

# Emotion Conditioned Creative Dialog Generation

**Khalid Alnajjar**

University of Helsinki  
Finland

`Khalid.alnajjar@helsinki.fi`

**Mika Hämäläinen**

University of Helsinki  
Finland

`mika.hamalainen@helsinki.fi`

## Abstract

We present a DialGPT based model for generating creative dialog responses that are conditioned based on one of the following emotions: *anger*, *disgust*, *fear*, *happiness*, *pain*, *sadness* and *surprise*. Our model is capable of producing a contextually apt response given an input sentence and a desired emotion label. Our model is capable of expressing the desired emotion with an accuracy of 0.6. The best performing emotions are *neutral*, *fear* and *disgust*. When measuring the strength of the expressed emotion, we find that *anger*, *fear* and *disgust* are expressed in the most strong fashion by the model.

## 1 Introduction

Dialog systems and different kinds of natural language interfaces are all around us. When we seek information or contact customer support, we are increasingly more often first greeted by a bot rather than a person. Bots are typically stiff and not life-like making communication with them an awkward experience. This is because their goal oriented nature typically sets some constraints in terms of how creative a system can be.

We can perceive a gap between dialog systems that are designed to convey a certain message in a goal oriented fashion and dialog systems that generate chit-chat. Chit-chat can be generated rather freely because any topical response is valid, where as a certain degree of factual correctness is to be expected from a goal oriented system.

In this paper, we seek to bring the two lines of dialog generation research closer together. We implement a system that can generate topical responses (in the sense of chit-chat) with a fixed emotional content. Thus the goal of the system is to convey a desired emotion in its response, no matter what the actual textual content ended up being. This means that part of the semantics of the out is fixed, while

a part as to how to contextually adapt the emotional content. is still up to computational creativity.

We base our experiments on a recently published dialog dataset that contains sentiment annotations<sup>1</sup>. The dataset is based on a video game called Fallout New Vegas and it has dialog where each line of the dialog is annotated as containing one of the following emotions: *anger*, *disgust*, *fear*, *happiness*, *neutral*, *pain*, *sadness* or *surprise*. This dataset makes for an optimal training data for the task we seek to solve.

## 2 Related Work

In terms of computational creativity and natural language generation, there are several papers out there that present work conducted on a variety of different creative language generation tasks such as poem generation (Hegade et al., 2021; Hämäläinen et al., 2022), humor generation (Weller et al., 2020; Alnajjar and Hämäläinen, 2021), news generation (Shu et al., 2021; Koppatz et al., 2022) and story generation (Vicente et al., 2018; Concepción et al., 2019). In this section, we will take a closer look at the work conducted on dialog generation.

Xie and Pu (2021) present work on generating empathetic dialog. Their model deals with the following categories of empathetic intent: *questioning*, *agreeing*, *acknowledging*, *sympathizing*, *encouraging*, *consoling*, *suggesting* and *wishing*. They base their model on the transformer architecture and they use RoBERTa for input encoding. In addition, they train a classifier that predicts the salient empathetic intent.

Dialog generation has also been tackled in a context-controlled and topic-aware manner (Ling et al., 2021). Their model consists of four parts: a hierarchical context encoder, a context dependent topic representation module, a context guided topic

<sup>1</sup><https://zenodo.org/record/6990638>

	Anger	Disgust	Fear	Happy	Neutral	Pained	Sad	Surprised
Sentences	3 335	932	1 620	4 029	8 802	994	1 055	1 649

Table 1: Number of sentences per emotion

Prompt	Response	Emotion
I hear you’ve been causing trouble.	Oh yeah? Fuck off, asshole.	Anger
I was hoping you’d be that stupid.	What? Hey, guys! Help me here!	Surprise
I’ve dealt with those newcomers.	I take care of those who help with that. Here, you earned it	Happiness

Table 2: Examples of the training data

transition module and a joint attention based response decoder.

Chen et al. (2021) present their work on multi-turn dialog generation. They use a cross-hierarchical encoder that encodes a sentence for an answer selector model, after this a response generator model is used to generate the final output. The initial encoding is done by a transformer based model while the final generation is done by an LSTM model.

Dialog adaptation has been studied before in the context of video games (Hämäläinen and Alnajjar, 2019). The authors use an LSTM model to paraphrase the syntax of existing dialog to introduce diversity and a word2vec model to adapt the meaning of the sentence towards a desired player attribute.

### 3 Creativity and Emotion

There are several takes on creativity in a computational setting. In this section, we cover some of these theoretical ways of understanding computational creativity. Theoretical foundation has been, for a long time, at the very core of computational creativity to combat systems that do mere generation. That is generation for the sake of outputting something by any means necessary.

A computationally creative system should exhibit skill, imagination and appreciation according to Colton (2008). He argues that all of these three components are a strict requirement for creativity to exist in a system. Skill refers to the system’s capability of producing a creative artifact whereas appreciation means that the system should also know why its creation is good. Imagination requires the system to be capable of garnering a lot of diverse output for one input.

Creativity can also be modeled through the FACE theory (Colton et al., 2011). This theory

states creativity comes from the interplay between framing, aesthetics, concept and expression. Expressions are the creative output produced by the system. The system itself is called concept. Aesthetics is similar to appreciation in the previous theory; it means that the system should be able to appreciate the creative value of its output. Framing highlights the fact that creativity does not take place in a vacuum but is presented in a context. In our case, framing would be the entire dialog between a human user and the machine.

Boden (1998) identifies three types of creativity; exploratory creativity, transformational creativity and combinatory creativity. In combinatory creativity, a system forms new artifacts by combining old ones in novel ways. In exploratory creativity, a system is conducting a search in a conceptual space discovering new creative artifacts. A system that can achieve transformational creativity can change its search space.

A system is considered to be autonomously creative if it can change its own standards without being explicitly told to do so Jennings (2010). The change cannot occur at random either because a simple random change at random intervals would otherwise be enough to satisfy the criterion.

Emotion is considered as a higher level cognitive phenomenon than a feeling (see Shouse 2005). Feelings are seen to be universally felt in a similar fashion as a response to some external or internal stimulus. Emotions, on the other hand, are culturally and socially represented and their existence in all cultures in a similar way is not a given thing (see Lim 2016).

One way of seeing emotions is that they rely on affect (see Russell 2003), which is a state our mind is continuously in. An affect can move across two axes: positive-negative and arousal-relaxed. The affect we feel is contextually resolved to a higher level emotion based on the context we find

ourselves in. For instance, a high arousal and negativity could be interpreted either as anger or disgust among others.

Ekman (1992) has identified six basic emotions: *anger, disgust, fear, happiness, sadness* and *surprise*. These emotions are also present in our dataset, which makes this theory optimal to build upon. The basic emotions are considered to be universal across cultures based on studies conducted on facial expressions.

## 4 Data

Our dataset consists of dialog in English where each line is annotated with an emotion label. Table 1 shows the data size and how many sentences were there in the dataset for each emotion. The dialog is from a video game called *Fallout New Vegas*<sup>2</sup> by Obsidian Entertainment. The game is set in a post-apocalyptic world inhabited by creatures that were born as a result of nuclear radiation such as mutants, ghouls and feral ghouls in addition to humans.

The game dialog consists mostly of player prompts that result in a response by an NPC (non-player character). In Table 2 we can see some examples of the dialog. In our case, we train our model by using the prompt and the emotion label as an input to predict the response.

*Fallout New Vegas* is a relatively large video game because it is an open world RPG (role-playing game) where the player can roam freely from one place to another. As a result of this, the game has a variety of different scripted characters which means that there is no bias in terms of having the same characters speaking with each other all the time. There is, however, a bias in the topic of conversations given that they take place in a fictional world. Many of the dialogs deal with fictional places, characters, items and so on.

## 5 Dialog Generation

In this section, we outline our approach to emotionally conditioned dialog generation. We conduct our experiment using Transformers (Wolf et al., 2020) and Datasets (Lhoest et al., 2021) Python libraries. We base our model on a pretrained model called DialoGPT medium<sup>3</sup>.

DialoGPT (Zhang et al., 2020) is based on the GPT-2 (Radford et al., 2019) architecture, which

in turn is based on the generic transformer language model (Vaswani et al., 2017). The transformer model leverages a stack of masked multi-head self-attention layers to train on large datasets. DialoGPT employs a maximum mutual information scoring function (Li et al., 2015; Zhang et al., 2018) during the training phase optimizing the reward with a policy gradient (Williams, 1992) with a sample averaged baseline (Zhang et al., 2018).

We changed the format of the data to be as “EMOTION<sub>1</sub>: SENTENCE<sub>1</sub>. EMOTION<sub>2</sub>: SENTENCE<sub>2</sub> [EOS]“, where EMOTION<sub>*n*</sub> indicates the emotion of the *n*th sentence in the conversation. This way, the model is exposed to emotional knowledge regarding the preceding locution of the response, as well. We use 90% of the data for training and 10% for validation. We train the model for 5 epochs.

## 6 Results and Evaluation

In order to evaluate our model, we sample dialog from an unrelated open-world RPG called *The Witcher 3: Wild Hunt*<sup>4</sup> by CD Projekt. The dialog is extracted using a w3strings decoder tool<sup>5</sup>. By using dialog from a different video game, we can see whether our model works in a different domain. The *Witcher 3* is set to a medieval fantasy world inhabited by magical creatures. Some of the generated output can be seen in Table 3.

We sample randomly 15 sentences for each emotion, and produce emotion conditioned replies. In the input, “EMOTION<sub>1</sub>:“ was not present as the dialogs in *The Witcher* did not include emotion labels. These are then evaluated by crowd-workers on a crowd-sourcing platform called Appen<sup>6</sup>. For every output we evaluate, we show the input sentence and the output to the user and ask them if the output expresses the desired emotion as a simple yes/no question. In case, the user selected “yes”, we also asked how strongly the line expressed the desired emotion on a 5-point scale (the stronger, the higher), the scale is 0-based from 0 to 4. Each generated output was evaluated by 5 different judges.

All in all, the judges agreement with the emotion label was **0.6**. In Figure 1, we can see the judges’ agreement per emotion label. The worst performing emotion is *pained* and the best performing ones are *neutral, fear* and *disgust*. *Sad, surprise* and

<sup>2</sup><https://fallout.bethesda.net/en/games/fallout-new-vegas>

<sup>3</sup><https://huggingface.co/microsoft/DialoGPT-medium>

<sup>4</sup><https://www.thewitcher.com/en/witcher3>

<sup>5</sup><https://www.nexusmods.com/witcher3/mods/1055/>

<sup>6</sup><https://appen.com/>

Prompt	Emotion	Output
Doesn't look like the talkative type.	Anger	What the fuck are you looking at?
Wait a minute! I ain't done nothin' to you!	Fear	I don't want any trouble, okay?
Asked you a question. How many?	Sad	I'm sorry, but I don't have the time to talk right now.
Need coin, fast? Those blasted nonhumans at the Vivaldi Bank say your credit's no good? Come see me!	Surprise	I've never heard of that before.

Table 3: Examples of the results




Figure 1: Percentage of "yes" answers per emotion

*anger* are all rated expressing the emotion over half of the time. These results are based on the aggregated results provided by Appen.




Figure 2: Strength of each emotion

Figure 2 shows how strong the judges rated each emotion to be. 0 indicates not strong at all and 4 very strong. As we can see, all emotions achieve above mid-level performance. *Anger*, *fear* and *disgust* being the most strongly expressed emotions. All in all, based on the results, it seems that the model is capable of producing emotional responses in a suitably strong manner.

## 7 Conclusions

Emotion conditioned dialog generation remains a challenge. In this paper, we have shown the vi-

ability of using a DialoGPT model for this task. The results are promising especially given that the evaluation was run on out-of-the-domain data.

The video game that was used in this paper also comes with audio files for each line of dialog. This is an interesting direction for future research because we could generate emotion conditioned dialog with audio as well. The same sentence can be said with multiple different tones and intonations to express different emotions. Building a system that can express emotion in generated speech as well would have great application potentials in speech oriented dialog systems.

## Acknowledgments

This work was partially financed by the Society of Swedish Literature in Finland with funding from Enhancing Conversational AI with Computational Creativity, and by the Ella and Georg Ehrnrooth Foundation for Modelling Conversational Artificial Intelligence with Intent and Creativity.

## References

- Khalid Alnajjar and Mika Hämmäläinen. 2021. When a computer cracks a joke: Automated generation of humorous headlines. In *Proceedings of the 12th International Conference on Computational Creativity (ICCC 2021)*. Association for Computational Creativity.
- Margaret A Boden. 1998. Creativity and artificial intelligence. *Artificial intelligence*, 103(1-2):347–356.
- Xiuying Chen, Zhi Cui, Jiayi Zhang, Chen Wei, Jianwei Cui, Bin Wang, Dongyan Zhao, and Rui Yan. 2021. Reasoning in dialog: Improving response generation by context reading comprehension. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 12683–12691.
- Simon Colton. 2008. Creativity versus the perception of creativity in computational systems. In *AAAI spring symposium: creative intelligent systems*, volume 8, page 7. Palo Alto, CA.

- Simon Colton, John William Charnley, and Alison Pease. 2011. Computational creativity theory: The face and idea descriptive models. In *ICCC*, pages 90–95. Mexico City.
- Eugenio Concepción, Pablo Gervás, and Gonzalo Méndez. 2019. Evolving the ines story generation system: From single to multiple plot lines. In *ICCC*, pages 220–227.
- P Ekman. 1992. Are there basic emotions? *Psychological review*, 99(3):550–553.
- Mika Hämmäläinen and Khalid Alnajjar. 2019. Creative contextual dialog adaptation in an open world rpg. In *Proceedings of the 14th International Conference on the Foundations of Digital Games*, pages 1–7.
- Mika Hämmäläinen, Khalid Alnajjar, and Thierry Poibeau. 2022. Modern french poetry generation with roberta and gpt-2. In *13th International Conference on Computational Creativity (ICCC) 2022*.
- Prakash Hegade, Rajaram M Joshi, Vibha G Hegde, Tejaswini Kale, and Srushti Basavaraddi. 2021. Po-miner: A web mining poem generator and its security model. *SN Computer Science*, 2(5):1–14.
- Kyle E Jennings. 2010. Developing creativity: Artificial barriers in artificial intelligence. *Minds and Machines*, 20(4):489–501.
- Maximilian Koppatz, Khalid Alnajjar, Mika Hämmäläinen, and Thierry Poibeau. 2022. Automatic generation of factual news headlines in finnish. In *15th International Natural Language Generation Conference (INLG)*.
- Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Mario Šaško, Gunjan Chhablani, Bhavitvya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Clément Delangue, Théo Matussière, Lysandre Debut, Stas Bekman, Pierric Cistac, Thibault Goehringer, Victor Mustar, François Lagunas, Alexander Rush, and Thomas Wolf. 2021. [Datasets: A community library for natural language processing](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 175–184, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2015. A diversity-promoting objective function for neural conversation models. *arXiv preprint arXiv:1510.03055*.
- Nangyeon Lim. 2016. Cultural differences in emotion: differences in emotional arousal level between the east and the west. *Integrative medicine research*, 5(2):105–109.
- Yanxiang Ling, Fei Cai, Xuejun Hu, Jun Liu, Wanyu Chen, and Honghui Chen. 2021. Context-controlled topic-aware neural response generation for open-domain dialog systems. *Information Processing & Management*, 58(1):102392.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- James A Russell. 2003. Core affect and the psychological construction of emotion. *Psychological review*, 110(1):145.
- Eric Shouse. 2005. Feeling, emotion, affect. *M/c journal*, 8(6).
- Kai Shu, Yichuan Li, Kaize Ding, and Huan Liu. 2021. Fact-enhanced synthetic news generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13825–13833.
- 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*, 30.
- Marta Vicente, Cristina Barros, and Elena Lloret. 2018. Statistical language modelling for automatic story generation. *Journal of Intelligent & Fuzzy Systems*, 34(5):3069–3079.
- Orion Weller, Nancy Fulda, and Kevin Seppi. 2020. Can humor prediction datasets be used for humor generation? humorous headline generation via style transfer. In *Proceedings of the Second Workshop on Figurative Language Processing*, pages 186–191.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Yubo Xie and Pearl Pu. 2021. Empathetic dialog generation with fine-grained intents. In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 133–147.
- Yizhe Zhang, Michel Galley, Jianfeng Gao, Zhe Gan, Xiujun Li, Chris Brockett, and Bill Dolan. 2018. Generating informative and diverse conversational responses via adversarial information maximization.

*Advances in Neural Information Processing Systems*,  
31.

Yizhe Zhang, Siqu Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and William B 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.