

Predicting Clinical Trial Results by Implicit Evidence Integration

Qiao Jin^{1*}, Chuanqi Tan², Mosha Chen², Xiaozhong Liu³, Songfang Huang²

¹School of Medicine, Tsinghua University

²Alibaba Group

³Indiana University Bloomington

jqa14@mails.tsinghua.edu.cn

{chuanqi.tcq, chenmosha.cms, songfang.hsf}@alibaba-inc.com

liu237@indiana.edu

Abstract

Clinical trials provide essential guidance for practicing Evidence-Based Medicine, though often accompanying with unendurable costs and risks. To optimize the design of clinical trials, we introduce a novel Clinical Trial Result Prediction (CTRP) task. In the CTRP framework, a model takes a PICO-formatted clinical trial proposal with its background as input and predicts the result, i.e. how the Intervention group compares with the Comparison group in terms of the measured Outcome in the studied Population. While structured clinical evidence is prohibitively expensive for manual collection, we exploit large-scale unstructured sentences from medical literature that implicitly contain PICOs and results as evidence. Specifically, we pre-train a model to predict the disentangled results from such implicit evidence and fine-tune the model with limited data on the downstream datasets. Experiments on the benchmark Evidence Integration dataset show that the proposed model outperforms the baselines by large margins, e.g., with a 10.7% relative gain over BioBERT in macro-F1. Moreover, the performance improvement is also validated on another dataset composed of clinical trials related to COVID-19.

1 Introduction

Shall COVID-19 patients be treated with hydroxychloroquine? In the era of Evidence-Based Medicine (EBM, Sackett 1997), medical practice should be guided by well-designed and well-conducted clinical research, such as randomized controlled trials. However, conducting clinical trials is expensive and time-consuming. Furthermore, inappropriately designed studies can be devastating in a pandemic: a high-profile Remdesivir clinical trial fails to achieve statistically significant conclusions (Wang et al., 2020b), partially because it

* Work done during internship at Alibaba DAMO Academy.

does not attain the predetermined sample size when “competing with” other inappropriately designed trials that are unlikely to succeed or not so urgent to test (e.g.: physical exercises and dietary treatments). Therefore, it is crucial to carefully design and evaluate clinical trials before conducting them.

Proposing new clinical trials requires support from previous evidence in medical literature or practice. For example, the World Health Organization (WHO) has launched a global megatrial, Solidarity (WHO, 2020), to prioritize clinical resources by recommending only four most promising therapies¹. The rationale for this suggestion comes from the integration of evidence that they might be effective against coronaviruses or other related organisms in laboratory or clinical studies (Peymani et al., 2016; Sheahan et al., 2017; Morra et al., 2018). However, manual integration of evidence is far from satisfying, as one study reports that about 86.2% of clinical trials fail (Wong et al., 2019) and even some of the Solidarity therapies do not get expected results (Mehra et al., 2020).

To assist clinical trial designing, we introduce a novel task: Clinical Trial Result Prediction (CTRP), i.e. predicting the results of clinical trials without actually doing them (§3). Figure 1 shows the architecture of the CTRP task. We define the input to be a clinical trial proposal², which contains free-texts of a Population (e.g.: “COVID-19 patients with severe symptoms”), an Intervention (e.g.: “Active remdesivir (i.v.)”), a Comparator (e.g.: “Placebos matched remdesivir”) and an Outcome (e.g.: “Time to clinical improvement”), i.e. a PICO-formatted query (Huang et al., 2006), and the background of the proposed trial. The output is the trial Result, denoting how (higher, lower, or no difference) I

¹Remdesivir, lopinavir/ritonavir, interferon beta-1a and chloroquine/hydroxychloroquine.

²The proposals need to be registered and approved before the clinical trials are conducted.

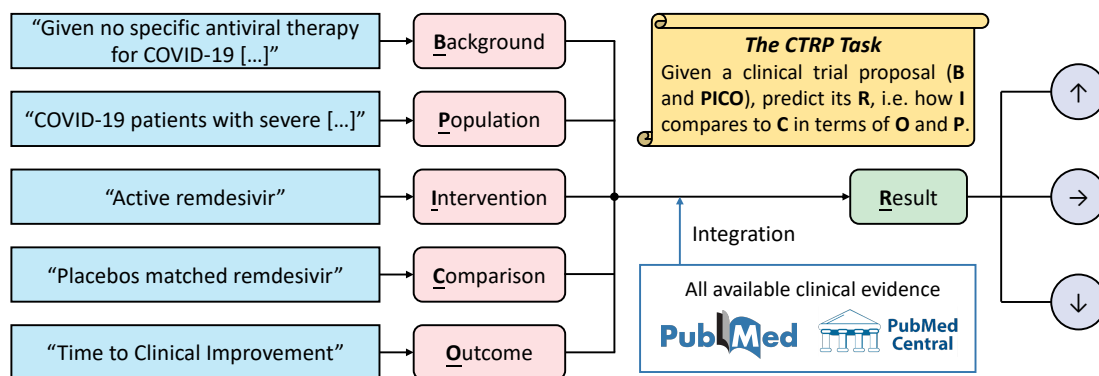


Figure 1: Architecture of the proposed Clinical Trial Result Prediction (CTRP) task.

compares to C in terms of O for P.

One particular challenge of this task is that evidence is entangled with other free-texts in the literature. Prior works have explored *explicit* methods for evidence integration through a pipeline of retrieval, extraction and inference on structured {P,I,C,O,R} evidence (Wallace et al., 2016; Singh et al., 2017; Jin and Szolovits, 2018; Lee and Sun, 2018; Nye et al., 2018; Marshall et al., 2017; Lehman et al., 2019; DeYoung et al., 2020; Zhang et al., 2020). However, they are limited in scale since getting domain-specific supervision for all clinical evidence is prohibitively expensive.

In this work, we propose to *implicitly* learn from such evidence by pre-training, instead of relying on *explicit* evidence with purely supervised learning. There are more than 30 million articles in PubMed³, which stores almost all available medical evidence and thus is an ideal source for learning. We collect 12 million sentences from PubMed abstracts and PubMed Central⁴ (PMC) articles with comparative semantics, which is commonly used to express clinical evidence (§4.1). P, I, C, O, and R are entangled with other free-texts in such sentences, which we denote as implicit evidence. Unlike previous efforts that seek to disentangle all of PICO and R, we only disentangle R out of the implicit evidence using simple heuristics (§4.2). For better learning the ordering function of I/C conditioned on P and O, we also use adversarial examples generated by reversing both the entangled PICO and the R in the pre-training (§4.3). Then, we pre-train a transformer encoder (Vaswani et al., 2017) to predict the disentangled R from the implicit evidence, which still contains PICO (§5.1). The model is named EBM-Net to reflect its utility for Evidence-Based

Medicine. Finally, we fine-tune the pre-trained EBM-Net on downstream datasets of the CTRP task (§5.2), which are typically small in scale (§6).

To evaluate model performance, we introduce a benchmark dataset, Evidence Integration (§6.1), by re-purposing the *evidence inference* dataset (Lehman et al., 2019; DeYoung et al., 2020). Experiments show that our pre-trained EBM-Net outperforms the baselines (§6.2) by large margins (§6.3). Clustering analyses indicate that EBM-Net can effectively learn quantitative comparison results (§6.4). In addition, the EBM-Net model is further validated on a dataset composed of COVID-19 related clinical trials (§6.5).

Our contribution is two-fold. First, we propose a novel and meaningful task, CTRP, to predict clinical trial results before conducting them. Second, unlike previous efforts that depend on structured data to understand the totality of clinical evidence, we heuristically collect unstructured textual data, i.e. implicit evidence, and utilize large-scale pre-training to tackle the proposed CTRP task. The datasets and codes are publicly available at <https://github.com/Alibaba-NLP/EBM-Net>.

2 Related Works

Predicting Clinical Trial Results: Most relevant works typically use only specific types or sources of information for prediction (e.g.: chemical structures (Gayvert et al., 2016), drug dosages or routes (Holford et al., 2000, 2010)). Gayvert et al. (2016) predicts clinical trial results based on chemical properties of the candidate drugs. Clinical trial simulation (Holford et al., 2000, 2010) applies pharmacological models to predict the results of a specific intervention with different procedural factors, such as doses and sampling intervals. Some use closely related report information, e.g.: interim

³<https://pubmed.ncbi.nlm.nih.gov/>

⁴<https://www.ncbi.nlm.nih.gov/pmc/>

analyses (Broglia et al., 2014) or phase II data for just phase II trials (De Ridder, 2005). Our task is (1) more generalizable, since all potential PICO elements can be represented by free-texts and thus modeled in our work; and (2) aimed at evaluating new clinical trial proposals.

Explicit Evidence Integration: It depends on the existence of structured evidence, i.e.: $\{P, I, C, O, R\}$ (Wallace, 2019). Consequently, collecting such explicit evidence is vital for further analyses, and is also the objective for most relevant works: Some seek to find relevant papers through retrieval (Lee and Sun, 2018); many works are aimed at extracting PICO elements from published literature (Wallace et al., 2016; Singh et al., 2017; Marshall et al., 2017; Jin and Szolovits, 2018; Nye et al., 2018; Zhang et al., 2020); the *evidence inference* task extracts R for a given ICO query using the corresponding clinical trial report (Lehman et al., 2019; DeYoung et al., 2020). However, since getting expert annotations is expensive, these works are typically limited in scale, with only thousands of labeled instances. Few works have been done to utilize the automatically collected structured data for analyses. In this paper, we adopt an end-to-end approach, where we use large-scale pre-training to *implicitly* learn from free-text clinical evidence.

3 The CTRP Task

The CTRP task is motivated to evaluate clinical trial proposals by predicting their results before actually conducting them, as discussed in §1. Therefore, we formulate the task to take as input exactly the information required for proposing a new clinical trial: free-texts of a background description and a PICO query to be investigated. Formally, we denote the strings of the input background as B and PICO elements as $P, I, C,$ and O , respectively. The task output is defined as one of the three possible comparison results: higher (\uparrow), no difference (\rightarrow), or lower (\downarrow) measurement O in intervention group I than in comparison group C for population P . We denote the result as R , and:

$$R(B, P, I, C, O) = \begin{cases} \uparrow & O(I) > O(C) \mid P \\ \downarrow & O(I) < O(C) \mid P \\ \rightarrow & O(I) \sim O(C) \mid P \end{cases}$$

Main metrics include accuracy and 3-way macro-averaged F1. We also use 2-way (\uparrow, \downarrow) macro-averaged F1 to evaluate human expectations (§6.2).

4 Implicit Evidence Integration

In this section, we introduce the Implicit Evidence Integration, which is used to collect pre-training data for comparative language modeling (§5.1).

Instead of collecting explicit evidence with structured $\{B, P, I, C, O - R\}$ information, we utilize a simple observation to collect evidence *implicitly*: clinical evidence is naturally expressed by comparisons, e.g.: “Blood oxygen is higher in the intervention group than in the placebo group”. Free-texts of P, I, C, O and R are entangled with other functional words that connect these elements in such comparative sentences, where R is a free-text version of the structured result R (e.g.: $R = \text{“higher ... than”}$ translates into $R = \uparrow$). We call these sentences entangled implicit evidence and denote them as $E_{ent} = \{PICO R\}$. Then, we disentangle R out of the E_{ent} by heuristics, getting R and the left $E_{dis} = \{PICO\}$. We also include adversarial instances generated from the original ones. Several examples are shown in Table 1.

Details of implicit evidence collection, disentanglement, and adversarial data generation are introduced in §4.1, §4.2 and §4.3, respectively.

4.1 Collection of Implicit Evidence

We collect implicit evidence from PubMed abstracts and PMC articles⁵, where most of the clinical evidence is published. PubMed contains more than 30 million abstracts, and PMC has over 6 million full-length articles. Each abstract is chunked into a background/method section and a result/conclusion section: For the unstructured abstracts, sentences before the first found implicit evidence are included in the background/method section. For the semi-structured abstracts where each section is labeled with a section name, the chunking is done by mapping the section name to either background/method or result/conclusion.

Sentences in abstract result/conclusion sections and main texts that express comparative semantics (Kennedy, 2004) are collected as implicit evidence. They are identified by a pattern detection heuristic, similar to the keyword method described in Jindal and Liu (2006): For expressions of superiority (\uparrow) and inferiority (\downarrow), we detect morpheme patterns of [*more/less/-er ... than ...*]. For expression of equality (\rightarrow), we detect morpheme patterns of [*similar ... to ...*] and [*no difference ... between ... and ...*]. The background/method section serves as the

⁵Articles in downstream experiment datasets are excluded.

E_{ent}	E_{dis}	R	r	R
“Our results also showed that serum TSH levels were slightly higher in the chloroquine group than in the placebo group.”	“Our results also showed that serum TSH levels were [MASK] in the chloroquine group [MASK] in the placebo group.”	“ slightly higher ... than ”	[HIGHER]	↑
“In conclusion, there is no difference between IFN treatment and supportive treatment for MERS patients in terms of mortality.”	“In conclusion, there is [MASK] IFN treatment [MASK] supportive treatment for MERS patients in terms of mortality.”	“ no difference between ... and ”	[NODIFF]	→
“Levels of viral antigen staining in lung sections of GS-5734-treated animals were significantly lower as compared to vehicle-treated animals.”	“Levels of viral antigen staining in lung sections of GS-5734-treated animals were [MASK] vehicle-treated animals.”	“ significantly lower as compared to ”	[LOWER]	↓

Table 1: Several examples of implicit evidence. Red, violet and blue denote superiority, equality and inferiority.

corresponding B for the collected implicit evidence. These sentences are denoted as E_{ent} , which contain entangled PICO-R.

We have collected 11.8 million such sentences. Among them, 2.4 million (20.2%), 3.5 million (29.9%) and 5.9 million (49.9%) express inferiority, equality and superiority respectively.

4.2 Disentanglement of Implicit Evidence

To disentangle the free-text result R from implicit evidence E_{ent} , we mask out the detected morphemes that express comparative semantics (e.g.: “higher than”) as well as other functional tokens that might be exploited by the model to predict the result (e.g.: p values). This generates the masked out result R and the left part E_{dis} ($\{PICO\}$) from E_{ent} ($\{PICOR\}$), i.e.: $R + E_{dis} = E_{ent}$. R is mostly a phrase with a central comparative adjective/adverb (e.g.: “significantly smaller than”) and can be directly mapped to R (↓ for the same example).

Nevertheless, R contains richer information than the sole change direction because of the central adjective/adverb. To utilize such information, we map free-texts of R to a finer-grained result label $r \in \mathcal{C}$ instead of the 3-way direction, where $\mathcal{C} = \{[POORER], [LONGER], [SLOWER], \dots\}$ is a manually-curated vocabulary for such labels and $|\mathcal{C}| = 34$. Each element can be mapped to its antonym in \mathcal{C} by a reversing function Rev: e.g.: $Rev([SMALLER]) = [GREATER]$ and $Rev([NODIFF]) = [NODIFF]$. This enables us to generate adversarial examples used below.

4.3 Adversarial Data Generation

We generate adversarial examples from the original ones using a simple rule of ordering: if the result r holds for the comparison I/C conditioned on P

and O, the reversed result $Rev(r)$ must hold for the reversed comparison C/I on the same condition. This is similar to generate adversarial examples for natural language inference task by logic rules (Minervini and Riedel, 2018; Wang et al., 2019).

However, Since $E_{dis} = \{PICO\}$ is only partially disentangled and P, I, C, O are still in their free-text forms, we cannot explicitly reverse I/C and generate such examples. As an alternative, we reverse the entire sentence order while keeping the word order between any two masked phrases in E_{dis} , getting E_{rev} . For example, if:

$E_{dis} =$ “[Levels of viral antigen staining in lung sections of GS-5734-treated animals were]₀ [MASK] [vehicle-treated animals]₁.”

and $r = [LOWER]$, then the reversed evidence is:

$E_{rev} =$ “[Vehicle-treated animals]₁ [MASK] [levels of viral antigen staining in lung sections of GS-5734-treated animals were]₀.”

and $Rev(r) = [HIGHER]$. This implicitly reverses the ordering direction of I and C without changing the P and O.

5 EBM-Net

We introduce the EBM-Net model in this section. Similar to BERT (Devlin et al., 2019), EBM-Net is essentially a transformer encoder (Vaswani et al., 2017), and follows the pre-training – fine-tuning approach: We pre-train EBM-Net by Comparative Language Modeling (CLM, §5.1) that is designed to learn the conditional ordering function of I/C. The pre-trained EBM-Net is fine-tuned to solve the CTRP task on downstream datasets (§5.2).

5.1 Comparative Language Modeling

We show the CLM architecture in Figure 2. CLM is adapted from the masked language modeling used

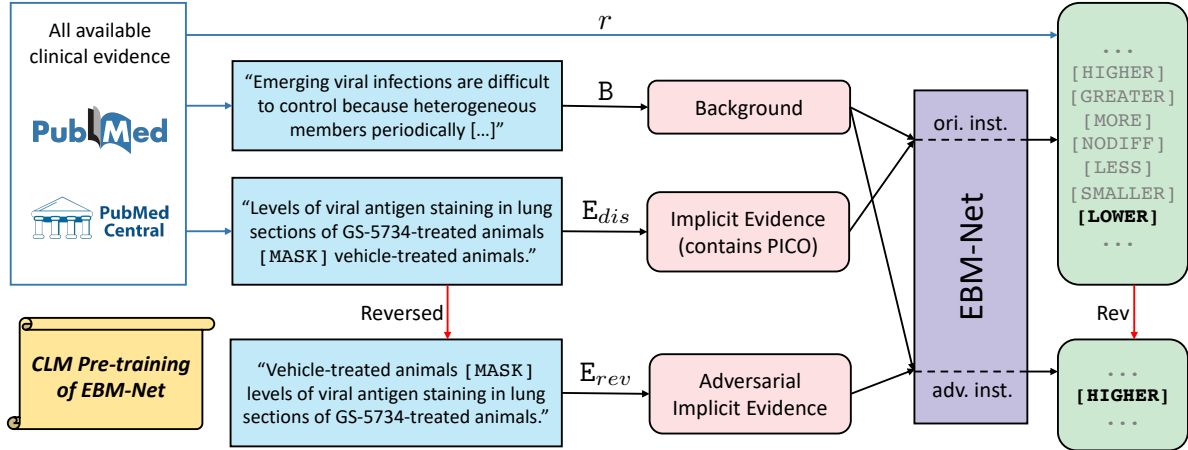


Figure 2: Architecture of CLM pre-training for EBM-Net. (ori.: original; adv.: adversarial; inst.: instance)

in BERT (Devlin et al., 2019), but differentiates from it in that: (1) EBM-Net masks out phrases R that suggest comparative results and predicts a specific set of comparative labels C ; (2) EBM-Net is also pre-trained on adversarial examples generated by comparison rules from the original examples.

During pre-training, EBM-Net takes as input the concatenation of background B and the corresponding partially disentangled implicit evidence E , i.e.: $\text{Input} = [[\text{CLS}], B, [\text{SEP}], E, [\text{SEP}]]$, where $[\text{CLS}]$ and $[\text{SEP}]$ are the special classification and separation tokens used in the original BERT and $E \in \{E_{dis}, E_{rev}\}$. B and E are associated with two different segment types. The special $[\text{MASK}]$ tokens are only used as placeholders for the masked out R . $[\text{CLS}]$ hidden state of the EBM-Net is used to predict the CLM label r with a linear layer followed by a softmax output unit:

$$\hat{r} = \text{SoftMax}(W_1 h_{[\text{CLS}]} + b_1) \in [0, 1]^{|C|}$$

We minimize the cross-entropy between the estimated \hat{r} and the empirical r distribution.

At input-level, the adversarial examples only differ from their original examples in word orders between E_{dis} and E_{rev} . However, their labels are totally reversed from r to $\text{Rev}(r)$. By regularizing the model to learn such conditional ordering function, CLM prevents the pre-trained model from learning unwanted and possibly biased co-occurrences between evident elements and their results.

5.2 CTRP Fine-tuning

During fine-tuning, EBM-Net takes as input the $[[\text{CLS}], B, [\text{SEP}], E_{exp}, [\text{SEP}]]$, where E_{exp} denotes the explicit evidence in the downstream

datasets of the proposed CTRP task. For example, $E_{exp} = [I, [\text{SEP}], O, [\text{SEP}], C]$ on the Evidence Integration dataset (§6.1). The sequence of PICO elements in E_{exp} can be tuned empirically. EBM-Net learns from scratch another linear layer that maps from the predicted CLM label probabilities \hat{r} to 3-way result label R logits. The final predictions are made by a softmax output unit:

$$\hat{R} = \text{SoftMax}(W_2 \hat{r} + b_2) \in [0, 1]^3$$

Cross-entropy between the estimated \hat{R} and the empirical R distribution is minimized in fine-tuning.

5.3 Configuration

The transformer weights of EBM-Net ($L=12$, $H=768$, $A=12$, $\#\text{Params}=110\text{M}$) are initialized with BioBERT (Lee et al., 2020), a variant of BERT that is also pre-trained on PubMed abstracts and PMC articles. The maximum sequence lengths for B , E_{dis} , E_{rev} , E_{exp} are 256, 128, 128, and 128, respectively. We use Adam optimizer (Kingma and Ba, 2014) to minimize the cross-entropy losses. EBM-Net is implemented using Huggingface’s Transformers library (Wolf et al., 2019) in PyTorch (Paszke et al., 2019). Pre-training on 12M implicit evidence takes about 1k Tesla P100 GPU hours.

6 Experiments

6.1 The Evidence Integration Dataset

The Evidence Integration dataset serves as a benchmark for our task. We collect this dataset by repurposing the *evidence inference* dataset (Lehman et al., 2019; DeYoung et al., 2020), which is essentially a machine reading comprehension task for

extracting the structured result (i.e.: R) of a given structured ICO query⁶ from the corresponding clinical trial report article. Since clinical trial reports already contain free-text result descriptions (i.e.: R) of the given ICO, solving the original task does not require the integration of previous clinical evidence. To test such capability for our proposed CTRP task, we remove the result/conclusion part and only keep the background/method part in the input clinical trial report. 34.6% tokens of the original abstracts are removed on average and the remained are used as the clinical trial backgrounds.

Specifically, input of the Evidence Integration dataset includes free texts of ICO elements \mathbb{I} , \mathbb{C} and \mathbb{O} which are the same as the original *evidence inference* dataset, and their clinical trial backgrounds \mathbb{B} . The output is the comparison result R . Following the original dataset split, there are 8,164 instances for training, 1,002 for validation, and 965 for test.

We also do experiments under the adversarial setting, where adversarial examples generated by reversing both the \mathbb{I}/\mathbb{C} order and the R label (similar to §5.1) are added. This setting is used to test model robustness under adversarial attack.

6.2 Compared Methods

We compare to a variety of methods, ranging from trivial ones like Random and Majority to the state-of-the-art BioBERT model. Two major approaches in open-domain question answering (QA) are tested as well: the knowledge base (KB) approach (MeSH ontology) and the text/retrieval approach (Retrieval + Evidence Inference), since solving our task also requires reasoning over a large external corpus. Finally, we introduce some ablation settings and the evaluation of human expectations.

Random: we report the expected performance of randomly predicting the result for each instance.

Majority: we report the performance of predicting the majority class (\rightarrow) for all test instances.ⁿ

Bag-of-Words + Logistic Regression: we concatenate the TF-IDF weighted bag-of-word vectors of \mathbb{B} , \mathbb{P} , \mathbb{I} , \mathbb{C} and \mathbb{O} as features and use logistic regression for learning.

MeSH Ontology: Since no external KB is available for our task, we use the training set as an internal alternative: we map the \mathbb{I} , \mathbb{C} and \mathbb{O} of the test

⁶P is not included in the original dataset as the background of the trial report contains it.

instances to terms in the Medical Subject Headings (MeSH)⁷ ontology by string matching. MeSH is a controlled and hierarchically-organized vocabulary for describing biomedical topics. Then, we find their nearest labeled instances in the training set, where the distance is defined by:

$$d(i, j) = \sum_{e \in \{\mathbb{I}, \mathbb{C}, \mathbb{O}\}} \min \text{TreeDist}(m_i^e, m_j^e)$$

m_i^e and m_j^e are MeSH terms identified in ICO element e of instance i and j , respectively. TreeDist is defined as the number of edges between two nodes on the MeSH tree. The majority label of the nearest training instances is used as the prediction.

Retrieval + Evidence Inference: State-of-the-art method on the *evidence inference* dataset (DeYoung et al., 2020) is a pipeline based on SciBERT (Beltagy et al., 2019): (1) find the exact evidence sentences in the clinical trial report for the given ICO query, using a scoring function derived from a fine-tuned SciBERT; and (2) predict the result R based on the found evidence sentences and the given ICO query by fine-tuning another SciBERT.

Our task needs an additional retrieval step to find relevant documents that might contain useful results of similar trials, as the input trial background does not contain the result information for the given ICO query. Documents are retrieved from the entire PubMed and PMC using a TF-IDF matching between their indexed MeSH terms and the MeSH terms identified in the ICO queries. We then apply the pipeline described above on the retrieved documents. This baseline is similar to but more domain-specific than BERTserini (Yang et al., 2019).

BioBERT: For this setting, we feed BioBERT with similar input to EBM-Net as is described in §5 and fine-tune it to predict the R using its special [CLS] hidden state.

Ablations: We conduct two sets of ablation experiments with EBM-Net: (1) Pre-training level, where we exclude the adversarial examples in pre-training, to analyze the utility of CLM against traditional LM. (2) Input level, where we exclude different input elements (\mathbb{B} , \mathbb{I} , \mathbb{C} , \mathbb{O}) to study their relative importance.

Human Expectations: We define the expected result (R_e) of a clinical trial (e.g.: $R_e = \downarrow$ for $\mathbb{O} =$ “mortality rate”) as the Human Expectation (HE),

⁷<https://www.nlm.nih.gov/mesh>

Model	Standard Evidence Integration			Adversarial Evidence Integration			$ \Delta $
	Accuracy	F1 (3-way)	F1 (2-way)	Accuracy	F1 (3-way)	F1 (2-way)	
Majority (\rightarrow)	41.76	19.64	–	41.76	19.64	–	–
Random (expected)	33.33	32.77	30.62	33.33	32.77	30.62	–
BoW + Logistic Regression	43.73	41.04	35.84	41.97	39.87	34.01	4.0
MeSH Ontology	38.55	36.33	31.01	34.46	33.19	34.77	10.6
Retrieval + Evidence Inference (DeYoung et al., 2020)	50.57	49.91	48.30	50.62	50.13	48.46	0.0
BioBERT (Lee et al., 2020)	55.96	54.33	51.98	53.11	52.84	51.59	5.1
EBM-Net (ours)	61.35	60.15	59.42	59.59	59.36	58.67	2.7
w/o adversarial pre-training	60.73	59.04	58.52	58.91	58.81	58.34	3.0
w/o B (background)	55.65	54.31	52.48	53.83	53.32	51.26	3.3
w/o I (intervention)	59.59	58.59	58.08	57.30	56.74	54.87	3.8
w/o C (comparison)	57.72	56.77	56.15	57.51	57.10	55.47	0.4
w/o O (outcome)	48.91	44.88	39.57	47.31	46.40	43.66	3.3
Human Expectations	56.79	–	68.86	56.79	–	68.86	–

Table 2: Main results on the benchmark Evidence Integration dataset. $|\Delta|$ denotes the absolute value of relative accuracy decrease from the standard to the adversarial setting. All numbers are percentages. (w/o: without)

which is the underlying motivation for conducting the corresponding trial. Generally, $R_e \in \{\uparrow, \downarrow\}$ since significant results are expected. To make fair comparisons, we use the 2-way macro-average F1: $F1(2\text{-way}) = (F1(\uparrow) + F1(\downarrow))/2$ as a main metric for evaluations of HE. HE performance is an overestimation of human performance: main biases are due to the shift of input trial distribution from the targeted proposal stage to the actual report stage, which contains fewer trials with unexpected results.

6.3 Main Results

Table 2 shows the main results on the Evidence Integration dataset, where accuracy and F1 (3-way) are used to compare model performance and F1 (2-way) is used for evaluating human expectations.

Results show that EBM-Net outperforms other baselines by large margins in both standard and adversarial settings. While being the strongest baseline, BioBERT is 10.7% relatively lower in macro-F1 (54.33% v.s. 60.15%) and 9.6% relatively lower in accuracy (55.96% v.s. 61.35%) than EBM-Net. The open-domain QA baselines perform even worse: for the MeSH Ontology method, the internal KB of only 8k entries is far from complete; for the Retrieval + Evidence Inference method, the PICO queries are so specific that no exactly relevant evidence can be found in other trials and retrieving only a few trials has limited utilities.

We use $|\Delta|$, the absolute value of relative accuracy decrease to measure model robustness under adversarial attacks. The higher the $|\Delta|$, the more vulnerable a model is. BioBERT has about twice

as much (5.1% v.s. 2.7%) $|\Delta|$ in the adversarial setting as EBM-Net does. It suggests that EBM-Net is more robust to adversarial attacks, which is a vital property for healthcare applications. EBM-Net without adversarial pre-training is less robust than EBM-Net as well (3.0% v.s. 2.7%), but not as vulnerable as BioBERT, indicating that robustness can be learned by pre-training with original implicit evidence to some extent and further consolidated by the adversarial evidence.

Unsurprisingly, EBM-Net with full input consistently outperforms all input-level ablations. Among them, O is the most important input element as the performance decreases dramatically on its ablation. This is expected as O is the standard of comparisons. B is the second most important element, since B contains methodological details of how the clinical trials will be conducted, which is also vital for result prediction. The performance does not decrease as much without I or C, since there is redundant information of them in B.

On the one hand, the accuracy of EBM-Net surpasses that of HE, mainly because the latter is practically a 2-way classifier. On the other hand, HE outperforms EBM-Net in terms of 2-way F1, but is still unsatisfying (68.86%). This suggests that the proposed CTRP task is hard and there is still room for further improvements.

6.4 Discussions

We study how different numbers of pre-training and fine-tuning instances influence the EBM-Net performance, in comparison to the BioBERT. Figure 3




Figure 3: **Left:** EBM-Net 3-way macro-F1 v.s. pre-training sizes compared to BioBERT; **Right:** EBM-Net and BioBERT 3-way macro-F1 v.s. fine-tuning sizes.

shows the results: **(Left)** The final performance of EBM-Net improves log-linearly as the pre-training dataset size increases, suggesting that there can be further improvements if more data is collected for pre-training but the marginal utility might be small. EBM-Net surpasses BioBERT when pre-trained by about 50k to 100k instances of implicit evidence, which are 5 to 10 times as many as the fine-tuning instances. **(Right)** EBM-Net is more robust in a few-shot learning setting: using only 10% of the training data, EBM-Net outperforms BioBERT fine-tuned with 100% of the training data. From zero-shot⁸ to using all the training data, EBM-Net improves only by 26.6% relative F1 (from 47.52% to 60.15%) while BioBERT improves largely by 60.0% relative F1 (from 32.77% to 54.33%).

We use t-SNE (Maaten and Hinton, 2008) to visualize the test instance representations derived from EBM-Net [CLS] hidden state in Figure 4. It shows that EBM-Net effectively learns the relationships between comparative results: the points cluster into three results (\uparrow , \downarrow , \rightarrow). While there is a clear boundary between the \downarrow cluster (dashed-blue circle) and the \uparrow cluster (dashed-red circle), the boundaries between the \rightarrow cluster (dashed-black circle) and the other two are relatively vague. It suggests that the learnt manifold follows a quantitatively continuous “ $\downarrow - \rightarrow - \uparrow$ ” pattern.

Out of the 373 mistakes EBM-Net makes on the test set, significantly less (11.8%, $p < 0.001$ by permutation test) predictions are opposite to the ground-truth (e.g.: predicting \uparrow when the label is \downarrow), also suggesting that EBM-Net effectively learn the relationship between comparison results. In addition, we notice that there is a considerable proportion of instances whose results are not predictable

⁸Zero-shot performance of BioBERT is defined as the expected results from random predictions.




Figure 4: T-SNE visualizations of EBM-Net representations of Evidence Integration test set instances. Red colored \blacktriangle , blue colored \blacktriangledown , and green colored \bullet refer to the corresponding R equaling \uparrow , \downarrow and \rightarrow , respectively.

without their exact reports. For example, some \uparrow and \downarrow differ only quantitatively, e.g.: “4% lidocaine” and “2% lidocaine”, and modeling such differences is beyond the scope of our task.

6.5 Validation on COVID-19 Clinical Trials

For analyzing COVID-19 related clinical trials, we further pre-train EBM-Net on the COVID-19 dataset (Wang et al., 2020a)⁹, also using the comparative language modeling (§5.1). It leads to a COVID-19 specific EBM-Net that is used in this section.

We use leave-one-out validation to evaluate EBM-Net on the 22 completed clinical trials in COVID-evidence¹⁰, which is an expert-curated database of available evidence on interventions for COVID-19. Again, EBM-Net outperforms BioBERT by a large margin (59.1% v.s. 50.0% accuracy). Expectedly, their 3-way F1 results (45.5% v.s. 36.1%) are close to those in the zero-shot learning setting since not many trials have finished. Accuracy and 2-way F1 performance of HE are 54.5% and 68.9%, and are close to those in Table 2. These further confirm the performance improvement of EBM-Net and the difficulty of the CTRP task.

7 Conclusions

In this paper, we introduce a novel task, CTRP, to predict clinical trial results without actually doing them. Instead of using structured evidence that is prohibitively expensive to annotate, we heuristically collect 12M unstructured sentences

⁹The 05/12/2020 version.

¹⁰covid-evidence.org (visited on 05/18/2020).

as implicit evidence, and use large-scale CLM pre-training to learn the conditional ordering function required for solving the CTRP task. Our EBM-Net model outperforms other strong baselines on the Evidence Integration dataset and is also validated on COVID-19 clinical trials.

Acknowledgement

We would like to thank the anonymous reviewers of EMNLP 2020 for their constructive comments. We are also grateful for Ning Ding, Yuxuan Lai, Yijia Liu, Yao Fu, Kun Liu and Rui Wang for helpful discussions at Alibaba DAMO Academy.

References

- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. Scibert: A pretrained language model for scientific text. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3606–3611.
- Kristine R Broglio, David N Stivers, and Donald A Berry. 2014. Predicting clinical trial results based on announcements of interim analyses. *Trials*, 15(1):73.
- Filip De Ridder. 2005. Predicting the outcome of phase iii trials using phase ii data: a case study of clinical trial simulation in late stage drug development. *Basic & clinical pharmacology & toxicology*, 96(3):235–241.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jay DeYoung, Eric Lehman, Ben Nye, Iain J. Marshall, and Byron C. Wallace. 2020. [Evidence inference 2.0: More data, better models](#).
- Kaitlyn M Gayvert, Neel S Madhukar, and Olivier Elemento. 2016. A data-driven approach to predicting successes and failures of clinical trials. *Cell chemical biology*, 23(10):1294–1301.
- N Holford, SC Ma, and BA Ploeger. 2010. Clinical trial simulation: a review. *Clinical Pharmacology & Therapeutics*, 88(2):166–182.
- N. H. G. Holford, H. C. Kimko, J. P. R. Monteleone, and C. C. Peck. 2000. [Simulation of clinical trials](#). *Annual Review of Pharmacology and Toxicology*, 40(1):209–234. PMID: 10836134.
- Xiaoli Huang, Jimmy Lin, and Dina Demner-Fushman. 2006. Evaluation of pico as a knowledge representation for clinical questions. In *AMIA annual symposium proceedings*, volume 2006, page 359. American Medical Informatics Association.
- Di Jin and Peter Szolovits. 2018. [PICO element detection in medical text via long short-term memory neural networks](#). In *Proceedings of the BioNLP 2018 workshop*, pages 67–75, Melbourne, Australia. Association for Computational Linguistics.
- Nitin Jindal and Bing Liu. 2006. Mining comparative sentences and relations. In *Proceedings of the 21st National Conference on Artificial Intelligence - Volume 2, AAAI’06*, page 1331–1336. AAAI Press.
- Christopher Kennedy. 2004. Comparatives, semantics of. *Concise Encyclopedia of Philosophy of Language and Linguistics*, pages 68–71.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Grace E. Lee and Aixin Sun. 2018. [Seed-driven document ranking for systematic reviews in evidence-based medicine](#). In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR ’18*, page 455–464, New York, NY, USA. Association for Computing Machinery.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240.
- Eric Lehman, Jay DeYoung, Regina Barzilay, and Byron C. Wallace. 2019. [Inferring which medical treatments work from reports of clinical trials](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3705–3717, Minneapolis, Minnesota. Association for Computational Linguistics.
- Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605.
- Iain Marshall, Joël Kuiper, Edward Banner, and Byron C. Wallace. 2017. [Automating biomedical evidence synthesis: RobotReviewer](#). In *Proceedings of ACL 2017, System Demonstrations*, pages 7–12, Vancouver, Canada. Association for Computational Linguistics.
- Mandeep R Mehra, Sapan S Desai, Frank Ruschitzka, and Amit N Patel. 2020. Hydroxychloroquine or chloroquine with or without a macrolide for treatment of covid-19: a multinational registry analysis. *The Lancet*.

- Pasquale Minervini and Sebastian Riedel. 2018. Adversarially regularising neural nli models to integrate logical background knowledge. In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 65–74.
- Mostafa Ebraheem Morra, Le Van Thanh, Mohamed Gomaa Kamel, Ahmed Abdelmotaleb Ghazy, Ahmed MA Altibi, Lu Minh Dat, Tran Ngoc Xuan Thy, Nguyen Lam Vuong, Mostafa Reda Mostafa, Sarah Ibrahim Ahmed, et al. 2018. Clinical outcomes of current medical approaches for middle east respiratory syndrome: A systematic review and meta-analysis. *Reviews in medical virology*, 28(3):e1977.
- Benjamin Nye, Junyi Jessy Li, Roma Patel, Yinfei Yang, Iain Marshall, Ani Nenkova, and Byron Wallace. 2018. A corpus with multi-level annotations of patients, interventions and outcomes to support language processing for medical literature. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 197–207, Melbourne, Australia. Association for Computational Linguistics.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. *Pytorch: An imperative style, high-performance deep learning library*. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- P Peymani, S Ghavami, B Yeganeh, R Tabrizi, S Sabour, B Geramizadeh, MR Fattahi, SM Ahmadi, and KB Lankarani. 2016. Effect of chloroquine on some clinical and biochemical parameters in non-response chronic hepatitis c virus infection patients: pilot clinical trial. *Acta bio-medica: Atenei Parmensis*, 87(1):46.
- David L Sackett. 1997. Evidence-based medicine. In *Seminars in perinatology*, pages 3–5. Elsevier.
- Timothy P Sheahan, Amy C Sims, Rachel L Graham, Vineet D Menachery, Lisa E Gralinski, James B Case, Sarah R Leist, Krzysztof Pyrc, Joy Y Feng, Iva Trantcheva, et al. 2017. Broad-spectrum antiviral gs-5734 inhibits both epidemic and zoonotic coronaviruses. *Science translational medicine*, 9(396).
- Gaurav Singh, Iain J. Marshall, James Thomas, John Shawe-Taylor, and Byron C. Wallace. 2017. A neural candidate-selector architecture for automatic structured clinical text annotation. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM '17*, page 1519–1528, New York, NY, USA. Association for Computing Machinery.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Byron C. Wallace. 2019. What does the evidence say? models to help make sense of the biomedical literature. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 6416–6420. International Joint Conferences on Artificial Intelligence Organization.
- Byron C. Wallace, Joël Kuiper, Aakash Sharma, Mingxi (Brian) Zhu, and Iain J. Marshall. 2016. Extracting pico sentences from clinical trial reports using supervised distant supervision. *Journal of Machine Learning Research*, 17(132):1–25.
- Haohan Wang, Da Sun, and Eric P Xing. 2019. What if we simply swap the two text fragments? a straightforward yet effective way to test the robustness of methods to confounding signals in nature language inference tasks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7136–7143.
- Lucy Lu Wang, Kyle Lo, Yoganand Chandrasekhar, Russell Reas, Jiangjiang Yang, Darrin Eide, Kathryn Funk, Rodney Kinney, Ziyang Liu, William Merrill, et al. 2020a. Cord-19: The covid-19 open research dataset. *arXiv preprint arXiv:2004.10706*.
- Yeming Wang, Dingyu Zhang, Guanhua Du, Ronghui Du, Jianping Zhao, Yang Jin, Shouzhi Fu, Ling Gao, Zhenshun Cheng, Qiaofa Lu, et al. 2020b. Remdesivir in adults with severe covid-19: a randomised, double-blind, placebo-controlled, multicentre trial. *The Lancet*.
- WHO. 2020. Solidarity clinical trial for covid-19 treatments.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.
- Chi Heem Wong, Kien Wei Siah, and Andrew W Lo. 2019. Estimation of clinical trial success rates and related parameters. *Biostatistics*, 20(2):273–286.
- Wei Yang, Yuqing Xie, Aileen Lin, Xingyu Li, Luchen Tan, Kun Xiong, Ming Li, and Jimmy Lin. 2019. End-to-end open-domain question answering with BERTserini. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 72–77, Minneapolis, Minnesota. Association for Computational Linguistics.

Tengteng Zhang, Yiqin Yu, Jing Mei, Zefang Tang, Xiang Zhang, and Shaochun Li. 2020. Unlocking the power of deep pico extraction: Step-wise medical ner identification. *arXiv preprint arXiv:2005.06601*.

A Collection of Implicit Evidence

We show several mapping examples from section names to background/method or result/conclusion in Table 4. The implicit evidence collection algorithm is shown in Algorithm 1, which uses the evidence detection algorithm described in Algorithm 2. Dataset statistics are shown in Table 3.

B Disentanglement of Implicit Evidence

Disentanglement of implicit evidence seeks to mask out the R in the implicit evidence and map it to r , which is a finer-grained label of comparison results. We show the distribution of the collected r in Figure 5.

When disentangling the free-text result R , we also mask out the following words:

- Numbers (e.g.: p-values);
- tokens within parentheses (e.g.: interpretations of results);
- Adverb before the central comparative adv/adj (e.g.: “significantly”).

which can be exploited by the model to predict the r during pre-training.

C Comparative Language Modeling

In Table 5, we show several originally collected instances E_{dis} , r , together with their adversarially reversed instances E_{rev} , $Rev(r)$.

D The Evidence Integration Dataset

In Table 6, we show several examples of Evidence Integration instances and their corresponding ones in *evidence inference*. The Evidence Integration dataset statistics are also shown in Table 3.

E Settings

We show the searched hyper-parameters and their chosen values of different experiments in Table 7. We manually tune the hyper-parameters where the best combination is chosen based on macro-F1 metric in the validation set. The number of hyper-parameter search trials is about 100. Training time

of EBM-Net is about 3 min/epoch of standard Evidence Integration and 6 min/epoch of adversarial Evidence Integration on 2 Tesla P100 GPUs at the optimal hyper-parameter setting, and the Inference time is about 100 instances/s on 1 Tesla P100 GPU.

All evaluation metrics are calculated by the `sklearn.metrics` package in Python.

F Results

Validation results of EBM-Net in comparison to BioBERT are shown in Table 8. Results on 22 completed clinical trials in COVID-evidence dataset are shown in Table 9.

Statistic	Pre-training (original)	Standard Evidence Integration (training)
Avg. Length of B	129.0	182.8
Avg. Length of E	16.8	16.7*
Avg. Length of I	NA	5.7
Avg. Length of C	NA	3.7
Avg. Length of O	NA	5.3
% of \uparrow	49.9	31.4
% of \rightarrow	29.9	44.3
% of \downarrow	20.2	24.3

Table 3: This table shows the statistics of the pre-training and the downstream Evidence Integration dataset. *Concatenation of I, C, O and two [SEP] tokens.

Section Name	Section Type
“BACKGROUND”	background/method
“METHODS”	background/method
“OBJECTIVES”	background/method
“INTRODUCTION”	background/method
“DESIGN”	background/method
“RESULTS”	result/conclusion
“CONCLUSIONS”	result/conclusion
“FINDINGS”	result/conclusion
“DISCUSSIONS”	result/conclusion
“SIGNIFICANCE”	result/conclusion

Table 4: Several examples of section name to background/method or result/conclusion mapping.




Figure 5: The frequency (in log scale) of different comparative language modeling labels. Red, green and blue colors denote the corresponding label expressing superiority, equality and inferiority, respectively.

E_{dis}	r	E_{rev}	$Rev(r)$
“Our results also showed that serum TSH levels were [MASK] in the chloroquine group [MASK] in the placebo group.”	[HIGHER]	“In the placebo group [MASK] in the chloroquine group [MASK] our results also showed that serum TSH levels were.”	[LOWER]
“In conclusion, there is [MASK] IFN treatment [MASK] supportive treatment for MERS patients in terms of mortality.”	[NODIFF]	“Supportive treatment for MERS patients in terms of mortality [MASK] IFN treatment [MASK] in conclusion, there is”	[NODIFF]
“Levels of viral antigen staining in lung sections of GS-5734-treated animals were [MASK] vehicle-treated animals.”	[LOWER]	“Vehicle-treated animals [MASK] levels of viral antigen staining in lung sections of GS-5734-treated animals were.”	[HIGHER]

Table 5: Several examples of comparative language modeling instances.

Algorithm 1 Implicit Evidence Collection

Input: An abstract A which is a sequence of words (from PubMed).

Output: A list of collected implicit evidence E and its background B from the input article.

$E \leftarrow []$

$B \leftarrow ""$

if A is structured **then**

 let $A = [[S_1, L_1], [S_2, L_2], \dots]$ where S_i is the i -th sentence and L_i is its section type label

for $[S_i, L_i]$ in A **do**

if $L_i = \text{background/method}$ **then**

$B \leftarrow B + S_i$

else

if $\text{EvidenceDetector}(S_i) \neq \text{False}$ **then**

$E.\text{append}(\text{EvidenceDetector}(S_i))$

end if

end if

end for

else

 let $A = [S_1, S_2, \dots]$ where S_i is the i -th sentence

$BG \leftarrow \text{True}$

 # BG controls whether a sentence is included in the background

for S_i in A **do**

if $\text{EvidenceDetector}(S_i) \neq \text{False}$ **then**

$BG \leftarrow \text{False}$

$E.\text{append}(\text{EvidenceDetector}(S_i))$

else

if BG **then**

$B \leftarrow B + S_i$

end if

end if

end for

end if

if Full article F is available in PMC **then**

 let $F = [s_1, s_2, \dots]$ where s_i is the i -th sentence

for s_i in A **do**

if $\text{EvidenceDetector}(s_i) \neq \text{False}$ **then**

$E.\text{append}(\text{EvidenceDetector}(s_i))$

end if

end for

end if

return E, B

Algorithm 2 EvidenceDetector

Input: A sentence S which is a sequence of words; Sets of adj/adv that suggest superiority (\mathcal{H}), inferiority (\mathcal{L}) and equality (\mathcal{E})

Output: False if the sentence is not a piece of implicit evidence; Central comparison word c and the change direction R if the sentence is a piece of implicit evidence.

if “than” in S **then**

if $S \cap \mathcal{H} \neq \emptyset$ and $S \cap \mathcal{L} = \emptyset$ **then**

$c \leftarrow S \cap \mathcal{H}$

$R \leftarrow “\uparrow”$

return c, R

else if $S \cap \mathcal{H} = \emptyset$ and $S \cap \mathcal{L} \neq \emptyset$ **then**

$c \leftarrow S \cap \mathcal{L}$

$R \leftarrow “\downarrow”$

return c, R

else

return False

end if

else if $S \cap \mathcal{E} \neq \emptyset$ **then**

$c \leftarrow S \cap \mathcal{E}$

$R \leftarrow “\rightarrow”$

return c, R

else

return False

end if

B	I	C	O	R
<p>“Background: Self-management programs for patients with heart failure can reduce hospitalizations and mortality. However, no programs have analyzed their usefulness for patients with low literacy. We compared the efficacy of a heart failure self-management program designed for patients with low literacy versus usual care. Methods: We performed a 12-month randomized controlled trial. From November 2001 to April 2003, we enrolled participants aged 30–80, who had heart failure and took furosemide. Intervention patients received education on self-care emphasizing daily weight measurement, diuretic dose self-adjustment, and symptom recognition and response. Picture-based educational materials, a digital scale, and scheduled telephone follow-up were provided to reinforce adherence. Control patients received a generic heart failure brochure and usual care. Primary outcomes were combined hospitalization or death, and heart failure-related quality of life. Results: 123 patients (64 control, 59 intervention) participated; 41% had inadequate literacy. Patients in the intervention group had a lower rate of hospitalization or death (crude incidence rate ratio (IRR) = 0.69; CI 0.4, 1.2; adjusted IRR = 0.53; CI 0.32, 0.89). This difference was larger for patients with low literacy (IRR = 0.39; CI 0.16, 0.91) than for higher literacy (IRR = 0.56; CI 0.3, 1.04), but the interaction was not statistically significant. At 12 months, more patients in the intervention group reported monitoring weights daily (79% vs. 29%, $p < 0.0001$). After adjusting for baseline demographic and treatment differences, we found no difference in heart failure-related quality of life at 12 months (difference = -2; CI -5, +9). Conclusion: A primary care-based heart failure self-management program designed for patients with low literacy reduces the risk of hospitalizations or death.”</p>	<p>“Follow-up and thorough education on self-care”</p>	<p>“Standard information about self-care”</p>	<p>“Knowledge ↑ about heart failure”</p>	
<p>“Background: Afghanistan’s national guidelines recommend chloroquine for the treatment of Plasmodium vivax infection, the parasite responsible for the majority of its malaria burden. Chloroquine resistance in P. vivax is emerging in Asia. Therapeutic responses across Afghanistan have not been evaluated in detail. Methods: Between July 2007 and February 2009, an open-label, randomized controlled trial of chloroquine and dihydroartemisinin-piperaquine in patients aged three months and over with slide-confirmed P. vivax mono-infections was conducted. Consistent with current national guidelines, primaquine was not administered. Subjects were followed up daily during the acute phase of illness (days 0-3) and weekly until day 56. The primary endpoint was the overall cumulative parasitological failure rate at day 56 after the start of treatment, with the hypothesis being that dihydroartemisinin-piperaquine was non-inferior compared to chloroquine ($\delta = 5\%$ difference in proportion of failures). Results: Of 2,182 individuals with positive blood films for P. vivax, 536 were enrolled in the trial. The day-28 cure rate was 100% in both treatment groups. Parasite clearance was more rapid with dihydroartemisinin-piperaquine than chloroquine. At day 56, there were more recurrent infections in the chloroquine arm (8.9%, 95% CI 6.0-13.1%) than the dihydroartemisinin-piperaquine arm (2.8%, 95% CI 1.4-5.8%), a difference in cumulative recurrence rate of 6.1% (2-sided 90%CI +2.6 to +9.7%). The log-rank test comparing the survival curves confirmed the superiority of dihydroartemisinin-piperaquine over chloroquine ($p = 0.003$). Multivariate analysis showed that a lower initial haemoglobin concentration was also independently associated with recurrence. Both regimens were well tolerated and no serious adverse events were reported. Conclusions: Chloroquine remains an efficacious treatment for the treatment of vivax malaria in Afghanistan. In a setting where radical therapy cannot be administered, dihydroartemisinin-piperaquine provides additional benefit in terms of post-treatment prophylaxis, reducing the incidence of recurrence from 4-8 weeks after treatment. Trial Registration: The trial was registered at ClinicalTrials.gov under identifier NCT00682578.”</p>	<p>“Dihydroartemisinin - piperaquine”</p>	<p>“Chloroquine”</p>	<p>“Parasite clearance at day 2”</p>	<p>→</p>

Table 6: Several examples of dataset instances. Strickethroughed texts are in the original *evidence inference* dataset but not in the Evidence Integration dataset.

Hyper-parameter	Pre-training	Standard Evid. Integ.	Adversarial Evid. Integ.
# Training Epochs	4	8, 10, 12, 16, 20, 24	8, 10, 12, 16, 20 , 24
Optimizer	Adam	Adam	Adam
Adam Learning Rate	5e-5	2e-5, 3e-5, 4e-5 5e-5	2e-5, 3e-5, 4e-5 5e-5
Adam Epsilon	1e-8	1e-8	1e-8
Learning Rate Schedule	Cosine	Cosine , Linear	Cosine , Linear
Warm-up Steps	10000	400, 600 , 800, 1000, 1200	400, 600 , 800, 1000, 1200
Max. Gradient Norm	1.0	1.0	1.0
Batch Size	144	96	96
I,C,O Sequence	NA	I-O-C , I-C-O, O-I-C	I-O-C

Table 7: Hyper-parameters of different experiments. Searched hyper-parameters are listed and the bolded ones denote the optimal. We have not tuned the hyper-parameters of the pre-training due to high computation costs.

Model	Standard Evidence Integration			Adversarial Evidence Integration		
	Accuracy	F1 (3-way)	F1 (2-way)	Accuracy	F1 (3-way)	F1 (2-way)
BioBERT	54.29	53.24	51.17	54.24	53.61	51.35
EBM-Net (ours)	59.78	58.99	57.95	59.48	58.94	57.41

Table 8: Validation results on the benchmark Evidence Integration dataset. All numbers are percentages.

Model	COVID-evidence		
	Acc	F1 (3)	F1 (2)
Majority (\rightarrow)	45.45	20.83	–
Random (expected)	33.33	31.58	28.54
BoW + Logistic Regression	45.45	32.00	20.00
BioBERT	50.00	36.11	30.43
EBM-Net (ours)	59.09	45.51	30.77
Human Expectations	54.54	–	68.87

Table 9: Prediction performance on the completed COVID-19 clinical trials. All numbers are percentages.