

# Closed-Book Training to Improve Summarization Encoder Memory

Yichen Jiang and Mohit Bansal

UNC Chapel Hill

{yichenj, mbansal}@cs.unc.edu

## Abstract

A good neural sequence-to-sequence summarization model should have a strong encoder that can distill and memorize the important information from long input texts so that the decoder can generate salient summaries based on the encoder’s memory. In this paper, we aim to improve the memorization capabilities of the encoder of a pointer-generator model by adding an additional ‘closed-book’ decoder without attention and pointer mechanisms. Such a decoder forces the encoder to be more selective in the information encoded in its memory state because the decoder can’t rely on the extra information provided by the attention and possibly copy modules, and hence improves the entire model. On the *CNN/Daily Mail* dataset, our 2-decoder model outperforms the baseline significantly in terms of ROUGE and METEOR metrics, for both cross-entropy and reinforced setups (and on human evaluation). Moreover, our model also achieves higher scores in a test-only DUC-2002 generalizability setup. We further present a memory ability test, two saliency metrics, as well as several sanity-check ablations (based on fixed-encoder, gradient-flow cut, and model capacity) to prove that the encoder of our 2-decoder model does in fact learn stronger memory representations than the baseline encoder.

## 1 Introduction

Text summarization is the task of condensing a long passage to a shorter version that only covers the most salient information from the original text. Extractive summarization models (Jing and McKeown, 2000; Knight and Marcu, 2002; Clarke and Lapata, 2008; Filippova et al., 2015) directly pick words, phrases, and sentences from the source text to form a summary, while an abstractive model generates (samples) words from a fixed-size vocabulary instead of copying from text directly.

<p><b>Original Text (truncated):</b> a family have claimed the body of an infant who was discovered deceased and buried on a sydney beach last year , in order to give her a proper funeral . on november 30 , 2014 , two young boys were playing on maroubra beach when they uncovered the body of a baby girl buried under 30 centimetres of sand . now locals filomena d'alessandro and bill green have claimed the infant 's body in order to provide her with a fitting farewell . 'we're local and my husband is a police officer and he's worked with many of the officers investigating it , ' ms d'alessandro told daily mail australia . scroll down for video . a sydney family have claimed the body of a baby girl who was found buried on maroubra beach ( pictured ) on november 30 , 2014 . filomena d'alessandro and bill green have claimed the infant 's remains , who they have named lily grace , in order to provide her with a fitting farewell . ' above all as a mother i wanted to do something for that little girl , ' she added . since january the couple , who were married last year and have three children between them , have been trying to claim the baby after they heard police were going to give her a ' destitute burial ' ...</p>
<p><b>Reference summary:</b> sydney family claimed the remains of a baby found on maroubra beach . filomena d'alessandro and bill green have vowed to give her a funeral . the baby 's body was found by two boys , buried in sand on november 30 . the infant was found about 20-30 metres from the water 's edge . police were unable to identify the baby girl or her parents .</p>
<p><b>Pointer-Generator baseline:</b> a sydney family have claimed the body of a baby girl was found buried on maroubra beach on november 30 , 2014 . locals filomena d'alessandro and bill green have claimed the infant 's body in order to provide her with a fitting farewell . now locals have claimed the infant 's body in order to provide her with a fitting farewell .</p>
<p><b>Pointer-Generator + closed-book decoder:</b> two young boys were playing on maroubra beach when they uncovered the body of a baby girl buried under 30 centimetres of sand . now locals filomena d'alessandro and bill green have claimed the infant 's body in order to provide her with a fitting farewell . ' above all as a mother i wanted to do something for that little girl , ' she added .</p>

Figure 1: Baseline model repeats itself twice (italic), and fails to find all salient information (highlighted in red in the original text) from the source text that is covered by our 2-decoder model. The summary generated by our 2-decoder model also recovers most of the information mentioned in the reference summary (highlighted in blue in the reference summary).

The last few years have seen significant progress on both extractive and abstractive approaches, of which a large number of studies are fueled by neural sequence-to-sequence models (Sutskever et al., 2014). One popular formulation of such models is an RNN/LSTM encoder that encodes the source passage to a fixed-size memory-state vector, and another RNN/LSTM decoder that generates the summary from this memory state. This paradigm is enhanced by attention mechanism (Bahdanau et al., 2015) and pointer network (Vinyals et al., 2015), such that the decoder can refer to (and weigh) all the encod-

ing steps’ hidden states or directly copy words from the source text, instead of relying solely on encoder’s final memory state for all information about the source passage. Recent studies (Rush et al., 2015; Nallapati et al., 2016; Chopra et al., 2016; Zeng et al., 2016; Gu et al., 2016b; Gulcehre et al., 2016; See et al., 2017) have demonstrated success with such seq-attention-seq and pointer models in summarization tasks.

While the advantage of attention and pointer models compared to vanilla sequence-to-sequence models in summarization is well supported by previous studies, these models still struggle to find the most salient information in the source text when generating summaries. This is because summarization, being different from other text-to-text generation tasks (where there is an almost one-to-one correspondence between input and output words, e.g., machine translation), requires the sequence-attention-sequence model to additionally decide where to attend and where to ignore, thus demanding a strong encoder that can determine the importance of different words, phrases, and sentences and flexibly encode salient information in its memory state. To this end, we propose a novel 2-decoder architecture by adding another ‘closed book’ decoder without attention layer to a popular pointer-generator baseline, such that the ‘closed book’ decoder and pointer decoder share an encoder. We argue that this additional ‘closed book’ decoder encourages the encoder to be better at memorizing salient information from the source passage, and hence strengthen the entire model. We provide both intuition and evidence for this argument in the following paragraphs.

Consider the following case. Two students are learning to do summarization from scratch. During training, both students can first scan through the passage once (encoder’s pass). Then student A is allowed to constantly look back (attention) at the passage when writing the summary (similar to a pointer-generator model), while student B has to occasionally write the summary without looking back (similar to our 2-decoder model with a non-attention/copy decoder). During the final test, both students can look at the passage while writing summaries. We argue that student B will write more salient summaries in the test because s/he learns to better distill and memorize important information in the first scan/pass by not looking back at the passage in training.

In terms of back-propagation intuition, during the training of a seq-attention-seq model (e.g., See et al. (2017)), most gradients are back-propagated from the decoder to the encoder’s hidden states through the attention layer. This encourages the encoder to correctly encode salient words at the corresponding encoding steps, but does not make sure that this information is not forgotten (overwritten in the memory state) by the encoder afterward. However, for a plain LSTM (closed-book) decoder without attention, its generated gradient flow is back-propagated to the encoder through the memory state, which is the only connection between itself and the encoder, and this, therefore, encourages the encoder to encode only the salient, important information in its memory state. Hence, to achieve this desired effect, we jointly train the two decoders, which share one encoder, by optimizing the weighted sum of their losses. This approximates the training routine of student B because *the sole encoder has to perform well for both decoders*. During inference, we only employ the pointer decoder due to its copying advantage over the closed-book decoder, similar to the situation of student B being able to refer back to the passage during the test for best performance (but is still trained hard to do well in both situations). Fig. 1 shows an example of our 2-decoder summarizer generating a summary that covers the original passage with more saliency than the baseline model.

Empirically, we test our 2-decoder architecture on the *CNN/Daily Mail* dataset (Hermann et al., 2015; Nallapati et al., 2016), and our model surpasses the strong pointer-generator baseline significantly on both ROUGE (Lin, 2004) and METEOR (Denkowski and Lavie, 2014) metrics, as well as based on human evaluation. This holds true both for a cross-entropy baseline as well as a stronger, policy-gradient based reinforcement learning setup (Williams, 1992). Moreover, our 2-decoder models (both cross-entropy and reinforced) also achieve reasonable improvements on a test-only generalizability/transfer setup on the DUC-2002 dataset.

We further present a series of numeric and qualitative analysis to understand whether the improvements in these automatic metric scores are in fact due to the enhanced memory and saliency strengths of our encoder. First, by evaluating the representation power of the encoder’s final memory state after reading long passages (w.r.t.

the memory state after reading ground-truth summaries) via a cosine-similarity test, we prove that our 2-decoder model indeed has a stronger encoder with better memory ability. Next, we conduct three sets of ablation studies based on fixed-encoder, gradient-flow cut, and model capacity to show that the stronger encoder is the reason behind the significant improvements in ROUGE and METEOR scores. Finally, we show that summaries generated by our 2-decoder model are qualitatively better than baseline summaries as the former achieved higher scores on two saliency metrics (based on cloze-Q&A blanks and a keyword classifier) than the baseline summaries, while maintaining similar length and better avoiding repetitions. This directly demonstrates our 2-decoder model’s enhanced ability to memorize and recover important information from the input document, which is our main contribution in this paper.

## 2 Related Work

### **Extractive and Abstractive Summarization:**

Early models for automatic text summarization were usually extractive (Jing and McKeown, 2000; Knight and Marcu, 2002; Clarke and Lapata, 2008; Filippova et al., 2015). For abstractive summarization, different early non-neural approaches were applied, based on graphs (Gianakopoulos, 2009; Ganesan et al., 2010), discourse trees (Gerani et al., 2014), syntactic parse trees (Cheung and Penn, 2014; Wang et al., 2013), and a combination of linguistic compression and topic detection (Zajic et al., 2004). Recent neural-network models have tackled abstractive summarization using methods such as hierarchical encoders and attention, coverage, and distraction (Rush et al., 2015; Chopra et al., 2016; Nallapati et al., 2016; Chen et al., 2016; Takase et al., 2016) as well as various initial large-scale, short-length summarization datasets like DUC-2004 and Gigaword. Nallapati et al. (2016) adapted the *CNN/Daily Mail* (Hermann et al., 2015) dataset for long-text summarization, and provided an abstractive baseline using attentional sequence-to-sequence model.

**Pointer Network for Summarization:** Pointer networks (Vinyals et al., 2015) are useful for summarization models because summaries often need to copy/contain a large number of words that have appeared in the source text. This provides the advantages of both *extractive* and *abstractive* ap-

proaches, and usually includes a gating function to model the distribution for the extended vocabulary including the pre-set vocabulary and words from the source text (Zeng et al., 2016; Nallapati et al., 2016; Gu et al., 2016b; Gulcehre et al., 2016; Miao and Blunsom, 2016). See et al. (2017) used a soft gate to control model’s behavior of copying versus generating. They further applied coverage mechanism and achieved the state-of-the-art results on *CNN/Daily Mail* dataset.

**Memory Enhancement:** Some recent works (Wang et al., 2016; Xiong et al., 2018; Gu et al., 2016a) have studied enhancing the memory capacity of sequence-to-sequence models. They studied this problem in Neural Machine Translation by keeping an external memory state analogous to data in the Von Neumann architecture, while the instructions are represented by the sequence-to-sequence model. Our work is novel in that we aim to improve the internal long-term memory of the encoder LSTM by adding a closed-book decoder that has no attention layer, yielding a more efficient internal memory that encodes only important information from the source text, which is crucial for the task of long-document summarization.

**Reinforcement Learning:** Teacher forcing style maximum likelihood training suffers from exposure bias (Bengio et al., 2015), so recent works instead apply reinforcement learning style policy gradient algorithms (REINFORCE (Williams, 1992)) to directly optimize on metric scores (Henß et al., 2015; Paulus et al., 2018). Reinforced models that employ this method have achieved good results in a number of tasks including image captioning (Liu et al., 2017; Ranzato et al., 2016), machine translation (Bahdanau et al., 2016; Norouzi et al., 2016), and text summarization (Ranzato et al., 2016; Paulus et al., 2018).

## 3 Models

### 3.1 Pointer-Generator Baseline

The pointer-generator network proposed in See et al. (2017) can be seen as a hybrid of *extractive* and *abstractive* summarization models. At each decoding step, the model can either sample a word from its vocabulary, or copy a word directly from the source passage. This is enabled by the attention mechanism (Bahdanau et al., 2015), which includes a distribution  $a_i$  over all encoding steps, and a context vector  $c_t$  that is the weighted sum of encoder’s hidden states. The attention mechanism

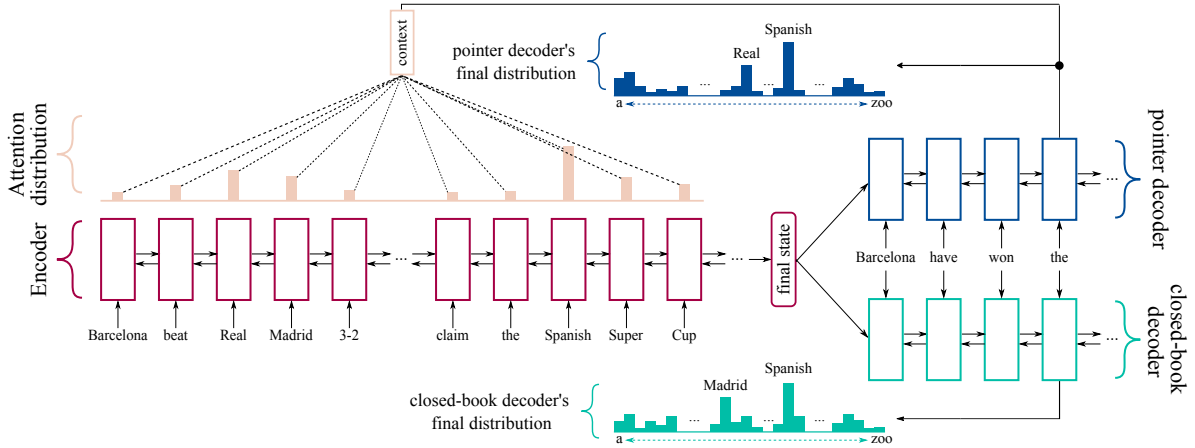


Figure 2: Our 2-decoder summarization model with a pointer decoder and a closed-book decoder, both sharing a single encoder (this is during training; next, at inference time, we only employ the memory-enhanced encoder and the pointer decoder).

is modeled as:

$$\begin{aligned} e_i^t &= v^T \tanh(W_h h_i + W_s s_t + b_{attn}) \\ a_i^t &= \text{softmax}(e_i^t); \quad c_t = \sum_i a_i^t h_i \end{aligned} \quad (1)$$

where  $v$ ,  $W_h$ ,  $W_s$ , and  $b_{attn}$  are learnable parameters.  $h_i$  is encoder's hidden state at  $i^{\text{th}}$  encoding step, and  $s_t$  is decoder's hidden state at  $t^{\text{th}}$  decoding step. The distribution  $a_i^t$  can be seen as the amount of attention at decode step  $t$  towards the  $i^{\text{th}}$  encoder state. Therefore, the context vector  $c_t$  is the sum of the encoder's hidden states weighted by attention distribution  $a^t$ .

At each decoding step, the previous context vector  $c_{t-1}$  is concatenated with current input  $x_t$ , and fed through a non-linear recurrent function along with the previous hidden state  $s_{t-1}$  to produce the new hidden state  $s_t$ . The context vector  $c_t$  is then calculated according to Eqn. 1 and concatenated with the decoder state  $s_t$  to produce the logits for the vocabulary distribution  $P_{vocab}^t$  at decode step  $t$ :  $P_{vocab}^t = \text{softmax}(V_2(V_1[s_t, c_t] + b_1) + b_2)$ , where  $V_1$ ,  $V_2$ ,  $b_1$ ,  $b_2$  are learnable parameters. To enable copying out-of-vocabulary words from source text, a pointer similar to Vinyals et al. (2015) is built upon the attention distribution and controlled by the generation probability  $p_{gen}$ :

$$\begin{aligned} p_{gen}^t &= \sigma(U_c c_t + U_s s_t + U_x x_t + b_{ptr}) \\ P_{attn}^t(w) &= p_{gen}^t P_{vocab}^t(w) + (1 - p_{gen}^t) \sum_{i:w_i=w} a_i^t \end{aligned}$$

where  $U_c$ ,  $U_s$ ,  $U_x$ , and  $b_{ptr}$  are learnable parameters.  $x_t$  and  $s_t$  are the input token and decoder's state at  $t^{\text{th}}$  decoding step.  $\sigma$  is the sigmoid function. We can see  $p_{gen}$  as a soft gate that controls

the model's behavior of copying from text with attention distribution  $a_i^t$  versus sampling from vocabulary with generation distribution  $P_{vocab}^t$ .

### 3.2 Closed-Book Decoder

As shown in Eqn. 1, the attention distribution  $a_i$  depends on decoder's hidden state  $s_t$ , which is derived from decoder's memory state  $c_t$ . If  $c_t$  does not encode salient information from the source text or encodes too much unimportant information, the decoder will have a hard time to locate relevant encoder states with attention. However, as explained in the introduction, most gradients are back-propagated through attention layer to the encoder's hidden state  $h_t$ , not directly to the final memory state, and thus provide little incentive for the encoder to memorize salient information in  $c_t$ .

Therefore, to enhance encoder's memory, we add a closed-book decoder, which is a unidirectional LSTM decoder without attention/pointer layer. The two decoders share a single encoder and word-embedding matrix, while out-of-vocabulary (OOV) words are simply represented as [UNK] for the closed-book decoder. The entire 2-decoder model is represented in Fig. 2. During training, we optimize the weighted sum of negative log likelihoods from the two decoders:

$$\begin{aligned} \mathcal{L}_{XE} &= \frac{1}{T} \sum_{t=1}^T -((1 - \gamma) \log P_{attn}^t(w|x_{1:t}) \\ &\quad + \gamma \log P_{cbdec}^t(w|x_{1:t})) \end{aligned} \quad (2)$$

where  $P_{cbdec}$  is the generation probability from the closed-book decoder. The mix ratio  $\gamma$  is tuned on the validation set.

	ROUGE			MTR
	1	2	L	Full
PREVIOUS WORKS				
* (Nallapati16)	35.46	13.30	32.65	
pg (See17)	36.44	15.66	33.42	16.65
OUR MODELS				
pg (baseline)	36.70	15.71	33.74	16.94
pg + cbdec	38.21	16.45	34.70	18.37
RL + pg	37.02	15.79	34.00	17.55
RL + pg + cbdec	<b>38.58</b>	<b>16.57</b>	<b>35.03</b>	<b>18.86</b>

Table 1: ROUGE F1 and METEOR scores (*non-coverage*) on CNN/Daily Mail test set of previous works and our models. ‘pg’ is the pointer-generator baseline, and ‘pg + cbdec’ is our 2-decoder model with closed-book decoder(cbdec). The model marked with \* is trained and evaluated on the anonymized version of the data.

### 3.3 Reinforcement Learning

In the reinforcement learning setting, our summarization model is the policy network that generates words to form a summary. Following Paulus et al. (2018), we use a self-critical policy gradient training algorithm (Rennie et al., 2016; Williams, 1992) for both our baseline and 2-decoder model. For each passage, we sample a summary  $y^s = w_{1:T+1}^s$ , and greedily generate a summary  $\hat{y} = \hat{w}_{1:T+1}$  by selecting the word with the highest probability at each step. Then these two summaries are fed to a reward function  $r$ , which is the ROUGE-L scores in our case. The RL loss function is:

$$\mathcal{L}_{RL} = \frac{1}{T} \sum_{t=1}^T (r(\hat{y}) - r(y^s)) \log P_{attn}^t(w_{t+1}^s | w_{1:t}^s) \quad (3)$$

where the reward for the greedily-generated summary ( $r(\hat{y})$ ) acts as a baseline to reduce variance. We train our reinforced model using the mixture of Eqn. 3 and Eqn. 2, since Paulus et al. (2018) showed that a pure RL objective would lead to summaries that receive high rewards but are not fluent. The final mixed loss function for RL is:  $\mathcal{L}_{XE+RL} = \lambda \mathcal{L}_{RL} + (1 - \lambda) \mathcal{L}_{XE}$ , where the value of  $\lambda$  is tuned on the validation set.

## 4 Experimental Setup

We evaluate our models mainly on *CNN/Daily Mail* dataset (Hermann et al., 2015; Nallapati et al., 2016), which is a large-scale, long-paragraph summarization dataset. It has online news articles (781 tokens or ~40 sentences on average) with paired human-generated summaries (56 tokens or 3.75 sentences on average). The

	ROUGE			MTR
	1	2	L	Full
PREVIOUS WORKS				
pg (See17)	39.53	17.28	36.38	18.72
RL* (Paulus17)	39.87	15.82	36.90	
OUR MODELS				
pg (baseline)	39.22	17.02	35.95	18.70
pg + cbdec	40.05	17.66	36.73	19.48
RL + pg	39.59	17.18	36.16	19.70
RL + pg + cbdec	<b>40.66</b>	<b>17.87</b>	<b>37.06</b>	<b>20.51</b>

Table 2: ROUGE F1 and METEOR scores (*with-coverage*) on the *CNN/Daily Mail* test set. Coverage mechanism (See et al., 2017) is used in all models except the RL model (Paulus et al., 2018). The model marked with \* is trained and evaluated on the anonymized version of the data.

	ROUGE			MTR
	1	2	L	Full
pg (See17)	37.22	15.78	33.90	13.69
pg (baseline)	37.15	15.68	33.92	13.65
pg + cbdec	37.59	16.84	34.43	13.82
RL + pg	39.92	16.71	36.13	15.12
RL + pg + cbdec	<b>41.48</b>	<b>18.69</b>	<b>37.71</b>	<b>15.88</b>

Table 3: ROUGE F1 and METEOR scores on DUC-2002 (test-only transfer setup).

entire dataset has 287,226 training pairs, 13,368 validation pairs and 11,490 test pairs. We use the same version of data as See et al. (2017), which is the original text with no preprocessing to replace named entities. We also use DUC-2002, which is also a long-paragraph summarization dataset of news articles. This dataset has 567 articles and 1~2 summaries per article.

All the training details (e.g., vocabulary size, RNN dimension, optimizer, batch size, learning rate, etc.) are provided in the supplementary materials.

## 5 Results

We first report our evaluation results on *CNN/Daily Mail* dataset. As shown in Table 1, our 2-decoder model achieves statistically significant improvements<sup>1</sup> upon the pointer-generator baseline (pg), with +1.51, +0.74, and +0.96 points advantage in ROUGE-1, ROUGE-2 and ROUGE-L (Lin, 2004), and +1.43 points advantage in METEOR (Denkowski and Lavie, 2014). In the reinforced setting, our 2-decoder model still maintains significant ( $p < 0.001$ )

<sup>1</sup>Our improvements in Table 1 are statistically significant with  $p < 0.001$  (using bootstrapped randomization test with 100k samples (Efron and Tibshirani, 1994)) and have a 95% ROUGE-significance interval of at most  $\pm 0.25$ .

<p><b>Reference summary:</b> mitchell moffit and greg brown from asapscience present theories. different personality traits can vary according to expectations of parents. beyoncé, hillary clinton and j. k. rowling are all oldest children.</p>
<p><b>Pointer-Gen baseline:</b> the kardashians are a strong example of a large celebrity family where the siblings share very different personality traits. on asapscience on youtube, the pair discuss how being the first, middle, youngest, or an only child affects us.</p>
<p><b>Pointer-Gen + closed-book decoder:</b> the kardashians are a strong example of a large celebrity family where the siblings share very different personality traits. on asapscience on youtube, the pair discuss how being the first, middle, youngest, or an only child affects us. the personality traits are also supposedly affected by whether parents have high expectations and how strict they were.</p>

Figure 3: The summary generated by our 2-decoder model covers salient information (highlighted in red) mentioned in the reference summary, which is not presented in the baseline summary.

Model	Score
2-Decoder Wins	49
Pointer-Generator Wins	31
Non-distinguishable	20

Table 4: Human evaluation for our 2-decoder model versus the pointer-generator baseline.

advantage in all metrics over the pointer-generator baseline.

We further add the coverage mechanism as in See et al. (2017) to both baseline and 2-decoder model, and our 2-decoder model (pg + cbdec) again receives significantly higher<sup>2</sup> scores than the original pointer-generator (pg) from See et al. (2017) and our own pg baseline, in all ROUGE and METEOR metrics (see Table 2). In the reinforced setting, our 2-decoder model (RL + pg + cbdec) outperforms our strong RL baseline (RL + pg) by a considerable margin (stat. significance of  $p < 0.001$ ). Fig. 1 and Fig. 3 show two examples of our 2-decoder model generating summaries that cover more salient information than those generated by the pointer-generator baseline (see supplementary materials for more example summaries).

We also evaluate our 2-decoder model with coverage on the DUC-2002 test-only generalizability/transfer setup by decoding the entire dataset with our models pre-trained on *CNN/Daily Mail*, again achieving decent improvements (shown in Table 3) over the single-decoder baseline as well as See et al. (2017), in both a cross-entropy and a reinforcement learning setup.

<sup>2</sup>All our improvements in Table 2 are statistically significant with  $p < 0.001$ , and have a 95% ROUGE-significance interval of at most  $\pm 0.25$ .

	similarity
pg (baseline)	0.817
pg + cbdec ( $\gamma = \frac{1}{2}$ )	0.869
pg + cbdec ( $\gamma = \frac{2}{3}$ )	<b>0.889</b>
pg + cbdec ( $\gamma = \frac{3}{4}$ )	0.872
pg + cbdec ( $\gamma = \frac{10}{11}$ )	0.860

Table 5: Cosine-similarity between memory states after two forward passes.

## 5.1 Human Evaluation

We also conducted a small-scale human evaluation study by randomly selecting 100 samples from the *CNN/DM* test set and then asking human annotators to rank the baseline summaries versus the 2-decoder’s summaries (randomly shuffled to anonymize model identity) according to an overall score based on readability (grammar, fluency, coherence) and relevance (saliency, redundancy, correctness). As shown in Table 4, our 2-decoder model outperforms the pointer-generator baseline (stat. significance of  $p < 0.03$ ).

## 6 Analysis

In this section, we present a series of analysis and tests in order to understand the improvements of the 2-decoder models reported in the previous section, and to prove that it fulfills our intuition that the closed-book decoder improves the encoder’s ability to encode salient information in the memory state.

### 6.1 Memory Similarity Test

To verify our argument that the closed-book decoder improves the encoder’s memory ability, we design a test to numerically evaluate the representation power of encoder’s final memory state. We perform two forward passes for each encoder (2-decoder versus pointer-generator baseline). For the first pass, we feed the entire article to the encoder and collect the final memory state; for the second pass we feed the ground-truth summary to the encoder and collect the final memory state. Then we calculate the cosine similarity between these two memory-state vectors. For an optimal summarization model, its encoder’s memory state after reading the entire article should be highly similar to its memory state after reading the ground truth summary (which contains all the important information), because this shows that when reading a long passage, the model is only encoding important information in its memory and

forgets the unimportant information. The results in Table 5 show that the encoder of our 2-decoder model achieves significantly ( $p < 0.001$ ) higher article-summary similarity score than the encoder of a pointer-generator baseline. This observation verifies our hypothesis that the closed-book decoder can improve the memory ability of the encoder.

## 6.2 Ablation Studies and Sanity Check

**Fixed-Encoder Ablation:** Next, we conduct an ablation study in order to prove the qualitative superiority of our 2-decoder model’s encoder to the baseline encoder. To do this, we train two pointer-generators with randomly initialized decoders and word embeddings. For the first model, we restore the pre-trained encoder from our pointer-generator baseline; for the second model, we restore the pre-trained encoder from our 2-decoder model. We then fix the encoder’s parameters for both models during the training, only updating the embeddings and decoders with gradient descent. As shown in the upper half of Table 6, the pointer-generator with our 2-decoder model’s encoder receives significantly higher ( $p < 0.001$ ) scores in ROUGE than the pointer-generator with baseline’s encoder. Since these two models have the exact same structure with only the encoders initialized according to different pre-trained models, the significant improvements in metric scores suggest that our 2-decoder model does have a stronger encoder than the pointer-generator baseline.

**Gradient-Flow-Cut Ablation:** We further design another ablation test to identify how the gradients from the closed-book decoder influence the entire model during training. Fig. 4 demonstrates the forward pass (solid line) and gradient flow (dashed line) between encoder, decoders, and embeddings in our 2-decoder model. As we can see, the closed-book decoder only depends on the word embeddings and encoder. Therefore it can affect the entire model during training by influencing either the encoder or the word-embedding matrix. When we stop the gradient flow between the encoder and closed-book decoder (① in Fig. 4), and keep the flow between closed-book decoder and embedding matrix (② in Fig. 4), we observe non-significant improvements in ROUGE compared to the baseline. On the other hand, when we stop the gradient flow at ② and keep ①, the improvements are statistically significant ( $p < 0.01$ ) (see the lower half

	ROUGE		
	1	2	L
FIXED-ENCODER ABLATION			
pg baseline’s encoder	37.59	16.27	34.33
2-decoder’s encoder	<b>38.44</b>	<b>16.85</b>	<b>35.17</b>
GRADIENT-FLOW-CUT ABLATION			
pg baseline	37.73	16.52	34.49
stop ①	37.72	16.58	34.54
stop ②	<b>38.35</b>	<b>16.79</b>	<b>35.13</b>

Table 6: ROUGE F1 scores of ablation studies, evaluated on *CNN/Daily Mail* validation set.




Figure 4: Solid lines represent the forward pass, and dashed lines represent the gradient flow in back-propagation. For the two ablation tests, we stop the gradient at ① and ② respectively.

of Table 6). This proves that the gradients back-propagated from closed-book decoder to the encoder can strengthen the entire model, and hence verifies the gradient-flow intuition discussed in introduction (Sec. 1).

**Model Capacity:** To validate and sanity-check that the improvements are the result of the inclusion of our closed-book decoder and not due to some trivial effects of having two decoders or larger model capacity (more parameters), we train a variant of our model with two duplicated (initialized to be different) attention-pointer decoders. We also evaluate a pointer-generator baseline with 2-layer encoder and decoder (pg-2layer) and increase the LSTM hidden dimension and word embedding dimension of the pointer-generator baseline (pg-big) to exceed the total number of parameters of our 2-decoder model (34.5M versus 34.4M parameters). Table 7 shows that neither of these variants can match our 2-decoder model in terms of ROUGE and METEOR scores, and hence proves that the improvements of our model are indeed because of the closed-book decoder rather than due to simply having more parameters.<sup>3</sup>

**Mixed-loss Ratio Ablation:** We also present eval-

<sup>3</sup>It is also important to point out that our model is not a 2-decoder ensemble, because we use only the pointer decoder during inference. Therefore, the number of parameters used for inference is the same as the pointer-generator baseline.

	ROUGE		
	I	2	L
pg baseline	37.73	16.52	34.49
pg + ptrdec	37.66	16.50	34.47
pg-2layer	37.92	16.48	34.62
pg-big	38.03	16.71	34.84
pg + cbdec	<b>38.87</b>	<b>16.93</b>	<b>35.38</b>

Table 7: ROUGE F1 and METEOR scores of sanity check ablations, evaluated on *CNN/DM* validation set.

	ROUGE		
	I	2	L
$\gamma = 0$	37.73	16.52	34.49
$\gamma = 1/2$	38.09	16.71	34.89
$\gamma = 2/3$	<b>38.87</b>	<b>16.93</b>	<b>35.38</b>
$\gamma = 5/6$	38.21	16.69	34.81
$\gamma = 10/11$	37.99	16.39	34.7

Table 8: ROUGE F1 scores on *CNN/DM* validation set, of 2-decoder models with different values of the closed-book-decoder:pointer-decoder mixed loss ratio.

uation results (on the *CNN/Daily Mail* validation set) of our 2-decoder models with different closed-book-decoder:pointer-decoder mixed-loss ratio ( $\gamma$  in Eqn. 2) in Table 8. The model achieves the best ROUGE and METEOR scores at  $\gamma = \frac{2}{3}$ . Comparing Table 8 with Table 5, we observe a similar trend between the increasing ROUGE/METEOR scores and increasing memory cosine-similarities, which suggests that the performance of a pointer-generator is strongly correlated with the representation power of the encoder’s final memory state.

### 6.3 Saliency and Repetition

Finally, we show that our 2-decoder model can make use of this better encoder memory state to summarize more salient information from the source text, as well as to avoid generating unnecessarily lengthy and repeated sentences besides achieving significant improvements on ROUGE and METEOR metrics.

**Saliency:** To evaluate saliency, we design a keyword-matching test based on the original *CNN/Daily Mail* cloze blank-filling task (Hermann et al., 2015). Each news article in the dataset is marked with a few cloze-blank keywords that represent salient entities, including names, locations, etc. We count the number of keywords that appear in our generated summaries, and found that the output of our best teacher-forcing model (pg+cbdec with coverage) contains 62.1% of those keywords, while the output provided by See et al. (2017) has only 60.4% covered. Our reinforced 2-decoder model (RL + pg + cbdec) further increases this percentage to 66.2%. The full com-

	saliency 1	saliency 2
pg (See17)	60.4%	27.95%
our pg baseline	59.6%	28.95%
pg + cbdec	62.1%	29.97%
RL + pg	62.5%	30.17%
RL + pg + cbdec	<b>66.2%</b>	<b>31.40%</b>

Table 9: Saliency scores based on *CNN/Daily Mail* cloze blank-filling task and a keyword-detection approach (Pasunuru and Bansal, 2018). All models in this table are trained with coverage loss.

	3-gram	4-gram	5-gram	sent
pg (baseline)	13.20%	12.32%	11.60%	8.39%
pg + cbdec	<b>9.66%</b>	<b>9.02%</b>	<b>8.55%</b>	<b>6.72%</b>

Table 10: Percentage of repeated 3, 4, 5-grams and sentences in generated summaries.

parison is shown in the first column of Table 9. We also use the saliency metric in Pasunuru and Bansal (2018), which finds important words detected via a keyword classifier (trained on the SQuAD dataset (Rajpurkar et al., 2016)). The results are shown in the second column of Table 9. Both saliency tests again demonstrate our 2-decoder model’s ability to memorize important information and address them properly in the generated summary. Fig. 1 and Fig. 3 show two examples of summaries generated by our 2-decoder model compared to baseline summaries.

**Summary Length:** On average, summaries generated by our 2-decoder model have 66.42 words per summary, while the pointer-generator-baseline summaries have 65.88 words per summary (and the same effect holds true for RL models, where there is less than 1-word difference in average length). This shows that our 2-decoder model is able to achieve higher saliency with similar-length summaries (i.e., it is not capturing more salient content simply by generating longer summaries).

**Repetition:** We observe that our 2-decoder model can generate summaries that are less redundant compared to the baseline, when both models are not trained with coverage mechanism. Table 10 shows the percentage of repeated n-grams/sentences in summaries generated by the pointer-generator baseline and our 2-decoder model.

**Abstractiveness:** Abstractiveness is another major challenge for current abstractive summarization models other than saliency. Since our baseline is an abstractive model, we measure the percentage of novel n-grams (n=2, 3, 4) in our generated summaries, and find that our 2-decoder model generates 1.8%, 4.8%, 7.6% novel n-grams while



our baseline summaries have 1.6%, 4.4%, 7.1% on the same test set. Even though generating more abstractive summaries is not our focus in this paper, we still show that our improvements in metric and saliency scores are not obtained at the cost of making the model more extractive.

#### 6.4 Discussion: Connection to Multi-Task Learning

Our 2-decoder model somewhat resembles a Multi-Task Learning (MTL) model, in that both try to improve the model with extra knowledge that is not available to the original single-task baseline. While our model uses MTL-style parameter sharing to introduce extra knowledge from the same dataset, traditional Multi-Task Learning usually employs additional/out-of-domain auxiliary tasks/datasets as related knowledge (e.g., translation with 2 language-pairs). Our 2-decoder model is more about how to learn to do a single task from two different points of view, as the pointer decoder is a hybrid of extractive and abstractive summarization models (primary view), and the closed-book decoder is trained for abstractive summarization only (auxiliary view). The two decoders share their encoder and embeddings, which helps enrich the encoder’s final memory state representation.

Moreover, as shown in Sec. 6.2, our 2-decoder model (pg + cbdec) significantly outperforms the 2-duplicate-decoder model (pg + ptrdec) as well as single-decoder models with more layers/parameters, hence proving that our design of the auxiliary view (closed-book decoder doing abstractive summarization) is the reason behind the improved performance, rather than some simplistic effects of having a 2-decoder ensemble or higher #parameters.

### 7 Conclusion

We presented a 2-decoder sequence-to-sequence architecture for summarization with a closed-book decoder that helps the encoder to better memorize salient information from the source text. On *CNN/Daily Mail* dataset, our proposed model significantly outperforms the pointer-generator baselines in terms of ROUGE and METEOR scores (in both a cross-entropy (XE) setup and a reinforcement learning (RL) setup). It also achieves improvements in a test-only transfer setup on the DUC-2002 dataset in both XE and RL cases. We further showed that our 2-decoder model indeed

has a stronger encoder with better memory capabilities, and can generate summaries with more salient information from the source text. To the best of our knowledge, this is the first work that studies the “representation power” of the encoders final state in an encoder-decoder model. Furthermore, our simple, insightful 2-decoder architecture can also be useful for other tasks that require long-term memory from the encoder, e.g., long-context QA/dialogue and captioning for long videos.

### 8 Acknowledgement

We thank the reviewers for their helpful comments. This work was supported by DARPA (YFA17-D17AP00022), Google Faculty Research Award, Bloomberg Data Science Research Grant, Nvidia GPU awards, and Amazon AWS. The views contained in this article are those of the authors and not of the funding agency.

### References

- D. Bahdanau, K. Cho, and Y. Bengio. 2015. Neural Machine Translation by Jointly Learning to Align and Translate. In *Third International Conference on Learning Representations*.
- Dzmitry Bahdanau, Philemon Brakel, Kelvin Xu, Anirudh Goyal, Ryan Lowe, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2016. An actor-critic algorithm for sequence prediction. In *International Conference on Learning Representations*.
- Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. Scheduled sampling for sequence prediction with recurrent neural networks. In *Advances in Neural Information Processing Systems*, pages 1171–1179.
- Qian Chen, Xiaodan Zhu, Zhenhua Ling, Si Wei, and Hui Jiang. 2016. Distraction-based neural networks for modeling documents. In *IJCAI*.
- Jackie Chi Kit Cheung and Gerald Penn. 2014. Unsupervised sentence enhancement for automatic summarization. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 775–786.
- Sumit Chopra, Michael Auli, and Alexander M Rush. 2016. Abstractive sentence summarization with attentive recurrent neural networks. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 93–98.
- James Clarke and Mirella Lapata. 2008. Global inference for sentence compression: An integer linear

- programming approach. *Journal of Artificial Intelligence Research*, 31:399–429.
- Michael Denkowski and Alon Lavie. 2014. Meteor universal: Language specific translation evaluation for any target language. In *Proceedings of the ninth workshop on statistical machine translation*, pages 376–380.
- Bradley Efron and Robert J Tibshirani. 1994. *An introduction to the bootstrap*. CRC press.
- Katja Filippova, Enrique Alfonseca, Carlos A Colmenares, Lukasz Kaiser, and Oriol Vinyals. 2015. Sentence compression by deletion with lstms. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 360–368.
- Kavita Ganesan, ChengXiang Zhai, and Jiawei Han. 2010. Opinosis: a graph-based approach to abstractive summarization of highly redundant opinions. In *Proceedings of the 23rd international conference on computational linguistics*, pages 340–348. ACL.
- Shima Gerani, Yashar Mehdad, Giuseppe Carenini, Raymond T Ng, and Bitan Nejat. 2014. Abstractive summarization of product reviews using discourse structure. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, volume 14, pages 1602–1613.
- George Giannakopoulos. 2009. Automatic summarization from multiple documents. *Ph. D. dissertation*.
- Jiatao Gu, Baotian Hu, Zhengdong Lu, Hang Li, and Victor OK Li. 2016a. Guided sequence-to-sequence learning with external rule memory. In *International Conference on Learning Representations*.
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor OK Li. 2016b. Incorporating copying mechanism in sequence-to-sequence learning. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*.
- Çaglar Gulcehre, Sungjin Ahn, Ramesh Nallapati, Bowen Zhou, and Yoshua Bengio. 2016. Pointing the unknown words. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*.
- Stefan Henß, Margot Mieskes, and Iryna Gurevych. 2015. A reinforcement learning approach for adaptive single-and multi-document summarization. In *GSCL*, pages 3–12.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Advances in Neural Information Processing Systems*, pages 1693–1701.
- Hongyan Jing and Kathleen R. McKeown. 2000. Cut and paste based text summarization. In *Proceedings of the 1st North American Chapter of the Association for Computational Linguistics Conference, NAACL 2000*, pages 178–185, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Kevin Knight and Daniel Marcu. 2002. Summarization beyond sentence extraction: A probabilistic approach to sentence compression. *Artificial Intelligence*, 139(1):91–107.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out: Proceedings of the ACL-04 workshop*, volume 8. Barcelona, Spain.
- S. Liu, Z. Zhu, N. Ye, S. Guadarrama, and K. Murphy. 2017. Improved Image Captioning via Policy Gradient optimization of SPIDeR. In *International Conference on Computer Vision*.
- Yishu Miao and Phil Blunsom. 2016. Language as a latent variable: Discrete generative models for sentence compression. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*.
- Ramesh Nallapati, Bowen Zhou, Çağlar Gülçehre, Bing Xiang, et al. 2016. Abstractive text summarization using sequence-to-sequence rnns and beyond. In *Computational Natural Language Learning*.
- Mohammad Norouzi, Samy Bengio, Navdeep Jaitly, Mike Schuster, Yonghui Wu, Dale Schuurmans, et al. 2016. Reward augmented maximum likelihood for neural structured prediction. In *Advances In Neural Information Processing Systems*, pages 1723–1731.
- Ramakanth Pasunuru and Mohit Bansal. 2018. Multi-reward reinforced summarization with saliency and entailment. In *Proceedings of the 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2018. A deep reinforced model for abstractive summarization. In *International Conference on Learning Representation*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2016. Sequence level training with recurrent neural networks. In *International Conference on Learning Representations*.

- Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jarret Ross, and Vaibhava Goel. 2016. Self-critical sequence training for image captioning. In *Computer Vision and Pattern Recognition (CVPR)*.
- Alexander M Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. In *Empirical Methods in Natural Language Processing*.
- 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*. Association for Computational Linguistics.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Sho Takase, Jun Suzuki, Naoaki Okazaki, Tsutomu Hirao, and Masaaki Nagata. 2016. Neural headline generation on abstract meaning representation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1054–1059.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. In *Advances in Neural Information Processing Systems*, pages 2692–2700.
- Lu Wang, Hema Raghavan, Vittorio Castelli, Radu Florian, and Claire Cardie. 2013. A sentence compression based framework to query-focused multi-document summarization. In *ACL*.
- Mingxuan Wang, Zhengdong Lu, Hang Li, and Qun Liu. 2016. Memory-enhanced decoder for neural machine translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.
- Hao Xiong, Zhongjun He, Xiaoguang Hu, and Hua Wu. 2018. Multi-channel encoder for neural machine translation. In *AAAI*.
- David Zajic, Bonnie Dorr, and Richard Schwartz. 2004. Bbn/umd at duc-2004: Topiary. In *Proceedings of the HLT-NAACL 2004 Document Understanding Workshop, Boston*, pages 112–119.
- Wenyuan Zeng, Wenjie Luo, Sanja Fidler, and Raquel Urtasun. 2016. Efficient summarization with read-again and copy mechanism. *arXiv preprint arXiv:1611.03382*.