Proceedings of Machine Learning for Healthcare 2017 JMLR W&C Track Volume 68

Diagnostic Inferencing via Improving Clinical Concept Extraction with Deep Reinforcement Learning: A Preliminary Study

Yuan Ling YUAN.LING@PHILIPS.COM
Sadid A. Hasan SADID.HASAN@PHILIPS.COM
Vivek Datla VIVEK.DATLA@PHILIPS.COM
Ashequl Qadir ASHEQUL.QADIR@PHILIPS.COM
Kathy Lee KATHY.LEE_1@PHILIPS.COM
Joey Liu JOEY.LIU@PHILIPS.COM
Oladimeji Farri DIMEJI.FARRI@PHILIPS.COM
Artificial Intelligence Laboratory, Philips Research North America
Cambridge, MA, USA

Abstract

Clinical diagnostic inferencing is a complex task, which often requires significant medical research and investigation based on an underlying clinical scenario. This paper proposes a novel approach by formulating the task as a reinforcement learning problem such that the system can infer the most probable diagnoses by optimizing clinical concept extraction from a free text case narrative via leveraging relevant external evidence. Such a formulation is deemed to be suitable due to the inherent complexity of the task and unavailability of sufficient annotated data. During training, the agent tries to learn the optimal policy through iterative search and consolidation of the most relevant clinical concepts that best describe a correct diagnosis. A deep Q-network architecture is trained to optimize a reward function that measures the accuracy of the candidate diagnoses and clinical concepts. Our preliminary experiments on the TREC CDS dataset demonstrate the effectiveness of our system over non-reinforcement learning-based strong baselines.

1. Introduction

Clinical diagnostic inferencing is a challenging task. For example, given a clinical case (past medical history, current signs and symptoms etc.), the clinician administers appropriate medical tests or procedures, infers the accurate diagnosis, and prescribes the best-possible treatment plan based on his/her experience or up-to-date knowledge/evidence obtained through substantial research on relevant external resources (Pelaccia et al., 2011; Kushniruk, 2001; Norman et al., 2007). In routine and emergency clinical practice, summaries of patient history, physical findings and lab results are assimilated by a typically multi-tasking clinician in a high-intensity environment. This cognitive activity is necessary for the clinician to arrive at his/her overall assessment of the patient condition towards administering appropriate treatment (Reilly et al., 2013).

This paper considers the challenge of inferring the diagnoses of a patient condition from a given medical case narrative. Our work mainly focuses on reducing the cognitive burden of assimilating the complex and diverse information in clinical reports within the Electronic Health Record (EHR) towards supporting prompt and accurate decision-making for a described patient scenario. We envisage that this work would lead to the busy clinician considering relevant differential diagnoses that could otherwise be ignored due to inadvertent diagnostic errors (Nendaz and Perrier, 2012; Graber et al., 2012; Berge and Mamede, 2013). In addition, relatively lower-skilled clinicians e.g. nurse practitioners can use the proposed system as a source of second opinion before contacting a physician towards accurately diagnosing and managing their patients.

Prior works that build Artificial Intelligence (AI) systems to support clinical decision making, mostly use structured clinical data (e.g. physiological signals, vital signs, lab tests etc.) stored in the EHR and follow a supervised approach to predict the probable diagnoses (Lipton et al., 2015; Choi et al., 2015, 2016). In contrast, we explore the discriminatory capability of the unstructured free text clinical narratives to infer the possible diagnoses based on an underlying clinical scenario. In recent Text REtrieval Conference (TREC) Clinical Decision Support (CDS) tracks[1], clinical diagnosis inferencing from free text clinical narratives has been showcased as a significant milestone in clinical question answering and a path to improving the accuracy of relevant biomedical article retrieval (Roberts et al., 2015, 2016; Goodwin and Harabagiu, 2016).

The efficacy of existing supervised models largely depends on the size of the annotated datasets used for training. Creation of labeled datasets requires expert-derived annotations, which are typically very expensive to obtain. Moreover, these models lack the ability to capture the underlying uncertainties related to generating differential diagnoses (Richardson et al., 1999) and linguistic complexities (Seidel et al., 2015) of a clinical scenario as they consider medical codes and a finite number of diagnoses for prediction labels. To overcome these issues, we formulate the diagnostic inferencing problem as a sequential decision making process using deep reinforcement learning by leveraging a small amount of labeled data with appropriate evidence garnered from relevant external resources.

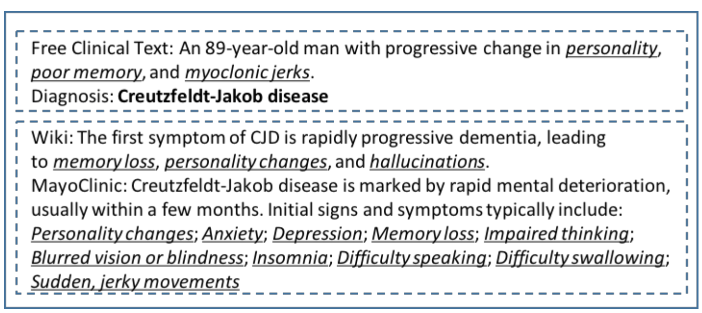

Figure 1: Evidences to refine clinical concepts.

Extracting appropriate clinical concepts from free clinical text is a critical first step for diagnosis inference. Existing clinical concept extraction tools mainly depend on large corpora of labeled examples and knowledge bases (Jonnalagadda et al., 2012; Kim et al., 2015; Chalapathy et al., 2016). Moreover, they are limited to the original content of the text

1. http://www.trec-cds.org/

as they do not consider evidence from external free text resources. Hence, clinical concepts extracted by these tools often lack aspects related to in-domain normalization, which may have a negative impact on the downstream clinical diagnostic inferencing task. For example, in Figure 1, "*personality*", "*poor memory*", and "*myoclonic jerks*" are the clinical concepts extracted from the free text clinical narrative. However, since concepts may be represented using different word-phrase combinations in the literature, without further processing of these concepts e.g. paraphrasing, it is difficult to map them to well-known concepts in a knowledge base and infer the correct diagnosis.

External resources such as relevant pages from Wikipedia and MayoClinic (see details in Section 4) can serve as the evidence to improve the original extracted concepts using one of the following ways: mapping of incomplete concepts to corresponding expressive concepts e.g. *personality* → *personality changes*, paraphrasing the concepts e.g. *poor memory* → *memory loss*, and supplementing with additional concepts e.g. *hallucinations* can be added to the relevant concept pool from the Wikipedia page.

In this paper, we propose a novel clinical diagnosis inferencing approach that uses a deep reinforcement learning technique to incrementally learn about the most appropriate clinical concepts that best describe the correct diagnosis by using evidences gathered from relevant external resources. During training, the agent tries to learn the optimal policy through iterative search and consolidation of the most relevant clinical concepts related to the given patient condition. A deep Q-network architecture (Mnih et al., 2015) is trained to optimize a reward function that measures the accuracy of the candidate diagnoses and clinical concepts. Our preliminary experiments on the TREC CDS dataset (Roberts et al., 2015) demonstrate the effectiveness of our system over various non-reinforcement learning-based baselines.

In the next sections, we present related work in the literature, describe the datasets, external resources, and the proposed clinical diagnosis inferencing approach, followed by the detailed experimental setup, results and discussion. Finally, we conclude the paper in Section 7.

2. Related Work

Addressing inference tasks generally requires significant contributions from domain experts and access to a variety of resources (Ferrucci et al., 2013; Lally et al., 2014) e.g. structured knowledge bases (KBs) (Yao and Van Durme, 2014; Bao et al., 2014; Dong et al., 2015). However, KBs have certain limitations such as knowledge incompleteness, sparsity, and fixed schema (Socher et al., 2013; West et al., 2014; Bordes et al., 2014), which motivate researchers to use unstructured textual resources like Wikipedia for various related tasks (Katz et al., 2005; Wu and Weld, 2010; Miller et al., 2016; Chen et al., 2017). In this paper, we also leverage the power of unstructured knowledge sources to address clinical diagnosis inferencing.

Previous research on clinical diagnosis inferencing mostly utilize structured clinical data e.g. physiological signals, vital signs, lab tests, and other variables (Lipton et al., 2015; Choi et al., 2015, 2016). EHRs typically store such data along with unstructured text documents that contain a relatively more complete picture of associated clinical events.

Recently, diagnosis inferencing from unstructured clinical text has gained much attention among AI researchers, mainly motivated by the TREC CDS tracks (Simpson et al., 2014; Roberts et al., 2015, 2016; Goodwin and Harabagiu, 2016; Zheng and Wan, 2016; Balaneshin-kordan and Kotov, 2016; Prakash et al., 2017; Ling et al., 2017). Although the main task in the CDS track was to retrieve relevant biomedical articles given a clinical scenario, researchers explored the task of diagnosis inferencing from clinical narratives, especially inspired by the pilot task in 2015[2] that investigated the impact of diagnostic information on retrieving relevant biomedical articles (Roberts et al., 2015, 2016).

Existing works on diagnostic inferencing mostly propose supervised classification models using various neural network architectures (Lipton et al., 2015; Choi et al., 2015; Prakash et al., 2017). However, such models heavily rely on large labeled data, and lack the ability to capture inherent ambiguities and complexities of a clinical scenario. Moreover, they are limited by the number of diagnosis labels and the use of medical codes to simplify the computational and linguistic difficulties of a clinical case. In contrast, we propose a novel approach for clinical diagnosis inferencing that formulates the task as a reinforcement learning problem to alleviate such issues. Other works have explored graph-based reasoning methods to incorporate relevant medical concepts and their associations (Shi et al., 2017; Geng and Zhang, 2014; Goodwin and Harabagiu, 2016; Zheng and Wan, 2016; Ling et al., 2017).

There exist works that use reinforcement learning for related clinical decision support tasks by mainly focusing on other modalities e.g. medical imaging (Netto et al., 2008) or specific domain-dependent use cases and clinical trials (Poolla, 2003; Shortreed et al., 2011; Zhao et al., 2011), but not for inferencing diagnosis. Recent works have shown the utility of deep reinforcement learning techniques for challenging tasks like playing games and entity extraction via utilizing external evidence (Mnih et al., 2015; Narasimhan et al., 2015, 2016). To the best of our knowledge, our work is the first to explore deep reinforcement learning for clinical diagnosis inference using unstructured text data from EHR.

3. Datasets

3.1 TREC CDS

We use the 2015 TREC CDS track dataset (Roberts et al., 2015) to conduct our experiments. This dataset contains 30 topics, where each topic is a medical case narrative that describes a patient scenario. Each topic contains "*description*", "*summary*", and "*diagnosis*" fields. "*description*" includes a comprehensive description of the patient's situation, whereas "*summary*" contains an abridged version of the most important information. The associated ground truth "*diagnosis*" is used to compute the goal-oriented reward for the agent. A topic example is partially shown in Table 1. We use the first 20 topics ("summary") for training and the rest for testing our system.

2. http://www.trec-cds.org/2015.html

Clinical Narratives (partially shown)
Description:
An 89-year-old man was brought to the emergency department by his wife and son after six months of progressive changes in cognition and personality. He began to have poor memory, difficulty expressing himself, and exhibited unusual behaviors, such as ...
Summary:
An 89-year-old man with progressive change in personality, poor memory, and myoclonic jerks.
Diagnosis:
Creutzfeldt-Jakob disease

Table 1: An example *TREC CDS* 2015 topic.

3.2 HumanDx

HumanDx[3] is a project to foster integrating efforts to map health problems to their possible diagnoses. We curate a list of diagnosis-findings relationships from HumanDx and create a dataset with entries related to 30 diagnoses (based on the TREC CDS dataset) with their associated findings (i.e. signs and symptoms). This dataset is used as the ground truth clinical concepts associated with each diagnosis (from TREC CDS), which are utilized to compute the intermediate (concept-level) rewards for the agent. An example diagnosis-symptoms pair is shown in Table 2.

HumanDx Example (partially shown)
Related Findings (Symptom Concepts):
agitation, fecal incontinence, costa rica, focal weakness, digital rectal examination, disorientation, ataxia, alert and oriented, ...
Diagnosis:
Creutzfeldt-Jakob disease

Table 2: An example diagnosis-symptoms pair from *HumanDx*.

4. External Sources for Evidence

Our work relies on external knowledge sources to provide candidate diagnoses for the sentences from a clinical narrative. We use two external knowledge sources: Wikipedia pages and MayoClinic pages. We index Wikipedia and MayoClinic using Elasticsearch[4]. As an example, Wikipedia and MayoClinic pages for the diagnosis "Creutzfeldt-Jakob disease" are partially displayed in Table 3.

4.1 Wikipedia

We select 37,245 Wikipedia pages under the clinical medicine category[5] in this study. Each page title is used as the diagnosis name and the texts from the *Signs and Symptoms* subsection are used as an evidence for relevant clinical concepts during the reinforcement learning process. As shown in Table 3, "*Signs and Symptoms*" section describes symptoms

3. https://www.humandx.org/
4. https://www.elastic.co/
5. https://en.wikipedia.org/wiki/Category:Clinical_medicine

External Knowledge Sources (partially shown)
Wikipedia - "Signs and Symptoms" Section:
The first symptom of CJD is usually rapidly progressive dementia, leading to memory loss, personality changes, and hallucinations. Other frequently occurring features include anxiety, depression, paranoia, obsessive-compulsive symptoms, and psychosis ...
MayoClinic - "Symptoms" Section:
Creutzfeldt-Jakob disease is marked by rapid mental deterioration, usually within a few months. Initial signs and symptoms typically include: Personality changes, Anxiety, Depression, Memory loss, Impaired thinking, Blurred vision or blindness, Insomnia, Difficulty speaking, Difficulty swallowing, Sudden, jerky movements ...
Diagnosis:
Creutzfeldt-Jakob disease

Table 3: *Wikipedia* and *MayoClinic* pages for "Creutzfeldt-Jakob disease".

of "Creutzfeldt-Jakob disease". These symptoms have a higher chance of appearing in a clinical narrative if the documented diagnosis is "Creutzfeldt-Jakob disease".

4.2 MayoClinic

The MayoClinic[6] disease corpus contains 1,117 pages, which include sections of Symptoms, Causes, Risk Factors, Treatments and Drugs, Prevention, etc. Each MayoClinic page title is regarded as a diagnosis. We select sentences from the *Symptoms* section as the external source of evidence for reinforcement learning. The MayoClinic corpus is also used as the diagnoses set for computing goal-oriented reward and retrieving final diagnosis results.

5. Reinforcement Learning Framework

We model the integration of external evidences for clinical diagnostic inferencing as a Markov Decision Process (MDP) (Bellman, 1957; Sutton and Barto, 1998). At each MDP step, one external article is considered from the evidence pool based on the current clinical concepts, where the evidence pool contains knowledge sources, such as Wikipedia articles (details in Section 4). Then, a new state s is generated that essentially comprises the agent's confidence on the current clinical concepts in describing a target diagnosis. In a state s, the agent takes an action a to get to the next state, $s' = s + a$. A reward function $r(s, a)$ is used to estimate the reward at each state s after taking an action a. We formulate our problem by estimating a state-action value function $Q(s, a)$, which determines the optimal action a to take in a state s using the Q-learning technique (Watkins and Dayan, 1992). The Q-function is approximated using a deep Q-network (DQN) architecture (Mnih et al., 2015). The trained DQN agent takes state s and reward r as input, and outputs an action a. Our reinforcement learning framework is shown in Figure 2. The algorithm of the MDP framework for clinical diagnosis inferencing is presented in Algorithm 1. For the training phase, the steps in the algorithm for each clinical narrative are run for multiple epochs. During the testing stage, each clinical narrative is processed only once in a single epoch.

6. http://www.mayoclinic.org/diseases-conditions

The next subsections provide details on the state, actions, and the reward function of the MDP framework.

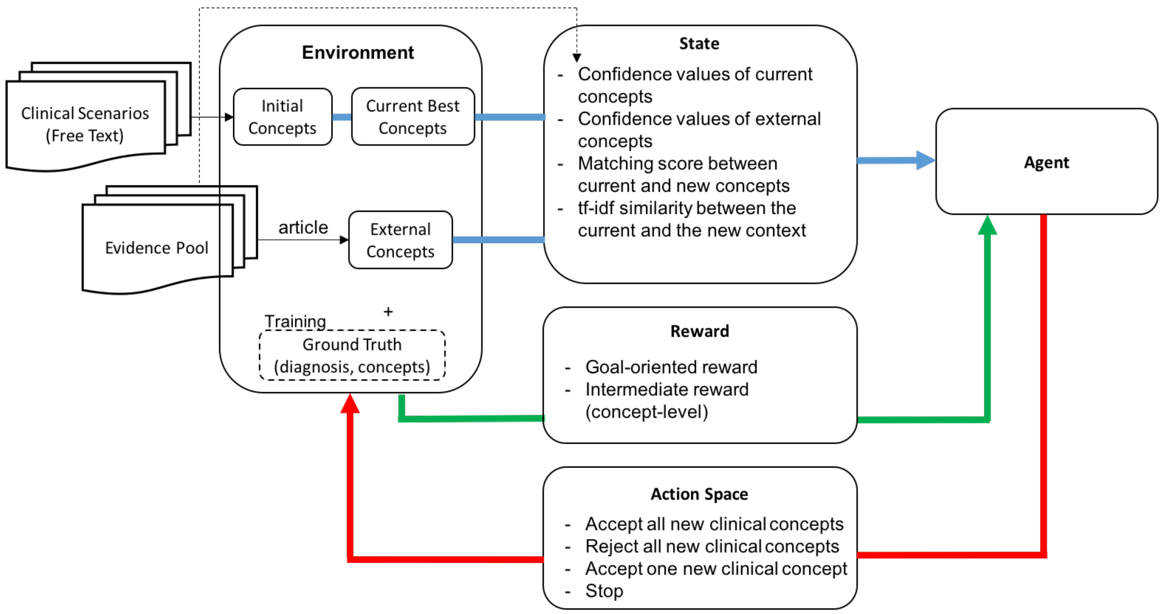

Figure 2: Reinforcement learning framework.

Algorithm 1: MDP framework

Input : Clinical narratives $N = \{N_1, N_2, ..., N_n\}$
Output: Improved clinical concepts C_i^{cur} towards correct diagnosis for each N_i

for *each clinical narrative* N_i *in* N **do**
    Retrieve top K Wikipedia articles Y_0;
    Retrieve top K MayoClinic articles Y_1;
end
for N_i *in* N **do**
    Extract entities C from N_i;
    $C^{cur} \leftarrow C$;
    $q \leftarrow 0$, $r \leftarrow 0$ //query type, reward;
    for $i = 1, K$ **do**
        Pop next article y from Y_q;
        Extract entities C^{new} from Y_q;
        Generate state vector v based on comparison between C^{cur} and C^{new};
        Send (v, r) to DQN agent, and get actions q and d from agent;
        if action $d ==$ "stop" **then break**;
        Update C^{cur} according to d;
        Calculate reward value r;
    end
    Send (v, r) to DQN agent.
end

5.1 Environment, State & Actions

Given a clinical narrative and a collection of external resources, the task is to find a set of appropriate clinical concepts and a correct diagnosis. We use MetaMap (Aronson, 2006) to extract the basic clinical concepts from clinical narratives and external articles. Since MetaMap returns various types of concepts, we only select symptom related concepts (Sondhi et al., 2012). Our evidence pool is built by querying articles from the Wikipedia and MayoClinic corpora (details in Section 4). For a given clinical scenario, we extract a set of initial concepts, which are then used as queries to search for the relevant list of external articles.

We consider each state s as a continuous real-valued vector comprising a set of current clinical concepts $C^{cur} = \{c_1^{cur}, \ldots, c_N^{cur}\}$ and a set of new concepts $C^{new} = \{c_1^{new}, \ldots, c_N^{new}\}$ extracted from an external article. The current and new clinical concepts are aligned as: 1) exact match between current and new concepts, 2) high similarity match based on a clinical phrase embedding-based similarity score, and 3) direct alignment to a current null value. The state vector essentially represents the similarity between the current clinical concepts set C^{cur} and the new concept set C^{new} including their context in the external article. The state vector encodes information about the current and new concepts as follows (Narasimhan et al., 2016): 1) confidence values of current concepts, 2) confidence values of external concepts, 3) matching scores between current and new concepts calculated by exploiting a phrase embedding model, and 4) tf-idf similarity between the current and the new context (window size = 5).

At each state, the DQN agent outputs an action from a set of actions: 1) accept all new clinical concepts, 2) reject all new clinical concepts; 3) accept one new clinical concept, and 4) stop.

5.2 Reward Function

Our reward function considers two types of rewards: 1) goal-oriented reward, r_{goal} that measures the quality of diagnoses outputs, and 2) intermediate (concept-level) reward, r_{inter} that evaluates the accuracy of extracted concepts. The combined reward function is denoted by:

$$r = w \times r_{goal} + (1 - w) \times r_{inter}, \tag{1}$$

where w is the weight factor empirically set to assign importance to a certain type of reward. We set $w = 0.9$ for our experiments where both rewards are used. A set of clinical concepts C is used as queries to search the MayoClinic corpus, and retrieve a ranked list of candidate diagnoses. According to the rank of the correct diagnosis among the candidate list, we get a Mean Reciprocal Rank (MRR) score. r_{goal} is computed by: $r_{goal} = MRR_{C^{new}} - MRR_{C^{cur}}$, where $MRR_C = \frac{1}{rank_C}$, and $rank_C$ represents the rank of the correct diagnosis among the list of candidate diagnoses, where the concept set C is used as the search queries.

We use HumanDx as our external source to compute the intermediate reward, r_{inter}. We use the correct diagnosis to search HumanDx and get a list of associated clinical concepts. For a clinical concept set C, we count the number of concepts (N_C) occurred in the HumanDx concept set. r_{inter} is computed as: $r_{inter} = N_{C^{new}} - N_{C^{cur}}$.

5.3 DQN Architecture

In order to learn the Q-value, the iterative updates are derived from the Bellman equationSutton and Barto (1998):

$$Q_{i+1}(s,a) = E[r + \gamma max_{a'} Q_i(s',a')|s,a], \tag{2}$$

where γ is a discount factor for the future rewards and the expectation is over the whole training process. It is impractical to maintain the Q-values for all possible state-action pairs. Mnih et al. (2015) proposed a deep Q-network (DQN) architecture, which approximates the Q-value function and predicts $Q(s,a)$ for all possible actions. We follow the DQN architecture in Narasimhan et al. (2015) to fit our problem formulation.

6. Experimental Setup

6.1 Evaluation Metrics

We use MRR, Recall at 5 (R@5) and R@10 to evaluate our models. MRR is computed by: $MRR = \frac{1}{|Q|}\sum_{i=1}^{|Q|}\frac{1}{rank_i}$, where $|Q|$ is the total number of topics, and $rank_i$ is the rank of the correct diagnosis for the i^{th} topic. R@5 is calculated as the number of topics with a correct diagnosis ranked within the top 5 diagnoses, divided by the total number of topics. R@10 is computed similarly.

6.2 Baselines

6.2.1 Baseline1

Our first baseline is a concept graph-based method proposed by Ling et al. (2017), which leverages external evidence without the reinforcement learning framework. It includes four steps: 1) extracting source concepts from the clinical narratives using MetaMap, 2) iteratively finding relevant evidence concepts (candidate diagnoses) from a structured knowledge base, UMLS (Bodenreider, 2004) and unstructured corpora (Wikipedia, MayoClinic, and both), 3) encoding both source and evidence concepts in a weighted graph via a regularizer-enhanced skip-gram model (Bordes et al., 2014; Mikolov et al., 2013a,b), and 4) ranking the relevant evidence concepts based on their association with the source concepts.

6.2.2 Baseline2

This method directly uses the clinical concepts extracted by MetaMap to query the external corpora to retrieve the correct diagnoses.

6.3 Results and Analyses

Table 4 presents the evaluation results. There are five parts in the results: 1) *Baseline1:* This part shows how external evidence can have an impact on the task without the reinforcement learning framework. We show the results for using different knowledge sources: Wikipedia, MayoClinic, and both, 2) *Baseline2:* This method (*Without RL*) does not use reinforcement learning and directly uses clinical concepts to retrieve the correct diagnoses, 3) *RL-Concept Variations:* Here, *Description* refers to the system that uses case narratives

	MRR	R@5	R@10		MRR	R@5	R@10
Baseline1				RL - *Evidence Variations*			
Wikipedia	26.01	44.78	44.78	Wikipedia	49.55	**70.00**	**90.00**
MayoClinic	32.64	47.62	51.00	MayoClinic	**51.45**	40.00	80.00
Both	32.29	48.00	48.00	Both	45.59	40.00	70.00
Baseline2				RL - *Reward Variations*			
Without RL	59.06	50.00	60.00	Goal-oriented	**53.26**	**60.00**	**80.00**
RL - *Concept Variations*				Intermediate	36.43	40.00	60.00
Description	52.90	50.00	50.00	Both	45.59	40.00	70.00

Table 4: Evaluation results.

from the available topic descriptions as input to the MetaMap concept extraction module (Roberts et al., 2015), 4) *RL-Evidence Variations:* Here, we compare the results of using different external sources inside the reinforcement learning framework, and 5) *RL-Reward Variations:* Here, we compare the results of using different reward functions.

Generally, our system has better performance than the baseline systems. We can see that by only using the Wikipedia source, baseline1 has the worst result. On the other hand, using MayoClinic yields better results. For the RL-variations experiments, we see that the Description systems introduce some noises to yield worse results. So, for the later variation experiments (part 4 and part 5), we only use the case summaries from the topic narratives. For other experiments with reinforcement learning, we see that using Wikipedia as external sources achieves the best R@5 and R@10 scores, while using MayoClinic as external evidence sources achieves the best MRR score. Furthermore, using goal-oriented reward achieves the best MRR, R@5 and R@10 scores. Our best system's improvement over the best baseline for recall@10 is statistically significant ($p = 0.03$), but recall@5 has a higher p value ($p = 0.16$). The best baseline method's better MRR result is also not statistically significant over our model ($p = 0.74$). Table 5 shows some examples where our proposed system improved the clinical concepts along with the output diagnosis ranking.

7. Conclusion and Future Work

We proposed a deep reinforcement learning-based framework that learns to diagnose from free text clinical narratives by utilizing external evidences. Our preliminary experiments on the TREC CDS dataset showed that our system achieves better results than two non-reinforcement learning-based baselines and exploring external evidence from Wikipedia and MayoClinic along with both goal-oriented and intermediate reward functions leads to improved clinical concepts that best describe a correct diagnosis. In future, we would like to explore the performance of our system on a larger EHR dataset such as MIMIC-III (Multi-parameter Intelligent Monitoring in Intensive Care) database (Johnson et al., 2016).

Initial Concepts	Final Concepts	Diagnosis	Rank Improved (%)
null	fatigue	dengue	43.86
personality, poor memory	personality change, memory loss	Creutzfeldt-Jakob disease	43.75
joint inflammation, null	joint pains, fainting	rheumatic fever	74.00

Table 5: Improved clinical concepts and diagnoses ranking with reinforcement learning.

References

Alan R Aronson. Metamap: Mapping text to the umls metathesaurus. *Bethesda, MD: NLM, NIH, DHHS*, pages 1–26, 2006.

Saeid Balaneshin-kordan and Alexander Kotov. Optimization method for weighting explicit and latent concepts in clinical decision support queries. In *Proceedings of the 2016 ACM on International Conference on the Theory of Information Retrieval*, pages 241–250. ACM, 2016.

Junwei Bao, Nan Duan, Ming Zhou, and Tiejun Zhao. Knowledge-based question answering as machine translation. *Cell*, 2(6), 2014.

Richard Bellman. A Markovian Decision Process. *Journal of Mathematics and Mechanics*, 6:679–684, 1957.

K. Berge and S. Mamede. Cognitive diagnostic error in internal medicine. *European Journal of Internal Medicine*, 24, 2013.

Olivier Bodenreider. The unified medical language system (umls): integrating biomedical terminology. *Nucleic acids research*, 32(suppl 1):D267–D270, 2004.

Antoine Bordes, Sumit Chopra, and Jason Weston. Question answering with subgraph embeddings. *arXiv preprint arXiv:1406.3676*, 2014.

Raghavendra Chalapathy, Ehsan Zare Borzeshi, and Massimo Piccardi. Bidirectional lstm-crf for clinical concept extraction. In *Proceedings of the Clinical Natural Language Processing Workshop*, pages 7–12. ACL, 2016.

Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. Reading wikipedia to answer open-domain questions. *arXiv preprint arXiv:1704.00051*, 2017.

Edward Choi, Mohammad Taha Bahadori, and Jimeng Sun. Doctor ai: Predicting clinical events via recurrent neural networks. *arXiv preprint arXiv:1511.05942*, 2015.

Edward Choi, Mohammad Taha Bahadori, Andy Schuetz, Walter F. Stewart, and Jimeng Sun. Retain: Interpretable predictive model in healthcare using reverse time attention mechanism. *CoRR*, abs/1608.05745, 2016.

Li Dong, Furu Wei, Ming Zhou, and Ke Xu. Question answering over freebase with multi-column convolutional neural networks. In *Proceedings of Association for Computational Linguistics*, pages 260–269, 2015.

David Ferrucci, Anthony Levas, Sugato Bagchi, David Gondek, and Erik T Mueller. Watson: beyond jeopardy! *Artificial Intelligence*, 199:93–105, 2013.

Shichao Geng and Qin Zhang. Clinical diagnosis expert system based on dynamic uncertain causality graph. In *Information Technology and Artificial Intelligence Conference (ITAIC), 2014 IEEE 7th Joint International*, pages 233–237. IEEE, 2014.

Travis R Goodwin and Sanda M Harabagiu. Medical question answering for clinical decision support. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, pages 297–306. ACM, 2016.

M. L. Graber, S. Kissam, V. L. Payne, A. N. Meyer, A. Sorensen, and N. Lenfestey. Cognitive interventions to reduce diagnostic error: a narrative review. *BMJ Quality & Safety*, 21, 2012.

Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. Mimic-iii, a freely accessible critical care database. *Scientific data*, 3, 2016.

Siddhartha Jonnalagadda, Trevor Cohen, Stephen Wu, and Graciela Gonzalez. Enhancing clinical concept extraction with distributional semantics. *J. of Biomedical Informatics*, 45(1):129–140, 2012.

Boris Katz, Gregory Marton, Gary C Borchardt, Alexis Brownell, Sue Felshin, Daniel Loreto, Jesse Louis-Rosenberg, Ben Lu, Federico Mora, Stephan Stiller, et al. External knowledge sources for question answering. In *TREC*, 2005.

Youngjun Kim, Ellen Riloff, and John F. Hurdle. A study of concept extraction across different types of clinical notes. In *AMIA 2015, American Medical Informatics Association Annual Symposium*, 2015.

Andre W. Kushniruk. Analysis of complex decision-making processes in health care: Cognitive approaches to health informatics. *Journal of Biomedical Informatics*, 34:365–376, 2001.

Adam Lally, Sugato Bachi, Michael A Barborak, David W Buchanan, Jennifer Chu-Carroll, David A Ferrucci, Michael R Glass, Aditya Kalyanpur, Erik T Mueller, J William Murdock, et al. Watsonpaths: scenario-based question answering and inference over unstructured information. *Yorktown Heights: IBM Research*, 2014.

Yuan Ling, Yuan An, and Sadid A. Hasan. Improving clinical diagnosis inference through integration of structured and unstructured knowledge. In *Proceedings of the 1st Workshop on Sense, Concept and Entity Representations and their Applications*, pages 31–36, Valencia, Spain, 2017. Association for Computational Linguistics.

Zachary C Lipton, David C Kale, Charles Elkan, and Randall Wetzell. Learning to diagnose with lstm recurrent neural networks. *arXiv preprint arXiv:1511.03677*, 2015.

Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013a.

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, 2013b.

Alexander Miller, Adam Fisch, Jesse Dodge, Amir-Hossein Karimi, Antoine Bordes, and Jason Weston. Key-value memory networks for directly reading documents. *CoRR*, abs/1606.03126, 2016.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

Karthik Narasimhan, Tejas D. Kulkarni, and Regina Barzilay. Language understanding for text-based games using deep reinforcement learning. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 1–11, 2015.

Karthik Narasimhan, Adam Yala, and Regina Barzilay. Improving information extraction by acquiring external evidence with reinforcement learning. *arXiv preprint arXiv:1603.07954*, 2016.

Mathieu Nendaz and Arnaud Perrier. Diagnostic errors and flaws in clinical reasoning: mechanisms and prevention in practice. *Swiss Medical Weekly*, 142, 2012.

Stelmo Magalhães Barros Netto, Anselmo Cardoso de Paiva, Areolino de Almeida Neto, Aristofanes Correa Silva, and Vanessa Rodrigues Coelho Leite. *Application on Reinforcement Learning for Diagnosis Based on Medical Image*. INTECH Open Access Publisher, 2008.

Geoff Norman, Meredith Young, and Lee Brooks. Non-analytical models of clinical reasoning: the role of experience. *Medical Education*, 41(12):1140–1145, 2007.

Thierry Pelaccia, Jacques Tardif, Emmanuel Triby, and Bernard Charlin. An analysis of clinical reasoning through a recent and comprehensive approach: the dual-process theory. *Medical Education Online*, 16(0), 2011.

Radhika Poolla. A reinforcement learning approach to obtain treatment strategies in sequential medical decision problems. *Graduate Theses and Dissertations, University of South Florida*, 2003.

Aaditya Prakash, Siyuan Zhao, Sadid A. Hasan, Vivek Datla, Kathy Lee, Ashequl Qadir, Joey Liu, and Oladimeji Farri. Condensed Memory Networks for Clinical Diagnostic Inferencing. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, pages 3274–3280, 2017.

J. B. Reilly, A. R. Ogdie, J. M. Feldt, and J. S. Myers. Teaching about how doctors think: a longitudinal curriculum in cognitive bias and diagnostic error for residents. *BMJ Quality & Safety*, 22, 2013.

W. S. Richardson, M. C. Wilson, G. H. Guyatt, D. J. Cook, and J. Nishikawa. Users' guides to the medical literature: XV. How to use an article about disease probability for differential diagnosis. Evidence-Based Medicine Working Group. *JAMA : The Journal of the American Medical Association*, 281(13):1214–1219, 1999.

Kirk Roberts, Matthew S Simpson, Ellen Voorhees, and William R Hersh. Overview of the trec 2015 clinical decision support track. In *TREC*, 2015.

Kirk Roberts, Matthew S. Simpson, Dina Demner-Fushman, Ellen M. Voorhees, and William R. Hersh. State-of-the-art in biomedical literature retrieval for clinical cases: a survey of the TREC 2014 CDS track. *Information Retrieval Journal*, 19(1-2):113–148, 2016.

Bastian M. Seidel, Steven Campbell, and Erica Bell. Evidence in clinical reasoning: a computational linguistics analysis of 789,712 medical case summaries 1983–2012. *BMC Medical Informatics and Decision Making*, 15(1), 2015.

Longxiang Shi, Shijian Li, Xiaoran Yang, Jiaheng Qi, Gang Pan, and Binbin Zhou. Semantic health knowledge graph: Semantic integration of heterogeneous medical knowledge and services. *BioMed Research International*, 2017, 2017.

Susan M Shortreed, Eric Laber, Daniel J Lizotte, T Scott Stroup, Joelle Pineau, and Susan A Murphy. Informing sequential clinical decision-making through reinforcement learning: an empirical study. *Machine learning*, 84(1-2):109–136, 2011.

Matthew S Simpson, Ellen M Voorhees, and William Hersh. Overview of the TREC 2014 Clinical Decision Support Track. In *TREC*, 2014.

Richard Socher, Danqi Chen, Christopher D Manning, and Andrew Ng. Reasoning with neural tensor networks for knowledge base completion. In *Advances in Neural Information Processing Systems*, pages 926–934, 2013.

Parikshit Sondhi, Jimeng Sun, Hanghang Tong, and ChengXiang Zhai. Sympgraph: A framework for mining clinical notes through symptom relation graphs. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1167–1175. ACM, 2012.

Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction* . The MIT Press, Cambridge, Massachusetts, London, England, 1998.

Christopher J. C. H. Watkins and Peter Dayan. Q-learning. In *Machine Learning*, pages 279–292, 1992.

Robert West, Evgeniy Gabrilovich, Kevin Murphy, Shaohua Sun, Rahul Gupta, and Dekang Lin. Knowledge base completion via search-based question answering. In *Proceedings of the 23rd international conference on World wide web*, pages 515–526. ACM, 2014.

Fei Wu and Daniel S. Weld. Open information extraction using wikipedia. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, ACL '10, pages 118–127, Stroudsburg, PA, USA, 2010. Association for Computational Linguistics. URL `http://dl.acm.org/citation.cfm?id=1858681.1858694`.

Xuchen Yao and Benjamin Van Durme. Information extraction over structured data: Question answering with freebase. In *ACL (1)*, pages 956–966. Citeseer, 2014.

Yufan Zhao, Donglin Zeng, Mark A Socinski, and Michael R Kosorok. Reinforcement learning strategies for clinical trials in nonsmall cell lung cancer. *Biometrics*, 67(4):1422–1433, 2011.

Ziwei Zheng and Xiaojun Wan. Graph-based multi-modality learning for clinical decision support. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, pages 1945–1948. ACM, 2016.