

Kansas Food Pantry Assistant Project: A Pilot Study

I. PROJECT OVERVIEW

Many Kansans experience food insecurity yet struggle to access reliable, timely information about food resources in their communities. Existing food pantry directories are often static, difficult to navigate, or out of date—creating barriers for individuals who may be in urgent need, have limited time, or lack advanced digital skills. This project addresses these challenges by developing an AI-powered Pantry Assistant chatbot that allows users to find food resources through simple, natural conversation rather than manual searching.

The goal of the project is to create an interactive system that helps community members easily identify nearby food pantries and other food opportunities across Kansas. By allowing users to ask questions in everyday language—such as where to find food today or which locations are open nearby—the chatbot guides individuals to relevant, up-to-date resources in a more intuitive and accessible way.

Behind the scenes, the system combines curated food resource data with modern AI techniques to ensure responses are accurate and useful. It draws from a structured dataset of food pantries across the state, retrieves relevant information based on user needs, and generates clear, conversational responses. The chatbot is delivered through a simple, web-based interface designed to be easy to use for a wide range of users, including those with limited technical experience.

By shifting from static lists to an interactive, conversational approach, this project aims to improve food navigation for residents while supporting organizations and stakeholders seeking better insight into food accessibility across Kansas.

II. DATASET

A. Selecting the Best Source

For sourcing food pantry-related data, we have considered the following three websites:

- **Kansas Food Source** (<https://kansasfoodsource.org/>): 814 food pantry listings in Kansas.
- **Food Pantries** (<https://www.foodpantries.org/st/kansas>): 116 food pantry listings in Kansas (which is just 14% of Kansas Food Source).
- **Free Food** (<https://www.freefood.org/s/kansas>): 81 food pantry listings in Kansas (which is just 10% of Kansas Food Source).

We conducted a case study to qualitatively assess the three websites. For the city of Wichita, which is the city with the highest listed number of pantries across all three websites, the three aforementioned candidate websites had 46, 28, and 27 pantries listed respectively.

The Food Pantries website consists of a few pantries which don't exist in the Kansas Food Source website, but a lot of the information like contact numbers for a few pantries proved out to be wrong, once it was crosschecked with information from 211, Google Maps, MapQuest, and Yelp. Other examples include missing/incomplete information for a few pantries too. The Free Food website contains more complete information compared to the Food Pantries website. However, the Kansas Food Source website turned out to be the most comprehensive source among all three. Along with the accuracy of the information, under each pantry, there are exhaustive information listed, which includes the full address, contact number, hours of operation, and ID requirements. While even this source is not 100% accurate since it is manually updated after a certain interval of time, from the above analysis, we conclude that Kansas Food Source is the most accurate and comprehensive among the three candidate websites.

B. Web Scraping Data from the Chosen Source

We selected **Kansas Food Source**, a website recommended by the Kansas Food Bank for locating pantries statewide. Note that, the last date of scraping and access was on **December 17th, 2025**, so the chatbot is limited to the data updated until that day. Additionally, we utilized the U.S. Census API to obtain a list of all counties in Kansas, which helped automate scraping by county on Kansas Food Source. The Census API provides an endpoint for the 2020 Decennial Census that returns all counties for a given state. We used Python 3.11.13 for this task. The code was run on Kaggle Notebook.

C. Data Pre-processing and Indexing

We used a parser to clean the scraped HTML tags and return structured data for each pantry: name, address, zipcode, city, county, phone number, open days/hours, requirements, notes, etc. We then saved the metadata in JSON format, which was used later on by our chatbot to retrieve relevant pantries based on the user query.

Embeddings of the pantry metadata were generated using a sentence-transformers model ("all-MiniLM-L6-v2") and FAISS was used to store these embeddings in a three-dimensional vector space for semantic search.

III. SYSTEM PIPELINE AND SETUP

System Pipeline

The system follows a modular pipeline designed to support accurate, up-to-date food resource information through a conversational interface.

- **Data Collection:** Periodically gathers publicly available information on food pantries and related food opportunities from trusted web sources (e.g., Kansas Food Source).
- **Data Preparation:** Cleans, standardizes, and organizes collected information into a structured dataset to support reliable search and use.
- **Data Quality Enhancement:** Identifies missing or incomplete fields (e.g., hours, contact details) and supplements records using additional web context and related entries.
- **Conversational Access:** Enables users to query food resources using natural language through a simple, web-based **Gradio chat interface**, and receive relevant, data-grounded responses.

System Setup

The conversational component is designed to balance usability, accuracy, and transparency while minimizing the risk of misinformation.

- **Conversational Engine:** Uses large language models (LLMs) to generate clear, user-friendly responses grounded in verified pantry data.
- **LLM Variants:** Incorporates two LLMs for comparison: (i) **Llama 3.1 8B Instruct**, an open-source instruction-tuned model accessed via a HuggingFace pipeline, and (ii) **GPT-3.5 Turbo**, a widely used proprietary model known for its efficiency.
- **Controlled Response Generation:** Combines each user query with a small set of relevant pantry records and explicit system instructions to ensure accurate, non-fabricated responses that consistently include essential details (e.g., hours and contact information).
- **Deployment and Feedback:** Deploys the chatbot via the **Gradio web interface** and evaluates performance by comparing system outputs with human answers and collecting feedback from community partners, including the Health and Wellness Coalition of Wichita and the Sedgwick County Health Department's CHIP Food Access group.

IV. EVALUATION OF THE PILOT STUDY

A. Creating a Test Set of User Queries

To assess the reliability and usability of our chatbot in its pilot stage, we constructed a list of **50 queries** that can be broadly divided into the following categories. Examples have been included with the categories:

- **Queries asking for food in a specific location (city/county):** "Where can I find food pantries in Sedgwick County?"
- **Ambiguous queries not pertaining to any specific location:** "Give me a list of places I can get free food close to me."
- **Queries which mention a location but demand specific days/hours:** "Are there any places to get free food close to Derby on Wednesdays?"

- **Looking for pantries in a specific location that don't need any ID:** "Please help me find places that can give me free food close to me in Maize. What documents should I bring with me?"
- **Mentioning a specific address to look for pantries nearby:** "I am at 4912 E. New Jersey Dr Wichita KS 67210. Where can I find food near me?"
- **Trying to recall a specific pantry the user went to before:** "Where is the pantry I went to in 67214 on 21st Street?"
- **Asking for the address of a certain pantry:** "Where is the Open Door food pantry?"

B. Performance of the Chatbot

Using the test queries above, we evaluated the performance of our chatbot on the metric 'Accuracy', which measures if the chatbot gave accurate information in reply to the user query.

Llama 3.1 8B Instruct accurately gave responses to **23** of the 50 queries, while **GPT 3.5 Turbo** answered **26** user queries correctly. In short, the accuracy of the chatbot is **46%** and **52%** respectively, implemented with the two LLMs used.

V. ERROR ANALYSIS AND INTERPRETATION OF ACCURACY

The observed accuracy of approximately **50%** should be interpreted in the context of this system's current scope and the goals of a pilot study. A detailed error analysis indicates that the majority of incorrect responses arise from a small number of well-defined failure modes, rather than from pervasive system unreliability.

The largest source of error occurs for queries that specify *location constraints*, particularly at the **county level**. In many cases, the system fails to correctly resolve county names, leading to retrieval of irrelevant pantry records and subsequent recommendations outside the intended geographic area. Because the evaluation set is heavily skewed toward location-specific queries, these errors disproportionately affect the overall accuracy score. Importantly, the system performs consistently well when users specify *city names* or *ZIP codes*, indicating that the limitation is localized to county-level entity resolution rather than a general inability to process geographic information.

A second contributing factor involves *underspecified proximity queries* (e.g., "pantries near me"). The current system does not yet implement clarification or follow-up dialogue to solicit missing location information. As a result, such queries are evaluated as incorrect even though the failure stems from an unimplemented dialogue management capability rather than incorrect reasoning or retrieval.

Additional errors stem from the use of semantic retrieval for eligibility-related constraints, such as *ID required* versus *no ID required*. While semantic search enables flexible natural language interaction, it can conflate semantically opposing conditions in the absence of explicit constraint modeling, leading to incorrect pantry recommendations in a limited number of cases.

Finally, a portion of errors is attributable to data quality issues in the scraped dataset. Because retrieved records are directly used to ground response generation, inconsistencies in field normalization (e.g., address information appearing in auxiliary fields such as *Notes*) can prevent relevant pantries from being retrieved for certain queries. These issues reflect limitations of the underlying data rather than failures of the conversational model itself.

Taken together, these findings indicate that the reported accuracy primarily reflects known, addressable limitations in geographic resolution, dialogue handling, constraint modeling, and dataset normalization. As such, the current score should be viewed as a conservative baseline for this pilot system, with substantial room for improvement through targeted refinements rather than fundamental redesign.

VI. LIMITATIONS AND PATH FORWARD

The chatbot architecture is intentionally modular, consisting of two tightly coupled components: a *retrieval module* and a *language model*. The retrieval module interprets user queries, identifies relevant constraints (e.g., location, hours, eligibility requirements), and selects appropriate food pantry records from a structured dataset. The language model then generates conversational responses grounded in this retrieved information. As a result, overall system performance is fundamentally bounded by the quality of retrieval. Regardless of the language model employed, inaccurate or incomplete retrieval inevitably limits response accuracy.

This design choice is reflected in the experimental results. While a proprietary model such as GPT-3.5 Turbo offers modest improvements over an open-source alternative (Llama 3.1 8B Instruct), the gains are incremental, underscoring that retrieval quality—not model capacity—is the dominant performance bottleneck. Consequently, the primary focus of

future development will be to strengthen the retrieval logic to robustly handle edge cases, complex constraints, and diverse user phrasing, thereby enabling the language model to operate on consistently relevant and reliable inputs.

Beyond retrieval, several targeted improvements will further enhance system reliability and usability. First, the evaluation query set should be expanded to better reflect the range of questions real users are likely to ask. Second, automated evaluation metrics should be introduced to complement accuracy-based measures, such as the number of conversational turns required to resolve a query and system response latency, providing a more holistic assessment of user experience. Third, the data collection pipeline should be strengthened to ensure consistent field normalization (e.g., addresses, hours, and eligibility information), reducing errors propagated from the underlying dataset.

From a deployment perspective, extending the current web interface beyond short-term hosting constraints will be essential for sustained community use. In addition, relying on a single data source, while appropriate for an initial prototype, highlights the need for multi-source integration and periodic updates to ensure information remains current and trustworthy. Finally, a structured user study will be critical for validating real-world usability, uncovering subtle interaction failures, and guiding iterative refinement.

While this work represents an early-stage pilot, the results demonstrate the feasibility of deploying a conversational, data-driven assistant for food resource navigation using largely unstructured public data. Despite current limitations, the system already exhibits the core functionality required for a 24/7, on-demand support tool. The identified limitations are well-scoped, addressable, and primarily engineering-focused, positioning this project for meaningful impact with continued development and stakeholder engagement.