Exploring the Design and Implementation Implications of a Dialogue Mood Logger on a Voice User Interface

Ciara Gilsenan, BA

A Dissertation

Presented to the University of Dublin, Trinity College

in partial fulfilment of the requirements for the degree of

Master of Science in Computer Science (Intelligent Systems)

Supervisor: Dr. Gavin Doherty, Robert Bowman

August 2022

Declaration

I, the undersigned, declare that this work has not previously been submitted as an exercise for a degree at this, or any other University, and that unless otherwise stated, is my own work.

Ciara Gilsenan

August 19, 2022

Permission to Lend and/or Copy

I, the undersigned, agree that Trinity College Library may lend or copy this thesis upon request.

Ciara Gilsenan

August 19, 2022

Exploring the Design and Implementation Implications of a Dialogue Mood Logger on a Voice User Interface

Ciara Gilsenan, Master of Science in Computer Science

University of Dublin, Trinity College, 2022

Supervisor: Dr. Gavin Doherty, Robert Bowman

Since the pandemic there has been a sharp global increase in anxiety and depression among people, and there is still an ongoing issue with the insufficient amount of support people can get access to. Due to this lack of support there has been much research conducted in the past on the benefits of self tracking and on the capabilities that voice assistants and conversational agents have in helping treat mental health conditions. This research presents a dialogue based mood logger on Google Assistant which facilitates daily mood tracking and allows users to learn and reflect on their mood data through analyses generated by the mood logger. The mood logger was developed using the Google Actions Console and Firebase. The dialogue and user experience was evaluated through a study in which the mood logger was deployed to 5 participants. Feedback from the participants was collected via a questionnaire. Findings from the study indicated that participants found the dialogue approach effective and time efficient. The findings also indicated that most participants found the analysis given from the mood logger informative and gave them insight into how daily logging can help them understand more about their wellbeing. Unfortunately there were some issues with the technology used for the mood logger, and in the future it may have to be transferred to another voice assistant. Discussions on the findings of the study as well as the design, implementation and implications of the dialogue mood logger will be discussed in this research.

Acknowledgments

I would like to thank my supervisors Dr. Gavin Doherty and Robert Bowman for their consistent guidance and encouragement over this past year, without which this dissertation would not have been possible.

Ciara Gilsenan

University of Dublin, Trinity College
August 2022

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Motivation

In 2019 1 in 8 people suffered from a mental health condition (who, 2022b), and according to the WHO, since the COVID 19 pandemic there has been a global increase of 25% in the prevalence of anxiety and depression. This was mainly due to the increased stress on people's lives caused by the pandemic including financial stress, work stress, grieving loved ones, fear of infection, etc (who, 2022a). This along with the decreased hospital and GP visits due to lockdown caused a significant gap in the support available for people with mental health conditions. Mental health services such as suicide prevention and counselling sessions were affected in countries around the globe. In 2021 the WHO reported that this had improved, however many people still don't have access to effective support and care (who, 2022a). As a result many people are looking for support online through online counselling and/or workshops. One technique used to help treat mental health conditions alongside counselling is Mood Logging or Mood Tracking.

1.1.1 Mood Logging

Mood Logging is an exercise used alongside many treatments for Cognitive Behavioural Therapy (CBT). CBT is an effective treatment for many conditions such as anxiety, depression, addictions, eating disorders, etc. It can allow people to analyse patterns in their thinking, emotions or behaviours and create practices to change them (Cherry, 2022). Mood logging involves recording your mood and any external factors that may have influenced it at set intervals e.g. once a day. It allows people to uniquely identify patterns in their stressors and/or triggers in their life which influence their anxiety and depression symptoms, and in turn people can make changes to their lifestyle in order manage their symptoms better.

There are many ways to track your mood. People can write it in a notebook or daily planner, which could also encourage them to journal which is another helpful tool for mental health support. Another way is through charting which is useful for people looking for patterns in their mood. Another way is with illustrations or graphics which can help people who prefer to take in information visually (Cherry, 2021). A more popular way of mood logging these days is online via websites or mobile apps. People can try a few ways to track mood and see which best suits them and their lifestyle. However these strategies of mood logging may not be for everyone. People may not have time in their busy schedule to write it down or to color it in a chart, and due to the large amount of mood logging apps and websites, it could be difficult to distinguish the good and bad ones.

1.1.2 Voice User Interfaces

With the increased use and interest in voice user interfaces (VUIs) there have been many cases of exploratory research being conducted on the effect they can have on mental health treatment. With their ability to perform tasks and have conversations with users, they can provide conversational experiences in which people can check in and examine how they are feeling. These conversations with VUIs can allow people to be more expressive and honest about how they feel, resulting in more tailored and accurate support (Maharjan et al., 2021).

1.2 Research Aim

The aim of this research is to develop a mood logging application on a voice user interface (VUI) that can allow users to log their mood daily. It also would provide weekly analysis on their mood data which would hopefully help users understand more about what impacts their mental health. A study would also be conducted to evaluate the user experience and dialogue approach of the mood logger. This study would hopefully give insight into how people view conversational experiences with voice assistants about their mood.

1.3 Outline

Chapter 2 will describe relevant research done on tracking for mental health, conversational agents, voice assistants and online mental health support. Chapter 3 will discuss the motivations and main design goals of the mood logger. Chapter 4 will describe in detail the architecture and conversational flow of the mood logger, followed by its implementation using Google Actions and Firebase. Chapter 5 follows this with evaluation of the resulting mood logger with screenshots of its look on a mobile device. Chapter 6 will discuss the study conducted to review the user experience of the mood logger with potential users, followed by Chapter 7 which will evaluate and analyse the results of the study. This research will then conclude in Chapter 8 which discusses implications, limitations and future work.

Chapter 2

Relevant Work

2.1 Tracking for Mental Health

There has been plenty of research done over the years that has looked into the benefits of tracking daily habits for your mental health. A paper by Kelley et al. (2017) looked into how self tracking can be important to the mental wellbeing of students. They collected data by conducting two studies on self tracking. The first study focused on understanding health professional's expert opinion on self tracking for students. The findings from this study indicated that common data types that are useful to track include exercise, sleep quantity, academic workload, and class attendance. The second study by Kelley et al. (2017) focused on student's perspectives on self tracking in general and in particular for mental health. According to their findings, workouts, steps, weight, and sleep were among the top tracked data types by students. The paper also pointed out that students felt like they understood more about themselves through self tracking and that it even motivated them to achieve their goals. This paper offers plenty of foundational information into what one could track and how it can be beneficial to their lives, especially with mental health.

One point to make about self tracking is that it can be difficult for people to recognise patterns in their data if they are tracking multiple things at once. A paper by Bentley et al. (2013) worked on an app that combined multiple tracking data types into one platform and provided users with analyses on their data. The goal of the app was to encourage the users to change their behaviours based on the analyses given. The app connected to multiple sensors and trackers so that it could simultaneously receive multiple streams of data. Some of the data was automatically tracked, e.g., locations and calendar events, while others came from sensors such as steps count a sleep measured on a Fitbit. Other data types such as mood and food intake were manually tracked by the users. Bentley et al. (2013) conducted a study with 60 participants that used the app for 90 days. Their findings reported that the app was able to deliver accurate observations which lead to the participants being able to understand more about themselves. Participants found it interesting to learn more about themselves as they ongoingly received observations. This would hopefully be achieved by this research with the mood logger.

Another study involving mood tracking done by Matthews and Doherty (2011) was conducted during the development of a mental health symptom tracking application for adolescents with mental health conditions. The goal of the system was to encourage adolescents to track their symptoms to be viewed together with their therapist in a clinical setting. These are to improve on other traditional "homework" given by clinicians that have a low completion rate (Matthews and Doherty, 2011). The application was developed for mobile devices, and allowed users to track mood, energy, sleep and a diary entry. It also provided visualisations of their data using graphs. Matthews and Doherty (2011) conducted a pilot study done with 9 adolescent participants that were in a public mental health clinic who used this digital mood diary between 2 therapy sessions. The results of the pilot reported a high adherence rate, showing that the adolescents were able to engage with the mood diary, and therapists reported that it helped the adolescents engage more in their conversations during sessions.

The research presented here can show how tracking mood and other daily habits can increase one's understanding of themselves and can benefit people and their therapists during treatment.

2.2 Voice User Interfaces & Conversational Agents

Over the years there has been a rise of use in voice user interfaces (VUI) such as Google Assistant/Google Now, Siri, Alexa, etc. Their ability to perform tasks and make life easier without the need for reading and typing makes them attractive to users. With this increased interest there is much potential in the abilities of these VUIs. They can help with voice biometrics to improve security, instant customer support from various businesses, voice assisted emailing and shopping, and much more (Palmer, 2022).

However these voice assistants or agents need to be able to facilitate functional conversations in order to carry out these tasks. There are various aspects of conversation to consider when developing conversational agents. A paper by Clark et al. (2019) investigated the design of conversations in voice agents and the challenges that come with it. Most conversations with Intelligent Personal Assistants (IPA) are task oriented, and the paper proposes that IPAs with a more human-like emotion in their conversations could lead to a stronger bond between the human and the agent, resulting in better quality conversations particularly about sensitive subjects such as healthcare. When conducting a study to see what aspects of conversation are important, Clark et al. (2019) split these into two categories - social and transactional. Upon investigating their findings, they found that participants did not consider important characteristics in conversation between humans such as active listening, vulnerability and trust to be the same in conversations between a human and an agent. They believed that conversations with agents are for transactional purposes and that conversations with agents do not necessarily need to mimic conversations between humans (Clark et al., 2019). The paper concludes that conversations with agents could become a new category of conversation style.

Conversational agents have also been developed in research for reflection. Kocielnik et al. (2018b) developed a conversational system 'Reflection Companion' which was designed to allow users to reflect on their self-tracked physical activity data. This reflection was provided through daily mini-dialogues about the user's Fitbit data. Their aim was to encourage users to reflect more about their own data and make changes based on the reflection. Kocielnik et al. (2018b) conducted a study with 33 Fitbit users and their findings indicated that these mini-dialogues can provide effective ways for people to self reflect. Another paper by Kocielnik et al. (2018a) discussed the development of a voice based and chat based conversational agent - Robota to facilitate activity journaling and self reflection for the workplace. Their aim was to investigate the effect voice based and chat based conversations with the agent have on self reflection. Robota was deployed to 10 participants that used it for three weeks. The findings of the paper indicated that participants appreciated the self reflection. The paper also indicated that voice interactions with Robota were more engaging and personal which could also enhance reflection (Kocielnik et al., 2018a). With this research people are finding that there are more ways for these conversational agents to be beneficial to people's health and wellbeing.

2.3 Online Mental Health Support

With the ongoing problem of lack of support for mental health in the healthcare industry (who, 2022b), a lot of people turn to online support. There have been an abundance of services out there that can offer mental health support. A lot of this support is offered throught apps. An example is a mood tracking app Daylio[1] which has over 10 million downloads. This is just one of the many mood tracking and/or mental health support apps that one can download. With all these apps to choose from, it could be difficult to distinguish the effective ones that have backed up their services through research from the

[1] `https://play.google.com/store/apps/details?id=net.daylio&hl=en_IE&gl=US`

apps that just use that as a marketing tool. A study by Marshall et al. (2019) collected apps in the Google and Apple Apps stores that claimed to have conducted research or included input from mental health experts as part of their development process. They examined a shortlist of 293 apps and found that just over 3% of the apps shown that they had research to back up their effectiveness. This research included pilot studies and controlled trials. 30% of the apps claimed to have received input from mental health professionals. According to Marshall et al. (2019) two thirds of their shortlisted apps were developed without any input from mental health experts. This shows that it is important that expert opinions are taken into account when developing online mental health support, if not it could have a negative effect on the people using these services if they are not equipped to handle sensitive information in the proper way.

Another paper by Balaskas et al. (2021) examined anxiety related apps that offer cognitive behavioral therapy (CBT) and looked into their content delivery and user engagement. 36 apps from both Apple and Google Play Store were tested. It was found that these CBT apps offered low intensity exercises for support and a small number of apps offered a structured treatment plan. These apps also don't offer as much clinician input as they should. In terms of engagement, a large portion of the apps used self-monitoring and visualisation but very few used other engagement techniques such as points, badges and level upgrades (Balaskas et al., 2021). The paper states that more tailoring options to individual users should be explored to help enhance the CBT treatment offered by these apps.

Engagement appears to be a big design goal for online mental health support. Companies and organisations can develop apps and put a lot of research into it, however it won't do any good if users aren't engaged enough to commit to it. Doherty et al. (2012) developed a platform SilverCloud which was designed to provide more engaging online support for depression. The platform was developed with design strategies including personal tailoring to the user, providing interactive elements, facilitating contact with therapists and peers. The paper focused on a CBT programme for depression, and conducted a study to evaluate the system. Their findings indicated that there was an high level of engagement among participants and even an improvement of depression symptoms. This research shows that these design strategies can change how online support is developed and that applications like these can be beneficial to use in addition to face-to-face therapy.

2.4 Mental Health Support via Conversational Agents

With the increased use of voice assistants/voice based conversational agents, more exploratory research is being done to investigate their capabilities in terms of mental health support. A paper by Maharjan et al. (2021) conducted a study to review two designs of a conversational agent named Sofia. The agent collects responses of the World Health Organization-Five Well-Being Index (WHO-5) questionnaire, one design collects discrete responses while the other collects open-ended responses, both using the same questions. The questionnaire acts as a self report for health and wellbeing and is usually completed on pen and paper. Similar to the research in this paper, Sofia was developed using Dialogflow and deployed to Google Home. The aim of the study was to explore the impacts of using a conversational agent to carry out the WHO-5. The findings of the study indicate that the conversational agent can be a useful medium for carrying out the self reporting questionnaire (Marshall et al., 2019). Participants found that the discrete design of Sofia was more habitable, however the open-ended design does all the agent to get a better understanding of the participant's wellbeing. This paper shows that conversational agents can be a useful tool in collecting information about a person's wellbeing.

Chapter 3

Design

The two most popular development toolkits available for developing programs for VUIs are Amazon's Alexa Skills Kit and Actions for Google Assistant. Google Assistant was chosen for this research as it was the most accessible.

3.1 Design Goals

There were initially two main design goals for the mood logger - to give the user the ability to log their mood daily, and the ability to receive an analysis on their mood. There were also two additional goals - the ability to send data to the user's GP, and the ability to view a log from a given day. These goals are discussed further below.

3.1.1 Log

In a mood logger it is important for the logging process to be simple and quick, therefore the dialogue in the voice mood logger had to facilitate this. The conversation had to be simple enough so the users would want to go back to it the next day. If they knew it would only take a few sentences and less than a minute, they could be more motivated to stay consistent with it. If it were too long, then the users might not be motivated to do it again.

What should be logged

As well as mood it was important to also log other variables representing other aspects of the users lives, that way the mood logger could find connections between people's day to day habits and their mood. There was a plethora of variables to choose from. The paper by Kelley et al. (2017) which looked into self tracking stated that the top 5 data types that their participants track are: steps, workouts, weight, sleep and eating habits. Additional information tracked was water intake, heart rate, phone usage, etc. This gave some insight into what kind of data that the mood logger could collect. It was important to select data types to track that would not end up being harmful to a user's mental health. Kelley et al. (2017) mentions that tracking some data might have some negative impacts, an example used was someone who has an eating disorder might get obsessed with tracking food. The same could be said for somebody with body-image issues obsessively tracking their weight or someone with anxiety obsessively tracking their heart rate. Therefore it was decided that these data types should not be recorded in this mood logger.

However the other data types mentioned seemed to work well. Sleep would be an excellent variable to track as it is known that sleep quality can influence people's moods. Since steps and working out were the top two tracked data types above, activity/exercise would also be interesting to track as it is known to elevate mood. Bentley et al. (2013) also tracked steps and sleep during their study, therefore these two data types would be useful. Although it was not one of the top things to track, water intake can be an important component in the state of one's mental health. Dehydration can effect one's bodily functions, including the ability to keep your mental health in check, therefore drinking enough water can be important to make sure there is less chance of a decline in one's wellbeing (Northstar Transitions, 2021). With all this in mind it was decided that **sleep, activity & water intake** would be the most suitable to track alongside mood in this mood logger.

3.1.2 Analysis

Analysing mood data can be a very useful tool when looking into what can effect your mental health. It can provide important information about what effects a person's mood that they might not have realised before. The goal of the analysis is to allow the user's to reflect on their mood by providing observations and patterns in the user's logs. These observations could shed light on what can trigger a person's anxiety symptoms and in turn the person could find ways to avoid these triggers, and produce healthier habits that influence good moods. This would be done on a weekly basis to provide new observations on the new data that the user would have logged in that past week. Bentley et al. (2013) were able to provide observations on their participants tracking to the point where behavioral changed occurred. This can hopefully be achieved with the mood logger. Some observations that the user could infer from the analyses could be:

- I am more anxious when I don't exercise.

- I don't sleep enough.
- I feel better when I drink water.

The dialogue the mood logger would use to deliver the analyses would not be as certain as these ones listed above, as the mood logger wouldn't be 100% accurate in its observations, since only the user knows truly how these things effect their mood. The mood logger would therefore use a different dialogue communicating what *might* effect the user's mood, rather than what *certainly* effects it. Examples of that would be:

- You are *more likely* to be anxious when you don't get enough sleep.
- You are *more likely* to feel less anxious when you exercise.

3.1.3 Send to GP

This tool would be useful to users when working with their therapist or GP on their mental health. Their GP may want to look at what they are logging each day to see if they are progressing in their treatment. GPs can also use this to monitor their patient's reaction to their medications such as antidepressants. If they see through the mood logs that they are not progressing or having a bad reaction, then they can make adjustments to the treatment. Similar work was seen in Kelley et al. (2017) who asked students who track daily habits how they would feel about sharing what they track with health professionals. 27% responded that it helped them better communicate with their doctors and helped them better understand each other.

3.1.4 View Log

This could be useful to the user in a case where they had a bad week and they wanted to know what they logged on one of the days but couldn't remember specifically what occurred; they could simply ask the mood logger what they logged that day and it would replay it back to them.

Chapter 4

Implementation

4.1 Overview

4.1.1 Architecture

The mood logger was implemented and deployed using the Google Actions Console[1] and Firebase[2]. The architecture can be seen in Figure 4.1. From looking at Figure 4.1 there were three main components that worked together. The Google Actions Console was used to build and deploy the mood logger, the Firebase cloud functions were used as a back-end which dealt with the logic and background functionalities, the Firestore Database contained the mood logs and analysis for each user. This chapter contains more details about each of these components.

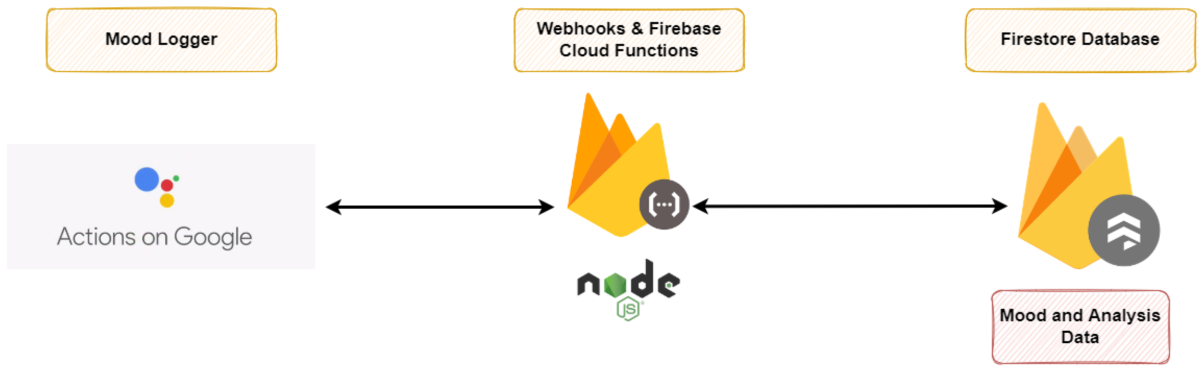

Figure 4.1: Mood Logger Technical Architecture

[1] https://developers.google.com/assistant/console
[2] https://firebase.google.com/

4.1.2 Flow

The mood logger can be started on Google Assistant by opening it and saying **"Talk to my Mood Logger"**, which invokes the Google Action. When it starts, it asks the user what it would like to do, and presents four options - logging your mood, viewing an analysis, viewing a log, or sending some data. The user then must select a task to continue. Once the task is complete the mood logger would ask the user if they would like to do any other task, if so they would go back to task selection, otherwise the conversation would end. An illustration of this can be seen in Figure 4.2.

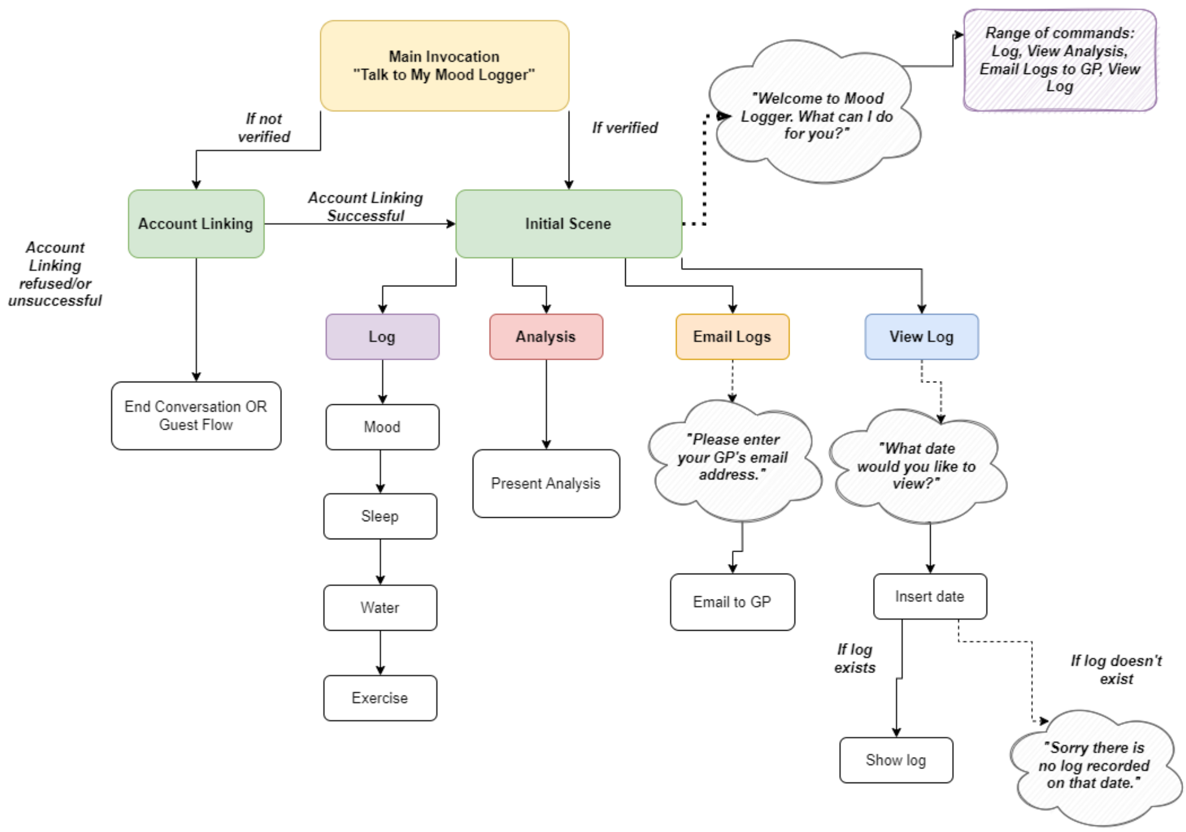

Figure 4.2: Mood Logger Flow

4.2 Google Actions Console

The Google Actions Console is a web tool that allows developers to create and deploy an action to Google Assistant. These actions can offer various functionalities such as connecting and opening apps on specified pages with just your voice [3], providing con-

[3]https://developers.google.com/assistant/app

versational experiences and games[4], and connecting to smart home devices[5]. The mood logger is a conversational action as it does not connect to any other apps or devices. A screenshot of the console can be seen in Figure 4.3.

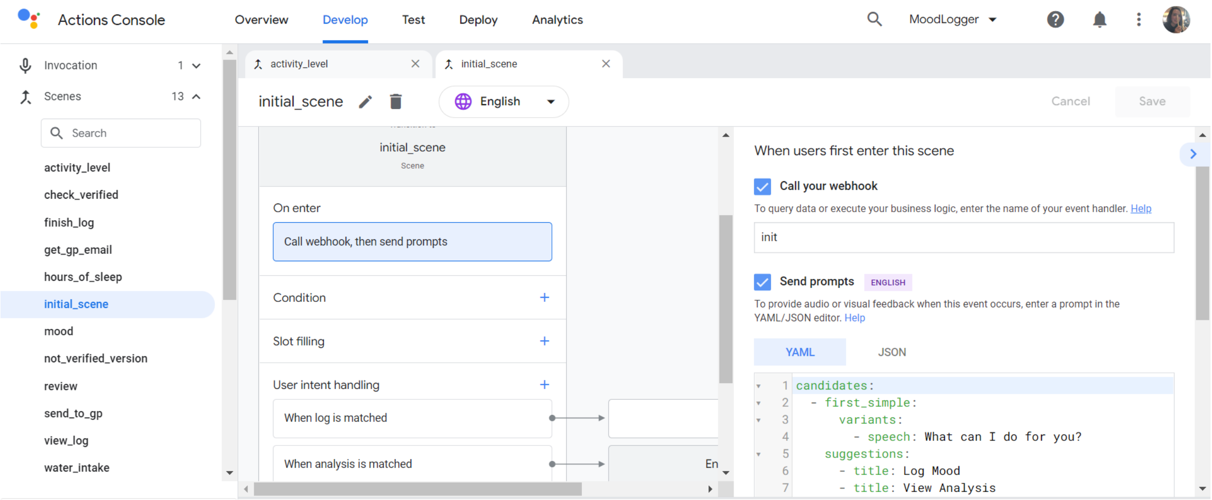

Figure 4.3: Screenshot of Actions Console

Conversational Actions allow developers to create and deploy additional functionalities to their Google Assistant. These can be quite simple to create thanks to the assistant's **Natural Language Understanding** (NLU), with which developers can create their actions without worrying about how the program will parse the natural language of the user's inputs (Google Developers, 2021a). Examples of its use in the mood logger is discussed in Sections 4.2.2 and 4.2.3.

The conversational action is comprised of several components, each of which was used when developing the mood logger: **Scenes**, **Intents**, **Slots**, **Types**, and **Webhooks**. Development can either take place on the actions web console itself or locally in an IDE which would use a YAML[6] file structure. The web console was thought to have been a better choice as it provided better support and helped visualise the conversation as it was being built. The actions console also provides a simulator, seen in Figure 4.4, to test the action on, which was very helpful during development. The following subsections goes into detail on the different building blocks of the mood logger as a conversational action, followed by an example where all of these components are used in the mood logger in Section 4.2.7.

[4]https://developers.google.com/assistant/conversational
[5]https://developers.google.com/assistant/smarthome/overview
[6]https://yaml.org/

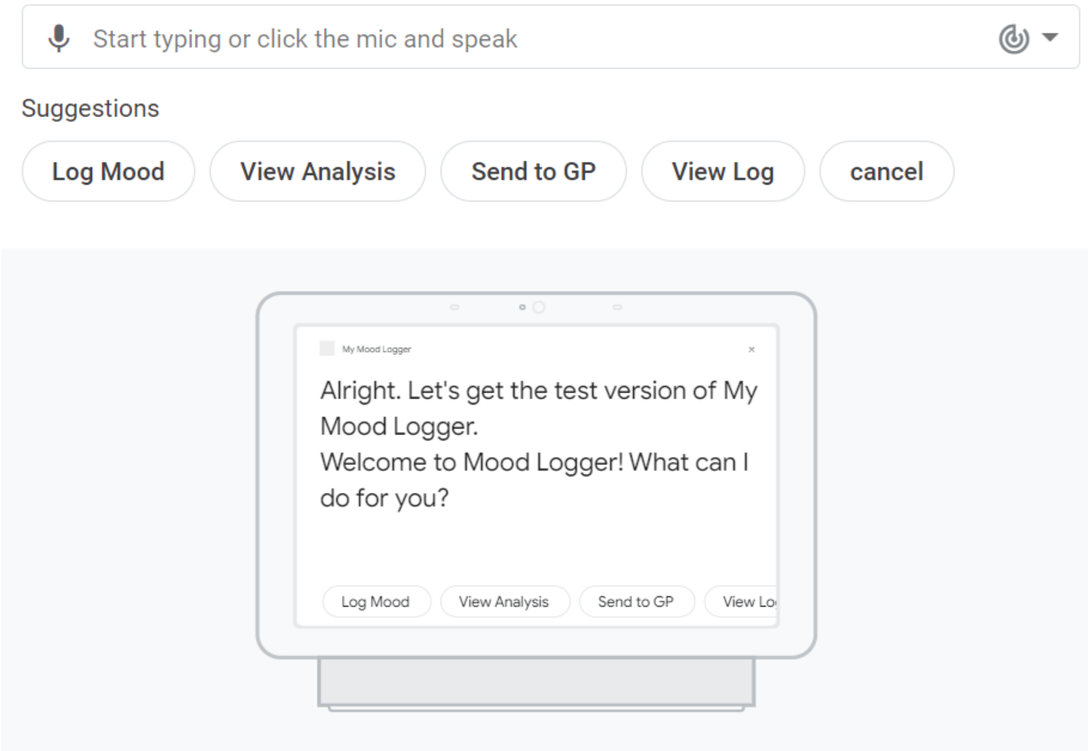

Figure 4.4: Testing Action in Simulator

4.2.1 Scenes

The structure of the mood logger's conversations can be described as a series on scenes. A scene is a building block in the conversational action representing the state of the action, and the action can transition between multiple scenes within one conversation (Google Developers, 2021b). Developers can configure the transitions of these scenes, allowing them to continue to the next scene or loop to the beginning of the scene when a certain condition is fulfilled. Within a scene the mood logger can prompt the user for a response, collect data, wait for a user response, call webhooks, etc. Google Actions Console also offers "System Scenes" such as **End Conversation** and **Push Notification**, however the majority of the scenes used in the action were custom scenes.

Each scene had a specific function and goal and when that goal was fulfilled it could transition to the next scene. An example goal used in a scene was getting the email address of the user's GP, which was fulfilled once the action identified an email address in the user's response. Another example from the mood logger was the "start" scene. When the mood logger is invoked, it presents the user with a scene offering four choices, as seen in the simulator in Figure 4.4. Once the user has responded with one of the specified choices, the mood logger can then transition to the corresponding scene. The mood logger has a scene for the collection of each variable during the logging process - mood, water intake, hours of sleep and activity of the day. There were also scenes for presenting the analysis of the mood data, sending the mood data to the user's GP and viewing a log.

4.2.2 Intents

An intent is a task to be executed by the action only when the user has expressed their interest in that task being done. Intents can be used within individual scenes or globally at any point during the action's execution. They are also used to invoke the action to start when a user initially opens their Assistant. Most intents used in an action are custom intents. When creating these, a set of training phrases must be provided which are used to create a language model. The intent's language model then fine-tunes Google Assistant's NLU to better understand the functionality of the action (Google Developers, 2022). Google Actions also offer "System Intents" for common phrases such *Yes* and *No*.

In the mood logger, intents were used in the initial scene where the four main tasks were presented for the user to choose from. An intent was set up for each task - *log mood, view analysis, view log* and *send to GP*. Therefore when one is matched it will start the corresponding task. For example, if the log intent is matched then the mood logger will prompt the user about their mood. Table 4.1 shows the training phrases used for the intents in the initial scene.

Intent	Training Phrases
Log	log, i want to log my mood, log today's mood, log mood, start log
Analysis	analysis, show analysis, I want my analysis, analysis please
View Log	view, i want to view a log, show me a log
Send to GP	I want to send to gp, I want to send my data to my gp, send, send to my gp

Table 4.1: Training Phrases for Intents in Initial Scene

4.2.3 Types and Slot Filling

Slot filling is a process in a scene where the Assistant wants to collect some data from the user. It starts with a prompt and continues to prompt the user until the correct data has been retrieved. The data retrieved from the user must be of a certain **Type**. In Google Actions, Types allow the Assistant to extract specific data from the user's input during slot filling. They can also be used in intents however this was not done for the mood logger. Google Actions offer system types for strings, numbers, timestamps, etc (Google Developers, 2021c).

When the mood logger prompts the user for their mood, water intake, hours of sleep and activity, the data is extracted via slot filling. Since the data for hours of sleep and water intake were integers, the ‘number’ system type was used for these slots. However the mood and activity slots were slightly more complex and needed custom types. These were defined as key-value pairs which worked well for these slots as they were able to sort the different moods and exercises into specific categories. For **mood** three categories were written - *good, okay, bad* and initialised as keys. The values for each of these keys were various synonyms of each of them, therefore the user could freely say what they were feeling and the mood logger would extract their mood from their response and sort it into its respective category. For example if the user responded “I felt very anxious today”, the mood logger would recognise “anxious” in that response as a synonym of *bad* and set that as the mood. Figure 4.5 shows the key value pairs for the mood type in the actions console. When being stored in the database the mood would be the respective category of the user’s input to make it homogeneous for conducting the analysis.

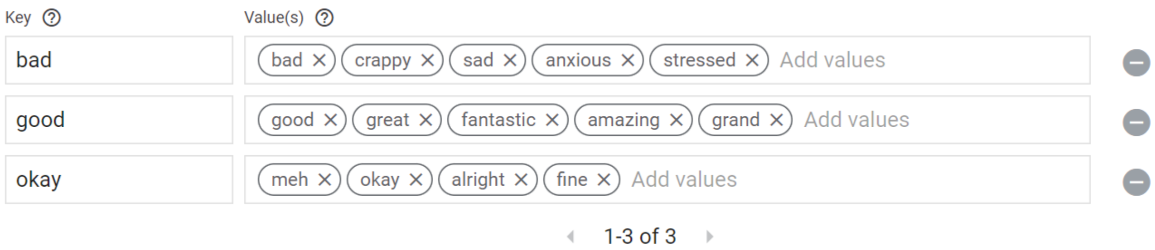

Figure 4.5: Mood Type Keys and Values

Similarly, **Activity** also had three categories - *not active, mild, moderate*. Mild activity represented activities such as walking, shopping, cleaning, etc. Moderate activity represented activities such as running, hiking, cycling, dancing, sports training, etc. Logging ‘not active’ usually indicated that the user worked from home or they had said they did not do any exercise that day. Figure 4.6 shows the key value pairs for the activity type in the actions console. Similar to the mood type, the respective category was stored during logging to help the analysis process.

Custom types could also be defined as regular expressions. A regular expression type was used in the Email GP scene if the user had not told the mood logger their GP’s email address. Slot filling would take place and a regular expression type was created to ensure the user entered the email address in the correct format[7].

[7] `https://howtodoinjava.com/java/regex/java-regex-validate-email-address/`

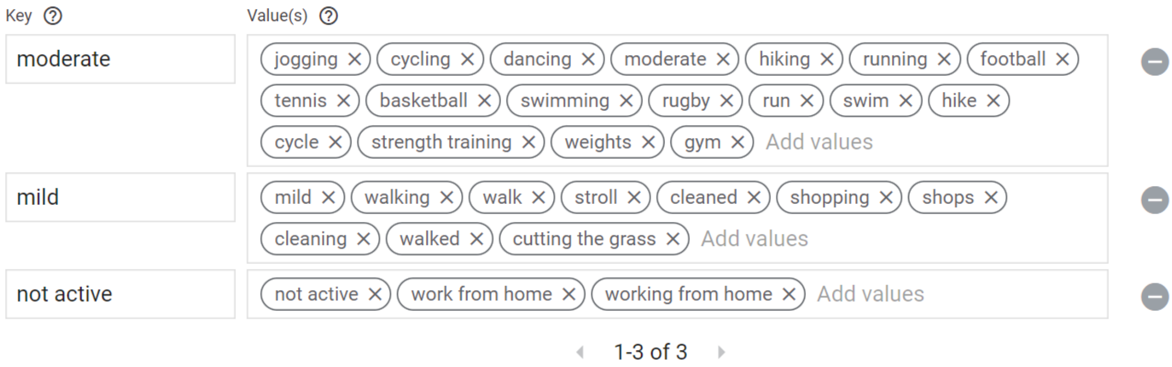

Figure 4.6: Activity Type Keys and Values

4.2.4 Webhooks

Webhooks handle the logical elements of the mood logger. They are triggered by the action and make HTTP requests to a given endpoint with a JSON payload. The endpoint acts as a webhook handler that deals with this payload (Google Developers, 2021d). Webhooks are an important element of the mood logger and were used in almost every scene. For example in the initial scene a webhook request is sent to check if the user has already logged that day. More importantly, a webhook request is sent after the logging process to store the logged data in the database. They can also be triggered when an intent is matched or when slot filling has finished. For example, when the user indicates that they want to send their data to their GP, a webhook request is sent to check whether the mood logger has the email address already. There are also some webhook requests after the slot filling process when logging to send some dynamic response to the user's logs.

These HTTP endpoints have to be deployed to the cloud in order for them to be reached. The Google Actions Console offers an inline cloud editor to write the code for these endpoints. However it proved difficult to use as there was no syntactic support and very little information was given when there was a deployment error. It was therefore a better choice to create a Firebase project for deploying these cloud functions. It provided easier deployment and there was more documentation available for connecting to the Firestore database. It was also easy to connect the project to the mood logger by providing the HTTP address in the actions console. In addition to easy deployment, Firebase also offers ESLint[8] to analyse and point out problems in the code, resulting in better and less error-prone functions. More details about the webhook handlers and cloud functions can be seen Section 4.3.

[8] https://eslint.org/

4.2.5 Session and User Storage

Google Actions also offers an easy way to keep track of variables in the mood logger. They can be saved in session storage which would then clear when the interaction with the mood logger finishes, or as user storage which would be accessible and usable across many conversations with the mood logger. An example when this is used in the mood logger is during the slot filling process when logging, in which the user's responses are saved in session storage, which can then be accessed in the webhook handlers in Firebase. User storage was used for variables that are unique to the user such as their name or email address. This information became accessible to the mood logger via Account Linking which is discussed in Section 4.2.6.

4.2.6 Account Linking

According to Google Actions, in order for the mood logger to have the ability to save and access user storage, the user would have to be verified. Verification is done through the process of Account Linking[9]. This enables the users to "sign up" to use the mood logger with their Google account. It also adds a layer of authentication for the user, assuring that their logs are not accessible to other users.

Account linking is done automatically by the Assistant once it is enabled in the actions console. A scene was created to check whether the user is verified, if so it continues as normal with the mood logger, otherwise it would change to an 'account linking scene'. It then provides information about the action to the user and they can either agree or disagree to use it. Unfortunately, since the functionality of the mood logger depends on the user agreeing to account linking, if they were to reject it they would not have the benefit of logging or getting analysis on their mood data. Once the user agrees to account linking, a token is saved in their user storage. This token contains information about the user's Google account such as their name, surname, & email address. The user's email address is the only variable from this token that the mood logger uses. Information on this is in Section 4.3.

4.2.7 All put together

To help visualise these components working together take a look at Figure 4.7 which shows the scene in the actions console for logging mood. On entering the scene, the mood logger calls a webhook to check whether the user has logged already that day by verifying a boolean in session storage. If they have not logged then it prompts the user about

[9] https://developers.google.com/assistant/identity/google-sign-in

their mood, which is when slot filling takes place. Figure 4.7 also shows that the variable *mood_today* is being assigned the user's response. In the *Condition* section of the scene it shows that it will continue to the next scene, which logs hours of sleep, once the slot filling has completed.

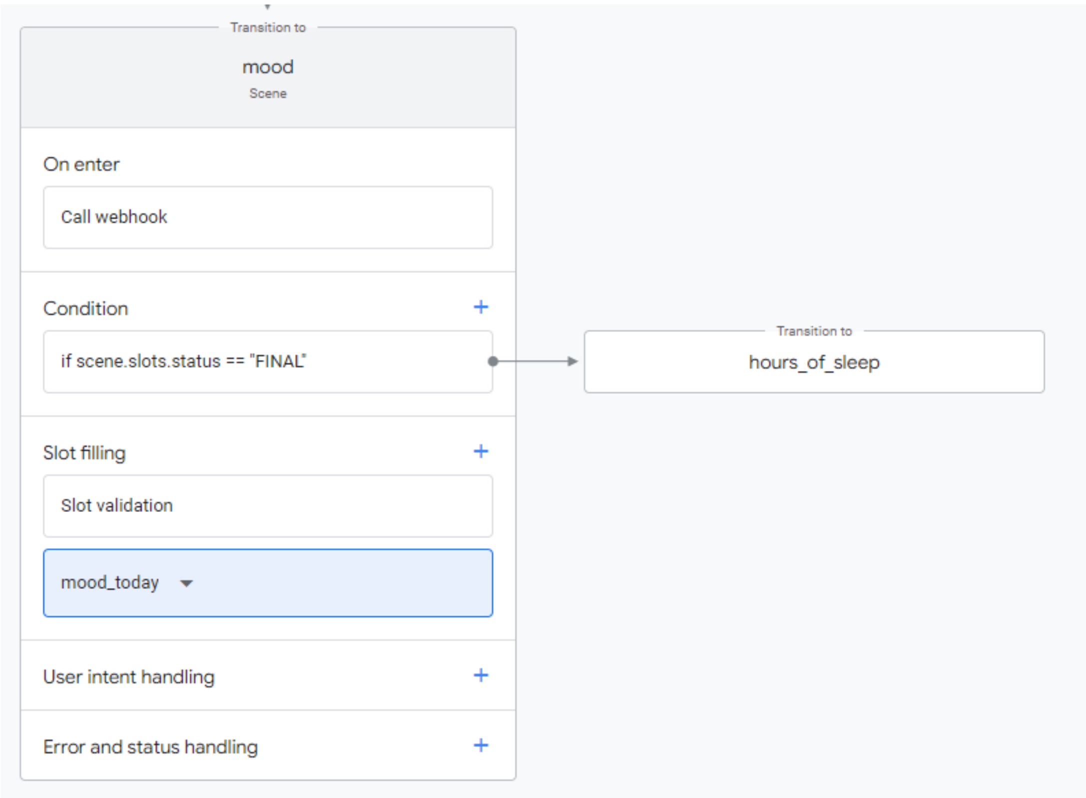

Figure 4.7: Scene that logs mood

4.3 Firebase

Firebase was used to write and deploy cloud functions that handle all the logic of the mood logger, and to store user logs and mood analyses. As mentioned in the previous section, it was the better choice compared to the inline cloud editor as it provided more support for programming and easier deployment.

4.3.1 Cloud Functions

Webhook Handlers

Webhook handlers were written to carry out logic using the payload sent in the webhook request from the actions console. The payload sent by the request contains a *conv* object, which is from the google assistant sdk for nodejs[10]. The payload is a JSON object containing the current session and user parameters from storage. Below is a list of all webhook handlers written for the mood logger and their descriptions. Figure 4.8 shows the **sendDataToDB** webhook handler described below. The handler includes the use of user and session storage saved by slot filling in the actions console.

init: Requested during the initial scene. This checks whether the user has their unique id (uid) set in user storage. If so then it retrieves the most recent analysis the mood logger has done from the database and saves it in the session storage as an array of strings, otherwise it is most likely a new user to the mood logger. It then checks whether the user's email address (from the token payload) is in the database, if not then it confirms that this is a new user, adds them to the database and initialises their uid in the user storage.

checkLog: Requested when the user intends to log their mood. This simply checks whether the user has logged that day already by searching in the database for a log on the current date, and sets the result as a boolean in session storage.

checkGPEmail: Requested when the user intends to send their GP their data. It checks in the database whether the user's GP's email address has been saved and sets the result as a boolean in session storage.

initLog: Requested on entering the mood scene, the handler checks the boolean initialised in *checkLog* and sends the corresponding prompt - *"How was your mood today?"* or *"Sorry, you have already logged your mood today."*

sendDataToDB: Requested after logging all variables, this formats the mood data that has been saved in session storage into a JSON and saves the new log in the database.

deliverAnalysis: Requested after the user intends to hear their analysis. This takes the analysis array in session storage, formats it into a string and sends returns it as a response.

setGPEmail: Requested when the user has entered their GP's email and wishes to send their mood data to them. This simply saves the entered email (saved in session storage) into the database.

[10] `https://github.com/actions-on-google/assistant-conversation-nodejs`

sendToGP: This retrieves the user's GP's email from the database and the calls a secondary function which sends the email. This function is described in the next section.

viewLog: Requested when the user intends to view a log. The handler retrieves the log from the date the user has entered (saved in session storage) from the database and sends it back as a response. If the handler cannot find a log on the given date it sends back *"Sorry there is no record of a log on that date"*.

```
app.handle("sendDataToDB", (conv) => {
  return db.collection("users").doc(conv.user.params.uid).collection("mood_data").add({
    date: admin.firestore.Timestamp.now(),
    mood: conv.session.params.mood_today,
    water: conv.session.params.water_intake,
    sleep: conv.session.params.hours_of_sleep,
    activity: conv.session.params.exercise,
  }).then((docRef) => {
    console.log("Added mood data with ID: ", docRef.id);
  });
});
```

Figure 4.8: Snippet of Webhook Handler sendDataToDB

Analysis Function

The analysis function provides an important functionality of the mood logger. It looks at a user's mood data, finds patterns and produces observations that the user can hopefully understand and learn from. Below explains the approach and implementation of the analysis.

Approach When formulating the approach to implementing the analysis, it was challenging since this was never done before and the analysis had to seem accurate to the user otherwise they wouldn't be interested in the mood logger. The goal was to produce analyses similar to Bentley et al. (2013) such as *"On days when you are busier you are happier"* and *"On Mondays you are less happy than you are the rest of the week"*. After some brainstorming the approach was developed. The idea was to select a variable and a threshold for it e.g. "sleep less than 7 hours", and retrieve each log that included this. The the most frequent mood logged within this data is calculated and this would create an observation such as - **"You are more likely to log bad moods when you sleep less than 7 hours."**. This was inspired from this article Santos (2022), in which the author analysed his data recorded in Daylio. He did much of his analysis based off of frequency which aided in constructing the method to implement the analysis for the mood logger.

Implementation The analysis function runs as a firebase scheduled function on a crob job of 7 days. The function randomly selects the threshold for each of the variables logged - sleep, water, activity. An example of which could be:

- Sleep less than 6 hours
- Moderately active
- Water intake more than 8 glasses

The thresholds are then passed into a secondary function to carry out the analysis. The function iterates through each user, retrieves their logs that match each threshold and then performs the analysis for each. The analysis is done by calculating the most frequently logged mood within each threshold. A helper function **getMostFreq** was written to return the most frequent element in a given array. An article by Demirkaya (2021) was useful when writing this function. A snippet of code of the analysis function is seen in Figure 4.9, which shows the code that analyses sleep and mood.

```
// sleep & mood
await users.doc(uid).collection("mood_data").where("sleep", lm, hours).get().then((snapshot) => {
  if (!snapshot.empty) {
    const mood = [];
    snapshot.forEach((doc) => {
      mood.push(doc.data().mood);
    });

    let symbol = "";
    if (lm == ">") {
      symbol = "more than";
    } else {
      symbol = "less than";
    }

    const topMood = getMostFreq(mood);
    analysis.push("You are more likely to log " + topMood + " when you get " + symbol + " " + hours + " hours of sleep.");
  }
});
```

Figure 4.9: Analysis Code - Sleep and Mood

Additionally to these thresholds, analysis is also carried out with the variables mixed, for example sleep and activity. This can produce observations such as **"When you sleep 7 hours, you are more likely to be moderately active."**. Another analysis calculated is the most frequent day of the week when a randomly chosen mood is logged, producing observations such as **"You usually log bad moods on Tuesday."** A code

```
// day & mood
await users.doc(uid).collection("mood_data").where("mood", "==", m).get().then((snapshot) => {
  if (!snapshot.empty) {
    const days = [];
    snapshot.forEach((doc) => {
      days.push(doc.data().date.toDate().toLocaleString("default", {weekday: "long"}));
    });

    const topDay = getMostFreq(days);
    analysis.push("On " + topDay + "s you usually log " + m + " moods");
  }
});
```

Figure 4.10: Analysis Code - Day and Mood

snippet of this can be seen in Figure 4.10. Observations can be made unrelated to mood also, such as days when the user is most active.

The analyses are collected in an array of strings which is then stored in the database. As mentioned in Section 4.3.1, the **init** webhook handler would retrieve these from the database when the mood logger starts and have them ready in session storage.

Email Function

The email function is a secondary function called from the webhook handler **sendToGP**. The parameters of this function are the GP's email, the user's email, the user's uid and the user's name. The function constructs the email in a HTML body, retrieves the user's logs and arranges them into an informative table. The email itself is sent with nodemailer[11]. Figure 4.11 shows the resulting email that is sent to the user's GP.

Dynamic Responses

An additional functionality implemented in the mood logger was dynamic responses during logging. These can improve the dialogue between the Assistant and the user as it wouldn't be saying the same thing each time a user logs, therefore enhancing the experience of the conversation. These dynamic responses were sorted into arrays for each variable logged, and a random response is selected depending on what the user inputs into the mood logger. For example there is a set of responses for good moods, bad moods, high & low water intake, a bad night's sleep, a good night's sleep, etc. If the user logs a good mood, a random response is selected from its corresponding array and sent back to the user. Table 4.2 shows some of the dynamic responses used.

In some cases the choosing the response to send back would depend on more than one variable logged from the user to ensure these responses don't negatively impact the

[11] `https://nodemailer.com/about/`

Google Assistant Mood Logger Data

Patient Name: Ciara Gilsenan

Patient Email: ciaragil98@gmail.com

Date	Mood	Water	Sleep	Activity
Mon Jul 04 2022 23:00:00	good	8	7	moderate
Thu May 26 2022 13:45:30	okay	3	6	mild
Mon Jul 11 2022 23:00:00	good	7	6	moderate
Sun Jul 24 2022 23:00:00	good	8	9	moderate
Fri Jul 29 2022 23:00:00	good	8	7	moderate
Fri Jun 24 2022 16:30:43	good	5	8	mild
Mon Jul 25 2022 23:00:00	okay	8	7	mild

Figure 4.11: Email sent from Mood Logger

Variable	Responses
Mood - Good	"That's great!", "Wonderful."
Mood - Bad	"That's too bad.", "I am sorry you are feeling this way."
Sleep $\geq$ 8 hours	"That sounds like a good night's sleep.", "You must be well rested."
Sleep $<$ 7 hours	"Hmm that doesn't sound like a lot."
Water $>$ 8 glasses	"Yay hydration!, "Nice work."
Water $<$ 7 glasses	"Hmm you may need to drink more water.", "That's all?"

Table 4.2: Mood Logger Dynamic Responses

conversation. For example, if the user logs a bad mood, but then logs more than 7 hours of sleep, it could be insensitive for the mood logger to reply *"You must be well rested"* as there would be a possibility the user had a bad night's sleep but still got 7 hours. Therefore if the user logs 7 or more hours of sleep, the webhook first checks what mood the user has logged before selecting a response.

Four additional webhook requests were written in Firebase for this - **reactToMood**, **reachToSleep**, **reactToWater**, **reactToActivity**.

4.3.2 Firestore

Firestore is Firebase's NoSQL database used to store user logs, user analyses, and the user's GP's email. It allows quick and easy querying, as well as the ability to structure the data however needed. For the mood logger the database is structured as a collection of user documents. The primary key of each user is their document id which is automatically generated by Firebase when creating a new user. This is also used as the user's uid which is stored in the mood logger in user storage. Each user document has a sub-collection for mood data and another for their analyses. It also has two fields for the user email and the GP email. Figure 4.12 shows how this structure looks in Firestore.

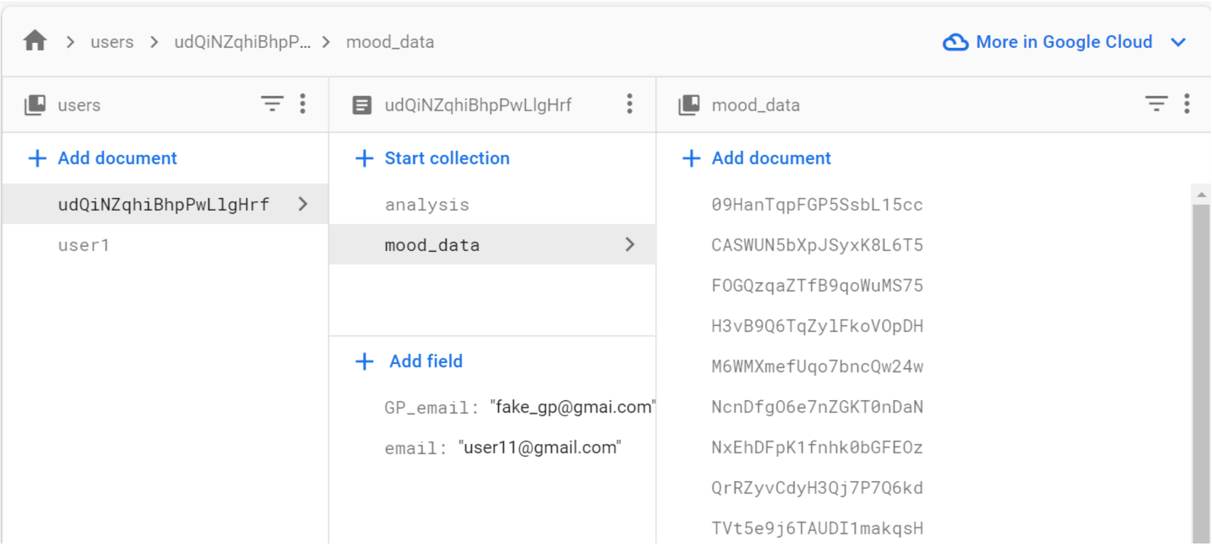

Figure 4.12: Firestore - User Document and Sub-collections

The logs are stored as documents containing 5 fields - activity, mood, sleep, water and the date of log. The analyses are stored as documents containing just two fields - the analysis array and the date of analysis. Figures 4.13 and 4.14 shows these in Firestore.

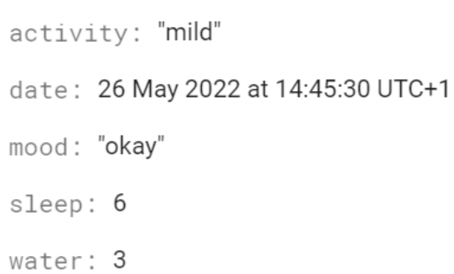

Figure 4.13: Firestore - Log Format

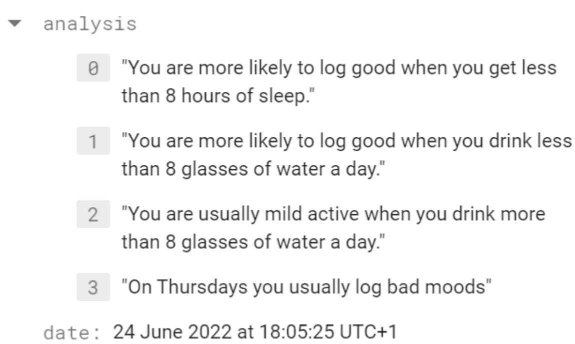

Figure 4.14: Firestore - Analysis Format

Chapter 5

Implementation - Evaluation

This chapter goes through the challenges and successes of implementing the mood logger described in Chapter 4.

5.1 Google Actions

Google Actions provided an engaging way to build conversational actions for Google Assistant. The console was simple to navigate, and once I understood how each of the components involved (intents, slot filling, scenes, etc) could work together, building the mood logger was quite interesting. Google Actions also provided documentation about these components as well as for account linking and deployment which was helpful.

5.2 Firebase

With Firebase it was initially difficult to connect the app to the Google Actions console. Particularly when trying to store the logs to the Firestore database there were issues with connecting to it and during run time the new logs were not appearing in the database. Unfortunately there was not much documentation available about using Firebase with Google Actions, therefore it was a slow start with a lot of trial and error before it became successful to store and retrieve logs from Firestore in real time with help from this tutorial[1]. Other elements of Firebase such as scheduling the cron job for the analysis function worked well, and the use of ESLint, which handled the code's syntax, lead to cleaner and better quality code.

[1] `https://firebase.google.com/codelabs/firestore-google-assistant`

5.3 Mood Logger Functionalities

5.3.1 Logging

The logging process was successfully built, achieving the short conversation that was planned in the design goals. The process involves asking four short questions for mood, sleep, water intake and activity, a snippet of this can be seen in Figure 5.1. Challenges for building this mainly involved how to store the logs during the conversation. Initially there were four webhook handlers, one for each variable, that would store the user's responses in user storage. However after implementing this it was found to be more efficient to save the responses in session storage, which took place automatically during slot filling, and have just one webhook request at the end of logging to save the responses in the database.

5.3.2 Analysis

The analysis function was successfully implemented as a weekly cron job in the Firebase app. An snippet of the result in the mood logger can be seen in Figure 5.2. The function was producing good analyses such as *"You are usually moderately active when you drink more than 8 glasses of water a day."* and *"You are more likely to log good when you get more than 8 hours of sleep"*. However some of the analyses being produced were not too accurate, example ones being *"You are more likely to log good moods when you sleep less than 8 hours"* or *"You are more likely to log good when you drink less than 6 glasses of water a day"*, since drinking less water and sleeping less would not usually lead to a better mood. However this is mainly due to the small amount of logs available in a short time before testing. If the mood logger were to be used consistently after this research, there would likely be some more accurate analyses being produced. This was the case for Bentley et al. (2013) that found that the more users logged the more accurate the analyses became.

5.3.3 Send to GP

The *sending to GP* functionality was successfully written in the Firebase app. It was completed with the help of Nodemailer and Sendinblue[2]. Upon request it sends the email to the user's GP, however the email can sometimes end up in spam or in the promotions folder. This unfortunately cannot be fixed via the code, however eventually the emails started coming into the primary folder. Figure 5.3 shows the mood logger's response to the user requesting this. The email itself can be seen again in Figure 4.11.

[2] `https://www.sendinblue.com/`

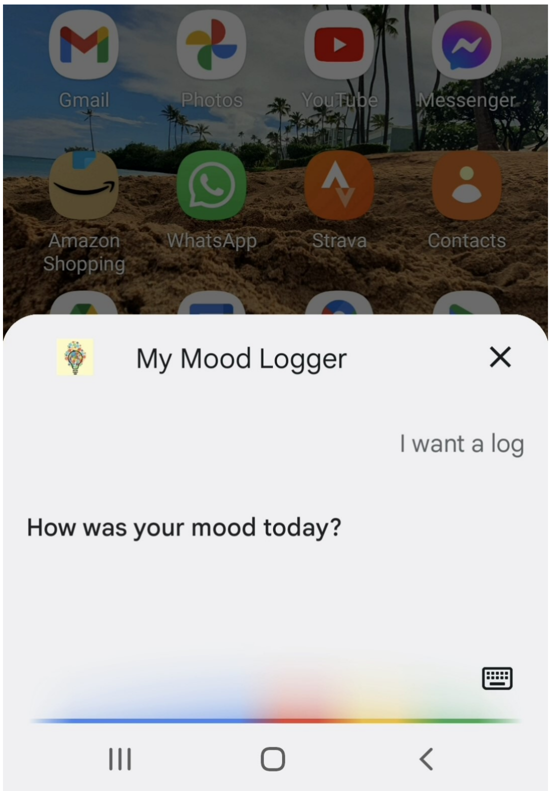

Figure 5.1: Logging

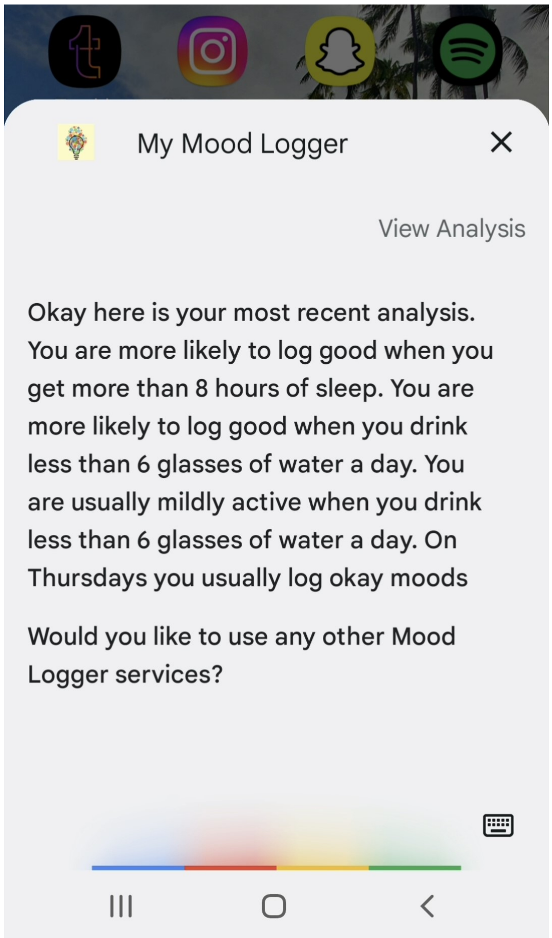

Figure 5.2: Receiving Analysis

5.3.4 View Log

Implementing the *view log* functionality had some challenges. The Actions Console has a system type for dates and it was initially unclear what format it could parse. After some testing it was found that it can parse responses such as:

- June 22nd 2021
- 22nd June 2022
- 22/6/21

However if the user responded "22nd of June 2022", the mood logger could not parse this. Another challenge in the **viewLog** webhook handler was specifying this date when querying Firestore for the log. Since the date objects stored in Firestore include the specific time of log as well as the date, it was difficult to retrieve a log since the time when querying it would have to be exactly the same as in the database. This was resolved with the help of the moment[3] package. Using this the query to Firestore was changed to find the log with a date between the start and end of the date given by the user, this way the date and time would be within the condition of the query and the correct log could be retrieved. For example if the user wanted the date “June 22nd 2022”, the query to Firestore would be between “June 22nd 2022 00:00:00” and “June 22nd 2022 23:59:59”. After implementing this the function was then successful and the result in the mood logger can be seen in Figure 5.4.

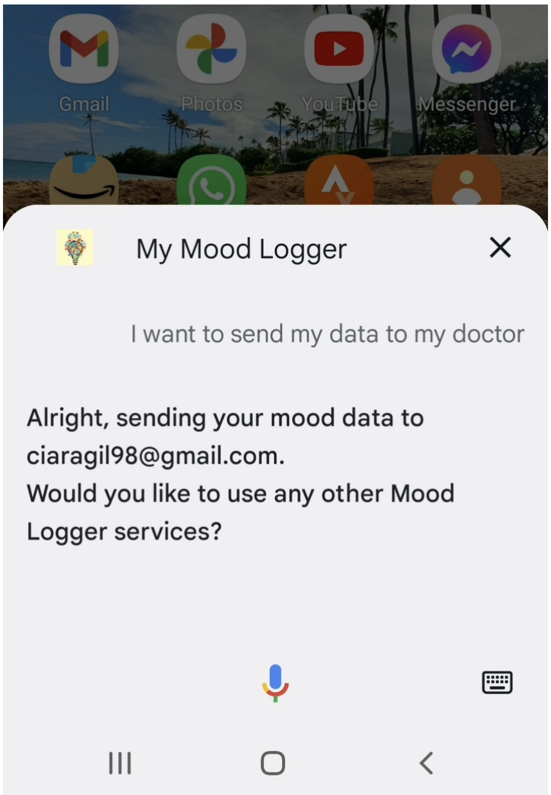

Figure 5.3: Sending to GP

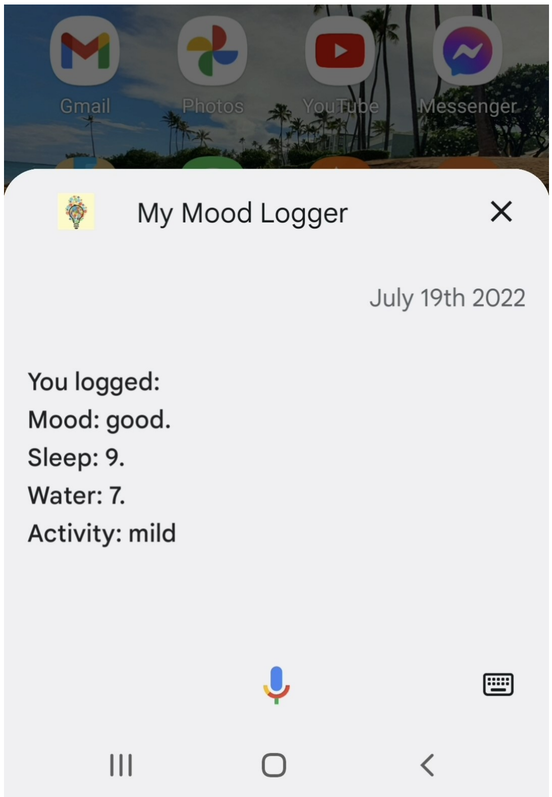

Figure 5.4: Viewing Log

[3]https://github.com/moment/moment

Chapter 6

The Study - Methodology

Along with the implementation of the mood logger, a study was conducted to explore and evaluate the user experience of the application. The study was submitted to the ethics committee and was approved, which allowed the study to collect participant's feedback.

6.1 Study Design

The purpose of the study was to evaluate the dialogue approach to mood logging and a dialogue delivery of mood analysis on fictional scenarios. It also aimed to investigate the impact of a voice based mood logger and whether participants would be encouraged to use this daily. The design of the study involved potential users interacting with the mood logger prototype and testing its capabilities, as well as a collecting data from a questionnaire.

6.2 Mood Logger for the Study

A separate version of the mood logger was developed for the study in order to protect the participants' personal data. It was similar to the one described in Chapter 4, however it only had two of its functionalities - logging and viewing an analysis. This version was also programmed to not store any information that the user inputted e.g., their logs. More on these functionalities are described in Section 6.4.2.

6.2.1 Deployment

This mood logger was also deployed to the participants through the google actions console. It offers three channels of deployment[1] - Alpha, Beta, and Production:

[1] `https://developers.google.com/assistant/console/publish`

- **Alpha**: for testing with a small group of users, does not require a review by Google before deployment.

- **Beta**: for testing with a slightly larger group of users, requires review by Google before deployment.

- **Production**: deployed to all Google Assistant users after a review by Google.

The mood logger for the study was deployed to the Beta channel which allowed it to go under review. This was chosen to ensure that there were no issues with the mood logger before distributing to participants, and to reassure the users that it was legitimate and not spam. The mood logger was approved shortly after submission for review and was then ready to be deployed to participants. This was done by adding their Gmail address to a *Beta Testers* list in the actions console, seen in Figure 6.1. Once added, the participants were sent an opt-in link provided by the actions console. They simply had to open the link and agree to be a tester, and then they could open Google Assistant on their device and start the mood logger.

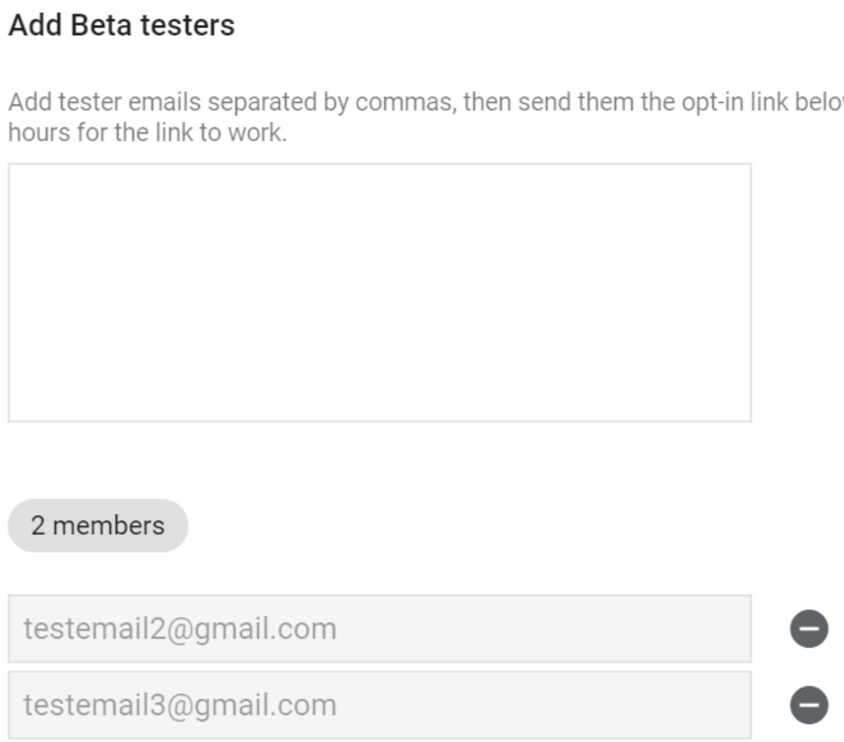

Figure 6.1: Adding testers to the mood logger for the study

6.3 Pilot Study

A pilot study took place with **two** participants in order to test its procedure. The goal of the pilot was to see how the deployment process works with the participants, if the mood logger behaves the same as it does in the testing simulator, and to make sure that the questionnaire responses are recorded correctly.

6.3.1 Findings from the Pilot

A main concern during the pilot was deployment. The mood logger worked without issue for one participant however the other participant had a hard time starting it. When invoking it, instead of the programmed greeting **"Welcome to the mood logger study"** it replied **"Alright. First, here is some details about how Google shares info with other services"**. The participant could not get past this to the mood logger. After researching why this could be the case, it was found that certain settings on the participant's device had to be enabled in order for the action to work for them.

- Web and App Activity must be turned on and a box labelled "Include Chrome history and activity from websites and apps that use Google services" must be ticked.
- Personal Results must be turned on.
- The language of Google Assistant must be set to either US English or UK English.

However after informing the participant to check and confirm these settings, the mood logger still would not start for them, which was strange. Therefore for the main study it was important to check with participants that the settings were enabled correctly before they opt-in. If it still did not work for some participants, an alternative way to test the mood logger was given by making them an *Actions Viewer* and giving them the link to the simulator that was shown in Figure 4.4.

Another observation learned from the pilot was to ensure that the participants gave their correct email address for deployment. The email had to be the same one that they would be logged into Google Assistant with. If they gave their work or college email instead, it would not work for them. Other than these issues everything else in the pilot went well. The questionnaire responses were recorded correctly and the mood logger itself worked the way it should.

6.4 Main Study

The main study started once the pilot was completed and analysed.

6.4.1 Participants

Five participants were recruited to take part in the study. The recruitment process involved posting in dissertation & research related groups on various social media platforms such as Facebook, LinkedIn and Reddit. The inclusion criteria for recruitment was having a Google account and access to Google Assistant as well as the ability to enable the necessary settings described in the previous section. Once participants registered their interest and gave their informed consent, the mood logger was deployed to them and they proceeded with the study. Demographic information about the participants can be seen in Table 6.1. The average age of the participants was 30, the youngest participant was 23 and the oldest participant was 51.

Previous Experience with Voice Assistants & Health Trackers

None of the participants had used a mood logger before this study. The majority of them occasionally used voice assistants such as Google Assistant, Alexa and Siri. The majority also didn't generally use health trackers such as Fitbits, step counters, heart rate trackers, etc. One participant said they use a tracker made by Apple.

6.4.2 Procedure

The procedure involved two interactions with the mood logger. The application prompts the user to choose which interaction to carry out, shown also in Figure 6.2. The participants could also restart the interactions at any point by saying phrases such as "go back", "restart", "start over".

- ***Interaction 1***: Participants logging mood, water intake, sleep and activity. This allowed the participants to see how mood logging would work in this dialogue format. None of the inputs from the user were stored.
- ***Interaction 2***: Participants were presented a fictional persona and an analysis of that persona. Three personas were written for the study and one is chosen as random for each participant. This interaction shows the participants what kind of analysis they would receive if they were using the mood logger daily. Below shows

Demographic	Frequency
Age	
20 - 24	2
25 - 29	1
30 +	2
Gender	
Female	3
Male	2
Qualifications	
Master's Degree	3
Ph.D	2
Degree Subject	
STEM (Science, Technology, Engineering, and Mathematics)	3
AHSS (Arts, Humanities and Social Sciences)	2
English Level	
Native	1
Near-Native	1
None-Native	3
Technology Skills	
Advanced	3
Intermediate	2

Table 6.1: Participants' Demographic Information

the personas and analyses given in the study and Figures 6.3 and 6.4 shows one of the personas being presented in the mood logger. The participants were assured before the study that any analysis presented to them was not from their own mood data, and that other participants were not able to access any information they enter into the application.

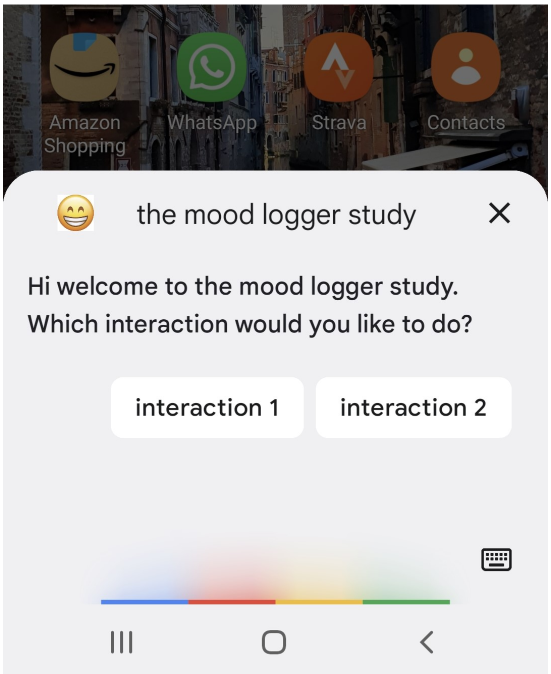

Figure 6.2: Mood Logger Study Prompt for Interactions

Interaction 2 Fictional Personas and Analyses

Jim is an accountant. He has been tracking his mood for about 6 months and finds it very helpful as his job can sometimes be quite stressful. After coming home from work he opens the mood logger and logs his mood for the day:

Mood: okay

Water: 5 glasses

Sleep: 6 hours

Activity: Moderate

After logging his mood. The mood logger notifies him that the monthly analysis of his mood data has been completed. This is what he receives:

- When you sleep less than 8 hours you are more likely to track "okay".
- You are more likely to drink more than 6 glasses of water when you exercise moderately.

- You are more likely to be at least moderately active on days you track mood "good".
- On Sunday you are less active than any other day of the week.

After receiving this information from the mood logger. Jim decides to have an early night to catch up on some sleep, hoping it will improve his mood for tomorrow.

Grace is a student. She is working a lot at the moment to prepare for exams. She started using the mood logger last month as she noticed that her mental health was declining and wanted to see if mood logging could help. After finishing her studying she logged her mood for the day.

Mood: bad

Water: 4 glasses

Sleep: 4 hours

Activity: mild

After logging her mood. The mood logger notifies her that the monthly analysis of her mood data has been completed. This is what she receives:

- You are more likely to have a poor mood when you get less than 8 hours sleep.
- You are more likely to have a good mood when you drink more than 6 glasses of water a day.
- You were less active this month than you were last month.

Grace realises that her poor sleep schedule and lack of activity might be hindering her ability to study well. She decides to increase her activity with daily walks as a study break, and to finish studying by 11pm so that she can get enough sleep.

Megan is a teacher. She tracks her mood regularly and finds the analysis/observations very helpful. She is currently trying to get at least 7 hours of sleep a night since last month the mood logger informed her that she is mostly tracking bad moods on days she is not getting enough sleep. After school she logs her mood:

Mood: good

Water: 8 glasses

Sleep: 7 hours

Activity: Moderate

After logging her mood. The mood logger notifies her that the monthly analysis of her mood data has been completed. This is what she receives:

- You track good on most days where you get at least 7 hours of sleep.
- You are less likely to track bad moods when you are moderately active.
- On Mondays you are most likely to track "meh".

Megan finds this analysis very intriguing, showing that a decent night sleep has been benefiting her mood greatly and will continue to do so. She will also continue to keep up with her activity as that also seems to benefit her mood.

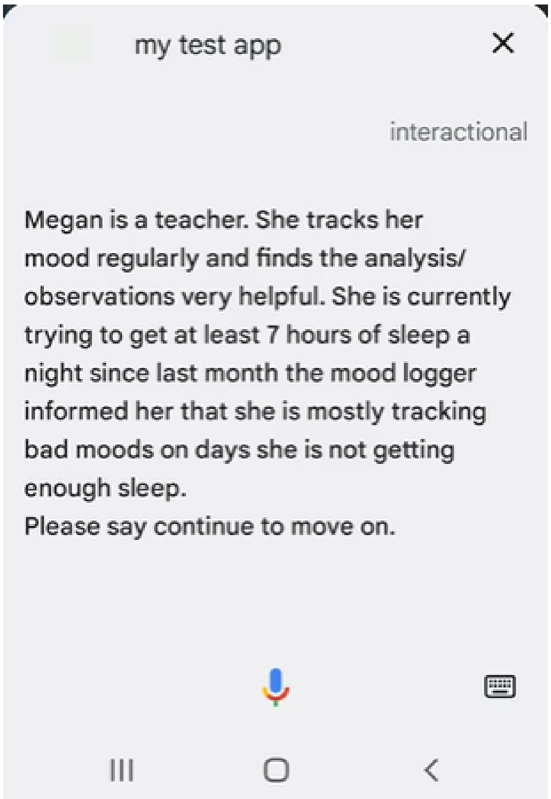

Figure 6.3: Interaction 2 part 1

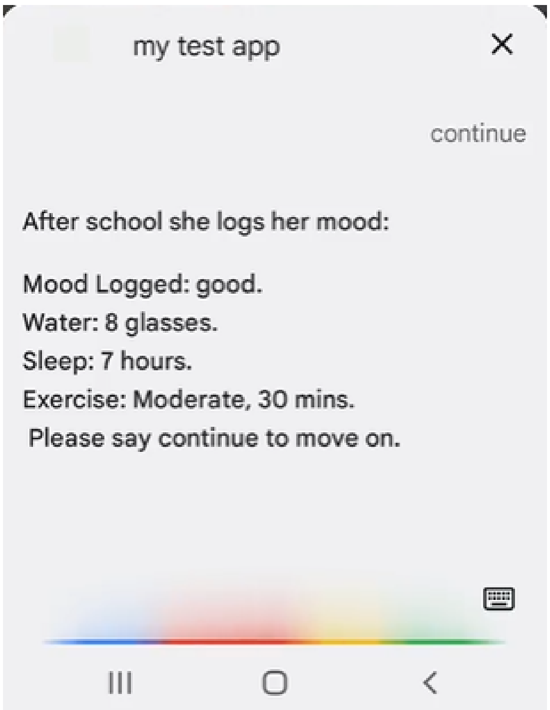

Figure 6.4: Interaction 2 part 2

6.4.3 Questionnaire

The participants could complete the interactions as many times as they pleased. Once they were finished they were given a short questionnaire that collected feedback on various aspects of the mood logger. The questions were formatted as multiple choice or 5 point rating scales with additional open ended questions. Questions covered:

- General use of voice assistants and health trackers.
- Ease of use of the mood logger.
- Likelihood of using the mood logger again.
- The types of data that was logged (water, sleep, activity).
- The mood logger's analysis.
- The dialogue of the mood logger.

6.4.4 Debriefing

Once the participants completed both interactions and the questionnaire, they were given some debriefing information about how they could delete the history of the interactions from their Google account should they wish to do so. Their email addresses were also removed from the *Beta Testers* list from Figure 6.1, so they no longer had access to it.

Chapter 7

The Study - Evaluation

7.1 Deployment

Deployment of the mood logger for the study was successful. Most of the participants were able to use it without the error that occurred in the pilot. However there were a few cases where the alternative method using the simulator had to be used. This still provided the same experience, however it would have been better to have them test it on their own devices to get a real feel for the application.

7.2 Results of the Study

This section examines the results of the questionnaire given to the participants following their interactions with the mood logger.

7.2.1 Analysis of the Study

Due to the small number of participants, it was not necessary to perform an in depth qualitative analysis on the responses. Therefore this section will describe the overall results of the questionnaire, with subsections covering its main topics. In the following subsections the participants are labelled as *"P1 - P5"*.

Mood logger Usability

Measured with a 5 point range from 'difficult to use' to 'easy to use', most participants scored the mood logger at 4 and above in its ease of use. On another 5 point scale from 'not likely' to 'likely', two out of the five participants rated that they were likely to use the mood logger again, seen in Figure 7.1. P4 commented *"It keeps track of how your* *day and what you do in it affects your mood and health overall. it's interesting to see how to improve one's health and mood lifestyle by creating good habits of sleep, drinking water and exercise."* On the other hand two participants rated neutral at 3 about using it again, P1 and P3 commenting *"The voice assistant could not understand what I was saying."* and *"Not a frequent user of Google Assistant"* respectively.

P2 did not have a good experience with the mood logger, saying *"It crashed several times, especially when the input was longer than just a few words. The speech recognition wasn't perfect. Longer statements were not possible as input ended whenever I made a short break (e.g. to breath)."* It is interesting that it worked well for most but not for this person. This can be because they might not have had the latest OS update or were not used to using Google Assistant. P1 and P2 both had issues with the Assistant's speech recognition. This could be because the mood logger's language was set to US and UK English since other dialects didn't seem necessary. From looking at the demographic information in Table 6.1 4 out 5 participants were not native English speakers, therefore it is possible that the Assistant could not distinguish some of the pronunciations of what the participants were inputting to the mood logger. Perhaps in the future it would be better to include all English dialects available in Google Actions - Australian, Canadian, Indian.

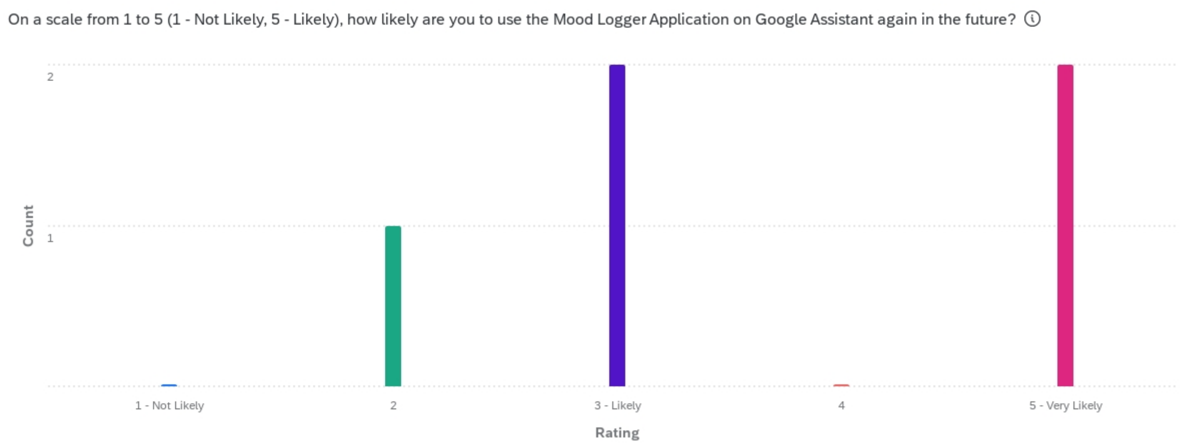

Figure 7.1: Questionnaire - Likelihood of using Mood Logger again.

Variables Logged alongside Mood

Most participants were comfortable logging data about their lifestyle - water intake, sleep and activity. This is good feedback since an important design goal discussed in Chapter 3 was to ensure that logging these variables did not have a negative influence on the participant's mood. None of the participants objected the possibility of the mood logger tracking more information in the future, which helps with any future development.

Analysis on Fictional Mood Data

For the study the participants were presented with the idea of a **monthly analysis**, which is different to the weekly analysis in the fully implemented version in Chapter 4. A weekly analysis in the full implementation was considered to be more informative to the user and at that point the study was already in motion, therefore it could not be changed.

Measured with a 5 point range from 'not great' to 'great', three out of five participants rated the mood logger's analysis at 5, seen in Figure 7.2. P1 commented *"It was very informative."* P3 commented *"It shows the possible results of good habits on your mood."* One participant rated it at 3 and P2 commented *"It seemed to be very generic"* which is understandable since an application like this is relatively new, therefore it would require further development to produce more complex analyses.

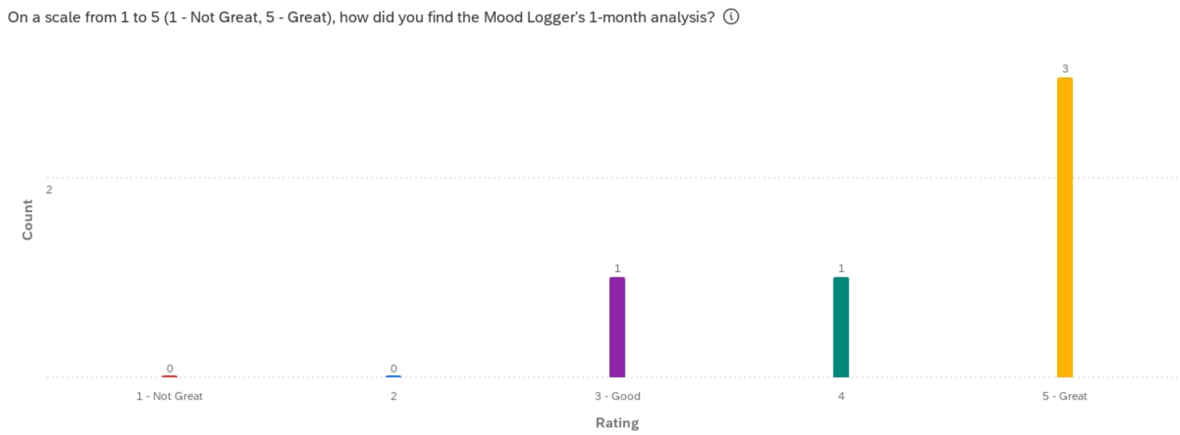

Figure 7.2: Questionnaire - Analysis

On another 5 point scale from 'not likely' to ' very likely', the majority rated 4 and above on the likelihood of being interested in receiving analysis on their own data should they use the mood logger in the future. P1 commented *"A monthly analysis would help me understand patterns in my data and get insights into my mood."* The participants understood from the analysis that mood logging can lead to understanding more about what influences your mood, and they can then build healthier habits from that. P1 and P3 wrote about the analysis *"Understand patterns of my mood, reflect on the month"* and *"Will make users feel optimistic about the results of their good habits"* respectively.

Conversational Properties

A portion of the questionnaire looked into the dialogue element of the mood logger. On a 5 point scale from "Not Good" to "Great", all participants rated 4 and above on the dialogue approach as seen in Figure 7.3, with 3 out of 5 saying that they prefer it to a text based mood logger such as journaling and charting. P1, P3 and P5 commented respectively *"Voice-based mood logging is more time efficient but I think that the quality of the assistant needs to be improved and be more accurate"*, *"Less time consuming and more consistent if its to be done daily"*, and *"Easier to talk than type"*.

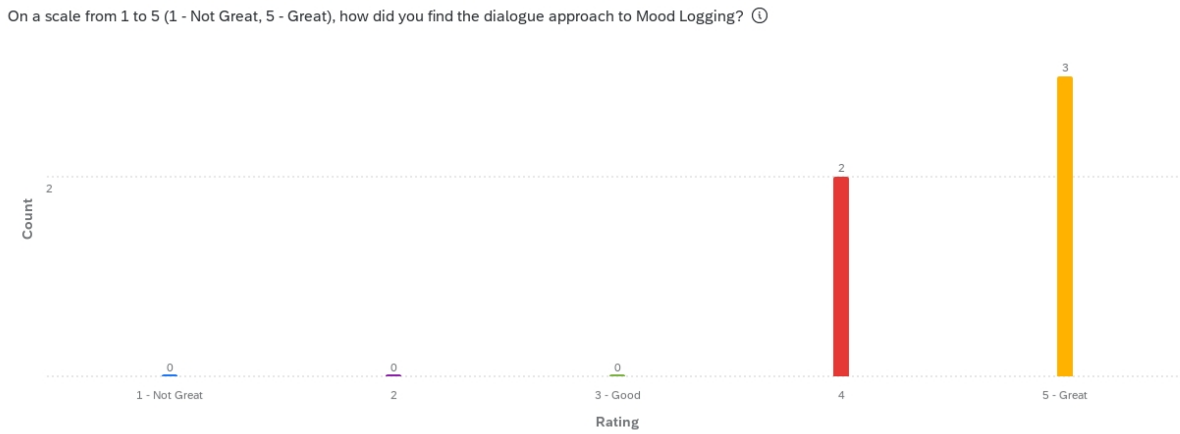

Figure 7.3: Questionnaire - Dialogue Element

P1 commented on how the Assistant's quality would need to be more accurate in the future. It's not clear here if P1 was referring to the dialogue itself or possibly the Assistant's speech recognition, however this combined with the comments in the usability section above show that there is some improvement needed in the application. The accuracy of the Assistant's speech recognition is not likely related to its Natural Language Understanding (NLU) since this deals with the Assistant's ability to parse and process natural language and not its ability to recognise the responses given by the user. The NLU is doing the best with the responses that it's been given, if the Assistant is recognising the inputs incorrectly then the NLU would not parse them in the way the user wants it to.

Chapter 8

Discussions & Conclusions

8.1 Discussion

8.1.1 What was Learned

During this research I have learned a lot about how to program and deploy applications for voice assistants which was very interesting work. In terms of implementation there was a lot to learn about how to build conversational actions for Google Assistant and to incorporate components such as intents and types, I had to take into account any possible phrases the user might say when using the mood logger so it was interesting to think about it that way. I learned how beneficial mood logging can be and how it can be an important contribution to the treatment of mental health conditions, specifically how suggesting something as simple as changing sleeping habits or adding an extra walk at the end of the day could have a big impact on your mood. I learned a lot about the process of conducting a study from start to finish as I had not done that before, including the process of applying for ethical approval and designing posts for recruitment.

8.1.2 Implications

This research is a stepping stone in developing dialogue based mood tracking applications on conversational agents and voice assistants. This adds on to previous research with voice assistants and mood tracking by also using dialogue to deliver mood data analysis. This research also offers insight to other possible ways to log mood with a simple conversation, and the range of possibly ways people can learn more about themselves and how their habits influence their mood and their symptoms. The study produced findings indicating that people liked the dialogue approach to mood logging, saying it was simple and time efficient. Participants also found the analysis informative, however there is room for improvement in its accuracy and complexity. A way to produce more informative and effective analyses could be to include more variables to log, that way more complex patterns could arise in the data. More complex analyses could also arise from simply being consistent with the mood logger as at the time of testing, there were not as many logs to use.

8.1.3 Limitations

Technical Limitations

As mentioned in Chapter 5 the main limitations during implementation was the lack of documentation on using Firebase with Google Actions, particularly when encountering errors during development it was difficult to find helpful solutions online and took some trial and error to resolve. Another unfortunate future limitation for this research is the deprecation of Google Actions in June 2023 which was announced in the Summer of 2022. If further development and deployment of the mood logger was to be considered, it would have to be transferred to another VUI such as Amazon's Alexa.

Limitations with the Study

Limitations during the study were mainly with recruitment, as it took substantially longer than planned to recruit 5 participants. This was most likely due to the lack of incentive for potential participants, therefore people were less likely to volunteer. There were also several instances where people did express their interest in the study but then did not follow up after being asked to give their informed consent to take part. The small number of participants was also due to time constraints of the study, if there was more time for the study, more participants could have been recruited which would have produced more user feedback. Another limitation was deployment, as mentioned in Chapter 7 some participants were unable to start the mood logger even after checking the appropriate settings were enables. They were still able to test the mood logger on an alternative platform, however testing it on their own devices would have been more effective in getting feedback.

8.1.4 Future Work

Since this is the first implementation of the mood logger, there is an abundance of further development that can be done. From the evaluation of the study in Chapter 7 some work can be done to improve the accuracy of the Google Assistant. Adding additional dialects to it such as Australian, Canadian and Indian English could help with this. Another option is to transfer and test the mood logger on multiple VUIs such as Alexa, Siri and Bixby to see if they offer a more efficient experience for the user. This might be inevitable should this work continue due to Google Actions being deprecated next year.

Additions can be made to the mood logger for new or potential users. Since a new user must sign in with their Google Account in order to use all the mood logger functionalities, a 'guest' version can be added to show potential users what it could do. This version would be similar to the study's version in which it would offer logging but would not store anything, and it could show potential users an example analysis. This way a user could try it out first before signing up to use it. Another option that could be implemented for users is the ability to delete their account and data upon request in the case that they no longer wish to use the mood logger.

A way to improve the email functionality is to have the mood logger automatically send the user's logs to their GP e.g. once a week. Therefore instead of the user specifically requesting for the data to be sent multiple times, they could instead tell the mood logger one time that they wish to send it every week so their GP can get regular updates.

A functionality that would have been a useful addition but could not be implemented due to time constraints was push notifications and reminders. Since Bentley et al. (2013) noted from their pilot study that reminders through silent notifications helped keep participants engaged and encouraged to log in more consecutive days, adding this to the mood logger would be helpful since it's likely users would forget to log from time to time. Notifying users when there was a new analysis on their data would also be a good addition, that way they know when to look at it.

There also could be some further work to do with the logging process. More variables could be included such as step count, menstrual cycle, medications, screen time, etc, so that users can get more information form the analyses. Participants were also not opposed to the idea of the mood logger tracking more information. However there is always a risk of adding a variable to log that could have a negative impact on the user, as mentioned in Chapter 3. A way to avoid this would be to allow new users to select what they want to log, this way the mood logger could be more tailored to the individual and people can avoid tracking information that they know might be harmful to them. This would also lead to more complex analyses being generated since there is more information to take into account. Of course, to help further development it would be useful to conduct another study with more participants who use the mood logger for a period of time e.g. 30 days. This would generate more user feedback on dialogue, what to log, analyses, etc, that future development can use to improve the mood logger further.

8.2 Conclusion

The aim of this research was to explore the design and implementation implications of a dialogue mood logger run on the voice user interface Google Assistant. This mood logger was implemented and deployed using Google Actions and Firebase, and was evaluated through a study which examined the mood logger's user experience and collected user feedback. The findings of the study indicated that the participants found the dialogue approach simple and time efficient. Participants were also interested in the analysis functionality of the mood logger, reporting that it was informative and they understood how daily logging and being consistent with good habits can have a beneficial impact on one's mood, similar to previous research done by Bentley et al. (2013) and Kelley et al. (2017). Additional findings from the study suggested that the usability of the mood logger could be improved in terms of the speech recognition which indicates that an alternative VUI could be tested to replace Google Assistant. This research offers a starting point into how mood logging can be done through a simple conversation with a VUI, and offers an abundance of possible modifications to improve its performance and efficiency. Future work can include adding more variables to log alongside mood which can lead to more complex analyses, as well as examining the technology used to facilitate the mood logger and possibly testing it on other VUIs.

Bibliography

2022a. COVID-19 pandemic triggers 25% increase in prevalence of anxiety and depression worldwide.

2022b. Mental disorders.

Andreas Balaskas, Stephen M Schueller, Anna L Cox, and Gavin Doherty. 2021. The Functionality of Mobile Apps for Anxiety: Systematic Search and Analysis of Engagement and Tailoring Features. *JMIR mHealth and uHealth*, 9(10):e26712.

Frank Bentley, Konrad Tollmar, Peter Stephenson, Laura Levy, Brian Jones, Scott Robertson, Ed Price, Richard Catrambone, and Jeff Wilson. 2013. Health Mashups. *ACM Transactions on Computer-Human Interaction*, 20(5):1–27.

K Cherry. 2021. What Is a Mood Tracker?

K. Cherry. 2022. How Cognitive Behavior Therapy Works.

Leigh Clark, Nadia Pantidi, Orla Cooney, Philip Doyle, Diego Garaialde, Justin Edwards, Brendan Spillane, Emer Gilmartin, Christine Murad, Cosmin Munteanu, Vincent Wade, and Benjamin R. Cowan. 2019. What Makes a Good Conversation? *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems.*

Bilge Demirkaya. 2021. Easiest way to Find the Most Frequent Element in Array — JavaScript in Plain English.

Gavin Doherty, David Coyle, and John Sharry. 2012. Engagement with online mental health interventions: An exploratory clinical study of a treatment for depression. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '12, page 1421–1430, New York, NY, USA. Association for Computing Machinery.

Google Developers. 2021a. Conversational Actions — .

Google Developers. 2021b. Scenes — Conversational Actions — .

Google Developers. 2021c. Types — Conversational Actions — .

Google Developers. 2021d. Webhooks — Conversational Actions — .

Google Developers. 2022. Intents — Conversational Actions — .

Christina Kelley, Bongshin Lee, and Lauren Wilcox. 2017. Self-tracking for mental wellness: Understanding expert perspectives and student experiences. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, page 629–641, New York, NY, USA. Association for Computing Machinery.

Rafal Kocielnik, Daniel Avrahami, Jennifer Marlow, Di Lu, and Gary Hsieh. 2018a. Designing for Workplace Reflection. *Proceedings of the 2018 Designing Interactive Systems Conference.*

Rafal Kocielnik, Lillian Xiao, Daniel Avrahami, and Gary Hsieh. 2018b. Reflection Companion. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(2):1–26.

Raju Maharjan, Darius Adam Rohani, Per Bækgaard, Jakob Bardram, and Kevin Doherty. 2021. Can we talk? Design Implications for the Questionnaire-Driven Self-Report of Health and Wellbeing via Conversational Agent. *CUI 2021 - 3rd Conference on Conversational User Interfaces.*

Jamie M. Marshall, Debra A. Dunstan, and Warren Bartik. 2019. The Digital Psychiatrist: In Search of Evidence-Based Apps for Anxiety and Depression. *Frontiers in Psychiatry*, 10.

Mark Matthews and Gavin Doherty. 2011. In the mood: Engaging teenagers in psychotherapy using mobile phones. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, page 2947–2956, New York, NY, USA. Association for Computing Machinery.

Northstar Transitions. 2021. Can Drinking Water Help Improve Mental Health?

Mike Palmer. 2022. Rise of the Voice User Interface.

Juan De Dios Santos. 2022. I've Tracked My Mood for Over 1000 Days: A Data Analysis.

Appendix

Demo Video: https://drive.google.com/file/d/1qgcSigmtSGEX6q1AQG4KBZimf0VZPYo9/view?usp=sharing

Firebase Github: https://github.com/CiaraG98/MoodLoggerBackend

Google Actions Github: https://github.com/CiaraG98/MoodLoggerAction

Questionnaire Attached Below

Introduction

Introduction

Thank you for testing the mood logging application on Google Assistant!

This questionnaire will ask you for your opinion on different aspects of the mood logger than you tested.

All of these questions are optional and you can exit the questionnaire any time without saving.

Don't forget to copy you anonymous ID.

Anonymous ID

Before you start

Here is your anonymous ID: ${e://Field/Random%20ID}

Please copy this number and save it somewhere. You can use this if you would like to withdraw your survey response.

Demographic Information

First some demographic information

What gender do you identify as?

- ○ Female
- ○ Male
- ○ Non-Binary
- ○ Prefer not to specify
- ○ [ ] Prefer to self-describe

What is your age?

What is the highest degree you have completed or will complete?

- Bachelor's Degree
- Master's Degree
- Ph.D. or higher
- Other

What subject do you study?

- Arts Humanities and Social Sciences (AHSS)
- Science, Technology, Engineering, and Mathematics (STEM)
- Health Sciences (HS)
- Other

What level of English do you have?

- Native
- Near-Native
- Non-Native

How would you describe your technology skills?

- Advanced
- Intermediate
- Beginner

Mood Logger

Mood Logger

Have you ever used an application like this before?

- Yes
- No

What application was it?

After trying the mood logger, do you see yourself using an application like this in the future?

- Yes
- No
- Maybe

How often would you say you use voice assistants?

- Every day
- Occasionally
- Never

Which voice assistant/s do you use?

How often would you say that you use health trackers?

- Every day
- Occasionally
- Never

What health tracker/s do you use?

On a scale from 1 to 5 (1 - Not Good, 5 - Good), how did you find the Mood Logger application on Google Assistant overall?

1 2

3 4

5

On a scale from 1 to 5 (1 - Difficult to use, 5 - Easy to use), how would you describe the use of the application?

1 2

3 4

5

Were you comfortable logging information about your lifestyle? (exercise, sleep, workload)

- Yes
- No

On a scale from 1 to 5 (1 - Not Likely, 5 - Likely), how likely are you to use the Mood Logger Application on Google Assistant again in the future?

1 2

3 4

5

Can you explain why?

Would you prefer the mood logger to be able to log more data about your lifestyle, e.g. daily steps, menstrual cycle, medications, weight, phone usage, etc.

- No
- Maybe
- Yes

Do you think mood logging can be beneficial to people's mental health?

- Yes
- No

Analysis

1-Month Analysis

On a scale from 1 to 5 (1 - Not Great, 5 - Great), how did you find the Mood Logger's 1-month analysis?

1 2

3 4

5

Can explain say why?

On a scale from 1 to 5 (1 - Not Likely, 5 - Very Likely), if you used this application daily in the future, would you likely also want to receive the analysis monthly?

1 2

3 4

5

Can you explain why?

Do you think this analysis could be an important part of Mood Logging?

- Yes
- No

Can you explain why?

Dialogue Approach

On a scale from 1 to 5 (1 - Not Great, 5 - Great), how did you find the dialogue approach to Mood Logging?

1 2

3 4

5

How did you find the length of the conversation?

Too Short Fine Too Long

Which of the following do you think would be a more effective form of mood logging?

Voice-based Mood Logging

Text-based Mood Logging

Can you explain why?