

Agent-based Decision Support for Pain Management in Primary Care Settings

Xu Guo¹, Han Yu^{1,2}, Chunyan Miao^{1,2} and Yiqiang Chen^{3,4}

¹School of Computer Science and Engineering, Nanyang Technological University (NTU)

²Joint NTU-UBC Research Centre of Excellence in Active Living for the Elderly (LILY)

³School of Computer Science and Technology, University of the Chinese Academy of Sciences

⁴Institute of Computing Technology, Chinese Academy of Sciences

xu008@e.ntu.edu.sg, {han.yu, ascymiao}@ntu.edu.sg, yqchen@ict.ac.cn

Abstract

The lack of systematic pain management training and support among primary care physicians (PCPs) limits their ability to provide quality care for patients with pain. Here, we demonstrate an Agent-based Clinical Decision Support System to empower PCPs to leverage knowledge from pain specialists. The system learns a general-purpose representation space on patients, automatically diagnoses pain, recommends therapy and medicine, and suggests a referral program to PCPs in their decision-making tasks.

1 Introduction

Pain is a major health problem around the world, causing negative economic, social and productivity impact [Goldberg and McGee, 2011]. Specialist care is required for effective intervention as pain is often accompanied by physiological and psychological disorders [Caudill, 2016]. Before patients are referred to specialists, primary care physicians (PCPs) are often consulted first. PCPs work with a consistent patient base to provide point-of-care interventions. However, in a recent survey of 500 PCPs, more than 70% reported a lack of confidence in treating patients with pain [Simon, 2012]. The lack of systematic training on pain management proves to be a major barrier in providing effective pain management in primary care settings. PCPs can benefit from insight and expertise possessed by pain specialists if there is a system to help them tap into this resource.




Figure 1: The system diagram of ACDSS.

In this paper, we showcase the Agent-based Clinical Decision Support System (ACDSS)¹ to help PCPs deliver high quality care to patients with pain with a software agent as an intelligent user interface [Yu *et al.*, 2007; Yu *et al.*, 2010; Yu *et al.*, 2011; Lin *et al.*, 2015]. The core of decision support for different tasks in ACDSS (Figure 1) is built upon an anonymized Electronic Health Records (EHRs) database from the pain department of Tan Tock Seng Hospital, Singapore. The guidelines² for each learning task are automatically extracted from free-text clinical notes using natural language processing (NLP) phrase detection tools. Specifically, our system incorporates four task-oriented modules for general diagnosis, general therapy, prescription and patient referral, respectively. The system receives a patient's information in terms of pain conditions and health status, and transforms the textual data into fixed-length parameter representations for privacy protection. The agent supports PCP decision-making based on given patient representations, and updates its learning model based on PCPs' acceptance of recommendations.

2 The AI Engine

In clinical settings, several barriers continue to impede the effective implementation of clinical decision support systems (CDSS). Among the challenges, knowledge representation and reasoning are the areas most in need of optimization [Kong *et al.*, 2008]. When dealing with diseases with high-dimensional attributes, sparsity and a large volume of data, traditional methods (e.g. case-based reasoning [Blanco *et al.*, 2013]) incur heavy computation overhead and suffer from poor generalization performance. Against this background, there is a clear need for an efficient and general knowledge representation to support real-time clinical decision-making.

To this end, we propose a data-driven agent-based approach for PCP decision support system in four pain management related tasks: 1) providing general diagnosis, 2) proposing general therapies, 3) recommending prescriptions, and 4) referring patients to specialists. We achieve this objective by employing an unsupervised feature learning method called deep patient [Miotto *et al.*, 2016] to derive general-purpose

¹<https://guoxuxu.github.io/pain/video.html>

²<https://guoxuxu.github.io/pain/GLS.html>

patient representations from EHR data. The technique is to train a set of stacked denoising auto-encoders on patient descriptions, and then evaluate them by fitting Naive Bayesian classifiers predicting the labels to the learned representations. The labels, as the learning objectives in each task setting, are automatically extracted from free-text clinical notes using phrase detection techniques, and then organized into dictionaries (). For example, in task 3, the agent may suggest ‘gabapentine’ and ‘tramadol’ for a patient with low back pain symptoms.

Patient referral can be treated as an optimization task which needs to jointly consider factors such as patient healthcare, degrees of pain, specialist suitability, timeslots available for both the patient and the suitable specialists, etc., in order to efficiently utilize the limited pain specialist resources while providing high quality services to patients. Previous task allocation optimization approaches [Pan *et al.*, 2016; Yu *et al.*, 2015; Yu *et al.*, 2016; Yu *et al.*, 2017a; Yu *et al.*, 2017b; Yu *et al.*, 2019] can be leveraged.

We evaluate this patient representation approach based on our proprietary database for over 500 patients with pain. To measure the multi-label classifiers, we use the Macro-averaging method (equation 1) as the evaluation metric, which averages the accuracies of each classifier on all classes. The testset ratios are all 0.3 and the performance of the proposed approach for the four tasks are shown in Figure 2.

$$Accuracy^{macro}(D) = \frac{\sum_{c_i \in C} Accuracy(D, c_i)}{|C|} \quad (1)$$

It can be observed from Figure 2(a) that Task 4, referring patients to specialists, achieved the highest accuracy; while Task 3, prescription recommendation, achieved the lowest accuracy. One possible explanation is that a medication may be suitable for multiple pain symptoms and, therefore, may require additional information beside patient descriptions. From Figure 2(b), the recall rate for predicting therapies significantly decreases when the number of labels to be predicted exceeds 4. This implies that if a patient has more complex symptoms, the relationships between a patient and the therapy categories are harder to learn.

In general, the patient representation space derived by the proposed approach has useful potential for pain management in primary care settings. It can be applied to a wide range of tasks. For example, we achieved personalization treatment recommendation by searching for similar previous cases using the unsupervised patient representation.




Figure 2: Performance evaluation based on practical tasks.

3 The ACDSS System

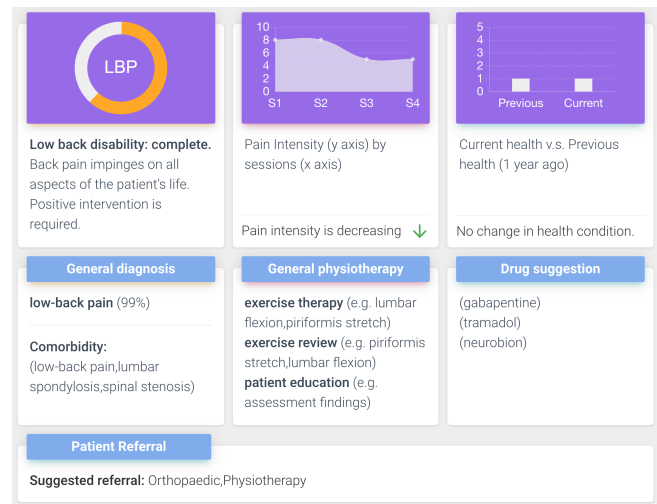


Figure 3: A screenshot of the functionality of the ACDSS system.

ACDSS is designed to improve data visualization, promote knowledge transfer and provide decision support for pain management in primary care settings.

The system consists of four decision support engines, each focusing on one of the four tasks mentioned above while sharing the same patient representation space (Figure 3). When a PCP enters information about a patient, the patient’s statistics are visualized at first, and then the text data (e.g., a patient’s demographics, descriptions of health state and pain symptoms, etc.) are transformed into parameters by the patient representation model for privacy protection. The agent plays its role based on the task-oriented model to make predictions or suggestions for the given patient representation.

In the diagnostic task, the agent predicts the types of pain and complications with its confidence level indicated. In the general therapy recommendation task, the agent first predicts which categories of the physiotherapy required, then searches for similar cases to find specific treatment details and presents them as a whole to PCPs. In the prescription recommendation task and the patient referral recommendation task, the agent makes suggestions in the same way based on the guidelines of each task. The recommended solutions that are accepted by the PCPs will be used to update the learning model for the specific task.

Acknowledgements

This research is supported by Nanyang Technological University, Nanyang Assistant Professorship (NAP); the Singapore Ministry of Education Academic Research Fund Tier 1 (Grant No. 2017-T1-001-270); the National Key Research and Development Plan of China (No. 2017YFB1002802); Beijing Natural Science Foundation (No 4194091); and Beijing Municipal Science & Technology Commission No. Z171100000117017.

References

- [Blanco *et al.*, 2013] Xiomara Blanco, Sara Rodríguez, Juan M Corchado, and Carolina Zato. Case-based reasoning applied to medical diagnosis and treatment. In *Distributed computing and artificial intelligence*, pages 137–146. Springer, 2013.
- [Caudill, 2016] Margaret A Caudill. *Managing pain before it manages you*. Guilford Publications, 2016.
- [Goldberg and McGee, 2011] Daniel S Goldberg and Summer J McGee. Pain as a global public health priority. *BMC public health*, 11(1):770, 2011.
- [Kong *et al.*, 2008] Guilan Kong, Dong-Ling Xu, and Jian-Bo Yang. Clinical decision support systems: a review on knowledge representation and inference under uncertainties. *International Journal of Computational Intelligence Systems*, 1(2):159–167, 2008.
- [Lin *et al.*, 2015] Han Lin, Jinghua Hou, Han Yu, Zhiqi Shen, and Chunyan Miao. An agent-based game platform for exercising people’s prospective memory. In *WI-IAT*, pages 235–236, 2015.
- [Miotto *et al.*, 2016] Riccardo Miotto, Li Li, Brian A Kidd, and Joel T Dudley. Deep patient: an unsupervised representation to predict the future of patients from the electronic health records. *Scientific reports*, 6:26094, 2016.
- [Pan *et al.*, 2016] Zhengxiang Pan, Han Yu, Chunyan Miao, and Cyril Leung. Efficient collaborative crowdsourcing. In *Thirtieth AAAI Conference on Artificial Intelligence*, pages 4248–4249, 2016.
- [Simon, 2012] Lee S Simon. Relieving pain in america: A blueprint for transforming prevention, care, education, and research. *Journal of pain & palliative care pharmacotherapy*, 26(2):197–198, 2012.
- [Yu *et al.*, 2007] Han Yu, Zhiqi Shen, and Chunyan Miao. Intelligent software agent design tool using goal net methodology. In *IAT*, pages 43–46, 2007.
- [Yu *et al.*, 2010] Han Yu, Yundong Cai, Zhiqi Shen, Xuehong Tao, and Chunyan Miao. Agents as intelligent user interfaces for the net generation. In *Proceedings of the 15th international conference on Intelligent user interfaces*, pages 429–430. ACM, 2010.
- [Yu *et al.*, 2011] Han Yu, Zhiqi Shen, Chunyan Miao, and Ah-Hwee Tan. A simple curious agent to help people be curious. In *AAMAS*, pages 1159–1160, 2011.
- [Yu *et al.*, 2015] Han Yu, Chunyan Miao, Zhiqi Shen, Cyril Leung, Yiqiang Chen, and Qiang Yang. Efficient task sub-delegation for crowdsourcing. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- [Yu *et al.*, 2016] Han Yu, Chunyan Miao, Cyril Leung, Yiqiang Chen, Simon Fauvel, Victor R Lesser, and Qiang Yang. Mitigating herding in hierarchical crowdsourcing networks. *Scientific reports*, 6(1):4, 2016.
- [Yu *et al.*, 2017a] Han Yu, Chunyan Miao, Yiqiang Chen, Simon Fauvel, Xiaoming Li, and Victor R Lesser. Algorithmic management for improving collective productivity in crowdsourcing. *Scientific reports*, 7(1):12541, 2017.
- [Yu *et al.*, 2017b] Han Yu, Zhiqi Shen, Simon Fauvel, and Lizhen Cui. Efficient scheduling in crowdsourcing based on workers’ mood. In *2017 IEEE International Conference on Agents (ICA)*, pages 121–126. IEEE, 2017.
- [Yu *et al.*, 2019] Han Yu, Chunyan Miao, Yongqing Zheng, Lizhen Cui, Simon Fauvel, and Cyril Leung. Ethically aligned opportunistic scheduling for productive laziness. *AAAI/ACM conference on Artificial Intelligence, Ethics, and Society*, 2019.