

# Health-LLM: Personalized Retrieval-Augmented Disease Prediction System

Qinkai Yu<sup>5,\*</sup>, Mingyu Jin<sup>1,\*</sup>, Dong Shu<sup>3,\*</sup>, Chong Zhang<sup>2,\*</sup>, Lizhou Fan<sup>4</sup>, Wenyue Hua<sup>1</sup>,  
Suiyuan Zhu<sup>7</sup>, Yanda Meng<sup>5</sup>, Zhenting Wang<sup>1</sup>, Mengnan Du<sup>6</sup>, Yongfeng Zhang<sup>1</sup>

<sup>1</sup> Rutgers University, <sup>2</sup> University of Liverpool, <sup>3</sup> Northwestern University,  
<sup>4</sup> University of Michigan, <sup>5</sup> University of Exeter, <sup>6</sup> New Jersey Institute of Technology,  
<sup>7</sup> New York University

## Abstract

Recent advancements in artificial intelligence (AI), especially large language models (LLMs), have significantly advanced healthcare applications and demonstrated potential in intelligent medical treatment. To promote professional and personalized healthcare, we propose an innovative framework, **Health-LLM**, which combines large-scale feature extraction and trade-off scoring of medical knowledge. Compared to traditional health management applications, our system has three main advantages: (1) It integrates health reports and medical knowledge into a large model to ask relevant questions to the Large Language Model for disease prediction; (2) It leverages a retrieval augmented generation (RAG) mechanism to enhance feature extraction; (3) It incorporates a semi-automated feature updating framework that can merge and delete features to improve the accuracy of disease prediction. We experimented with a large number of health reports to assess the effectiveness of the Health-LLM system. The results indicate that the proposed system surpasses the existing ones and has the potential to advance disease prediction and personalized health management significantly.

## 1 Introduction

The integration of AI into healthcare, notably through large language models (LLMs) such as GPT-3.5 (Rasmy et al., 2021) and GPT-4 (Achiam et al., 2023), has reshaped the field of health management. Recent studies highlight the crucial role of LLM in using machine learning to improve healthcare outcomes (Biswas, 2023; Singhal et al., 2022). Advancements in AI for healthcare demonstrate a shift towards models that handle complex medical data and offer improved precision. Nonetheless, traditional health management methods often struggle with the constraints imposed by static data and uniform standards, making them unfit to fully meet individual needs (Uddin et al.,

2019; López-Martínez et al., 2020; Beam and Kohane, 2018; Ghassemi et al., 2020). The health reports of patients offer a wealth of data, this information has the potential to predict future health issues and tailor health recommendations, but the difficulty lies in transforming these extensive data into practical insights.

This study focuses on the Clinical Prediction with Large Language Models (CPLLM) approach (Shoham and Rappoport, 2023), which showcases the superior predictive capabilities of LLMs fine-tuned on clinical data. In particular, we propose an innovative system, **Health-LLM**, utilizing data analytics, machine learning, and medical knowledge for comprehensive health management. The system can provide users with personalized health recommendations based on predicted health risks (see Figure 1), ultimately helping prevent future health complications.

Specifically, the system uses the Llama Index (Liu, 2022) framework to analyze the information from the patient’s health report. Then it assigns different scores to these features by the Llama Index, which is prepared with professional medical information by RAG. The scoring method asks the language model questions about the patient’s condition. Our system also incorporates automated feature engineer technology (He et al., 2021) to perform iterative optimization to extract important features and stable weights and scores. Finally, the system is trained based on the XGBoost model (Chen and Guestrin, 2016) to make early predictions of existing diseases and provide personalized health recommendations to individuals.

We compare the performance of our system with traditional methods (including Pretrain-BERT (Devlin et al., 2018), TextCNN (Chen, 2015), Hierarchical Attention (Yang et al., 2016), Text BiLSTM with Attention (Liu and Guo, 2019), RoBERT (Liu et al., 2019)) as well as mainstream large-scale language models (GPT-3.5, GPT-4, LLaMA-2 (Tou-

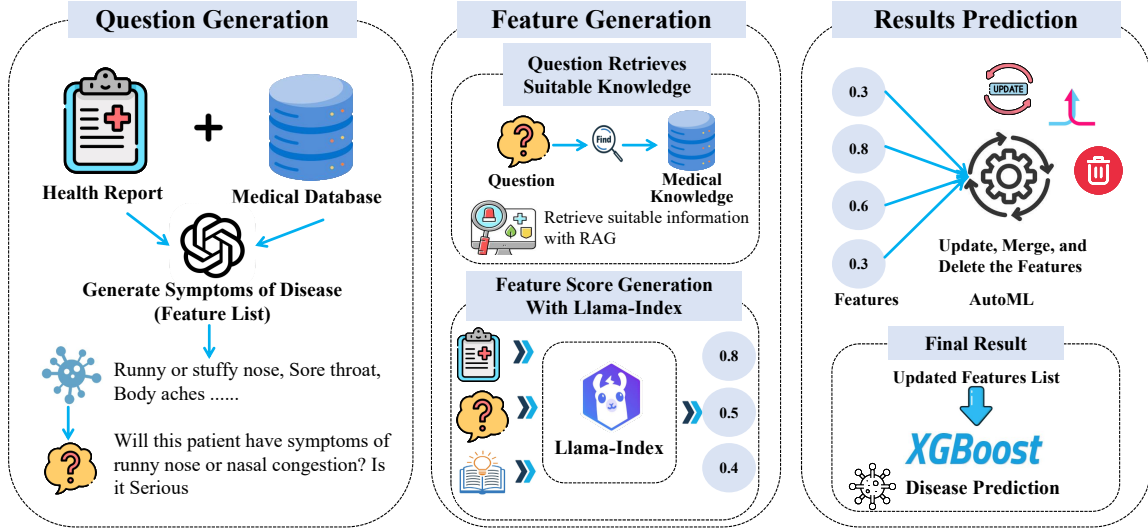


Figure 1: The overall workflow of the Health-LLM: from feature extraction to XGBoost prediction.

vron et al., 2023)) in three different settings (zero-shot, few-shot, and information retrieval) to show the effectiveness of our system. Among them, the accuracy of GPT-4 combined with information retrieval by retrieval augmented generation (RAG) for disease diagnosis is 0.68, and the F1 score is 0.71, while our system has achieved an accuracy of 0.833 and an F1 score of 0.762, respectively. Our key contributions are as follows:

- We propose an innovative Health-LLM framework that combines large-scale feature extraction, precise scoring of medical knowledge using the Llama Index structure, and machine learning techniques to enable personalized disease prediction from patient health reports.
- Our proposed Health-LLM framework achieves state-of-the-art performance on disease prediction tasks, surpassing existing methods like GPT-4 and fine-tuned LLaMA-2 models as demonstrated through extensive experiments.

## 2 Background

### 2.1 AI for Health Management

AI is revolutionizing healthcare through machine learning and relevant methods to enhance healthcare outcomes. This evolution is significantly driven by the emergence of LLMs, as seen in studies such as (Biswas, 2023; Singhal et al., 2022; Rasmy et al., 2021). These models are vital in clinical applications, including disease prediction and diagnosis. The intersection of AI and healthcare has seen notable progress, fueled by the availability of extensive health datasets and the advancement

of sophisticated LLMs. Recent research, such as (Wang et al., 2023), demonstrates the immense potential of LLMs in the healthcare sector, where they are used to understand and generate health reports and evaluate various health situations.

A key development in this field is the Clinical Prediction with Large Language Models (CPLLM), which highlights the potential of LLMs fine-tuned on clinical data (Shoham and Rappoport, 2023). CPLLM, using historical diagnosis records, has shown superiority over traditional models such as logistic regression and even advanced models such as Med-BERT (Rasmy et al., 2021) in predicting future disease diagnoses. Another significant advancement in AI for health is the COAD framework (Wang et al., 2023), which addresses the limitations of previous Transformer-based automatic diagnosis methods. AMIE (Articulate Medical Intelligence Explorer) is a medical knowledge graph created by Google (Tu et al., 2024). It extracts and stores medical knowledge, including information about diseases, symptoms, and treatments. These advancements indicate a trend in AI for health, shifting towards models that effectively manage the complexity and subtleties of medical data. The progress made by CPLLM, COAD and AMIE underscores the transformative impact these technologies can have on healthcare, enhancing precision, efficiency, and personalization in patient care.

### 2.2 Retrieval Augmented Generation

The retrieval augmented generation (RAG) (Lewis et al., 2020) method is a natural language processing model that combines retrieval and generation

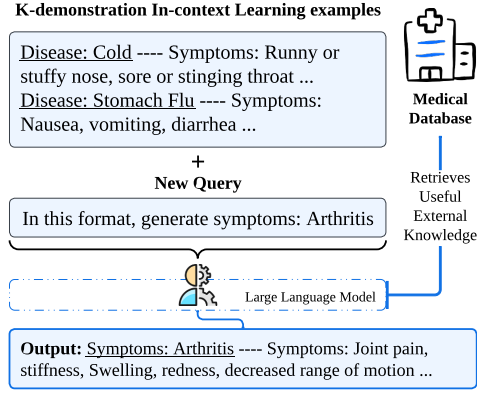


Figure 2: In-context learning workflow of symptom features generation in Health-LLM.

components to handle knowledge-intensive tasks. The method consists of two stages: Retrieval and Generation. During the retrieval phase, RAG employs the Dense Passage Retrieval (DPR) system to retrieve the most relevant documents from a large-scale document database that answers the input question. The input question is encoded as a vector, which is then compared with the document vectors in the database to locate the most relevant documents. The main idea behind this method is to use a large-scale document collection to enhance the generation model’s ability and improve the model’s efficiency in dealing with complex and knowledge-dependent problems.

### 3 Implementation

#### 3.1 In-context Learning for Symptom Features Generation

In the initial phase of our implementation, we systematically extract symptom features from a range of diseases by harnessing the in-context learning capabilities of LLMs. Figure 2 shows an example of the corresponding workflow. We first prompt the model with a series of examples, such as "disease: cold, symptoms: runny or stuffy nose, sore or tingling throat, cough, sneeze", to teach it the pattern generated by the symptom profile. Leveraging this in-context learning paradigm, our system is ready to take new query inputs and efficiently produce symptom descriptors for various diseases in a batch processing mode.

To enhance the precision of this generative process, we integrate a supplementary medical knowledge base, employing a RAG mechanism for enriched knowledge retrieval. Through the guidance of the system prompts, the LLM is asked to answer

questions about the extracted symptoms. Since LLM may lack specialized knowledge in medicine, we provide contextual information for these questions in an embedded form. Thus, we utilize advanced RAG technology to synchronize our queries with the knowledge base. In particular, RAG helps identify and retrieve the three most relevant pieces of information that align with the symptoms mentioned in the question. Our system then extracts these pieces of information and seamlessly integrates them to enrich the input prompts for the models.

#### 3.2 Assigning Score by Llama Index

To integrate LLMs from different sources, we adopt the Llama Index framework (Liu, 2022) for question answering (QA) in health reports. This approach allows us to take full advantage of advanced natural language processing models to extract features, a key step in our ability to predict disease and provide health advice.

To ensure accurate and domain-specific responses, we then prompt the system with our high-quality generated queries relevant to the health issue. The LLM then assigns a confidence score between 0 and 1, indicating the system’s perception of the health issue’s severity level. For instance, the question "Does this person have good sleeping habits?" might receive a response like "Sleep: 0.6," suggesting moderately positive sleeping patterns. This numerical confidence score then becomes a key attribute in our classification model.

The Llama index serves to streamline document-based QA through a strategic "search-then-synthesize" approach. The process unfolds as follows. Initially, health report documents are curated and formatted into plain text. These documents are then segmented into smaller manageable text blocks. Each block is processed through a text-embedding interface, transforming it into a vector representation that is subsequently stored within a vector database; here, OpenAI’s embeddings can be utilized for this transformation (OpenAI, 2021).

When receiving questions, the system converts them into vectors to facilitate search within the vector database, aiming to identify the most relevant text block(s). The identified text block is then amalgamated with the query to formulate a refined request. This newly created request is sent to the OpenAI API for processing. In total, we create a list of 152 questions and an area knowledge

database. All the scores resulting from answering these questions will be used as features to enter the downstream machine learning model. For example: Sleep: 0.6, which means that the person's sleep condition is okay.

### 3.3 Predictive Model and Health Advice

In our quest to develop a robust disease classification system, we have established a comprehensive framework that encompasses 61 disease labels, ranging from common ailments such as insomnia and indigestion to more complex conditions like endocrine disorders. We fit the features with XGBoost, which allows the XGBoost model to fit the feature scores extracted from the Llama index, and learn the feature representation of each disease under the Llama index. The results of XGBoost fitting in this framework are multiple binary combinations, with "0" indicating no associated disease and "1" indicating an associated disease. In addition, we have also classified certain diseases (e.g., fatty liver) at a finer level of granularity. For example, "0" indicates mild fatty liver and "1" indicates severe fatty liver.

The significance of domain-specific knowledge incorporation has become clear in our quest to improve disease prediction through feature preprocessing. Addressing this, we use Context-Aware Automated Feature Engineering (CAAFE), which utilizes the power of LLMs (Hollmann et al., 2023) to generate features iteratively with semantic relevance informed by the dataset's context. We have evolved from a semi-automated system to a fully automated one. We now use LLMs to autonomously craft feature and dataset descriptions, thereby streamlining the feature engineering process and enriching our models with contextually meaningful data. After the prediction is completed, we will query the LLM for the last time with the predicted diseases, and let it generate targeted health suggestions based on the health report and professional knowledge.

### 3.4 Interact with Health-LLM

Users have two ways to interact with the Health-LLM system. First, the user can submit his/her health report to the system and get the prediction and the health advice. Second, the user can describe his/her health problem and ask the system, the system will record their dialog and save as a health dialog. The dialog messages are similar to the health reports in the training process. Health-

LLM generates the corresponding disease predictions and gives personal suggestions based on the dialog messages following the process shown in Figure 1. The large language model responsible for the dialog and giving personal suggestions in Health-LLM is GPT-4 turbo. In brief, Health-LLM provides diagnosis and personalized advice through interaction with the user.

## 4 A Case Study

In this section, we describe the entire process by which our system accurately predicts a patient's disease and compare its performance with, for example, using a large model alone. Our case is under dataset IMCS-21 (Table 1) (Chen et al., 2023). **Settings.** To our knowledge, IMCS-21 is currently the only dataset on telemedicine consultation in China. The data set contains a total of 4,116 annotated samples with 164,731 utterances across 10 pediatric diseases, including *bronchitis*, *fever*, *diarrhoea*, *upper respiratory infection*, *dyspepsia*, *cold*, *cough*, *jaundice*, *constipation* and *bronchopneumonia*. Each dialogue contains an average of 40 utterances, 523 Chinese characters (580 characters if including self-report) and 26 entities. To prevent data leakage, we pre-processed the data by using GPT4 to eliminate content containing specific medical conditions from the data conversations and transformed the conversations into the form of an Electronic Patient Record (EPR). We use two metrics: ACC(Accuracy), F1(Macro F1-score), to evaluate the prediction quality.

### 4.1 Health-LLM Diagnostic Test and Comparative Experiment

**Health-LLM Diagnostic Test.** Trained on the given dataset IMCS-21, our system is capable of taking complex patient inputs and provide medical suggestions. We provide an example of the user-system interaction in Table 4 in the Appendix. To employ this system, the user inputs their information such as physical conditions and disease symptoms. The system predicts potential diseases for the user based on prior experience and provides the appropriate diagnosis, as well as follow-up recommendations. We give an example in Table 1. Our system plays the role of a doctor who guides the user step by step in describing his or her real symptoms. The output of the system is the specific name of the disease and reasonable suggestions. Our validation process uses conversations from the



Table 1: An example of the interactive process in case study.

User Inputs and Conversation with Health-LLM	
User:	In the morning, the stools are normal and thin, and they are a little beaded at night.
Health-LLM:	Hello, how long has this situation? Is the diet regular? Can you provide me with basic information such as your age and gender? What is the feeling of abdomen?
User:	It has been four or five days. I have a stool when I ate Amoshilin in the first rhinitis. After a few days of laxative spirit, it is like this now. Haven't eaten food and cold.
Prediction and Personal Advice	
Health-LLM:	Gastrointestinal dysfunction and Diarrhea. Choose light meals and avoid stimulating foods. You don't need to take medicine. You are welcome. Wish you health and happiness!

Table 2: Comparing with existing methods on the Health-LLM diagnostic test.

Models	Accuracy	F1
Pretrain-Bert with XGBoost	0.326	0.303
GPT-3.5(zero-shot)	0.333	0.361
GPT-3.5(few shot)	0.381	0.349
GPT-4(zero-shot)	0.390	0.312
TextCNN	0.437	0.429
GPT-3.5(few shot with information retrieval)	0.451	0.451
Hierarchical Attention	0.495	0.477
Text BiLSTM with Attention	0.512	0.500
RoBERT	0.585	0.543
GPT-4(few shot)	0.620	0.671
GPT-4(few shot with information retrieval)	0.680	0.718
Fintuned-LLaMA-2-7B	0.710	0.593
Fintuned-LLaMA-2-13B	0.730	0.671
<b>Health-LLM (Ours)</b>	<b>0.833</b>	<b>0.762</b>

validation set as input to simulate user adoption.

**Comparative Experiment.** Our comparative analysis involves two open-source models, GPT-3.5 and GPT-4. We deploy these models in a zero-shot setting to diagnose cases. Following this, we will transition to a few-shot context, providing each model with a selection of prior predictions to enhance its diagnostic capabilities. Subsequently, we will integrate supplementary medical knowledge into the larger models to refine their predictions further. The culmination of this process will involve leveraging the dataset we use in Health-LLM to fine-tune LLaMA-2 (Touvron et al., 2023), with the objective of observing the impact on its performance. For extended texts and chapter-level analysis in the traditional method, researchers have introduced a framework known as Hierarchical Attention (Yang et al., 2016) and RoBERTa (Liu et al., 2019), which excels in handling multi-classification challenges in case texts. The conventional approach fails in accurately diagnosing diseases from medical reports, as it lacks the deep text comprehension exhibited by larger models. Additionally, it struggles to encapsulate all the pertinent information within lengthy paragraphs. This approach ranks impressively, just behind advanced models like GPT-4 and fine-tuned Llama-2 in terms of performance. The comparative test results are shown in Table 2.

Experiments showed that our system achieved an overall prediction accuracy of 83% for the disease, with F1 score of 0.762. Health-LLM achieves the best performance compared to direct inference using other large models. **Ablation Study.** In this section, we verify the effectiveness of each component of our Health-LLM. The results are displayed in Table 3. The experimental results have two implications. First, it is necessary to index professional healthcare data to ensure the accuracy of diagnostic reasoning. Our diagnostic reasoning accuracy has improved significantly in the group with indexed professional healthcare knowledge. Second, our data processing and feature extraction with CAAFE is effective.

Table 3: The result of Health-LLM diagnostic test on IMCS-21 dataset.

Model	Acc	F1
Health-LLM without Retrieval	0.78	0.714
Health-LLM without CAAFE	0.77	0.721
Health-LLM	0.83	0.762

## 5 Conclusions

We present a novel system, Health-LLM, that combines large-scale feature extraction, precise scoring of medical knowledge, and machine learning techniques to make better use of patient health reports. It improves the prediction of potential diseases compared to GPT-3.5, GPT-4, and fine-tuned LLaMA-2. Our system is capable to predict potential diseases and provide customized health advice to individuals. In addition to its practical advancements, Health-LLM also serves as a proof-of-concept LLM-based system in healthcare, highlighting the huge potential of further development of AI applications in the health field.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Andrew L Beam and Isaac S Kohane. 2018. Big data and machine learning in health care. *Jama*, 319(13):1317–1318.
- Som S Biswas. 2023. Role of chat gpt in public health. *Annals of biomedical engineering*, 51(5):868–869.
- Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794.
- Wei Chen, Zhiwei Li, Hongyi Fang, Qianyu Yao, Cheng Zhong, Jianye Hao, Qi Zhang, Xuanjing Huang, Jiajie Peng, and Zhongyu Wei. 2023. A benchmark for automatic medical consultation system: frameworks, tasks and datasets. *Bioinformatics*, 39(1):btac817.
- Yahui Chen. 2015. Convolutional neural network for sentence classification. Master’s thesis, University of Waterloo.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Marzyeh Ghassemi, Tristan Naumann, Peter Schulam, Andrew L Beam, Irene Y Chen, and Rajesh Ranganath. 2020. A review of challenges and opportunities in machine learning for health. *AMIA Summits on Translational Science Proceedings*, 2020:191.
- Xin He, Kaiyong Zhao, and Xiaowen Chu. 2021. Autotml: A survey of the state-of-the-art. *Knowledge-Based Systems*, 212:106622.
- Noah Hollmann, Samuel Müller, and Frank Hutter. 2023. Gpt for semi-automated data science: Introducing caafe for context-aware automated feature engineering. *arXiv preprint arXiv:2305.03403*.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Gang Liu and Jiabao Guo. 2019. Bidirectional lstm with attention mechanism and convolutional layer for text classification. *Neurocomputing*, 337:325–338.
- Jerry Liu. 2022. Llamaindex, 11 2022. URL [https://github.com/jerryliu/llama\\_index](https://github.com/jerryliu/llama_index). (Cited on page 8).
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. *Roberta: A robustly optimized bert pretraining approach*.
- Fernando López-Martínez, Edward Rolando Núñez-Valdez, Vicente García-Díaz, and Zoran Bursac. 2020. A case study for a big data and machine learning platform to improve medical decision support in population health management. *Algorithms*, 13(4):102.
- OpenAI. 2021. *Embedding models*.
- Laila Rasmy, Yang Xiang, Ziqian Xie, Cui Tao, and Degui Zhi. 2021. Med-bert: pretrained contextualized embeddings on large-scale structured electronic health records for disease prediction. *NPJ digital medicine*, 4(1):86.
- Ofir Ben Shoham and Nadav Rappoport. 2023. *Cpllm: Clinical prediction with large language models*.
- Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, et al. 2022. Large language models encode clinical knowledge. *arXiv preprint arXiv:2212.13138*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Tao Tu, Anil Palepu, Mike Schaeckermann, Khaled Saab, Jan Freyberg, Ryutaro Tanno, Amy Wang, Brenna Li, Mohamed Amin, Nenad Tomasev, et al. 2024. Towards conversational diagnostic ai. *arXiv preprint arXiv:2401.05654*.
- Shahadat Uddin, Arif Khan, Md Ekramul Hossain, and Mohammad Ali Moni. 2019. Comparing different supervised machine learning algorithms for disease prediction. *BMC medical informatics and decision making*, 19(1):1–16.
- Huimin Wang, Wai-Chung Kwan, Kam-Fai Wong, and Yefeng Zheng. 2023. Coad: Automatic diagnosis through symptom and disease collaborative generation. *arXiv preprint arXiv:2307.08290*.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. Hierarchical attention networks for document classification. In *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: human language technologies*, pages 1480–1489.

## Appendix

Table 4: Example of one health report we have made by dataset IMCS-21

IMCS-21
-Hello, there is a pain around the navel, I don't know what's going on (female, 29 years old)
-Hello, how long has this situation?
-Two or three days.
-It hurts, and it will not hurt for a while.
-Have you used any medicine? Have you ever done?
-There is no medication and no examination.
-Is the stool normal?
-normal.
-Are there any other symptoms? Do you want to vomit?
-No.
-Is it faint pain?
-Once appetite, a little bit bloated.
-It may be gastrointestinal dysfunction.
-Yes, faint pain.
-Eat some song Meibing Try.relative.
-It felt like a needle was tied, and it was fine for a few seconds.
-It feels that the problem is not particularly big.
-Try it if you take the medicine I said.
-Is there a compound fairy crane grassyitis tablet at home?
-This is mainly to treat diarrhea. Don't eat without diarrhea. Oh, alright.
-Uh-huh.
-Which aspect can cause gastrointestinal function?
-It may be that the diet may be a mental factor or the autoimmune system.
-It may also be the digestive system or autoimmune system. relative.
-OK
-Uh-huh.
-It seemed a bit like a diarrhea. After eating 1 at noon, I wanted to pull it after a while, and I was a little bit pulled.
-You can eat the medicine you said.
-Oh well.
-Uh-huh.