Human-Intelligent Agent Interaction for behaviour change

CM3203 Final Report

Conor Kemp | C1821404

Abstract

Human agents are capable of giving advice to other humans through the use of their natural intelligence. Some 'AI assistants', such as Alexa, are examples of Intelligent Agents (IA) because they use sensors to receive requests and then accomplish their goals by returning results. These IAs are also capable of providing advice to a human, which results in the question: Are humans more likely to accept or refuse an IA's advice in comparison to a human agent?

The aim of this project is to identify what physical and interactive attributes are crucial for an IA to possess in order to be considered trustworthy for providing behaviour change advice. Identifying and cataloguing the impact of these attributes will allow developers to implement IA systems that are deemed trustworthy by a user. This study involved the development of an IA, with variable attributes, for a week-long study with six participants who would be given advice by the IA. The verdict of the advice (if the advice was followed) would then be used to determine if the IA was deemed trustworthy.

The results of the study concluded that participants would prefer to receive advice twice per day and would prefer the gender of an IA to be neutral. Participants also preferred to receive a varied amount of advice and receiving the same advice everyday caused participants feel unlikely to follow the advice given. There was some evidence that receiving the same advice everyday caused participants to be more likely to question the Analysis method of the IA. However, further testing on this should be done to corroborate it.

Acknowledgements

I would like to thank my supervisor Dr Parisa Eslambolchilar for her support throughout this process. I would also like to thank the participants of the study for their patience and interest in this project. Without them this would not have been possible.

Table of Contents

Table of Figures and Tables

Introduction

Within this study Intelligent agent (IA) will be referred to as “a system that acts intelligently” (David Poole et al 1998) and the attributes of an IA would be the key characteristics that define it. IA would be a more apt name for systems that are often called Artificial Intelligence (AI); systems like Siri, Google Assistant and Alexa. These systems were developed with the aim of aiding a user in their day-to-day life (Google 2016). The marketing of these products referred to them as AI. However, they were implementations of IA because they act intelligently, using sub-fields of AI such as machine learning or natural language processing (NLP) to process and complete user requests (Google 2016).

The aim of this project was to identify what physical and interactive attributes were crucial for an Intelligent Agent to possess in order to be deemed trustworthy. By identifying these key attributes, IA systems can be designed to include these attributes and thus be more trustworthy to a user. Being trustworthy would be beneficial to an IA system because a user will interact and use the system more often as well as be more likely to follow advice given by the IA system.

Two attributes were analysed in this study period, the Analysis method and Advice type. The Analysis method was the method in which the user thought that the IA was using to generate advice. The Analysis method compared machine learning against traditional programming. Although the advice generated by the IA always used traditional programming when providing advice. The second attribute that was tested was the Advice type. This involved the comparison of advice given to the participants with and without justification, to see whether advice was more likely to be followed if justification was given.

Another two attributes were analysed using the participants’ answers to the questions in the questionnaires given at the beginning and end of the study period. These attributes were the Gender of the IA and the Interaction amount.

Background

Existing research has been done on the field of advice given by a non-human entity, such as determining when advice from a non-human entity was not followed (Prahl and Swol 2017) or determining when advice from an AI was deemed preferable to advice from a human (Longoni and Cian 2020). However, a study on the Analysis and Advice type attributes that an IA's should possess in order to be deemed trustworthy by a user has not been seen previously in the literature. The closest study to this was by Liao et al (2021), which investigated an IA's ability to justify its decision making. This study was sufficiently different to the current study as it involved the testing of a value driven agent whereas the IA in the current study had no success criteria, with its only aim being to provide advice to the user.

The market split in the UK in 2020 was 49.88% iOS and 49.87% Android (Shanhong 2021). However, the market share worldwide is split 70.93% Android and 27.47% iOS (O'Dea 2021). This information was used to decide on the mobile platform used to create the application used in this study.

Trust is a field that is well investigated in Psychology, from investigating the difference between trust and trustworthiness (Ben-Ner and Putterman 2001; Hardin 2002), to investigating how we measure trust (Ben-Ner and Halldórsson 2010; Özer et al 2018). Identifying trust in an IA has been found to be difficult because human trust is complicated. Trust has multiple dimensions: who trusts, who is trusted and what the goal of this trust is (Tschopp and Ruef 2020a). A user might not trust the IA in its entirety but might trust them in a limited scope. Therefore, to identify trust in an IA would be better to ask if the user trusted the IA in narrower scope such as "Do you trust the IA to do X?" (Tschopp and Ruef 2020a) or "Did you trust the advice given to you by the IA?". This was used when writing the questions for the questionnaires used in the study [Figure 1 and 2; Appendix 1 and 2].

When talking about trusting computers it is often said that transparency is key. However, the opposite would likely be the case as when a computer would be designed to be transparent one would be giving up on trust and instead gaining control (Botsman 2020). In order for a user to trust an IA, they need to learn how to think critically about IAs, such as how far they should trust IAs and where to set boundaries (Tschopp and Ruef 2020b). A user's trust in an IA has to be earned by the IA (Tschopp and Ruef 2020b). The IA should be transparent where applicable, so that the user can begin to trust the IA by having some control due to understanding how and why their data might be being used. Ultimately if a user is to fully trust an IA, then the user will need to take steps towards trusting the IA, sometimes with blind faith.

One hypothesis that will be tested in this study will be if the users have a preference for or against the idea of machine learning being used to provide advice. There will likely be no discernible difference for the Analysis attribute (Machine Learning vs Traditional Programming). The Advice type attribute (With Justification vs No Justification) however will likely see a stark difference, with IA's having the "With Justification" attribute seeing more advice given being followed when compared to the IA's with the "Without Justification" attribute. Another hypothesis is that participants will also be more likely to prefer a female IA when compared to a male or neutral IA. The final hypothesis for this project is that participants will want to receive advice from the IA once per day. These should be shown in participant answers to the questionnaires.

- Question 1
    - How likely are you to follow the advice given by this Intelligent Agent?
        - Very Likely
        - Likely
        - Unlikely
        - Very Unlikely
- Question 2
    - How often would you like to receive advice?
        - Once Per Day
        - Twice Per Day
        - Four Times Per Day
- Question 3
    - What gender would you prefer the Intelligent Agent to be?
        - Male
        - Female
        - Neutral
- Question 4
    - What method of communication would you prefer the Intelligent Agent to use?
        - Text
        - Speech
        - Text & Speech
- Question 5
    - Are you likely to follow advice given by this Intelligent Agent if a reason for the advice is given?
        - Yes
        - No
- Question 6
    - Do you trust applications that have machine learning implemented? (Machine learning refers to algorithms that improve themselves overtime through experience and data acquisition.)
        - Yes
        - No

[Figure 1. First Questionnaire. This is an example of the questions that the participants of the study were given in the first questionnaire at the start of the study period. The questionnaire was provided to the user within the Guidance after they had activated their assigned IA]

- Question 1
  - How likely would you to continue to follow the advice given by this Intelligent Agent?
    - Very Likely
    - Likely
    - Unlikely
    - Very Unlikely
- Question 2
  - How often would you have liked to receive advice from the Intelligent Agent during the study?
    - Once a week
    - Once Per Day
    - Twice Per Day
    - Four Times Per Day
- Question 3
  - What gender would you have preferred the Intelligent Agent to be over the course of the study period?
    - Male
    - Female
    - Neutral
- Question 4
  - What method of communication would you prefer the Intelligent Agent to have used during the study period?
    - Text
    - Speech
    - Text & Speech
- Question 5
  - Did you trust the advice given by the Intelligent Agent?
    - Yes
    - No
- Question 6
  - Do you trust applications that have machine learning implemented? (Machine learning refers to algorithms that improve themselves overtime through experience and data acquisition.)
    - Yes
    - No
- Question 7
  - Did you determine that the Intelligent Agent was not using machine learning? (Machine learning refers to algorithms that improve themselves overtime through experience and data acquisition.)
    - Yes
    - No

- Question 8
  - How do you feel about the Intelligent Agent using traditional programming techniques instead of machine learning to provide advice? (Machine learning refers to algorithms that improve themselves overtime through experience and data acquisition.)
    - Comfortable
    - Uncomfortable
- Question 9
  - How did you feel about the device monitoring you 24/7?
    - Comfortable
    - Uncomfortable
- Question 10
  - How much of the advice do you think you followed?
    - None of it
    - Nearly none of it
    - Most of it
    - All of it
- Question 11
  - Would you prefer to follow advice given by a human over advice given by the Intelligent Agent in this application?
    - Yes
    - No

[Figure 2. Second Questionnaire. This is an example of the questions that the participants of the study were given in the second questionnaire at the end of the study period. The questionnaire was provided to the user after the study period had ended and as they were going to upload the usage data generated]

Approach

In order to achieve the project's aim, a platform needed to be developed to allow a user to interact with an IA. Therefore, a bespoke Android based mobile application called Guidance was created for use in this study [Appendix 5]. This application was designed to collect data and then use it to provide advice to the user. This advice would then be checked against additional data collected to determine if the advice given to the user was followed.

A key design point of this platform was that it needed to be modular, allowing for attributes to have multiple options implemented whose visibility could be changed depending on the specific attribute being tested. This would allow for multiple users to download a single application and activate different IAs with different attributes by entering an assigned passcode. Determining the key attributes would then be done by comparing each IA against each other. IAs that provided the most amount of followed advice would be the IAs that were the most trustworthy.

Guidance was implemented as an Android application because having a more accessible application allowed more participants to partake in the study. Furthermore, the Android Operating System (OS) supports installing third party applications without the use of the Google Play Store, meaning that participants could be emailed a link to download Guidance's Android Application Package (APK) file which could then be used to install Guidance. Apple's iOS does not support the installation of third-party applications that are not on the Apple app store (Nield 2020). To install a third-party application that is not on the Apple app store on an iOS device would either involve Jailbreaking the device or using TestFlight (Nield 2020). Jailbreaking is a technique that is not recommended as it potentially exposes the device to unauthorised access (Nield 2020). Installing a third-party application via TestFlight would involve the user downloading the TestFlight application and then being invited to a "beta test" of the application (Nield 2020), which would be more problematic than using an Android implementation.

The only inclusion and exclusion criteria for this project were that participants were required to own an Android phone with Android 7.0 or higher installed. All eligible participants that enquired between 12/04/2021 and 19/04/2021 regarding the study were forwarded the participant information sheet (PIS) [Appendix 3] and consent form [Appendix 4] and included in the study if they agreed to take part and signed the consent form.

Ethical approval was given for this project by the School of Computer Science & Informatics Research Ethics Committee. The SREC reference for this is: COMSC/Ethics/2021/036. Participants of the study were recruited through email and a post on Yammer. Emails were sent to various Cardiff University School offices, requesting the forwarding of a participation request. An example of the email template used is shown in figure 3. Participants were then given a week to read the PIS [Appendix 3] and sign the consent form [Appendix 4] for the project. The Yammer post followed the same template as the emails and is shown in figure 4. The project was then started with the participants that had signed the consent forms. The projects duration was 7 days.

Alongside the attribute testing, the participants of the study were also given two questionnaires, one at the start of the study [Figure 1] and one at the end [Figure 2]. These questionnaires were provided to the user within the application and aimed to identify if the usage data generated correlated with the participants answers to the questionnaires. The questionnaires also aimed to determine if the user's opinion of IAs changed over the study period as well as to determine the user's preference for the Gender and Interaction attributes.

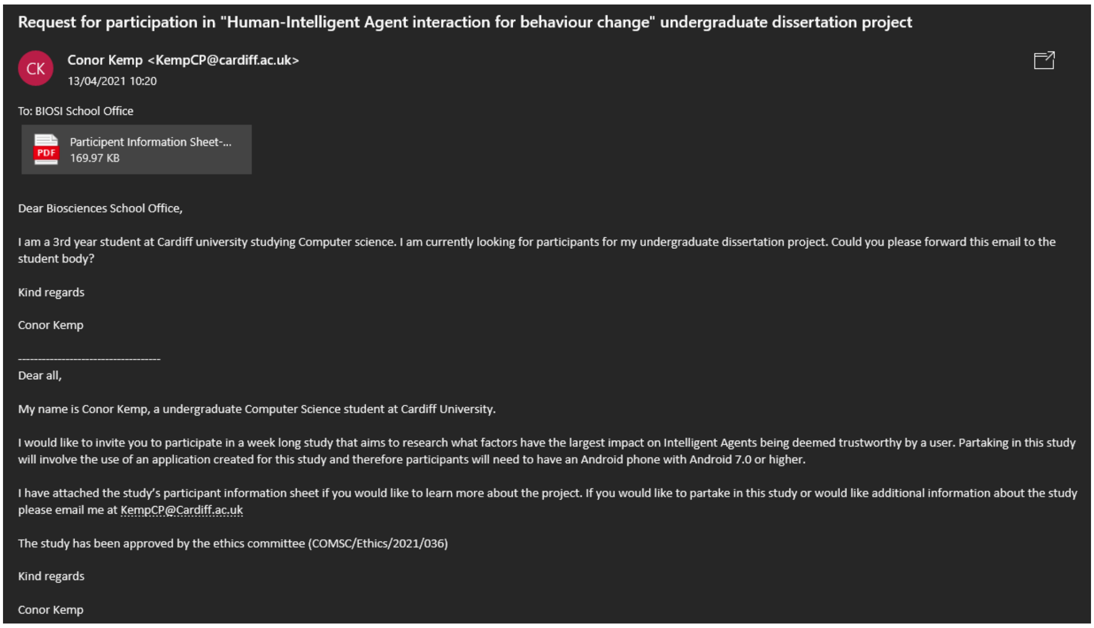

Request for participation in "Human-Intelligent Agent interaction for behaviour change" undergraduate dissertation project

CK Conor Kemp <KempCP@cardiff.ac.uk>
13/04/2021 10:20

To: BIOSI School Office

Participant Information Sheet-...
169.97 KB

Dear Biosciences School Office,

I am a 3rd year student at Cardiff university studying Computer science. I am currently looking for participants for my undergraduate dissertation project. Could you please forward this email to the student body?

Kind regards

Conor Kemp

Dear all,

My name is Conor Kemp, a undergraduate Computer Science student at Cardiff University.

I would like to invite you to participate in a week long study that aims to research what factors have the largest impact on Intelligent Agents being deemed trustworthy by a user. Partaking in this study will involve the use of an application created for this study and therefore participants will need to have an Android phone with Android 7.0 or higher.

I have attached the study's participant information sheet if you would like to learn more about the project. If you would like to partake in this study or would like additional information about the study please email me at KempCP@Cardiff.ac.uk

The study has been approved by the ethics committee (COMSC/Ethics/2021/036)

Kind regards

Conor Kemp

[Figure 3. Participation Request Email. Example of an email that was sent requesting a Cardiff University School office to forward to the student body. The lower half of the email was also posted on Yammer]

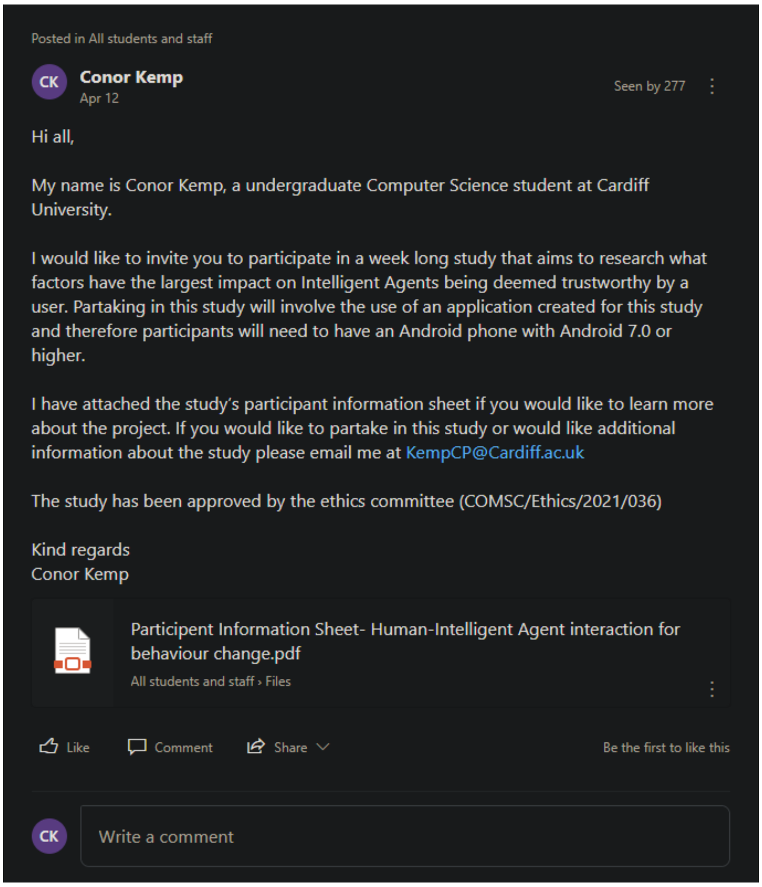

Posted in All students and staff

CK **Conor Kemp**
Apr 12

Seen by 277

Hi all,

My name is Conor Kemp, a undergraduate Computer Science student at Cardiff University.

I would like to invite you to participate in a week long study that aims to research what factors have the largest impact on Intelligent Agents being deemed trustworthy by a user. Partaking in this study will involve the use of an application created for this study and therefore participants will need to have an Android phone with Android 7.0 or higher.

I have attached the study's participant information sheet if you would like to learn more about the project. If you would like to partake in this study or would like additional information about the study please email me at KempCP@Cardiff.ac.uk

The study has been approved by the ethics committee (COMSC/Ethics/2021/036)

Kind regards
Conor Kemp

Participant Information Sheet- Human-Intelligent Agent interaction for behaviour change.pdf
All students and staff › Files

Like Comment Share Be the first to like this

CK Write a comment

[Figure 4. Yammer Participation Request Post. The Yammer post used to recruit participants for the study]

Specification & Design

Guidance's requirements were that it will need to be able to store data which will then be used by an IA to provide advice to the user. The data types stored were split into four different categories: Sensor, External, System and User Inputted. The Sensor category contained user generated data, collected through the device's sensors. This involved data that the user had created through their normal routine. The data types for this category were Step, Location and Ambient Temperature. The External category involved environmental data for each day. This data was collected through the use of third-party services such as OpenWeatherAPI (OpenWeather 2021). The data types for this category were Weather, External Temperature, Sunrise and Sunset time. The System category involved user generated data, collected automatically through the Android OS. The only data type in this category was Screentime usage data. Finally, the User Inputted category involved data types where the user inputted a rating directly. The data types in this category were the user's rating of their socialness and mood for that day.

The user had the option of enabling or disabling the storage of each data type. When a user enabled or disabled a data type it was recorded. Identifying which data types were deemed sensitive by the users would help developers because they could provide more information regarding the operation of these data types. This might then convince the user to enable the data type.

Guidance was designed to operate in the background of the user's device, only going into the foreground to display notifications to the user. These notifications either gave the user advice or notified the user that a data type was currently being stored by the application. An example of how the user would interact with Guidance is shown in figure 5.

The user interface (UI) of Guidance allowed the user to view the advice that they had been given by the IA, information about the IA's attributes, user information, the data types that were currently enabled or disabled, their advice ranking preferences, and input data relating to their socialness and mood (depending on whether either of the aforementioned data types were enabled). The user was able to view the advice for their current day on both the Main Menu screen and the Advice screen. The Advice screen also displayed the advice that they had received in the past and the advice that they would have received for a specific day in the future. The user was also be able to view the attributes of the IA that they had been assigned on the "Intelligent Agent Properties" screen.

If the user had enabled either the Socialness or Mood data types then they would have received a "Daily Question" notification each day, requesting they enter data relating to these data types. Clicking on this notification would have brought the user to the Daily Questions screen where they could then choose a rating of how they were feeling regarding their socialness and their mood. If the user had only selected one of these data types, then only that one will be displayed. If the user had not selected either of these data types then they would not receive a Daily Question notification and would not have been able to access the Daily Question screen. The process of how a user would interact with the system in this regard is shown in figure 6.

Initially the UI available to the user was going to contain five screens being a Main Menu, Settings, Data Types, Justification and Advice as shown in figure 7. However, the Settings and Justification screens were removed from the design of the application. Settings was no longer needed, as there were no settings to adjust within the application that would need to be in a dedicated Settings screen. Instead of a Settings screen there was a User Information screen [Figure 8] and a Ranking Preference screen [Figure 9]. This was done because understanding the gender and age of a user would be beneficial towards determining if there was a preference to following advice based on gender or age. Thus, the User Information screen would allow the user to optionally add their age and/or gender. Additionally, the User Information screen also allowed the user to optionally enter their name which was then used when providing advice, by addressing the user by their name. The Ranking Preference screen allowed the user to set a ranking preference for which data types took precedence. This was used by the IA when providing advice so that the advice given was relevant to the participant. Justification was removed from the application due to a design decision to have justification shown to the user in the advice message.

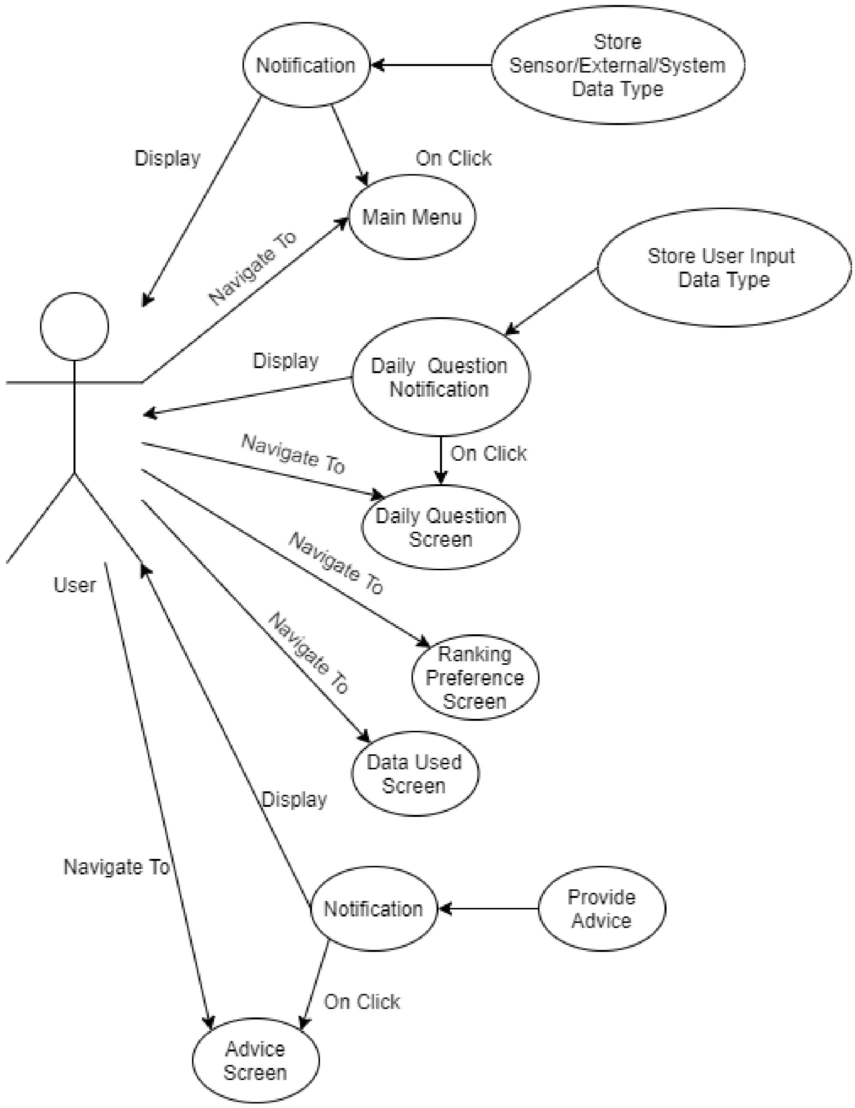

[Figure 5. User Interaction with Guidance application. UML diagram displaying why the user would receive notifications, what happens if the user clicks on a notification and how the user would access the screens normally]

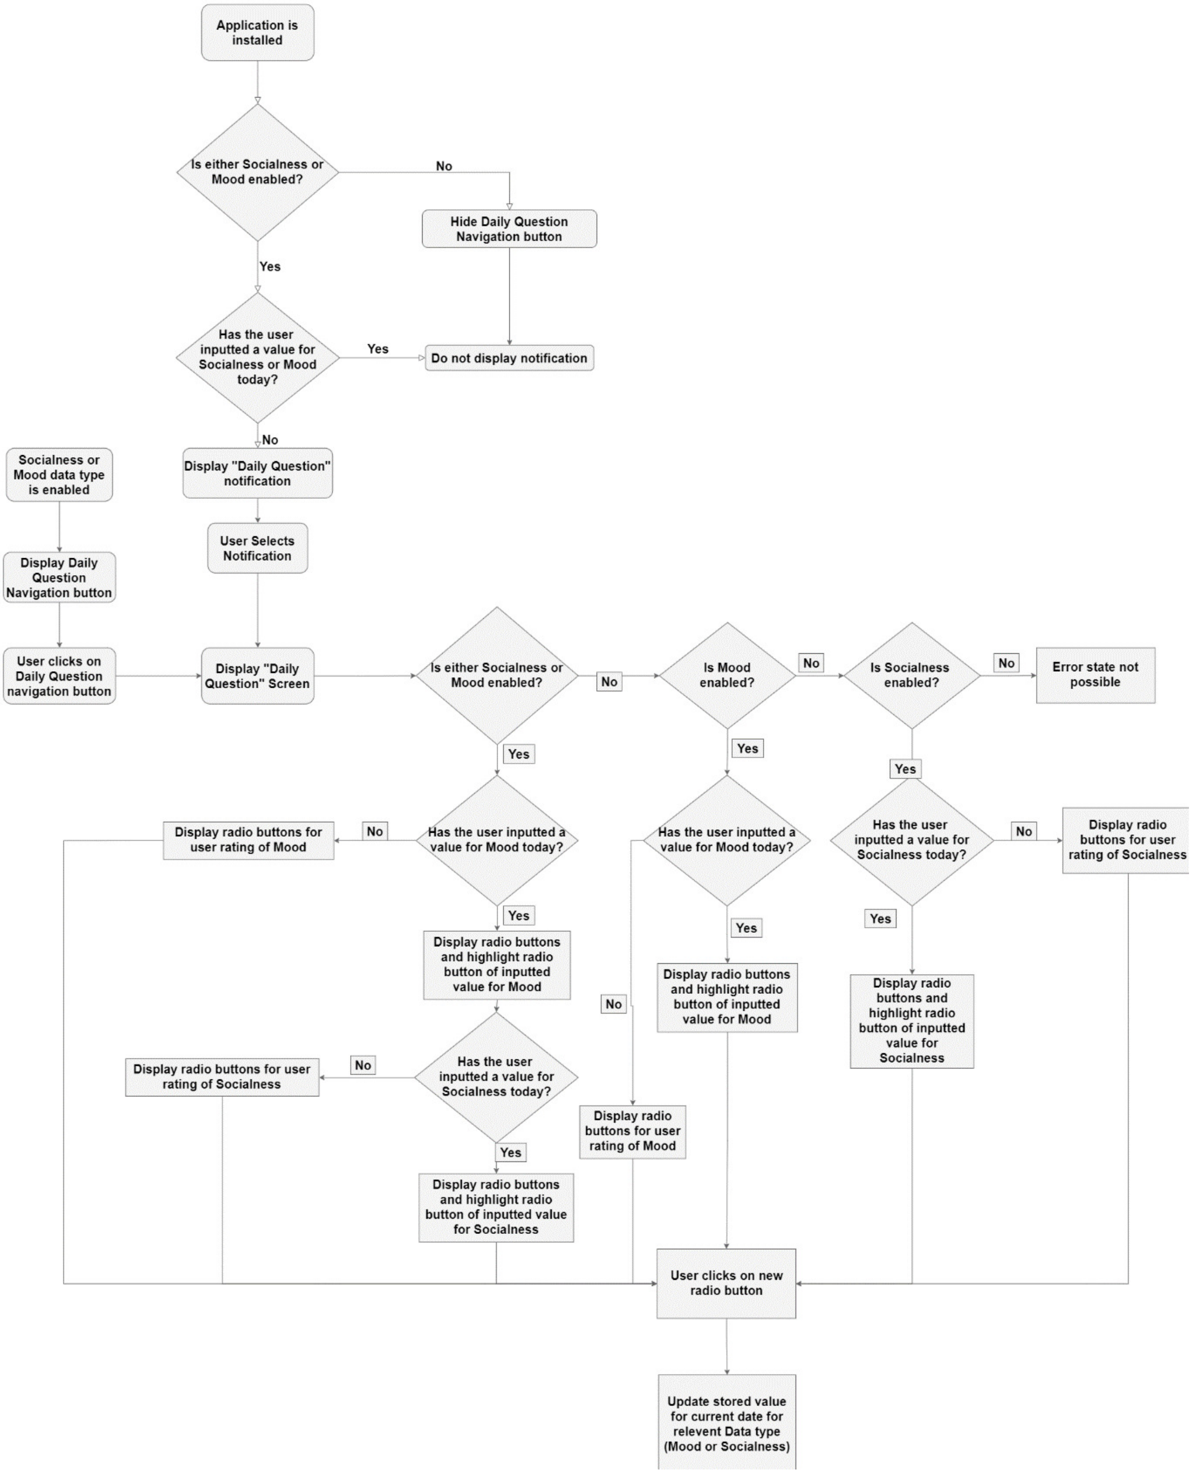

[Figure 6. Daily Question Screen flowchart. Flowchart displaying how the user will input data types of the User Inputted category]

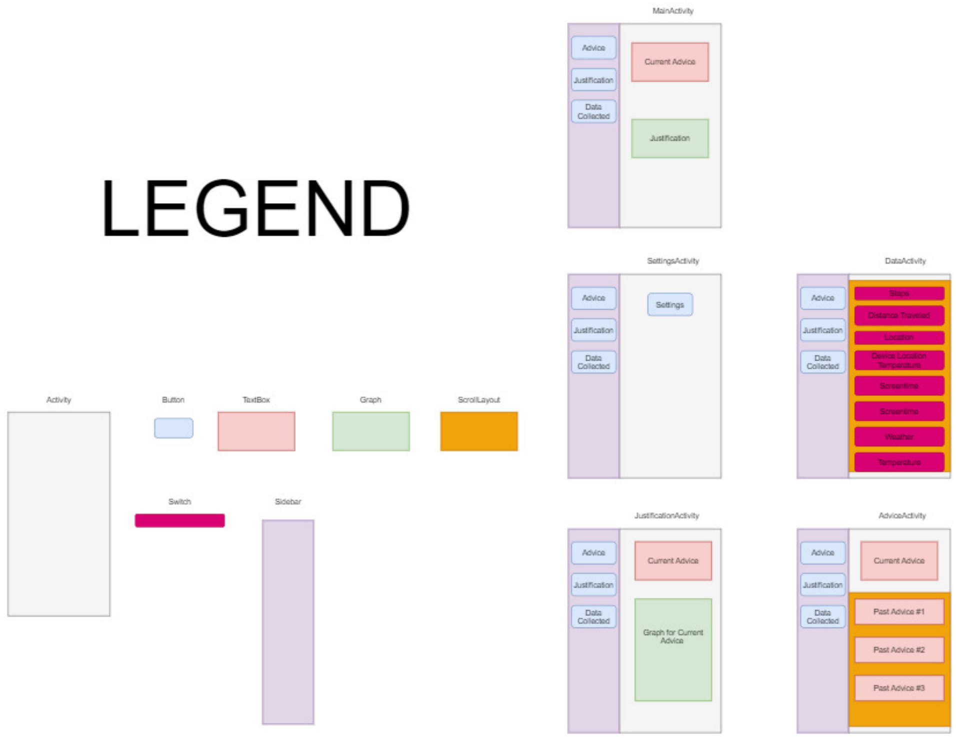

[Figure 7. Initial UI Design. Designs of the initial UI that will be used in the application]

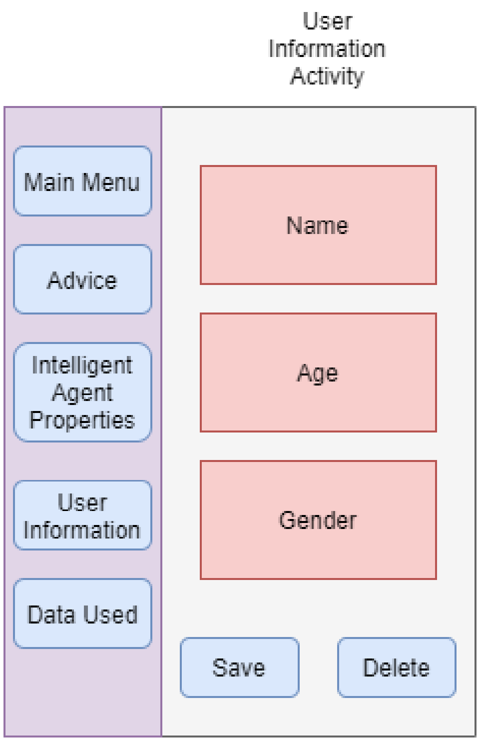

[Figure 8. User Information Activity. Design of the User Information screen]

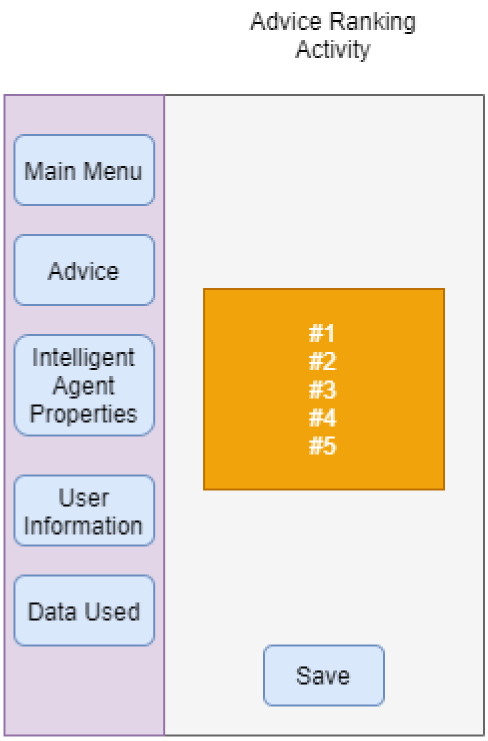

[Figure 9. Advice Ranking Activity. Design of the Advice Ranking screen]

Implementation

The source code for Guidance is available online as a GitHub repository [Appendix 5].

Guidance used a local storage system called MongoDB Realm to store all app related data. MongoDB Realm was chosen primarily due to its ease of use. Realm data models are defined as normal Java classes that extend the RealmObject class, meaning that each data type could have a model RealmObject class created for it that would then be used to create new objects which would then be stored in the Realm database. These models can then be queried using functions on the object which can then either return a generic RealmResults instance of the queried class (a list of the queried model) or a single object of the queried model. The alternative to MongoDB Realm would have been SQLite which is a relational database with 20 years of support (SQLite 2021). However, SQLite would have needed more complicated querying statements when compared to Realm. Furthermore, SQLite is slower when compared to Realm with regards to querying all records in the database (Lipika 2016). Guidance needed to query the data stored in the database often, meaning that slow queries would have reduced the effectiveness of the application, potentially causing annoyance to the participant of the study, which may have affected the results of the study. A problem that was encountered when first using MongoDB realm was that the documentation for the Android SDK version was lacking in detail. However, the documentation was updated providing more code snippets and examples of use cases. The best practices for implementing MongoDB Realm database models and queries may not have been observed when implementing Guidance. However, a strict schema was developed and followed when designing the RealmObjects used, with each RealmObject class having an accompanying supporting class for the queries needed to access and store data in the Realm Database.

Guidance could provide advice relating to 5 different data types, which are as follows; Step, Location, Screentime, Socialness and Mood. If there was no advice available for the user then an entry for the date was made with the advice type being "NoAdvice". The user had the option of establishing a ranking preference of the data types that they would prefer to receive advice about, through the use of the Advice Ranking Activity. When providing advice to the user, Guidance would use this ranking to then assess what advice it should give to the user. Once all the data types available have been assessed, Guidance would then provide advice to the user if possible, using preselected advice templates. If the user had entered their name into the application, it would address the user directly when providing advice.

The Weather, External Temperature and Sun data types were used as supplementary data when providing advice. They were used to determine a date when the advice could be followed. For the weather data type, this was done by checking the Weather data type Realm model for all dates after the current date whose weather type was "Clouds" or "Clear" and for the External Temperature data type this was done by checking the internal Weather Realm Model for all dates after the current date whose temperature was within a reasonable range (between 0^{o}C and 30^{o}C). For example, if IA determined that a user had been inside for several days in a row, but on the following day it would rain, the IA would query the Weather data category to determine when the next available clear day (no rain) was in which the user would be able to go outside.

One problem that was encountered when implementing Guidance was that one advice type that was planned to be offered related to the ambient temperature of the room where the device was located. The code for this was implemented and would record the ambient temperature of the room roughly every fifteen minutes. However, whilst the Android OS supported the use of an ambient temperature sensor, it was up to the manufacturer to actually implement the physical sensor onto the device

(Android Developers 2019). Therefore, most modern phones would not be able to make use of the ambient temperature functionality of the device. The last two Samsung phones that had an ambient temperature sensor listed on Gsmarena were the Samsung Galaxy S4 (Gsmarena. 2013a) and the Samsung Galaxy Note 3 (Gsmarena. 2013b) both of which were released in 2013 and were only supported up to Android 5.0.

Results and Evaluation

Disclaimer

The study originally started on the 19/04/2021. However, due to a bug present in advice functionality of Guidance, when the user had the Location data type enabled, it would cause Guidance to crash resulting in the user not receiving advice. This invalided the results of the study because there was no guarantee that all participates were receiving advice. Therefore, the study was restarted on the 24/04/2021 and continued until the 01/05/2021. The results from the 19/04/2021 – 24/04/2021 were discarded when the study was restarted.

Furthermore, when evaluating the usage data generated from the study it became apparent that the verdict of the advice given was not accessed. A closer inspection of the code revealed a problem with the service that was meant to give a verdict of the advice given by the IA, resulting in there being no verdict for any of the advice given by the IA. This means that the verdict of the advice given in this study cannot be used to determine the key attributes that an IA needs in order to be deemed trustworthy. However, the problems with the advice verdict functionality did not impact the users perceived working of the application and therefore the Analysis attribute can still be assessed.

Additionally, the participants' usage data regarding which data types the user enabled or disabled was not recorded. It is unknown why this was not recorded. Therefore, the data types deemed sensitive by the users cannot be ascertained.

Results

The figures used in this study are drawn from the participant usage data [Appendix 6]. This study had a total of six participants, three males and three females. The age range of the participants was 20-36.

Although the verdict of the advice given by the IAs in this study cannot be used to assess the key attributes the IA needed to have to be deemed trustworthy, the questionnaires can still be assessed. The participants in the study were given two questionnaires over the duration of the study period, one at the start of the study period [Figure 1] and one at the end [Figure 2].

The first question in questionnaire one aimed to determine how many participants would be likely to follow the advice given by and IA. The first question in questionnaire two aimed to determine if the user would continue to follow the advice given by the IA. Most participants said they were likely to follow the advice given by the IA [Figure 10, 11]. However, by the end of the study most participants stated that they were unlikely to continue to follow this advice [Figure 10]. Participants 2, 4 and 5 changed their answer regarding this at the end of the study. Participants 4 and 5 said they were unlikely to continue to follow the advice from the IA, whereas previously they had stated at the start of the study that they were likely to follow this advice [Figures 11, 12]. Participant 2 changed their answer from Very Likely to Likely [Figures 11, 12].

Figures 12 and 13, show that the participants who were unlikely to continue to follow the advice of the IA also think that they followed nearly none of the advice given, whereas the two participants who said that they were likely to continue to follow the advice given by the IA thought that they had followed most of the advice that they had been given.

Figures 14 shows that the majority users at the start and end of the study had a preference towards advice being given twice per day. Participant 2 changed their mind at the end of the study, from preferring advice to be given once per day to advice being given twice per day [Figure 15, 16].

Figure 17 shows that the majority of users had a preference for the IA's gender to be neutral. This did not change over the duration of the study [Figure 18, 19]. However, participant 1 changed their preference from a male IA to a female IA [Figure 18, 19]. At the start of the study all participants stated that their preferred the communication method with the IA would be text [Figure 20]. At the end of the study only participant 1 had changed their answer from text to text and speech [Figure 21].

The majority of participants stated that they were more likely to follow the advice given by the IA if a reason for the advice was given [Figure 22]. Only participant 6 stated that that they were not likely to follow the advice if a reason was given. Participant 6 also stated that they were unlikely to follow the advice given by the IA [Figure 11]. Figure 23 shows that participants 1,2,3 and 6 trusted the advice given by the IA, whereas participants 4 and 5 did not. All participants stated that they trusted applications that have machine learning implemented [Figure 24]. This did not change over the course of the study [Figure 25].

Participants 1,4 and 5 stated that they determined that the IA was not using machine learning [Figure 26], whereas participants 2, 3 and 6 stated they did not determine that the IA was not using machine learning [Figure 26]. Participants 1, 2 and 3 were assigned IAs that had the Traditional Programming for the Analysis attribute [Table 1], whereas participants 4,5 and 6 were told that they had Machine Learning for the Analysis attribute [Table 1], when in fact had Traditional Programming. All IAs given to participants in the study used Traditional Programming when providing advice to the participants. All participants stated that they were comfortable with the IA using traditional programming techniques instead of machine learning to provide advice to them [Figure 27]. The majority of the participants were comfortable with the device monitoring them continually [Figure 28]. The majority of participants also preferred to follow the advice given by the IA instead of advice given by a human [Figure 29].

The breakdown of the total amount of advice given by the IA was 64% Step, 18% Screentime, 6% Location and 12% No Advice, with no advice being given for the Socialness and Mood data types [Figure 30]. Some participants received advice regarding various data types [Table 2]. However, most participants received advice regarding one data type, the step data type [Table 2]. Several participants received varying amounts of advice [Table 2], with participant 3 receiving advice the most amount of advice and participant 2 receiving none [Table 2].

Only one participant used the Ranking Preference screen (Participant 4) to change their preferences. However, they did not adjust their ranking preference and instead increased their preferred step count [Table 3].

Evaluation

Figures 10 and 11 show that the majority of participants felt that they were more likely to follow the advice given by the IA at the start of the study, whereas at the end of the study they felt that they were unlikely to continue to follow the advice given by the IA [Figure 12]. This showed how the duration of the study impacted their opinion of the IA. Participants at the start of the study may have been more hopeful that they would follow the advice given by the IA, whereas after being given advice by the IA every day, they thought that they would be unlikely to continue to follow the advice given.

Interestingly, figure 13 shows that the participants who felt that they were unlikely to continue to follow the advice given by the IA (participants 1,4,5 and 6) felt that they had followed nearly none of the advice given, whereas the participants who said they were likely to continue to follow the advice (participants 2 and 3) felt that they had followed the majority of the advice. The majority of the advice given to participants 1,4,5 and 6 was Step advice [Table 2]. This is likely why they felt that they did not follow most of the advice given and why they felt that they were unlikely to continue to follow the advice given by the IA. Participant 3 received the least amount of Step advice, instead mostly receiving Screentime related advice and on three days no advice [Table 2]. This may be why they felt that they would be more likely to continue to follow the advice given by the IA because they were not consistently given advice from the same data category. Participant 2 did not receive any advice from the IA [Table 2]. Therefore, it is interesting that they thought that they would continue to follow advice despite receiving none. It is also interesting that they thought they had followed most of the advice despite receiving none. This could have been due to a problem with the application not recording the advice given to this participant.

The participants' opinion regarding the interaction amount was consistent throughout the study, with a majority preferring to receive advice twice per day [Figure 15, 16]. At the end of the study participant 2 stated that they would have preferred to receive advice twice per day. As previously mentioned, the usage data generated does not show that participant 2 received advice from the IA. Therefore, it may be that their opinion changed from advice once per day to advice twice per day because they did not receive any advice from the IA. The Interaction attribute for all of the IAs provided to the participants was set to once per day [Table 1]. This shows that the majority of users who participated in the study would have preferred to receive more advice from the IA, despite being unlikely to continue to follow the advice given by the IA. This seemed counterintuitive and did not correlate with the answers shown in figures 12 and 13 because most participants said that they were unlikely to continue to follow the advice given and thought that they had followed nearly none of the advice given during the study period. If a focus group discussion had been carried out at the end of the study period, this could have been discussed with the participants, to better understand their answers.

The Gender attribute for all of the IAs given to the participants was set to male [Table 1]. Figures 18 and 19 show that the majority of participants would have preferred neutrally gendered IA. This is contradictory to the hypothesis that most participants would have preferred a female IA. Only participant 1 changed their opinion regarding this, from preferring a male IA to preferring a female IA. This is likely due to the IA assigned to them being male and therefore at the end of the study they disliked the male IA and therefore would prefer to have a female IA.

The participants who were told that they were given the traditional programming for the Analysis attribute were participant 1, 2 and 3, whereas the participants who were told that they had received Machine Learning for the analysis attribute were participants 4, 5 and 6. As shown in figure 26, participants 2 and 3 stated that they did not determine that the IA was not using machine learning despite being given an IA that informed them that it was using traditional programming, whereas participants 4 and 5 determined that the IA was using traditional programming instead of machine learning. This was likely because the participants who were told that they had received machine learning were more critical of the advice that they were receiving when compared to the participants who had received traditional programming. Another likely option is that participants 1, 5 and 6 were given advice regarding one data type consistently and therefore may have discerned that the IA was using traditional programming. This was supported by participants 2 and 3 stating that they had not realised. However, due to participant 6 being unaware that the IA was using machine learning and participant 2 receiving no advice, more testing should be done on this to draw a conclusion.

As shown in figure 30 and table 2, the majority of the advice given by the IA was regarding the users' step count. The default ranking preference was set to the following ranking; Step, Location, Screentime, Socialness and Mood [Figure 31]. It is likely that Step advice was provided the majority of the time due to users not changing their ranking preference or their ideal step count. This may be because users were not aware that they were able to change their ranking preferences of the data types. Additionally, Screentime advice was likely the second most advice type given due to Location's similarity to the Step data type. If the participant had left their house, then they likely would reach their step goal. If they had not left their house then they would have been unlikely to achieve their Step goal and thus be given Step advice. By changing the default ranking preference to have Location higher than Step might allow for a more balanced advice spread. Constantly receiving advice regarding one data category would likely result in the user ignoring or dismissing the advice given by the IA. This is potentially shown in figures 12 and 13, because participants felt that were unlikely to continue to follow the advice given by the IA and felt that they had followed nearly none of the advice given. The participants who felt this way were participants 1,4,5 and 6 [Figure 12, 13], who mostly received advice from one data category [Table 2].

No advice was given for either of the Mood and Socialness data categories. This was likely due to the Step category being ranked higher than both the Mood and Socialness categories. Therefore, when the user might have been feeling unhappy or unsocial, the advice they could have received about this might not have been given and instead they might have received advice regarding a data category higher on the ranking preference list. A solution to this might be to change the default ranking preference to favour the Mood and Socialness data categories. However, it may be due to users not wanting to acknowledge that they might be unhappy or that they might not have socialised and therefore choosing a higher option on the Daily Questions screen for each of the categories.

A potential solution to the majority of the advice being step related advice could be to have a tutorial on the Ranking preference screen, explaining how to use the screen and what it does. This would have been useful because users might not have known that they could change the ranking preference of the advice. Throughout the study period, only participant 4 changed their ideal step count, but not the ranking preferences [Table 3].

A potential problem that was encountered during the project was with Guidance's implementation of the step counter for the Step data type. The step counter used in Guidance was executed as a service, meaning that if the user had battery optimisation enabled, then the Step Service may have been halted. This meant that the steps that the user may have done for the rest of the day would have overwritten the steps currently recorded. In hindsight a solution for this problem would have been to check if the current steps recorded for the day were less than the current count recorded by the sensor, if so then both were combined and then the combined value was stored.

Due to the small participant pool the usage data generated from this study would have been hard to draw a concrete conclusion from if the verdict of the advice had been assessed. Assessing both the Analysis and Advice attribute at the same time was a mistake. In hindsight only one attribute should have been tested.

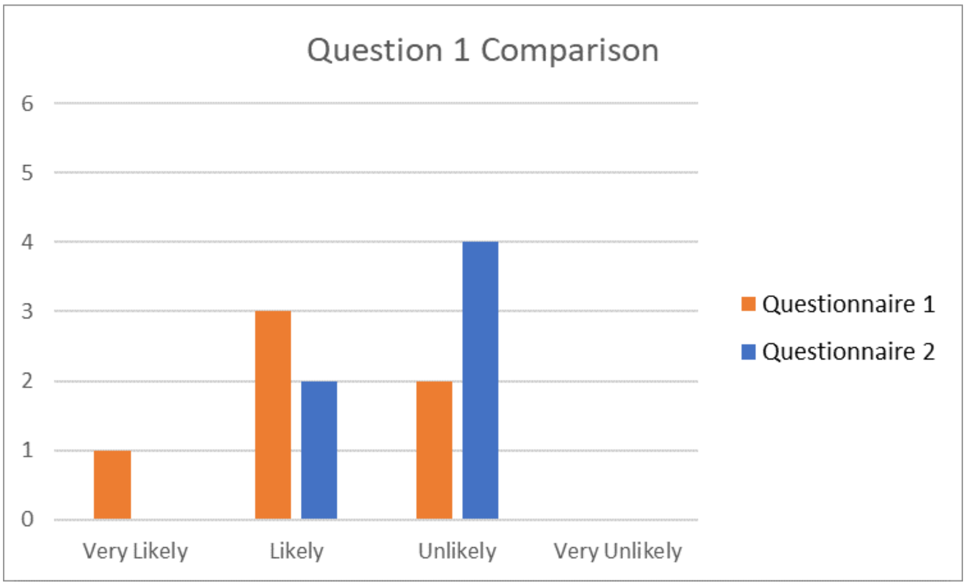

[Figure 10. Question 1 Comparison. Comparison of the questions "How likely are you to follow the advice given by this Intelligent Agent?" at the start of the study and "How likely would you to continue to follow the advice given by this Intelligent Agent?" from the end of the study]

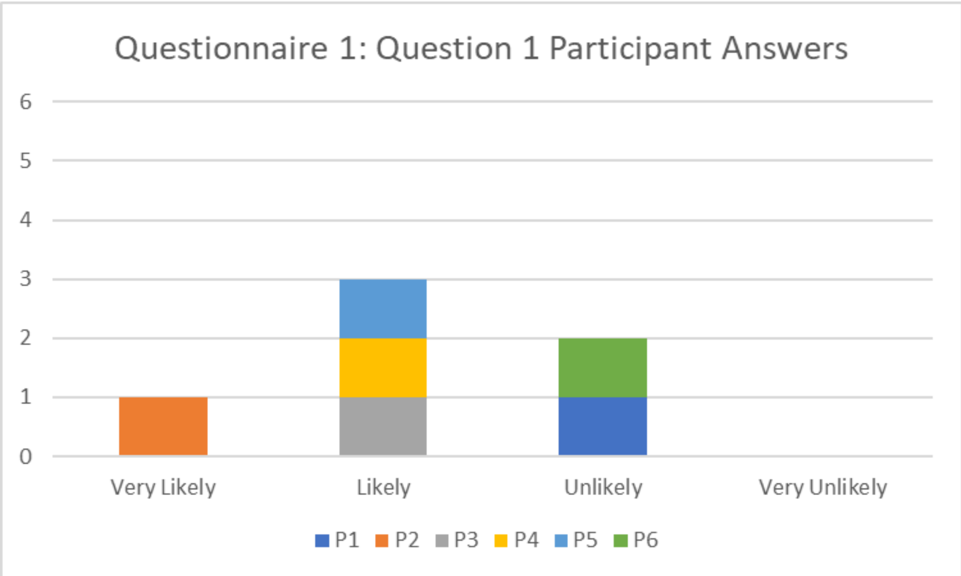

[Figure 11. Questionnaire 1: Question 1 Participant Answers. Participant answers for the question: How likely are you to follow the advice given by this Intelligent Agent? From questionnaire 1]

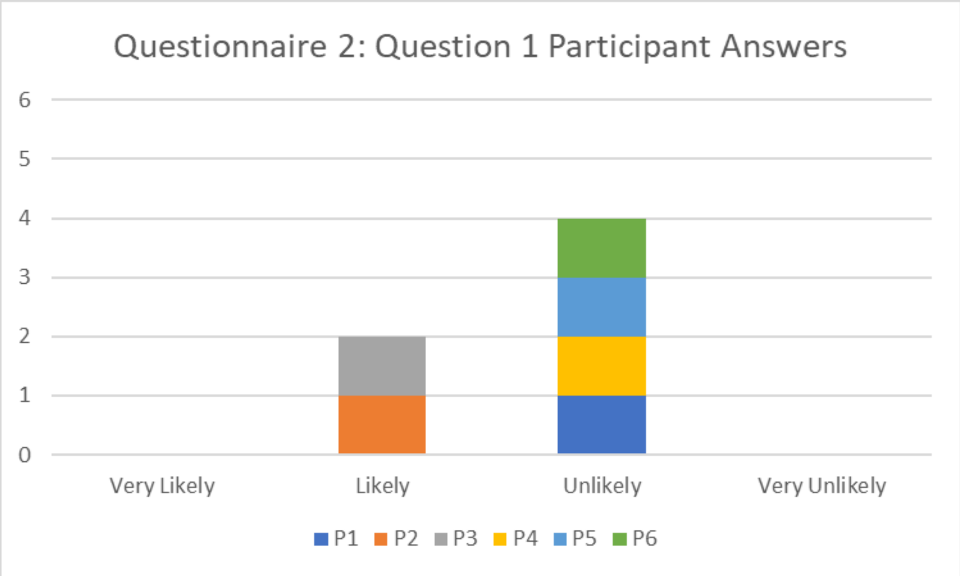

[Figure 12. Questionnaire 2: Question 1 Participant Answers. Participant answers for the question: How likely would you to continue to follow the advice given by this Intelligent Agent? From questionnaire 2]

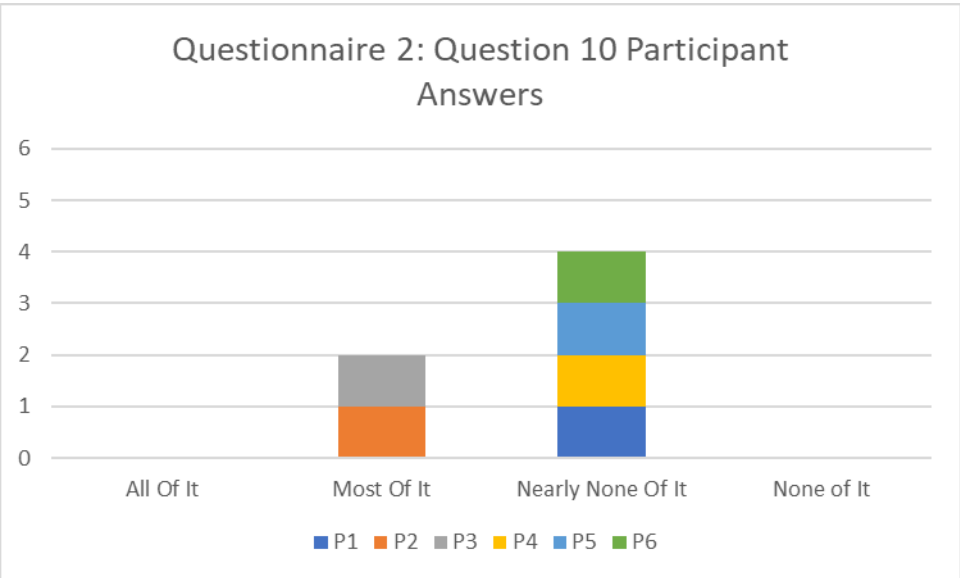

[Figure 13. Questionnaire 2: Question 10 Participant Answers. Participant answers for the question: How much of the advice do you think you followed? From questionnaire 2]

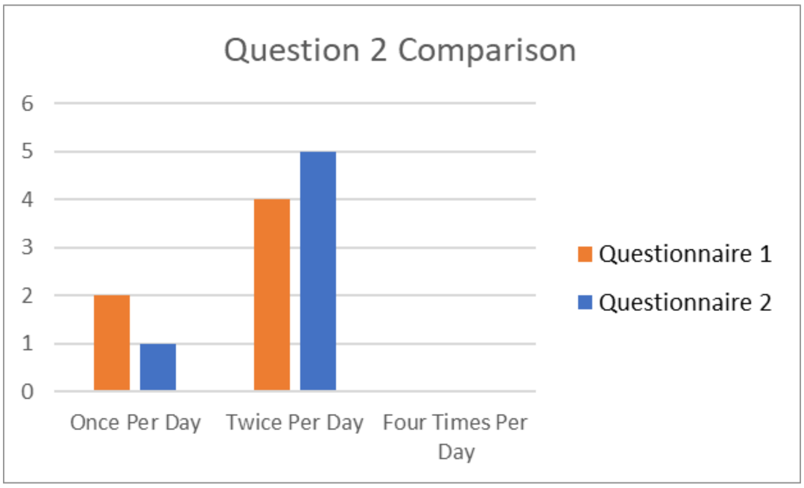

[Figure 14. Question 2 Comparison. Comparison of the questions "How often would you like to receive advice?" at the start of the study and "How often would you have liked to receive advice from the Intelligent Agent during the study?" from the end of the study]

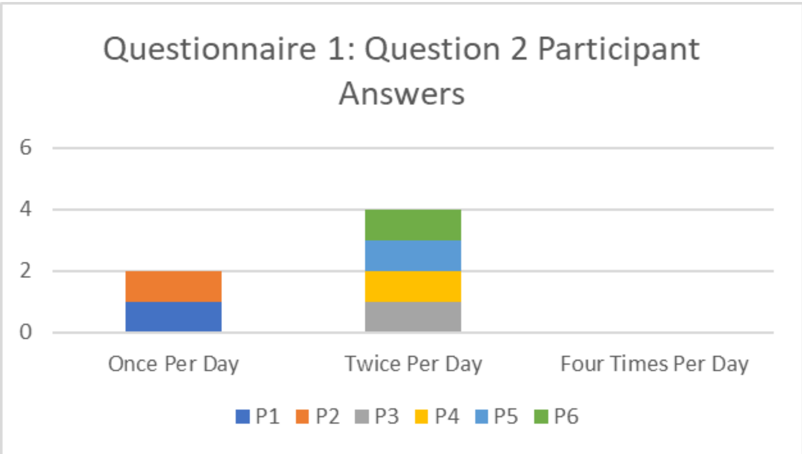

[Figure 15. Questionnaire 1: Question 2 Participant Answers. Participant answers for the question: How often would you like to receive advice? From questionnaire 1]

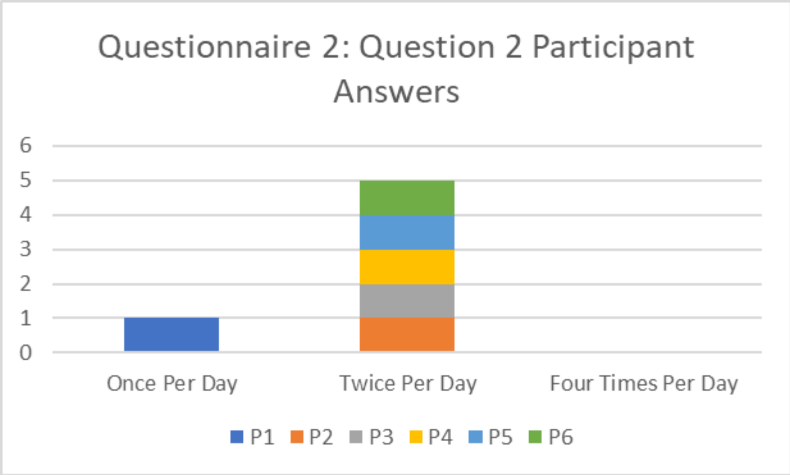

[Figure 16. Questionnaire 2: Question 2 Participant Answers. Participant answers for the question: How often would you have liked to receive advice from the Intelligent Agent during the study? From questionnaire 1]

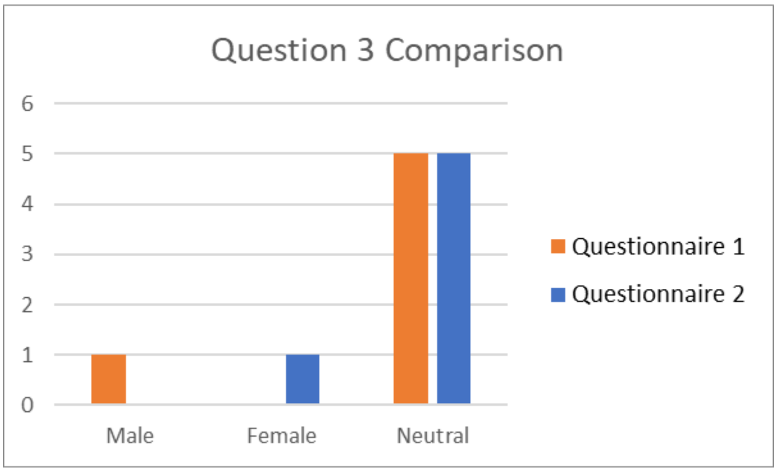

[Figure 17. Question 3 Comparison. Comparison of the questions “What gender would you prefer the Intelligent Agent to be?” at the start of the study and “What gender would you have preferred the Intelligent Agent to be over the course of the study period?” from the end of the study]

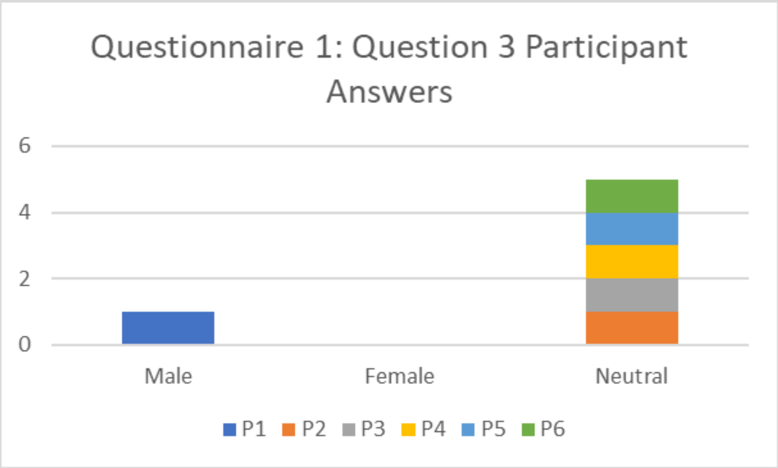

[Figure 18. Questionnaire 1: Question 3 Participant Answers. Participant answers for the question: What gender would you prefer the Intelligent Agent to be? From questionnaire 1]

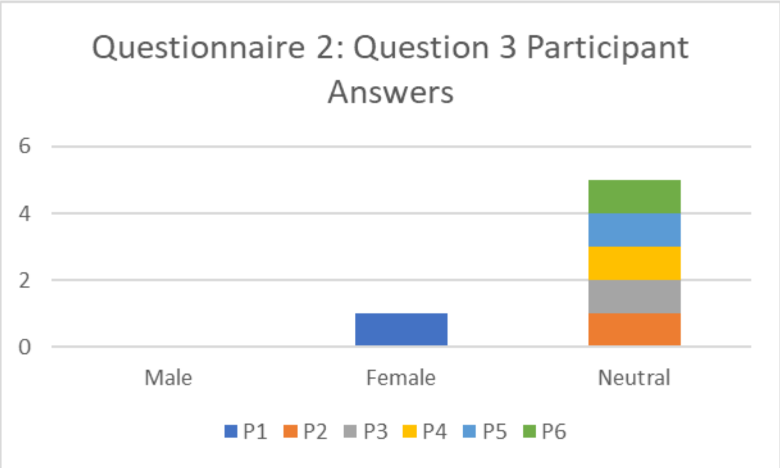

[Figure 19. Questionnaire 2: Question 3 Participant Answers. Participant answers for the question: What gender would you have preferred the Intelligent Agent to be over the course of the study period? From questionnaire 1]

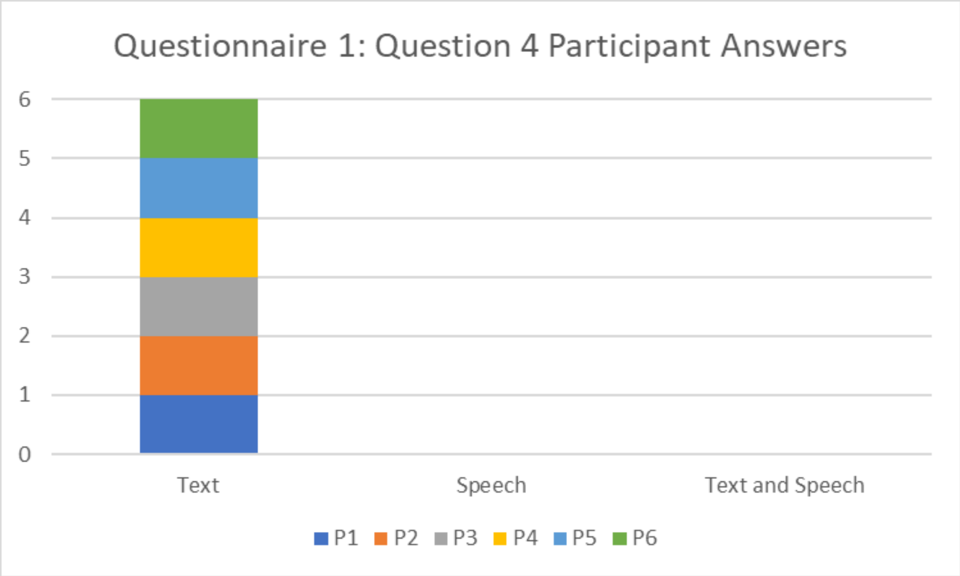

[Figure 20. Questionnaire 1: Question 4 Participant Answers. Participant answers for the question: What method of communication would you prefer the Intelligent Agent to use? From questionnaire 1]

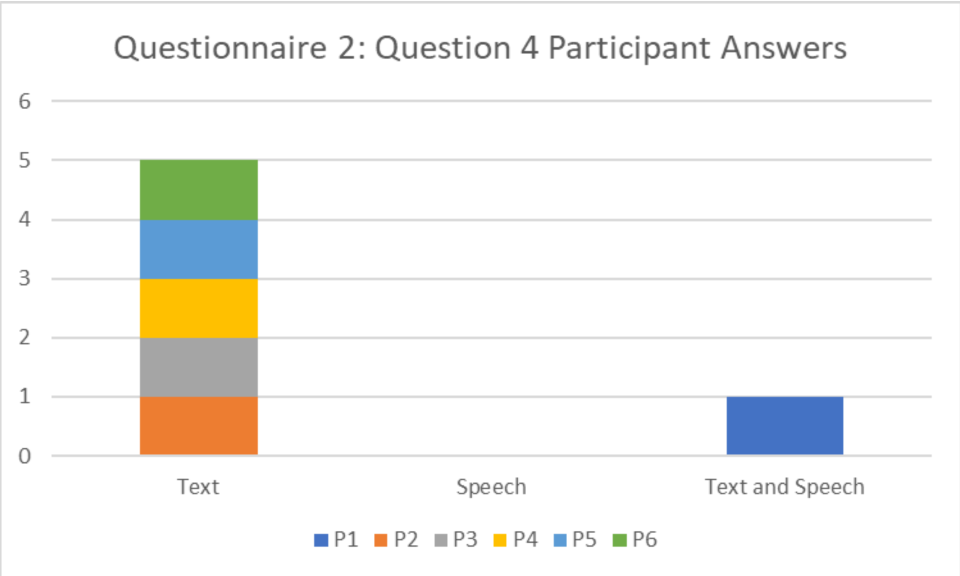

[Figure 21. Questionnaire 2: Question 4 Participant Answers. Participant answers for the question: What method of communication would you prefer the Intelligent Agent to have used during the study period? From questionnaire 1]

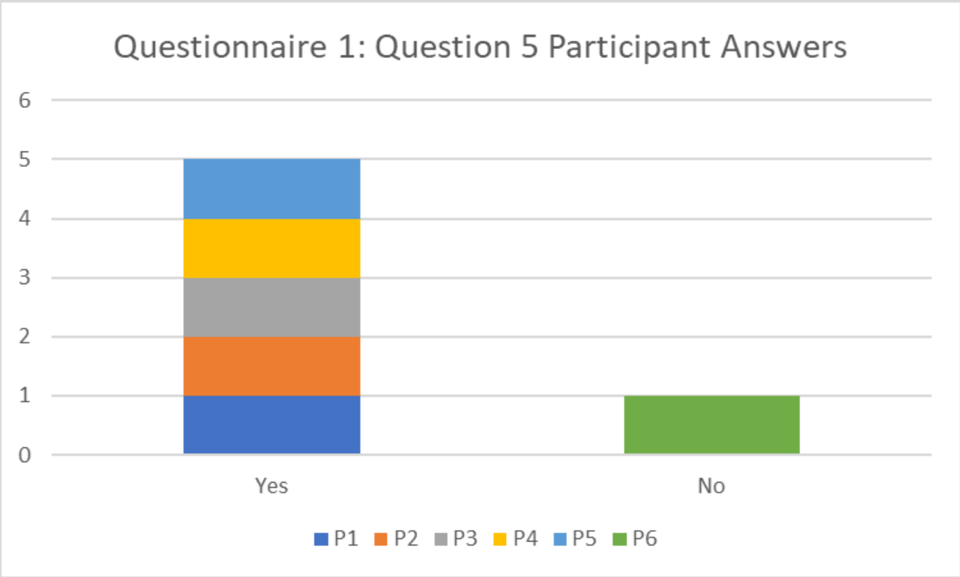

[Figure 22. Questionnaire 1: Question 5 Participant Answers. Participant answers for the question: Are you likely to follow advice given by this Intelligent Agent if a reason for the advice is given? From questionnaire 1]

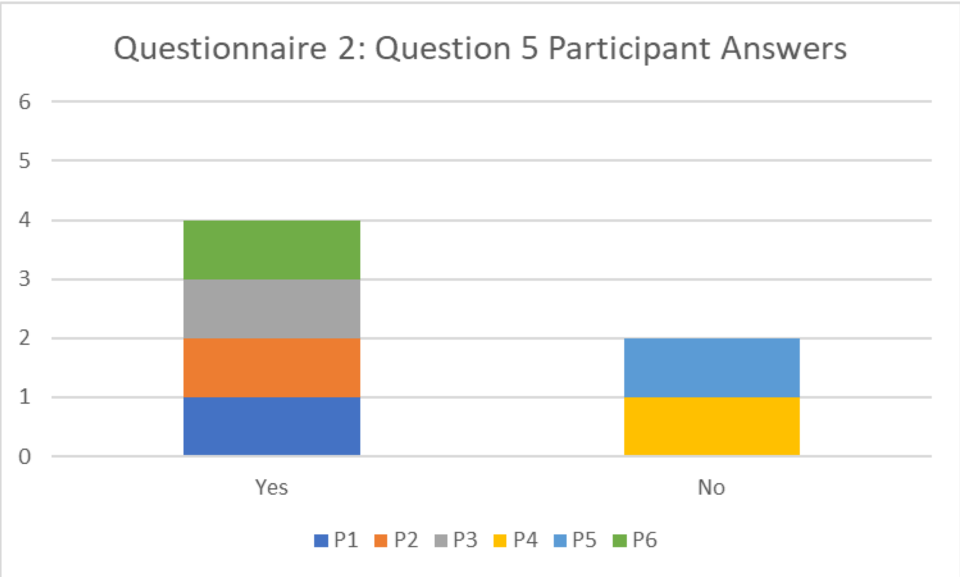

[Figure 23. Questionnaire 2: Question 5 Participant Answers. Participant answers for the question: Did you trust the advice given by the Intelligent Agent? From questionnaire 2]

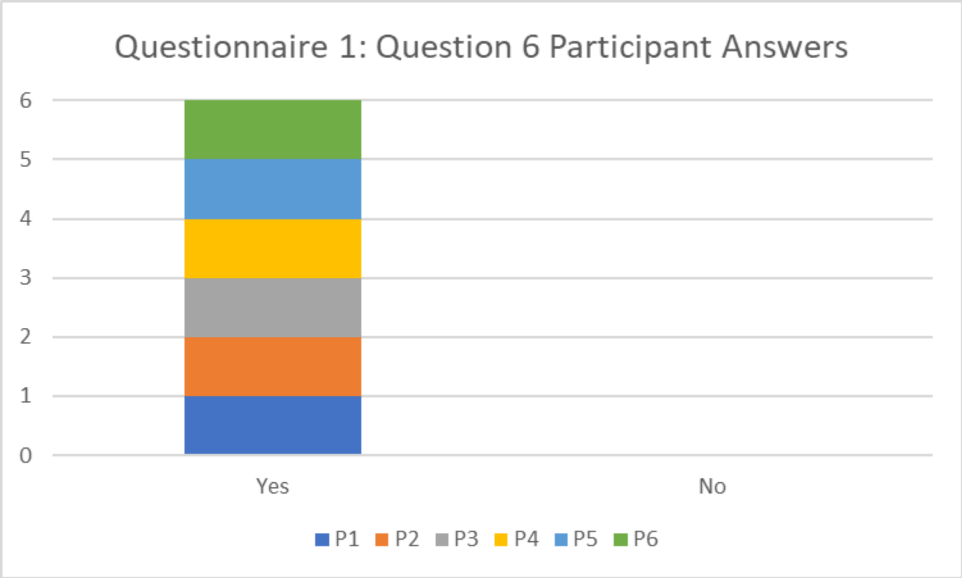

[Figure 24. Questionnaire 1: Question 6 Participant Answers. Participant answers for the question: Do you trust applications that have machine learning implemented? From questionnaire 2]

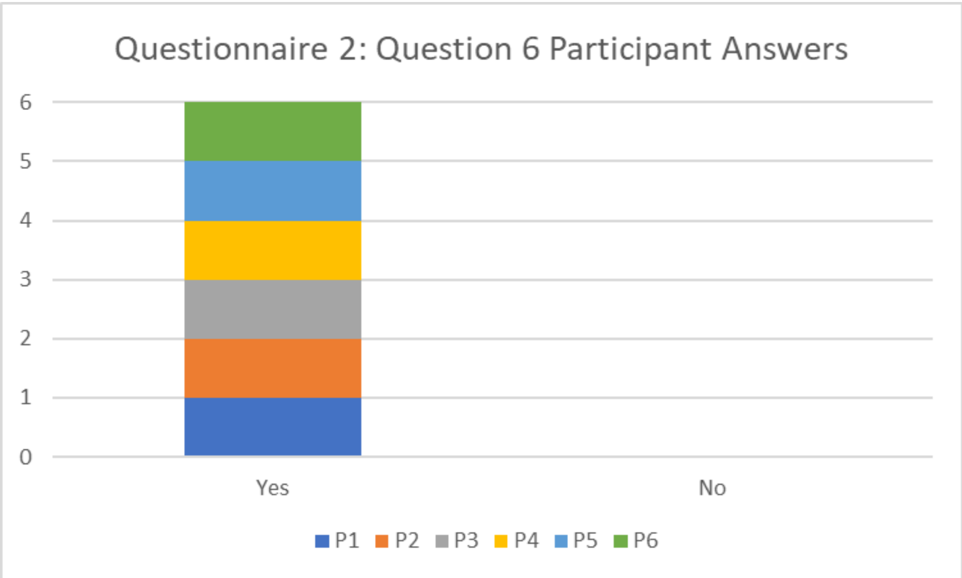

[Figure 25. Questionnaire 2: Question 6 Participant Answers. Participant answers for the question: Do you trust applications that have machine learning implemented? From questionnaire 2]

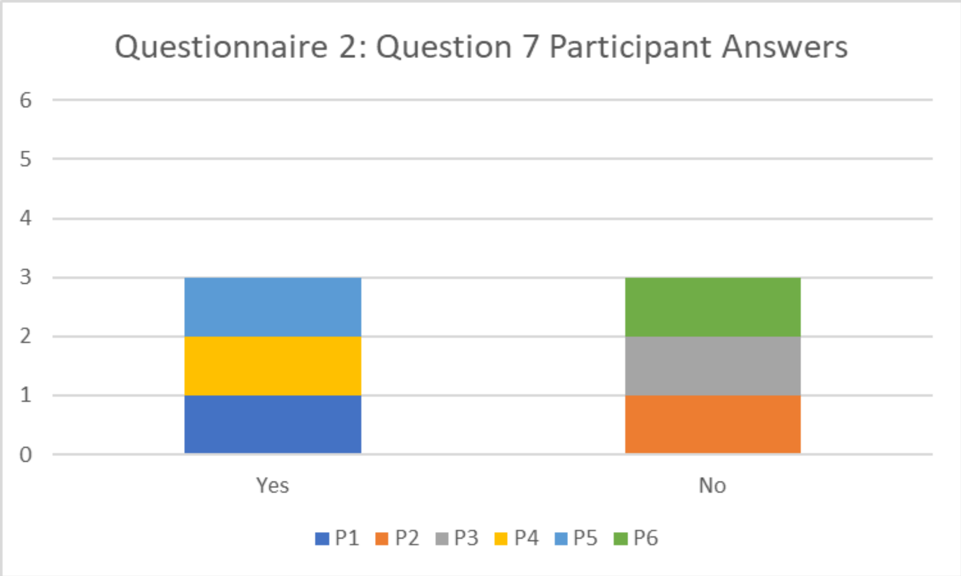

[Figure 26. Questionnaire 2: Question 7 Participant Answers. Participant answers for the question: Did you determine that the Intelligent Agent was not using machine learning? From questionnaire 2]

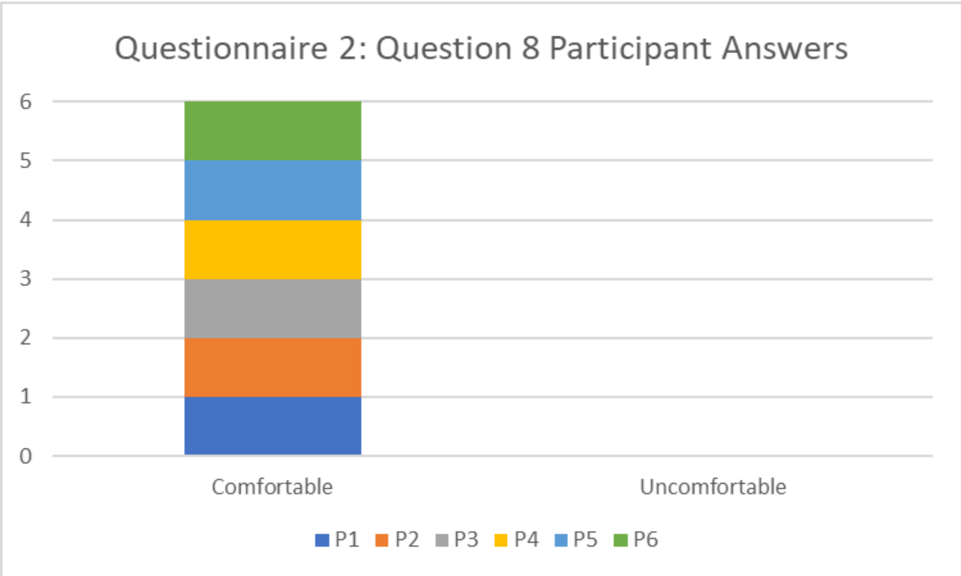

[Figure 27. Questionnaire 2: Question 8 Participant Answers. Participant answers for the question: How do you feel about the Intelligent Agent using traditional programming techniques instead of machine learning to provide advice? From questionnaire 2]

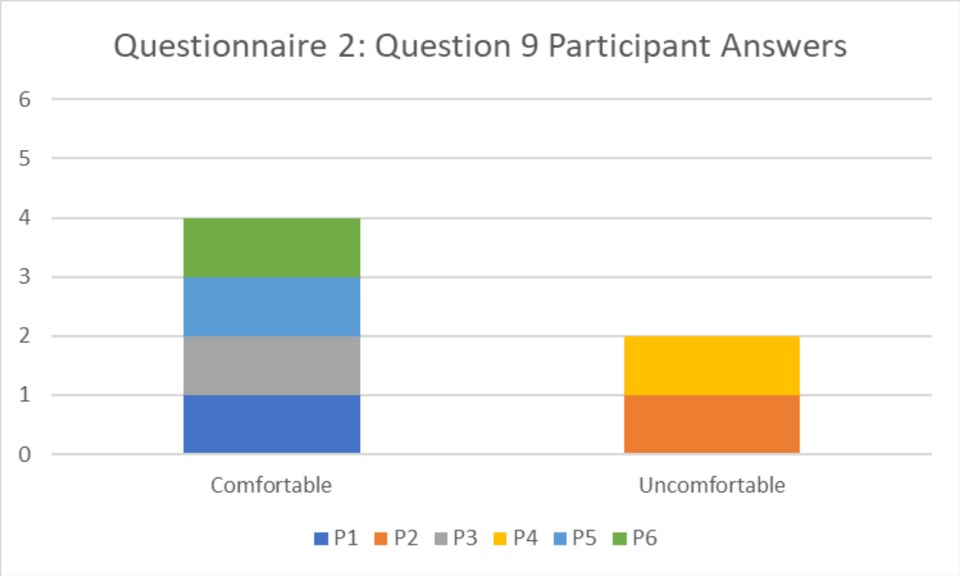

[Figure 28. Questionnaire 2: Question 9 Participant Answers. Participant answers for the question: How do you feel about the Intelligent Agent using traditional programming techniques instead of machine learning to provide advice? From questionnaire 2]

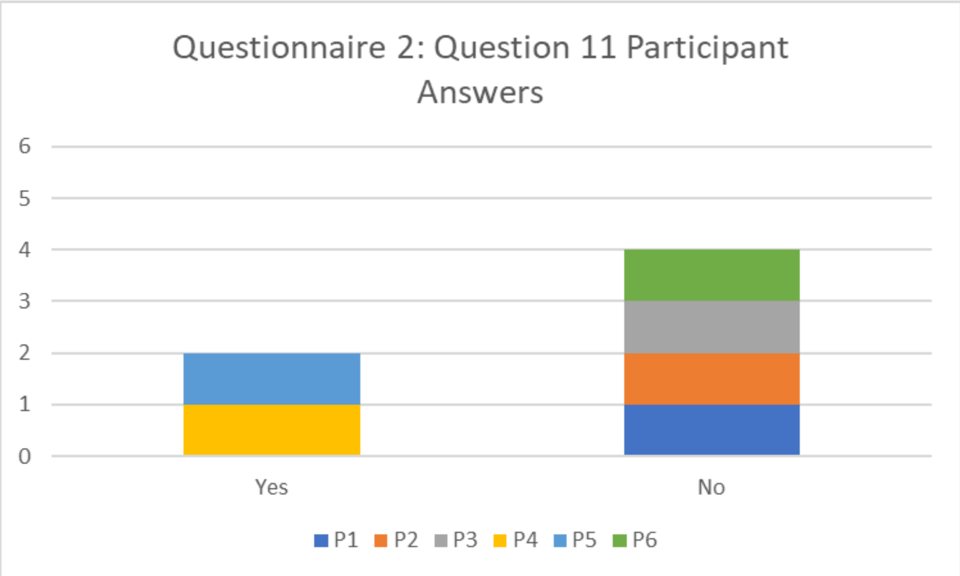

[Figure 29. Questionnaire 2: Question 11 Participant Answers. Participant answers for the question: Would you prefer to follow advice given by a human over advice given by the Intelligent Agent in this application? From questionnaire 2]

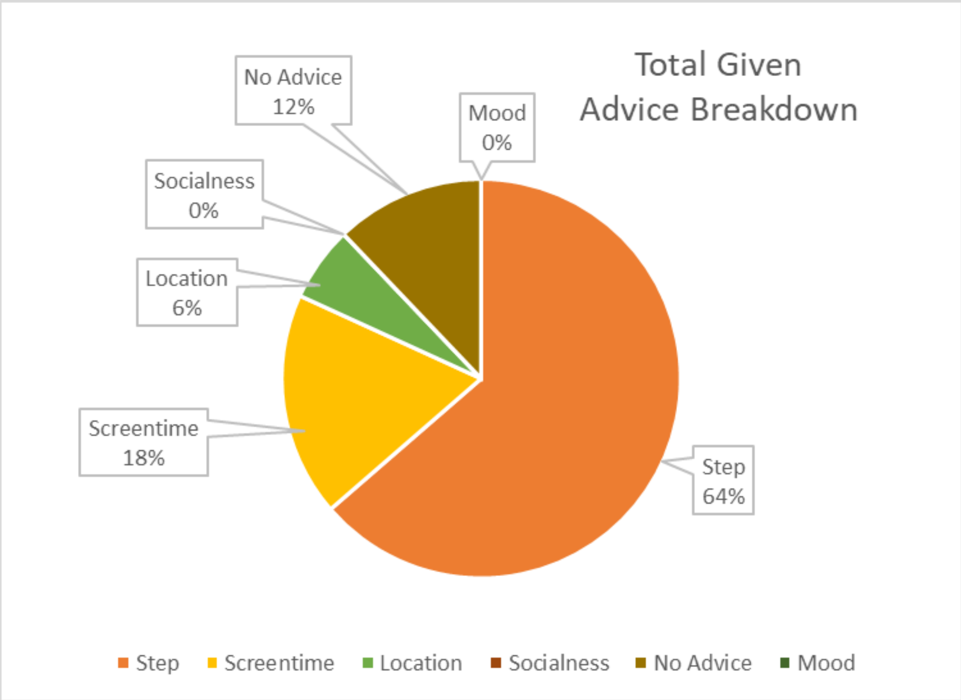

[Figure 30. Total Advice Given Breakdown. Breakdown of all the advice given to all participants in the study, please note one participant received no advice from the IA]

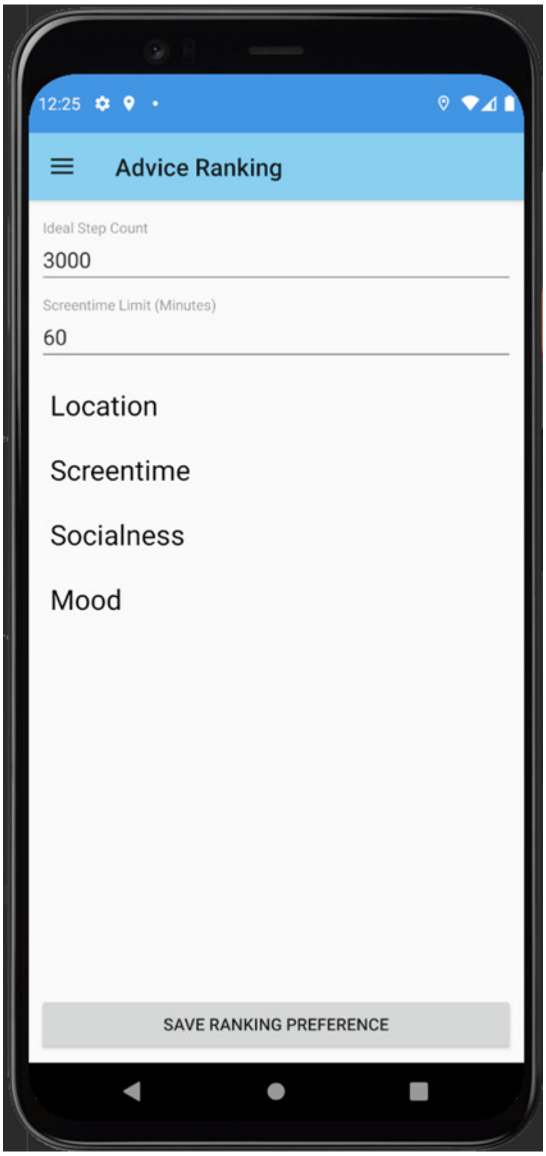

[Figure 31. Default Advice Ranking Screen. The default ranking of the data types within Guidance]

Participant	Analysis	Advice	Gender	Interaction	Output
1	Traditional Programming	No Justification	Male	Once Per Day	Text
2	Traditional Programming	With Justification	Male	Once Per Day	Text
3	Traditional Programming	With Justification	Male	Once Per Day	Text
4	Machine Learning	No Justification	Male	Once Per Day	Text
5	Machine Learning	With Justification	Male	Once Per Day	Text
6	Machine Learning	With Justification	Male	Once Per Day	Text

[Table 1. Participants Assigned Attributes. Table displaying the attributes of the IA that was provided to each user]

Participant Advice Breakdown	Participant 1	Participant 2	Participant 3	Participant 4	Participant 5	Participant 6
Step	5	0	1	5	6	4
Location	0	0	0	2	0	0
Screentime	0	0	4	0	0	2
Socialness	0	0	0	0	0	0
Mood	0	0	0	0	0	0
No Advice	0	0	3	0	0	1
Total	5	0	8	7	6	7

[Table 2. Participant Advice Breakdown. Table displaying the breakdown of the advice that each participant received]

Participant 4 Ranking Preference Change	Sun Apr 25 19:37:37 GMT+01:00 2021
Steps	0
Location	1
Screentime	2
Socialness	3
Mood	4
Ideal Step Count	10000
Screentime Limit	60

[Table 3. Participant 4 Ranking Preference Change. The ranking preference change participant 4 made on the 25/04/2021]

Future Work

Access to a larger participant base would allow the study to be repeated to corroborate the results. This would provide a more representative data set, potentially allowing for a conclusion to be drawn from the users perceived Analysis attribute. Additional attributes could also be tested at this time, these being; the gender of the IA (male vs female vs neutral), the output method used (speech vs text) and the interaction frequency. The option for the IA's gender to be changed and the adjustment of the advice frequency are currently implemented within Guidance. The advice output attribute currently has a framework built into Guidance, and would therefore be relatively easy to implement. More participants would mean that the results are more representative of the general population and would also help to eliminate anomalous results provided by atypical participants. Atypical participants being participants that do not follow typical behaviour such as a marathon runner who would be regularly outside for extended periods of time or someone who is immunocompromised and therefore cannot go outside. Increasing the duration of the study period could also be done at this time, potentially allowing for questionnaires to be provided to the participants throughout the study period, potentially showing if the users opinion of the IA changes over the course of the study period.

Another avenue of research that could be done would be determining if there was an age-based preference towards following advice given by an IA. This study could potentially be done by using the existing application or overhauling it to make it more accessible to an older generation. The advice types currently used in the study might not be applicable to an older generation. Therefore, a wider variety of advice would likely be needed in order to produce meaningful advice.

A long-term study of six to twelve months would provide insight into how a user interacts with the IA on a longer time scale. A long-term study would be more suited to having the IA be more integrated into the user's everyday life. This could be done by having multiple devices for the IA to receive data and interact with the participants, such as an application, a dedicated static device in the user's residence and a wearable device such as a FitBit or smart watch. A study of this length and magnitude would require a team of researchers to design and implement each section. Additionally, this study would benefit from external funding. Funding would allow for professional consultants to be hired. Consultants would allow for the IA's platforms to be assessed externally by a third party. This would need to be done at multiple stages in the project such as at the design, implementation and testing stages to ensure that nothing is overlooked. Additionally, external funding would allow for the devices used in the study to be provided to participants. This would increase the number of participants available to participate in the study because they would not need to provide their own devices. Furthermore, by providing the devices to the user it would ensure each separate device (application, wearable and static device) can communicate with each other without fault. Due to the long-term nature of this study, the participant's opinion of the IA could be assessed to determine if it changes over the course of the study. This can be done through the use of questionnaires being given at regular intervals.

Conclusions

The aim of this project was to identify what physical and interactive attributes were crucial for an IA to possess in order to be deemed trustworthy. This was partially accomplished in this study because the users' answers from the questionnaires used in the study have provided valuable insight regarding what interactive attributes the users would prefer the IA to have. This was shown when examining the user's preference for the IA's gender. The majority of participants would have preferred the IA's gender to be neutral. This disagrees with the hypothesis that users would prefer the gender of the IA to be female.

Participants in this study also had a preference for the Interaction attribute of the IA. The majority of the participants would have preferred to receive advice twice per day. However, the reasoning behind this is unclear. This disagrees with the hypothesis that users would want to receive advice once per day.

From the results of study, it was concluded that participants were more likely to want to follow advice if they received a varied amount of advice from the IA. If participants were given the same advice constantly, they would be unlikely to want to follow the advice given by the IA and therefore it should be ensured that the advice a user would receive from an IA is varied. This should be retested to confirm this conclusion because errors present in the application required the study to be restarted, which may have affected the user's opinion. The participants' answers regarding the Analysis attribute showed some meaning, however there is not enough data to draw a firm conclusion.

The main limitations of this project were that it was conducted with a small participant pool, over a short period of time, meaning that some attributes were unable to be tested adequately. Due to the problems present in the verdict functionality of the application and the small participant pool when restarting the project only the Analysis attribute should have been tested. If the verdict functionality of the application had been working, the participant pool would have made drawing a conclusion from the data set impossible.

Another limitation of the project was that it was conducted and implemented by one researcher, over a short period of time (3 months). This meant that there was no time for a rigorous testing period, using real world scenarios, which resulted in flaws being present in the application used in the study. This then led to the study needing to be restarted.

Additional work that could be done in the future would be to redo the project with additional participants to obtain more accurate results, which would be more representative of the general population. When redoing the study, the application should be improved upon to eliminate the bugs that were present in this study. If participant pool for the new study was large enough additional attributes could be tested, such as the Interaction attribute. The new study's duration could also be extended to assess how the user's opinion changes throughout the study, with additional questionnaires at regular intervals.

Reflection on Learning

One thing that I have learnt from this project is how to structure a MongoDB Realm database, how each model should be structured and how the relationships between models' functions. I have found that having utility classes with the sole purpose of providing an interface for another class or function to operate with a specific Realm model has helped to reduce the amount of code needed to operate the application. I think that I will be able to use and apply this knowledge throughout my programming career, as this knowledge can be applied to other situations where a class or function would need to communicate with a database or dataset. I will definitely make use of this knowledge in my personal projects, where I plan to use MongoDB Realm to store data.

Another lesson that I have learnt from this project is that code readability is crucial. I found that both class and function names needed to be kept consistent throughout the project to allow me to better understand each section when rereading the code, after working on another section of the application. I believe that this can be done through the use of a schema, which greatly improves the readability of code, potentially allowing it to be understood without the use of code comments. I did not start this project with a schema but instead implemented it midway through the project. I think that it would have been better to have created a schema at the start of the project but I am aware that this is not always possible. Therefore, I will ensure that for my next project I will assess if a schema is possible to be implemented at the start of the project, and if it is not possible, I will ensure that I reflect on the work done midway through the project and attempt to implement a schema, refactoring the code where necessary.

Code comments, especially within the utility classes, have allowed for rapid and easy development of new functionality. Therefore, in future projects, after implementing a section of the project I will ensure that each function is explained in the form of a code comment that gives a sufficient explanation to someone who is not versed in the project. In addition to this I have found that explaining the input variables and the return value (if there was a return value) of each function to be beneficial. This has especially been shown in the utility classes present in Guidance. The functions within the utility classes are similar to each other but often have slight differences, such as the scheduleJob function which has multiple different inputs depending on the needs of the job being scheduled. These might not be obvious to an unknowing reader. By explaining each input, the correct function can be used. Furthermore, I have found that having these utility functions, which perform the same task just with different inputs, benefit from having the same name. This has been especially useful when implementing the different job scheduling functions present in Guidance. Thus, I will continue to do this throughout my programming career.

Another thing that I have learnt is that not all Android users are well versed in the operations of their devices. After communicating with participants of the study, I found that some did not know how to access the application permissions for a specific application and that one user found that their phone denied them easy access to the downloaded APK file, requiring them to manually find the APK file in the File Explorer and then activate it. I believe that this is primarily due to an older version of Android and therefore I think that if I was to redo the study, I would likely limit the supported version to Android 9.0 and above. This would have the downside of limiting the participant pool but it would ensure that users have the same consistent experience because different versions of the Android OS require different implementations to preform the functionality currently present in Guidance.

A problem that was encountered during the project was with the Location data type. If this was enabled, after several days the application would crash when attempting to provide advice to the user.

This would then cause the Android OS to unscheduled the jobs scheduled by Guidance, resulting in the jobs used to store data not taking place. However, the main problem was that advice was not being provided to the user, meaning that the data collected from the study became invalidated. Therefore, the study needed to be restarted. I then fixed the problem and restarted the project to ensure that results obtained were representative of a real-life situation. When restarting the project, the previous study data was overwritten. If I create a new study I will learn from my mistakes and ensure that the application has a better testing phase before being deployed. One way in which I would ensure that there is a better testing phase would be by the use of alpha testers who would test pilot the application in real world scenarios to identify potential flaws within it.

However, if given the chance to redo the study, I would definitely change some aspects of the study. The main aspect of the study that I would change would depend on the number of participants available. Having more participants would allow more attributes to be tested. For example, if there were more than 20 participants available, I would likely test three different attributes against each other. These attributes would be the Analysis method, Advice type and IA Gender. However, if only 6 participants were available, I would only test the Advice type attribute, which would compare if users are more likely to follow advice if justification for the advice is given. In hindsight I should have done this when I restarted the study as it would have been difficult to compare the Analysis attribute and the Advice type attribute against each other. It would have been difficult to compare the Analysis and Advice attributes against each other because there only would have been one participant on either side of the testing pool for these attributes, meaning that any conclusion made would likely be made with incomplete results.

At the end of the study when evaluating the results, I found that all of the advice given to participants had no verdict. By inspecting the functionality responsible in Guidance for creating the verdict of the results, I found several bugs which prevented the advice given to participants from being assigned a verdict. The main problem within the code was with the updateAdviceUsageData function which was responsible for updating the AdviceUsageData realm model with the verdict. The verdict was unable to be updated because the query was being filtered by the AdviceUsageData's "dateTime" field. The AdviceUsageData did not have a "dateTime" field and therefore no results were being presented to the function to update the verdict. This was fixed by changing the query to only look for a specific object ID and to filter it to the first one it found with that ID. Further testing should be done on this to ensure that all future advice given can be updated with a verdict. I have learnt that I should have tested this section more thoroughly before moving on to the implementation of the next section of the application. In the future I will ensure that each section that is implemented is tested thoroughly through the use of automated test scripts. I will also ensure that these scripts are run after every section is implemented to diagnose potential problems with the application as they occur.

I had prior experience in Android app development and wanted to take advantage of this knowledge when creating the Guidance application, rather than starting from scratch with development of an iOS application. I think that this was the correct decision to make as my familiarity with Android development allowed the application to be implemented faster. This also allowed me to increase my knowledge of the Android ecosystem. By understanding how each version of Android differs, I will be able to make more informed design and implementation decisions when creating new Android applications. When developing the application, I needed to interact with cloud service applications to upload the usage data generated from the study. This experience has allowed me to understand the authentication process used within cloud service applications which will aid me in the future when designing applications that will need to interface with them.

Glossary

Attribute: Key characteristics that define an Intelligent Agent

Advice Attribute: If users will take advice at face value or does additional information (justification) for the advice help

Agent: Something that acts in an environment

Analysis Attribute: User’s perceived method in which the Intelligent Agent creates advice

Data Type: Refers to a type of data that is collected by Guidance to be used when providing advice to the user.

Interaction Attribute: The amount of advice the Intelligent Agent provides to the participant

Gender Attribute: The gender of the Intelligent Agent

Guidance: The name of the application used in the study

Intelligent Agent: A system that acts intelligently

Machine Learning (Analysis attribute): The use of an artificially made system which learns as it provides advice

No Justification (Advice Attribute): No reasoning is given to the user when providing advice

Output Attribute: The method in which the Intelligent Agent communicates advice to the user

Traditional Programming (Analysis attribute): The use of human made code to provide advice to the user

With Justification (Advice attribute): A reason for the advice given is provided

Table of Abbreviations

AI – Artificial Intelligence

APK file – Android Application Package

IA- Intelligent Agent

iOS- iPhone Operating System

NLP – Natural Language Processing

OS- Operating System

Appendices

Appendix 1: Questionnaire One: See supporting PDF document: “Questionnaire 1- Human-Intelligent Agent Interaction for behaviour change”

Appendix 2: Questionnaire Two: See supporting PDF document: “Questionnaire 2- Human-Intelligent Agent Interaction for behaviour change”

Appendix 3: Participant information sheet. See supporting PDF document: “Participant Information Sheet- Human-Intelligent Agent interaction for behaviour change”

Appendix 4: Participant consent form. See supporting PDF document: “Consent Form- Human-Intelligent Agent Interaction for behaviour change”

Appendix 5: Guidance application source code. The GitHub repository for Guidance can be found at: https://github.com/11kempCp/Guidance

Appendix 6: Participant Usage Data: See folder in supporting documents: “Participant Usage Data”

Appendix 7: Research Training Certificate: See supporting PDF document: “RIOTP- Conor Kemp”

References

Android Developers. 2019. Sensors Overview. Available at: https://developer.android.com/guide/topics/sensors/sensors_overview [Accessed: 05/05/2021]

Ben-Ner A. Halldórsson F. 2010. Trusting and trustworthiness: What are they, how to measure them, and what affects them. Journal of Economic Psychology. DOI: https://doi.org/10.1016/j.joep.2009.10.001

Ben-Ner A. & Putterman, L. (2001). Trusting and trustworthiness. Boston Law Review, 81, 523–551.

Botsman R. 2020. How to think, not what to think. [Linkedin] Date Written: 27/01/2020. Available at: https://www.linkedin.com/pulse/introducing-rethink-rachel-rachel-botsman [Accessed: 09/05/2021]

David Poole et al. 1998. Computational Intelligence: A Logical Approach. 1st Ed. New York: Oxford University Press

Google. 04/10/2016. #madebygoogle. Available at: https://www.youtube.com/watch?v=q4y0KOeXViI [Accessed 05/05/2021]

Gsmarena. 2013a. Samsung Galaxy Note 3. https://www.gsmarena.com/samsung_galaxy_note_3-5665.php [Accessed: 05/05/2021]

Gsmarena. 2013b. Samsung Galaxy S4. https://www.gsmarena.com/samsung_i9500_galaxy_s4-5125.php [Accessed: 05/05/2021]

Hardin, R. 2002. Trust and trustworthiness. New York. Russell Sage Foundation.

Liao B. et al. Representation, justification, and explanation in a value-driven agent: an argumentation-based approach. AI Ethics 1, 5–19 (2021). https://doi.org/10.1007/s43681-020-00001-8

Lipika G. 2016. How Realm is Better As Compared To SQLite?. Available at: http://blogs.quovantis.com/how-realm-is-better-as-compared-to-sqlite/#:~:text=Realm%20does%20not%20use%20SQLite,a%20transactional%20SQL%20database%20engine. [Accessed: 05/05/2021]

Longoni C. Cian L. 2020. Artificial Intelligence in Utilitarian vs. Hedonic Contexts: The “Word-of-Machine” Effect. Journal of Marketing. DOI: https://doi.org/10.1177/0022242920957347

Nield David. 2020. How to Install Apps From Outside Your Phone's App Store. Available at: https://www.wired.com/story/install-apps-outside-app-store-sideload/ [Accessed: 05/05/2021]

O’Dea S. 2021. Mobile operating systems' market share worldwide from January 2012 to January 2021. Available at: https://www.statista.com/statistics/272698/global-market-share-held-by-mobile-operating-systems-since-2009/ [Accessed: 05/05/2021]

OpenWeather. 2021. Weather API. Available at: https://openweathermap.org/api [Accessed: 05/05/2021]

Özer et al. 2018. Trust and Trustworthiness. The Handbook of Behavioral Operations. 489-523. DOI: https://doi.org/10.1002/9781119138341.ch14

Prahl A, Van Swol L. Understanding algorithm aversion: When is advice from automation discounted?. Journal of Forecasting. 2017. https://doi.org/10.1002/for.2464

Shanhong L. 2021. Market share of mobile operating systems in the United Kingdom (UK) from 2010 to 2020. Available at: https://www.statista.com/statistics/487373/market-share-mobile-operating-systems-uk/#:~:text=iOS%20had%20the%20highest%20usage,point%20lower%20with%2049.87%20percent.[Accessed: 05/05/2021]

SQLite. 2021. History Of SQLite Releases. Available at: https://www.sqlite.org/chronology.html [Accessed: 05/05/2021]

Tschopp M. Ruef M. 2020a. AI & Trust - Stop asking how to increase trust in AI.

Tschopp M. Ruef M. 2020b. Trust and AI -Three Wrong Questions