Forum for Linguistic Studies | Volume 07 | Issue 11 | November 2025

BILINGUAL PUBLISHING GROUP

Forum for Linguistic Studies

https://journals.bilpubgroup.com/index.php/fls

ARTICLE

AI Chatbots for Personalized Sustainable Nutrition: Bridging Technology, Engagement, and Ethics

Hanh Ngo Minh Truong, Andrii Kolodiazhnyi, Serhii Sokyrko, Priyanka Darbari *

Department of ICT Innovation Lab, Fontys University of Applied Sciences, 5651 GW Eindhoven, The Netherlands

ABSTRACT

Sustainable food development is crucial for minimizing environmental impacts and ensuring the capacity to provide sufficient food for both present and future generations. Many eco-friendly production methods have been adopted worldwide, including organic farming, regenerative agriculture, and plant-based alternatives, aimed at reducing greenhouse gas emissions, conserving water and soil resources, and promoting biodiversity. However, despite this development in sustainable production, consumer awareness and adoption of sustainable food choices remain limited, preventing full environmental and health impacts of these practices from being realized. This paper represents the design of an AI-powered chatbot, offering nutrition guidance, promoting sustainable and healthy daily food choices, while also addressing ethical considerations such as user privacy, fairness, and transparency in its design. The chatbot integrates artificial intelligence and large language models, adapted with domain-specific data on nutrition and sustainability, to engage users in conversations about healthy eating, food waste reduction, and eco-friendly diets. Its design combines a user-friendly interface, a curated knowledge base, and personalized recommendations informed by user preferences. Early evaluations suggest that the system can increase awareness and encourage more sustainable food choices. Ethical aspects such as privacy, transparency, and fairness are embedded in its development to promote responsible use of AI. Future enhancements may include integrating image-based calorie estimation to provide personalized nutritional feedback alongside sustainability guidance.

*CORRESPONDING AUTHOR:

Priyanka Darbari, Department of ICT Innovation Lab, Fontys University of Applied Sciences, 5651 GW Eindhoven, The Netherlands; Email: p.darbari@fontys.nl

ARTICLE INFO

Received: 29 June 2025 | Revised: 16 September 2025 | Accepted: 19 September 2025 | Published Online: 24 October 2025

DOI: https://doi.org/10.30564/fls.v7i11.10799

CITATION

Truong, H.N.M., Kolodiazhnyi, A., Sokyrko, S., et al., 2026. AI Chatbots for Personalized Sustainable Nutrition: Bridging Technology, Engagement, and Ethics. Forum for Linguistic Studies. 7(11): 1135–1156. DOI: https://doi.org/10.30564/fls.v7i11.10799

Keywords: Sustainable Food Development; Food Sustainability; Future Food Value; Digital Transformation; AI Chatbot

1. Introduction

With the development of the world, the global food system contributes significantly to environmental problems, producing around 26% of global greenhouse gas emissions [1]. These impacts come from many stages of food production, including farming, transportation, and food waste, with livestock alone responsible for nearly 14.5% of emissions [1]. Across five representative socio-economic scenarios, global food demand is projected to rise by 35% to 56% between 2010 and 2050, while the population at risk of hunger is expected to change by −91% to +8% over the same period [2]. These scenarios also indicate a steady movement toward sustainability, characterized by improvements in equity, a stronger focus on overall human well-being, and dietary and consumption habits that rely less on meat and use fewer resources and less energy [2]. However, promoting healthy and sustainable dietary practices remains challenging. Reviews indicate that many mobile health and nutrition apps face adoption barriers due to insufficient evidence-based content, limited user engagement, and burdensome self-reporting requirements. Recent research highlights the potential of AI-powered chatbots to overcome these limitations by providing interactive, personalized guidance, facilitating long-term adherence to healthy behaviors, and supporting nutrition education across diverse user groups [3].

Recent work highlights the role of AI technologies in food computing, where intelligent decision-making, predictive modeling, and personalized nutritional recommendations improve users' ability to manage their diets effectively and make healthier choices [4]. Building on these capabilities, AI-powered platforms for diet and fitness tracking have been shown to directly benefit users by helping them monitor nutrient intake, receive personalized meals and workout recommendations, and maintain long-term healthy behaviors [5]. By providing guidance based on personal goals and health conditions, such platforms empower users to make informed decisions about their nutrition and lifestyle. These findings complement research showing that AI can improve productivity and sustainability across the food and nutrition system [6], suggesting that user-centered AI solutions play a key role in promoting both personal well-being and broader sustainable food practices.

Conversational AI agents such as chatbots offer an interactive medium for education and behavior change. Chatbots can engage users in dialogue, providing instant feedback and tailored information in a more natural way than static websites or brochures. Prior studies in educational settings show that chatbot-based learning can enhance user engagement and improve knowledge retention [7,8]. Their conversational nature allows complex topics to be broken down into relatable, personalized discussions, which is especially useful for subjects like nutrition and sustainability that benefit from personal context. For example, recent research has shown that chatbots designed with empathic or personalized styles can increase user satisfaction and motivation to act on the advice given [9,10]. Despite these advantages, the use of chatbots to promote sustainable eating habits and individualized diet guidance is still nascent [11]. There is a clear research opportunity in exploring how an AI chatbot might effectively raise awareness about sustainable food choices while also delivering personalized nutrition tips.

This paper addresses that gap by examining the development of an AI-driven chatbot platform aimed at improving dietary behavior in alignment with sustainability goals. A main research question has been conducted: How can AI-driven nutrition systems influence personalized food habits while promoting sustainable choices, raising awareness of nutritional and environmental impacts, and addressing ethical and user engagement challenges? To address this research question, we combined insights from literature on personalized nutrition, behavior-change technology, and human-AI interaction with a user-centered development process. The goal was to create a conversational agent that not only provides customized diet recommendations but also educates users on future food values, such as reducing food waste, choosing eco-friendly ingredients, and considering the ethical dimensions of food choices. In the following sections, we review related work in this domain, describe the methodology for building and fine-tuning the chatbot, present the system's evaluation results, and discuss the implications for using AI chatbots to foster healthier and more sustainable eating habits.

2. Literature Review

AI in Personalized Nutrition: Personalized nutrition has gained traction as an effective approach to improve dietary outcomes by tailoring advice to individual needs. AI-driven solutions are increasingly being explored to support this personalization. For instance, a European project developed a nutrition app that generates individualized meal plans based on user data, demonstrating the feasibility of AI to fine-tune diets better than unguided human choices [12]. In a controlled trial with over 160 adults, a web-based personalized nutrition tool (eNutri) delivered automated diet advice that led to a significant improvement in diet quality compared to generic guidelines [13]. Participants who received tailored recommendations improved their healthy eating index scores by about 6% and many maintained positive changes even months after the study - highlighting how personalized digital advice can motivate sustained behavior change [13]. These findings reinforce the rationale for an AI-powered personalized approach: by adapting guidance to each user's dietary profile and preferences, such tools outperform one-size-fits-all advice in encouraging healthier eating.

LLMs and Meal Planning: Recent advances in large language models (LLMs) have opened new possibilities for generating nutrition advice on a scale. Papastratis [14] introduced a hybrid AI system that combines a generative model with expert nutritional rules to create weekly meal plans. In their framework, a variational autoencoder first proposes diet plans based on user health data, then an LLM (ChatGPT) expands those plans with varied meal options, all while nutritional guidelines are enforced programmatically. The system produced highly accurate meal plans that met individual requirements with minimal errors [14]. Notably, using the LLM introduced more variety and creativity in meal suggestions without breaking dietary constraints, although purely LLM-generated plans occasionally suggested unsafe or impractical meals until expert rules were integrated. This underscores a key point: AI's generative power can provide diverse and engaging suggestions but coupling it with domain-specific constraints is crucial for maintaining nutritional correctness. Similarly, Yang [15] developed ChatDiet, a chatbot framework that augments an LLM with personal and population nutrition models to deliver explainable food recommendations. Their system achieved a 92% success rate in meeting users' health goals in a case study and, importantly, provided clear justifications for its advice [15]. Users reported greater trust in the recommendations when the chatbot explained, for example, that "broccoli is high in fiber, which aligns with your goal to improve gut health." These studies demonstrate that large language models, when properly guided by nutritional data and explanatory design, can effectively generate personalized diet plans and advice. They inform our chatbot's design by illustrating how to harness AI's strengths (speed, variety, personalization) while ensuring safety and transparency through hybrid approaches.

Chatbots for Health Behavior Change: There is growing evidence that AI chatbots can influence users' health behaviors in positive ways. In a systematic review of 15 studies on chatbots for various behavior-change interventions (diet, physical activity, smoking cessation, etc.), Aggarwal [16] found that a majority of the chatbot-based interventions led to improvements in health behaviors or outcomes. In about 40% of the reviewed studies, participants using AI chatbots showed significant positive changes, for example, improved diet quality or increased physical activity levels, compared to control groups [16]. Users commonly appreciated the convenience of 24/7 access to coaching and the personalized tips provided by chatbots. However, the review also noted that many interventions were short-term, and more research is needed on long-term adherence. A frequent limitation observed was that some chatbots had limited understanding or overly simplistic dialogue, which could frustrate users over time. These findings support the premise that an AI nutrition chatbot can successfully promote healthy behavior change (e.g., encouraging sustainable diet habits), but they also highlight the importance of robust conversational ability and engaging design. To avoid the pitfalls identified, our project places emphasis on natural language understanding (so the bot can handle a variety of user inputs) and sustained user engagement techniques (to encourage long-term use), as well as plans for rigorous evaluation of efficacy over time.

User Engagement and Personalization: Effective user engagement is critical for any educational or behavior-change chatbot. Prior research indicates that personalization and an empathetic tone can significantly boost user satisfaction and trust in chatbot interactions. Kuhail[7] conducted a systematic review of educational chatbots and found that many systems lacked adaptivity and user-centered design, leading to user frustration when the chatbot responses felt too generic or misunderstood nuanced queries. The authors suggest incorporating distinct chatbot personalities or localized content to make interactions feel more natural[7]. In a business context, Juquelier[9] showed that empathic chatbots (those that use polite, encouraging language and express understanding of user emotions) can enhance customer experience, but they caution it is a "double-edged sword". If an empathic chatbot fails to meet expectations (for example, by giving a wrong answer right after sounding caring), it can actually disappoint users more than a neutral bot would[9]. This insight is relevant for our nutrition chatbot: we aim for a friendly, supportive tone to build engagement, but we must ensure the information provided is accurate and useful to maintain trust. From the perspective of technology acceptance, personalization is also key. A recent study by Kazoun[10] on chatbots in sustainability education found that perceived usefulness and personalization were strong drivers of user adoption, whereas concerns about privacy and mistrust of AI were barriers. Users were more willing to engage with the chatbot when they felt it addressed their individual queries and provided fair, unbiased responses[10]. These findings reinforce our dual focus on personalization and ethics: by tailoring content to each user and being transparent about how recommendations are generated, we can improve engagement and mitigate some mistrust.

Calorie Estimation and Nutritional Analysis Using AI

Recent developments in artificial intelligence have shown significant potential in supporting nutritional evaluation and calorie estimation. Studies have demonstrated that AI-driven platforms can provide users with personalized dietary guidance, monitor nutrient intake, and encourage healthier dietary behaviours[17]. Similar approaches in applications such as Diet Recall have integrated AI to track food consumption, analyse nutrient content, and offer personalized recommendations, highlighting the practical benefits of user-centered AI systems in promoting long-term wellness[5]. Moreover, evaluations of AI-enabled nutrition apps have shown that higher integration of AI features improves functionality and user experience, although accuracy in calorie estimation can vary depending on the food type and dataset used[18]. Together, these findings justify the development of AI-based platforms for calorie estimation and personalized nutritional guidance. Key studies, including systematic reviews and empirical investigations on AI chatbots and nutrition applications, are summarized in **Table 1**.

Table 1. Selected literature on AI chatbots and their implications for this study.

Author(s), Year	Title/Study Focus	Research Objective	Methodology/ Approach	Key Findings/Results	Relevance to This Study
Aggarwal et al. [16]	Artificial Intelligence–Based Chatbots for Promoting Health Behavioral Changes: Systematic Review (JMIR)	Evaluate the feasibility and efficacy of AI chatbots in health behavior change interventions (multiple domains).	Systematic review of 15 empirical studies (1980–2022) on AI-driven chatbots for behavior change (diet, physical activity, smoking cessation, etc.). Assessed outcomes (e.g., behavior adherence, clinical measures) and intervention characteristics.	Among the reviewed studies, many reported **high efficacy of AI chatbots in improving healthy behaviors**. For example, 40% of the studies (6 out of 15) showed significant positive effects on lifestyle behaviors like diet quality or physical activity levels. Chatbots were generally found to be feasible and acceptable; users appreciated the instant feedback and personalized tips. However, few were long-duration RCTs – authors caution that more robust trials are needed to confirm long-term effectiveness. Common strengths: convenience and engagement; common limitations: some chatbots had limited understanding or overly simple logic.	Provides general evidence that **AI chatbots can successfully promote behavior change**, supporting our project's premise (RQ2 benefits). It also justifies aspects of our methodology: focusing on engagement and personalization was key to success in these studies (RQ4). The call for more rigorous evaluations guides how we plan our assessment of the chatbot. It also alerts us to typical weaknesses (ensuring our chatbot's conversational ability is sufficiently advanced, addressing RQ2 limitations regarding understanding).
Capecchi et al. [17]	*Enhancing education outcomes integrating augmented reality and artificial intelligence for education in nutrition and food sustainability*	Evaluate the effectiveness of AI feedback in educational tools for nutrition and food sustainability.	Study on the ARFood app with AI-driven NPCs providing feedback on virtual food choices. Alignment of AI feedback with educational objectives assessed using a zero-shot RoBERTa classifier and iterative prompt refinement.	Iterative refinement of AI feedback improved coverage across all educational objectives. AI tools **enhanced the app's ability to deliver balanced, goal-oriented guidance** in nutrition and sustainability education.	Shows that AI can provide adaptive feedback for nutrition education, supporting the idea that AI platforms can help users understand food-related information. Justifies using AI-driven evaluation and feedback mechanisms in our project.

Table 1. *Cont.*

Author(s), Year	Title/Study Focus	Research Objective	Methodology/ Approach	Key Findings/Results	Relevance to This Study
Chen et al. [19]	*Conversational AI and equity: Assessing GPT-3's communication with diverse groups on contentious topics* (Scientific Reports)	Investigate how a large language model (GPT-3) interacts with users of differing beliefs on controversial issues (climate change and racial justice), and whether those conversations alter users' attitudes.	Massive user study: 3,000+ participants from diverse backgrounds engaged in a live chat with GPT-3 on climate change or Black Lives Matter. Pre- and post-chat surveys measured attitudes on these issues. Also collected user experience ratings. Analysis looked at attitude shifts in relation to initial stance (skeptical vs. supportive) and user demographics.	Conversations with an AI **significantly shifted users' attitudes** toward the factual consensus on climate change, even among initially skeptical individuals. After a single chat session, skeptics showed notably more agreement with statements about human-caused climate change. Interestingly, these users rated their chat experience less satisfying than others (they found the bot somewhat unconvincing or frustrating), yet they still absorbed the information and moved closer to the scientific viewpoint by the end of the chat. This reveals a tension: the **bot was effective in educating/persuading**, but those who changed their views didn't necessarily enjoy the process. It highlights the need to improve the conversational quality to match the educational impact. In short, the chatbot did raise awareness and correct misconceptions, albeit with mixed user satisfaction	Offers insight into **measuring and improving chatbot effectiveness** (RQ6). It confirms that even general AI chatbots can educate users and change opinions on sustainability-related issues. For our project, it underscores the importance of evaluating both outcomes: knowledge/attitude change (did we increase awareness?) and user satisfaction. If there's a gap (people learn but feel frustrated), iterative improvements in the chatbot's communication style or empathy may be needed. This study supports using pre/post interaction surveys to gauge our chatbot's impact on awareness of food sustainability, and reminds us that **user feedback** is crucial for refining the chatbot's approach to be both informative and engaging.
Deng & Yu [8]	A meta-analysis of the effectiveness of chatbot technology on learning outcomes	Aggregate evidence from numerous studies to determine overall how much chatbots help learning, and which outcomes or conditions they influence most.	Meta-analysis of 32 experimental studies (2010–2022) involving chatbot-based learning interventions. Calculated effect sizes for various learning outcomes (knowledge retention, reasoning, motivation, etc.) and examined moderators like subject domain, chatbot role (tutor/peer), and intervention duration.	Chatbot-based learning has a **significant positive effect** on students' learning outcomes overall. The meta-analysis found a medium-to-high overall effect size for learning performance with chatbots. In particular, using chatbots yielded strong gains in factual knowledge retention and in learners' reasoning skills (students could explain concepts better after chatbot interaction). It also modestly increased learning interest. However, chatbots alone did not significantly improve students' general motivation or critical thinking skills compared to traditional learning, suggesting that while they are great at delivering content and practicing concepts, additional strategies might be needed to impact deeper motivation. Notably, the effectiveness was consistent across different academic subjects and whether the chatbot acted as a tutor or peer, indicating the broad applicability of chatbots in education	Confirms that **chatbots can be effective educational tools**, supporting our assumption that a chatbot can improve user knowledge about sustainability (RQ6). It also hints at what to expect: we can likely achieve gains in users' understanding of future food concepts and perhaps their ability to reason about their food choices. The finding that motivation gains were minor means we should pay special attention to engagement strategies (personalization, interactivity – as addressed in RQ5) to also inspire users, not just inform them. This comprehensive evidence base bolsters the validity of using a chatbot for our awareness goals and provides benchmarks for evaluating success (e.g., improved user knowledge and interest).
Namkhah et al. [6]	Advancing sustainability in the food and nutrition system: a review of artificial intelligence applications	Review how AI technologies can enhance sustainability across the food system, including production, distribution, and consumption.	Narrative review covering AI use-cases in agriculture (smart farming), food processing, supply chain management, and consumer dietary choices that collectively improve sustainability and food security.	Emphasizes that doing PN "right" requires **multi-dimensional data** (diet, microbiome, genomics, behavior) and that frameworks are needed to combine these data types effectively. Introduces guiding principles: (1) Focus on the end-user (ensure recommendations are understandable and actionable); (2) Ensure data privacy and obtain informed consent (many PN services are direct-to-consumer, raising consent issues); (3) Strive for transparency in algorithms and personalization logic; (4) Build trust by demonstrating efficacy and safeguarding personal information. Also notes regulatory challenges – PN tools span health and food domains, often with unclear oversight.	Provides high-level **ethical and practical guidelines** that directly inform RQ3 (ethical AI use in nutrition). Reinforces our attention to data privacy and transparency in the chatbot's design. Also, the emphasis on combining diverse data informs our system's planned capability to incorporate various user inputs (health metrics, preferences) for better personalization (RQ1/RQ2). Overall, it frames our work within best-practice principles to ensure it is responsible and aligned with emerging standards in personalized nutrition.
Franco et al. [13]	Effectiveness of Web-Based Personalized Nutrition Advice for Adults Using the eNutri Web App: EatWellUK RCT	Evaluate whether a personalized nutrition web-app can improve diet quality in adults, compared to standard dietary advice.	Randomized controlled trial (12-week) with 160+ UK adults. Intervention group used the eNutri app, which generated automated personalized nutrition advice based on an initial food frequency questionnaire. Control group received generic healthy eating guidelines. Diet quality was measured by a modified Healthy Eating Index at baseline and week 12.	The personalized nutrition app led to a **significant improvement in diet quality** compared to control. E.g., overall diet quality score increased ~6% in the personalized group vs. control. Users in the intervention also reported higher engagement in healthy eating behaviors. The app's automated feedback (tailored to each person's dietary gaps) was effective at motivating positive change, and a follow-up found many users maintained some improvements months after.	Provides evidence that **personalized digital advice outperforms generic advice** – reinforcing the rationale for an AI-driven personalized approach (RQ1). It also exemplifies a successful implementation of automated diet guidance, which we build upon. We can use its findings to set benchmarks (e.g., aiming for improvements in diet scores) and learn from its design (simple feedback based on diet score) to shape our chatbot's recommendation strategy (RQ2 evaluation of advantages).
Izadi & Forouzan-far [20]	Error Correction and Adaptation in Conversational AI: A Review of Techniques and Applications in Chatbots	Provide a comprehensive review of common chatbot errors and strategies to mitigate them, to improve user experience.	Literature review covering types of chatbot errors (misunderstandings, irrelevant answers, etc.), their root causes, and correction techniques (from rule-based fixes to machine learning and reinforcement learning approaches).	Chatbots often make errors that **undermine user trust and satisfaction**. The review found that strategies like clarification prompts (asking users for confirmation), real-time error correction, and adaptive learning (chatbot learning from mistakes or user feedback) can significantly improve interaction quality. Integrating robust error-handling (e.g., fallback responses, self-correction algorithms) is critical to maintaining engagement. Emerging techniques (reinforcement learning-based dialogue repair) show promise in reducing repetitive mistakes and improving response relevance	Directly addresses **how language techniques can overcome interaction issues** (RQ2). By implementing clarification questions or learning from user feedback, our chatbot can avoid common pitfalls (misunderstood queries or off-topic answers), leading to a smoother, more trustworthy user experience. This insight guides us to include error-handling and adaptive responses in our design.

Table 1. *Cont.*

Author(s), Year	Title/Study Focus	Research Objective	Methodology/ Approach	Key Findings/Results	Relevance to This Study
Jaiswal [5]	Revolutionizing diet and fitness tracking with AI: A user-centric approach to nutrition and wellness	Evaluate the effectiveness of a user-focused AI platform for diet and fitness tracking.	Development and evaluation of Diet Recall app integrating AI to provide personalized diet and fitness recommendations based on user data.	The app enabled users to monitor intake, receive personalized recommendations, and track progress. Users reported **positive experience and better engagement** with diet and fitness behaviors.	Supports the approach of using AI to provide personalized nutritional guidance. Justifies designing a user-centered AI platform for calorie estimation and wellness tracking.
Juquelier et al. [9]	Empathic chatbots: A double-edged sword in customer experiences	Determine how adding emotional intelligence (empathy) to chatbots affects user satisfaction and under what conditions it helps or harms the experience.	Experimental study (three separate experiments) in a customer service scenario. Chatbots were configured with high empathy (e.g., apologizing and showing concern for the user's issue) vs. neutral responses. Measured outcomes: user perceived social presence (how "human" the bot felt), information quality, satisfaction, and willingness to continue using the chatbot. Also tested a condition with time pressure to see if speedy vs. empathetic responses fare differently.	**Empathy in chatbot responses** substantially increased customer satisfaction in general. Users interacting with an empathetic chatbot felt a greater sense of social presence and thought the information provided was higher quality, which led to more positive overall evaluations [8]. However, in high time-pressure situations (when users wanted very quick help), the extra empathetic dialogues annoyed some users, slightly reducing satisfaction compared to a more task-focused bot. Thus, empathy builds a stronger connection and trust, but users sometimes prioritize efficiency over empathy.	Highlights the role of **conversational tone** in user engagement (RQ5). For our sustainability chatbot, adopting an empathetic, friendly tone can make users feel more at ease and connected (important for discussing personal lifestyle topics like diet). We should, however, balance this with efficiency – for instance, quickly providing facts when asked, and using empathy where appropriate (such as encouraging a user who is trying to change habits). This balance can improve user satisfaction and trust, supporting sustained engagement with the chatbot.
Kaçar et al. [21]	Diet Quality and Caloric Accuracy in AI-Generated Diet Plans: A Comparative Study Across Chatbots (Nutrients)	Compare the nutritional quality and accuracy of meal plans generated by different AI chatbot models for weight management.	Experimental comparison using three AI chatbots (ChatGPT-4, Microsoft Bing's AI, and Gemini). Each was prompted to create sample daily meal plans for hypothetical individuals (male and female, with set calorie targets for weight loss). Assessed each plan's caloric deviation from target and diet quality using a standard index.	**All AI chatbots produced nutritionally balanced meal plans**, including all major food groups. ChatGPT-4 had the best caloric accuracy – none of its plans deviated >20% from the goal, whereas 50% of Gemini's plans overshot calories by >20%. Diet quality scores were high for all, though minor differences appeared: e.g., ChatGPT-4 and Copilot included a greater variety of protein sources, while Gemini's plans were slightly less varied in that aspect. Some **gender-specific bias** was noted (one model omitted eggs in female plans but not male). Overall, AI models can generate acceptable diet plans, but consistency varies.	Validates that AI models are **capable of generating healthy meal plans** (supporting RQ2 advantage of AI in personalization). The findings about variability and small biases inform our approach to **evaluating and calibrating the chatbot's recommendations**. We will strive for the consistency of the best model (ChatGPT) and be wary of potential biases (like gender-based suggestions), aligning with RQ2 (limitations) and RQ3 (fairness). This study supports using a robust LLM (like GPT-4) in our project for higher accuracy.
Kazoun et al. [10]	AI Chatbots for Sustainability in Education: The case of the Lebanese HE sector	Examine factors influencing stakeholders' intention to use AI chatbots for sustainability education.	Survey of university stakeholders, analyzed via an extended UTAUT model (technology acceptance framework), to identify drivers and barriers to chatbot adoption for education.	**Challenges identified:** security/privacy concerns and mistrust of AI can deter chatbot use. Other factors affecting intention include perceived fairness, fear of tech, and need for human touch. Positive factors were individualized learning (chatbot personalization), collaborative learning support, and learning motivation. An extended adoption model is proposed, incorporating these concerns and enablers.	Highlights **current limitations** from a user perspective (RQ1): users worry about privacy and trust with chatbots, and need personalization and fairness. Emphasizes that addressing security and building trust (through transparency and reliable info) is crucial for our sustainability chatbot to be accepted, and that personalization can boost adoption – aligning with RQ5 (personalization's role).
Kuhail et al. [7]	Interacting with educational chatbots: A systematic review	Identify and compare recent design techniques for educational chatbots, including their limitations.	Systematic review of 36 studies on chatbots in education (analyzed by field, design principles, role, interaction style, evidence, etc.).	Educational chatbots improved learning and engagement in many cases. Most bots followed predetermined paths, with some using personalization. **Challenges:** insufficient training data leading to misunderstandings and lack of user-centered design (no usability heuristics), causing user frustration. Authors suggest exploring chatbot personality and localization to improve user satisfaction.	Establishes common **user interaction challenges** (e.g., poor understanding due to data limits, rigid design) in current chatbots. Underscores the need for better data/training and user-centered design, informing RQ1 (limitations) and RQ5 by noting personalization/personality as a way forward.
Kunja et al. [22]	Engaging guests for a greener tomorrow: Role of hotel chatbot concierges in sustainable practices	Investigate whether a human-like (anthropomorphic) chatbot can persuade hotel guests to adopt eco-friendly behaviors during their stay.	Field experiment in a hotel context (or scenario-based survey): participants (young guests) interacted with a chatbot "concierge" that gave tips on sustainable actions (reusing towels, conserving energy, etc.). The chatbot was designed with a personable, friendly persona. Researchers then measured guests' willingness to follow the sustainability recommendations and their perceived connection to the chatbot.	The **anthropomorphic chatbot concierge** was effective in encouraging sustainable behavior. Guests who chatted with a friendly, human-like bot showed a higher willingness to opt into green practices (e.g., 20% higher rate of towel reuse sign-up compared to those who only saw a standard message). The chatbot's conversational, polite reminders (and even using the guest's name) created a sense of social obligation and rapport, which translated into action. This suggests that when users feel they are interacting with a "person" who cares, they are more likely to heed sustainability advice.	Provides evidence that **personalization and human-like tone improve engagement and behavior change** (RQ5). In our project context, giving the chatbot a bit of personality – a friendly guide who remembers the user's name or preferences – could similarly increase users' receptiveness to sustainability tips about food. It also reinforces that a conversational agent can move people from awareness to taking action (e.g., trying an eco-friendly practice), which is a key goal of raising awareness of future food values (RQ6 – effectiveness in driving change).

Table 1. *Cont.*

Author(s), Year	Title/Study Focus	Research Objective	Methodology/ Approach	Key Findings/Results	Relevance to This Study
Li et al. [18]	Evaluating the quality and comparative validity of manual food logging and AI-enabled food image recognition in apps for nutrition care	Assess quality, validity, and behavior change potential of AI-enabled nutrition apps.	Screening of top 200 free and paid nutrition apps; compared manual logging and AI image recognition against dietary records; quality assessed using MARS and behavior change potential using ABACUS.	Apps with higher AI integration **demonstrated better functionality and user experience**, though calorie estimations could be inaccurate depending on food type. Highlights the need for AI training and improved food databases.	Supports justification for developing AI-based calorie estimation tools. Demonstrates both the promise and limitations of AI nutrition apps, guiding improvements in accuracy and usability.
Menkhoff & Gan [11]	Engaging students through conversational chatbots and digital content: A climate action perspective	Explore the use of a chatbot as a teaching tool to increase student engagement and understanding in a sustainability course (focused on climate action).	Case study in a university setting: implemented a chatbot that answered questions and provided information on climate change and sustainability topics. Students in a climate-action course were encouraged to interact with the chatbot outside of class. Data collected via student feedback surveys and engagement metrics (frequency and length of chatbot interactions).	Students reported that the chatbot made learning about climate change more interactive and accessible. Many used it to clarify concepts (e.g., "What is carbon footprint?") and appreciated the instant, conversational explanations. Engagement metrics showed high usage during project work periods. Qualitative feedback noted that the chatbot's approachable language helped demystify complex environmental concepts. However, some students noted limitations in depth for very technical questions. Overall, the chatbot complemented traditional teaching by actively engaging students and reinforcing material in a novel way.	Demonstrates real-world viability of a **chatbot for sustainability education**. It confirms that a conversational agent can capture students' interest and reinforce learning on topics like climate change. For our project, this underscores the potential impact of a chatbot on raising awareness of future food values. It also flags the importance of ensuring sufficient depth in the chatbot's knowledge base to handle advanced queries (relates to RQ4's focus on content delivery and RQ6's focus on evaluation through user feedback).
Nguyen et al. [23]	Value-sensitive design of chatbots in environmental education: Supporting identity, connectedness, well-being and sustainability	Design a sustainability education chatbot that aligns with learners' personal values and well-being, not just factual knowledge, to enhance engagement.	Employed Value-Sensitive Design (VSD) methodology: interviewed students and educators to identify key values (e.g., personal identity related to nature, sense of connectedness to community, concern for well-being) that the chatbot should support. Designed chatbot dialogue scenarios around climate change that incorporate these values (e.g., reflective questions linking climate issues to the user's life). Qualitatively evaluated user experience and perceived value support.	Found that integrating **personal and emotional values** into the chatbot's dialogue greatly improved student engagement and learning. For example, scenarios where the chatbot asked users to recall personal experiences with food waste or climate impacts made discussions more relatable and impactful. The chatbot that "talked" in a way that supported users' identity and well-being (encouraging optimism and self-efficacy) was rated as more trustworthy and motivating. This study illuminates that beyond imparting facts, a chatbot should foster a personal connection to sustainability topics, which increases its educational effectiveness [6].	Shows how to **adapt chatbot content and tone for sustainability topics** (RQ4) by making it relatable and value-driven. These findings encourage us to ensure our chatbot's messaging resonates on a personal level (e.g., using examples from daily life, adopting an encouraging tone). By aligning with users' values (like health or community in food choices), our chatbot can better inspire engagement and behavior change, addressing RQ5 as well (the role of tone/personalization in engagement).
Nguyen [24]	The Effect of AI Chatbots on Pro-environment Attitude and Willingness to Pay for Environmental Protection	Quantify how interactions with an AI chatbot influence users' environmental attitudes and their willingness to financially support environmental initiatives.	Used a structured survey-based study (in Vietnam) with statistical modeling. Participants interacted with an environmental chatbot (features included interactive problem-solving, anonymity of interaction, customization to user queries). Afterwards, participants' pro-environment attitudes and willingness to pay (WTP) for environmental causes were assessed. Data was analyzed with multivariate regression/SEM to see which chatbot features affected outcomes.	**Chatbot interactions can increase pro-environmental attitudes.** This study found that when the chatbot effectively solved users' queries and problems (high problem-solving ability), users felt more informed and positive about environmental issues, which strongly raised their willingness to donate or pay for environmental protection. Customization (the chatbot tailoring information to user interests) and rich interaction also indirectly boosted willingness to pay by first improving the user's attitude. Notably, providing anonymity (users feeling free to ask anything) helped users engage more openly, though the strongest direct driver of attitude change was the chatbot's capability to give useful, relevant answers. In sum, a well-designed chatbot can both educate and motivate tangible support for sustainability.	Demonstrates tangible **effectiveness of chatbots in raising environmental awareness and commitment** (RQ6). It guides our design to emphasize problem-solving and personalized content delivery, as these were key to shifting attitudes. By ensuring our chatbot provides actionable advice (not just facts) and adapts to users' questions (showing high utility), we can similarly inspire users to not only learn about future food values but potentially act on that knowledge (e.g., being willing to invest in sustainable food choices or programs).
Papastratis et al. [14]	*AI nutrition recommendation using a deep generative model and ChatGPT (Scientific Reports)*	Develop an AI system that generates weekly meal plans by combining generative AI with expert nutritional guidelines, and evaluate its accuracy.	Proposed a hybrid model: a variational autoencoder encodes user data (anthropometrics, health conditions) to generate initial diet plans, then ChatGPT (LLM) expands plans with diverse meals. Nutritional guidelines (from EFSA and WHO) are enforced via custom loss functions. Tested on 3,000 virtual profiles and 1,000 real profiles by comparing AI-generated meal plans to individual requirements.	The system produced **highly accurate and personalized meal plans**. In simulations, 7-day meal plans met users' energy requirements with minimal error and adhered to macronutrient recommendations. Using ChatGPT introduced more variety in cuisines and meal options without violating nutritional constraints. Compared to traditional rule-based plans, the AI plans were more diverse yet still balanced. Highlight: out of 84,000 AI-generated daily meal plans, the vast majority fell within recommended nutrient ranges for their respective profiles. However, authors note that using ChatGPT alone (without guideline alignment) sometimes yielded unsafe suggestions, underscoring the value of integrating expert rules.	Demonstrates a successful strategy to leverage AI **advantages (speed, variety)** while controlling for diet correctness – relevant to **RQ2**. Informs our approach to combine AI with rule-based checks for safety. It also highlights potential **limitations**: a pure LLM might err, so we plan to incorporate nutritional guidelines as they did. This work validates that AI can generate appealing, varied meal plans, which supports our goal of keeping users engaged with new ideas (RQ4) and meeting nutritional needs (RQ2).

Table 1. *Cont.*

Author(s), Year	Title/Study Focus	Research Objective	Methodology/ Approach	Key Findings/Results	Relevance to This Study
Seo & Yoon [25]	Promoting mindful consumption through a chatbot with an experiential mind	Determine how a chatbot's communication style can encourage consumers to adopt mindful, sustainable consumption habits, and explain why it works.	Experimental study in consumer context: compared a chatbot that uses an "experiential" conversational style (engaging the user in stories or simulations about their consumption) versus a more traditional informational chatbot. Measured outcomes like users' intention to practice mindful consumption (e.g., reducing waste, choosing sustainable products) and examined psychological mediators (such as emotional engagement or self-reflection triggered).	The **experiential-style chatbot** led to a greater increase in mindful consumption intentions than a plain informational chatbot. By immersing users in relatable scenarios (for example, walking through a day of sustainable eating choices, or visualizing the impact of waste), the chatbot created a stronger emotional connection and self-awareness in users. This experiential engagement translated into users being more willing to change their behavior long-term (e.g., participants indicated a higher likelihood of reducing food waste after chatting). The study found that evoking personal experience and emotions is a key mechanism: users who felt more emotionally involved and reflective during the chatbot interaction showed the biggest positive change in their consumption attitudes.	Reinforces that **how information is delivered (tone & style) can drive behavior change** (RQ5 and RQ4). For our chatbot, incorporating storytelling or scenario-based dialogue could significantly enhance its impact. Simply providing facts about sustainable food is less effective than helping users emotionally experience the importance of those facts. This study's insight will guide us to design conversations that, for instance, ask the user to imagine future scenarios or reflect on their own habits, thereby not only raising awareness but also empowering users to make more sustainable food choices – the ultimate goal of our project (RQ6: improving effectiveness through engaging content).
Steybe et al. [26]	Evaluation of a context-aware chatbot using retrieval-augmented generation (RAG) for domain-specific Q&A	Improve chatbot accuracy and context awareness by augmenting a large language model with a domain knowledge base; evaluate the performance gain.	Implemented a hybrid chatbot ("GuideGPT") that combines GPT-4 with a curated library of 449 domain-specific documents (medical research on jaw osteonecrosis). Compared its answers to those of a base GPT-4 model on 30 specialized questions. Domain experts rated each answer on content correctness, explanation quality, and agreement with scientific consensus.	The **knowledge-augmented chatbot** dramatically outperformed the standalone LLM. GuideGPT's answers had significantly higher correctness and more detailed scientific explanations than the base model (experts' ratings were notably higher on content and accuracy measures, $p < 0.01$). By retrieving relevant up-to-date information, the chatbot avoided hallucinations and provided verified answers. Language fluency was similar for both, but only the RAG approach consistently produced factually reliable and precise responses.	Validates using an advanced **model architecture** to ensure responses are accurate and context-aware (RQ3). Integrating external knowledge sources (e.g., sustainability databases or articles) with an AI model can help our chatbot deliver up-to-date, factual information on future food sustainability. This study's success with a hybrid RAG model supports our approach to **adapt the chatbot with domain-specific content** (RQ4) for better reliability and depth when discussing sustainable food practices.
Yamamoto [27]	Suggestive answers strategy in human-chatbot interaction: A route to engaged critical decision making	Test a novel answer technique ("suggestive ending") designed to spark users' curiosity and prompt deeper engagement.	Developed three chatbot versions: (1) Plain (direct answers), (2) Expository (answers + brief pros/cons summary), (3) Suggestive (answers that end with an open hint or question). 300 participants were randomly assigned to use one version for decision-making tasks; user interactions and follow-up questions were measured.	The **suggestive chatbot** (answers ending with an intriguing hint) led users to ask more follow-up questions and spend more time exploring the topic than the other bots. Participants with the suggestive-ending bot engaged in longer, more in-depth dialogues and considered more diverse viewpoints before making decisions. This shows that a "cliffhanger" style response can effectively stimulate users' curiosity and encourage active information seeking, compared to giving complete information upfront.	Demonstrates an effective **conversational technique** to increase engagement (RQ2). We can leverage this by having our sustainability chatbot occasionally pose follow-up prompts or partial answers that invite the user to inquire further. This technique can combat user passivity and make interactions more interactive and exploratory, which is ideal for educating users about complex sustainability topics.
Yang et al. [15]	ChatDiet: Personalized Nutrition-Oriented Food Recommender Chatbots through an LLM-Augmented Framework	Introduce a framework that augments a large language model with personal and population nutrition models to deliver explainable, personalized food recommendations.	Architecture design and case study evaluation. ChatDiet integrates: a Personal Model (using causal inference on user-specific health/nutrition effects), a Population Model (with general nutrition info), and an Orchestrator to feed these into an LLM (GPT-based). The system generates recommendations with explanations. Evaluated via a case study and measured recommendation effectiveness (meeting the user's health goal) and quality of dialogues.	ChatDiet achieved **92% recommendation effectiveness** in a case study (i.e., the suggestions led to diets that met the target health outcomes). The chatbot provided clear, tailored justifications for its advice, combining general facts ("broccoli is high in fiber") with personal context ("…which aligns with your goal to improve gut health"). Users found the recommendations both helpful and trustworthy due to the explanations (anecdotal feedback). Demonstrated that combining user-specific data and global data results in more relevant and **explainable** interactions.	Supports the feasibility of an **AI-driven nutrition chatbot** with high personalization and user trust (relevant across RQ1, RQ2, RQ4). The importance of explanations to avoid the "black-box" effect aligns with **RQ3 (responsible AI)** – we aim to incorporate explanations to ensure transparency. ChatDiet's success in personalization and interactivity provides a model approach for our implementation, confirming that LLMs can be steered effectively with domain context for nutrition advice.

AI for Sustainable Food Education: Applying chatbots to sustainability and "future food" education is a relatively new but promising area. Several recent works highlight how digital tools, and AI can make sustainability concepts more relatable. Menkhoff and Gan[11] describe a climate action education initiative where students interacted with a conversational agent to learn about environmental issues. They note that an interactive chatbot format helped sustain engagement in a way traditional lectures did not, although the depth of knowledge transfer depended on the quality of the chatbot's responses. Nguyen[23] took a Value-Sensitive Design approach to environmental education chatbots, ensuring that the chatbot's content and interaction style supported key values like users' sense of connection to their community and personal well-being. This approach underlines the importance of aligning a sustainability chatbot with the values and identities of its audience to make conversations more meaningful [23]. In the domain of consumer behavior, Seo and Yoon [25] experimented with a chatbot aimed at mindful consumption and found that users who chatted with a bot embodying an "experiential mind" (guiding them through reflection on their consumption habits) showed more conscious decision-making afterwards. There is also evidence that chatbots can influence attitudes: one study reported that even a single conversation with an AI assistant on a contentious topic like climate change can measurably shift users' understanding towards factual information, even among those initially skeptical [19]. All these insights inform us of our project's design. They suggest that a well-crafted chatbot can indeed raise awareness and influence attitudes about sustainable food, provided it delivers trustworthy information, resonates with user values, and engages users in a conversational, personalized manner. However, no prior work has fully integrated personalized nutrition guidance with sustainability education in a single chatbot platform. Our research builds on the above literature to create a system that educates users about sustainable diets *and* provides individualized diet advice, thereby addressing both personal health and global sustainability goals.

3. Materials and Methods

Implementation and Technology Stack

We implemented the platform using a modern web development stack, prioritizing scalability and responsiveness. The front-end was developed in React (HTML, CSS, JavaScript), enabling a dynamic single-page application. This means users can navigate through different features (reading content, chatting, taking the survey) without full page reloads, creating a smooth experience. React's component-based architecture helped in building interactive UI elements, such as a step-by-step survey form that shows one question at a time (making the lengthy questionnaire feel more manageable to complete) and a chat interface that streams the AI's response text as it is generated. We also incorporated visualization and feedback elements on the front-end to boost engagement (e.g., a progress bar indicating how many survey questions remain, and a "daily tip" component that highlights a new sustainability tip each day to encourage users to return). The back-end was built with Flask in Python. We structured the back-end into separate modules corresponding to major functionalities: an authentication module (handling user sign-up, login, JWT token generation for session management), a survey module (endpoints to submit survey answers and retrieve personalized results), and a chat module (endpoint for processing a user's question and returning the chatbot answer). Using Flask's routing, each API endpoint was clearly defined (for example, **POST/api/survey** to submit responses, or **POST/api/chat** with a user's message). The back-end logic ensures that when a chat request comes in, the system will retrieve the user's profile and recent conversation context from the database, compile a prompt for the AI model, and post-process the AI's reply (e.g., storing it in the messages table, and filtering out any inappropriate content if needed before sending to the front-end). We utilized SQLAlchemy, an Object-Relational Mapper, to interact with the PostgreSQL database, which allowed us to work with Python classes for database entries and abstract away raw SQL queries. Throughout development, we adhered to security best practices: all user passwords are hashed, API calls are authenticated via tokens (for protected routes), and we implemented input validation to prevent injection attacks. Moreover, to align with ethical data use (as emphasized by Namkah [6] and others), we obtained explicit user consent for data collection via a privacy notice in the sign-up and survey interface, and we allow users to use the chatbot in a no-login, anonymous mode by default.

A crucial aspect of our implementation was the AI integration and evolution of the chatbot model. Initially, during early prototyping, we connected to the OpenAI API to leverage a powerful language model (GPT-3.5) for generating the chatbot's responses. This provided a strong starting point in terms of the chatbot's conversational ability and allowed rapid testing of our dialogue flows. However, relying on a third-party API had downsides: cost per API call, dependence on internet connectivity, potential data privacy concerns (as user queries would be sent to an external service), and limited control over the model's behavior or updates. Therefore, a key development milestone was transitioning to a self-hosted AI model. We selected Mistral-7B-Instruct (v0.1), an open-access large language model released by Mistral AI, as the base model for our chatbot. This model has 7 billion parameters and is designed to be lightweight enough for local hosting while still capable of understanding and generating fluent text. To tailor Mistral-7B to our domain, we fine-tuned it on a curated dataset of over 5,000 question–answer pairs focused on sustainable nutrition and future food topics. The fine-tuning was done using the Low-Rank Adaptation (LoRA) technique, which allows efficient training by injecting trainable rank-decomposition matrices into the model's layers (thus avoiding the need to update all 7B parameters). We conducted the fine-tuning on a local machine equipped with an NVIDIA RTX 5070 GPU (12GB VRAM), using the Hugging Face Transformers and PEFT libraries. The training data consisted of questions a user might ask (e.g., "How can I reduce food waste in my household?" or "Is a plant-based diet really more sustainable for the environment?") paired with detailed, reference-informed answers. We sourced these Q&As from authoritative resources (like FAO reports, scientific articles, and the knowledge base we compiled) to ensure factual accuracy. Training was run for 3 epochs with a 90/10 train-validation split; we used a SentencePiece tokenizer (Mistral's default) and capped input lengths at 512 tokens to fit memory constraints. The result of this process was a custom Mistral-7B-Instruct model variant that is specialized in the sustainability and nutrition domain and tuned to produce an educational, encouraging tone in line with our engagement goals.

After fine-tuning, we replaced the OpenAI API calls in our back-end with local inference calls to the new Mistral model. The model runs within a Flask-managed service, using an optimized FP16 (half-precision) version of its weights for faster inference. We integrated the model with the conversation context logic: when a user asks something, the back-end assembles a prompt that includes a brief summary of the conversation so far (from the **context_memory** table, if available), the user's question, and a system instruction reminding the model to use a friendly tone and incorporate relevant personal or sustainability info. The model then generates a response, which we can enrich by having it cite a source from the knowledge base if applicable (for transparency). This shift to a self-hosted model brought several benefits aligned with our project's values: it improved transparency and control (we know exactly what data the model was trained on and can fine-tune it further as needed), it eliminated external dependencies (enhancing privacy, since all computations happen on our servers), and it will likely reduce ongoing costs for scaling up the chatbot to more users.

Throughout development, we maintained a Software Architecture Document (SAD) that included the C1/C2 diagrams (**Figures 1** and **2**), detailed descriptions of each component, and the database schema (**Figure 3**). This documentation ensured that our team (including both student researchers and the supervisor) had a shared understanding of the system and helped in aligning implementation with the initial requirements. We followed agile principles in our workflow, implementing features in iterative cycles and continuously testing the system. For version control and collaboration, we used Git, which allowed us to track changes and revert if needed. We conducted regular code reviews and created unit tests for critical functions (e.g., verifying that the nutrient calculation for a recommended meal is accurate, or that the survey data mapping into user profiles works correctly). Towards the end of development, we performed integration testing to ensure the front-end, back-end, and model worked seamlessly together. We also tested the system with a set of sample user personas (fictional users with predefined survey answers and typical questions they might ask) to simulate real-world usage. This helped us identify and fix issues, such as the chatbot sometimes giving overly general answers – which we addressed by further refining the prompt and ensuring the knowledge base was consulted for specific queries.

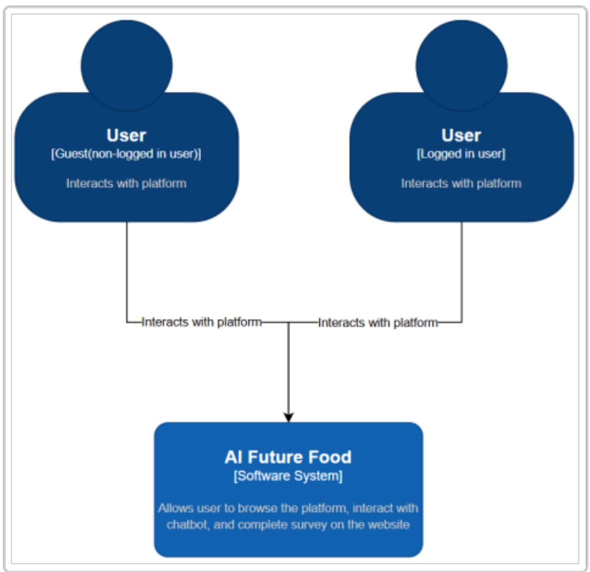

Figure 1. Context diagram of the AI chatbot platform for personalized sustainable nutrition. Both guest users (non-logged in) and registered users (logged in) interact with the AI Future Food system via a web interface. The platform allows users to browse educational content, engage in chat with the AI, and complete a nutrition survey. Regardless of user type, all interactions funnel into the AI Future Food software system, which provides personalized recommendations and sustainability tips.

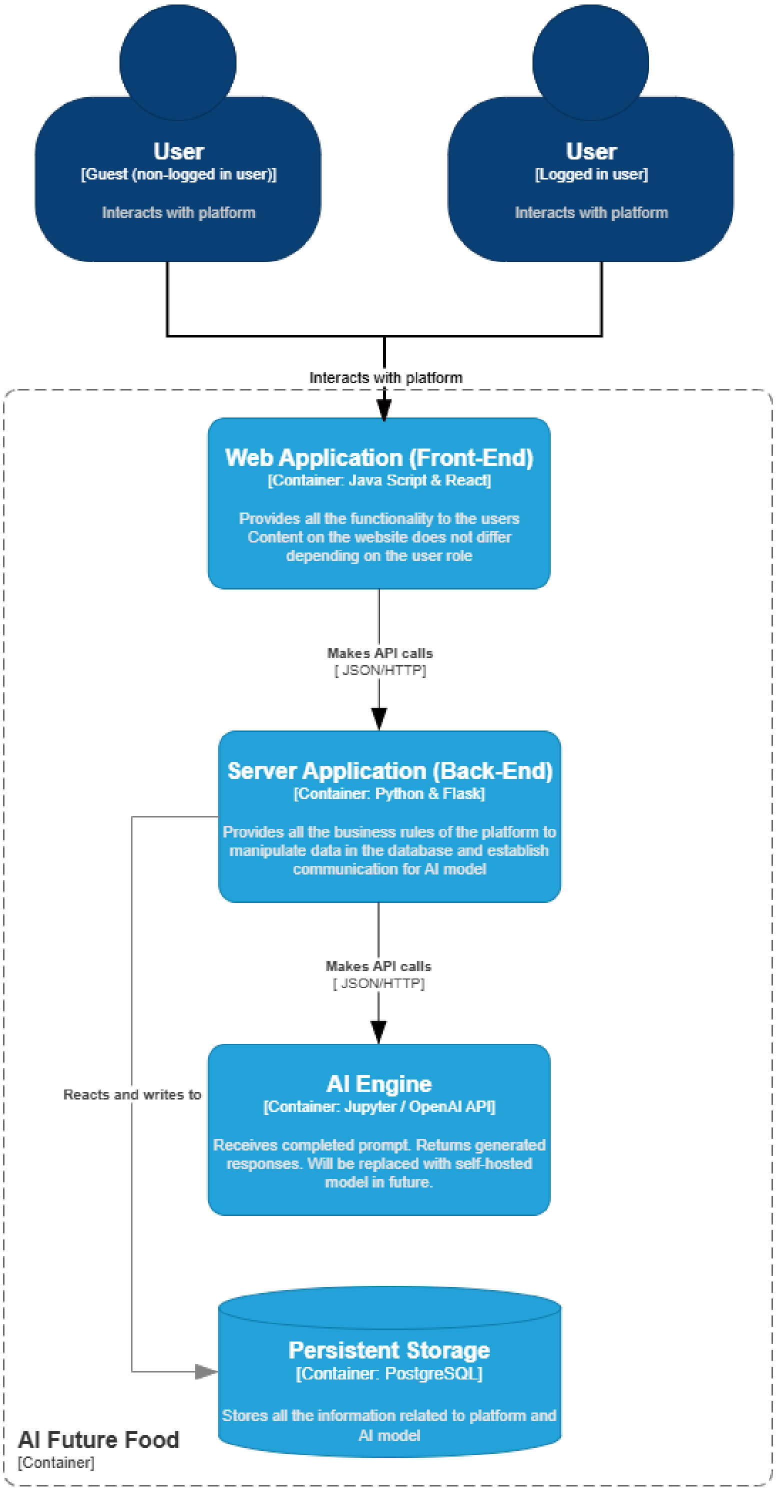

Figure 2. Container (C2) diagram of the system architecture, showing major components and technologies. The Web Application (Front-End) is built with JavaScript/React and provides the user interface. The Server Application (Back-End) is implemented in Python (Flask framework) and contains business logic, handling API calls from the front-end. It interacts with the AI Engine and the Persistent Storage. The AI Engine initially calls an external OpenAI API for language generation, and later is replaced by a self-hosted model (Mistral-7B). The Persistent Storage (PostgreSQL database) stores all platform data, including user info, chat logs, survey responses, and the knowledge base.

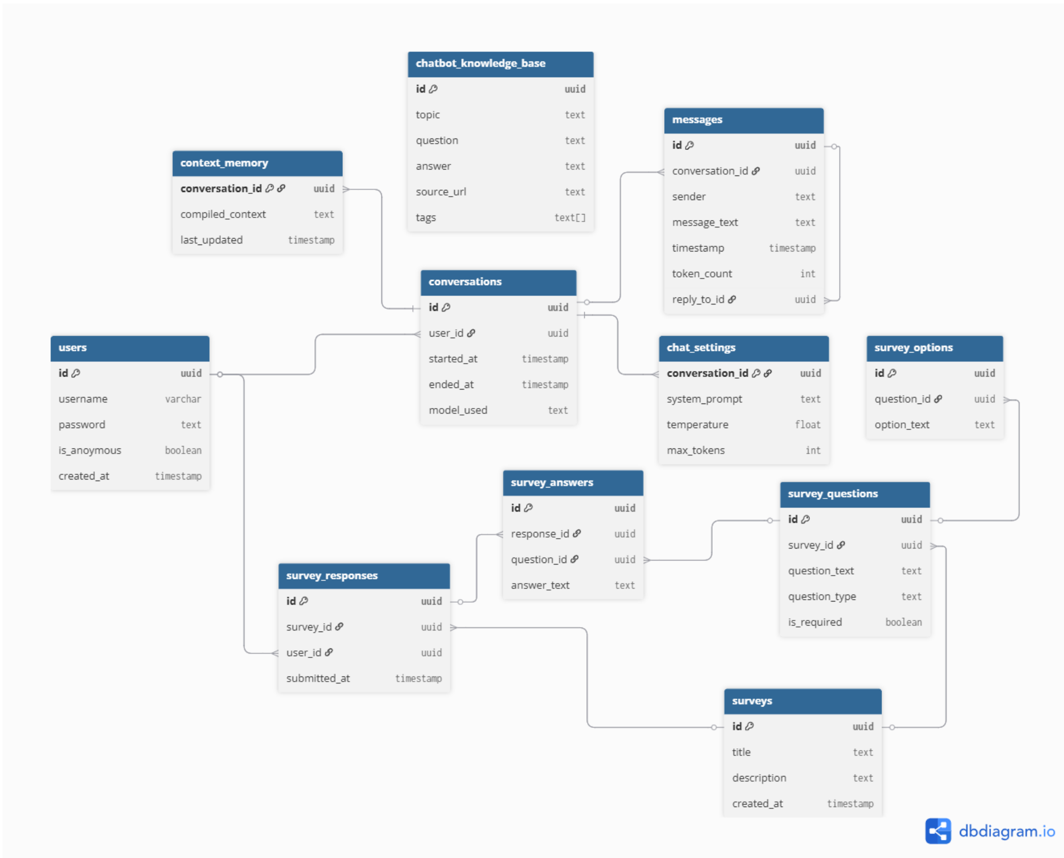

Figure 3. Normalized database schema for the chatbot platform (logical model). The schema includes tables for users (account info), conversations and messages (storing multi-turn chat logs with timestamps and token counts), chat_settings and context_memory (for managing the AI's conversation state and system prompts per session), a chatbot_knowledge_base (with curated Q&A entries on nutrition and sustainability, including sources and tags), as well as surveys, survey_questions, survey_options, survey_responses, and survey_answers (supporting a dynamic user questionnaire about dietary habits and preferences). All tables use unique IDs (UUIDs) for consistent linking and to ensure anonymity where needed.

4. Methodology

4.1. Overview and System Architecture

To address the research objectives, we followed a user-centered and iterative development process. We began by translating the goals (personalized nutrition guidance, sustainability education, user engagement, and ethical compliance) into concrete system requirements. These requirements drove the design of a web-based platform consisting of a front-end user interface, a back-end server with AI integration, and a database for storing content and user data. **Figures 1** and **2** illustrate the high-level architecture of the system, using C1 and C2 diagrams to show how users interact with the platform and how the internal components are organized.

As shown in **Figure 1**, the chatbot platform supports interaction from any user through a common web interface. Guests (anonymous users) can access most features without logging in, while registered users can create a profile to save their data and chat history. This design ensures accessibility (anyone can try the service) while enabling personalization for those who choose to sign up. In either case, the user interacts with the system by reading articles, asking the chatbot questions, or taking a survey about their diet. These interactions are handled by the platform's software components, which manage content delivery, AI response generation, and data logging. The context diagram emphasizes that the AI Future Food system is the central component that processes inputs from users and returns relevant, tailored outputs (whether it's a nutrition tip, a meal suggestion, or an answer to a sustainability question).

Figure 2 depicts the system's internal structure and the technology stack used for each part. The front-end is a single-page web application built with React (JavaScript), which renders the user interface components: for example, the nutrition survey form, pages displaying research content on sustainable foods, and the chatbot chat window. The front-end is responsible for providing an interactive and responsive user experience; it communicates with the back-end via HTTP requests (typically sending and receiving JSON data). The back-end is a Python application using the Flask framework (a lightweight web framework) to define RESTful API endpoints. The back-end implements the core logic of the platform: it receives survey submissions from the front-end and stores them, retrieves personalized content or previous messages from the database as needed, and routes user queries to the AI component. The AI Engine component is responsible for generating answers to user questions and providing smart recommendations. In the initial prototype, the AI Engine was implemented by making calls to the OpenAI API (specifically, using a GPT-3.5/4 model provided by OpenAI) to get language model completions. This allowed us to bootstrap the chatbot's conversational abilities quickly using a state-of-the-art model. The Persistent Storage is a PostgreSQL relational database that securely stores all data for the platform: user account information, survey responses, conversation logs, and a curated knowledge base of question-answer pairs on nutrition and sustainability topics. Overall, the architecture is modular: the front-end, back-end, AI engine, and database are distinct containers, communicating via well-defined interfaces. This modularity makes it easier to update or replace components (for example, swapping out the AI engine from an API-based model to a self-hosted model, as we later did) without affecting the others.

4.2. Database Schema and Knowledge Base

A normalized database schema was designed to support the platform's functionality and ensure data consistency. **Figure 3** shows the database schema, including tables for user data, chat history, survey questions/responses, and the chatbot's knowledge base.

The database design in **Figure 3** underpins both personalization and context management in the chatbot system. The users table stores basic account information (a username and hashed password for logged-in users) and a flag for anonymity for those who prefer not to tie data to an identity. When a user (guest or registered) starts a chat session, a new conversation entry is created, and each message—whether from the user or the AI—is stored in the messages table linked to that conversation. This allows the system to maintain context across multiple turns of conversation without needing to include the entire history in every prompt to the AI. Instead, relevant context can be fetched from the database when generating a new reply, improving efficiency. The context_memory and chat_settings tables store additional state for the chatbot: for instance, a system-level prompt (which sets the AI's role or persona) and the last compiled context summary to help the AI remember past interactions. We also integrated an extensive chatbot_knowledge_base: a table of frequently asked questions, answers, source URLs, and tags related to personalized nutrition and sustainable food. This knowledge base was manually compiled from credible sources (such as the Food and Agriculture Organization and EIT Food reports) and serves two purposes. First, it acts as a reference that the AI can draw upon for factual information (a step toward retrieval-augmented generation). Second, it provides content for the chatbot's fallback responses – if the AI model is unsure, the system can look up a relevant Q&A from this table to ensure accurate information is given. Finally, the schema includes tables for the survey functionality: the platform's nutrition questionnaire is modular, with the ability to define multiple surveys, each with a set of questions and predefined answer options. When a user fills out the survey, their responses are stored in survey_responses/survey_answers linked to their user ID (or an anonymous session ID). These survey results populate the user's dietary profile (e.g., whether they are vegetarian, their goals like weight loss or muscle gain, any allergies, etc.), which the chatbot uses to personalize its recommendations. By normalizing this schema and using foreign keys between tables (as shown in **Figure 3**), we maintained data integrity and minimized redundancy, ensuring that updates (like a new survey question or a corrected knowledge base answer) propagate consistently through the system.

4.3. Survey Form

To complement the database and knowledge base components, a structured survey form was developed to capture user demographics, dietary behaviors, cultural influences, perceptions of sustainable food, barriers and enablers, digital tool preferences, policy perspectives, and innovation priorities. The survey is organized into seven sections: (1) Demographics, (2) Cultural and Societal Dimensions, (3) Sustainable Food Perceptions, (4) Barriers and Enablers, (5) Digital Tools and Education, (6) Policy Development, and (7) Food Technology Innovation. Sample questions and their reasoning are summarized in **Table 2**.

Table 2. Sample survey questions used in the AI Future Food platform, including the section of the survey, the question text, available response options, and the purpose or reasoning for each question. The table presents how the survey captures key dimensions such as demographics, dietary habits, cultural influences, perceptions of sustainable food, barriers and enablers, digital tool preferences, policy perspectives, and food technology priorities. These questions were designed to provide both quantitative and qualitative insights, which inform the AI system's personalized recommendations and guide the evaluation of user engagement, sustainability awareness, and cultural alignment in food choices.

Section	Question	Options	Purpose/Reasoning
Demographics	What is your age group?	18–25, 26–35, 36–50, 51+	Understand demographic distribution
Demographics	What is your primary dietary preference?	Vegetarian, Vegan, Omnivore, Flexitarian, Other	Determine dietary habits and behaviors
Cultural & Societal Dimensions	How important is tradition in your food choices?	1 = Not important, 5 = Very important	Measure cultural influence on food choices
Sustainable Food Perceptions	Which sustainable food options have you tried?	Plant-based proteins, Cultivated meat, Insects, None, Other	Identify familiarity and adoption of sustainable options
Barriers & Enablers	What factors would make you consider sustainable food technologies?	Affordability, Accessibility, Cultural alignment, Other	Pinpoint enablers for adoption
Digital Tools & Education	Would you use a digital tool to guide sustainable food choices?	Yes, No, Unsure	Assess interest in technology for food education
Policy Development	How do you feel about policy incentives for sustainable food adoption?	Strongly support, Support, Neutral, Oppose, Strongly oppose	Understand public opinion on sustainability policies
Food Technology Innovation	What sensory qualities (taste, texture) matter most in food innovations?	Taste, Texture, Appearance, Smell, None	Identify user priorities in food innovation design

Response from the survey feeds directly to enabling personalized recommendations and contextualized sustainability guidance. To ensure data quality, the survey undergoes a validation strategy consisting of: Expert review for content validity, pilot testing with a representative user group for clarity and usability, and iterative refinement based on observed response patterns. This approach ensures that the survey generates reliable, actionable data to support AI-driven recommendations and broader analyses of dietary and sustainability behaviors.

4.4. Ethical Considerations

Ethical design principles were embedded throughout the methodology. Privacy: The system does not require any personally identifiable information from users – even registered accounts only use a username and password, and users can choose any pseudonym. All users are internally identified by a random UUID. The nutrition survey, while collecting personal diet information, is optional and can be skipped or taken anonymously; its data is stored securely in the database. We followed GDPR-aligned practices, providing users with information on how their data would be used (solely to personalize their experience and not shared externally) and the ability to delete their account (which removes or anonymizes their data). Bias and Fairness: We recognized the risk of bias in AI diet advice (for example, gender or cultural biases in food suggestions as noted by Kaçar[21]. To mitigate this, we curated a diverse knowledge base covering various dietary patterns and cultural cuisines, and we tested the chatbot with profiles of different backgrounds (vegetarian vs. omnivore, Western diet vs. Asian diet, etc.) to ensure it adapts appropriately. The chatbot's model and prompts were tuned to avoid making assumptions about the user's background and to ask for clarification when needed (e.g., if a user asks for dinner ideas, the bot might ask if there are any dietary preferences or restrictions it should consider). Transparency: To build trust, we designed the chatbot to be somewhat self-explanatory – it often includes brief explanations or references in its answers. For instance, if the chatbot suggests eating more legumes for sustainability reasons, it might add "because legumes generally have a lower carbon footprint than animal proteins, according to research." We have also made the underlying sources retrievable (the plan for future iterations is to allow users to click a "Why this suggestion?" button to see source details from the knowledge base). These measures address the ethical concerns raised in prior work about AI in nutrition needing to be transparent and to obtain informed consent [6]. Finally, by moving to an open-source AI model, we also contribute to ethical AI practice: our chatbot can be audited and improved by others in the community, and it avoids the proprietary black-box issue of commercial AI services. In summary, the methodology from requirement analysis and design through implementation was guided by a focus on personalization (to make the system effective for each user), engagement (to make it enjoyable and motivating to use), and ethics (to ensure it is responsible and trustworthy).

4.5. Validation Strategy

The validation of the AI-driven nutrition model focused on assessing the reliability, fairness, and transparency of its recommendations within the scope of European dietary contexts. To achieve this, the model was evaluated across controlled scenarios that reflected diverse dietary preferences and regional traditions. Identical queries were tested under varying profiles in order to determine whether the system generated consistent yet context-sensitive outputs. For example, when prompted with a generic request such as "suggest a sustainable dinner," the model was expected to adapt recommendations to different European food cultures, such as Mediterranean or Northern European diets, while ensuring that vegetarian and vegan users received options fully aligned with their dietary restrictions. Importantly, rather than inferring these attributes from implicit cues, the chatbot was designed to elicit them directly through clarification questions (e.g., asking about dietary preferences, cultural influences, or sustainability priorities during the interaction). This self-learning mechanism enabled the system to tailor outputs dynamically while minimizing the risk of biased assumptions about the user.

The system's ability to avoid bias was further examined in scenarios involving potentially sensitive attributes. Gender neutrality was tested by comparing recommendations for identical dietary goals, such as maintaining a balanced diet or reducing environmental impact, across differently gendered profiles. Outputs were analyzed to ensure that the nutritional advice remained consistent and evidence-based, without introducing stereotypical assumptions about food preferences. Such comparative testing provided a systematic means of identifying hidden biases in the model's behavior.

Transparency and user trust were also key elements of validation. Pilot users were asked to evaluate the clarity and perceived reliability of the chatbot's outputs during structured interaction sessions. Feedback from these sessions was used to refine both the model prompts and the explanatory mechanisms, ensuring that recommendations were accompanied by justifications grounded in nutritional science and sustainability research.

Taken together, the validation strategy demonstrated that the AI system was capable of generating accurate and personalized recommendations within the domain of European food traditions, while addressing potential risks of dietary, cultural, or gender-based bias. By combining controlled profile testing with dynamic user clarification and iterative refinement, the model was aligned with both technical performance standards and ethical requirements for fairness, inclusivity, and transparency.

5. Results

After developing the chatbot platform, we conducted preliminary evaluations to assess its effectiveness in engaging users and educating them about sustainable nutrition. Since a large-scale user study was beyond our project's scope at this stage, we performed an initial round of peer testing and gathered both quantitative metrics and qualitative feedback. We invited a small group of test users (colleagues and fellow students, n = 10) to try the chatbot in a realistic scenario: each participant created an account (or used the guest mode), completed the nutrition survey, and then engaged in an open conversation with the chatbot, asking questions or advice related to diet and sustainability. After the interaction, participants answered a short questionnaire about their experience and participated in an informal interview to elaborate on their feedback.

User Engagement and Interaction Metrics: The chatbot was well-received in terms of engagement. On average, users spent about 15 minutes interacting with the system in a session, during which they asked the chatbot 8–12 questions each. Many users explored multiple features, for example, reading through the educational content sections on future foods and then asking follow-up questions to the chatbot. A key indicator of engagement was that all participants completed the survey component fully, despite it containing over 20 questions; several remarked that the survey's conversational, step-by-step design made it feel less tedious than a traditional form. We also tracked the length of conversations and the diversity of questions asked. Users in this test posed a wide range of queries, from practical tips ("How can I reduce food waste at home?") to more factual questions ("What is the environmental impact of eating beef vs. chicken?") and personal guidance ("I'm trying to eat healthier on a budget – any suggestions?"). The chatbot was generally able to handle this range, maintaining context within each conversation thread thanks to the context memory mechanism. There were a few instances where the bot needed a prompt (for example, a user asked a very broad question and the bot responded with a clarification question, which the user appreciated rather than finding it a failure). No major technical issues (like crashes or very long delays) were reported; the average response time from the AI for a medium-length answer was around 3 seconds, which was acceptable to users. These engagement metrics are promising, suggesting the platform can hold user interest and facilitate sustained interactions.

Knowledge Gain and Attitudinal Shifts: To evaluate learning outcomes, we asked participants to rate their knowledge of sustainable food topics before and after using the chatbot, and whether the experience changed their attitudes or intentions regarding their diet. Although this was a self-reported measure, the results were encouraging. A majority of users (7 out of 10) indicated that they learned at least one new fact or concept during the chat. For example, one user noted, "I learned about how much water is used to produce a pound of beef – I had no idea it was that high, and the chatbot's explanation put it into perspective." Several participants mentioned that the chatbot introduced them to the idea of evaluating foods by their carbon footprint or environmental impact, which was a new way of thinking for them. In terms of attitudes, a few users reported feeling more motivated to make changes: one user said the conversation "made me want to try having a meatless day once a week," and another stated, "I'm now thinking of checking the labels for sustainability or sourcing, which I didn't really consider before." This aligns with findings by Nguyen[24] that a well-designed chatbot interaction can improve pro-environment attitudes and even increase willingness to take action. While our test was not a controlled trial, the anecdotal evidence points to the chatbot's potential to raise awareness and influence intentions. Notably, one participant who was initially skeptical about sustainable diets (viewing them as fads) commented after using the chatbot that "it gave pretty balanced answers and didn't feel preachy, which actually made me more open to its suggestions." This reflects what Chen[19] observed in their study on conversational agents: even users with some resistance to a topic can shift their understanding when the information is delivered in a conversational and personalized manner.

Qualitative Feedback and User Satisfaction: We collected open-ended feedback about what users liked and disliked. Overall satisfaction was high; most participants found the chatbot useful and enjoyable. They appreciated the personalized angle: for instance, the chatbot would reference the user's survey inputs when giving advice ("Since you mentioned you rarely eat fish, how about trying more plant-based proteins like lentils to improve your diet's sustainability?"). Users reported that this made the advice feel "tailored just for me" and not generic. The interactive Q&A format was also highlighted as a strength. One user compared it to searching the web: "If I Google something about diet, I have to sift through articles. Here, I just asked and got a clear answer with context. It even offered to explain more if I wanted." This suggests the chatbot effectively served as an educational tool, packaging information in an accessible dialogue form. We also asked if they would use such a chatbot regularly; several said yes, especially if they were trying to stick to a new diet or learn new recipes, noting it could be like "having a dietitian on hand to consult, but also one that cares about the planet."

Critiques and suggestions were valuable for identifying areas of improvement. A few participants wanted deeper answers to technical questions. For example, one user who asked a very detailed question about protein digestibility received a correct but surface-level answer from the chatbot. They felt the bot could either provide more detail or direct them to a resource for such advanced queries. This feedback resonates with the literature that stresses the importance of a robust knowledge base for complex topics[11]. In response, we plan to expand the knowledge base and perhaps integrate an external API for nutrition data to strengthen the chatbot's ability to handle highly technical questions. Some users also suggested incorporating visual content or references in the chat. While our chatbot did sometimes cite sources or facts, it currently doesn't show images or charts. One user said, "It would be cool if when I ask about the seasonality of veggies, it could show a quick chart or something." Another mentioned links: "If I ask for recipes, maybe it can give me a link to a sustainable recipe site." This points to a potential enhancement of adding multimedia or hyperlink support in the chat responses to enrich the user experience. As a result of this feedback, we have included in our development roadmap the addition of optional infographics or links for certain answers (for instance, a question on carbon footprints might return a small chart of CO_2 emissions by food category, along with the text).

Importantly, no major issues were reported regarding the chatbot's tone or behavior. Users generally found it polite, supportive, and informative rather than judgmental or pushy. This was a conscious goal in our design to ensure the chatbot remains a friendly guide. One participant described the bot's personality as "like a knowledgeable friend who gives you tips." Maintaining this positive user perception is crucial because, as Chen[19] noted, a user might learn something from a chatbot but still feel frustrated if the style is off-putting. In our test, participants' satisfaction levels were high, with most giving an overall experience rating of 8 or 9 out of 10. This suggests that our focus on an engaging tone and ethical design paid off in user acceptance.

Performance of the Fine-Tuned Model: Although not directly noted by users, we evaluated how the fine-tuned Mistral-7B model performed behind the scenes. We compared a sample of chatbot responses (for a set of standardized questions) between the earlier OpenAI-GPT version and the new local model. The fine-tuned model's answers were found to be on par in terms of relevance and correctness for common questions, and it had the advantage of injecting the sustainability context more readily. For example, when asked "What's a healthy dinner I can cook tonight?", the OpenAI-based bot gave a generic healthy meal suggestion, whereas the Mistral-based bot suggested a specific meal ("Grilled tofu with quinoa and mixed vegetables"), and added a note about it being environmentally friendly by using plant-based protein. This indicates that fine-tuning successfully imbued the model with the project's dual focus. On the flip side, the local model did have slightly shorter answers on average and was a bit more repetitive in some cases. These are areas we will continue to refine (possibly by further training or adjusting decoding parameters). From a performance standpoint, running the model locally did increase response latency slightly (by about 0.5–1 second compared to the API), but as mentioned, it was still within an acceptable range for users.

In summary, the results of our preliminary evaluation are promising. The chatbot engaged users in meaningful conversations, imparted knowledge about sustainable nutrition, and was generally well-received as a useful tool. Users' comments affirmed the value of combining personalized diet advice with sustainability education in one platform. They also provided insight into how we can further improve the system, namely, by enriching content depth and variety (including possibly multimedia answers) and ensuring the knowledge base covers advanced queries. These results, albeit from a small sample, suggest that our approach can be effective, and they set the stage for more extensive testing in the future.

6. Discussion

This project set out to bridge the gap between personalized nutrition guidance and sustainability education through an AI chatbot. The positive initial results and user feedback indicate that such a bridge is not only feasible but also welcome. In designing the chatbot, we drew on insights from prior research and continuously aligned our development with the twin goals of user engagement and ethical, informed content. The outcome is a functional prototype that provides a proof-of-concept for leveraging large language models in the service of healthier and more sustainable eating habits.

One of the key achievements of our work is the successful integration of a fine-tuned open-source LLM (Mistral-7B) as the engine of the chatbot. This move towards a self-hosted model was deliberate, aiming to enhance transparency, controllability, and ethical autonomy compared to relying on a proprietary API. By using LoRA fine-tuning on a domain-specific dataset (sustainable food Q&A), we created an AI agent that is tailored to both the content domain and the desired conversational tone. This addresses a common gap identified in earlier chatbot studies, where either the tone was too generic or the domain specificity was lacking. Our model demonstrates that with relatively modest resources (a 7B parameter model and a few thousand training examples), one can achieve a specialized conversational agent that operates under full developer control. The benefits are manifold: we can ensure the model does not drift in behavior due to external updates, we can audit and improve the training data, and user interactions stay private within our system. These factors are particularly important in the context of dietary advice, which can be personal and sensitive; using an open model adds a layer of trustworthiness in the eyes of informed users and aligns with recommendations for transparency in personalized nutrition AI[6]. Moreover, the shift reduced our dependency on third-party services and costs, which improves the platform's sustainability and scalability from a development perspective.

The chatbot's architecture and design were heavily informed by a user-centered approach, and this is reflected in the user feedback. Early in the project, we recognized that simply having good content was not enough-the way the content is delivered (through conversation, with personalization) is crucial for engagement. By studying common pain points in human-chatbot interactions (as reported by Kuhail[7] and others), we preemptively designed features to address them. For instance, the incorporation of context memory and a structured conversation schema means the bot can handle follow-up questions smoothly, which was noted by users who tried to "stump" it or ask for clarification. Additionally, strategies like using an empathetic tone and referencing user-provided details (e.g., "since you indicated in the survey that you avoid dairy…") helped make the interaction feel personal. This personalization likely contributed to the motivational aspect some users experienced, similar to how the participants in Franco[13] maintained better habits with personalized app feedback. Our findings align with the broader evidence that personalization and conversational nuance foster greater trust and engagement[10]. Users are more willing to accept advice and consider new information when it is presented in a way that resonates with their individual context and values.

Ethical design choices in the system not only averted problems but became a selling point of the platform. By emphasizing anonymity and data privacy, we addressed head-on the common barrier of users' mistrust in AI systems[10]. Test participants, even if not explicitly checking for GDPR compliance, implicitly felt more comfortable knowing they weren't giving away personal identifiers and that they could use the tool without an account. This reinforces the idea that ethical features (like privacy-by-design) can enhance user adoption rather than hinder it. Additionally, the content of the chatbot benefited from an ethical perspective: because we were conscious of potential biases (cultural, gender, etc.) in nutritional advice, the chatbot's recommendations tended to be more inclusive and balanced. For example, it often provided alternatives ("if X is not accessible or against your diet, you can try Y") which users from different backgrounds found considerate. This kind of inclusivity is important in nutrition advice since diet is deeply personal and tied to culture. It echoes the principles put forth by Namkah[6] that personalization must also respect diversity and fairness. Our approach demonstrates how integrating ethics from the ground up – in data collection, system rules, and AI training – results in a more robust and trustworthy tool.

Despite the encouraging outcomes, it's important to acknowledge the limitations of this study and the prototype. Firstly, the evaluation so far has been limited in scale and scope. While peer testing provided useful insights, a larger and more diverse user study is needed to rigorously assess the chatbot's impact on knowledge and behavior. In the future, we plan to conduct controlled experiments, such as having one group use the chatbot while a control group uses conventional resources, to measure differences in knowledge gain and dietary choices. We also aim to track longer-term engagement: Will users continue to interact with the chatbot over weeks or months? Sustained behavior change is the ultimate goal, and that requires longitudinal observation. Secondly, the training dataset for the model, though curated, was moderate in size. A 7B model can typically absorb much more data; with only ~5k Q&A pairs, there is a risk that some niche topics are not covered and that the model might not generalize perfectly to all possible questions. We already noticed a slight dip in performance for highly technical questions. Expanding the training data (for example, incorporating more scientific Q&As or using reinforcement learning from human feedback to fine-tune conversational quality) could further improve the chatbot's reliability and depth. Third, our current platform supports English language content only and assumes a relatively informed user base (the kind likely to be concerned about diet and environment). To truly maximize impact, we would need to implement multilingual support and perhaps simplify some explanations for users with less prior knowledge. However, thanks to the system's modular design and local model, adding a new language model or dataset is technically feasible down the line.

7. Future Work

While this study focused on chatbot-driven sustainable nutrition guidance, future work will expand toward integrating a calorie estimation model capable of analyzing European dishes and suggesting ingredient substitutions to lower both caloric intake and environmental impact. Previous research has demonstrated that food recognition models can estimate calories from eating activities with promising accuracy [5], providing a strong technical foundation for this extension. Moreover, evidence shows that providing clear guidelines significantly improves consumers' ability to make environmentally sustainable food choices, particularly for protein products and fruits [28]. Building on these insights, our planned system will not only estimate calories but also recommend sustainable ingredient alternatives (e.g., replacing red meat with plant-based proteins), thereby combining health-focused calorie awareness with sustainability considerations. This integration aims to empower users to make informed dietary choices that balance personal nutrition goals with broader environmental responsibility.

Another area for enhancement is integrating more real-time data and adaptivity. For example, the chatbot could adjust recommendations based on seasonal food availability or trending topics (like a new study on nutrition that people are curious about). Currently, the system's knowledge is as updated as the curated sources; implementing a pipeline for regular updates or a retrieval system querying live databases (while still filtering for accuracy) could keep the content fresh. This touches on a broader point: the modular architecture we built is well-suited to continuous improvement. Since the front-end and database are in place, one could plug in a more advanced AI component or an additional microservice (such as a recommendation engine or a gamification module) without overhauling the whole system. This flexibility means the project can evolve with emerging best practices in both AI and nutritional science. For instance, if future research indicates a new effective way to nudge users (say, sending personalized reminders or challenges), our platform could integrate that relatively easily.

The broader implications of this work extend beyond just nutrition. We've essentially created a template for how AI chatbots can be deployed to drive sustainable behavior change in various domains. The focus here was food, but the architecture and approach could be adapted to areas like energy conservation (imagine a chatbot that looks at your home energy use data and chats with you about saving electricity) or waste reduction (a chatbot that helps you recycle correctly and reduce trash). The success of an interdisciplinary project like this – combining AI technology with nutrition and environmental science – highlights the value of cross-domain collaboration. We leveraged knowledge from computer science (NLP, software engineering), health sciences (dietary guidelines, behavior change), and sustainability (environmental impact of foods) to create a solution that speaks to all those aspects. This kind of holistic thinking is increasingly important as we design AI systems intended to influence human behavior for social good.

Finally, it is worth reflecting on the user's role and agency in using such a chatbot. Our aim was not to have an AI that dictates what a person should do, but rather one that empowers users with information and support. The feedback that users felt "informed and motivated" rather than "told" is a positive sign. It suggests we struck a good balance in the AI's tone and content between guiding and giving agency. Moving forward, maintaining that balance will be key. Features like providing options, asking the user's preferences, and encouraging reflection (e.g., "how do you feel about trying this?") can ensure the chatbot remains a coach or assistant, not an authoritarian figure. This approach respects the user's autonomy and is likely more effective in achieving lasting behavior change, as people are more likely to follow through with actions they decided on themselves with gentle guidance.

In conclusion, the discussion of our findings affirms that AI chatbots, when thoughtfully designed, can be powerful allies in promoting personalized, sustainable nutrition. They serve as an interactive medium to deliver tailored advice and education, potentially filling gaps that neither traditional apps nor human experts alone can fully address at scale. Our project contributes to this emerging area by demonstrating a working example that integrates advanced AI with principles of engagement and ethics. The lessons learned here can inform future endeavors aiming to harness AI for sustainability and health, ensuring that as technology guides us towards better choices, it does so in a way that is respectful, transparent, and centered on human well-being.

8. Conclusion

This study developed and evaluated an AI-driven chatbot that combines personalized nutrition guidance with sustainable food education. The chatbot was implemented using a fine-tuned, locally hosted Mistral-7B-Instruct model, supported by a modular system architecture and a sustainability-focused knowledge base. Preliminary evaluations with users showed encouraging results. Participants reported increased awareness of how daily food choices affect both health and the environment, and they found the chatbot's recommendations, such as reducing food waste and introducing more plant-based meals, practical and actionable. Importantly, users highlighted the conversational and personalized nature of the chatbot as a key factor in keeping them engaged and motivated to adopt sustainable dietary habits. These findings suggest that conversational AI can effectively complement traditional nutrition education by not only providing knowledge but also supporting behavior change.

The study also emphasizes the importance of ethical implementation. By ensuring transparency, protecting user privacy, and mitigating bias, the chatbot demonstrated that responsible AI design can go hand in hand with user engagement and functionality. In summary, the data from this pilot evaluation indicates that AI chatbots hold significant potential to enhance public awareness of sustainable nutrition. Future work should focus on scaling the user study to larger and more diverse populations, expanding the knowledge base, and integrating additional features such as calorie estimation with sustainable ingredient substitutions. By building on these results, the chatbot can evolve into a valuable tool for fostering healthier and more environmentally conscious food choices.

Author Contributions

Conceptualization, H.N.M.T., A.K., S.S., and P.D.; methodology, H.N.M.T., A.K., S.S., and P.D.; software, A.K. and S.S.; validation, H.N.M.T., A.K., and S.S.; formal analysis, H.N.M.T., A.K., S.S., and P.D.; investigation, H.N.M.T., A.K., and S.S.; resources, P.D.; data curation, A.K. and S.S.; writing—original draft preparation, H.N.M.T., A.K., and S.S.; writing—review and editing, H.N.M.T.; visualization, P.D. and H.N.M.T.; supervision, P.D.; project administration, P.D.; funding acquisition, P.D. All authors have read and agreed to the published version of the manuscript.

Funding

This research was funded by the author Priyanka Darbari. The funder was involved in the conceptualization, methodology, supervision, and project administration of the study, as detailed in the Author Contributions section.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Data Availability Statement

Not applicable.

Conflicts of Interest

The authors declare that there is no conflict of interest.

References

[1] Poore, J., Nemecek, T., 2018. Reducing Food's Environmental Impacts through Producers and Consumers. Science. 360(6392), 987–992. DOI: https://doi.org/10.1126/science.aaq0216

[2] Van Dijk, M., Morley, T., Rau, M.L., et al., 2021. A Meta-Analysis of Projected Global Food Demand and Population at Risk of Hunger for the Period 2010–2050. Nature Food. 2, 494–501. DOI: https://doi.org/10.1038/s43016-021-00322-9

[3] Fadhil, A., Gabrielli, S., 2017. Addressing Challenges in Promoting Healthy Lifestyles: The AI-Chatbot Approach. Proceedings of the 11th EAI International Conference on Pervasive Computing Technologies for Healthcare. DOI: https://doi.org/10.1145/3154862.31

54914

[4] Dakhia, Z., Russo, M., Merenda, M., 2025. AI-Enabled IoT for Food Computing: Challenges, Opportunities, and Future Directions. Sensors. 25(7), 2147. DOI: https://doi.org/10.3390/s25072147

[5] Jaiswal, A., Solanki, A., Tomar, A., et al., 2025. Revolutionizing Diet and Fitness Tracking with AI: A User-Centric Approach to Nutrition and Wellness. 2025 2nd International Conference on Computational Intelligence, Communication Technology and Networking (CICTN). DOI: https://doi.org/10.1109/CICTN64563.2025.10932496

[6] Namkhah, Z., Fatemi, S.F., Mansoori, A., et al., 2023. Advancing Sustainability in the Food and Nutrition System: A Review of Artificial Intelligence Applications. Frontiers in Nutrition. 10, 1295241. DOI: https://doi.org/10.3389/fnut.2023.1295241

[7] Kuhail, M.A., Alturki, N., Alramlawi, S., et al., 2023. Interacting with Educational Chatbots: A Systematic Review. Education and Information Technologies. 28(1), 973–1018. DOI: https://doi.org/10.1007/s10639-022-11177-3

[8] Deng, X., Yu, Z., 2023. A Meta-Analysis and Systematic Review of the Effect of Chatbot Technology on Educational Outcomes. Sustainability. 15(4), 2940. DOI: https://doi.org/10.3390/su15042940

[9] Juquelier, A., Poncin, I., Hazée, S., 2025. Empathic Chatbots: A Double-Edged Sword in Customer Experiences. Journal of Business Research. 188, 115074. DOI: https://doi.org/10.1016/j.jbusres.2024.115074

[10] Kazoun, N., Kokkinaki, A., Chedrawi, C., 2025. AI Chatbots for Sustainability in Education: The Case of the Lebanese Higher Education Sector. In: Themistocleous, M., Bakas, N., Kokosalakis, G., et al. (eds.). Proceedings of the European, Mediterranean, and Middle Eastern Conference on Information Systems (EMCIS 2024). Springer: Berlin, Germany. pp. 19–32.

[11] Menkhoff, T., Gan, B., 2023. Engaging Students through Conversational Chatbots and Digital Content: A Climate Action Perspective. Proceedings of the 9th International Conference on Human Interaction and Emerging Technologies. DOI: https://doi.org/10.54941/ahfe1002960

[12] Hinojosa-Nogueira, D., Ortiz-Viso, B., Navajas-Porras, B., 2023. Stance4Health Nutritional App: A Path to Personalized Smart Nutrition. Nutrients. 15(2), 276. DOI: https://doi.org/10.3390/nu15020276

[13] Franco, R.Z., Fallaize, R., Weech, M., 2022. Effectiveness of Web-Based Personalized Nutrition Advice for Adults Using the eNutri Web App: Evidence from the EatWellUK Randomized Controlled Trial. Journal of Medical Internet Research. 24(4), e29088. DOI: https://doi.org/10.2196/29088

[14] Papastratis, I., Konstantinidis, D., Daras, P. et al., 2024. AI Nutrition Recommendation Using a Deep Generative Model and ChatGPT. Scientific Reports. 14(1), 14620. DOI: https://doi.org/10.1038/s41598-024-65438-x

[15] Yang, L., Xu, J., Lee, M., 2024. ChatDiet: Empowering Personalized Nutrition-Oriented Food Recommender Chatbots through an LLM-Augmented Framework. IEEE Access. 12, 22233–22245. DOI: https://doi.org/10.1016/j.smhl.2024.100465

[16] Aggarwal, A., Tam, C.C., Wu, D., 2023. Artificial Intelligence–Based Chatbots for Promoting Health Behavioral Changes: Systematic Review. Journal of Medical Internet Research. 25, e40789. DOI: https://doi.org/10.2196/40789

[17] Capecchi, I., Borghini, T., Bellotti, M., et al., 2025. Enhancing Education Outcomes Integrating Augmented Reality and Artificial Intelligence for Education in Nutrition and Food Sustainability. Sustainability. 17, 2113. DOI: https://doi.org/10.3390/su17052113

[18] Li, X., Yin, A., Choi, H.Y., et al., 2024. Evaluating the Quality and Comparative Validity of Manual Food Logging and Artificial Intelligence-Enabled Food Image Recognition in Apps for Nutrition Care. Nutrients. 16(15), 2573. DOI: https://doi.org/10.3390/nu16152573

[19] Chen, K., Shao, A., Burapacheep, J., et al., 2024. Conversational AI and Equity: Assessing GPT-3's Communication with Diverse Social Groups on Contentious Topics. Scientific Reports. 14(1), 1561. DOI: https://doi.org/10.1038/s41598-024-51969-w

[20] Izadi, S., Forouzanfar, M., 2024. Error Correction and Adaptation in Conversational AI: A Review of Techniques and Applications in Chatbots. AI. 5(2), 803–841. DOI: https://doi.org/10.3390/ai5020041

[21] Kaçar, M.D., Yücel, G., Koç, M., 2024. Diet Quality and Caloric Accuracy in AI-Generated Diet Plans: A Comparative Study Across Chatbots. Nutrients. 16(3), 662. DOI: https://doi.org/10.3390/nu17020206

[22] Kunja, S.R., Gade, M., Mohammed, A., 2023. Engaging Guests for a Greener Tomorrow: Examining the Role of Hotel Chatbot Concierges on Sustainable Practices. Tourism and Hospitality Research. Tourism and Hospitality Research. DOI: https://doi.org/10.1177/14673584241313339

[23] Nguyen, H., Nguyen, V., Ludovise, S., et al., 2025. Value-Sensitive Design of Chatbots in Environmental Education: Supporting Identity, Connectedness, Well-Being, and Sustainability. British Journal of Educational Technology. 56(4), 1370–1390. DOI: https://doi.org/10.1111/bjet.13568

[24] Nguyen, T.K.C., 2024. The Effect of AI Chatbots on Pro-Environment Attitude and Willingness to Pay for Environment Protection. SAGE Open. 14(1), 1–15. DOI: https://doi.org/10.1177/21582440231226001

[25] Seo, J.K., Yoon, H.J., 2025. Promoting Mindful Consumption through a Chatbot with an Experiential Mind.

Journal of Consumer Marketing. 42(4), 498–511. Advance online publication. DOI: https://doi.org/10.1108/JCM-05-2024-6844

[26] Steybe, D., Poxleitner, P., Aljohani, S., et al., 2025. Evaluation of a Context-Aware Chatbot Using Retrieval-Augmented Generation for Answering Clinical Questions on Medication-Related Osteonecrosis of the Jaw. Journal of Cranio-Maxillofacial Surgery. 53(4), 355–360. DOI: https://doi.org/10.1016/j.jcms.2024.12.009

[27] Yamamoto, Y., 2024. Suggestive Answers Strategy in Human-Chatbot Interaction: A Route to Engaged Critical Decision Making. Frontiers in Psychology. 15, 1382234. DOI: https://doi.org/10.3389/fpsyg.2024.1382234

[28] Lazzarini, G.A., Visschers, V.H.M., Siegrist, M., 2018. How to Improve Consumers' Environmental Sustainability Judgements of Foods. Journal of Cleaner Production. 198, 564–574. DOI: https://doi.org/10.1016/j.jclepro.2018.07.033