



A Holonic Multi-agent Based Diagnostic Decision Support System for Computer-Aided History and Physical Examination

Zohreh Akbari¹(✉) and Rainer Unland^{1,2}

¹ Institute for Computer Science and Business Information Systems (ICB),
University of Duisburg-Essen, Essen, Germany

{zohreh.akbari, rainer.unland}@icb.uni-due.de

² Department of Information Systems, Poznan University of Economics,
Poznan, Poland

Abstract. Available Medical Diagnosis systems (MDSs), including the state of the art, mainly focus on finding the perfect link between the patient's medical history and their health knowledge. However, no matter how powerful they are in performing this action, it is always possible that the final strong deduction is based on some incomplete input. Prior to this process, a physician should literally perform the Differential Diagnosis (DDx), in which (S)he carefully listens to the symptoms explained by the patient, considers some potential diagnoses and then tries to gather enough evidence and supporting information to shrink the probability of the other candidates. In a patient encounter, this method is used in a process called the History and Physical examination (H&P). Only physicians and in some institutions, in order to compensate the shortage of the physicians, some specially trained nurses are qualified to perform this process. A system capable of guiding a focus H&P, however, will allow less experienced nurses to perform this process, and furthermore, can provide second opinions in critical cases. The DDx domain is in fact a holonic domain; hence, a MDS with holonic architecture could be able to perform this process. As the Holonic Medical Diagnosis System (HMDS), tends to cover the stages in the H&P, this system could also be added to available MDSs in order to provide them with the essential comprehensive input and allow their integration in the clinical workflow. This paper will demonstrate the performance of this system and concentrates on its learning process.

Keywords: Holonic Multi-Agent System (HMAS) · Machine Learning (ML) Medical Diagnosis System (MDS) · Reward engineering

1 Introduction

A Diagnostic Decision Support Systems (DDSS) or Medical Diagnosis System (MDS) is a specific type of Clinical Decision Support Systems (CDSSs) that is developed to provide an ordered list of potential diagnoses for given signs and symptoms. The physician then takes the suggested diagnoses and the supportive information into

consideration, and if necessary orders further tests to narrow down this list [1]. In fact, these systems can save the physician's time, and remind ignored critical possibilities.

Available systems, including the state of the art, mainly focus on finding the perfect link between the given input and their health knowledge. However, prior to this process there should be a precise method to guide the user in providing the right, all-encompassing input. This is similar to what a physician does when listening to a patient. (S)he would carefully listen to the symptoms explained by the patient, considers some potential diagnoses and then tries to gather enough evidence and supporting information to shrink the probability of the other candidates. This method is called the Differential Diagnosis (DDx). In medicine, a DDx is the distinguishing of a particular disease or condition from others presenting similar symptoms [2], which helps to reduce diagnosis errors [3]. In a patient encounter, this method is used in a process called the History and Physical examination (H&P). The H&P is a critical component of a patient encounter in which information relevant to a present complaint is obtained [4]. The H&P report includes the Chief Complaint (CC), the Medical History (Hx), the Physical Exam (PE) and the Assessment and Plan (DDx & P). As mentioned, the state of the art of the MDSs merely covers the last part and takes the result of the other parts as input. This is exactly one of the limitations of these systems, and the reason why they cannot be well integrated into the clinical workflow of the hospitals.

The state of the art of the MDSs includes the IBM Watson Health [5] and the Isabel [6]. The IBM Watson Health is claimed to be the most powerful Artificial Intelligence (AI) based system capable of performing medical diagnosis. This system is literally a cognitive computer system that has the capability of understanding the natural language. Although Watson is very powerful, its success is very much depending on the quality of the input, i.e. the output of already performed H&P, and it is actually best used in order to create patient-specific treatment plans. Literally Watson has never been involved in the medical diagnostic process, but only in improving the diagnosis and assisting with identifying treatment options for patients who have already been diagnosed [7]. Isabel on the other hand, is a knowledge-based CDSS that facilitates diagnostic reminders and DDx [6]. A systematic review of DDx generators was conducted in [8], according to which Isabel was associated with the highest rates of diagnosis retrieval. However, as stated in [9] Isabel is still too slow and its accuracy drops significantly if only limited information is available [10].

As mentioned, no matter how powerful the MDSs are in connecting the patient's medical history to medical knowledge; it is always possible that the final strong deduction is based on some incomplete input. It is also discussed in [11] that: "No matter how good you are at diagnosing and treating, unless you asked the right questions in a timely manner, all the knowledge in the world won't be helpful." Hence, a focused H&P is the key to a flawless diagnosis. As the shortage of physicians is worsening in the recent years, roles such as Physician Assistant (PA) and Nurse Practitioner (NP) have been introduced in order to ease the problem. PAs and NPs are qualified to perform the H&P step, diagnose medical problems and carry out necessary treatments mainly under the supervision of a physician. Undertaking the H&P step, they would help the doctors to be able to see more patients within a certain period of

time, as they would just need to review and assess the already prepared H&P report. A system capable of guiding a focus H&P, however, will allow less experienced nurses to perform this process and can also provide second opinions in critical cases. Such system could also be added to available MDSs in order to provide them with the essential comprehensive input.

As the DDX domain meets the characteristics of holonic domains (see Sect. 2), the general research question here is whether a holonic multi-agent architecture could in practice support the implementation of the DDX. Hence, in fact, the development of a Holonic Multi-Agent System (HMAS), which is capable of performing DDX, is the practical contribution of this work. On the other hand, as is the case with intelligent systems, the introduction of the Machine Learning (ML) techniques that support the functionality of this system, is the theoretical contribution of this work. In this regard, the sub-question being answered in this paper is how we can provide the system with the right learning data and more particularly unbiased feedback.

Hence, the rest of the paper is organized as follows. Section 2 briefly describes the holonic nature of DDX, and in Sect. 3, the Holonic Medical Diagnosis System (HMDS) is introduced (for a more comprehensive description of the system please refer to [12]). Section 4 describes the ML techniques that have been adapted and applied to the system, and Sect. 5 concentrates on the rewarding process. Section 6 concludes with a summary and some directions for the future researches.

2 The Holonic Nature of DDX

The domain attributes that indicate the appropriateness of a multi-agent based solution are presented in [13]. These criteria are clearly met by the DDX problem. Firstly, the environment is open, dynamic, uncertain and complex. Moreover, considering the doctors with different specialties and subspecialties, a MDS can be naturally modelled as societies of cooperating agents, and hence here the agents are natural metaphors. This also implies the distribution of data, control and expertise.

An interesting overview of Multi-Agent System (MAS) architectures is presented in [14]. One of the well-known architectures is the HMAS architecture. The term holon was first introduced in [15] in order to name recursive and self-similar structures. A holonic agent of a well-defined software architecture may join several other holonic agents to form a super-holon, which is then seen as a single holonic agent with the same software architecture [16]. A holon can be a member of several super-holons, if the super-holons' goals do not contradict each other or the holon is indifferent to those conflicting goals [16]. The organizational structure of a holonic society, i.e. the holarchy, offers advantages such as robustness, efficiency, and adaptability [16]. Obviously, all these advantages do not mean that this architecture is more effective than the others, however, specifically for the domains, it is meant for. As stated in [16] such domains should (1) involve actions of different granularities, (2) induce abstraction levels, (3) demand decomposable problem settings, (4) require efficient intra-holonic communications, (5) include cooperative elements, and (6) may urge real-time actions.

The DDX problem can be decomposed recursively into subproblems of weighting the probability of the presence of the possible diseases. These subproblems may induce different abstraction levels and can be of different granularities. According to the nature of DDX, the problem solvers are collaborative and those dealing with similar diseases need to have more communications, which are to be conducted in timely manner. Hence, this domain meets the characteristics of holonic domains. Another attractive characteristic of HMASs is the support for self-organization, which is the autonomous continuous arrangement of the parts of a system in such a way that best matches its objectives. The HMDS also takes advantage of this characteristic (see Sect. 4).

3 The Holonic Medical Diagnosis System (HMDS)

The HMDS, first introduced in [17], is a HMAS, consisting of medical experts (holons). This structure is applied to this system in a way that supports DDX.

3.1 The Architecture of the HMDS

In general, a MDS either rely on highly smart deliberative agents as one extreme or on a large set of comparatively simple (reactive) agents as the other extreme. The HMDS as a HMAS realizes an improved version of the latter approach. It consists of two types of agents: comparatively simple Disease Representative Agents (DRA) as the system fundamentals on the lowest level and more sophisticated Disease Specialist Agents (DSA) as decision makers on the higher levels of the system [12] (see Fig. 1).

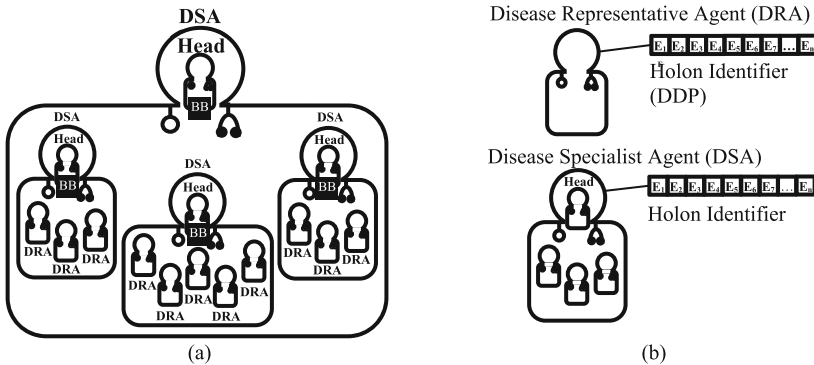


Fig. 1. (a) DRAs and DSAs in HMDS (BB: Blackboard) (b) Holon identifier

DRAs are atomic agents, and form the leaves of the holarchy. Each DRA is an expert on a specific disease and maintains a pattern store that contains the Disease Description Pattern (DDP) – an array of possible signs, symptoms, and test results, i.e. the holon identifier. Thus, in order to join the diagnosis process, these agents only need to perform some kind of pattern matching, i.e. calculating their Euclidean distance to

the diagnosis request. DSAs are holons consisting of numbers of DRAs and/or DSAs with similar symptoms; i.e., representing similar diseases. This encapsulation, in fact, enables the implementation of the DDX. The DSAs are designed as moderated groups [16], where agents give up part of their autonomy to their super-holon. To this end, for each DSA, a head is defined, which provides an interface and represents the super-holon to its environment. This head is created for the lifetime of the super-holon, based on agent cloning [18]. For this purpose, each head is capable of cloning, i.e., creating a copy of its code, and passing the relevant information to the new agent [12]. In this instance, the holon identifier will be the average of the holon identifiers of the members.

The holarchy has one root, in fact a DSA, which will play the role of the common and exclusive interface to the outside world. The system starts with this DSA, takes all the DRAs as its members, and then let the DSAs form automatically. Although this process is based on the affinity and satisfaction [19], and at the beginning, no information about the satisfaction factor is available, it is still possible to initially form the DSAs based on the affinity, i.e., the similarity between the diseases¹. To this end, the mentioned DSA accepts the DRAs, as its members, clusters them, and defines for each of the clusters, i.e. super-holons, a head. This is repeated recursively until no further clustering is necessary (see Sect. 4.1). This can be performed once as the system is being defined and accelerate the self-organization. Later on, the system can still refine its architecture using its self-organization technique (see Sect. 4).

3.2 The Functionality of the HMDS

In principle, the HMDS works as follows: The head of the holarchy receives the diagnosis request as a specific combination of signs, symptoms and medical test results and places it on its blackboard². Each agent of the system that has knowledge of this blackboard, i.e. any member of this super-holon, can read the messages on this blackboard. A DRA's reaction to a request message is to send back its similarity to the request. However, based on the provided information a DSA can decide whether it wants to join the diagnosis process or not. This will actually control the data flow in the holarchy. This decision is made based on some simple statistical information about the DSA's members. Considering the holon identifier of the members, the head will check if the request would be an outlier, and in negative cases it will decide to join the diagnosis process. This means that it will read the information from its head's blackboard and will place it on its own blackboard. Then the same process repeats recursively until the request reaches the final level of the holarchy. Results obtained by participating agents now flow the other way round from bottom to the top of the holarchy and are sorted according to their similarity and frequency. Literally, each agent will send its top diagnoses together with all of the signs, symptoms or test results that are relevant from its point of view, but their presence or absence has not been specified. This implies that originally not

¹ The Euclidean distance between the DDPs of the DRAs.

² In HMDS, the communications are solely done via the blackboard systems [30] of the DSAs.

provided relevant information might be requested from the user in a second step [12]. Section 3.3 demonstrates the functionality of the system.

3.3 The Examples of the System Simulations

The simulations presented in this paper have been conducted using the GAMA platform, and the disease-related data have been gathered from Mayo Clinic website [20]. For the following cases, a small system with only 20 diseases has been simulated.

Case 1 (Diagnosis of Lung Cancer). The first example is based on an actual H&P report, which is provided in [21] for study purposes. In our simulation, a hierarchy with three levels is formed. The chief complaint is entered into the system and the system's reaction is monitored using the actual report as a benchmark. The H&P report includes:

Chief Complaint (CC): Shortness of Breath (SOB)

History of Present Illness (HPI): Chest pain, Chills, Cough, Fever, History of breathing troubles, History of cancer in family, History with tobacco, Night Sweats, Productive cough, Vomit, Weight loss, Wheezing

Review of Systems (ROS): Anxiety, Fainting, Fatigue, Weakness, Heart palpitations and arrhythmias, Changes in skin color, Changes in appetite

Physical Examination (PE): Cyanosis, Edema, Swollen lymph nodes

Diagnostic Tests: Blood test, CT scan

Assessment and Plan: (1) Asthma (Asthma tests), (2) Lung Cancer (X-ray/CT scan, Biopsy), (3) Pneumonia (Blood Test), (4) Sarcoidosis (Blood Test), (5) Tuberculosis (PPD: Purified Protein Derivative skin test for tuberculosis).

In HMDS, entering the SOB, the DSA of the pulmonary diseases will be activated and the suggested signs and symptoms to be checked includes: Anxiety, Chest discomfort, Chills, Cough, Cyanosis, Diarrhea, Fainting, Fatigue and Weakness, Fever, Presence of frequent respiratory infections, Heart palpitations and arrhythmias, Hoarseness, Itching, Loss of appetite, Nausea and vomit, Night sweats, Phlegm, Sweats, Edema, Swollen lymph nodes, Weight loss, and Wheezing. These signs and symptoms very much match the ones mentioned in HPI, ROS and PE sections of the original H&P report. After entering the value of these signs and symptoms according to their presence or absence, the final DDX list would be: (1) Asthma, (2) lung cancer, (3) Pulmonary Edema, (4) Tuberculosis, (5) Sarcoidosis, (6) Pneumonia, (7) Bronchitis, (8) Pulmonary Embolism, (9) COPD, and (10) Lymphoma. The suggested medical tests would be asthma tests, x-ray/CT scan, sputum cytology, biopsy, pulse oximetry, arterial blood gas analysis, and sputum test for tuberculosis. These results match the actual H&P to a considerable degree and could even be improved through learning. In this case, the CT scan showed a large mediastinal mass, and the final diagnosis was lung cancer.

Case 2 (Metastatic Lung Cancer to Common Bile Duct Cancer). The HMDS also acts well in the presence of multiple diseases, e.g. the metastasis cases. This example is extracted from a medical paper in [22]. The signs and symptoms in this case included abdominal pain, coarse breath sounds, dry cough, jaundice, and shortness of breath; and the final Diagnosis was metastatic lung cancer to common bile duct cancer, with

suggested medical tests to be: Blood test, CT scan, ERCP (Endoscopic Retrograde Cholangiopancreatography), and Biopsy. Giving these symptoms to the HMDS two different DSAs will be activated: the DSA of pulmonary diseases and the DSA of Hepatology and Gastrointestinal Disorders. Their super-holon will then put the output of both members in order. The DDx list will be: (1) Bile Duct Cancer, (2) Cholangitis, (3) Asthma, (4) Lung Cancer, (5) Hepatitis B, (6) Pulmonary Edema, (7) PSC, (8) Pulmonary Embolism, (9) Bronchitis, (10) Lymphoma. This DDx list actually includes the bile duct cancer as the first and the lung cancer as the fourth possible diagnosis, and therefore the possibility of Metastasis can be clearly mentioned to the doctor.

4 Learning in HMDS

As mentioned in Sect. 3.1, the initial holarchy of the system can be created using clustering in different levels of the holarchy. Clustering is an unsupervised learning technique and hence doesn't require any learning feedback. After having the initial holarchy, however, it is still essential to support the system in learning and updating based on the new observations. Medical knowledge demonstrates a steady upward growth, and diagnosis is also very much affected by the geographical regions. Hence, in order to adapt and improve the behavior of the system, we need to: (1) Update the medical knowledge based on the new instances, (2) Refine the holarchy according to the experience and the feedback. In HMDS, holon identifiers are updated applying the exponential smoothing method, as a supervised learning method, and the self-organization of the holarchy is supported by Q-learning, as a reinforcement learning technique. Since both methods require feedback from the environment, it is clear that the quality of the feedback is of central importance. The rest of this section covers the above mentioned techniques and provides an experimental example, however, the discussion on the learning feedback can be followed in Sect. 5.

4.1 The Machine Learning Techniques Used in HMDS

Clustering. The Density-Based Spatial Clustering of Applications with Noise (DBSCAN) [23] is one of the best algorithms for clustering in HMDS. In [24] a simple and effective method for automatic detection of the input parameter of DBSCAN is presented, which helps best to deal with the complexity of the problem at hand.

Exponential Smoothing. Exponential smoothing is a very popular scheme for producing smoothed time series [25]. Using this technique, the past observations are assigned exponentially decreasing weights and recent ones are given relatively higher weights:

$$s_t = \alpha \cdot x_t + (1 - \alpha) \cdot s_{t-1} \quad (1)$$

where α is the smoothing factor, and $0 < \alpha < 1$. As it is clear, the smoothed statistic s_t is a simple weighted average of the current observation x_t and the previous smoothed statistic s_{t-1} . In HMDS, this method is used in order to update the holon identifiers.

Holonic-Q-Learning (HQL). The Holonic-QL is a Q-learning technique introduced for self-organization in the HMDS [12]. In HQL, the Q-value is in fact measuring how good it is for a holon to be a member, of another holon. In this case, the states are the existing holons $\{h_i\}$ and action h_i indicates becoming a sub-holon of holon i :

$$Q_t(\text{sub}(h), h) \leftarrow (1 - \alpha_t)(1 - \alpha_t)Q_{t-1}(\text{sub}(h), h) + \alpha_t(R_t(\text{sub}(h), h) + \gamma \operatorname{argmax}_{Q_{t-1}(h, \text{sup}(h))} (Q_{t-1}(h, \text{sup}(h)).\text{Aff}(\text{sub}(h), \text{sup}(h)))) \quad (2)$$

where, $\alpha_t = \frac{1}{1 + \text{visits}_t(\text{sub}(h), h)}$, $\gamma \in [0, 1)$ is the discount factor, $\text{Aff}(\text{sub}(h), h) = 1 - \frac{d(\text{sub}(h), h)}{\max d(\text{sub}(h), h)}$, and the reward is calculated by the head of a super-holon (Sect. 5). For more information regarding the HQL and the proof of convergence please refer to [12].

4.2 An Example of Learning in HMDS

Case 3 (A New DDX for Arthritis). In this simulation, the system covers 45 diseases in a holarchy with four levels. Here, again, a real case is used, in which the Madelung-Launois-Bensaude disease (MLB) is suggested as a new DDX for arthritis [26]. MLB is a disease that causes the concentration of fatty tissue in proximal upper body. In 2008, for the first time some instances of this disease have been observed with distal fatty tissue that were misdiagnosed first as arthritis, which normally includes joint pain and joint swelling. In this experiment, the new observations will be given to the HMDS, which so far does not consider the MLB disease with arthritis for the reason of DDX. The system should then be able to come to the same conclusion as in [26] and add the DRA acting for the MLB disease to the super-holon containing the DRA acting for arthritis. Essentially, if an agent is not involved in a diagnosis process it will not receive any reward and hence its Q-value will remain same during that round. As the agent participates in a diagnosis process, it will be rewarded and hence its Q-value will be updated. In case the Q-value of any of the members of the super-holon is getting close to be a noise (close to lower three-sigma limit), the agent will start exploring new opportunities to join new super-holons. One promising approach for this agent is to try to become a member of those super-holons that were activated at the same time with its current super-holon. This will guarantee that the agent would have some common interests with the members of its new super-holon(s).

Figure 2 shows the changes in the Q-values of the different DRAs in case 3. Entering the distal MLB instances into the system, the Q-value of the DRA acting for the MLB disease will get closer to the lower outlier threshold. At this point, the super-holon of the arthritis disease will also be activated and as the DRA acting for the MLB disease is looking for a chance to join some new super-holons, it will be informed and will try to become a member of this super-holon. Considering the holon identifier of its members, this super-holon will then check whether the DRA acting for the MLB disease would be an outlier and since the case is negative, it will accept this new




Fig. 2. Changes in the Q-values of the different DRAs in case 3

member. At this stage, the Q-value will be an accumulated value, calculated based on the rewards received from both super-holons. As it can be followed on the diagram, since this value is now greater than the average of the Q-values in both super-holons, the DRA acting for MLB will stop exploring at this point.

5 Reward Engineering in HMDS

It should be noted that the terms “feedback” and “reward” have different definitions here. A feedback is the final diagnosis, suggested by the physician for a given diagnosis request. However, a reward is a numerical value, which is calculated using a reward function that considers the feedback. Agents receive their rewards from their environments. The HMDS has distributed environments, i.e. the highest holon receives the feedback from the outside world, however, the environment for the rest of the holons is their own super-holon, from which they receive their rewards.

In the HMDS, the system will announce the result in form of a list of top 10 possible diseases, and will then receive a feedback, i.e. the final diagnosis, from the physician, and would consider the physician diagnosis for the reason of reward calculation. However, the system cannot simply take this diagnosis to calculate the rewards, but should have a strong policy to reduce the possibility of errors and biased diagnoses. The system is meant to be used in hospitals, and hence, is actually multi-user and receives the feedback from different physicians. Moreover, the system will consider counterfactual learning, which refers to the ability to learn from forgone outcomes, i.e. the outcome of the option(s) that were not chosen [27]. The use of this method literally improves decision-making when evidence is solid but causal links are unclear [28]. In order to design a mechanism to consider the other options next to the physician’s final diagnosis, the probability of selecting the options is estimated with a combination of ϵ -greedy and softmax rules. Using this method the physician’s diagnosis is given the highest selection probability, i.e. $1 - \epsilon$, and all the others are ranked

and weighted using the softmax rule. For this reason, we have used the most common softmax method, which uses a Gibbs or Boltzmann distribution, and have adapted it to our problem as follows:

$$p(d_i) = \varepsilon \frac{e^{simF_i/\tau}}{\sum_{j=0}^9 e^{simF_j/\tau}} \quad (3)$$

where d_i for $0 \leq i \leq 9$ is a disease in the output list, $simF_i$ is the product of the similarity of d_i to the diagnosis request³ and its frequency, and τ is a positive number called temperature. Low temperatures cause a greater difference in the selection probability of the diseases.

After clarifying the problem with the feedback, next step will be to design the reward protocol. The term “reward engineering”, first coined by Dewey in [29], refers to the engineering of the agent’s environment in order to make the reward assignment more reliable. According to [29], in reward engineering the agent’s goal is not changed, however, the environment is being partially designed so that reward maximization leads to desirable behavior. For this reason, the human factor may even be removed from the loop and the reward may be assigned via an automatic mechanism.

In the HMDS, the holon which is on the top level of the holarchy will receive the final diagnosis⁴ and will place it on its blackboard. Consequently, all the members will be informed and the DSAs would repeat the same action, i.e. place the final diagnosis on their blackboard, until the announcement reaches the DRAs⁵. The DSAs, which have participated in the diagnosis process, are now responsible for calculating the rewards for each of their members. In such a super-holon, the DRAs and the DSAs, which have not participated in the diagnosis process, will respond to their head by sending back their similarities to the final diagnosis. The DSAs, which have originally participated in the diagnosis process, however, will respond by sending back the highest reward they have given to their members. These values are called the suggested rewards. Considering the suggested rewards, the super-holon will take the highest value as its highest reward and then assign rewards to its members based on the three-sigma rule:

$$r_i = \begin{cases} 0, & suggR_i \text{ is an outlier} \\ r^*(3\sigma - |\bar{r} - suggR_i|)/3\sigma, & \text{else} \end{cases} \quad (4)$$

where $suggR_i$ is the actual value suggested by the i – th member, \bar{r} is the mean of the suggested rewards, r^* is the highest reward, and σ is the standard deviation of the suggested rewards. Hence, the value of the rewards, except for the case of a penalty, is a fraction of the r^* (always between zero and one), and implicitly considers the

³ The Euclidean distance between the corresponding DDP and the diagnosis request.

⁴ The holon identifier of the DRA acting for the final diagnosis, which is already updated considering the new observation.

⁵ If a super-holon, which has originally not participated in the diagnosis process, realizes that the disease given as the final diagnosis matches one of its members, it will assign the reward -1 (penalty) to this member.

satisfaction and affinity factors both. Greater reward values indicate that the problem is more relevant to the super-holon, and thus incompatibilities will be penalized more. Hence, in case a member of a super-holon is close to be a noise, its assigned reward and consequently Q-value will differ greatly from the average rewards and Q-values. This would then lead to higher probabilities of exploration (An agent will decide to explore new opportunities if its Q-value is close to be an outlier in its super-holon).

6 Conclusion

This paper explained that one of the main limitations of the available MDSs is the lack of support for the H&P examination, which prevents their integration in the clinical workflow. The H&P examination uses the DDx method and this domain is in fact a holonic domain, hence, a MDS with holonic architecture could be able to perform the DDx process. The HMDS tends to cover the stages in the H&P process, and hence allow less experienced nurses to perform this process, and can also provide a second opinion in critical cases. This research also discussed that this system should be supported with appropriate ML techniques in order to maintain and improve its functionality. Along with the learning methods, the quality of the learning data, i.e. the feedback is also of central importance. In HMDS, the final report is given to a physician and s(he) makes the final diagnosis, however, in rewarding stage the system cannot simply take this diagnosis, but should have a strong policy to reduce the possibility of errors and biased diagnoses. For this purpose, the system, which is meant to be used in hospitals is actually considered as a multi-user system and hence receives the feedback from different physicians. Moreover, it will apply counterfactual learning and considers the option(s) that were not chosen. This paper suggests a combination of ϵ -greedy and softmax rules for this purpose. As demonstrated in this paper, the simulation of the system is now in progress and future work will include the complete simulation and validation of the system based on the reward engineering proposed in this work.

References

1. Berner, E.: Clinical Decision Support Systems: Theory and Practice. Springer, New York (2016). <https://doi.org/10.1007/978-0-387-38319-4>
2. Merriam-Webster: Differential Diagnosis. <https://www.merriam-webster.com/dictionary/differential%20diagnosis>
3. Maude, J.: Differential diagnosis: the key to reducing diagnosis error, measuring diagnosis and a mechanism to reduce healthcare costs. *Diagnosis* **1**(1), 107–109 (2014)
4. Segen, J.: Concise Dictionary of Modern Medicine. McGraw-Hill, New York (2006)
5. IBM Watson Health. <https://www.ibm.com/watson/health/>
6. Fisher, H., Tomlinson, A., Ramnarayan, P., Britto, J.: ISABEL: support with clinical decision making. *Pediatr. Nurs.* **15**(7), 34–35 (2003)
7. Saxena, M.: IBM Watson Progress and 2013 Roadmap. https://www.slideshare.net/manoj saxena2/ibm-watson-progress-and-roadmap-saxena/7-Watson_Healthcare_Products_1H_2013. Accessed 22 Mar 2018

8. Riches, N., Panagioti, M., Alam, R., Cheraghi-Sohi, S., Campbell, S., Esmail, A., Bower, P.: The effectiveness of electronic differential diagnoses (DDX) generators: a systematic review and meta-analysis. *PLoS One* **11**(3), e0148991 (2016)
9. Yuan, M.: Watson and healthcare: how natural language processing and semantic search could revolutionize clinical decision support. <https://www.ibm.com/developerworks/library/os-ind-watson/>. Accessed 22 Mar 2018
10. Graber, M.L., Mathew, A.: Performance of a web-based clinical diagnosis support system for internists. *J. Gen. Intern. Med.* **23**(1), 37–40 (2008)
11. Husain, I.: Why IBM's artificial intelligence "Watson" could not replace a physician. In: *iMedicalApps*. <https://www.imedicalapps.com/2011/02/ibm-watson-replace-physician-artificial-intelligence/>. Accessed 22 Mar 2018
12. Akbari, Z., Unland, R.: A holonic multi-agent system approach to differential diagnosis. In: Berndt, J.O., Petta, P., Unland, R. (eds.) *MATES 2017. LNCS (LNAI)*, vol. 10413, pp. 272–290. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-64798-2_17
13. Wooldridge, M.: Chapter 9: Methodologies. In: *An Introduction to MultiAgent Systems*. Wiley (2009)
14. Lavendelis, E., Grundspenkis, J.: Open holonic multi-agent architecture for intelligent tutoring system development. In: *Proceedings of the IADIS International Conference on Intelligent Systems and Agents* (2008)
15. Koestler, A.: *The Gost in the Machine*. Hutchinson, Paris (1967)
16. Gerber, C., Siekmann, J., Vierke, G.: Holonic multi-agent systems. Technical report DFKI-RR-99-03, German Research Centre for Artificial Intelligence (1999)
17. Unland, R.: A holonic multi-agent system for robust, flexible, and reliable medical diagnosis. In: Meersman, R., Tari, Z. (eds.) *OTM 2003. LNCS*, vol. 2889, pp. 1017–1030. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-39962-9_97
18. Shehory, O., Sycara, K., Chalasani, P., Jha, S.: Agent cloning: an approach to agent mobility and resource allocation. *IEEE Commun. Mag.* **36**(7), 58–67 (1998)
19. Rodriguez, S.: From analysis to design of holonic multi-agent systems: a framework, methodological guidelines and applications. Ph.D. thesis, University of Technology of Belfort-Montbéliard (2005)
20. Mayo Clinic. <https://www.mayoclinic.org/>
21. University of North Carolina - School of Medicine, History and Physical Examination (H&P) Examples. <https://www.med.unc.edu/medclerk/education/grading/history-and-physical-examination-h-p-examples>
22. Cochrane, J.: Metastatic lung cancer to the common bile duct presenting as obstructive jaundice. *J. Hepatol. Gastrointest. Disord.* **2**(121) (2016). <https://www.omicsonline.org/open-access/metastatic-lung-cancer-to-the-common-bile-duct-presenting-as-obstructivejaundice-jhgd-1000121.php?aid=69570>
23. Ester, M., Kriegel, H.-P., Sander, J., Xu, X.: A density-based algorithm for discovering clusters in large spatial databases with noise. In: Simoudis, E., Han, J., Fayyad, U. (eds.): *Proceedings of the 2nd International Conference on Knowledge Discovery and Data Mining (KDD-1996)*, pp. 226–231 (1996)
24. Akbari, Z., Unland, R.: Automated determination of the input parameter of the DBSCAN based on outlier detection. In: *Artificial Intelligence Applications and Innovations, IFIP Advances in Information and Communication Technology*, vol. 475, pp. 280–291 (2016)
25. NIST/SEMATECH e-Handbook of Statistical Methods. <http://www.itl.nist.gov/div898/handbook/>. Accessed Jan 2018
26. Lemaire, O., Paul, C., Zabraniecki, L.: Distal Madelung-Launois-Bensaude disease: an unusual differential diagnosis of acromalic arthritis. *Clin. Exp. Rheumatol.* **26**, 351–353 (2008)

27. Palminteri, S., Lefebvre, G., Kilford, E., Blakemore, S.-J.: Confirmation bias in human reinforcement learning: evidence from counterfactual feedback processing. *PLoS Comput. Biol.* **13**(8), e1005684 (2017)
28. Kayes, D., Anna, K.: *The Learning Advantage: Six Practices of Learning-Directed Leadership*. Springer, London (2011). <https://doi.org/10.1057/9780230305595>
29. Dewey, D.: Reinforcement learning and the reward engineering principle. In: 2014 AAAI Spring Symposium Series (2014)
30. Corkill, D.: Blackboard systems. *AI Experts* **6**(9), 40–47 (1991)