

# A Cross-Domain Transferable Neural Coherence Model

Peng Xu<sup>1</sup>, Hamidreza Saghir<sup>1</sup>, Jin Sung Kang<sup>1</sup>, Teng Long<sup>1</sup>, Avishek Joey Bose<sup>\*1,2</sup>,  
Yanshuai Cao<sup>1</sup>, Jackie Chi Kit Cheung<sup>1,2,3</sup>

<sup>1</sup>Borealis AI

{peng.z.xu, hamidreza.saghir, jinsung.kang, leo.long, yanshuai.cao}@borealisai.com

<sup>2</sup>McGill University

{joey.bose, jcheung}@cs.mcgill.ca

<sup>3</sup>Canada CIFAR Chair, Mila

## Abstract

Coherence is an important aspect of text quality and is crucial for ensuring its readability. One important limitation of existing coherence models is that training on one domain does not easily generalize to unseen categories of text. Previous work (Li and Jurafsky, 2017) advocates for generative models for cross-domain generalization, because for discriminative models, the space of incoherent sentence orderings to discriminate against during training is prohibitively large. In this work, we propose a local discriminative neural model with a much smaller negative sampling space that can efficiently learn against incorrect orderings. The proposed coherence model is simple in structure, yet it significantly outperforms previous state-of-art methods on a standard benchmark dataset on the Wall Street Journal corpus, as well as in multiple new challenging settings of transfer to unseen categories of discourse on Wikipedia articles.

## 1 Introduction

Coherence is a discourse property that is concerned with the logical and semantic organization of a passage, such that the overall meaning of the passage is expressed fluidly and clearly. It is an important quality measure for text generated by humans or machines, and modelling coherence can benefit many applications, including summarization, question answering (Verberne et al., 2007), essay scoring (Miltsakaki and Kukich, 2004; Burstein et al., 2010) and text generation (Park and Kim, 2015; Kiddon et al., 2016; Holtzman et al., 2018).

The ability to generalize to new domains of text is desirable for NLP models in general. Besides the practical reason of avoiding costly retraining

on every new domain, for coherence modelling, we would also like our model to make decisions based on the semantic relationships between sentences, rather than simply overfit to the structural cues of a specific domain.

The standard task used to test a coherence model in NLP is sentence ordering, for example, to distinguish between a coherently ordered list of sentences and a random permutation thereof. Earlier work focused on feature engineering, drawing on theories such as Centering Theory (Grosz et al., 1995) and Rhetorical Structure Theory (Thompson and Mann, 1987) to propose features based on local entity and lexical transitions, as well as more global concerns regarding topic transitions (Elsner et al., 2007). With the popularization of deep learning, the focus has shifted towards specifying model architectures, including a number of recent models that rely on distributed word representations used in a deep neural network (Li and Jurafsky, 2017; Nguyen and Joty, 2017; Logeswaran et al., 2018).

One key decision which forms the foundation of a model is whether it is discriminative or generative. Discriminative models depend on contrastive learning; they use automatic corruption methods to generate incoherent passages of text, then learn to distinguish coherent passages from incoherent ones. By contrast, generative approaches aim at maximising the likelihood of the training text, which is assumed to be coherent, without seeing incoherent text or explicitly incorporating coherence into the optimization objective.

It has been argued that neural discriminative models of coherence are prone to overfitting on the particular dataset and domain that they are designed for (Li and Jurafsky, 2017), possibly due to the expressive nature of functions learnable by a neural network. Another potential problem for discriminative models raised by Li and Jurafsky is

\*Work done while the author was an intern at Borealis AI.

that there are  $n!$  possible sentence orderings for a passage with  $n$  sentences, thus the sampled negative instances can only cover a tiny proportion of this space, limiting the performance of such models. There is thus an apparent association between discriminative models and high performance on a specific narrow domain.

We argue in this paper that there is, in fact, nothing inherent about discriminative models that cause previous systems to be brittle to domain changes. We demonstrate a solution to the above problems by combining aspects of previous generative and discriminative models to produce a system that works well in both in-domain and cross-domain settings, despite being a discriminative model overall.

Our method relies on two key ideas. The first is to reexamine the operating assumption of previous work that a global, passage-level model is necessary for good performance. While it is true that coherence is a property of a passage as a whole, capturing long-term dependencies in sequences remains a fundamental challenge when training neural networks in practice (Trinh et al., 2018). On the other hand, it is *plausible* that much of global coherence can be decomposed into a series of local decisions, as demonstrated by foundational theories such as Centering Theory. Our hypothesis in this work is that there remains much to be learned about local coherence cues which previous work has not fully captured and that these cues make up an essential part of global coherence. We demonstrate that such is the case.

Our model thus take neighbouring pairs of sentences as inputs, for which the space of negatives is much smaller and can therefore be effectively covered by sampling other sentences in the same document. Surprisingly, adequately modelling local coherence alone significantly outperforms previous approaches, and furthermore, local coherence captures text properties that are domain-agnostic, generalizing much better in open-domain settings to unseen categories of text.

Our second insight is that the superiority of previous generative approaches in cross-domain settings can be effectively incorporated into a discriminative model as a pre-training step. We show that generatively pre-trained sentence encoders enhance the performance of our discriminative local coherence model.

We demonstrate the effectiveness of our ap-

proach on the Wall Street Journal (WSJ) benchmark dataset, as well as on three challenging new evaluation protocols using different categories of articles drawn from Wikipedia that contain increasing levels of domain diversity. We show that our discriminative model significantly outperforms strong baselines on all datasets tested. Finally, hypothesis testing shows that the coherence scores from our model have a significant statistical association with the “rewrite” flag for regular length Wikipedia articles, demonstrating that our model prediction aligns with human judgement of text quality.

To summarize, our contributions are:

1. We correct the misconception that discriminative models cannot generalize well for cross-domain coherence scoring, with a novel local discriminative neural model.
2. We propose a set of cross-domain coherence datasets with increasingly difficult evaluation protocols.
3. Our new method outperforms previous methods by a significant margin on both the previous closed domain WSJ dataset as well as on all open-domain ones, setting the new state-of-the-art for coherence modelling.
4. Even with the simplest sentence encoder, averaged GloVe, our method frequently outperforms previous methods, while it can gain further accuracy by using stronger encoders.

## 2 Related Work

Barzilay and Lapata (2008) introduced the entity grid representation of a document, which uses the local syntactic transitions of entity mentions to model discourse coherence. Three tasks for evaluation were introduced for evaluation: discrimination, summary coherence rating, and readability assessment. Many models were proposed to extend this model (Eisner and Charniak, 2011; Feng and Hirst, 2012; Guinaudeau and Strube, 2013), including models relying on HMMs (Louis and Nenkova, 2012) to model document structure.

Driven by the success of deep neural networks, many neural models were proposed in the past few years. Li and Hovy (2014) proposed a neural clique-based discriminative model to compute the coherence score of a document by estimating a coherence probability for each clique of  $L$  sentences. Nguyen and Joty (2017) proposed a neural entity

grid model with convolutional neural network that operates over the entity grid representation. Mohiuddin et al. (2018) extended this model for written asynchronous conversations. Both methods rely on hand-crafted features derived from NLP pre-processing tools to enhance the original entity grid representation. We take a different approach to feature engineering in our work, focusing on the effect of supervised or unsupervised pre-training.

Li and Jurafsky (2017) was the first work to use generative models to model coherence and proposed to evaluate the performance of coherence models in an open-domain setting. Most recently, Logeswaran et al. (2018) used an RNN based encoder-decoder architecture to model the coherence which can also be treated as the generative model. One obvious disadvantage of generative models is that they maximize the likelihood of training text but never see the incoherent text. In other words, to produce a binary classification decision about coherence, such a generative model only sees data from one class. As we will demonstrate later in the experiments, this puts generative models at a disadvantage comparing to our local discriminative model.

### 3 Background: Generative Coherence Models

To understand the advantages of our local discriminative model, we first introduce the previous global passage-level generative coherence models. We will use “passage” and “document” interchangeably in this work, since all the models under consideration work in the same way for a full document or a passage in document.

Generative models are based on the idea that in a coherent passage, subsequent sentences should be predictable given their preceding sentences, and vice versa. Let us denote the corpus as  $\mathcal{C} = \{d_k\}_{k=1}^N$ , which consists of  $N$  documents, with each document  $d_k$  comprised of a sequence of sentences  $\{s_i\}$ . Formally, generative coherence models are trained using a log-likelihood objective as follows (with some variations according to the specific model):

$$\max_{\theta} \sum_{d \in \mathcal{C}} \sum_{s \in d} \log p(s|c_s; \theta), \quad (1)$$

where  $c_s$  is the context of the sentence  $s$  and  $\theta$  represents the model parameters.  $c_s$  can be chosen as the next or previous sentence (Li and Juraf-

sky, 2017), or all previous sentences (Logeswaran et al., 2018).

There are two hidden assumptions behind this maximum likelihood approach to coherence. First, it assumes that conditional log likelihood is a good proxy for coherence. Second, it assumes that training can well capture the long-range dependencies implied by the generative model.

Conditional log likelihood essentially measures the compressibility of a sentence given the context; i.e., how predictable  $s$  is given  $c_s$ . However, although incoherent next sentence is generally not predictable given the context, the inverse is not necessarily true. In other words, a coherent sentence does not need to have high conditional log-likelihood, as log likelihood can also be influenced by other factors such as fluency, grammaticality, sentence length, and the frequency of words in a sentence. Second, capturing long-range dependencies in neural sequence models is still an active area of research with many challenges (Trinh et al., 2018), hence there is no guarantee that maximum likelihood learning can faithfully capture the inductive bias behind the first assumption.

### 4 Our Local Discriminative Model

We propose the local coherence discriminator model (LCD) whose operating assumption is that the global coherence of a document can be well approximated by the average of coherence scores between consecutive pairs of sentences. Our experimental results later will validate the appropriateness of this assumption. For now, this simplification allows us to cast the learning problem as discriminating consecutive sentence pairs  $(s_i, s_{i+1})$  in the training documents (assumed to be coherent) from incoherent ones  $(s_i, s')$  (negative pairs to be constructed).

**Training objective:** Formally, our discriminative model  $f_{\theta}(\cdot, \cdot)$  takes a sentence pair and returns a score. The higher the score, the more coherent the input pair. Then our training objective is:

$$\mathcal{L}(\theta) = \sum_{d \in \mathcal{C}} \sum_{s_i \in d} \mathbb{E}_{p(s'|s_i)} [L(f_{\theta}(s_i, s_{i+1}), f_{\theta}(s_i, s'))] \quad (2)$$

where  $\mathbb{E}_{p(s'|s_i)}$  denotes expectation with respect to negative sampling distribution  $p$  which could be conditioned on  $s_i$ ; and  $L(\cdot, \cdot)$  is a loss function that takes two scores, one for a positive pair and one for a negative sentence pair.

**Loss function:** The role of the loss function is to encourage  $f^+ = f_\theta(s_i, s_{i+1})$  to be high while  $f^- = f_\theta(s_i, s')$  to be low. Common losses such as margin or log loss can all be used. Through experimental validation, we found that margin loss to be superior for this problem. Specifically,  $L$  takes on the form:  $L(f^+, f^-) = \max(0, \eta - f^+ + f^-)$  where  $\eta$  is the margin hyperparameter.

**Negative samples:** Technically, we are free to choose any sentence  $s'$  to form a negative pair with  $s_i$ . However, because of potential differences in genre, topic and writing style, such negatives might cause the discriminative model to learn cues unrelated to coherence. Therefore, we only select sentences from the same document to construct negative pairs. Specifically, suppose  $s_i$  comes from document  $d_k$  with length  $n_k$ , then  $p(s'|s_i)$  is a uniform distribution over the  $n_k - 1$  sentences  $\{s_j\}_{j \neq i}$  from  $d_k$ . For a document with  $n$  sentences, there are  $n - 1$  positive pairs, and  $(n - 1) * (n - 2) / 2$  negative pairs. It turns out that the quadratic number of negatives provides a rich enough learning signal, while at the same time, is not too prohibitively large to be effectively covered by a sampling procedure. In practice, we sample a new set of negatives each time we see a document, hence after many epochs, we can effectively cover the space for even very long documents. Section 5.7 discusses further details on sampling.

#### 4.1 Model Architecture

The specific neural architecture that we use for  $f_\theta$  is illustrated in Figure 1. We assume the use of some pre-trained sentence encoder, which is discussed in the next section.

Given an input sentence pair, the sentence encoder maps the sentences to real-valued vectors  $S$  and  $T$ . We then compute the concatenation of the following features: (1) concatenation of the two vectors  $(S, T)$ ; (2) element-wise difference  $S - T$ ; (3) element-wise product  $S * T$ ; (4) absolute value of element-wise difference  $|S - T|$ . The concatenated feature representation is then fed to a one-layer MLP to output the coherence score.

In practice, we make our overall coherence model bidirectional, by training a forward model with input  $(S, T)$  and a backward model with input  $(T, S)$  with the same architecture but separate parameters. The coherence score is then the average from the two models.

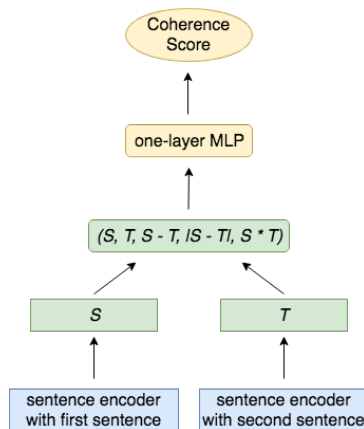


Figure 1: Generic architecture for our proposed model.

#### 4.2 Pre-trained Generative Model as the Sentence Encoder

Our model can work with any pre-trained sentence encoder, ranging from the most simplistic average GloVe (Pennington et al., 2014) embeddings to more sophisticated supervised or unsupervised pre-trained sentence encoders (Conneau et al., 2017). As mentioned in the introduction, since generative models can often be turned into sentence encoder, generative coherence model can be leveraged by our model to benefit from the advantages of both generative and discriminative training, similar to (Kiros et al., 2015; Peters et al., 2018). After initialization, we freeze the generative model parameters to avoid overfitting.

In Section 5, we will experimentally show that while we do benefit from strong pre-trained encoders, the fact that our local discriminative model improves over previous methods is independent of the choice of sentence encoder.

### 5 Experiments

#### 5.1 Evaluation Tasks

Following Nguyen and Joty (2017) and other previous work, we evaluate our models on the discrimination and insertion tasks. Additionally, we evaluate on the paragraph reconstruction task in open-domain settings, in a similar manner to Li and Jurafsky (2017).

In the *discrimination* task, a document is compared to a random permutation of its sentences, and the model is considered correct if it scores the original document higher than the permuted one. Twenty permutations are used in the test set in accordance with previous work.

In the *insertion* task, we evaluate models based on their ability to find the correct position of a sentence that has been removed from a document. To measure this, each sentence in a given document is relocated to every possible position. An insertion position is selected for which the model gives the highest coherence score to the document. The insertion score is then computed as the average fraction of sentences per document reinserted into their original position.

In the *reconstruction* task, the goal is to recover the original correct order of a shuffled paragraph given the starting sentence. We use beam search to drive the reconstruction process, with the different coherence models serving as the selection mechanism for beam search. We evaluate the performance of different models based on the rank correlation achieved by the top-1 reconstruction after search, averaged across different paragraphs.

For longer documents, since a random permutation is likely to be different than the original one at many places, the discrimination task is easy. Insertion is much more difficult since the candidate documents differ only by the position of one sentence. Reconstruction is also hard because small errors accumulate.

## 5.2 Datasets and Protocols

**Closed-domain:** The single-domain evaluation protocol is done on the Wall Street Journal (WSJ) portion of Penn Treebank (Table 2), similar to previous work (Nguyen and Joty, 2017)<sup>1</sup>.

**Open-domain:** Li and Jurafsky (2017) first proposed open-domain evaluation for coherence modelling using Wikipedia articles, but did not release the dataset<sup>2</sup>.

Hence, we create a new dataset based on Wikipedia and design three cross-domain evaluation protocols with increasing levels of difficulty. Based on the ontology defined by DBpedia<sup>3</sup>, we choose seven different categories under the domain *Person* and three other categories from irrelevant domains. We parse all the articles in these categories and extract paragraphs with more than 10 sentences to be used as the passages for training and evaluation. The statistics of this dataset

<sup>1</sup>Since the preprocessing pipeline is different, we generate the random permutations by ourselves.

<sup>2</sup>A version of preprocessed data consisting of integer indices is available online, but it is not usable without important preprocessing details.

<sup>3</sup><http://mappings.dbpedia.org/server/ontology/classes/>

is summarized in Table 1. The three settings with increasing level of hardness are as follows:

1. **Wiki-A(rticle)** randomly split all paragraphs of the seven categories under *Person* into training part and testing part;
2. **Wiki-C(ategory)** hold out paragraphs in one category from *Person* for evaluation and train on the remaining categories in *Person*;
3. **Wiki-D(omain)** train on all seven categories in *Person*, and evaluate on completely different domains, such as Plant, Institution, CelestialBody, and even WSJ.

Wiki-A setting is essentially the same protocol as the open domain evaluation as the one used in (Li and Jurafsky, 2017). Importantly, there is no distribution drift (up to sampling noise) between training and testing. Thus, this protocol only tests whether the coherence model is able to capture a rich enough set of signal for coherence, and does not check whether the learned cues are specific to the domain, or generic semantic signals. For example, cues based on style or regularities in discourse structure may not generalize to different domains. Therefore, we designed the much harder Wiki-C and Wiki-D to check whether the coherence models capture cross-domain transferrable features. In particular, in the Wiki-D setting, we even test whether the models trained on *Person* articles from Wikipedia generalize to WSJ articles.

Domain	Category	# Paras	Avg. # Sen.
Person	Artist	9553	11.87
	Athlete	23670	12.26
	Politician	2420	11.62
	Writer	3310	11.83
	MilitaryPerson	6428	11.90
	OfficeHolder	6578	11.54
	Scientist	2766	11.77
Species	Plant	3100	12.26
Organization	Institution	5855	11.58
Place	CelestialBody	414	11.55

Table 1: Statistics of the Wiki Dataset.

## 5.3 Baselines

We compared our proposed model LCD against two document-level discriminative models: (1) Clique-based discriminator **Clique-Discr.** (Li and Hovy, 2014; Li and Jurafsky, 2017) with window size 3 and 7. (2) Neural entity grid model **Grid-CNN** and **Extended Grid-CNN** (Nguyen

and Joty, 2017); And three generative models: (3) **Seq2Seq** (Li and Jurafsky, 2017); (4) **Vae-Seq2Seq** (Li and Jurafsky, 2017)<sup>4</sup>; (5) **LM**, an RNN language model, and use the difference between conditional log likelihood of a sentence given its preceding context, and the marginal log likelihood of the sentence. All the results are based on our own implementations except **Grid-CNN** and **Extended Grid-CNN**, for which we used code from the authors.

We compare these baselines to our proposed model with three different encoders:

1. **LCD-G**: use averaged GloVe vectors (Pennington et al., 2014) as the sentence representation;
2. **LCD-I**: use pre-trained InferSent (Conneau et al., 2017) as the sentence encoder;
3. **LCD-L**: apply max-pooling on the hidden state of the language model to get the sentence representation.

#### 5.4 Results on Domain-specific Data

	<b>Discr.</b>	<b>Ins.</b>
Clique-Discr. (3)	70.91	11.53
Clique-Discr. (7)	70.30	5.01
Grid-CNN	85.57 (85.13)	23.12
Extended Grid-CNN	88.69 (87.51)	25.95
Seq2Seq	86.95	27.28
Vae-Seq2Seq	87.01	26.73
LM	86.50	26.33
LCD-G	92.51	30.30
LCD-I	94.54	32.34
LCD-L	<b>95.49</b>	<b>33.79</b>

Table 2: Accuracy of **Discrimination** and **Insertion** tasks evaluated on WSJ. For Grid-CNN and Extended Grid-CNN, the numbers outside brackets are taken from the corresponding paper, and numbers shown in the bracket are based on our experiments with the code released by the authors.

We first evaluate the proposed models on the Wall Street Journal (WSJ) portion of Penn Treebank (Table 2). Our proposed models perform significantly better than all other baselines, even if we use the most naïve sentence encoder, *i.e.*, averaged GloVe vectors. Among all the sentence encoders,

<sup>4</sup>As the authors did not release their code, so we tried our best at replicating their model. Some important implementation details are missing from (Li and Jurafsky, 2017), so we cannot guarantee exactly the same setup as in (Li and Jurafsky, 2017).

**LM** trained on the local data in an unsupervised fashion performs the best, better than **InferSent** trained on a much larger corpus with supervised learning. In addition, combining the generative model **LM** with our proposed architecture as the sentence encoder improves the performance significantly over the generative model alone.

#### 5.5 Results on Open-Domain Data

Clique-Discr. (3)	76.17
Clique-Discr. (7)	73.86
Seq2Seq	86.63
Vae-Seq2Seq	82.40
LM	93.83
LCD-G	91.32
LCD-I	94.01
LCD-L	<b>96.01</b>

Table 3: Accuracy of discrimination task under Wiki-A

We next present results in the more challenging open-domain settings. Tables 3, 4, and 5 present results on the discriminative task under the Wiki-A, Wiki-C, Wiki-D settings. We do not report results of the neural entity grid models, since these models heavily depend on rich linguistics features from a preprocessing pipeline, but we cannot obtain these features on the Wiki datasets with high enough accuracy using standard preprocessing tools. As in the closed-domain setting, our proposed models outperform all the baselines for almost all tasks even with the averaged GloVe vectors as the sentence encoder. Generally, **LCD-L** performs better than **LCD-I**, but their performances are comparable under Wiki-D setting. This result may be caused by the fact that **InferSent** is pre-trained on a much larger dataset in a supervised way, and generalizes better to unseen domains.

As the Wiki-A setting is similar to the open-domain setting proposed by Li and Jurafsky (2017), and we also have similar observations as stated in their papers. The generative models perform quite well under this setting and applying them on top of our proposed architecture as the sentence encoder further enhances their performances, as illustrated in Table 3. However, as observed in Tables 4 and 5, the generative models do not generalize as well into unseen categories, and perform even worse in unseen domains. We emphasize that a protocol like Wiki-A or simi-

Model	Artist	Athlete	Politician	Writer	MilitaryPerson	OfficeHolder	Scientist	Average
Clique-Discr. (3)	73.01	68.90	73.82	73.28	72.86	73.74	74.56	72.88
Clique-Discr. (7)	71.26	66.56	73.72	72.01	72.67	72.62	71.86	71.53
Seq2Seq	82.72	73.45	84.88	85.99	81.40	83.25	85.27	82.42
Vae-Seq2Seq	82.58	74.14	84.70	84.94	81.07	82.66	85.09	82.17
LM	88.18	78.79	88.95	90.68	87.02	87.35	91.92	87.56
LCD-G	89.66	86.06	90.98	90.26	89.23	89.86	90.64	89.53
LCD-I	92.14	89.03	93.23	92.07	91.63	92.39	93.03	91.93
LCD-L	<b>93.54</b>	<b>90.13</b>	<b>94.04</b>	<b>93.68</b>	<b>93.20</b>	<b>93.01</b>	<b>94.81</b>	<b>93.20</b>

Table 4: Accuracy of discrimination task under Wiki-C setting.

Model	Plant	Institution	CelestialBody	WSJ	Average
Clique-Discr. (3)	66.14	66.51	60.38	64.71	64.44
Clique-Discr. (7)	65.47	69.14	61.44	66.66	65.68
Seq2Seq	82.58	80.86	69.44	74.62	76.88
Vae-Seq2Seq	81.90	78.00	69.10	73.27	75.57
LM	81.88	83.82	74.78	79.78	80.07
LCD-G	86.57	86.10	79.16	82.51	83.59
LCD-I	<b>89.07</b>	88.58	80.41	<b>83.27</b>	85.33
LCD-L	88.83	<b>89.46</b>	<b>81.31</b>	82.23	<b>85.48</b>

Table 5: Accuracy of discrimination task under Wiki-D setting.

lar setup considered in Li and Jurafsky (2017) is insufficient for evaluating open domain performance. Because difficulties in open domain coherence modelling lie not only in the variety of style and content in the dataset, but also in the fact that training set cannot cover all potential variation there is in the wild, making cross domain generalization a critical requirement.

## 5.6 Paragraph Order Reconstruction Results

Model	Wiki-D (CelestialBody)	Wiki-A
Seq2Seq	0.2104	0.2119
LM	0.1656	0.1420
LCD-I	<b>0.2507</b>	0.2744
LCD-L	0.2326	<b>0.2900</b>

Table 6: Kendall’s tau for re-ordering on Wiki-A/D

As shown by the discrimination and insertion tasks, Seq2Seq and LM are the stronger baselines, so for paragraph reconstruction, we compare our method to them, on two cross domain settings, the simpler Wiki-A and the harder Wiki-D. We report the reconstruction quality via Kendall’s tau rank correlation in Table 6, which shows that our method is superior by a significant margin.

## 5.7 Hyperparameter Setting and Implementation Details

In this work, we search through different hyperparameter settings by tuning on the development data of the WSJ dataset, then apply the same setting across all the datasets and protocols. The

fact that one set of hyperparameters tuned on the closed-domain setting works across all protocols, including open-domain ones, demonstrates the robustness of our method.

The following hyperparameter settings are chosen: Adam optimizer (Kingma and Ba, 2014) with default settings and learning rate 0.001, and no weight decay; the number of hidden state  $d_h$  for the one-layer MLP as 500, input dropout probability  $p_i$  as 0.6, hidden dropout probability  $p_h$  as 0.3; the margin loss was found to be superior to log loss, and margin of 5.0 was selected. In addition, we use early-stopping based on validation accuracy in all runs.

Furthermore, during training, every time we encounter a document, we sample 50 triplets  $(s_i, s_{i+1}, s')$ ’s, where  $(s_i, s_{i+1})$ ’s form positive pairs while  $(s_i, s')$ ’s form negative pairs. So effectively, we resample sentences so that documents are trained for the same number of steps regardless of the length. For all the documents including the permuted ones, we add two special tokens to indicate the start and the end of the document.

## 5.8 Analysis

### 5.8.1 Ablation Study

To better understand how different design choices affect the performance of our model, we present the results of an ablation study using variants of our best-performing models in Table 7. The protocol used for this study is Wiki-D with CelestialBody and Wiki-WSJ, the two most challenging

datasets in all of our evaluations.

The first variant uses a unidirectional model instead of the default bidirectional mode with two separately trained models. The second variant only uses the concatenation of the two sentence representations as the features instead of the full feature representation described in Section 4.1.

Model	CelestialBody	Wiki-WSJ
LCD-L	81.31	82.23
no bidirectional	80.33	82.30
no extra features	79.28	79.84

Table 7: Ablation study: **Discr.** under Wiki-D

As we can see, even our ablated models still outperform the baselines, though performance drops slightly compared to the full model. This demonstrates the effectiveness of our general framework for modelling coherence.

### 5.8.2 Effect of Sample Coverage




Figure 2: Discrimination accuracy on CelestialBody and Wiki-WSJ with different portions of all valid samples. The x axis is in log-scale.

Previous work raised concerns that negative sampling cannot effectively cover the space of negatives for discriminative learning (Li and Jurafsky, 2017). Fig. 2 shows that for our local discriminative model, there is a diminishing return when considering greater coverage beyond certain point (20% on these datasets). Hence, our sampling strategy is more than sufficient to provide good coverage for training.

### 5.8.3 Comparison with Human Judgement

To evaluate how well our coherence model aligns with human judgements of text quality, we compare our coherence score to Wikipedia’s article-level “rewrite” flags. This flag is used for articles that do not adhere to Wikipedia’s style guidelines,

which could be due to other reasons besides text coherence, so this is an imperfect proxy metric. Nevertheless, we aim to demonstrate a potential correlation here, because carelessly written articles are likely to be both incoherent and in violation of style guidelines. This setup is much more challenging than previous evaluations of coherence models, as it requires the comparison of two articles that could be on very different topics.

For evaluation, we want to verify whether there is a difference in average coherence between articles marked for rewrite and articles that are not. We select articles marked with an article-level rewrite flag from Wikipedia, and we sample the non-rewrite articles randomly. We then choose articles that have a minimum of two paragraphs with at least two sentences. We use our model trained for the Wiki-D protocol, and average its output scores per paragraph, then average these paragraph scores to obtain article-level scores. This two-step process ensures that all paragraphs contribute roughly equally to the final coherence score. We then perform a one-tailed *t*-test for the mean coherence scores between the rewrite and no-rewrite groups.

We find that among articles of a typical length between 2,000 to 6,000 characters (Wikipedia average length c. 2,800 characters), the average coherence scores are 0.56 (marked for rewrite) vs. 0.79 (not marked) with a p-value of .008. For longer articles of 8,000 to 14,000 characters, the score gap is smaller (0.60 vs 0.64), and p-value is 0.250. It is possible that in the longer marked article, only a subportion of the article is incoherent, or that other stylistic factors play a larger role, which our simple averaging does not capture well.

## 6 Conclusion

In this paper, we examined the limitations of two general frameworks for coherence modelling; i.e. , passage-level discriminative models and generative models. We propose a simple yet effective local discriminative neural model which retains the advantages of generative models while addressing the limitations of both kinds of models. Experimental results on a wide range of tasks and datasets demonstrate that the proposed model outperforms previous state-of-the-art methods significantly and consistently on both domain-specific and open-domain datasets.



## Acknowledgements

We thank all the anonymous reviewers for their valuable inputs.

## References

- Regina Barzilay and Mirella Lapata. 2008. Modeling local coherence: An entity-based approach. *Computational Linguistics*, 34(1):1–34.
- Jill Burstein, Joel Tetreault, and Slava Andreyev. 2010. Using entity-based features to model coherence in student essays. In *Human language technologies: The 2010 annual conference of the North American chapter of the Association for Computational Linguistics*, pages 681–684. Association for Computational Linguistics.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loic Barrault, and Antoine Bordes. 2017. Supervised learning of universal sentence representations from natural language inference data. *EMNLP*.
- Micha Eisner and Eugene Charniak. 2011. Extending the entity grid with entity-specific features. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: short papers-Volume 2*, pages 125–129. Association for Computational Linguistics.
- Micha Elsner, Joseph Austerweil, and Eugene Charniak. 2007. A unified local and global model for discourse coherence. In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Proceedings of the Main Conference*, pages 436–443.
- Vanessa Wei Feng and Graeme Hirst. 2012. Extending the entity-based coherence model with multiple ranks. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 315–324. Association for Computational Linguistics.
- Barbara J Grosz, Scott Weinstein, and Aravind K Joshi. 1995. Centering: A framework for modeling the local coherence of discourse. *Computational linguistics*, 21(2):203–225.
- Camille Guinaudeau and Michael Strube. 2013. Graph-based local coherence modeling. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 93–103.
- Geoffrey E Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R Salakhutdinov. 2012. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*.
- Ari Holtzman, Jan Buys, Maxwell Forbes, Antoine Bosselut, David Golub, and Yejin Choi. 2018. Learning to write with cooperative discriminators. *arXiv preprint arXiv:1805.06087*.
- Chloé Kiddon, Luke Zettlemoyer, and Yejin Choi. 2016. Globally coherent text generation with neural checklist models. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 329–339.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- 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.
- Jiwei Li and Eduard Hovy. 2014. A model of coherence based on distributed sentence representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2039–2048.
- Jiwei Li and Dan Jurafsky. 2017. Neural net models for open-domain discourse coherence. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 198–209.
- Lajanugen Logeswaran, Honglak Lee, and Dragomir Radev. 2018. Sentence ordering and coherence modeling using recurrent neural networks. *AAAI*.
- Annie Louis and Ani Nenkova. 2012. A coherence model based on syntactic patterns. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 1157–1168. Association for Computational Linguistics.
- Eleni Miltsakaki and Karen Kukich. 2004. Evaluation of text coherence for electronic essay scoring systems. *Natural Language Engineering*, 10(1):25–55.
- Tasnim Mohiuddin, Shafiq Joty, and Dat Tien Nguyen. 2018. Coherence modeling of asynchronous conversations: A neural entity grid approach. *arXiv preprint arXiv:1805.02275*.
- Dat Tien Nguyen and Shafiq Joty. 2017. A neural local coherence model. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1320–1330.
- Cesc C Park and Gunhee Kim. 2015. Expressing an image stream with a sequence of natural sentences. In *Advances in Neural Information Processing Systems*, pages 73–81.

- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *NAACL*.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. *arXiv preprint arXiv:1703.01365*.
- Sandra A Thompson and William C Mann. 1987. Rhetorical structure theory. *IPRA Papers in Pragmatics*, 1(1):79–105.
- Trieu H Trinh, Andrew M Dai, Thang Luong, and Quoc V Le. 2018. Learning longer-term dependencies in rnns with auxiliary losses. *arXiv preprint arXiv:1803.00144*.
- Suzan Verberne, Lou Boves, Nelleke Oostdijk, and Peter-Arno Coppen. 2007. Evaluating discourse-based answer extraction for why-question answering. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 735–736. ACM.