

Design of an LLM-Powered AI Assistant Chatbot for Nonprofit Trust Management

S. Ian Steve Waugh, R. Priya Professor

Department of Computer Application - PG
VISTAS, Chennai

Abstract

This study presents the design and development of a human-centred AI assistant chatbot powered by Large Language Models (LLMs), tailored for nonprofit trust management systems. Method: Leveraging Lang Chain and OpenAI 's GPT-4 API, the chatbot system integrates with a Fast API backend, React Native frontend, and MongoDB database. It is structured to deliver modularity, real-time interaction, and data-driven responses in a transparent and scalable framework. Human-centred principles were prioritized during design, inspired by Shneiderman 's vision of application focused AI. **Results:** The chatbot demonstrated a 92% query resolution accuracy in test environments, with an average backend response time of under 150 milliseconds. **Feedback from usability testing confirmed ease of navigation and improved donor engagement. Conclusion:** The proposed system confirms the effectiveness of LLM-driven AI assistants in nonprofit platforms and showcases how modern AI frameworks like Lang Chain and OpenAI 's GPT-4 can simplify donation workflows, enhance trust, and improve transparency. **Impact:** This chatbot model offers a blueprint for NGOs and nonprofit organizations to integrate intelligent support systems that reduce manual workload, increase donor trust, and scale communication outreach.

Keywords: AI assistant, ChatGPT, Charity Donation, Conversational AI, Data-Driven Interaction, Fast API, Generative AI, Human-Centred AI, Lang Chain, LLM, MongoDB, Natural Language Processing, Nonprofit Systems, OpenAI API.

I. INTRODUCTION

THE recent evolution of Large Language Models (LLMs), such as OpenAI's GPT-4o mini, has fundamentally transformed how intelligent systems interact with users in natural language. However, simply accessing LLMs through APIs does not guarantee effective communication, contextual consistency, or tailored responses in real-world systems. This is where Lang Chain—a leading LLM orchestration framework—plays a pivotal role.

While LLMs have seen broad adoption in commercial tools, their integration within nonprofit systems remains under explored. Most charity platforms lack intelligent interfaces for real-time, personalized engagement with donors and sponsors, which leads to reduced transparency, lower retention, and missed opportunities for trust building.

To address these limitations, we present a Fast API-powered AI assistant chatbot integrated with Lang Chain and OpenAI's GPT-4o mini API, purpose-built for nonprofit trust management. Lang Chain enables dynamic chaining of LLM logic with external context, using prompt templates and memory modules to deliver precise, context-aware answers. Combined with structured backend services and real-time data from MongoDB, this system allows donors to ask natural questions, receive transparent responses, and interact meaningfully with the charity platform.

Our design adheres to Ben Shneiderman's application-centred AI vision [1], which promotes building AI systems that are useful, explainable, and supportive of human goals—rather than focusing solely on human mimicry. Our chatbot is not designed to replicate human traits or behaviour but rather to serve as a reliable, efficient digital assistant that enhances donor experience, simplifies information access, and strengthens the operational backbone of charity platforms. This work demonstrates how modern tools like Lang Chain can turn raw LLM power into actionable, domain-specific value in social impact environments.

II. LITERATURE SURVEY

A. Human-Centred AI and Design Principles

B. Shneiderman, "Design Lessons From AI's Two Grand Goals: Human Emulation and Useful Applications," highlights two contrasting visions of AI—human emulation versus useful applications. The paper strongly advocates for designing practical AI systems that empower users, which directly aligns with our chatbot's utility driven approach.

B. Shneiderman, "Human-Centred Artificial Intelligence: Reliable, Safe & Trustworthy," presents a framework for AI that prioritizes reliability and transparency. This reinforces the importance of our human-centred design in building trust with nonprofit donors.

G. Marcus and E. Davis, "Rebooting AI," critiques the overpromising of black-box models and stresses the need for transparent, modular systems. Our architecture, using Lang Chain and explainable AI, follows this line of reasoning. [4] B. Friedman and D. G. Hendry, "Value Sensitive Design," argues that human values should be embedded in system design from the beginning. Our application integrates ethical data access, user control, and transparency for this reason.

R. Ramesh et al., "Conversational Agent for Mental Health Support," demonstrates the effectiveness of domain-specific chatbots in sensitive applications. It supports our use of LLMs for trust-based and emotionally aware user interaction.

J. Heer, "Agency plus Automation," discusses how AI tools should be designed to complement human decision-making, not replace it. This perspective aligns with our assistant acting as a supportive guide for donors.

M. R. Endsley, "Lessons from Human-Automation Research," emphasizes the importance of situational awareness and shared control in intelligent systems. Our real-time ETL strategy and transparent prompts support this design. [8] G. Klein et al., "Ten Challenges for Making

Automation a 'Team Player'," identifies collaboration, communication, and mutual predictability as essential for successful human agent teaming. Our architecture reflects these qualities through its structured data flow and interactive feedback design.

B. Related Work

The need for human-centred AI interfaces in mission-driven sectors remains significantly under explored. While commercial platforms like Alexa, Siri, and Google Assistant [6] have demonstrated mainstream success in natural language interfaces, the nonprofit and trust management domains have lagged in adoption.

Our study bridges this gap by building upon foundational work in human-centred AI [1][2], modular architectures [3], and value-sensitive design [5]. Prior research in mental health agents [8] and domain-specific conversational tools [7] validates the feasibility and effectiveness of tailoring LLM-based assistants for targeted user groups. Moreover, the emergence of LLM orchestration frameworks like Lang Chain [11] enables developers to deliver transparent, prompt driven logic integrated with real-time backend infrastructure. This

research uniquely combines these insights to create an ethical, transparent, and scalable assistant tailored to the needs of nonprofit organizations and their donors.

III. METHODS AND PROCEDURES

A. System Architecture

System Architecture Figure 1 illustrates the complete architecture of our AI assistant system, which connects the mobile frontend, backend services, Lang Chain orchestration, and GPT-4o mini. This layered design enables real-time, context-driven interactions with the LLM, making the system responsive and user-friendly.

Each part of the system was developed using modern, open-source technologies to ensure scalability, modularity, and maintainability. The web-based admin panel is developed using React.js, allowing authorized users to manage data such as student profiles, fundraising campaigns, and news articles in real-time. The mobile application, built with React Native, is used by sponsors to access and interact with the system's content seamlessly. Fast API manages backend routing and API exposure, while Lang Chain integrates LLM-based reasoning logic, including prompt templates and memory chains. MongoDB, paired with Beanie ODM, provides a flexible and performant data layer.




Figure 1. System Architecture Flow.

This architectural design strongly aligns with Ben Shneiderman 's vision of application-centred AI [1]. It

places human roles at the core of data management through an intuitive web interface, while also delivering accessible and explainable AI services to end users via mobile. The combination of real-time admin interaction, transparent context curation, and contextual LLM responses reflects a human-empowering AI philosophy rather than black-box automation.

Table 1. Technology Stack Overview.

Component	Technology Used	Purpose
Mobile App	React Native	Donor-side interaction
Admin Panel	React.js	Admin-side content/data management
Backend	Fast API	API development & routing
Database	MongoDB + Beanie ODM	Real-time data handling
LLM	GPT-4o mini + Lang Chain	Context-aware AI responses

B. Implementation

To provide the LLM with domain-specific knowledge, we use a file-based context system. The file app/aiChatContent.txt is an Extract Transform-Load (ETL) export that runs hourly and is generated from real-time user entries in the web admin panel. All key collections such as employee profile, students profile, fund raising, news articles, help a child, sponsors profile, and about us are included in this export. These collections are updated by authorized staff through the admin UI and are continuously refined to support chatbot accuracy. The structured .txt file serves as the primary context input for the chatbot. Lang Chain uses this input along with its Prompt Template and LLM Chain components to dynamically generate responses tailored to the user query.

The ETL-generated .txt file contains a flattened summary of all key collections in the system. It includes structured, readable entries from:

Table 2. Collections Used in the ETL Pipeline.

Collection Name	Purpose / Data Type
employee_profile	Admin/staff details
students_profile	Student records and sponsorship info
fund_raising	Campaign details and goals
news_articles	News, awareness, gallery, and videos
help_a_child	Promotional and emotional engagement content
sponsors_profile	Donor identities and their contributions
about_us	Mission/vision content for AI responses

IV. RESULTS

The system was deployed in a controlled environment for testing with simulated sponsor data.

- Backend API latency: Averaged 130–150 Ms
- Chatbot LLM response accuracy: 92% for static and dynamic queries using GPT-4o mini
- MongoDB data consistency: Verified Realtime sync with mobile UI
- User satisfaction (n=10): Average score of

4.6/5 on chatbot clarity and ease of use

Interaction logs revealed the chatbot could handle complex prompts such as “show me children from the latest campaign” or “how much did I donate last month?” with reliability.

Additionally, over 100 simulated queries were tested across key interaction categories. The majority (35%) focused on student sponsorship, followed by fundraising inquiries (30%), and informational/help queries (20%). The chatbot responded accurately to 92% of the prompts, with fallback messages triggered in only 6% of cases due to ambiguous or malformed input. An hourly ETL process remained consistent across a 7-day trial, ensuring reliable context delivery.

User interviews revealed that the chatbot helped donors feel more connected and informed, especially when reviewing student profiles or tracking donations. The mobile-first experience made access convenient, and the clarity of responses reduced the need for manual clarification from administrators.




Figure 2. Query Distribution by Type .

These visualizations provide a clear breakdown of the chatbot 's tested capabilities and outcomes. Figure 2 illustrates the distribution of user queries across key operational categories, reflecting donor interests and platform utility. Figure 3 complements this by showing the effectiveness of the AI assistant in handling those queries, highlighting high accuracy and minimal fallback rates during evaluation.

V. DISCUSSION

The integration of Lang Chain and GPT-4o mini within a structured charity workflow significantly enhances user engagement, transparency, and trust core values highlighted in Shneiderman 's vision of application-centred AI [1]. By providing Realtime, personalized interactions, the system successfully bridges the gap between technology and human-centred communication, promoting donor empathy and organizational transparency. The modular backend and RESTful interfaces employed by the system simplify deployment and ensure maintainability, essential for nonprofit organizations that typically operate with limited resources. While highly effective, the current implementation does present several limitations, including:

- Lack of multilingual support, potentially limiting global accessibility and donor engagement.

- Dependency on hourly ETL cycles for context updates, introducing minor delays in reflecting real-time data changes.
- Limited scope in handling donor questions beyond the application's current knowledge domain, leading occasionally to fallback responses.
- Absence of offline capabilities, meaning users require active internet connectivity to interact with the chatbot.
- Potential need for more interactive AI features to handle complex donor inquiries

and dynamic user scenarios effectively.




Figure 3. Response Accuracy Breakdown

Nevertheless, the architecture is designed with extensibility in mind, enabling straightforward enhancements to address these limitations in future iterations. From a deployment and operational standpoint, the architecture is notably cost-effective, leveraging open-source tools such as React, Fast API, MongoDB, and Lang Chain, as well as utilizing GPT-4o mini for efficient, low-cost inference. This approach significantly reduces the financial and technical barriers to adopting advanced AI-driven tools, aligning well with nonprofit operational constraints and facilitating scalable, sustainable solutions.

VI. CONCLUSION

This study successfully demonstrated the value and effectiveness of a human-centred, LLM driven chatbot designed specifically for nonprofit trust management systems. The integration of Lang Chain and OpenAI 's GPT-4o mini API provided a transparent, responsive, and contextually rich communication tool, significantly enhancing donor engagement and operational transparency. This aligns strongly with the principles of

application-focused AI proposed by B. Shneiderman[1], underscoring practical utility, human empowerment, and system explainability.

Future work includes:

- Introducing multilingual support, enabling the chatbot to reach and effectively serve a global audience.
- Developing and integrating more advanced, user-centric LLMs to improve personalization and handle a broader range of donor queries effectively.
- Implementing interactive AI features to proactively inform donors of specific benefits and impacts associated with their contributions.
- Expanding open-source availability with additional customizable features, making it easier for NGOs to adopt and adapt the chatbot system to meet evolving organizational needs.
- Providing visual dashboards for administrators, enhancing insights into donor interactions and engagement.

REFERENCES

1. B. Shneiderman, "Design Lessons From AI 's Two Grand Goals: Human Emulation and Useful Applications," IEEE Trans. Technol. Soc., vol. 1, no. 2, pp. 73–82, Jun. 2020.
2. B. Shneiderman, "Human-centred artificial intelligence: Reliable, safe & trustworthy," Int. J. Human Computer. Interact., vol. 36, no. 6, pp. 495– 504, Mar. 2020.
3. G. Marcus and E. Davis, Rebooting AI: Building Artificial Intelligence We Can Trust, Penguin Random House, New York, NY, USA, 2019.
4. B. Friedman and D. G. Hendry, Value Sensitive Design: Shaping Technology With Moral Imagination. Cambridge, MA, USA: MIT Press, 2019.
5. M. McTear, Z. Callejas, and D. Griol, "The Conversational Interface: Talking to Smart Devices," Springer, 2016.
6. R. Ramesh et al., "Design and Evaluation of a Conversational Agent for Mental Health Support," Proc. CHI, 2021.
7. MongoDB Documentation. [Online]. Available: <https://www.mongodb.com/docs/>, 2024.
8. Fast API Documentation. [Online]. Available: <https://fastapi.tiangolo.com>, 2024.
9. Lang Chain Documentation. [Online]. Available: <https://docs.langchain.com>, 2024.
10. [9] Lang Chain Documentation. [Online]. Available: <https://docs.langchain.com>, 2024.
11. [10] Beanie ODM. [Online]. Available: <https://roman-right.github.io/beanie/>, 2024.

13. OpenAI API Documentation. [Online]. Available: <https://platform.openai.com/docs/>, 2024.
14. J. Heer, "Agency plus Automation: Designing artificial intelligence into interactive systems," Proc. Nat. Acad. Sci. USA, vol. 116, no. 6, pp. 1844–1850, 2019.