# A Knowledge Graph Based Health Assistant

**Lin Bo, Wenjuan Luo, Zang Li, Xiaoqing Yang, Han Zhang, Daxin Zheng**
AI Labs, Didi Chuxing
Beijing, China
{bolinlin,luowenjuan,lizang,yangxiaoqing,arsenezhang,zhengdaxin}@didiglobal.com

## Abstract

With the rapid development of artificial intelligence, especially natural language processing technologies, chatbots have been designed and implemented for various use cases, ranging from e-commerce customer service to travel reservation agents. In the area of healthcare, it is of great importance to develop an intelligent personalized health assistant chatbot that can help people analyze their medical symptoms and pay more attention to their health. In this paper, we introduce a medical knowledge-based dialogue system that acts as a health assistant which provides medical advice and nutrition suggestions based on user's disease symptoms reported via chat. With the construction of a medical knowledge graph, our chatbot system can ask about relevant symptoms and achieve a more accurate search range of possible diseases. With well-designed algorithms, it efficiently matches symptoms against its medical knowledge base and recommend a suitable diagnosis. Our health assistant chatbot was launched in 2018 and has served hundreds of thousands of users who are drivers working in Didi Chuxing (DiDi) online ride-sharing platform.

## 1 Introduction

As the internet economy booms, tens of millions of drivers work in the online ride-sharing platforms for better job flexibility. Epidemical report [9] reveals that some characteristics of occupational exposure in drivers (e.g. stress, traffic pollutants, unchanging postures, exposure to noise) may lead to certain medical issues such as cardiovascular diseases and musculoskeletal disorders. As pointed out in the Sustainable Development Goals for the period 2015-2030 [6], while providing job opportunities, sustainable economic growth requires to create safe and decent working conditions that allow people to have quality jobs that simulate the economy. In order to enhance working environment and lower health risk for ride-share drivers, we develop a personalized health assistant chatbot by leveraging advanced technologies in artificial intelligence and machine learning. This healthcare chatbot is able to recognize symptoms through conversations via text-messaging and search across its medical knowledge base to recommend solutions to health issues. As the chatbot offers immediate responses to common health problems, it makes professional medical knowledge more accessible.

The health assistant chatbot is developed based on medical knowledge graph, which is tailored to focus more on occupational diseases of drivers. We adopt technologies of knowledge graph and dialogue system to build a task-oriented chatbot that can recognize symptoms described by user and respond in real-time with a suitable diagnosis, like the way a doctor will do. In Section 2 we describe the structure of the chatbot system and the construction of the medical knowledge graph. Experiment results of relation classification are shown in Section 3. We discuss the potential of artificial intelligence technologies and some of the possible limitations in Section 4.

## 2 Methodology

### 2.1 Medical Knowledge-based Dialogue System

The architecture of the medical knowledge-based dialogue system is illustrated in Figure 1. It consists of three components:

AI for Social Good workshop at NeurIPS (2019), Vancouver, Canada.

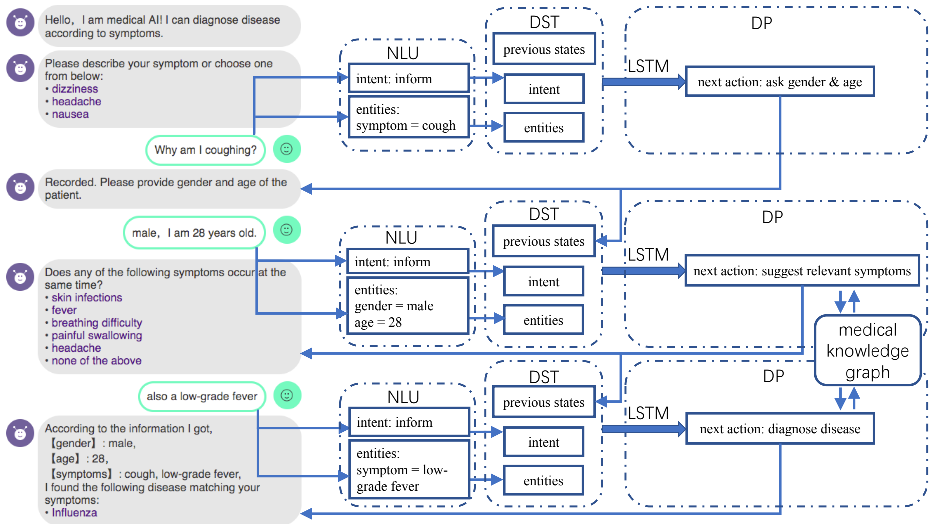


Figure 1: Architecture of the medical knowledge-based dialogue system. On the left panel, we show an example conversation with the user consulting for his symptoms.

- Natural Language Understanding (NLU): This module analyzes the user's text-input and extracts useful information. Specifically it can be divided into two parts: an intent classifier using convolutional-neural-network [5] and an entity extractor using conditional random field approach [3]. The entity set we defined for our healthcare chatbot contains user's basic information (i.e. gender, age) and disease symptoms.
- Dialogue state tracking (DST): It memorizes the intent and entities extracted at each dialogue turn during current conversation.
- Dialogue policy (DP): This module determines the next action of the chatbot and sends a proper response to the user, either asking for more information or providing a suitable diagnosis according to the medical knowledge graph.

At each dialogue turn, the chatbot system updates its dialogue state with the identified intent and entities. The next chatbot action is then determined by long short-term memory (LSTM) network [4] with sequence of historical dialogue states and actions as the model input. With disease symptoms and relevant information collected through multi-turn conversation with the user, the chatbot searches across its medical knowledge base and recommend an appropriate diagnosis. Moreover, in order to narrow the search range of possible diseases, relevant symptoms need to be confirmed and clarified by the user. Equipped with medical knowledge graph, the chatbot system is able to identify symptoms that frequently co-occur with the ones reported by the user. In the situation where there is not sufficient symptoms information to determine a reliable diagnosis, the chatbot will ask about relevant symptoms in the next dialogue turn to clarify user's health condition.

### 2.2 Medical Knowledge Graph Based Question Answering

Traditional approaches of developing knowledge based question answering (KBQA) systems could be classified into three categories [2]: rule based [8], keyword based [11], and synonym based [10]. Rule based approaches only understand a small set of "canned" questions by hand-crafted rules, which not only leads to low recall but also is very labour consuming. On the other hand, keyword based and synonym based approaches, although perform better than rule based systems at understanding user intents, still cannot fully understand the questions. To avoid these deficits, [2] designs to represent user questions by templates extracted from knowledge base. Different from all these methods, we use multi-round dialogue interaction mechanism to understand the user's question as shown in Fig. 1. To assist the question answering system, we construct a medical knowledge graph as illustrated in the following.

#### 2.2.1 Medical Knowledge Graph Construction and Completion

The construction of a knowledge graph requires multiple triplets in the form of (S, P, O), where S stands for Subject, P stands for Predicate, and O stands for Object. The technologies involved include

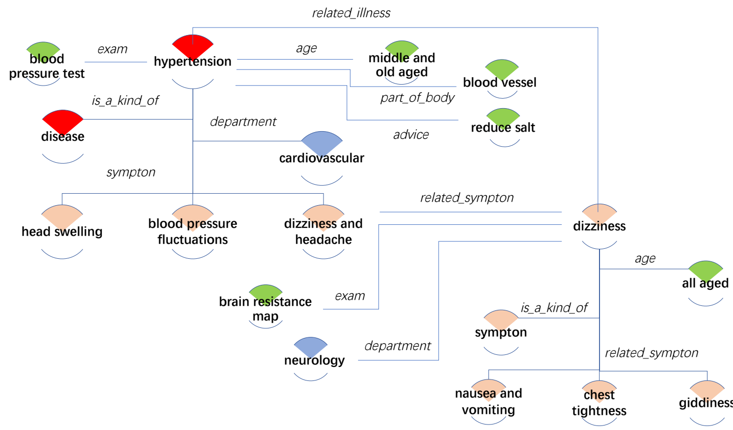


Figure 2: Subgraph of Medical Knowledge Base

web crawlers, data cleansing, structured information extraction, entity identification and relation classification. Specifically, we selectively construct a medical knowledge base that includes common diseases and symptoms of drivers and corresponding medical examinations and diet recommendations. In detail, our medical knowledge graph contains 5,863 disease or symptom entities, 26 attributes, and 39,112 SPO triplets. Fig 2 illustrates a subgraph of our knowledge base. Red nodes indicate diseases and pink nodes stand for symptoms. Blue nodes denote the diseases/symptoms' medical specialties, and green nodes represent other attributes of the entities.

As some attributes of entities are missing, especially the disease prevention and treatment advice, we use search engine to get information in order to fill these attributes. In detail, we specify the entity as the keyword, and obtain related corpora. After data cleansing, we perform entity identification and relation classification to generate SPO candidates. Finally, the set of SPOs are manually selected to construct the knowledge graph. We present experimental results of relation classification in Sec 3.

### 2.2.2 Question Answering

QA over a knowledge base gives out accurate and concise results, provided that natural language questions can be understood and mapped precisely to structured queries over the knowledge base [2]. In our system, question answering system mainly includes three steps: problem understanding, query construction and answer calibration. Firstly, problem understanding consists of entity identification and predicate recognition. Secondly, the query construction is the process of translating user question into query clause in the knowledge system. For example, suppose the back-end knowledge base is built on ES (Elastic Search), one needs to rewrite the user question into a ES query statement. Finally, answer calibration refers to re-ranking the relevant results and generating final answers.

## 3 Experiments

In this section, we conduct experiments to testify the effectiveness of our methods. Specifically for relation classification, we design our model based on a bi-directional GRU with attention(BIGRU+ATT) [12] and XGBOOST [1], where BIGRU with attention is used to train sentence embedding with respect to relation. XGBOOST utilizes the aforementioned sentence embedding and pre-trained entity embedding for final relation classification. Entity embeddings are pre-trained via word2vec [7] from a large Chinese wikipedia dataset, consisting of 65,001,293 triplets and corresponding sentence descriptions. For entity not present in the pre-training corpus, the entity embedding are computed as follows,

$$E_{emb} = \sum_{i=1}^{N} W_{i,emb}/N \tag{1}$$

where $E_{emb}$ denotes the embedding vector of entity $E$, and $W_{i,emb}$ is the embedding vector of word $W_i$, while $W_i$ is the $i$th word in the segmented word list of the entity. Table 1 presents the results and Fig 3 gives the confusion matrix.

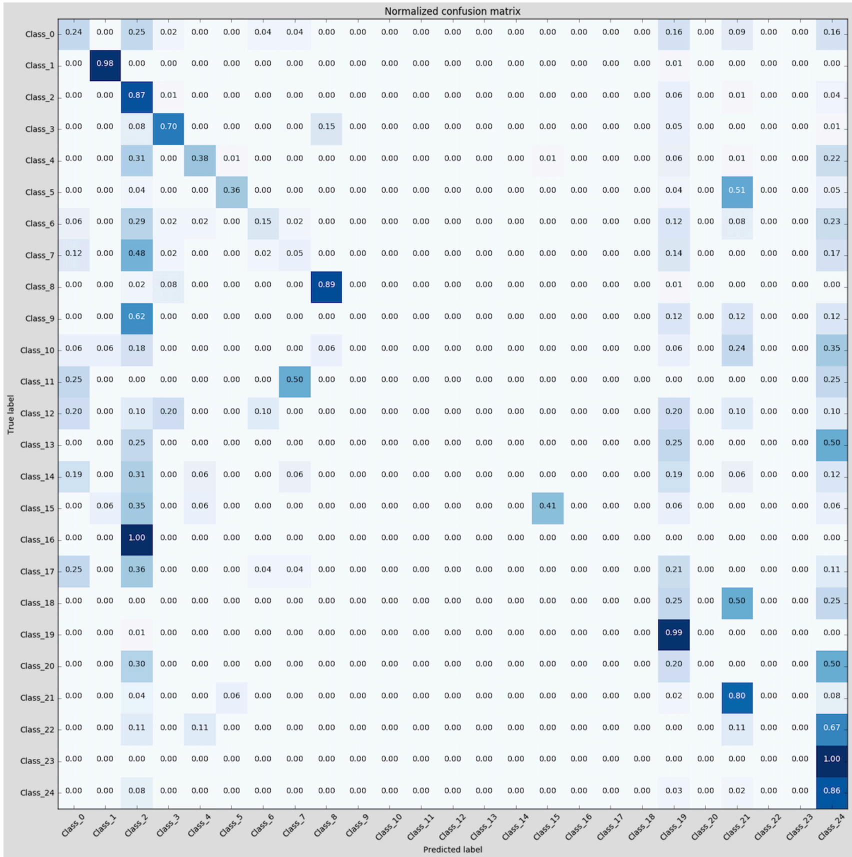


Figure 3: Confusion Matrix of Relation Classification

Table 1: Results for Relation Classification

| Method | avergePrecision | averageRecall | averageF1 |
|---|---|---|---|
| BIGRU+ATT | 0.6478 | 0.7790 | 0.7074 |
| Ours | **0.8820** | **0.8806** | **0.8813** |

## 4 Discussion and Conclusion

As language is the most natural embodiment of human intelligence, natural language processing is always considered as the most important technology for building up intelligent chatbots. Meanwhile, knowledge graph that can help chatbot memorize, associate and reason about semantic connections of entities, bridges the gap from perceived intelligence to cognitive intelligence. Note that logical reasoning remains a challenge in knowledge graph, in our case, only simple secondary reasoning is performed, such as from "symptom" to "disease" then to related "symptom". With multi-turn interactions with the user, our chatbot collects symptoms and offers preliminary medical advice, however it is not intended to replace professional medical advice, diagnosis or treatment. Regarding critical medical issues the chatbot not only responds with detailed knowledge of the disease but also directs the user to seek a doctor by providing contact information of doctor's office.

The health assistant chatbot has been deployed on DiDi's mobile application and WeChat official accounts platform since 2018 and has served hundreds of thousands of users who are ride-share drivers. It is integrated into a driver assistant application which also includes skills like weather forecast, gas station navigation, etc., and is one of the most popular sub-modules. It serves drivers with $\sim$8,000 conversations about symptom consultation daily. This healthcare chatbot, under the premise of enhancing health awareness of the drivers, also demonstrates the potential of artificial intelligence technologies in making ride-share driver a safe and respected profession.

## References


[1] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794. ACM, 2016.

[2] Wanyun Cui, Yanghua Xiao, Haixun Wang, Yangqiu Song, Seung-won Hwang, and Wei Wang. Kbqa: Learning question answering over qa corpora and knowledge bases. *Proc. VLDB Endow.*, 10(5):565–576, January 2017.

[3] Jana Diesner and Kathleen M. Carley. Conditional random fields for entity extraction and ontological text coding. *Computational and Mathematical Organization Theory*, 14(3):248–262, 9 2008.

[4] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

[5] Yoon Kim. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*, 2014.

[6] Bandy X Lee, Finn Kjaerulf, Shannon Turner, Larry Cohen, Peter D Donnelly, Robert Muggah, Rachel Davis, Anna Realini, Berit Kieselbach, Lori Snyder MacGregor, et al. Transforming our world: implementing the 2030 agenda through sustainable development goal indicators. *Journal of public health policy*, 37(1):13–31, 2016.

[7] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013.

[8] Shiyan Ou, Constantin Orasan, Dalila Mekhaldi, and Laura Hasler. Automatic question pattern generation for ontology-based question answering. *Proceedings of the Twenty-First International FLAIRS Conference (2008)*, pages 183–188.

[9] F Ronchese and M Bovenzi. [occupational risks and health disorders in transport drivers]. *Giornale Italiano Di Medicina Del Lavoro Ed Ergonomia*, 34(3):352, 2012.

[10] Christina Unger, Lorenz Bühmann, Jens Lehmann, Axel-Cyrille Ngonga Ngomo, Daniel Gerber, and Philipp Cimiano. Template-based question answering over rdf data. In *Proceedings of the 21st International Conference on World Wide Web*, WWW '12, pages 639–648, New York, NY, USA, 2012. ACM.

[11] Philipp Unger, Christinai ; Cimiano. Pythia: Compositional meaning construction for ontology-based question answering on the semantic web. *Proc. 16th International Conference on Applications of Natural Language to Information Systems, NLDB 2011*, pages 153–160.

[12] Peng Zhou, Wei Shi, Jun Tian, Zhenyu Qi, Bingchen Li, Hongwei Hao, and Bo Xu. Attention-based bidirectional long short-term memory networks for relation classification. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 207–212, 2016.