

Profiling News Discourse Structure Using Explicit Subtopic Structures Guided Critics

Prafulla Kumar Choubey and Ruihong Huang

Department of Computer Science and Engineering

Texas A&M University

(prafulla.choubey, huangrh)@tamu.edu

Abstract

We present an actor-critic framework to induce subtopical structures in a news article for news discourse profiling. The model uses multiple critics that act according to known subtopic structures while the actor aims to outperform them. The content structures constitute sentences that represent latent subtopic boundaries. Then, we introduce a hierarchical neural network that uses the identified subtopic boundary sentences to model multi-level interaction between sentences, subtopics, and the document. Experimental results and analyses on the NewsDiscourse corpus show that the actor model learns to effectively segment a document into subtopics and improves the performance of the hierarchical model on the news discourse profiling task¹.

1 Introduction

News discourse profiling is a discourse processing task that aims to classify sentences in news articles into different content types, where each content type characterizes the specific discourse role of a sentence in describing a news story (Choubey et al., 2020). It is vital to effectively contextualize the occurrence of a news event, which has been shown useful for extracting event structures from a document (Choubey et al., 2020; Choubey and Huang, 2021). Furthermore, this task is likely to benefit a range of other NLP applications that require deep story-level text understanding, such as text summarization and complex question answering.

As the discourse roles are interpreted with respect to the main event, the current approach for discourse profiling uses a hierarchical neural network model (Choubey et al., 2020) that relies on a sentence-level encoder to obtain sentence embeddings, that heed to the local context, and a document embedding to obtain the underlying main

topic. The hierarchical model, intuitively, provides a mechanism for capturing both global and local dependencies among sentences and the main topic. However, the model is completely unaware of the underlying content organization structures that are used while producing news reports. Besides, squeezing document-level features into a single vector provides limited space to learn effective document representation and model its interaction with the sentences.

To extend the modeling capabilities and incorporate document-level content organization structures in the hierarchical model, we propose to decompose a document into latent subtopics by identifying subtopic boundary sentences, and model two levels of interactions—between sentences and a subtopic, and between subtopics and the document. We hypothesize that learning subtopic representations allows the model to focus on the locally relevant sentences, independent of main content, and identify its fine-grained discourse function within local subtopical context. Further interactions between subtopics and the document vector helps to determine the broader role of a subtopic with respect to the main content. For instance, in news document in Figure 1, we can identify discourse role of sentence *S7* by combining the two levels of information. First, sentences *S6* and *S8* describe events that happened years before the main event which can be modelled through interaction between document and subtopic embedding corresponding to [*S6-S8*]. Then, events in *S7* has temporal proximity with the events in sentences *S6* and *S8* which can be modelled through interaction between sentence *S7* embedding and subtopic embedding.

Several past works have independently studied the subtopic structures, and a document can exhibit multiple subtopic structures depending on the used segmentation criteria. In this paper, we consider two subtopic structures: 1) broad-genre topic segments generated by the TextTiling algorithm; and

¹Code and data are available at https://github.com/prafulla77/Discourse_Profiling_RL_EMNLP21Findings

S1	Police in South Australia (SA) on Thursday charged Geoffrey Adams over the 45-year-old cold case murder of his wife Colleen. M1	Events after Arrest
S2	Officers arrested Adams, aged 70, at his home in Wallaroo, 156 km north of SA's capital Adelaide, on Wednesday evening and charged him with murder. M1	
S3	Shortly after making the arrest, police travelled with Adams to a property in Maitland, 50 km south, where the couple lived with their two children at the time of Colleen's disappearance in 1973. C2	
S4	Footage captured by local news showed Adams in the backyard of the property with officers where he pointed to a concrete slab before becoming visually emotional. C2	
S5	Police on Thursday began work to recover Colleen Adams' remains from that site. C2	
S6	Colleen Adams was last seen alive on Nov. 22, 1973, at the Maitland property where Geoffrey Adams claimed she left with an unidentified middle-aged woman at 7 a.m. local time having told him that she was leaving him. D1	History
S7	She was officially reported missing by her mother the following month. D1	
S8	The case was declared a major crime in 1979 but police never made an arrest. D1	
S9	SA Police on Sunday announced that the case would be reviewed as part of Operation Persist, saying that they "have not given up hope of providing her family with some answers about her disappearance." D4	
S10	"We have an open mind as to what's happened and how it's happened, but we are of the opinion she has been murdered. D4	
S11	We don't hold out hope that she's just missing and will turn up," Michael Newbury, a sergeant with the Major Crime investigation unit, said on Sunday. D4	

Figure 1: An example document annotated with three different subtopic structures. The first is based on TextTiling (Hearst, 1997) and is shown with the black-solid line ([S1-S8],[S9-S11]). The second structure is based on locally inverted pyramid structure (discussed in § 5.2) and is shown through red-dashed lines ([S1-S5],[S6-S8],[S9-S11]). The third, shown by colored boxes, segments document based on the temporal position where the first segment (S1, S2) focuses on the main event, second segment (S3, S4, S5) describes events following the main event, third segment (S6, S7, S8) describes historical events and the last segment (S9, S10, S11) again covers current context.

2) news genre-specific inverted pyramid structure identified through a plausible rules-based surrogate (Section 5). Depending on the document, one of these subtopic structures may be more suitable for identifying discourse role labels. Also depending on the document, the most suitable subtopic structure may not be strictly the same as one of the known subtopic structures. For instance, in Figure 1, we show three subtopic structures, namely, TextTiling, inverted pyramid structure, and a new subtopic structure that is based on temporal frames. Here, three subtopic structures only partially overlap with each other. Notably, sentences in each of the segments obtained by considering *temporal frames* exhibit homogeneous discourse role labels and could be the most suitable structure to consider.

In this work, we limit ourselves to indirectly using known explicit subtopic structures as critics in a new variant of *actor-critic* model that selects between the standard REINFORCE (Williams, 1992) algorithm or imitation learning for training actor. Specifically, when subtopic structure identified by actor performs better than all known explicit subtopic structures, we baseline the standard REINFORCE (Williams, 1992) algorithm with the

average of reward obtained by all explicit subtopic structures. On the other hand, if one of the explicit subtopic structures perform better than the actor, we force the actor to imitate that subtopic structure. Intuitively, this allows the actor model to learn to identify subtopic boundary sentences most suitable for a given document that perform better than or at least comparable to a known explicit subtopic structure on discourse profiling task.

Experimental results on the NewsDiscourse corpus show that modeling latent subtopic structures in a hierarchical discourse model improves its performance by 2.6 and 1.3 points on average macro and micro F1 scores respectively. The improvement is consistent to different model initialization and shows that modeling the underlying sequential content organization structure enables the system to better predict content types for individual sentences.

2 Related Work

Theory on News Content Organization The theoretical studies on the organization of different news elements (Van Dijk, 1985; Myers and Simms, 1989; Bell, 1998; Schokkenbroek, 1999; Ytreberg,

2001; Mani et al., 2005; Saleh, 2014) have been extensively explored through case-studies in journalism and discourse. For instance, Pottker (2003) and Filak (2019) studied the widely used inverted pyramid structure that follows the standard relevance ordering. Specific to the content structures of hard-news reports, White (1997) observed that the generic hard-news exhibit non-linear structure where sub-components such as consequences, causes, contextualization or other supportive information possess **orbital** relationship with the main event. Further, these sub-components are organized around the main event with repetitions of the most newsworthy event, i.e. main content. This aligns with our inverted pyramid structure based on rule of relevance ordering, where a main sentence following non-main sentences indicates the segment boundaries. Besides, feature news also follow well-defined content organization structures such as having an introductory anecdotes or back-grounding which is taken care by our rules to separate historical and anecdotal contents from the relevance ordering.

Neural Models for Discourse Modeling Deep neural networks have been successfully explored for modeling discourse (Ji and Eisenstein, 2014a,b; Becker et al., 2017), including hierarchical models (Li et al., 2016b; Liu and Lapata, 2017; Dai and Huang, 2018) to induce hierarchical structure. Morey et al. (2017), however, found that some of the improvements from neural models on RST parsing are attributed to differences in evaluation procedures. Nonetheless, Morey et al. (2017) concluded that neural models are more effective in modeling discourse, though the relative error reduction rates are lower than reported. The better discourse modeling capabilities of neural models are also evident from their widespread adoption in follow-up works such as Lin et al. (2019); Zhang et al. (2020) and Koto et al. (2021).

For automatic text segmentation, multitude of approaches such as lexical overlap, bayesian learning or dynamic programming (Hearst, 1997; Choi, 2000; Utiyama and Isahara, 2001; Eisenstein and Barzilay, 2008; Du et al., 2013) have been proposed. The recent works rely on neural network models to learn different aspects of text segmentation such as coherence and cohesion (Wang et al., 2017; Sehikh et al., 2017; Bahdanau et al., 2016a; Arnold et al., 2019).

Reinforcement Learning for NLP Applications

Reinforcement learning has been frequently used for sequence generation tasks to mitigate exposure bias or to directly optimize task-specific evaluation metrics such as BLEU score (Ranzato et al., 2015; Henß et al., 2015; Bahdanau et al., 2016b; Paulus et al., 2017; Fedus et al., 2018). In addition, RL have been explored for range of NLP tasks such as question-answering (Xiong et al., 2017), dialog generation (Li et al., 2016a), text summarization (Chen and Bansal, 2018), knowledge-graph reasoning (Lin et al., 2018) and relation extraction (Qin et al., 2018). To the best of our knowledge, we are the first to explore RL techniques for exposing underlying content organization structures in news articles as well as using a linguistically motivated critic to reduce the variance of *reinforce* algorithm.

3 Task Description

News discourse profiling categorizes sentences in news articles into eight schematic categories that are defined following the news content schemata proposed by Van Dijk (Teun A, 1986; Van Dijk, 1988a,b; Choubey et al., 2020). The eight content types describe the common discourse roles of sentences in telling a news story. Specifically, *Main Event (M1)* sentences introduce the main event relating to the major subjects of a news article. *Consequence (M2)* sentences describe consequence events immediately triggered by the main event. *Previous Events (C1)* sentences describe the recent events that act as possible causes or preconditions for the main event while *Current Context (C2)* sentences describe remaining context informing contents.

News articles may also describe past events that precede the main events in months and years (*Historical Event (D1)*) or unverifiable situations that are often fictional or personal accounts of incidents of an unknown person (*Anecdotal Event (D2)*). Lastly, opinionated contents including reactions from immediate participants, experts, known personalities as well as journalists or news sources are covered in the *Evaluation (D3)* category, except speculations and projected consequences that are labeled as *Expectation (D4)*.

4 Model

We model the discourse profiling task as a two step process. Given a news document $X : \{H, x_1, x_2, \dots, x_n\}$ comprising of headline H and

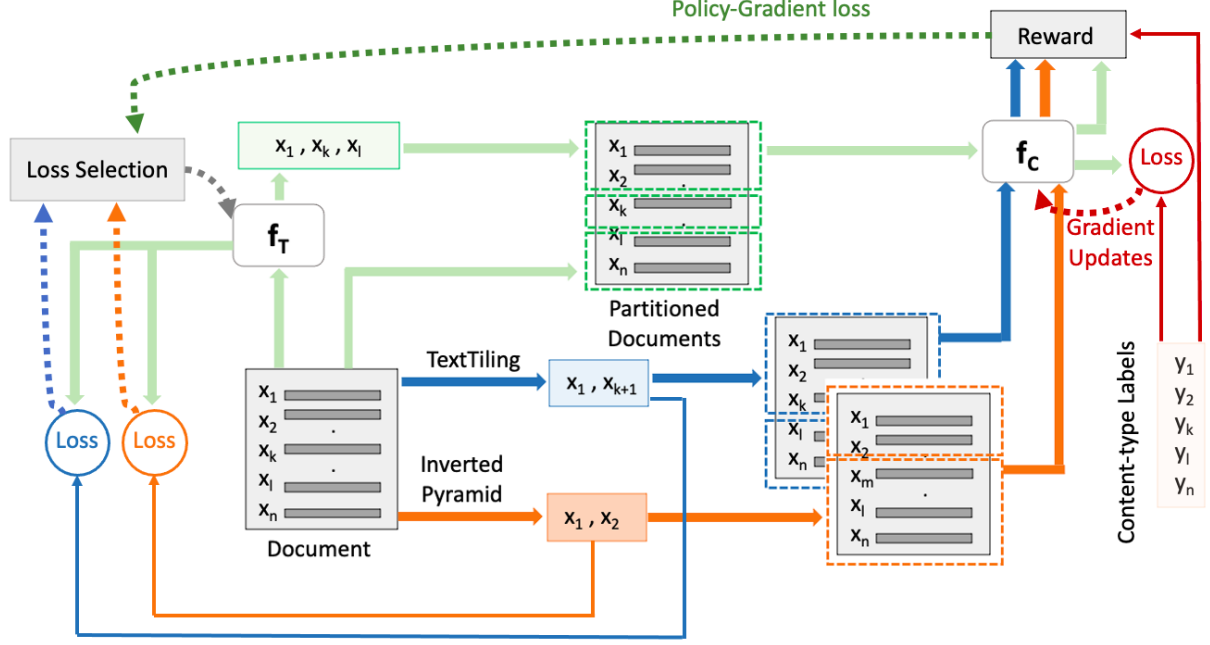


Figure 2: Neural-Network Architecture, including Gradient Flow Paths, for Incorporating Document-level Content Structures in a Discourse Profiling System

n sentences with their content-type labels $Y : \{y_1, y_2, \dots, y_n\}$, our main goal is to learn a model $f : X \rightarrow Y$ that classifies each sentence x_i in the document X to its content type y_i . In the first step, a latent function $f_T : X \rightarrow T \in \{1, 2, \dots, n\}^k$, a classifier, is used to identify k subtopic boundary sentences in the document. These boundary sentences are used to partition documents into multiple subtopics. In the second step, a classification function $f_C : [X, T] \rightarrow Y$ combines the output of latent function f_T with the sentences X in document to perform final content-type classification (Figure 2). Overall, the model consists of a sentence encoder, a biLSTM-based (Hochreiter and Schmidhuber, 1997) hierarchical encoder to obtain contextualized sentence representations used by both f_T and f_C , and a pointer decoder network to select subtopic boundary sentences that is exclusively used by f_T .

4.1 Hierarchical Sentence Encoder

We use a hierarchical encoder to learn context-aware sentence representations. Given a word sequence x_i represented by $\{w_{i1}, w_{i2}, \dots, w_{im}\}$, we first transform the sequence to contextualized word embeddings E_i using the pre-trained ELMo (Peters et al., 2018). Then, we use a word-level biLSTM layer over E_i to obtain hidden state representations H_i and take their weighted aver-

age to obtain the local sentence embedding S^L . Weights for hidden states are obtained using a two-layered feed-forward neural network. Finally, we apply another sentence-level biLSTM over the sequence of headline and sentence embeddings $\{H^L, S_1^L, S_2^L, \dots, S_n^L\}$ to obtain the contextualized sentences representations S^C that are later used in both the sub-modules f_T and f_C (eq. 1).

$$\begin{aligned}
 [E_{i1}, E_{i2}, \dots, E_{im}] &= \text{ELMo}([w_{i1}, w_{i2}, \dots, w_{im}]) \\
 [H_{i1}, H_{i2}, \dots, H_{im}] &= \text{biLSTM}^L([E_{i1}, E_{i2}, \dots, E_{im}]) \\
 \alpha_i[k] &= W_{\alpha 1}(\tanh(W_{\alpha 2}E_{ik} + b_{\alpha 2})) + b_{\alpha 1} \in R \\
 A_i &= \text{softmax}(\alpha_i) \in R^m \quad (1) \\
 S_i^L &= \sum_k A_i[k]H_{ik} \in R^{2d_{rnn}} \\
 [H^C, S_1^C, \dots, S_n^C] &= \text{biLSTM}^C([H^L, S_1^L, \dots, S_n^L])
 \end{aligned}$$

4.2 Identifying Subtopic Boundary Sentences through Pointer Decoder Network

Given the sentence embeddings from hierarchical sentence encoders, we use an LSTM decoder based pointer network (Vinyals et al., 2015) to identify subtopic boundary sentences. We initialize decoder's hidden states with the document encoding ($d_{k=1}^h = D$ from eq. 3) and start decoding using embedding corresponding to the first sentence (S_1^C in eq. 1). At decoding step k , we calculate the subtopic boundary sentence probability following

eq. 2, where T_{k-1} is the index of $(k-1)^{th}$ boundary sentence.

$$\begin{aligned} d_k^h &= LSTMCell(d_{k-1}^h, S_{T_{k-1}}^C) \\ u_i^k &= [W_p^1(S_i^C) * W_p^2(d_k^h); W_p^1(S_i^C) - W_p^2(d_k^h)] \\ score_i^k &= \begin{cases} v_p^T \tanh(u_i^k), & i > T_{k-1} \\ -\infty, & \text{otherwise} \end{cases} \quad (2) \\ p(T_k|T_1, \dots, T_{k-1}; H^C, \dots, S_n^C) &= softmax(score^k) \end{aligned}$$

The pointer-decoder network (eq. 2) along with contextualized sentence encoder (eq. 1) constitute our f_T model. Note that we do not a priori know the number of subtopic boundary sentences in a document. Therefore, we append a special sentence “*eod*” at the end of each document, which when sampled as subtopic boundary sentence indicates the end of subtopic boundary sentence decoding.

4.3 Discourse Profiling

Given the list of subtopic boundary sentences T_L and contextualized sentence representations $[H^C, S_1^C, \dots, S_n^C]$, we use scalar soft-attentions (α_s) over sentence representation, as described in eq. 3, to learn local subtopic (T) and global document (D) representations. Finally, we combine sentence, local subtopic and document representations through element-wise product and differences (u_i) and use a two-layered feed-forward neural network to predict the labels. The networks defined in eq. 3 together with the contextualized sentence encoding network in equation 1 make the discourse profiling network f_C .

$$\begin{aligned} \alpha_s[i] &= W_{s1}(\tanh(W_{s2}S_i^C + b_{s2})) + b_{s1} \in R \\ A_T &= softmax(\alpha_s[T_L[j] : T_L[j+1]]) \in R^{T_L[j]-T_L[j+1]} \\ T &= \sum_{k=T_L[j]}^{T_L[j+1]} A_T[k].S^C[k] \in R^{2d_{rnn}} \\ A_s &= softmax(\alpha_s) \in R^n \\ D &= \sum_i A_s[i].H_s[i] \in R^{2d_{rnn}} \\ u_i &= [S_i^C - T; S_i^C * T; T - D; T * D] \in R^{8d_{rnn}} \\ \hat{y}_i &= softmax(W_{c1}(\tanh(W_{c2}u_i + b_{c2})) + b_{c1}) \in R^9 \quad (3) \end{aligned}$$

4.4 Learning f_T through Subtopic Structures-guided Critic

Our goal is to train the neural network-based subtopic boundary sentence scorer f_T model using indirect supervision derived from the performance

of f_C on discourse profiling task. Intuitively, REINFORCE algorithm (Williams, 1992), that has shown success in a range of NLP tasks, offers a suitable mechanism to train our f_T model. However, the vanilla reinforce is known to suffer from the problem of high variance. In addition, vanilla reinforce is incapable of inducing any known subtopic structure into the f_T . Therefore, we propose a new variation to the *actor-critic* (Konda and Tsitsiklis, 2000) model which defines multiple critics, each using a known subtopic structure, and trains in either of the imitation or reinforcement learning mode, depending on the performance of f_C classifier with known subtopic structures or the subtopic boundary sentences predicted by f_T .

Specifically, we consider f_T as the actor network that samples subtopic boundary sentences (T_S) following eq. 2. Then, we use sampled subtopic boundary sentences to partition the news document and use eq. 3 to identify content-types $\hat{Y} : \{\hat{y}_1, \dots, \hat{y}_n\}$ for all the sentences. We calculate the average of micro and macro F1 scores of the predicted content types \hat{Y} and use that as the reward R_A for our actor network. Following the same steps with reference subtopic boundary sentences T_R^j that are derived from a known subtopic structures (j^{th}), we also obtain the reward R_C^j for all our critics. Next, if the actor’s reward exceeds all the critics rewards ($R_A > R_C^j \forall j$), we use the reinforcement learning formulation and train f_T using the L_{RL} loss in eq. 4. Alternatively, if actor’s reward is lower than any of the critic, we use imitation learning with the cross-entropy loss (L_{IL}) based on the critic with maximum reward. The discourse profiling classifier f_C is trained using the standard cross-entropy loss on discourse profiling task (L_C).

$$\begin{aligned} L_{RL} &= (R_A - \bar{R}_C) \left(\sum_i -\log \frac{\exp(T_S[i])}{\sum_{T_k \in T_S[i-1:]} \exp(T_k)} \right) \\ \bar{R}_C &= \sum_{j=1}^J R_C^j / J \\ L_{IL} &= \sum_i -\log \frac{\exp(T_R[i])}{\sum_{T_k \in T_R[i-1:]} \exp(T_k)} \\ T_R &= argmax_{T_R^j} (R_C^j) \\ L_C &= \sum_i \sum_{c \in labels} -y_i^c \log(\hat{y}_i^c) \quad (4) \end{aligned}$$

At every iteration, the RL loss term forces the f_C model to perform at least as good as the model with a known subtopic structure and its reference

subtopic boundary sentences T_R that yield highest critic reward. The f_C model thus converges to parameters that obtain higher reward than its counterpart with reference subtopic boundary sentences. For the f_T model, if it chooses good subtopic boundary sentences T_S , that give higher reward than T_R , it further increases the likelihood for T_S . However, if it chooses bad T_S , the imitation learning loss forces the model to mimic subtopic structure with highest reward.

5 Document-level Content Organization Structures

We experiment with two subtopic structures: 1) broad-genre topic segments and 2) news genre-specific inverted pyramid structure. Note that both topic segments and inverted pyramid structures are automatically identified through statistical model or a plausible rules-based surrogate, which may not perfectly represent subtopic boundaries.

5.1 TextTiling

TextTiling (Hearst, 1997) is a paragraph-level model of discourse structure based on the notion of subtopic shift. It uses lexical co-occurrence and distribution patterns to divide a document into a sequence of topically coherent segments. It is a widely used algorithm to find subtopics segments in text and presents an effective representation for document-level content organization structures. We use the implementation provided with *nlTK* (Bird et al., 2009) in our experiments.

5.2 Locally Inverted Pyramid

We also consider the inverted pyramid structure (Pottker, 2003), that is most often used in news media (Dai et al., 2018). It organizes the news content in decreasing order of relevance, placing the most relevant information at the top and then arranging the remaining details in an decreasing order of relevance. While the inverted pyramid is a global content organization structure, we made a simplifying assumption that a document consists of smaller sequences of segments that locally follow the inverted pyramid structure. We identify a sentence as representing a subtopic boundary if it breaks the non-increasing relevance order of preceding sentences, i.e. its relevance lies higher to its preceding sentences on the relevance scale. Given that the relevance order of sentences is not always

aligned with their textual order, it provides an accessible proxy to define subtopical boundaries.

Specifically, since the eight discourse content types align with the relevance order of content in a document, with main sentence ($M1$) being the most relevant and central to the document, followed by immediate consequences ($M2$), causes and general context ($C1, C2$) and then opinions and expectations ($D3, D4$), it allows us to use the content types of sentences to extract subtopic boundary sentences and partition a document into smaller subtopical segments. For instance, a main event sentence following context-informing or supportive contents will make the main event sentence a subtopic boundary sentence. With the above rationales, we first identify the first sentence of a document as a subtopic boundary sentence. Then, given a document and content labels $(x_i, y_i) \in X$, we identify new subtopic boundary sentences x_i following the rules defined in Algorithm 1. Note that we dissociated historical ($D1$) and anecdotal ($D2$) content types from the relevance ordering as they are frequently used to set the tone for a news article or to highlight main argument with personal experiences or historical events.

Algorithm 1 Rules to identify subtopic boundary sentences for the Inverted Pyramid structure

- 1: $x_i \in \{M1, M2\}$ and $x_{i-1} \in \{C1 - D4\}$
 - 2: $x_i \in \{C1, C2\}$ and $x_{i-1} \in \{D1 - D4\}$
 - 3: $x_i \in \{D1, D2\}$ and $x_{i-1} \notin \{D1, D2\}$
 - 4: $x_i \notin \{D1, D2\}$ and $x_{i-1} \in \{D1, D2\}$
-

6 Evaluation

6.1 Dataset

We evaluate our content organization structure-aware model on the NewsDiscourse Corpus². It consists of 802 English news articles taken from three different news sources, NYT, Xinhua and Reuters, and covers business, crime, disaster and politics domains. Each sentence in a news document is annotated with one of the eight discourse content types (described in §3), and additionally speech and non-speech labels.³ We used 502 documents for training, 100 documents for validation

²The dataset was obtained from https://github.com/prafulla77/Discourse_Profiling

³Speech labels are not related to the discourse profiling structure in a news document. Therefore, in our experiments, we only focus on classifying a sentence into one of the eight discourse content types.

and 200 documents for testing, following the standard splits provided with the dataset. Models are evaluated on F1 score for each content type as well as micro F1 and macro P/R/F1 scores using the implementation provided by the scikit-learn (Pedregosa et al., 2011) library.

6.2 Baseline Models

Hierarchical: uses the hierarchical neural network architecture as proposed by Choubey et al. (2020) to learn sentence and document encodings and model associations between each sentence and the document encoding.

Self-Critic: uses the output of f_T network under the test-time inference algorithm (Rennie et al., 2017) to identify subtopic boundary sentences. Specifically, we take the argmax over the probability $p(T_k|T_1, \dots, T_{k-1}; H^C, \dots, S_n^C)$ (eq. 2) at k^{th} decoding step to identify the k^{th} subtopic boundary sentence. This model learns to build content structures entirely from the indirect supervision signal obtained from the f_C model, average of micro and macro F1 scores on the discourse profiling task.

6.3 Subtopic Structure-aware Models

TextTiling: directly uses the output of TextTiling model to partition documents and is trained on the loss (L_C) for discourse profiling task alone. The model does not include pointer network for identifying subtopic boundaries. Besides, it uses all other neural components and is structurally similar to RL-based models.

Joint-IP: learns to jointly identify subtopic boundary sentences, defined with rules in §5.2 to induce local inverted pyramid structure, and predict content types. It replaces pointer network with a two layer feed-forward neural network to identify subtopic boundary sentences. The model is trained on average of L_C and a binary cross entropy loss over subtopic boundary sentences.

RL-TT and RL-IP: use single critic defined through TextTiling and locally inverted pyramid subtopic structure respectively.

RL-IP/TT: uses two critics, first defined through TextTiling and second defined through locally inverted pyramid subtopic structure.

6.4 Implementation Details

Both biLSTMs in hierarchical sentence encoder have hidden dimension of 512. Decoder LSTM

Models	Macro			Micro
	P	R	F1	F1
Hierarchical	55.60	51.10	51.70	58.24
Self-Critic	58.61	50.09	51.87	57.65
TextTiling	53.72	52.13	51.47	57.62
Joint-IP	55.74	51.34	52.45	58.65
RL-TT	57.67	52.91	53.02	58.12
RL-IP	56.04	53.76	54.15	59.07
RL-IP/TT	56.42	55.20	54.42	59.21

Table 1: Results for the best-performing systems on validation dataset.

has hidden dimension of 1024. All two-layered feed forward networks use 1024 hidden units, including all networks used for calculating scalar attention weights as well as networks used to predict subtopic boundary sentences and content-type labels. All models use fixed word-embeddings and are trained using Adam (Kingma and Ba, 2014) optimizer with the learning rate of 5e-5 and dropout rate (Srivastava et al., 2014) of 0.4 on the output activations of both BiLSTMs and all neural layers. The models are trained for 15 epochs and we use the epoch yielding the best validation performance.

Consistent with the experimental setup used by Choubey et al. (2020), we run each neural model ten times with random seeds and report the average performance. As reinforcement learning or neural networks in general are sensitive to random seeds, analyzing average results alleviates the influence of randomness and provides stable empirical results. Learning rate and dropout rate are identified using grid search. First, we search learning and dropout rates from [1e-3, 5e-4, 1e-4, 5e-5] and [0.4, 0.5, 0.6] respectively using the hierarchical model. Then, both learning and dropout rates are kept constant for all models. Each training run takes ~1200 seconds without any major increase in training time from the RL component. We use one document per training iteration. All experiments are performed on NVIDIA GTX 1080 Ti 11GB using PyTorch 1.2.0+cu92 (Paszke et al., 2019) and AllenNLP 0.8.3 (Gardner et al., 2017).

7 Results and Analysis

Tables 1 and 2 compare all models on the validation and test datasets respectively. First, on the validation dataset, the best performing self-critic, TextTiling, and joint-IP models perform similar to the hierarchical model. Only the joint-IP model could

Models	M1	M2	C1	C2	D1	D2	D3	D4	Macro			Micro
	F1								P	R	F1	F1
Hierarchical	49.6	27.9	22.5	58.1	64.1	48.1	67.4	57.6	56.9	53.7	54.4(± 0.80)	60.9(± 0.70)
Self-Critic	51.5	29.4	27.2	58.2	61.4	55.3	67.5	59.7	59.0	55.1	56.1(± 0.49)	61.6(± 0.71)
TextTiling	50.7	31.2	26.1	57.6	61.1	52.3	66.5	58.7	58.5	54.0	55.5(± 0.98)	60.6(± 1.40)
Joint-IP	52.2	27.6	27.9	58.5	62.7	52.0	67.3	59.4	59.0	54.2	55.8(± 0.56)	61.4(± 0.70)
RL-TT	51.8	29.2	28.5	57.9	63.2	55.7	67.5	60.1	59.1	55.4	56.6(± 0.46)	61.7(± 0.61)
RL-IP	52.0	28.1	28.9	58.7	62.6	56.4	67.4	60.6	59.3	55.3	56.7(± 0.37)	61.9(± 0.38)
RL-IP/TT	52.6	28.7	26.6	58.0	63.5	59.2	68.3	60.6	58.7	56.4	57.0(± 0.38)	62.2(± 0.59)

Table 2: Performance of different systems on test dataset. All results correspond to average of 10 training runs with random seeds. In addition, we report standard deviation for both macro and micro F1 scores. Statistical significance tests show that both the macro and micro F1 scores for RL-IP/TT model are significantly better compared to the hierarchical, self-critic, TextTiling and joint-IP models with $p < 0.05$ on paired t test (Dietterich, 1998). Similarly, the macro F1 scores for RL-TT and RL-IP models are significantly better compared to the hierarchical, TextTiling and Joint-IP models with $p < 0.05$.

obtain consistent improvements over the hierarchical model, improving macro and micro F1 scores by 0.75 and 0.41 points respectively. However, on the test dataset, based on the average performance of 10 runs, all three models outperform the hierarchical model. TextTiling and joint-IP models, that directly use explicit subtopic structures, yield 1.1 and 1.4 points improvement in the average macro F1 score respectively. The self-critic model that directly learns subtopic boundary sentences through reward defined using its performance on the discourse profiling task performs better than the models with explicit subtopic structures, outperforming hierarchical model by 1.7 and 0.7 points in average macro and micro F1 scores. The higher average performance for self-critic model, as evident from results in Table 2, can be partly attributed to the lower accuracy of the subtopic structure identification models. In addition, self-critic framework allows the model to identify subtopic boundary sentences by directly optimizing numerical performance thereby identifying the subtopic structures that are optimal with respect to the parameters of content-type classification model.

This is also evident from the improved performance of the three actor-critic models, RL-TT, RL-IP and RL-IP/TT. Actor-critic based learning helps models to learn to identify subtopic boundary sentences that are useful for news discourse profiling but not strictly same as the used subtopic structures. Overall, our best performing RL-IP/TT model yields 2.72 and 0.97 points higher macro F1 and micro F1 score over the hierarchical model on the validation set, with comparable margin of improvement on the average F1 scores on test dataset.

	RM-IP	RM-TT	IP-TT	RM-TT-IP
Overlap	324	236	139	83

Table 3: Subtopic boundary sentences overlap between TextTiling and inverted pyramid subtopic structures and RL-IP/TT model on validation dataset. There are total 952 subtopic boundary sentences identified by RL-IP/TT model, and 589 and 540 subtopic boundary sentences identified by inverted pyramid and TextTiling structures respectively.

The consistent improvement from different models provides evidence that learning latent content organization structure to further segment documents and modeling both local subtopic representations and global document representations leads to better main topic induction and achieve improved content-type classification performance.

Between TextTiling and inverted pyramid subtopic structures, we observe that the latter performs better in both joint learning and actor-critic learning frameworks. This is expected since TextTiling is a broad-genre subtopic structure while inverted pyramid structure is specific to the news articles. Further, our rules (§5.2) used to build local inverted pyramid structures directly correlate with the content types of sentences. It is also worth noting that jointly using TextTiling and inverted pyramid structures as critics performs better than each structure when used individually.

7.1 Distributions of subtopic boundary sentences

In Table 3, we examine distributional overlap among subtopic boundary sentences identified by RL-IP/TT model (RM) and the TextTiling (TT) and

	RM	IP	TT
Temporal frames	13	18	7

Table 4: Subtopic boundary sentences overlap between Temporal-frames based subtopic structure and TextTiling, inverted pyramid, and RL-IP/TT model on a subset of 10 documents from validation dataset. There are total 68 subtopic boundary sentences identified by RL-IP/TT model, and 79, 52 and 58 subtopic boundary sentences identified by inverted pyramid, TextTiling, and temporal frames-based structures respectively.

inverted pyramid (IP) subtopic structures on the validation dataset. We observe that subtopic boundary sentences identified by RL-IP/TT model exhibit higher overlap with the inverted pyramid (324) than the TextTiling structure (236). Interestingly, models based on inverted pyramid structure obtain better performance than the TextTiling-based models (Tables 1 and 2). The higher overlap for inverted pyramid structure corroborates its greater effectiveness in inducing appropriate subtopic structure for discourse profiling.

In addition, we manually annotated a subset of 10 documents from the validation dataset with subtopic structure based on temporal frames (as shown in Figure 1). As shown in Table 4, subtopic boundary sentences identified by RL-IP/TT model exhibit overlap with the temporal frames-based subtopic structure. This is not implausible given overlap between temporal frame-based subtopic structure and inverted pyramid, and TextTiling, as noted from Table 4. In a nutshell, subtopic boundary sentences identified by different subtopic structures exhibit partial overlap, and by using multiple critics to guide actor network in our actor critic formulation, we can enable the model to learn subtopic structure that is not necessarily identical to the used critics but more effective in profiling discourse structure for a given document.

8 Conclusion

We have presented a document-level content-organization structures aware neural network model for discourse profiling. We explored actor-critic learning based frameworks to induce subtopic structures in a news document. Then, we model two levels of interactions - between sentences and the local subtopic representation, and between subtopic representations and the document - that consistently outperformed the previous best hierarchical model. For future work, we intend to experi-

ment with new modeling techniques to incorporate explicit subtopic structures. Further, we plan to extend topical structures to other discourse structures such as rhetorical relations and model their inter-dependencies with the discourse profiling task.

9 Acknowledgements

We gratefully acknowledge support from National Science Foundation via the awards IIS-1942918 and IIS-2127746. We would also like to thank the anonymous reviewers for their feedback.

References

- Sebastian Arnold, Rudolf Schneider, Philippe Cudré-Mauroux, Felix A. Gers, and Alexander Löser. 2019. [SECTOR: A neural model for coherent topic segmentation and classification](#). *Transactions of the Association for Computational Linguistics*, 7:169–184.
- Dzmitry Bahdanau, Philemon Brakel, Kelvin Xu, Anirudh Goyal, Ryan Lowe, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2016a. [An actor-critic algorithm for sequence prediction](#). *arXiv e-prints*, abs/1607.07086.
- Dzmitry Bahdanau, Philemon Brakel, Kelvin Xu, Anirudh Goyal, Ryan Lowe, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2016b. [An actor-critic algorithm for sequence prediction](#). *arXiv preprint arXiv:1607.07086*.
- Maria Becker, Michael Staniek, Vivi Nastase, Alexis Palmer, and Anette Frank. 2017. Classifying semantic clause types: Modeling context and genre characteristics with recurrent neural networks and attention. In *Proceedings of the 6th Joint Conference on Lexical and Computational Semantics (* SEM 2017)*, pages 230–240.
- Allan Bell. 1998. The discourse structure of news stories. In *Approaches to media discourse*, pages 64–104.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. "O'Reilly Media, Inc."
- Yen-Chun Chen and Mohit Bansal. 2018. [Fast abstractive summarization with reinforce-selected sentence rewriting](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 675–686, Melbourne, Australia. Association for Computational Linguistics.
- Freddy Y. Y. Choi. 2000. [Advances in domain independent linear text segmentation](#). In *1st Meeting of the North American Chapter of the Association for Computational Linguistics*.

- Prafulla Kumar Choubey and Ruihong Huang. 2021. [Automatic data acquisition for event coreference resolution](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1185–1196, Online. Association for Computational Linguistics.
- Prafulla Kumar Choubey, Aaron Lee, Ruihong Huang, and Lu Wang. 2020. [Discourse as a function of event: Profiling discourse structure in news articles around the main event](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5374–5386, Online. Association for Computational Linguistics.
- Zeyu Dai and Ruihong Huang. 2018. [Improving implicit discourse relation classification by modeling inter-dependencies of discourse units in a paragraph](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 141–151, New Orleans, Louisiana. Association for Computational Linguistics.
- Zeyu Dai, Himanshu Taneja, and Ruihong Huang. 2018. [Fine-grained structure-based news genre categorization](#). In *Proceedings of the Workshop Events and Stories in the News 2018*, pages 61–67, Santa Fe, New Mexico, U.S.A. Association for Computational Linguistics.
- Thomas G Dietterich. 1998. Approximate statistical tests for comparing supervised classification learning algorithms. *Neural computation*, 10(7):1895–1923.
- Lan Du, Wray Buntine, and Mark Johnson. 2013. [Topic segmentation with a structured topic model](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 190–200, Atlanta, Georgia. Association for Computational Linguistics.
- Jacob Eisenstein and Regina Barzilay. 2008. [Bayesian unsupervised topic segmentation](#). In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 334–343, Honolulu, Hawaii. Association for Computational Linguistics.
- William Fedus, Ian Goodfellow, and Andrew M Dai. 2018. Maskgan: better text generation via filling in the_. *arXiv preprint arXiv:1801.07736*.
- Vincent F Filak. 2019. *Convergent Journalism: An Introduction: Writing and Producing Across Media*. Routledge.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke S. Zettlemoyer. 2017. [Allennlp: A deep semantic natural language processing platform](#).
- Marti A. Hearst. 1997. [Text tiling: Segmenting text into multi-paragraph subtopic passages](#). *Computational Linguistics*, 23(1):33–64.
- Stefan Henß, Margot Mieskes, and Iryna Gurevych. 2015. A reinforcement learning approach for adaptive single-and multi-document summarization. In *GSCL*, pages 3–12.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Yangfeng Ji and Jacob Eisenstein. 2014a. [Representation learning for text-level discourse parsing](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13–24, Baltimore, Maryland. Association for Computational Linguistics.
- Yangfeng Ji and Jacob Eisenstein. 2014b. [Representation learning for text-level discourse parsing](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13–24, Baltimore, Maryland. Association for Computational Linguistics.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Vijay R Konda and John N Tsitsiklis. 2000. Actor-critic algorithms. In *Advances in neural information processing systems*, pages 1008–1014.
- Fajri Koto, Jey Han Lau, and Timothy Baldwin. 2021. [Top-down discourse parsing via sequence labelling](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 715–726, Online. Association for Computational Linguistics.
- Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. 2016a. [Deep reinforcement learning for dialogue generation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202, Austin, Texas. Association for Computational Linguistics.
- 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, Austin, Texas. Association for Computational Linguistics.
- Xi Victoria Lin, Richard Socher, and Caiming Xiong. 2018. [Multi-hop knowledge graph reasoning with reward shaping](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3243–3253, Brussels, Belgium. Association for Computational Linguistics.

- Xiang Lin, Shafiq Joty, Prathyusha Jwalapuram, and M Saiful Bari. 2019. [A unified linear-time framework for sentence-level discourse parsing](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4190–4200, Florence, Italy. Association for Computational Linguistics.
- Yang Liu and Mirella Lapata. 2017. [Learning contextually informed representations for linear-time discourse parsing](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1289–1298, Copenhagen, Denmark. Association for Computational Linguistics.
- Inderjeet Mani, James Pustejovsky, and Robert Gaizauskas. 2005. *The language of time: a reader*. Oxford University Press on Demand.
- Mathieu Morey, Philippe Muller, and Nicholas Asher. 2017. [How much progress have we made on RST discourse parsing? a replication study of recent results on the RST-DT](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1319–1324, Copenhagen, Denmark. Association for Computational Linguistics.
- Jack Elliott Myers and Michael Simms. 1989. *The Longman dictionary of poetic terms*. Addison-Wesley Longman Ltd.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2017. A deep reinforced model for abstractive summarization. *arXiv preprint arXiv:1705.04304*.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Horst Pottker. 2003. News and its communicative quality: The inverted pyramid—when and why did it appear? *Journalism Studies*, 4(4):501–511.
- Pengda Qin, Weiran Xu, and William Yang Wang. 2018. [Robust distant supervision relation extraction via deep reinforcement learning](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2137–2147, Melbourne, Australia. Association for Computational Linguistics.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2015. Sequence level training with recurrent neural networks. *arXiv preprint arXiv:1511.06732*.
- Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. 2017. Self-critical sequence training for image captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7008–7024.
- Naveed Saleh. 2014. *The Complete Guide to Article Writing: How to Write Successful Articles for Online and Print Markets*. "F+ W Media, Inc."
- Christina Schokkenbroek. 1999. News stories: structure, time and evaluation. *Time & Society*, 8(1):59–98.
- Imran Sehikh, Dominique Fohr, and Irina Illina. 2017. Topic segmentation in asr transcripts using bidirectional rnns for change detection. In *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 512–518. IEEE.
- 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.
- Van Dijk Teun A. 1986. News schemata. *Studying writing: linguistic approaches*, 1:155–186.
- Masao Utiyama and Hitoshi Isahara. 2001. [A statistical model for domain-independent text segmentation](#). In *Proceedings of the 39th Annual Meeting of the Association for Computational Linguistics*, pages 499–506, Toulouse, France. Association for Computational Linguistics.
- Teun A Van Dijk. 1985. Structures of news in the press. *Discourse and communication: New approaches to the analysis of mass media discourse and communication*, 10:69.
- Teun A Van Dijk. 1988a. News analysis. *Case Studies of International and National News in the Press*. New Jersey: Lawrence.
- Teun A Van Dijk. 1988b. *News as discourse*. Hillsdale, NJ, US: Lawrence Erlbaum Associates, Inc.

- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. [Pointer networks](#). In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.
- Liang Wang, Sujian Li, Yajuan Lv, and Houfeng Wang. 2017. [Learning to rank semantic coherence for topic segmentation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1340–1344, Copenhagen, Denmark. Association for Computational Linguistics.
- Peter White. 1997. Death, disruption and the moral order: the narrative impulse in mass-media ‘hard news’ reporting. *Genres and institutions: Social processes in the workplace and school*, 101:133.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.
- Caiming Xiong, Victor Zhong, and Richard Socher. 2017. Dcn+: Mixed objective and deep residual coattention for question answering. *arXiv preprint arXiv:1711.00106*.
- Espen Ytreberg. 2001. Moving out of the inverted pyramid: narratives and descriptions in television news. *Journalism Studies*, 2(3):357–371.
- Longyin Zhang, Yuqing Xing, Fang Kong, Peifeng Li, and Guodong Zhou. 2020. [A top-down neural architecture towards text-level parsing of discourse rhetorical structure](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6386–6395, Online. Association for Computational Linguistics.