

Cross-Task Instance Representation Interactions and Label Dependencies for Joint Information Extraction with Graph Convolutional Networks

Minh Van Nguyen, Viet Dac Lai and Thien Huu Nguyen

Department of Computer and Information Science

University of Oregon, Eugene, Oregon, USA

{minhmv, vietl, thien}@cs.uoregon.edu

Abstract

Existing works on information extraction (IE) have mainly solved the four main tasks separately (entity mention recognition, relation extraction, event trigger detection, and argument extraction), thus failing to benefit from inter-dependencies between tasks. This paper presents a novel deep learning model to simultaneously solve the four tasks of IE in a single model (called FourIE). Compared to few prior work on jointly performing four IE tasks, FourIE features two novel contributions to capture inter-dependencies between tasks. First, at the representation level, we introduce an interaction graph between instances of the four tasks that is used to enrich the prediction representation for one instance with those from related instances of other tasks. Second, at the label level, we propose a dependency graph for the information types in the four IE tasks that captures the connections between the types expressed in an input sentence. A new regularization mechanism is introduced to enforce the consistency between the golden and predicted type dependency graphs to improve representation learning. We show that the proposed model achieves the state-of-the-art performance for joint IE on both monolingual and multilingual learning settings with three different languages.

1 Introduction

Information Extraction (IE) is an important and challenging task in Natural Language Processing (NLP) that aims to extract structured information from unstructured texts. Following the terminology for IE in the popular ACE 2005 program (Walker et al., 2006), we focus on four major IE tasks in this work: entity mention extraction (EME), relation extraction (RE), event trigger detection (ETD), and event argument extraction (EAE).

Given an input sentence, a vast majority of prior work has solved the four tasks in IE independently at both instance and task levels (called independent




Figure 1: A sentence example with the annotations for the four IE tasks. Blue words corresponds to entity mentions while red words are event triggers. Also, orange edges represent relations while green edges indicate argument roles.

prediction models). First, at the instance level, each IE task often requires predictions/classifications for multiple instances in a single input sentence. For instance, in RE, one often needs to predict relations for every pair of entity mentions (called relation instances) in the sentence while multiple word spans in the sentence can be viewed as multiple instances where event type predictions have to be made in ETD (trigger instances). As such, most prior work on IE has performed predictions for instances in a sentence separately by treating each instance as one example in the dataset (Zhou et al., 2005; Nguyen and Grishman, 2015a; Santos and Guimaraes, 2015; Chen et al., 2015; Nguyen and Grishman, 2015b; Lai et al., 2020). Second, at the task level, prior work on IE tends to perform the four tasks in a pipelined approach where outputs from one task are used as inputs for other tasks (e.g., EAE is followed by EME and ETD) (Li et al., 2013; Chen et al., 2015; Veyseh et al., 2020c).

Despite its popularity, the main issue of the independent prediction models is that they suffer from the error propagation between tasks and the failure to exploit the cross-task and cross-instance inter-dependencies within an input sentence to improve the performance for IE tasks. For instance, such systems are unable to benefit from the dependency that the *Victim* of a *Die* event has a high chance to

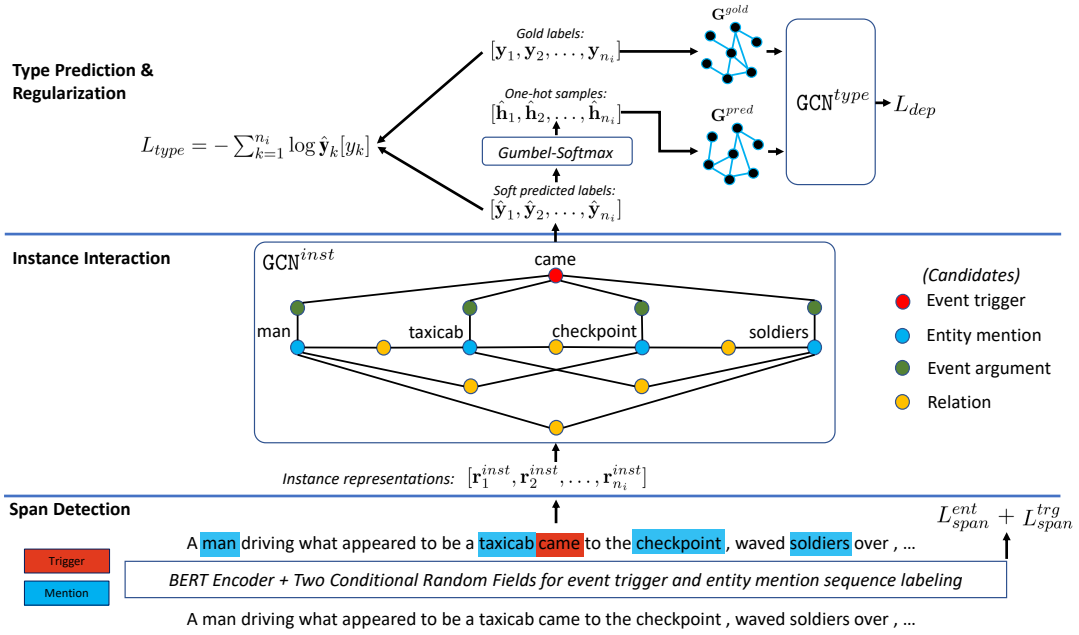


Figure 2: Overall architecture of our proposed model. At the representation level, GCN^{inst} is used to enrich the representations for instances of the four tasks. At the label level, GCN^{type} is responsible for capturing the connections between the types in the dependency graphs, thus helping the model learn the structural difference between the gold graph G^{gold} and the predicted graph G^{pred} .

also be the *Victim* of an *Attack* event in the same sentence (i.e., type or label dependencies). To address these issues, some prior work has explored joint inference models where multiple tasks of IE are performed simultaneously for all task instances in a sentence, using both feature-based models (Roth and Yih, 2004; Li et al., 2013; Miwa and Sasaki, 2014; Yang and Mitchell, 2016) and recent deep learning models (Miwa and Bansal, 2016; Zhang et al., 2019). However, such prior work has mostly considered joint models for a subset of the four IE tasks (e.g., EME+RE or ETD+EAE), thus still suffering from the error propagation issue (with the missing tasks) and failing to fully exploit potential inter-dependencies between the four tasks. To this end, this work aims to design a single model to simultaneously solve the four IE tasks for each input sentence (joint four-task IE) to address the aforementioned issues of prior joint IE work.

Few recent work has considered joint four-task IE, using deep learning to produce state-of-the-art (SOTA) performance for the tasks (Wadden et al., 2019; Lin et al., 2020). However, there are still two problems that hinder further improvement of such models. First, at the instance level, an important component of deep learning models for joint IE involves the representation vectors of the instances that are used to perform the corresponding prediction tasks for IE in an input sentence (called

predictive instance representations). For joint four-task IE, we argue that there are inter-dependencies between predictive representation vectors of related instances for the four tasks that should be modeled to improve the performance for IE. For instance, the entity type information encoded in the predictive representation vector for an entity mention can constrain the argument role that the representation vector for a related EAE instance (e.g., involving the same entity mention and some event trigger in the same sentence) should capture and vice versa. As such, prior work for joint four-task IE has only computed predictive representation vectors for instances of the tasks independently using shared hidden vectors from some deep learning layer (Wadden et al., 2019; Lin et al., 2020). Although this shared mechanism helps capture the interaction of predictive representation vectors to some extent, it fails to explicitly present the connections between related instances of different tasks and encode them into the representation learning process. Consequently, to overcome this issue, we propose a novel deep learning model for joint four-task IE (called *FourIE*) that creates a graph structure to explicitly capture the interactions between related instances of the four IE tasks in a sentence. This graph will then be consumed by a graph convolutional network (GCN) (Kipf and Welling, 2017; Nguyen and Grishman, 2018) to enrich the representation vector

for an instance with those from the related (neighboring) instances for IE.

Second, at the task level, existing joint four-task models for IE have only exploited the cross-task type dependencies in the decoding step to constrain predictions for the input sentence (by manually converting the type dependency graphs of the input sentence into global feature vectors for scoring the predictions in the beam search-based decoding) (Lin et al., 2020). The knowledge from cross-task type dependencies thus cannot contribute to the training process of the IE models. This is unfortunate as we expect that deeper integration of this knowledge into the training process could provide useful information to enhance representation learning for IE tasks. To this end, we propose to use the knowledge from cross-task type dependencies to obtain an additional training signal for each sentence to directly supervise our joint four-task IE model. In particular, our motivation is that the types expressed in a sentence for the four IE tasks can be organized into a dependency graph between the types (global type dependencies for the sentence). As such, in order for a joint model to perform well, the type dependency graph generated by its predictions for a sentence should be similar to the dependency graph obtained from the golden types (i.e., a global type constraint on the predictions in the training step). A novel regularization term is thus introduced into the training loss of our joint model to encode this constraint, employing another GCN to learn representation vectors for the predicted and golden dependency graphs to facilitate the graph similarity promotion. To our knowledge, this is the first work that employs global type dependencies to regularize joint models for IE.

Finally, our extensive experiments demonstrate the effectiveness of the proposed model on benchmark datasets in three different languages (e.g., English, Chinese, and Spanish), leading to state-of-the-art performance on different settings.

2 Problem Statement and Background

Problem Statement: The joint four-task IE problem in this work takes a sentence as the input and aims to jointly solve four tasks EAE, ETD, RE, and EAE using an unified model. As such, the goal of EME is to detect and classify entity mentions (names, nominals, pronouns) according to a set of predefined (semantic) entity types (e.g., *Person*). Similarly, ETD seeks to identify and classify event

triggers (verbs or normalization) that clearly evoke an event in some predefined set of event types (e.g., *Attack*). Note that event triggers can involve multiple words. For RE, its concern is to predict the semantic relationship between two entity mentions in the sentence. Here, the set of relations of interest is also predefined and includes a special type of *None* to indicate *no-relation*. Finally, in EAE, given an event trigger, the systems need to predict the roles (also in a predefined set with a special type *None*) that each entity mention plays in the corresponding event. Entity mentions are thus also called event argument candidates in this work. Figure 1 presents a sentence example where the expected outputs for each IE task are illustrated.

Graph Convolutional Networks (GCN): As GCNs are used extensively in our model, we present their computation process in this section to facilitate the discussion. Given a graph $\mathbf{G} = (\mathbf{V}, \mathbf{E})$ where $\mathbf{V} = \{v_1, \dots, v_u\}$ is the node set (with u nodes) and \mathbf{E} is the edge set. In GCN, the edges in \mathbf{G} are often captured via the adjacency matrix $\mathbf{A} \in \mathbb{R}^{u \times u}$. Also, each node $v_i \in \mathbf{V}$ is associated with an initial hidden vector \mathbf{v}_i^0 . As such, a GCN model involves multiple layers of abstraction in which the hidden vector \mathbf{v}_i^l for the node $v_i \in \mathbf{V}$ at the l -th layer is computed by ($l \geq 1$):

$$\mathbf{v}_i^l = \text{ReLU}\left(\frac{\sum_{j=1}^u \mathbf{A}_{ij} \mathbf{W}^l \mathbf{v}_j^{l-1} + \mathbf{b}^l}{\sum_{j=1}^u \mathbf{A}_{ij}}\right)$$

where \mathbf{W}^l and \mathbf{b}^l are trainable weight and bias at the l -th layer. Assuming N GCN layers, the hidden vectors for the nodes in \mathbf{V} at the last layer $\mathbf{v}_1^N, \dots, \mathbf{v}_u^N$ would capture richer and more abstract information for the nodes, serving as the outputs of the GCN model. This process is denoted by: $\mathbf{v}_1^N, \dots, \mathbf{v}_u^N = \text{GCN}(\mathbf{A}; \mathbf{v}_1^0, \dots, \mathbf{v}_u^0; N)$.

3 Model

Given an input sentence $\mathbf{w} = [w_1, w_2, \dots, w_n]$ (with n words), our model for joint four-task IE on \mathbf{w} involves three major components: (i) Span Detection, (ii) Instance Interaction, and (iii) Type Dependency-based Regularization.

3.1 Span Detection

This component aims to identify spans of entity mentions and event triggers in \mathbf{w} that would be used to form the nodes in the interaction graph between different instances of our four IE tasks for \mathbf{w} . As such, we formulate the span detection

problems as sequence labeling tasks where each word w_i in \mathbf{w} is associated with two BIO tags to capture the span information for entity mentions and event triggers in \mathbf{w} . Note that we do not predict entity types and event types at this step, leading to only three possible values (i.e., B, I, and O) for the tags of the words.

In particular, following (Lin et al., 2020), we first feed \mathbf{w} into the pre-trained BERT encoder (Devlin et al., 2019) to obtain a sequence of vectors $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n]$ to represent \mathbf{w} . Here, each vector \mathbf{x}_i serves as the representation vector for the word $w_i \in \mathbf{w}$ that is obtained by averaging the hidden vectors of the word-pieces of w_i returned by BERT. Afterward, \mathbf{X} is fed into two conditional random field (CRF) layers to determine the best BIO tag sequences for event mentions and event triggers for \mathbf{w} , following (Chiu and Nichols, 2016). As such, the Viterbi algorithm is used to decode the input sentence while the negative log-likelihood losses are employed as the training objectives for the span detection component of the model. For convenience, let L_{span}^{ent} and L_{span}^{trg} be the negative log-likelihoods of the gold tag sequences for entity mentions and event triggers (respectively) for \mathbf{w} . These terms will be included in the overall loss function of the model later.

3.2 Instance Interaction

Based on the tag sequences for \mathbf{w} from the previous component, we can obtain two separate span sets for the entity mentions and event triggers in \mathbf{w} (the golden spans are used in the training phase to avoid noise). For the next computation, we first compute a representation vector for each span (i, j) ($1 \leq i \leq j \leq n$) in these two sets by averaging the BERT-based representation vectors for the words in this span (i.e., $\mathbf{x}_i, \dots, \mathbf{x}_j$). For convenience, let $\mathbf{R}^{ent} = \{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_{n_{ent}}\}$ ($n_{ent} = |\mathbf{R}^{ent}|$) and $\mathbf{R}^{trg} = \{\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_{n_{trg}}\}$ ($n_{trg} = |\mathbf{R}^{trg}|$) be the sets of span representation vectors for the entity mentions and event triggers in \mathbf{w} ¹. The goal of this component is to leverage such span representation vectors to form instance representations and enrich them with instance interactions to perform necessary predictions in IE.

Instance Representation. Prediction instances in our model amount to the specific objects that we need to predict a type for one of the four IE

¹We will also refer to entity mentions and event triggers interchangeably with their span representations \mathbf{e}_i and \mathbf{t}_i in this work.

tasks. As such, the prediction instances for EME and ETD, called entity and trigger instances, correspond directly to the entity mentions and event triggers in \mathbf{R}^{ent} and \mathbf{R}^{trg} respectively (as we need to predict the entity types for $\mathbf{e}_i \in \mathbf{R}^{ent}$ and the event types for $\mathbf{t}_i \in \mathbf{R}^{trg}$ in this step). Thus, we also use \mathbf{R}^{ent} and \mathbf{R}^{trg} as the sets of initial representation vectors for the entity/event instances for EME and ETD in the following. Next, for RE, the prediction instances (called relation instances) involve pairs of entity mentions in \mathbf{R}^{ent} . To obtain the initial representation vector for a relation instance, we concatenate the representation vectors of the two corresponding entity mentions, leading to the set of representation vectors \mathbf{rel}_{ij} for relation instances: $\mathbf{R}^{rel} = \{\mathbf{rel}_{ij} = [\mathbf{e}_i, \mathbf{e}_j] \mid \mathbf{e}_i, \mathbf{e}_j \in \mathbf{R}^{ent}, i < j\}$ ($|\mathbf{R}^{rel}| = n_{ent}(n_{ent} - 1)/2$). Finally, for EAE, we form the prediction instances (called argument instances) by pairing each event trigger in \mathbf{R}^{trg} with each entity mention in \mathbf{R}^{ent} (for the argument role predictions of the entity mentions with respect to the event triggers/mentions). By concatenating the representation vectors of the paired entity mentions and event triggers, we generate the initial representation vectors \mathbf{arg}_{ij} for the corresponding argument instances: $\mathbf{R}^{arg} = \{\mathbf{arg}_{ij} = [\mathbf{t}_i, \mathbf{e}_j] \mid \mathbf{t}_i \in \mathbf{R}^{trg}, \mathbf{e}_j \in \mathbf{R}^{ent}\}$ ($|\mathbf{R}^{arg}| = n_{trg}n_{ent}$)². We also use the prediction instances and their representation vectors interchangeably in this work.

Instance Interaction. The initial representation vectors for the instances so far do not explicitly consider beneficial interactions between related instances. To address this issue, we explicitly create an interaction graph between the prediction instances for the four IE tasks to connect related instances to each other. This graph will be consumed by a GCN model to enrich instance representations with interaction information afterward. In particular, the node set \mathbf{N}^{inst} in our instance interaction graph $\mathbf{G}^{inst} = \{\mathbf{N}^{inst}, \mathbf{E}^{inst}\}$ involves all prediction instances for the four IE tasks, i.e., $\mathbf{N}^{inst} = \mathbf{R}^{ent} \cup \mathbf{R}^{trg} \cup \mathbf{R}^{rel} \cup \mathbf{R}^{arg}$. The edge set \mathbf{E}^{inst} then captures instance interactions by connecting the instance nodes in \mathbf{N}^{inst} that involve the same entity mentions or event triggers (i.e., two instances are related if they concern the same entity mention or event trigger). As such, the edges in \mathbf{E}^{inst} are created as follows:

²In our implementation, \mathbf{R}^{rel} and \mathbf{R}^{arg} are transformed into vectors of the same size with those in \mathbf{R}^{ent} and \mathbf{R}^{trg} (using one-layer feed forward networks) for future computation.

(i) An entity instance node \mathbf{e}_i is connected to all relation instance nodes of the forms $\mathbf{rel}_{ij} = [\mathbf{e}_i, \mathbf{e}_j]$ and $\mathbf{rel}_{ki} = [\mathbf{e}_k, \mathbf{e}_i]$ (sharing entity mention \mathbf{e}_i).

(ii) An entity instance node \mathbf{e}_j is connected to all argument instance nodes of the form $\mathbf{arg}_{ij} = [\mathbf{t}_i, \mathbf{e}_j]$ (sharing entity mention \mathbf{e}_j).

(iii) A trigger node \mathbf{t}_i is connected to all argument instance nodes of the form $\mathbf{arg}_{ij} = [\mathbf{t}_i, \mathbf{e}_j]$ (i.e., sharing event trigger \mathbf{t}_i).

GCN. To enrich the representation vector for an instance in \mathbf{N}^{inst} with the information from the related (neighboring) nodes, we feed \mathbf{G}^{inst} into a GCN model (called GCN^{inst}). For convenience, we rename the initial representation vectors of all the instance nodes in \mathbf{N}^{inst} by: $\mathbf{N}^{inst} = \{\mathbf{r}_1, \dots, \mathbf{r}_{n_i}\}$ ($n_i = |\mathbf{N}^{inst}|$). Also, let $\mathbf{A}^{inst} \in \{0, 1\}^{n_i \times n_i}$ be the adjacency matrix of the interaction graph \mathbf{G}^{inst} where $\mathbf{A}_{ij}^{inst} = 1$ if the instance nodes \mathbf{r}_i and \mathbf{r}_j are connected in \mathbf{G}^{inst} or $i = j$ (for self-connections). The interaction-enriched representation vectors for the instances in \mathbf{N}^{inst} are then computed by the GCN^{inst} model: $\mathbf{r}_1^{inst}, \dots, \mathbf{r}_{n_i}^{inst} = \text{GCN}^{inst}(\mathbf{A}^{inst}; \mathbf{r}_1, \dots, \mathbf{r}_{n_i}; N_i)$ where N_i is the number of layers for the GCN^{inst} model.

Type Embedding and Prediction. Finally, the enriched instance representation vectors $\mathbf{r}_1^{inst}, \dots, \mathbf{r}_{n_i}^{inst}$ will be used to perform the predictions for the four IE tasks. In particular, let $t_k \in \{ent, trg, rel, arg\}$ be the corresponding task index and y_k be the ground-truth type (of the task t_k) for the prediction instance \mathbf{r}_k in \mathbf{N}^{inst} . Also, let $\mathcal{T} = \mathcal{T}^{ent} \cup \mathcal{T}^{trg} \cup \mathcal{T}^{rel} \cup \mathcal{T}^{arg}$ be the union of the possible entity types (in \mathcal{T}^{ent} for EME), event types (in \mathcal{T}^{trg} for ETD), relations (in \mathcal{T}^{rel} for RE), and argument roles (in \mathcal{T}^{arg} for EAE) in our problem ($y_k \in \mathcal{T}^{t_k}$). Note that \mathcal{T}^{rel} and \mathcal{T}^{arg} contain the special types *None*. To prepare for the type predictions and the type dependency modeling in the next steps, we associate each type in \mathcal{T} with an embedding vector (of the same size as \mathbf{e}_i and \mathbf{t}_i) that is initialized randomly and updated during our training process. For convenience, let $\mathcal{T} = [\bar{\mathbf{t}}_1, \dots, \bar{\mathbf{t}}_{n_t}]$ where $\bar{\mathbf{t}}_i$ is used interchangeably for both a type and its embedding vector in \mathcal{T} (n_t is the total number of types). As such, to perform the prediction for an instance \mathbf{r}_k in \mathbf{N}^{inst} , we compute the dot products between \mathbf{r}_k^{inst} and each type embedding vectors in $\mathcal{T}^{t_k} \cap \mathcal{T}$ to estimate the possibilities that \mathbf{r}_k has a type in \mathcal{T}^{t_k} . Afterward, these scores are normalized by the softmax function to obtain the probability distribution $\hat{\mathbf{y}}_k$ over the possible types in

\mathcal{T}^{t_k} for \mathbf{r}_k : $\hat{\mathbf{y}}_k = \text{softmax}(\mathbf{r}_k^{inst} \bar{\mathbf{t}}^T | \bar{\mathbf{t}} \in \mathcal{T}^{t_k} \cap \mathcal{T})$. In the decoding phase, the predicted type \hat{y}_k for \mathbf{r}_k is obtained via the argmax function (greedy decoding): $\hat{y}_k = \text{argmax} \hat{\mathbf{y}}_k$. The negative log-likelihood over all the prediction instances is used to train the model: $L_{type} = -\sum_{k=1}^{n_i} \log \hat{\mathbf{y}}_k[y_k]$.

3.3 Type Dependency-based Regularization

In this section, we aim to obtain the type dependencies across tasks and use them to supervise the model in the training process (to improve the representation vectors for IE). As presented in the introduction, our motivation is to generate global dependency graphs between types of different IE tasks for each input sentence whose representations are leveraged to regularize the model during training. In particular, starting with the golden types $\mathbf{y} = y_1, y_2, \dots, y_{n_i}$ and the predicted types $\hat{\mathbf{y}} = \hat{y}_1, \hat{y}_2, \dots, \hat{y}_{n_i}$ for the instance nodes in \mathbf{N}^{inst} , we build two dependency graphs \mathbf{G}^{gold} and \mathbf{G}^{pred} to capture the global type dependencies for the tasks (called the golden and predicted dependency graphs respectively). Afterward, to supervise the training process, we seek to constrain the model so the predicted dependency graph \mathbf{G}^{pred} is similar to the golden graph \mathbf{G}^{gold} (i.e., using the dependency graphs as the bridges to inject the global type dependency knowledge in \mathbf{G}^{gold} into the model).

Dependency Graph Construction. Both \mathbf{G}^{gold} and \mathbf{G}^{pred} involve the types of all the four IE tasks in \mathcal{T} as the nodes. To encode the type dependencies, the connections/edges in \mathbf{G}^{gold} are computed based on the golden types $\mathbf{y} = y_1, y_2, \dots, y_{n_i}$ for the instance nodes in \mathbf{N}^{inst} as follows:

(i) For each relation instance node $\mathbf{r}_k = [\mathbf{e}_i, \mathbf{e}_j] \in \mathbf{N}^{inst}$ that has the golden type $y_k \neq \text{None}$, the relation type node y_k is connected to the nodes of the golden entity types for the corresponding entity mentions \mathbf{e}_i and \mathbf{e}_j (called **entity_relation type edges**).

(ii) For each argument instance node $\mathbf{r}_k = [\mathbf{t}_i, \mathbf{e}_j]$ that has the role type $y_k \neq \text{None}$, the role type node y_k is connected to both the node for the golden event type of \mathbf{t}_i (called **event_argument type edges**) and the node for the golden entity type of \mathbf{e}_j (called **entity_argument type edges**).

The same procedure can be applied to build the predicted dependency graph \mathbf{G}^{pred} based on the predicted types $\hat{\mathbf{y}} = \hat{y}_1, \hat{y}_2, \dots, \hat{y}_{n_i}$. Also, for convenience, let \mathbf{A}^{gold} and \mathbf{A}^{pred} (of size $n_t \times n_t$) be the binary adjacency matrices of \mathbf{G}^{gold} and \mathbf{G}^{pred}

(including the self-loops) respectively.

Regularization. In the next step, we obtain the representation vectors for the dependency graphs \mathbf{G}^{gold} and \mathbf{G}^{pred} by feeding them into a GCN model (called GCN^{type}). This GCN model has N_t layers and uses the initial type embeddings $\mathcal{T} = [\bar{\mathbf{t}}_1, \dots, \bar{\mathbf{t}}_{n_t}]$ as the inputs. In particular, the outputs of GCN^{type} for the two graphs involve $\bar{\mathbf{t}}_1^{gold}, \dots, \bar{\mathbf{t}}_{n_t}^{gold} = \text{GCN}^{type}(\mathbf{A}^{gold}; \bar{\mathbf{t}}_1, \dots, \bar{\mathbf{t}}_{n_t}; N_t)$ and $\bar{\mathbf{t}}_1^{pred}, \dots, \bar{\mathbf{t}}_{n_t}^{pred} = \text{GCN}^{type}(\mathbf{A}^{pred}; \bar{\mathbf{t}}_1, \dots, \bar{\mathbf{t}}_{n_t}; N_t)$ that encode the underlying information for the type dependencies presented in \mathbf{G}^{gold} and \mathbf{G}^{pred} . Finally, to promote the similarity of the type dependencies in \mathbf{G}^{gold} and \mathbf{G}^{pred} , we introduce the mean square difference between their GCN^{type} -induced representation vectors into the overall loss function for minimization: $L_{dep} = \sum_{i=1}^{n_t} \|\bar{\mathbf{t}}_i^{gold} - \bar{\mathbf{t}}_i^{pred}\|_2^2$.

Our final training loss is thus: $L = L_{span}^{ent} + L_{span}^{trg} + L_{type} + \lambda L_{dep}$ (λ is a trade-off parameter). **Approximating \mathbf{A}^{pred} .** We distinguish two types of parameters in our model so far, i.e., the parameters used to compute instance representations, e.g., those in BERT and \mathbf{G}^{inst} (called θ^{inst}), and the parameters for type dependency regularization, i.e., those for the type embeddings $\bar{\mathbf{t}}_1, \dots, \bar{\mathbf{t}}_{n_t}$ and \mathbf{G}^{type} (called θ^{dep}). As such, the current implementation only enables the training signal from L_{dep} to back-propagate to the parameters θ^{dep} and disallows L_{dep} to influence the instance representation-related parameters θ^{inst} . To enrich the instance representation vectors with type dependency information, we expect L_{dep} to be deeper integrated into the model by also contributing to θ^{inst} . To achieve this goal, we note that the block of back-propagation between L_{dep} and θ^{inst} is due to their only connection in the model via the adjacency matrix \mathbf{A}^{pred} , whose values are either one or zero. As such, the values in \mathbf{A}^{pred} are not directly dependent on any parameter in θ^{inst} , making it impossible for the back-propagation to flow. To this end, we propose to approximate \mathbf{A}^{pred} with a new matrix $\hat{\mathbf{A}}^{pred}$ that directly involves θ^{inst} in its values. In particular, let \mathbf{I}^{inst} be the index set of the non-zero cells in \mathbf{A}^{pred} : $\mathbf{I}^{inst} = \{(i, j) | \mathbf{A}_{ij}^{pred} = 1\}$. As the elements in \mathbf{I}^{inst} are determined by the indexes i_1, \dots, i_{n_i} in \mathcal{T} of the predicted types $\hat{y}_1, \hat{y}_2, \dots, \hat{y}_{n_i}$ (respectively), we also seek to compute the values for the approximated matrix $\hat{\mathbf{A}}^{pred}$ based on such indexes. Accordingly, we first define the matrix $\mathbf{B} = \{b_{ij}\}_{i,j=1..n_t}$ where the element b_{ij} at the i -th row and j -th col-

umn is set to $b_{ij} = i * n_t + j$. The approximated matrix $\hat{\mathbf{A}}^{pred}$ is then obtained by:

$$\hat{\mathbf{A}}^{pred} = \sum_{(i,j) \in \mathbf{I}^{inst}} \exp(-\beta(\mathbf{B} - in_t - j)^2) \quad (1)$$

Here, $\beta > 0$ is a large constant. For each element $(i, j) \in \mathbf{I}^{inst}$, all the elements in the matrix $(\mathbf{B} - in_t - j)^2$ are strictly positive, except for the element at (i, j) , which is zero. Thus, with a large value for β , the matrix $\exp(-\beta(\mathbf{B} - in_t - j)^2)$ has the value of one at cell (i, j) and nearly zero at other cells. Consequently, the values of $\hat{\mathbf{A}}^{pred}$ at the positions in \mathbf{I}^{inst} are close to one while those at other positions are close to zero, thus approximating our expected matrix \mathbf{A}^{pred} and still directly depending on the indexes i_1, \dots, i_{n_t} .

Addressing the Discreteness of Indexes.

Even with the approximation $\hat{\mathbf{A}}^{pred}$, the back-propagation still cannot flow from L_{dep} to θ^{inst} due to the block of the discrete and non-differentiable index variables i_1, \dots, i_{n_t} . To address this issue, we propose to apply the Gumbel-Softmax distribution (Jang et al., 2017) that enables the optimization of models with discrete random variables, by providing a method to approximate one-hot vectors sampled from a categorical distribution with continuous ones.

In particular, we first rewrite each index i_k by: $i_k = \mathbf{h}_k \mathbf{c}_k^T$, where \mathbf{c}_k is a vector whose each dimension contains the index of a type in \mathcal{T}^{t_k} in the joint type set \mathcal{T} , and \mathbf{h}_k is the binary one-hot vector whose dimensions correspond to the types in \mathcal{T}^{t_k} . \mathbf{h}_k is only turned on at the position corresponding to the predicted type $\hat{y}_k \in \mathcal{T}^{t_k}$ (indexed at i_k in \mathcal{T}). In our current implementation, \hat{y}_k (thus the index i_k and the one-hot vector \mathbf{h}_k) is obtained via the argmax function: $\hat{y}_k = \text{argmax} \hat{\mathbf{y}}_k$, which causes the discreteness. As such, the Gumbel-Softmax distribution method helps to relax argmax by approximating \mathbf{h}_k with a sample $\hat{\mathbf{h}}_k = \hat{h}_{k,1}, \dots, \hat{h}_{k,|\mathcal{T}^{t_k}|}$ from the Gumbel-Softmax distribution:

$$\hat{h}_{k,j} = \frac{\exp((\log(\pi_{k,j}) + g_j)/\tau)}{\sum_{j'=1}^{|\mathcal{T}^{t_k}|} \exp((\log(\pi_{k,j'}) + g_{j'})/\tau)} \quad (2)$$

where $\pi_{k,j} = \hat{\mathbf{y}}_{k,j} = \text{softmax}_j(\mathbf{r}_k^{inst} \bar{\mathbf{t}}^T | \bar{\mathbf{t}} \in \mathcal{T}^{t_k} \cap \mathcal{T})$, $g_1, \dots, g_{|\mathcal{T}^{t_k}|}$ are the i.i.d samples drawn from Gumbel(0,1) distribution (Gumbel, 1948): $g_j = -\log(-\log(u_j))$ ($u_j \sim \text{Uniform}(0, 1)$), and τ is the temperature parameter. As $\tau \rightarrow 0$, the

sample $\hat{\mathbf{h}}_k$ would become close to our expected one-hot vector \mathbf{h}_k . Finally, we replace \mathbf{h}_k with the approximation $\hat{\mathbf{h}}_k$ in the computation for i_k : $i_k = \hat{\mathbf{h}}_k \mathbf{c}_k^T$ that directly depends on \mathbf{r}_k^{inst} and is applied in $\hat{\mathbf{A}}^{pred}$. This allows the gradients to flow from L_{dep} to the parameters θ^{inst} and completes the description of our model.

4 Experiments

Datasets. Following the prior work on joint four-task IE (Wadden et al., 2019; Lin et al., 2020), we evaluate our joint IE model (FourIE) on the ACE 2005 (Walker et al., 2006) and ERE datasets that provide annotation for entity mentions, event triggers, relations, and argument roles. In particular, we use three different versions of the ACE 2005 dataset that focus on three major joint inference settings for IE: (i) **ACE05-R** for joint inference of EME and RE, (ii) **ACE05-E** for joint inference of EME, ETD and EAE, and (iii) **ACE05-E+** for joint inference of the four tasks EME, ETD, RE, and EAE. ACE05-E+ is our main evaluation setting as it fits to our model design with the four IE tasks of interest.

Datasets	Split	sents	ents	rels	events
ACE05-R	Train	10,051	26,473	4,788	-
	Dev	2,424	6,362	1,131	-
	Test	2,050	5,476	1,151	-
ACE05-E	Train	17,172	29,006	4,664	4,202
	Dev	923	2,451	560	450
	Test	832	3,017	636	403
ACE05-E+	Train	19,240	47,525	7,152	4,419
	Dev	902	3,422	728	468
	Test	676	3,673	802	424
ERE-EN	Train	14,219	38,864	5,045	6,419
	Dev	1,162	3,320	424	552
	Test	1,129	3,291	477	559
ACE05-CN	Train	6,841	29,657	7,934	2,926
	Dev	526	2,250	596	217
	Test	547	2,388	672	190
ERE-ES	Train	7,067	11,839	1,698	3,272
	Dev	556	886	120	210
	Test	546	811	108	269

Table 1: Numbers of sentences (i.e., **sents**), entity mentions (i.e., **ents**), relations (i.e., **rels**), and events (i.e., **events**) in the datasets.

For ERE, following (Lin et al., 2020), we combine the data from three datasets for English (i.e., LDC2015E29, LDC2015E68, and LDC2015E78) that are created under the Deep Exploration and Filtering of Test (DEFT) program (called **ERE-EN**). Similar to ACE05-E+, ERE-EN is also used to evaluate the joint models on four IE tasks.

To demonstrate the portability of our model to other languages, we also apply FourIE to the joint

four-IE datasets on Chinese and Spanish. Following (Lin et al., 2020), we use the ACE 2005 dataset for the evaluation on Chinese (called **ACE05-CN**) and the ERE dataset (LDC2015E107) for Spanish (called **ERE-ES**).

To ensure a fair comparison, we adopt the same data pre-processing and splits (train/dev/test) in prior work (Lin et al., 2020) for all the datasets. As such, ACE05-R, ACE05-E, ACE05-E+, and ACE05-CN involve 7 entity types, 6 relation types, 33 event types, and 22 argument roles while ERE-ES and ERE-EN include 7 entity types, 5 relation types, 38 event types, and 20 argument roles. The statistics for the datasets are shown in Table 1.

Hyper-parameters and Evaluation Criteria. We fine-tune the hyper-parameters for our model using the development data. The suggested values are shown in the appendix. To achieve a fair comparison with (Lin et al., 2020), we employ the *bert-large-cased* model for the English datasets and *bert-multilingual-cased* model for the Chinese and Spanish datasets. Finally, we follow the same evaluation script and correctness criteria for entity mentions, event triggers, relations, and argument as in prior work (Lin et al., 2020). The reported results are the average performance of 5 model runs using different random seeds.

Performance Comparison. We compare the proposed model **FourIE** with two prior models for joint four-task IE: (i) **DyGIE++** (Wadden et al., 2019): a BERT-based model with span graph propagation, and (ii) **OneIE** (Lin et al., 2020): *the current state-of-the-art (SOTA) model* for joint four-task IE based on BERT and type dependency constraint at the decoding step. Table 2 presents the performance (F1 scores) of the models on the test data of the English datasets. Note that in the tables, the prefixes “Ent”, “Trg”, “Rel”, and “Arg” represent the extraction tasks for entity mentions, event triggers, relations, and arguments respectively while the suffixes “-I” and “-C” correspond to the identification performance (only concerning the offset correctness) and identification+classification performance (evaluating both offsets and types).

As can be seen from the table, FourIE is consistently better than the two baseline models (DyGIE++ and OneIE) across different datasets and tasks. The performance improvement is significant for almost all the cases and clearly demonstrates the effectiveness of the proposed model.

Finally, Table 3 reports the performance of

Datasets	Task	DyGIE++	OneIE	FourIE	$\Delta\%$
ACE05-R	Ent-C	88.6	88.8	88.9	0.1
	Rel-C	63.4	67.5	68.9 †	1.4
ACE05-E	Ent-C	89.7	90.2	91.3 †	1.1
	Trg-I	-	78.2	78.3	0.1
	Trg-C	69.7	74.7	75.4 †	0.7
	Arg-I	53.0	59.2	60.7 †	1.5
	Arg-C	48.8	56.8	58.0 †	1.2
ACE05-E+	Ent-C	-	89.6	91.1 †	1.5
	Rel-C	-	58.6	63.6 †	5.0
	Trg-I	-	75.6	76.7 †	1.1
	Trg-C	-	72.8	73.3 †	0.5
	Arg-I	-	57.3	59.5 †	2.2
ERE-EN	Arg-C	-	54.8	57.5 †	2.7
	Ent-C	-	87.0	87.4	0.4
	Rel-C	-	53.2	56.1 †	2.9
	Trg-I	-	68.4	69.3 †	0.9
	Trg-C	-	57.0	57.9 †	0.9
ERE-EN	Arg-I	-	50.1	52.2 †	2.1
	Arg-C	-	46.5	48.6 †	2.1

Table 2: F1 scores of the models on the test data of English datasets. Δ indicates the performance difference between FourIE and OneIE. Rows with † designate the significant improvement ($p < 0.01$) of FourIE over OneIE.

FourIE and OneIE on the Chinese and Spanish datasets (i.e., ACE05-CN and ERE-ES). In addition to the monolingual setting (i.e., trained and evaluated on the same languages), following (Lin et al., 2020), we also evaluate the models on the multilingual training settings where ACE05-CN and ERE-ES are combined with their corresponding English datasets ACE05-E+ and EAE-EN (respectively) to train the models (for the four IE tasks), and the performance is then evaluated on the test sets of the corresponding languages (i.e., ACE05-CN and ERE-ES). It is clear from the table that FourIE also significantly outperforms OneIE across nearly all the different setting combinations for languages, datasets and tasks. This further illustrates the portability of FourIE to different languages.

Test Data	Train Data	Task	OneIE	FourIE	$\Delta\%$
ACE05-CN	ACE05-CN	Ent-C	88.5	88.7	0.2
		Rel-C	62.4	65.1 †	2.7
		Trg-C	65.6	66.5 †	0.9
		Arg-C	52.0	54.9 †	2.9
	ACE05-CN ACE05-E+	Ent-C	89.8	89.1	-0.7
		Rel-C	62.9	65.9 †	3.0
		Trg-C	67.7	70.3 †	2.6
ERE-ES	ERE-ES	Arg-C	53.2	56.1 †	2.9
		Ent-C	81.3	82.2 †	0.9
		Rel-C	48.1	57.9 †	9.8
	ERE-ES ERE-EN	Trg-C	56.8	57.1	0.3
		Arg-C	40.3	42.3 †	2.0
		Ent-C	81.8	82.7 †	0.9
		Rel-C	52.9	59.1 †	6.2
ERE-EN	Trg-C	59.1	61.3 †	2.2	
	Arg-C	42.3	45.4 †	3.1	

Table 3: F1 scores on Chinese and Spanish test sets. † marks the significant improvement ($p < 0.01$) of FourIE over OneIE.

Effects of GCN^{inst} and GCN^{type} . This section evaluates the contributions of the two important components in our proposed model FourIE, i.e., the instance interaction graph with GCN^{inst} and the type dependency graph with GCN^{type} . In particular, we examine the following ablated/varied models for FourIE: (i) “**FourIE-GCN^{inst}**”: this model excludes the instance interaction graph and the GCN model GCN^{inst} from FourIE so the initial instance representations \mathbf{r}_k are directly used to predict the types for the instances (replacing the enriched vectors \mathbf{r}_k^{inst}), (ii) “**FourIE-GCN^{type}**”: this model eliminates the type dependency graph and the GCN model GCN^{type} (thus the loss term L_{dep} as well) from FourIE, (iii) “**FourIE-GCN^{inst}-GCN^{type}**”: this model removes both the instance interaction and type dependency graphs from FourIE, (iv) “**FourIE-GCN^{type}+TDDecode**”: this model also excludes GCN^{type} ; however, it additionally applies the global type dependencies features to score the joint predictions for the beam search in the decoding step (the implementation for this beam search is inherited from (Lin et al., 2020) for a fair comparison), and (v) “**FourIE- $\hat{\mathbf{A}}^{pred}$** ”: instead of employing the approximation matrix $\hat{\mathbf{A}}^{pred}$ in FourIE, this model directly uses the adjacency matrix \mathbf{A}^{pred} in the L_{dep} regularizer (L_{dep} thus does not influence the instance representation-related parameters θ^{inst}). Table 4 shows the performance of the models on the development dataset of ACE05-E+ for four IE tasks.

Models	Ent-C	Rel-C	Trg-C	Arg-C
FourIE	89.6	64.3	71.0	59.0
FourIE-GCN^{inst}	89.1	62.3	70.3	57.5
FourIE-GCN^{type}	88.5	61.8	69.9	56.6
FourIE-GCN^{inst}-GCN^{type}	88.2	59.3	68.9	56.1
FourIE-GCN^{type}+TDDecode	88.8	59.6	70.8	56.8
FourIE-$\hat{\mathbf{A}}^{pred}$	89.0	62.3	70.2	57.6

Table 4: F1 scores of the models on the ACE05-E+ dev data.

The most important observation from the table is that both GCN^{inst} and GCN^{type} are necessary for FourIE to achieve the highest performance for the four IE tasks. Importantly, replacing GCN^{type} in FourIE with the global type dependency features for decoding (i.e., “**FourIE-GCN^{type}+TDDecode**”) as in (Lin et al., 2020) or eliminating the approximation $\hat{\mathbf{A}}^{pred}$ for L_{dep} produces inferior performance, especially for relation and argument extraction. This clearly demonstrates the benefits for deeply integrating knowledge from type dependen-

cies to influence representation learning parameters with L_{dep} for joint four-task IE.

Contributions of Type Dependency Edges. Our type dependency graphs \mathbf{G}^{gold} and \mathbf{G}^{pred} involves three categories of edges, i.e., entity_relation, entity_argument, and event_argument type edges. Table 5 presents the performance of FourIE (on the development data of ACE05-E+) when each of these edge categories is excluded from our type dependency graph construction.

Models	Ent-C	Rel-C	Trg-C	Arg-C
FourIE	89.6	64.3	71.0	59.0
FourIE - entity_relation	88.7	61.9	71.0	57.5
FourIE - entity_argument	89.3	63.2	70.0	56.9
FourIE - event_argument	89.5	64.1	69.8	57.7

Table 5: F1 scores of the ablated models for type dependency edges on the ACE05-E+ dev data.

The table clearly shows the importance of different categories of type dependency edges for FourIE as the elimination of any category would generally hurt the performance of the model. In addition, we see that the contribution level of the type dependency edges intuitively varies according to the tasks of consideration. For instance, entity_relation type edges are helpful mainly for entity mention, relation and argument extraction. Finally, an error analysis is conducted in the appendix to provide insights about the benefits of the type dependency graphs \mathbf{G}^{gold} and \mathbf{G}^{pred} for FourIE (i.e., by comparing the outputs of FourIE and “**FourIE-GCN^{type}**”).

5 Related Work

The early joint methods for IE have employed feature engineering to capture the dependencies between IE tasks, including Integer Linear Programming for Global Constraints (Roth and Yih, 2004; Li et al., 2011), Markov Logic Networks (Riedel et al., 2009; Venugopal et al., 2014), Structured Perceptron (Li et al., 2013, 2014; Miwa and Sasaki, 2014; Judea and Strube, 2016), and Graphical Models (Yu and Lam, 2010; Yang and Mitchell, 2016).

Recently, the application of deep learning has facilitated the joint modeling for IE via shared parameter mechanisms across tasks. These joint models have focused on different subsets of the IE tasks, including EME and RE (Zheng et al., 2017; Katiyar and Cardie, 2017; Bekoulis et al., 2018; Fu et al., 2019; Luan et al., 2019; Sun et al., 2019; Veyseh et al., 2020b,a), event and temporal RE (Han et al., 2019), and ETD and EAE (Nguyen et al., 2016; Zhang et al., 2019; Nguyen and Nguyen, 2019).

However, none of these work has explored joint inference for four IE tasks EME, ETD, RE, and EAE as we do. The two most related works to ours include (Wadden et al., 2019) that leverages the BERT-based information propagation via dynamic span graphs, and (Lin et al., 2020) that exploits BERT and global type dependency features to constrain the decoding step. Our model is different from these works in that we introduce a novel interaction graph for instance representations for four IE tasks and a global type dependency graph to directly inject the knowledge into the training process.

6 Conclusion

We present a novel deep learning framework to jointly solve four IE tasks (EME, ETD, RE, and EAE). Our model attempts to capture the interdependencies between instances of the four tasks and their types based on instance interaction and type dependency graphs. GCN models are employed to induce representation vectors to perform type predictions for task instances and regularize the training process. The experiments demonstrate the effectiveness of the proposed model, leading to SOTA performance over multiple datasets on English, Chinese, and Spanish. In the future, we plan to extend the model to include more IE tasks (e.g., coreference resolution).

Acknowledgments

This research has been supported by the Army Research Office (ARO) grant W911NF-21-1-0112. This research is also based upon work supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via IARPA Contract No. 2019-19051600006 under the Better Extraction from Text Towards Enhanced Retrieval (BETTER) Program. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ARO, ODNI, IARPA, the Department of Defense, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein. This document does not contain technology or technical data controlled under either the U.S. International Traffic in Arms Regulations or the U.S. Export Administration Regulations.

References

- Giannis Bekoulis, Johannes Deleu, Thomas Demeester, and Chris Develder. 2018. Adversarial training for multi-context joint entity and relation extraction. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.
- Yubo Chen, Liheng Xu, Kang Liu, Daojian Zeng, and Jun Zhao. 2015. Event extraction via dynamic multi-pooling convolutional neural networks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*.
- Jason PC Chiu and Eric Nichols. 2016. Named entity recognition with bidirectional lstm-cnns. In *Transactions of the Association for Computational Linguistics*.
- 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*.
- Tsu-Jui Fu, Peng-Hsuan Li, and Wei-Yun Ma. 2019. GraphRel: Modeling text as relational graphs for joint entity and relation extraction. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- Emil Julius Gumbel. 1948. Statistical theory of extreme values and some practical applications: a series of lectures. In *US Government Printing Office*.
- Rujun Han, Qiang Ning, and Nanyun Peng. 2019. Joint event and temporal relation extraction with shared representations and structured prediction. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*.
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical reparameterization with gumbel-softmax. In *Proceedings of the 5th International Conference on Learning Representations*.
- Alex Judea and Michael Strube. 2016. Incremental global event extraction. In *Proceedings of the 26th International Conference on Computational Linguistics*.
- Arzoo Katiyar and Claire Cardie. 2017. Going out on a limb: Joint extraction of entity mentions and relations without dependency trees. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*.
- Thomas N. Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *Proceedings of the 5th International Conference on Learning Representations*.
- Viet Dac Lai, Tuan Ngo Nguyen, and Thien Huu Nguyen. 2020. Event detection: Gate diversity and syntactic importance scores for graph convolution neural networks. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Qi Li, S. Anzaroot, W. Lin, Xiang Li, and H. Ji. 2011. Joint inference for cross-document information extraction. In *Proceedings of the 20th ACM international conference on Information and knowledge management*.
- Qi Li, Heng Ji, Yu Hong, and Sujian Li. 2014. Constructing information networks using one single model. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*.
- Qi Li, Heng Ji, and Liang Huang. 2013. Joint event extraction via structured prediction with global features. In *Proceedings of the 51th Annual Meeting of the Association for Computational Linguistics*.
- Ying Lin, Heng Ji, Fei Huang, and Lingfei Wu. 2020. A joint neural model for information extraction with global features. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Yi Luan, Dave Wadden, Luheng He, Amy Shah, Mari Ostendorf, and Hannaneh Hajishirzi. 2019. A general framework for information extraction using dynamic span graphs. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Makoto Miwa and Mohit Bansal. 2016. End-to-end relation extraction using LSTMs on sequences and tree structures. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*.
- Makoto Miwa and Yutaka Sasaki. 2014. Modeling joint entity and relation extraction with table representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*.
- Thien Huu Nguyen, Kyunghyun Cho, and Ralph Grishman. 2016. Joint event extraction via recurrent neural networks. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics*.
- Thien Huu Nguyen and Ralph Grishman. 2015a. Event detection and domain adaptation with convolutional neural networks. In *The 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*.
- Thien Huu Nguyen and Ralph Grishman. 2015b. Relation extraction: Perspective from convolutional neural networks. In *Proceedings of the 1st NAACL*

- Workshop on Vector Space Modeling for NLP (VSM)*.
- Thien Huu Nguyen and Ralph Grishman. 2018. Graph convolutional networks with argument-aware pooling for event detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Trung Minh Nguyen and Thien Huu Nguyen. 2019. One for all: Neural joint modeling of entities and events. In *Proceedings of the Association for the Advancement of Artificial Intelligence (AAAI)*.
- Sebastian Riedel, Hong-Woo Chun, Toshihisa Takagi, and Jun'ichi Tsujii. 2009. A markov logic approach to bio-molecular event extraction. In *Proceedings of the BioNLP 2009 Workshop Companion Volume for Shared Task*.
- Dan Roth and Wen-tau Yih. 2004. A linear programming formulation for global inference in natural language tasks. In *Proceedings of the Eighth Conference on Computational Natural Language Learning*.
- Cicero dos Santos and Victor Guimaraes. 2015. Boosting named entity recognition with neural character embeddings. In *Proceedings of the Fifth Named Entity Workshop*.
- Changzhi Sun, Yeyun Gong, Yuanbin Wu, Ming Gong, Daxin Jiang, Man Lan, Shiliang Sun, and Nan Duan. 2019. Joint type inference on entities and relations via graph convolutional networks. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- Deepak Venugopal, Chen Chen, Vibhav Gogate, and Vincent Ng. 2014. Relieving the computational bottleneck: Joint inference for event extraction with high-dimensional features. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*.
- Amir Pouran Ben Veyseh, Franck Deroncourt, Dejing Dou, and Thien Huu Nguyen. 2020a. Exploiting the syntax-model consistency for neural relation extraction. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Amir Pouran Ben Veyseh, Franck Deroncourt, My Thai, Dejing Dou, and Thien Huu Nguyen. 2020b. Multi-view consistency for relation extraction via mutual information and structure prediction. In *Proceedings of the Association for the Advancement of Artificial Intelligence (AAAI)*.
- Amir Pouran Ben Veyseh, Tuan Ngo Nguyen, and Thien Huu Nguyen. 2020c. Graph transformer networks with syntactic and semantic structures for event argument extraction. In *Proceedings of the Findings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP Findings)*.
- David Wadden, Ulme Wennberg, Yi Luan, and Hananeh Hajishirzi. 2019. Entity, relation, and event extraction with contextualized span representations. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- Christopher Walker, Stephanie Strassel, Julie Medero, and Kazuaki Maeda. 2006. Ace 2005 multilingual training corpus. In *Technical report, Linguistic Data Consortium*.
- Bishan Yang and Tom M. Mitchell. 2016. Joint extraction of events and entities within a document context. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Xiaofeng Yu and Wai Lam. 2010. Jointly identifying entities and extracting relations in encyclopedia text via a graphical model approach. In *Proceedings of the 23th International Conference on Computational Linguistics*.
- Junchi Zhang, Yanxia Qin, Yue Zhang, Mengchi Liu, and Donghong Ji. 2019. Extracting entities and events as a single task using a transition-based neural model. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*.
- Suncong Zheng, Feng Wang, Hongyun Bao, Yuexing Hao, Peng Zhou, and Bo Xu. 2017. Joint extraction of entities and relations based on a novel tagging scheme. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*.
- GuoDong Zhou, Jian Su, Jie Zhang, and Min Zhang. 2005. Exploring various knowledge in relation extraction. In *Proceedings of the 43rd annual meeting of the association for computational linguistics*.

A Hyper-parameters

We fine-tune the hyper-parameters for our model FourIE using the development data of the ACE05-E+ dataset (our main and largest evaluation dataset). The selection criteria is based on the average F1 scores of the four IE tasks of consideration (EME, ETD, RE, and EAE). The following values are suggested by the fine-tuning: $2e-5$ for the learning rate of BertAdam for the optimizer; 10 for the batch size; $N_i = 2$ and $N_t = 3$ for the numbers of layers in the GCN models G^{inst} and G^{pred} respectively; 300 hidden units for all the layers of the feed forward networks GCN models, and type embeddings; $\beta = 100$ for the constant in the approximation $\hat{\mathbf{A}}^{pred}$; $\tau = 0.1$ for the temperature parameter; and $\lambda = 0.5$ for the trade-off parameter in the loss function. To achieve a consistency, we apply the same hyper-parameters from this fine-tuning for other datasets.

B Analysis

Analysis. To better understand the contribution of the knowledge from the type dependency graphs \mathbf{G}^{gold} and \mathbf{G}^{pred} for FourIE on EAE, we analyze the set of all the argument instances on the ACE05-E+ development set (called \mathcal{A}) that FourIE can successfully predict the argument roles while “FourIE-GCN^{type}” fails to do so. In particular, we find three major categories of the instances in \mathcal{A} that highlight the benefits of the type dependency graphs:

(i) One-edge constraints (accounting for 28.9% of \mathcal{A}): The incorrect argument role predictions of “FourIE-GCN^{type}” for these instances violate the constraint on the possible argument roles of event types. As FourIE does not have this issue, it suggests that FourIE can learn and enforce those constraints (i.e., from the event_argument edges of \mathbf{G}^{gold}) from the training. For instance, in the sentence “... *the United States upped its military presence, deploying more missile-firing warships to the Red Sea*”, both FourIE and “FourIE-GCN^{type}” can recognize “*deploying*” as an event trigger of type *Transport*. However, regarding the entity mention “*Red Sea*”, FourIE correctly assigns the *Destination* role for the *Transport* event while “FourIE-GCN^{type}” incorrectly considers it as the role *Place* (an invalid role for the event type *Transport*).

(ii) Two-edge constraints (representing 36.5% of \mathcal{A}): The predictions from “FourIE-GCN^{type}” in this category involve argument roles that are never assigned to an entity mention of some entity type

in an event mention/trigger of some event type. FourIE can avoid this issue as it can recognize these constraints from the combinations of two neighboring edges (i.e., an event_argument and entity_argument edge). For example, in the sentence “... *the tanks and Bradley fighting vehicles ... backed by the Apache attack helicopters ... punched through the Republican Guard defenses ...*”, both FourIE and “FourIE-GCN^{type}” can detect “*helicopters*” as an entity mention of type *Vehicle* which is an argument for the “*Attack*” event triggered by “*punched through*”. However, “FourIE-GCN^{type}” incorrectly predicts the argument role of “*Attacker*” for “*helicopters*” while FourIE can successfully return the *Instrument* in this case. In fact, we cannot find any *Vehicle* entity that plays the *Attacker* role in an *Attack* event in the training data, providing a useful information for FourIE to learn and fix the error.

(iii) Four-edge constraints (accounting for 19.2% of \mathcal{A}): The failure of “FourIE-GCN^{type}” for the instances in this category can be fixed if the model exploits the co-occurrence of event types and argument roles in the same sentences. In particular, for two event mentions with related event types in the same sentences, an entity mention that plays some role in one event tends to also play some related role in the other event. These co-occurrence can be captured via two event_argument edges and two entity_argument edges (sharing the same entity type) in the type dependency graphs of FourIE to address the issue. Consider an example sentence: “*Two 13-year-old children were among those killed in the Haifa bus bombing, Israeli public radio said ...*”. Both FourIE and “FourIE-GCN^{type}” can identify the *Person* entity mention “*those*” as the argument of role *Victim* for the *Die* event triggered by “*killed*”. However, regarding the *Attack* event triggered by “*bombing*”, only FourIE can correctly predict “*those*” as an argument of role *Target*. This success can be attributed to the ability of FourIE to learn the co-occurrence that an entity mention has a higher chance to play the role *Target* in an *Attack* event if it also has a role of *Victim* for a *Die* event mentioned in the same sentence.

Finally, the instances in the remaining 15.4% of \mathcal{A} tend to involve more complicated constraints/dependencies that cannot be associated with any of the three categories above.