

Modulating Language Models with Emotions

Ruibao Liu¹ Jason Wei² Chenyan Jia³ Soroush Vosoughi¹
¹Dartmouth College ²ProtagoLabs ³University of Texas at Austin
ruibo.liu.gr@dartmouth.edu jason@protagolabs.com
chenyanjia@utexas.edu soroush@dartmouth.edu

Abstract

Generating context-aware language that embodies diverse emotions is an important step towards building empathetic NLP systems. In this paper, we propose a formulation of modulated layer normalization—a technique inspired by computer vision—that allows us to use large-scale language models for emotional response generation. In automatic and human evaluation on the MojiTalk dataset, our proposed modulated layer normalization method outperforms prior baseline methods while maintaining diversity, fluency, and coherence. Our method also obtains competitive performance even when using only 10% of the available training data.

1 Introduction

Building interactive systems that can understand and express human emotions has been a long-term goal of artificial intelligence (Shen and Feng, 2020; Huang et al., 2018; Salovey and Sluyter, 1997). Given a context, an intelligent agent ought to be able to generate responses that not only consider the context but also reflect a specified emotion, a task called *emotional response generation*. One common representation of emotions is through *emojis*, which often convey the underlying emotions in an utterance (Zhou and Wang, 2018). Table 1 shows an example generation in this formulation.

To tackle this problem, prior work has proposed a number of different models, including variants of sequence-to-sequence (Seq2Seq) models (Serban et al., 2016; Li et al., 2016a), variational autoencoders (VAE) (Gu et al., 2019; Shen et al., 2017; Zhao et al., 2017) and adversarial networks (Kong et al., 2019; Li et al., 2017). Their generated responses are often dull or generic, partially due to the limited training data for diverse emotions (Li et al., 2017). More recent studies have tried to

Context: good game start morning off tigers v eagles.

Emotion	Response
😊	good luck to all the eagles
👏	i m not a tigers fan but we ve got a win
💔	we ve got to wait for tommorrow for the game
👉	hope you enjoyed the match with your team

Table 1. Example generation of our method for four different emojis. Context is an actual random tweet, and emotion is specified by emojis.

pre-train language models (LMs) on specific domain data to pivot generation towards certain direction (Gao et al., 2020; Zhang et al., 2020; Yang et al., 2020; Keskar et al., 2019). However, training a LM from scratch can be costly, and collecting sufficient pre-training data in diverse emotions is also challenging, especially for low-resource emotions (Yang et al., 2019a).

In this work, we present a simple and easy-to-deploy technique that can enable pre-trained large-scale LMs to generate fine-grained emotional responses. Specifically, we inject emotional signals specified by 64 commonly used emojis via Modulated Layer Normalization (Mod-LN), a technique widely adopted in computer vision but whose potential has not been well studied yet in NLP. The main advantages of our method are:

- Instead of designing or re-training models from scratch, our method is plug-and-play. In this work, we show its effectiveness on BERT (2019) and GPT-2 (2019), but one can easily extend our method to other Transformer-based LMs.
- By fully exploiting the transfer learning ability of pre-trained LMs, we achieve comparable emotional response generation performance as prior best-performing work with only 10% of the training data, which is especially beneficial for low-resource scenarios.

2 Approach

Given a context text and a specified emoji as a target emotion, we aim to generate responses that both reflect the emotion associated with the emoji and the semantic information in the context. In this work, we demonstrate how to inject target emotions through a modulation module of layer normalization (§2.1). We also provide data preparation and model adaptation strategies on two typical LMs (BERT and GPT-2) to aid reproduction and extension (§2.2).

2.1 Modulated Layer Normalization

Layerwise-normalization (LN) is commonly used in Transformer-based (Vaswani et al., 2017) language models (LMs) (Devlin et al., 2019; Radford et al., 2019; Yang et al., 2019b) to stabilize hidden state dynamics and reduce training time (Ba et al., 2016). In the vanilla implementation (Figure 1(a)), data are normalized by their own mean μ and standard deviation σ without relying on external inputs.




Figure 1: Overview of (a) Vanilla Layer Normalization (LN) and (b) Modulated Layer Normalization (Mod-LN) in Transformer-based LMs. The modulation module in Mod-LN uses two Multi-Layer Perceptrons (MLPs) that each have two sets of dense layers. It uses the external emotion input to modulate regularization towards a certain emotion c . FFN: Feed-Forward Network. Self-Attn: Multi-head Self-Attention blocks.

In contrast to vanilla LN that only regularizes data itself, Mod-LN introduces an external modulation module shared across the whole dataset, which is independent of the individual data samples and able to modulate the regularization towards external inputs c (Figure 1 (b)). Specifically, for an input hidden state tensor x in layer l , it is normalized by Mod-LN as

$$x = \text{MLP}_\gamma^{(l)}(c) \cdot \frac{x - \mu}{\sigma + \epsilon} + \text{MLP}_\beta^{(l)}, \quad (1)$$

where ϵ is the smoothing parameter to avoid dividing by zero. $\text{MLP}_\gamma^{(l)}$ and $\text{MLP}_\beta^{(l)}$ are two trainable modulation modules for a certain layer l . They are computed by

$$\text{MLP}_\gamma^{(l)}(c) = W_\gamma^{(l,2)} \cdot \text{Swish}(W_\gamma^{(l,1)}c), \quad (2)$$

$$\text{MLP}_\beta^{(l)}(c) = W_\beta^{(l,2)} \cdot \text{Swish}(W_\beta^{(l,1)}c + b), \quad (3)$$

where $W^{(l,1)}$ and $W^{(l,2)}$ are dense layers belonging to layer l , with weights size of $[64, \frac{1}{2} \cdot \text{dim}_h]$ and $[\frac{1}{2} \cdot \text{dim}_h, \text{dim}_h]$ respectively¹. Dense layers connect 64 emoji classes to the output hidden states of the language model, and b is a bias added to γ . We use the Swish activation (Ramachandran et al., 2017), which has been shown to outperform ReLU (Xu et al., 2015) on several challenging datasets. Though conceptually simple, such MLP based modules have been shown to be a faster and more efficient alternative to vanilla dot product self-attention in NLP (Tay et al., 2021) and CV (Tolstikhin et al., 2021). Our work uses MLPs as a plug-and-play modulator rather than a replacement for self-attentions, allowing us to shift the hidden states towards a given target emotion.

2.2 Data Preparation and Model Adaptation

For the text input, we concatenate ground-truth context with corresponding response as a whole input to feed into LMs. We add a pre-defined separator token ([SEP] for BERT and [UNK] for GPT-2) between context and response, to make LMs aware of the range of each part. We also pad both context and response to a max sequence length with the padding token.

Encoder-Decoder models have been successful in many text-to-text generation tasks, such as question answering (Chen et al., 2017; Seo et al., 2017), news summarization (Chopra et al., 2016; Rush et al., 2015), and style transfer (Li et al., 2018; Liu et al., 2021). For the response generation task, the encoder encodes the context text into a fixed-length vector in latent space, while the decoder decodes the generated response tokens step-by-step, given the encoded context vector and the ground truth token from the previous step; this method is also known as teacher-forcing (Zhang et al., 2019c; Cho et al., 2014).

In this work, we consider leveraging the transfer learning power of large-scale LMs—using LMs

¹The hidden size dim_h of bert-large-uncased and GPT-2 medium model are both 1024.

as encoder and decoder—to better capture the complicated relationship between context and response (Rothe et al., 2020). Auto-regressive LMs (ARLMs), such as GPT-2 are trained to iteratively predict the next step token given the past, while Masked Language Models (MLM), such as BERT, are trained to predict missing tokens given both the preceding and subsequent text. In contrast to the uni-directional attention flow in ARLM, the attention flow of MLM is bi-directional, and thus if we directly use MLM as decoder, the prediction of tokens in the response will also attend to (i.e., have the context of) future tokens; this could potentially lead to exposure bias (Schmidt, 2019). Inspired by recent text-to-text LMs such as T5 (Raffel et al., 2020) and BART (Lewis et al., 2019), for MLM decoder, we modify the original bi-directional attention mask to make it uni-directional.

We experiment with two encoder-decoder models built on MLM and ARLM: 1) BERT-to-BERT: using bi-directional BERT as both encoder and decoder, but forcing the decoder BERT to attend to past context with uni-directional mask, and 2) GPT2-to-GPT2: using uni-directional GPT-2 as both encoder and decoder.

3 Experimental Setup

Dataset. For all the experiments, we use the MojiTalk (Zhou and Wang, 2018) dataset, a large Twitter conversation corpus ($N \approx 700k$) of responses that each contain one or more of 64 popular emojis. Following the original paper, we split the corpus into training, validation, and test sets of 596,959, 32,600, and 32,600 conversation pairs, respectively. We fine-tune the two LM-based encoder-decoder models on this dataset and generate responses given contexts and all possible emotions using top- k random decoding (Fan et al., 2018) on a machine with four RTX 2080 GPUs².

Models. We evaluate three models in total. We take the Reinforced Conditional Variational AutoEncoders (R-CVAE) model from Zhou and Wang (2018) as **Baseline** (current best-performing model on 64-emoji controlled response generation), **Mod-LN MLM**: BERT-to-BERT (large, uncased) + Mod-LN, and **Mod-LN ARLM**: GPT2-to-GPT2 (large) + Mod-LN.

²We choose $k = 10$ for a balance of generation diversity and readability through empirical observation.

Model	Emoji Acc (%)		
	Hits@1	Hits@3	Hits@5
Baseline: R-CVAE			
w/. 10% train data	13.4	27.1	33.6
w/. 100% train data	26.2	44.2	53.4
Mod-LN MLM			
w/. 10% train data	20.5	47.4	59.1
w/. 100% train data	33.6	56.8	72.2
Mod-LN ARLM			
w/. 10% train data	27.9	43.4	64.1
w/. 100% train data	34.4	60.3	82.5

Table 1: Accuracy of emotional response judged by DeepMoji on classifying emotions in responses generated by R-CVAE, MLM (BERT) with Mod-LN, and ARLM (GPT-2) with Mod-LN.

4 Evaluation

Good emotional responses should accurately reflect the intended emotion, be diverse, and have coherent language. We thus evaluate three aspects of generated responses: emotion control (§4.1), response diversity (§4.2), and coherence and fluency (§4.3). We also use Amazon Mechanical Turk (MTurk) to run a manual evaluation of emotion control and readability in generated responses (§4.4).

4.1 Emotion Control

First, we evaluate whether intended emotions were reflected in the responses generated by various models. We choose DeepMoji (Felbo et al., 2017)³ as the judgment classifier. DeepMoji was trained on a large-scale emoji dataset containing 1,246 million tweets and 64 distinct emojis, and as far as we know, is state-of-the-art for 64-emoji classification tasks. Since the meanings of different emojis can overlap with subtle differences, we compute Hits@ k ($k = \{1, 3, 5\}$) classification accuracy (Gao et al., 2020) to describe the performance of models in different criteria. As shown in Table 1, our proposed models outperform R-CVAE with a large margin. Of note, LM-based models reveal more robust performance in extreme data scarcity cases: our models achieve comparable performance with R-CVAE even when using only 10% of the training data. Between BERT and GPT-2, GPT-2 shows superior performance, partially because its weights are from auto-regressive pre-training.

³We chose the official implementation by huggingface: <https://github.com/huggingface/torchMoji>.

Model	TTR-1	TTR-2	Avg. len	%stop
Human Reference	0.059	0.43	11.7	50.4
Baseline: R-CVAE				
w/. 10% train data	0.034	0.24	8.6	60.1
w/. 100% train data	0.051	0.33	9.2	59.3
Mod-LN MLM				
w/. 10% train data	0.054	<u>0.43</u>	18.2	49.3
w/. 100% train data	<u>0.059</u>	0.39	14.3	49.2
Mod-LN ARLM				
w/. 10% train data	0.056	0.38	15.9	48.7
w/. 100% train data	0.057	0.40	12.5	48.5

Table 2: Lexical diversity of generated responses from various models. TTR-1/TTR-2: unigram/bigram type-token ratio; Avg. len: average number of tokens in generated responses; %stop: average percent of stop words among all tokens in the generated responses.

4.2 Generation Diversity

As shown in Table 2, we evaluate the diversity of responses generated by each model in terms of unigram and bigram type-token ratios, average length, and percent of stop words in generated responses, with values for the human-generated responses shown for reference. As measured by the type-token ratio for both uni- and bi-grams, our proposed models generate more diverse responses. In addition, compared with the R-CVAE, the responses generated by our models are longer and use fewer stop words. The advance can be attributed to the using of large-scale language models as base models.

4.3 Fluency and Coherence

Moreover, we evaluate the fluency and coherence of machine-generated text. For fluency, we trained a standalone language model on the human-generated responses using KenLM (Heafield, 2011) to measure the perplexity of generated texts. To evaluate coherence between the context and the generated responses, we compute the similarity between the generated text and human-generated responses using BERTScore (Zhang et al., 2019b), with the human-generated responses as reference. We configure the BERTScore using 24-layer RoBERTa-large (Liu et al., 2019) as for English tasks. Table 3 shows these results. For perplexity and BERTScore, our Mod-LN models outperform the R-CVAE in both 10% and 100% training data cases.

Model (vs Ref.)	PPL	BERTScore (%)		
		Precision	Recall	F1
Baseline: R-CVAE				
w/. 10% train data	121.18	74.9	83.0	76.7
w/. 100% train data	92.64	80.8	80.8	80.8
Mod-LN MLM				
w/. 10% train data	79.24	78.4	80.1	78.8
w/. 100% train data	50.72	82.9	84.1	83.5
Mod-LN ARLM				
w/. 10% train data	51.55	83.7	80.7	83.2
w/. 100% train data	<u>36.31</u>	84.7	86.2	<u>85.4</u>

Table 3: Fluency as measured by perplexity (PPL) and coherence as measured by BERTScore of generated responses from various models. Ref.: Human-generated responses.

4.4 Human Evaluation

In total 120 MTurk participants manually evaluated the emotion control and readability of responses from our proposed models and the original human-generated reference data. The average age of participants was 38.40 years-old (SD = 12.26, Median=34.50). More than half (65.8%) of participants were male, and 34.2% were female. The average completion time of each survey was 4.53 minutes. Participants were paid \$1 per survey, averaging to more than \$13 per hour wage for each participant, significantly above the U.S. federal minimum wage.

Procedure Each participant was assigned to read five randomly selected context-response pairs without being informed of the sources of the responses. They were asked to rate 1) emotion control: “How well the emotion conveyed in the response agrees with the specified emoji? (1-very well to 7-not at all)”, and 2) readability: “Please rate the readability of the response on a 7-point scale. (1-very low to 7-very high)”. The readability measure included five items adapted from a previous study (Graefe et al., 2018), specifically, well-written, concise, comprehensive, coherent, and clear. Since the five measures had very high agreement (Cronbach’s⁴ $\alpha = .91$), we average the five measures into one as a general readability index.

Results The participant’s averaged ratings (μ) and Standard Errors (SE) are reported in Table 4.

⁴Cronbach’s alpha is a measure of internal consistency between sets of items.

Response Source	Annotator Ratings: μ (SE)	
	EMO CTRL.	READABILITY
Human Reference	5.62 (0.10)	5.34 (0.10)
Baseline: R-CVAE	4.80 (0.10)	4.67 (0.10)
Mod-LN MLM	5.43 (0.10)	5.20 (0.10)
Ablation: Vanilla GPT-2	4.98 (0.10)	4.64 (0.10)
Mod-LN ARLM	5.40 (0.10)	5.32 (0.10)

Table 4: Humans manually evaluated the emotional control and readability of responses from the original data (human reference), Baseline and proposed models on a 7-point scale (1: low quality, 7: high quality). We also take the generative LM: vanilla GPT-2, as the ablation reference.

As shown in the table, the standard error of the mean among all annotators is .10, which is very low for a 7-point scale, indicating large agreement between annotators. Responses generated by Mod-LN MLM (BERT), Mod-LN ARLM (GPT-2), and the human-generated references had no statistically significant differences in emotion control and readability. All were rated significantly higher than plain GPT-2 and R-CVAE in both emotion control and readability ($p < .001$ for one-way repeated measures ANOVA). We also conducted pairwise multiple comparisons in our analysis as post-hoc analysis. In terms of emotion control, both of our two proposed models and original reference data were rated significantly better than vanilla GPT-2 ($p < .007$). For readability, both our models, vanilla GPT-2, and original reference data were rated significantly more readable than R-CVAE ($p < .001$).

5 Related Work

Emotional Text Generation. VAE-based models (Park et al., 2018; Shen et al., 2017; Zhao et al., 2017; Serban et al., 2017), adversarial networks (Kong et al., 2019; Li et al., 2017; Yu et al., 2017) and reinforcement learning systems (Li et al., 2019, 2016b) have dominated sentiment-aware dialogue models. Other methods have been developed using LSTM (Song et al., 2019) and GRU (Wei et al., 2019; Zhou et al., 2018). All these methods, however, are built on relatively coarse emotion types, partially due to the limited modeling ability of RNNs. Our model outperforms current state-of-the-art R-CVAE (Zhou and Wang, 2018) in the same 64-emoji settings.

Modulated Normalization. Though not common in NLP, modulated normalization has been previously used in computer vision. In addition to work mentioned in the introduction (De Vries et al., 2017), adversarial networks such as CGAN (Miyato and Koyama, 2018), self-attention GAN (Zhang et al., 2019a) and Style GAN (Karras et al., 2019) have used modulated normalization to inject external signal into their models. In NLP, previous studies have tried to modulate normalization for classification tasks (Houlsby et al., 2019) and multilingual machine translation (Bapna and Firat, 2019), however, both these methods require architecture-level modifications. Our method, on the other hand, is plug-and-play, requiring minimal modifications to the architecture and thus easier to deploy for a diverse set of applications.

6 Conclusions

We have proposed a modulated layer normalization approach to generating responses of varying specified emotions. Our approach allows us to leverage large pre-trained models, while remaining simple and easily-extendable. In empirical experiments, our approach substantially outperforms prior work and achieves comparable results using only 10% of the available training data, all while maintaining diversity, fluency, and coherence.

References

- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. [Layer normalization](#). In *Advances in Neural Information Processing Systems, Deep Learning Symposium*.
- Ankur Bapna and Orhan Firat. 2019. [Simple, scalable adaptation for neural machine translation](#). 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 1538–1548, Hong Kong, China. Association for Computational Linguistics.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. [Reading wikipedia to answer open-domain questions](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. [Learning phrase representations using RNN encoder–decoder for statistical machine translation](#). In *Proceedings of*

- the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 1724–1734, Doha, Qatar. Association for Computational Linguistics.
- Sumit Chopra, Michael Auli, and Alexander M. Rush. 2016. [Abstractive sentence summarization with attentive recurrent neural networks](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 93–98, San Diego, California. Association for Computational Linguistics.
- Harm De Vries, Florian Strub, Jérémie Mary, Hugo Larochelle, Olivier Pietquin, and Aaron C Courville. 2017. [Modulating early visual processing by language](#). In *Advances in Neural Information Processing Systems*, pages 6594–6604.
- 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.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018. [Hierarchical neural story generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898.
- Bjarke Felbo, Alan Mislove, Anders Søgaard, Iyad Rahwan, and Sune Lehmann. 2017. [Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1615–1625, Copenhagen, Denmark. Association for Computational Linguistics.
- Xiang Gao, Yizhe Zhang, Michel Galley, Chris Brockett, and Bill Dolan. 2020. [Dialogue response ranking training with large-scale human feedback data](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 386–395, Online. Association for Computational Linguistics.
- Andreas Graefe, Mario Haim, Bastian Haarmann, and Hans-Bernd Brosius. 2018. Readers’ perception of computer-generated news: Credibility, expertise, and readability. *Journalism*, 19(5):595–610.
- Xiaodong Gu, Kyunghyun Cho, Jung Woo Ha, and Sunghun Kim. 2019. [Dialogwae: Multimodal response generation with conditional wasserstein auto-encoder](#). In *7th International Conference on Learning Representations, ICLR 2019*.
- Kenneth Heafield. 2011. [Kenlm: Faster and smaller language model queries](#). In *Proceedings of the sixth workshop on statistical machine translation*, pages 187–197. Association for Computational Linguistics.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for NLP](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799. PMLR.
- Chenyang Huang, Osmar R Zaiane, Amine Trabelsi, and Nouha Dziri. 2018. [Automatic dialogue generation with expressed emotions](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 49–54.
- Tero Karras, Samuli Laine, and Timo Aila. 2019. [A style-based generator architecture for generative adversarial networks](#). In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4401–4410.
- Nitish Shirish Keskar, Bryan McCann, Lav Varshney, Caiming Xiong, and Richard Socher. 2019. CTRL - A Conditional Transformer Language Model for Controllable Generation. *arXiv preprint arXiv:1909.05858*.
- Xiang Kong, Bohan Li, Graham Neubig, Eduard Hovy, and Yiming Yang. 2019. [An adversarial approach to high-quality, sentiment-controlled neural dialogue generation](#). *arXiv preprint arXiv:1901.07129*.
- 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*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016a. [A diversity-promoting objective function for neural conversation models](#). In *Proceedings of NAACL-HLT*, pages 110–119.
- Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. 2016b. [Deep reinforcement learning for dialogue generation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202, Austin, Texas. Association for Computational Linguistics.
- Jiwei Li, Will Monroe, Tianlin Shi, Sébastien Jean, Alan Ritter, and Dan Jurafsky. 2017. [Adversarial learning for neural dialogue generation](#). pages 2157–2169.
- Juncen Li, Robin Jia, He He, and Percy Liang. 2018. [Delete, retrieve, generate: a simple approach to sentiment and style transfer](#). 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 1865–1874, New Orleans, Louisiana. Association for Computational Linguistics.
- Ziming Li, Julia Kiseleva, and Maarten de Rijke. 2019. [Dialogue generation: From imitation learning to inverse reinforcement learning](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6722–6729.
- Ruibo Liu, Chenyan Jia, and Soroush Vosoughi. 2021. A transformer-based framework for neutralizing and reversing the political polarity of news articles. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1):1–26.
- 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*.
- Takeru Miyato and Masanori Koyama. 2018. [cgans with projection discriminator](#). In *International Conference on Learning Representations*.
- Yookoon Park, Jaemin Cho, and Gunhee Kim. 2018. [A hierarchical latent structure for variational conversation modeling](#). 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 1792–1801.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#). *OpenAI Blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21:1–67.
- Prajit Ramachandran, Barret Zoph, and Quoc V Le. 2017. [Searching for activation functions](#). *arXiv preprint arXiv:1710.05941*.
- Sascha Rothe, Shashi Narayan, and Aliaksei Severyn. 2020. [Leveraging pre-trained checkpoints for sequence generation tasks](#). *Transactions of the Association for Computational Linguistics*, 8:264–280.
- Alexander M Rush, Sumit Chopra, and Jason Weston. 2015. [A neural attention model for abstractive sentence summarization](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 379–389.
- Peter Ed Salovey and David J Sluyter. 1997. *Emotional development and emotional intelligence: Educational implications*. Basic Books.
- Florian Schmidt. 2019. [Generalization in generation: A closer look at exposure bias](#). In *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pages 157–167, Hong Kong. Association for Computational Linguistics.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. [Bidirectional attention flow for machine comprehension](#). In *5th International Conference on Learning Representations, ICLR 2017*.
- Iulian V Serban, Alessandro Sordoni, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2016. [Building end-to-end dialogue systems using generative hierarchical neural network models](#).
- Iulian Vlad Serban, Alessandro Sordoni, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017. [A hierarchical latent variable encoder-decoder model for generating dialogues](#). In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Lei Shen and Yang Feng. 2020. [CDL: Curriculum dual learning for emotion-controllable response generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 556–566, Online. Association for Computational Linguistics.
- Xiaoyu Shen, Hui Su, Yanran Li, Wenjie Li, Shuzi Niu, Yang Zhao, Akiko Aizawa, and Guoping Long. 2017. [A conditional variational framework for dialog generation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 504–509, Vancouver, Canada. Association for Computational Linguistics.
- Zhenqiao Song, Xiaoqing Zheng, Lu Liu, Mu Xu, and Xuanjing Huang. 2019. [Generating responses with a specific emotion in dialog](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3685–3695, Florence, Italy. Association for Computational Linguistics.
- Yi Tay, Dara Bahri, Donald Metzler, Da-Cheng Juan, Zhe Zhao, and Che Zheng. 2021. [Synthesizer: Rethinking self-attention in transformer models](#). In *Proceedings of the 37th International Conference on Machine Learning (ICML 2021)*.
- Ilya Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, et al. 2021. [Mlp-mixer: An all-mlp architecture for vision](#).
- 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 I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.

- Wei Wei, Jiayi Liu, Xianling Mao, Guibing Guo, Feida Zhu, Pan Zhou, and Yuchong Hu. 2019. [Emotion-aware chat machine: Automatic emotional response generation for human-like emotional interaction](#). In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pages 1401–1410.
- Bing Xu, Naiyan Wang, Tianqi Chen, and Mu Li. 2015. [Empirical evaluation of rectified activations in convolutional network](#). *arXiv preprint arXiv:1505.00853*.
- Ze Yang, Wei Wu, Can Xu, Xinnian Liang, Jiaqi Bai, Liran Wang, Wei Wang, and Zhoujun Li. 2020. [StyleDGPT: Stylized response generation with pre-trained language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1548–1559, Online. Association for Computational Linguistics.
- Ze Yang, Wei Wu, Jian Yang, Can Xu, and Zhoujun Li. 2019a. [Low-resource response generation with template prior](#). 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 1886–1897, Hong Kong, China. Association for Computational Linguistics.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019b. [Xlnet: Generalized autoregressive pretraining for language understanding](#). In *Advances in neural information processing systems*, pages 5754–5764.
- Lantao Yu, Weinan Zhang, Jun Wang, and Yong Yu. 2017. [Seqgan: Sequence generative adversarial nets with policy gradient](#). In *Thirty-first AAAI conference on artificial intelligence*.
- Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. 2019a. [Self-attention generative adversarial networks](#). In *International Conference on Machine Learning*, pages 7354–7363.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2019b. [Bertscore: Evaluating text generation with BERT](#). *CoRR*, abs/1904.09675.
- Wen Zhang, Yang Feng, Fandong Meng, Di You, and Qun Liu. 2019c. [Bridging the gap between training and inference for neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4334–4343, Florence, Italy. Association for Computational Linguistics.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. [DIALOGPT : Large-scale generative pre-training for conversational response generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278, Online. Association for Computational Linguistics.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. [Learning discourse-level diversity for neural dialog models using conditional variational autoencoders](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 654–664, Vancouver, Canada. Association for Computational Linguistics.
- Hao Zhou, Minlie Huang, Tianyang Zhang, Xiaoyan Zhu, and Bing Liu. 2018. [Emotional chatting machine: Emotional conversation generation with internal and external memory](#). In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Xianda Zhou and William Yang Wang. 2018. [Mojitalk: Generating emotional responses at scale](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1128–1137.