EDITOR JOANNA F. DEFRANCO
Pennsylvania State University; jfd104@psu.edu

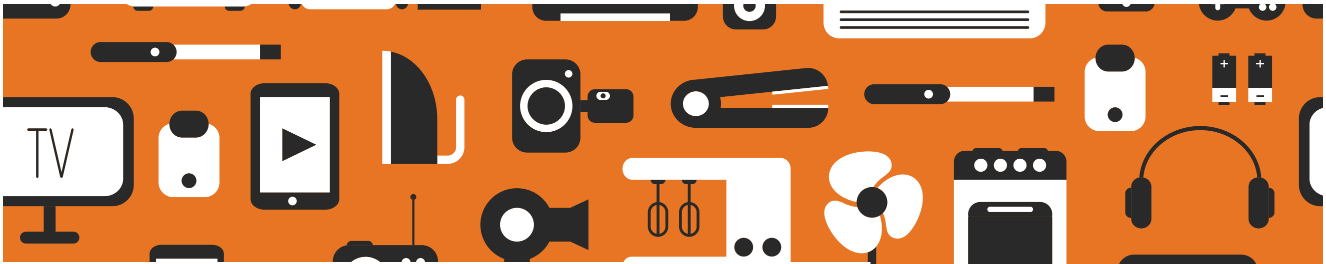

Alexa, M.D.

Sarah Jane Mee and Salam Daher, New Jersey Institute of Technology

In the health-care field, intelligent personal assistants (IPAs) are commonly used as diagnostic guides for health-care professionals or support for patients. This article expands the use of IPAs in health care by demonstrating using Alexa to train health-care personnel.

The Internet of Things (IoT) is at the forefront of new developments in software and technology. An intelligent personal assistant (IPA) is a type of IoT device that can respond to voice commands, complete basic tasks, and mimic the structure of human conversation. Popular IPAs include Amazon's Alexa, Apple's Siri, Google's Google Assistant, and Microsoft's Cortana. These services are used for daily tasks such as getting directions and making grocery lists as well as more complex tasks in health care and business. Within the health-care field, IPAs are often used as knowledge databases or diagnostic guides for health-care professionals and patients. For example, university researchers developed an Alexa skill to deliver interventions to metastatic breast cancer patients.[10] Another organization created an Alexa skill to help patients inquire about their own health-care plan and prescriptions.[11]

In this article, we explore the technical capabilities of Amazon's Alexa to create a customizable training model for health-care training. Human conversation is needed to train health-care providers (e.g., doctors, nurses) in scenarios such as a review of systems where they ask the patient questions about body systems to obtain his or her medical history and current health status. We developed an Alexa skill that can respond to user prompts using dynamic content provided by a subject-matter expert, initiate actions/questions, and provide custom feedback. We combined Alexa with dynamic content using the Google Sheets and Google Drive application programming interfaces (APIs), offered through the Google Cloud Platform. We describe the development process and present the technology's capabilities, limitations, and tips for optimization. This software opens the door to future research into the effectiveness and practicality of using this technology as a training tool for health-care students and professionals.

IPAs

Voice-driven communication with devices is a common feature in the current world of technology. IPAs such as Amazon's Alexa, Apple's Siri, Google's Google Assistant, and Microsoft's Cortana allow people to communicate with their devices in a hands-free manner and complete tasks

Digital Object Identifier 10.1109/MC.2021.3126528
Date of current version: 14 February 2022

such as getting directions, searching for recipes, setting reminders, messaging, checking the weather, and controlling smart devices. These intelligent devices are at the leading edge of technology advancements in artificial intelligence, speech recognition, and natural language processing and are being widely implemented in fields such as business, health care, and education. Although still relatively new, voice-focused IPAs provide more flexibility and applicability than the digital assistants of the early 2000s.

An intelligent personal assistant is a type of IoT device that can respond to voice commands, complete basic tasks, and mimic the structure of human conversation.

Uses of IPAs by the general public

The Pew Research Center (2020) reported that 46% of American adults use digital assistants, with 42% of the public using them on smartphones, 14% on a computer or tablet, and 8% on a standalone device as their medium.[1] Three categories of the most common voice commands are music, hands-free searching, and IoT control. IPA usage for IoT control involves manipulating devices such as smart thermostats, lights, speakers, and cameras. Other popular command categories include volume control, weather reports, timers, and alarms.[2]

IPA usage in health care

Although growing, the usage of IPAs in the health-care field is much less than in other areas like business and education.[3–4] Within the health-care field, IPAs can be used to help patients follow a medical routine, track eating and sleeping habits, and monitor signals such as heart rate and blood pressure.[4] General users tend to ask IPAs about illness symptoms or medication information, while broader uses in health care involve aiding users in scheduling doctors' appointments, setting and executing reminders for medication, and maintaining a health schedule.[5]

Customizable IPAs

IPAs come with a wide range of built-in skills, but they are often inflexible and difficult to customize.[6] Customizable content is extremely important for training health-care providers. There is literature and anecdotal evidence from health-care educators that there is a need for flexible and customizable IPAs with skills to assist users in creating personalized content.[1–6]

We created a review of systems' Alexa skill, where the user can ask Alexa (that is, the patient) a series of questions to obtain past and current health information. We combined Alexa and Google APIs to create custom content for Alexa to add variability to Alexa's responses, keep score of users' questions, have time-triggered events during skill execution, and provide custom performance feedback.

METHODOLOGY

We used Amazon Echo Show 8 (First Generation, 2019 Release) to gain access to the Alexa software, but any device with Alexa capabilities can be used. An Amazon Developer Account grants access to the Alexa Skills Kit and Developer Console,[7] which is needed to build a custom, Alexa-hosted Python skill and interface with Google Sheets. We used the Google Cloud Platform to access the Google Developer Console[8] and enabled the Google Drive and Google Sheets APIs on our project. Finally, we imported the gspread library[9] to provide functionality for opening, reading from, and writing to Google spreadsheets (see Figure 1).

Custom content for Alexa

Each question that a user can ask Alexa is an intent, and the response to an intent is managed by an intent handler in the back end. We edited the content of the intents in the front end and changed Alexa's responses in the back end to customize the questions and answers that Alexa can interpret in this skill.

Adding variability to responses

When responses are hard coded, Alexa responds to user prompts with the same answer every time. To add variability, we used the Google Sheets and Google Drive APIs to make our spreadsheet an answer bank that Alexa could read from. We listed each possible response in individual cells and assigned them group numbers based on what question they correlated to. We iterated through cells based on group number and picked a random response for Alexa to output. This allowed Alexa to give different responses for the same user prompt.

As we added more responses to our spreadsheet, the skill would often timeout or be unable to loop through too many lines of the sheet. When looping through lines of the spreadsheet in an intent handler, a timeout occurs after 65–70 lines. When the loop includes actions, such as copying or writing, a timeout occurs after around 30 lines. Additionally, when reading the cells of two columns within the loop, the limit is reached at 15–20 lines. To optimize the performance, we put all responses in a single cell, delimited by hash marks and numbers. We read this cell at the beginning of each skill execution and parsed the responses within the individual intent handlers. This allowed us to bypass looping through the rows of the spreadsheet and the timeout issues.

We created a Google Apps Script in Google Sheets to automatically format the plain text content directly inputted from health-care educators into the optimized format needed for the skill. This permitted the user to keep the customizability of his or her content while enabling him or her to remain separated from the inner workings of the skill.

Finally, to keep track of Alexa's current responses, we created a second sheet within our spreadsheet and wrote to the sheet every time Alexa outputted a response. This let us easily see what Alexa said last and keep track of the delivered responses.

Timer-triggered responses

We investigated having Alexa produce a response that was not prompted by a user utterance. We chose to have Alexa "sneeze" when the user is 2 min into skill execution by taking note of the starting time of the skill and comparing it to the current time in each intent handler. When the starting time and the current time were 2 min apart, we had Alexa output a prerecorded sneeze audio from the Alexa Skills Kit Sound Library.

Keeping score

We explored how to keep track of which questions the user asked by creating a session and persistent attributes within our skill, which allowed us to store data between intents and skill sessions. Each intent had associated attributes that incremented each time a user asked that specific question. Furthermore, we created a global score attribute that was an aggregate of all the other scores. This let us keep track of which questions the user had or had not asked.

We added weighted scores for individual intents to enable the user to make some questions more important than others. We allowed some intents to increment their scores by more than one each time they were triggered, resulting in some intents contributing more to the aggregate score than others, thereby creating weighted scores.

Additionally, we created a function that permits the user to save his or her score for a single run of the skill. We created a third sheet within our spreadsheet that acted as a database of user scores. When a user asks to save his or her score, the values of each attribute are written to the sheet. These user scores are associated with the name of the user, which the user is prompted to say at the beginning of skill execution. Each row in the database corresponds to an individual run of the skill, meaning that all entries are unique regardless of duplicate usernames. This provides for the concept of user accounts, where the user can have a name and score associated with his or her attempt, which can be saved for later access.

> IPAs come with a wide range of built-in skills, but they are often inflexible and difficult to customize.

Finally, providing custom performance feedback and after-action review (AAR) to users is essential for training. We structured the feedback to revolve around individual and aggregate scores. If the score for a question was greater than zero, the user had successfully asked that question and received positive feedback. Likewise, if the score for a question was zero, the user failed to ask that question and received negative feedback. The AAR was stored in the spreadsheet in the same one-cell formatting as the responses. Educators can input the general feedback topic in the spreadsheet, such as

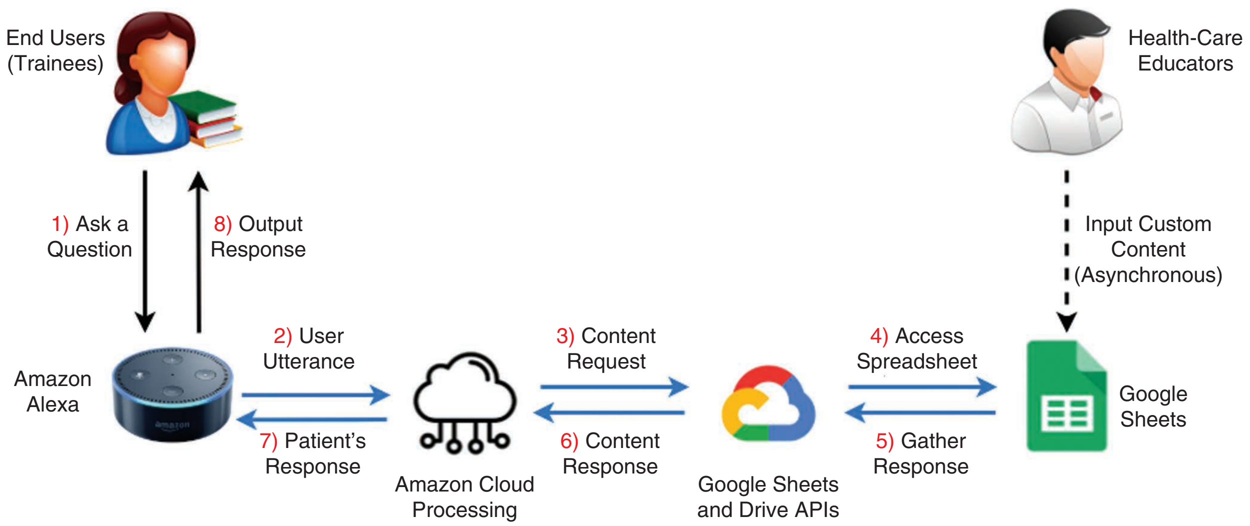

FIGURE 1. An overview of skill execution. API: application programming interface.

"medical history," and the skill will tell the user whether they successfully covered that topic based on the score for that intent and his or her overall score.

Master spreadsheet

Educators need to be able to input their own custom content and save individual user scores. We created a master spreadsheet that contains four-digit codes and links to another spreadsheet to handle multiple people or organizations using the same skill. Each code-link pair corresponds to a different organization, allowing each party to have his or her own spreadsheet to input content in and use for the skill. When the skill begins execution, the user is prompted to say his or her four-digit code. The skill then searches the master spreadsheet for the corresponding spreadsheet link and uses that new spreadsheet for skill execution. The master spreadsheet provides educators with their own separate spreadsheets, which have customizable content and an exclusive database for user scores.

Providing custom performance feedback and after-action review to users is essential for training.

DISCUSSION

The outcome of our exploration into the capabilities of Alexa and the Google APIs is a flexible and customizable review of systems' skill that allows users to ask personalized questions to an Amazon Alexa. The Google Sheets API connection enabled us to use a spreadsheet within the skill as an answer bank to add variability to Alexa's responses as a database to store user scores during skill execution and to keep track of Alexa's current response. Within the spreadsheet, we optimized reading and writing from the sheet by inserting all responses into one single cell to circumvent timeout problems. The skill keeps track of which questions the user has and has not asked as well as provides the user with the opportunity to save his or her score to a database and receive customized performance feedback. We also enabled a timer-triggered Alexa response, which caused Alexa to "sneeze" during skill execution. Finally, educators can have their own individualized spreadsheets and databases, which can be navigated to through the master spreadsheet.

Limitations

Limitations between Alexa and the Google APIs include Google Sheets imposing a limit of 50,000 characters per cell and request limitations (namely, 500 requests per 100 s per project, and 100 requests per 100 s per use). In addition, there are limitations when iterating through lines of the spreadsheet that depend on what is happening in the body of the loop. Actions such as copying and writing are more costly than simply reading through lines. Finally, the master spreadsheet has a limit of approximately 35–37 times, meaning that the skill can read and open a new spreadsheet link around 35–37 times before it times out. When a user code is implemented or the loop is storing the values of each cell, the limit tightens to 20–25 lines.

Future work

The implications of this technology are very widespread. With a lack of IPAs in the health-care field, technology that allows users to customize content is extremely valuable. Future plans include testing this technology with a user study to explore usability with health-care educators, and training effectiveness for health-care students. Additionally, we are interested in trying to combine this Alexa skill with the Unity platform to create a virtual health-care simulation that involves both graphics and speech.

We explored the technical capabilities of Amazon's Alexa in combination with Google Sheets and Google Drive APIs to create a customizable health-care training model. We created a custom skill that allows the user to input personalized dynamic content, initiate questions, add variability to Alexa's responses, keep a database of user scores, have time-triggered responses within skill execution, and receive feedback on his or her performance. The skill was designed as a review of systems' training models for health-care students; however, it has many additional possible applications in the health-care field. This technology explores new functions of intelligent personal agents in the health-care training industry. C

REFERENCES

1. "Voice assistants used by 46% of Americans, mostly on smartphones," Pew Research Center, Aug. 25, 2020. [Online]. Available: https://www.pewresearch.org/fact-tank/2017/12/12/nearly-half-of-americans-use-digital-voice-assistants-mostly-on-their-smartphones/ (accessed Sep. 9, 2021).
2. T. Ammari, J. Kaye, J. Y. Tsai, and F. Bentley, "Music, search, and IoT: How people (really) use voice assistants," *ACM Trans. Comput.-Human Interact.*, vol. 26, no. 3, pp. 1–28, Apr. 2019, doi: 10.1145/3311956.
3. L. Laranjo et al., "Conversational agents in healthcare: A systematic review," *J. Amer. Med. Inform. Assoc.*, vol. 25, no. 9, pp. 1248–1258, Jul. 2018, doi: 10.1093/jamia/ocy072.
4. G. Riccardi, "Towards healthcare personal agents," in *Proc. Workshop*

on *Roadmapping the Future Multimodal Interact. Res. Including Bus. Opportunities Challenges (RFMIR '14)*, 2014, pp. 53–56, doi: 10.1145/2666253.2666266.

5. B. Kinsella, "Voice assistant consumer adoption report for healthcare 2019," Voicebot.ai, Mar. 3, 2020. [Online]. Available: https://voicebot.ai/voice-assistant-consumer-adoption-repor-for-health care-2019/ (accessed Sep. 9, 2021).
6. M. Campbell and M. Jovanovic, "Conversational artificial intelligence: Changing tomorrow's health care today," *Computer*, vol. 54, no. 8, pp. 89–93, Aug. 2021, doi: 10.1109/MC.2021.3083155.
7. "Create Alexa skills kit | Amazon Alexa voice development," Amazon. [Online]. Available: https://developer.amazon.com/en-US/alexa/alexa-skills-kit (accessed Sep. 9, 2021).
8. Google Cloud Platform. [Online]. Available: https://console.developers.google.com/ (accessed Sep. 9, 2021).
9. A. Burnashev, gspread. [Online]. Available: https://docs.gspread.org/en/latest/ (accessed Sep. 9, 2021).
10. J. Hallman, "Amazon Alexa skill offers supportive care to breast cancer patients," Pennsylvania State Univ., University Park, PA, USA, Apr. 27, 2021. [Online]. Available: https://tinyurl.com/2wk79dd7
11. L. Dyrda, "ChristianaCare launches Alexa skill for home health: 4 Details," *Beckers Hospital Rev.*, Oct. 21, 2020. [Online]. Available: https://tinyurl.com/4ruzxc6r

SARAH JANE MEE is a computer science and information technology student at New Jersey Institute of Technology, Newark, New Jersey, 07103, USA. Contact her at sm2778@njit.edu.

SALAM DAHER is an assistant professor of informatics and computer science at New Jersey Institute of Technology, Newark, New Jersey, 07103, USA. She is also a courtesy faculty member in the College of Nursing at the University of Central Florida, Orlando, Florida, 32816, USA. She is the corresponding author for this article. Contact her at salam.daher@njit.edu.

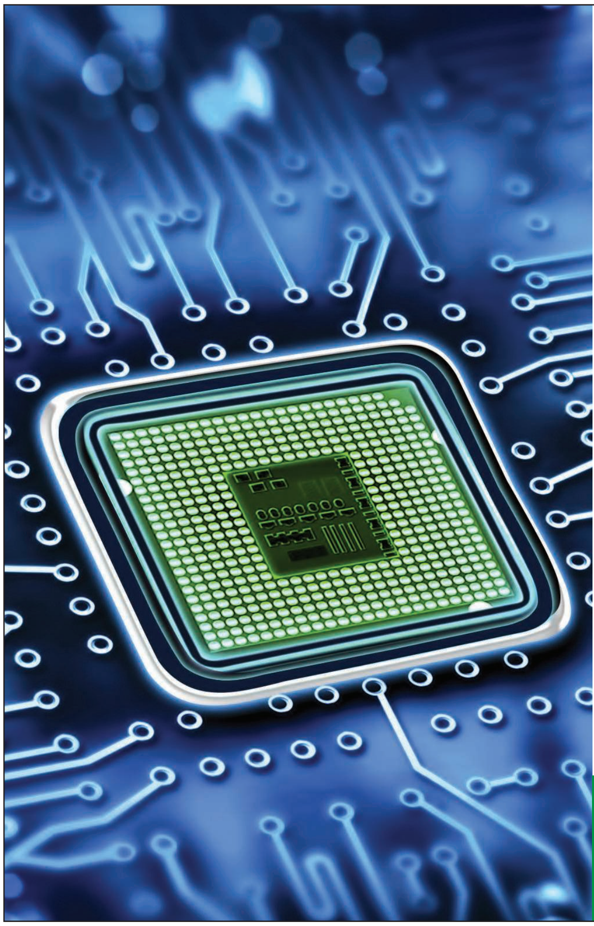

Digital Object Identifier 10.1109/MC.2022.3142897