

Neural Abstractive Summarization with Structural Attention

Tanya Chowdhury¹, Sachin Kumar² and Tanmoy Chakraborty¹

¹IIIT-Delhi, India

²Carnegie Mellon University, USA

{tanya14109,tanmoy}@iiitd.ac.in, sachink@andrew.cmu.edu

Abstract

Attentional, RNN-based encoder-decoder architectures have achieved impressive performance on abstractive summarization of news articles. However, these methods fail to account for long term dependencies within the sentences of a document. This problem is exacerbated in multi-document summarization tasks such as summarizing the popular opinion in threads present in community question answering (CQA) websites such as Yahoo! Answers and Quora. These threads contain answers which often overlap or contradict each other. In this work, we present a hierarchical encoder based on structural attention to model such inter-sentence and inter-document dependencies. We set the popular pointer-generator architecture and some of the architectures derived from it as our baselines and show that they fail to generate good summaries in a multi-document setting. We further illustrate that our proposed model achieves significant improvement over the baselines in both single and multi-document summarization settings – in the former setting, it beats the best baseline by 1.31 and 7.8 ROUGE-1 points on CNN and CQA datasets, respectively; in the latter setting, the performance is further improved by 1.6 ROUGE-1 points on the CQA dataset.

1 Introduction

Sequence-to-sequence (seq2seq) architectures with attention have led to tremendous success in many conditional language generation tasks such as machine translation [Sutskever *et al.*, 2014; Bahdanau *et al.*, 2015; Luong *et al.*, 2015], dialog generation [Vinyals and Le, 2015], abstractive summarization [Rush *et al.*, 2015; See *et al.*, 2017; Paulus *et al.*, 2018] and question answering [Yin *et al.*, 2016]. Attention models that compute a context vector based on the entire input text at every decoder step have especially benefited abstractive summarization. This is because – (1) LSTM encoders struggle to encode long documents into just one vector from which one can decode the summary, and (2) generating summary involves a lot of copying from the source text which attention models easily facilitate that can

also be made part of the training process [See *et al.*, 2017; Nallapati *et al.*, 2016].

Studies in abstractive summarization majorly focus on generating news summaries, and the CNN/Dailymail dataset [Hermann *et al.*, 2015; Nallapati *et al.*, 2016] is largely used as a benchmark. This is a very clean dataset with fluent and properly structured text collected from a restricted domain, and it hardly features any repetitive information. Therefore, the current models which have been built on this dataset do not need to account too much for repeated or contradicting information. Additionally, these models are trained on truncated documents (up to 400 words), and increasing the word length leads to a drop in performance of these models [See *et al.*, 2017]. This is because either LSTM is unable to encode longer documents accurately or in this dataset, the first few sentences contain most of the summary information. However, real-world documents are usually much longer, and their summary information is spread across the entire document rather than in the first few sentences. In this paper, we explore one such dataset dealing with community question answering (CQA) summarization.

Multi-document summarization (MDS) is a well-studied problem with possible applications in summarizing tweets [Cao *et al.*, 2016], news from varying sources [Yasunaga *et al.*, 2017], reviews in e-commerce services [Ganesan *et al.*, 2010], and so on. CQA services such as *Quora*, *Yahoo! Answers*, *Stack Overflow* help in curating information which may often be difficult to obtain directly from the existing resources on the web. However, a majority of question threads on these services harbor unstructured, often repeating or contradictory answers. Additionally, the answers vary in length from being a few words to a thousand words long. A large number of users with very different writing styles contribute to form these knowledge bases, which results in a diverse and challenging corpora. We define the task of *CQA summarization* as follows – given a list of such question threads, summarize the popular opinion reflected in them. This can be viewed as a multi-document summarization task [Chowdhury and Chakraborty, 2019].

We envision a solution to this problem by incorporating structural information present in language in the summarization models. Recent studies have shown promise in this direction. For example, [Fernandes *et al.*, 2018] add graph neural networks (GNNs) [Li *et al.*, 2015] to the seq2se en-

coder model for neural abstractive summarization. [Liu *et al.*, 2019] model documents as multi-root dependency trees, each node representing a sentence; and pick the induced tree roots to be summary sentences in an extractive setting. Recently, [Song *et al.*, 2018] attempt to preserve important structural dependency relations obtained from syntactic parse trees of the input text in the summary by using a structure infused copy mechanism. However, obtaining labeled data to explicitly model such dependency relations is expensive. Here, we attempt to mitigate this issue by proposing a structural encoder for summarization based on the prior work on structure attention networks [Liu and Lapata, 2018] which implicitly incorporates structural information within end-to-end training.

Our major contributions in this work are three-fold:

1. We enhance the pointer-generator architecture by adding a structural attention based encoder to implicitly capture long term dependency relations in summarization of lengthy documents.
2. We further propose a hierarchical encoder with multi-level structural attention to capture document-level discourse information in the multi-document summarization task.
3. We introduce multi-level contextual attention in the structural attention setting to enable word level copying and to generate more abstractive summaries, compared to similar architectures.

We compare our proposed solution against the popular pointer-generator model [See *et al.*, 2017] and a few other summarization models derived from it in both single and multi-document settings. We show that our structural encoder architecture beats the strong pointer-generator baseline by 1.31 ROUGE-1 points on the CNN/Dailymail dataset and by 7.8 ROUGE-1 points on the concatenated CQA dataset for single document summarization (SDS). Our hierarchical structural encoder architecture further beats the concatenated approach by another 1.6 ROUGE-1 points on the CQA dataset. A qualitative analysis of the generated summaries observe considerable qualitative gains after inclusion of structural attention. Our structural attention based summarization model is one of the few abstractive approaches to beat extractive baselines in MDS. The code is public at <https://bit.ly/35i7q93>.

2 Models

In this section, we first describe the seq2seq based pointer-generator architecture proposed by [See *et al.*, 2017; Nallapati *et al.*, 2016]. We then elaborate on how we incorporate structural attention into the mechanism and generate non-projective dependency trees to capture long term dependencies within documents. Following this, we describe our hierarchical encoder architecture and propose multi-level contextual attention to better model discourse for MDS.

2.1 Pointer-Generator Model

The pointer-generator architecture (PG) [See *et al.*, 2017] serves as the starting point for our approach. PG is inspired

from the LSTM encoder-decoder architecture [Nallapati *et al.*, 2016] for the task of abstractive summarization. Tokens are fed to a bidirectional encoder to generate hidden representations h , and a unidirectional decoder is used to generate the summary. At the decoding step t , the decoder uses an attention mechanism [Bahdanau *et al.*, 2015] to compute an attention vector a^t for every encoder hidden state. These attention vectors are pooled to compute a context vector c_t for every decoder step. The context vector c_t at that step along with the hidden representations h are passed through a feed-forward network followed by softmax to give a multinomial distribution over the vocabulary $p_{\text{vocab}}(w)$.

Additionally, [See *et al.*, 2017] introduce a copy mechanism [Vinyals *et al.*, 2015] in this model. At each decoding step t , it predicts whether to generate a new word or copy the word based on the attention probabilities a^t using a probability value p_{gen} . This probability is a function on the decoder input x_t , context vector c_t , and decoder state s_t . p_{gen} is used as a soft switch to choose between copying and generation as shown below:

$$p(w) = p_{\text{gen}}p_{\text{vocab}}(w) + (1 - p_{\text{gen}}) \sum_{i:w_i=w} a_i^t$$

This model is trained by minimizing the negative log-likelihood of the predicted sequence normalized by the sequence length: $L = -\frac{1}{T} \sum_{t=0}^T \log p(w_t^*)$, where T is the length of the generated summary. Once L gets converged, an additional coverage loss is used to train the model further, which aims to minimize repetitions in the generated summaries. We refer the readers to [See *et al.*, 2017] for more details.

2.2 Structural Attention

We observe that while the pointer-generator architecture performs well on the news corpora, it fails to generate meaningful summaries on complex datasets. We hypothesize that taking into account intra-document discourse information will help in generating better summaries. We propose an architecture that implicitly incorporates such information using document structural representation [Liu and Lapata, 2018; Kim *et al.*, 2017] to include richer structural dependencies in the end-to-end training.

Proposed Model

We model our document as a non-projective dependency parse tree by constraining inter-token attention as weights of the dependency tree. We use the Matrix tree theorem [Koo *et al.*, 2007; Tutte, 1984] to carry out the same.

As shown in Figure 1, we feed our input tokens (w_i) to a bi-LSTM encoder to obtain hidden state representations h_i . We decompose the hidden state vector into two parts: d_i and e_i , which we call the structural part and the semantic part, respectively:

$$[e_i, d_i] = h_i \tag{1}$$

For every pair of two input tokens, we transform their structural parts d and try to compute the probability of a parent-child relationship edge between them in the dependency tree. For tokens j and k , this is done as:

$$u_j = \tanh(W_p d_j); \quad u_k = \tanh(W_c d_k)$$

where W_p and W_c are learned. Next, we compute an inter-token attention function f_{jk} as follows:

$$f_{jk} = u_k^T W_a u_j$$

where W_a is also learned. For a document with K tokens, f is a $K \times K$ matrix representing inter-token attention. We model each token as a node in the dependency tree and define the probability of an edge between tokens at positions j and k , $P(z_{jk} = 1)$, which is given as,

$$A_{jk} = \begin{cases} 0 & \text{if } j = k \\ \exp(f_{jk}) & \text{otherwise} \end{cases}$$

$$L_{jk} = \begin{cases} \sum_{j'=1}^K A_{j'k} & \text{if } j = k \\ -A_{jk} & \text{otherwise} \end{cases}$$

$$f_j^r = W_r d_j$$

$$\bar{L}_{jk} = \begin{cases} \exp(f_j^r) & j = 1 \\ L_{jk} & j > 1 \end{cases}$$

$$P(z_{jk} = 1) = (1 - \delta(j, k)) A_{jk} \bar{L}_{kk}^{-1} - (1 - \delta(j, 1)) A_{jk} \bar{L}_{kj}^{-1}$$

where $\delta(x, y) = 1$ when $x = y$. We denote $P(z_{jk} = 1)$ by a_{jk} (structural attention). Let a_j^r be the probability of the j^{th} token to be the root:

$$a_j^r = \exp(W_r d_j) \bar{L}_{j1}^{-1} \quad (2)$$

We use this (soft) dependency tree formulation to compute a structural representation r for each encoder token as,

$$s_i = a_i^r e_{\text{root}} + \sum_{k=1}^n a_{ki} e_k; \quad c_i = \sum_{k=1}^n a_{ik} e_i$$

$$r_i = \tanh(W_r [e_i, s_i, c_i])$$

Thus, for encoder step i , we now obtain a structure infused hidden representation r_i . We then compute the contextual attention for each decoding time step t as,

$$e_{\text{struct}_i}^t = v^t \tanh(W_r r_i + W_s s_t + b_{\text{attn}})$$

$$a_{\text{struct}_i}^t = \text{softmax}(e_{\text{struct}_i}^t)$$

Now, using $a_{\text{struct}_i}^t$, we can compute a context vector similar to standard attentional model by weighted sum of the hidden state vectors as $c_{\text{struct}_i}^t = \sum_{i=1}^n a_{\text{struct}_i}^t h_i$. At every decoder time step, we also compute the basic contextual attention vector a^t (without structure incorporation), as discussed previously. We use $c_{\text{struct}_i}^t$ to compute P_{vocab} and p_{gen} . We, however, use the initial attention distribution a_t to compute $P(w)$ in order to facilitate token level pointing.

2.3 Hierarchical Model

The structural attention model, while efficient, requires $\mathcal{O}(K^2)$ memory to compute attention, where K is the length of the input sequence making it memory intensive for long documents. Moreover, in CQA, answers reflect different opinions of individuals. In such a case, concatenating answers results in user specific information loss. To model discourse structure better in case of conflicting (where one answer contradicts other), overlapping or varying opinion, we




Figure 1: Structural attention based encoder architecture. Tokens (w_i) are independently fed to a bi-LSTM encoder, and hidden representations h_i are obtained. These representations are split into structural (d_i) and semantic (e_i) components. The structural component is used to build a non-projective dependency tree. The inter-token attention (a_{jk}) is computed based on the marginal probability of an edge between two nodes in the tree ($P(z_{jk} = 1)$).

introduce a hierarchical encoder model based on structural attention (Figure 2). We feed each answer independently to a bi-LSTM encoder and obtain token level hidden representations h_{idx, tid_x} where idx is the document (or answer) index and tid_x is the token index. We then transform these representations into structure infused vectors r_{idx, tid_x} as described previously. For each answer in the CQA question thread, we pool the token representations to obtain a composite answer vector r_{idx} . We consider three types of pooling: *average*, *max* and *sum*, out of which *sum* pooling performs the best after initial experiments. We feed these structure infused answer embeddings to an answer-level bi-LSTM encoder and obtain higher-level encoder hidden states h_{idx} from which we calculate structure infused embeddings g_{idx} . At the decoder time step t , we calculate contextual attention at answer as well as token level as follows:

$$e_{\text{ans}_{idx}}^t = v^t \tanh(W_g g_{idx} + W_s s_t + b_{\text{attn}})$$

$$a_{\text{ans}}^t = \text{softmax}(e_{\text{ans}}^t)$$

$$e_{\text{token}_{idx, tid_x}}^t = v^t \tanh(W_h h_{idx, tid_x} + W_s s_t + b_{\text{attn}})$$

$$a_{\text{token}}^t = \text{softmax}(e_{\text{token}}^t)$$

We use the answer-level attention distribution a_{ans}^t to compute the context vector at every decoder time step which we use to calculate p_{vocab} and p_{gen} as described before.

To enable copying, we use a_{token} . The final probability of predicting word w is given by,

$$p(w) = p_{\text{gen}} p_{\text{vocab}}(w) + (1 - p_{\text{gen}}) \sum_{i:w_i=w} a_{\text{token}_i}^t$$

3 Datasets

We primarily use two datasets to evaluate the performance:

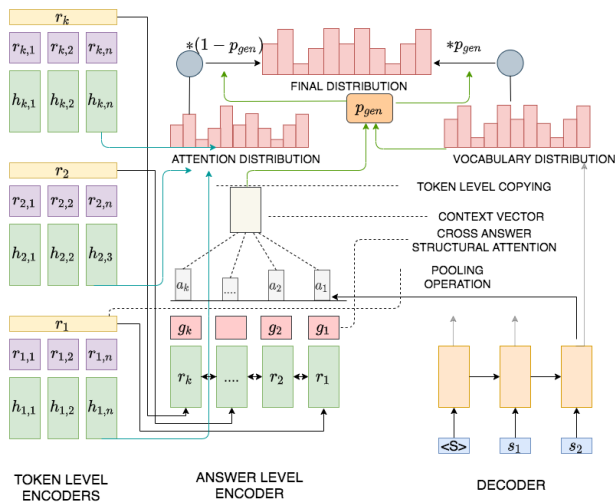


Figure 2: Hierarchical encoder with structural attention + multi level contextual attention for multi-document summarization. First, token-level structured representations are computed. For each answer, these representations are pooled (r_i) and fed to an answer-level encoder. Structural infused answer level embeddings (g_i) are calculated across answers. At every decoder time step, we calculate the context vector and probability p_{gen} based on structural contextual attention. We calculate the copying vector (source attention) by projecting token-level attention onto the vocabulary.

(i) The **CNN/Dailymail**¹ dataset [Hermann *et al.*, 2015; Nallapati *et al.*, 2016] is a news corpora containing document-summary pairs. Bulleted extracts from the CNN and Dailymail news pages online are projected as summaries to the remaining documents. The scripts released by [Nallapati *et al.*, 2016] are used to extract approximately 250k training pairs, 13k validation pairs and 11.5k test pairs from the corpora. We use the non-anonymized form of the data to provide fair comparison to the experiments, conducted by [See *et al.*, 2017]. It is a factual, English corpora with an average document length being 781 tokens and an average summary length being 56 tokens. We use two versions of this dataset – one with 400 word articles, and the other with 800 word articles. Most research reporting results on this dataset are on CNN/Dailymail-400. We also consider a 800 token version of this dataset as longer articles harbor more intra-document structural dependencies which would allow us to better demonstrate the benefits of structure incorporation. Moreover, longer documents resemble real-world datasets. (ii) We also use the **CQA dataset**² [Chowdhury and Chakraborty, 2019] which is generated by filtering the Yahoo! Answers L6 corpora to find question threads where the best answer can serve as a summary for the remaining answers. The authors use a series of heuristics to arrive at a set of 100k question thread-summary pairs. The summaries are generated by modifying best answers and selecting most question-relevant sentences from them. The remaining answers serve as candidate documents for summarization mak-

¹<https://github.com/abisee/cnn-dailymail>

²<https://bitbucket.org/tanya14109/cqasumm/src/master/>

ing up a large-scale, diverse and highly abstract dataset. On an average, the corpus has 12 answers per question thread, with 65 words per answer. All summaries are truncated at hundred words. We split the 100k dataset into 80k training instances, 10k validation and 10k test instances. We additionally extract the upvote information corresponding to every answer from the L6 dataset and assume that upvotes have a high correlation with the relative importance and relevance of an answer. We then rank answers in decreasing order of upvotes before concatenating them as required for several baselines. Since Yahoo! Answers is an unstructured and unmoderated question-answer repository, this has turned out to be a challenging summarization dataset [Chowdhury and Chakraborty, 2019]. Additionally, we include analysis on MultiNews [Fabri *et al.*, 2019], a news based MSD corpora to aid similar studies. It is the first large scale MSD news dataset consisting of 56,216 article summary pairs, crowd-sourced from various different news websites.

4 Competing Methods

We compare the performance of the following SDS and MDS models.

- **Lead3**: It is an extractive baseline where the first 100 tokens of the document (in case of SDS datasets) and concatenated ranked documents (in case of MDS datasets) are picked to form the summary.
- **KL-Summ**: It is an extractive summarization method introduced by [Haghighi and Vanderwende, 2009] that attempts to minimize KL-Divergence between candidate documents and generated summary.
- **LexRank**: It is an unsupervised extractive summarization method [Erkan and Radev, 2004]. A graph is built with sentences as vertices, and edge weights are assigned based on sentence similarity.
- **TextRank**: It is an unsupervised extractive summarization method which selects sentences such that the information being disseminated by the summary is as close as possible to the original documents [Mihalcea and Tarau, 2004].
- **Pointer-Generator (PG)**: It is a supervised abstractive summarization model [See *et al.*, 2017], as discussed earlier. It is a strong and popularly used baseline for summarization.
- **Pointer-Generator + Structure Infused Copy (PG+SC)**: Our implementation is similar to one of the methods proposed by [Song *et al.*, 2018]. We explicitly compute the dependency tree of sentences and encode a structure vector based on features like POS tag, number of incoming edges, depth of the tree, etc. We then concatenate this structural vector for every token to its hidden state representation in pointer-generator networks.
- **Pointer-Generator+MMR (PG+MMR)**: It is an abstractive MDS model [Lebanoff *et al.*, 2018] trained on the CNN/Dailymail dataset. It combines Maximal Marginal Relevance (MMR) method with pointer-

Parameter	Value
Vocabulary size	50,000
Input embedding dim	128
Training decoder steps	100
Learning Rate	0.15
Optimizer	Adagrad
Adagrad Init accumulator	0.1
Max gradient norm (for clipping)	2.0
Max decoding steps (for BS decoding)	120
Min decoding steps (for BS decoding)	35
Beam search width	4
Weight of coverage loss	1
GPU	GeForce 2080 Ti

Table 1: Parameters common to all PG-based models.

Method	CNN/Dailymail-400			CNN/Dailymail-800		
	R-1	R-2	R-L	R-1	R-2	R-L
Lead3	40.34	17.70	36.57	40.34	17.70	36.57
KL-Summ	30.50	11.31	28.54	28.42	10.87	26.06
LexRank	34.12	13.31	31.93	32.36	11.89	28.12
TextRank	31.38	12.29	30.06	30.24	11.26	27.92
PG	39.53	17.28	36.38	36.81	15.92	32.86
PG+SC	39.82	17.68	36.72	37.48	16.61	33.49
PG+Transformers	39.94	36.44	36.61	-	-	-
PG+SA	40.02	17.88	36.71	38.15	16.98	33.20

Table 2: Performance on two versions of the CNN/Dailymail dataset based on ROUGE-1 (R-1), ROUGE-2 (R-2) and ROUGE-L (R-L) scores.

Method	CQASUMM			MultiNews		
	R-1	R-2	R-L	R-1	R-2	R-L
Lead3	8.7	1.0	5.2	39.41	11.77	14.51
KL-Summ	24.4	4.3	13.9	-	-	-
LexRank	28.4	4.7	14.7	38.27	12.70	13.20
TextRank	27.8	4.8	14.9	38.40	13.10	13.50
PG	23.2	4.0	14.6	41.85	12.91	16.46
PG + SC	25.5	4.2	14.8	-	-	-
PG+MMR	16.2	3.2	10	36.42	9.36	13.23
Hi-MAP	30.9	4.69	14.9	43.47	14.87	17.41
PG+SA	29.4	4.9	15.3	43.24	13.44	16.9
PG+HSA	31.0	5.0	15.2	43.49	14.02	17.21

Table 3: Performance on the CQA and MultiNews datasets.

generator networks, and shows significant performance on DUC-04 and TAC-11 datasets.

- **Hi-MAP:** It is an abstractive MDS model by [Fabbri *et al.*, 2019] extending PG and MMR.
- **Pointer-Generator + Structural Attention (PG+SA):** It is the model proposed in this work for SDS and MDS tasks. We incorporate structural attention with pointer generator networks and use multi-level contextual attention to generate summaries.
- **Pointer-Generator + Hierarchical Structural Attention (PG+HSA):** We use multi-level structural attention to additionally induce a document-level non-projective dependency tree to generate more insightful summaries.

5 Experimental Results

5.1 Quantitative Analysis

We compare the performance of the models on the basis of ROUGE-1, 2 and L F1-scores [Lin, 2004] on the CNN/Dailymail (Table 2), CQA and Multinews (Table 3) datasets. We observe that infusion of structural information leads to considerable gains over the basic PG architecture on both the datasets. Our approach fares better than explicitly incorporating structural information [Song *et al.*, 2018]. The effects of incorporating structure is more significant in the CQA dataset (+7.8 ROUGE-1 over PG) as compared to the CNN/Dailymail dataset (+0.49 ROUGE-1 over PG). The benefit is comparatively prominent in CNN/Dailymail-800 (+1.31 ROUGE-1 over PG) as compared to CNN-Dailymail-400 (+0.49 ROUGE-1 over PG).

5.2 Qualitative Analysis

We ask human evaluators³ to compare CQASUMM summaries generated by competing models on the ground of content and readability. Here we observe a significant gain with structural attention. PG has difficulty in summarizing articles with repetitive information and tends to assign lower priority to facts repeating across answers. Methods like LexRank, on the other hand, tend to mark these facts the most important. Incorporating structure solves this problem to some extent (due to the pooling operations). We find that PG sometimes picks sentences containing opposing opinions for the same summary. We find this occurrence to be less frequent in the structural attention models. This phenomenon is illustrated with an instance from the CQA dataset in Table 4.

5.3 Model Diagnosis and Discussion

Tree depth. The structural attention mechanism has a tendency to induce shallow parse trees with high tree width. This leads to highly spread out trees especially at the root.

MDS baselines. While supervised abstractive methods like PG significantly outperform unsupervised non-neural methods in SDS, we find them to be inferior or similar in performance in MDS. This has also been reported by recent MDS related studies such as [Lebanoff *et al.*, 2018] where LexRank is shown to significantly beat basic PG on the DUC-04 and TAC-11 datasets. This justifies the choice of largely unsupervised baselines in recent MDS related studies such as [Yasunaga *et al.*, 2017; Lebanoff *et al.*, 2018].

Dataset layout. We observe that while increasing the length of the input documents lowers ROUGE score for the CNN/Dailymail dataset, it boosts the same in the CQA dataset. This might be attributed to the difference in information layout in the two datasets. The CNN/Dailymail has a very high Lead-3 score signifying summary information concentration within the first few lines of the document.

Coverage with structural attention. Training a few thousand iterations with coverage loss is known to significantly reduce repetition. However, in the structural attention models, while training with the CQA dataset we observe that soon

³The evaluators were experts in NLP, and their age ranged between 20-30 years.

after coverage loss has converged, on further training the repetition in generated summaries starts increasing again. Future work would include finding a more appropriate coverage function for datasets with repetition.

6 Related Work

Neural Abstractive Summarization

PG-based models. Since the studies conducted by [Nallapati *et al.*, 2016; See *et al.*, 2017], many approaches to neural abstractive summarization have extended them in different ways. [Cohan *et al.*, 2018] extend the approach proposed by [See *et al.*, 2017] with a hierarchical encoder to model discourse between various sections of a scientific paper using PubMed dataset.

MDS. Most studies on MDS were performed during DUC (2001-2007) and TAC (2008-2011) workshops. Opinosis [Ganesan *et al.*, 2010] uses word-level opinion graphs to find cliques and builds them into summaries in highly redundant opinion sentences. Recently, [Lebanoff *et al.*, 2018] adapt the single document architecture to a multi-document setting by using maximal marginal relevance method to select representative sentences from the input and feed to the encoder-decoder architecture. [Nema *et al.*, 2017] propose a query based abstractive summarization approach where they first encode both the query and the document, and at each decoder time step, calculate attention based on both.

Structure Infused Document Modeling

There have been numerous studies to incorporate structure in document representations by adding syntactic parsing to pipeline like architectures. However, external parser computation is costly and does not scale well. [Tang *et al.*, 2015a; Tang *et al.*, 2015b] are one of the first to obtain document representations by first computing sentence representations and hierarchically aggregating them. [Yang *et al.*, 2016] propose a model to implicitly add structural information within end-to-end training by assigning attention scores to each sentence according to the context at that state. [Kim *et al.*, 2017] encoded structural information as graphical models within deep neural networks to facilitate end-to-end training. They use two graphical attention structures – linear chain conditional random fields and graph-based parsing models, and show how both of them can be implemented as layers of a neural network. Recent work by [Balachandran *et al.*, 2020] proposes a similar idea by incorporating both latent and explicit sentence dependencies into single document summarization architecture. They further explicitly induce structure by injecting a coreferring mention graph across sentences.

7 Conclusion

In this work, we proposed an approach to incorporate structural attention within end-to-end training in summarization networks. We achieved a considerable improvement in terms of ROUGE scores compared to our primary baseline model on both the CNN/Dailymail and CQA datasets. We also introduced multi-level contextual attention which helped in generating more abstract summaries. The analysis hinted that incorporation of some form of structural attention might be

Question	What came first? The chicken or the egg.?
CQA Summary	That one zygote cell divided to produce the first true chicken . Prior to that first true chicken zygote , there were only non-chickens . the zygote cell is the only place where dna mutations could produce a new animal , and the zygote cell is housed in the chicken 's egg . <i>So, the egg must have as first.</i>
Ans. 1	<i>The rooster! and he's STILL crowing about it ... always talking about this chick and that chick you know the type.</i>
Ans. 2	the chicken because it eventually evolved and layed eggs.
Ans. 3	the egg came first, well this is just my theory but i believe the egg came first because the chicken is a <i>cross bred animal</i> or just an evolved creature of its earlier ancestors see the species before the chicken was evolving and <i>eventually the species started mating and the next generation looked different so then there were chickens</i> im not saying im right i mean im no scientist but i like to think im pretty close
Ans. 4	<i>the egg</i> of course...its a lot easier for an chicken egg to be a mutant of some other egg laying animal
Ans. 5	The Chicken Comes Before <i>The Egg</i> . God Made The Chicken Before The Egg.
Ans. 6	<i>An egg containing a mutation that we now know as a chicken...</i>
Pointer-Generator	it is a retorical example . <i>egg containing changes in their dna.</i> in an animal like a chicken, dna from the male sperm cell and a female ovum meet and combine . that is , two <i>non-chickens mated</i> and the dna in their new zygote contained the mutation (s) that <i>produced the first true chicken .</i>
Pointer-Generator+ Structure	an egg!! The question has not specified chicken egg! <i>An egg containing a mutation we today no as a chicken.</i> It could be a dinosaur egg, dinosaur have been around for millions of years much longer than chickens!!!.
Pointer-Generator + Hierarchical Structure	The egg. the chicken is a cross bred animal . its a lot easier for a chicken egg to be a mutant of some other animal. Eventually the species started mating and in the next generation, there were chickens.

Table 4: A *Yahoo!* Answers question thread with a few of its answers. Answers 3, 4 and 6 support the ‘egg theory’, and Answers 2, 5 support the ‘chicken theory’. Answer 1 is there for humor. Majority of the answers in the original document support the ‘egg theory’. The PG summary seems to be inconsistent, and can be seen to support both the theories within the same answer. We see that while both of our proposed models unilaterally support the ‘egg theory’, the answer of the hierarchical model is framed better.

the key to achieve significant improvement compared to the extractive counterpart in the complex multi-document summarization task.

References

[Bahdanau *et al.*, 2015] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning

- to align and translate. In *ICLR*, 2015.
- [Balachandran *et al.*, 2020] Vidhisha Balachandran, Artidoro Pagnoni, Jay Yoon Lee, Dheeraj Rajagopal, Jaime Carbonell, and Yulia Tsvetkov. Structsum: Incorporating latent and explicit sentence dependencies for single document summarization. *arXiv preprint arXiv:2003.00576*, 2020.
- [Cao *et al.*, 2016] Ziqiang Cao, Chengyao Chen, Wenjie Li, Sujian Li, Furu Wei, and Ming Zhou. Tgsum: Build tweet guided multi-document summarization dataset. In *AAAI*, pages 1–8, 2016.
- [Chowdhury and Chakraborty, 2019] Tanya Chowdhury and Tanmoy Chakraborty. Cqasumm: Building references for community question answering summarization corpora. In *CODS-COMAD*, pages 18–26. ACM, 2019.
- [Cohan *et al.*, 2018] Arman Cohan, Franck Dernoncourt, Doo Soon Kim, Trung Bui, Seokhwan Kim, Walter Chang, and Nazli Goharian. A discourse-aware attention model for abstractive summarization of long documents. In *NAACL*, pages 615–621, 2018.
- [Erkan and Radev, 2004] Günes Erkan and Dragomir R Radev. Lexrank: Graph-based lexical centrality as salience in text summarization. *Journal of artificial intelligence research*, 22:457–479, 2004.
- [Fabbri *et al.*, 2019] Alexander R Fabbri, Irene Li, Tianwei She, Suyi Li, and Dragomir R Radev. Multi-news: a large-scale multi-document summarization dataset and abstractive hierarchical model. *arXiv preprint arXiv:1906.01749*, 2019.
- [Fernandes *et al.*, 2018] Patrick Fernandes, Miltiadis Allamanis, and Marc Brockschmidt. Structured neural summarization. *arXiv preprint arXiv:1811.01824*, 2018.
- [Ganesan *et al.*, 2010] Kavita Ganesan, ChengXiang Zhai, and Jiawei Han. Opinosis: A graph based approach to abstractive summarization of highly redundant opinions. In *COLING*, pages 340–348, 2010.
- [Haghighi and Vanderwende, 2009] Aria Haghighi and Lucy Vanderwende. Exploring content models for multi-document summarization. In *NAACL*, pages 362–370. Association for Computational Linguistics, 2009.
- [Hermann *et al.*, 2015] Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. Teaching machines to read and comprehend. In *NIPS*, pages 1693–1701, 2015.
- [Kim *et al.*, 2017] Yoon Kim, Carl Denton, Luong Hoang, and Alexander M Rush. Structured attention networks. *arXiv preprint arXiv:1702.00887*, 2017.
- [Koo *et al.*, 2007] Terry Koo, Amir Globerson, Xavier Carreras Pérez, and Michael Collins. Structured prediction models via the matrix-tree theorem. In *EMNLP-CoNLL*, pages 141–150, 2007.
- [Lebanoff *et al.*, 2018] Logan Lebanoff, Kaiqiang Song, and Fei Liu. Adapting the neural encoder-decoder framework from single to multi-document summarization. In *EMNLP*, pages 4131–4141, 2018.
- [Li *et al.*, 2015] Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493*, 2015.
- [Lin, 2004] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Proc. ACL workshop on Text Summarization Branches Out*, pages 1–10, 2004.
- [Liu and Lapata, 2018] Yang Liu and Mirella Lapata. Learning structured text representations. *TACL*, 6:63–75, 2018.
- [Liu *et al.*, 2019] Yang Liu, Ivan Titov, and Mirella Lapata. Single document summarization as tree induction. In *NAACL*, pages 1745–1755, Minneapolis, Minnesota, June 2019.
- [Luong *et al.*, 2015] Thang Luong, Hieu Pham, and Christopher D. Manning. Effective approaches to attention-based neural machine translation. In *EMNLP*, pages 1412–1421, 2015.
- [Mihalcea and Tarau, 2004] Rada Mihalcea and Paul Tarau. Text-rank: Bringing order into text. In *EMNLP*, pages 1–10, 2004.
- [Nallapati *et al.*, 2016] Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Çağlar Gulçehre, and Bing Xiang. Abstractive text summarization using sequence-to-sequence RNNs and beyond. In *SIGNLL*, pages 280–290, 2016.
- [Nema *et al.*, 2017] Preksha Nema, Mitesh M. Khapra, Anirban Laha, and Balaraman Ravindran. Diversity driven attention model for query-based abstractive summarization. In *ACL*, pages 1063–1072, 2017.
- [Paulus *et al.*, 2018] Romain Paulus, Caiming Xiong, and Richard Socher. A deep reinforced model for abstractive summarization. In *ICLR*, 2018.
- [Rush *et al.*, 2015] Alexander M. Rush, Sumit Chopra, and Jason Weston. A neural attention model for abstractive sentence summarization. In *EMNLP*, page 379–389, 2015.
- [See *et al.*, 2017] Abigail See, Peter J. Liu, and Christopher D. Manning. Get to the point: Summarization with pointer-generator networks. In *ACL*, pages 1073–1083, 2017.
- [Song *et al.*, 2018] Kaiqiang Song, Lin Zhao, and Fei Liu. Structure-infused copy mechanisms for abstractive summarization. In *COLING*, pages 1717–1729, 2018.
- [Sutskever *et al.*, 2014] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *NIPS*, pages 3104–3112, 2014.
- [Tang *et al.*, 2015a] Duyu Tang, Bing Qin, and Ting Liu. Document modeling with gated recurrent neural network for sentiment classification. In *EMNLP*, pages 1422–1432, 2015.
- [Tang *et al.*, 2015b] Duyu Tang, Bing Qin, and Ting Liu. Learning semantic representations of users and products for document level sentiment classification. In *ACL-IJCNLP*, pages 1014–1023, 2015.
- [Tutte, 1984] William Thomas Tutte. Graph theory, vol. 21. *Encyclopedia of Mathematics and its Applications*, 1984.
- [Vinyals and Le, 2015] Oriol Vinyals and Quoc V. Le. A neural conversational model. *CoRR*, abs/1506.05869, 2015.
- [Vinyals *et al.*, 2015] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer networks. In *NIPS*, pages 2692–2700, 2015.
- [Yang *et al.*, 2016] Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. Hierarchical attention networks for document classification. In *NAACL*, pages 1480–1489, 2016.
- [Yasunaga *et al.*, 2017] Michihiro Yasunaga, Rui Zhang, Kshitij Meelu, Ayush Pareek, Krishnan Srinivasan, and Dragomir Radev. Graph-based neural multi-document summarization. In *CoNLL*, pages 452–462, 2017.
- [Yin *et al.*, 2016] Jun Yin, Xin Jiang, Zhengdong Lu, Lifeng Shang, Hang Li, and Xiaoming Li. Neural generative question answering. In *Proceedings of the Workshop on Human-Computer Question Answering*, *ACL*, pages 36–42, San Diego, California, June 2016.