

# Lessons from A Large Language Model-based Outdoor Trail Recommendation Chatbot with Retrieval Augmented Generation

Julia Ann Mathew and Suining He  
{julia.ann\_mathew, suining.he}@uconn.edu  
Ubiquitous & Urban Computing Lab  
School of Computing  
University of Connecticut

## Abstract

The increasing popularity of outdoor recreational activities (such as hiking and biking) has boosted the demand for a conversational AI system to provide informative and personalized suggestion on outdoor trails. Challenges arise in response to (1) how to provide accurate outdoor trail information via conversational AI; and (2) how to enable usable and efficient recommendation services. To address above, this paper discusses the preliminary and practical lessons learned from developing Judy, an outdoor trail recommendation chatbot based on the large language model (LLM) with retrieval augmented generation (RAG). To gain concrete system insights, we have performed case studies with the outdoor trails in Connecticut (CT), US. We have conducted web-based data collection, outdoor trail data management, and LLM model performance studies on the RAG-based recommendation. Our experimental results have demonstrated the accuracy, effectiveness, and usability of Judy in recommending outdoor trails based on the LLM with RAG.

## CCS Concepts

- Computing methodologies → Artificial intelligence.

## Keywords

Outdoor Trail Recommendation, Large Language Model, Retrieval Augmented Generation

## 1 Introduction

The rising demand for the outdoor recreational activities has created the need for the outdoor trail recommendation [9]. Such recommendation aims to consolidate scattered knowledge such as trail location, difficulty, length, and permitted activities (e.g., biking), and provide accurate recommendation about trails for safe and pleasant activity planning.

To meet this need, this prototype study discusses the practical lessons and insights from the development of Judy, a outdoor trail recommendation chatbot based on conversational AI. Toward development and deployment of Judy, we have investigated the following two questions. *First*, how can we provide accurate outdoor trail knowledge in response to complex user query? Existing rule-based chatbots for recommendation largely follow predefined workflows. While they can handle specific and structured recommendation tasks, these chatbots could not handle complex or customized queries that deviate from their predefined rules, limiting their recommendation capability. General-purpose chatbots based

on conversational AI, on the other hand, may lack the *domain-specific knowledge* needed for effective outdoor trail recommendation. *Second*, how can we enable usable and efficient recommendation services? Existing outdoor trail recommendation platforms such as AllTrails [11] offer static information. However, these platforms largely lack the *conversational interactivity*. Therefore, the users often need to shift through multiple trail reviews to plan their trips.

Toward addressing the above questions, we have performed real-world outdoor trail data studies (over 260 trails in Connecticut), and developed the system of Judy as the large language model (LLM) recommendation chatbot with retrieval augmented generation (RAG) [10]. In particular, we have implemented RAG to retrieve the most relevant trail information from a constructed trail database, such that Judy can provides accurate outdoor trail information. We have designed practical performance enhance mechanisms (such as caching) to enhance retrieval efficiency, and enable usable recommendation services. We have performed real-world data analysis based on outdoor trail data in Connecticut (CT), US, and conducted extensive experimental studies to validate the accuracy, efficiency, and usability of Judy in recommending the outdoor trails.

The rest of the paper is organized as follows. We first overview the system designs in Sec. 2. Then we present the design of Judy with the LLM chatbot and RAG design in Sec. 3. We present the experimental studies and lessons learned in Sec. 4, and conclude the paper in Sec. 5.

## 2 System Overview

The designs of Judy, as showcased in Fig. 1, consist of the following two phases: (1) *Data Preparation & Preprocessing*: In this phase, we pre-process and prepare the basic outdoor trail information (including the trail name, length, and difficult), and establish the MySQL database hosted on the Amazon RDS to streamline data access and the subsequent retrieval augmented generation. (2) *User Query & Recommendation with RAG*: Given the user query of the outdoor trail, the LLM interprets the user query, and formulates a corresponding SQL query that fetches the relevant data from the structured database. The SQL query result, along with the user original question, are fed to the LLM. The LLM leverages the above contexts to generate a structured and user-friendly response regarding the outdoor trail recommendation.

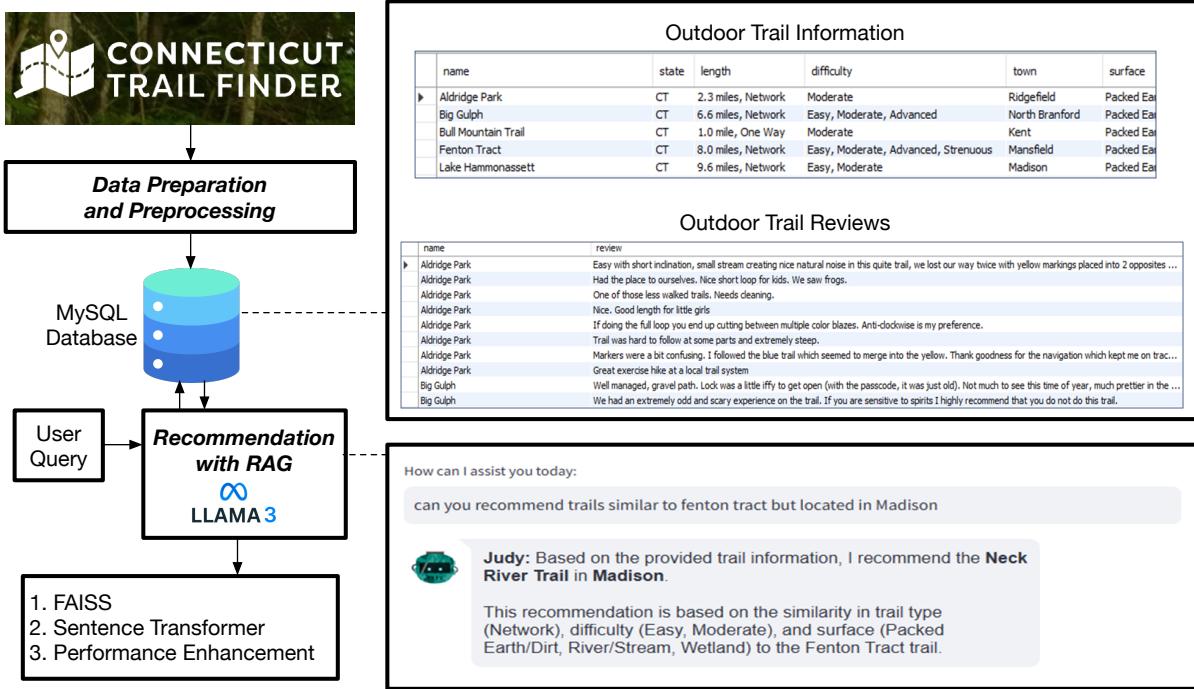


Figure 1: Overview of Judy and samples of our outdoor trail information and collected reviews.

### 3 Core Designs of Judy

This section presents the core designs of Judy. We first present the data preparation and preprocessing in Sec. 3.1, followed by the retrieval augmented generation designs in Sec. 3.2.

#### 3.1 Data Preparation and Preprocessing

In developing Judy, we have crawled the basic information of outdoor trails a credible platform called CT Trail Finder, such as length, difficulty, and pet permissions. To perform further natural language analysis, we have collected reviews for these trails from the Google Reviews and TrailLink [12]. To prepare our analysis, we have performed data pre-processing. We have filtered away irrelevant trail reviews, removed emojis, excessive line spaces, asterisks, and other unnecessary characters. We have performed standardized abbreviations (such as replacing “I’ve” with “I have”) to ensure the review consistency. We have re-structured the reviews such that formatting is consistent across all reviews.

In our Judy development, we have adopted the Selenium and BeautifulSoup for web scraping to gather trail-related data such as reviews, and trail features. To prepare the database, we have used MySQL, which is hosted on Amazon RDS, to form a relational database to store structured data, such as trail information and reviews.

#### 3.2 User Query and Recommendation with RAG

The user queries in our current system designs typically pertain to features of the outdoor trail, such as the trail’s length, location, difficulty level, permitted activities, trail accessibility, and other attributes.

The LLM will determine whether the user’s query is recommendation-relevant. If yes, Judy will leverage the LLM to generate a personalized response. If the query is not recommendation-relevant, Judy assesses whether it can be answered using the existing table schema. If the LLM in Judy the user query can be resolved using the pre-defined schema, Judy will perform the SQL query, and retrieve the necessary information regarding the outdoor trail. If the user query requires more nuanced insights, such as opinions or user experiences, Judy will invoke the RAG function to process the reviews of the relevant trail to construct an accurate and contextually rich response. Examples of such questions include “*what do people say about the scenery on Aldridge trail*”, “*did visitors encounter wildlife at Big Gulph trail*”, “*how crowded is the Pine Hill trail usually*”, and “*what do reviews say about biking on the Windsor Locks Canal trail*”.

In terms of the RAG function, Judy will retrieve the relevant reviews and their corresponding embeddings. We have evaluated three kinds of sentence embeddings – that is, Ollama (omic-embed-text) [4], and two kinds of pre-trained Sentence Transformers [2]. One Sentence Transformer is trained question-answer (QA) pairs (multi-qa-mpnet-base-cos-v1) [8], and maps sentences and paragraphs to a 512-dimensional dense vector space. Another Sentence Transformer is the distilled version of multilingual Universal Sentence Encoder (distiluse-base-multilingual-cased-v2) [1, 6]. This transformer maps sentences and paragraphs to a 768-dimensional dense vector space and was designed for semantic search.

In our experimental studies, we have observed that both the Sentence Transformers exhibited significantly faster response times compared to the Ollama embedding. Specifically, we have structured and created a total of 25 user queries (the same set of queries

in Sec. 3) with the ground-truth answers for our testing evaluation. The Sentence Transformer pre-trained on QA pairs (`multi-qa-mpnet-base-cos-v1`) has delivered the fastest average response time at 0.73s, followed by 2.89s by the distilled multilingual Universal Sentence Encoder (`distiluse-base-multilingual-cased-v2`), and 4.69s by Ollama embeddings (`nomic-embed-text`).

In our Judy, we adopt the Facebook AI Similarity Search (FAISS) [3] to rank the relevance of the reviews to the user’s query. We find the top  $k$  relevant reviews based on their similarity scores (cosine similarity based on a dot product on normalized embeddings). These selected relevant reviews and the user’s query are fed to the LLM, and Judy synthesizes them and generates a comprehensive and contextually appropriate responses.

In our Judy development, we have used Llama3 [13] for Judy’s conversational abilities, natural language understanding, and response generation. We have adopted LangChain to form the conversational flow and integrates Llama3 with MySQL for enhanced conversational functionality. We have adopted FAISS for similarity searches and nearest neighbor queries on trail reviews dataset for quick information retrieval.

In our experimental studies, we have observed the significantly long response time of using RAG from our outdoor trail database. To further enhance the efficiency of the RAG-based approach, we have implemented caching for the embeddings and reviews of the queried trails.

## 4 Experiment and Preliminary Insights

We find the measure of outdoor trail recommendation matching based on the percentage of the correct answers against all the questions. We measure the recommendation matching via the semantic similarities across the generated responses and ground-truth answers using the pre-trained sentence transformer (`multi-qa-MiniLM-L6-cos-v1`) [7]. We used and tested a total of 25 user queries (the same set of queries in Sec. 3) with the ground-truth answers for our testing evaluation. We also compare Judy with the same LLM backbone based on SQL only (without RAG).

*Recommendation Matching:* In our experiment, we have observed that Judy achieved a recommendation matching of 96%, which outperforms the 88% accuracy of Judy without RAG. Judy with RAG retrieves the top 5 most relevant reviews to answer a user’s question. Judy therefore provides concise and accurate responses. In contrast, Judy without RAG needs to process all the review related to a specific outdoor trail through SQL, which can overwhelm the LLM and result in reduced response performance.

*RAG Designs:* To determine the optimal value for  $k$  (the number of the outdoor trail reviews sent to the LLM backbone), we have evaluated impacts of different  $k$ ’s on efficiency and recommendation matching. In Fig. 2(a), we have shown the performance of Judy with the embeddings using the sentence transformer pre-trained on QA pairs (`multi-qa-mpnet-base-cos-v1`). We can see that  $k = 5$  yields an average response time of 0.73s and recommendation matching of 96%. We have also observed that when  $k = 10$ , the response time increases to 1.17 seconds, but the accuracy remains the same at 96%. The increase in response time is due to processing of more trail reviews, which adds computational overhead without improving accuracy.

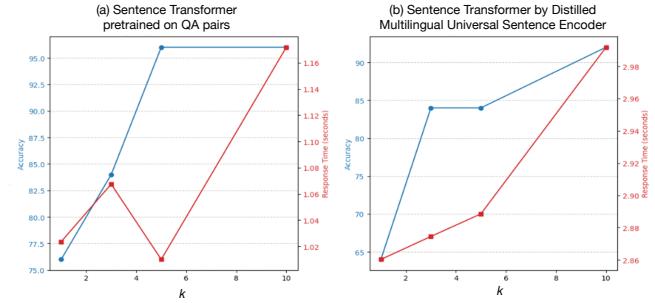


Figure 2: Sensitivity studies over  $k$  for Judy with two sentence transformers: (a) pretrained on QA pairs; and (b) distilled multilingual universal sentence encoder.

## 5 Conclusion and Future Work

We have developed Judy, a large language model (LLM) based outdoor trail recommendation chatbot based on retrieval augmented generation (RAG). We have performed data collection and case studies with the outdoor trails in Connecticut, US. Our data analysis and experimental results have revealed the lessons learned in implementing RAG-based recommendation, and demonstrated the accuracy, effectiveness, and usability of our chatbot in recommending outdoor trails.

In our future studies, we will perform more in-depth studies on: (1) integrating more data sources (e.g., weather, social network, city planning [5, 14, 15]) to enhance the outdoor trail recommendation and usability; and (2) involving more user studies to gain user-centered insights into their receptivity and acceptability about Judy.

## Acknowledgment

This project is supported, in part, by the National Science Foundation (NSF) under Grants 2303575 and 2239897, and the Connecticut Division of Emergency Management & Homeland Security (DEMHS) Hazard Mitigation Grant Program (HMGP). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

## References

- [1] Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, et al. 2018. Universal sentence encoder. *arXiv preprint arXiv:1803.11175* (2018).
- [2] R Devika, Subramaniyaswamy Vairavasundaram, C Sakthi Jay Mahenthir, Vijayakumar Varadarajan, and Ketan Kotecha. 2021. A deep learning model based on BERT and sentence transformer for semantic keyphrase extraction on big social data. *IEEE Access* 9 (2021), 165252–165261.
- [3] Matthijs Douze, Alexandre Guizhong, Chengqi Deng, Jeff Johnson, Gergely Szilvassy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. 2024. The faiss library. *arXiv preprint arXiv:2401.08281* (2024).
- [4] S Hariprasath, Arjun Paramarthalingam, Shanmugam Sundaramurthy, and Stefano Cirillo. 2024. A Study on Word Embeddings in Local LLM-based Chatbot Applications. In *2024 International Conference on Innovation and Intelligence for Informatics, Computing, and Technologies (3ICT)*, IEEE, 232–237.
- [5] Suining He and Kang G. Shin. 2022. Distribution Prediction for Reconfiguring Urban Dockless E-Scooter Sharing Systems. *IEEE Transactions on Knowledge and Data Engineering* 34, 12 (2022), 5722–5740. doi:10.1109/TKDE.2021.3062074
- [6] HuggingFace. 2025. HuggingFace Transformers (`distiluse-base-multilingual-cased-v2`). <https://huggingface.co/sentence-transformers/distiluse-base-multilingual-cased-v2>.
- [7] HuggingFace. 2025. HuggingFace Transformers (`multi-qa-MiniLM-L6-cos-v1`). <https://huggingface.co/sentence-transformers/multi-qa-MiniLM-L6-cos-v1>.

- [8] HuggingFace. 2025. HuggingFace Transformers (multi-qa-mpnet-base-cos-v1). <https://huggingface.co/sentence-transformers/multi-qa-mpnet-base-cos-v1>.
- [9] Iustina Ivanova and Mike Wald. 2023. Recommender systems for outdoor adventure tourism sports: hiking, running and climbing. *Human-Centric Intelligent Systems* 3, 3 (2023), 344–365.
- [10] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Kütter, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Proc. NeurIPS* 33 (2020), 9459–9474.
- [11] Klemen Prah. 2024. Creating vector profile graph in ArcGIS Pro based on hiking trail tracking with AllTrails app. *Dela* 62 (2024), 135–155.
- [12] Rails to Trails Conservancy. 2025. TrailLink: Bike, Run, Walk. <https://www.traillink.com/>.
- [13] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971* (2023).
- [14] Xi Yang, Suining He, and Mahan Tabatabaie. 2023. Equity-Aware Cross-Graph Interactive Reinforcement Learning for Bike Station Network Expansion. In *Proc. ACM SIGSPATIAL* (Hamburg, Germany) (*SIGSPATIAL '23*). Article 45, 12 pages.
- [15] Xi Yang, Jiachen Wang, Han Song, and Suining He. 2024. Micromobility Flow Prediction: A Bike Sharing Station-level Study via Multi-level Spatial-Temporal Attention Neural Network. *ACM SIGKDD International Workshop on Urban Computing (UrbComp)*.