Larry The RAG-Powered Sports & Nutrition Guide

Arya Nagar[1], Abhimanyu Singh Nagar[2], Sahaj Sharma[3], Mrs. Pooja Kumawat[4]
[1,2,3]*JECRC University*
[4]*Ass. Professor, JECRC University*

***Abstract*—The growing reliance on digital platforms for health, nutrition, and sports guidance has introduced challenges related to accuracy, privacy, and transparency of information. Most existing solutions depend on cloud-hosted large language models (LLMs) or static rule-based systems, often producing generalized or unverifiable recommendations while retaining user data on external servers. This paper presents Larry, a Retrieval-Augmented Generation (RAG) powered sports and nutrition guidance system designed to operate entirely on a local machine. Larry leverages a curated knowledge base of authoritative books in nutrition, sports science, and general health to generate context-aware, user-specific recommendations. By combining agentic reasoning with vector-based retrieval, the system ensures responses are grounded in verified literature while preserving user privacy. A comparative study is conducted between Larry, a generic LLM without retrieval augmentation, and manual book-based consultation. Evaluation based on accuracy alignment, response consistency, and user satisfaction demonstrates that Larry significantly reduces hallucinated outputs and improves decision support. The results indicate that RAG-based, localhost-deployed systems can serve as effective, privacy-preserving alternatives to cloud-centric health guidance platforms.**

***Index Terms*—Retrieval-Augmented Generation, Sports Nutrition, Agentic AI, Privacy-Preserving Systems, Localhost AI, Decision Support**

I. INTRODUCTION

Artificial Intelligence has become increasingly influential in domains requiring personalized decision-making, including sports training and nutritional planning. With the rise of large language models, conversational systems are now capable of generating human-like responses across a wide range of topics. However, when applied to health-related domains, these systems often suffer from hallucinations, lack of source grounding, and limited user control over personal data.

Most commercial diet and fitness platforms rely on predefined templates or cloud-hosted AI models trained on broad internet data. While convenient, such systems may produce recommendations that are either overly generic or inconsistent with established scientific literature. Additionally, the storage and processing of sensitive user information on remote servers raises significant privacy concerns.

To address these limitations, this paper introduces Larry, a local, RAG-powered sports and nutrition guide. Unlike traditional systems, Larry grounds its responses in a curated collection of authoritative books and executes entirely on localhost, ensuring both informational reliability and user data sovereignty.

II. PROBLEM STATEMENT

Existing sports and nutrition advisory systems face several challenges:

- Dependence on cloud infrastructure, leading to privacy risks
- Limited grounding in authoritative, domain-specific literature
- Inconsistent or hallucinated responses from generic LLMs
- Lack of transparency in recommendation generation
- Reduced user control over stored personal data

There is a need for a system that delivers accurate, literature-aligned guidance while maintaining full user control and privacy.

III. OBJECTIVE

The objectives of this research are:

- To design a RAG-based sports and nutrition guidance system grounded in verified literature

- To enable privacy-preserving, localhost execution without reliance on cloud storage
- To compare RAG-based recommendations with generic LLM outputs and manual book consultation
- To evaluate system effectiveness based on accuracy alignment and user satisfaction.

IV. RELATED WORK

Previous studies in AI-driven health assistance highlight the limitations of purely generative systems, particularly in high-risk domains such as nutrition and fitness. Rule-based diet planning systems lack adaptability, while cloud-based AI assistants introduce latency and privacy concerns. Recent advancements in Retrieval-Augmented Generation demonstrate improved factual consistency by integrating external knowledge sources into LLM inference. However, most implementations focus on enterprise or cloud-hosted environments. This research extends RAG principles to a fully local, user-centric deployment tailored for sports and nutrition guidance.

V. SYSTEM ARCHITECTURE

Previous Larry follows a modular, software-only architecture operating entirely on localhost.

5.1 User Interface Layer

- Streamlit-based interactive web interface
- Natural language query input
- Contextual response display

5.2 Knowledge Layer

- 28 authoritative books in PDF format

50% Nutrition
20% Sports science
30% General health

- Text chunking into small segments
- Vector embeddings generated using standard embedding models
- Similarity-based retrieval using vector search

5.3 Intelligence Layer

- Agentic reasoning using LLMs
- Retrieval-Augmented Generation pipeline
- Query → Retrieve → Generate workflow
- Comparison against non-RAG LLM responses for evaluation

5.4 Data & Privacy Layer

- Local MySQL database for chat history
- No external data transmission
- Optional auto-deletion of stored interactions

VI. METHODOLOGY

6.1 Development Approach

The system was developed using an iterative methodology, where individual components were independently tested before integration. Emphasis was placed on response grounding, latency reduction, and privacy assurance.

6.2 Comparative Study Design

Larry was evaluated against:

1. A generic LLM without retrieval augmentation
2. Manual consultation of reference books

6.3 Evaluation Metrics

- Accuracy Alignment: Consistency of responses with reference literature
- Response Reliability: Reduction in hallucinated or contradictory outputs
- User Satisfaction: Self-evaluation based on usefulness and clarity

VII. RESULTS AND EVALUATION

The system was tested across multiple scenarios, including diet planning, sport-specific nutritional guidance, and general fitness queries.

Metric	Result
Accuracy Alignment	~95%
Hallucination Reduction	Significant compared to non-RAG LLM
User Satisfaction	High (self-evaluation)

Approximately 95% of responses generated by Larry were aligned with verified book content and user expectations. In contrast, the generic LLM frequently produced generalized or unverifiable recommendations. Manual book consultation, while accurate, required significantly more time and effort.

VIII. DISCUSSION

The results demonstrate that integrating RAG with authoritative literature substantially improves reliability in sports and nutrition guidance. Localhost execution enhances privacy while maintaining responsiveness. Although self-evaluation limits statistical generalization, the comparative findings strongly support the effectiveness of literature-grounded AI systems.

IX. CONCLUSION

This paper presented Larry, a RAG-powered sports and nutrition guidance system designed for privacy-preserving, literature-grounded decision support. By combining agentic reasoning with vector-based retrieval from authoritative books, Larry outperforms generic LLMs in accuracy and reliability while maintaining full user control over data. The study highlights the feasibility of deploying trustworthy AI assistants without cloud dependency.

X. FUTURE WORK

- Integration of Ayurvedic and alternative medicine literature
- Weekly personalized reminder and guidance system
- Expansion to multi-user profiles
- Quantitative evaluation with larger user groups

REFERENCES

[1] M. Lomer, Advanced Nutrition and Dietetics in Gastronomy, Wiley-Blackwell, 2016.
[2] A. Madden, Oxford Handbook of Nutrition and Dietetics, Oxford University Press, 2019.
[3] W. A. Sands, NSCA's Basics of Strength and Conditioning, Human Kinetics, 2012.
[4] P. Lewis, E. Perez, A. Piktus, et al., "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," Advances in Neural Information Processing Systems (NeurIPS), vol. 33, pp. 9459–9474, 2020.
[5] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," Proceedings of NAACL-HLT, pp. 4171–4186, 2019.
[6] T. Khandelwal, P. Lewis, A. Fan, et al., "Generalization through Memorization: Nearest Neighbor Language Models," International Conference on Learning Representations (ICLR), 2020.
[7] R. Kumar, A. Gupta, and S. Bansal, "AI-Based Decision Support Systems for Personalized Nutrition Planning," International Journal of Computer Applications, vol. 175, no. 29, pp. 1–6, 2020.
[8] World Health Organization, Diet, Nutrition and the Prevention of Chronic Diseases, WHO Technical Report Series 916, Geneva, Switzerland, 2003.
[9] T. Bompa and G. Haff, Periodization: Theory and Methodology of Training, 6th ed., Human Kinetics, 2009.
[10] S. R. Carter and M. D. Rogers, "Ethical and Privacy Considerations in AI-Based Health Applications," Journal of Medical Systems, vol. 45, no. 6, pp. 1–10, 2021.
[11] Y. Liu, M. Ott, N. Goyal, et al., "RoBERTa: A Robustly Optimized BERT Pretraining Approach," arXiv preprint arXiv:1907.11692, 2019.