Published in final edited form as:

ACM Trans Comput Healthc. 2021 March ; 2(2): . doi:10.1145/3423209.

Attention-Gated Graph Convolutions for Extracting Drug Interaction Information from Drug Labels

TUNG TRAN, RAMAKANTH KAVULURU

University of Kentucky, United States

HALIL KILICOGLU

National Library of Medicine, United States

Abstract

Preventable adverse events as a result of medical errors present a growing concern in the healthcare system. As drug-drug interactions (DDIs) may lead to preventable adverse events, being able to extract DDIs from drug labels into a machine-processable form is an important step toward effective dissemination of drug safety information. Herein, we tackle the problem of jointly extracting mentions of drugs and their interactions, including interaction *outcome*, from drug labels. Our deep learning approach entails composing various intermediate representations, including graph-based context derived using graph convolutions (GCs) with a novel attention-based gating mechanism (holistically called GCA), which are combined in meaningful ways to predict on all subtasks jointly. Our model is trained and evaluated on the 2018 TAC DDI corpus. Our GCA model in conjunction with transfer learning performs at 39.20% F1 and 26.09% F1 on entity recognition (ER) and relation extraction (RE), respectively, on the first official test set and at 45.30% F1 and 27.87% F1 on ER and RE, respectively, on the second official test set. These updated results lead to improvements over our prior best by up to 6 absolute F1 points. After controlling for available training data, the proposed model exhibits state-of-the-art performance for this task.

Additional Key Words and Phrases:

Neural networks; multi-task learning; relation extraction; drug-drug interactions

1 INTRODUCTION

Preventable adverse events (AEs) are negative consequences of medical care resulting in injury or illness in a way that is generally considered avoidable. According to a report [16]

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

corresponding author: R. Kavuluru, University of Kentucky, Department of Internal Medicine, Division of Biomedical Informatics, 725 Rose St, Lexington, KY 40536; ramakanth.kavuluru@uky.edu;
Author's current address: H. Kilicoglu, School of Information Sciences, University of Illinois at Urbana-Champaign, 501 E Daniel St., Champaign, IL, 61820.

by the Department of Human and Health Services, based on an analysis of hospital visits by over a million Medicare beneficiaries, about *one in seven* hospital visits was associated with an AE, with 44% being considered clearly or likely preventable. Overall, AEs were responsible for an estimated US \$324 million in Medicare spending for the studied month of October 2008. Preventable AEs thus introduce a growing concern in the modern healthcare system as they represent a significant fraction of hospital admissions and play a significant role in increased healthcare costs. Alarming, preventable AEs have been cited as the eighth leading cause of death in the U.S., with an estimated fatality rate between 44,000 and 98,000 each year [14]. As drug-drug interactions (DDIs) may lead to a variety of preventable AEs, being able to extract DDIs from prescription drug labels¹ is an important effort toward effective dissemination of drug safety information. This includes extracting information such as adverse drug reactions and DDIs as indicated by drug labels. The U.S. Food and Drug Administration (FDA), for example, has recently begun to transform Structured Product Labeling (SPL) documents into a computer-readable format, encoded in national standard terminologies, that will be made available to the medical community and the public [6]. The initiative to develop a database of structured drug safety information that can be indexed, searched, and sorted is an important milestone toward a *fully automated* health information exchange system.

To aid in this effort, we propose a supervised deep learning model able to tackle the problem of DDI extraction in an *end-to-end* fashion. Most prior efforts assume all drug entities are known ahead of time (more in Section 2) such that the DDI extraction task reduces to a simpler *binary relation classification* task. We propose a system able to identify drug mentions in addition to their interactions. Concretely, the system takes as input the textual content of the label (indicating dosage and drug safety precautions) of a target drug and, as output, identifies mentions of other drugs that interact with the target drug. Thus, only one of the two interacting drugs is known beforehand (i.e., the “label drug”), while the other (i.e., the “precipitating drug,” or simply precipitant) is an unknown that our model is expected to extract. Along with identifying precipitants, we also determine the type of interaction associated with each precipitant; that is, whether the interaction is designated as being pharmacodynamic (PD) or pharmacokinetic (PK). In pharmacology, PD interactions are associated with a consequence on the organism, while PK interactions are associated with changes in how one or both of the interacting drugs is absorbed, transported, distributed, metabolized, and excreted when used jointly. Beyond identifying the interaction type, it is also important to identify the outcome or consequence of an interaction. As defined, PK consequence can be captured using a small fixed vocabulary, while identifying PD effects is a much more contrived process. The latter involves additionally identifying spans of text corresponding to a mention of a PD effect and linking each identified PD precipitant to one or more PD effects. We provide a more formal description of the task in Section 3.1. Figure 1 shows two simple examples illustrating the extracted outcomes for a PD and a PK interaction.

¹Drug labels are documents that communicate important usage information for a particular drug, including dosage, warnings and precautions, and potential adverse reactions including those stemming from drug-drug interactions.

To address this end-to-end variant of DDI extraction, we propose a multi-task joint learning architecture wherein various intermediate hidden representations are composed and are then combined to produce predictions for each subtask. These intermediate encodings include sequence-based contextual representations based on bidirectional Long Short-Term Memory (BiLSTM) networks and graph-based representations based on graph convolution (GC) networks. GCs over dependency parse trees are useful for capturing long-distance syntactic dependencies. We innovate on conventional GCs with a sigmoid gating mechanism derived via additive attention, referred to as Graph Convolution with Attention-Gating (GCA), which determines whether or not (and to what extent) information propagates between source and target nodes corresponding to edges in the dependency tree. The attention component controls information flow by producing a sigmoid gate (corresponding to a value in [0, 1]) for each edge based on an attention-like mechanism that measures relevance between node pairs. Intuitively, some dependency edges are more relevant than others; for example, *negations* or *adjectives* linked to important nouns via dependency edges may have a large influence on the overall meaning of a sentence, while *articles*, such as “the,” “a,” and “an,” have little or no influence comparatively. A standard GC would compose all source nodes with equal weighting, while the GCA would be more selective by possibly assigning a higher sigmoid value to *negations/adjectives* and a lower sigmoid value to *articles*.

We train and evaluate our model on the Text Analysis Conference (TAC) 2018 dataset for DDI extraction from drug labels [6]. The training data contains 22 drug labels, referred to as TR22, with gold standard annotations. As training data is scarce, we additionally propose a transfer learning step whereby the model is first trained on external data for extracting DDIs including the NLM-DDI CD corpus² and SemEval-2013 Task 9 dataset [9]; we refer to these as NLM180 and DDI2013, respectively. Two official test sets of 57 and 66 drug labels, referred to as Test Set 1 and 2, respectively, with gold standard annotations are used strictly for evaluation. Table 1 contains more information about these datasets and their characteristics. The contributions of this study are as follows:

- We show that our proposed graph convolution neural network with attention gating (GCA) improves over the standard graph convolution network with gains that are statistically significant.
- We show that our GCA-based model with transfer learning, via integration of deep contextualized representations and pretraining on external data, offers statistically significant gains of over four absolute F1 points compared with the prior best model [29]. This earlier model is from a prior effort³ that is solely based on BiLSTMs.
- We additionally show that the GCA-based model is highly complementary of simple BiLSTM models; that is, by combining the two via ensembling, we improve over the prior best by more than 7 absolute F1 points in overall performance with gains that are statistically significant.

²<https://lhce-brat.nlm.nih.gov/NLMDDICorpus.htm>.

³Tran et al. [29] was published as part of the non-refereed TAC; this study is an extension of our original report.

Among comparable methods, our GCA-based method exhibits state-of-the-art performance on all metrics, on both publicly available test sets, after controlling for available training data. Our code⁴ is made publicly available on GitHub.

2 RELATED WORKS

Prior studies on DDI extraction have focused primarily on binary relation extraction where drug entities are known during test time and the learning objective is reduced to a simpler *relation classification* (RC) task. In RC, pairs of known drug entities occurring in the same sentence are assigned a label, from a fixed set of labels, indicating relation type (including the *none* or *null* relation). Typically, no preliminary drug entity recognition or additional consequence prediction step is required. In this section, we cover prior relation extraction methods for DDI as well as participants of the initial TAC DDI challenge.

2.1 Deep Learning Methods for DDI Extraction

State-of-the-art methods for DDI extraction typically involve some variant of convolutional neural networks (CNNs) or recurrent neural networks (RNNs), or a hybrid of the two. Many studies utilize the dependency parse structure of an input sentence to capture long-distance dependencies, which has previously been shown to improve performance in general relation extraction tasks [31] and those in the biomedical domain [19, 21]. Liu et al. [20] first proposed the use of standard CNNs for DDI extraction. Their approach involved convolving over an input sentence with drug entities bound to generic tokens in conjunction with so-called *position vectors*. Position vectors are used to indicate the offset between a word and each drug of the pair and provide additional spatial features. Improvements were attained, in a follow-up study, by instead convolving over the shortest dependency path between the candidate drug pair [19]. Zhao et al. [32] introduced an enhanced version of the CNN-based method by deploying word embeddings that were pretrained on syntactic parses, part-of-speech embeddings, and traditional handcrafted features. Suárez-Paniagua et al. [26] instead focused on fine-tuning various hyperparameter settings including word and position vector dimensions and convolution filter sizes for improved performance. Kavuluru et al. [11] introduced the first neural architecture for DDI extraction based on hierarchical RNNs, wherein hidden intermediate representations are composed in a sequential fashion with cyclic connections, with character and word-level input. Sahu and Anand [24] experimented with various ways of composing the output of a bidirectional LSTM network including max-pooling and attention pooling. Lim et al. [18] proposed a *recursive* neural network architecture using recurrent units called TreeLSTMs to produce meaningful intermediate representations that are composed based on the structure of the dependency parse tree of a sentence. Asada et al. [2] demonstrated that combining representations of a CNN over the input text and graph convolutions over the molecular structure of the target drug pair (as informed by an external drug database) can result in improved DDI extraction performance. More recently, Sun et al. [27] proposed a hybrid RNN/CNN method by convolving over the contextual representations produced by a preceding RNN.

⁴https://github.com/bionlproc/gcn_ddi.

2.2 TAC 2018 DDI Track

TAC is a series of workshops organized by NIST aimed at encouraging research in natural language processing (NLP) by providing large test collections along with a standard evaluation procedure. The “DDI Extraction from Drug Labels” track [6] is established with the goal of transforming the contents of drug labels into a machine-processable format with linkage to standard terminologies. Tang et al. [28] placed first in the challenge using an encoder/decoder architecture to jointly identify precipitants and their interaction types and a rule-based system to determine interaction outcome. In addition to the provided training data, they downloaded and manually annotated a collection of 1,148 sentences to be used as external training data. Tran et al. [29] placed second in the challenge using a BiLSTM for joint entity recognition and interaction type prediction, followed by a CNN with two separate dense output layers (one for PK and one for PD) for outcome prediction. Dandala et al. [5] placed third in the challenge using a BiLSTM (with CRFs) with part-of-speech and dependency features as input for entity recognition. Next, an Attention-LSTM model was used to detect relations between recognized entities. The embeddings were pretrained on a corpus of FDA-released drug labels and used to initialize the model. NLM180 was used for training with TR22 serving as the development set. Other participants proposed systems involving similar approaches including BiLSTMs and CNNs as well as traditional linear and rule-based methods. We note that our method is the first to explore graph convolutions and transfer learning based on deep contextualized representation pretraining for this particular problem.

3 MATERIALS AND METHODS

We begin by formally describing the end-to-end task in Section 3.1. Next, we describe our approach to framing and modeling the problem (Section 3.2), the proposed network architecture (Section 3.4), the data used for transfer learning (Section 3.5), and our model-ensembling approach (Section 3.6). Finally, in Section 3.7, we describe the method for model evaluation.

3.1 Task Description

Herein, we describe the end-to-end task of automatically detecting drugs and their interactions, including the outcome of identified interactions, as conveyed in drug labels. We first define drug label as a collection of sections (e.g., DOSAGE & ADMINISTRATION, CONTRAINDICATIONS, and WARNINGS) where each section contains one or more sentences. Each sentence is annotated with a list of zero or more **mentions** and **interactions**. The overall task, in essence, involves fundamental language processing techniques including named entity recognition (NER) and relation extraction (RE). The first subtask of NER is focused on identifying *mentions* in the text corresponding to precipitants, interaction triggers, and interaction effects. Precipitating drugs (or simply precipitants) are defined as substances, drugs, or a drug class involved in an interaction. The second subtask of RE is focused on identifying sentence-level *interactions*; specifically, the goal is to identify the interacting precipitant, the type of the interaction, and the outcome of the interaction. The interaction outcome depends on the interaction type as follows. PD interactions are associated with a specified *effect* corresponding to a span within the text that describes the

outcome of the interaction. Figure 1 features a simple example of a PD interaction that is extracted from the drug label for **Adenocard**, where the precipitant is *digitalis* and the effect is “ventricular fibrillation.” Naturally, it is possible for a precipitant to be involved in multiple PD interactions. PK interactions, on the other hand, are associated with a label from a fixed vocabulary of National Cancer Institute (NCI) Thesaurus codes indicating various levels of increase/decrease in functional measurements. For example, consider the sentence “There is evidence that treatment with phenytoin leads to decreased intestinal absorption of furosemide, and consequently to lower peak serum furosemide concentrations.” Here, *phenytoin* is involved in a PK interaction with the label drug, *furosemide*, and the type of PK interaction is indicated by the NCI Thesaurus code C54615, which describes a decrease in the maximum serum concentration ($C_{\max}$) of the label drug. Lastly, *unspecified* (UN) interactions are interactions with an outcome that is not explicitly stated in the text and is typically indicated through cautionary remarks.

3.2 Joint Modeling Approach

Since only precipitants are annotated in the ground truth, we model the task of precipitant recognition and interaction type prediction jointly. We accomplish this by reducing the problem to a sequence tagging problem via a novel NER tagging scheme. That is, for each precipitant drug, we additionally encode the associated interaction type. Hence, there are three possible precipitant tags: **DYN**, **KIN**, and **UN** for precipitants with pharmacodynamic, pharmacokinetic, and unspecified interactions, respectively. Two more tags, **TRI** and **EFF**, are added to further identify mentions of *triggers* and *effects* concurrently. To properly identify boundaries, we employ the BILOU encoding scheme [23]. In the BILOU scheme, *B*, *I*, and *L* tags are used to indicate the beginning, inside, and last token of a multi-token entity, respectively. The *U* tag is used for unit-length entities, while the *O* tag indicates that the token is outside of an entity span. As a preprocessing step, we identify the label drug in the sentence, if it is mentioned, and bind it to a generic entity token (e.g., “LABELDRUG”). We also account for indirect mentions of the label drug, such as the generic version of a brand-name drug, or cases where the label drug is referred to by its drug class. To that end, we built a lexicon of drug names mapped to alias using NLM’s Medical Subject Heading (MeSH) tree as a reference. Table 2 shows how the tagging scheme is applied to a simple example.

Once we have identified the precipitant (as well as triggers/effects) and the interaction type for each precipitant, we subsequently predict the outcome or consequence of the interaction (if any). To that end, we consider all entity spans annotated with **KIN** tags and assign them a label from a static vocabulary of 20 NCI concept codes corresponding to PK consequence (i.e., multi-class classification). Likewise, we consider all entity spans annotated with **DYN** tags and link them to mention spans annotated with **EFF** tags; we accomplish this via binary classification of all pairwise combinations. For entity spans annotated with **UN** tags, no additional outcome prediction is needed.

3.3 Notations and Neural Building Blocks

In this section, we describe notations used in the remainder of this study. In addition, we provide a generic definition of the canonical CNN and BiLSTM networks that are later used

as building blocks in model construction. For ease of notation, we assume fixed sentence length n and word length $\hat{n}$; in practice, we set n and $\hat{n}$ to be the maximum sentence/word length and zero-pad shorter sentences/words. Moreover, we use square brackets with matrices to indicate a row indexing operation; for example, $X[j]$ denotes the vector corresponding to the j^{th} row of matrix X .

Henceforth, the abstract function $f_{\text{CNN}}^{w, d_{\text{out}}}(\cdot): \mathbb{R}^{n \times d_{\text{in}}} \mapsto \mathbb{R}^{d_{\text{out}}}$ is used to represent the CNN that convolves on a window of size w in a sentence with n words, mapping an $n \times d_{\text{in}}$ matrix to a vector representation of length d_{out} , where d_{in} is the word embedding size. This is an abstraction of the canonical CNN for NLP first proposed by Kim [13] and is defined as follows. First, we denote the convolution operation $\star$ as the sum of the element-wise products of two matrices. That is, for two matrices A and B of the same dimensions, $A \star B = \sum_j \sum_k A_{j,k} \cdot B_{j,k}$. Suppose the input is a sequence of vector representations $\mathbf{x}^1, \dots, \mathbf{x}^n \in \mathbb{R}^{d_{\text{in}}}$; the output representation $\mathbf{g} \in \mathbb{R}^{d_{\text{out}}}$ is defined such that

$$\mathbf{g}_k = \max(f_{\text{convolve}}(k, \mathbf{x}^1, \dots, \mathbf{x}^w), \dots, f_{\text{convolve}}(k, \mathbf{x}^{n-w+1}, \dots, \mathbf{x}^n)) \text{ for } k = 1, \dots, d_{\text{out}},$$

given a convolution function f_{convolve} that *convolves* over a contiguous window of size w n , defined as

$$f_{\text{convolve}}(k, \mathbf{v}^1, \dots, \mathbf{v}^w) = \text{ReLU}\left(W^k \star \begin{pmatrix} \mathbf{v}^1 \\ \vdots \\ \mathbf{v}^w \end{pmatrix} + b^k\right),$$

where $\mathbf{v}^1, \dots, \mathbf{v}^w \in \mathbb{R}^{d_{\text{in}}}$ are input vectors; $W^k \in \mathbb{R}^{w \times d_{\text{in}}}$ and $b^k \in \mathbb{R}$ for $k = 1, \dots, d_{\text{out}}$, are network parameters (corresponding to a set of d_{out} convolutional filters); and $\text{ReLU}(x) = \max(0, x)$ is the linear rectifier activation function. Here, d_{out} is a hyperparameter that determines the number of convolutional filters and thus the size of the final feature vector. In the study, we denote the convolution as an abstract function $f_{\text{CNN}}^{w, d_{\text{out}}}(\cdot): \mathbb{R}^{n \times d_{\text{in}}} \mapsto \mathbb{R}^{d_{\text{out}}}$ that convolves on a window of size w and maps an $n \times d_{\text{in}}$ matrix to a vector representation of length d_{out} .

Likewise, we represent the BiLSTM network as an abstract function

$f_{\text{BiLSTM}}^{d_{\text{out}}}(\cdot): \mathbb{R}^{n \times d_{\text{in}}} \mapsto \mathbb{R}^{n \times d_{\text{out}}}$ that maps a sequence of n input vectors (e.g., word embeddings) of d_{in} size (as an $n \times d_{\text{in}}$ matrix) to a corresponding sequence of n output context vectors of d_{out} size (as an $n \times d_{\text{out}}$ matrix). Let $\overrightarrow{\text{LSTM}}$ and $\overleftarrow{\text{LSTM}}$ represent an LSTM composition in the forward and backward direction. Suppose the input is a sequence of vector representations $\mathbf{x}^1, \dots, \mathbf{x}^n \in \mathbb{R}^{d_{\text{in}}}$; the output of a standard bidirectional LSTM network (BiLSTM) is a matrix $H \in \mathbb{R}^{n \times d_{\text{out}}} = (\mathbf{h}^1, \dots, \mathbf{h}^n)^{\top}$ such that

$$\overrightarrow{\mathbf{h}}^i = \overrightarrow{\text{LSTM}}(\mathbf{x}^i),$$

$$\overleftarrow{\mathbf{h}}^i = \overleftarrow{\text{LSTM}}(\mathbf{x}^i),$$

$$\mathbf{h}^i = \overrightarrow{\mathbf{h}}^i \parallel \overleftarrow{\mathbf{h}}^i, \quad \text{for } i = 1, \dots, n,$$

where $\parallel$ is the vector concatenation operator and $\mathbf{h}^i \in \mathbb{R}^{d_{\text{out}}}$ represents the context centered at the i^{th} word. Here, d_{out} is a hyperparameter that determines the size of the context embeddings. In the study, we denote the BiLSTM network as an abstract function $f_{\text{BiLSTM}}^{d_{\text{out}}}(\cdot): \mathbb{R}^{n \times d_{\text{in}}} \mapsto \mathbb{R}^{n \times d_{\text{out}}}$ that maps a sequence of n input vectors (e.g., word embeddings) of d_{in} size (as an $n \times d_{\text{in}}$ matrix) to a corresponding sequence of n output context vectors of d_{out} size (as an $n \times d_{\text{out}}$ matrix).

3.4 Neural Network Architecture and Training Details

We begin by describing how the three types of intermediate representations are composed. The construction of word-, context-, and graph-based representations are described in Sections 3.4.1, 3.4.2, and 3.4.3, respectively. Next, we describe the predictive components of the network that share and utilize the intermediate representations. In Section 3.4.4, we describe the sequence-labeling component of the network used to extract drugs and their interactions. In Section 3.4.5, we describe the component for predicting the interaction outcome. An overview of the architecture is shown in Figure 2. Lastly, we describe the model configuration and training process in Section 3.4.6.

3.4.1 Word-Level Representation.—Suppose the input is a sentence of length n represented by a sequence of word indices $w_1, \dots, w_n$ into the vocabulary $\mathcal{V}^{\text{Word}}$. Each word is mapped to a word embedding vector via embedding matrices $E^{\text{Word}} \in \mathbb{R}^{|\mathcal{V}^{\text{Word}}| \times \delta}$ such that δ is a hyperparameter that determines the size of word embeddings. In addition to word embeddings, we employ character-CNN-based representations as commonly observed in recent neural NER models [4]. Character-based models capture morphological features and help generalize to out-of-vocabulary words. For the proposed model, such representations are composed by convolving over character embeddings of size π using a window of size 3, producing η feature maps; the feature maps are then max-pooled to produce η -length feature representations. Correspondingly, we denote $E^{\text{Char}} \in \mathbb{R}^{|\mathcal{V}^{\text{Char}}| \times \pi}$ as the embedding matrix given the character vocabulary $\mathcal{V}^{\text{Char}}$; the character-level embedding matrix $C^i \in \mathbb{R}^{\hat{n} \times \pi}$ for the word at position i is

$$C^i = \begin{pmatrix} E^{\text{Char}}[c_i, 1] \\ \vdots \\ E^{\text{Char}}[c_i, \hat{n}] \end{pmatrix},$$

where c_{ij} for $1 \leq i \leq n$, $1 \leq j \leq \hat{n}$, represents the j^{th} character index of the i^{th} word. The word-level representation $R^{\text{word}} \in \mathbb{R}^{n \times (\delta + \eta)}$ is a concatenation of character-based word embeddings and pretrained word embeddings along the feature dimension; formally,

$$R^{\text{Word}} = \begin{pmatrix} E^{\text{Word}}[w_1] \parallel f_{\text{CNN}}^{3, \eta}(C^1) \\ \vdots \\ E^{\text{Word}}[w_n] \parallel f_{\text{CNN}}^{3, \eta}(C^n) \end{pmatrix}.$$

3.4.2 Context-Based Representation.—We compose context-based representation by simply processing the word-level representation with a BiLSTM layer as is common practice; concretely, $R^{\text{Context}} = f_{\text{BiLSTM}}^{\rho}(R^{\text{Word}})$, where ρ is a hyperparameter that determines the size of the context embeddings.

3.4.3 Graph-Based Representation.—In addition to the sequential nature of LSTMs, we propose an alternative and complementary graph-based approach for representing context using GC networks. Typically composed on dependency parse trees, graph-based representations are useful for relation extraction as they capture long-distance relationships among words of a sentence as informed by the sentence’s syntactic dependency structure. While GCs are typically applied repeatedly, our initial cross-validation results indicate that single-layered GCs are sufficient and deep GCs typically resulted in performance degradation; moreover, Zhang et al. [31] report good performance with similarly shallow GC layers. Hence, the following formulation describes a *single-layered* GC network, with an additional attention-based sigmoid gating mechanism, which we holistically refer to as a GCA network. Initially motivated in Section 1, the GCA improves on conventional GCs with a sigmoid-gating mechanism derived via an alignment score function associated with *additive* attention [3]. The sigmoid “gate” determines whether or not (and to what extent) information is propagated based on a learned alignment function that conceives a “relevance” score between a *source* and *target* node (more later).

As a preprocessing step, we use a dependency parsing tool to generate the projective dependency tree for the input sentence. We represent the dependency tree as an $n \times n$ adjacency matrix A , where $A_{ij} = A_{ji} = 1$ if there is a dependency relation between words at positions i and j . This matrix controls the flow of information between pairs of words corresponding to connected nodes in the dependency tree (ignoring dependency *type*); however, it is also important for the *existing* information of each node to carry over on each application of the GC. Hence, as with prior work [31], we use the modified version $\tilde{A} = A + I$, where I is the identity matrix to allow for self-loops in the GC network. The graph-based representation $R^{\text{Graph}} \in \mathbb{R}^{n \times \beta}$ is composed such that

$$R^{\text{Graph}}[i] = \tanh \left(\sum_{j=1}^n \tilde{A}_{i,j} W^{\text{Graph}} R^{\text{Context}}[j] + \mathbf{b}^{\text{Graph}} \right),$$

where $W^{\text{Graph}} \in \mathbb{R}^{\beta \times \rho}$, $\mathbf{b}^{\text{Graph}} \in \mathbb{R}^{\beta}$ are network parameters; $\tanh(\cdot)$ is the hyperbolic tangent activation function; and β is a hyperparameter that determines the hidden GC layer size. Thus, information propagated from source nodes $j = 1, \dots, n$ to target node i , based on the summation of intermediate representations, is unweighted and shares equal importance.

As stated previously, we propose to extend the standard GC by adding an attention-based sigmoid gating mechanism to control the flow of information via the gating matrix $G \in \mathbb{R}^{n \times n}$. We define G such that

$$G_{i,j} = \sigma(\mathbf{v} \cdot \mathbf{a}^{i,j}) \quad \text{for } i = 1, \dots, n, j = 1, \dots, n,$$

where $\mathbf{v} \in \mathbb{R}^{\alpha}$ is a network parameter and $\mathbf{a}^{i,j} \in \mathbb{R}^{\alpha}$ is the hidden attention layer composed as a function of the context representation at source node i and target node j ; concretely,

$$\mathbf{a}^{i,j} = \tanh(W^{\text{Source}} R^{\text{Context}}_{[i]} + W^{\text{Target}} R^{\text{Context}}_{[j]} + \mathbf{b}^{\text{Attn}}),$$

where $W^{\text{Source}}, W^{\text{Target}} \in \mathbb{R}^{\alpha \times \rho}$, and $\mathbf{b}^{\text{Attn}} \in \mathbb{R}^{\alpha}$ are network parameters and α is a hyperparameter that determines hidden attention layer size. Intuitively, the network learns the relevance of node i to node j via the attention $\mathbf{a}^{i,j}$ and outputs a between 0 and 1 at gate $G_{i,j}$. Gate $G_{i,j}$ controls the flow of information from node i to j , where 0 indicates no information is passed and 1 indicates that all information is passed. To integrate the gating mechanism, we simply redefine $\tilde{A} = (A + I) \times G$. In the next two sections, we show how the intermediate representations are used for end-task prediction.

3.4.4 Sequence Labeling.—The sequence labeling (SL) task for detecting precipitant drugs and their interaction type is handled by a bidirectional LSTM trained on a combination of the two types of losses: conditional random fields (CRFs) and softmax cross entropy (SCE). Using CRFs results in choosing a globally optimal assignment of tags to the sequence, whereas a standard softmax at the output of each step may result in less globally consistent assignments (e.g., an L tag following an O tag) but better local or partial assignments. We begin by introducing a bidirectional LSTM layer that processes the various intermediate representations. The new representation, $R^{\text{SL}} \in \mathbb{R}^{n \times \gamma}$, is defined such that

$$R^{\text{SL}} = f_{\text{BLSTM}}^Y \begin{pmatrix} R^{\text{Word}}_{[1]} \| R^{\text{Context}}_{[1]} \| R^{\text{Graph}}_{[1]} \\ \vdots \\ R^{\text{Word}}_{[n]} \| R^{\text{Context}}_{[n]} \| R^{\text{Graph}}_{[n]} \end{pmatrix},$$

where γ is a hyperparameter that determines the hidden layer size. While R^{Graph} is based on R^{Context} and R^{Context} is based on R^{Word} , we observed that combining these intermediate representations (manifesting at varying depths in the architecture) resulted in improved sequence-labeling performance according to preliminary experiments and prior results from Tran et al. [29]. As with residual networks [8], they additionally provide a kind of shortcut or “skip connection” over intermediate layers.

Given a set of n_{tag} possible tags, we compose an $n \times n_{\text{tag}}$ score matrix Y (where $Y_{i,t}$ represents the score of the t^{th} tag at position i) such that $Y[i] = W^{\text{Out}} R^{\text{SL}}[i] + \mathbf{b}^{\text{Out}}$, where $W^{\text{Out}} \in \mathbb{R}^{n_{\text{tag}} \times \gamma}$, $\mathbf{b}^{\text{Out}} \in \mathbb{R}^{n_{\text{tag}}}$ are network parameters. Given example x and the truth tag assignment as a matrix $\bar{Y}$ where rows are one-hot vectors over all possible tags, the SCE loss is

$$\ell_{\text{SCE}}(x, \tilde{A}, \bar{Y}; \theta) = - \sum_{i=1}^n \sum_{t=1}^{n_{\text{tags}}} \bar{Y}_{i,t} \log \left(\frac{\exp(Y_{i,t})}{\sum_{k=1}^{n_{\text{tags}}} \exp(Y_{i,k})} \right),$$

where $\bar{Y}_{i,t} \in \{0, 1\}$ indicates whether the tag t is assigned at position i and θ is the set of all network parameters. Next, we define the CRF loss as commonly used with LSTM-based models for entity recognition. We learn a transition score matrix $M \in \mathbb{R}^{n_{\text{tag}} \times n_{\text{tag}}}$, inferred from the training data, such that M_{ij} is the transition score from tag i to tag j . Given an example x as a sequence of word indices $w_1, \dots, w_n$ and candidate tag sequence $\bar{y}$ as a sequence of tag indices $s_1, \dots, s_n$, the tag assignment score (t-score) is defined as

$$\text{t-score}(x, \tilde{A}, \bar{y}; \hat{\theta}) = \text{t-score}(w_1, \dots, w_n, \tilde{A}, s_1, \dots, s_n; \hat{\theta}) = \sum_{i=1}^n (Y_{i, s_i} + M_{s_{i-1}, s_i}),$$

where $\hat{\theta} = \theta \cup \{M\}$. Intuitively, this score summarizes the likelihood of observing a transition from tag s_{i-1} to tag s_i in addition to the likelihood of emitting tag s_i given the semantic context for $i = 1, \dots, n$. Thus, Y is treated as a matrix of emission scores for the CRF. For an example with input x and truth tag assignment $\bar{y}$, the loss is computed as the negative log-likelihood of the tag assignment as informed by the normalized tag assignment score, or

$$\ell_{\text{CRF}}(x, \tilde{A}, \bar{y}; \hat{\theta}) = - \log \frac{\exp(\text{t-score}(x, \tilde{A}, \bar{y}; \hat{\theta}))}{\sum_{y \in \mathcal{S}} \exp(\text{t-score}(x, \tilde{A}, y; \hat{\theta}))},$$

where $\mathcal{S}$ is the set of all possible tag assignments. The final per-example loss for sequence labeling is simply a summation of the two losses: $\mathcal{L}_{\text{L}} = \mathcal{L}_{\text{CE}} + \mathcal{L}_{\text{CRF}}$. During testing, we use the Viterbi algorithm [30], a dynamic programming approach, to decode and identify the globally optimal tag assignment.

3.4.5 Consequence Prediction.—Once precipitants (and corresponding interaction types) have been identified, we perform so-called consequence prediction (CP) for all precipitant drugs identified as participating in PD or PK interactions. The classification task of CP takes as input the target sentence and two candidate entities that are referred to as the *subject* and *object* entities. Here the *subject* is always a precipitating drug; on the other hand, the *object* designation depends on the type of interaction (more later). First, we define the representation matrix for CP as $R^{\text{CP}} \in \mathbb{R}^{n \times (\rho + \beta)}$, where

$$R^{CP} = \begin{pmatrix} R^{Context_{[1]}} \parallel R^{Graph_{[1]}} \\ \vdots \\ R^{Context_{[n]}} \parallel R^{Graph_{[n]}} \end{pmatrix}.$$

We process the matrix via convolutions of windows sizes 3, 4, and 5 and concatenate the results to produce the final feature vector g^{CP} . In addition to CNN features, we map entities to their graph-based context features and append it to g^{CP} , which has been previously shown to work well in a similar architecture [17]. Concretely, the final feature vector is

$$g^{CP} = f_{CNN}^{3,\mu}(R^{CP}) \parallel f_{CNN}^{4,\mu}(R^{CP}) \parallel f_{CNN}^{5,\mu}(R^{CP}) \parallel R^{CP}_{[t_{Sub}]} \parallel R^{CP}_{[t_{Obj}]}$$

with $g^{CP} \in \mathbb{R}^{3\mu + 2(\rho + \beta)}$, where μ , as a hyperparameter, is the number of CNN filters per convolution and t_{Sub} and t_{Obj} are the position index of the last word (typically the “head” word) of the subject and object, respectively.

The actual entities determined to be the subject/object pair are based on the interaction type; for PD interactions, the *subject* is the precipitant drug and the *object* is some candidate *effect* mention. For PK interactions, however, the subject is the precipitant drug but the object is chosen to be the closest (based on character-offset) mention of the drug label with respect to the target precipitant drug. We found this appropriate based on manual review of the data, as the NCI code being assigned depends highly on whether the increase/decrease in functional measurements is with respect to the label drug or the precipitant drug. In case the label drug is not mentioned, a generic “null” vector is used to represent the object.

When performing sequence labeling, we pass in the entire dependency tree encoded as the matrix $\tilde{A}$. However, when performing consequence prediction and both entities are non-null, we pass in a *pruned* version of the entire tree that is tailored to the entity pair. We apply the same pruning strategy proposed by Zhang et al. [31], wherein for a pair of subject and object entities (corresponding to t_{Sub} and t_{Obj}), we keep only nodes either along or *within one hop* of the shortest dependency path. This prevents distant and irrelevant portions of the dependency tree from influencing the model while retaining important modifying and negating terms. Thus, the notation $\tilde{A}_{Sub \leftrightarrow Obj}$ is used to denote the pruned version of $\tilde{A}$ as a function of the entity pair indicated by t_{Sub} and t_{Obj} . For PK interactions where the label drug is not mentioned, we simply pass in the un-pruned dependency tree.

To determine whether there is a PD interaction between a pair of entities, we employ a standard binary classification output layer. Concretely, for example sentence $\hat{x}$ and output $y \in \{0, 1\}$, the probability of a PD interaction between the entity pair is $q = \text{sigmoid}(\mathbf{w}^{PD} \cdot \mathbf{g}^{CP} + b^{PD})$, where $\mathbf{w}^{PD} \in \mathbb{R}^{3\mu + 2(\rho + \beta)}$ and $b^{PD} \in \mathbb{R}$ are network parameters. The associated binary cross-entropy loss is

$$\mathcal{L}_{PD}(x, \tilde{A}_{Sub \leftrightarrow Obj}, \hat{y}; \theta) = -(\hat{y} \log q + (1 - \hat{y}) \log(1 - q)),$$

where $\hat{y} \in \{0, 1\}$ indicates the ground truth. For PK interactions, we instead use a softmax function to produce a probability distribution, represented as vector $\mathbf{q} \in \mathbb{R}^{20}$, over the 20 labels corresponding to NCI Thesaurus codes. Concretely, the predicted probability of label j is $q_j = \exp(y_j^{\text{PK}}) / \exp(\sum_{k=1}^{20} y_k^{\text{PK}})$, where $\mathbf{y}^{\text{PK}} = \mathbf{W}^{\text{PK}} \mathbf{g}^{\text{CP}} + \mathbf{b}^{\text{PK}}$ and $\mathbf{W}^{\text{PK}} \in \mathbb{R}^{20 \times [3\mu + 2(\rho + \beta)]}$ and $\mathbf{b}^{\text{PK}} \in \mathbb{R}^{20}$ are network parameters. Given a one-hot vector $\bar{\mathbf{y}} \in \mathbb{R}^{20}$ indicating the ground truth, the associated softmax cross-entropy loss is

$$\ell_{\text{PK}}(x, \tilde{A}_{\text{Sub} \leftrightarrow \text{Obj}}, \bar{\mathbf{y}}; \theta) = - \sum_{j=1}^{20} \bar{y}_j \log q_j.$$

The loss for a batch of examples is simply the sum of its constituent example-based losses.

3.4.6 Neural Network Configuration and Training Details.—For each training iteration, we randomly sample 10 sentences from the training data. These are re-composed into three sets of task-specific examples $\mathcal{S}$, $\mathcal{D}$, and $\mathcal{X}$ corresponding to the tasks of sequence labeling, PD prediction, and PK prediction, respectively. Unlike our prior work, in which the subtasks were trained in an interleaved fashion, we train on all three objectives jointly. Here, we dynamically switch between one of four training objective losses based on whether there are available training examples (in the batch and for the current iteration) for each task. The final training loss is then

$$\ell = \begin{cases} \sum_{x \in \mathcal{S}} \ell_{\text{SL}}(x) + \sum_{x \in \mathcal{D}} \ell_{\text{PD}}(x) + \sum_{x \in \mathcal{X}} \ell_{\text{PK}}(x) & \text{if } |\mathcal{D}| > 0 \text{ and } |\mathcal{X}| > 0, \\ \sum_{x \in \mathcal{S}} \ell_{\text{SL}}(x) + \sum_{x \in \mathcal{X}} \ell_{\text{PK}}(x) & \text{if } |\mathcal{D}| = 0 \text{ and } |\mathcal{X}| > 0, \\ \sum_{x \in \mathcal{S}} \ell_{\text{SL}}(x) + \sum_{x \in \mathcal{D}} \ell_{\text{PD}}(x) & \text{if } |\mathcal{D}| > 0 \text{ and } |\mathcal{X}| = 0, \\ \sum_{x \in \mathcal{S}} \ell_{\text{SL}}(x) & \text{otherwise.} \end{cases}$$

We train the network for a maximum of 10,000 iterations, checkpointing and evaluating every 100 iterations on a validation set of sentences from four held-out drug labels. Only the checkpoint that performed best on the validation set is kept for test time evaluation. The choice of hyperparameters is shown in Table 3; discrete numbered parameters corresponding to embedding or hidden size were chosen from $\{10, 25, 50, 100, 200, 400\}$ based on random search and optimized by assessing 11-fold cross-validation performance on TR22. The learning and dropout rates are set to typical default values. We used Word2Vec embeddings pretrained on the corpus of PubMed abstracts [22]. All other variables are initialized using values drawn from a normal distribution with a mean of 0 and standard deviation of 0.1 and further tuned during training. Words were tokenized on both spaces and punctuation marks; punctuation tokens were kept as is common practice for NER-type systems. For dependency parsing, we use SyntaxNet, which implements the transition-based neural model by Andor et al. [1]. We trained the aforementioned parser, using default settings, on the GENIA corpus [12] and use it to obtain projective dependency parses for each example.

3.5 Transfer Learning with Network Pretraining

An obstacle in solving this flavor of DDI extraction as a machine learning problem is the high potential for overfitting given the sparse nature of the output space, which is further intensified by the scarce availability of high-quality training data. As quality training data is expensive and requires domain expertise, we propose to use a transfer learning approach, where the model is pretrained on external data, as follows. First, we pretrain on the DDI2013 dataset, which contains strictly binary relation DDI annotations and no interaction consequence annotation. Hence, DDI2013 is only used to train the sequence labeling objective $\mathcal{L}_L(x)$. Next, we pretrain on NLM180, a collection of 180 drug labels annotated in a comparable format to TR22 but that follows a different set of guidelines and lacks comprehensive interaction consequence annotation. Finally, we fine-tune for the target task by training on the official TR22 dataset.

Translating NLM180 and DD2013 to the TAC 2018 format is an imperfect process given structural (breadth and depth of annotations) and semantic (guidelines in addition to annotator experience and vision) differences. For example, differences in how entity boundaries are annotated, such as whether or not modifier terms should be kept as part of a named entity, may have a large impact on model performance. Hence, we expect the translated versions of NLM180 and DDI2013 to be very noisy as training examples for the target task. We describe the translation process for DDI2013 in Sections 3.5.1 and 3.5.2. We provide summary statistics about these datasets in Table 1.

In light of recent breakthroughs across many biomedical tasks as a result of recent advances in deep contextualized representations using transformers (i.e., BERT [7]) for biomedicine (i.e., BioBERT [15]), we include additional experiments with our models where the contextual features from pretrained BioBERT models are used as drop-in replacements for context representations. As an implementation detail, these representations are generated by the BioBERT base-with-casing model pretrained on PubMed articles that are publicly available.⁵ As the BERT architecture uses a type of subword tokenization, we map subwords back to their original tokens so that units of representation align with our word-level tokenization scheme.

3.5.1 NLM180 Mapping Scheme.—In NLM180, there is no distinction between triggers and effects; moreover, PK effects are limited to *coarse*-grained (binary) labels corresponding to an *increase* or *decrease* in function measurements. Hence, a direct mapping from NLM180 to the TR22 annotation scheme is impossible. As a compromise, NLM180 “triggers” were mapped to TR22 *triggers* in the case of unspecified and PK interactions. For PD interactions, we instead mapped NLM180 “triggers” to TR22 *effects*, which we believe to be appropriate based on our manual analysis of the data. Since we do not have both *trigger* and *effect* for every PD interaction, we opted to ignore trigger mentions altogether in the case of PD interactions to avoid introducing mixed signals. While trigger recognition has no bearing on relation extraction performance, this policy has the effect of reducing the recall upper bound on NER by about 25% based on early cross-validation results. To

⁵<https://github.com/dmis-lab/biobert>.

overcome the lack of fine-grained annotations for PK outcome in NLM180, we deploy the well-known bootstrapping approach [10] to incrementally annotate NLM180 PK outcomes using TR22 annotations as a starting point. To mitigate the problem of semantic drift, we reannotated by hand iterative predictions that were not consistent with the original NLM180 coarse annotations (i.e., active learning [25]).

3.5.2 DDI2013 Mapping Scheme.—The DDI2013 dataset contains annotations that are incomplete with respect to the target task; specifically, annotations are limited to typed binary relations between any two mentioned drugs in the sentence (and not necessarily between a mentioned drug and the label drug) without outcome or consequence prediction. In DDI2013, there are four types of interactions: *mechanism*, *effect*, *advice*, and *int*. The *mechanism* type indicates that a PK mechanism is being discussed; *effect* indicates that the consequence of a PD interaction is being discussed; *advice* indicates suggestions regarding the handling of the drugs; and *int* is an interaction without any specific additional information. We translate the annotation by first applying a filtering step on all interactions such that it conforms to the target task; namely, we filter in a way that only interactions involving the label drug are kept. The nonlabel drug entity is then annotated as a precipitant with an interaction tag based on the following mapping scheme. Entities involved in a *mechanism* relation with the drug label are treated as **KIN** precipitants; likewise, entities in *effect* and *advice* relations are treated as **DYN** precipitants and *int* relations are treated as **UNK** precipitants. As there is no consequence annotation, the mapped examples are used to train the sequence labeling objective but *not* the other objective.

3.6 Voting-Based Ensembling

Our prior effort [29] showed that model ensembling resulted in optimal performance for this task. Hence, model ensembling remains a key component of the proposed model. Our ensembling method is based on ensembling over 10 models, each trained with randomly initialized weights and a random development split. Intuitively, models collectively “vote” on predicted annotations that are kept and annotations that are discarded. A unique annotation (entity or relation) has one vote for each time it appears in one of the 10 model prediction sets. In terms of implementation, unique annotations are incrementally added (to the final prediction set) in order of descending vote count; subsequent annotations that conflict (i.e., overlap based on character offsets) with existing annotations are discarded. Hence, we loosely refer to this approach as “voting-based” ensembling.

3.7 Model Evaluation

We used the official evaluation metrics for NER and relation extraction based on the standard precision, recall, and F1 micro-averaged over exactly matched entity/relation annotations. We use the strictest matching criteria corresponding to the official “primary” metric (of the TAC DDI task), as opposed to the “relaxed” metric that ignores mention and interaction *type*. Concretely, the matching criteria for entity recognition consider entity bounds as well as the type of the entity. The matching criteria for relation extraction comprehensively considers precipitant drugs and, for each, the corresponding interaction type and interaction outcome. As relation extraction evaluation takes into account the bounds of constituent entity predictions, relation extraction performance is

heavily reliant on entity recognition performance. On the other hand, we note that while NER evaluation considers *trigger* mentions, *triggers* are ignored when evaluating relation extraction performance. Two test sets of 57 and 66 drug labels, referred to as Test Set 1 and 2, respectively, with gold standard annotations are used for evaluation.

Next, we discuss the differences between these test sets. As shown in Table 1, Test Set 1 closely resembles TR22 with respect to the sections that are annotated. However, Test Set 1 is more sparse in the sense that there are more sentences per drug label (144 vs. 27), with a smaller proportion of those sentences having gold annotations (23% vs. 51%). Test Set 2 is unique in that it contains annotations from only two sections, namely DRUG INTERACTIONS and CLINICAL PHARMACOLOGY, the latter of which is not represented in TR22 (nor Test Set 1). Lastly, TR22, Test Set 1, and Test Set 2 all vary with respect to the distribution of interaction types, with TR22, Test Set 1, and Test Set 2 containing a higher proportion of PD, UN, and PK interactions, respectively. Model performance is assessed using a single overall metric defined as the average of *entity recognition* and *relation extraction* performance across both test sets.

4 RESULTS AND DISCUSSION

In order to assess model performance with confidence intervals and draw conclusions based on statistical significance, we perform a technique called *bootstrap ensembling* proposed by Kavuluru et al. [11]. That is, for each neural network (NN), we train a pool of 30 models each with a different set of randomly initialized weights and training-development set split. Performance of the NN is evaluated based on computing the 95% confidence interval around the mean F1 of $N = 100$ ensembles, where each ensemble is assembled from a set of 10 models randomly sampled from the pool. This approach allows us to better assess average performance which is a non-trivial task given the high variance nature of models learned with limited training data. Our method for model ensembling (by “voting”) is described in Section 3.6.

We present the main results of this study in Table 4, where we compare our prior efforts using strictly BiLSTMs (BL) and our current best results with graph convolutions and pretrained BERT embeddings (GCA/B). BL with TR22 and NLM180 as training data corresponds to our prior best at 28.16% overall F1, while GCA/B with TR22, NLM180, and DDI2013 as training data represents our current best at 33.55% overall F1. Here, we observe a ≈ 5 -point gain in overall F1 (statistically significant at 95% confidence level based on nonoverlapping confidence intervals), with most gains owing to a substantial improvement in entity recognition performance. We note that GCA is more precision focused, while BL is more recall focused; moreover, GCA/B tends to exhibit better performance on Test Set 1, while BL tends to exhibit better performance on Test Set 2. This hints that the two architectures are highly complementary and may work well in combination. Indeed, when combined via ensembling, we observe a major performance gain across almost all measures. Here, for each ensemble, we sample five models from each pool of models (GCA/B and BL) for a total of 10 models to ensure that results remain comparable. The resulting hybrid model exhibits the best performance overall, improving over the prior best by 2 points and over the current best by 6 points in overall F1 at 35.45%. These differences are statistically

significant at the 95% confidence level. Next, we highlight that a main benefit of the GCA model is that it operates well with very small amounts of training data, as evident by the almost 2-absolute-point improvement over the BiLSTM model when trained solely on TR22. These gains tend to be less notable when we involve examples from NLM180 and DDI2013. Lastly, we note that GCA (graph convolution with attention gating) performs better than the standard GC (graph convolution *without* attention gating) by 2 absolute points in overall F1 with improvements that are consistent across all metrics. We present a comparison of our results with other works in Table 5. We omit results by Tang et al. [28] as they are not directly comparable to ours given the stark difference in available training data. When training on strictly TR22 and NLM180 (thus being comparable to most prior work), our model exhibits state-of-the-art performance across all metrics on either test set.

We present Figures 3 and 4 to illustrate error cases to be discussed later in Section 5. In addition to actual and predicted annotations, these figures include a sigmoid gating activity visualization for edges in the dependency tree. The visualization serves two purposes. First, it confirms the intuition for this particular design, and second, it provides a means to interpret model decisions. That is, we can observe the importance of each edge in the dependency tree as deemed by the network for a particular example. In Figure 3, for example, we can observe that for the target word “digoxin” (which is a precipitant, the second occurrence in the sentence), the words “use,” “concomitantly,” and “with” show very high activity. Likewise, signal flow from “hemodynamic” to “effects” is strong, and vice versa. Less important words such as articles appear to receive less incoming activity overall, even through self-loops.

It is quite clear that the absolute performance scores reported here are nowhere near the high scores needed to truly automate the process of extracting interactions from drug labels. While results are state of the art when compared to several other efforts in the field, it is still necessary to address the relatively low evaluation measures and highlight the practical implications of these results. To begin, we note that the evaluation metric used has very strict matching criteria (as deemed appropriate by the shared task organizers), leading to greater penalties for inexact predictions. While excellent for fine-grained comparisons between models, these measurements may not be meaningful in isolation. In real-world applications, an evaluation measure that is more lenient on partial matching may provide a more useful evaluation of the model if the end-task does not require, for example, perfect entity-boundary detection. A model that is able to identify the concept “contraceptive” instead of a more comprehensive version, for example, “oral contraceptive,” may still be useful. It would also allow flexibility in cases where there is semantic ambiguity (even among trained annotators) and accommodate occasional cases of annotation errors occurring in the test set (more in Section 5). Besides the strict matching criteria, the sheer complexity of the task and lack of large amounts of high-quality training data may also be contributing to low scores. Here, we additionally note that this extraction problem and shared task arose as a real-world need from FDA/NLM researchers to begin the task of transforming SPL product labels into a structured format. While the outputs of these models are not perfect, the mentions/interactions extracted are useful as they can expedite the process by providing human annotators with a reasonable starting point. In the end, as this exhaustive form of

DDI extraction is still new, our methods and results serve as a strong baseline effort in this domain.

5 ERROR ANALYSIS

In this section, we perform error analysis to identify challenging cases typically resulting in erroneous predictions by the model. One major source of difficulty for the model is boundary detection in cases of multiword entities. Errors of this type are especially prominent in case of *effect* mentions, which may manifest as potentially long noun phrases. Phrases with conjunctions or punctuation marks (or a combination of) may also present an obstacle for the model; for example, an *effect* expressed as “serious and/or life-threatening reaction” may instead be predicted as simply “life-threatening reaction.” Figure 3 shows a general case of this error where the model recognizes “potentiation of adverse hemodynamic effects” as the effect while the ground truth identifies the effect as simply “adverse hemodynamic effects.” This leads to both a false positive and a false negative for both the NER and the RE evaluation. We note that, given the potentially limitless ways an *effect* may be expressed, any disagreement among annotators (for cases beyond those addressed in annotator guidelines) during the initial annotation process will lead to inconsistent ground-truth data and thus negatively affect downstream model performance. As an example, consider the following two sentences that appear in TR22: “Co-administration of SAMSCA with potent CYP3A inducers ..” and “For patients chronically taking potent inducers of CYP3A, ..” Here, one sentence is annotated such that *potent* is included as part of the precipitant expression, while another is annotated such that this modifier is excluded.

Mixed signals and noisy labels in general tend to be an issue especially when there is limited training data as deep learning models are prone to overfitting. When evaluating on purely effect mentions, we obtain a micro-F1 score of 66% (54% Precision, 87% Recall). However, the micro-F1 is 87% when ignoring the starting boundary offset and 86% when ignoring the ending boundary offset during evaluation, corresponding to roughly a 20 absolute micro-F1 gain in performance. When applying the same looser evaluation criteria to triggers and precipitants, the gains are only $\approx 6\%$ and $\approx 5\%$, respectively. Thus, there is immense potential for improving entity recognition of *effect* mentions if we can better handle boundary detection, possibly via rule-based methods or postprocessing adjustments, with the added benefit of improving consequence prediction performance for PD interactions.

Precipitants interacting with the label drug being mentioned multiple times may also cause issues for the model. As an example, consider the sentence presented in Figure 3. Our model identifies both mentions of the precipitant “Digoxin” as being involved in an interaction with the drug Savella; however, the ground truth more specifically recognizes the second mention as the sole precipitant. This results in an additional false positive with respect to both NER and RE evaluation. Lastly, there are cases where the model will mistake a mention subtly referring to the label drug as a precipitant. This is a common occurrence in cases where the label drug is not referred to by name, but by a class of drugs. Typically, identifying a mention as a reference to the drug label beforehand will disqualify it from being predicted as a precipitant. While we do use a lexicon of drug names mapped to drug synonyms and drug classes to identify these indirect mentions, it is not exhaustive for all

drugs. For example, within the label of the drug Lexapro, consider the sentence “Altered anticoagulant effects, including increased bleeding, have been reported when SSRIs and SNRIs are coadministered with warfarin.” Here, the model recognized SSRI and SNRI as precipitants. This is incorrect, however, as Lexapro is an SSRI and these mentions are more than likely referring to Lexapro. Without this information, the model likely assumes that it is an implicit case where the label drug is not mentioned and therefore assume all drug mentions are precipitants. Hence, curating a more exhaustive lexicon for indirect mentions of the label drug will improve overall performance.

Lastly, we describe a source of difficulty stemming from incorrectly classifying interaction types. Figure 4 presents an example sentence where our model mistakes PK for PD interactions and a *trigger* mention for an *effect* mention. As PD and PK interactions tend to frequently co-occur with *effect* and *trigger* mentions, respectively, predicted annotations tend to be polarized toward one pair (PD with *effect*) or the other (PK with *trigger*). Hence, differentiating between types of interactions for each recognized precipitant is another interesting class of error. Among all correctly recognized precipitants (based purely on boundary detection), we analyzed cases where one type of interaction, among PD, PK, and UN, is mistaken for another via the confusion matrix in Table 6. Clearly, many errors are due to cases where (1) we mistake *unspecified* precipitants for PD precipitants and (2) we mistake PK precipitants for *unspecified* precipitants. We conjecture that making precise implicit connections (not only whether there is evidence in the form of trigger words or phrases but also whether the evidence concerns the particular precipitant) is highly nontrivial. Likely, this aspect may be improved by inclusion of more high-quality training data. Confusion between *trigger* and *effect* mentions is less concerning; among more than 1,000 cases, there are six cases where we mistake *effect* for *trigger* and 20 cases where we mistake *trigger* for *effect*.

6 CONCLUSION

In this study, we proposed an end-to-end method for extracting mentions of drugs and their interactions from drug labels, including interaction outcome in the case of PK and PD interactions. The method involved composing various intermediate representations including sequential and graph-based context, where the latter is produced using a novel attention-gated version of the graph convolution over dependency parse trees. The so-called graph convolution with attention-gating, along with transfer learning via serial pretraining using other annotated DDI datasets including DDI2013, resulted in an improvement over our original TAC challenge entry by up to 6 absolute F1 points overall. Among comparable studies (based on training data composition), our method exhibits state-of-the-art performance across all metrics and test sets. Future work will focus on curating more quality training data and leveraging semisupervised methods to overcome the scarcity in training data.

ACKNOWLEDGMENTS

We thank anonymous reviewers for their constructive criticism and suggestions toward improvement of the manuscript.

This research was partially conducted during TT's participation in the Lister Hill National Center for Biomedical Communications (LHNCBC) Research Program in Medical Informatics for Graduate students at the U.S. National Library of Medicine, National Institutes of Health. HK is supported by the intramural research program at the U.S. National Library of Medicine. RK and TT were supported by the U.S. National Library of Medicine through grant R21LM012274.

REFERENCES

- [1]. Andor Daniel, Alberti Chris, Weiss David, Severyn Aliaksei, Presta Alessandro, Ganchev Kuzman, Petrov Slav, and Collins Michael. 2016. Globally normalized transition-based neural networks. arXiv preprint arXiv:1603.06042 (2016).
- [2]. Asada Masaki, Miwa Makoto, and Sasaki Yutaka. 2018. Enhancing drug-drug interaction extraction from texts by molecular structure information. In Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers). 680–685.
- [3]. Bahdanau Dzmitry, Cho Kyunghyun, and Bengio Yoshua. 2015. Neural machine translation by jointly learning to align and translate. In Proceedings of 3th International Conference on Learning Representations (ICLR'15).
- [4]. Chiu Jason P. C. and Nichols Eric. 2016. Named entity recognition with bidirectional LSTM-CNNs. Transactions of the Association for Computational Linguistics 4 (2016), 357–370.
- [5]. Dandala Bharath, Mahajan Diwakar, and Poddar Ananya. 2018. IBM research system at TAC 2018: Deep learning architectures for drug-drug interaction extraction from structured product labels. In Proceedings of the 2018 Text Analysis Conference (TAC'18).
- [6]. Demner-Fushman Dina, Fung Kin Wah, Do Phong, Boyce Richard D., and Goodwin Travis. 2018. Overview of the TAC 2018 drug-drug interaction extraction from drug labels track. In Proceedings of the 2018 Text Analysis Conference (TAC'18).
- [7]. Devlin Jacob, Chang Ming-Wei, Lee Kenton, and Toutanova Kristina. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. Proceedings of the 17th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies.
- [8]. He Kaiming, Zhang Xiangyu, Ren Shaoqing, and Sun Jian. 2016. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 770–778.
- [9]. Herrero-Zazo María, Segura-Bedmar Isabel, Paloma Martínez, and Thierry Declerck. 2013. The DDI corpus: An annotated corpus with pharmacological substances and drug–drug interactions. Journal of Biomedical Informatics 46, 5 (2013), 914–920. [PubMed: 23906817]
- [10]. Jones Rosie, McCallum Andrew, Nigam Kamal, and Riloff Ellen. 1999. Bootstrapping for text learning tasks. In IJCAI-99 Workshop on Text Mining: Foundations, Techniques and Applications, Vol. 1.
- [11]. Kavuluru Ramakanth, Rios Anthony, and Tran Tung. 2017. Extracting drug-drug interactions with word and character-level recurrent neural networks. In 5th IEEE International Conference on Healthcare Informatics (ICHI'17). IEEE, 5–12.
- [12]. Kim J-D, Ohta Tomoko, Tateisi Yuka, and Tsujii Jun'ichi. 2003. GENIA corpus—A semantically annotated corpus for bio-textmining. Bioinformatics 19, Suppl_1 (2003), i180–i182. [PubMed: 12855455]
- [13]. Kim Yoon. 2014. Convolutional neural networks for sentence classification. In Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP'14). Association for Computational Linguistics, Doha, Qatar, 1746–1751. <http://www.aclweb.org/anthology/D14-1181>.
- [14]. Kohn Linda T., Corrigan Janet M., and Donaldson Molla S.. 2000. To Err Is Human: Building a Safer Health System. Vol. 6. National Academies Press.
- [15]. Lee Jinhyuk, Yoon Wonjin, Kim Sungdong, Kim Donghyeon, Kim Sunkyu, So Chan Ho, and Kang Jaewoo. 2020. BioBERT: A pre-trained biomedical language representation model for biomedical text mining. Bioinformatics 36, 4 (2020), 1234–1240. [PubMed: 31501885]
- [16]. Levinson Daniel R.. 2010. Adverse events in hospitals: National incidence among Medicare beneficiaries. Department of Health and Human Services Office of the Inspector General.

- [17]. Li Fei, Zhang Meishan, Fu Guohong, and Ji Donghong. 2017. A neural joint model for entity and relation extraction from biomedical text. *BMC Bioinformatics*18, 1 (2017), 198. [PubMed: 28359255]
- [18]. Lim Sangrak, Lee Kyubum, and Kang Jaewoo. 2018. Drug drug interaction extraction from the literature using a recursive neural network. *PloS One*13, 1 (2018), e0190926. [PubMed: 29373599]
- [19]. Liu Shengyu, Chen Kai, Chen Qingcai, and Tang Buzhou. 2016. Dependency-based convolutional neural network for drug-drug interaction extraction. In *2016 IEEE International Conference on Bioinformatics and Biomedicine (BIBM'16)*. IEEE, 1074–1080.
- [20]. Liu Shengyu, Tang Buzhou, Chen Qingcai, and Wang Xiaolong. 2016. Drug-drug interaction extraction via convolutional neural networks. *Computational and Mathematical Methods in Medicine*2016 (2016), 1–8.
- [21]. Luo Yuan, Uzuner Özlem, and Szolovits Peter. 2016. Bridging semantics and syntax with graph algorithms—State-of-the-art of extracting biomedical relations. *Briefings in Bioinformatics*18, 1 (2016), 160–178. [PubMed: 26851224]
- [22]. Pyysalo Sampo, Ginter Filip, Moen Hans, Salakoski Tapio, and Ananiadou Sophia. 2013. Distributional semantics resources for biomedical text processing. In *Proceedings of 5th International Symposium on Languages in Biology and Medicine*. 39–44.
- [23]. Ratinov Lev and Roth Dan. 2009. Design challenges and misconceptions in named entity recognition. In *Proceedings of the 13th Conference on Computational Natural Language Learning*. Association for Computational Linguistics, 147–155.
- [24]. Sahu Sunil Kumar and Anand Ashish. 2018. Drug-drug interaction extraction from biomedical texts using long short-term memory network. *Journal of Biomedical Informatics* 86 (2018), 15–24. [PubMed: 30142385]
- [25]. Settles Burr. 2012. Active learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*6, 1 (2012), 1–114.
- [26]. Suárez-Paniagua Víctor, Segura-Bedmar Isabel, and Martínez Paloma. 2017. Exploring convolutional neural networks for drug–drug interaction extraction. *Database*2017 (2017), 1–15.
- [27]. Sun Xia, Dong Ke, Ma Long, Sutcliffe Richard, He Feijuan, Chen Sushing, and Feng Jun. 2019. Drug-drug interaction extraction via recurrent hybrid convolutional neural networks with an improved focal loss. *Entropy*21, 1 (2019), 37.
- [28]. Tang Siliang, Zhang Qi, Zheng Tianpeng, Zhou Mengdi, Chen Zhan, Shen Lixing, Ren Xiang, Zhuang Yueting, Pu Shiliang, and Wu Fei Wu. 2018. Two step joint model for drug drug interaction extraction. In *Proceedings of the 2018 Text Analysis Conference (TAC'18)*.
- [29]. Tran Tung, Kavuluru Ramakanth, and Kilicoglu Halil. 2018. A multi-task learning framework for extracting drugs and their interactions from drug labels. In *Proceedings of the 2018 Text Analysis Conference (TAC'18)*.
- [30]. Viterbi Andrew. 1967. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory*13, 2 (1967), 260–269.
- [31]. Zhang Yuhao, Qi Peng, and Manning Christopher D.. 2018. Graph convolution over pruned dependency trees improves relation extraction. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.
- [32]. Zhao Zhehuan, Yang Zhihao, Luo Ling, Lin Hongfei, and Wang Jian. 2016. Drug drug interaction extraction from biomedical literature using syntax convolutional neural network. *Bioinformatics*32, 22 (2016), 3444–3453. [PubMed: 27466626]

CCS Concepts:

• **Information systems** → **Information extraction**; • **Computing methodologies** → *Multi-task learning, Neural networks*;

Fig. 1. Simple examples illustrating the end-to-end DDI extraction task. We first (1) identify mentions including precipitants; for each precipitant, we (2) determine the type of interaction and, based on interaction type, (3) determine the interaction outcome. In the case of PD interactions (left), the outcome corresponds to one of the previously identified *effect* spans. In the case of PK interactions (right), the outcome corresponds to an NCI Thesaurus code indicating the type and level of increases/decreases in functional measurements.

Fig. 2. Overview of the neural network architecture for a simplified example from the drug label Adenocard. Here, the ground truth indicates that *digitalis* is a pharmacodynamic precipitant associated with the effect “ventricular fibrillation.” The PK predictive component is omitted given there are no precipitants involved in a PK interaction.

Fig. 3. An example sentence from the drug label for Savella along with the resulting prediction and ground-truth labels. Red arrows indicate interaction outcome.

Fig. 4. An example sentence from the drug label for Aubagio along with the resulting prediction and ground-truth labels. Red arrows indicate interaction outcome, where C54357 is a PK label corresponding to the NCI Thesaurus code for “Increased Concomitant Drug Level.”

Table 1.

Characteristics of Various Datasets

	*DDI2013	*NLM180	TR22	Test Set 1	Test Set 2
Number of drug labels	715	180	22	57	66
Total number of sentences	6,489	5,757	603	8,195	4,256
Number of sentences per drug label (average)	9	32	27	144	64
Number of words per sentence (average)	21	23	24	22	23
Proportion of sentences with <i>annotations</i>	70%	27%	51%	23%	23%
Number of mentions per <i>annotated</i> sentence (average)	2.3	4.0	3.8	3.7	3.6
Proportion of mentions that are precipitant	100%	57%	53%	56%	55%
Proportion of mentions that are trigger	–	20%	28%	30%	33%
Proportion of mentions that are effect	–	23%	19%	14%	12%
Proportion of interactions that are pharmacodynamic	14%	47%	49%	33%	28%
Proportion of interactions that are pharmacokinetic	9%	25%	21%	28%	47%
Proportion of interactions that are unspecified	77%	28%	30%	39%	25%

* Statistics for NLM180 and DDI2013 were computed on mapped examples (based on our own annotation mapping scheme) and not based on the original data.

Table 2.

Example of the Sequence Labeling Scheme for the Sentence in Figure 1, Where LABELDRUG Is Substitute for Adenocard

0	0	0	0	0	0	0	U-DYN
The	use	of	LABELDRUG	in	patients	receiving	digitalis
0	0	0	B-TRI	L-TRI	B-EFF	L-EFF	0
may	be	rarely	associated	with	ventricular	fibrillation	.

Table 3.

Model Configuration Obtained Through Random Search over 11-Fold Cross-Validation of TR22 (Training Data)

Setting	Value	Setting	Value
Learning Rate	0.001	Context Embedding Size (ρ)	100
Dropout Rate	0.5	GC Hidden Size (β)	100
Character Embedding Size (π)	25	GC Attention Size (α)	25
Character Representation Size (η)	50	Sequence LSTM Hidden Size (γ)	200
Word Embedding Size (δ)	200	Outcome CNN Filter Count (μ)	50

Table 4.

Main Results Based on 95% Confidence Interval around Mean Precision, Recall, and F1 Based on Evaluating
N = 100 Ensembles for Each Model

Method	Training Data	Test 1/Entity			Test 1/Relation			Test 2/Entity			Test 2/Relation			Overall		
		P (%)	R (%)	F (%)	P (%)	R (%)	F (%)	P (%)	R (%)	F (%)	P (%)	R (%)	F (%)	P (%)	R (%)	F (%)
BL	TR22	23.82	42.04	30.39	14.74	18.38	16.35	26.15	39.69	31.51	12.48	15.43	13.79	19.30 ± 0.12	28.88 ± 0.12	23.01 ± 0.06
GCA	TR22	32.87	32.35	32.59	22.70	13.95	17.27	38.82	31.31	34.65	19.26	11.63	14.49	28.41 ± 0.09	22.31 ± 0.16	24.75 ± 0.11
BL ⁽¹⁾	TR22 + NLM180	27.05	39.87	32.22	19.94	22.20	21.00	32.49	41.92	36.60	21.82	23.93	22.82	25.32 ± 0.09	31.98 ± 0.11	28.16 ± 0.06
GCA	TR22 + NLM180	38.30	31.20	34.38	27.97	15.14	19.63	44.13	31.18	36.53	31.79	15.76	21.06	35.55 ± 0.18	23.32 ± 0.20	27.90 ± 0.17
GCA/B	TR22 + NLM180	45.30	34.90	39.41	32.23	18.92	23.83	49.90	32.77	39.55	32.20	17.23	22.43	39.91 ± 0.22	25.96 ± 0.22	31.30 ± 0.20
BL	TR22 + NLM180 + DDI2013	29.27	41.93	34.47	22.93	25.42	24.11	38.73	43.79	41.10	27.11	27.32	27.21	29.51 ± 0.10	34.61 ± 0.10	31.72 ± 0.06
GC ⁽²⁾	TR22 + NLM180 + DDI2013	38.85	36.30	37.52	29.82	18.59	22.88	43.74	34.88	38.80	31.14	16.40	21.48	35.89 ± 0.20	26.54 ± 0.20	30.17 ± 0.19
GCA	TR22 + NLM180 + DDI2013	41.58	38.24	39.83	31.84	20.49	24.93	47.54	36.12	41.04	32.07	17.81	22.90	38.26 ± 0.16	28.17 ± 0.12	32.18 ± 0.12
GCA/B	TR22 + NLM180 + DDI2013	47.61	37.91	42.20	34.99	20.86	26.12	52.12	34.77	41.70	34.78	18.53	24.17	42.37 ± 0.23	28.02 ± 0.20	33.55 ± 0.20
GCA/B + BL ⁽³⁾	TR22 + NLM180 + DDI2013	36.88	45.09	40.56	28.22	25.13	26.57	46.99	44.59	45.75	32.98	25.76	28.91	36.27 ± 0.17	35.14±0.15	35.45 ± 0.12

⁽¹⁾Our original challenge submission using a BiLSTM-based approach and trained on only TR22 and NLM180.

⁽²⁾For reference, we include an evaluation of the standard GC without attention gating.

⁽³⁾Our current best is a combination of GCA (with BERT embeddings) and BL by ensembling.

Table 5.

Comparison of Our Method with Comparable (Based on Training Data) Methods, Based on *Single Runs*, among Teams in the Top 3

Method	Training Data	Test 1/Entity			Test 1/Relation			Test 2/Entity			Test 2/Relation			Overall		
		P (%)	R (%)	F (%)	P (%)	R (%)	F (%)	P (%)	R (%)	F (%)	P (%)	R (%)	F (%)	P (%)	R (%)	F (%)
Dandala et al. [5]	TR22 + NLM180	41.94	23.19	29.87	25.24	16.10	19.66	44.61	29.31	35.38	22.99	16.83	19.43	33.70	21.36	26.09
Tran et al. [29]	TR22 + NLM180	29.50	37.45	33.00	22.08	21.13	21.59	36.68	40.02	38.28	22.53	21.13	23.55	27.70	29.93	29.11
GCA/B + BL (Ours)	TR22 + NLM180	35.51	43.43	39.07	27.02	24.05	25.45	43.27	44.24	43.75	30.39	25.22	27.57	34.05	34.23	33.96

Table 6.

Confusion Matrix for Interaction Type

		<u>Predicted</u>		
		PD	PK	UN
<u>Actual</u>	PD	788	37	68
	PK	57	353	147
	UN	170	10	599