

# DisSent: Learning Sentence Representations from Explicit Discourse Relations

Allen Nie<sup>\*1</sup> Erin D. Bennett<sup>\*2</sup> Noah D. Goodman<sup>1,2</sup>  
<sup>1</sup>Department of Computer Science <sup>2</sup>Department of Psychology  
Stanford University

anie@cs.stanford.edu {erindb,ngoodman}@stanford.edu

## Abstract

Learning effective representations of sentences is one of the core missions of natural language understanding. Existing models either train on a vast amount of text, or require costly, manually curated sentence relation datasets. We show that with dependency parsing and rule-based rubrics, we can curate a high quality sentence relation task by leveraging explicit discourse relations. We show that our curated dataset provides an excellent signal for learning vector representations of sentence meaning, representing relations that can only be determined when the meanings of two sentences are combined. We demonstrate that the automatically curated corpus allows a bidirectional LSTM sentence encoder to yield high quality sentence embeddings and can serve as a supervised fine-tuning dataset for larger models such as BERT. Our fixed sentence embeddings achieve high performance on a variety of transfer tasks, including SentEval, and we achieve state-of-the-art results on Penn Discourse Treebank’s implicit relation prediction task.

## 1 Introduction

Developing general models to represent the meaning of a sentence is a key task in natural language understanding. The applications of such general-purpose representations of sentence meaning are many — paraphrase detection, summarization, knowledge-base population, question-answering, automatic message forwarding, and metaphoric language, to name a few.

We propose to leverage a high-level relationship between sentences that is both frequently and systematically marked in natural language: the *discourse relations* between sentences. Human writers naturally use a small set of very common transition words between sentences (or sentence-like

phrases) to identify the relations between adjacent ideas. These words, such as *because*, *but*, and *and*, which mark the conceptual relationship between two sentences, have been widely studied in linguistics, both formally and computationally, and have many different names. We use the name “discourse markers”.

Learning flexible meaning representations requires a sufficiently demanding, yet tractable, training task. Discourse markers annotate deep conceptual relations between sentences. Learning to predict them may thus represent a strong training task for sentence meanings. This task is an interesting intermediary between two recent trends. On the one hand, models like InferSent (Conneau et al., 2017) are trained to predict entailment—a strong conceptual relation that must be hand annotated. On the other hand, models like BERT (Devlin et al., 2018) are trained to predict random missing words in very large corpora (see Table 1 for the data requirements of the models we compare). Discourse marker prediction may permit learning from relatively little data, like entailment, but can rely on naturally occurring data rather than hand annotation, like word-prediction.

We thus propose the DisSent task, which uses the Discourse Prediction Task to train sentence embeddings. Using a data preprocessing procedure based on dependency parsing, we are able to automatically curate a sizable training set of sentence pairs. We then train a sentence encoding model to learn embeddings for each sentence in a pair such that a classifier can identify, based on the embeddings, which discourse marker was used to link the sentences. We also use the DisSent task to fine-tune larger pre-trained models such as BERT.

We evaluate our sentence embedding model’s performance on the standard fixed embedding evaluation framework developed by Conneau et al. (2017), where during evaluation, the sentence em-

\* equal contribution

bedding model’s weights are not updated. We further evaluate both the DisSent model and a BERT model fine-tuned on DisSent on two classification tasks from the Penn Discourse Treebank (PDTB) (Rashmi et al., 2008).

We demonstrate that the resulting DisSent embeddings achieve comparable results to InferSent on some evaluation tasks, and superior on others. The BERT model fine-tuned on the DisSent tasks achieved state-of-the-art on PDTB classification tasks compared to other fine-tuning strategies.

## 2 Discourse Prediction Task

Hobbs (1985) argues that discourse relations are always present, that they fall under a small set of categories, and that they compose into parsable structures. We draw inspiration from Rhetorical Structure Theory (RST) (Mann and Thompson, 1988), which deals with the general task of segmenting natural text into elementary discourse units (EDUs) (Carlson and Marcu, 2001) and parsing into complex discourse structures (e.g. Lin et al. 2019). However, for our task, we narrow our scope to a small subset of especially straightforward discourse relations. First, we restrict our interest to only a subset of EDUs (sentence-like text fragments) that can be interpreted as grammatically complete sentences in isolation. This includes EDUs that appear as full sentences in the original text, as well as subordinate clauses with overt subjects and finite verb phrases. Second, we focus here on explicit discourse markers between adjacent sentences (or EDUs), rather than implicit relations between a sentence (or EDU) and the related discourse. This is a significant simplification from related work in discourse theory, e.g. describing the wealth of complex structures a discourse can take (Webber et al., 2003) or compiling a comprehensive set of discourse relations (Rashmi et al., 2008; Hobbs, 1979, 1985; Jasin-skaja and Karagjosova, 2015; Knott, 1996). We are able to make this simplification because our goal is not to annotate natural text, but to curate a set of sentence pairs for a particular set of discourse relations.

With this focus in mind, we propose a new task for natural language understanding: discourse marker prediction. Given two sentences in our curated corpus (which may have been full sentences in the original text or may have been subclauses), the model must predict which discourse marker

Task	# of examples	# of tokens
SNLI + MNLI	0.9M	16.3M
DisSent Books 5	3.2M	63.5M
SkipThought	—	800M
BERT MLM/NSP	—	3300M

Table 1: Training data size (in millions) in each pre-training task. DisSent Books 5 only uses 5 discourse markers instead of all.

Marker	Extracted Pairs	Percent (%)
but	1,028,995	21.86
and	1,020,316	21.68
as	748,886	15.91
when	527,031	11.20
if	472,852	10.05
before	218,305	4.64
because	167,358	3.56
while	161,818	3.44
though	104,218	2.21
after	95,847	2.04
so	76,940	1.63
although	37,511	0.80
then	16,429	0.35
also	16,365	0.35
still	13,421	0.29
Total	4,706,292	100.0

Table 2: Number of pairs of sentences extracted from BookCorpus for each discourse marker and percent of each marker in the resulting dataset.

was used by the author to link the two ideas. For example, “She’s late to class \_\_\_\_\_ she missed the bus” would likely be completed with *because*, but “She’s sick at home \_\_\_\_\_ she missed the class” would likely be completed with *so*, and “She’s good at soccer \_\_\_\_\_ she missed the goal” would likely be completed with *but*. These pairs have similar syntactic structures and many words in common, but the meanings of the component sentences lead to strong intuitions about which discourse marker makes the most sense. Without a semantic understanding of the sentences, we would not be able to guess the correct relation. We argue that success at choosing the correct discourse marker requires a representation that reflects the full meaning of a sentence.

We note that perfect performance at this task is impossible for humans (Malmi et al., 2018), because different discourse markers can easily appear in the same context. For example, in some cases, markers are (at least close to) synonymous with one another (Knott, 1996). Other times, it is possible for multiple discourse markers to link the same pair of sentences and change the inter-

pretation. (In the sentence “Bob saw Alice was at the party, (then|so|but) he went home,” changing the discourse marker drastically changes our interpretation of Bob’s goals and feeling towards Alice.) Despite this ceiling on absolute performance, a discourse marker can frequently be inferred from the meanings of the sentences it connects, making this a useful training task.

### 3 Model

#### 3.1 Sentence Encoder Model

We adapt the best architecture from Conneau et al. (2017) as our sentence encoder. This architecture uses a standard bidirectional LSTM (Graves et al., 2013), followed by temporal max-pooling to create sentence vectors. We parameterize the BiLSTM with the different weights  $\theta_1$  and  $\theta_2$  to reflect the asymmetry of sentence processing. We then concatenate the forward and backward encodings.

We apply global max pooling to construct the encoding for each sentence. That is, we apply an element-wise max operation over the temporal dimension of the hidden states. Global max pooling builds a sentence representation from all time steps in the processing of a sentence (Collobert and Weston, 2008; Conneau et al., 2017), providing regularization and shorter back-propagation paths.

$$\begin{aligned} \vec{h}_t &= \text{LSTM}_t(w_1, \dots, w_t | \theta_1) \\ \overleftarrow{h}_t &= \text{LSTM}_t(w_T, \dots, w_t | \theta_2) \\ h_t &= [\vec{h}_t; \overleftarrow{h}_t] \\ s_i &= \text{MaxPool}(h_1, \dots, h_T) \end{aligned} \quad (1)$$

Our objective is to predict the discourse relations between two sentences from their vectors,  $s_i$  where  $i \in \{1, 2\}$ . Because we want generally useful sentence vectors after training, the learned computation should happen before the sentences are combined to make a prediction. However, some non-linear interactions between the sentence vectors are likely to be needed. To achieve this, we include a fixed set of common pair-wise vector operations: subtraction, multiplication, and average.

$$\begin{aligned} s_{\text{avg}} &= \frac{1}{2}(s_1 + s_2) \\ s_{\text{sub}} &= s_1 - s_2 \\ s_{\text{mul}} &= s_1 * s_2 \\ S &= [s_1, s_2, s_{\text{avg}}, s_{\text{sub}}, s_{\text{mul}}] \end{aligned} \quad (2)$$

Finally we use an affine fully-connected layer to project the concatenated vector  $S$  down to a lower dimensional representation, and then project it down to a vector of label size (the number of discourse markers). We use softmax to compute the probability distribution over discourse relations.

#### 3.2 Fine-tuning Model

Sentence relations datasets can be used to provide high-level training signals to fine-tune other sentence embedding models. In this work, we fine-tune BERT (Devlin et al., 2018) on the DisSent task and evaluate its performance on the PDTB implicit relation prediction task. We use the BERT-base model which has a 12-layer Transformer encoder. We directly use the [CLS] token’s position as the embedding for the entire sentence pair.

After training BERT-base model on the DisSent task, we continue to fine-tune BERT-base model on other evaluation tasks to see if training on DisSent tasks provides additional performance improvement and learning signal for the BERT-base model.

### 4 Data Collection

We present an automatic way to collect a large dataset of sentence pairs and the relations between them from natural text corpora using a set of explicit discourse markers and universal dependency parsing (Schuster and Manning, 2016).

#### 4.1 Corpus and Discourse Marker Set

For training and evaluation datasets, we collect sentence pairs from BookCorpus (Zhu et al., 2015), text from unpublished novels (*Romance*, *Fantasy*, *Science fiction*, and *Teen* genres), which was used by Kiros et al. (2015) to train their SkipThought model. We identified common discourse markers, choosing those with a frequency greater than 1% in PDTB. Our final set of discourse markers is shown in Table 2 and we experiment with three subsets of discourse markers (ALL, 5, and 8), shown in Table 4.

#### 4.2 Dependency Parsing

Many discourse markers in English occur almost exclusively between the two statements they connect, but for other discourse markers, their position relative to their connected statements can vary (e.g. Figure 1). For this reason, we use the Stanford CoreNLP dependency parser (Schuster and

S1	marker	S2
Her eyes flew up to his face.	and	Suddenly she realized why he looked so different.
The concept is simple.	but	The execution will be incredibly dangerous.
You used to feel pride.	because	You defended innocent people.
I'll tell you about it.	if	You give me your number.
Belter was still hard at work.	when	Drade and barney strolled in.
We plugged bulky headsets into the dashboard.	so	We could hear each other when we spoke into the microphones.
It was mere minutes or hours.	before	He finally fell into unconsciousness.
And then the cloudy darkness lifted.	though	The lifeboat did not slow down.

Table 3: **Example pairs** from our Books 8 dataset.

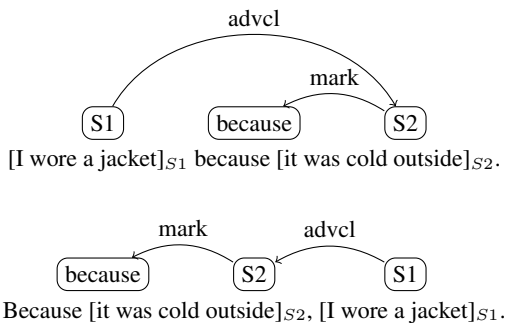


Figure 1: **Dependency patters for extraction:** While the relative order of a discourse marker (e.g. *because*) and its connected sentences is flexible, the dependency relations between these components within the overall sentence remains constant. See Appendix A.1 for dependency patterns for other discourse markers.

Manning, 2016) to extract the appropriate pairs of sentences (or sentence-like EDUs) for a discourse marker, in the appropriate conceptual order. Each discourse marker, when it is used to link two statements, is parsed by the dependency parser in a systematic way, though different discourse markers may have different corresponding dependency patterns linking them to their statement pairs.<sup>1</sup>

Within the dependency parse, we search for the governor phrase (which we call “S2”) of the discourse marker and check for the appropriate dependency relation. If we find no such phrase, we reject the example entirely (thus filtering out polysemous usages, like “that’s *so* cool!” for the discourse marker *so*). If we find such an S2, we search for “S1” within the same sentence (SS). Searching for this relation allows us to capture pairs where the discourse marker starts the sentence and connects the following two clauses (e.g. “*Because* [it was cold outside]<sub>S2</sub>, [I wore a jacket]<sub>S1</sub>.”). If a sentence in the corpus contains only a discourse marker and S2, we assume the

<sup>1</sup> See Appendix A.1 for more details on dependency-based extraction.

discourse marker links to the immediately previous sentence (IPS), which we label S1.

For some markers, we further filter based on the order of the sentences in the original text. For example, the discourse marker *then* always appears in the order “S1, then S2”, unlike *because*, which can also appear in the order “Because S2, S1”. Excluding proposed extractions in an incorrect order makes our method more robust to incorrect dependency parses.

### 4.3 Training Dataset

Using these methods, we curated a dataset of 4,706,292 pairs of sentences for 15 discourse markers. Examples are shown in Table 3. We randomly divide the dataset into train/validation/test set with 0.9, 0.05, 0.05 split. The dataset is inherently unbalanced, but the model is still able to learn rarer classes quite well (see Appendix A.4 for more details on the effects of class frequencies). Our data are publicly available at <https://github.com/windweller/DisExtract>.

## 5 Related Work

Current state of the art models either rely on completely supervised learning through high-level classification tasks or unsupervised learning.

Supervised learning has been shown to yield general-purpose representations of meaning, training on semantic relation tasks like Stanford Natural Language Inference (SNLI) and MultiNLI (Bowman et al., 2015; Williams et al., 2018; Conneau et al., 2017). Large scale joint supervised training has also been explored by Subramanian et al. (2018), who trained a sentence encoding model on five language-related tasks. These supervised learning tasks often require human annotations on a large amount of data which are costly to obtain. Our discourse prediction approach extends

these results in that we train on semantic relations, but we use dependency patterns to automatically curate a sizable dataset.

In an unsupervised learning setting, SkipThought (Kiros et al., 2015) learns a conditional joint probability distribution for the next sentence. ELMo (Peters et al., 2018) uses a BiLSTM to predict the missing word using the masked language modeling (MLM) objective. OpenAI-GPT2 (Radford et al., 2019) directly predicts the next word. BERT (Devlin et al., 2018) uses MLM as well as predicting whether the next sequence comes from the same document or not. Despite the overwhelming success of these models, Phang et al. (2018) shows that fine-tuning these models on supervised learning datasets can yield improved performance over difficult natural language understanding tasks.

Jernite et al. (2017) have proposed a model that also leverages discourse relations. They manually categorize discourse markers based on human interpretations of discourse marker similarity, and the model predicts the category instead of the individual discourse marker. Their model also trains on auxiliary tasks, such as sentence ordering and ranking of the following sentence and must compensate for data imbalance across tasks. Their data collection methods only allow them to look at paragraphs longer than 8 sentences, and sentence pairs with sentence-initial discourse markers, resulting in only 1.4M sentence pairs from a much larger corpus. Our proposed model extracts a wider variety of sentence pairs, can be applied to corpora with shorter paragraphs, and includes no auxiliary tasks.

## 6 Experiments

For all our models, we tuned hyperparameters on the validation set, and report results from the test set. We use stochastic gradient descent with initial learning rate 0.1, and anneal by the factor of 5 each time validation accuracy is lower than in the previous epoch. We train our fixed sentence encoder model for 20 epochs, and use early stopping to prevent overfitting. We also clip the gradient norm to 5.0. We did not use dropout in the fully connected layer in the final results because our initial experiments with dropout showed lower performance when generalizing to SentEval. We experimented with both global mean pooling and global max pooling and found the later to perform

Label	Discourse Markers
Books 5	<i>and, but, because, if, when</i>
Books 8	<i>and, but, because, if, when, before, though, so</i>
Books ALL	<i>and, but, because, if, when, before, though, so, as, while, after, still, also, then, although</i>

Table 4: **Discourse marker sets** used in our experiments. Books ALL contains 4.7M sentence pairs, Books 8 contains 3.6M, and Books 5 contains 3.2M.

much better at generalization tasks. All models we report used a 4096 hidden state size. We are able to fit our model on a single Nvidia Titan X GPU.

**Fine-tuning** We fine-tune the BERT-base model on the DisSent tasks with 2e-5 learning rate for 1 epoch because all DisSent tasks corpora are quite large and fine-tuning for longer epochs did not yield improvement. We fine-tune BERT on other supervised learning datasets for multiple epochs and select the epoch that provides the best performance on the evaluation task. We find that fine-tuning on MNLI is better than on SNLI or both combined. This phenomenon is also discussed in Phang et al. (2018).

**Discourse Marker Set** We experimented with three subsets of discourse markers, shown in Table 4. We first trained over all of the discourse markers in our ALL marker set. The model achieved 67.5% test accuracy on this classification task. Overall we found that markers with similar meanings tended to be confusable with one another. A more detailed analysis of the model’s performance on this classification task is presented in Appendix A.4.

Because there appears to be intrinsic conceptual overlap in the set of ALL markers, we experimented on different subsets of discourse markers. We choose sets of 5 and 8 discourse markers that were both non-overlapping and frequent. The set of sentence pairs for each smaller dataset is a strict subset of those in any larger dataset. Our chosen sets are shown in Table 4.

**Marked vs Unmarked Prediction Task** Adjacent sentences will always have a relationship, but some are marked with discourse markers while others are not. Humans have been shown to perform well above chance at guessing whether a discourse marker is marked vs. unmarked (Patterson and Kehler, 2013; Yung et al., 2017), indicating a

systematicity to this decision.

We predict that high quality sentence embeddings will contain useful information to determine whether a discourse relation is explicitly marked. Furthermore, success at this task could help natural language generation models to generate more human-like long sequences.

To test this prediction, we create an additional set of tasks based on Penn Discourse Treebank (Rashmi et al., 2008). This hand-annotated dataset contains expert discourse relation annotations between sentences. We collected 34,512 sentences from PDTB<sup>2</sup> (see Appendix), where 16,224 sentences are marked with implicit relation type, and 18,459 are marked with explicit relation type.

**Implicit Relation Prediction Task** Sporleder and Lascarides (2008) have argued that sentence pairs with explicitly marked relations are qualitatively different from those where the relation is left implicit. However, despite such differences, Qin et al. (2017) were able to use an adversarial network to leverage explicit discourse data as additional training to increase the performance on the implicit discourse relation prediction task. We use the same dataset split scheme for this task as for the implicit vs explicit task discussed above. Following Ji and Eisenstein (2015) and Qin et al. (2017), we predict the 11 most frequent relations. There are 13,445 pairs for training, and 1,188 pairs for evaluation.

**SentEval Tasks** We evaluate the performance of generated sentence embeddings from our fixed sentence encoder model on a series of natural language understanding benchmark tests provided by Conneau et al. (2017). The tasks we chose include sentiment analysis (MR, SST), question-type (TREC), product reviews (CR), subjectivity-objectivity (SUBJ), opinion polarity (MPQA), entailment (SICK-E), relatedness (SICK-R), and paraphrase detection (MRPC). These are all classification tasks with 2-6 classes, except for relatedness, for which the model predicts human similarity judgments.

## 6.1 Results

**Training Task** On the discourse marker prediction task used for training, we achieve high levels of test performance for all discourse markers. (Though it is interesting that *because*, perhaps the

<sup>2</sup> <https://github.com/cgpotts/pdtb2>

Model	All		Books 8		Books 5	
	F1	Acc	F1	Acc	F1	Acc
GloVe-bow	17.1	41.8	27.6	47.3	41.7	52.5
Ngram-bow	28.1	51.8	44.0	58.1	54.1	63.3
BiLSTM	47.2	67.5	64.4	73.5	72.1	77.3
BERT	60.1	77.5	76.2	82.9	82.6	86.1

Table 5: **Discourse classification task performance:** Unweighted average F1 across discourse markers on the test set, and overall accuracy. Ngram-bow is a bag-of-words model built on mixture of ngram features. GloVe-bow averages word embedding with correction to frequency (Arora et al., 2017). BiLSTM is the DisSent sentence encoder model. BERT is finetuned on all of the DisSent tasks.

conceptually deepest relation, is also systematically the hardest for our model.) The larger the set of discourse markers, the more difficult the task becomes, and we therefore see lower test accuracy despite larger dataset size. We conjecture that as we increase the number of discourse markers, we also increase the ambiguity between them (semantic overlap in discourse markers’ meanings), which may further explain the drop in performance. The training task performance for each subset is shown in Table 5. We provide per-discourse-marker performance in the Appendix.

**Discourse Marker Set** Varying the set of discourse markers doesn’t seem to help or hinder the model’s performance on generalization tasks. Top generalization performance on the three sets of discourse markers is shown in Table 6. Similar generalization performance was achieved when training on 5, 8, and all 15 discourse markers.

The similarity in generalization performance across discourse sets shows that the top markers capture most relationships in the training data.

**Marked vs Unmarked Prediction Task** In determining whether a discourse relation is marked or unmarked, DisSent models outperform InferSent and SkipThought (as well as previous approaches on this task) by a noticeable margin. Much to our surprise, fine-tuned BERT models are not able to perform better than the BiLSTM sentence encoder model. We leave explorations of this phenomenon to future work. We report the results in Table 7 under column MVU.

**Implicit Discourse Relation Task** Not surprisingly, DisSent task provided the much needed distant supervision to classify the types of implicit

Model	MR	CR	SUBJ	MPQA	SST	TREC	SICK-R	SICK-E	MRPC
Self-supervised training methods									
DisSent Books 5 <sup>†</sup>	80.2	85.4	93.2	90.2	82.8	91.2	0.845	83.5	76.1
DisSent Books 8 <sup>†</sup>	79.8	85.0	93.4	<u>90.5</u>	83.9	93.0	<u>0.854</u>	83.8	<u>76.1</u>
DisSent Books ALL <sup>†</sup>	80.1	84.9	<u>93.6</u>	90.1	<u>84.1</u>	<b>93.6</b>	0.849	83.7	75.0
Disc BiGRU	—	—	88.6	—	—	81.0	—	—	71.6
Unsupervised training methods									
FastSent	70.8	78.4	88.7	80.6	—	76.8	—	—	72.2
FastSent + AE	71.8	76.7	88.8	81.5	—	80.4	—	—	71.2
Skipthought	76.5	80.1	93.6	87.1	82.0	92.2	0.858	82.3	73.0
Skipthought-LN	79.4	83.1	93.7	89.3	82.9	88.4	0.858	79.5	—
Supervised training methods									
DictRep (bow)	76.7	78.7	90.7	87.2	—	81.0	—	—	—
InferSent	81.1	86.3	92.4	90.2	<b>84.6</b>	88.2	0.884	86.1	76.2
Multi-task training methods									
LSMTL	<b>82.5</b>	<b>87.7</b>	<b>94.0</b>	<b>90.9</b>	83.2	93.0	<b>0.888</b>	<b>87.8</b>	<b>78.6</b>

Table 6: **SentEval Task Results Using Fixed Sentence Encoder.** We report the best results for generalization tasks. <sup>†</sup> indicates models that we trained. FastSent, FastSent + AE (Hill et al., 2016), SkipThought (Kiros et al., 2015), SkipThought-LN, DictRep (bow), and InferSent are reported from Conneau et al. (2017). LSMTL is reported from Subramanian et al. (2018). Globally best results are shown in **bold**, best DisSent results are underlined.

discourse relations much better than InferSent and SkipThought. DisSent outperforms word vector models evaluated by Qin et al. (2017), and is only 3.3% lower than the complex state of the art model that uses adversarial training designed specifically for this task. When we fine-tune BERT models on the DisSent corpora, we are able to outperform all other models and achieve state-of-the-art result on this task. We report the results in Table 7 under column IMP.

**SentEval Tasks** Results of our models, and comparison to other approaches, are shown in Table 6. Despite being a much simpler task than SkipThought and allowing for much more scalable data collection than InferSent, DisSent performs as well or better than these approaches on most generalization tasks.

DisSent and InferSent do well on different sets of tasks. In particular, DisSent outperforms InferSent on TREC (question-type classification). InferSent outperforms DisSent on the tasks most similar to its training data, SICK-R and SICK-E. These tasks, like SNLI, were crowdsourced, and seeded with images from Flickr30k corpus (Young et al., 2014).

Although DisSent is trained on a dataset derived from the same corpus as SkipThought, DisSent almost entirely dominates SkipThought’s performance across all tasks. In particular, on the SICK

Model	IMP	MVU
Sentence Encoder Models		
SkipThought (Kiros et al., 2015)	9.3	57.2
InferSent (Conneau et al., 2017)	39.3	84.5
Patterson and Kehler (2013)	—	86.6
Fine-tuned Models		
DisSent Books 5	40.7	86.5
DisSent Books 8	41.4	<b>87.9</b>
DisSent Books ALL	<b>42.9</b>	87.6
Fine-tuned Models		
BERT	52.7	80.5
BERT + MNLI	53.7	80.7
BERT + SNLI + MNLI	51.3	79.8
BERT + DisSent Books 5	<b>54.7</b>	81.6
BERT + DisSent Books 8	52.4	80.6
BERT + DisSent Books ALL	53.2	<b>81.8</b>
Previous Single Task Models		
Word Vectors (Qin et al., 2017)	36.9	74.8
Lin et al. (2009) + Brown Cluster	40.7	—
Adversarial Net (Qin et al., 2017)	<b>46.2</b>	—

Table 7: **Discourse Generalization Tasks using PDTB:** We report test accuracy for sentence embedding and state-of-the-art models.

dataset, DisSent and SkipThought perform similarly on the relatedness task (SICK-R), but DisSent strongly outperforms SkipThought on the entailment task (SICK-E). This discrepancy highlights an important difference between the two models. Whereas both models are trained to, given a particular sentence, identify words that appear near that sentence in the corpus, DisSent focuses

on learning specific kinds of relationships between sentences – ones that humans tend to explicitly mark. We find that reducing the model’s task to only predicting a small set of discourse relations, rather than trying to recover all words in the following sentence, results in better features for identifying entailment and contradiction without losing cues to relatedness.

Overall, on the evaluation tasks we present, DisSent performs on par with previous state-of-the-art models and offers advantages in data collection and training speed.

## 7 Extraction Validation

We evaluate our extraction quality by comparing the manually extracted and annotated sentence pairs from Penn Discourse Treebank (PDTB) to our automatic extraction of sentence pairs from the source corpus Penn Treebank (PTB). On the majority of discourse markers, we can achieve a relatively high extraction precision.

We apply our extraction pipeline on raw PTB dataset because we want to see how well our pipeline converts raw corpus into a dataset. Details of our alignment procedure is described in Appendix A.2. Overall, even though we cannot construct the explicit discourse prediction section of the PDTB dataset perfectly, training with imprecise extraction has little impact on the sentence encoder model’s overall performance.

We compute the extraction precision as the percentage of PTB extracted pairs that can be successfully aligned to PDTB. In Figure 2, we show that extraction precision varies across discourse markers. Some markers have higher quality (e.g. *because*, *so*) and some lower quality (e.g. *and*, *still*).

We show in Figure 3 that we tend to have low distances overall for the successfully aligned pairs. That is, whenever our extraction pipeline yields a match, the dependency parsing patterns do extract high quality training pairs.

## 8 Discussion

**Implicit and explicit discourse relations** We focus on explicit discourse relations for training our embeddings. Another meaningful way to exploit discourse relations in training is by leveraging implicit discourse signals. For instance, [Jernite et al. \(2017\)](#) showed that predicting sentence ordering could help to generate meaningful sentence embeddings. But adjacent sentences can be




Figure 2: **Extraction error rates:** proportion of unalignable extracted pairs per discourse marker.




Figure 3: **Extraction quality for aligned pairs:** Distances from aligned extracted pairs to nearest gold pair.

related to one another in many different, complicated ways. For example, sentences linked by contrastive markers, like *but* or *however* are likely expressing different or opposite ideas.

Identifying other features of natural text that contain informative signals of discourse structure and combining these with explicit discourse markers is an appealing direction for future research.

**Multilingual generalization** In principle, the DisSent model and extraction methods would apply equally well to multilingual data with minimal language-specific modifications. Within universal dependency grammar, discourse markers across languages should correspond to structurally similar dependency patterns. Beyond dependency parsing and minimal marker-specific pattern development (see Appendix A.1), our extraction method is automatic, requiring no annotation of the original dataset, and so any large dataset of raw text in a language can be used.



## 9 Conclusion

We present a *discourse marker prediction* task for training and fine-tuning sentence embedding models. We train our model on this task and show that the resulting embeddings lead to high performance on a number of established tasks for sentence embeddings. We fine-tune larger models on this task and achieve state-of-the-art on the PDTB implicit discourse relation prediction.

A dataset for this task is easy to collect relative to other supervised tasks. It provides cheap and noisy but strong training signals. Compared to unsupervised methods that train on a full corpus, our method yields more targeted and faster training. Encouragingly, the model trained on discourse marker prediction achieves comparable generalization performance to other state of the art models.

## Acknowledgement

We thank Chris Potts for the discussion on Penn Discourse Treebank and the share of preprocessing code. We also thank our anonymous reviewers and the area chair for their thoughtful comments and suggestions. The research is based upon work supported by the Defense Advanced Research Projects Agency (DARPA), via the Air Force Research Laboratory (AFRL, Grant No. FA8650-18-C-7826). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA, the AFRL or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon.

## References

Sanjeev Arora, Yingyu Liang, and Tengyu Ma. 2017. [A simple but tough-to-beat baseline for sentence embeddings](#). *International Conference on Learning Representations*.

Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics.

Lynn Carlson and Daniel Marcu. 2001. [Discourse tagging reference manual](#). *ISI Technical Report ISI-TR-545*, 54:56.

Ronan Collobert and Jason Weston. 2008. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the 25th international conference on Machine learning*, pages 160–167. ACM.

Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. [Supervised learning of universal sentence representations from natural language inference data](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680, Copenhagen, Denmark. Association for Computational Linguistics.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [Bert: Pre-training of deep bidirectional transformers for language understanding](#). *arXiv preprint arXiv:1810.04805*.

Alex Graves, Navdeep Jaitly, and Abdel-rahman Mohamed. 2013. Hybrid speech recognition with deep bidirectional lstm. In *Automatic Speech Recognition and Understanding (ASRU), 2013 IEEE Workshop on*, pages 273–278. IEEE.

Felix Hill, Kyunghyun Cho, and Anna Korhonen. 2016. [Learning distributed representations of sentences from unlabelled data](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1367–1377, San Diego, California. Association for Computational Linguistics.

Jerry R Hobbs. 1979. [Coherence and coreference](#). *Cognitive science*, 3(1):67–90.

Jerry R Hobbs. 1985. [On the coherence and structure of discourse](#). *Tech. Rep*, CSLI-85-37.

Katja Jasinskaja and Elena Karagjosova. 2015. Rhetorical relations. *The companion to semantics*. Oxford: Wiley.

Yacine Jernite, Samuel R. Bowman, and David Sonntag. 2017. [Discourse-Based Objectives for Fast Unsupervised Sentence Representation Learning](#). *arXiv:1705.00557 [cs, stat]*. ArXiv: 1705.00557.

Yangfeng Ji and Jacob Eisenstein. 2015. [One vector is not enough: Entity-augmented distributed semantics for discourse relations](#). *Transactions of the Association for Computational Linguistics*, 3:329–344.

Ryan Kiros, Yukun Zhu, Ruslan R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Skip-thought vectors. In *Advances in neural information processing systems*, pages 3294–3302.

Alistair Knott. 1996. A data-driven methodology for motivating a set of coherence relations.

- Xiang Lin, Shafiq Joty, Prathyusha Jwalapuram, and Saiful Bari. 2019. [A unified linear-time framework for sentence-level discourse parsing](#).
- Ziheng Lin, Min-Yen Kan, and Hwee Tou Ng. 2009. Recognizing implicit discourse relations in the penn discourse treebank. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 1-Volume 1*, pages 343–351. Association for Computational Linguistics.
- Eric Malmi, Daniele Pighin, Sebastian Krause, and Mikhail Kozhevnikov. 2018. [Automatic prediction of discourse connectives](#). In *Proceedings of the 11th Language Resources and Evaluation Conference*, Miyazaki, Japan. European Language Resource Association.
- William C Mann and Sandra A Thompson. 1988. Rhetorical structure theory: Toward a functional theory of text organization. *Text-Interdisciplinary Journal for the Study of Discourse*, 8(3):243–281.
- Gary Patterson and Andrew Kehler. 2013. Predicting the presence of discourse connectives. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 914–923.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Jason Phang, Thibault Févry, and Samuel R Bowman. 2018. [Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks](#). *arXiv preprint arXiv:1811.01088*.
- Lianhui Qin, Zhisong Zhang, Hai Zhao, Zhiting Hu, and Eric Xing. 2017. [Adversarial connective-exploiting networks for implicit discourse relation classification](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1006–1017, Vancouver, Canada. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#).
- Prasad Rashmi, Dinesh Nihkil, Lee Alan, Miltsakaki Eleni, Robaldo Livio, Joshi Aravind, Webber Bonnie, et al. 2008. The penn discourse treebank 2.0. In *Lexical Resources and Evaluation Conference*.
- Sebastian Schuster and Christopher D Manning. 2016. Enhanced english universal dependencies: An improved representation for natural language understanding tasks. In *Lexical Resources and Evaluation Conference*.
- Caroline Sporleder and Alex Lascarides. 2008. Using automatically labelled examples to classify rhetorical relations: An assessment. *Natural Language Engineering*, 14(3):369–416.
- Sandeep Subramanian, Adam Trischler, Yoshua Bengio, and Christopher J Pal. 2018. [Learning general purpose distributed sentence representations via large scale multi-task learning](#). In *International Conference on Learning Representations*.
- Bonnie Webber, Matthew Stone, Aravind Joshi, and Alistair Knott. 2003. Anaphora and discourse structure. *Computational linguistics*, 29(4):545–587.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78.
- Frances Yung, Kevin Duh, Taku Komura, and Yuji Matsumoto. 2017. A psycholinguistic model for the marking of discourse relations. *Dialogue & Discourse*, 8(1):106–131.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. [Aligning books and movies: Towards story-like visual explanations by watching movies and reading books](#). In *Proceedings of the IEEE international conference on computer vision*, pages 19–27.

## A Appendix

### A.1 Details on Dependency Based Sentence Extraction

While universal dependency grammar provides enough information to identify discourse markers and their connected statements, different discourse markers are parsed with different dependency relations. For each discourse marker of interest, we identify the appropriate dependency pattern (see Figure 4).

We excluded any pair where one of the sentences was less than 5 or more than 50 words long and any pairs where one of the sentences was more than 5 times the length of the other.

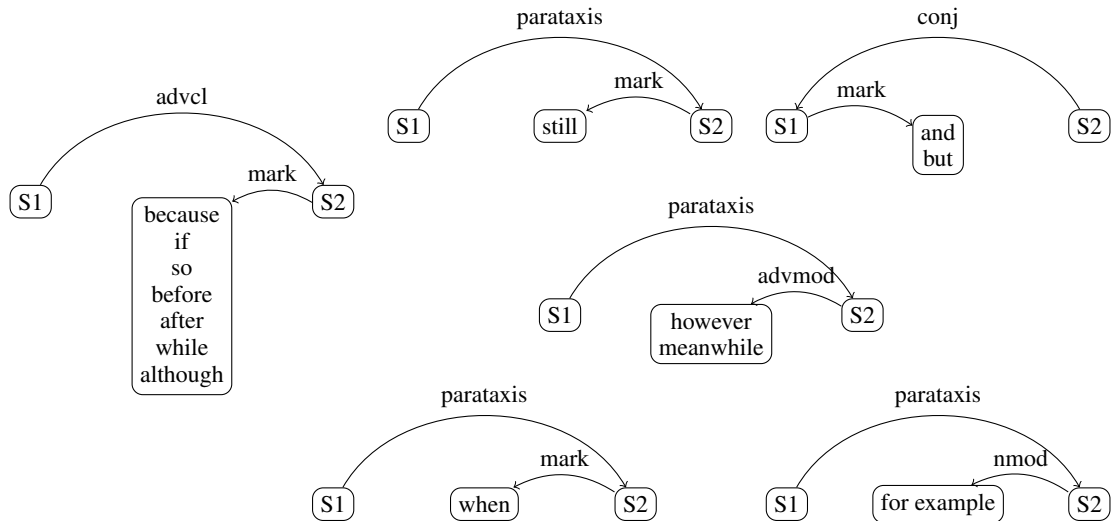


Figure 4: Dependency patterns used for extraction for each discourse marker.

Dependency parsing allows us to design our extraction method such that each S1 and S2 is interpretable as a full sentence in isolation, and the appropriate conceptual relation holds between the pair. However, occasionally we get ungrammatical sentences or the wrong pair of sentences for a relation. This incorrect extraction can happen in several ways. First, we might choose grammatical but incorrect pairs. Rashmi et al. (2008) found that 61% of discourse markers appear in the same sentence (SS) with both S1 and S2, and another 30% link S2 to the immediate predecessor (IPS). For the remaining examples (non-adjacent previous sentence NAPS - 9%, or following sentence FS - less than 1%), our method incorrectly extracts an IPS pair. Second, not all parses are correct (e.g. “Himself close his eyes.” was extracted due to an incorrect parse). Finally, even with correct parses, some extracted sentences are nonsensical or ungrammatical out of context due to implicit subjects, unresolved pronouns, or marked embedded clauses. Fortunately these errors were relatively rare, and many could be avoided simply by enforcing that the extracted sentences each have a main verb and satisfy a minimum length. Overall this method extracts high-quality sentence pairs with appropriately labeled relations.

## A.2 Procedures in Extraction Validation

We preprocess the PTB sentences by limiting the vocabulary size to 10,000 and tokenizing numbers. Then we run our extraction pipeline on the preprocessed PTB. We apply the same preprocessing to the PDTB sentences.

We refer to the gold sentence pair from the PDTB as (G1, G2), and our extracted sentence pair from the PTB as (S1, S2). We first obtain the minimum of S1-G1 distance and S2-G2 distance over all gold pairs. If this distance is smaller than 0.7, we consider the corresponding gold pair to be an alignment for this extracted pair.

Given an aligned pair ((G1, G2), (S1, S2)), we measure the extraction quality by computing the average of normalized G1-S1 and G2-S2 distance. We compute this distance for all pairs and all discourse markers.

We analyze our extraction quality in two steps: align sentence pairs from the two datasets and then calculate extraction quality on each aligned pair. In the alignment step, for each extracted pair, we calculate its distance to all pairs from PDTB using the normalized Levenshtein distance:

$$d(s_1, s_2) = \frac{\text{Levenshtein}(s_1, s_2)}{\max(\text{len}(s_1), \text{len}(s_2))} \quad (3)$$

## A.3 Implicit vs. Explicit Prediction Task Setup

For each pair of connected sentences, whose relation type has been labeled in PDTB, the discourse relation between them may have been explicitly marked (via a discourse relation word) or not. We can pose the task of a binary classification of whether the sentence pair appeared as explicitly or implicitly marked, given only the two sentences and no additional information. We evaluate DisSent and InferSent sentence embedding models and a word vector baseline on this task.

We follow Patterson and Kehler (2013)’s pre-processing. The dataset contains 25 sections in total. We use sections 0 and 1 as the development set, sections 23 and 24 for the test set, and we train on the remaining sections 2-22.

This task is different from the setting in Patterson and Kehler (2013). We do not allow the classifier to access the underlying discourse relation type and we only provide the individual sentence embeddings as input features. In contrast, Patterson and Kehler (2013) used a variety of discrete features provided by the PDTB dataset for their classifier, including the hand-annotated relation types.

#### A.4 Classification Performance

To investigate the qualitative relations among our largest set of discourse markers, the ALL marker set, we build a confusion matrix of the test set classifications. Figure 5 reflects classification performance for the model trained on the full dataset, that we later show generalization results for. This model is clearly influenced by frequency, such that it tends to misclassify infrequent discourse markers as frequent ones. However, deviations from the effect of frequency appear to be semantically meaningful.

Classifications errors are much more common for semantically similar discourse marker pairs than would be expected from frequency alone. The most common confusion is when the synonymous marker *although* is mistakenly classified as *but*. The temporal relation markers *before*, *after* and *then*, intuitively very similar discourse markers, are rarely confused for anything but each other. The fact that they are indeed confusable may reflect the tendency of authors to mark temporal relation primarily when it is ambiguous.

Figure 6 reflects a model trained on a balanced subset of our training set. When the model can no longer rely on base rates of discourse markers to make judgments, overall accuracy drops from 68% to 47%. However inspecting the matrices shows very similar confusability, suggesting that training on unbalanced data does not greatly decrease sensitivity to non-frequency predictors.

To more quantitatively represent the connection between what the two models learn, we compute the correlation between the balanced confusions and the *residuals* of the unbalanced confusions (when predicted linearly from log fre-

quency). These residuals account for 64% of the variance in the balanced confusions ( $R^2 = 0.6431$ ,  $F(1, 223) = 401.8$ ,  $p < .001$ ). That is, we come close to predicting the balanced confusions from the unbalanced ones.




Figure 5: **Confusion Matrix** trained on the ALL dataset extracted from BookCorpus. Each cell represents the proportion of instances of the actual discourse marker misclassified as the classified discourse marker. This proportion is log-transformed to highlight small differences. Discourse markers are arranged in order of frequency from left (least frequent) to right (most frequent).

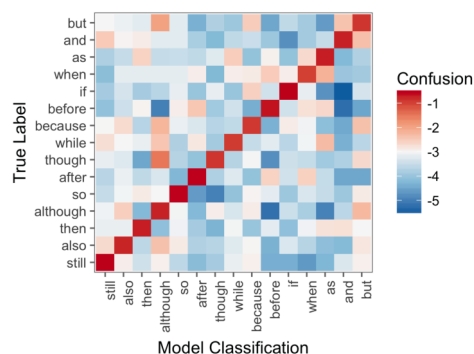


Figure 6: **Balanced Classifier Confusion Matrix** trained on a balanced subset of the ALL dataset where discourse markers are capped at 13,421 occurrences each. Each cell represents the proportion of instances of the actual discourse marker misclassified as the classified discourse marker. This proportion is log-transformed to highlight small differences. Discourse markers are arranged in order of frequency from left (least frequent) to right (most frequent).

#### A.5 Baseline performance on training task

As a reference point for training task performance we present baseline performance. Note that a model which simply chose the most common class would perform with 21.79% accuracy on the ALL task, 28.35% on the BOOKS 8 task, and 31.87%

on the BOOKS 5 task. Using either unigram, bigram and trigram bag of words or Arora et al. (2017)’s baseline sentence representations as features to a logistic regression results in much lower performance than our DisSent classifier. Table 9 shows the precision and recall for the bag-of-words model. Table 10 shows the precision and recall for the Arora et al. (2017) embeddings.

Marker	All		Books 8		Books 5	
	Prec	Rec	Prec	Rec	Prec	Rec
and	71.8	78.2	78.3	78.5	80.6	79.4
but	71.4	73.2	72.3	79.1	75.3	79.9
because	44.9	36.2	50.1	36.9	54.8	37.7
if	79.1	75.0	77.5	79.6	80.7	81.4
when	60.5	61.8	71.2	74.0	76.9	77.2
so	49.3	48.0	55.8	46.1	—	—
though	48.0	29.7	61.0	38.8	—	—
before	65.0	60.9	76.6	63.5	—	—
as	68.0	76.5	—	—	—	—
while	45.6	35.9	—	—	—	—
after	55.5	41.9	—	—	—	—
although	24.4	6.7	—	—	—	—
still	42.0	20.9	—	—	—	—
also	36.1	13.6	—	—	—	—
then	30.9	11.7	—	—	—	—
Avg	66.7	68.0	73.6	73.3	77.5	77.4
Accuracy	67.5		73.5		77.3	

Table 8: **DisSent model performance:** Test recall / precision for each discourse marker on the classification task, weighted average precision and recall across discourse markers, and overall accuracy.

## A.6 Embedding dimensions of models

DisSent uses a BiLSTM encoder with 4096 hidden state dimensions. InferSent (Conneau et al., 2017) uses 4096 embedding dimensions. Disc BiGRU (Jernite et al., 2017) hidden state has 512 dimensions. FastSent and FastSent + AE (Hill et al., 2016) have 500 dimensions. SkipThought (Kiros et al., 2015) and SkipThought-LN (Conneau et al., 2017) models trained on 600-dimension word embeddings and produced 2400-dimension sentence embeddings. DictRep (bow) is from Conneau et al. (2017). LSMTL (Subramanian et al., 2018) uses 2048-dimension bi-directional GRU as encoder, and trained on 512 dimension word embeddings.

## A.7 Limitations of evaluation

The generalization tasks that we (following Conneau et al. (2017)) use to compare models focus on sentiment, entailment, and similarity. These are narrow operational definitions of semantic mean-

Marker	All		Books 8		Books 5	
	Prec	Rec	Prec	Rec	Prec	Rec
and	60.1	65.0	65.6	70.1	70.1	71.4
but	49.9	65.3	55.2	69.4	59.7	69.9
because	34.7	10.2	42.1	11.1	42.8	10.6
if	54.6	56.9	58.8	56.4	64.4	60.0
when	43.2	40.1	52.1	52.2	58.4	54.3
so	35.5	11.3	38.5	11.0	—	—
though	40.6	20.8	56.2	25.2	—	—
before	47.8	29.1	56.6	35.4	—	—
as	51.9	63.1	—	—	—	—
while	33.4	11.6	—	—	—	—
after	41.0	17.6	—	—	—	—
although	11.9	0.4	—	—	—	—
still	34.7	2.5	—	—	—	—
also	16.7	0.4	—	—	—	—
then	36.2	2.1	—	—	—	—
Average	40.2	40.3	46.2	44.5	53.3	50.7
Accuracy	51.8		58.1		63.3	

Table 9: **Ngram Bag-of-words baseline sentence embeddings performance on DisSent training task:** test recall / precision for each discourse marker on the classification task, and overall accuracy. Average metric reports the weighted average of all classes.

ing. A model that generates meaningful sentence embeddings should excel at these tasks. However, success at these tasks does not necessarily imply that a model has learned a deep semantic understanding of a sentence.

Sentiment classification, for example, in many cases only requires the model to understand local structures. Text similarity can be computed with various textual distances (e.g., Levenshtein or Jaro distance) on bag-of-words, without a compositional representation of the sentence. Thus, the ability of our, and other, models to achieve high performance on these metrics may reflect a competent representation sentence meaning; but more rigorous tests are needed to understand whether these embeddings capture sentence meaning in general.

Marker	All		Books 8		Books 5	
	Prec	Rec	Prec	Rec	Prec	Rec
and	46.9	59.4	52.9	63.6	58.0	64.3
but	38.1	57.9	43.5	62.3	48.9	62.4
because	24.1	0.5	20.2	0.3	27.7	0.47
if	41.8	37.1	46.2	37.9	50.5	38.2
when	36.8	25.8	45.6	40.0	58.3	41.3
so	37.0	2.5	39.5	2.9	—	—
though	27.2	1.4	29.7	1.3	—	—
before	42.0	10.0	48.8	11.8	—	—
as	43.4	55.6	—	—	—	—
while	29.1	3.4	—	—	—	—
after	37.1	4.8	—	—	—	—
although	0.0	0.0	—	—	—	—
still	0.0	0.0	—	—	—	—
also	0.0	0.0	—	—	—	—
then	0.0	0.0	—	—	—	—
Avg	50.1	51.1	57.5	56.4	63.0	62.2
Accuracy	41.8		47.3		52.5	

Table 10: **Corrected GloVe Bag-of-words sentence embeddings performance on DisSent training task:** test recall / precision for each discourse marker on the classification task, and overall accuracy. Average metric reports the weighted average of all classes.