

SentBS: Sentence-level Beam Search for Controllable Summarization

Chenhui Shen^{*1,2} Liying Cheng^{*1,3} Lidong Bing^{†1} Yang You² Luo Si¹

¹DAMO Academy, Alibaba Group ²National University of Singapore

³Singapore University of Technology and Design

{chenhui.shen, liying.cheng}@alibaba-inc.com

{l.bing, luo.si}@alibaba-inc.com youy@comp.nus.edu.sg

Abstract

A wide range of control perspectives have been explored in controllable text generation. Structure-controlled summarization is recently proposed as a useful and interesting research direction. However, current structure-controlling methods have limited effectiveness in enforcing the desired structure. To address this limitation, we propose a sentence-level beam search generation method (SentBS), where evaluation is conducted throughout the generation process to select suitable sentences for subsequent generations. We experiment with different combinations of decoding methods to be used as sub-components by SentBS and evaluate results on the structure-controlled dataset MReD. Experiments show that all explored combinations for SentBS can improve the agreement between the generated text and the desired structure, with the best method significantly reducing the structural discrepancies suffered by the existing model, by approximately 68%.¹

1 Introduction

Controllable text generation is receiving increasing attention due to its wide range of applications. Depending on the use cases, the controllable generation tasks may focus on a wide range of control perspectives, such as entities (Narayan et al., 2022; Fan et al., 2018a), aspects (Hayashi et al., 2021), and keywords (Wang et al., 2021; He et al., 2020). Recently, Shen et al. (2022) propose a sentence-level labeled meta-review dataset, MReD, for the controllable summarization task from a new control perspective that focuses on controlling the structure of the output summary. The input consists of several reviews on the same research paper, and a control sequence specifying the desired summary

structure. For instance, with a control sequence of “abstract | strength | decision”, the generated output should be composed of a sentence that summarizes the contents of the paper, followed by a sentence discussing the strengths, then the last sentence giving the final decision.

Previous controllable summarization models are commonly fine-tuned on pre-trained transformer architectures (Vaswani et al., 2017) such as BART (Lewis et al., 2020) and Pegasus (Zhang et al., 2020a), with the control signals merged into the text input or prompts (Shen et al., 2022; Narayan et al., 2022; He et al., 2020; Keskar et al., 2019; Fan et al., 2018a). Previous works mainly focus on improving the summary’s similarity with the gold reference, leaving room for further improvement on the controllability. In particular, the best-performing model on the recently released MReD dataset still generates around 29% of the sentences that do not follow the control structure², which is far from satisfactory.

In this paper, we explore how to enhance the structure-controllability in summarization. Specifically, we notice the following possible pitfalls in the existing summarization models. First, those models usually treat generation as a standalone process, which continuously generates the tokens solely based on the logits predictions, without stopping to reconsider whether the generated sequences satisfy the control signals. Moreover, autoregressive models can suffer from error propagation in generation due to self-attention (Vaswani et al., 2017). Therefore, if the previous sequences are not well-controlled, subsequent generations may deviate further from the desired output. Motivated by this, we propose the **S**entence-level **B**eam **S**earch (SentBS) method to address the identified issues. For generating each sentence, SentBS first pro-

*Chenhui and Liying are under the Joint PhD Program between Alibaba and their corresponding universities.

† Corresponding author.

¹Our code and data will be available at <https://github.com/Shen-Chenhui/SentBS>.

²For instance, the generated sentence may discuss the weakness of the paper even though it corresponds to a “strength” control label.

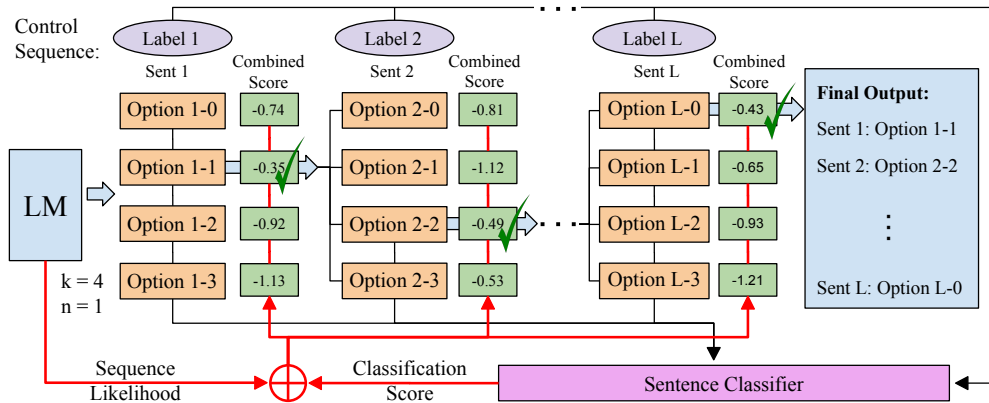


Figure 1: Illustration of SentBS. The score values are for illustration purposes only. For simplicity, we only illustrate for $k = 4$ and $n = 1$.

duces multiple sentence options, evaluates and selects the best sentence according to both the control structure as well as the model’s log-likelihood, then continues the generation for the next sentence.

Experiments show that SentBS can significantly improve the model’s structure-controllability. In particular, our best setting removes up to 68% of control mistakes produced by the existing model on MReD without compromising the summarization quality. The human evaluation further proves that SentBS significantly improves the fluency of summaries.

To summarize, our main contributions are: (1) To the best of our knowledge, we are the first to conduct sentence-by-sentence controlled generation for structure-controllable summarization. (2) We propose SentBS, which conducts text generation with continuous evaluation on the sentence level with respect to the control requirements. This method can be easily applied to existing autoregressive models. (3) Experiments show that SentBS significantly increases the model’s structure controllability while preserving the summarization quality.

2 Related Work

Conditional Text Generation. Large pretrained language models have shown impressive performance (Lewis et al., 2020; Zhang et al., 2020a) on generation tasks. Many controllable generation tasks (Shen et al., 2022; Narayan et al., 2022; Chia et al., 2022; He et al., 2020; Cheng et al., 2020b; Keskar et al., 2019; Fan et al., 2018a) make use of these models by merging the control signals into either the sources or targets. Our method differs from these approaches by breaking the controlled generation process down to the sentence level, and we make explicit use of the control signals during

the generation phase.

Decoding Methods. Common decoding methods include beam search (Meister et al., 2020; Stahlberg and Byrne, 2019; Graves, 2012), nucleus sampling (Holtzman et al., 2019), and beam sampling (Caccia et al., 2019) (see more details in Appendix A). In our SentBS method, we use these methods³ as sub-components for generating multiple sentence options. Hopkins and Kiela (2017), Ghazvininejad et al. (2016) and Zhang and Lapata (2014) also works on decoding methods that enforce linguistic constraints on the generation structure. However, none of these works use a sentence-by-sentence generation and evaluation strategy based on whether the semantic content of the generation meets the control requirements.

3 Sentence-Level Beam Search

Given the nature of autoregressive models, errors from previously generated tokens at inference time can be easily propagated to affect the subsequent tokens. To better control the structure, we propose SentBS, to evaluate and select the outputs at the sentence level during generation. SentBS can be easily applied to existing generation models during inference only. In this section, we explain how SentBS leverages the control sequence where each control label corresponds to the desired sentence category. Please see more details on the task and dataset in Section 4.1.

3.1 Method Details

We illustrate the generation process of SentBS in Figure 1. Given a generative model, a control sequence consisting of several labels, and the con-

³https://huggingface.co/docs/transformers/internal/generation_utils.

catenated review texts, SentBS first generates k sentence options in parallel (e.g., Option 1-0/1-1/1-2/1-3) using multiple decoding methods such as beam search, nucleus sampling, and beam sampling⁴. We calculate a combined score (e.g., Option 1-1 has the best combined score of -0.35) for each sentence by adding its normalized sequence likelihood with a classifier-predicted probability score of the sentence belonging to the required category (see classifier details in Appendix B). According to the combined score, we select the top n sentences and feed them individually into the decoder as prompts⁵ for generating the next sentence.

We generate k subsequent sentence options for each of the n prompts, resulting in a total of $k \cdot n$ sentences. These newly generated sentences will be concatenated after the corresponding prompts to form new prompts for the next generation, and new scores will be calculated to select the next top n prompts. In particular, the sequence likelihood score is recalculated for the full sequence. The same generation process continues until all sentences required in the control sequence are produced. This generation process is similar to beam search, except that instead of merely selecting tokens based on log-likelihood, we also conduct the selection on the sentence level based on both the control requirements and the sequence likelihood.

3.2 Method Applications

Although SentBS evaluates the structural requirements on a sentence level, it can be easily extended to segment-level control methods where each control label corresponds to one or more sentences of the same category. In this case, only the first generated sentence has an explicit label, whereas the subsequent sentences can either have the same label as their previous sentences, or carry the next label in the control sequence. Therefore, after generating the first sentence, we can again apply SentBS to the generation of subsequent sentences by comparing the probabilities corresponding to the two allowed labels. For instance, given a control sequence of “abstract | strength | decision”, after generating the first “abstract” sentence, we look at the classification score on both “abstract” and “strength” for the second sentence, and assign the label with higher

⁴We modify these decoding methods to generate at the sentence level.

⁵Note that unlike Gehman et al. (2020) and Sheng et al. (2019) that use user-designed prompts, we use the model-generated sentences for continued generation.

probability to it. Thus, each sentence option for the second sentence is assigned a label of either “abstract” or “strength”, depending on the generated content. In the same manner, we can generate the subsequent sentences until the model gives the stop signal (i.e., the “eos” token).

Following the logic above, we can further extend SentBS to other controlled generation tasks on a paragraph level (e.g. style transfer), as long as the control requirements can be evaluated on the sentence level. For instance, for converting non-formal text to formal text, we can score each sentence according to a combined score of sequence log-likelihood and the degree of formality. Nevertheless, this is beyond the scope of this paper and we leave the investigation of such applications to future work.

4 Experiments

4.1 Task and Dataset

We carry out experiments on the benchmark MReD dataset (Shen et al., 2022). Each sentence in the target meta-review is annotated with a category label according to its main function from the following: “abstract”, “strength”, “weakness”, “suggestion”, “rating summary”, “rebuttal process”, “ac disagreement”, “decision”, and “misc” (see Appendix C).

This dataset comes with two control task settings for structure-controllable summarization: *Sent-Ctrl*, which uses each label in the control sequence to indicate a category for one summary sentence in the summary, and *Seg-Ctrl*, which uses each label in the control sequence to represent a segment of summary sentences of the same category. The ultimate goal for both settings is to generate meta-review passages derivable from the reviews while at the same time obeying the required structure.

4.2 Evaluation Metrics

We evaluate the summarization quality by reporting the F1 scores of the standard Rouge-1/2/L (Lin, 2004) and BERTScore⁶ (Zhang et al., 2020b). In terms of the structural agreement, we report both the human and automatic evaluation results. For the human evaluation, we use “structure similarity” following (Shen et al., 2022), while the automatic evaluation is conducted by passing the summary into a LSTM-CRF classifier (Lample et al., 2016) (see appendix D). We also aggregate the total edit distance between the summary structure and the

⁶https://github.com/Tiiiger/bert_score

	Structure↑	Edits↓	BERTScore↑	R ₁ / R ₂ / R _L ↑
Sent-Ctrl (Shen et al., 2022)				
- Reported	0.706	-	-	38.73 / 10.82 / 23.05
- Reproduced	0.737	1470.0	0.8631	38.97 / 11.19 / 23.48
Sent-Ctrl + SentBS (ours)				
Nucleus sampling ($k = 4$)	0.852	785.0	0.8627	36.50 / 9.11 / 21.76
Nucleus sampling ($k = 5$)	0.887	600.3	0.8637	36.96 / 9.28 / 22.25
Nucleus sampling ($k = 6$)	0.891	570.3	0.8634	36.79 / 9.21 / 22.12
Nucleus sampling ($k = 7$)	0.893	567.7	0.8642	37.00 / 9.34 / 22.30
Nucleus sampling ($k = 8$)	0.885	605.3	0.8642	37.07 / 9.62 / 22.47
Beam Sampling ($k = 4$)	0.771	1142.0	0.8627	37.44 / 10.80 / 23.11
Beam Sampling ($k = 5$)	0.803	993.7	0.8626	37.82 / 10.87 / 23.09
Beam Sampling ($k = 6$)	0.802	998.7	0.8621	37.71 / 10.86 / 22.93
Beam Sampling ($k = 7$)	0.795	1011.0	0.8617	37.88 / 10.84 / 22.83
Beam Sampling ($k = 8$)	0.794	1007.0	0.8613	37.67 / 10.77 / 22.76
Beam search + Nucleus sampling ($k = 4$)	0.874	665.0	0.8644	38.23 / 10.47 / 23.22
Beam search + Nucleus sampling ($k = 5$)	0.887	609.7	0.8645	38.37 / 10.54 / 23.27
Beam search + Nucleus sampling ($k = 6$)	0.894	569.3	0.8648	38.36 / 10.67 / 23.35
Beam search + Nucleus sampling ($k = 7$)	0.904	519.3	0.8648	38.30 / 10.62 / 23.37
Beam search + Nucleus sampling ($k = 8$)	0.915	467.7	0.8647	38.32 / 10.66 / 23.42
Beam search + Beam sampling + Nucleus sampling ($k = 4$)	0.839	833.7	0.8645	38.33 / 10.96 / 23.48
Beam search + Beam sampling + Nucleus sampling ($k = 5$)	0.868	700.3	0.8644	38.39 / 10.95 / 23.54
Beam search + Beam sampling + Nucleus sampling ($k = 6$)	0.883	634.0	0.8644	38.32 / 10.90 / 23.45
Beam search + Beam sampling + Nucleus sampling ($k = 7$)	0.885	607.7	0.8643	38.38 / 10.85 / 23.45
Beam search + Beam sampling + Nucleus sampling ($k = 8$)	0.893	573.0	0.8640	38.43 / 10.88 / 23.38
Seg-Ctrl (Shen et al., 2022)				
-Reported	0.623	-	-	36.38 / 10.04 / 21.90
-Reproduced	0.755	855.0	0.8604	36.69 / 10.44 / 22.29
Seg-Ctrl + SentBS (ours)	0.887	394.3	0.8601	36.75 / 10.35 / 22.51

Table 1: Main results. We divide the table into 2 sections using double horizontal lines. The top section shows the Sent-Ctrl baseline and Sent-Ctrl with various settings of SentBS, and the bottom section shows the Seg-Ctrl baseline and Seg-Ctrl with SentBS. For the latter section, we use the “Beam search + Beam sampling + Nucleus sampling” setting and $k = 8$ for SentBS.

gold structure for the full test set to quantitatively evaluate the control performance. This additional metric is named “total edits”.

4.3 Experimental Settings

We follow Shen et al. (2022) and reproduce the results using the same reported settings (see Appendix E) as baselines. We use SentBS during inference for *Sent-Ctrl* and *Seg-Ctrl*. Since Shen et al. (2022) uses beam search with a beam size of 4 during inference, we set $n = 4$ for SentBS to maintain 4 sentences for each generation step as well. We experiment with different combinations of beam search, nucleus sampling, and beam sampling as decoding sub-components, and explore various k values up to 8 (see more in Appendix F). We set *top-p* for nucleus sampling (see Appendix A) to 0.9. All reported results for SentBS are the average of 3 runs to mitigate the uncertainty caused by sampling methods.

4.4 Main Results

4.4.1 SentBS for *Sent-Ctrl*

In Table 1, we first reproduce the “Sent-Ctrl” baseline. We discover that the generated structures hardly follow the corresponding control sequences in a strict sense (see Appendix Table 5).

Next, we evaluate SentBS with different decod-

ing methods and various k values on the *Sent-Ctrl* model. Besides SentBS with single decoding methods (“Nucleus sampling” and “Beam sampling”), we also explore combined decoding methods: “Beam search + Nucleus sampling”, where we use beam search to generate 1 sentence option, nucleus sampling for the rest; and “Beam search + Beam sampling + Nucleus sampling”, where we use beam search for 1 option, beam sampling for $\frac{k}{2}$ options⁷ and nucleus sampling for the rest.

SentBS with single decoding methods. In general, both single-method settings achieves better generation structures as compared to the *Sent-Ctrl* baseline. For nucleus sampling, the generation quality improves with increasing k values (except for $k = 6$), and the structure similarity is reasonably good for $k > 4$. With more sampled sentences, SentBS has more candidates to choose from to better suit the quality and structural requirements. However, beam sampling has mediocre structure scores regardless of k values. This is because beam sampling only keeps the top k sampled outputs with the best sequence likelihood. Therefore, text diversity may suffer and there may not be sufficient sentences of different label categories for SentBS

⁷If $\frac{k}{2}$ is not an integer, we use the nearest smaller integer of this value. For instance, if $k = 7$, we use beam sampling to generate 3 options.

Beam Size	Structure \uparrow	Edits \downarrow	BERTScore \uparrow	R ₁ / R ₂ / R _L \uparrow
4	0.737	1470	0.8631	38.97 / 11.19 / 23.48
5	0.732	1424	0.8625	38.81 / 10.92 / 23.34
6	0.723	1456	0.8618	38.67 / 10.85 / 23.19
7	0.727	1427	0.8613	38.55 / 10.73 / 23.00
8	0.722	1448	0.8608	38.29 / 10.68 / 22.90

Table 2: Beam Search generation results on MReD using Sent-Ctrl with increasing beam sizes.

to choose from. Nevertheless, neither nucleus sampling nor beam sampling produces sufficiently fluent text as compared to *Sent-Ctrl* (according to Rouge 1/2/L scores). Nucleus sampling produces much less fluent generations as compared to beam sampling, possibly caused by accidentally sampling some of the less likely tokens which are most probably discarded by beam sampling.

SentBS with multiple decoding methods. Various decoding methods combined can produce a good range of sentence options for both quality and diversity, so that SentBS can generate a much better structure with good text quality. In particular, “Beam search + Nucleus sampling” consistently achieves the best BertScores under different k values, and has the best structure scores for $k > 5$. At $k = 8$, more than 68% of the total edits from *Sent-Ctrl* can be avoided. “Beam search + Beam sampling + Nucleus sampling” has better Rouge 1/2/L performances than “Beam search + Nucleus sampling”, and has the best Rouge-L score when $k = 5$. The essential difference between these two settings is that, for the same k values, the latter replaces some of the former’s sentence options from nucleus sampling with beam sampling. Therefore, it is not surprising that the latter has a slightly worse structure. Nevertheless, SentBS with either of the combined decoding strategies far exceeds *Sent-Ctrl* in terms of control, while at the same time achieving higher contextual BERTScores and comparable Rouge-L scores, although slightly lower Rouge-1/2 scores. For an abstractive dataset like MReD, contextual similarity and Rouge-L may be more important metrics to evaluate fluency.

Effect of Increasing k Values The widely-used beam search method usually experiences worse generation quality with a large beam size. On MReD, we also observe that the conventional beam search method experiences consistent decreases in the BertScores and the Rouge 1/2/L scores with increasingly large beam numbers (Table 2). Nevertheless, it is exciting to see that SentBS exhibits relatively stable or even improving generation qual-

	Sent-Ctrl	SentBS
Fluency	0.520	0.780*
Content Relevance	0.680	0.700
Structure Similarity	0.699	0.838**
Decision Correctness	0.740	0.680

Table 3: Human evaluation. * indicates for p-value < 0.05 , ** for p-value < 0.0001 by Welsh’s t-test.

ity and structure with increasing k values (Table 1). Therefore, it may be easier to apply SentBS to other datasets with less need to tune the k values.

4.4.2 SentBS for *Seg-Ctrl*

We also report the results of applying SentBS on the “Seg-Ctrl” setting in Table 1. For simplification, we show the results of SentBS with the “Beam search + Beam sampling + Nucleus sampling” decoding strategy at $k = 8$. Again, SentBS helps to significantly improve the generation structure, while achieving good generation quality.

4.5 Human Evaluation

We manually evaluate the “Beam search + Beam sampling + Nucleus sampling” setting of SentBS at $k = 8$ and the *Sent-Ctrl* baseline on 50 random test instances. 2 Human judges are asked to judge the summary in terms of fluency, content relevance, structure similarity, and decision correctness. All scores are in the range of 0 to 1 (see Appendix G).

From Table 3, we can see that SentBS is significantly better than *Sent-Ctrl* in terms of fluency and structure similarity. SentBS is also more preferred (although not significantly) by human judges for content relevance. However, it has a lower (not significant) decision correctness. The model may sample tokens for the wrong decisions, while still satisfying the control requirement. One way to mitigate this impact is to train the classifier to differentiate various decisions and explicitly use the correct decision in the control sequence.

5 Conclusions

In this work, We refocus the attention of controllable summarization from the similarity with gold toward the model’s controllability. We propose a simple and effective method, SentBS, to enhance the structure-controllability of the Transformers BART model trained on MReD. With extensive experiments and human evaluation, we show that without the need for retraining, SentBS can significantly improve the controllability of summary structure while achieving good generation quality.

Limitations

SentBS requires a combination of maximization-based and sampling-based decoding strategies to work well. It requires an excessive amount of sentences not used in the final output to be generated. This wastes a lot of computational resources and makes the generation inefficient. For instance, the standard beam search for *Sent-Ctrl* requires around 1h to complete. When we use $k = 4$ and $n = 4$ for nucleus sampling and beam sampling, the decoding time is extended to around 5h. For the setting of $k = 8$ and $n = 4$ for combined decoding methods, the time required for generating the full test set is 16h. This is also due to the fact that the Huggingface Transformers model does not accept variable-length prompts for parallel decoding, so the generation of sentence options for each possible prompt needs to be carried out one by one. We leave the investigation of the above issues for future work.

So far, we have shown that the performance of SentBS is highly dependent on the nature and combination of the decoding methods. We will leave the exploration of how to better select and combine decoding methods to future work.

References

- David H Ackley, Geoffrey E Hinton, and Terrence J Sejnowski. 1985. A learning algorithm for boltzmann machines. *Cognitive science*.
- Jianzhu Bao, Bin Liang, Jingyi Sun, Yice Zhang, Min Yang, and Ruifeng Xu. 2021. Argument pair extraction with mutual guidance and inter-sentence relation graph. In *Proceedings of EMNLP*.
- Jianzhu Bao, Jingyi Sun, Qinglin Zhu, and Ruifeng Xu. 2022. Have my arguments been replied to? argument pair extraction as machine reading comprehension. In *Proceedings of ACL*.
- Chaitanya Bhatia, Tribikram Pradhan, and Sukomal Pal. 2020. Metagen: An academic meta-review generation system. In *Proceedings of ACM-SIGIR*.
- Massimo Caccia, Lucas Caccia, William Fedus, Hugo Larochelle, Joelle Pineau, and Laurent Charlin. 2019. Language gans falling short. In *ICLR*.
- Liyang Cheng, Lidong Bing, Qian Yu, Wei Lu, and Luo Si. 2020a. Ape: argument pair extraction from peer review and rebuttal via multi-task learning. In *Proceedings of EMNLP*.
- Liyang Cheng, Dekun Wu, Lidong Bing, Yan Zhang, Zhanming Jie, Wei Lu, and Luo Si. 2020b. Ent-desc: Entity description generation by exploring knowledge graph. In *Proceedings of EMNLP*.
- Liyang Cheng, Tianyu Wu, Lidong Bing, and Luo Si. 2021. Argument pair extraction via attention-guided multi-layer multi-cross encoding. In *Proceedings of ACL-IJCNLP*.
- Yew Ken Chia, Lidong Bing, Soujanya Poria, and Luo Si. 2022. Relationprompt: Leveraging prompts to generate synthetic data for zero-shot relation triplet extraction. In *Findings of ACL*.
- Angela Fan, David Grangier, and Michael Auli. 2018a. Controllable abstractive summarization. In *Proceedings of WNGT*.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018b. Hierarchical neural story generation. In *Proceedings of ACL*.
- Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A Smith. 2020. Realtocixityprompts: Evaluating neural toxic degeneration in language models. In *Findings of EMNLP*.
- Marjan Ghazvininejad, Xing Shi, Yejin Choi, and Kevin Knight. 2016. Generating topical poetry. In *Proceedings of EMNLP*.
- Alex Graves. 2012. Sequence transduction with recurrent neural networks. *arXiv preprint arXiv:1211.3711*.
- Hiroaki Hayashi, Prashant Budania, Peng Wang, Chris Ackerson, Raj Neervannan, and Graham Neubig. 2021. Wikiasp: A dataset for multi-domain aspect-based summarization. *TACL*.
- Junxian He, Wojciech Kryściński, Bryan McCann, Nazneen Rajani, and Caiming Xiong. 2020. Ctrlsum: Towards generic controllable text summarization. *arXiv preprint arXiv:2012.04281*.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. The curious case of neural text degeneration. In *Proceedings of ICLR*.
- Jack Hopkins and Douwe Kiela. 2017. Automatically generating rhythmic verse with neural networks. In *Proceedings of ACL*.
- Xinyu Hua, Mitko Nikolov, Nikhil Badugu, and Lu Wang. 2019. Argument mining for understanding peer reviews. In *Proceedings of NAACL*.
- Dongyeop Kang, Waleed Ammar, Bhavana Dalvi, Madeleine van Zuylen, Sebastian Kohlmeier, Eduard Hovy, and Roy Schwartz. 2018. A dataset of peer reviews (peerread): Collection, insights and nlp applications. In *Proceedings of NAACL*.
- Nitish Shirish Keskar, Bryan McCann, Lav R Varshney, Caiming Xiong, and Richard Socher. 2019. Ctrl: A conditional transformer language model for controllable generation. *arXiv preprint arXiv:1909.05858*.

- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. In *Proceedings of NAACL-HLT*.
- 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 ACL*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*.
- Clara Meister, Tim Vieira, and Ryan Cotterell. 2020. Best-first beam search. *TACL*.
- Shashi Narayan, Gonçalo Simões, Yao Zhao, Joshua Maynez, Dipanjan Das, Michael Collins, and Mirella Lapata. 2022. A well-composed text is half done! composition sampling for diverse conditional generation. In *Proceedings of ACL*.
- Chenhui Shen, Liying Cheng, Ran Zhou, Lidong Bing, Yang You, and Luo Si. 2022. Mred: A meta-review dataset for structure-controllable text generation. *Findings of ACL*.
- Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng. 2019. The woman worked as a babysitter: On biases in language generation. In *Proceedings of EMNLP-IJCNLP*.
- Felix Stahlberg and Bill Byrne. 2019. On nmt search errors and model errors: Cat got your tongue? In *Proceedings of EMNLP-IJCNLP*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*.
- Yufei Wang, Ian Wood, Stephen Wan, Mark Dras, and Mark Johnson. 2021. Mention flags (mf): Constraining transformer-based text generators. In *Proceedings of ACL-IJCNLP*.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020a. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *Proceedings of ICML*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020b. [Bertscore: Evaluating text generation with bert](#). In *ICLR*.
- Xingxing Zhang and Mirella Lapata. 2014. Chinese poetry generation with recurrent neural networks. In *Proceedings of EMNLP*.

A Decoding Strategies

Maximization-based decoding methods such as beam search (Meister et al., 2020; Stahlberg and Byrne, 2019; Graves, 2012) and greedy search (similar to beam search with a beam size of 1) work well for optimizing the generation quality, whereas diversity-based methods such as sampling (Holtzman et al., 2019; Fan et al., 2018b; Ackley et al., 1985) focus on improving the generation diversity. To mitigate the trade-off between generation quality and diversity, Holtzman et al. (2019) propose nucleus sampling, where the generator samples only from a nucleus of tokens with cumulative probability larger than a specified $top-p$ value, and Caccia et al. (2019) propose beam sampling to integrate sampling with beam search.

B Evaluation during Generation

For evaluating the sequence likelihood, we obtain the log-likelihood of each generated token from the decoder, and then calculate the normalized sequence log-likelihood by averaging the log-likelihood for all tokens in the sequence. For evaluating the degree to which the generated sentence satisfies the control requirement, we obtain the category log-likelihood for each sentence from a trained classifier. The trained sentence classifier⁸ is based on the Roberta-Large architecture⁹. We fine-tune the model using the training set of labeled meta-reviews in the MReD dataset and obtain an accuracy of 0.853 on the test set.

Since both the sequence log-likelihood and classification log-likelihood are derived from probabilities between 0 - 1 and are of the same scale, we didn't experiment with different weightings but simply add them together. Nevertheless, it would be very straightforward to implement additional weightings to the evaluation process.

C The MReD dataset

Data from the peer review domain has come increasingly popular for research (Hua et al., 2019; Kang et al., 2018; Cheng et al., 2020a; Bhatia et al., 2020; Bao et al., 2021; Cheng et al., 2021; Bao et al., 2022). The MReD dataset consists of 7,089 fully annotated meta-reviews on ICLR 2018-2021

⁸We cannot directly apply the LSTM-CRF tagging model provided with MReD by Shen et al. (2022), which requires the full passage in order to output a sequence of labels

⁹<https://huggingface.co/roberta-large>

papers with their corresponding reviews. It is collected from the OpenReview¹⁰ portal.

We use the provided filtered version of MReD dataset¹¹ with the *Sent-Ctrl* and the *Seg-Ctrl* settings for all our experiments. Specifically, this version contains 5,354 examples for training, 665 examples for validation, and 674 examples for testing.

The detailed definitions for the categories are as follows:

- abstract: A piece of summary about the contents of the submission.
- strength: Opinions about the submission's strengths.
- weakness: Opinions about the submission's weaknesses.
- rating summary: A summary about reviewers' rating scores or decisions
- ac disagreement: Area chair (AC) shares different opinions to reviewers.
- rebuttal process: Contents related to authors' rebuttal with respect to reviews or discussions between reviewers in the rebuttal period.
- suggestion: Concrete suggestions for improving the submission.
- decision: Final decision (i.e., accept or reject) on the submission.
- misc: None of the above, such as courtesy expressions.

D Evaluation of Structural Controllability

For evaluation of structure similarity, Shen et al. (2022) has manually labeled the sentence category sequence for randomly selected test outputs as the prediction structures. Treating each category label as a single token, they calculate a normalized edit distance and use 1 to subtract this value for the final similarity score. Similarly, we use the same method and calculate structure similarity as:

$$1 - \frac{\text{edit_distance}(c, r)}{\max[\text{len}(c), \text{len}(r)]}$$

where r stands for the gold reference, c stands for the candidate prediction, and $\text{len}(c)$ stands for the total number of labels in the candidate.

¹⁰<https://openreview.net/>

¹¹https://github.com/Shen-Chenhui/MReD/tree/master/summarization/abstractive/filtered_controlled_data

	Structure \uparrow	Edits \downarrow	BERTScore \uparrow	R ₁ /R ₂ /R _L \uparrow
STS	0.683	2064	0.8589	37.17/10.53/22.38
ITSP	0.841	719	0.8624	37.95/10.62/23.30

Table 4: Results for STS and ITSP.

In addition, Shen et al. (2022) has trained a LSTM-CRF(Lample et al., 2016) classifier on the training set’s target meta-review passages to predict the sentence categories. A high classification accuracy of 0.8583 is reported on the test split, making it possible to evaluate the output structure automatically using this trained classifier. To do so, we first use the nltk sentence tokenizer to break down the outputs into sentences, then use the above LSTM-CRF sequential tagging model to predict the sentence label sequence.

Nevertheless, the normalized edit distance cannot quantitatively show how frequently the generations disobey the control signals. Therefore, we count the total edit distance for the whole test set and name this metric “total edits”.

E Experimental Settings

We reproduce the results of *Sent-Ctrl* and *Seg-Ctrl* by training MREd on the state-of-the-art “bart-

¹²<https://huggingface.co/facebook/bart-large-cnn>

large-cnn” model¹². Following Shen et al. (2022), we set source truncation to 2048, learning rate to $5e-5$, and use the Adam optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$ and no warm-up nor weight decay. The decoding method used is beam search with a beam size of 4. The model is trained for 3 epochs on a single Tesla V100 GPU.

F Choice of k and n

Inspired by beam search, we explore k up to 8 with $n = 4$ for our SentBS since it is similar to beam search to a certain degree. The Huggingface BART model uses a default beam size of 4 for beam search decoding, and up to 2 times the beam size of tokens are considered during the generation process.

G Human Evaluation

For fluency and relevance evaluations, the judges are provided with the summaries from the two models in a random order, and asked to give a score of 1 to their preferred generation and 0 to the other. They may give both summaries a score of 1 if the two are equally good.

For structure similarity and decision correctness evaluations, we follow Shen et al. (2022) and ask the judges to predict the control structure and decision from the given summary, which are then evaluated against the gold for the final results.

Control Sequence	Generated Passage
abstract strength strength	<p>Sent-Ctrl: [This paper proposes a new inductive bias to neural network architecture - namely a divide and conquer strategy know from algorithmics.]←ABSTRACT [The reviewers agree that the paper is well written and easy to follow.]←STRENGTH [The experimental section should be extended, as currently the only convincing success story lies in convex hull experiments, and the generalization scores themselves are pretty bad thus it is not clear if this can be called a success story.]←WEAKNESS</p> <p>SentBS: [This paper proposes a new inductive bias to neural network architecture - namely a divide and conquer strategy know from algorithmics.]←ABSTRACT [The paper is well written and clear.]←STRENGTH [The proposed approach is novel and interesting.]←STRENGTH</p>
abstract strength weakness weakness decision	<p>Sent-Ctrl: [This paper proposes a method for learning branching strategies within branch-and-bound solvers. <sep>]←ABSTRACT [Strengths: <sep> - The proposed method is novel and interesting, and the paper is well-written <sep>]←STRENGTH [Weaknesses: (1) the experimental setup itself is questionable.]←WEAKNESS [The authors consider two setups for the SCIP solver, "clean", with depth-first-search node selection and all other functions disabled (whatever that means), and "default".]←WEAKNESS [Some experiments are conducted under the "clean" setup (Table 1, Figure 2 and 3) while other experiments are made under "default" setup.]←WEAKNESS [(2) the presented experimental results are inconsistent with the literature, which I find suspicious the overall method description comprises several blind spots and fallacious arguments.]←WEAKNESS [(3) the authors claim their method results in non-myopic policies, without defining what they mean by that.]←WEAKNESS [I agree with the reviewers and recommend rejection of the paper, and encourage the authors to revise the paper based on the reviewers' feedback and resubmit to a different venue.]←DECISION</p> <p>SentBS: [The paper proposes a novel framework for learning branching strategies within branch-and-bound solvers based on a simplified graph-convolutional network (GCN) combined with a novelty-search evolutionary strategy (NS-ES) for training, and a new representation of B&B trees for computing novelty scores.]←ABSTRACT [The idea is interesting.]←STRENGTH [However, the experimental setup is questionable and the presented experimental results are inconsistent with the literature.]←WEAKNESS [In particular, the reviewers found several arguments to be fallacious.]←WEAKNESS [The paper is mostly based on the work from Gasse et al., reuses the same code, (almost) the same problem benchmarks, and some portions of text which are identical.]←WEAKNESS [Overall, this is a good paper, but not yet ready for publication.]←DECISION</p>

Table 5: Examples of outputs from *Sent-Ctrl*. We use red to highlight locations where generation deviate from the control labels, and cyan otherwise. Note that a <sep> is generated to represent a newline.