Towards Building a Virtual Assistant Health Coach

1st Itika Gupta
Department of Computer Science
University of Illinois at Chicago
Chicago, USA
igupta5@uic.edu

2nd Barbara Di Eugenio
Department of Computer Science
University of Illinois at Chicago
Chicago, USA
bdieugen@uic.edu

3rd Brian Ziebart
Department of Computer Science
University of Illinois at Chicago
Chicago, USA
bziebart@uic.edu

4th Bing Liu
Department of Computer Science
University of Illinois at Chicago
Chicago, USA
liub@uic.edu

5th Ben Gerber
Institute for Health Research and Policy
University of Illinois at Chicago
Chicago, USA
bgerber@uic.edu

6th Lisa Sharp
Institute for Health Research and Policy
University of Illinois at Chicago
Chicago, USA
sharpl@uic.edu

7th Rafe Davis
Institute for Health Research and Policy
University of Illinois at Chicago
Chicago, USA
rdavis5@uic.edu

8th Aiswarya Baiju
Department of Computer Science
University of Illinois at Chicago
Chicago, USA
abaiju2@uic.edu

***Abstract*—In this paper, we present our data collection process, annotation schemas and agreement results for extracting health goals from SMS conversations between a health coach and the patients. We also discuss our preliminary results for automatically detecting topic boundaries in health coaching dialogues. This is our first step towards building an autonomous virtual assistant health coach that learns from expert demonstration to interact with patients via SMS.**

***Index Terms*—Virtual Assistant, Health Coach, SMART Goal Setting**

I. Introduction

Health coaching is an established process for improving poor health behaviors by providing education on health-related topics, setting personalized and realizable health-related goals, monitoring and encouraging progress towards those goals, and sequencing or refining a progression of health goals over time. Automated coaching systems have been developed with the aim of improving health coaching accessibility for millions of people who could benefit from them. The simplest methods involve programmable prompting devices [1], which remind the participant at a pre-specified time in a pre-specified manner. These methods have been shown to be effective, particularly for medicine adherence [2]. WebDietAID [3] is a more sophisticated interactive online nutrition counseling service that seeks to imitate "the types of intervention performed by the nutrition counselor." An influential strand of work on conversational agents, namely, systems that participate in an interaction with users, has been conducted by Bickmore and his group [4]–[6]. However, these systems are typically pre-programmed and often limited to predefined scripts of communicative actions, making them inflexible for automatically responding to situations beyond those anticipated by the system designer and for fully personalizing to the individual participant. Motivated by the limitations of existing systems, we aim to develop an autonomous health coaching system that learns to automate an increasing amount of the coach-participant interactions without programmatically defining appropriate behaviors. We believe the ability to recognize proposed health goals in dialogues will help us to progress towards our aim of building such a system.

In this paper, we will first talk about our data collection process, two annotation schemas, and inter-annotator agreement results. We will then use these annotations to build a model for topic boundary detection. Lastly, we will discuss our future plans for building the complete pipeline for extracting goals.

II. Data Collection and Annotations

A key requirement for enabling a system to effectively learn from demonstration is access to human demonstrations from situations that are similar to those that the autonomous system will encounter. Therefore, we recruited a health coach and 28 patients to collect conversation data, where the health coach helped the patients to set SMART (Specific, Measurable, Attainable, Realistic, Time Bound) goals on a weekly basis for one month. The patients were given Fitbit Alta to monitor their progress. The coach used patients' online Fitbit accounts to monitor each patient's progress. Only one patient did not complete the study. We have a corpus of 2858 messages. Among these messages, approximately 53% messages were sent by the coach and 47% messages by the patients. Below we will look into two types of annotation schemas designed

This work is supported by National Science Foundation, Early-concept Grants for Exploratory Research (EAGER) program (Award Abstract 1650900)

by us to help automate the extraction of health goal set by the user.

A. SMART Goal Annotation

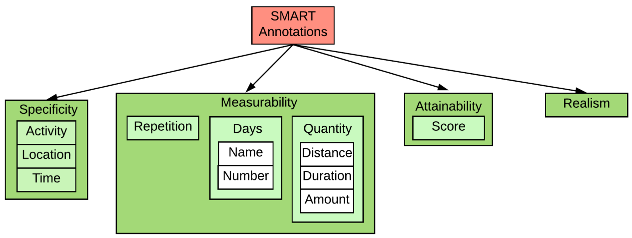

Fig. 1. SMART goal annotation schema

Annotation schema for annotating Specificity, Measurability, Attainability and Realism of a goal is shown in Fig 1. We didn't annotate for Timeliness as the patients set a new goal every week. Each of the SMART annotation category can have some additional optional tags such as *previous* when referring to previous week's goal, *accomplished* or *remaining* to annotate progress, *update* to add another goal to an existing one and *other* for anything else not covered by our schema. In order to perform inter-annotator agreement, two annotators annotated four unseen patients' data (447 messages) and achieved kappa score of above 0.69 on each individual {S, M, A, R} category.

B. Stages and Phases Annotations

In order to understand the structure of health coaching dialogues, we annotate the conversations for stages and phases. Stages capture the higher level information of coaching dialogue and comprises of *goal setting* and *goal action* stages. Where as phases capture the lower level information within each stage. Goal setting stage consists of *goal identification*, *goal refining*, *goal negotiation*, *anticipate barrier* and *solve barrier* phases. Goal action stage consists of the same phases plus an additional *follow up* phase. Two annotators annotated four unseen patients' data (398 messages) for all the possible stages-phases, and achieved an inter-annotator agreement of 0.93.

III. Topic Boundary Detection

Coaching dialogues can be very complex as it might involve a lot of negotiation between the patient and the coach before a goal is finalized. The conversation might also involve updates to the goal in case the patient encounters a barrier such as weather change, illness, or work responsibilities. Therefore, the task of goal extraction becomes harder if the information is not localized at one place in the dialogue. Topic segmentation is one of the approaches to deal with such dialogues and has been around for almost two decades [7], [8]. Therefore, we used our stages and phases annotations from previous step, and formulated the problem of topic boundary detection as a binary classification problem. A message is assigned to class 1 if it marks the beginning of a new topic, else it is assigned to class 0. We ended up with 558 samples in class 1 and 2280 in class 0. As most of the researches has shown the usefulness of lexical and syntactic features for detecting coherence between two utterances, we used the following features to train our classifier: sender of the message, message length, Part-Of-Speech tags, time difference between two utterances, presence of stop words, and unigrams after lemmatization.

IV. Results

We divided the patients' data into training and testing using 80-20 rule, and performed 5-fold cross validation. We achieved an average F-score of 0.87 on our data. In order to test the performance of our model on an external data, we transcribed 5 videos on SMART health coaching from YouTube. It was a total of 241 utterances. Even though the data was collected from a completely different source and were from face-to-face interactions, we achieved an F-score of 0.77. We found from our experiments that Logistic Regression performed the best out of all other machine learning algorithms such as Support Vector Machines, Decision Tress, K-NN, and Naive Bayes. We expected Decision Trees to perform better as the data seemed to partition well along the feature space. But we found that the same reason caused Decision Trees to perform poorly as it over-fitted the training data, where as Logistic Regression was able to generalize the model well.

V. Future Work

In our next steps, we plan to build a pipeline where first all the possible SMART goals in a conversation are detected, and then topic segmentation module is used to detect topic boundaries. Then we will use language model to predict the most likely stage and phase tags for each of these topic boundaries, and use them to figure out the most recent and up to date goal of the patient. This pipeline will also help in future to keep track of dialogue states in a real time conversation.

References

[1] T.A. Schollmeyer and J.C. Elsmore. Pharmacist-programmable medication prompting system and method, 1985. US Patent 4,504,153.

[2] A.S. Andrade, H.F. McGruder, A.W. Wu, S.A. Celano, R.L. Skolasky Jr, O.A. Selnes, I.C. Huang, and J.C. McArthur. A programmable prompting device improves adherence to highly active antiretroviral therapy in HIV-infected subjects with memory impairment. Clinical Infectious Diseases, 41(6), pp.875-882, 2005.

[3] A. Riva, C. Smigelski and R. Friedman, 2000. WebDietAID: an interactive Web-based nutritional counselor. In Proceedings of the AMIA Symposium (p. 709). American Medical Informatics Association.

[4] D. Schulman, T.W. Bickmore and C.L. Sidner. An Intelligent Conversational Agent for Promoting Long-Term Health Behavior Change Using Motivational Interviewing. In AAAI Spring Symposium: AI and Health Communication, 2011.

[5] J. Ren, D. Schulman, B. Jack, and T.W. Bickmore. Supporting longitudinal change in many health behaviors. In CHI14 Extended Abstracts on Human Factors in Computing Systems, pages 16571662. ACM, 2014.

[6] T.W. Bickmore, D. Schulman, and C.L. Sidner. A reusable framework for health counseling dialogue systems based on a behavioral medicine ontology. Journal of Biomedical Informatics, 44(2):183197, 2011.

[7] J. Arguello and C. Rose . Topic segmentation of dialogue. In Proceedings of the HLT-NAACL 2006 Workshop on Analyzing Conversations in Text and Speech, pages 4249. Association for Computational Linguistics, 2006

[8] N. Boufaden, G. Lapalme, and Y. Bengio. Topic segmentation: A first stage to dialog-based information extraction. In In Natural Language Processing Pacific Rim Symposium, NLPRS01. Citeseer, 2001