

Improved Document Modelling with a Neural Discourse Parser

Fajri Koto Jey Han Lau Timothy Baldwin

School of Computing and Information Systems
The University of Melbourne

ffajri@student.unimelb.edu.au, jeyhan.lau@gmail.com, tbaldwin@unimelb.edu.au

Abstract

Despite the success of attention-based neural models for natural language generation and classification tasks, they are unable to capture the discourse structure of larger documents. We hypothesize that explicit discourse representations have utility for NLP tasks over longer documents or document sequences, which sequence-to-sequence models are unable to capture. For abstractive summarization, for instance, conventional neural models simply match source documents and the summary in a latent space without explicit representation of text structure or relations. In this paper, we propose to use neural discourse representations obtained from a rhetorical structure theory (RST) parser to enhance document representations. Specifically, document representations are generated for discourse spans, known as the elementary discourse units (EDUs). We empirically investigate the benefit of the proposed approach on two different tasks: abstractive summarization and popularity prediction of online petitions. We find that the proposed approach leads to improvements in all cases.

1 Introduction

Natural language generation and document classification have been widely conducted using neural sequence models based on the encoder–decoder architecture. The underlying technique relies on the production of a context vector as the document representation, to estimate both tokens in natural language generation and labels in classification tasks. By combining recurrent neural networks with attention (Bahdanau et al., 2015), the model is able to learn contextualized representations of words at the sentence level. However, higher-level concepts, such as discourse structure beyond the sentence, are hard for an RNN to learn, especially for longer documents. We hypothesize that NLP

tasks such as summarization and document classification can be improved through the incorporation of discourse information.

In this paper, we propose to incorporate latent representations of discourse units into neural training. A discourse parser can provide information about the document structure as well as the relationships between discourse units. In a summarization scenario, for example, this information may help to remove redundant information or discourse disfluencies. In the case of document classification, the structure of the text can provide valuable hints about the document category. For instance, a scientific paper follows a particular discourse narrative pattern, different from a short story. Similarly, we may be able to predict the societal influence of a document such as a petition document, in part, from its discourse structure and coherence.

Specifically, discourse analysis aims to identify the organization of a text by segmenting sentences into units with relations. One popular representation is Rhetorical Structure Theory (RST) proposed by Mann and Thompson (1988), where the document is parsed into a hierarchical tree, where leaf nodes are the segmented units, known as Entity Discourse Units (EDUs), and non-terminal nodes define the relations.

As an example, in Figure 1 the two-sentence text has been annotated with discourse structure based on RST, in the form of 4 EDUs connected with discourse labels `attr` and `elab`. Arrows in the tree capture the nuclearity of relations, wherein a “satellite” points to its “nucleus”. The *Nucleus* unit is considered more prominent than the *Satellite*, indicating that the *Satellite* is a supporting sentence for the *Nucleus*. Nuclearity relationships between two EDUs can take the following three forms: Nucleus–Satellite, Satellite–Nucleus, and Nucleus–Nucleus. In this

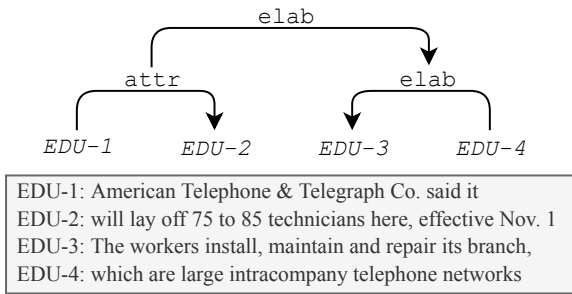


Figure 1: An example of a discourse tree, from (Yu et al., 2018); *elab* = elaboration; *attr* = attribute.

work, we use our reimplementa-tion of the state of the art neural RST parser of Yu et al. (2018), which is based on eighteen relations: *purp*, *cont*, *attr*, *evid*, *comp*, *list*, *back*, *same*, *topic*, *mann*, *summ*, *cond*, *temp*, *eval*, *text*, *cause*, *prob*, *elab*.¹

This research investigates the impact of dis-course representations obtained from an RST parser on natural language generation and docu-ment classification. We primarily experiment with an abstractive summarization model in the form of a pointer-generator network (See et al., 2017), focusing on two factors: (1) whether summariza-tion benefits from discourse parsing; and (2) how a pointer-generator network guides the summa-rization model when discourse information is pro-vided. For document classification, we investigate the content-based popularity prediction of online petitions with a deep regression model (Subrama-nian et al., 2018). We argue that document struc-ture is a key predictor of the societal influence (as measured by signatures to the petition) of a docu-ment such as a petition.

Our primary contributions are as follows: (1) we are the first to incorporate a neural discourse parser in sequence training; (2) we empirically demonstrate that a latent representation of dis-course structure enhances the summaries gener-ated by an abstractive summarizer; and (3) we show that discourse structure is an essential factor in modelling the popularity of online petitions.

2 Related Work

Discourse parsing, especially in the form of RST parsing, has been the target of research over a long period of time, including pre-neural feature engi-

neering approaches (Hernault et al., 2010; Feng and Hirst, 2012; Ji and Eisenstein, 2014). Two approaches have been proposed to construct dis-course parses: (1) bottom-up construction, where EDU merge operations are applied to single units; and (2) transition parser approaches, where the discourse tree is constructed as a sequence of parser actions. Neural sequence models have also been proposed. In early work, Li et al. (2016a) applied attention in an encoder-decoder frame-work and slightly improved on a classical feature-engineering approach. The current state of the art is a neural transition-based discourse parser (Yu et al., 2018) which incorporates implicit syntax features obtained from a bi-affine dependency parser (Dozat and Manning, 2017). In this work, we employ this discourse parser to generate dis-course representations.

2.1 Discourse and Summarization

Research has shown that discourse parsing is valu-able for summarization. Via the RST tree, the salience of a given text can be determined from the nuclearity structure. In extractive summarization, Ono et al. (1994), O’Donnell (1997), and Marcu (1997) suggest introducing penalty scores for each EDU based on the nucleus-satellite structure. In recent work, Schrimpf (2018) utilizes the *topic* relation to divide documents into sentences with similar topics. Every chunk of sentences is then summarized in extractive fashion, resulting in a concise summary that covers all of the topics dis-cussed in the passage.

Although the idea of using discourse infor-mation in summarization is not new, most work to date has focused on extractive summarization, where our focus is abstractive summarization. Gerani et al. (2014) used the parser of Joty et al. (2013) to RST-parse product reviews. By extract-ing graph-based features, important aspects are identified in the review and included in the sum-mary based on a template-based generation frame-work. Although the experiment shows that the RST can be beneficial for content selection, the proposed feature is rule-based and highly tailored to review documents. Instead, in this work, we extract a latent representation of the discourse di-rectly from the Yu et al. (2018) parser, and incor-porate this into the abstractive summarizer.

¹The details of each relation can be found on the RST website <http://www.sfu.ca/rst/index.html>

2.2 Discourse Analysis for Document Classification

Bhatia et al. (2015) show that discourse analyses produced by an RST parser can improve document-level sentiment analysis. Based on DPLP (Discourse Parsing from Linear Projection) — an RST parser by Ji and Eisenstein (2014) — they recursively propagate sentiment scores up to the root via a neural network.

A similar idea was proposed by Lee et al. (2018), where a recursive neural network is used to learn a discourse-aware representation. Here, DPLP is utilized to obtain discourse structures, and a recursive neural network is applied to the doc2vec (Le and Mikolov, 2014) representations for each EDU. The proposed approach is evaluated over sentiment analysis and sarcasm detection tasks, but found to not be competitive with benchmark methods.

Our work is different in that we use the latent representation (as distinct from the decoded discrete predictions) obtained from a neural RST parser. It is most closely related to the work of Bhatia et al. (2015) and Lee et al. (2018), but intuitively, our discourse representations contain richer information, and we evaluate over more tasks such as popularity prediction of online petitions.

3 Discourse Feature Extraction

To incorporate discourse information into our models (for summarization or document regression), we use the RST parser developed by Yu et al. (2018) to extract shallow and latent discourse features. The parser is competitive with other traditional parsers that use heuristic features (Feng and Hirst, 2012; Li et al., 2014; Ji and Eisenstein, 2014) and other neural network-based parsers (Li et al., 2016b).

3.1 Shallow Discourse Features

Given a discourse tree produced by the RST parser (Yu et al., 2018), we compute several shallow features for an EDU: (1) the nuclearity score; (2) the relation score for each relation; and (3) the node type and that of its sibling.

Intuitively, the nuclearity score measures how informative an EDU is, by calculating the (rela-

tive) number of ancestor nodes that are nuclei:²

$$\frac{\sum_{x \in \text{ancestor}(e)} \mathbb{1}_{\text{nucleus}}(x)}{h(\text{root})}$$

where e is an EDU; $h(x)$ gives the height from node x ;³ and $\mathbb{1}_{\text{nucleus}}(x)$ is an indicator function, i.e. it returns 1 when node x is a nucleus and 0 otherwise.

The relation score measures the importance of a discourse relation to an EDU, by computing the (relative) number of ancestor nodes that participate in the relation:

$$\frac{\sum_{x \in \text{ancestor}(e)} \mathbb{1}_{r_j}(x) h(x)}{\sum_{x \in \text{ancestor}(e)} h(x)}$$

where r_j is a discourse relation (one of 18 in total).

Note that we weigh each ancestor node here by its height; our rationale is that ancestor nodes that are closer to the root are more important. The formulation of these shallow features (nuclearity and relation scores) are inspired by Ono et al. (1994), who propose a number of ways to score an EDU based on the RST tree structure.

Lastly, we have 2 more features for the node type (nucleus or satellite) of the EDU and its sibling. In sum, our shallow feature representation for an EDU has 21 dimensions: 1 nuclearity score, 18 relation scores, and 2 node types.

3.2 Latent Discourse Features

In addition to the shallow features, we also extract latent features from the RST parser.

In the RST parser, each word and POS tag of an EDU span is first mapped to an embedding and concatenated to form the input sequence $\{x_1^w, \dots, x_m^w\}$ (m is number of words in the EDU). Yu et al. (2018) also use syntax features ($\{x_1^s, \dots, x_m^s\}$) from the bi-affine dependency parser (Dozat and Manning, 2017). The syntax features are the output of the multi-layer perceptron layer (see Dozat and Manning (2017) for full details).

The two sequences are then fed to two (separate) bi-directional LSTMs and average pooling is applied to learn the latent representation for an

²The ancestor nodes of an EDU are all the nodes traversed in its path to the root.

³Note that tree height is computed from the leaves, and so the *height* of the root node is equivalent to the *depth* of a leaf node.

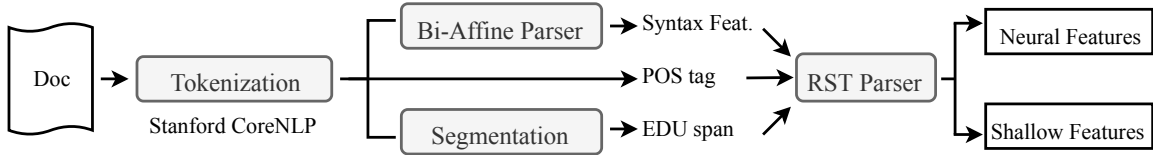


Figure 2: Pipeline of RST feature extraction

EDU:

$$\begin{aligned} \{h_1^w, \dots, h_m^w\} &= \text{Bi-LSTM}_1(\{x_1^w, \dots, x_m^w\}) \\ \{h_1^s, \dots, h_m^s\} &= \text{Bi-LSTM}_2(\{x_1^s, \dots, x_m^s\}) \\ h^e &= \text{Avg-Pool}(\{h_1^w, \dots, h_m^w\}) \oplus \\ &\quad \text{Avg-Pool}(\{h_1^s, \dots, h_m^s\}) \end{aligned}$$

where \oplus denotes the concatenate operation.

Lastly, Yu et al. (2018) apply another bi-directional LSTM over the EDUs to learn a contextualized representation:

$$\{f_1, \dots, f_n\} = \text{Bi-LSTM}(\{h_1^e, \dots, h_n^e\})$$

We extract the contextualized EDU representations ($\{f_1, \dots, f_n\}$) and use them as latent discourse features.

3.3 Feature Extraction Pipeline

In Figure 2, we present the feature extraction pipeline. Given an input document, we use Stanford CoreNLP to tokenize words and sentences, and obtain the POS tags.⁴ We then parse the processed input with the bi-affine parser (Dozat and Manning, 2017) to get the syntax features.

The RST parser (Yu et al., 2018) requires EDU span information as input. Previous studies have generally assumed the input text has been pre-processed to obtain EDUs, as state-of-the-art EDU segmentation models are very close to human performance (Hernault et al., 2010; Ji and Eisenstein, 2014). For our experiments, we use the pre-trained EDU segmentation model of Ji and Eisenstein (2014) to segment the input text to produce the EDUs.

Given the syntax features (from the bi-affine parser), POS tags, EDU spans, and tokenized text, we feed them to the RST parser to extract the shallow and latent discourse features.

We re-implemented the RST Parser in *PyTorch* and were able to reproduce the results reported in the original paper. We train the parser on the same

⁴<https://stanfordnlp.github.io/CoreNLP/>

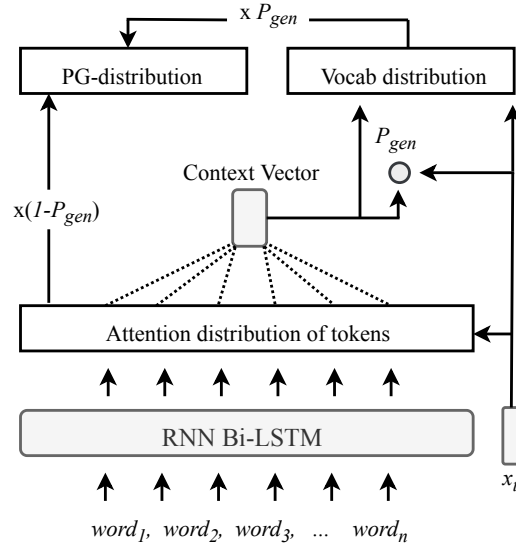


Figure 3: Architecture of the pointer-generator network (See et al., 2017).

data (385 documents from the Wall Street Journal), based on the configuration recommended in the paper.

To generate syntax features, we re-train an open-source bi-affine model, and achieve over 95% unlabelled and labelled attachment scores.⁵ Source code used in our experiments is available at: <https://github.com/fajri91/RSTExtractor>.

4 Abstractive Summarization

Abstractive summarization is the task of creating a concise version of a document that encapsulates its core content. Unlike extractive summarization, abstractive summarization has the ability to create new sentences that are not in the original document; it is closer to how humans summarize, in that it generates paraphrases and blends multiple sentences in a coherent manner.

Current sequence-to-sequence models for abstractive summarization work like neural machine translation models, in largely eschewing symbolic

⁵<https://github.com/XuezheMax/NeuroNLP2>

analysis and learning purely from training data. Pioneering work such as [Rush et al. \(2015\)](#), for instance, assumes the neural architecture is able to learn main sentence identification, discourse structure analysis, and paraphrasing all in one model. Studies such as [Gehrmann et al. \(2018\)](#); [Hsu et al. \(2018\)](#) attempt to incorporate additional supervision (e.g. content selection) to improve summarization. Although there are proposals that extend sequence-to-sequence models based on discourse structure — e.g. [Cohan et al. \(2018\)](#) include an additional attention layer for document sections — direct incorporation of discourse information is rarely explored.

[Hare and Borchardt \(1984\)](#) observe four core activities involved in creating a summary: (1) topic sentence identification; (2) deletion of unnecessary details; (3) paragraph collapsing; and (4) paraphrasing and insertion of connecting words. Current approaches ([Nallapati et al., 2016](#); [See et al., 2017](#)) capture topic sentence identification by leveraging the pointer network to do content selection, but the model is left to largely figure out the rest by providing it with a large training set, in the form of document–summary pairs. Our study attempts to complement the black-box model by providing additional supervision signal related to the discourse structure of a document.

4.1 Summarization Model

Our summarization model is based on the pointer–generator network ([See et al., 2017](#)). We present the architecture in [Figure 3](#), and summarize it as follows:

$$\{h_i\} = \text{Bi-LSTM}_1(\{w_i\}) \quad (1)$$

$$e_i^t = v^\top \tanh(W_h h_i + W_s s_t + b_e) \quad (2)$$

$$a^t = \text{softmax}(e^t)$$

$$h_t^* = \sum_i a_i^t h_i$$

$$P_{voc} = \text{softmax}(V'(V[s_t, h_t^*] + b_v) + b'_v)$$

$$p_{gen} = \sigma(w_{h^*}^\top h_t^* + w_s^\top s_t + w_x^\top x_t + b_g)$$

where $\{h_i\}$ are the encoder hidden states, $\{w_i\}$ are the embedded encoder input words, s_t is the decoder hidden state, and x_t is the embedded decoder input word.

The pointer–generator network allows the model to either draw a word from its vocabulary (generator mode), or select a word from the input document (pointer mode). p_{gen} is a scalar denoting the probability of triggering the generator

mode, and P_{voc} gives us the generator mode’s vocabulary probability distribution. To get the final probability distribution over all words, we sum up the attention weights and P_{voc} :

$$P(w) = p_{gen} P_{voc}(w) + (1 - p_{gen}) \sum_{i:w_i=w} a_i^t$$

To discourage repetitive summaries, [See et al. \(2017\)](#) propose adding coverage loss in addition to the cross-entropy loss:

$$c^t = \sum_{t'=0}^{t-1} a^{t'}$$

$$e_i^t = v^\top \tanh(W_h h_i + W_s s_t + W_c c_i^t + b_e)$$

$$\text{covloss}_t = \sum_i \min(a_i^t, c_i^t)$$

Intuitively, the coverage loss works by first summing the attention weights over all words from previous decoding steps (c^t), using that information as part of the attention computation (e_i^t), and then penalising the model if previously attended words receive attention again (covloss_t). [See et al. \(2017\)](#) train the model for an additional 3K steps with the coverage penalty after it is trained with cross-entropy loss.

4.2 Incorporating the Discourse Features

We experiment with several simple methods to incorporate the discourse features into our summarization model. Recall that the discourse features (shallow or latent) are generated for each EDU, but the summarization model operates at the word level. To incorporate the features, we assume each word within an EDU span receives the same discourse feature representation. Henceforth we use g and f to denote shallow and latent discourse features.

Method-1 (M1): Incorporate the discourse features in the Bi-LSTM layer ([Equation \(1\)](#)) by concatenating them with the word embeddings:

$$\{h_i\} = \text{Bi-LSTM}_1(\{w_i \oplus f_i\}); \text{ or}$$

$$\{h_i\} = \text{Bi-LSTM}_1(\{w_i \oplus g_i\})$$

Method-2 (M2): Incorporate the discourse features by adding another Bi-LSTM:

$$\{h'_i\} = \text{Bi-LSTM}_2(\{h_i \oplus f_i\}); \text{ or}$$

$$\{h'_i\} = \text{Bi-LSTM}_2(\{h_i \oplus g_i\})$$

Method	F1			Recall			Precision		
	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L
PG	36.82	15.92	33.57	37.36	16.10	34.05	38.72	16.86	35.32
+M1-latent	37.76	16.51	34.48	40.15	17.52	36.65	37.90	16.64	34.61
+M1-shallow	37.45	16.23	34.22	40.15	17.38	36.68	37.34	16.24	34.13
+M2-latent	38.04	16.73	34.83	38.92	17.05	35.62	39.54	17.51	36.23
+M2-shallow	37.15	16.13	33.96	38.52	16.68	35.21	38.19	16.67	34.91
+M3-latent	37.04	16.05	33.86	37.52	16.22	34.29	38.95	16.98	35.63
+M3-shallow	37.09	16.15	33.95	39.05	16.97	35.73	37.62	16.46	34.45
PG+Cov	39.32	17.22	36.02	40.33	17.61	36.93	40.82	17.99	37.42
+M1-latent	40.06	17.63	36.70	44.44	19.53	40.69	38.60	17.05	35.39
+M1-shallow	39.78	17.50	36.50	43.50	19.08	39.89	38.94	17.22	35.75
+M2-latent	40.00	17.62	36.72	43.53	19.17	39.94	39.28	17.37	36.09
+M2-shallow	39.58	17.30	36.36	44.00	19.19	40.38	38.40	16.87	35.31
+M3-latent	39.23	17.00	36.00	42.95	18.54	39.37	38.29	16.69	35.16
+M3-shallow	39.57	17.31	36.28	43.85	19.14	40.17	38.37	16.86	35.20

Table 1: Abstractive summarization results.

Method-3 (M3): Incorporate the discourse features in the attention layer (Equation (2)):

$$e_i^t = v^T \tanh(W_h h_i + W_s s_t + W_f f_i + b_e); \text{ or}$$

$$e_i^t = v^T \tanh(W_h h_i + W_s s_t + W_g g_i + b_e)$$

4.3 Data and Result

We conduct our summarization experiments using the anonymized CNN/DailyMail corpus (Nallapati et al., 2016). We follow the data preprocessing steps in See et al. (2017) to obtain 287K training examples, 13K validation examples, and 11K test examples.

All of our experiments use the default hyperparameter configuration of See et al. (2017). Every document and its summary pair are truncated to 400 and 100 tokens respectively (shorter texts are padded accordingly). The model has 256-dimensional hidden states and 128-dimensional word embeddings, and vocabulary is limited to the most frequent 50K tokens. During test inference, we similarly limit the length of the input document to 400 words and the length of the generated summary to 35–100 words for beam search.

Our experiment has two pointer-generator network baselines: (1) one without the coverage mechanism (“PG”); and (2) one with the coverage mechanism (“PG+Cov”; Section 4.1). For each baseline, we incorporate the latent and shallow discourse features separately in 3 ways (Section 4.2), giving us 6 additional results.

We train the models for approximately 240,000–270,000 iterations (13 epochs). When we include the coverage mechanism (second baseline), we train for an additional 3,000–3,500 iterations using the coverage penalty, following See et al. (2017).

We use ROUGE (Lin, 2004) as our evaluation metric, which is a standard measure based on overlapping n-grams between the generated summary and the reference summary. We assess unigram (R-1), bigram (R-2), and longest-common-subsequence (R-L) overlap, and present F1, recall and precision scores in Table 1.

For the first baseline (PG), we see that incorporating discourse features consistently improves recall and F1. This observation is consistent irrespective of how (e.g. M1 or M2) and what (e.g. shallow or latent features) we add. These improvements do come at the expense of precision, with the exception of M2-latent (which produces small improvements in precision). Ultimately however, the latent features are in general a little better, with M2-latent producing the best results based on F1.

We see similar observations for the second baseline (PG+Cov): recall is generally improved at the expense of precision. In terms of F1, the gap between the baseline and our models is a little closer, and M1-latent and M2-latent are the two best performers.

Example-1
Reference: nigel short said women should accept they're ` hard-wired very differently ' made comments when explaining why there were so few women in chess . female chess players reacted angrily to mr short 's statements last night .
PG + Cov + M1-latent: nigel short , 49 , said women should accept they were ` hard-wired very differently ' and were n't as adept at playing chess as men . mr short , who was the first english player to play a world chess championship match , made the comments when explaining why there were few women in competitive chess .
PG + Cov: nigel short has controversially claimed women were not suited to playing chess because it required logical thinking . but critics immediately pointed out that he lost a high-profile game against a woman grandmaster . he told new in chess magazine : ` why should -lsb- men and women -rsb- function in the same way ? '
Example-2
Reference: deva joseph hit problems when she could n't fit handbag inside suitcase . 14-year-old left in floods of tears after flight to spain left without her . offered to pay for bag to go in hold but was told she needed a credit card . easyjet said it should have made an exception to policy of accepting cash .
PG + Cov + M1-latent: deva joseph , 14 , was left stranded at a busy airport after being told she would not be allowed to take her flight home to spain . the teenager offered to pay for the second item to be put in the hold but was told only credit cards would be accepted - even though she is too young to have one . a schoolgirl was left stranded at a busy airport after easyjet refused to let her board its plane - because she was carrying two pieces of hand luggage .
PG + Cov: deva joseph , 14 , was stranded at a busy airport after being told she would not be allowed to take her flight home to spain . teenager offered to pay for the second item to be put in the hold but she is too young to have one .

Figure 4: Comparison of summaries between our model and the baseline.

4.4 Analysis and Discussion

We saw previously that our models generally improve recall. To better understand this, we present 2 examples of generated summaries, one by the baseline (“PG+Cov”) and another by our model (“M1-latent”), in Figure 4. The highlighted words are overlapping words in the reference. In the first example, we notice that the summary generated by our model is closer to the reference, while the baseline has other unimportant details (e.g. *he told new in chess magazine : ‘ why should -lsb- men and women -rsb- function in the same way ?*). In the second example, although there are more overlapping words in our model’s summary, it is a little repetitive (e.g. first and third paragraph) and less concise.

Observing that our model generally has better recall (Table 1) and its summaries tend to be more verbose (e.g. second example in Figure 4), we calculated the average length of generated summaries for PG+Cov and M1-latent, and found that they are of length 55.2 and 64.4 words respectively. This suggests that although discourse information

helps the summarization model overall (based on consistent improvement in F1), the negative side effect is that the summaries tend to be longer and potentially more repetitive.

5 Petition Popularity Prediction

Online petitions are open letters to policy-makers or governments requesting change or action, based on the support of members of society at large. Understanding the factors that determine the popularity of a petition, i.e. the number of supporting signatures it will receive, provides valuable information for institutions or independent groups to communicate their goals (Proskurnia et al., 2017).

Subramanian et al. (2018) attempt to model petition popularity by utilizing the petition text. One novel contribution is that they incorporate an auxiliary ordinal regression objective that predicts the scale of signatures (e.g. 10K vs. 100K). Their results demonstrate that the incorporation of auxiliary loss and hand-engineered features boost performance over the baseline.

In terms of evaluation metric, Subramanian et al. (2018) use: (1) mean absolute error (MAE); and (2) mean absolute percentage error (MAPE), calculated as $\frac{100}{n} \sum_{i=1}^n \frac{|\hat{y}_i - y_i|}{y_i}$, where n is the number of examples and \hat{y}_i (y_i) the predicted (true) value. Note that in both cases, lower numbers are better.

Similar to the abstractive summarization task, we experiment with incorporating the discourse features of the petition text to the petition regression model, under the hypothesis that discourse structure should benefit the model.

5.1 Deep Regression Model

As before, our model is based on the model of Subramanian et al. (2018). The input is a concatenation of the petition’s title and content words, and the output is the log number of signatures. The input sequence is mapped to GloVe vectors (Pennington et al., 2014) and processed by several convolution filters with max-pooling to create a fixed-width hidden representation, which is then fed to fully connected layers and ultimately activated by an exponential linear unit to predict the output. The model is optimized with mean squared error (MSE). In addition to the MSE loss, the authors include an auxiliary ordinal regression objective that predicts the scale of signatures (e.g. {10, 100, 1000, 10000, 100000}), and found that it

improves performance. Our model is based on the best model that utilizes both the MSE and ordinal regression loss.

5.2 Incorporating the Discourse Features

We once again use the methods of incorporation presented in Section 4.2. As the classification model uses convolution networks, only Method-1 is directly applicable.

We also explore replacing the convolution networks with a bidirectional LSTM (“Bi-LSTM w/ GloVe”), based on the idea that recurrent networks are better at capturing long range dependencies between words and EDUs. For this model, we test both Method-1 and Method-2 to incorporate the discourse features.⁶

Lastly, unlike the summarization model that needs word level input (as the pointer network requires words to attend to in the source document), we experiment with replacing the input words with EDUs, and embed the EDUs with either the latent (“Bi-LSTM w/ latent”) or the shallow (“Bi-LSTM w/ shallow”) features.

5.3 Data, Result, and Discussion

We use the US Petition dataset from (Subramanian et al., 2018).⁷ In total we have 1K petitions with over 12M signatures after removing petitions that have less than 150 signatures. We use the same train/dev/test split of 80/10/10 as Subramanian et al. (2018).

We present the test results in Table 2. We tune the models based on the development set using MAE, and find that most converge after 8K–10K iterations of training. We are able to reproduce the performance of the baseline model (“CNN w/ GloVe”), and find that once again, adding the shallow discourse features improves results.

Next we look at the LSTM model (“Bi-LSTM w/ GloVe”). Interestingly, we found that replacing the CNN with an LSTM results in improved MAE, but worse MAPE. Adding discourse features to this model generally has marginal improvement in all cases.

When we replace the word sequence with EDUs (“Bi-LSTM w/ latent” and “Bi-LSTM w/ shallow”), we see that the latent features outperform the shallow features. This is perhaps unsurprising,

⁶Our LSTM has 200 hidden units, and uses a dropout rate of 0.3, and L2 regularization.

⁷The data is collected from <https://petitions.whitehouse.gov>.

Model	MAE	MAPE
CNN w/ GloVe	1.16	14.38
+ M1-latent	1.15	14.66
+ M1-shallow	1.12⁽¹⁾	14.19
Bi-LSTM w/ GloVe	1.14	14.57
+ M1-latent	1.13	14.39
+ M1-shallow	1.13	14.25
+ M2-latent	1.12	14.02
+ M2-shallow	1.13	14.20
Bi-LSTM w/ latent	1.11⁽²⁾	13.91
Bi-LSTM w/ shallow	1.15	14.67

Table 2: Average petition regression performance over 3 runs (noting that lower is better for both MAE and MAPE). One-sided t-tests show that both (1) and (2) are significantly better than the baseline ($p < 0.05$ and $p < 0.005$, resp.).

given that the shallow discourse features have no information about the actual content, and are unlikely to be effective when used in isolation without the word features.

6 Conclusion and Future Work

In this paper, we explore incorporating discourse information into two tasks: abstractive summarization and popularity prediction of online petitions. We experiment with both hand-engineered shallow features and latent features extracted from a neural discourse parser, and found that adding them generally benefits both tasks. The caveat, however, is that the best method of incorporation and feature type (shallow or latent) appears to be task-dependent, and so it remains to be seen whether we can find a robust universal approach for incorporating discourse information into NLP tasks.

References

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. *International Conference on Learning Representations*.

Parminder Bhatia, Yangfeng Ji, and Jacob Eisenstein. 2015. Better document-level sentiment analysis from rst discourse parsing. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2212–2218.

Arman Cohan, Franck Dernoncourt, Doo Soon Kim,

- Trung Bui, Seokhwan Kim, Walter Chang, and Nazli Goharian. 2018. A discourse-aware attention model for abstractive summarization of long documents. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 615–621, New Orleans, Louisiana.
- Timothy Dozat and Christopher D. Manning. 2017. Deep biaffine attention for neural dependency parsing. In *Proceedings of the 2016 international conference on learning representations*, pages 1–8.
- Vanessa Wei Feng and Graeme Hirst. 2012. Text-level discourse parsing with rich linguistic features. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics*, pages 60–68.
- Sebastian Gehrmann, Yuntian Deng, and Alexander M Rush. 2018. Bottom-up abstractive summarization. In *Proceedings of Empirical Methods in Natural Language Processing*, pages 4098–4109.
- Shima Gerani, Yashar Mehdad, Giuseppe Carenini, Raymond T. Ng, and Bitia Nejat. 2014. Abstractive summarization of product reviews using discourse structure. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1602–1613.
- Victoria C. Hare and Kathleen M. Borchardt. 1984. Direct instruction of summarization skills. *Journal Reading Research Quarterly, Wiley, International Reading Association*, pages 62–78.
- Hugo Hernault, Helmut Prendinger, David A. duVerle, and Mitsuru Ishizuka. 2010. Hilda: A discourse parser using support vector machine classification. *Dialogue and Discourse*, 1(3):1–33.
- Wan Ting Hsu, Chieh-Kai Lin, Ming-Ying Lee, Kerui Min, Jing Tang, and Min Sun. 2018. A unified model for extractive and abstractive summarization using inconsistency loss. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, pages 132–141.
- Yangfeng Ji and Jacob Eisenstein. 2014. Representation learning for text-level discourse parsing. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, pages 13–24.
- Shafiq R. Joty, Giuseppe Carenini, Raymond T. Ng, and Yashar Mehdad. 2013. Combining intra- and multi-sentential rhetorical parsing for document-level discourse analysis. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 486–496.
- Quoc V. Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *Proceedings of The 31st International Conference on Machine Learning*, pages 1188–1196.
- Kangwook Lee, Sanggyu Han, and Sung-Hyon Myaeng. 2018. A discourse-aware neural network-based text model for document-level text classification. *Journal of Information Science*, 44.
- Jiwei Li, Rumeng Li, and Eduard H. Hovy. 2014. Recursive deep models for discourse parsing. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2061–2069.
- Qi Li, Tianshi Li, and Baobao Chang. 2016a. Discourse parsing with attention-based hierarchical neural networks. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 362–371.
- Qi Li, Tianshi Li, and Baobao Chang. 2016b. Discourse parsing with attention-based hierarchical neural networks. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 362–371.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text Summarization Branches Out: Proceedings of the ACL-04 Workshop*, pages 74–81.
- 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*, pages 243–281.
- Daniel Marcu. 1997. From discourse structures to text summaries. In *Proceedings of of ACL Workshop on Intelligent Scalable Text Summarisation*, pages 82–88.
- Ramesh Nallapati, Bowen Zhou, Cicero Nogueira dos santos, Caglar Gulcehre, and Bing Xiang. 2016. Abstractive text summarization using sequence-to-sequence rnns and beyond. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 280–290.
- Michael O’Donnell. 1997. Variable-length on-line document generation. In *Proceedings of the 6th European Workshop on Natural Language Generation*, pages 82–91.
- Kenji Ono, Kazuo Sumita, and Seiji Miike. 1994. Abstract generation based on rhetorical structure extraction. In *Proceedings of the 15th conference on Computational linguistics*, pages 344–348.
- Jeffrey Pennington, Richard Socher, and Christopher D. 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.
- Julia Proskurnia, Przemyslaw A. Grabowicz, Ryota Kobayashi, Carlos Castillo, Philippe Cudré-Mauroux, and Karl Aberer. 2017. Predicting the success of online petitions leveraging multidimensional time-series. In *WWW ’17 Proceedings of the*

26th International Conference on World Wide Web, pages 755–764.

Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. In *Proceedings of Empirical Methods in Natural Language Processing*, pages 379–389.

Natalie M. Schrimpf. 2018. Using rhetorical topics for automatic summarization. *Proceedings of the Society for Computation in Linguistics*, 1(1):125–135.

Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 1073–1083.

Shivashankar Subramanian, Timothy Baldwin, and Trevor Cohn. 2018. Content-based popularity prediction of online petitions using a deep regression model. In *ACL 2018: 56th Annual Meeting of the Association for Computational Linguistics*, volume 2, pages 182–188.

Nan Yu, Meishan Zhang, and Guohong Fu. 2018. Transition-based neural RST parsing with implicit syntax features. In *roceedings of the 27th International Conference on Computational Linguistics*, pages 559–570.