A Reinforcement Learning Approach to Improve User Achievement of Health-Related Goals

Diogo Martinho[1(✉)], João Carneiro[1], José Neves[2], Paulo Novais[2], Juan Corchado[3], and Goreti Marreiros[1]

[1] Research Group on Intelligent Engineering and Computing for Advanced Innovation and Development (GECAD), Institute of Engineering, Polytechnic of Porto, Porto, Portugal
{diepm,jrc,mgt}@isep.ipp.pt
[2] ALGORITMI Centre, University of Minho, Guimarães, Portugal
jneves@di.uminho.pt
[3] BISITE Digital Innovation Hub, University of Salamanca. Edificio Multiusos, Salamanca, Spain
corchado@usal.es

Abstract. The demand and interest for personalized, efficient, and inexpensive healthcare solutions has significantly increased over the last decade to overcome the major limitations of existing traditional healthcare approaches. This new trend relies on the definition of intelligent mechanisms that can persuade the end-user to achieve health-related outcomes and ultimately improve his health condition and well-being. In this sense, the work here proposed explores a Multi-Agent System composed by personal agents that follow user preferences and a coaching agent which relies on a reinforcement learning approach to identify the most impactful messages to persuade a certain agent to follow established health-related goals. To validate the proposed system, a set of simulations were performed considering different types of persuasive messages and we were able to identify the most adequate sequence of messages that can persuade different users to achieve health-related goals based on their preferences.

Keywords: Reinforcement learning · Healthcare systems · Multi-agent systems

1 Introduction

We are currently witnessing a shift in the healthcare paradigm with all the noticeable advancements in technology. In a world where the costs of living are increasing at a daunting rate derived from increasing taxes, lower and lower incomes, pandemic effects, etc. [1–3], and at the same time we observe a significant burden of the healthcare providers [4], the reality is that the quality of traditional healthcare methods fall short of required to correctly support patients and their needs during their daily lives [5]. We think less and less about disease-centered solutions and the focus has diverged towards a patient-centered reality where the patient and the underlying needs and preferences become key factors to treat and manage his/her own health condition effectively [6]. In this sense, we observe strong efforts in the literature to develop personalized healthcare solutions that

G. Marreiros et al. (Eds.): EPIA 2021, LNAI 12981, pp. 266–277, 2021.
https://doi.org/10.1007/978-3-030-86230-5_21

can support patients during their daily lives and in return improve their health condition and well-being. Such solutions rely on the intelligent use of both novel technological features and successful engagement strategies that can persuade the patient to be an active player and in return guide him/her to follow healthier behaviors [7–9].

We observe several works in this area already combining engagement strategies such as Persuasion and Behavior Change Theories [10–12] with Artificial Intelligence methods such as Natural Language Processing [13, 14], Computer Vision [15], Fuzzy Logic [16, 17], Reinforcement Learning (RL) [18], etc. Among these works, we highlight the use of Reinforcement Learning as one of the most effective approaches to draw healthier outcomes. In the work of [18], the authors identified several domains of application of RL in the area of Healthcare. They referred to works done in dynamic treatment regimens which includes treatment and management of chronic diseases such as diabetes, depression, and cancer, or diseases that require critical care such as anesthesia or sepsis. They also referred to the use of RL to process both structured and unstructured data which can range from medical image processing to free text analysis. Finally, the authors also pointed towards other interesting and general domains such as using RL for healthcare resource scheduling, use of RL in surgery procedures, and use of RL in the healthcare management with the inclusion of healthcare plans to manage physical activities. In the work of [19] the authors also discussed the many domains of application of RL in the area of healthcare and suggested the use of RL in the development of dialogue systems which includes the development of multi-agent systems with conversational agents that can rely on RL mechanisms to identify the best way to interact with the patient.

In this work we follow the conceptual ideas first discussed in [20] and propose a Multi Agent System which considers both agents that support the patient and follow his/her preferences and needs and a coaching agent which persuades other agents to accomplish health-related goals by using a RL strategy to identify the best sequence of messages to exchange with those agents (based on their preferences and needs). To validate our proposed model, we selected different simulation scenarios which include different types of messages that can be exchanged with the user (by the personal agent) and we observed how easily the coaching agent was able to identify the best sequence of action compared to a "normal" coaching agent that did not make use of this approach and attempted to persuade other agents with random sequences of messages. We were able to observe significant advantages over the use of RL for this purpose and we also identified other benefits of our proposal more related with the generic structure that was applied in the RL process.

2 Proposed Model

The main goal of the proposed work is to identify the most impactful sequence of interactions that should happen between the healthcare system and a certain patient based on his preferences and needs. As such we introduce a Multi-Agent System composed by two main entities which will interact with each other and understand how to interact with a patient. These entities are referred as Personal Agent and Coaching Agent (Fig. 1).

It is important to note that this interaction happens always between one PA, that represents a patient, and the CA. For several patients, a new instantiation for both PA and CA specific for that patient would be necessary.

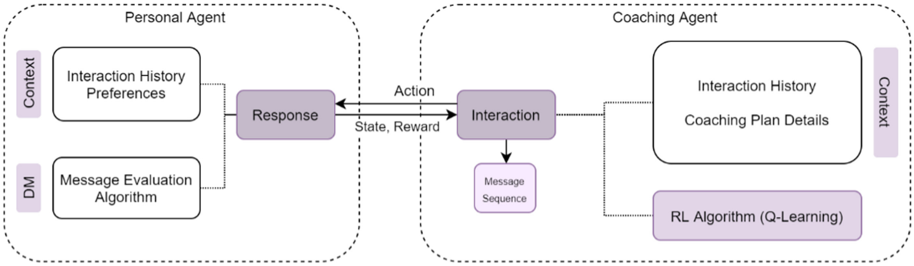

Fig. 1. Multi-agent system architecture

2.1 Personal Agent

The first entity of the proposed system is the Personal Agent (PA), and as the name refers this agent will interact with the person (patient) and exchange messages throughout the day to motivate him/her to follow health-related goals. To do so, the PA will follow patient preferences regarding context variables (for example, preferences regarding which types of messages are most suited at a certain time of the day, which days of the week the person prefers to do physical activities, which days of the week the person prefers to rest, previously accepted or rejected goals based on the messages that were sent to that person, etc.). Besides that, the PA will also have its own internal Message Evaluation algorithm to decide whether the person is more or less likely to accept or reject a certain message to follow a health-related goal (also combining the information from the context). This evaluation algorithm can be configured to be as simple as to follow a set of preferences to a more complex and intelligent process that can combine these preferences and other information obtained from the context and past interactions with the user to better judge whether a certain interaction is more likely to succeed or not. In the case of this study and as explained in more detail in the Discussion section, the first case was considered.

2.2 Coaching Agent

The second entity of the proposed system is the Coaching Agent (CA) and has the main task of selecting the most appropriate messages to send to the patient. To do so, the CA will process context information which also includes details from all previous interactions with that patient (information related to goals that were achieved, how many messages were exchanged before the patient achieved a new goal, how many goals were achieved during the current week/month, etc.) and details related with the current coaching plan given to the user. The structure of a coaching plan has been previously introduced in [21] and the main idea is to establish health-related goals that can be improved depending on the patient progress. This information is then combined with the execution of the RL algorithm to identify the best sequence of messages to exchange with a certain patient. This process is iterative and depending on the response obtained from the PA side, the RL algorithm will be improved for the next iteration until a correct sequence of action is identified for a certain patient. Below we introduce the general definition of RL as well as the model-free implementation of RL with the Q-Learning algorithm that is used by the CA.

Reinforcement Learning. Reinforcement learning is a machine learning process in which an *agent* has different goals related to a specific *environment* and must decide which *action* is the most adequate to achieve a goal according to a given *state* of the environment. This goal can be represented by a value (*reward*) and the agent should attempt to maximize this value when choosing which actions to take (*policy*) given a state signal. Over the long run the agent will identify the optimal policy which corresponds to the policy that provides the highest reward value.

In this work, each action is represented as a set of n messages, or in other words, as $M = \{m_1, m_2, \ldots, m_n\}$, that can be sent to a certain patient before he achieves (or not) an established health-related goal. An interaction history will be stored with the result of the interaction after sending M at certain time (or stage) and the corresponding reward value (which in this case indicates whether the patient achieved or not and established health-related goal). This flow corresponds to a Markov Decision Process, in which the CA and the environment (in our case the environment equivalent to the PA) interact with each other during a set of Y stages. At a stage $y \in Y$, the CA is presented with a state $s_y \in S$ from a set of states and must select an action $a_y \in A(s_y)$ from a set of possible actions available for that state.After selecting that action, the CA receives a reward r_{y+1} resulting from the action taken and is then presented with a new state s_{y+1}. With this result, the CA may change (or not) the preferred policy for the next interaction, and over the long run the RL problem will be "solved" with the CA finding the policy that achieves the highest reward value. The expected return (the sum of rewards) for the CA starting in state s, taking action a, and then following policy ρ thereafter is called as the action-value function for policy ρ. In the context of Q-learning, this function is also called as Q-function and is denoted as:

$$Q_y^{\rho}(s, a) \doteq E_{\rho}\left[\sum_{j=y}^{Y} r_{j+1} \middle| s_y = s, a_y = a\right]$$

It should be noted that we opted to use Q-Learning method in the proposed work, in comparison to other RL methods such as SARSA because of the off-policy characteristics of the Q-Learning method. As a result, it is possible to evaluate and improve policies that are different from the selected policy in each stage and therefore it is possible to maintain a continuous exploration process in which the CA can select different policies while learning what is the optimal policy in a certain environment.

3 Results and Discussion

To validate the proposed model, we defined 3 different simulation scenarios and compared the results according to 3 different levels of complexity. In the first scenario we tested our model by considering 2 different types of persuasive messages; in the second scenario 4 different persuasive messages were considered; and in the third scenario 6 different persuasives messages were considered. For simulation purposes we do not define each persuasive message but instead refer as Message 1, Message 2, Message 3, etc., but these message could be easily adapted to refer to some of the most known and used persuasive models that were mentioned in the introduction of this study, such as the six principles of persuasion by Cialdini [22, 23] or the Transtheoretical Model of Behavior Change [24], among others [10].

Each scenario was tested with a time length of 30 days and we studied both the accuracy of each set of messages and corresponding interactions as well as the total number of health-related goals that were achieved after accepting (or not) the messages sent to the PA. Furthermore, 10 simulations were performed for each scenario and the results were compared in terms of average values obtained.

To validate our model, we compared the approach here proposed in which the CA uses the RL algorithm to decide what messages should be sent to the PA with an approach where the CA does not use the RL algorithm and instead attempts to send a random set of messages to the PA.

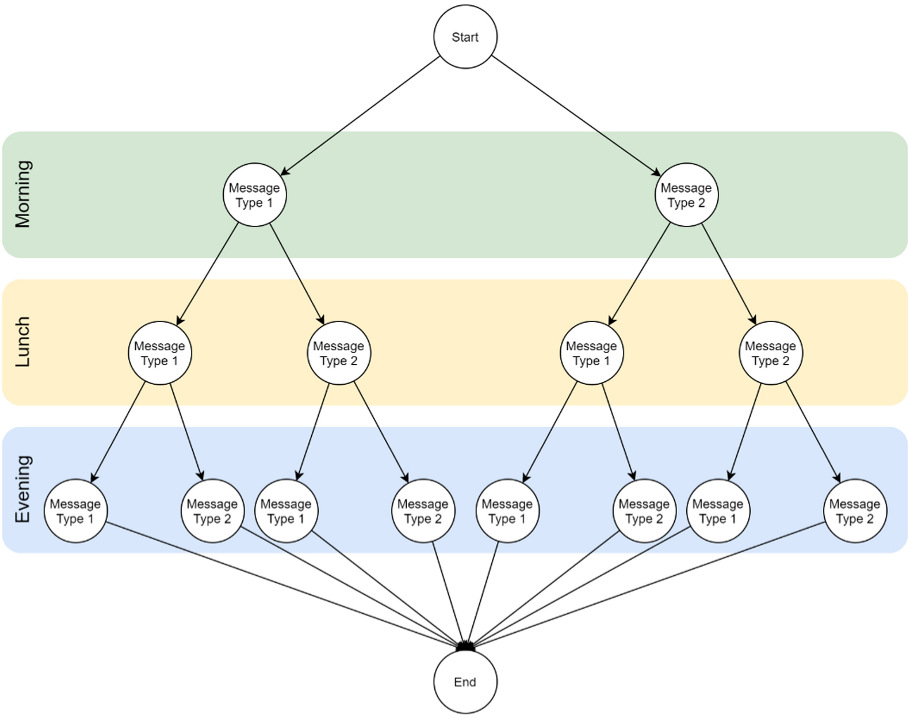

Fig. 2. Message flow example (2 message types)

For all the scenarios simulated, we considered that the maximum number of interactions between the CA and PA would be of three messages per day, and that a goal would be achieved only after the PA accepted all three messages exchanged. We chose this number specifically with given evidence on the literature [25] that recommends the exchange of messages at certain periods of the day, such as morning, lunch, and early evening, and these interactions would represent those moments of the day. Furthermore, the CA was also allowed to exchange repeated types of messages during a day. To exemplify this, in Fig. 2 it is shown the total number of possible combinations of messages that can be sent during a day when two types of messages are considered.

For this example, the total number of combinations corresponds to 8 possible outcomes, and for the remaining scenarios corresponds to a total of n^3 possible outcomes, with n being the number of different types of persuasive messages considered.

The evaluation algorithm defined for the PA was simplified to only select a random sequence of types of messages as the most impactful, which in turn would persuade the patient to achieve a health-related goal. This simplification of the evaluation process was considered since the main goal of this work is to test the performance of the CA using the RL algorithm and, as such, the way the PA could be modeled to evaluate each received message in a more or less intelligent way falls out of the scope of this study.

The average results obtained regarding the number of accepted goals and messages exchanged are presented in Table 1.

After performing all simulations in each scenario, it is possible to observe that the use of RL compared to a Random approach obtained far better results regarding both the total number of messages exchanged and the total number of goals completed after 30 days.

Table 1. Simulation results

Day	Scenario 1 – 2 Messages				Scenario 2 – 4 Messages				Scenario 3 – 6 Messages			
	Goals – RL	Goals – RA	Msgs – RL	Msgs – RA	Goals – RL	Goals – RA	Msgs – RL	Msgs – RA	Goals – RL	Goals – RA	Msgs – RL	Msgs – RA
1	0,1	0,2	1,5	1,7	0	0	1,2	1,1	0	0	1.2	1.2
2	0,6	0,3	4,2	3,5	0	0	2,9	2,5	0	0	2.5	2.5
3	1,3	0,3	7,2	4,9	0,3	0	5,1	4	0.1	0	4.4	3.7
4	2,3	0,4	10,2	6,4	0,9	0	7,8	5,5	0.2	0	7	4.8
5	3,3	0,5	13,2	8,1	1,6	0	10,7	6,8	0.6	0	9.8	6
6	4,3	0,6	16,2	10	2,4	0	13,6	8,1	1	0	12.7	7.2
7	5,3	0,7	19,2	11,7	3,2	0	16,6	9,7	1.6	0	15.6	8.2
8	6,3	0,8	22,2	13,2	4,1	0	19,6	10,7	2.2	0	18.6	9.2
9	7,3	1	25,2	15	5	0	22,6	12,2	3.1	0	21.6	10.2
10	8,3	1	28,2	16,2	6	0	25,6	13,6	4	0	24.6	11.6
11	9,3	1	31,2	17,6	7	0	28,6	14,8	4.9	0	27.6	12.9
12	10,3	1,2	34,2	19,6	8	0	31,6	15,8	5.8	0	30.6	14
13	11,3	1,5	37,2	21,4	9	0	34,6	17	6.8	0	33.6	15.6
14	12,3	1,6	40,2	22,8	10	0	37,6	18,3	7.8	0	36.6	16.6
15	13,3	1,6	43,2	24,4	11	0	40,6	19,8	8.8	0	39.6	17.7
16	14,3	1,6	46,2	25,6	12	0	43,6	21,1	9.8	0	42.6	18.7

(*continued*)

Table 1. (*continued*)

Day	Scenario 1 – 2 Messages				Scenario 2 – 4 Messages				Scenario 3 – 6 Messages			
	Goals – RL	Goals – RA	Msgs – RL	Msgs – RA	Goals – RL	Goals – RA	Msgs – RL	Msgs – RA	Goals – RL	Goals – RA	Msgs – RL	Msgs – RA
17	15,3	1,7	49,2	27,3	13	0	46,6	22,4	10.8	0	45.6	19.9
18	16,3	1,9	52,2	29,1	14	0	49,6	23,6	11.8	0	48.6	21
19	17,3	2	55,2	30,4	15	0	52,6	25,2	12.8	0	51.6	22.2
20	18,3	2	58,2	32,1	16	0	55,6	26,3	13.8	0	54.6	23.2
21	19,3	2,3	61,2	34,4	17	0	58,6	27,3	14.8	0	57.6	24.3
22	20,3	2,4	64,2	36,5	18	0,1	61,6	28,5	15.8	0	60.6	25.5
23	21,3	2,6	67,2	38,2	19	0,1	64,6	29,5	16.8	0.1	63.6	26.7
24	22,3	2,7	70,2	39,8	20	0,2	67,6	30,8	17.8	0.1	66.6	28
25	23,3	3	73,2	41,5	21	0,2	70,6	32	18.8	0.1	69.6	29.2
26	24,3	3,3	76,2	43,5	22	0,2	73,6	33,1	19.8	0.1	72.6	30.7
27	25,3	3,5	79,2	45,4	23	0,2	76,6	34,6	20.8	0.1	75.6	31.9
28	26,3	3,7	82,2	47,4	24	0,2	79,6	35,8	21.8	0.2	78.6	33.4
29	27,3	3,8	85,2	48,9	25	0,2	82,6	37,2	22.8	0.2	81.6	34.7
30	28,3	3,9	88,2	50,6	26	0,2	85,6	38,4	23.8	0.2	84.6	35.7

Looking at the total number of achieved goals, in the first scenario, both agents achieved the highest average values compared to the remaining scenarios. However, while the Random Agent was only able to obtain an average of 4 goals accepted at the end of 30 days, the RL Agent was able to obtain an average exceeding 28 goals accepted. In the second and third scenario, the Random Agent obtained even worse results and was only able to persuade the PA to achieve a health-related goal in two simulations for each scenario. On the other hand, the RL Agent still obtained very satisfactory results and only had a slight decrease in the total number of achieved goals to an average of 26 goals accepted in the second scenario and nearly 24 goals in the third scenario. Regarding the number of messages exchanged to persuade the PA to achieve a goal during each day, both the RL Agent and Random Agent also obtained the best results in the first scenario, with an average of nearly 88 messages exchanged between the PA and the RL Agent and 51 Messages between the Random Agent and the PA. In the second and third scenarios, a significant decrease was observed for the Random Agent (only nearly 39 messages exchanged after 30 days in the second scenario and only nearly 36 messages exchanged in the third scenario). A less significant decrease was observed in the case of the RL Agent (nearly 86 messages exchanged after 30 days in the second scenario and 85 messages exchanged in the third scenario). These results show a clear superior performance between the RL Agent and the Random Agent, however it still was not clear if the accuracy of all the interactions performed between the CA and the PA were also superior in the case of the interactions between the RL Agent and the PA and the Random Agent and the PA. Therefore, the system was also evaluated regarding the number of messages that were exchanged per day between each of these two entities and the number of messages that were actually successful to persuade the PA to achieve a health-related goal. These results are shown in Fig. 3, Fig. 4 and Fig. 5 respectively.

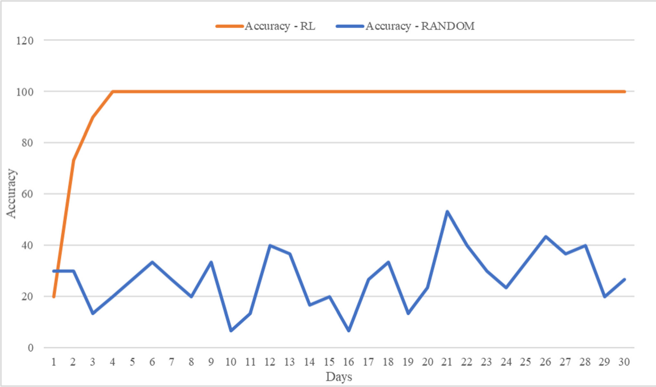

Fig. 3. Average accuracy results (1st scenario)

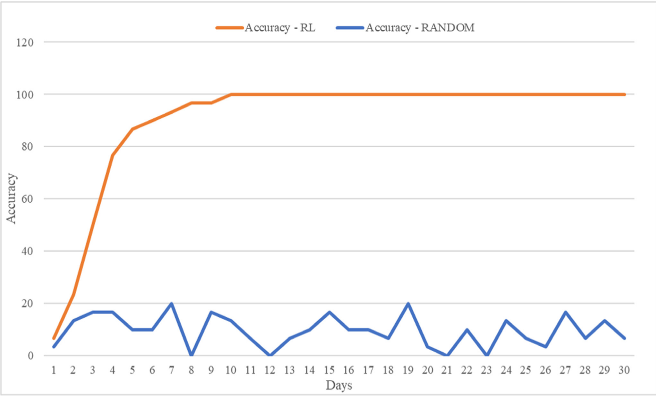

Fig. 4. Average accuracy results (2nd scenario)

A linear increase on the accuracy value was observed from the interactions between the RL Agent and PA in each of the three scenarios with a peak of 100% accuracy achieved by the 4th day in the first scenario, 10th day in the second scenario and 13th day in the third scenario. This accuracy value is supported by the fact of how the Message Evaluation Algorithm was configured for the PA in this study. Since we are only considering a set of Messages the RL Agent learns which set and sequence of messages is most persuading and then repeats the same set for the remaining days in each scenario. On the other hand, the highest accuracy value obtained from the interactions between the Random Agent and PA was of slightly over 55% in the first scenario. In the remaining scenarios, the peak observed had a value lower than 20%. These results also reveal a clear performance with the use of the RL method. In fact, despite the scenario complexity it was still possible

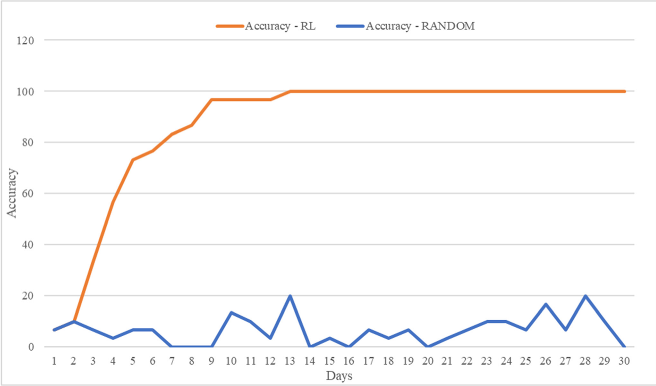

Fig. 5. Average accuracy results (3rd scenario)

to observe that the RL Agent was able to obtain very similar, yet high results in terms of accuracy, number of accepted goals and number of messages exchanged. These numbers are even more significant if we compare them with a less intelligent approach such as the use of a random strategy to persuade the user to follow health-related goals.

Although these results are very positive, they are still lacking and can be further improved. First, the proposed model and the performed study was made on the sole basis that we were trying to identify what was the best approach to interact with a certain user during a day and that way be able to persuade him to achieve a certain health-related goal. This approach does not tell us however, what the best approach is to interact with a certain user, over the time, on each day. In other words, we do not know whether a certain set of messages work best on a day and a completely different set of messages work best on the next day or next week because the PA was always modelled to keep the same initial preferences that were configured. Therefore, additional improvements could and should be considered in the continuation of this work to have a model that can learn with the interactions between the user not only during a day but also take into account the context to make a better judgement of the ideal interaction at any given moment in time. Second, regarding the RL method considered (in this case and as we explained, it was opted to use the Q-Learning method), it could and should also be improved to support scenarios of even higher complexity. Even though we were able to obtain very fast results in each simulation, we are aware that the algorithm performance will decrease significantly as the complexity of the scenario increases and in a scenario with dozens of different types of messages consider, this would force to incorporate a different approach to measure the ideal set of messages to exchange with the user. Such approach could consider the use of deep Q-Learning with neural networks to approximate the Q-Function that was presented. Third, the proposed model should be further evaluated considering users with different preferences and ways of living (or in other words, improve the current Evaluation Algorithm). Our goal is to have a system that is prepared to deal with any type of user, whether they like to perform physical exercise during the day or work out on weekends or after eating a more indulging meal, etc. By doing so, the system will be prepared to interact with any user regardless of age or gender and correctly support the user during his/her daily life. Finally, although the proposed Multi-Agent approach currently considers only two types of agents, we intend to incorporate a third agent (Checker Agent) that will monitor the user progress in real time and acquire health inputs from devices such as Smart Bands. By doing so, we will be able to evaluate the accuracy of each interaction even more precisely and understand how easily (or not) a health-related goal was achieved after exchanging a set of messages.

4 Conclusions and Future Work

The rising costs of living and the overall demand for answers to relieve the healthcare burden has led towards a reality in which traditional healthcare no longer is adequate and instead new solutions are being developed to support people throughout their daily lives from a patient-centered perspective in which the person himself/herself becomes the key factor to manage and improve his/her health condition. Personalizing and enhancing the support provided based on the preferences, needs and the way people behave is essential to accomplish this goal, and as such, emerging healthcare systems must be able to understand what the best approach is when interacting with the patient to lead him/her towards improving his/her health condition. The work here presented explores a Reinforcement Learning approach to identify the best way to interact with a patient based on his/her preferences and needs, experience from previous interactions under a perspective of motivating the patient to accomplish health-related goals. To do we incorporated this approach in a Multi-Agent System in which we consider agents that represent patients and their preferences and an agent that attempts to persuade those agents using a combination of messages that best fits the preferences of those patients. To validate our model, we can a set of simulations with different scenarios and levels of complexity and obtained better results compared to a less intelligent approach. As future work we intend to improve the proposed model according to different points that were discussed in this study and which are mostly related with the ability to consider higher complexity scenarios and patients with more complex behaviors and ways of living. Additionally, the inclusion of additional agents that can monitor the health condition of the patient will be ideal to measure the accuracy and the impact of each interaction performed between the proposed system and the patient. With these points in mind, it will be possible to present an intelligent approach that correctly support the patient despite of his/her age, gender, or any demographical aspects but instead operate based on how that person behaves throughout his/her daily life and at the same time lead him/her to accomplish health-related goals result in healthier lifestyles.

Acknowledgments. The work presented in this paper has been developed under the EUREKA - ITEA3 Project PHE (PHE-16040), and by National Funds through FCT (Fundação para a Ciência e a Tecnologia) under the under the project UIDB/00760/2020 and by NORTE-01-0247-FEDER-033275 (AIRDOC - "Aplicação móvel Inteligente para suporte individualizado e monitorização da função e sons Respiratórios de Doentes Obstrutivos Crónicos") by NORTE 2020 (Programa Operacional Regional do Norte).

References

1. Li, Y., Mutchler, J.E.: Older adults and the economic impact of the COVID-19 pandemic. J. Aging Soc. Policy **32**, 477–487 (2020)
2. Clarke, L.: An introduction to economic studies, health emergencies, and COVID-19. J. Evid. Based Med. **13**, 161–167 (2020)
3. Laborde, D., Martin, W., Vos, R.: Impacts of COVID-19 on global poverty, food security, and diets: Insights from global model scenario analysis. Agri. Econ. (2021)
4. Papoutsi, E., Giannakoulis, V.G., Ntella, V., Pappa, S., Katsaounou, P.: Global burden of COVID-19 pandemic on healthcare workers. Eur. Respir. Soc. (2020)
5. Miller, I.F., Becker, A.D., Grenfell, B.T., Metcalf, C.J.E.: Disease and healthcare burden of COVID-19 in the United States. Nat. Med. **26**, 1212–1217 (2020)
6. Chawla, N.V., Davis, D.A.: Bringing big data to personalized healthcare: a patient-centered framework. J. Gen. Intern. Med. **28**, 660–665 (2013)
7. Costa, A., Heras, S., Palanca, J., Jordán, J., Novais, P., Julián, V.: Argumentation schemes for events suggestion in an e-Health platform. In: International Conference on Persuasive Technology, pp. 17–30. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-55134-0_2
8. Costa, A., Heras, S., Palanca, J., Novais, P., Julián, V.: A persuasive cognitive assistant system. In: International Symposium on Ambient Intelligence, pp. 151–160. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-40114-0_17
9. Costa, A., Heras, S., Palanca, J., Novais, P., Julián, V.: Persuasion and recommendation system applied to a cognitive assistant. ADCAIJ: Adv. Distrib. Comput. Artif. Intell. J. Salamanca **5**, (2016). (ISSN: 2255–2863)
10. Alkiş, N., Findik-Coşkunçay, D.: Use of persuasion strategies in mobile health applications. In: Sezgin, E., Yildirim, S., Yildirim, S.Ö., Sumuer, E. (eds.) Current and Emerging mHealth Technologies, pp. 9–21. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-73135-3_2
11. Qasim, M.M., Ahmad, M., Omar, M.: Persuasive strategies in mobile healthcare: a systematic literature review. Literatures **87**, 7 (2017)
12. de Vries, R.A.J., Truong, K.P., Zaga, C., Li, J., Evers, V.: A word of advice: how to tailor motivational text messages based on behavior change theory to personality and gender. Pers. Ubiquit. Comput. **21**(4), 675–687 (2017). https://doi.org/10.1007/s00779-017-1025-1
13. Donadello, I., Dragoni, M., Eccher, C.: Explaining reasoning algorithms with persuasiveness: a case study for a behavioural change system. In: Proceedings of the 35th Annual ACM Symposium on Applied Computing, pp. 646–653 (2020)
14. Anselma, L., Mazzei, A.: Building a persuasive virtual dietitian. In: Informatics, p. 27. Multidisciplinary Digital Publishing Institute (2020)
15. Leo, M., Carcagnì, P., Mazzeo, P.L., Spagnolo, P., Cazzato, D., Distante, C.: Analysis of facial information for healthcare applications: a survey on computer vision-based approaches. Information **11**, 128 (2020)
16. Sengan, S., Kamalam, G., Vellingiri, J., Gopal, J., Velayutham, P., Subramaniyaswamy, V.: Medical information retrieval systems for e-Health care records using fuzzy based machine learning model. Microprocess. Microsyst. 103344 (2020)
17. Mardani, A., et al.: Application of decision making and fuzzy sets theory to evaluate the healthcare and medical problems: a review of three decades of research with recent developments. Expert Syst. Appl. **137**, 202–231 (2019)
18. Yu, C., Liu, J., Nemati, S.: Reinforcement learning in healthcare: a survey. arXiv preprint arXiv:1908.08796 (2019)

19. Coronato, A., Naeem, M., De Pietro, G., Paragliola, G.: Reinforcement learning for intelligent healthcare applications: a survey. Artif. Intell. Med. **109**, 101964 (2020)
20. Martinho, D., Carneiro, J., Novais, P., Neves, J., Corchado, J., Marreiros, G.: A conceptual approach to enhance the well-being of elderly people. In: Moura Oliveira, P., Novais, P., Reis, L.P. (eds.) EPIA 2019. LNCS (LNAI), vol. 11805, pp. 50–61. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-30244-3_5
21. Martinho, D., Vieira, A., Carneiro, J., Martins, C., Almeida, A., Marreiros, G.: A definition of a coaching plan to guide patients with chronic obstructive respiratory diseases. In: Rocha, Á., Adeli, H., Reis, L.P., Costanzo, S., Orovic, I., Moreira, F. (eds.) WorldCIST 2020. AISC, vol. 1161, pp. 54–64. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45697-9_6
22. Cialdini, R.B.: The Psychology of Persuasion. New York (1993)
23. Cialdini, R.B., Cialdini, R.B.: Influence: The Psychology of Persuasion. Collins, New York (2007)
24. Prochaska, J.O., DiClemente, C.C.: Stages and processes of self-change of smoking: toward an integrative model of change. J. Consult. Clin. Psychol. **51**, 390 (1983)
25. Bidargaddi, N., et al.: To prompt or not to prompt? A microrandomized trial of time-varying push notifications to increase proximal engagement with a mobile health app. JMIR mHealth uHealth **6**, e10123 (2018)