

Attention Head Masking for Inference Time Content Selection in Abstractive Summarization

Shuyang Cao and Lu Wang

Computer Science and Engineering

University of Michigan

Ann Arbor, MI

{caoshuy, wangluxy}@umich.edu

Abstract

How can we effectively inform content selection in Transformer-based abstractive summarization models? In this work, we present a simple-yet-effective *attention head masking* technique, which is applied on encoder-decoder attentions to pinpoint salient content at inference time. Using attention head masking, we are able to reveal the relation between encoder-decoder attentions and content selection behaviors of summarization models. We then demonstrate its effectiveness on three document summarization datasets based on both *in-domain* and *cross-domain* settings. Importantly, our models outperform prior state-of-the-art models on CNN/Daily Mail and New York Times datasets. Moreover, our inference-time masking technique is also data-efficient, requiring less than 20% of the training samples to outperform BART fine-tuned on the full CNN/DailyMail dataset.

1 Introduction

Large pre-trained Transformers have achieved state-of-the-art results on various summarization datasets with a fine-tuning phase to streamline the summarization pipeline (Lewis et al., 2020; Yan et al., 2020). Yet, it is still unclear *how one can use large models more effectively for abstractive summarization*. For example, prior work shows that informing content selection via attention weight updating in recurrent neural networks can further boost summarizer performance (Gehrmann et al., 2018). However, with multi-heads attentions at all layers in Transformers (Vaswani et al., 2017), highlighting salient content becomes non-trivial.

In this work, we propose an **inference-time attention head masking mechanism** that works on encoder-decoder attentions to underscore salient content from the source and improve the quality of abstractive summaries. Based on this mechanism, we first demonstrate the relation between encoder-decoder attentions and content selection behaviors,

on three summarization datasets of CNN/DailyMail (CNN/DM), New York Times (NYT), and XSum. Second, we study whether multiple heads at the same layer collectively guide the summarization. Partial masking is found to be most effective, indicating a strong collaborative effect and the importance of head selection.

Based on these observations, we evaluate attention head masking on summarization benchmarks with salience labels provided by externally trained content selectors. On all three datasets, our model consistently outperforms fine-tuned BART (Lewis et al., 2020) and several top performing Transformer-based abstractive summarization models (Zhang et al., 2019b; Yan et al., 2020). Summaries generated by our model are also considered to have better informativeness by human judges. Moreover, we illustrate that attention head masking is **data-efficient**: on CNN/DM, BART fine-tuned on less than 20% of the training data outperforms a version trained on the full set. Finally, we show that our method is effective under a **cross-domain** setting. With a content selector trained on NYT, BART fine-tuned on CNN/DM gains more than three points of ROUGE scores when tested on NYT articles.¹

2 Related Work

Large Pre-trained Models for Summarization.

Many recent advancements in text summarization have been achieved by large pre-trained language models (Zhang et al., 2019a; Liu and Lapata, 2019; Song et al., 2019; Zhang et al., 2019b). In particular, BART has demonstrated impressive performance on summarization, and is used as the base model in this work. Nonetheless, all prior attempts take pre-trained models as is and conduct fine-tuning on target datasets, without knowing if

¹Our code is available at: https://shuyangcao.github.io/projects/inference_head_masking.

it is the most effective usage. In contrast, we bring insights into the relation between attentions and content selection via masking operations to further improve summarization performance.

Content Selection for Abstractive Summarization. Content selection is a crucial step, where salient information is first detected and then summarized into concise abstracts (Chen and Bansal, 2018; Xu and Durrett, 2019). To minimize the propagation of selection errors, content selection is modeled as an extra component and learned within an end-to-end trained model (Zhou et al., 2017; Li et al., 2018; Gehrmann et al., 2018). To the best of our knowledge, we are the first to apply masks on selected layers and attention heads in Transformers for content selection in summarization. Moreover, our masking mechanism is only activated during inference, without any model modification.

Analyzing Multi-head Attentions has attracted growing interests in the NLP community (Clark et al., 2019; Kovaleva et al., 2019). Among the work that is relevant to encoder-decoder attentions, Michel et al. (2019) and Voita et al. (2019) observe that only a small portion of heads is relevant for translation and encoder-decoder attentions tend to be more important than self-attentions. Meanwhile, word alignments for machine translation are induced from encoder-decoder attention weights (Li et al., 2019; Kobayashi et al., 2020). However, none of prior work employs attentions to improve generation quality. As far as we are aware, this is the first work that studies the content selection effects of encoder-decoder attentions and uses them to guide better summary generation.

3 Attention Head Masking

We adopt large pre-trained sequence-to-sequence Transformer models (BART, specifically) for abstractive summarization. Transformer is built with **multi-head attentions**. Attentions are computed per step based on a query \mathbf{q} along with the key and value matrices, \mathbf{K} and \mathbf{V} :

$$\text{Attention}(\mathbf{q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{q}\mathbf{K}^T}{\sqrt{d_k}} + \mathbf{m}\right)\mathbf{V} \quad (1)$$

where d_k is a scaling factor and \mathbf{m} is for padding or masking future tokens (when the value is $-\infty$).

Masking Operation. We propose attention head masking in encoder-decoder attentions, which blocks attentions to unimportant tokens, to better

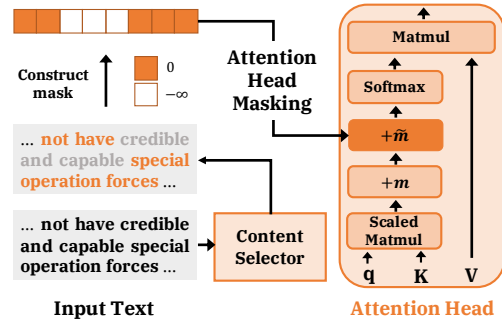


Figure 1: Illustration of attention head masking ($\tilde{\mathbf{m}}$).

concentrate multi-head attentions on salient input tokens. Importantly, it is activated during inference. Concretely, we add an $\tilde{\mathbf{m}}$ inside the softmax operator of Eq. 1, with implementation displayed in Fig. 1. The size of $\tilde{\mathbf{m}}$ is the same as the input length. If the i -th token is tagged as salient, the corresponding element in $\tilde{\mathbf{m}}$ is set to 0 (attendable to the attention heads), and $-\infty$ otherwise (hidden from these heads). The saliency labels can be predicted by an externally trained content selector.

4 Encoder-decoder Attentions and Content Selection

In this section, we first probe into the content selection behavior of each single head (§ 4.1), and then study the synergism among heads at the same layer (§ 4.2). In § 4.3, we analyze the attentions' focus.

Our analysis is conducted on CNN/DM (Hermann et al., 2015), NYT (Consortium and Company, 2008), and XSum (Narayan et al., 2018). We follow Lewis et al. (2020) for data preprocessing and train/validation/test splits on CNN/DM and XSum, and adopt the setups in Paulus et al. (2018) for NYT, except that we keep entities and numbers. The number of samples in training, validation, and test set are: 287,188, 13,367 and 11,490 for CNN/DM; 588,909, 32,716 and 32,703 for NYT; 204,045, 11,332 and 11,334 for XSum.

For experiments in this section, we create an *analysis set* of 1,000 random samples from the validation split of each dataset to reduce computational cost.

4.1 Content Selection Effects

First, we study the feasibility of using encoder-decoder attentions to inform content selection and subsequently boost summary informativeness. Concretely, we apply attention head masking based on oracle content selection labels (henceforth **or-**




Figure 2: ROUGE-1 F1 improvement with oracle masks for each head at each layer on the analysis set of CNN/DM. Overall, top layers see greater improvement than bottom layers. Layer 1 is the bottom layer connected with the word embeddings.

acle masking). Oracle labels are constructed by aligning a reference summary to the source article, where we iteratively find the longest common subsequences between the two.

Taking a fine-tuned BART model, we apply oracle masking on each head at each layer when decoding on the analysis set. The ROUGE score obtained in this setting is denoted as r_{ora} . We then apply uniform encoder-decoder attention weights over the source to build a baseline that mimics no content selection, inspired by Wiegrefe and Pinter (2019). This yields a ROUGE score of r_{uni} . The **content select effect** per head can thus be calculated as the ROUGE improvement, i.e., $r_{ora} - r_{uni}$.

Overall, it is more effective to constrain attentions to salient content at the top layers, according to the results on CNN/DM in Fig. 2. Specifically, with oracle masking, the top layer yields the most ROUGE-1 improvement. We observe similar trends on NYT and XSum (figures are in Appendix C). This indicates the feasibility of **leveraging attention head masking to improve summary informativeness**.

4.2 Synergism Analysis

Next, we study whether masking multiple heads can further boost content selection and whether they form synergy. On the left of Fig. 3, we show content selection effect by gradually applying oracle masking on more heads at each layer, with heads sorted based on individual ROUGE improvements. Notably, the most ROUGE-1 improvement is achieved by masking 15 (out of 16) heads at the top layer, suggesting a **strong collaborative effect on content selection by masking multiple heads**.

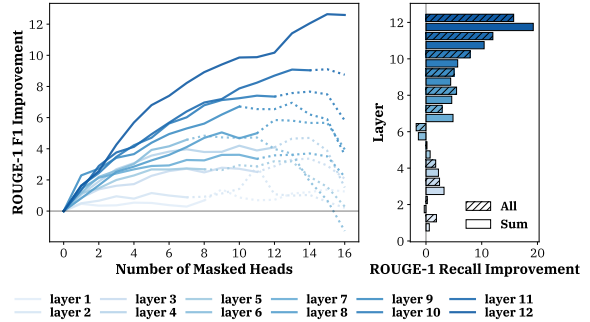


Figure 3: **[Left]** ROUGE-1 F1 improvement by incrementally applying oracle masking to the next head with most ROUGE improvement per layer on CNN/DM. Dotted lines indicate that the newly masked heads do not have individual ROUGE improvements. **[Right]** ROUGE-1 recall improvement by masking all heads vs. sum of improvement by masking each head separately on CNN/DM. Best viewed in color.

We further compare the ROUGE score gains between oracle masking on all heads and the sum of individual effects, illustrated on the right of Fig. 3. The discrepancies between the two values suggest that the heads may not be independent at pinpointing salient content. In Appendix D, we reach similar results on NYT and XSum.

Based on the above observations, we argue that **it is necessary to select layers and heads accordingly to achieve the best content selection effect**, with more summarization results reported in § 5.

4.3 Attention Focus

We further provide a fine-grained study on what types of words the heads attend to. Concretely, we consider each word generated during decoding, denoted as y . Given an attention head, we follow the highest attention weight to identify the input word x (“**attende**”). We study several categories of attendee x : (1) word in the reference (SALIENT); (2) CONTENT word; (3) the FIRST and LAST words in the document. For SALIENT and CONTENT, we further consider two subcategories: $x = y$ (COPY) and $x \neq y$ (NON-COPY). We then tally the occurrences of each type of attendees per head at each layer on the analysis set.

We show the percentages of COPY and NON-COPY SALIENT attendees, COPY CONTENT attendees, and FIRST attendees on CNN/DM in Fig. 4. As can be seen, top layers tend to focus on input tokens that will be generated as is, while bottom layers attend to salient words that are not used for current generation. Additionally, bottom layers fre-

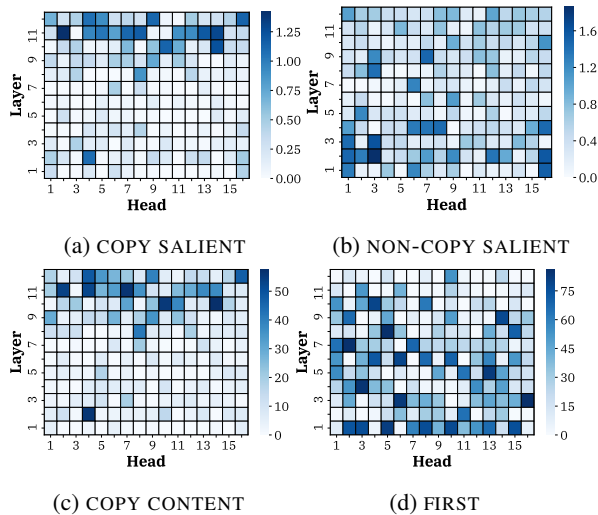


Figure 4: COPY and NON-COPY SALIENT attendee word percentages on the analysis set of CNN/DM. Top layers focus on words to be “copied”, while bottom layers attend to the broader salient context.

quently attend to the first token of the document, where bottom layers are more likely performing context gathering. On NYT and XSum (figures are in Appendix E), similar trends are observed except that the FIRST attendees are more focused by the top layers on NYT articles, where many of them start with all capitalized words.

5 Summarization Results with Attention Head Masking

In this section, we show how to leverage attention head masking and a content selector to improve summary informativeness on three datasets. We first train a binary sequence tagger for each dataset to label salient tokens in the source, used for **system masking** for attention heads. Our sequence tagger is a RoBERTa (Liu et al., 2019) encoder followed by a double layer multilayer perceptron (MLP) with a hyperbolic tangent activation function in between. To obtain the probability for each token, the MLP output is further fed into a sigmoid activation function. Details for training and decoding are in Appendix A.

The decision boundary for the sequence tagger is selected according to the F1 score calculated between the predicted tags and the ground-truth labels on the validation set. We search for the best decision boundary from 0.1 to 0.4, with a step size of 0.01. The final decision boundaries used for taggers trained on CNN/DM, NYT, XSum are 0.20, 0.24, and 0.18, achieving ROUGE-1 F1 of 43.70, 44.10, and 31.56, respectively.

Model	R-1	R-2	R-L
BERTSUM (Liu and Lapata, 2019)	42.13	19.60	39.18
UNILM (Dong et al., 2019)	43.33	20.21	40.51
PEGASUS (Zhang et al., 2019b)	44.17	21.47	41.11
PROPHETNET (Yan et al., 2020)	44.20	21.17	41.39
BART (ours)	44.19	21.20	40.98
+ attention head masking (ours)	45.54*	22.24*	42.44*

(a) CNN/DailyMail

Model	R-1	R-2	R-L
BOTTOMUP (Gehrmann et al., 2018)	47.38	31.23	41.81
DCA (Celikyilmaz et al., 2018)	48.08	31.19	42.33
BERTSUM (Liu and Lapata, 2019)	49.02	31.02	45.55
ASGARD (Huang et al., 2020)	51.29	34.97	48.26
BART (ours)	53.00	36.31	48.90
+ attention head masking (ours)	53.52*	36.69	49.24

(b) New York Times

Model	R-1	R-2	R-L
BERTSUM (Liu and Lapata, 2019)	38.81	16.50	31.27
PEGASUS	47.21	24.56	39.25
BART (ours)	45.36	22.30	37.11
+ attention head masking (ours)	45.35	22.31	37.15

(c) XSum

Table 1: Automatic evaluation with ROUGE. *: significantly better than BART with approximate randomization test ($p < 0.005$). Our method outperforms BART and previous models on CNN/DM and NYT.

To select which heads at which layers to mask, we employ a greedy selection strategy. On the analysis set, we gradually apply system masking on four heads with most ROUGE improvement according to the study in § 4.1, and we select the heads that achieve the highest sum of ROUGE-1 F1 and ROUGE-2 F1. We apply four heads each time to reduce computational cost of hyperparameter searching. Heads selected for each dataset are in Appendix B.

In-domain Results. Table 1 shows that applying our attention head masking technique on BART obtains significantly better results on CNN/DM and NYT, compared to several top performing abstractive summarization models trained with large Transformers. The improvement is more pronounced for CNN/DM than the other two datasets. We believe this is due to the difference in abstractiveness among the three datasets. CNN/DM has more extractive summaries compared to the other datasets (Grusky et al., 2018), suggesting attention head masking is more effective on extractive datasets. Notably, PEGASUS is pre-trained with 3.8TB of news articles, the BART model used in our work is only pre-trained with 160GB of a com-

	w/ masking	w/o masking	Tie
Informativeness	36.0%	19.3%	44.7%
Faithfulness	10.0%	7.3%	82.7%

Table 2: Percentages of summaries with and without attention head masking favored by annotators on informativeness and faithfulness. The Krippendorff’s α for informativeness and faithfulness are 0.30 and 0.47.




Figure 5: Results on CNN/DM with different sizes of training data. Our method consistently improves the summarizer.

bination of news, books, stories, and web text. The large size of the pre-training data might be a big contributor to the better performance by PEGASUS on XSum.

For **human evaluation**, we hire three fluent English speakers to rate 50 pairs of summaries generated with and without attention head masking based on BART for **informativeness** and **faithfulness**. Informativeness measures how well the summary captures salient content from the article, while faithfulness indicates whether the summary correctly reflects the content in the source article. The annotators are asked to determine if attention head masking improves any of the two aspects. As shown in Table 2 where all ratings by three judges are considered, summaries generated with attention head masking are considered to have better informativeness, but no substantial improvement on faithfulness is observed.

Limited Training Data. Next, we study if our masking technique is still effective if given limited training samples. We use the limited training samples to train both the summarizer and the content selector. As can be seen in Fig. 5, our masking technique consistently increases ROUGE scores with varying amounts of training data. Notably, our model trained on only 30K samples (with attention head masking) outperforms the model trained on the full dataset, suggesting that directly informing content selection is more data-efficient than model fine-tuning on more summaries.

Selector Training Data	R-1	R-2	R-L
No masking	31.11	14.68	28.19
10K	34.98	17.95	31.87
100K	34.71	17.70	31.61
589K (full)	35.13	18.07	32.03

Table 3: Results on NYT summaries generated by BART trained on CNN/DM, with masks predicted by content selectors trained on different sizes of NYT data.

Cross-domain Results. Finally, we show results on NYT using BART fine-tuned on CNN/DM, with system masks predicted by a tagger trained on different sizes of NYT samples (Table 3). Using a selector trained with only 10k of target domain samples, we already significantly improve the performance by BART trained on CNN/DM only.

6 Conclusion

We propose attention head masking that constrains encoder-decoder attention heads to attend to salient tokens, to inform content selection in abstractive summarization. With this technique, we first demonstrate the relation between encoder-decoder attentions and content selection behaviors. With system masks predicted by external content selectors, we show that attention head masking can consistently improve ROUGE scores over competitive summarization models on three benchmarks. Summaries generated with attention head masking are also preferred by human judges more frequently. Additional experiments demonstrate that our method is more data-efficient and effective on both in-domain and cross-domain settings.

Acknowledgements

This research is supported in part by National Science Foundation through Grant IIS-1813341, and by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via contract # FA8650-17-C-9116. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein. We thank three anonymous reviewers for their constructive suggestions.

References

- Asli Celikyilmaz, Antoine Bosselut, Xiaodong He, and Yejin Choi. 2018. [Deep communicating agents for abstractive summarization](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1662–1675, New Orleans, Louisiana. Association for Computational Linguistics.
- Yen-Chun Chen and Mohit Bansal. 2018. [Fast abstractive summarization with reinforce-selected sentence rewriting](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 675–686, Melbourne, Australia. Association for Computational Linguistics.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. [What does BERT look at? an analysis of BERT’s attention](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286, Florence, Italy. Association for Computational Linguistics.
- Linguistic Data Consortium and New York Times Company. 2008. *The New York Times Annotated Corpus*. LDC corpora. Linguistic Data Consortium.
- Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. [Unified language model pre-training for natural language understanding and generation](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 13063–13075. Curran Associates, Inc.
- Sebastian Gehrmann, Yuntian Deng, and Alexander Rush. 2018. [Bottom-up abstractive summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4098–4109, Brussels, Belgium. Association for Computational Linguistics.
- Max Grusky, Mor Naaman, and Yoav Artzi. 2018. [Newsroom: A dataset of 1.3 million summaries with diverse extractive strategies](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 708–719, New Orleans, Louisiana. Association for Computational Linguistics.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. [Teaching machines to read and comprehend](#). In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 1693–1701. Curran Associates, Inc.
- Luyang Huang, Lingfei Wu, and Lu Wang. 2020. [Knowledge graph-augmented abstractive summarization with semantic-driven cloze reward](#).
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Goro Kobayashi, Tatsuki Kuribayashi, Sho Yokoi, and Kentaro Inui. 2020. [Attention is not only a weight: Analyzing transformers with vector norms](#).
- Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. 2019. [Revealing the dark secrets of BERT](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4365–4374, Hong Kong, China. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Wei Li, Xinyan Xiao, Yajuan Lyu, and Yuanzhuo Wang. 2018. [Improving neural abstractive document summarization with explicit information selection modeling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1787–1796, Brussels, Belgium. Association for Computational Linguistics.
- Xintong Li, Guanlin Li, Lema Liu, Max Meng, and Shuming Shi. 2019. [On the word alignment from neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1293–1303, Florence, Italy. Association for Computational Linguistics.
- Yang Liu and Mirella Lapata. 2019. [Text summarization with pretrained encoders](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3730–3740, Hong Kong, China. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *arXiv preprint arXiv:1907.11692*.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. [Are sixteen heads really better than one?](#) In *Advances in Neural Information Processing Systems*,

- volume 32, pages 14014–14024. Curran Associates, Inc.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. [Don't give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807, Brussels, Belgium. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of NAACL-HLT 2019: Demonstrations*.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2018. [A deep reinforced model for abstractive summarization](#). In *International Conference on Learning Representations*.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2019. [MASS: Masked sequence to sequence pre-training for language generation](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5926–5936, Long Beach, California, USA. PMLR.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30, pages 5998–6008. Curran Associates, Inc.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Senrich, and Ivan Titov. 2019. [Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5797–5808, Florence, Italy. Association for Computational Linguistics.
- Sarah Wiegrefe and Yuval Pinter. 2019. [Attention is not not explanation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 11–20, Hong Kong, China. Association for Computational Linguistics.
- Jiacheng Xu and Greg Durrett. 2019. [Neural extractive text summarization with syntactic compression](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3292–3303, Hong Kong, China. Association for Computational Linguistics.
- Yu Yan, Weizhen Qi, Yeyun Gong, Dayiheng Liu, Nan Duan, Jiusheng Chen, Ruofei Zhang, and Ming Zhou. 2020. [Prophetnet: Predicting future n-gram for sequence-to-sequence pre-training](#).
- Haoyu Zhang, Jingjing Cai, Jianjun Xu, and Ji Wang. 2019a. [Pretraining-based natural language generation for text summarization](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 789–797, Hong Kong, China. Association for Computational Linguistics.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J. Liu. 2019b. [Pegasus: Pre-training with extracted gap-sentences for abstractive summarization](#).
- Qingyu Zhou, Nan Yang, Furu Wei, and Ming Zhou. 2017. [Selective encoding for abstractive sentence summarization](#). *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.

A Training and Decoding Settings

When training the sequence taggers, we minimize the average binary cross-entropy of each token’s selection probability relative to the ground-truth label. The parameters of the RoBERTa encoder are fixed. We set the learning rate to 5×10^{-4} and batch size to 128. Unless specified, all the models in this paper are trained with Adam (Kingma and Ba, 2015) optimizer and training will be stopped if there is no improvement on the validation set for 2 consecutive epochs.

For BART models, we follow the instructions provided by Fairseq (Ott et al., 2019) to set the training hyperparameters on CNN/DM and XSum. We use the same hyperparameters for CNN/DM and NYT, except that we adopt a linear learning rate decay of 30,000 steps in total for NYT.


During test, we use a beam size of 5, 5, 6 for CNN/DM, NYT, and XSum, respectively. To reduce computational cost, we use beam size 1 for our analysis experiments on all datasets. The length penalties are 2.0, 1.5 and 1.0 for CNN/DM, NYT, and XSum, following Lewis et al. (2020). We set the minimal and maximal lengths during decoding as: 55 and 140 for CNN/DM, 0 and 140 for NYT, and 10 and 60 for XSum.

B Head Selection


For CNN/DM, we apply masking to all heads at layer 1. The ROUGE-1/2/L F1 on the analysis set are 36.43/16.02/33.59.

For NYT, we apply masking to 12 heads at layer 3. The indices of heads are: 3, 4, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15. The ROUGE-1/2/L F1 on the analysis set are 55.27/39.20/48.16.

For XSum, we apply masking to 12 heads at layer 3. The indices of heads are: 1, 2, 3, 4, 6, 7,



(a) XSum



(b) NYT

Figure 6: ROUGE-1 improvement with oracle masks for each head at each layer on the analysis sets of XSum and NYT.

8, 9, 11, 13, 14, 15. The ROUGE-1/2/L F1 on the analysis set are 45.77/22.82/37.60.

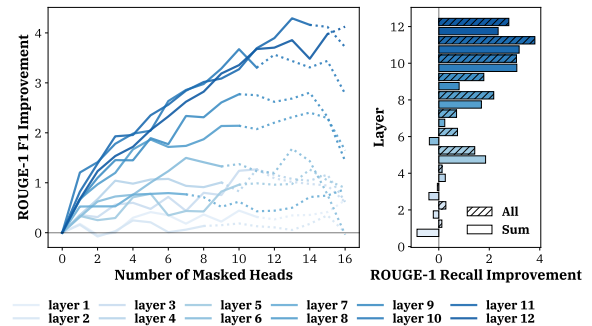
C Content Selection Effects on XSum and NYT

The content selection effects for BART models fine-tuned on XSum and NYT, measured by the ROUGE improvement from the uniform attention weight setting to the oracle masking setting, are shown in Fig. 6.

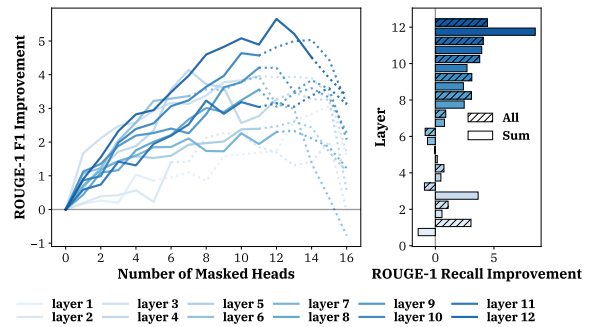
On all three datasets, it is more effective to constrain attentions to salient content at the top layers. Especially, the top layer yields the most ROUGE-1 improvement. Moreover, the ROUGE improvement by a specific head varies among different datasets.

D Additional Results for Synergism Analysis

We show the synergism analysis for models fine-tuned on XSum and NYT in Fig. 7. They both echo the observation on CNN/DM that multiple heads



(a) XSum



(b) NYT

Figure 7: [Left] ROUGE-1 F1 improvement by incrementally applying oracle masking to the next head with most ROUGE improvement per layer on XSum and NYT. Dotted lines indicate that the newly masked heads do not have individual ROUGE improvements. [Right] ROUGE-1 recall improvement by masking all heads vs. sum of improvement by masking each head separately on XSum and NYT. Better displayed with color.

have strong collaborative effects and heads may not be independent at pinpointing different salient content.

E Attention Focus

We show the percentages of each type of attendees on the analysis sets of XSum, NYT, and CNN/DM in Fig. 8, Fig. 9, and Fig. 10, respectively. We find that heads have similar focus for salient words, content words, and the last word across different datasets. Interestingly, the attention focus for the first word on NYT is different from other datasets. On NYT, many articles start with all capitalized words, which might become the focus of some heads.

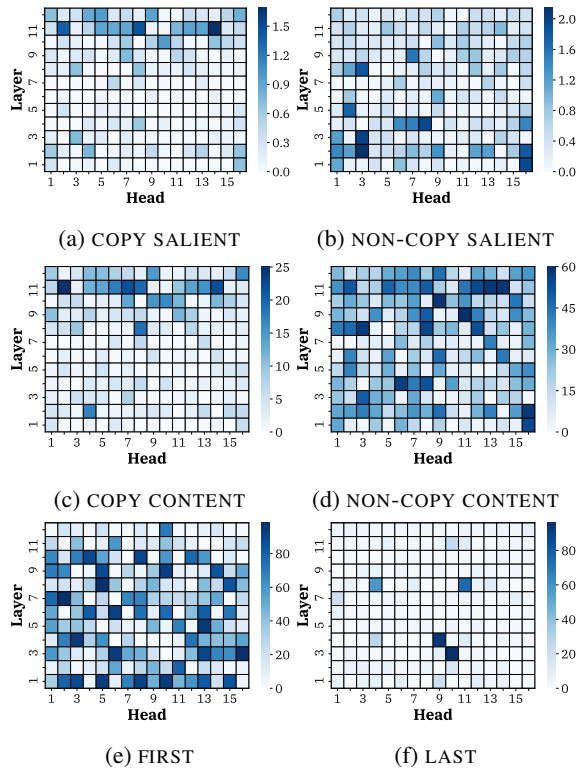


Figure 8: Percentages of COPY SALIENT, NON-COPY SALIENT, COPY CONTENT, NON-COPY CONTENT, FIRST and LAST attendees for each head at each layer on the analysis set of XSum.

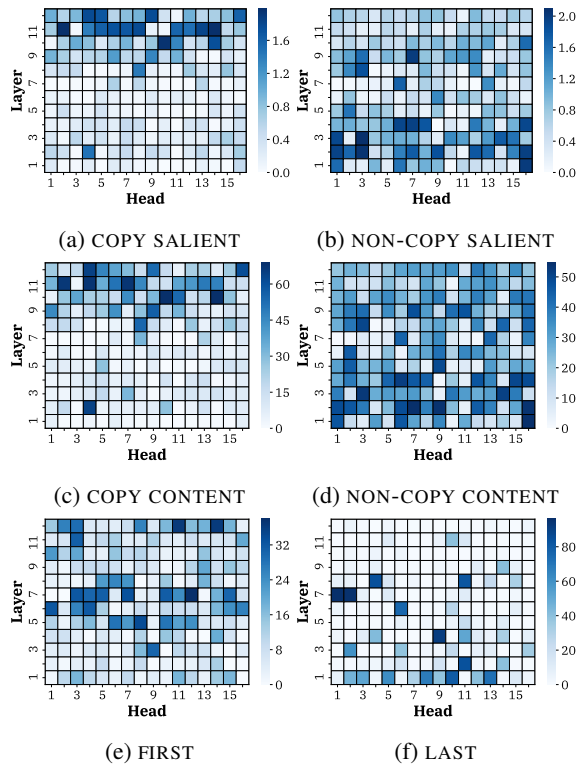


Figure 9: Percentages of COPY SALIENT, NON-COPY SALIENT, COPY CONTENT, NON-COPY CONTENT, FIRST and LAST attendees for each head at each layer on the analysis set of NYT.

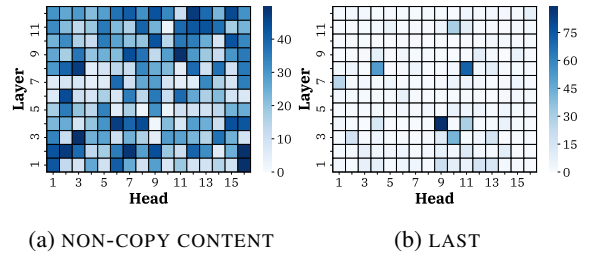


Figure 10: Percentages of NON-COPY CONTENT and LAST attendees for each head at each layer on the analysis set of CNN/DM.