

Creative Natural Language Generation

Tuhin Chakrabarty¹ Vishakh Padmakumar² He He² Nanyun Peng³

¹Columbia University

²New York University

³ University of California, Los Angeles

tuhin.chakr@cs.columbia.edu, vp1271@nyu.edu, hhe@cs.nyu.edu, vnpeng@ucla.edu

1 Introduction

Large language models such as GPT-3 (Brown et al., 2020), BART, (Lewis et al., 2019) etc., have advanced the state of the art in several natural language generation tasks such as text summarization (Zhang et al., 2020) and machine translation (Liu et al., 2020). However when it comes to open-ended tasks with a focus on creativity such as generating stories (Fan et al., 2018a), poetry (Ghazvininejad et al., 2016), or various forms of figurative language (Chakrabarty et al., 2021), these state-of-the-art language models are often found to be inadequate.

The principal reason for this is that, in addition to composing grammatical and fluent sentences to articulate the intended content, these tasks usually also require extensive world and common-sense knowledge, as well as discourse modeling, to make sure the outputs maintain long-term coherence while remaining creative. It should also be noted that current approaches to text generation for specialized tasks require lots of training data for supervision. However, most existing corpora for creative forms of text are limited in size. Even if such a corpus existed for creative tasks, learning the distribution of existing data and sampling from it will unlikely lead to truly novel, creative output. Creative composition requires *deviating* from the norm, whereas standard generation approaches seek to mimic the norm.

This tutorial aims to bring awareness of the important and emerging research area of open-domain creative generation, with a focus on language generation while also touching on multi-modal generation (e.g., image captioning, visual metaphors, and visual story generation). It targets natural language processing (NLP) and artificial intelligence (AI) researchers as well as creative writing practitioners who are interested in building systems that are capable of emulating as well as augmenting human

creativity.

In particular, we will review recent studies on creative language generation both at the sentence level as well as longer forms of text. We will provide the audiences with a holistic view of 1) the importance and challenges of building creative language generation systems; 2) methods for different forms of creative language generation such as story (Yang et al., 2022; Yao et al., 2019), poetry (Tian and Peng, 2022), humor (He et al., 2019; Mittal et al., 2022), metaphors (Chakrabarty et al., 2021; Stowe et al., 2021; Chakrabarty et al., 2020b), sarcasm (Chakrabarty et al., 2020a), and hyperbole (Tian et al., 2021) 3) how can models for creativity infer user intention and preferences, allow for fine-grained control, and take (natural language) feedback? In particular, how could the recent advancement of AI shape the future workforce for creativity? We will conclude the tutorial by outlining future research directions in this area.

2 Tutorial Outline

In this tutorial, we will review the history of creative language generation both in shorter and longer forms. Then, we will move to the recent advances in creative language generation that employ transformer-based language models as well as external world knowledge from existing resources. We will also touch upon how much creativity can we elicit from larger models like GPT3 (Brown et al., 2020) and where they are still lacking. Finally, we will discuss the real-world implications of creative language generation and how humans can interact or collaborate with these models to satisfy their specific needs. In particular, we will present recent community efforts in the following topics:

1. Reviewing the history of creative language generation and how neural methods have shown considerable improvements over prior

approaches.

2. Introductions to contemporary methods for creative text generation along three main axes content planning for long-form creative text generation, figurative language generation with commonsense knowledge and the surprisal or twist factor which we term the XFactor in creative NLG
3. Discussion on how large-pretrained language models such as GPT-3 can perform creative language generation tasks and what are some of its benefits and where we can still have targeted improvements.
4. Introduction to the challenges in evaluating creative text. What are the possible dangers of relying on crowd workers from Amazon Mechanical Turk (Karpinska et al., 2021; Clark et al., 2021)? What are the tradeoffs of using expert vs crowd worker evaluation of creativity in language generation (Chakrabarty et al., 2023a)?
5. Examining how advances in creative NLG have opened up directions of research in the co-creative domain. How do amateur and skilled writers benefit from these models? How do these models fit into existing creative writing workflows? And how does this technology need to improve to become more impactful and useful to end users?
6. Lessons learned open challenges, and discussion about how to build robust, reliable, and useful systems for creative language.

3 History

Due to the lack of vast research on creative language generation and its importance in training and testing generative models, it is necessary to have a cutting-edge tutorial on an emerging and timely topic. We are unaware of any tutorials on the exact same topic in the past 4 years' ACL/EMNLP/EACL/NAACL conferences, with the only exception of the ACL 2020 tutorial (Mou and Vechtomova, 2020) on Stylized Text Generation: Approaches and Applications. The tutorial was mainly about style transfer. While there are some overlaps between style transfer and creative language generation, we believe our tutorial will benefit the audiences in terms of learning the vast

landscape of creative language generation in the age of pre-trained language models. Finally, our tutorial will also touch upon human-AI collaboration for creativity as well as creativity for vision and language tasks which has not been touched upon in prior tutorials.

4 Prerequisite Knowledge

Our target audience is general NLP conference attendances; therefore, no specific knowledge is assumed of the audience except basic machine learning and NLP background:

- Familiar with common natural language processing concepts (e.g., word representation, syntax, semantics) as found in an introductory NLP course.
- Familiar with the problems/setups of (open-domain) generation and creative forms of text such as story, poetry, metaphors etc
- Has basic knowledge about machine learning models such as deep neural networks, classifiers, and pre-trained models such as BERT (Devlin et al., 2019), DALLE (Ramesh et al., 2021, 2022), GPT2 (Radford et al., 2019), GPT3 (Brown et al., 2020) BART (Lewis et al., 2020).

5 Tutorial Content

This tutorial presents a systematic overview of the history and the frontier of creative language generation. We will also introduce methods for sentence level and longer forms of creative language generation, and careful consideration in designing the evaluation of model outputs as well as how LLM's can aid in providing assistance during the process of creative writing. We will then do deep dives. The detailed contents are outlined below.

5.1 Motivation, History, and Challenges [20 mins]

We will first motivate the importance of the problem by looking into works from psychology that examine what it means to be creative (Torrance, 1966) and then demonstrating practical applications of models that can produce creative outputs. Then, we will outline the challenges of building and evaluating creative generation models and systems. We will also include a brief introduction to the history of creative language generation and how

many of the challenges encountered by the community when developing contemporary language models share parallels with those faced by researchers working on these problems prior to the advent of statistical and neural techniques in NLP.

5.2 Recent Methods for Creative Generation [75 mins]

We detail various contemporary methods for creative text generation along three main axes charting progress in each. [VP: This could use a bit more punch]

Content Planning - “Austen’s Plots” [30 min]

In this section, we will discuss how approaches to control the content of the generated text by sketching a plan (Yao et al., 2019) has enabled pre-trained language models to generate higher quality stories with coherent plot lines (Goldfarb-Tarrant et al., 2020; Rashkin et al., 2020) as well as poetry with form constraints like sonnets (Tian and Peng, 2022). We then discuss the recent phase shift to adapting this style of content planning to large language models such as GPT3 to generate even longer, yet coherent, stories (over 1000 words) via recursive prompting (Yang et al., 2022, 2023).

Figurative Language Generation with Commonsense Knowledge - “The Bard’s Metaphors” [30 minutes]

Pre-trained language models typically excel at understanding the literal meaning of the text and generating responses accordingly. However, when it comes to creative tasks, they often struggle to effectively employ figurative language, which is essential for adding depth and nuance to the text. We will discuss how incorporating commonsense knowledge from external sources (Bosse-lut et al., 2019) enables models to better generate similes and metaphors (Chakrabarty et al., 2020b; Stowe et al., 2021; Chakrabarty et al., 2021) and sarcasm (Chakrabarty et al., 2020a). Finally, we examine how chain-of-thought prompting can elicit better figurative language understanding that was learned during the pre-training of large language models resulting in opportunities to generate higher quality illustrations for the same (Chakrabarty et al., 2022b).

The X-factor - “Dickens’ Twist” [15 minutes]

Finally, there is the ineffable quality of creative writing which grips the reader to keep turning the page. While this element is most challenging to recreate from language models, we discuss works

that attempt to do so by learning word-level relationships to generate puns (He et al., 2019) and break down intangible qualities such as humor into their basic principles for modeling (Tian et al., 2022).

5.3 Challenges in Evaluation of Creative NLG outputs [20 mins]

As the community makes progress in improving the various elements of the creative generation process, benchmarking progress becomes more challenging. One of the common practices in evaluating creative output is relying on crowd worker judgments from platforms such as Amazon Mechanical Turk. However, there are multiple challenges in these evaluations (Karpinska et al., 2021; Clark et al., 2021) such as crowd-workers spending limited time on reading and evaluating outputs, under-specified instructions for evaluation, variability in judgments across the same set of workers across different times, constructing proper qualification tests, setting up proper wages for crowd workers. A more promising alternative is to look into how experts might be better suited to evaluating outputs from creative NLG systems because their expectations might differ from amateur crowd workers (Clark and Smith, 2021). We discuss how recent work has delved back into the fundamentals of creativity to design evaluation axes based on the Torrance tests of creative thinking (Torrance, 1966) and measure these using expert judgments (Chakrabarty et al., 2023a).

5.4 Human AI Collaboration for Creativity [30 mins]

Recent developments in natural language generation (NLG) using large language models have brought us closer than ever to the goal of building AI-powered creative writing tools. In this section, we will discuss the potential of NLG to have a significant impact in the creative writing domain—especially with respect to brainstorming, generation of story details, and research assistance (Chakrabarty et al., 2023b). We will focus on different interaction interfaces for AI-assisted creativity, the extent to which they understand user intent, and finally, whether the human-AI collaboration improves the final creative output. We will end this section with the positives as well as limitations of current models as identified by expert and professional writers.

5.5 Conclusion, Future Directions, and Discussion [25 min]

We will conclude the tutorial by discussing future directions to build impactful, reliable and useful systems for creative language generation.

6 Tutorial Coverage and Suggested Reading List

While the tutorial will include our own work (Yao et al., 2019; He et al., 2019; Mittal et al., 2022; Goldfarb-Tarrant et al., 2020; Chakrabarty et al., 2020b, 2021; Akoury et al., 2020; Stowe et al., 2021; Tian et al., 2021; Tian and Peng, 2022; Padmakumar and He, 2022; Chakrabarty et al., 2022a; Yang et al., 2022), we anticipate that roughly 40% of the tutorial content will be pulled from work by other researchers in NLP and machine learning communities include but not limited to (Ghazvininejad et al., 2016; Fan et al., 2018b, 2019; Van de Cruys, 2020; Riedl and Young, 2010; Lin and Riedl, 2021; Brahman and Chaturvedi, 2020; Mirowski et al., 2023; Clark et al., 2021). A more comprehensive list of related papers will be provided before the tutorial.

7 Tutorial Instructors

Our instructors consist of experts who have conducted research in different aspects related to the tutorial topic.

Nanyun (Violet) Peng Nanyun (Violet) Peng is an Assistant Professor in the Department of Computer Science at the University of California Los Angeles. She received her Ph.D. in Computer Science from Johns Hopkins University. Her research focuses on the generalizability of NLP technologies, with applications to creative language generation, low-resource information extraction, and zero-shot cross-lingual transfer. Her works have won the Outstanding Paper Award at NAACL 2022, the Best Paper Award at AACL 2022 Deep Learning on Graphs workshop, and have been featured an IJCAI 2022 early career spotlight. She has given a tutorial at NAACL 2018 on information extraction.

Tuhin Chakrabarty Tuhin Chakrabarty is a Ph.D. candidate in Computer Science at Columbia University and a part of the Natural Language Processing group, where he is advised by Smaranda Muresan. His research is supported by the Columbia Center of Artificial Intelligence & Technology (CAIT) and Amazon Science Ph.D. Fellow-

ship. He was also a fellow at The New York Times R&D team working on Natural Language Generation. His overarching research question centers around how we can use large language models for creativity. He has published several papers in various NLP conferences and journals including ACL, NAACL, TACL and EMNLP.

He He He He is an Assistant Professor of Computer Science and the Center for Data Science at New York University. She is affiliated with the CILVR Lab, the Machine Learning for Language Group, and the Alignment Research Group. Her research focuses on building intelligent systems that can communicate with humans effectively and enable individuals to achieve their goals. Today's systems are often opaque, brittle, and difficult to control, which limits their usefulness in human-centered applications. To make them our trustworthy collaborators, her research aims to (i) understand the computational foundation of generalization in novel scenarios, and (ii) build interactive systems that align with users' goals. She has given a tutorial at EMNLP 2021 on robustness and adversarial examples in NLP.

Vishakh Padmakumar Vishakh Padmakumar is a Ph.D. student in Data Science at New York University advised by He He. His research is broadly in the field of natural language processing and human-AI collaboration with a focus on collaborative text generation for creative writing tasks and other interactive settings. Prior to this, he was a Graduate Research Associate at the NYU Center for Social Media and Politics working on political stance classification and multimodal content sharing in online disinformation campaigns. He has published papers at several NLP and machine learning venues including ACL, EMNLP, and ICML and was the chair of the ACL 2023 Student Research Workshop.

References

- Nader Akoury, Shufan Wang, Josh Whiting, Stephen Hood, Nanyun Peng, and Mohit Iyyer. 2020. Storium: A dataset and evaluation platform for machine-in-the-loop story generation. *arXiv preprint arXiv:2010.01717*.
- Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. 2019. **COMET: Commonsense transformers for automatic knowledge graph construction**. In *Proceedings of the 57th Annual Meeting of the Association for*

- Computational Linguistics*, pages 4762–4779, Florence, Italy. Association for Computational Linguistics.
- Faeze Brahman and Snigdha Chaturvedi. 2020. Modeling protagonist emotions for emotion-aware storytelling. *arXiv preprint arXiv:2010.06822*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Tuhin Chakrabarty, Debanjan Ghosh, Smaranda Muresan, and Nanyun Peng. 2020a. R³: Reverse, retrieve, and rank for sarcasm generation with commonsense knowledge. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7976–7986, Online. Association for Computational Linguistics.
- Tuhin Chakrabarty, Philippe Laban, Divyansh Agarwal, Smaranda Muresan, and Chien-Sheng Wu. 2023a. Art or artifice? large language models and the false promise of creativity. *arXiv preprint arXiv:2309.14556*.
- Tuhin Chakrabarty, Smaranda Muresan, and Nanyun Peng. 2020b. Generating similes effortlessly like a pro: A style transfer approach for simile generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6455–6469, Online. Association for Computational Linguistics.
- Tuhin Chakrabarty, Vishakh Padmakumar, Faeze Brahman, and Smaranda Muresan. 2023b. Creativity support in the age of large language models: An empirical study involving emerging writers.
- Tuhin Chakrabarty, Vishakh Padmakumar, and He He. 2022a. Help me write a poem: Instruction tuning as a vehicle for collaborative poetry writing. *arXiv preprint arXiv:2210.13669*.
- Tuhin Chakrabarty, Arkady Saakyan, Olivia Winn, Artemis Panagopoulou, Yue Yang, Marianna Apidianaki, and Smaranda Muresan. 2022b. I spy a metaphor: Large language models and diffusion models co-create visual metaphors. Preprint under review.
- Tuhin Chakrabarty, Xurui Zhang, Smaranda Muresan, and Nanyun Peng. 2021. MERMAID: Metaphor generation with symbolism and discriminative decoding. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4250–4261, Online. Association for Computational Linguistics.
- Elizabeth Clark, Tal August, Sofia Serrano, Nikita Haduong, Suchin Gururangan, and Noah A Smith. 2021. All that’s human’s not gold: Evaluating human evaluation of generated text. *arXiv preprint arXiv:2107.00061*.
- Elizabeth Clark and Noah A Smith. 2021. Choose your own adventure: Paired suggestions in collaborative writing for evaluating story generation models. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3566–3575.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018a. Hierarchical neural story generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898, Melbourne, Australia. Association for Computational Linguistics.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018b. Hierarchical neural story generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2019. Strategies for structuring story generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2650–2660, Florence, Italy. Association for Computational Linguistics.
- Marjan Ghazvininejad, Xing Shi, Yejin Choi, and Kevin Knight. 2016. Generating topical poetry. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*.
- Seraphina Goldfarb-Tarrant, Tuhin Chakrabarty, Ralph Weischedel, and Nanyun Peng. 2020. Content planning for neural story generation with aristotelian rescoring. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4319–4338.
- He He, Nanyun Peng, and Percy Liang. 2019. Pun generation with surprise. 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 1734–1744, Minneapolis, Minnesota. Association for Computational Linguistics.
- Marzena Karpinska, Nader Akoury, and Mohit Iyyer. 2021. The perils of using mechanical turk to evaluate open-ended text generation. In *Proceedings of the*

- 2021 *Conference on Empirical Methods in Natural Language Processing*, pages 1265–1285.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- 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.
- Zhiyu Lin and Mark O. Riedl. 2021. Plug-and-blend: A framework for controllable story generation with blended control codes. In *Proceedings of the 2021 AAAI Conference on AI and Interactive Digital Entertainment*.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. [Multilingual denoising pre-training for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Piotr Mirowski, Kory W. Mathewson, Jaylen Pittman, and Richard Evans. 2023. [Co-writing screenplays and theatre scripts with language models: Evaluation by industry professionals](#). In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, New York, NY, USA. Association for Computing Machinery.
- Anirudh Mittal, Yufei Tian, and Nanyun Peng. 2022. [AmbiPun: Generating humorous puns with ambiguous context](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1053–1062, Seattle, United States. Association for Computational Linguistics.
- Lili Mou and Olga Vechtomova. 2020. [Stylized text generation: Approaches and applications](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts*, pages 19–22, Online. Association for Computational Linguistics.
- Vishakh Padmakumar and He He. 2022. [Machine-in-the-loop rewriting for creative image captioning](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 573–586, Seattle, United States. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR.
- Hannah Rashkin, Asli Celikyilmaz, Yejin Choi, and Jianfeng Gao. 2020. [PlotMachines: Outline-conditioned generation with dynamic plot state tracking](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4274–4295, Online. Association for Computational Linguistics.
- Mark O Riedl and Robert Michael Young. 2010. Narrative planning: Balancing plot and character. *Journal of Artificial Intelligence Research*, 39:217–268.
- Kevin Stowe, Tuhin Chakrabarty, Nanyun Peng, Smaranda Muresan, and Iryna Gurevych. 2021. [Metaphor generation with conceptual mappings](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6724–6736, Online. Association for Computational Linguistics.
- Yufei Tian and Nanyun Peng. 2022. [Zero-shot sonnet generation with discourse-level planning and aesthetics features](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3587–3597, Seattle, United States. Association for Computational Linguistics.
- Yufei Tian, Divyanshu Sheth, and Nanyun Peng. 2022. [A unified framework for pun generation with humor principles](#).
- Yufei Tian, Arvind krishna Sridhar, and Nanyun Peng. 2021. [HypoGen: Hyperbole generation with commonsense and counterfactual knowledge](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1583–1593, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- E Paul Torrance. 1966. Torrance tests of creative thinking. *Educational and Psychological Measurement*.
- Tim Van de Cruys. 2020. [Automatic poetry generation from prosaic text](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2471–2480, Online. Association for Computational Linguistics.

- Kevin Yang, Dan Klein, Nanyun Peng, and Yuandong Tian. 2023. [DOC: Improving long story coherence with detailed outline control](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3378–3465, Toronto, Canada. Association for Computational Linguistics.
- Kevin Yang, Yuandong Tian, Nanyun Peng, and Dan Klein. 2022. [Re3: Generating longer stories with recursive reprompting and revision](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Lili Yao, Nanyun Peng, Ralph Weischedel, Kevin Knight, Dongyan Zhao, and Rui Yan. 2019. [Plan-and-write: Towards better automatic storytelling](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7378–7385.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. [Pegasus: Pre-training with extracted gap-sentences for abstractive summarization](#). In *International Conference on Machine Learning*, pages 11328–11339. PMLR.