Virtual agents supporting novice and expert professionals in healthcare

Dorian Miller
IBM Research-Almaden
San Jose, CA 95120, USA
millerbd@us.ibm.com

Robert J. Moore
IBM Research-Almaden
San Jose, CA 95120, USA
rjmoore@us.ibm.com

CHI 2020 Extended Abstracts, April 25–30, 2020, Honolulu, HI, USA.

ACM ISBN 978-1-4503-6819-3/20/04.
DOI: https://doi.org/10.1145/3334480.XXXXXXX
*update the above block & DOI per your rightsreview confirmation (provided after acceptance)

Abstract
Technology has advanced to make virtual conversational agents popular for general content. Research is taking on more advanced agents. We are building two agents for the healthcare domain. We need to invent new techniques for one agent to support users with varying degrees of expertise. We outline an approach for the system to capture the knowledge of the healthcare domain. Then we outline interaction designs for the agents to support doctors, nurses, pharmacist, benefits managers, and analysts in support of treating the patient.

Author Keywords
Conversational UX design; UX design; Conversation Analysis; Human-Computer Interaction; Business Intelligence; Healthcare.

CSS Concepts
• Human-centered computing, Human computer interaction (HCI) Interaction techniques •Applied computing; Life and medical sciences; Health care information systems

Introduction
How should virtual conversational agents, such as chatbots or voice assistants, handle users of varying degrees of knowledge and expertise? We explore this

question by designing conversational agents to support healthcare professionals, for example, who treat patients. Our new design builds on two virtual healthcare agents we are researching as part of product development in the IBM Watson Health organization.

Designing conversational interfaces for an interdisciplinary team of healthcare professionals is challenging. To illustrate there are significant asymmetries in knowledge across the interdisciplinary team. On the healthcare provider side of treating patients, doctors, nurses, and pharmacists collaborate to treat a patient. Each role brings its own specialized knowledge. Doctors have diagnosis training, nurses have care details, and pharmacists have in-depth drug information. People from these roles communicate with others in the same role or in different roles. Participants tailor the conversation according to each other's displayed understandings. Conversation analysts call this "recipient design" [1].

On the healthcare payer side, an insurance adjuster processes claims for our patients' treatment. Here a benefits manager and analyst work together. The benefits manager has expertise in claims and transactions and the analyst specializes in data science and analytical techniques. The benefits manager and analyst analyze the claims of a population, plan policies and plan offerings that are beneficial to all parties. Similar to the providers, the payers will adapt their communication to their displayed understandings.

We are currently building and studying two conversational agents that provide information to healthcare professionals, rather to patients. The Micromedex agent [2] is for healthcare providers to get answers to questions about drugs and conditions, such as "what drugs treat a diabetes?" or "what is the adult dosage of Tazarotene for psoriasis?". The Payer agent [3] lets providers ask questions related to the analysis of claims data, for example, "what are the top drugs in the class for diabetes?" Both virtual agents share common medical knowledge about drugs.

The state of our conversational agents is that they support the most common language of the users. Our initial prototypes used more specialized language, however, in our user studies, some participants unfamiliar with the language were confused. To better enable mutual understanding between user and agent, we employ the Natural Conversation Framework (NCF) from IBM Research [4], which provides a rich set of mechanisms for users to get clarifying information. Because the "repair" patterns of the NCF, such as the user paraphrase request, example request and definition request, are not available from most chatbots or voice assistants, we must initially teach users that our agents can do this.

Our goal

While our agents currently respond to all users in the same way, our goal is to enable them to tailor their responses to different kinds of users. In other words, we aim to implement a simplified version of what conversation analysts call "recipient design" [1], or the tailoring of an utterance to its particular recipient. Rather than this level of granularity, we plan to design role-tailored responses, using different vocabulary, to the same inquiries and requests for the various professional roles of our users, including doctors, nurses, pharmacists, benefits managers, and analysts. This means the agent supports multiple domains of knowledge and vocabularies.

As one user demonstrates expertise, the agent adjusts its utterances and actions for that expertise but handles other knowledge areas at the default, non-expert level. Technically this simply requires a method for identifying the user's professional role and the use of a corresponding context variable to select the appropriate, role-specific response.

In the next section, we describe four major research areas involved in accomplishing this goal. We aim to have a repeatable approach. Given the approach new conversational agents can be efficiently created in the healthcare domain. The approach applies to other areas involving more than one domain of knowledge.

Representing different levels of knowledge

We propose that the knowledge is divided into separate ontologies. An ontology in computer science is a formal representation of a domain of knowledge containing entities and their relationships. Corresponding to an ontology is a system that answers questions related to the particular domain knowledge. Input to the system is a series of questions. For each question the system executes a query against the ontology model to retrieve an answer. The ontology and question-answer system can be expanded to support a wider range of questions.

An ontology should capture the language users mention in their utterances. We propose the level of expertise of the language be annotated; for example, "measure" and "dimension" are language used by analysts when discussing business intelligence, but not by claims managers. Capturing the knowledge in modules enables them to be assembled depending on the virtual agent's requirements. In our examples, the drug knowledge ontology can be reused by both agents.

1 A: For what dimension do you want to see the drug costs?

2 U: What is a dimension?

3 A: A dimension is categorical attribute of the data.

4 U: I do not know which dimension.

5 A: Ok, instead, how would you like rank the costs?

6 U: Show me the plans.

7 A: OK, here is the chart…

Figure 1: Agent adjusting level of expertise for the user

Conversational agents displaying different levels of knowledge

A conversational agent interprets a user's utterances and responds. A user with expertise in an area may use expert terminology or layperson vocabulary. At minimum the agent needs to understand both kinds of language to perform competently. Beyond that, the agent can identify the expert language and use it as a heuristic to identify that the user has an expertise. Figure 1 shows how the agents changes to layperson's terminology. Besides individual words, the phrasing of a user's question also indicates expertise. Here are two examples:

Layperson: "What are the top drugs?"
Data Scientist: "Show me the cost by drugs."

While the layperson's utterance uses a simple construction, "What are...," the data scientist's uses a more specialized format, measure by dimension, that is cost and drug product, respectively. Once the agent identifies a user has an expert, it can then adapt its responses to suit how an expert would talk about the topic. Usually the expert's language can express an idea in a more concise way.

Agents guiding users through recommendation

In the previous section we discussed a turn between the user's utterance and the conversational agent's response. More generally, the user and agent have a session of several turns, covering multiple questions. At the end of each question, the agent has the opportunity to suggest follow-up topics. These suggestions can be tailored to the user.

If the user asks general questions, the agent could make suggestions on how to elaborate the information. For example, the user asks, "What are my admissions". The answer could be a single number. The agent can suggest other details to add; for example, maybe the number of admissions could be plotted by the most common conditions patients are admitted to a hospital or the names of the hospitals with the most admissions.

If the user asks a more detailed question, the agent can suggest other follow ups. The suggestions could be tailored to the role of a user. A business manager might be looking for insights, while an analyst is collecting the data to make a recommendation.

1 U: Show me the drug costs by plan for 2017.

2 A: Ok, here is the chart <graphic>

3 A: Do you want to add a filter or drill down?

4 U: Filter for insulin.

5 A: Here is the updated chart <graphic>

Figure 2: Agent adjusting level of expertise for the user

Steering user into known knowledge area

Users starting with an agent need to become familiar with the agent's capabilities. A benefits manager, for example is looking for insights into restructuring a healthcare plan and might ask "what changes should I make to the plan?" Our agent helps the user get data necessary for the user to apply human insight. So that agent needs to direct the user to a supported topic. The agent could continue to engage so that the user describes details that are in scope, for example, in the Eliza system [5]. Then the agent can guide the user through the data.

Challenges beyond the initial ideas

The previous sections outline aspects of a complete solution. Other issues remain. The conversational agent needs to smoothly switch its level of expertise. We believe heuristics will be data used to detect when to switch. The switch needs to happen when the system has enough confidence. It should not switch too early, too late, or seesaw between expert and lay levels.

The user's clarifying questions and other repairs are also an indication that the virtual agent may be talking over his or her head. If the user, asks for example "what's a dimension?" it is an indication that the user is not familiar with business intelligence jargon and instead requires a lay vocabulary.

Conclusion

Virtual conversational agents are on an exciting frontier to make information more accessible to a larger group of users. Our research is enabling agents to support more in-depth conversations about the healthcare domain. The approaches of capturing knowledge and interaction patterns generalize to other domains.

References

[1] Emanuel A. Schegloff, 2007. Sequence Organization in Interaction: A Primer in Conversation Analysis. Cambridge University Press.

[2] Abdul Quamar, Chuan Lei, Dorian Miller, Fatma Özcan, Fatma, Jeff Kreulen, Robert J. Moore, Vasilis Efthymiou: An Ontology-Based Conversation System for Knowledge Bases. SIGMOD 2020.

[3] Dorian Miller, Rebecca Niehus, Robert J. Moore: Improving the user experience for healthcare professionals using a conversational agent to complete business intelligence analysis. AHFE 2020.

[4] Robert J. Moore, Raphael Arar. 2019. Conversational UX Design: A Practitioner's Guide to the Natural Conversation Framework. Association for Computing Machinery.

[5] Joseph Weizenbaum. 1966. ELIZA—a computer program for the study of natural language communication between man and machine. Communications of the ACM.