Intelligent Multi-Purpose Healthcare Bot Facilitating Shared Decision Making

Mohammad Mehdi Afsar
University of Calgary
Calgary, Alberta
afsar@ieee.org

ABSTRACT

Patient decision aids (PtDAs) have been promoted to facilitate personalized information retrieval and decision support; nonetheless, although promoted for more than 20 years, they have generally failed to gain a foothold in the general delivery of healthcare. Intelligent interactive agent technologies could address the design features necessary to facilitate support and shared-decision making. In this thesis, we develop and build a PtDA for Prostate cancer using intelligent agent technology. The proposed system, called ALAN, has a multi-layered architecture with three layers. While the first layer (User-Interface) is responsible to effectively interact with users (patients and physicians), the bottom layer (Data) handles requests regarding storing and retrieving the data. Unlike most existing bots, our core objective is to enable ALAN with learning abilities, which can evolve in the course of time and improve its behaviour with minimum distraction of the user. To this end, reinforcement learning and deep learning algorithms are employed in the main layer, i.e., Analytical Decision Making. This research is expected to have impact on delivery of personalized healthcare.

KEYWORDS

Shared-decision making; Patient decision aids; Multi-agent systems; Chatbots; Reinforcement learning.

ACM Reference Format:
Mohammad Mehdi Afsar. 2019. Intelligent Multi-Purpose Healthcare Bot Facilitating Shared Decision Making. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 3 pages.

1 INTRODUCTION

There are many decisions in healthcare where the scientific evidence supporting the *best* choice is insufficient or equivocal. These decisions are *preference-sensitive*, because the ultimate choice should depend on patients' preferences for the benefits, harms, and scientific uncertainties associated with the medically reasonable alternatives [10]. For preference-sensitive decisions, *shared-decision making* is advocated as a way of engaging and informing patients to make a value-based choice and delivering high quality care [5].

Patient decision aids (PtDAs) have been promoted to facilitate shared-decision making, and there has been considerable scientific study into the use of these instruments for this purpose [7]. PtDAs provide patients with balanced information about their options, while eliciting their preferences and any areas of cognitive dissonance. Nonetheless, despite being introduced over 20 years ago, PtDAs have generally failed to gain a foothold in the delivery of healthcare. Common PtDAs barriers, such as a lack of applicability due to patient characteristics and/or a specific clinical situation [7], are further compounded by healthcare providers who are overworked [6], time-constrained, and/or improperly trained on how to administer and use PtDAs in their clinical practice [3].

Recent advances in intelligent agent technology (i.e., "chatbots") could address the design features necessary to support a "next generation" of PtDAs, and thus facilitate improved shared-decision making using agent-based systems. Agents are software entities that can interact with each other and human counterparts to solve problems; to share expertise; to work in parallel or sequence on common problems; to represent multiple viewpoints and the knowledge of multiple experts; and to compete for limited resources.

In this thesis, we utilize the intelligent agent technology to facilitate and enhance shared-decision making in healthcare. We propose ALAN, an Artificial intelligence- and Learning-based multi-Agent system for Next-generation PtDAs. ALAN is composed of three layers, namely the User-Interface (UI), Analytical Decision-Making (ADM), and Data layers. The UI has an interactive interface to communicate with users in a natural way, and its main task is to receive the request through natural language and show the response in the most logical form. On the other hand, the ADM is behind the scene and responsible for learning, making important decisions, and analysis. In the bottom, all the data is stored and retrieved in the data layer. Unlike existing virtual assistants and chatbots that are mostly deterministic [8], the main feature of ALAN is its ability in learning and making smart decisions. In this regard, the Learning agent, located in the ADM, plays a critical role and is equipped with Reinforcement Learning (RL) algorithms.

2 THE PROPOSED SYSTEM

As mentioned earlier, ALAN has a multi-layered architecture with three layers. We have devised incremental development and prototyping approach for ALAN. The first prototype includes the development of the UI and Data layers. The UI (the top layer) is responsible to effectively interact with users (patients and physicians) and is composed of three agents, i.e., Human Interaction (HI), Personal Assistant (PA), and Survey agents. The main duties of the HI can be summarized as receiving requests from the users, deciding to query which agent based on the request of users, and showing the returned answer from other agents in an appropriate way. The PA, on the other hand, has four main tasks, namely handling referrals, booking appointments, billing services, and external communications. Also, the Survey agent is in charge of gathering data from

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

patients and uploading this data to the system. As a proof of concept, we have already developed a voice-activated interactive prototype system that understands the conversation between a patient and a physician during the patient's visit session and fetches information from a private cloud setup and presents it to them without interfering with the conversation. Note that we have used IBM Watson [1] for speech recognition. Moreover, we have developed a web-based application, called PROCheck, to have an interactive environment with patients. The Survey agent pushes surveys to patients through PROCheck and the filled out surveys are saved to the database again using the Survey agent.

On the other hand, the bottom layer is the Data layer, which is composed of Database agent and Ontology agent. While the Database agent is in charge of handling common database operations, including data retrieval, storing, deletion, and search, the Ontology agent is responsible to modify the ontology tree, which keeps high-level and conceptual information in an hierarchical manner and also information for mapping different data in the database. We already have a mechanism in place to collect patient information using instruments, such as OAB-V8 and the Expanded Prostate Index Composite-26 (EPIC-26) domain scores. We leverage this and other data sources (provided by the Prostate Cancer Centre and Community Health Sciences Department, etc.) for intelligent data aggregation, predictive analytics, and personalized patient care. Not only our database contains patient information, but it also keeps other types of data, including the opinion of patients about their disease collected from online discussion forums, description of medical terms and diseases, training data for our Learning agent, and different documents. We have already developed a mechanism to crawl online discussion forums in order to collect the opinion of patients about their diseases. This data is particularly useful when used by Natural Language Processing (NLP) techniques to figure out relationships between patient pre- and post-treatment conditions and to find patterns for improved healthcare.

The second prototype includes the development of the ADM layer (the middle layer), which is composed of three agents, namely Learning, Decision Making (DM), and Analytics agents. The DM agent is queried by other agents when they need to come up with a decision about an important case. Machine learning and data fusion techniques are used in this agent to make the decision. A good example to illustrate the importance of the Decision-making agent is our previous work, LOWA [2], which was used in classifying breast tumor classification. The Analytics agent is mainly responsible for data analysis. Different machine learning algorithms and prediction models are used to analyze the data. For instance, when a referral letter is received and processed by the PA, the Analytics agent is used to predict the disease and risk levels based on disease parameters extracted from the referral letter received. Fig. 1 depicts the multi-layered architecture of ALAN.

Our major objective and novel research part of this study is to enable ALAN to preserve the conversation in the most natural way possible and to improve its behaviour (answers or services) during the time and with minimum distraction of the user. To this end, we will employ RL and deep learning algorithms for the Learning agent. To explain the importance and procedure of the Learning agent, consider the following scenario. ALAN is queried by a physician to provide the possible treatment for Prostate cancer

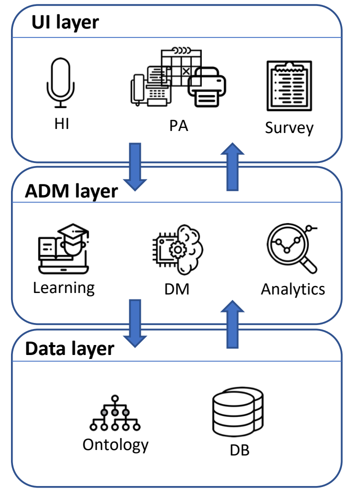

Figure 1: The multi-layered architecture of ALAN

patients. The system explores its repository, pulls the treatment types and shows the list. The very first option for this scenario is to use NLP techniques, such as *lexical matching* and *semantic search*. These techniques are helpful for a very small repository where the similarity between documents is very low; however, when we started to grow our repository and tested different queries, the accuracy of these techniques considerably diminished. More importantly, our target is to have an interactive search based on the opinion of the physician where there is no need to a rich dataset to train our model. This is mainly because the number of training dataset for this purpose is scarce [4]. Accordingly, we utilize RL algorithms.

We model our information retrieval problem as a Markov Decision Process (MDP) where it learns to optimize the search using local repository and the Internet search, if required. To represent the MDP, we use a tuple $< S, A, R, T >$, where S is the set of all states, A is the set of all actions, R is the reward function, and T is the transition function. The state is the confidence of our algorithm in finding the similar documents. The action is the decision of what to do next according to the current state. The agent can choose to do a new search or to return a set of documents it thinks they are relevant. The main strategy in selecting rewards is to maximize the accuracy of search while minimizing the number of feedback the agent needs from the user. To solve this problem, we plan to try different RL algorithms, including Q-Learning and PPO (proximal policy optimization) [9].

3 CONCLUSION

We propose ALAN, a multi-purpose healthcare bot for a next-generation PtDA. ALAN utilizes both multi-layer architecture and multi-agent system technologies to provide a broad range of medical services. The main feature of ALAN is its ability in learning and improving its behaviour and services in the course of time. In the future, our focus is mainly on the Learning agent and employing RL algorithms.

REFERENCES

[1] [n. d.]. IBM Watson. https://www.ibm.com/watson/.. ([n. d.]). Published 2018. Accessed December, 2018.

[2] Emad A Mohammed, Christopher T Naugler, and Behrouz H Far. 2016. Breast tumor classification using a new OWA operator. *Expert Systems with Applications* 61 (2016), 302–313.

[3] Benjamin Moulton and Jaime S King. 2010. Aligning ethics with medical decision-making: the quest for informed patient choice. *The Journal of Law, Medicine & Ethics* 38, 1 (2010), 85–97.

[4] Karthik Narasimhan, Adam Yala, and Regina Barzilay. 2016. Improving information extraction by acquiring external evidence with reinforcement learning. *arXiv preprint arXiv:1603.07954* (2016).

[5] Annette M O Connor, Hilary A Llewellyn-Thomas, Ann Barry Flood, et al. 2004. Modifying unwarranted variations in health care: shared decision making using patient decision aids. *Health Affairs* 23 (2004), 63–72.

[6] Annette M O Connor, Hilary A Llewellyn-Thomas, Ann Barry Flood, et al. 2004. Modifying unwarranted variations in health care: shared decision making using patient decision aids. *HEALTH AFFAIRS-MILLWOOD VA THEN BETHESDA MA-* 23 (2004), VAR–63.

[7] Annette M O'Connor, Carol L Bennett, Dawn Stacey, Michael Barry, Nananda F Col, Karen B Eden, Vikki A Entwistle, Valerie Fiset, Margaret Holmes-Rovner, Sara Khangura, et al. 2009. Decision aids for people facing health treatment or screening decisions. *Cochrane database of systematic reviews* 3 (2009).

[8] Elahe Paikari and André van der Hoek. 2018. A framework for understanding chatbots and their future. In *Int. Workshop on Cooperative and Human Aspects of Software Engineering*. 13–16.

[9] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).

[10] John E Wennberg, Elliott S Fisher, and Jonathan S Skinner. 2003. Geography and the debate over medicare reform. *Health Affairs* (2003), 96–114.