



DivineLens: A Mythological Chatbot for Everyday Queries

Ms. T. Vijaya Laxmi¹, J. Ruthvika², K. Sri Vignesh³, Shaik Saman⁴

¹Assistant professor, Department of Information Technology, Matrusri Engineering College.

^{2,3,4,5}Students of Department of Information Technology, Matrusri Engineering College

Abstract :

In today's fast-paced world, individuals frequently seek guidance for emotional, ethical, and decision-making challenges. While modern conversational AI systems provide generic responses, they often lack domain-specific wisdom and contextual relevance. This paper presents DivineLens, an AI-powered mythological chatbot designed to provide personalized guidance by leveraging ancient Hindu scriptures, including the Bhagavad Gita, Ramayana, and Mahabharata. The system utilizes a hybrid Retrieval-Augmented Generation (RAG) framework that combines semantic vector search with structured scriptural datasets to retrieve relevant knowledge. DivineLens interprets user-expressed life problems and generates explainable responses consisting of scriptural references, contextual meanings, real-world connections, and actionable guidance. Unlike traditional question-answering systems, the proposed approach emphasizes emotion-aware reasoning, multi-source knowledge integration, and explainability. Experimental case studies demonstrate the system's ability to deliver meaningful and contextually aligned responses, bridging ancient wisdom with modern-day challenges. Furthermore, the system enhances user trust by providing transparent reasoning through explicit verse references and contextual explanations. The integration of multiple scriptures enables diverse perspectives, improving the depth and relevance of guidance. This work highlights the potential of combining AI with cultural knowledge systems to create intelligent, human-centric advisory tools.

Keywords:

Mythological Chatbot, Retrieval-Augmented Generation, Spiritual AI, Explainable AI, Natural Language Processing, Hybrid Retrieval

I. Introduction

In recent years, Artificial Intelligence (AI) has significantly advanced the development of

conversational systems capable of understanding and responding to human queries. However, most existing systems focus on general-purpose knowledge and lack the ability to provide deep, meaningful, and context-aware guidance for personal life challenges.

Ancient Indian scriptures such as the Bhagavad Gita, Ramayana, and Mahabharata contain profound philosophical insights addressing emotional, ethical, and existential dilemmas. Despite their relevance, accessing and interpreting this knowledge remains challenging for modern users due to language complexity and contextual depth.

To address this gap, this paper introduces DivineLens, a mythological chatbot designed to map everyday human problems to scriptural wisdom. Unlike conventional chatbots, DivineLens functions as a guidance system rather than a simple question-answering tool, offering empathetic, explainable, and actionable responses.

The key contributions of this work include:

- A hybrid RAG-based framework integrating multiple Hindu scriptures.
- A life-problem-to-scripture mapping mechanism for contextual understanding.
- An explainable response generation structure including meaning, connection, and guidance.
- A multi-source retrieval system combining semantic and contextual relevance.

II. Related Work :

Recent advancements in Natural Language Processing (NLP) and Large Language Models (LLMs) have significantly transformed the development of intelligent conversational systems. Transformer-based architectures such as BERT and GPT have enabled machines to understand contextual semantics, generate coherent responses, and perform complex reasoning tasks. These

capabilities have led to the emergence of domain-specific question-answering (QA) systems that aim to retrieve and generate knowledge from structured and unstructured textual sources.

One prominent direction of research focuses on applying deep learning models to extract knowledge from ancient and context-rich texts. Pandey et al. developed a question-answering system over ancient scriptures using a hybrid architecture combining BERT embeddings, attention mechanisms, and recurrent neural networks [1]. Their system demonstrated high accuracy in capturing semantic relationships within complex textual data, highlighting the effectiveness of transformer-based models in handling traditional knowledge sources. However, their work primarily focuses on extractive QA and does not extend to interpretative or advisory responses.

Similarly, Vivechan AI presents a system that leverages Large Language Models in combination with FAISS-based vector search to retrieve relevant passages from multiple Indian scriptures [2]. By indexing large textual corpora and performing similarity-based retrieval, the system enables efficient access to scriptural knowledge. While this approach improves retrieval performance and supports multi-source querying, it mainly emphasizes information extraction rather than providing structured, user-centric guidance.

The integration of Retrieval-Augmented Generation (RAG) has further enhanced the reliability and factual grounding of AI systems. MufassirQAS employs a RAG-based architecture that combines vector databases with LLMs to minimize hallucinations and ensure that responses are grounded in authoritative religious texts [3]. This approach improves transparency by linking generated responses with source references. Additionally, NirvanaNavigator extends the RAG paradigm by incorporating hybrid retrieval mechanisms and conversational memory, allowing for personalized and context-aware spiritual guidance [4]. These systems demonstrate the potential of combining retrieval and generation but still lack a structured framework for mapping real-life user problems to domain-specific wisdom.

Another area of research focuses on improving question answering over structured and complex documents. PDFTriage introduces a method for handling long and structured documents by

leveraging document metadata such as sections, pages, and tables to enhance retrieval accuracy [5]. This work highlights the importance of preserving document structure during retrieval. Furthermore, transformer-based models have been applied to classify and organize mythological texts, as demonstrated by Paul et al., who developed a topic classification framework for Indian mythology using BERT-based architectures [6]. Such classification techniques facilitate better organization and retrieval of large-scale textual data but do not directly address response generation or user interaction.

Despite these advancements, several limitations remain in existing systems. Most approaches focus on fact-based question answering and lack the ability to provide meaningful, context-aware guidance tailored to individual user needs. Additionally, many systems rely on a single source of knowledge, which restricts the diversity and richness of retrieved information. Another critical limitation is the lack of explainability, as many models generate responses without clearly indicating the reasoning or source behind them. Furthermore, existing systems do not effectively bridge the gap between modern human problems—such as emotional stress, decision-making, and ethical dilemmas—and ancient philosophical teachings.

To overcome these challenges, DivineLens introduces a novel framework that integrates multi-scripture retrieval, emotional understanding, and explainable response generation. Unlike traditional QA systems, the proposed approach interprets user queries as life problems and maps them to relevant scriptural insights across multiple sources. By combining hybrid retrieval techniques with a structured response format that includes interpretation, contextual connection, and actionable guidance, DivineLens provides a more human-centric and meaningful interaction model. This positions the system as not just an information retrieval tool, but an intelligent advisory system that bridges the gap between ancient wisdom and contemporary life challenges.

III. System Overview and Architecture

DivineLens is an AI-powered mythological chatbot designed to provide personalized guidance for everyday life challenges by leveraging ancient Hindu scriptures. The system integrates modern Natural Language Processing (NLP) techniques with



traditional knowledge sources to map user-expressed problems to relevant scriptural insights. The overall framework is based on a hybrid Retrieval-Augmented Generation (RAG) pipeline enhanced with query expansion, emotion-aware processing, and cross-encoder re-ranking.

When a user submits a query describing a real-life problem, the system processes the input through multiple stages, including semantic understanding, multi-source retrieval, and structured response generation. Unlike conventional chatbots, DivineLens ensures that responses are grounded in authentic scriptural knowledge while maintaining contextual relevance and explainability.

The architecture of DivineLens follows a modular pipeline consisting of interconnected components that handle input processing, retrieval, and response generation. The system workflow is illustrated as a sequence of layered operations:

Input Processing and Emotion Analysis :

The user query is first processed through a text normalization step and passed to an emotion classification model. The system uses a transformer-based emotion detection model to identify the emotional state of the user (e.g., stress, anger, confusion). This emotional context is later used to guide the tone and structure of the generated response.

Query Expansion Module:

To improve retrieval performance, the system applies a query expansion technique using a predefined dictionary of related terms. This step enriches the original query with semantically relevant keywords, enabling better matching with scriptural content.

Embedding and Vector Retrieval:

The expanded query is converted into a vector representation and passed to a vector database (ChromaDB). The system performs dense retrieval by identifying passages that are semantically similar to the query. This allows the model to retrieve contextually relevant information even when exact keywords are not present.

Hybrid Retrieval Mechanism:

In addition to semantic retrieval, the system employs BM25-based keyword retrieval using the HybridRetriever module. This step captures lexical similarities between the query and documents. The results from dense retrieval and sparse retrieval are combined to form a hybrid set of candidate passages.

Cross-Encoder Re-ranking:

The combined candidate documents are further refined using a cross-encoder model. The re-ranking process evaluates each query-document pair and assigns relevance scores, ensuring that the most contextually appropriate passages are selected. This step significantly improves the accuracy and quality of retrieved information.

Context Aggregation:

Top-ranked passages from multiple scriptures (Bhagavad Gita, Ramayana, and Mahabharata) are aggregated to form a unified context. This multi-source aggregation provides diverse perspectives and enriches the response generation process.

Prompt Builder and RAG Pipeline:

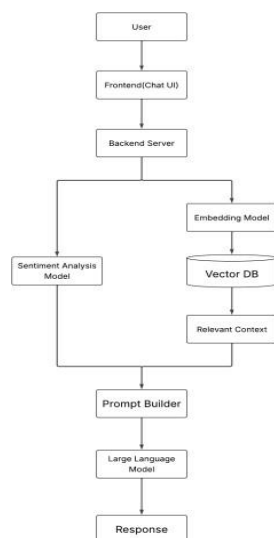
The aggregated context, along with the original user query and detected emotional state, is passed to a prompt builder. The prompt is structured to guide the Large Language Model (LLM) in generating meaningful and grounded responses. This stage represents the core of the Retrieval-Augmented Generation pipeline.

Response Generation and Structuring:

The LLM generates the final response based on the constructed prompt. The output is structured into multiple components:

- Empathetic acknowledgment of the user's situation
- Scriptural references (source, chapter, verse)
- Meaning and interpretation of retrieved passages
- Contextual connection to the user's problem
- Actionable guidance or advice

This structured format enhances interpretability, user trust, and practical usefulness.



IV. Methodology

The methodology of DivineLens focuses on transforming user-expressed life problems into meaningful and contextually grounded guidance using a multi-stage Retrieval-Augmented Generation (RAG) pipeline. Unlike traditional question-answering systems, the proposed approach integrates query expansion, hybrid retrieval, cross-encoder re-ranking, and emotion-aware response generation to enhance both relevance and interpretability.

4.1 Query Understanding and Emotion Detection

The system processes user input as a description of a real-life situation rather than a factual query. A transformer-based emotion classification model is used to identify the emotional state of the user, such as stress, anger, or confusion. This emotional context helps guide the tone and style of the generated response, making it more empathetic and user-centric.

4.2 Query Expansion

To improve retrieval performance, the system expands the original query using a predefined dictionary of related terms. This step enriches the query with additional context and ensures that relevant scriptural passages can be retrieved even when the exact words do not match the stored data.

4.3 Dense Retrieval using Vector Embeddings

The expanded query is converted into a vector representation and used to search within a vector

database. This process retrieves passages that are semantically similar to the user's query, enabling the system to identify relevant content even when there is no direct keyword overlap.

4.4 Sparse Retrieval using BM25

In parallel with semantic retrieval, the system applies a keyword-based retrieval method using BM25. This approach focuses on matching important words in the query with those in the documents, ensuring that lexically relevant passages are also considered.

4.5 Hybrid Retrieval and Fusion

The results from both semantic (dense) retrieval and keyword-based (sparse) retrieval are combined to form a unified set of candidate passages. This hybrid approach improves the overall retrieval quality by capturing both contextual meaning and exact keyword matches.

4.6 Cross-Encoder Re-ranking

The combined candidate passages are further refined using a cross-encoder model. This model evaluates the relevance of each passage with respect to the query and assigns a score. The passages are then ranked, and the most relevant ones are selected. This step improves precision by considering deeper contextual relationships.

4.7 Multi-Source Context Aggregation

The top-ranked passages are collected from multiple scriptures, including the Bhagavad Gita, Ramayana, and Mahabharata. This aggregation ensures that the system provides diverse perspectives and enhances the richness of the generated response.

4.8 Retrieval-Augmented Generation

The aggregated context, along with the original user query and detected emotional state, is used to construct a prompt for the language model. The model generates responses based on both the query and the retrieved scriptural content, ensuring that the output is grounded in authentic knowledge.

4.9 Explainable Response Structuring

The generated response is structured into multiple components to improve clarity and interpretability:



Meaning: Explanation of the retrieved scriptural content

Connection: How the content relates to the user's problem

Guidance: Practical advice derived from the interpretation

This structured approach enhances user trust and ensures that the reasoning behind the response is clear.

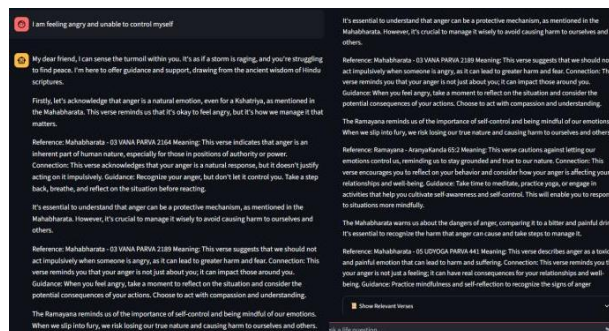
V. Results and Discussion

The performance of DivineLens was evaluated using sample user queries related to real-life problems such as career confusion and anger management. The system was tested to determine whether it could provide meaningful, relevant, and explainable guidance using scriptural knowledge.

For a query related to career confusion, the system retrieved relevant verses from the Bhagavad Gita and Mahabharata. It explained the idea of duty and purpose, connected it to the user's situation, and suggested focusing on consistent effort and clarity. Similarly, for a query related to anger, the system retrieved passages from the Ramayana and Mahabharata, explained the effects of anger, and provided practical guidance such as reflection and self-control.

The results show that the system effectively connects modern life problems with scriptural teachings. The responses are structured, easy to understand, and include meaningful explanations. The use of multiple scriptures improves the depth of guidance, while emotion detection helps generate more empathetic responses.

Overall, DivineLens demonstrates its ability to provide useful and context-aware guidance by combining AI techniques with traditional knowledge. However, the system depends on the available dataset, and future improvements can focus on expanding the knowledge base and introducing quantitative evaluation methods.



VI. Conclusion and Future Scope

This paper presented DivineLens, a mythological chatbot designed to provide guidance for everyday life problems by leveraging ancient Hindu scriptures. The system integrates modern AI techniques such as Retrieval-Augmented Generation, hybrid retrieval, and emotion-aware processing to generate meaningful and context-aware responses. By mapping user queries to relevant teachings from the Bhagavad Gita, Ramayana, and Mahabharata, DivineLens functions as a guidance-oriented system rather than a traditional question-answering chatbot.

The results demonstrate that the system can generate structured, explainable, and relevant responses that connect ancient wisdom with modern challenges. The combination of multi-source retrieval, re-ranking, and emotion-aware processing enhances both the quality and personalization of responses, making the system more engaging and user-centric.

In the future, DivineLens can be extended to support multiple languages, allowing users from diverse linguistic backgrounds to interact with the system more effectively. The knowledge base can also be expanded by incorporating scriptures from other religions, enabling a more inclusive and comparative guidance system. Additionally, the system can be enhanced by introducing character-specific responses, where guidance is provided from the perspective of specific figures such as Krishna, Rama, or other key characters, making the interaction more immersive and personalized. These improvements can further strengthen DivineLens as a comprehensive and culturally inclusive AI-based guidance platform.





VII. References

- [1] A. K. Pandey et al., “Extractive Question Answering over Ancient Scriptures Texts using Generative AI and NLP Techniques,” IEEE Access, 2024.
- [2] O. Soni et al., “Vivechan AI: Extracting Wisdom from Ancient Indian Texts Through LLM,” in Proceedings of the International Conference on Cognitive & Cloud Computing (IC3Com), 2024.
- [3] A. Y. Alan et al., “Improving LLM Reliability with RAG in Religious Question-Answering: MufassirQAS,” Turkish Journal of Engineering, vol. 9, no. 3, pp. 544–559, 2025.
- [4] Vijay R. et al., “NirvanaNavigator: An AI-Powered Chatbot for Spiritual Insights,” International Research Journal of Modernization in Engineering Technology and Science (IRJMETS), vol. 7, no. 8, 2025.
- [5] J. Saad-Falcon et al., “PDFTriage: Question Answering over Long, Structured Documents,” in Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP), 2024.
- [6] A. Paul et al., “Transformer-based Pauranic Topic Classification in Indian Mythology,” Sādhanā, vol. 49, 2024.
- [7] P. Lewis et al., “Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks,” in Advances in Neural Information Processing Systems (NeurIPS), 2020.
- [8] J. Devlin et al., “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” in Proceedings of NAACL-HLT, 2019.