

A Latent Variable Recurrent Neural Network for Discourse Relation Language Models

Yangfeng Ji
Georgia Institute of Technology
Atlanta, GA 30308, USA
jiyfeng@gatech.edu

Gholamreza Haffari
Monash University
Clayton, VIC, Australia
gholamreza.haffari
@monash.edu

Jacob Eisenstein
Georgia Institute of Technology
Atlanta, GA 30308, USA
jacobe@gatech.edu

Abstract

This paper presents a novel latent variable recurrent neural network architecture for jointly modeling sequences of words and (possibly latent) discourse relations between adjacent sentences. A recurrent neural network generates individual words, thus reaping the benefits of discriminatively-trained vector representations. The discourse relations are represented with a latent variable, which can be predicted or marginalized, depending on the task. The resulting model can therefore employ a training objective that includes not only discourse relation classification, but also word prediction. As a result, it outperforms state-of-the-art alternatives for two tasks: implicit discourse relation classification in the Penn Discourse Treebank, and dialog act classification in the Switchboard corpus. Furthermore, by marginalizing over latent discourse relations at test time, we obtain a discourse informed language model, which improves over a strong LSTM baseline.

1 Introduction

Natural language processing (NLP) has recently experienced a neural network “tsunami” (Manning, 2016). A key advantage of these neural architectures is that they employ discriminatively-trained distributed representations, which can capture the meaning of linguistic phenomena ranging from individual words (Turian et al., 2010) to longer-range linguistic contexts at the sentence level (Socher et al., 2013) and beyond (Le and Mikolov, 2014). Because they are discriminatively trained, these meth-

ods can learn representations that yield very accurate predictive models (e.g., Dyer et al, 2015).

However, in comparison with the probabilistic graphical models that were previously the dominant machine learning approach for NLP, neural architectures lack flexibility. By treating linguistic annotations as random variables, probabilistic graphical models can marginalize over annotations that are unavailable at test or training time, elegantly modeling multiple linguistic phenomena in a joint framework (Finkel et al., 2006). But because these graphical models represent uncertainty for every element in the model, adding too many layers of latent variables makes them difficult to train.

In this paper, we present a hybrid architecture that combines a recurrent neural network language model with a latent variable model over shallow discourse structure. In this way, the model learns a discriminatively-trained distributed representation of the local contextual features that drive word choice at the intra-sentence level, using techniques that are now state-of-the-art in language modeling (Mikolov et al., 2010). However, the model treats shallow discourse structure — specifically, the relationships between pairs of adjacent sentences — as a latent variable. As a result, the model can act as both a discourse relation classifier and a language model. Specifically:

- If trained to maximize the conditional likelihood of the discourse relations, it outperforms state-of-the-art methods for both implicit discourse relation classification in the Penn Discourse Treebank (Rutherford and Xue, 2015) and dialog act classification in Switch-

board (Kalchbrenner and Blunsom, 2013). The model learns from both the discourse annotations as well as the language modeling objective, unlike previous recursive neural architectures that learn only from annotated discourse relations (Ji and Eisenstein, 2015).

- If the model is trained to maximize the joint likelihood of the discourse relations and the text, it is possible to marginalize over discourse relations at test time, outperforming language models that do not account for discourse structure.

In contrast to recent work on continuous latent variables in recurrent neural networks (Chung et al., 2015), which require complex variational autoencoders to represent uncertainty over the latent variables, our model is simple to implement and train, requiring only minimal modifications to existing recurrent neural network architectures that are implemented in commonly-used toolkits such as Theano, Torch, and CNN.

We focus on a class of *shallow discourse relations*, which hold between pairs of adjacent sentences (or utterances). These relations describe how the adjacent sentences are related: for example, they may be in CONTRAST, or the latter sentence may offer an answer to a question posed by the previous sentence. Shallow relations do not capture the full range of discourse phenomena (Webber et al., 2012), but they account for two well-known problems: implicit discourse relation classification in the Penn Discourse Treebank, which was the 2015 CoNLL shared task (Xue et al., 2015); and dialog act classification, which characterizes the structure of interpersonal communication in the Switchboard corpus (Stolcke et al., 2000), and is a key component of contemporary dialog systems (Williams and Young, 2007). Our model outperforms state-of-the-art alternatives for implicit discourse relation classification in the Penn Discourse Treebank, and for dialog act classification in the Switchboard corpus.

2 Background

Our model scaffolds on recurrent neural network (RNN) language models (Mikolov et al., 2010), and recent variants that exploit multiple levels of linguistic detail (Ji et al., 2015; Lin et al., 2015).

RNN Language Models Let us denote token n in a sentence t by $y_{t,n} \in \{1 \dots V\}$, and write $\mathbf{y}_t = \{y_{t,n}\}_{n \in \{1 \dots N_t\}}$ to indicate the sequence of words in sentence t . In an RNN language model, the probability of the sentence is decomposed as,

$$p(\mathbf{y}_t) = \prod_n^{N_t} p(y_{t,n} | \mathbf{y}_{t,<n}), \quad (1)$$

where the probability of each word $y_{t,n}$ is conditioned on the entire preceding sequence of words $\mathbf{y}_{t,<n}$ through the summary vector $\mathbf{h}_{t,n-1}$. This vector is computed recurrently from $\mathbf{h}_{t,n-2}$ and from the *embedding* of the current word, $\mathbf{X}_{y_{t,n-1}}$, where $\mathbf{X} \in \mathbb{R}^{K \times V}$ and K is the dimensionality of the word embeddings. The language model can then be summarized as,

$$\mathbf{h}_{t,n} = \mathbf{f}(\mathbf{X}_{y_{t,n}}, \mathbf{h}_{t,n-1}) \quad (2)$$

$$p(y_{t,n} | \mathbf{y}_{t,<n}) = \text{softmax}(\mathbf{W}_o \mathbf{h}_{t,n-1} + \mathbf{b}_o), \quad (3)$$

where the matrix $\mathbf{W}_o \in \mathbb{R}^{V \times K}$ defines the *output embeddings*, and $\mathbf{b}_o \in \mathbb{R}^V$ is an offset. The function $\mathbf{f}(\cdot)$ is a deterministic non-linear transition function. It typically takes an element-wise non-linear transformation (e.g., tanh) of a vector resulting from the sum of the word embedding and a linear transformation of the previous hidden state.

The model as described thus far is identical to the recurrent neural network language model (RNNLM) of Mikolov et al. (2010). In this paper, we replace the above simple hidden state units with the more complex Long Short-Term Memory units (Hochreiter and Schmidhuber, 1997), which have consistently been shown to yield much stronger performance in language modeling (Pham et al., 2014). For simplicity, we still use the term RNNLM in referring to this model.

Document Context Language Model One drawback of the RNNLM is that it cannot propagate long-range information between the sentences. Even if we remove sentence boundaries, long-range information will be attenuated by repeated application of the non-linear transition function. Ji et al. (2015) propose the Document Context Language Model (DCLM) to address this issue. The core idea is to represent context with *two* vectors: $\mathbf{h}_{t,n}$, representing *intra*-sentence word-level context, and \mathbf{c}_t , representing *inter*-sentence context. These two vectors




Figure 1: A fragment of our model with latent variable z_t , which only illustrates discourse information flow from sentence $(t-1)$ to t . The information from sentence $(t-1)$ affects the distribution of z_t and then the words prediction within sentence t .

$$p(y_{t,n+1} | z_t, \mathbf{y}_{t,<n}, \mathbf{y}_{t-1}) = g\left(\underbrace{\mathbf{W}_o^{(z_t)} \mathbf{h}_{t,n}}_{\substack{\text{relation-specific} \\ \text{intra-sentential context}}} + \underbrace{\mathbf{W}_c^{(z_t)} \mathbf{c}_{t-1}}_{\substack{\text{relation-specific} \\ \text{inter-sentential context}}} + \underbrace{\mathbf{b}_o^{(z_t)}}_{\text{bias}} \right) \quad (4)$$

Figure 2: Per-token generative probabilities in the discourse relation language model

are then linearly combined in the generation function for word $y_{t,n}$,

$$p(y_{t,n} | \mathbf{y}_{t,<n}, \mathbf{y}_{<t}) = \text{softmax}(\mathbf{W}_o \mathbf{h}_{t,n-1} + \mathbf{W}_c \mathbf{c}_{t-1} + \mathbf{b}_o), \quad (5)$$

where \mathbf{c}_{t-1} is set to the last hidden state of the previous sentence. Ji et al. (2015) show that this model can improve language model perplexity.

3 Discourse Relation Language Models

We now present a probabilistic neural model over sequences of words and shallow discourse relations. Discourse relations z_t are treated as latent variables, which are linked with a recurrent neural network over words in a *latent variable recurrent neural network* (Chung et al., 2015).

3.1 The Model

Our model (see Figure 1) is formulated as a two-step generative story. In the first step, context information from the sentence $(t-1)$ is used to generate the discourse relation between sentences $(t-1)$ and t ,

$$p(z_t | \mathbf{y}_{t-1}) = \text{softmax}(\mathbf{U} \mathbf{c}_{t-1} + \mathbf{b}), \quad (6)$$

where z_t is a random variable capturing the discourse relation between the two sentences, and \mathbf{c}_{t-1} is a vector summary of the contextual information from sentence $(t-1)$, just as in the DCLM (Equation 5). The model maintains a default context vector \mathbf{c}_0 for the first sentences of documents, and treats it as a parameter learned with other model parameters during training.

In the second step, the sentence \mathbf{y}_t is generated, conditioning on the preceding sentence \mathbf{y}_{t-1} and the discourse relation z_t :

$$p(\mathbf{y}_t | z_t, \mathbf{y}_{t-1}) = \prod_n^{N_t} p(y_{t,n} | \mathbf{y}_{t,<n}, \mathbf{y}_{t-1}, z_t), \quad (7)$$

The generative probability for the sentence \mathbf{y}_t decomposes across tokens as usual (Equation 7). The per-token probabilities are shown in Equation 4, in Figure 2. Discourse relations are incorporated by parameterizing the output matrices $\mathbf{W}_o^{(z_t)}$ and $\mathbf{W}_c^{(z_t)}$; depending on the discourse relation that holds between $(t-1)$ and t , these matrices will favor different parts of the embedding space. The bias term $\mathbf{b}_o^{(z_t)}$ is also parametrized by the discourse relation, so that each relation can favor specific words.

Overall, the joint probability of the text and discourse relations is,

$$p(\mathbf{y}_{1:T}, \mathbf{z}_{1:T}) = \prod_t^T p(z_t | \mathbf{y}_{t-1}) \times p(\mathbf{y}_t | z_t, \mathbf{y}_{t-1}). \quad (8)$$

If the discourse relations z_t are not observed, then our model is a form of latent variable recurrent neural network (LVRNN). Connections to recent work on LVRNNs are discussed in § 6; the key difference is that the latent variables here correspond to linguistically meaningful elements, which we may wish to predict or marginalize, depending on the situation.

Parameter Tying As proposed, the Discourse Relation Language Model has a large number of parameters. Let K , H and V be the input dimension,

hidden dimension and the size of vocabulary in language modeling. The size of each prediction matrix $\mathbf{W}_o^{(z)}$ and $\mathbf{W}_c^{(z)}$ is $V \times H$; there are two such matrices for each possible discourse relation. We reduce the number of parameters by factoring each of these matrices into two components:

$$\mathbf{W}_o^{(z)} = \mathbf{W}_o \cdot \mathbf{V}^{(z)}, \quad \mathbf{W}_c^{(z)} = \mathbf{W}_c \cdot \mathbf{M}^{(z)}, \quad (9)$$

where $\mathbf{V}^{(z)}$ and $\mathbf{M}^{(z)}$ are relation-specific components for intra-sentential and inter-sentential contexts; the size of these matrices is $H \times H$, with $H \ll V$. The larger $V \times H$ matrices \mathbf{W}_o and \mathbf{W}_c are shared across all relations.

3.2 Inference

There are two possible inference scenarios: inference over discourse relations, conditioning on words; and inference over words, marginalizing over discourse relations.

Inference over Discourse Relations The probability of discourse relations given the sentences $p(\mathbf{z}_{1:T} \mid \mathbf{y}_{1:T})$ is decomposed into the product of probabilities of individual discourse relations conditioned on the adjacent sentences $\prod_t p(z_t \mid \mathbf{y}_t, \mathbf{y}_{t-1})$. These probabilities are computed by Bayes' rule:

$$p(z_t \mid \mathbf{y}_t, \mathbf{y}_{t-1}) = \frac{p(\mathbf{y}_t \mid z_t, \mathbf{y}_{t-1}) \times p(z_t \mid \mathbf{y}_{t-1})}{\sum_{z'} p(\mathbf{y}_t \mid z', \mathbf{y}_{t-1}) \times p(z' \mid \mathbf{y}_{t-1})}. \quad (10)$$

The terms in each product are given in Equations 6 and 7. Normalizing involves only a sum over a small finite number of discourse relations. Note that inference is easy in our case because all words are observed and there is no probabilistic coupling of the discourse relations.

Inference over Words In discourse-informed language modeling, we marginalize over discourse relations to compute the probability of a sequence of sentence $\mathbf{y}_{1:T}$, which can be written as,

$$p(\mathbf{y}_{1:T}) = \prod_t \sum_{z_t} p(z_t \mid \mathbf{y}_{t-1}) \times p(\mathbf{y}_t \mid z_t, \mathbf{y}_{t-1}), \quad (11)$$

because the word sequences are observed, decoupling each z_t from its neighbors z_{t+1} and z_{t-1} . This decoupling ensures that we can compute the

overall marginal likelihood as a product over local marginals.

3.3 Learning

The model can be trained in two ways: to maximize the joint probability $p(\mathbf{y}_{1:T}, \mathbf{z}_{1:T})$, or to maximize the conditional probability $p(\mathbf{z}_{1:T} \mid \mathbf{y}_{1:T})$. The joint training objective is more suitable for language modeling scenarios, and the conditional objective is better for discourse relation prediction. We now describe each objective in detail.

Joint likelihood objective The joint likelihood objective function is directly adopted from the joint probability defined in Equation 8. The objective function for a single document with T sentences or utterances is,

$$\ell(\boldsymbol{\theta}) = \sum_t^T \log p(z_t \mid \mathbf{y}_{t-1}) + \sum_n^{N_t} \log p(y_{t,n} \mid \mathbf{y}_{t,<n}, \mathbf{y}_{t-1}, z_t), \quad (12)$$

where $\boldsymbol{\theta}$ represents the collection of all model parameters, including the parameters in the LSTM units and the word embeddings.

Maximizing the objective function $\ell(\boldsymbol{\theta})$ will jointly optimize the model on both language and discourse relation prediction. As such, it can be viewed as a form of multi-task learning (Caruana, 1997), where we learn a shared representation that works well for discourse relation prediction and for language modeling. However, in practice, the large vocabulary size and number of tokens means that the language modeling part of the objective function tends to dominate.

Conditional objective This training objective is specific to the discourse relation prediction task, and based on Equation 10 can be written as:

$$\begin{aligned} \ell_r(\boldsymbol{\theta}) = & \sum_t^T \log p(z_t \mid \mathbf{y}_{t-1}) + \log p(\mathbf{y}_t \mid z_t, \mathbf{y}_{t-1}) \\ & - \log \sum_{z'} p(z' \mid \mathbf{y}_{t-1}) \times p(\mathbf{y}_t \mid z', \mathbf{y}_{t-1}) \end{aligned} \quad (13)$$

The first line in Equation 13 is the same as $\ell(\boldsymbol{\theta})$, but the second line reflects the normalization over all

possible values of z_t . This forces the objective function to attend specifically to the problem of maximizing the conditional likelihood of the discourse relations and treat language modeling as an auxiliary task (Collobert et al., 2011).

3.4 Modeling limitations

The discourse relation language model is carefully designed to decouple the discourse relations from each other, after conditioning on the words. It is clear that text documents and spoken dialogues have sequential discourse structures, and it seems likely that modeling this structure could improve performance. In a traditional hidden Markov model (HMM) generative approach (Stolcke et al., 2000), modeling sequential dependencies is not difficult, because training reduces to relative frequency estimation. However, in the hybrid probabilistic-neural architecture proposed here, training is already expensive, due to the large number of parameters that must be estimated. Adding probabilistic couplings between adjacent discourse relations $\langle z_{t-1}, z_t \rangle$ would require the use of dynamic programming for both training and inference, increasing time complexity by a factor that is quadratic in the number of discourse relations. We did not attempt this in this paper; we do compare against a conventional HMM on the dialogue act prediction task in § 5.

Ji et al. (2015) propose an alternative form of the document context language model, in which the contextual information c_t impacts the hidden state h_{t+1} , rather than going directly to the outputs y_{t+1} . They obtain slightly better perplexity with this approach, which has fewer trainable parameters. However, this model would couple z_t with *all* subsequent sentences $y_{>t}$, making prediction and marginalization of discourse relations considerably more challenging. Sequential Monte Carlo algorithms offer a possible solution (de Freitas et al., ; Gu et al., 2015), which may be considered in future work.

4 Data and Implementation

We evaluate our model on two benchmark datasets: (1) the Penn Discourse Treebank (Prasad et al., 2008, PDTB), which is annotated on a corpus of Wall Street Journal articles; (2) the Switchboard di-

ologue act corpus (Stolcke et al., 2000, SWDA), which is annotated on a collections of phone conversations. Both corpora contain annotations of discourse relations and dialogue relations that hold between adjacent spans of text.

The Penn Discourse Treebank (PDTB) provides a low-level discourse annotation on written texts. In the PDTB, each discourse relation is annotated between two argument spans, $Arg1$ and $Arg2$. There are two types of relations: explicit and implicit. Explicit relations are signalled by discourse markers (e.g., “*however*”, “*moreover*”), and the span of $Arg1$ is almost totally unconstrained: it can range from a single clause to an entire paragraph, and need not be adjacent to either $Arg2$ nor the discourse marker. However, automatically classifying these relations is considered to be relatively easy, due to the constraints from the discourse marker itself (Pitler et al., 2008). In addition, explicit relations are difficult to incorporate into language models which must generate each word exactly once. On the contrary, *implicit* discourse relations are annotated only between adjacent sentences, based on a semantic understanding of the discourse arguments. Automatically classifying these discourse relations is a challenging task (Lin et al., 2009; Pitler et al., 2009; Rutherford and Xue, 2015; Ji and Eisenstein, 2015). We therefore focus on implicit discourse relations, leaving to the future work the question of how to apply our modeling framework to explicit discourse relations. During training, we collapse all relation types other than implicit (explicit, ENTREL, and NOREL) into a single dummy relation type, which holds between all adjacent sentence pairs that do not share an implicit relation.

As in the prior work on first-level discourse relation identification (e.g., Park and Cardie, 2012), we use sections 2-20 of the PDTB as the training set, sections 0-1 as the development set for parameter tuning, and sections 21-22 for testing. For pre-processing, we lower-cased all tokens, and substituted all numbers with a special token “*NUM*”. To build the vocabulary, we kept the 10,000 most frequent words from the training set, and replaced low-frequency words with a special token “*UNK*”. In prior work that focuses on detecting individual relations, balanced training sets are constructed so that

there are an equal number of instances with and without each relation type (Park and Cardie, ; Biran and McKeown, 2013; Rutherford and Xue, 2014). In this paper, we target the more challenging multi-way classification problem, so this strategy is not applicable; in any case, since our method deals with entire documents, it is not possible to balance the training set in this way.

The Switchboard Dialog Act Corpus (SWDA) is annotated on the Switchboard Corpus of human-human conversational telephone speech (Godfrey et al., 1992). The annotations label each utterance with one of 42 possible speech acts, such as AGREE, HEDGE, and WH-QUESTION. Because these speech acts form the structure of the dialogue, most of them pertain to both the preceding and succeeding utterances (e.g., AGREE). The SWDA corpus includes 1155 five-minute conversations. We adopted the standard split from Stolcke et al. (2000), using 1,115 conversations for training and nineteen conversations for test. For parameter tuning, we randomly select nineteen conversations from the training set as the development set. After parameter tuning, we train the model on the full training set with the selected configuration. We use the same preprocessing techniques here as in the PDTB.

4.1 Implementation

We use a single-layer LSTM to build the recurrent architecture of our models, which we implement in the CNN package.¹ Our implementation is available on <https://github.com/jiyfeng/drlm>. Some additional details follow.

Initialization Following prior work on RNN initialization (Bengio, 2012), all parameters except the relation prediction parameters \mathbf{U} and \mathbf{b} are initialized with random values drawn from the range $[-\sqrt{6/(d_1 + d_2)}, \sqrt{6/(d_1 + d_2)}]$, where d_1 and d_2 are the input and output dimensions of the parameter matrix respectively. The matrix \mathbf{U} is initialized with random numbers from $[-10^{-5}, 10^{-5}]$ and \mathbf{b} is initialized to $\mathbf{0}$.

Learning Online learning was performed using AdaGrad (Duchi et al., 2011) with initial learning

rate $\lambda = 0.1$. To avoid the exploding gradient problem, we used norm clipping trick with a threshold of $\tau = 5.0$ (Pascanu et al., 2012). In addition, we used value dropout (Srivastava et al., 2014) with rate 0.5, on the input \mathbf{X} , context vector \mathbf{c} and hidden state \mathbf{h} , similar to the architecture proposed by Pham et al. (2014). The training procedure is monitored by the performance on the development set. In our experiments, 4 to 5 epochs were enough.

Hyper-parameters Our model includes two tunable hyper-parameters: the dimension of word representation K , the hidden dimension of LSTM unit H . We consider the values $\{32, 48, 64, 96, 128\}$ for both K and H . For each corpus in experiments, the best combination of K and H is selected via grid search on the development set.

5 Experiments

Our main evaluation is discourse relation prediction, using the PDTB and SWDA corpora. We also evaluate on language modeling, to determine whether incorporating discourse annotations at training time and then marginalizing them at test time can improve performance.

5.1 Implicit discourse relation prediction on the PDTB

We first evaluate our model with implicit discourse relation prediction on the PDTB dataset. Most of the prior work on first-level discourse relation prediction focuses on the “one-versus-all” binary classification setting, but we attack the more general four-way classification problem, as performed by Rutherford and Xue (2015). We compare against the following methods:

Rutherford and Xue (2015) build a set of feature-rich classifiers on the PDTB, and then augment these classifiers with additional automatically-labeled training instances. We compare against their published results, which are state-of-the-art.

Ji and Eisenstein (2015) employ a recursive neural network architecture. Their experimental setting is different, so we re-run their system using the same setting as described in § 4.

¹<https://github.com/clab/cnn>

Model	Accuracy	Macro F_1
<i>Baseline</i>		
1. Most common class	54.7	—
<i>Prior work</i>		
2. (Rutherford and Xue, 2015)	55.0	38.4
3. (Rutherford and Xue, 2015) with extra training data	57.1	40.5
4. (Ji and Eisenstein, 2015)	56.4	40.0
<i>Our work - DRLM</i>		
5. Joint training	57.1	40.5
6. Conditional training	59.5*	42.3

* significantly better than lines 2 and 4 with $p < 0.05$

Table 1: Multiclass relation identification on the first-level PDTB relations.

Results As shown in Table 1, the conditionally-trained discourse relation language models (DRLM) outperforms all alternatives, on both metrics. While the jointly-trained DRLM is at the same level as the previous state-of-the-art, conditional training on the same model provides a significant additional advantage, indicated by a binomial test.

5.2 Dialogue Act tagging

Dialogue act tagging has been widely studied in both NLP and speech communities. We follow the setup used by Stolcke et al. (2000) to conduct experiments, and adopt the following systems for comparison:

Stolcke et al. (2000) employ a hidden Markov model, with each HMM state corresponding to a dialogue act.

Kalchbrenner and Blunsom (2013) employ a complex neural architecture, with a convolutional network at each utterance and a recurrent network over the length of the dialog. To our knowledge, this model attains state-of-the-art accuracy on this task, outperforming other prior work such as (Webb et al., 2005; Milajevs and Purver, 2014).

Results As shown in Table 2, the conditionally-trained discourse relation language model (DRLM) outperforms all competitive systems on this task. A binomial test shows the result in line 6 is significantly better than the previous state-of-the-art (line 4). All comparisons are against published results, and Macro- F_1 scores are not available. Accuracy

1. Model	Accuracy
<i>Baseline</i>	
2. Most common class	31.5
<i>Prior work</i>	
3. (Stolcke et al., 2000)	71.0
4. (Kalchbrenner and Blunsom, 2013)	73.9
<i>Our work - DRLM</i>	
5. Joint training	74.0
6. Conditional training	77.0*

* significantly better than line 4 with $p < 0.01$

Table 2: The results of dialogue act tagging.

is more reliable on this evaluation, since no single class dominates, unlike the PDTB task.

5.3 Discourse-aware language modeling

As a joint model for discourse and language modeling, DRLM can also function as a language model, assigning probabilities to sequences of words while marginalizing over discourse relations. To determine whether discourse-aware language modeling can improve performance, we compare against the following systems:

RNNLM+LSTM This is the same basic architecture as the RNNLM proposed by (Mikolov et al., 2010), which was shown to outperform a Kneser-Ney smoothed 5-gram model on modeling Wall Street Journal text. Following Pham et al. (2014), we replace the Sigmoid nonlinearity with a long short-term memory (LSTM).

DCLM We compare against the Document Context Language Model (DCLM) of Ji et al. (2015). We use the “context-to-output” variant, which is identical to the current modeling approach, except that it is not parametrized by discourse relations. This model achieves strong results on language modeling for small and medium-sized corpora, outperforming RNNLM+LSTM.

Results The perplexities of language modeling on the PDTB and the SWDA are summarized in Table 3. The comparison between line 1 and line 2 shows the benefit of considering multi-sentence context information on language modeling. Line 3 shows that adding discourse relation information

Model	PDTB			SWDA		
	K	H	PPLX	K	H	PPLX
<i>Baseline</i>						
1. RNNLM	96	128	117.8	128	96	56.0
2. DCLM	96	96	112.2	96	96	45.3
<i>Our work</i>						
3. DRLM	64	96	108.3	128	64	39.6

Table 3: Language model perplexities (PPLX), lower is better. The model dimensions K and H that gave best performance on the dev set are also shown.

yields further improvements for both datasets. We emphasize that discourse relations in the test documents are marginalized out, so no annotations are required for the test set; the improvements are due to the disambiguating power of discourse relations in the training set.

Because our training procedure requires discourse annotations, this approach does not scale to the large datasets typically used in language modeling. As a consequence, the results obtained here are somewhat academic, from the perspective of practical language modeling. Nonetheless, the positive results here motivate the investigation of training procedures that are also capable of marginalizing over discourse relations at training time.

6 Related Work

This paper draws on previous work in both discourse modeling and language modeling.

Discourse and dialog modeling Early work on discourse relation classification utilizes rich, hand-crafted feature sets (Joty et al., 2012; Lin et al., 2009; Sagae, 2009). Recent representation learning approaches attempt to learn good representations jointly with discourse relation classifiers and discourse parsers (Ji and Eisenstein, 2014; Li et al., 2014). Of particular relevance are applications of neural architectures to PDTB implicit discourse relation classification (Ji and Eisenstein, 2015; Zhang et al., 2015; Braud and Denis, 2015). All of these approaches are essentially classifiers, and take supervision only from the 16,000 annotated discourse relations in the PDTB training set. In contrast, our approach is a probabilistic model over the entire text.

Probabilistic models are frequently used in dia-

log act tagging, where hidden Markov models have been a dominant approach (Stolcke et al., 2000). In this work, the emission distribution is an n -gram language model for each dialogue act; we use a conditionally-trained recurrent neural network language model. An alternative neural approach for dialogue act tagging is the combined convolutional-recurrent architecture of Kalchbrenner and Blunsom (2013). Our modeling framework is simpler, relying on a latent variable parametrization of a purely recurrent architecture.

Language modeling There are an increasing number of attempts to incorporate document-level context information into language modeling. For example, Mikolov and Zweig (2012) introduce LDA-style topics into RNN based language modeling. Sordani et al. (2015) use a convolutional structure to summarize the context from previous two utterances as context vector for RNN based language modeling. Our models in this paper provide a unified framework to model the context and current sentence. Wang and Cho (2015) and Lin et al. (2015) construct bag-of-words representations of previous sentences, which are then used to inform the RNN language model that generates the current sentence. The most relevant work is the Document Context Language Model (Ji et al., 2015, DCLM); we describe the connection to this model in § 2. By adding discourse information as a latent variable, we attain better perplexity on held-out data.

Latent variable neural networks Introducing latent variables to a neural network model increases its representational capacity, which is the main goal of prior efforts in this space (Kingma and Welling, 2014; Chung et al., 2015). From this perspective, our model with discourse relations as latent variables shares the same merit. Unlike this prior work, in our approach, the latent variables carry a linguistic interpretation, and are at least partially observed. Also, these prior models employ continuous latent variables, requiring complex inference techniques such as variational autoencoders (Kingma and Welling, 2014; Burda et al., 2016; Chung et al., 2015). In contrast, the discrete latent variables in our model are easy to sum and maximize over.

7 Conclusion

We have presented a probabilistic neural model over sequences of words and shallow discourse relations between adjacent sequences. This model combines positive aspects of neural network architectures with probabilistic graphical models: it can learn discriminatively-trained vector representations, while maintaining a probabilistic representation of the targeted linguistic element: in this case, shallow discourse relations. This method outperforms state-of-the-art systems in two discourse relation detection tasks, and can also be applied as a language model, marginalizing over discourse relations on the test data. Future work will investigate the possibility of learning from partially-labeled training data, which would have at least two potential advantages. First, it would enable the model to scale up to the large datasets needed for competitive language modeling. Second, by training on more data, the resulting vector representations might support even more accurate discourse relation prediction.

Acknowledgments

Thanks to Trevor Cohn, Chris Dyer, Lingpeng Kong, and Quoc V. Le for helpful discussions, and to the anonymous reviewers for their feedback. This work was supported by a Google Faculty Research award to the third author. It was partially performed during the 2015 Jelinek Memorial Summer Workshop on Speech and Language Technologies at the University of Washington, Seattle, and was supported by Johns Hopkins University via NSF Grant No IIS 1005411, DARPA LORELEI Contract No HR0011-15-2-0027, and gifts from Google, Microsoft Research, Amazon and Mitsubishi Electric Research Laboratory.

References

Yoshua Bengio. 2012. Practical recommendations for gradient-based training of deep architectures. In *Neural Networks: Tricks of the Trade*, pages 437–478. Springer.

Or Biran and Kathleen McKeown. 2013. In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 69–73, Sophia, Bulgaria.

Chloé Braud and Pascal Denis. 2015. Comparing word representations for implicit discourse relation classi-

fication. In *Proceedings of Empirical Methods for Natural Language Processing (EMNLP)*, pages 2201–2211, Lisbon, September.

Yuri Burda, Roger Grosse, and Ruslan Salakhutdinov. 2016. Importance weighted autoencoders. In *Proceedings of the International Conference on Learning Representations (ICLR)*.

Rich Caruana. 1997. Multitask learning. *Machine learning*, 28(1):41–75.

Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron Courville, and Yoshua Bengio. 2015. A recurrent latent variable model for sequential data. In *Neural Information Processing Systems (NIPS)*, Montréal.

R. Collobert, J. Weston, L. Bottou, M. Karlen, K. Kavukcuoglu, and P. Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research*, 12:2493–2537.

João FG de Freitas, Mahesan Niranjan, Andrew H. Gee, and Arnaud Doucet. Sequential monte carlo methods to train neural network models. *Neural computation*, 12(4):955–993.

John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *The Journal of Machine Learning Research*, 12:2121–2159.

Chris Dyer, Miguel Ballesteros, Wang Ling, Austin Matthews, and Noah A. Smith. 2015. Transition-based dependency parsing with stack long short-term memory. In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 334–343, Beijing, China.

Jenny Rose Finkel, Christopher D Manning, and Andrew Y Ng. 2006. Solving the problem of cascading errors: Approximate bayesian inference for linguistic annotation pipelines. In *Proceedings of Empirical Methods for Natural Language Processing (EMNLP)*, pages 618–626.

John J Godfrey, Edward C Holliman, and Jane McDaniel. 1992. Switchboard: Telephone speech corpus for research and development. In *ICASSP*, volume 1, pages 517–520. IEEE.

Shixiang Gu, Zoubin Ghahramani, and Richard E Turner. 2015. Neural adaptive sequential monte carlo. In *Neural Information Processing Systems (NIPS)*, Montréal.

Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.

Yangfeng Ji and Jacob Eisenstein. 2014. Representation learning for text-level discourse parsing. In *Proceedings of the Association for Computational Linguistics (ACL)*, Baltimore, MD.

- Yangfeng Ji and Jacob Eisenstein. 2015. One vector is not enough: Entity-augmented distributional semantics for discourse relations. *Transactions of the Association for Computational Linguistics (ACL)*, June.
- Yangfeng Ji, Trevor Cohn, Lingpeng Kong, Chris Dyer, and Jacob Eisenstein. 2015. Document context language models. In *International Conference on Learning Representations, Poster Paper*, volume abs/1511.03962.
- Shafiq Joty, Giuseppe Carenini, and Raymond Ng. 2012. A novel discriminative framework for sentence-level discourse analysis. In *Proceedings of Empirical Methods for Natural Language Processing (EMNLP)*.
- Nal Kalchbrenner and Phil Blunsom. 2013. Recurrent convolutional neural networks for discourse compositionality. In *Proceedings of the Workshop on Continuous Vector Space Models and their Compositionality*, pages 119–126, Sofia, Bulgaria, August. Association for Computational Linguistics.
- Diederik P Kingma and Max Welling. 2014. Auto-encoding variational bayes. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Quoc Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- Jiwei Li, Rumeng Li, and Eduard Hovy. 2014. Recursive deep models for discourse parsing. In *Proceedings of Empirical Methods for Natural Language Processing (EMNLP)*.
- Ziheng Lin, Min-Yen Kan, and Hwee Tou Ng. 2009. Recognizing implicit discourse relations in the penn discourse treebank. In *Proceedings of Empirical Methods for Natural Language Processing (EMNLP)*, pages 343–351, Singapore.
- Rui Lin, Shujie Liu, Muyun Yang, Mu Li, Ming Zhou, and Sheng Li. 2015. Hierarchical recurrent neural network for document modeling. In *Proceedings of Empirical Methods for Natural Language Processing (EMNLP)*, pages 899–907, Lisbon, September.
- Christopher D. Manning. 2016. Computational linguistics and deep learning. *Computational Linguistics*, 41(4).
- Tomas Mikolov and Geoffrey Zweig. 2012. Context dependent recurrent neural network language model. In *Proceedings of Spoken Language Technology (SLT)*, pages 234–239.
- Tomas Mikolov, Martin Karafiát, Lukas Burget, Jan Cernocký, and Sanjeev Khudanpur. 2010. Recurrent neural network based language model. In *INTER-SPEECH*, pages 1045–1048.
- Dmitrijs Milajevs and Matthew Purver. 2014. Investigating the contribution of distributional semantic information for dialogue act classification. In *Proceedings of the 2nd Workshop on Continuous Vector Space Models and their Compositionality (CVSC)*, pages 40–47.
- Joonsuk Park and Claire Cardie. Improving implicit discourse relation recognition through feature set optimization. In *Proceedings of the 13th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 108–112, Seoul, South Korea, July. Association for Computational Linguistics.
- Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. 2012. On the difficulty of training recurrent neural networks. *arXiv preprint arXiv:1211.5063*.
- Vu Pham, Théodore Bluche, Christopher Kermorvant, and Jérôme Louradour. 2014. Dropout improves recurrent neural networks for handwriting recognition. In *Frontiers in Handwriting Recognition (ICFHR), 2014 14th International Conference on*, pages 285–290. IEEE.
- Emily Pitler, Mridhula Raghupathy, Hena Mehta, Ani Nenkova, Alan Lee, and Aravind Joshi. 2008. Easily identifiable discourse relations. In *Proceedings of the International Conference on Computational Linguistics (COLING)*, pages 87–90, Manchester, UK.
- Emily Pitler, Annie Louis, and Ani Nenkova. 2009. Automatic sense prediction for implicit discourse relations in text. In *Proceedings of the Association for Computational Linguistics (ACL)*, Suntec, Singapore.
- Rashmi Prasad, Nikhil Dinesh, Alan Lee, Eleni Milt-sakaki, Livio Robaldo, Aravind Joshi, and Bonnie Webber. 2008. The Penn Discourse Treebank 2.0. In *Proceedings of LREC*.
- Attapol T Rutherford and Nianwen Xue. 2014. Discovering implicit discourse relations through brown cluster pair representation and coreference patterns. In *Proceedings of the European Chapter of the Association for Computational Linguistics (EACL)*.
- Attapol Rutherford and Nianwen Xue. 2015. Improving the inference of implicit discourse relations via classifying explicit discourse connectives. pages 799–808, May–June.
- Kenji Sagae. 2009. Analysis of discourse structure with syntactic dependencies and data-driven shift-reduce parsing. In *Proceedings of the 11th International Conference on Parsing Technologies (IWPT'09)*, pages 81–84, Paris, France, October. Association for Computational Linguistics.
- Richard Socher, Alex Perelygin, Jean Y Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank.

- In *Proceedings of Empirical Methods for Natural Language Processing (EMNLP)*, Seattle, WA.
- Alessandro Sordani, Michel Galley, Michael Auli, Chris Brockett, Yangfeng Ji, Meg Mitchell, Jian-Yun Nie, Jianfeng Gao, and Bill Dolan. 2015. A neural network approach to context-sensitive generation of conversational responses. In *Proceedings of the North American Chapter of the Association for Computational Linguistics (NAACL)*, Denver, CO, May.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958.
- Andreas Stolcke, Klaus Ries, Noah Coccaro, Elizabeth Shriberg, Rebecca Bates, Daniel Jurafsky, Paul Taylor, Rachel Martin, Carol Van Ess-Dykema, and Marie Meteer. 2000. Dialogue act modeling for automatic tagging and recognition of conversational speech. *Computational linguistics*, 26(3):339–373.
- Joseph Turian, Lev Ratinov, and Yoshua Bengio. 2010. Word representations: A simple and general method for semi-supervised learning. pages 384–394.
- Tian Wang and Kyunghyun Cho. 2015. Larger-context language modelling. *arXiv preprint arXiv:1511.03729*.
- Nick Webb, Mark Hepple, and Yorick Wilks. 2005. Dialogue act classification based on intra-utterance features. In *Proceedings of the AAAI Workshop on Spoken Language Understanding*.
- Bonnie Webber, Markus Egg, and Valia Kordoni. 2012. Discourse structure and language technology. *Journal of Natural Language Engineering*, 1.
- Jason D Williams and Steve Young. 2007. Partially observable markov decision processes for spoken dialog systems. *Computer Speech & Language*, 21(2):393–422.
- Nianwen Xue, Hwee Tou Ng, Sameer Pradhan, Rashmi Prasad, Christopher Bryant, and Attapol T Rutherford. 2015. The CoNLL-2015 shared task on shallow discourse parsing. In *Proceedings of the Conference on Natural Language Learning (CoNLL)*.
- Biao Zhang, Jinsong Su, Deyi Xiong, Yaojie Lu, Hong Duan, and Junfeng Yao. 2015. Shallow convolutional neural network for implicit discourse relation recognition. In *Proceedings of Empirical Methods for Natural Language Processing (EMNLP)*, pages 2230–2235, Lisbon, September.