

# Clinical Reading Comprehension: A Thorough Analysis of the emrQA Dataset

Xiang Yue

Bernal Jimenez Gutierrez

Huan Sun

The Ohio State University

{yue.149, jimenezgutierrez.1, sun.397}@osu.edu

## Abstract

Machine reading comprehension has made great progress in recent years owing to large-scale annotated datasets. In the clinical domain, however, creating such datasets is quite difficult due to the domain expertise required for annotation. Recently, Pampari et al. (2018) tackled this issue by using expert-annotated question templates and existing i2b2 annotations to create *emrQA*, the first large-scale dataset for question answering (QA) based on clinical notes. In this paper, we provide an in-depth analysis of this dataset and the clinical reading comprehension (CliniRC) task. From our qualitative analysis, we find that (i) *emrQA* answers are often incomplete, and (ii) *emrQA* questions are often answerable without using domain knowledge. From our quantitative experiments, surprising results include that (iii) using a small sampled subset (5%-20%), we can obtain roughly equal performance compared to the model trained on the entire dataset, (iv) this performance is close to human expert's performance, and (v) BERT models do not beat the best performing base model. Following our analysis of the *emrQA*, we further explore two desired aspects of CliniRC systems: the ability to utilize clinical domain knowledge and to generalize to unseen questions and contexts. We argue that both should be considered when creating future datasets.<sup>1</sup>

## 1 Introduction

Medical professionals often query over clinical notes in Electronic Medical Records (EMRs) to find information that can support their decision making (Demner-Fushman et al., 2009; Rosenbloom et al., 2011; Wang et al., 2018). One way to facilitate such information seeking activities is to build a natural language question answering (QA) system that can extract precise answers from clinical notes (Cairns et al., 2011; Cao et al., 2011; Wren, 2011; Abacha and Demner-Fushman, 2016, 2019).

<sup>1</sup>Our code is available at <https://github.com/xiangyue9607/CliniRC>.

RECORD #992321, Date: 2145-09-22
<b>Context:</b> ... <u>For HTN control, pt was given HCTZ</u> and lopressor which sufficiently controlled his BP. Pt was sent home on HCTZ 25mg daily and atenolol 50mg daily.
...
ADDITIONAL COMMENTS: 1.) <u>Take hydrochlorothiazide 25mg daily</u> and atenolol 50mg daily for your blood pressure. You should also take aspirin 81mg daily.
<b>Question:</b> Why has the patient been prescribed hctz? <b>Answer:</b> <u>For HTN control, pt was given HCTZ and lopressor which sufficiently</u>
<b>Question:</b> What was the dosage prescribed of hydrochlorothiazide? <b>Answer:</b> <u>ADDITIONAL COMMENTS: 1.) Take hydrochlorothiazide 25mg daily and atenolol 50mg daily for your</u>

Figure 1: Examples from the *emrQA* dataset: Part of a clinical note as *context* and 2 *question-answer* pairs. Due to the original *emrQA* generation issues, oftentimes answers are incomplete or contain irrelevant parts to the questions (the underlined parts are what we think the most relevant to the questions).

Machine reading comprehension (RC) aims to automatically answer questions based on a given document or text corpus and has drawn wide attention in recent years. Many neural models (Cheng et al., 2016; Wang et al., 2017; Wang and Jiang, 2017; Seo et al., 2017; Chen et al., 2017; Devlin et al., 2019) have achieved very promising results on this task, owing to large-scale QA datasets (Hermann et al., 2015; Rajpurkar et al., 2016; Trischler et al., 2017; Joshi et al., 2017; Yang et al., 2018). Unfortunately, clinical reading comprehension (CliniRC) has not observed as much progress due to the lack of such QA datasets.

In order to create QA pairs on clinical texts, annotators must have considerable medical expertise and data handling must be specifically designed to address ethical issues and privacy concerns. Due to these requirements, using crowdsourcing like in the open domain to create large-scale clinical QA

datasets becomes highly impractical (Wei et al., 2018).

Recently, Pampari et al. (2018) found a smart way to tackle this issue and created *emrQA*, the first large-scale QA dataset on clinical texts. Instead of relying on crowdsourcing, *emrQA* was semi-automatically generated based on annotated question templates and existing annotations from the n2c2 (previously called i2b2) challenge datasets<sup>2</sup>. Example QA pairs from the dataset are shown in Figure 1.

In this paper, we aim to gain a deep understanding of the ClinIRC task and conduct a thorough analysis of the *emrQA* dataset. We first explore the dataset directly by carrying out a meticulous qualitative analysis on randomly-sampled QA pairs and we find that: 1) Many answers in the *emrQA* dataset are incomplete and hence are hard to read and ineffective for training (§3.1). 2) Many questions are simple: More than 96% of the examples contain the same key phrases in both questions and answers. Though Pampari et al. (2018) claims that 39% of the questions may need knowledge to answer, our error analysis suggests only a very small portion of the errors (2%) made by a state-of-the-art reader might be due to missing external domain knowledge (§3.2).

Following our qualitative analysis of the *emrQA* dataset, we conduct a comprehensive quantitative analysis based on state-of-the-art readers and BERT models (BERT-base (Devlin et al., 2019) as well as its biomedical and clinical versions: BioBERT (Lee et al., 2019) and ClinicalBERT (Alsentzer et al., 2019)) to understand how different systems behave on the *emrQA* dataset. Surprising results include: 1) Using a small sampled subset (5%-20%), we can obtain roughly equal performance compared to the model trained on the entire dataset, suggesting that many examples in the dataset are redundant (§4.1). 2) The performance of the best base model is close to the human expert’s performance<sup>3</sup> (§4.2). 3) The performance of BERT models is around 1%-5% worse than the best performing base model (§4.3).

After completing our analysis of the dataset, we explore two potential needs for systems doing ClinIRC: 1) The need to represent and use clinical domain knowledge effectively (§5.1) and 2) the need to generalize to unseen questions and contexts (§5.2). To investigate the first one, we analyze sev-

<sup>2</sup><https://portal.dbmi.hms.harvard.edu/projects/n2c2-nlp/>

<sup>3</sup>Which is obtained by comparing *emrQA* answers to answers created by our medical experts on sampled QA pairs.

	Medication	Relation
# Question	222,957	904,592
# Context	261	423
# Question Template	80	139
Question: avg. tokens	8.00	7.91
Answers: avg. tokens	9.47	10.41
Context: avg. tokens	1062.66	889.23

Table 1: Statistics of two major subsets, *Medication* and *Relation*, of the *emrQA* dataset.

eral types of clinical questions that require domain knowledge and can frequently appear in the real clinical setting. We also carry out an experiment showing that adding knowledge explicitly yields around 5% increase in F1 over the base model when tested on samples that we created by altering the original questions to involve semantic relations. To study generalizability, we ask medical experts to create new questions based on the unseen clinical notes from MIMIC-III (Johnson et al., 2016), a freely accessible critical care database. We find that the performance of the best model trained on *emrQA* drops by 40% under this new setting, showing how critical it is for us to develop more robust and generalizable models for the ClinIRC task.

In summary, given our analysis of the *emrQA* dataset and the task in general, we conclude that future work still needs to create better datasets to advance ClinIRC. Such datasets should be not only large-scale, but also less noisy, more diverse, and allow researchers to directly evaluate a system’s ability to encode domain knowledge and to generalize to new questions and contexts.

## 2 Overview of the *emrQA* dataset

Similar to the open-domain reading comprehension task, the Clinical Reading Comprehension (ClinIRC) task is defined as follows:

**Definition 2.1.** Given a patient’s clinical note (context)  $C = \{c_1, \dots, c_n\}$  and a question  $Q = \{t_1, \dots, t_m\}$ , the ClinIRC task aims to extract a continuous span  $A = \{c_i, c_{i+1}, \dots, c_{i+k}\} (1 \leq i \leq i+k \leq n)$  from the context as the answer, where  $c_i, t_j$  are tokens.

The *emrQA* dataset (Pampari et al., 2018) was semi-automatically generated from expert-annotated question templates and existing i2b2 annotations. More specifically, clinical question templates were first created by human experts. Then, manual annotations from the medication information extraction, relation learning, and coreference

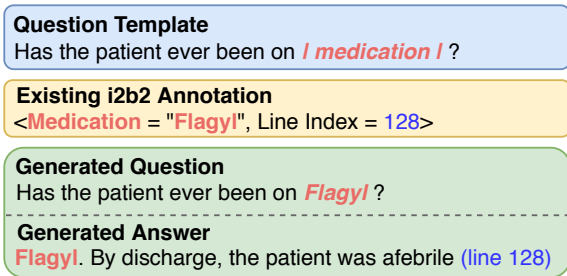


Figure 2: An example to illustrate how emrQA generates QA pairs.

resolution i2b2 challenges were re-framed into answers for the question templates. After linking question templates to i2b2 annotations, the gold annotation entities were used to both replace placeholders in the question templates and extract the sentence around them as answers. An example of this generation process can be seen in Figure 2.

The emrQA dataset contains 5 subsets: *Medication*, *Relation*, *Heart Disease*, *Obesity* and *Smoking*, which were generated from 5 i2b2 challenge datasets respectively. The answer format in each dataset is different. For the *Obesity* and *Smoking* datasets, answers are categorized into 7 classes and the task is to predict the question’s class based on the context. For the *Medication*, *Relation*, and *Heart Disease* datasets, answers are usually short snippets from the text accompanied by a longer span around it which we refer to as an evidence. The short snippet is a single entity or multiple entities while the evidence contains the entire line around those entities in the clinical note. For questions that cannot be answered via entities, only the evidence is provided as an answer. Given that some questions do not have short answers and that entire evidence spans are usually important for supporting clinical decision making (Demner-Fushman et al., 2009), we treat the *answer evidence*<sup>4</sup> as our answer just as is done in (Pampari et al., 2018).

In this work, we mainly focus on the *Medication* and *Relation* datasets because (1) they make up 80% of the entire emrQA dataset and (2) their format is consistent with the span extraction task, which is more challenging and meaningful for clinical decision making support. We filter the answers whose lengths (number of tokens) are more than 20. The detailed statistics of the two datasets are shown in Table 1.

<sup>4</sup>For simplicity, we use “answer” directly henceforth.

Metric	Medication	Relation
Quality Score	3.92	4.75
EM	26.0	92.0
F1	74.7	95.4

Table 2: An estimate of the quality of answers in the *Medication* and *Relation* datasets based on the analysis of our randomly sampled 50 questions for each dataset. Quality scores are the average of two human annotators’ (maximum: 5). EM and F1 scores are calculated between *human-labeled* answers v.s. *emrQA* answers.

### 3 In-depth Qualitative Analysis

In this section, we carry out an in-depth analysis of the emrQA dataset. We aim to examine (1) the quality and (2) level of difficulty for the generated QA pairs in the emrQA dataset.

#### 3.1 How clean are the emrQA answers?

Since the emrQA dataset was created via a generation framework unlike human-labeled or crowdsourcing datasets, the quality of the datasets remains largely unknown. In order to use this dataset to explore the ClinRC task, it is essential to determine whether it is meaningful.

In order to do this, we randomly sample 50 QA pairs from the *Medication* and the *Relation* datasets respectively. Since some questions share the same answer due to automatic generation, we make sure all the samples have different answers.

Since the questions were generated from expert created templates, most of them are human-readable and unambiguous. We therefore mainly focus on evaluating answer quality. We ask two human experts to score each answer from 1 to 5 depending on the relevance of the answer to the question (1: irrelevant or incorrect; 2: missing key parts; 3: contains key parts but is not human-readable or contains many irrelevant parts; 4: contains key parts and is only missing a few parts or has a few irrelevant extra segments; 5: perfect answer). We also ask human annotators to label the gold answers and then calculate the Exact Match (EM) and F1 score (F1) of the emrQA answers v.s. human gold answers. The answer quality score, EM and F1 in both datasets, are shown in Table 2.

The scores of the *Medication* dataset are low since most of the answers are broken sentences or contain unnecessary segments. For instance, in the Figure 2 example, the correct answer should be “*Clindamycin was changed to Flagyl*”, how-

Error Type	Question	emrQA Answers	Prediction	Error Ratio	
				Medication	Relation
<b>Span mismatch - include key info</b>	Does she have a history of known drug allergies?	ALLERGIES: He had no known drug allergies	He had no known drug allergies	78%	66%
<b>Span mismatch - miss key info</b>	What is the current dose of lasix?	MEDS: K-Dur 20 BID, Nexium 20, lasix 160 BID	BID	4%	0%
<b>Ambiguous questions</b>	What is the patient's low history?	At the time of discharge, her potassium had been low despite repletion	11) Low grade, anemia	8%	4%
<b>Incorrect golds</b>	What is the patient's incisions status?	Wash incisions with warm water and gentle soap	Do not apply lotions, creams, ointments or powders to incision	2%	2%
<b>False negatives</b>	Is there a mention of fluid in the record?	There is some fluid, or mucosal thickening in the ethmoid and sphenoid sinuses	The amount of fluid layering at the apices and the pleural spaces appear slightly decreased	2%	18%
<b>May need external knowledge</b>	What treatment has the patient had for his CAD?	CAD s/p CABG 2003 s/p	Pt's vancomycin was stopped after 14 days of treatment	2%	2%
<b>Others</b>	Is the patient's right hand ganglion cyst well-controlled?	right hand ganglion cyst removed	x 3 right hand ganglion cyst	4%	8%

Table 3: Error analysis on 50 sampled questions from the *Medication* and *Relation* dev sets respectively. Example question, ground truth and prediction from either *Medication* or *Relation* are given for each type of error.

ever, the emrQA answer misses important parts “*Clindamycin was changed to*” and contains irrelevant parts “*By discharge, the patient was afebrile*”. These issues are common in the *Medication* dataset and make it difficult to train a good system. To understand why the generated answers contain such noise, we explored the “*i2b2 2009 Medication*” challenge dataset which was used to create these QA pairs. We found that most documents in this dataset contain many complete sentences split into separate lines. Since the i2b2 annotation are token based and the emrQA obtains full lines around the token as evidence spans, these lines often end up being broken sentences. We tried to relabel the answers with existing sentence segmentation tools and heuristic measures but found that it is very challenging to obtain concise and complete text spans as answers.

Compared with the *Medication* dataset, the answer quality of the *Relation* dataset is much better. In most cases, the answers are complete and meaningful sentences with no unnecessary parts.

### 3.2 How challenging are the emrQA pairs?

Another observation from the 50 samples is that 96% of the answers in the *Medication* dataset and 100% of the answers in the *Relation* dataset contain the key phrase in the question. This is due to the generation procedure illustrated in Figure 2. In this

example, the key phrase or entity (“*Flagyl*”) in the question is also included in the answer. This undoubtedly makes the answer easier to extract as long as the model can recognize significant words and do “word matching”.

To further explore how much clinical language understanding is needed and what kind of errors do the state-of-the-art reader make, we conduct error analysis using DocReader (Chen et al., 2017) (also used in (Pampari et al., 2018)) on the emrQA dataset. More specifically, we randomly sample 50 questions that are answered incorrectly by the model (based on exact match metric) from the *Medication* and *Relation* dev set respectively<sup>5</sup>. The results are shown in Table 3 (examples for each error type are also given for better understanding).

Since emrQA answers are often incomplete in the dataset, we deem *span mismatch* errors acceptable as long as the predictions include the key part of the ground truths. Surprisingly, *span mismatch-include key info* errors, along with *ambiguous questions*, *incorrect golds* and *false negatives* (the prediction is correct but it is not in the emrQA answers) errors, which are caused by the dataset itself, account for 90% of total errors, suggesting that the accuracy of these models is even higher than we report.

<sup>5</sup>Note that these 100 samples are sampled from errors, which are different from the previously sampled ones.



Another interesting finding from the error analysis is that to our surprise, only a very small amount (2%) of errors may have been caused by a lack of external domain knowledge while Pampari et al. (2018) claim that 39% of the questions in the emrQA dataset need domain knowledge. This surprising result might be due to: (1) neural models being able to encode relational or associative knowledge from the text corpora as has also been reported in recent studies (Petroni et al., 2019; Bouraoui et al., 2020), and (2) questions and answers sharing key phrases (as we mentioned earlier in §3.1) in many samples, making it more likely that fewer questions need external knowledge to be answered than previously reported.

## 4 Comprehensive Quantitative Analysis

In this section, we conduct comprehensive experiments on the emrQA dataset with state-of-the-art readers and recently dominating BERT models. Full experimental settings are described in Appendix A.

### 4.1 How redundant are the emrQA pairs?

Though there are more than 1 million questions in the emrQA dataset (as shown in Table 1), many questions and their patterns are very similar since they are generated from the same question templates. This observation leads to a natural question: *do we really need so many questions to train an CliniRC system?* If many questions are similar to each other, it is very likely that using a sampled subset can achieve roughly the same performance that is based on the entire dataset.

To verify our hypothesis, we first split the two datasets into train, dev, and test set with the proportion of 7:1:2 w.r.t. the contexts (full statistics are shown in Appendix Table A1). Then we randomly sample {5%, 10%, 20%, 40%, 60%} and {1%, 3%, 5%, 10%, 15%}<sup>6</sup> of the QA pairs in each document (context) of the *Medication* and the *Relation* training sets respectively. We run DocReader (Chen et al., 2017) on the sampled subsets and evaluate them on the same dev and test set.

As shown in Figure 3, using 20% of the questions in the *Medication* and 5% of the questions in the *Relation* dataset can achieve roughly the same performance as using the entire training sets.

<sup>6</sup>The sampling percentage of the *Relation* dataset is smaller than the *Medication* dataset since the former one has more QA pairs (roughly 4 times).




Figure 3: Impact of *training size* on the performance of DocReader (Chen et al., 2017) based on the *Medication* and *Relation* dataset.

These verify our hypothesis, and illustrate learning a good and robust reader system based on the emrQA dataset does not need so many question-answer pairs. While deep models are often data-hungry, it does not mean more data can always lead to better performance. In addition to the training size, diversity should also be considered as another important criterion for data quality.

In the following experiments, we use the sampled subsets (20% for *Medication* and 5% for *Relation*) considering the time and memory cost as well as performance.

### 4.2 Little room for improvement

Since the answers in emrQA are often incomplete, the performance of a model is more appropriately reflected by its F1 score. As shown in Table 2, we obtain F1 scores of 74% and 95% on two datasets respectively when we test human-labeled answers against the emrQA answers on a sampled dataset. We can see from Table 4 that the best performing reader, DocReader, achieves around 70% and 94% F1 performance on the *Medication* and *Relation* test set respectively, which are very close to the human performance just described. Though designing more complex and advanced models may achieve better scores, such scores are obtained w.r.t. noisy emrQA answers and may not translate meaningfully to real cases.

### 4.3 BERT does not always win

BERT models have achieved very promising results recently in various NLP tasks including RC (Devlin et al., 2019). We follow their experiment setting of BERT for doing reading comprehension on the SQuAD (Rajpurkar et al., 2016) dataset. To our surprise, as shown in Table 4, BERT models (BERT-base, its biomedical version BioBERT (Lee et al., 2019), and its clinical version ClinicalBERT

Model	Medication				Relation			
	Dev		Test		Dev		Test	
	EM	F1	EM	F1	EM	F1	EM	F1
BiDAF (Seo et al., 2017)	25.50	68.13	23.35	67.18	81.51	90.84	82.74	91.27
DocReader (Chen et al., 2017)	29.20	72.78	25.68	70.45	86.43	94.44	86.94	94.85
QANet (Yu et al., 2018)	27.67	69.40	24.74	67.34	82.41	90.61	82.68	91.56
BERT-base (Devlin et al., 2019)	26.62	68.75	24.00	67.49	80.17	90.01	83.29	92.38
BioBERT (Lee et al., 2019)	27.81	71.90	24.75	69.97	81.57	91.38	83.61	92.62
ClinicalBERT (Alsentzer et al., 2019)	27.14	71.84	24.06	69.05	83.12	91.96	85.33	93.06

Table 4: Overall performance of all models on the *Medication* and *Relation* dataset. All numbers are percentages.

(Alsentzer et al., 2019)) do not dominate as they do in the open-domain RC tasks. The reasons may be three-fold: 1) BERT benefits the most from large training corpora. The training corpora of BERT-base and BioBERT are Wikipedia + BookCorpus (Zhu et al., 2015) and PubMed articles respectively, both of which may have different vocabularies and use different language expressions from clinical texts. Though ClinicalBERT was pretrained on MIMIC-III (Johnson et al., 2016) clinical texts, the training size of the corpus (~50M words) is far less than that used in BERT (~3300M words), which may make the model less powerful as it is on the open-domain tasks. 2) Longer Contexts. As can be seen from Table 1, the number of tokens in the contexts is commonly larger than open-domain RC datasets like SQuAD (~1000 v.s.~116 avg). We suspect that long contexts might make it more challenging to model sequential information. For sequences that are longer than the *max length* of the BERT model, they are truncated into a set of short sequences, which may hinder the model from capturing long dependencies (Dai et al., 2019) and global information in the entire document. 3) Easy Questions. Another possible reason might be the question patterns are too easy and a simpler reader with far less parameters can learn the patterns and obtain satisfying performance.

Additionally, to further evaluate the models in the fine-grained level, inspired by (Gururangan et al., 2018), we partition the *Medication* and *Relation* test sets into Easy and Hard subsets using a base model. The details of Easy/Hard splits can be found in Appendix C. As can be seen from Table A4, most of the questions in the two datasets are easy, which indicates the emrQA dataset might not be challenging for the current QA models. More difficult datasets are needed to advance the Clinical Reading Comprehension task.

## 5 Desiderata in Real-World CliniRC

Following our analysis of the emrQA dataset, we further study two aspects of clinical reading comprehension systems that we believe are crucial for their real-world applicability: the need to encode clinical domain knowledge and to generalize to unseen questions and documents.

### 5.1 External domain knowledge is needed

So far, we have shown that domain knowledge may not be very useful for models answering questions in the *emrQA dataset*; however, we argue that systems in real-world CliniRC need to be able to encode and use clinical domain knowledge effectively.

Clinical text often contains high variability in many domain-specific words due to abbreviations and synonyms. The presence of different aliases in the question and context can make it difficult for a model to represent semantics accurately and choose the correct span. Besides, medical domain-specific relations (e.g., *treats*, *caused by*) and hierarchical relations (e.g., *isa*) between medical concepts would be likely to appear. The process followed to generate the current emrQA dataset leads to these problems being largely under-represented, even though they can be very common in real cases. We use the following 3 examples as representatives to illustrate the real cases we may encounter.

**Synonym.** For example, for the question in Figure 2, “*Has this patient ever been on Flagyl?*”, it is easy for the model to answer since “*Flagyl*” appears in the context. However, if we change “*Flagyl*” to its synonyms “*Metronidazole*” (which may not appear in training) in the question, it is hard for the reader to extract the correct answer, as it is not possible for model to capture the semantic meaning of “*Metronidazole*” as “*Flagyl*”.

**Clinical Relations.** Another example is the ques-

tion shown in Figure 1, “Why has the patient been prescribed *hctz*?”. Currently, machines can easily find the answer since keyword “*hctz*” is mentioned in the answer. However, given a situation where the drug “*hctz*” does not appear in the local context of “*HTN*”, our model may have a better chance to extract the correct answers if it stores the relation “(*hctz*, *treats*, *HTN*)”.

**Hierarchical Relation.** For the question “Is there a history of mental illness?”, it is more likely that the medical report describes a specific type of psychological condition rather than mention the general phrase “*mental illness*” since clinical support require specifics. To obtain the correct answer in this case “*Depression with previous suicidal ideation.*”, encoding the relation “(*depression*, *isa*, *mental illness*)” would probably help the model make a correct prediction.

These three cases help illustrate how complex medical relations affect the real ClinIRC task. Without leveraging external domain knowledge, it is difficult for models to capture the semantic relations necessary to resolve such cases.

In order to verify our claim quantitatively, we select *synonym* as a representative relation type and manipulate each question by replacing its entities with plausible synonyms or abbreviations. We then introduce external domain knowledge into current models and compare their performance against base models on these augmented questions.

More specifically, we first detect entities in the questions and link them to a medical knowledge base (KB): UMLS (Bodenreider, 2004) using a biomedical and clinical text NLP pipeline tool, *ScispaCy* (Neumann et al., 2019). Synonyms of detected entities are then retrieved from UMLS and used to replace the original mention. We filter the questions that do not contain entities or that contain entities with no synonyms. We focus on the *Relation* dataset and only modify the questions in the dev and test set; the questions in the training set are not modified. Finally, we get 69,912 and 125,338 questions in the dev and test set.

We then introduce a simple Knowledge Incorporation Module (KIM) to evaluate the usefulness of external domain knowledge. Formally, given a question  $q : \{w_1^q, w_2^q, \dots, w_l^q\}$  and its context  $c : \{w_1^c, w_2^c, \dots, w_m^c\}$ , where  $w_i^q, w_j^c$  are words (tokens), all the words can be mapped to  $d_1$  dimensional vectors via a word embedding matrix  $E_w \in \mathbb{R}^{d_1 \times |\mathcal{V}|}$ , where  $\mathcal{V}$  denotes the word vocab-




Figure 4: Performances of DocReader and DocReader + Knowledge Incorporation Module (KIM) on our created questions modified from the *Relation* dataset.

ulary. So we have  $q : \mathbf{w}_1^q, \dots, \mathbf{w}_l^q \in \mathbb{R}^{d_1}$  and  $c : \mathbf{w}_1^c, \dots, \mathbf{w}_m^c \in \mathbb{R}^{d_1}$ .

We then detect entities  $\{e_1^q, e_2^q, \dots, e_n^q\}$  in the question and entities  $\{e_1^c, e_2^c, \dots, e_o^c\}$  in the context and map them to a medical knowledge base (KB), UMLS (Bodenreider, 2004) using scispacy (Neumann et al., 2019). Note that  $l$  is not equal to  $n$  and  $m$  is not equal to  $o$ , since not every token can be mapped to an entity in KB. For entities that contain multiple words, we align them to the first token, same as the alignment used in (Zhang et al., 2019). We then map detected entities to  $d_2$  dimensional vectors  $\{e_1^q, e_2^q, \dots, e_n^q\}$  and  $\{e_1^c, e_2^c, \dots, e_o^c\}$  via an entity embedding matrix  $E_e \in \mathbb{R}^{d_2 \times |\mathcal{U}|}$ , which is pretrained on the entire UMLS KB using the knowledge embedding method TransE (Bordes et al., 2013).  $\mathcal{U}$  denotes the entity vocabulary.

We merge the word embeddings with entity embeddings to feed them into a Multi-layer Perceptron (MLP):

$$\begin{aligned} \mathbf{h}_i^q &= \sigma(\mathbf{W}_c \mathbf{w}_i^q + \mathbf{W}_e \mathbf{e}_i^q + \mathbf{b}) \\ \mathbf{h}_j^c &= \sigma(\mathbf{W}_c \mathbf{w}_j^c + \mathbf{W}_e \mathbf{e}_j^c + \mathbf{b}) \end{aligned} \quad (1)$$

where  $\sigma$  is activation function,  $W_c, W_e, b$  are trainable parameters and  $h_i^q, h_j^c$  denote the integrated embeddings that contain information from both the word  $c_j$  and the entity  $e_j$  in the question and context respectively. For the word that is not mapped to an entity,  $e_j$  will be set to  $\mathbf{0}$ . The merged embeddings are used as the input to the base reader.

As shown in Figure 4, by adding a basic Knowledge Incorporation Module to the base model, we obtain around 5% increase of F1 score on the manipulated questions in the test set. This suggests that for questions that involve relations between medical concepts, external domain knowledge may be quite important.

Model	Existing Questions		Paraphrased Questions		New Questions		Overall		emrQA Relation	
	EM	F1	EM	F1	EM	F1	EM	F1	EM	F1
DocReader	58.33	71.62	38.09	57.28	29.41	35.35	40.00	53.27	86.94	94.85
ClinicalBERT	58.33	73.12	38.09	62.04	23.53	48.79	38.00	60.19	85.33	93.06

Table 5: Results of models when tested on new questions and unseen clinical notes (not in emrQA, but from MIMIC-III dataset). Performance drops around 40% compared with previously reported on the *Relation* test set, highlighting generalizability as an essential future direction for CliniRC.

## 5.2 Generalizing to unseen questions and documents

The aim of CliniRC is to build robust QA systems for doctors to retrieve information buried in clinical texts. When deploying a CliniRC system to a new environment (e.g., a new set of clinical records, a new hospital, etc.), it is infeasible to create new QA pairs for training every time. Thus, an ideal CliniRC system is able to generalize to unseen documents and questions after being fully trained.

To test the generalizability of models trained on emrQA (we focus on the *Relation* dataset here), our medical experts created 50 new questions that were not present in the emrQA dataset and extracted answers from unseen patient notes in the MIMIC-III (Johnson et al., 2016) dataset. This dataset consists of three types of questions: 12 questions were made from emrQA question templates but contain entities which do not appear in the training set (e.g., “How was the diagnosis of acute cholecystitis made?” was created from the template “How was the diagnosis of |problem| made?”). The other 38 questions have different forms from existing question templates: 21 paraphrase existing questions from emrQA (e.g., “Was an edema found in the physical exam?”) was paraphrased from “Does he have any evidence of |problem| in |test|?”) and 17 are completely semantically different from the ones in the emrQA dataset (e.g., “What chemotherapy drugs are being administered to the patient?”).

As could be expected, we see in Table 5 that the more the new questions deviate from the original emrQA, the more the models struggle to answer them. We observe a performance drop of roughly 20% compared to the *Relation* test set on questions made from emrQA templates using MIMIC III clinical notes which were not in the original dataset. For question that are more significantly different, we notice an approximate 40% and 60% loss in F1 score when predicting paraphrased questions and entirely new questions respectively. This steep drop in performance for these new settings, espe-

cially for paraphrased and new questions, shows how much work there is to be done on this front and highlights generalizability as an important future direction in CliniRC. We also notice that ClinicalBERT works slightly better than the base model DocReader. The reason might be ClinicalBERT was pretrained on the MIMIC-III dataset, which might help the model have a better understanding of the context.

**Summary.** Based on these two aspects and our previous thorough analysis of the emrQA dataset, it is clear that better datasets are needed to advance CliniRC. Such datasets should be not only large-scale, but also less noisy, more diverse, and moreover allow researchers to systematically evaluate a model’s ability to encode domain knowledge and to generalize to new questions and contexts.

## 6 Related Work

We present a brief overview of open-domain, biomedical and clinical question answering tasks, which are most related to our work:

**Question Answering (QA)** aims to automatically answer questions asked by humans based on external sources, such as Web (Sun et al., 2016), knowledge base (Yih et al., 2015; Sun et al., 2015) and free text (Chen et al., 2016). As an important type of QA, reading comprehension intends to answer a question after reading the passage (Hirschman et al., 1999). Recently, the release of large-scale RC datasets, such as CNN & Daily Mail (Hermann et al., 2015), Stanford Question-Answering Dataset (SQuAD) (Rajpurkar et al., 2016, 2018) makes it possible to solve RC tasks by building deep neural models (Hermann et al., 2015; Wang and Jiang, 2017; Seo et al., 2017; Chen et al., 2017).

More recently, contextualized word representations and pretrained language models, such as ELMo (Peters et al., 2018), GPT (Radford et al., 2018), BERT (Devlin et al., 2019), have been



demonstrated to be very useful in various NLP tasks including RC. By seeing diverse contexts in large corpora, these pretrained language models can capture the rich semantic meaning and produce more accurate and precise representations for words given different contexts. Even a simple classifier or score function built upon these pretrained contextualized word representations perform well in extracting answer spans (Devlin et al., 2019).

**Biomedical and Clinical QA.** Due to the lack of large-scale annotated biomedical or clinical data, QA and RC systems in these domains are often rule-based and heuristic feature-based (Lee et al., 2006; Niu et al., 2006; Athenikos and Han, 2010).

In recent years, BioASQ challenges (Tsatsaronis et al., 2012) proposed the Biomedical Semantic QA task, where the participants need to respond to each test question with relevant articles, snippets and exact answers. Šuster and Daelemans (2018) use summary points of clinical case reports to build a large-scale cloze-style dataset (CliCR), which is similar to the style of CNN & Daily Mail dataset. Jin et al. (2019b) presents PubMedQA, which extracts question-style titles and their corresponding abstracts as the questions and contexts respectively. A few QA pairs are annotated by human experts and most of them are annotated based a simple heuristic rule with “yes/no/maybe”.

Due to the great power of contextualized word representations, pretrained language models also have been introduced to biomedical and clinical domain, e.g., BioELMo (Jin et al., 2019a), BioBERT (Lee et al., 2019), and ClinicalBERT (Alsentzer et al., 2019). They adopt similar architectures of the original models but pretrained on the medical and clinical corpus, such as PubMed articles and MIMIC-III (Johnson et al., 2016) clinical notes.

## 7 Conclusion

We study the Clinical Reading Comprehension (CliniRC) task with the recently created emrQA dataset. Our qualitative and quantitative analysis as well as exploration of the two desired aspects of CliniRC systems show that future clinical QA datasets should not only be large-scale but also less noisy and more diverse. Moreover, questions that involve complex relations and are across different domains should be included, and then more advanced external knowledge incorporation methods as well as domain adaptation methods can be carefully designed and systematically evaluated.

## Acknowledgments

We thank our medical experts for their annotations. We thank Ping Zhang, Changchang Yin and anonymous reviewers for their helpful comments. This research was sponsored in part by the Patient-Centered Outcomes Research Institute Funding ME-2017C1-6413, the Army Research Office under cooperative agreements W911NF-17-1-0412, NSF Grant IIS1815674, and Ohio Supercomputer Center (Center, 1987). The views and conclusions contained herein are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notice herein.

## References

- Asma Ben Abacha and Dina Demner-Fushman. 2016. Recognizing question entailment for medical question answering. In *AMIA Annual Symposium Proceedings*, volume 2016, page 310.
- Asma Ben Abacha and Dina Demner-Fushman. 2019. A question-entailment approach to question answering. *BMC bioinformatics*, 20(1):511.
- Emily Alsentzer, John R Murphy, Willie Boag, Weihung Weng, Di Jin, Tristan Naumann, and Matthew McDermott. 2019. Publicly available clinical bert embeddings. *NAACL Clinical NLP Workshop 2019*.
- Sofia J Athenikos and Hyoil Han. 2010. Biomedical question answering: A survey. *Computer methods and programs in biomedicine*, 99(1):1–24.
- Olivier Bodenreider. 2004. The unified medical language system (umls): integrating biomedical terminology. *Nucleic acids research*, 32(suppl\_1):D267–D270.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *NeurIPS’13*, pages 2787–2795.
- Zied Bouraoui, Jose Camacho-Collados, and Steven Schockaert. 2020. Inducing relational knowledge from bert. *AAAI’20*.
- Brian L Cairns, Rodney D Nielsen, James J Masanz, James H Martin, Martha S Palmer, Wayne H Ward, and Guergana K Savova. 2011. The mipacq clinical question answering system. In *AMIA annual symposium proceedings*, volume 2011, page 171.
- YongGang Cao, Feifan Liu, Pippa Simpson, Lamont Antieau, Andrew Bennett, James J Cimino, John Ely,

- and Hong Yu. 2011. Askhermes: An online question answering system for complex clinical questions. *Journal of biomedical informatics*, 44(2):277–288.
- Ohio Supercomputer Center. 1987. [Ohio supercomputer center](#).
- Danqi Chen, Jason Bolton, and Christopher D Manning. 2016. A thorough examination of the cnn/daily mail reading comprehension task. In *ACL’16*, pages 2358–2367.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading wikipedia to answer open-domain questions. In *ACL’17*, pages 1870–1879.
- Jianpeng Cheng, Li Dong, and Mirella Lapata. 2016. Long short-term memory-networks for machine reading. In *EMNLP’16*, pages 551–561.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime G Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. Transformer-xl: Attentive language models beyond a fixed-length context. In *ACL’19*, pages 2978–2988.
- Dina Demner-Fushman, Wendy W Chapman, and Clement J McDonald. 2009. What can natural language processing do for clinical decision support? *Journal of biomedical informatics*, 42(5):760–772.
- Joshua C Denny, Anderson Spickard III, Kevin B Johnson, Neeraja B Peterson, Josh F Peterson, and Randolph A Miller. 2009. Evaluation of a method to identify and categorize section headers in clinical documents. *Journal of the American Medical Informatics Association*, 16(6):806–815.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT’19*, pages 4171–4186.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A Smith. 2018. Annotation artifacts in natural language inference data. In *NAACL-HLT’18*, pages 107–112.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *NeurIPS’15*, pages 1693–1701.
- Lynette Hirschman, Marc Light, Eric Breck, and John D Burger. 1999. Deep read: A reading comprehension system. In *ACL’1999*, pages 325–332.
- Qiao Jin, Bhuwan Dhingra, William Cohen, and Xinghua Lu. 2019a. Probing biomedical embeddings from language models. In *NAACL’19 RepEval2019 Workshop*, pages 82–89.
- Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. 2019b. Pubmedqa: A dataset for biomedical research question answering. In *EMNLP-IJCNLP’19*, pages 2567–2577.
- Alistair EW Johnson, Tom J Pollard, Lu Shen, H Lehman Li-wei, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. 2016. MIMIC-III, a freely accessible critical care database. *Scientific data*, 3:160035.
- Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. In *ACL’17*, pages 1601–1611.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2019. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*.
- Minsuk Lee, James Cimino, Hai Ran Zhu, Carl Sable, Vijay Shanker, John Ely, and Hong Yu. 2006. Beyond information retrieval—medical question answering. In *AMIA annual symposium proceedings*, volume 2006, page 469.
- Mark Neumann, Daniel King, Iz Beltagy, and Waleed Ammar. 2019. ScispaCy: Fast and robust models for biomedical natural language processing. In *ACL BioNLP Workshop 2019*.
- Yun Niu, Xiaodan Zhu, and Graeme Hirst. 2006. Using outcome polarity in sentence extraction for medical question-answering. In *AMIA Annual Symposium Proceedings*, volume 2006, page 599.
- Anusri Pampari, Preethi Raghavan, Jennifer Liang, and Jian Peng. 2018. emrqa: A large corpus for question answering on electronic medical records. In *EMNLP’18*, pages 2357–2368.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *NAACL-HLT’18*, pages 2227–2237.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language models as knowledge bases? In *EMNLP-IJCNLP’19*, pages 2463–2473.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. *Technical report, OpenAI*.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don’t know: Unanswerable questions for squad. In *ACL’18*, pages 784–789.

- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. In *EMNLP'16*, pages 2383–2392.
- S Trent Rosenbloom, Joshua C Denny, Hua Xu, Nancy Lorenzi, William W Stead, and Kevin B Johnson. 2011. Data from clinical notes: a perspective on the tension between structure and flexible documentation. *Journal of the American Medical Informatics Association*, 18(2):181–186.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hananneh Hajishirzi. 2017. Bidirectional attention flow for machine comprehension. In *ICLR'17*.
- Huan Sun, Hao Ma, Xiaodong He, Wen-tau Yih, Yu Su, and Xifeng Yan. 2016. Table cell search for question answering. In *WWW'16*, pages 771–782.
- Huan Sun, Hao Ma, Wen-tau Yih, Chen-Tse Tsai, Jingjing Liu, and Ming-Wei Chang. 2015. Open domain question answering via semantic enrichment. In *WWW'15*, pages 1045–1055.
- Simon Šuster and Walter Daelemans. 2018. Clicr: A dataset of clinical case reports for machine reading comprehension. In *NAACL-HLT'18*, pages 1551–1563.
- Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordani, Philip Bachman, and Kaheer Suleman. 2017. NewsQA: A machine comprehension dataset. In *ACL'17 RepLANLP Workshop*, pages 191–200.
- George Tsatsaronis, Michael Schroeder, Georgios Paliouras, Yannis Almirantis, Ion Androutsopoulos, Eric Gaussier, Patrick Gallinari, Thierry Artieres, Michael R Alvers, Matthias Zschunke, et al. 2012. Bioasq: A challenge on large-scale biomedical semantic indexing and question answering. In *2012 AAAI Fall Symposium Series*.
- Shuohang Wang and Jing Jiang. 2017. Machine comprehension using match-lstm and answer pointer. *ICLR'17*.
- Wenhui Wang, Nan Yang, Furu Wei, Baobao Chang, and Ming Zhou. 2017. Gated self-matching networks for reading comprehension and question answering. In *ACL'17*, pages 189–198.
- Yanshan Wang, Liwei Wang, Majid Rastegar-Mojarad, Sungrim Moon, Feichen Shen, Naveed Afzal, Sijia Liu, Yuqun Zeng, Saeed Mehrabi, Sunghwan Sohn, et al. 2018. Clinical information extraction applications: a literature review. *Journal of biomedical informatics*, 77:34–49.
- Qiang Wei, Amy Franklin, Trevor Cohen, and Hua Xu. 2018. Clinical text annotation—what factors are associated with the cost of time? In *AMIA Annual Symposium Proceedings*, volume 2018, page 1552.
- Jonathan D Wren. 2011. Question answering systems in biology and medicine—the time is now. *Bioinformatics*, 27(14):2025–2026.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *EMNLP'18*, pages 2369–2380.
- Wen-tau Yih, Ming-Wei Chang, Xiaodong He, and Jianfeng Gao. 2015. Semantic parsing via staged query graph generation: Question answering with knowledge base. In *ACL'15*, pages 1321–1331.
- Adams Wei Yu, David Dohan, Minh-Thang Luong, Rui Zhao, Kai Chen, Mohammad Norouzi, and Quoc V Le. 2018. Qanet: Combining local convolution with global self-attention for reading comprehension. *ICLR'18*.
- Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019. ERNIE: Enhanced language representation with informative entities. In *ACL'19*.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *ICCV'15*, pages 19–27.

	Medication	Relation
# Train (Q / C)	154,684 / 182	621,428 / 296
# Dev (Q / C)	23,081 / 26	101,700 / 42
# Test (Q / C)	45,192 / 53	181,464 / 85
<b>Total</b>	<b>222,957 / 261</b>	<b>904,592 / 423</b>

Table A1: Statistics of train, dev, test set of the *Medication* and *Relation* datasets.

## A Experimental Set-up

We split the two datasets *Medication* and *Relation* based on the documents (clinical texts) into train, dev, test with the ratio 7:1:2. The statistics are shown in Table A1.

We adopt Exact Match (EM) and F1 score (F1) as our evaluation metrics, same as the open-domain RC (Rajpurkar et al., 2016). We use SQuAD v1.1 official evaluation script<sup>1</sup> to evaluate all the models. All the models used in the paper, BiDAF<sup>2</sup>, DocReader<sup>3</sup>, QANet<sup>4</sup>, BERT<sup>5</sup>, BioBERT<sup>6</sup>, ClinicalBERT<sup>7</sup> are run based on the implementations listed here and strictly followed the instructions.

For reproducibility, we list all the key hyperparameters we use for each method in the Table A2.

We implement our Knowledge Incorporation Module based on DocReader implementations. Entity embeddings are pretrained using TransE (Bordes et al., 2013) with the dimension of 100. The hyperparameters are kept same as the DocReader. All the models are run on NVIDIA GeForce GTX 1080 GPUs. We save the best model (with the highest EM) on the dev set and use it for test set.

## B Performance on Shorter Contexts

Using the entire clinical record as the context might be too long for models to capture sequential information. We also try to split the entire record into different sections (e.g., “medical history”, “family history”) based on some heuristic measures. Specifically, in order to split the clinical notes into sections, we notice that most sections begin with easily identifiable headers. To detect these headers we use a combination of heuristics such as whether the line contains colons, all uppercase formatting

<sup>1</sup><https://rajpurkar.github.io/SQuAD-explorer/>

<sup>2</sup><https://github.com/allenai/bi-att-flow>

<sup>3</sup><https://github.com/facebookresearch/DrQA>

<sup>4</sup><https://github.com/BangLiu/QANet-PyTorch>

<sup>5</sup><https://github.com/google-research/bert>

<sup>6</sup><https://github.com/dmis-lab/biobert>

<sup>7</sup><https://github.com/EmilyAlsentzer/clinicalBERT>

Method	Hyper-parameters Setting
<b>DocReader</b>	epoch: 30; batch-size: 16; test-batch-size:16; dropout-rate:0.4; doc-layers: 3; question-layers: 3; grad-clipping: 10; tune-partial: 1000; max-len: 30; the others are set as default
	init_lr: 0.001; batch-size:6; num_epochs: 2; cluster: True; len-opt: True; word_count_th: 10; char_count_th: 50; sent_size_th: 4000; num_sents_th: 500; ques_size_th: 30; word_size_th: 30; para_size.th: 4000; the others are set as default;
<b>QANet</b>	batch-size: 4; lr: 0.001; grad-clip: 5; use-ema: True; epoch: 30; para_limit: 4000; ques_limit: 30; ans_limit: 30; char_limit: 40; num-head:1; the others are set as default;
	train_batch_size: 6; learning_rate: 3e-5; num_train_epochs: 3.0; max_seq_length: 384; doc_stride: 128; the others are set as default;

Table A2: Hyperparameters settings for all the methods used in the experiments.

Dataset	Model	Dev		Test	
		EM	F1	EM	F1
medication	DocReader	32.19	76.21	33.45	77.08
	ClinicalBERT	30.16	74.81	32.18	75.79
relation	DocReader	87.21	94.32	87.54	94.97
	ClinicalBERT	85.46	93.92	85.67	93.14

Table A3: Performance of the two models on the shorter context setting.

or phrases found in a list of clinical headers taken from SecTag (Denny et al., 2009). We then select the section that contains the answer as the context (~100 words avg). We select DocReader and ClinicalBERT as representative methods and re-run them on the modified shorter context. The results are shown in Table A3. The performance of the two models is improved compared with the performance of models built on the whole record (long context). However, ClinicalBERT still does not outperform DocReader in this setting, indicating that longer context may not explain why BERT models do not win on this dataset or that shortening context in a such manner might break long dependencies.

This experiment setting may also inspire future research on “Open Clinical Reading Comprehension”. Given that patients often have multiple clini-



Distribution of Easy/Hard Questions		Easy		Hard		Total			
		Medication	33,037 (73.1%)	12,155 (26.9%)	45,192 (100%)				
		Relation	165,271 (91.1%)	16,193 (8.9%)	181,464 (100%)				
Results	Medication	Model	Easy		Hard		Total		
			EM	F1	EM	F1	EM	F1	
	DocReader (Chen et al., 2017)	30.25	73.78	13.26	61.46	25.68	70.45		
	ClinicalBERT (Alsentzer et al., 2019)	28.25	72.02	12.64	60.98	24.06	69.05		
	Relation	Model	DocReader (Chen et al., 2017)	87.66	95.39	79.85	89.62	86.94	94.85
			ClinicalBERT (Alsentzer et al., 2019)	86.06	93.71	78.09	86.57	85.33	93.06

Table A4: Performance of DocReader and ClinicalBERT on the easy/hard questions split.

cal records, it may not be feasible to jointly use all of them as context for one question. Given multiple records for one patient (instead of just one) and a question, the model would first need to retrieve the most relevant paragraphs and do reading comprehension on each of them or find clever ways to merge them. Such a setting would be interesting for future ClinIRC datasets to explore.

### C Easy/Hard Questions Split

We partition the questions into Easy and Hard. Specifically, we first train a BiLSTM reader and do the prediction on the test set. We obtain the performance of each question template by averaging the performance of all the questions made by this template (such template and question mappings are included in the emrQA dataset). Question templates that obtain higher performance than the overall performance are labeled as "Easy" and "Hard" otherwise. Then we map the difficulty level of question templates back to each question. The reason why we focus on splitting on the question template level is that we can avoid some random noise (e.g., random errors produced by the model on some questions). Also, we release the difficulty level of each question template so that users can easily know which questions are easy or hard and do not need to run a base model to obtain such mappings again. Distributions of easy/hard questions and results of the two selected models are shown in Table A4.