

Foodbot: A Goal-Oriented Just-in-Time Healthy Eating Interventions Chatbot

Philips Kokoh Prasetyo
pprasetyo@smu.edu.sg
Singapore Management University

Palakorn Achananuparp
palakorna@smu.edu.sg
Singapore Management University

Ee-Peng Lim
eplim@smu.edu.sg
Singapore Management University

ABSTRACT

Recent research has identified a few design flaws in popular mobile health (mHealth) applications for promoting healthy eating lifestyle, such as mobile food journals. These include tediousness of manual food logging, inadequate food database coverage, and a lack of healthy dietary goal setting. To address these issues, we present Foodbot, a chatbot-based mHealth application for goal-oriented just-in-time (JIT) healthy eating interventions. Powered by a large-scale food knowledge graph, Foodbot utilizes automatic speech recognition and mobile messaging interface to record food intake. Moreover, Foodbot allows users to set goals and guides their behavior toward the goals via JIT notification prompts, interactive dialogues, and personalized recommendation. Altogether, the Foodbot framework demonstrates the use of open-source data, tools, and platforms to build a practical mHealth solution for supporting healthy eating lifestyle in the general population.

CCS CONCEPTS

• Human-centered computing → Personal digital assistants.

KEYWORDS

mHealth, chatbot, diet, self-tracking, food journal, goal-setting, just-in-time intervention, food recommendation, knowledge graph

ACM Reference Format:

Philips Kokoh Prasetyo, Palakorn Achananuparp, and Ee-Peng Lim. 2020. Foodbot: A Goal-Oriented Just-in-Time Healthy Eating Interventions Chatbot. In *14th EAI International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth '20)*, May 18–20, 2020, Atlanta, GA, USA. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3421937.3421960>

1 INTRODUCTION

Motivation. Our health begins with what we eat. Maintaining healthy eating habit is one of the key factors in reducing the risk of chronic preventable diseases, such as coronary vascular diseases and type-2 diabetes, and increasing life expectancy. This generally involves consuming a healthful and balanced diet and practicing portion control on a regular basis. Nevertheless, a lack of motivation, poor self-regulation, and personal biases tend to prevent people from successfully developing and maintaining healthy eating behavior.

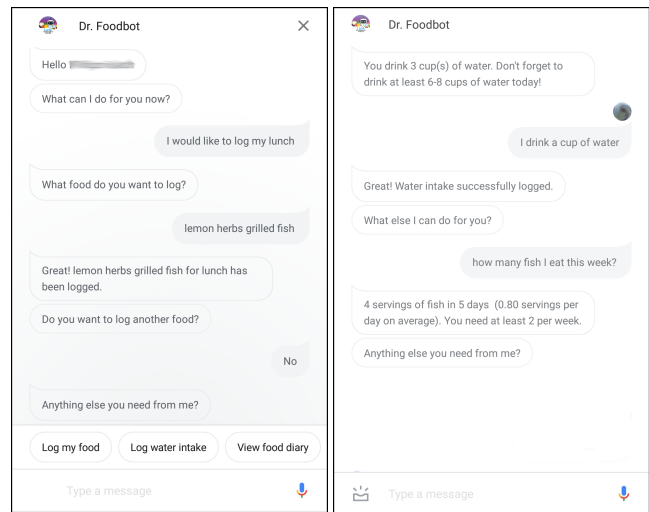
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PervasiveHealth '20, May 18–20, 2020, Atlanta, GA, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7532-0/20/05...\$15.00

<https://doi.org/10.1145/3421937.3421960>



(a) Food Logging

(b) Goal Tracking

Figure 1: Dialogue Turns in Foodbot

Subsequently, researchers have looked into mHealth and mobile sensing technologies as a cost-effective way to facilitating healthy eating lifestyle at a population scale. A popular type of mHealth applications (apps) for dietary self-tracking is mobile food journal, such as MyFitnessPal, in which users manually record their food entries by selecting items from a food database. However, previous studies [1, 5] recently identified design flaws in mobile food journal apps which lead to unintended and undesirable user behaviors.

Specifically, many users were frustrated by the tedious nature of manual data entry effort [5]. Next, most food databases used by mobile food journal apps consist of a combination of commercially available and user-contributed food and nutritional records. Even so, many users were not able to find many food items, such as those from local restaurants [5], due to poor data coverage. Lastly, unexpected lapses in healthy eating habits of mobile food journal users [1] suggested that mobile food journal apps tend to overly emphasize caloric and weight goals, whereas consumption of a healthful and balanced diet was not used as goals. Consequently, these design flaws limit the practicality and effectiveness of mHealth technologies in supporting healthy eating lifestyle.

Objectives. To address the aforementioned issues, we present Foodbot, a chatbot-based mHealth application for goal-oriented just-in-time healthy eating interventions. The main contribution of Foodbot is in its open-source design incorporating: (1) a natural-language user interface; (2) a large-scale food knowledge graph; and (3) goal-setting and JIT intervention and recommendation, to

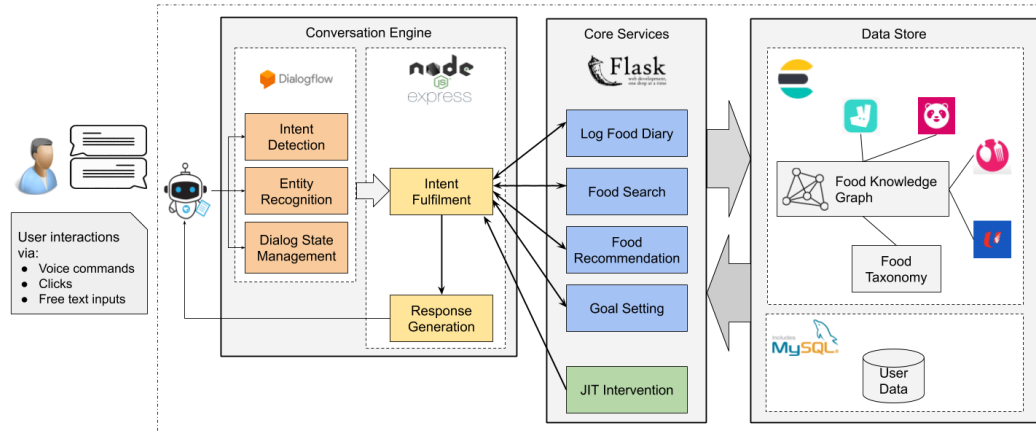


Figure 2: System Architecture

overcome the burden of manual dietary self-tracking and provide users with a personalized guidance toward a healthy eating lifestyle. Its behavior change design follows common techniques [13], such as *goal and planning* and *feedback and monitoring*.

As shown in Figure 1a, Foodbot utilizes automatic speech recognition and mobile messaging interface, ubiquitous in most smartphone platforms, to replace traditional food-journal user interface. Users simply dictate what they have eaten through simple dialogue turns. Next, Foodbot relies on a large-scale food knowledge graph to ensure sufficient coverage of food items. A simple key innovation of the food knowledge graph is in its data-driven and open-source approach to constructing an extensive database of locally available food items, sourcing from multiple online food delivery services. Together with the dietary self-tracking data, the food knowledge graph also enables Foodbot to provide personalized food recommendation to the users.

Lastly, Foodbot allows users to set specific dietary intake goals following evidence-based dietary guidelines [1], for example, a fish intake goal is at least two servings per week, a water intake goal is about 6-8 glasses per day, etc. After the goals are set, Foodbot will monitor users' progress toward those goals. As needed, JIT notification prompts will be sent to the users to remind them about their progress toward the goals. Users can also directly interact with Foodbot to inquire about their goals and request personalized food recommendations. The goal-tracking dialogue turns following a JIT prompt are shown in Figure 1b.

2 RELATED WORK

Early human-computer interaction (HCI) research has pioneered the applications of virtual agents in the health-related domains. Those studies typically focus on designing persuasive agents with human-like qualities for promoting physical activities [2–4, 15] and healthy eating [4, 12] in various populations. Recently, the widespread adoption of the smartphone & mobile apps and the substantial advances in wearable & mobile sensing, machine learning, and artificial intelligence (AI) technologies have renewed interests in conversational agents (commonly known as chatbot) as a promising solution to a population-scale and cost-effective delivery of

personalized health interventions. Several studies have explored the use of chatbot for monitoring and promoting behavior change in a wide variety of healthcare domains [14]. For health promotion, a new wave of data-centric chatbots for physical activities [7, 10] and healthy eating [6–8] have recently been explored.

For chatbots designed for supporting dietary monitoring and personalized feedback, our Foodbot is most similar to CoachAI platform [6, 7]. Both Foodbot and CoachAI employ similar behavior change techniques. However, CoachAI operates as a telemedicine platform, utilizing chatbot as a front-end user interface for monitoring and delivering interventions, thus requiring full supervision from healthcare practitioners. In contrast, all user interactions in Foodbot are completely automated, making it more scalable for population-wide interventions. Next, CoachAI uses wellness questionnaires to periodically collect overall dietary patterns from users, whereas Foodbot focuses on a food journal approach to dietary self-monitoring, which captures more granular data. Furthermore, Foodbot incorporates JIT mechanism in monitoring and providing personalized feedback.

The usage of a natural-language user interface for dietary self-tracking in Foodbot is similar to a nutrition dialogue system in [9], where a speech and language understanding model was trained to recognize 975 food concepts, including quantity, brand, and description. Foodbot depends on services provided by Google Assistant and Dialogflow, which were further tuned, to recognize 177,700 food entities from natural-language inputs.

3 THE FOODBOT SYSTEM

Foodbot is an open-source chatbot built on top of Dialogflow platform and accessible through Google Assistant. The Foodbot system, shown in Figure 2, consists of three main components, namely *conversation engine*, *core services*, and *data store*. Using Google Assistant, a user can interact with Foodbot via voice commands, clicks (on embedded user interface elements, e.g., cards and suggestion chips), and free-text inputs.

The conversation engine consists of: (1) Dialogflow services for intent detection, entity recognition, and dialog state management; and (2) a webhook server to bridge Dialogflow and core services.

With built-in automatic speech recognition service, the conversation engine first translates user-interaction events into specific service requests and then generates personalized responses back to the user. Each time the webhook server received messages from Dialogflow, it triggers an intent fulfillment to execute the corresponding core service to complete the user request provided that the information is complete. If the information given by the user is incomplete or ambiguous, the intent fulfillment will trigger response generation to send follow-up clarification questions.

The core services component consists of a collection of web services performing main functionalities, i.e., food logging, food search, food recommendation, and goal setting and tracking. The services are categorized into two types: request-based services (depicted as blue boxes in Figure 2), and JIT intervention services. The request-based services execute intent fulfillment requests. The JIT intervention relies on a job scheduler to invoke the services. The intervention triggers intent fulfillment to generate a response as notification prompt. Currently, JIT intervention delivers food recommendation and goal-tracking reminder.

Lastly, the data store component holds food knowledge graph and user data, including food journal data. The user data are stored in MySQL database and food knowledge graph in Elasticsearch.

3.1 Conversation Engine

The conversation engine is responsible for processing natural-language inputs, determining user intents and corresponding tasks, and managing dialog state. These are handled by three main modules in the Dialogflow platform. Given an input text (e.g., transcribed voice commands or free-form text messages), the intent detection module infers user intent from the input while, in parallel, the entity extraction module identifies relevant entities and information in the input. We train machine-learning models in Dialogflow by supplying 20-30 examples for each intent. Periodically, we add new examples gathered from past user interactions with Foodbot and retrain the models to improve the detection accuracy.

Next, the dialog state management module keeps track of conversation state between the user and Foodbot and determines follow-up intents based on contexts (e.g., previous responses). The intent fulfillment module acts as a router, bridging the user's natural language inputs and core services. Utilizing dialog state information together with user intent, context, and message, the intent fulfillment module decides whether Foodbot should ask a follow-up question or proceed to execute the user request. Lastly, the response generation module generates responses based on a pre-defined set of templates (i.e., canned responses) manually crafted for specific intents.

3.2 Core Services

The core services component provides back-end functionalities to the conversation engine. Utilizing food knowledge graph and user data, the core services component executes intent fulfillment requests and JIT intervention. Three main functionalities are offered through the following modules: (1) food logging; (2) personalized food recommendation; and (3) goal-setting and JIT intervention. The core services component also provides a general food item retrieval service based on food name, category, and location.

Food Logging. The food logging module provides basic services for food journaling. These services are essential to other modules, such as personalized recommendation and goal tracking, as they utilize the user's food log entries as the main input data.

The food logging module receives input messages from intent fulfillment containing food name and meal occasion (i.e., breakfast, lunch, dinner, or snack). The input food name will be matched with existing food entities in the food knowledge graph before a food log entry is recorded. If the exact match cannot be found, the user will be shown a list of top-15 most similar food items and asked to select one from the list. The user can also log water intake through Foodbot. Optionally, the user is able to view and modify past food log entries by accessing their food journal interface. To avoid lapses in recording food and water intakes, Foodbot reminds the user to log their food entries on a daily basis via notification prompts.

Personalized Recommendation. This module generates personalized food recommendation based on the user's past food intake data. The recommendations are delivered by either normal requests or JIT intervention. That is, users can directly ask Foodbot for food recommendation for a specific meal occasion. Before regular meal time, JIT intervention triggers food recommendation to be pushed to users. For example, Foodbot sends food recommendation for lunch to users at around 11 AM everyday.

Currently, the personalized recommendation module employs a simple rule-based algorithm to suggest food items for each meal occasion based on the user's recent and frequent consumption records. This method tends to yield higher recommendation accuracy than most state-of-the-art recommendation algorithms, but at the expense of novel choices [11]. When past consumption data are not available (e.g., for new users), the most globally popular items will be used for recommendations instead. To improve the user experience from personalized recommendation, more sophisticated algorithms and contextual information will be integrated into Foodbot in the future.

Goal-Setting and JIT Intervention. Users can ask Foodbot to keep track of a number of dietary intake goals. By default, the goal-setting module recommends a few intake goals according to evidence-based dietary guidelines [1]. The JIT intervention module identifies windows of opportunity to nudge the users toward their goals and avoids generating excessive prompts, which might lead to notification fatigue. Together, the goal-setting and JIT intervention modules provide *just-in-time*, *personalized*, and *actionable* feedback to users.

To determine progress toward the food intake goals from the user's food logging data, we utilize food category labels in the food knowledge graph. These labels correspond to basic dietary constituents, i.e., fruit & vegetable, red & processed meat, fish, added sugar, nut, and water. A serving amount is determined based on the frequency of relevant keywords in the user's food log entries matching the corresponding labels. For example, if the food journal contains two entries for *apple* and *banana*, the current progress for fruit & vegetable goal is 2 servings. At the moment, this simple approach is good enough for general behavior monitoring since Foodbot does not currently consider precise nutritional values, e.g., calories, as parts of the dietary goals.

Toward the end of the day, the JIT intervention module will send a notification prompt to the user if the gap between current and

target intakes exceeds certain thresholds. Additional conditions may also be applied to different type of goals. For example, for weekly fish intake goal, Foodbot will remind the user if no fish intake is observed before the dinner time on Wednesday (the middle of the week). At the end of the week, Foodbot will present the user with a weekly progress report to help the user reflect on their past consumption and, if needed, adjust their future behavior. If the user fails to achieve any of the goals, Foodbot will propose a new set of attainable goals.

3.3 Food Knowledge Graph

Food knowledge graph supports various core services' modules by providing a comprehensive knowledge base of local food items available from restaurants and grocery stores in Singapore (the location in which the target users of Foodbot reside). The food knowledge graph is constructed from an open-source online data about local foods and restaurants. To collect such raw data, we develop our own web crawlers to scrape and aggregate relevant information from the following online sources: (1) Online food delivery services Deliveroo and Foodpanda; (2) Burpple, a popular online social network for food and restaurant reviews; (3) Online grocery shopping site NTUC FairPrice, the largest local supermarket chains.

We believe that these sources contain a wide range of foods and beverages which are commonly consumed in Singapore. Given the initial crawled data, we perform data cleaning (e.g., removing duplicate records, text parsing and normalizing, etc.) and data integration to derive a unique set of food items representing entities in the knowledge graph. In total, there are 177,700 unique entities comprising: (a) 158,211 food items from local restaurants, food courts, and hawker centres in Singapore; (b) 15,588 packaged food items from grocery stores, and (c) 3,901 generic food items imported from food taxonomy, created by [16]. The inclusion of generic food items is to better handle naming variations.

Schematically, the food knowledge graph contains two types of entities representing food items and restaurants. Each type of entities has its own corresponding attributes. For example, food entity's attributes include name variations, descriptions, prices, and reviews, whereas restaurant entity's attributes include location, contact number, and cuisine type. We also enrich food entities with food item categories from food taxonomy [16], utilizing the categories for goal tracking service as explained in Section 3.2. Two types of linkages define relationships between food-restaurant entities (availability) and food-food entities (similarity).

4 CONCLUSION AND FUTURE WORK

We present an open-source chatbot for healthy eating intervention that runs on the Google Assistant platform. Powered by a large-scale food knowledge graph, Foodbot provides a natural-language user interface to reduce barriers in dietary self-tracking. In addition, it guides users toward setting realistic healthy eating goals according to evidence-based dietary guidelines and nudges them to achieve their goals via JIT personalized actionable feedback. Foodbot can be accessed at <https://research.larc.smu.edu.sg/foodbot>.

For future work, we plan to conduct user studies to evaluate the usability of Foodbot and, subsequently, the effectiveness of Foodbot in healthy lifestyle interventions in the general population. Next,

the food knowledge graph can be enriched with ingredients and nutritional information to improve the quality of food recommendation and expand the scope of goal-setting. Lastly, we wish to utilize natural-language generation models to improve response generation, enabling automatic generation of personalized and persuasive responses.

ACKNOWLEDGMENTS

This research is supported by the National Research Foundation, Singapore under its International Research Centres in Singapore Funding Initiative. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

REFERENCES

- [1] Palakorn Achananuparp, Ee-Peng Lim, and Vibhanshu Abhishek. 2018. Does journaling encourage healthier choices? Analyzing healthy eating behaviors of food journalers. In *Proc. of the 2018 International Conference on Digital Health (DH '18)*. 35–44.
- [2] I. M. Albaina, T. Visser, C. A. P. G. van der Mast, and M. H. Vastenburger. 2009. Flowie: A persuasive virtual coach to motivate elderly individuals to walk. In *Proc. of the 3rd International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth '09)*. 1–7.
- [3] Timothy W. Bickmore and Rosalind W. Picard. 2005. Establishing and Maintaining Long-Term Human-Computer Relationships. *ACM Trans. Comput.-Hum. Interact.* 12, 2 (June 2005), 293–327.
- [4] Timothy W. Bickmore, Daniel Schulman, and Candace L. Sidner. 2011. A Reusable Framework for Health Counseling Dialogue Systems Based on a Behavioral Medicine Ontology. *Journal of Biomedical Informatics* 44, 2 (2011), 183–197.
- [5] Felicia Cordeiro, Daniel A. Epstein, Edison Thomaz, Elizabeth Bales, Arvind K. Jagannathan, Gregory D. Abowd, and James Fogarty. 2015. Barriers and Negative Nudges: Exploring Challenges in Food Journaling. In *Proc. of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. 1159–1162.
- [6] Ahmed Fadhil and Silvia Gabrielli. 2017. Addressing Challenges in Promoting Healthy Lifestyles: The AI-Chatbot Approach. In *Proc. of the 11th EAI International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth '17)*. 261–265.
- [7] Ahmed Fadhil, Gianluca Schiavo, and Yunlong Wang. 2019. CoachAI: A Conversational Agent Assisted Health Coaching Platform. (2019). arXiv:cs.AI/1904.11961
- [8] Silvia Gabrielli, Kate Marie, and Carolina Della Corte. 2018. SLOWBot (Chatbot) Lifestyle Assistant. In *Proc. of the 12th EAI International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth '18)*. 367–370.
- [9] Mandy Korpusik, Salima Taylor, Sai Krupa Das, Cheryl Gilhooly, Susan Roberts, and James Glass. 2019. A Food Logging System for iOS with Natural Spoken Language Meal Descriptions (P21-009-19). *Current Developments in Nutrition* 3, Supplement_1 (06 2019).
- [10] Jan-Niklas Kramer, Florian Künzler, Varun Mishra, Bastien Pisset, David Kotz, Shawna Smith, Urte Scholz, and Tobias Kowatsch. 2019. Investigating Intervention Components and Exploring States of Receptivity for a Smartphone App to Promote Physical Activity: Protocol of a Microrandomized Trial. *JMIR Res Protoc* 8, 1 (31 Jan 2019), e11540.
- [11] Yue Liu, Helena Lee, Palakorn Achananuparp, Ee-Peng Lim, Tzu-Ling Cheng, and Shou-De Lin. 2019. Characterizing and Predicting Repeat Food Consumption Behavior for Just-in-Time Interventions. In *Proc. of the 9th International Conference on Digital Public Health (DPH2019)*. 11–20.
- [12] I. Mazzotta, F. de Rosis, and V. Carofiglio. 2007. Portia: A User-Adapted Persuasion System in the Healthy-Eating Domain. *IEEE Intelligent Systems* 22, 6 (Nov 2007), 42–51.
- [13] Susan Michie, Michelle Richardson, Marie Johnston, Charles Abraham, Jill Francis, Wendy Hardeman, Martin P. Eccles, James Cane, and Caroline E. Wood. 2013. The Behavior Change Technique Taxonomy (v1) of 93 Hierarchically Clustered Techniques: Building an International Consensus for the Reporting of Behavior Change Interventions. *Annals of Behavioral Medicine* 46, 1 (03 2013), 81–95.
- [14] Juanan Pereira and Oscar Diaz. 2019. Using Health Chatbots for Behavior Change: A Mapping Study. *Journal of Medical Systems* 43, 5 (04 Apr 2019), 135.
- [15] Daniel Schulman and Timothy Bickmore. 2009. Persuading Users through Counseling Dialogue with a Conversational Agent. In *Proc. of the 4th International Conference on Persuasive Technology (Persuasive '09)*. 1–8.
- [16] Ingmar Weber and Palakorn Achananuparp. 2016. Insights from Machine-Learned Diet Success Prediction. *Pacific Symposium on Biocomputing* 21 (2016), 540–51.