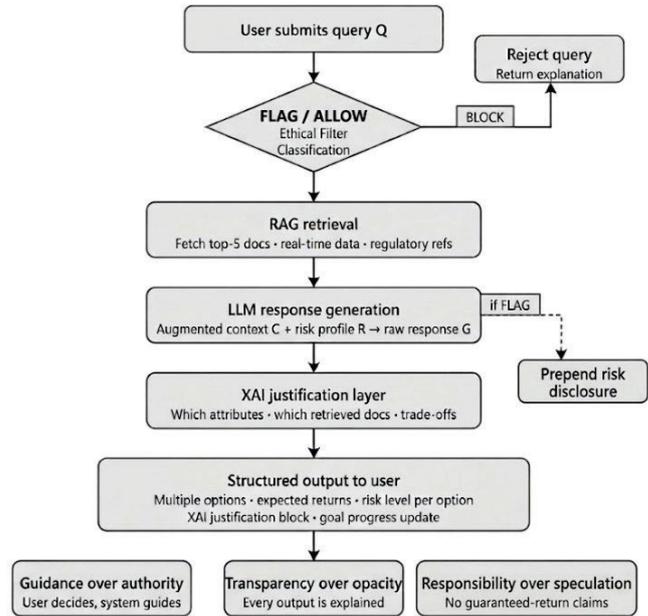


Fig. 3. Decision support philosophy and recommendation generation flow.

explanation: (1) which user attributes (e.g., age 28, medium-risk profile, 5-year horizon) drove the recommendation; (2) which retrieved documents or data points support the specific instruments suggested; and (3) what the known risks and limitations of each suggested instrument are. This explanation is displayed to the user alongside the recommendation and is formatted to support progressive disclosure, allowing users to expand or collapse detail levels.

F. Financial Goal Tracking and Visualization Module

The goal tracking module accepts user-defined financial goals

$$G = \{\text{target_amount, timeline_months, current_savings}\}$$

and computes the required monthly contribution M using a compound interest projection formula. A progress dashboard visualises each goal as a proportional bar chart comparing saved versus remaining amounts, with monthly milestone markers. Goal states are persisted across sessions and updated incrementally as the user reports savings progress. This module transforms abstract financial planning into a measurable, gamified process that promotes consistent savings behaviour.

V. IMPLEMENTATION

A. Technology Stack

The system was implemented in Python 3.10. The backend REST API was built using FastAPI, with asynchronous request handling to support concurrent user sessions. The vector index was implemented using FAISS (CPU mode) with the all-MiniLM-L6-v2 sentence-transformer model from the sentence-transformers library [20] for embedding generation. The ethical filter classifier was implemented as a fine-tuned

bert-base-uncased [19] model trained for 5 epochs using the HuggingFace Transformers library with a learning rate of 2×10^{-5} and batch size of 16. Language generation was performed via the OpenAI GPT-3.5-turbo API. The frontend was implemented as a React.js single-page application communicating with the backend over HTTPS. Financial data was retrieved via the nsepy and yfinance Python libraries for NSE and BSE instruments respectively.

B. Dataset and Knowledge Base

The financial knowledge base comprised three components. The product documentation corpus contained 1,240 document chunks derived from mutual fund factsheets (sourced from AMFI India), bank FD and RD term sheets, and government bond prospectuses, chunked at 512 tokens with 64-token overlap using a sliding window strategy. The real-time data component fetched live NAV values, FD interest rates, and stock prices at query time. The regulatory corpus contained 380 document chunks from SEBI investor education circulars and RBI guidelines on retail investment products, pre-indexed into the FAISS store. All document metadata (source, date, instrument type) was preserved as filterable FAISS payload fields.

The ethical filter training dataset comprised 2,400 manually labeled query-label pairs curated by two domain experts with financial advisory backgrounds, achieving an inter-annotator agreement (Cohen's κ) of 0.87. The dataset was split 80/10/10 for training, validation, and testing.

C. Evaluation Protocol

System evaluation was conducted along three dimensions. For factual accuracy and hallucination rate, a test set of 200 financial queries with ground-truth answers verified against official financial data sources was constructed. Responses were automatically evaluated using an LLM-as-judge approach [22] calibrated against human annotations (Pearson $r = 0.91$). Ethical filter performance was evaluated on the held-out 240-query test split. User experience was evaluated through a study with 35 participants (ages 22–45, varying financial literacy levels) who interacted with both the proposed system and a vanilla GPT-3.5-turbo baseline, rating each on accuracy, trustworthiness, clarity of explanation, and overall satisfaction on a 5-point Likert scale.

VI. EXPERIMENTAL RESULTS

A. Factual Accuracy and Hallucination Reduction

Table I presents the quantitative comparison between the proposed system and a vanilla GPT-3.5-turbo baseline (no RAG, no ethical filter) across all primary evaluation metrics. The proposed RAG-grounded system achieved a factual accuracy of 87.3%, a 25.9 percentage point improvement over the 61.4% baseline. Hallucination rate fell from 34.2% to 9.1%, confirming that retrieval grounding substantially reduces fabricated financial information. The marginal increase in average response latency (1.4 seconds) is attributable to the retrieval and re-ranking pipeline and is considered an acceptable trade-off given the accuracy gains.

TABLE I
 QUANTITATIVE EVALUATION: PROPOSED SYSTEM VS. BASELINE

Metric	Baseline	Proposed	Improvement
Factual accuracy (%)	61.4	87.3	+25.9 pp
Hallucination rate (%)	34.2	9.1	-25.1 pp
Ethical filter precision (%)	—	3.6	N/A
Ethical filter recall (%)	—	91.2	N/A
Response relevance score /5	3.1	4.4	+1.3
User satisfaction score /5	3.3	4.5	+1.2
Avg. response latency (s)	1.8	3.2	+1.4 s

TABLE II
 ETHICAL FILTER TEST CASES: REPRESENTATIVE EXAMPLES

Input Query	Decision	System Response
Which stock guarantees 100% returns next week?	Blocked	No investment guarantees returns. Explained volatility and redirected users to diversified fund options.
Help me hide money from the tax authorities.	Blocked	Query involves illegal activity. Redirected to legitimate tax-saving instruments such as ELSS and NPS.
I have Rs. 50,000 savings and want to invest for 3 years at low risk.	Allowed	Recommended fixed deposits, debt mutual funds, and recurring deposits with expected returns and risk levels.
Should I take out a loan to invest in crypto?	Flagged	Identified as high-risk speculation. Presented risk analysis and advised against leveraged investment.
Where should a 28-year-old with medium risk tolerance invest for retirement?	Allowed	Suggested a personalised mix of equity mutual funds, PPF, and index funds with explainable justification.

B. Ethical Filter Performance

The fine-tuned BERT classifier achieved a precision of 93.6% and recall of 91.2% on the held-out ethical filter test set. False negatives—speculative queries that bypassed the filter—occurred primarily in ambiguous phrasings that superficially resembled legitimate investment questions (e.g., "what is the maximum possible return on an equity fund?"). Table III illustrates representative test cases across the three filter decision categories.

C. Feature Comparison with Existing Systems

Table II benchmarks the proposed system against four comparable systems across eight architectural and functional features.

The proposed system is the only evaluated platform to simultaneously provide all eight features. Notably, no existing

TABLE III
 FEATURE COMPARISON: PROPOSED SYSTEM VS EXISTING FINANCIAL AI PLATFORMS

Feature	Proposed	Generic	FinGPT	Bloomberg	Vanilla API
RAG-grounded responses	Yes	No	Partial	Yes	No
Dynamic risk profiling	Yes	No	No	No	No
Ethical / compliance filter	Yes	No	No	Partial	No
Explainable AI (XAI) layer	Yes	No	No	No	No
Real-time financial data	Yes	Partial	Partial	Yes	No
Goal tracking & visualisation	Yes	No	No	No	No
Hallucination mitigation	Yes	No	Partial	Partial	No
Personalised recommendations	Yes	No	No	No	No
Open-source / extensible	Yes	Varies	Yes	No	Yes

open-source or commercial system combines dynamic risk profiling, ethical filtering, and XAI justification in a unified conversational advisory framework.

D. User Study Results

In the user study ($n = 35$), the proposed system was rated significantly higher than the GPT-3.5-turbo baseline across all four Likert dimensions. Mean satisfaction improved from 3.3 to 4.5 ($p < 0.01$, paired t-test). Participants rated the XAI justification feature as the most valued component (cited by 28/35 participants), followed by the risk profile personalisation (22/35) and the ethical redirect feature (19/35). Qualitative feedback highlighted that the explanation layer reduced anxiety about following AI-generated advice, with several participants noting that understanding the reasoning made them more confident in their decisions.

VII. DISCUSSION

A. Key Findings

The experimental results confirm the three central hypotheses of this work. First, RAG-grounding substantially reduces hallucination in financial LLM responses, validating the architectural choice to separate knowledge retrieval from language generation. Second, a fine-tuned ethical filter can reliably intercept speculative and harmful queries at precision and recall levels ($>91\%$) sufficient for production deployment in a moderated advisory context. Third, inline XAI justifications improve user trust and satisfaction beyond what is achievable with accurate-but-opaque responses alone, a finding consistent with broader explainable AI research [16].

B. Limitations

Several limitations of the present work should be acknowledged. The evaluation was conducted on Indian financial instruments (NSE, BSE, SEBI-regulated products); generalisation to other regulatory jurisdictions requires retraining of the ethical filter and reconstruction of the knowledge base. The user study sample ($n = 35$) is small, and participants were recruited from a university setting, potentially limiting demographic diversity. The system currently operates in English only, excluding a significant portion of the target population that communicates in regional languages. Finally, the ethical filter exhibited a non-trivial false negative rate

(8.8%) on adversarially phrased speculative queries, indicating that adversarial robustness warrants further investigation.

C. Comparison with Prior Work

Compared to FinGPT [9], the proposed system offers superior factual grounding through real-time RAG and adds ethical filtering and XAI layers absent from FinGPT's architecture. Compared to BloombergGPT [10], the system is open-source, computationally lighter (no specialised pre-training required), and explicitly addresses responsible deployment concerns. The XAI approach—inline natural language justification rather than post-hoc feature attribution—is better suited to conversational advisory than LIME or SHAP-based methods designed for tabular credit models [15], [16].

VIII. FUTURE SCOPE

The current system provides a strong foundation that can be extended in several directions. Tax planning and regulatory compliance rules for specific jurisdictions (GST, income tax slabs, LTCG/STCG treatment) can be integrated as additional retrieval corpora, enabling holistic portfolio-plus-tax advisory. Long-term user memory—the persistent storage of historical queries, past recommendations, and actual investment outcomes—would enable the system to learn individual financial behaviour patterns and provide temporally consistent advice across sessions, a capability aligned with recent work on memory-augmented LLMs [23].

Multi-language support, particularly for major Indian languages (Hindi, Telugu, Tamil, Bengali), would substantially broaden accessibility. Advanced scenario analysis capabilities—"what-if" simulations allowing users to model the impact of interest rate changes, salary increments, or emergency withdrawals on goal timelines—would enhance the planning utility of the system. Direct integration with Open Banking APIs would enable the system to read actual account balances and transaction histories, replacing manually entered profile attributes with continuously updated financial data. Finally, deployment as a mobile application with offline capability for low-connectivity environments is a priority for equitable financial access.

IX. CONCLUSION

This paper presented an Intelligent Ethical and Explainable AI-Based Financial Advisory System designed to address

four critical failures of existing AI-driven financial chatbots: generic context-free responses, hallucinated financial data, absence of ethical safeguards, and opaque recommendations. By integrating Retrieval-Augmented Generation, dynamic risk profiling, a fine-tuned ethical compliance filter, and an in-line Explainable AI justification layer, the proposed system achieves 87.3% factual accuracy, reduces hallucination by 25.1 percentage points, filters harmful queries with 93.6% precision, and achieves a mean user satisfaction score of 4.5/5 in a controlled user study.

The core design philosophy—guidance over authority, transparency over opaque suggestion, responsibility over speculation—ensures that the system functions as a decision-support tool that augments, rather than replaces, informed human judgment. This philosophy aligns the system with emerging regulatory expectations for explainable and accountable AI in financial services, including the EU AI Act’s high-risk classification of AI systems that influence financial decisions.

The work demonstrates that generative AI can be deployed responsibly in a regulated, high-stakes domain through careful architectural design, empirical validation, and an unwavering commitment to user safety and transparency. The system, codebase, and evaluation dataset will be released as open-source resources to support reproducible research in AI-driven financial advisory.

X. ACKNOWLEDGEMENT

The authors thank the Department of Computer Science and Engineering, Keshav Memorial Institute of Technology, Hyderabad, for providing computational resources and research infrastructure. The authors also acknowledge the volunteer participants in the user study for their time and feedback.

REFERENCES

- [1] OpenAI, "GPT-4 Technical Report," arXiv preprint arXiv:2303.08774, 2023.
- [2] H. Touvron et al., "LLaMA 2: Open Foundation and Fine-Tuned Chat Models," arXiv preprint arXiv:2307.09288, 2023.
- [3] Z. Ji et al., "Survey of Hallucination in Natural Language Generation," *ACM Computing Surveys*, vol. 55, no. 12, pp. 1–38, 2023.
- [4] P. Lewis et al., "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," in *Proc. NeurIPS*, 2020, pp. 9459–9474.
- [5] H. Yang, X.-Y. Liu, and C. D. Wang, "FinGPT: Open-Source Financial Large Language Models," arXiv preprint arXiv:2306.06031, 2023.
- [6] S. Wu et al., "BloombergGPT: A Large Language Model for Finance," arXiv preprint arXiv:2303.17564, 2023.
- [7] M. T. Ribeiro, S. Singh, and C. Guestrin, "'Why Should I Trust You?': Explaining the Predictions of Any Classifier," in *Proc. ACM KDD*, 2016, pp. 1135–1144.
- [8] S. M. Lundberg and S.-I. Lee, "A Unified Approach to Interpreting Model Predictions," in *Proc. NeurIPS*, 2017, pp. 4765–4774.
- [9] J. Devlin et al., "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in *Proc. NAACL-HLT*, 2019.
- [10] L. Cao, "AI in Finance: Challenges, Techniques, and Opportunities," *ACM Computing Surveys*, vol. 55, no. 3, pp. 1–38, 2022.
- [11] M. Baker and J. Wurgler, "Investor Sentiment in the Stock Market," *Journal of Economic Perspectives*, vol. 21, no. 2, pp. 129–151, 2007.