Operationalizing Behavior Change Techniques in Conversational Agents

Maria Inês Bastos[1][a], Ana Paula Cláudio[1][b], Isa Brito Félix[2][c], Mara Pereira Guerreiro[2][d], Maria Beatriz Carmo[1][e] and João Balsa[1][f]

[1] *LASIGE, Dep. Informática, Faculdade de Ciências, Universidade de Lisboa, Portugal*

[2] *Nursing Research, Innovation and Development Centre of Lisbon (CIDNUR), Nursing School of Lisbon, Lisbon, Portugal*

Keywords: Conversational Agents, Ontologies, Behavior Change Techniques, Dialogue Platforms.

Abstract: Departing from previous work on the use of well-established behavior change techniques in an mHealth intervention based on a conversational agent (CA), we propose in this contribution a new architecture for the design of behavior change CAs. This novel approach combines the use of an advanced natural language platform (Dialogflow) with the explicit representation, in an ontology, of how behavior change techniques can be operationalized. The integration of these two components is explained, as well as the most challenging aspect of using the advanced features of the platform in a way that allowed the agent to lead the dialogue flow, when needed. A successful proof of concept was built, which can be the basis for the development of advanced conversational agents, combining natural language tools with ontology-based knowledge representation.

1 INTRODUCTION

We present a novel approach to the development of conversational agents that combines the design and development features offered by advanced natural language tools with the use of knowledge needed to support agents in the pursuit of their goals.

Besides the importance of building more advanced and versatile conversational agents, the drive for this work comes from other two dimensions: the importance of the role of conversational assistants in healthcare (Guerreiro et al., 2021) and the incorporation of mechanisms that allow agents to induce behavior change in their interlocutor.

In this work, we focus on how to provide agents with the ability of operationalizing specific *behavior change techniques* (BCTs), which are components of an intervention designed to change behavior (Michie et al., 2013).

The fact that noncommunicable diseases, like Type 2 diabetes (T2D), account for a considerable number of deaths (71% of deaths worldwide in 2016 (World Health Organization, 2021)) is a key indicator of the importance of developing mechanisms that help people preventing and managing this type of diseases. Besides, the consistent growing availability of mobile devices led to the development of thousands of mHealth apps (mobile health applications); in a recent scoping review (Wattanapisit et al., 2020), the authors concluded that tasks such as 'disease-specific care' and 'health promotion' can be successfully fulfilled with the support of mHealth apps.

As behavior change interventions require multiple interactions with a patient/user, the idea of a relational agent (one that is designed to build and maintain long-term social-emotional relationships) (Bickmore et al., 2005) became crucial since the development of the first stage of our work (Félix et al., 2019; Balsa et al., 2020). In our previous work, we developed a rule-based prototype of a mobile application with an intelligent virtual assistant to be used in an intervention to promote the self-care of older people with T2D.

In order to extend this work so that it could be used in a wider range of situations (other chronic diseases, multi-morbidity, or targeting diverse types of users, for instance), we had to overcome two of its limitations: the input from the user had to be chosen from a limited set of options, and the fact that the BCTs were to a greater than ideal extent hard coded in the dialogues definition.

[a] https://orcid.org/0000-0002-5200-4981
[b] https://orcid.org/0000-0002-4594-8087
[c] https://orcid.org/0000-0001-8186-9506
[d] https://orcid.org/0000-0001-8192-6080
[e] https://orcid.org/0000-0002-4768-9517
[f] https://orcid.org/0000-0001-8896-8152

Bastos, M., Cláudio, A., Félix, I., Guerreiro, M., Carmo, M. and Balsa, J.
Operationalizing Behavior Change Techniques in Conversational Agents.
DOI: 10.5220/0010826800003116
In *Proceedings of the 14th International Conference on Agents and Artificial Intelligence (ICAART 2022) - Volume 1*, pages 216-224
ISBN: 978-989-758-547-0; ISSN: 2184-433X

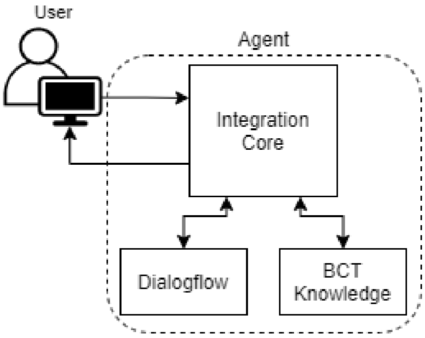

Figure 1: Simplified system's architecture.

Continuing this work, we developed, and present here, a new type of agent that allows us to overcome the above mentioned limitations, while illustrating it with BCTs operationalization.

As it is sketched in Figure 1, in which we present a simplified view of the architecture, the use of Dialogflow[1] allowed us to deal with the first limitation. To tackle the second limitation, we added a module that allows the explicit representation of knowledge on BCTs operationalization.

The core sections of this paper will describe in detail these components and how they interconnect. But first, in section 2, we highlight and discuss some related work. In section 3, we present the knowledge component of our agent, namely the ontology that allows us to better characterize BCTs operationalization. In section 4, we describe the dialogue engine and how Dialogflow was used. In section 5, we present the detailed architecture and explain how the agent works, illustrating it with a demo in section 6. Finally, in section 7, we present some conclusions and directions for future work.

2 RELATED WORK

Our work derives from contributions in several diverse scientific areas. In Figure 2 we schematize how those different areas directly or indirectly contribute to our work.

The widespread use of mobile equipment across people from all generations has been responsible for the dissemination of mobile applications with diverse purposes. Among these, a great number of developments were made in mHealth applications, the ones that somehow "support the achievement of health objectives" (World Health Organization, 2011). A research carried out by Bhuyan et al. concluded that 60% of the adults who were mHealth apps users considered them useful in achieving health behavior goals (Bhuyan et al., 2016). Moreover, mHealth applications are often used by people with little experience of technology, as noted by Zapata et al. (Zapata et al., 2015). So, mHealth offers a vast terrain of opportunity to design digital solutions as tools to help improve self-care, which can have a positive effect in individual health outcomes and reduce the burden on health systems. However, the conception of such applications entails high responsibility and the process must be grounded on valid scientific studies.

[1]A platform for the development of natural language conversation interfaces (https://cloud.google.com/dialogflow/)

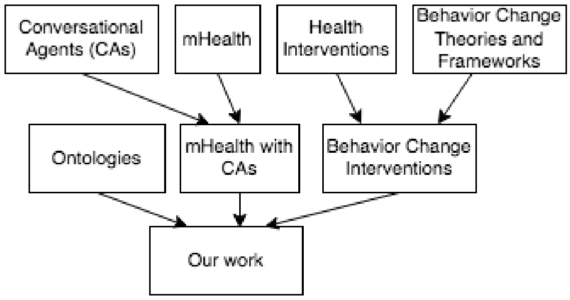

Figure 2: Context of our contribution.

The use of embodied conversational agents (ECA) playing the role of virtual assistants has proved to be well accepted by users in several health applications (Baptista et al., 2020; Gong et al., 2020). In particular, Baptista et al. conducted a study to evaluate the acceptability of an ECA (called Laura) to deliver self-management education and support for patients with T2D. The users are prompted to complete weekly interactive sessions with Laura. This virtual assistant provides education, feedback and motivational support for glucose level monitoring, taking medication, physical activity, healthy eating and foot care. The study had 66 respondents (mean age 55 years) and the results identified positive reaction of the majority of users for having a friendly, nonjudgmental, emotional and motivational support provided by a human-like character. Just around a third of them considered Laura not real, boring and annoying (Baptista et al., 2020). Users respond to Laura by choosing or by speaking out one of the options displayed in the screen. The authors refer a "sophisticated script logic" enabling the user to interact with Laura in several predetermined variations. The solution was implemented by a company that provides chatbots and no details are given about these scripts. A more recent 12 months trial with 187 adults with T2D (age around 57) study resorting to the same solution, performed during 12 months involving 187 adults suffering from T2D (mean age 57), found a successful adoption of the program and a significant improvement in participants' health-related quality of life (HRQOL) (Gong et al., 2020).

A recent systematic review about Artificial Intelligence-Based Conversational Agents for Chronic Conditions (Schachner et al., 2020) revealed the immaturity of the field, using the authors' own words. They concluded that there is a lack of evidence-based evaluation of the solutions, most of them quasi-experimental studies. These authors suggest that important future research would be the definition of the AI architecture that should be adopted and the adequate assessment process for the overall solution.

As mentioned before, our research started some time ago (Buinhas et al., 2019) with the development of a first prototype. Key project features were resorting to an anthropomorphic assistant with a relevant role in the interaction with the user, and to support the intervention design in a well-defined theoretical framework, the Behavior Change Wheel (Michie et al., 2014). This theoretical framework provides guidance to choose the more adequate behavior change techniques in specific contexts.

A BCT is an observable, replicable, and irreducible component of an intervention designed to change behavior. They have been organized in a taxonomy (BCTTv1) (Michie et al., 2013; Cane et al., 2015), which provides a definition for each of the 93 techniques. An example of a BCT is "problem solving"; it requires users to pinpoint factors influencing the behavior and subsequently select strategies to achieve it. After surveying 166 medication adherence apps to ascertain whether they incorporated BCTs, Morrissey et al. concluded that, from the 93 possible techniques only a dozen were found in the evaluated apps (Morrissey et al., 2016). This result clearly shows that more work is needed in incorporating evidence on BCTs in available applications.

Although conversational agents (CA) have been researched for decades, namely since the seminal ELIZA (Weizenbaum, 1966), the shift towards behavior change is more recent (for instance, to promote a healthier lifestyle – typically regarding physical activity and type of diet). As reported by (Kramer et al., 2020), CAs were found to have an important value regarding the use of persuasive communication in the health domain (Kramer et al., 2020), namely when targeting coaching tasks. As Zhang et al. point out, there is a "lack of understanding around theoretical guidance and practical recommendations on designing AI chatbots for lifestyle modification programs" (Zhang et al., 2020). In order to overcome this, these authors developed an AI chatbot behavior change model that has persuasive conversational capacity as a central component. Our work has this same feature, although the architecture we present goes beyond the coverage of just two specific topics, as Zhang and his colleagues do.

Regarding the use of ontologies, it is worth mentioning two recent works that, in a different way, stress the importance of its use in the context of our work.

Within the *Human Behaviour-Change Project* (Michie et al., 2017), an ontology is being developed for representing behavior change interventions (BCIs) and their evaluation (Michie et al., 2021). Although it focuses on the more general aspects that characterize BCIs, it clearly opens the possibility of a link to our work, via the common BCT concept that, having a different perspective, has in our work the grounds for the representation of its operationalization.

Also relevant to us is the recent work on the development of dialogue managers combining ontologies and planning. Teixeira et al. combine a conversational ontology and Artificial Intelligence planning to generate dialogue managers capable of performing goal-oriented dialogues in the health domain (Teixeira et al., 2021). As the authors note, in the health domain it is critical to have predictable and reliable systems, making a knowledge component crucial to guarantee that, even if complemented, as it is in our case, with other resources, namely for natural language understanding tasks.

2.1 Dialogue Engines

In order to create our conversational agent, we chose to adopt an existing tool, since many excellent tools are currently available. The two main features we wanted were: the possibility of having natural language interaction and some control of the dialogue flow, in order to address situations where the agent is the one responsible for leading the dialogue (and not the opposite, like what happens in question/answering contexts, for instance).

Several dialogue tools were analyzed and tested before choosing the most adequate. Due to space limitation, we will just enumerate the tools we considered: Twine[2], Yarn Spinner[3], WOOL[4], Watson Assistant[5], FAtiMA Toolkit[6], and Dialogflow.

As stated before, in this work, one of our main interests was to make the dialogue more dynamic.

[2]https://twinery.org/
[3]https://yarnspinner.dev/
[4]https://www.woolplatform.eu/
[5]https://www.ibm.com/cloud/watson-assistant
[6]https://fatima-toolkit.eu/

Dialogflow, being a natural language understanding platform, has proven to be quite useful in that aspect, making it easier to build and deliver the agent messages and capturing the user's response. In previous work, the dialogue portions corresponding to BCTs operationalization were initially created in YARN (Balsa et al., 2020). Besides the fact that defining dialogs with that type of tool is a time-consuming task, the definition of the whole logic of the interaction was also harder than what we have with Dialogflow. Besides, with this tool, the only user dialogue we have to provide is the training phrases, and the agent will learn from over time and usage.

3 AN ONTOLOGY OF BEHAVIOR CHANGE CONCEPTS

One of the main limitations of the antecedent agent was that BCT operationalization was too rigid, i.e. the dialogue flow was too dependent on the way the corresponding dialogues were defined. In order to make dialogues more natural and flexible, we decided to characterize BCT operationalization in a more general way. For that, we needed to identify the main concepts involved and the relations between them. We did this by defining an ontology representing the relevant behavior change intervention concepts. Since the operationalization of a BCT depends also on user specific information, in order to characterize a specific way of operationalizing a BCT we had to incorporate concepts related to the user's characteristics.

The resulting ontology can be divided in two main class entities: *Behavior Change Intervention* (BCI) and *User*. The first one, BCI, includes several classes representing both general concepts (like *BCT* or *Behavior determinant* – see below) and specific ones (like *Food Topic*, representing the relevant topics related to a healthy diet). As a starting point, our illustrative domain is healthy nutrition, as it should be an universal concern, independent of age or health condition. The *User* class includes concepts critical to the operationalization, like the age category (adult, senior, ...) or the identification of some risk condition.

Some of these entities are related by six object properties: *has active goals, has active topic, related to, targeted during, triggered by*, and *triggers*. For further characterization, there are also eight data properties: *BCT order*, *has age*, *has BMI* (Body Mass Index), *has competence score*, *has genre*, *has height*, *has risk level*, and *has weight*.

The classes that the current work focuses on and are: *Behavior Change Intervention* (an intervention that has the aim of influencing human behavior); *Behavior Change Technique* (an observable, replicable, and irreducible component of an intervention designed to alter or redirect causal processes that regulate behavior — an "active ingredient" on an intervention (Michie et al., 2013); *Behavior Determinant* (a factor that influences positively or negatively a behavior). For instance, a reason of non-adherence to the user's agreed goal – lack of motivation, forgetfulness, ...); *Behavior Goal* (a goal defined in terms of the behavior to be achieved); *Behavior Topic* (a topic that is targeted during the intervention, aimed to help achieve/maintain the desired behaviors); *Operationalization* (the act of delivering one or more BCTs, based on several conditions); **User** (a person who's being subjected to the Behavior Change Intervention).

3.1 How the Ontology Is Used

This ontology can be used for several purposes within the development of a dynamic Behavior Change Intervention, making it easier to connect the concepts involved and shape the structure of the intervention.

When creating an application aimed at influencing the human behavior, the use of an ontology can be useful to detach the logic of the application from the specific data (the individuals and their relationships).

In order to understand better, the following paragraphs describe a scenario where the ontology is used during the interaction with the user. Figure 3 shows the classes and relationships that are relevant to the scenario described.

During *Review Tasks* (one of the dialogue stages, as explained in the next section), the user is asked whether he/she completed or not a previously agreed goal. When the user answers negatively, the agent will ask about the reason for not meeting the goal, and the reason the user gives is denominated determinant. The determinants of non-adherence can be as simple as *lack of motivation*, or *reluctance in changing their habits*, so the individuals of the class Behavior Determinant are key words of such motives (so far, the determinants included are motivation, habits, difficulty, appetite, and loneliness). These determinants are examples and are not related with a particular behavior.

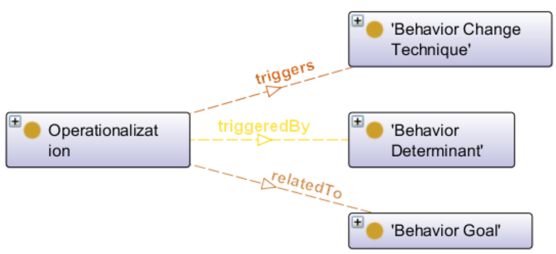

Figure 3: Ontology entities relevant to BCTs operationalization.

After the user gives an answer (providing a determinant), the agent queries the ontology to see which operationalization related to the active goal is triggered by that determinant. When the right operationalization is identified, the agent chooses the right BCT to trigger. If it is the case that the operationalization is sequential, the agent also determines the triggering order. After that, it searches the dialogue file and fetches the dialogue that is supposed to be delivered to complete the execution of a specific BCT.

4 DIALOGUE ENGINE

As stated before, we chose to use Dialogflow CX as the dialogue engine. In this section, we describe how our Dialogflow agent is built and used, along with the details regarding the training process and the way we deal with the situations where the agent takes control of the dialogue.

4.1 How Dialogflow Is Used

Dialogflow CX has a visual builder in its console, where conversation paths are graphed as a state machine model, making it easier to design, enhance, and maintain. Conversation states and state transitions are first-class types that provide explicit and powerful control over conversation paths. We can clearly define a series of steps that we want the end-user to go through. Following the work of Bickmore's relational agents group (Bickmore et al., 2005), dialogues are organized in eight steps: *opening, social talk, review tasks, assess, counseling, assign tasks, pre-closing* and *closing*. The visual state machine created for our agent is represented in Figure 4.

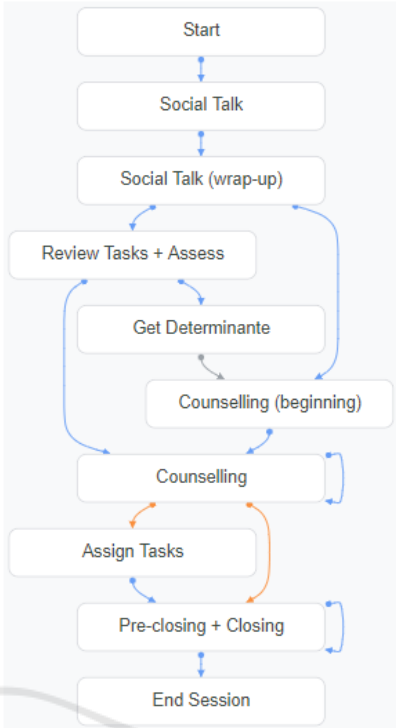

Figure 4: Conversation path of the Dialogflow agent.

Each state is represented by a *page*, that can be configured to collect information from the end-user that is relevant for the conversational state represented by the page. For simpler dialogue steps, such as opening, social talk, assess, pre-closing and closing, we configured the respective pages to give static text responses, since there are no parameters or external conditions that they depend on. For more complex dialogue steps, such as review tasks, counseling and assign tasks, we configured their respective pages to enable webhook, in order to provide a dynamic response based on external conditions.

To each page can be added *state handlers*, that are used to control the conversation by creating responses for the end-users and/or by transitioning the current page. There are two types of state handlers: *routes* and *event handlers*. Routes are called when an end-user input matches an intent (categorizes an end-user's intention for one conversation turn) and/or some condition.

Figure 5 shows the routes in our Counselling page. The first three routes (with the blue left border) are intent routes, which means that either one of those routes are called when the user input matches one of those intents. The last two routes (with the orange left border) are condition routes, that are called if the respective condition is satisfied. At the end of a route, there's a *Transition* field, which defines the next page in the conversation.

In other resources, and in Dialogflow ES, an intent usually contains a field with the training phrases and

Intent	Condition	Fulfillment	Transition
affirmative-answer			Counselling
negative-answer			Counselling
counselling			Counselling
	$session.params.counselling = "FINISH" AND $session.params.assignTasks = True		Assign Tasks
	$session.params.counselling = "FINISH" AND $session.params.assignTasks = False		Pre-closing + Closing

Figure 5: Counselling page.

a field with the agent response, that is sent back to the user if their input matches the training phrases for that intent. During an interaction, a user can say "yes" or "no" multiple times, which means that, with this type of intent, every time the user can say "yes" or "no" during the conversation, there would have to be an intent with the same training phrases but different agent response, depending on the context. In Dialogflow CX, an intent only contains the logic to detect what the user says (the training phrases), which means an intent can be reused in multiple places of a conversation.

4.2 Training Process

When an agent is trained, Dialogflow uses the training data to build machine learning models specifically for that agent. This training data primarily consists of intents, intent training phrases, and entities referenced in a flow, which are effectively used as machine learning data labels. However, agent models are built using parameter prompt responses, state handlers, agent settings, and many other pieces of data associated with the agent.

In the agent settings, it can be chosen to train the agent automatically or manually. By default, the training is executed automatically, showing a popup dialogue in the console every time there's an update of the flow.

4.3 Leading the Dialogue

Technologies like Dialogflow are generally used in virtual agents meant to handle questions from the user. In our case, it is important that the agent is the one taking control of the conversation, asking the questions, and delivering the content in a "doctor-patient" type of way, i.e. with the doctor leading the interaction.

On previous work, after the agent's response, the user interface had buttons for the user to choose their answer, and during longer explanations or a change in the dialogue phase, it would present a "Continue" button (Balsa et al., 2020). When the user has total control over the input, we can not expect them to casually answer that. The agent must seek the user's attention and interest, keeping the user engaged in the conversation.

In Dialogflow CX, all the agent responses are handled in page, and a page can have an entry fulfillment, and a static response for each route. The entry fulfillment is optional, and it is what the agent will respond to the end-user when a page initially becomes active. For each route added to a page, there is a fulfillment field, where it is possible to add several types of response messages, although, in this work we only used text response messages.

The entry fulfillment feature is a great advantage since it makes it possible for the agent to say something to the user without the need to have a previous input. This gives the agent more initiative and control over the conversation, being extremely useful between dialogue phases.

5 AGENT ARCHITECTURE

The architecture of our agent has comprises three main components: the *Core*, the *Dialogflow Engine*, and the *Ontology* (Figure 6). The *Core* controls the interface and, along with the *Dialogflow Engine*, controls the flow of execution and the speech of the agent. The *Ontology* holds external data that can be queried by the *Core* whenever necessary.

Figure 6 shows the steps that take place for one conversational turn of a session:

1. The end-user types something, known as user input.
2. The user interface (UI), responsible for the view provided to the user, receives the input and forwards it to the Dialogflow API in a detect intent request (handled by the Application Controller).
3. The Dialogflow API receives the detect intent request. It matches the user input to an intent or form parameter, sets parameters as needed, and updates the session state. In case it needs to call a webhook-enabled fulfillment, it sends a webhook request to the Webhook Service, otherwise it jumps straight to step 6.
4. The Webhook Service receives the webhook request and it takes any actions necessary, such as querying the ontology and/or fetching dialogue from external sources (JSON files).
5. The Webhook Service builds a response and sends a webhook response back to Dialogflow.
6. Dialogflow creates a detect intent response. If a webhook was called, it uses the response provided in the webhook response. If no webhook was called, it uses the static response defined in the Dialogflow Agent. The detect intent response is send to the user interface.
7. The user interface receives the detect intent response and forwards the text response to the end-user.
8. The end-user sees the response.

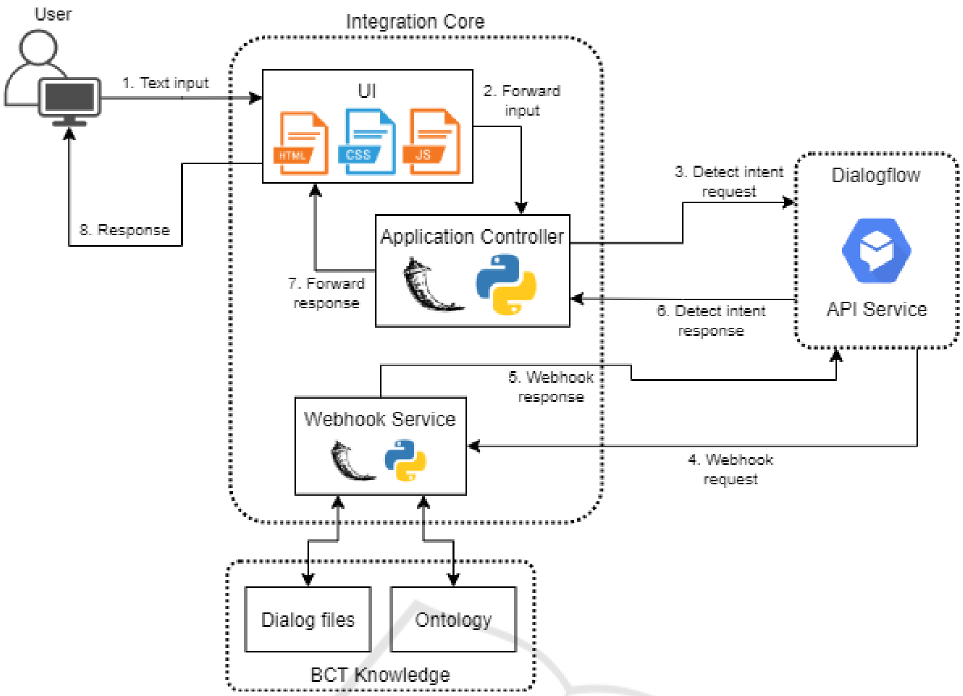

Figure 6: Agent architecture.

In short, when a user submits a message, it is sent to Dialogflow to detect the intent of the user. Dialogflow will process the text, then send back a fulfillment response (either static, or dynamic, by means of the Webhook Service).

6 DEMO

During a few dialogue steps, there are specific BCTs that are always operationalized in the same way, therefore being simpler to execute. For example, during *Review Tasks*, the BCT *Self-monitoring of behavior (2.3)*[7] is always executed and delivered in the same way: by collecting data on the user's behavior. To do that, our system takes advantage of the following ontology classes: *User* and *Behavior Goal*. Linking those classes is an object property, labeled *has active goals*, that connects one or more behavior goals to a specific user. That way, the system can easily get access to the goal that the user agreed on, and through its label, ask if they completed it or not.

During Assess and Counseling, the procedure is a bit more complex, since the BCTs operationalized on those dialogue steps are executed in several different ways, depending on more than just one condition.

[7]The numbers next to the names of the BCTs refer to the codes used in the taxonomy (Michie et al., 2013).

Following the case mentioned in the previous paragraph, after *Review Tasks*, comes *Assess*, and during that step, the system takes advantage of the following classes: *Operationalization, User, Behavior Goal, Behavior Determinant* and *Behavior Change Technique*. As was mentioned in Section 3.1, when the system gets the non-adherence determinant, it queries the ontology in order to find which operationalization, related to the active goal, is triggered by that given determinant. After that, the system accesses the object property *triggers*, to see which BCTs are triggered in that operationalization. In case the operationalization is of the *Sequential* type (a complex operationalization where the BCTs are delivered in sequential order), the system accesses the data property BCT order, to be able to deliver the BCTs in the suitable order.

Figure 7 illustrates an output demonstration of the examples given in the previous paragraphs. The agent asks the user if they completed the agreed goal ("have at least three main meals", highlighted in green on the first chat message on the left upper corner), to which the user answers negatively. After that, there is the operationalization of another simpler BCT, *Feedback on behavior (2.2)*, statically delivered by Dialogflow since it only depends on the intent ("Yes" or "No"). The user answers saying "I don't have motivation", to which the system catches the "motivation" behavior determinant. The operationalization related to the active goal, and triggered by motivation, is the individ-

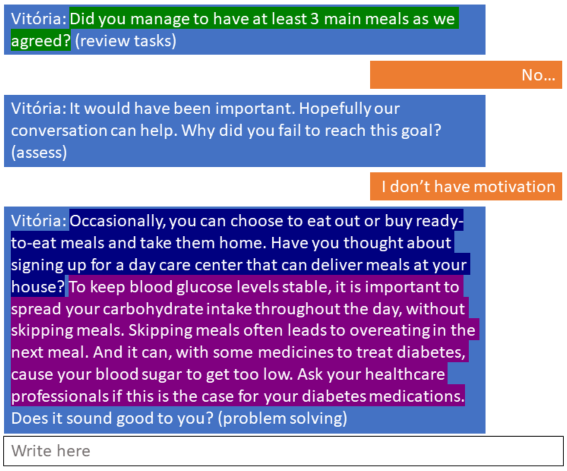

Figure 7: An example dialogue.

ual OP_Alim_Det1_2, which triggers the BCTs *Problem solving (1.2)* and *Information of health consequences (5.1)*, having "1.2, 5.1" as BCT order. Given those conditions, the system accesses the dialogue file and extracts the dialogue parts related to those BCTs, ordering them accordingly and sending them to the user. In the last chat message on the lower left corner of Figure 7, highlighted in dark blue is the dialogue delivering *Problem solving (1.2)*, and higlighted in purple is the dialogue part that delivers *Information of health consequences (5.1)*.

7 CONCLUSIONS

The main goal of this work was achieved with the design and implementation of a novel architecture for the development of conversational agents for behavior change interventions. This architecture allows the combination of an advanced dialogue engine with a learning ability (DialogFlow) with the representation of knowledge on the operationalization of behavior change techniques, by means of defining an ontology of behavior change intervention concepts.

The design was explained and the functioning of the agent illustrated.

Regarding future work, two immediate steps follow: the incorporation of additional knowledge so that the agent capacity can be enlarged; the additional training of the system with a set of dialogues that were developed in a previous work.

Additionally, we intend to incorporate in the design mechanisms that will allow the consideration of an ethical dimension in this type of agents. As recognized by Zhang and her colleagues (Zhang et al., 2020), the ethical dimensions regarding conversational agents development has been completely absent. But, in contexts where the goal is to induce behavior change in humans, the incorporation of ethical principles and the insurance of responsibility in the systems' designers is of paramount importance.

ACKNOWLEDGEMENTS

This work was supported by FCT through the LASIGE Research Unit, ref. UIDB/00408/2020 and ref. UIDP/00408/2020.

REFERENCES

Balsa, J., Félix, I., Cláudio, A. P., Carmo, M. B., Silva, I. C. e., Guerreiro, A., Guedes, M., Henriques, A., and Guerreiro, M. P. (2020). Usability of an Intelligent Virtual Assistant for Promoting Behavior Change and Self-Care in Older People with Type 2 Diabetes. *Journal of Medical Systems*, 44(7):130.

Baptista, S., Wadley, G., Bird, D., Oldenburg, B., Speight, J., and Group, T. M. D. C. R. (2020). Acceptability of an Embodied Conversational Agent for Type 2 Diabetes Self-Management Education and Support via a Smartphone App: Mixed Methods Study. *JMIR mHealth and uHealth*, 8(7):e17038.

Bhuyan, S. S., Lu, N., Chandak, A., Kim, H., Wyant, D., Bhatt, J., Kedia, S., and Chang, C. F. (2016). Use of Mobile Health Applications for Health-Seeking Behavior Among US Adults. *Journal of Medical Systems*, 40(6):153.

Bickmore, T. W., Caruso, L., Clough-Gorr, K., and Heeren, T. (2005). 'It's just like you talk to a friend' relational agents for older adults. *Interacting with Computers*, 17(6):711–735.

Buinhas, S., Cláudio, A. P., Carmo, M. B., Balsa, J., Cavaco, A., Mendes, A., Félix, I., Pimenta, N., and Guerreiro, M. P. (2019). Virtual Assistant to Improve Self-care of Older People with Type 2 Diabetes: First Prototype. In García-Alonso, J. and Fonseca, C., editors, *Gerontechnology*, Communications in Computer and Information Science, pages 236–248. Springer International Publishing.

Cane, J., Richardson, M., Johnston, M., Ladha, R., and Michie, S. (2015). From lists of behaviour change techniques (BCTs) to structured hierarchies: Comparison of two methods of developing a hierarchy of BCTs. *British Journal of Health Psychology*, 20(1):130–150.

Félix, I. B., Guerreiro, M. P., Cavaco, A., Cláudio, A. P., Mendes, A., Balsa, J., Carmo, M. B., Pimenta, N., and Henriques, A. (2019). Development of a Complex Intervention to Improve Adherence to Antidiabetic Medication in Older People Using an Anthropomorphic Virtual Assistant Software. *Frontiers in Pharmacology*, 10:680.

Gong, E., Baptista, S., Russell, A., Scuffham, P., Riddell, M., Speight, J., Bird, D., Williams, E., Lotfaliany, M., and Oldenburg, B. (2020). My Diabetes Coach, a Mobile App–Based Interactive Conversational Agent to Support Type 2 Diabetes Self-Management: Randomized Effectiveness-Implementation Trial. *Journal of Medical Internet Research*, 22(11):e20322.

Guerreiro, M. P., Angelini, L., Rafael Henriques, H., El Kamali, M., Baixinho, C., Balsa, J., Félix, I. B., Khaled, O. A., Carmo, M. B., Cláudio, A. P., Caon, M., Daher, K., Alexandre, B., Padinha, M., and Mugellini, E. (2021). Conversational Agents for Health and Wellbeing Across the Life Course: Protocol for an Evidence Map. *JMIR Research Protocols*, 10(9):e26680.

Kramer, L. L., ter Stal, S., Mulder, B. C., de Vet, E., and van Velsen, L. (2020). Developing Embodied Conversational Agents for Coaching People in a Healthy Lifestyle: Scoping Review. *Journal of Medical Internet Research*, 22(2):e14058.

Michie, S., Atkins, L., and West, R. (2014). *The behaviour change wheel : a guide to designing interventions*. Silverback.

Michie, S., Richardson, M., Johnston, M., Abraham, C., Francis, J., Hardeman, W., Eccles, M. P., Cane, J., and Wood, C. E. (2013). The Behavior Change Technique Taxonomy (v1) of 93 Hierarchically Clustered Techniques: Building an International Consensus for the Reporting of Behavior Change Interventions. *Annals of Behavioral Medicine*, 46(1):81–95.

Michie, S., Thomas, J., Johnston, M., Aonghusa, P. M., Shawe-Taylor, J., Kelly, M. P., Deleris, L. A., Finnerty, A. N., Marques, M. M., Norris, E., O'Mara-Eves, A., and West, R. (2017). The Human Behaviour-Change Project: harnessing the power of artificial intelligence and machine learning for evidence synthesis and interpretation. *Implementation Science*, 12(1):121.

Michie, S., West, R., Finnerty, A., Norris, E., Wright, A., Marques, M., Johnston, M., Kelly, M., Thomas, J., and Hastings, J. (2021). Representation of behaviour change interventions and their evaluation: Development of the Upper Level of the Behaviour Change Intervention Ontology. *Wellcome Open Research*, 5(123).

Morrissey, E. C., Corbett, T. K., Walsh, J. C., and Molloy, G. J. (2016). Behavior Change Techniques in Apps for Medication Adherence. *American Journal of Preventive Medicine*, 50(5):e143–e146.

Schachner, T., Keller, R., and v Wangenheim, F. (2020). Artificial Intelligence-Based Conversational Agents for Chronic Conditions: Systematic Literature Review. *Journal of Medical Internet Research*, 22(9):e20701.

Teixeira, M. S., Maran, V., and Dragoni, M. (2021). The interplay of a conversational ontology and AI planning for health dialogue management. In *Proceedings of the 36th Annual ACM Symposium on Applied Computing*, pages 611–619, Virtual Event Republic of Korea. ACM.

Wattanapisit, A., Teo, C. H., Wattanapisit, S., Teoh, E., Woo, W. J., and Ng, C. J. (2020). Can mobile health apps replace GPs? A scoping review of comparisons between mobile apps and GP tasks. *BMC Medical Informatics and Decision Making*, 20(1):5.

Weizenbaum, J. (1966). ELIZA—a computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1):36–45.

World Health Organization (2011). *mHealth : new horizons for health through mobile technologies*. World Health Organization, Geneva.

World Health Organization (2021). Noncommunicable diseases: Mortality. [online]. https://www.who.int/data/gho/data/themes/topics/topic-details/GHO/ncd-mortality [accessed 2021-09-28].

Zapata, B. C., Fernández-Alemán, J. L., Idri, A., and Toval, A. (2015). Empirical Studies on Usability of mHealth Apps: A Systematic Literature Review. *Journal of Medical Systems*, 39(2):1.

Zhang, J., Oh, Y. J., Lange, P., Yu, Z., and Fukuoka, Y. (2020). Artificial Intelligence Chatbot Behavior Change Model for Designing Artificial Intelligence Chatbots to Promote Physical Activity and a Healthy Diet: Viewpoint. *Journal of Medical Internet Research*, 22(9):e22845.