Deep Learning-enabled Natural Language Processing to Identify Directional Pharmacokinetic Drug-Drug Interactions

Joel Zirkle, Xiaomei Han, Rebecca Racz, Mohammadreza Samieegohar, Anik Chaturbedi, John Mann, Shilpa Chakravartula, and Zhihua Li
Division of Applied Regulatory Science, Office of Clinical Pharmacology, Office of Translational Sciences, Center for Drug Evaluation and Research, U.S. Food and Drug Administration

Abstract

During drug development, it is essential to gather information about the change in clinical exposure of a drug (the object) due to the pharmacokinetic (PK) drug-drug interactions (DDIs) with another drug (the precipitant). While many natural language processing (NLP) methods for DDI have been published, most of them were designed to evaluate if (and what kind of) DDI relationships exist in the text, without identifying the direction of DDI (object vs. precipitant drug). This makes it impossible to use these methods for automatic extraction of clinical exposure changes of a given object drug directly from literature or drug labels.

The Text Analysis Conference (TAC) DDI track 2019 corpus was reannotated for identifying the direction of a PK DDI and used to fine-tune a BioBERT language model on this task by following the training and validation steps prespecified by TAC. This initial attempt showed the model achieved an F-score of 0.82 in identifying sentences that contain PK DDI and an F-score of 0.97 in identifying object vs. precipitant drugs in those sentences.

Introduction

One example demonstrating the need for automatic methods for NLP of DDI information is the identification of change in clinical exposures of an object drug due to other precipitant drugs. This kind of PK DDI information is important in a clinical setting, but also critical during drug development. For example: in evaluating a drug's potential to cause QT prolongation, clinical and nonclinical studies are required by international regulatory guidelines to cover the high clinical exposure scenario. Given a specific drug of interest (the object drug), gathering information from existing biomedical literature and regulatory labels about all other drugs (precipitant drugs) that could change the object drug's clinical exposure through DDI is an important step towards establishing its high clinical exposure.

There have been several community-wide initiatives to develop NLP methods to extract DDIs from literature and drug labels, for example the Text Analysis Conference (TAC) DDI tracks in 2018 and 2019. However, these methods are usually only capable of identifying sentences that involve DDI and then classifying which kind of DDI is involved (e.g. PK or PD). For example, given the task of "identify all DDIs where clinical exposure of verapamil is changed by another drug from natural language text", most published methods can only perform sentence classification: screen all input sentences and identify those that describe DDI relations with verapamil. Verapamil (as a potent modulator of cytochrome P450 enzymes and P-glycoprotein) may be involved in many DDIs as either the precipitant or the object drug, and the given task requires an additional step to screen the DDI sentences with the "correct" direction: those that describe verapamil as the object drug. Here we provide a complete solution to the task of identifying precipitant vs object drugs in PK DDIs from natural language text (see Figure 1 for example sentences).

We have developed a deep-learning model that is capable of identifying the directionality of PK DDIs from natural language text (e.g. from FDA drug labels).

Example 1. Verapamil does not induce peripheral arterial spasm.

Example 2. Verapamil therapy may increase serum levels of cyclosporine.

Example 3. Therapy with rifampin may markedly reduce oral verapamil bioavailability.

Figure 1. Example sentences as input to model. Example 1: no PK DDI detected. Example 2: PK DDI detected with verapamil identified as precipitant and cyclosporine as object. Example 3: PK DDI detected with rifampin as precipitant and verapamil as object.

Materials and Methods

Our model is based on the neural network language model BERT and is designed to perform two sequential steps (Figure 2):

1. Sentence Classification: identify all sentences that involve PK DDI
2. Named Entity Recognition: label object drugs and precipitant drugs in sentences containing PK DDI.

We re-annotated training and testing/validation data as provided by the TAC 2019 DDI track. Specifically 381 structured product labels (21593 sentences) were annotated as either PK DDI or non-PK DDI as training data. Separately, 81 structured product labels (10592 sentences) were annotated for testing/validation. On top of sentence-level annotations, each of the PK DDI sentences also have entity-level annotations: precipitant and object entities.

The model was built using BioBERT-Large v1.1 and trained using Tensorflow v2, the resulting model we have designated as BioBERT_directionalDDI. Performance of our model was subsequently evaluated using the validation data. Multiple independent models were run from random seeds to ensure that model performance was not an outlier.

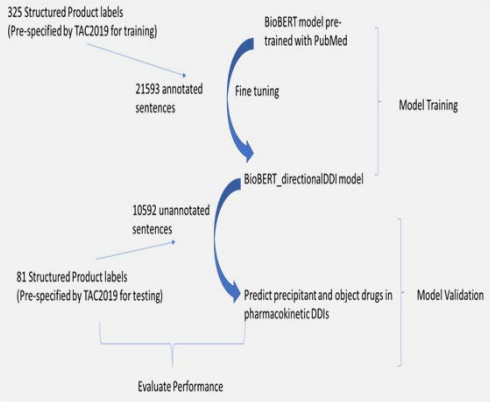

Figure 2. Training and validation procedure

Results and Discussion

Model performance was evaluated on the pre-specified validation data as provided by TAC 2019. As the model involves two sequential steps, performance metrics (F-score, precision and recall) are provided for each step.

For the sentence classification task, our model resulted in a precision of 82.7%, a recall of 80.6% and an F-score of 81.6% (Table 1). This suggests that, for all sentences that actually carry PK DDI information, about 81% will be correctly classified by the model. The remaining 19% will be mistakenly classified as non-PK DDI (meaning either no DDI or other types of DDI such as PD).

For the second step of identify object vs precipitant drugs, our model resulted in a precision of 100% for both object and precipitant entities (no false positives). The recall for object entities was 93.7% and for precipitant entities it was 94.6%. The F-score for object entities was 96.7% and for precipitant entities it was 97.2% (Table 2). Hence about 94% of all entities (object and precipitants) are correctly identified by the model. Such high precision and recall suggest that, given a PK DDI sentence, it is very likely that our model will correctly identify the object and precipitant drugs.

Recall	80.6%
Precision	82.7%
F score	81.6%

Table 1. Performance of the sentence classification step.

	Object	Precipitant
Recall	93.7%	94.6%
Precision	100%	100%
F score	96.8%	97.2%

Table 2. Performance of the named entity recognition step (identifying object and precipitant drugs in PK DDI sentences).

Example use-case: CiPA drugs

CiPA (Comprehensive in vitro Proarrhythmia Assay) is a cardiac safety paradigm. All FDA human prescription drug labels (over 46000 SPLs) were downloaded (as of 3/15/2023) and sentences were extracted. These sentences were fed into our BioBERT_directionalDDI model to find all sentences where verapamil is the object drug in a PK DDI (Table 3).

CiPA drug	total sentences	PK DDI sentences	object sentences
vandetanib	923	132	16
sotalol	16658	121	42
quinidine	31266	7682	1583
ibutilide	688	0	0
dofetilide	8346	645	601
disopyramide	6087	101	49
bepridil	149	0	0
terfenadine	6135	1098	629
risperidone	68070	3783	2250
pimozide	13493	2647	1818
ondansetron	77681	2511	1234
droperidol	1939	16	0
clozapine	16343	2774	1773
clarithromycin	90366	10775	2464
cisapride	6433	1167	1001
chlorpromazine	12207	1103	101
astemizole	2431	250	244
verapamil	43977	7691	1504
tamoxifen	27560	961	401
ranolazine	11979	1932	620
nitrendipine	190	0	0
nifedipine	40859	6850	3714
mexiletine	5359	725	252
metoprolol	97595	3986	2790
loratadine	5875	436	194
diltiazem	45605	10797	1830

Table 3. Our model applied to sentences extracted from all FDA drug labels on CiPA drugs. Columns: total sentences is the number of sentences that contain the drug name. PK DDI sentences are the number of sentences that contain PK DDI information, as identified by the model. Object sentences are the number of sentences that the model identified where the CiPA drug is the object in the PK DDI.

Conclusion

To identify the clinical exposure change due to PK DDI from a sentence there are naturally two steps:

1. Identify sentences that carry DDI information in the PK category
2. Identify the precipitant and objects drugs in those sentences.

Due to the lack of an open-source model to accomplish this task, we have developed a model that is capable of reading natural language sentences and detected which sentences contain PK DDI information and then identify which drugs in those sentences are object drugs and which are precipitant drugs. Our model and training and validation data can be found at: https://github.com/FDA/Neural-Networks-based-Natural-Language-Processing