

Neural Ranking Models for Temporal Dependency Structure Parsing

Yuchen Zhang
Brandeis University
yuchenz@brandeis.edu

Nianwen Xue
Brandeis University
xuen@brandeis.edu

Abstract

We design and build the first neural temporal dependency parser. It utilizes a neural ranking model with minimal feature engineering, and parses time expressions and events in a text into a temporal dependency tree structure. We evaluate our parser on two domains: news reports and narrative stories. In a parsing-only evaluation setup where gold time expressions and events are provided, our parser reaches 0.81 and 0.70 f-score on unlabeled and labeled parsing respectively, a result that is very competitive against alternative approaches. In an end-to-end evaluation setup where time expressions and events are automatically recognized, our parser beats two strong baselines on both data domains. Our experimental results and discussions shed light on the nature of temporal dependency structures in different domains and provide insights that we believe will be valuable to future research in this area.

1 Introduction

Temporal relation classification is important for a range of NLP applications that include but are not limited to story timeline construction, question answering, summarization, etc. Most work on temporal information extraction models the task as a pair-wise classification problem (Bethard et al., 2007; Chambers et al., 2007; Chambers and Jurafsky, 2008; Ning et al., 2018a): given an individual pair of time expressions and/or events, the system predicts whether they are temporally related and which specific relation holds between them. An alternative approach is to model the temporal relations in a text as a temporal dependency structure (TDS) for the entire text (Kolomiyets et al., 2012). Such a temporal dependency structure has the advantage that (1) it can be easily used to infer additional temporal relations between time expressions and/or events that are not directly connected

via the transitivity properties of temporal relations, (2) it is computationally more efficient because a model does not need to consider all pairs of time expressions and events in a text, and (3) it is easier to use for downstream applications such as timeline construction.

However, most existing automatic systems are pair-wise models trained with traditional statistical classifiers using a large number of manually crafted features (Bethard et al., 2017). The few exceptions include the work of Kolomiyets et al. (2012), which describes a temporal dependency parser based on traditional feature-based classifiers, and Dligach et al. (2017), which describes a system using neural network based models to classify individual temporal relations. More recently, a semi-structured approach has also been proposed (Ning et al., 2018b).

In this work, taking advantage of a newly available data set annotated with temporal dependency structures – the Temporal Dependency Tree (TDT) Corpus¹ (Zhang and Xue, 2018), we develop a neural temporal dependency structure parser using minimal hand-crafted linguistic features. One of the advantages of neural network based models is that they are easily adaptable to new domains, and we demonstrate this advantage by evaluating our temporal dependency parser on data from two domains: news reports and narrative stories. Our results show that our model beats a strong logistic regression baseline. Direct comparison with existing models is impossible because the only similar dataset used in previous work that we are aware of is not available to us (Kolomiyets et al., 2012), but we show that our models are competitive against similar systems reported in the literature.

The main contributions of this work are:

- We design and build the first end-to-end neu-

¹https://github.com/yuchenz/structured_temporal_relations_corpus

ral temporal dependency parser. The parser is based on a novel neural ranking model that takes a raw text as input, extracts events and time expressions, and arranges them in a temporal dependency structure.

- We evaluate the parser by performing experiments on data from two domains: news reports and narrative stories, and show that our parser is competitive against similar parsers. We also show the two domains have very different temporal structural patterns, an observation that we believe will be very valuable to future temporal parser development.

The rest of the paper is organized as follows. Since temporal structure parsing is a relatively new task, we give a brief problem description in §2. We describe our end-to-end pipeline system in §3. Details of the neural ranking model are discussed in §4. In §5 we present and discuss our experimental results, and error analysis are presented in §6. In §7 we discuss related work and situate our work in the broader context, and we conclude our paper in §8.

2 Problem Description

In this section we give a brief description of the temporal dependency parsing task (more details in Zhang and Xue (2018)). In a temporal structure parsing task, a text is parsed into a dependency tree structure that represents the inherent temporal relations among time expressions and events in the text. The nodes in this tree are mostly time expressions and events which are represented as contiguous spans of words in the text. They can also be pre-defined meta nodes, which serve as reference times for other time expressions and events, and they constitute the top-most part of the tree. For example, *Past_Ref*, *Present_Ref*, *Future_Ref*, and *Document Creation Time (DCT)* are all pre-defined meta nodes. The edges in the tree represent temporal relations between each parent-child pair. The temporal relations can be one of *Includes*, *Before*, *Overlap*, and *After*, or *Depend-on* which holds between two time expressions. Unlike syntactic dependency parsing where each word in a sentence is a node in the dependency structure, in a temporal dependency structure only some of the words in a text are nodes in the structure. Therefore, this process naturally falls into two stages: first time expression and event recognition, and then temporal relation parsing. Fig-

ure 1 is an example temporal dependency tree for a news report paragraph.

Due to the fact that different types of time expressions and events behave differently in terms of what can be their antecedents, and recognition of these types can be helpful for determining temporal relations, finer classifications of time expressions and events are also defined. Time expressions are further classified into *Vague Time*, *Absolute Concrete Time*, and *Relative Concrete Time*, according to whether or not the time expression can be located on the timeline, and whether or not the interpretation of its temporal location depends on another time expression. Events are further classified into *Eventive Event*, *State*, *Habitual Event*, *Completed Event*, *Ongoing Event*, *Modalized Event*, *Generic Habitual*, and *Generic State*, according to the eventuality type of the event. Our experiments will show that these fine-grained classifications are very helpful for the overall temporal structure parsing accuracy.

3 A Pipeline System

We build a two-stage pipeline system to tackle this temporal structure parsing problem. The first stage performs event and time expression identification. In this stage, given a text as input, spans of words that indicate events or time expressions are identified and categorized. We model this stage as a sequence labeling process. A standard Bi-LSTM sequence model coupled with BIO labels is applied here. Word representations are the concatenation of word and POS tag embeddings.

The second stage performs the actual temporal structure parsing by identifying the antecedent for each time expression and event, and identifying the temporal relation between them. In this stage, given events and time expressions identified in the first stage as input, the model outputs a temporal dependency tree in which each child node is an event or time expression that is temporally related to another event or time expression or pre-defined meta node as its parent node. This stage is modeled as a ranking process: for each node, a finite set of neighboring nodes are first selected as its candidate parents. These candidates are then ranked with a neural network model and the highest ranking candidate is selected as its parent. We use a ranking model because it is simple, more intuitive and easier to train than a traditional transition-based or graph-based model, and

minimized:

$$L = -\log P(y_i | D_k) \\ = -\log \frac{\exp[s_{i,y_i}]}{\sum_{y'_i} \exp[s_{i,y'_i}]}$$

where s_{i,y'_i} is the score for child-candidate pair (i, y'_i) as described in §4.1.

4.3 Decoding

During decoding, the parser constructs the temporal dependency tree incrementally by identifying the parent node for each event or time expression in textual order. To ensure the output parse is a valid dependency tree, two constraints are applied in the decoding process: (i) there can only be one parent for each node, and (ii) descendants of a node cannot be its parent to avoid cycles. Candidates violating these constraints are omitted from the ranking process.²

4.4 Temporal Relation Labeling

The neural model described above generates an unlabeled temporal dependency tree, with each parent being the most salient reference time for the child. However it doesn't model the specific temporal relation (e.g. "before", "overlap") between a parent and a child. We extend this basic architecture to both identify parent-child pairs and predict their temporal relations. In this new model, instead of ranking child-candidate pairs (i, y'_i) , we rank child-candidate-relation tuples (i, y'_i, l_k) , where l_k is the k th relation in the pre-defined set of possible temporal relation labels L . We compute this ranking by re-defining the pair score s_{i,y'_i} . Here, pair score s_{i,y'_i} is no longer a scalar score but a vector \mathbf{s}_{i,y'_i} of size $|L|$, where $\mathbf{s}_{i,y'_i}[k]$ is the scalar score for y'_i being the parent of i with temporal relation l_k . Accordingly, the lengths of \mathbf{c}_i and \mathbf{o}_i are *number of candidates* $\times |L|$. Finally, the tuple (i, y'_i, l_k) associated with the highest score in \mathbf{o}_i predicts that y'_i is the parent for i with temporal relation label l_k .

4.5 Variations of the Basic Neural Model

4.5.1 Linguistically Enriched Models

A variation of the basic neural model is a model that takes a few linguistic features as input ex-

²An alternative decoding approach would be to perform a global search for a Maximum Spanning Tree. However, due to the nature of temporal structures, our greedy decoding process rarely hits the constraints.

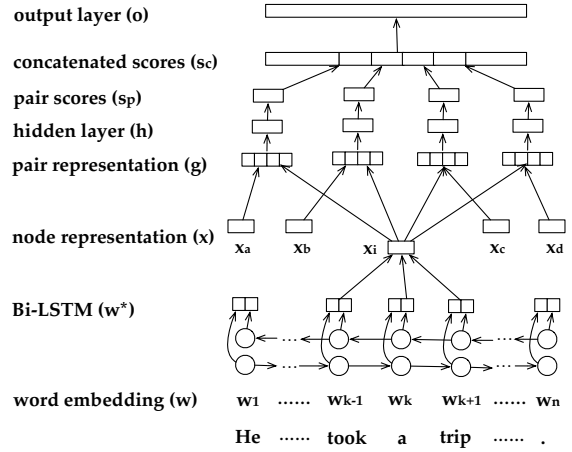


Figure 2: Neural Ranking Model Architecture. x_i is the current child node, and x_a, x_b, x_c, x_d are the candidate parent nodes for x_i . Arrows from Bi-LSTM layer to x_a, x_b, x_c, x_d are not shown.

PLICITLY. In this model, we extend the pair representation \mathbf{g}_{i,y'_i} with local features: $\mathbf{g}_{i,y'_i} = [\mathbf{x}_i, \mathbf{x}_{y'_i}, \phi_{i,y'_i}]$.

Time and event type feature: Stage 1 of the pipeline not only extracts text spans that are time expressions or events, but also labels them with pre-defined categories of different types of time expressions and events. Readers are referred to Zhang and Xue (2018) for the full category list. Through a careful examination of the data, we notice that time expressions or events are selective as to what types of time expression or events can be their parent. In other words, the category of the child time expression or event has a strong indication on which candidate can be its parent. For example, a time expression's parent can only be another time expression or a pre-defined meta node, and can never be an event; and an eventive event's parent is almost certainly another eventive event, and is highly unlikely to be a stative event. Therefore, we include the time expression and event type information predicted by stage 1 in this model as a feature. More formally, we represent a time/event type as a fixed-length embedding \mathbf{t} , and concatenate it to the pair representation $\mathbf{g}_{i,y'_i} = [\mathbf{x}_i, \mathbf{x}_{y'_i}, \mathbf{t}_i, \mathbf{t}_{y'_i}]$.

Distance features: Distance information can be useful for predicting the parent of a child. Intuitively, candidates that are closer to the child are more likely to be the actual parent. Through data examination, we also find that a high percentage of nodes have parents in close proximity. Therefore,

we include two distance features in this model: the node distance between a candidate and the child nd_{i,y'_i} , and whether they are in the same sentence ss_{i,y'_i} . One-hot representations are used for both features to represent according conditions listed in Table 1.

conditions for feature nd_{i,y'_i} :
$i.node_id - y'_i.node_id = 1$
$i.node_id - y'_i.node_id > 1$ and $i.sent_id = y'_i.sent_id$
$i.node_id - y'_i.node_id > 1$ and $i.sent_id \neq y'_i.sent_id$
$i.node_id - y'_i.node_id < 1$
conditions for feature ss_{i,y'_i} :
$i.sent_id = y'_i.sent_id$
$i.sent_id \neq y'_i.sent_id$

Table 1: Conditions for node distance and same sentence features.

The final pair representation for our linguistically enriched model is as follows:

$$g_{i,y'_i} = [\mathbf{x}_i, \mathbf{x}_{y'_i}, \mathbf{t}_i, \mathbf{t}_{y'_i}, nd_{i,y'_i}, ss_{i,y'_i}]$$

4.5.2 Attention Model on Time and Event Representation

In the basic neural model, a straight-forward sum-pooling is used as the multi-word time expression and event representation. However, multi-word event expressions usually have meaning-bearing head words. For example, in the event “took a trip”, “trip” is more representative than “took” and “a”. Therefore, we add an attention mechanism (Bahdanau et al., 2014) over the Bi-LSTM output vectors in each multi-word expression to learn a task-specific notion of headedness (Lee et al., 2017):

$$\alpha_t = \tanh(\mathbf{W} \cdot \mathbf{w}_t^*)$$

$$w_{i,t} = \frac{\exp[\alpha_t]}{\sum_{k=START(i)}^{END(i)} \exp[\alpha_k]}$$

$$\hat{\mathbf{x}}_i = \sum_{t=START(i)}^{END(i)} w_{i,t} \cdot \mathbf{w}_t^*$$

where $\hat{\mathbf{x}}_i$ is a weighted sum of Bi-LSTM output vectors in span i . The weights $w_{i,t}$ are automatically learned. The final pair representation for our attention model is as follows:

$$g_{i,y'_i} = [\mathbf{x}_i, \mathbf{x}_{y'_i}, \mathbf{t}_i, \mathbf{t}_{y'_i}, nd_{i,y'_i}, ss_{i,y'_i}, \hat{\mathbf{x}}_i, \hat{\mathbf{x}}_{y'_i}]$$

This model variation is also beneficial in an end-to-end system, where time expression and event spans are automatically extracted in Stage 1. When extracted spans are not guaranteed correct time expressions and events, an attention layer on

a slightly larger context of an extracted span has a better chance of finding representative head words than a sum-pooling layer strictly on words within a event or time expression span.

5 Experiments

5.1 Data

All of our experiments are conducted on the datasets described in Zhang and Xue (2018). This is a temporal dependency structure corpus in Chinese. It covers two domains: news reports and narrative fairy tales. It consists of 115 news articles sampled from Chinese TempEval2 datasets (Verhagen et al., 2010) and Chinese Wikipedia News³, and 120 fairy tale stories sampled from Grimm Fairy Tales⁴. 20% of this corpus, distributed evenly on both domains, are double annotated with high inter-annotator agreements. We use this part of the data as our development and test datasets (10% documents for development and 10% for testing), and the remaining 80% as our training dataset.

5.2 Baseline Systems

We build two baseline systems to compare with our neural model. The first is a simple baseline which links every time expression or event to its immediate previous time expression or event. According to our data, if only position information is considered, the most likely parent for a child is its immediate previous time expression or event. This baseline uses the most common temporal relation edge label in the training datasets, i.e. “overlap” for news data, and “before” for grimm data.

The second baseline is a more competitive baseline for stage 2 in the pipeline. It takes the output of the first stage as input, and uses a similar ranking architecture but with logistic regression classifiers instead of neural classifiers. The purpose of this baseline is to compare our neural models against a traditional statistical model under otherwise similar settings. We conduct robust feature engineering on this logistic regression model to make sure it is a strong benchmark to compete against. Table 2 lists the features and feature combinations used in this model.

³<https://zh.wikinews.org>

⁴https://www.grimmstories.com/zh/grimm_tonghua/index

time type and event type features:
$i.type$ and $y'_i.type$
if $i.type = absolute\ time$ and $y'_i.type = root$
if $i.type = time$ and $y'_i.type = root$
are $i.type$ and $y'_i.type$ time, eventive, or stative
are $i.type$ and $y'_i.type$ root, time, or event
are $i.type$ and $y'_i.type$ root, time, eventive, or stative
if $i.type = y'_i.type = event$ and $\hat{y}.type = state$,
for all \hat{y} between i and y'_i
distance features:
if $i.sent_id = y'_i.sent_id$
$i.node_id - y'_i.node_id$
if $i.node_id - y'_i.node_id = 1$
combination features:
if $i.type = state$ and $i.sent_id \neq y'_i.sent_id$
if $i.type = state$ and $i.node_id - y'_i.node_id = 1$
if $i.type = y'_i.type = event$ and
$i.node_id - y'_i.node_id = 1$
if $i.type = state$ and $y'_i.type = event$ and
$i.node_id - y'_i.node_id = 1$ and
$i.node_id.in_sent = 1$ and
$i.sent_id \neq 1$
other features:
if i and y'_i are in quotation marks

Table 2: Features in the logistic regression system.

5.3 Evaluation

We perform two types of evaluations for our systems. First, we evaluate the stages of the pipeline and the entire pipeline, i.e. end-to-end systems where both time expression and event recognition, as well as temporal dependency structures are automatically predicted. Our models are compared against the two strong baselines described in §5.2. These evaluations are described in §5.3.1.

The second evaluation focuses only on the temporal relation structure parsing part of our pipeline (i.e. Stage 2), using gold standard time expression and event spans and labels. Since most previous work on temporal relation identification use gold standard time expression and event spans, this evaluation gives us some sense of how our models perform against models reported in previous work even though a strict comparison is impossible because different data sets are used. These evaluations are described in §5.3.2.

All neural networks in this work are implemented in Python with the DyNet library (Neubig et al., 2017). The code is publicly available⁵.

For Stage 1, all models are trained with Adam optimizer with early stopping and learning rate 0.001. The dimensions of word embeddings, POS tag embeddings, Bi-LSTM output vectors, and MLP hidden layers are tuned on the dev set to 256,

⁵https://github.com/yuchenz/tdp_ranking

evaluated label	news			grimm		
	p	r	f	p	r	f
all span	.81	.74	.78	.83	.74	.78
time	.83	.81	.82	.97	.62	.76
event	.81	.73	.77	.83	.74	.78

Table 3: Stage 1 cross-validation on span detection and binary time/event recognition.

time/event type	news	grimm
vague time	.77	.82
concrete absolute	.67	-
concrete relative	.75	-
event	.61	.77
state	.65	.61
completed	.62	.26
modalized	.46	.31

Table 4: Stage 1 (time/event type recognition) cross-validation f1-scores on the full set.

32, 256, and 256 respectively. POS tags in Stage 1 are acquired using the joint POS tagger from Wang and Xue (2014). The tagger is trained on Chinese Treebank 7.0 (Xue et al., 2010). For Stage 2, the dimensions of word embeddings, time/event type embeddings, Bi-LSTM output vectors, and MLP hidden layers are tuned on the dev set to 32, 16, 32, and 32 respectively. The optimizer is Adam with early stopping and learning rate 0.001.

5.3.1 End-to-End System Evaluation

Stage 1: Time and Event Recognition For Stage 1 in the pipeline, we perform BIO tagging with the full set of time expression and event types (i.e. a 11-way classification on all extracted spans). Extracted spans will be nodes in the final dependency tree, and time/event types will support features in the next stage. We evaluate Stage 1 performance using 10-fold cross-validation of the entire data set. We use the “exact match” evaluation metrics for BIO sequence labeling tasks, and compute precision, recall, and f-score for each label type.

We first ignore fine-grained time/event types and only evaluate unlabeled span detection and time/event binary classification to show how well our system identify events and time expressions, and how well our system distinguishes time expressions from events. Table 3 shows the cross-validation results on these two evaluations. Span detection and event recognition show similar per-

	model	news				grimm			
		unlabeled f		labeled f		unlabeled f		labeled f	
		dev	test	dev	test	dev	test	dev	test
temporal relation parsing with gold spans	Baseline-simple	.64	.68	.47	.43	.78	.79	.39	.39
	Baseline-logistic	.81	.79	.63	.54	.74	.74	.60	.63
	Neural-basic	.78	.75	.67	.57	.72	.74	.60	.63
	Neural-enriched	.80	.78	.67	.59	.76	.77	.63	.65
	Neural-attention	.83	.81	.76	.70	.79	.79	.66	.68
end-to-end systems with automatic spans	Baseline-simple	.39	.40	.26	.25	.44	.47	.27	.25
	Baseline-logistic	.36	.34	.24	.22	.43	.49	.33	.37
	Neural-basic	.37	.36	.21	.23	.42	.45	.33	.35
	Neural-enriched	.51	.52	.32	.35	.44	.49	.33	.37
	Neural-attention	.54	.54	.36	.39	.44	.49	.35	.39

Table 5: Stage 2 results (f-scores) with gold spans and timex/event labels (top), and automatic spans and timex/event labels generated by stage 1 (bottom). Best performances are in bold.

formance on both news and narrative domains. Time expressions have a higher recognition rate than events in news data, which is consistent with the observation that time expressions usually have a more limited vocabulary and more strict lexical patterns. On the other hand, due to the scarcity of time expressions in the Grimm data, time expression recognition in this domain has a very high precision but low recall, which results in a much lower f-score than news.

Labeled full set evaluation results on time/event type classification are reported in Table 4. Time expressions have higher recognition rates than events on both domains, and dominant event types (“event”, “state”, etc.) have higher and more stable recognition rates than other types. Event types with very few training instances, such as “modalized event” (<7%), achieve lower and more unstable recognition rates. Other types with less than 2% instances achieve close to 0 recognition f-scores, and are not reported in this table.

Stage 2: Temporal Dependency Parsing For Stage 2 in the pipeline, we conduct experiments on the five systems described above: a simple baseline, a logistic regression baseline, a basic neural model, a linguistically enriched neural model, and an attention neural model. All models are trained on automatically predicted spans of time expressions and events, and time/event types generated by Stage 1 using 10-fold cross-validation, with gold standard edges (and edge labels) mapped onto the automatic spans. Evaluations in Stage 2 are against gold standard spans and edges, and evaluation metrics are precision, recall, and

f-score on $\langle child, parent \rangle$ tuples for unlabeled trees, and $\langle child, relation, parent \rangle$ triples for labeled trees.

Bottom rows in Table 5 report the end-to-end performance of our five systems on both domains. On both labeled and unlabeled parsing, our basic neural model with only lexical input performs comparable to the logistic regression model. And our enriched neural model with only three simple linguistic features outperforms both the logistic regression model and the basic neural model on news, improving the performance by more than 10%. However, our models only slightly improve the unlabeled parsing over the simple baseline on narrative Grimm data. This is probably due to (1) it is a very strong baseline to link every node to its immediate previous node, since in a narrative discourse linear temporal sequences are very common; and (2) most events breaking the temporal linearity in a narrative discourse are implicit stative descriptions which are harder to model with only lexical and distance features. Finally, attention mechanism improves temporal relation labeling on both domains.

5.3.2 Temporal Relation Evaluation

To facilitate comparison with previous work where gold events are used as parser input, we report our results on temporal dependency parsing with gold time expression and event spans in Table 5 (top rows). These results are in the same ballpark as what is reported in previous work on temporal relation extraction. The best performance in Kolomiyets et al. (2012) are 0.84 and 0.65 f-scores for unlabeled and labeled parses, achieved

by temporal structure parsers trained and evaluated on narrative children’s stories. Our best performing model (Neural-attention) reports 0.81 and 0.70 f-scores on unlabeled and labeled parses respectively, showing similar performance. It is important to note, however, that these two works use different data sets, and are not directly comparable. Finally, parsing accuracy with gold time/event spans as input is substantially higher than that with predicted spans, showing the effects of error propagation.

6 Error Analysis

We perform error analysis on the output of our best model (Neural-attention) on the development data sets. We focus on analyzing our neural ranking model (i.e. Stage 2), with gold time expression and event spans and labels as input.

First, we look at errors by the types of antecedents. Most events in both news and grimm data depend on their immediate previous event or time expression as their reference time parent. 71% of the events in the news data set and 78% of the events in the Grimm data have the immediate previous node as their antecedent. The confusion matrix in Table 6 illustrates how strongly this bias affects our models. Our model learns the bias and incorrectly links around half of the events (47% in news and 46% in grimm) to their immediate previous node when the *correct* temporal dependency is further back in the text.

	news			grimm		
	pre	far	total	pre	far	total
pre	317	11	328	750	60	810
far	65	72	137	104	122	226
total	382	83	465	854	182	1036

Table 6: Parent node confusion matrix. Rows are gold parents and columns are automatically parsed parents. “pre” means the parent is the immediate previous node of the child event, “far” means the parent is further back from the child event.

Second, we look at errors in temporal relation labels. Considering only correctly recognized parent-child pairs, we draw a confusion matrix as in Table 7. Our data has very few *after* relations in both domains, which explains why the model has difficulty identifying this relation. There are also very few *include* and *depend-on* relations in the Grimm data, however they are identified with a

news	be	af	ov	in	de	total
before	1	0	21	2	0	24
after	1	0	1	0	0	2
overlap	1	0	295	0	0	296
include	0	0	4	52	0	56
depend-on	0	0	0	0	117	117
total	3	0	321	54	117	495
grimm	be	af	ov	in	de	total
before	367	0	55	0	0	422
after	1	0	2	0	0	3
overlap	74	1	314	0	0	389
include	3	0	0	10	0	13
depend-on	0	0	0	0	12	12
total	445	1	371	10	12	839

Table 7: Temporal relation confusion matrix. Rows are gold relation labels and columns are automatic relation labels. “be, af, ov, in, de” stand for “before, after, overlap, include, and depend-on”.

relatively high accuracy. This is probably because, according to the temporal dependency structure design (Zhang and Xue, 2018), these relations hold only between restricted pairs of parent and child: *include* requires a time expression parent and an event child, and *depend-on* requires that the parent be the root. The main confusion among temporal relations is between *before* and *overlap*. In news data, with a high occurrence of *overlap* relations (60% *overlap* and 5% *before*), most *before* parents are wrongly recognized as *overlap*. Grimm data has a more balanced distribution of these two temporal relations (46% *overlap* and 50% *before*), however, 13% *before* and 17% *overlap* are wrongly labeled as the other.

7 Related Work

7.1 Related Work on Temporal Relation Modeling

There is a significant amount of research on temporal relation extraction (Bethard et al., 2007; Bethard, 2013; Chambers and Jurafsky, 2008; Chambers et al., 2014; Ning et al., 2018a). Most of the previous work models temporal relation extraction as pair-wise classification between individual pairs of events and/or time expressions. Some of the models also add a global reasoning step to local pair-wise classification, typically using Integer Linear Programming, to exploit the transitivity property of temporal relations (Chambers and Jurafsky, 2008). Such a pair-wise clas-

sification approach is often dictated by the way the data is annotated. In most of the widely used temporal data sets, temporal relations between individual pairs of events and/or time expressions are annotated independently of one another (Pustejovsky et al., 2003; Chambers et al., 2014; Styler IV et al., 2014; O’Gorman et al., 2016; Mostafazadeh et al., 2016).

Our work is most closely related to that of Kolomiyets et al. (2012), which also treats temporal relation modeling as temporal dependency structure parsing. However, their dependency structure, as described in Bethard et al. (2012), is only over events, excluding time expressions which are an important source of temporal information, and it also excludes *states* (stative events), which makes the temporal dependency structure incomplete. Moreover, their corpus only consists of data in the narrative stories domain. We instead choose to develop our model based on the data set described in Zhang and Xue (2018), which introduces a more comprehensive and linguistically grounded annotation scheme for temporal dependency structures. This structure includes both events and time expressions, and uses the linguistic notion of *temporal anaphora* to guide the annotation of the temporal dependency structure. Since in this temporal dependency structure each parent-child pair is considered to be an instance of temporal anaphora, the parent is also called the *antecedent* and the child is also referred to as the *anaphor*. Their corpus consists of data from two domains: news reports and narrative stories.

More recently, Ning et al. (2018b) proposed a semi-structured approach to model temporal relations in a text. Based on the observation that not all pairs of events have well-defined temporal relations, they propose a multi-axis representation in which well-defined temporal relations only hold between events on the same axis. The temporal relations between events in a text form multiple disconnected subgraphs. Like other work before them, their annotation scheme only covers events, to the exclusion of time expressions.

7.2 Related Work on Neural Dependency Parsing

Most prior work on neural dependency parsing is aimed at syntactic dependency parsing, i.e. parsing a sentence into a dependency tree that represents the syntactic relations among the words.

Recent work on dependency parsing typically uses transition-based or graph-based architectures combined with contextual vector representations learned with recurrent neural networks (e.g. Bi-LSTMs) (Kiperwasser and Goldberg, 2016).

Temporal dependency parsing is, however, different from syntactic dependency parsing. In temporal dependency parsing, for each event or time expression, there is more than one other event or time expression that can serve as its reference time, while the most closely related one is selected as the gold standard reference time parent. This naturally falls into a ranking process where all possible reference times are ranked and the best is selected. In this sense our neural ranking model for temporal dependency parsing is closely related to the neural ranking model for coreference resolution described in Lee et al. (2017), both of which extract related spans of words (entity mentions for coreference resolution, and events or time expressions for temporal dependency parsing). However, our temporal dependency parsing model differs from Lee et al.’s coreference model in that the ranking model for coreference only needs to output the best candidate for each individual pairing and cluster all pairs that are coreferent to each other. In contrast, our ranking model for temporal dependency parsing needs to rank not only the candidate antecedents but also the temporal relations between the antecedent and the anaphor. In addition, the model also adds connectivity and acyclic constraints in the decoding process to guarantee a tree-structured output.

8 Conclusion and Future Work

In this paper, we present the first end-to-end neural temporal dependency parser. We evaluate the parser with both gold standard and automatically recognized time expressions and events. In both experimental settings, the parser outperforms two strong baselines and shows competitive results against prior temporal systems.

Our experimental results show that the model performance drops significantly when automatically predicted event and time expressions are used as input instead of gold standard ones, indicating an error propagation problem. Therefore, in future work we plan to develop joint models that simultaneously extract events and time expressions, and parse their temporal dependency structure.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Steven Bethard. 2013. Cleartk-timeml: A minimalist approach to tempeval 2013. In *Second Joint Conference on Lexical and Computational Semantics (*SEM)*, volume 2, pages 10–14.
- Steven Bethard, Oleksandr Kolomiyets, and Marie-Francine Moens. 2012. Annotating story timelines as temporal dependency structures. In *Proceedings of the eighth international conference on language resources and evaluation (LREC)*, pages 2721–2726. ELRA.
- Steven Bethard, James H Martin, and Sara Klingsstein. 2007. Timelines from text: Identification of syntactic temporal relations. In *Semantic Computing, 2007. ICSC 2007. International Conference on*, pages 11–18. IEEE.
- Steven Bethard, Guergana Savova, Martha Palmer, and James Pustejovsky. 2017. Semeval-2017 task 12: Clinical tempeval. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 565–572, Vancouver, Canada. Association for Computational Linguistics.
- Nathanael Chambers, Taylor Cassidy, Bill McDowell, and Steven Bethard. 2014. Dense event ordering with a multi-pass architecture. *Transactions of the Association for Computational Linguistics*, 2:273–284.
- Nathaniel Chambers and Daniel Jurafsky. 2008. Jointly combining implicit constraints improves temporal ordering. In *EMNLP-2008*.
- Nathaniel Chambers, Shan Wang, and Daniel Jurafsky. 2007. Classifying temporal relations between events. In *ACL-2007*.
- Dmitriy Dligach, Timothy Miller, Chen Lin, Steven Bethard, and Guergana Savova. 2017. Neural temporal relation extraction. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, volume 2, pages 746–751.
- Eliyahu Kiperwasser and Yoav Goldberg. 2016. Simple and accurate dependency parsing using bidirectional lstm feature representations. *arXiv preprint arXiv:1603.04351*.
- Oleksandr Kolomiyets, Steven Bethard, and Marie-Francine Moens. 2012. Extracting narrative timelines as temporal dependency structures. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*, pages 88–97. Association for Computational Linguistics.
- Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. 2017. End-to-end neural coreference resolution. *arXiv preprint arXiv:1707.07045*.
- Nasrin Mostafazadeh, Alyson Grealish, Nathanael Chambers, James Allen, and Lucy Vanderwende. 2016. Caters: Causal and temporal relation scheme for semantic annotation of event structures. In *Proceedings of the The 4th Workshop on EVENTS: Definition, Detection, Coreference, and Representation, San Diego, California, June. Association for Computational Linguistics*.
- Graham Neubig, Chris Dyer, Yoav Goldberg, Austin Matthews, Waleed Ammar, Antonios Anastasopoulos, Miguel Ballesteros, David Chiang, Daniel Clothiaux, Trevor Cohn, Kevin Duh, Manaal Faruqi, Cynthia Gan, Dan Garrette, Yangfeng Ji, Lingpeng Kong, Adhiguna Kuncoro, Gaurav Kumar, Chaitanya Malaviya, Paul Michel, Yusuke Oda, Matthew Richardson, Naomi Saphra, Swabha Swayamdipta, and Pengcheng Yin. 2017. Dynet: The dynamic neural network toolkit. *arXiv preprint arXiv:1701.03980*.
- Qiang Ning, Hao Wu, Haoruo Peng, and Dan Roth. 2018a. Improving temporal relation extraction with a globally acquired statistical resource. *arXiv preprint arXiv:1804.06020*.
- Qiang Ning, Hao Wu, and Dan Roth. 2018b. A multi-axis annotation scheme for event temporal relations. *arXiv preprint arXiv:1804.07828*.
- Tim O’Gorman, Kristin Wright-Bettner, and Martha Palmer. 2016. Richer event description: Integrating event coreference with temporal, causal and bridging annotation. *Computing News Storylines*, page 47.
- James Pustejovsky, Patrick Hanks, Roser Sauri, Andrew See, Robert Gaizauskas, Andrea Setzer, Dragomir Radev, Beth Sundheim, David Day, Lisa Ferro, et al. 2003. The timebank corpus. In *Corpus linguistics*, volume 2003, page 40. Lancaster, UK.
- William F Styler IV, Steven Bethard, Sean Finan, Martha Palmer, Sameer Pradhan, Piet C de Groen, Brad Erickson, Timothy Miller, Chen Lin, Guergana Savova, et al. 2014. Temporal annotation in the clinical domain. *Transactions of the Association for Computational Linguistics*, 2:143–154.
- Marc Verhagen, Roser Sauri, Tommaso Caselli, and James Pustejovsky. 2010. Semeval-2010 task 13: Tempeval-2. In *Proceedings of the 5th international workshop on semantic evaluation*, pages 57–62. Association for Computational Linguistics.
- Zhiguo Wang and Nianwen Xue. 2014. Joint pos tagging and transition-based constituent parsing in chinese with non-local features. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 733–742.

Nianwen Xue, Zixin Jiang, Xiuhong Zhong, Martha Palmer, Fei Xia, Fu-Dong Chiou, and Meiyu Chang. 2010. Chinese treebank 7.0. *Linguistic Data Consortium, Philadelphia*.

Yuchen Zhang and Nianwen Xue. 2018. Structured interpretation of temporal relations. In *Proceedings of 11th edition of the Language Resources and Evaluation Conference (LREC-2018)*, Miyazaki, Japan.