

# Unsupervised Extractive Summarization using Pointwise Mutual Information

**Vishakh Padmakumar**  
New York University  
vishakh@nyu.edu

**He He**  
New York University  
hehe@cs.nyu.edu

## Abstract

Unsupervised approaches to extractive summarization usually rely on a notion of sentence importance defined by the semantic similarity between a sentence and the document. We propose new metrics of relevance and redundancy using pointwise mutual information (PMI) between sentences, which can be easily computed by a pre-trained language model. Intuitively, a relevant sentence allows readers to infer the document content (high PMI with the document), and a redundant sentence can be inferred from the summary (high PMI with the summary). We then develop a greedy sentence selection algorithm to maximize relevance and minimize redundancy of extracted sentences. We show that our method outperforms similarity-based methods on datasets in a range of domains including news, medical journal articles, and personal anecdotes.

## 1 Introduction

Modern neural network-based approaches to summarization require a large amount of document-summary pairs that are usually unavailable outside of the news domain. For example, summarization datasets of personal narratives and office meetings contain only a few hundred examples (Ouyang et al., 2017; Carletta et al., 2005). In this work, we tackle the problem of unsupervised extractive summarization which aims to select important sentences from the document. While there exists extensive prior work (Radev et al., 2000; Mihalcea and Tarau, 2004; Liu and Lapata, 2019; Zheng and Lapata, 2019), most approaches rely on the assumption that important sentences are *similar* to other sentences in the document. However, it is unclear if similarity-based features lead to meaningful content selection (Kedzie et al., 2018).

Inspired by recent work on formalizing the notion of importance in summarization (Peyrard,

2019), we propose metrics for *relevance* and *redundancy* based on pointwise mutual information (PMI). Intuitively, a relevant summary allows the reader to maximally infer the document content, and a summary has minimal redundancy if each sentence in it provides additional information. Therefore, we measure the relevance of a summary by its PMI with the document. High relevance means that the probability of the document increases conditioning on the summary. Similarly, we measure redundancy by PMI of sentence pairs *within the summary*. A sentence is redundant if seeing other sentences significantly increases its probability.

Based on the new metrics, we design a simple sentence extraction algorithm. We estimate the PMI of sentence pairs by a pre-trained language model fine-tuned on in-domain documents. We then use a simple sequential sentence selection algorithm for extractive summarization, which greedily maximizes relevance and minimizes redundancy.

Experimental results show that our algorithm outperforms similarity-based methods across multiple domains, including news, personal stories, and medical articles.<sup>1</sup>

## 2 Relevance and Redundancy

We begin by formalizing relevance and redundancy for summarization. Consider a document of  $n$  sentences  $D = \{d_1, \dots, d_n\}$  and a summary of  $m$  sentences  $S = \{s_1, \dots, s_m\}$ . We would like to measure the relevance of  $S$  to  $D$  and the redundancy of  $S$ .

**Relevance.** Relevance measures how well the summary condenses the original text such that we can infer its key content. Specifically, a summary

<sup>1</sup>Our code and pretrained models are available at <https://github.com/vishakhpk/mi-unsup-summ>.

sentence is relevant if observing it reduces our uncertainty about (unseen) sentences in the document. For example, the summary may contain the main link in the thread of conversation in the document. We thus quantify the relevance of a summary sentence  $s$  to a document sentence  $d$  by their PMI:

$$\text{Rel}(s, d) \stackrel{\text{def}}{=} \text{pmi}(s; d), \quad (1)$$

which measures the dependence between  $s$  and  $d$ . A positive score means that  $s$  and  $d$  are very likely to co-occur, thus seeing one implies the other. A zero score means that  $s$  and  $d$  are independent. A negative score means that  $s$  and  $d$  are unlikely to co-occur, e.g. contradicting sentences, thus such a summary sentence is discouraged. We further define the relevance of a summary  $S$  to the document  $D$  by the sum of sentence-level relevance:

$$\text{Rel}(S, D) \stackrel{\text{def}}{=} \sum_{s \in S} \sum_{d \in D} \text{Rel}(s, d). \quad (2)$$

**Redundancy.** Redundancy measures how much overlap exists among the summary sentences. It is typically measured by the semantic similarity between two sentences. However, even if two sentences express different meanings, there is redundancy if one is entailed by the other. For example, consider:

1. “Michelle, of South Shields, Tyneside, says she feels like a new woman after dropping from dress size 30 to size 12.”
2. “Michelle weighed 25st 3lbs when she joined the group in April 2013 and has since dropped to 12st 10lbs.”

Though expressing different information, both imply Michelle’s weight loss.

Given a summary sentence, we want to assign a score proportional to the amount of information in the sentence which is already present in the rest of the summary. Therefore, we quantify the redundancy of a sentence  $s$  given another sentence  $s'$  by their dependence in terms of PMI:

$$\text{Red}(s, s') \stackrel{\text{def}}{=} \text{pmi}(s; s'). \quad (3)$$

Similarly, the redundancy of a summary  $S$  is defined by the total redundancy of all sentence pairs:

$$\text{Red}(S) \stackrel{\text{def}}{=} \sum_{i=1}^m \sum_{j=i+1}^m \text{pmi}(s_i; s_j). \quad (4)$$

**Estimate PMI.** By definition,  $\text{pmi}(s; d) = \log \frac{p(s|d)}{p(d)} = \log \frac{p(d|s)}{p(s)}$ . Since both  $s$  and  $d$  are sentences, we use a language model,  $p_{\text{LM}}$ , to estimate the probabilities. Conditional probabilities are computed by considering the condition sentence as the prefix.

Note that while PMI can be computed in two equivalent ways according to the definition, the estimates from a language model do not guarantee that  $\frac{p_{\text{LM}}(d|s)}{p_{\text{LM}}(d)} = \frac{p_{\text{LM}}(s|d)}{p_{\text{LM}}(s)}$ . Thus we choose to condition on the summary sentence:

$$\text{pmi}(s; d) \stackrel{\text{def}}{=} \log \frac{p_{\text{LM}}(d | s)}{p_{\text{LM}}(d)}. \quad (5)$$

This is consistent with our definition of relevance: *seeing the summary*, how well we can estimate the document content. For redundancy, we condition on the earlier sentence:

$$\text{pmi}(s_i; s_j) \stackrel{\text{def}}{=} \begin{cases} \log \frac{p_{\text{LM}}(s_j | s_i)}{p_{\text{LM}}(s_j)} & \text{if } i < j \\ \text{pmi}(s_j; s_i) & \text{otherwise} \end{cases}, \quad (6)$$

since a sentence is redundant if it can be inferred from previous sentences.

### 3 Sequential Sentence Extraction

Given relevance and redundancy defined above, we aim to select important sentences from the document that maximize relevance and minimize redundancy. We consider a weighted combination of the two criteria:

$$\max_{S \subseteq D} \lambda_1 \text{Rel}(S, D) + \lambda_2 \text{Red}(S) \text{ s.t. } |S| \leq k, \quad (7)$$

where  $|S|$  denotes the number of sentences in the summary. This is a combinatorial problem that is expensive to solve when  $k$  is large. Therefore, we solve it approximately by selecting sentences sequentially in a greedy fashion. Given the previously selected sentences, we select the next sentence from the document that maximally improves the objective (7) until  $k$  sentences are selected. Our full algorithm is shown in Algorithm 1, where  $\Delta(s)$  denotes the change incurred by the new summary  $S \cup \{s\}$ .

### 4 Experiments

**Datasets.** We evaluate on a diverse range of domains: (i) news articles: CNN-Dailymail (CNN-DM) (See et al., 2017) and XSum (Narayan et al., 2018); (ii) personal anecdotes: Reddit (Ouyang

---

**Algorithm 1:** ExtractSentences( $D, k$ )

---

```
 $S := \emptyset$ 
for  $j \leftarrow 1$  to  $k$  do
  for  $s \in D \setminus S$  do
     $\Delta(s) := \lambda_1 \text{Rel}(\{s\}, D) +$ 
       $\lambda_2 \sum_{s' \in S} \text{Red}(s', s)$ 
  end
   $s_j := \text{argmax}_{s \in D \setminus S} \Delta(s)$ 
   $S := S \cup \{s_j\}$ 
end
return  $S$ 
```

---

et al., 2017) and Reddit-TIFU (Kim et al., 2019); (iii) scientific articles: PubMed (Kedzie et al., 2018). Further, these datasets allow us to evaluate on highly abstractive summarization (XSum), small data (Reddit), and long documents (PubMed). Details of the datasets are shown in Appendix A.

**Baselines.** We compare our approach against the following unsupervised extraction methods: (i) **heuristic:** lead- $k$  which selects the first  $k$  sentences. (ii) **similarity-based:** TextRank (Barrios et al., 2016; Mihalcea and Tarau, 2004) and PacSum (Zheng and Lapata, 2019) which use graph-based selection with similarity metrics based on tf-idf sentence vectors and BERT embeddings, respectively.<sup>2</sup>

To ablate the contribution of PMI, we also include a variant of our algorithm which uses cosine similarity of tf-idf sentence representations to measure redundancy and relevance of sentence pairs. Additionally, we include two **reference** methods: oracle extraction<sup>3</sup> and the state-of-the-art supervised approaches: For CNN-Dailymail, XSum, Reddit-TIFU and Pubmed the results are from Zhang et al. (2019) and for Reddit from Ouyang (2019).

**Language model fine-tuning.** For our method, we use the pre-trained GPT-2 large model (Radford et al., 2019) to calculate the PMI. To adapt the language model to specific domains, we fine-tune it on the training documents (excluding the gold summaries) in each dataset. To make it better fit our task of estimating the probability of one sentence

given another, as described in Section 2, we fine-tune GPT-2 on two-sentence segments (as opposed to a long stream of tokens). Each segments consist of a pair of sentences from the document.<sup>4</sup>

**Implementation details.** We preprocess the documents with spaCy (Honnibal and Montani, 2017) to split the text into sentences. All hyperparameters are tuned on 200 randomly sampled document-summary pairs selected from the validation set to optimize the Rouge-1 F-measure, including  $\lambda_1$  and  $\lambda_2$  in our method which balances relevance and redundancy scores in Equation (7), the number of keywords in TextRank and the number of sentences to select for all extractive methods.

To select the values of  $\lambda_1$  and  $\lambda_2$  we run a grid search at intervals of 0.1 from -2 to 2 for both. For all datasets, the best weighting was 2 for relevance and -2 for redundancy.<sup>5</sup> The weights are intuitive because we want to maximize relevance ( $\lambda_1$ ) and minimize redundancy ( $\lambda_2$ ). For the Lead- $k$  baselines,  $k$  was 3 for CNN-Dailymail, Reddit-TIFU and XSum, 4 for Reddit and 9 for PubMed.

## 4.1 Results

We evaluate all methods on the five datasets using the Rouge-1/2/L (Lin and Rey, 2004) F-measure. Table Table 1 shows our main results.

**PMI vs similarity.** We first compare PMI and tf-idf in our framework. The results (Ours (PMI) vs Ours (tf-idf)) show that measuring relevance and redundancy using PMI is better than word overlap, especially on narratives (Reddit). The Reddit writing style is less reporting facts like in news and more describing a sequence of events, thus it is helpful to capture the dependence between events. We show an example contrasting the two metrics in Table 2. Further, our method achieves better or comparable results across all datasets compared to other similarity-based methods. The oracle extraction is predictably the best result across all datasets. Our extractive approach outperforms the supervised baseline on the smallest dataset, Reddit, demonstrating the utility of unsupervised approaches in this setting. More examples are shown in Appendix B.

---

<sup>2</sup>For TextRank, Zheng and Lapata (2019) showed that tf-idf works better than BERT embeddings so we follow the same.

<sup>3</sup>We select sentences greedily to optimize the Rouge-1 score against the reference (Nallapati et al., 2017).

---

<sup>4</sup>Our preliminary results show that this works slightly better than standard LM fine-tuning.

<sup>5</sup>Note that there are multiple optimal values, e.g. 1 and -1, the selected values are based on tie-breaking in our algorithm.

Dataset	CNN-DM			XSum			Reddit			Reddit-TIFU			PubMed		
	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L
Lead-k	39.69	17.22	24.82	<b>19.48</b>	2.62	12.49	26.52	7.05	18.80	15.4	2.24	10.81	37.34	10.54	18.00
TextRank	34.11	12.78	22.51	19.04	3.05	<b>12.64</b>	23.76	7.90	16.41	18.70	3.49	13.08	46.73	17.28	22.28
PacSum	<b>40.26</b>	<b>17.55</b>	<b>24.92</b>	19.44	2.71	12.44	-	-	-	-	-	-	-	-	-
Ours (tf-idf)	34.17	12.99	22.44	18.14	3.14	12.31	24.14	8.19	16.96	18.24	3.18	12.45	46.54	17.77	22.66
Ours (PMI)	36.68	14.52	23.32	19.07	<b>3.22</b>	12.47	<b>28.18</b>	<b>8.60</b>	<b>20.26</b>	<b>18.93</b>	<b>3.72</b>	<b>13.11</b>	<b>46.84</b>	<b>17.81</b>	<b>23.01</b>
Oracle Ext	53.53	29.56	37.21	31.85	7.51	20.87	36.63	15.24	27.87	34.96	10.75	24.91	54.88	23.72	28.30
Supervised	44.17	21.47	41.11	47.21	24.56	39.25	20.90	5.40	18.90	26.63	9.01	21.6	45.09	19.56	27.42

Table 1: Rouge-1/2/L F-Measure scores of unsupervised extractive methods, oracle extraction, and SOTA results using supervised learning on CNN-DM, XSum, Reddit, Reddit-TIFU and PubMed. The best results among unsupervised methods are in bold. Our PMI-based extractor outperforms similarity-based methods on non-news domains, and is comparable on news domains. (PacSum results on non-news domains are not reported as the released model is fine-tuned only on the news domain.)

Gold Summary	Ours(PMI)	PacSum
Cillian McCann was filmed by his mother Toni at seven weeks old. In the clip, the little boy can clearly be seen trying to speak to his family. After several attempts he manages to say "hello". The average child can say six words by the time they reach 18 months.	Whose adorable son Cillian said his first word at just seven weeks old. <b>In the video Cillian is seen struggling to get his word out, but with a bit of encouragement from his mother he finally says hello.</b> Toni says that Cillian was very alert from a young age and had been trying to make out words since he was just five weeks old.	Most parenting advice says you don't have to worry if your baby doesn't start speaking until around 18 months. The tiny tot, who is now nine weeks old, was filmed by his 36-year-old mother who says that she knew he had been trying to communicate for a while. <b>Cillian has three older sisters, Toni revealed that her little girls, Sophie(bottom right), Eva(bottom left) and Ellie(top), did not start talking at such an early age.</b>

Table 2: Example of summaries selected by PMI vs PacSum. PacSum selected the sentence (bolded) about the child's siblings. It mentions talking at an early age which is the main theme of the article. However, it does not inform us how the child started talking in the video. Our method using PMI selected sentences focusing on the content of the video and the child.

**Extraction algorithm.** Another important component in extractive methods is the sentence selection/decoding algorithm. The most common approach is to select sentence greedily according to certain objective. TextRank uses a graph-based method inspired by PageRank. However, in Table 1 we see that TextRank and Ours (tf-idf) (using greedy selection) achieve similar results, showing that the selection algorithm does not have a large impact, which is also found by Zheng and Lapata (2019).

**Ablation study.** To understand the contribution of relevance and redundancy in the proposed metric, we conducted an ablation study on CNN/DM and Reddit-TIFU. In Table 3, we see that relevance alone does well, but augmenting it with redundancy obtains the best performance across all metrics. Minimizing redundancy alone works poorly because it cannot identify important content.

**Position bias in PacSum.** One may wonder why PacSum and lead- $k$  significantly outperform other extractive methods on CNN-DM. We hypothe-

Method	CNN-DM			Reddit-TIFU		
	R-1	R-2	R-3	R-1	R-2	R-3
Only Rel.	35.17	13.79	22.91	18.45	3.67	12.90
Only Red.	23.85	6.47	15.48	16.41	2.35	11.66
Rel. and Red.	<b>36.68</b>	<b>14.52</b>	<b>23.32</b>	<b>18.93</b>	<b>3.72</b>	<b>13.11</b>

Table 3: Ablation Results on CNN-DM and Reddit-TIFU. We observe that the combination of relevance and redundancy yields the best performance across all metrics.

size that they take advantage of the lead bias on CNN/DM. Figure 1 shows histograms of the positions of summary sentences selected by our method and PacSum on CNN/DM. Notably, 82.3% of sentences selected by PacSum were in the first three, and this value drops to 21.4% in our method. This provides an empirical explanation as to why PacSum is so far ahead of the other extractive approaches on CNN-DM. The authors (Zheng and Lapata, 2019) also noted a drop in performance when positional information was removed in PacSum. In addition, their performance degrades on XSum (which doesn't suffer from lead bias). A

concurrent work (Xu et al., 2020) performed a similar analysis, observing the reliance on position information of PacSum in the news domain.




Figure 1: Histogram of the position of sentences selected by our method and PacSum on CNN/DM. PacSum uses position information which allows it to take advantage of the lead bias. In contrast, our method is position-agnostic but still captures the fact that earlier sentences are more important in news articles.

## 5 Related Work

**Similarity-based summarization.** Most unsupervised extractive summarization methods rely on sentence pair similarity as a proxy for importance (Zheng and Lapata, 2019; Mihalcea and Tarau, 2004; Barrios et al., 2016; Erkan and Radev, 2004) and use variants of the Pagerank (Page et al., 1999) algorithm to perform selection. We propose PMI as an alternative and compare it to a concurrent similarity based approach.

**Leveraging pretrained language models.** Scoring a sentence on its ability to predict subsequent sentences using a language model has been adapted for sentence summarization. (West et al., 2019). Zhou and Rush (2019) use two language models, a generic pretrained language model for contextual matching and a task specific one to enforce fluency. A concurrent work (Xu et al., 2020) used sentence level transformer self-attentions and probabilities to rank sentences for unsupervised extractive summarization. We use the language model to compute PMI, which then scores sentences on relevance and redundancy as criteria for selection.

**Diversity in content selection.** Maximal Marginal Relevance (Goldstein and Carbonell, 1998) has been used to produce summaries that prioritize diversity in selected content. A similarity metric is used to produce summaries based on similarity to a query while maintaining diversity among selected sentences in various domains (Chandu et al., 2017). This can be seen as

analogous to our comparison to our approach tf-idf based selection.

## 6 Conclusion and Discussion

We propose metrics for relevance and redundancy in summarization based on pointwise mutual information, and an unsupervised extractive summarization algorithm using pre-trained language models. We demonstrate the effectiveness of our method on both news and non-news domains. Supervised models often learn the lead bias in the datasets and degrade significantly when such cues are absent (Kedzie et al., 2018). Furthermore, even human evaluation of content selection has large variance (Nenkova et al., 2007; Chaganty et al., 2018). Our work is a first step towards formalizing a notion of importance that informs algorithm design in summarization. We believe it is important to have a better formalization of content importance in terms of both task definition/evaluation and modeling. We hope our results will spur more work in this direction.

## References

- Federico Barrios, Federico López, Luis Argerich, and Rosa Wachenchauzer. 2016. Variations of the similarity function of textrank for automated summarization. *arXiv preprint arXiv:1602.03606*.
- J. Carletta, S. Ashby, S. Bourban, M. Flynn, M. Guillemot, T. Hain, J. Kadlec, V. Karaiskos, W. Kraaij, M. Kronenthal, G. Lathoud, M. Lincoln, A. Lisowska, I. McCowan, W. Post, D. Reidsma, and P. Wellner. 2005. The AMI meeting corpus: A pre-announcement. *Springer Lecture Notes in Computer Science*, 3869.
- A. Chaganty, S. Mussmann, and P. Liang. 2018. The price of debiasing automatic metrics in natural language evaluation. In *Association for Computational Linguistics (ACL)*.
- Khyathi Chandu, Aakanksha Naik, Aditya Chandrasekar, Zi Yang, Niloy Gupta, and Eric Nyberg. 2017. [Tackling biomedical text summarization: OAQA at BioASQ 5B](#). In *BioNLP 2017*, pages 58–66, Vancouver, Canada,. Association for Computational Linguistics.
- Günes Erkan and Dragomir R Radev. 2004. Lexrank: Graph-based lexical centrality as salience in text summarization. *Journal of artificial intelligence research*, 22:457–479.
- Jade Goldstein and Jaime Carbonell. 1998. [Summarization: \(1\) using MMR for diversity-based reranking and \(2\) evaluating summaries](#). In *TIPSTER TEXT PROGRAM PHASE III: Proceedings of*

- a Workshop held at Baltimore, Maryland, October 13-15, 1998, pages 181–195, Baltimore, Maryland, USA. Association for Computational Linguistics.
- Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. To appear.
- C. Kedzie, K. McKeown, and D. III. 2018. Content selection in deep learning models of summarization. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Byeongchang Kim, Hyunwoo Kim, and Gunhee Kim. 2019. *Abstractive summarization of Reddit posts with multi-level memory networks*. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2519–2531, Minneapolis, Minnesota. Association for Computational Linguistics.
- C. Lin and M. Rey. 2004. Looking for a few good metrics: ROUGE and its evaluation. In *NTCIR Workshop*.
- Y. Liu and M. Lapata. 2019. Text summarization with pretrained encoders. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Rada Mihalcea and Paul Tarau. 2004. TextRank: Bringing order into text. In *Proceedings of the 2004 conference on empirical methods in natural language processing*, pages 404–411.
- Ramesh Nallapati, Feifei Zhai, and Bowen Zhou. 2017. Summarunner: A recurrent neural network based sequence model for extractive summarization of documents. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- S. Narayan, S. B. Cohen, and M. Lapata. 2018. Don't give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- A. Nenkova, R. J. Passonneau, and K. McKeown. 2007. The pyramid method: Incorporating human content selection variation in summarization evaluation. In *ACM Transactions on Speech and Language Processing*.
- J. Ouyang, S. Chang, and K. McKeown. 2017. Crowd-sourced iterative annotation for narrative summarization corpora. In *European Association for Computational Linguistics (EACL)*.
- Jessica Jin Ouyang. 2019. Adapting automatic summarization to new sources of information. Link: <https://doi.org/10.7916/d8-5nar-6b61>.
- Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab.
- Maxime Peyrard. 2019. A simple theoretical model of importance for summarization. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1059–1073.
- D. R. Radev, H. Jing, and M. Budzikowska. 2000. Centroid-based summarization of multiple documents: Sentence extraction, utility-based evaluation, and user studies. In *Proceedings of the NAACL-ANLP 2000 Workshop: Automatic Summarization*.
- A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever. 2019. Language models are unsupervised multitask learners.
- A. See, P. J. Liu, and C. D. Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *Association for Computational Linguistics (ACL)*.
- Peter West, Ari Holtzman, Jan Buys, and Yejin Choi. 2019. *Bottlesum: Unsupervised and self-supervised sentence summarization using the information bottleneck principle*.
- Shusheng Xu, Xingxing Zhang, Yi Wu, Furu Wei, and Ming Zhou. 2020. *Unsupervised extractive summarization by pre-training hierarchical transformers*.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J Liu. 2019. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. *arXiv preprint arXiv:1912.08777*.
- Hao Zheng and Mirella Lapata. 2019. Sentence centrality revisited for unsupervised summarization. *arXiv preprint arXiv:1906.03508*.
- Jiawei Zhou and Alexander M Rush. 2019. Simple unsupervised summarization by contextual matching. *arXiv preprint arXiv:1907.13337*.

## A Dataset Details

CNN/DM is known to have a very strong extractive Lead-3 baseline as is common in the news domain. XSum contains summaries of BBC news articles but is highly abstractive in nature. The Reddit dataset is a small corpus of around 500 personal stories shared on Reddit with abstractive and extractive summaries. For Reddit-TIFU, we use the TIFU-long subset as used in [Zhang et al. \(2019\)](#). The Reddit-TIFU didn't come with a train split and since we look at unsupervised methods, we used 200 pairs as validation data to decide parameters and report test results on the rest. The PubMed dataset contains longer medical journal articles with the corresponding abstracts functioning as the groundtruth summaries.

Dataset	Train	Validation	Test
CNN-DM	287,113	13,368	11,490
XSum	204,045	11,332	11,334
Reddit	404	24	48
Reddit-TIFU	-	200	41139
PubMed	21,250	1,250	2,500

Table 4: Size of train, validation and test splits of datasets used

## B Analysis of Examples

Table 5 shows some example summaries from the CNN-Dailymail validation set in comparison to extractive candidate summaries obtained for the corresponding documents using the baseline Lead-3 approach, our Interpolated PMI based approach and the PacSum approach ([Zheng and Lapata, 2019](#)) that uses sentence similarity to obtain state-of-the-art Rouge results on the dataset. These are shown to highlight the difference between using PMI and similarity for sentence selection. In the first example, the gold summary details about how medical information regarding two patients was leaked to sales representatives. The similarity based approach selects all three sentences associated with only one of the patients whereas the PMI based approach yields a summary that contains information about both. Once the first sentence concerning the first patient is selected, all sentences associated with it are penalised by a corresponding amount resulting in a more well rounded selection of information. Similarly, in the second example the sentence that details how the

child overcame his initial struggles to speak after some encouragement from his parent was only selected by the PMI approach. This summarises exactly what happens in the video clip being spoken about and the same point is highlighted even in the gold summary. The contents of the paragraph can be easily understood when the information about the clip is used as a context. The rest of the paragraph talks goes into how the child spoke faster than his siblings which explains the selection made by PacSum. The third example highlights the issue of PacSum being identical to the Lead-3 baseline by modelling position information present in the dataset. In the fourth example, the PacSum based approach selects two sentences with quotes that have negative connotations while the one selected by PMI about how the protagonist could not forgive himself could serve to better explain the need for an intervention on the Dr. Phil show.

The purpose of this is to highlight that the intangible nature of the definition of relevance. The content selected varies between PMI and sentence similarity and each might find an application in the right setting. It again highlights the need to consider what one expects from the summarisation task.

Gold	Interpolated PMI	PacSum	Lead3
<p>Tim Esworthy, 66, has a prosthetic limb after losing his leg in a workplace incident and said he had been targeted by cold callers selling products easing joint pain. Christine Lewis, 62, is wheelchair-bound following a brain haemorrhage. She is also on list of people obtained by the mail and has been targeted by stairlift salesmen.</p>	<p>Tim Esworthy, 66, from Colchester, was 'absolutely appalled' to find his private medical details had been sold. Christine Lewis, who is recovering from a brain haemorrhage she had 12 years ago, was on a list of people who have mobility problems obtained by the mail and has been targeted by stairlift salesmen cold calling her. "They shouldn't have my information, especially if they know I'm disabled because they are targeting me because they think I'm vulnerable."</p>	<p>Tim Esworthy, 66, from Colchester, was 'absolutely appalled' to find his private medical details had been sold. Retired financial services manager Tim Esworthy was 'absolutely appalled' to find his private medical details had been sold. They know they can target vulnerable people because they have their medical information.</p>	<p>Tim Esworthy, 66, from Colchester, was 'absolutely appalled' to find his private medical details had been sold. Case 1 : Pensioner who lost leg at work. Retired financial services manager Tim Esworthy was 'absolutely appalled' to find his private medical details had been sold.</p>
<p>Cillian McCann was filmed by his mother Toni at seven weeks old. In the clip, the little boy can clearly be seen trying to speak to his family. After several attempts he manages to say "hello". The average child can say six words by the time they reach 18 months.</p>	<p>Whose adorable son Cillian said his first word at just seven weeks old. In the video Cillian is seen struggling to get his word out, but with a bit of encouragement from his mother he finally says hello. Toni says that Cillian was very alert from a young age and had been trying to make out words since he was just five weeks old.</p>	<p>Most parenting advice says you don't have to worry if your baby doesn't start speaking until around 18 months. The tiny tot, who is now nine weeks old, was filmed by his 36-year-old mother who says that she knew he had been trying to communicate for a while. Cillian has three older sisters, Toni revealed that her little girls, Sophie(bottom right), Eva(bottom left) and Ellie(top), did not start talking at such an early age.</p>	<p>Most parenting advice says you don't have to worry if your baby doesn't start speaking until around 18 months. Whose adorable son Cillian said his first word at just seven weeks old. The tiny tot, who is now nine weeks old, was filmed by his 36-year-old mother who says that she knew he had been trying to communicate for a while.</p>
<p>Ashleigh Humphrys, 20, died in a hit-and-run early on Sunday morning. Police believe the driver of the car was heading to work. A man is assisting police with their investigations after the death. Ms Humphrys was walking home after celebrating her birthday with friends. A security guard ran police after she was walking disorientated. CCTV footage shows two taxis stop near her before she was struck and put hazard lights on. Then a car drove past the taxis, mounted the footpath before swerving back onto the road and driving off. A taxi is said to have been seized and police are talking to a person 'within the vicinity' at the time of the incident.</p>	<p>Brisbane woman Ashleigh Humphrys died in a hit-and-run incident after deciding to walk from Toowong to her Seventeen Mile Rocks home in Brisbane after having an argument with a friend while they were out celebrating her 20th birthday. Only moments later the guard, who was still on the phone to police while driving around trying to find Ms Humphrys, discovered her dead on the road at the city end of the western freeway. Just before Ms Humphrys was hit, CCTV footage shows two taxis stop near the woman and put their hazard lights on before a car drove past the taxis, mounted the footpath and then swerved back onto the road before driving off.</p>	<p>The driver of a car that hit and killed a young woman in the early hours on Sunday morning was on the way to work, police believe. Brisbane woman Ashleigh Humphrys died in a hit-and-run incident after deciding to walk from Toowong to her Seventeen Mile Rocks home in Brisbane after having an argument with a friend while they were out celebrating her 20th birthday. Now, after it was revealed that a man was assisting police with their investigations, officers have said they believe he was on his way to work and went to his shift as normal on Sunday, the Courier Mail reported.</p>	<p>The driver of a car that hit and killed a young woman in the early hours on Sunday morning was on the way to work, police believe. Brisbane woman Ashleigh Humphrys died in a hit-and-run incident after deciding to walk from Toowong to her Seventeen Mile Rocks home in Brisbane after having an argument with a friend while they were out celebrating her 20th birthday. Now, after it was revealed that a man was assisting police with their investigations, officers have said they believe he was on his way to work and went to his shift as normal on Sunday, the Courier Mail reported.</p>
<p>Dr. Phil McGraw staged a highly-charged intervention with Nick Gordon last Thursday. With his mother, Michelle, by his side a sobbing Gordon talked about missing Bobbi Kristina. Gordon is now in rehab after the intervention having been drinking heavily and taking xanax. Girlfriend Bobbi Kristina has been in a medically induced coma since January 31 and Gordon has not been allowed to see her. The dramatic intervention will air Wednesday on the Dr Phil show.</p>	<p>Amid scenes of high emotion, an often incoherent Gordon admitted drinking heavily and taking xanax, for which he has a prescription, in an attempt to deal with life since Bobbi Kristina was found face down and unresponsive in her bathtub on January 31. Breakdown: With his mother, Michelle, by his side Nick Gordon struggles to stay coherent as he is questioned by Dr Phil. According to his mother, Michelle, Gordon can not forgive himself for his 'failure' to revive Bobbi Kristina</p>	<p>Weeping and wailing Nick Gordon, the troubled fiancé of Bobbi Kristina Brown, has admitted that he has twice tried to kill himself and confessed: "I'm so sorry for everything." Asked if he still intended to kill himself he said: "If anything happens to Krissi I will." Amid scenes of high emotion, an often incoherent Gordon admitted drinking heavily and taking xanax, for which he has a prescription, in an attempt to deal with life since Bobbi Kristina was found face down and unresponsive in her bathtub on January 31.</p>	<p>Weeping and wailing Nick Gordon, the troubled fiancé of Bobbi Kristina Brown, has admitted that he has twice tried to kill himself and confessed: "I'm so sorry for everything." Gordon, 25, was speaking to Dr Phil McGraw in a dramatic intervention due to air on Wednesday, Daily Mail online can reveal. Asked if he still intended to kill himself he said: "If anything happens to Krissi I will."</p>

Table 5: Example summaries obtained from the CNN-Dailymail validation set compared to the corresponding extractive candidate summary obtained using Interpolated PMI, PacSum (State-of-the-art unsupervised summary using sentence similarity) and the Lead-3 Baseline