

# Affective Prompt-Tuning-Based Language Model for Semantic-Based Emotional Text Generation

Zhaodong Gu, South China University of Technology, China

Kejing He, South China University of Technology, China\*

 <https://orcid.org/0000-0003-4116-037X>

## ABSTRACT

The large language models based on transformers have shown strong text generation ability. However, due to the need for significant computing resources, little work has been done to generate emotional text using language models such as GPT-2. To address this issue, the authors proposed an affective prompt-tuning-based language model (APT-LM) equipped with an affective decoding (AD) method, aiming to enhance emotional text generation with limited computing resources. In detail, the proposed model incorporates the emotional attributes into the soft prompt by using the NRC emotion intensity lexicon and updates the additional parameters while freezing the language model. Then, it steers the generation toward a given emotion by calculating the cosine distance between the affective soft prompt and the candidate tokens generated by the language model. Experimental results show that the proposed APT-LM model significantly improves emotional text generation and achieves competitive performance on sentence fluency compared to baseline models across automatic evaluation and human evaluation.

## KEYWORDS

Affective Decoding, Discrete Emotion, Emotional Text Generation, Language Model, Prompt-Tuning

## INTRODUCTION

Artificial intelligence (AI) has numerous applications in various fields, including cloud computing (Bisht & Vampugani, 2022; Ilyas et al., 2022), intelligent systems (Casillo et al., 2022; Deveci et al., 2023), digital transformation (Gupta et al., 2023; Li et al., 2023), text detection (Yen et al., 2021; Zhang et al., 2023), and more. However, AI generally lacks the ability to express human emotions. Emotional intelligence is an important branch of artificial intelligence, which has been widely studied and explored in the field of natural language processing (NLP) (Barbosa et al., 2022; Chopra et al., 2022; Ismail et al., 2022). Emotional text generation, in particular, holds great potential for a variety of applications. Research shows that systems that can express emotions significantly improve user satisfaction (Prendinger & Ishizuka, 2005; Abo-Hammour et al., 2013; Arqub & Abo-Hammour,

DOI: 10.4018/IJSWIS.339187

\*Corresponding Author

This article published as an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>) which permits unrestricted use, distribution, and production in any medium, provided the author of the original work and original publication source are properly credited.

2014). In the field of dialogue systems, some studies have improved the generated responses by endowing the dialogue system with increased empathy towards human users (Colombo et al., 2019; Abo-Hammour et al., 2014; Abu Arqub et al., 2012). The controlled emotional text generation model not only enables a more meaningful dialogue between AI agents and humans, but also aims to establish emotional connections with readers. This model proves beneficial for conversation therapy robots, as it can produce suitable emotional responses according to the user's psychological state (Sarivougioukas & Vagelatos, 2022).

Despite the diverse range of applications for emotional text generation, how to better integrate emotions into the generation model is still a difficult problem. Conventional methods for emotional text generation primarily rely on discourse templates and manual rules, frequently demonstrating limitations in addressing complex situations. Deep learning based methods have gained widespread usage in NLP tasks since the proposal of the recurrent neural network (RNN) Encoder-Decoder model by Cho et al. (2014). Vaswani et al. (2017) proposed the transformer architecture, which has significantly advanced NLP tasks by enabling large-scale language model training on massive datasets (Radford et al., 2019). Although deep learning-based methods have made significant progress in classification tasks (Chen et al., 2022), generating emotional text remains a challenging task, in contrast to the success achieved in sentiment classification (Yang et al., 2019). In particular, generation models based on transformer architecture can leverage a significant amount of unlabeled data for training. Nevertheless, this approach makes it challenging to change the attributes of the generated text without either modifying the model's architecture or utilizing specific attribute data for fine-tuning (Keskar et al., 2019; Zhang et al., 2018). Therefore, it becomes more difficult to incorporate emotional attributes into the pretrained model. However, most of the emotional text generation models today are still based on the Seq2Seq model of the RNN architecture, failing to take advantage of the latest pretrained transformer models such as GPT.

Therefore, we propose an affective prompt-tuning based language model (denoted as APT-LM) in this paper. Different from fine-tuning, affective prompt-tuning only needs to fine-tune the extra emotional attributes parameters freezing the whole model parameters. In detail, we first select discrete emotional attributes from the NRC Emotion Intensity Lexicon and map them to continuous soft prompts through embedding layers of autoregressive language models such as GPT2. Then, during the training process, we only update these custom parameters. To enhance the emotional expression of generated text, we further propose a simple and effective decoding method called affective decoding (denoted as AD) based on APT. This method can first obtain the possible candidate sets of the token generated by the model in each step of the decoding process. Then it calculates the cosine distance between the candidate sets and the soft prompt. Finally, we obtain the generated text that can maintain semantic coherence and enhance emotional expression. Experiment results show that our model significantly outperforms other models in automatic and human evaluation, not only in terms of accuracy in generating corresponding emotional texts, but also in terms of emotional richness of texts and can maintain strong competitiveness in other metrics such as perplexity, distinct-1-2-3, grammar, and fluency. The case study demonstrates the flexibility and scalability of APT-LM, allowing the model to incorporate other attributes into the model and generate specific attribute texts by adjusting custom parameters. The proposed model holds significant industrial relevance and offers a wide range of applications, seamlessly integrating with diverse intelligent systems. It proves valuable in emotion analysis across social media platforms, enabling businesses and organizations to assess public sentiment regarding their products, services, or campaigns. Moreover, its utility extends to virtual therapy applications, where it delivers emotionally sensitive responses to users seeking mental health support or counseling (Zhang et al., 2023). The model's impact also extends to entertainment and gaming, interactive educational content, healthcare chatbots, smart assistants imbued with emotional intelligence, and personalized recommendation systems.

The main contributions of this paper can be summarized as follows:

- (1) Proposing the affective prompt-tuning based language model (APT-LM) to generate emotional text using state-of-the-art language model under limited computational resources.
- (2) Designing the affective decoding (AD) method to further enhance the emotional expression of the generated text while maintaining sentence fluency.
- (3) Conducting experiments on two public datasets has demonstrated that our model outperforms other baseline models significantly in emotional expression, as observed through both automatic evaluation and human evaluation.

## RELATED WORKS

### Efficient Training Models

Fine-tuning has been the prevalent technique for using pretrained language models since the release of GPT (Radford et al., 2018) and BERT (Devlin et al., 2019). However, it requires updating and storing all the model's parameters. Given the vast number of parameters in pretrained models, researchers are beginning to study how to give full play to the ability of the large pretrained models under the condition of limited hardware resources. These studies generally tend to freeze most of the pretrained parameters and only perform gradient updates in the customized parameters or decoding process. These techniques can be broadly categorized into two groups: lightweight fine-tuning and plug-and-play decoding.

On lightweight fine-tuning, Rebuffi et al. (2017) propose the concept of an "adapter," which inserts task-specific layers with a small number of parameters between each layer of the pretrained model. Houslyby et al. (2019) freeze the pretrained BERT-large model and add only 2-4% extra parameters with the "adapter" to achieve close performance on GLUE compared with the fine-tuning method. Li and Liang (2021) propose "prefix-tuning" to train pretrained models for natural language generation tasks instead of fine-tuning. First, each transformer layer begins with the addition of a small, continuously updated task-specific vector known as the prefix. Then, the model only optimizes the prefix during training. Based on "prefix-tuning," Lester et al. (2021) propose "prompt-tuning," which still uses task-specific parameters but only adds and tunes them on the input layer. It is noted that "prompt-tuning" uses the soft prompt (continuous vector) instead of the hard prompt, which is used in GPT-3 (Brown et al., 2020) for in-context learning or prompting with some discrete tokens. In comparison, "adapter" emphasizes minimizing architectural modifications and achieving task-specific adaptation through a constrained set of parameters in additional adapter layers. On the other hand, "prefix-tuning" prioritizes updating a task-specific vector without altering the model's architecture, and "prompt-tuning" goes a step further by incorporating task-specific parameters exclusively on the input layer using a soft prompt. Unlike Lester et al. (2021) using a masked language model, we explore using an autoregressive language model (Radford et al., 2019) for prompt-tuning in this paper.

On plug-and-play decoding, Dathathri et al. (2019) propose a model called PPLM, which integrates the pretrained LM with several simple attribute discriminators and iteratively updates the past activation functions (hidden states) to steer the generation without training the LM. Pascual et al. (2021) propose a model-agnostic method called Keyword2Text (K2T). It works with any autoregressive language model and doesn't need any discriminators. Specifically, when considering a topic or keyword as a hard constraint, K2T introduces an adjustment to the probability distribution across the vocabulary, favoring semantically similar words. It allows the model to control text generation without training the LM. PPLM's strength lies in its iterative refinement process, enabling fine-tuned control, but it requires training additional attribute discriminators for controlling text generation. In contrast, K2T does not necessitate extra training and only updates the probability distribution across the vocabulary rather than modifying the hidden states of the model.

## Emotional Text Generation Models

Many methods have been introduced to model the relationships between distinct emotions. We can generally classify these methods into two categories: discrete and dimensional emotion models. In the dimensional or continuous emotion model, emotions are mapped to several dimensions (Russell, 2003), including valence, arousal, and dominance. Valence stands for pleasure, with negative as unpleasant and positive as pleasant. Arousal evaluates the strength of emotion. Dominance stands for the extent to which a participant is in control of the emotional state. Colombo et al. (2019) introduce an emotion-driven dialog system, designed to produce controlled emotional responses through the utilization of a continuous representation of emotions, employing a valence-arousal-dominance (VAD) lexicon (Mohammad, 2018). In this work, although we follow the discrete emotion models, emotions are projected to a continuous embedding layer instead of the VAD space.

In discrete emotion models, basic emotions can be classified into various categories (Ekman, 1992; Plutchik, 2001), such as joy, sadness, and anger. Although discrete emotion models have been criticized for their different numbers of basic emotions in different models (Russell, 1994), the majority of emotional text generation models are still based on the discrete emotion model. Zhou et al. (2018) propose an emotional chatbot model (ECM), which generates emotional texts by employing emotional category embeddings, internal emotional memory, and external emotional memory. Ghosh et al. (2017) propose an affect-LM model that generates emotional dialogue texts across four specific emotion categories, each characterized by varying influence strengths. ECM employs gated recurrent unit (GRU) and attention to implement various modules for controlling the content and emotions of generated text. However, its complex mechanism can make model training and interpretation relatively challenging. On the other hand, affect-LM utilizes long short-term memory (LSTM) to easily and efficiently control the emotional attributes and intensity of generated text, with the aid of an additional design parameter. However, the above methods need to fine-tune all parameters and do not use the state-of-the-art pretrained language model, while our proposed method freezes the model parameters and only needs to update a small number of custom parameters to obtain competitive emotional text generation capability.

In order to better demonstrate the limitations of the previous work in this paper, we have summarized them as shown in the Table 1.

**Table 1. Limitations of the Previous Work**

Model	Method	Limitation
Efficient Training Models	Fine-tuning (Radford et al., 2018)	Need to train all parameters from scratch.
	Adapter (Rebuffi et al., 2017)	Change the architecture of the model and need add additional adapter layer.
	Prefix-tuning (Li & Liang, 2021)	Need to add prefix to every single Transformer layer.
	Prompt-tuning (Lester et al., 2021)	No explore the emotional text generation and nor the use autoregressive models for generation.
	PPLM (Dathathri et al., 2019)	Require training additional attribute discriminators for controlling text generation.
	K2T (Pascual et al., 2021)	Cannot be applied in emotional text generation.
Emotional Text Generation Models	Colombo et al. (2019)	Do not follow the Ekman's emotion theory.
	ECM (Zhou et al., 2018)	Do not use the language model like GPT-2
	Affect-LM (Ghosh et al., 2017)	Do not use the language model like GPT-2

## Proposed Method

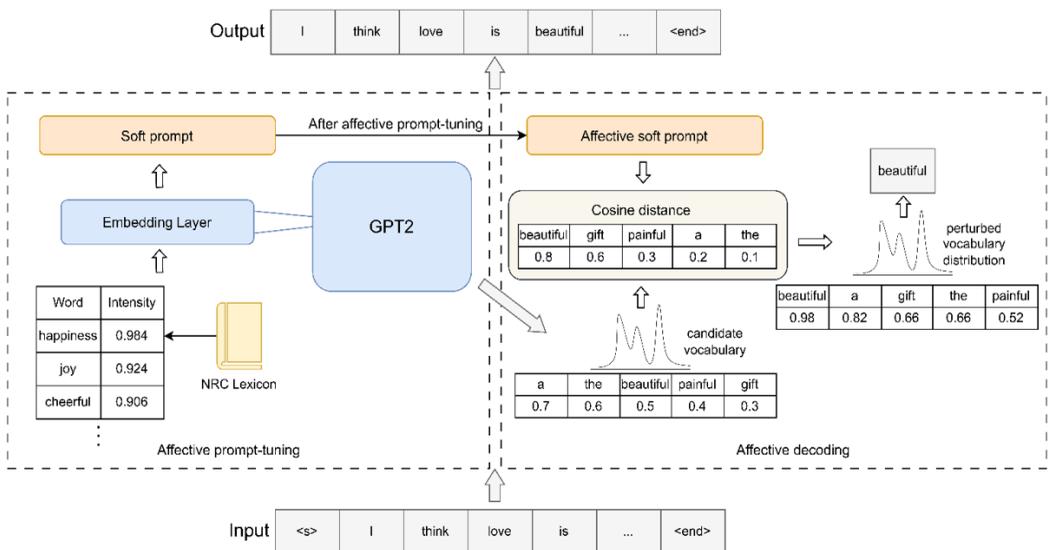
Our proposed model for emotional text generation based on prompt-tuning is illustrated in Figure 1, which mainly consists of two modules: affective prompt-tuning and affective decoding. The affective prompt-tuning first converts the discrete emotional tokens into a continuous space and integrates them into the soft prompt initialization process. Then, it updates the extra emotional attributes parameters through the soft prompt while freezing the whole model parameters during training. The affective decoding method enhances the emotional expression of the text by calculating the cosine distance between the candidate tokens and the soft prompt after affective prompt-tuning. With these modules, APT-LM can easily integrate with autoregressive language models, such as GPT-2, as it does not require changing the architecture of the model, only defining the emotion parameters to be trained through the embedding layer of the language model. APT-LM's architecture is designed to efficiently scale across diverse hardware configurations, making it adaptable to both smaller setups and more extensive computing environments. The model's flexibility in terms of fine-tuning and customization allows developers to tailor APT-LM to their specific use cases by designing different custom soft prompts. This adaptability enhances user experience by providing a more personalized and efficient solution that aligns with individual project requirements.

### Affective Prompt-Tuning

Assume we have an autoregressive language model based on transformer architecture, such as GPT. With a input sequence of tokens  $X = \{x_1, x_2, \dots, x_n\}$ , LM is trained to calculate the unconditional probability of the sequence  $P(X)$  and updates the model parameters  $\theta$ . The probability can be expressed as the product of conditional probabilities through the recursive application of the chain rule, represented as:

$$P_{\theta}(X) = \prod_i^n P_{\theta}(x_i | x_{<i}) \tag{1}$$

Figure 1. Overview of the Proposed Model for Emotional Text Generation



Usually, prompting is done by adding a sequence of tokens  $Pr$  in the front of the input sequence  $X$ . In GPT-3, the prompt can be represented by a sequence of tokens  $Pr = \{pr_1, pr_2, \dots, pr_n\}$ , which are all part of the model embedding table and are still affected by the model parameters  $\theta$ .

However, our method enables the prompt to have an independent emotional parameter  $\theta_{Pr}$  while freezing the parameter  $\theta$  of the original model. Then, our new condition generation model can be formulated as follows:

$$P_{\theta_{Pr}}(X) = \prod_i^n P_{\theta_{Pr}}(x_i | [Pr; x_{<i}]) \quad (2)$$

In the autoregressive transformer model, it is assumed that the activation at time step  $i$  is  $h_i = [h_i^{(1)}, \dots, h_i^{(n)}]$ , where  $h_i^{(j)}$  is the activation of the  $j$ -th transformer layer at time step  $i$ . Then, the model computes  $h_{<i}$  as a function of  $x_i$  and the past activations in its left context, as follows:

$$h_i = LM(x_i, h_{<i}) \quad (3)$$

where  $h_{<i}$  is generally used to calculate the distribution of the next token. The distribution can be formulated as  $P_{\theta_{Pr}}(x_i | h_{<i}) = \text{softmax}(W_{\theta_{Pr}} \cdot h_i^{(n)})$ . Therefore, we can transform discrete tokens into continuous word embedding space, and its effect will propagate from bottom to top to all transformer activation layers. Eventually, it affects the generated text from left to right.

We propose an efficient and simple method of initializing soft prompts with emotional attributes. Specifically, we apply the NRC Emotion Intensity Lexicon (Mohammad & Kiritchenko, 2018) for selecting words to initialize the affective soft prompts based on the descending emotional intensity of each word. This is because NRC Emotion Intensity Lexicon provides emotion classification and corresponding emotional intensity for each word, covering the six basic emotions of Ekman's emotion theory. In this paper, we focus on generating emotional text based on Ekman's emotion theory. As shown in Figure 1, we use the words with the emotion of joy in the NRC lexicon as an example. First, we map the discrete sequence of tokens  $Pr$  to the continuous embedding space through the embedding layer of the LM, and get the soft prompt  $Pr_d \in \mathbb{R}^{pr \times d}$ , where  $pr$  is the length of the prompt and  $d$  is the dimension of the embedding space. Second, we concatenate the soft prompt to the embedding of the input sequence  $X$ . Finally, we only update the parameters  $\theta_{Pr}$  when training the LM to generate output sequence  $Y$ , as follows:

$$\max \log P_{\theta_{Pr}}(Y | X) = \max \sum_i^n \log P_{\theta_{Pr}}(x_i | h_{<i}) \quad (4)$$

It is noted that APT-LM utilizes the generalization and knowledge obtained from pre-training language model for emotional text generation. If traditional fine-tuning methods are used to control LM for emotional text generation, a large amount of emotion annotated data needs to be used to update all parameters of LM. On one hand, the datasets are difficult to obtain due to labor costs, and on the other hand, training LM from scratch requires a large amount of computing resources. APT-LM can utilize the emotional information provided by the NRC lexicon and the embedding layer of LM to generate custom training parameters. During the training process, LM is frozen, thus maintaining the

advantage of generating text using LM. Updating the custom parameters ensures that LM is guided to generate emotional text.

## Affective Decoding

After prompt-tuning the LM with emotional attributes, the generation can be guided by an affective soft prompt. To enhance the emotional expression of the generated text, we propose affective decoding (AD). For each decoding step, AD aims to: (1) generate output from the most likely candidate set predicted by the LM and (2) generate output with a meaningful emotional association with the soft prompt. As a result, the generated text can increase the emotional expression while maintaining the semantic coherence associated with the prefix. Formally, our method can be formulated as follows:

$$sim = \max\{cos(h_w, h_{pr_j}) : 0 \leq j \leq len(Pr)\} \quad (5)$$

$$C(k) = P_{\theta_{pr}}(w | x_{<i>_i}) + \lambda \times sim \quad (6)$$

$$x_i = \arg \max C(k) \quad (7)$$

where  $W^{(k)}$  is the set of the most likely  $k$  tokens generated by the LM after affective prompt-tuning and  $P_{\theta_{pr}}(w | x_{<i>_i})$  is the probability of candidate  $w$  predicted by the LM.  $cos(h_w, h_{pr_j})$  measures the cosine similarity between the candidate token  $w$  and the affective soft prompt. This is due to the inclusion of emotional information in the affective soft prompt, which guides the language model (LM) in generating emotional texts. The cosine distance serves as a metric to gauge the similarity between the distribution of the next token and the affective soft prompt. Intuitively, the closer  $w$  is to the affective soft prompt, the more emotional information it contains, and thus the more likely it is to generate an emotionally rich response. The candidate representation  $h_w$  is computed by the LM given the concatenation of  $x_{<i>_i}$  and  $w$ . As depicted in Figure 1, we assume  $k$  is 5,  $\lambda$  is 0.6, and input is “I think love is.” First, APT-LM will generate 5 most likely candidate words as the next generation token, with generation weights of  $\{(a, 0.7), (the, 0.6), (beautiful, 0.5), (painful, 0.4), (gift, 0.3)\}$ . Second, by calculating the similarity between candidate words and emotional soft cues using equation (5), we can obtain:  $\{(beautiful, 0.8), (gift, 0.6), (painful, 0.3), (a, 0.2), (the, 0.1)\}$ . Finally, the final generated word weights can be obtained through equation (6), which are  $\{(beautiful, 0.98), (a, 0.82), (gift, 0.66), (the, 0.66), (painful, 0.52)\}$ . So, the next generation token is “beautiful,” resulting in the output text “I think love is beautiful.” The generation process will continue until the end condition is reached.

## Experimental Setup

### Datasets

We conduct experiments on two public datasets that include text and corresponding emotional labels. GoEmotions (Demszky et al., 2020) is a fine-grained multi-label English emotion dataset that consists of 58,000 Reddit comments with artificial emotion category tags. It includes 27 emotion categories, including one neutral category, and provides the mapping relationship of 27 emotion categories to Ekman’s six basic emotion categories. ISEAR (Scherer & Wallbott, 1994) is a single-label categorical emotion corpus that contains 7,666 sentences but covers only five basic emotions of Ekman’s emotion theory. We use only the train dataset covered with Ekman’s emotion theory for training. The statistics of datasets are shown in Table 2.

Table 2. Statistics of Datasets

Dataset	Types of emotion						overall
	anger	disgust	fear	joy	sadness	surprise	
GoEmotions	7,022	1,013	929	21,733	4,032	6,668	41,397
ISEAR	1,096	1,096	1,095	1,094	1,096	-	5,477

## Baseline Models

To the best of our knowledge, there is little research to address emotional text generation with prompting, and we focus on exploring the ability of the LM rather than RNN-based models for emotional text generation. Thus, we compare our method with the following baseline models:

**GPT2-Medium** (Radford et al., 2019): It is the 355M parameter version of GPT-2. To evaluate the effectiveness of the proposed method and model, we compare the vanilla GPT2-Medium with our proposed model.

**Affective-PPLM** (Goswamy et al., 2020): It serves as a plug-and-play extension of the PPLM framework proposed by Dathathri et al. (2019). It affords flexibility in selecting the base text generation model, specifying the emotion category from a range of eight basic emotions and offering fine-grained control over the intensity of emotion within each category.

**CTRL** (Keskar et al., 2019): It provides a conditional transformer language model with control codes. We add extra new codes for generating emotional text based on Ekman’s emotion theory by fine-tuning CTRL on GoEmotions and ISEAR datasets.

## Evaluation Metrics

In automatic evaluation, the evaluation of emotional text generation is still a challenging problem because there is no unified standard to evaluate the emotional factors well. Therefore, we intend to use the proportion of emotional tokens in the generated text and the external text emotion classifier for emotional evaluation. Then, we combine it with human evaluation to verify the effectiveness of our model.

### Automatic Evaluation

Perplexity (denoted as **PPL**) is commonly employed to assess the fluency of generated text. Lower perplexity is crucial in applications like machine translation, chatbots, and text generation, where natural and coherent language output is essential. A model with lower perplexity is more likely to generate text that is contextually appropriate and coherent. Distinct-N (denoted as **Dist-n**) quantifies the diversity of text in passages by assessing the count of distinct n-grams across all samples, divided by the total number of words (Li et al., 2016). In applications where variety in language is important, such as creative writing, content generation, or dialogue systems, a model with high Distinct-N is preferred. It ensures that the model can produce diverse and interesting output. Grammaticality (denoted as **Grammar**) verifies the grammatical correctness of the generated text (Warstadt et al., 2019). In applications where grammatical accuracy is crucial, such as content generation for professional use, educational tools, or formal communication, a model with high grammaticality is desired to ensure the quality of the generated content. Emotion intensity (denoted as **EmoInt**) measures the emotional expression of the generated text at the word level. In applications where emotional tone matters, like sentiment analysis, creative writing, or dialogue systems, a model with accurate emotion intensity can generate text that conveys the desired emotional tone effectively. Specifically, we calculate the emotion intensity of each word in the output based on the NRC lexicon and average the result to

obtain the emotion intensity of the output. It is noted that not every word exists in the NRC lexicon. For words not in the lexicon, we use the emotion word embeddings (EWE) (Agrawal et al., 2018) to calculate the cosine distance between the word and the corresponding emotion category. It is because the EWE has the capability to map emotionally similar words into proximate spaces while positioning emotionally dissimilar words at greater distances from each other. Accuracy (denoted as **Acc**) measures the emotional expression of the generated text at the sentence level. In emotion-sensitive applications, such as chatbots, virtual assistants, or customer support systems, accurate emotional expression is crucial for effective communication. A model with high accuracy in emotion expression is desirable for such applications. It evaluates whether the generated text accurately expresses the corresponding emotion. We use an external emotion classifier<sup>1</sup> on hugging face to verify whether the model is capable of generating text that corresponds to the given emotion. Its evaluation accuracy is 66%.

### Human Evaluation

Following Dathathri et al. (2019), we randomly select 20 prefixes to generate the continuation with different models and decoding methods (normal decoding and affective decoding). Then, we ask three human annotators to compare the generated text against two criteria: Fluency and AffectInt. AffectInt measures the emotional expression of a continuation, while Fluency evaluates the grammatical problem. According to the 5-point Likert scale theory, the scores assigned to all these evaluation metrics span from 1 to 5, corresponding to strongly disagree, disagree, not necessarily, agree, and strongly agree, respectively. During the evaluation process, all annotators were kept unaware of the model responsible for generating the continuation to ensure the integrity and validity of the results.

### Implementation Details

In this paper, we fine-tune CTRL on an open platform<sup>2</sup> that provides free computing resources with a Tesla V100. We adopt GPT2-Medium to generate emotional text and conduct other experiments on a RTX3060. For all models, we use the same 35 prefixes following Dathathri et al. (2019) to generate 10 continuations for each prefix, a total of 350 output and use both the top-k samples and the top-p samples during the decoding. The top-k is set to 10 and the top-p is set to 0.9. On affective prompt-tuning, the maximum length of the text is set to 64. Following the transformer fine-tuning configure<sup>3</sup>, the other parameters are the same as the default configure.

## EXPERIMENTAL RESULTS

### Automatic Evaluation Results

The evaluation results of different emotional generation models are reported in Table 3. The best result in each column is highlighted in bold. The vanilla GPT2-Medium performs the worst in emotional expression, with results of 0.152 and 0.274 for Acc and EmoInt metrics, respectively.

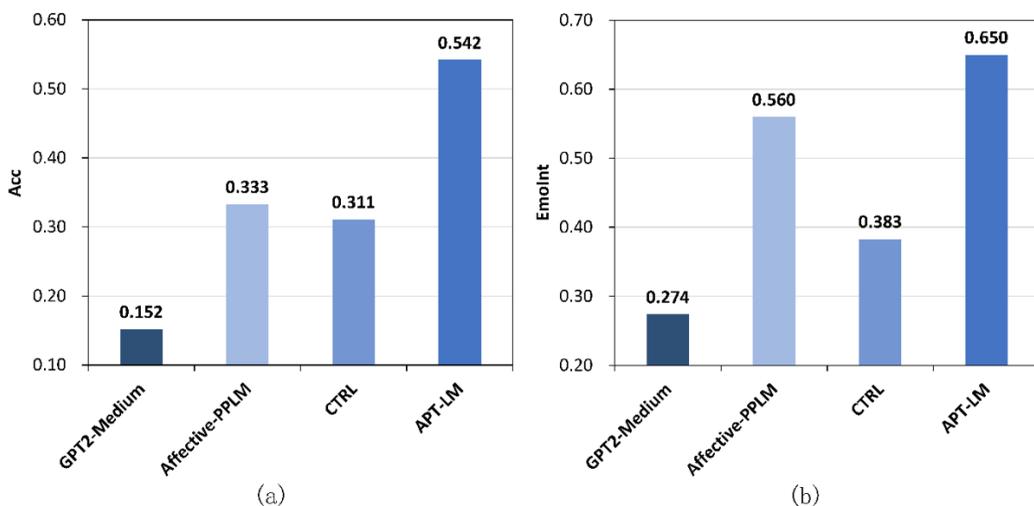
Table 3. Automatic Evaluation Results of Different Emotional Generation Models

Model		PPL	Dist-1	Dist-2	Dist-3	Grammar	Acc	EmoInt
GPT2-Medium		41.87	0.288	0.737	<b>0.906</b>	<b>0.797</b>	0.152	0.274
Affective-PPLM		<b>34.86</b>	<b>0.413</b>	<b>0.769</b>	0.863	0.789	0.333	0.560
CTRL	GoEmotions	36.25	0.214	0.542	0.663	0.692	0.336	0.397
	ISEAR	37.71	0.326	0.658	0.676	0.723	0.286	0.368
APT-LM	GoEmotions	38.79	0.209	0.637	0.836	0.746	<b>0.556</b>	0.646
	ISEAR	40.07	0.237	0.684	0.883	0.785	0.528	<b>0.654</b>

This indicates that a language model without fine-tuning or prompts struggles to leverage the emotional knowledge acquired during pre-training. Consequently, it faces difficulty accurately generating text corresponding to the required emotional categories, and the generated text contains less emotional information. Affective-PPLM shows superior performance in PPL, Dist-1, and Dist-2 metrics, which can be attributed to its use of a new loss term defined by Gaussian functions. This guides the model's generation towards the given emotion during the inference stage, improving the diversity and fluency of the generated text. However, it cannot accurately generate text corresponding to the given emotion, as reflected in its score of 0.333 for Acc metric. CTRL faces challenges in accurately and effectively generating emotional text as well. Its approach involves conditioning text generation on a control code, serving as an attribute variable that signifies a data source. Nonetheless, reliance on a specific control code may decrease sample diversity and hinder emotional text generation, as the generated samples tend to resemble the data source associated with the control code. Furthermore, CTRL provides each emotional control code for a large language model trained from scratch, which can be expensive. In contrast, our model requires updating only a small number of custom parameters to utilize the knowledge acquired by the language model during pre-training for emotional text generation.

As shown in Figure 2, we observe that APT-LM outperforms other benchmarks in terms of emotional metrics. APT-LM achieves the best average results on Acc and EmoInt metrics, with values of 0.542 and 0.650, respectively. Table 3 shows that, compared to the vanilla GPT2-Medium, APT-LM significantly improves emotional text generation while keeping competitive results in terms of PPL, Dist-1-2-3, and Grammar. This is because it leaves the parameters of the original language model unchanged, preserving the emotional knowledge acquired during its pre-training stage. Meanwhile, the affective soft prompt is appended at the beginning to the embedding of the input sentence, guiding the model to generate emotional text in a left-to-right fashion. Although Affective-PPLM shows improvement in emotion expression at the word level with an EmoInt metric of 0.560, it falls short in effectively conveying corresponding emotions at the sentence level. In contrast, APT-LM not only utilizes affective soft prompts to guide emotion expression at the sentence level but also strengthens the model's emotional expression at the word level through affective decoding. Moreover, CTRL requires expensive computation resource to train the model from scratch for emotional text generation, although it can improve sentence fluency to some extent. However,

Figure 2. Average Results of Emotional Metrics for Different Models



this undermines the generalization and knowledge that LM obtained during the pre-training stage. In contrast, in the case of limited computational resources, APT-LM can outperform CTRL by freezing the model parameters and updating additional affective parameters with affective prompt-tuning. This illustrates that APT-LM effectively employs custom parameters, specifically the affective soft prompt, to accurately convey emotions. Additionally, APT-LM employs affective decoding to assess the emotional similarity between the distribution of the next token and the affective soft prompt, thereby guiding the model to generate emotionally rich text. Also, we find that our model does not achieve the best performance on the Dist-N metrics, which is because the affective soft prompt steers the generation and the affective decoding reduces the probability distribution of other tokens that are not related to the corresponding emotion, resulting in the generation of less diverse but more emotional text. It is worth noting that, as shown in Table 3, APT-LM performs better in the PPL metric on the GoEmotions dataset compared to GPT2-Medium, with a score of 38.79, and still maintains strong competitiveness in text fluency. In conclusion, APT-LM focuses on assessing language models' proficiency in generating emotional texts. In comparison to alternative models, APT-LM excels not only in producing precise emotional texts, as indicated by the Acc metric, but also in incorporating a higher degree of emotional information, as reflected in the EmoInt metric. This underscores the effectiveness of our approach.

### Human Evaluation Results

In human evaluation, we focus on the fluency and affective intensity of the generated text. As shown in Table 4 and Figure 3, we observe that APT-LM outperforms GPT2-Medium, Affective-PPLM, and CTRL on AffectInt, with improvements of +0.92, +0.41, and +0.74, respectively. The observed improvements in AffectInt scores for APT-LM can be attributed to its specialized design for emotional text generation, the effective use of affective soft prompts and decoding mechanisms in comparison to other models. Furthermore, as shown in Table 4, APT-LM still outperforms other models in terms of Fluency, with a score of 3.63. Human annotators perceive that the text generated by our model conveys more emotional information and is closer to human natural language, contributing to APT-LM's superiority in the Fluency metric compared to other models. In conclusion, APT-LM is capable of generating coherent and emotional text, as it benefits from steering generation toward a given emotion by affective prompt-tuning (APT) and affective decoding (AD).

### Parameter Analysis

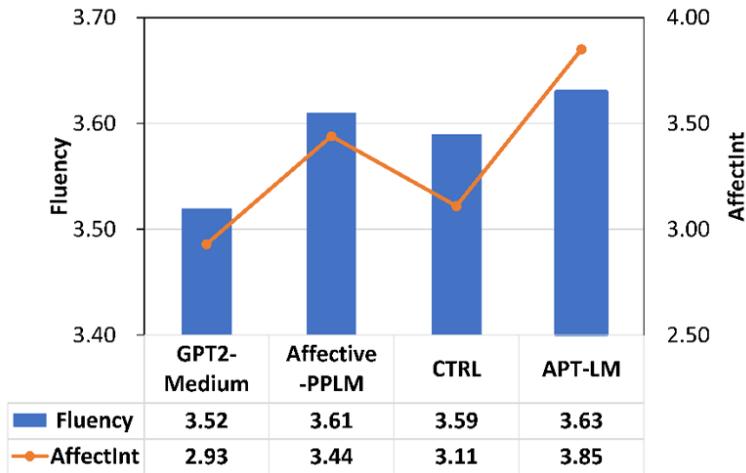
To investigate the impact of affective coefficient  $\lambda$  in equation (6) and the length of the soft prompt  $pr$  on the emotional expression capability of the model, we conduct parameter analysis experiments similar to automatic evaluation.  $\lambda$  can take values from  $\{0, 0.2, 0.4, 0.6, 0.8, 1\}$  and  $pr$  can be taken from  $\{1, 10, 20\}$ . We average the results of the two datasets and present the visual outcomes in terms of sentence fluency and emotional text generation.

As shown in Figure 4 (a), the PPL metric increases with the parameter  $\lambda$  increasing when the affective coefficient is greater than 0.6. When the length of the soft prompt is equal to 20, PPL is

Table 4. Human Evaluation Results of Different Emotional Generation Models

Model	Fluency	AffectInt
GPT2-Medium	3.52	2.93
Affective-PPLM	3.61	3.44
CTRL	3.59	3.11
APT-LM	<b>3.63</b>	<b>3.85</b>

Figure 3. Human Evaluation Results of Different Emotional Generation Models



relatively low. From Figure 4 (b) to (d), we observe that Distinct-N shows a downward trend with the increase of the affective coefficient (Distinct-N always reaches its maximum value when  $\lambda$  is 0). Similarly, the same phenomena can be found in Figure 4 (e) and the reason for this is the same as we mentioned in the analysis of automatic evaluation results, which is affective decoding makes the text generations less diverse but more emotional.

For the results in Figure 5, APT-LM presents better performance on both Acc and EmoInt metrics when  $pr = 20$ . Furthermore, when  $\lambda$  is equal to 0.6, it effectively enhances the model’s emotional expression capability without excessively shifting the tokens distribution over the vocabulary. Hence, the parameter  $\lambda$  is set to 0.6 and  $pr$  is set to 20 for a better generation in this paper.

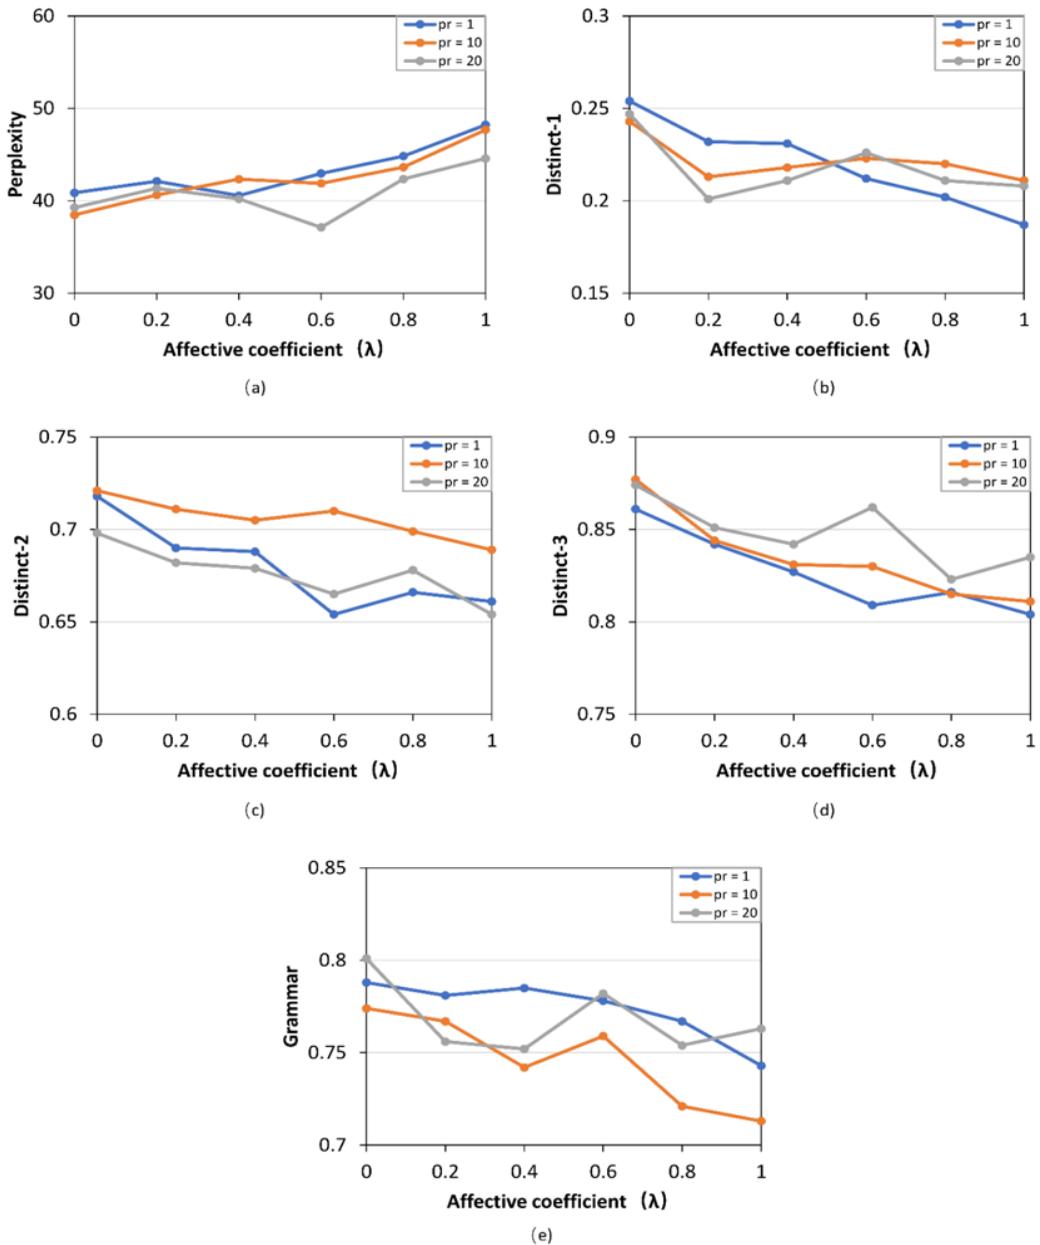
### Ablation Study

To further verify our proposed decoding method, we conduct ablation experiments for affective decoding while maintaining the parameters we discussed above. Table 5 shows the results of the ablation study, where w/o AD means we do not adopt affective decoding for the model. It is not surprising that APT-LM shows poorer performance in terms of text fluency and diversity, but stronger ability in emotional expression. This is because AD can further strengthen the ability of emotional text generation by shifting the token distribution over the vocabulary toward a given emotion. However, it is worth noting that the APT-LM achieves the best performance in terms of Fluency and AffectInt. This might be due to the fact that the texts generated by the APT-LM give more information about emotions, which is more in line with the language of humans.

### Case Study

Table 6 presents the emotional text generation results of different models. The prefix is underlined, and the words related to the given emotion are highlighted in bold. We set a maximum text length of 64. If the generated text exceeds this limit or if the model prematurely inserts a sentence-ending symbol, we use ellipses to indicate the truncation. For the joy emotion, the vanilla GPT2-Medium can generate a fluent but not emotional sentence. Affective-PPLM and CTRL generate weak emotional text compared to our model. Although the text generated by our model contains the emotional word “love” several times, it still maintains a high level of linguistic fluency. For the sadness emotion, we use the same prefix for comparison. It is clear that our model tends to generate more emotional expression and the Affective-PPLM fails to maintain good language

Figure 4. Effect of the Parameter  $\lambda$  and  $pr$  on Sentence Fluency



fluency with the same emotional word “prison”. The CTRL model produces a higher number of emotional statements compared to the vanilla GPT2-Medium, although it does not reach the level of emotional expressiveness demonstrated by our model. In conclusion, by using affective prompting and affective decoding, APT-LM produces responses that not only convey information but also resonate with the emotional undertones present in the input. Furthermore, APT-LM exhibits the capability to generate emotional text by adjusting the emotional information within soft prompt, encompassing emotions like joy and sadness. Theoretically, APT-LM can extend

Figure 5. Effect of the Parameter  $\lambda$  and  $pr$  on Emotional Text Generation

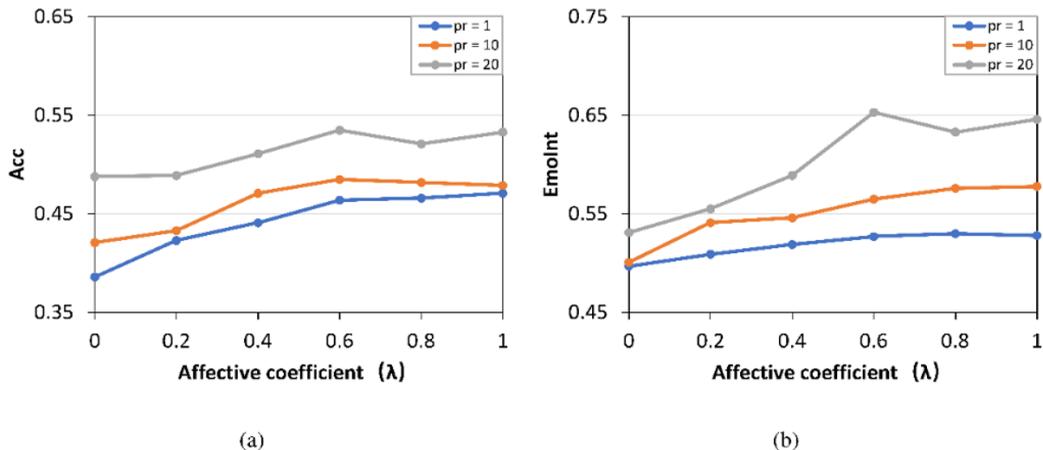


Table 5. Ablation Study on Different Datasets

Dataset	Model	Automatic Evaluation							Human Evaluation	
		PPL	Dist-1	Dist-2	Dist-3	Grammar	Acc	EmoInt	Fluency	AffectInt
GoEmotions	APT-LM	38.79	0.209	0.637	0.836	0.746	<b>0.556</b>	<b>0.646</b>	<b>3.62</b>	<b>3.88</b>
	w/o AD	<b>37.87</b>	<b>0.22</b>	<b>0.654</b>	<b>0.853</b>	<b>0.801</b>	0.469	0.537	3.55	3.65
ISEAR	APT-LM	40.07	0.237	0.684	0.883	0.785	<b>0.528</b>	<b>0.654</b>	<b>3.64</b>	<b>3.82</b>
	w/o AD	<b>38.21</b>	<b>0.240</b>	<b>0.697</b>	<b>0.891</b>	<b>0.821</b>	0.488	0.564	3.58	3.71

this proficiency to effectively incorporate other additional attribute information such as topic, sentiment, gender, and personality. This greatly enhances the scalability and generalization of the model, allowing APT-LM to be applied to more fields, such as recommendation systems and personalized AI. This methodology empowers the model to make efficient use of its potent language generation capabilities. Notably, APT-LM accomplishes this with minimal reliance on labeled data, underscoring its prowess in guiding controlled text generation while fully harnessing the robust text generation capabilities inherent in language models. This advancement contributes to a deeper comprehension of the emotional context inherent in human communication within the field of natural language understanding and processing.

## CONCLUSION

In this paper, we proposed an affective prompt-tuning based language model (APT-LM) and an affective decoding (AD) method for steering generation toward a given emotion. Due to the need for significant computing resources, little work has been done to generate emotional text using language model like GPT-2. With limited computational resources, APT-LM only updates the additional parameters while freezing the language model during training. Compared with the baseline model, our model shows significant enhancement in emotional text generation in both automatic and human evaluation, while maintaining competitiveness in sentence fluency and diversity. Moreover, we validated the effectiveness of AD through ablation experiments and demonstrated the powerful

Table 6. Case Study Results

Emotion	Model	Generation
Joy	GPT2-Medium	<u>In brief</u> : I will be making a statement on Wednesday and I will post it here on the site. You can also find me in Twitter, Facebook, YouTube and here on reddit. Please let me know if there is anything I've missed or can add. Thanks for reading...
	Affective-PPLM	<u>In brief</u> , the new law will allow people convicted of drug offenses to receive credit for any community <b>service</b> they have served. That could help to keep some drug offenders off welfare. The law's critics are hoping that the Supreme Court's landmark ruling will make a <b>victory celebration</b> more common...
	CTRL	<u>In brief</u> , I have to say that it is a very <b>good</b> book. It has some great tips on how to get the most out of your trip in an efficient way. The only thing missing from my opinion was more detail about what you can do when there are no other options for getting around...
	APT-LM	<u>In brief</u> , you can see this <b>kind of love</b> is one of the <b>best</b> ways to bring people to your <b>life</b> . <b>Love</b> is a <b>gift</b> that cannot be faked and is always <b>genuine</b> . So, why not use this <b>love</b> to build your <b>life</b> in the most <b>beautiful</b> way possible, like <b>loving</b> to learn! Here are 6 <b>Love Skills</b> that...
Sadness	GPT2-Medium	<u>In brief</u> , one can understand the necessity of a state that is willing to pay to protect a population against non-state <b>violence</b> . It is understandable why the state pays for it. A state is willing to accept a certain amount of income as payment for protection from <b>violence</b> , if that is the only means available to pay for...
	Affective-PPLM	<u>In brief</u> , when you are in the midst of a legal matter, the state attorney general is the <b>ultimate</b> law judge. If you're accused of a crime, the prosecutor is your trial lawyer and the state attorney is your jail. The criminal court is a judicial <b>prison, jail prison prison prison...</b>
	CTRL	<u>In brief</u> , the relationship was over. <b>Sad</b> but true. I'm so sorry for your <b>loss</b> and <b>sadness</b> you must feel right now. It is a very difficult thing to deal with, it's hard to understand what he went through the <b>pain of losing</b> his wife...
	APT-LM	<u>In brief</u> I am a doctor with about ten years experience in treating <b>cancer</b> . I also have some experience working on people <b>suffering</b> from <b>depression</b> and <b>anxiety</b> , and other health related <b>anxiety</b> as well. I have also worked on <b>patients suffering</b> from the effects of <b>chronic pain</b> and other medical conditions including <b>cancer</b> ...

emotional text generation ability of APT-LM through case study. Additionally, APT-LM has strong scalability and can be used in multiple applications by adjusting custom training parameters, such as chatbots, dialogue systems, creative writing, educational tools, sensory analysis, virtual assistants, and customer support systems. Although APT-LM demonstrates strong emotional text generation capabilities, it can only generate single attribute text at a time and may generate harmful or misleading information. For future work, we can enhance the generated text of APT-LM by employing multi-task learning, integrating additional attributes such as topic. Additionally, through joint training, we can train an auxiliary detoxification classifier to guide APT-LM in generating text that is harmless. Furthermore, delving into the interpretability of the proposed model represents a promising direction for future research.

**AUTHOR NOTE**

Correspondence concerning this article should be addressed to Kejing He, School of Computer Science and Engineering, South China University of Technology. E-mail: kejinghe@ieec.org

**DATA AVAILABILITY**

The data used to support the findings of this study are included within the article.

## **CONFLICT OF INTEREST**

The authors declares no conflicts of interest.

## **FUNDING STATEMENT**

This work was supported by the Fundamental Research Funds for the Central Universities (2022ZYGXZR004).

## REFERENCES

- Abo-Hammour, Z., Abu Arqub, O., Momani, S., & Shawagfeh, N. (2014). Optimization solution of Troesch's and Bratu's problems of ordinary type using novel continuous genetic algorithm. *Discrete Dynamics in Nature and Society*, 2014, 1–15. doi:10.1155/2014/401696
- Abo-Hammour, Z., Alsmadi, O., Momani, S., & Abu Arqub, O. (2013). A genetic algorithm approach for prediction of linear dynamical systems. *Mathematical Problems in Engineering*, 2013, 1–12. doi:10.1155/2013/831657
- Abu Arqub, O., Abo-Hammour, Z., Momani, S., & Shawagfeh, N. (2012). Solving singular two-point boundary value problems using continuous genetic algorithm. *Abstract and Applied Analysis*, 2012, 1–25. doi:10.1155/2012/205391
- Agrawal, A., An, A., & Papagelis, M. (2018). Learning emotion-enriched word representations. *Proceedings of the 27th international conference on computational linguistics*, 950–961.
- Arqub, O. A., & Abo-Hammour, Z. (2014). Numerical solution of systems of second-order boundary value problems using continuous genetic algorithm. *Information Sciences*, 279, 396–415. doi:10.1016/j.ins.2014.03.128
- Barbosa, A., Bittencourt, I. I., Siqueira, S. W., Dermeval, D., & Cruz, N. J. (2022). A context-independent ontological linked data alignment approach to instance matching. *International Journal on Semantic Web and Information Systems*, 18(1), 1–29. doi:10.4018/IJSWIS.295977
- Bisht, J., & Vampugani, V. S. (2022). Load and cost-aware min-min workflow scheduling algorithm for heterogeneous resources in fog, cloud, and edge scenarios. *International Journal of Cloud Applications and Computing*, 12(1), 1–20. doi:10.4018/IJCAC.2022010105
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., & Amodei, D. et al. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.
- Casillo, M., Gupta, B. B., Lombardi, M., Lorusso, A., Santaniello, D., & Valentino, C. (2022). Context aware recommender systems: A novel approach based on matrix factorization and contextual bias. *Electronics (Basel)*, 11(7), 1003. doi:10.3390/electronics11071003
- Chen, X., Li, J., & Zhang, Y. F. (2022). Multidirectional gradient feature with shape index for effective texture classification. *International Journal on Semantic Web and Information Systems*, 18(1), 1–19. doi:10.4018/IJSWIS.312183
- Cho, K., van Merriënboer, B., Gülçehre, Ç., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using RNN encoder-decoder for statistical machine translation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1724–1734. doi:10.3115/v1/D14-1179
- Chopra, M., Singh, S. K., Sharma, A., & Gill, S. S. (2022). A comparative study of generative adversarial networks for text-to-image synthesis. *International Journal of Software Science and Computational Intelligence*, 14(1), 1–12. doi:10.4018/IJSSCI.300364
- Colombo, P., Witon, W., Modi, A., Kennedy, J., & Kapadia, M. (2019). Affect-driven dialog generation. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 3734–3743.
- Dathathri, S., Madotto, A., Lan, J., Hung, J., Frank, E., Molino, P., Yosinski, J., & Liu, R. (2019). Plug and play language models: A simple approach to controlled text generation. *International Conference on Learning Representations*.
- Demszky, D., Movshovitz-Attias, D., Ko, J., Cowen, A., Nemade, G., & Ravi, S. (2020). Goemotions: A dataset of fine-grained emotions. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 4040–4054. doi:10.18653/v1/2020.acl-main.372
- Deveci, M., Gokasar, I., Pamucar, D., Zaidan, A. A., Wen, X., & Gupta, B. B. (2023). Evaluation of cooperative intelligent transportation system scenarios for resilience in transportation using type-2 neutrosophic fuzzy VIKOR. *Transportation Research Part A, Policy and Practice*, 172, 103666. doi:10.1016/j.tra.2023.103666

Devlin, J., Chang, M., Lee, K., & Toutanova, K. (2019). Bert: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 4171-4186.

Ekman, P. (1992). An argument for basic emotions. *Cognition and Emotion*, 6(3-4), 169-200. doi:10.1080/02699939208411068

Ghosh, S., Chollet, M., Laksana, E., Morency, L.-P., & Scherer, S. (2017). Affect-Im: A neural language model for customizable affective text generation. *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, 1, 634-642. doi:10.18653/v1/P17-1059

Goswamy, T., Singh, I., Barkati, A., & Modi, A. (2020). Adapting a language model for controlled affective text generation. *Proceedings of the 28th International Conference on Computational Linguistics*, 2787-2801. doi:10.18653/v1/2020.coling-main.251

Gupta, B. B., Gaurav, A., Panigrahi, P. K., & Arya, V. (2023). Analysis of artificial intelligence-based technologies and approaches on sustainable entrepreneurship. *Technological Forecasting and Social Change*, 186, 122152. doi:10.1016/j.techfore.2022.122152

Houlsby, N., Giurgiu, A., Jastrzebski, S., Morrone, B., De Laroussilhe, Q., Gesmundo, A., Attariyan, M., & Gelly, S. (2019). Parameter-efficient transfer learning for NLP. *International Conference on Machine Learning*, 2790-2799.

Ilyas, Q. M., Ahmad, M., Rauf, S., & Irfan, D. (2022). RDF query path optimization using hybrid genetic algorithms: Semantic web vs. data-intensive cloud computing. *International Journal of Cloud Applications and Computing*, 12(1), 1-16. doi:10.4018/IJCAC.2022010101

Ismail, S., Shishtawy, T. E., & Alsammak, A. K. (2022). A new alignment word-space approach for measuring semantic similarity for Arabic text. *International Journal on Semantic Web and Information Systems*, 18(1), 1-18. doi:10.4018/IJSWIS.297036

Lester, B., Al-Rfou, R., & Constant, N. (2021). The power of scale for parameter-efficient prompt tuning. *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 3045-3059. doi:10.18653/v1/2021.emnlp-main.243

Li, J., Galley, M., Brockett, C., Gao, J., & Dolan, B. (2016). A diversity-promoting objective function for neural conversation models. *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 110-119. doi:10.18653/v1/N16-1014

Li, S., Gao, L., Han, C., Gupta, B., Alhalabi, W., & Almakdi, S. (2023). Exploring the effect of digital transformation on firms' innovation performance. *Journal of Innovation & Knowledge*, 8(1), 100317. doi:10.1016/j.jik.2023.100317

Li, X. L., & Liang, P. (2021). Prefix-tuning: Optimizing continuous prompts for generation. *Proceedings of the 59th annual meeting of the association for computational linguistics and the 11th international joint conference on natural language processing*, 1, 4582-4597.

Mohammad, S. (2018). Obtaining reliable human ratings of valence, arousal, and dominance for 20,000 English words. *Proceedings of the 56th annual meeting of the association for computational linguistics*, 1, 174-184. doi:10.18653/v1/P18-1017

Mohammad, S., & Kiritchenko, S. (2018). Understanding emotions: A dataset of tweets to study interactions between affect categories. *Proceedings of the 11th international conference on language resources and evaluation*.

Pascual, D., Egressy, B., Meister, C., Cotterell, R., & Wattenhofer, R. (2021). A plug-and-play method for controlled text generation. *Findings of the Association for Computational Linguistics: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 3973-3997.

Plutchik, R. (2001). The nature of emotions: Human emotions have deep evolutionary roots, a fact that may explain their complexity and provide tools for clinical practice. *American Scientist*, 89(4), 344-350. doi:10.1511/2001.28.344

Prendinger, H., & Ishizuka, M. (2005). The empathic companion: A character-based interface that addresses users' affective states. *Applied Artificial Intelligence*, 19(3-4), 267-285. doi:10.1080/08839510590910174

- Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). Improving language understanding by generative pre-training. *OpenAI*, *1*(6), 5.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). Language models are unsupervised multitask learners. *OpenAI*, *1*(8), 9.
- Rebuffi, S.-A., Bilen, H., & Vedaldi, A. (2017). Learning multiple visual domains with residual adapters. *Proceedings of the 31st international conference on neural information processing systems*, 506-516.
- Russell, J. A. (1994). Is there universal recognition of emotion from facial expression? A review of the cross-cultural studies. *Psychological Bulletin*, *115*(1), 102–141. doi:10.1037/0033-2909.115.1.102 PMID:8202574
- Russell, J. A. (2003). Core affect and the psychological construction of emotion. *Psychological Review*, *110*(1), 145–172. doi:10.1037/0033-295X.110.1.145 PMID:12529060
- Sarivougioukas, J., & Vagelatos, A. (2022). Fused contextual data with threading technology to accelerate processing in home UbiHealth. *International Journal of Software Science and Computational Intelligence*, *14*(1), 1–14. doi:10.4018/IJSSCI.285590
- Scherer, K. R., & Wallbott, H. G. (1994). Evidence for universality and cultural variation of differential emotion response patterning. *Journal of Personality and Social Psychology*, *66*(2), 310–328. doi:10.1037/0022-3514.66.2.310 PMID:8195988
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
- Warstadt, A., Singh, A., & Bowman, S. R. (2019). Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics*, *7*, 625–641. doi:10.1162/tacl\_a\_00290
- Yang, M., Yin, W., Qu, Q., Tu, W., Shen, Y., & Chen, X. (2019). Neural attentive network for cross-domain aspect-level sentiment classification. *IEEE Transactions on Affective Computing*, *12*(3), 761–775. doi:10.1109/TAFCC.2019.2897093
- Yen, S., Moh, M., & Moh, T.-S. (2021). Detecting compromised social network accounts using deep learning for behavior and text analyses. *International Journal of Cloud Applications and Computing*, *11*(2), 97–109. doi:10.4018/IJCAC.2021040106
- Zhang, H., Song, H., Li, S., Zhou, M., & Song, D. (2023). A survey of controllable text generation using transformer-based pre-trained language models. *ACM Computing Surveys*, *56*(3), 1–37. doi:10.1145/3617680
- Zhang, Q., Guo, Z., Zhu, Y., Vijayakumar, P., Castiglione, A., & Gupta, B. B. (2023). A deep learning-based fast fake news detection model for cyber-physical social services. *Pattern Recognition Letters*, *168*, 31–38. doi:10.1016/j.patrec.2023.02.026
- Zhang, Z., Li, C., Gupta, B. B., & Niu, D. (2018). Efficient compressed ciphertext length scheme using multi-authority CP-ABE for hierarchical attributes. *IEEE Access : Practical Innovations, Open Solutions*, *6*, 38273–38284. doi:10.1109/ACCESS.2018.2854600
- Zhou, H., Huang, M., Zhang, T., Zhu, X., & Liu, B. (2018). Emotional chatting machine: Emotional conversation generation with internal and external memory. *Proceedings of the 32nd AAAI conference on artificial intelligence and 30th innovative applications of artificial intelligence conference and 8th AAAI symposium on educational advances in artificial intelligence*, 730-738. doi:10.1609/aaai.v32i1.11325

## ENDNOTES

- <sup>1</sup> <https://huggingface.co/j-hartmann/emotion-english-distilroberta-base>
- <sup>2</sup> <https://openi.pcl.ac.cn/>
- <sup>3</sup> [https://github.com/huggingface/transformers/blob/main/examples/pytorch/language-modeling/run\\_clm\\_no\\_trainer.py](https://github.com/huggingface/transformers/blob/main/examples/pytorch/language-modeling/run_clm_no_trainer.py)