

## Article

# EvoText: Enhancing Natural Language Generation Models via Self-Escalation Learning for Up-to-Date Knowledge and Improved Performance

Zhengqing Yuan <sup>1,\*</sup> , Huiwen Xue <sup>2,†</sup>, Chao Zhang <sup>1,†</sup> and Yongming Liu <sup>1,\*</sup> <sup>1</sup> School of Artificial Intelligence, Anhui Polytechnic University, Wuhu 241009, China<sup>2</sup> School of Optoelectronic Science and Engineering, Soochow University, Suzhou 215031, China

\* Correspondence: zhengqingyuan@ieee.org (Z.Y.); liuyongming1015@163.com (Y.L.)

† These authors contributed equally to this work.

**Abstract:** In recent years, pretrained models have been widely used in various fields, including natural language understanding, computer vision, and natural language generation. However, the performance of these language generation models is highly dependent on the model size and the dataset size. While larger models excel in some aspects, they cannot learn up-to-date knowledge and are relatively difficult to relearn. In this paper, we introduce EvoText, a novel training method that enhances the performance of any natural language generation model without requiring additional datasets during the entire training process (although a prior dataset is necessary for pretraining). EvoText employs two models: *G*, a text generation model, and *D*, a model that can determine whether the data generated by *G* is legitimate. Initially, the fine-tuned *D* model serves as the knowledge base. The text generated by *G* is then input to *D* to determine whether it is legitimate. Finally, *G* is fine-tuned based on *D*'s output. EvoText enables the model to learn up-to-date knowledge through a self-escalation process that builds on a priori knowledge. When EvoText needs to learn something new, it simply fine-tunes the *D* model. Our approach applies to autoregressive language modeling for all Transformer classes. With EvoText, eight models achieved stable improvements in seven natural language processing tasks without any changes to the model structure.

**Keywords:** training system; fine-tuning; BERT; GPT

**Citation:** Yuan, Z.; Xue, H.; Zhang, C.; Liu, Y. EvoText: Enhancing Natural Language Generation Models via Self-Escalation Learning for Up-to-Date Knowledge and Improved Performance. *Appl. Sci.* **2023**, *13*, 4758. <https://doi.org/10.3390/app13084758>

Academic Editors: Ahmed Rafea and Julian Szymanski

Received: 10 March 2023

Revised: 3 April 2023

Accepted: 7 April 2023

Published: 10 April 2023



**Copyright:** © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

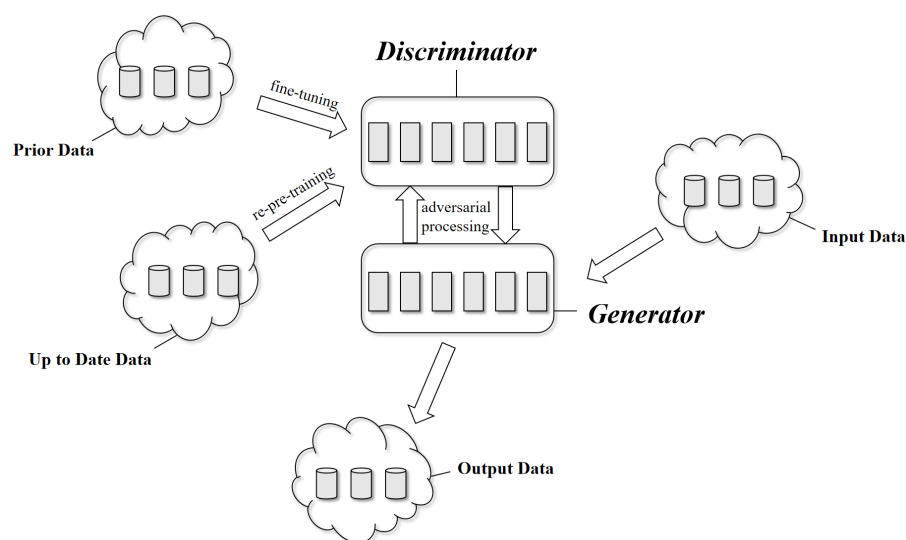
Pre-training models have shown great promise in natural language processing, with the Transformer model [1] proposing an encoder–decoder architecture based solely on the self-attention mechanism, enabling the construction of large-scale models that can be pretrained on vast amounts of data. Language models [2–4] can be broadly categorized into two types: autoregressive language modeling and autoencoder language modeling. autoregressive language models, such as ELMO [5], GPT [6], and T5 [7], predict the next possible word based on the preceding context, making them well-suited for generative tasks. On the other hand, autoencoder language models, such as BERT [8] and RoBERTa [9], predict intermediate words based on context and are better suited for natural language understanding tasks.

In recent years, generative models, including VAE [10], GAN [11], and DDPM [12], have made significant progress in computer vision. However, natural language generation presents unique challenges due to its discrete and temporal nature. To address these challenges, a common approach is to use unsupervised training of large language models, such as GPT-3 [13], which has 175 billion parameters. Despite their potential, training these models can be challenging due to their size, and further training once deployed can be difficult. As a result, a zero-shot approach is often adopted, which does not require fine-tuning the model for specific downstream tasks. However, this approach has limitations;

large language models may not perform as well as smaller models with fine-tuning for certain tasks that rely heavily on supervised learning. It is also akin to using a computer that has not been upgraded for an extended period, and its obsolescence is only a matter of time. Therefore, there is an urgent need for novel approaches that can balance the benefits and limitations of large language models to improve their performance and longevity. While reinforcement learning from human feedback (RLHF) [14] is one such approach that holds promise, it is still limited by the availability and quality of human feedback data, and it may not always generalize well to other tasks.

This paper introduces a novel training process for pretrained models, which we call **EvoText**. The proposed EvoText method can continuously learn from new data without suffering from the limitations of unsupervised learning or requiring additional datasets for fine-tuning. Specifically, we merge the input text and the generated text and then use a natural language understanding model to label the data for supervised fine-tuning of the generative model. Simply retraining the smaller discriminator model is required to enable the generative model to acquire up-to-date knowledge. To address the issues of natural language understanding errors and overfitting, we adopt a small learning rate and epoch size during fine-tuning. Unlike GAN models, we do not modify the parameters of the natural language understanding model during training, and only fine-tune it when necessary to incorporate new knowledge. The contributions of this work are as follows:

- EvoText partially mitigates the problem of low-quality samples generated by the generative model.
- EvoText improves the model's performance without additional data during the training process. (Note that while additional datasets are used in the system for warm-up and learning up-to-date data, as illustrated in Figure 1, only the data generated by the generator are used in the crucial training process.)
- EvoText enables continuous and sustainable improvement in seven natural language processing tasks, including natural language understanding and natural language generation, without altering the model's structure.
- The proposed method achieves results comparable to those of large generative networks, even with relatively limited computational resources.
- We will make the source code for EvoText publicly available on GitHub (<https://github.com/DLYuanGod/Auto-learning> (accessed on 9 March 2023)).



**Figure 1.** EvoText full process.

The novelty of this paper lies in the proposed method for enhancing the performance of natural language generation models without the need for additional datasets during training. With this approach, EvoText allows the generative model to acquire up-to-date

knowledge with minimal additional cost. This is achieved by updating the knowledge base through retraining the smaller discriminator model, as opposed to retraining the entire model with additional data.

## 2. Background

In this section, we briefly overview the language model, Transformer, the GPT series of models, the BERT series of models, and the GAN model.

### 2.1. Language Modeling

In this subsection, we briefly describe the implementation of the two language modelings.

#### 2.1.1. Autoregressive Language Modeling

Autoregressive language modeling [15,16] is a type of language modeling used in natural language processing. It involves predicting the next token in a sequence based on the previous tokens. Given a sequence of  $N$  tokens,  $(t_1, t_2, \dots, t_N)$ , the probability of a token  $t_k$  is modeled by calculating the conditional probability of  $t_k$  given all preceding tokens,  $(t_1, t_2, \dots, t_{k-1})$ , using the following formula:

$$P(t_1, t_2, \dots, t_N) = \prod_{k=1}^N P(t_k | t_1, t_2, \dots, t_{k-1}) \quad (1)$$

This Formula (1) calculates the joint probability of all the tokens in the sequence. The product of all the conditional probabilities is taken from the first token to the last token. This means that to predict the probability of a token, we need to know the probability of all preceding tokens. In the case of the backward model, the token after  $t_k$  needs to be calculated.

#### 2.1.2. Autoencoder Language Modeling

Autoencoder language modeling (ALM) [17,18] is a type of language modeling that involves predicting a target token in a sequence of tokens, while considering the probabilities of all preceding and following tokens. This can be represented mathematically as predicting  $t_k$  in the sequence  $(t_1, t_2, \dots, t_k, \dots, t_N)$ , where the probabilities of  $t_{1:k-1}$  and  $t_{k+1:N}$  are also calculated.

ALM is a powerful modeling method that can be applied to natural language understanding tasks. By considering the entire sentence when predicting a target token, ALM can capture the semantic and syntactic relationships between words in the sentence. This makes it particularly useful for tasks such as language generation, summarization, and machine translation.

The concept of ALM is closely related to the masked language model (MLM), which is used as a pretraining strategy in the popular BERT model. In MLM, a percentage of the tokens in a sequence are randomly masked out, and the model is trained to predict the missing tokens while considering the context of the surrounding tokens. This approach is similar to ALM, as it also involves predicting a target token while considering the context of the surrounding tokens.

### 2.2. Transformer

The Transformer model [3,19–23] is modeled and applied to natural language processing tasks using only self-attentive mechanisms.

#### 2.2.1. Transformer Encoder

Transformer Encoder takes a sequence of tokens as input, which are first processed through a word embedding and positional embedding layer. The resulting vector dimension is called  $d_{model}$ .

Next, the Transformer Encoder uses a self-attentive mechanism to compute the output tokens. This mechanism involves creating three copies of the input token, which are referred to as  $Q$ ,  $K$ , and  $V$ . Each of these copies is used in the attention calculation, which computes the weights between all pairs of tokens in the input sequence. The attention calculation formula is shown below:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$

Here,  $d_k$  represents the vector dimension processed by each attention head. The value of  $d_k$  is equal to  $d_{model}$  divided by the number of attention heads used in the multi-head attention mechanism.

Finally, the values calculated for each attention head are concatenated and passed through a multi-layer perceptron (MLP) layer to produce the final output tokens. Overall, the self-attentive mechanism used in the Transformer Encoder allows the model to capture complex relationships between tokens in the input sequence and produce high-quality representations for natural language processing tasks.

### 2.2.2. Transformer Decoder

The decoder module is calculated in the same way as the encoder module, except that masking is added to the multi-headed attention mechanism to mask tokens that have not yet been generated.

## 2.3. GPT Series of Models

### 2.3.1. GPT

The Generative Pre-training Transformer (GPT) [6] was introduced by Radford et al. in 2018 as an improvement on the Transformer model, which had been mainly used for natural language understanding tasks. GPT was the first model to apply a pretrained Transformer model to natural language processing.

GPT uses a multi-layer Transformer Decoder for the language model, which consists of 12 blocks of Transformer Decoders [24] with up to 117 million parameters. This approach allows GPT to generate high-quality natural language text by predicting the next word in a sequence of words.

Overall, GPT has been shown to achieve impressive results on a range of natural language processing tasks, such as text classification, language translation, and text generation. Its success has led to the development of larger and more powerful Transformer-based models, such as GPT-2 and GPT-3, which continue to push the boundaries of natural language processing.

### 2.3.2. GPT-2 and GPT-3

GPT-2 [25] and GPT-3 [13] represent advanced versions of the original Generative Pre-training Transformer (GPT) model. They employ 48 and 96-layer Transformer Decoder stacks, respectively, with a significantly larger number of parameters: 1.5 billion and 175 billion.

In addition to their larger model sizes, both GPT-2 and GPT-3 are trained on larger and more diverse datasets, enabling them to capture more complex and nuanced patterns in natural language. As a result, these models have achieved state-of-the-art performance on a range of natural language processing tasks, including language translation, question answering, and text generation.

The success of GPT-2 and GPT-3 highlights the tremendous potential of pretrained Transformer models in natural language processing research. With continued advances in this field, we can expect even more powerful language models in the future.

#### 2.4. BERT Series of Models

BERT [8,26], RoBERTa [9], ALBERT [27], XLNET [28], TinyBERT [29], and ELECTRA [30] are all state-of-the-art natural language understanding models that employ the Transformer Encoder layer. While there are slight differences in their training approaches, they all share the same underlying architecture.

The architecture of the Transformer Encoder layer is particularly well-suited for natural language understanding tasks, as it allows the model to capture long-range dependencies between words and phrases in the text. By leveraging self-attention mechanisms, these models can dynamically weigh the importance of different words in the input sequence, enabling them to extract more nuanced and complex representations of language.

#### 2.5. GAN

The main idea of generative adversarial networks (GAN) [11,31,32] is to build two models, a generator ( $G$ ) model and a discriminator ( $D$ ) model. During training, the  $G$  model tries to improve its manufacturing process to create realistic outputs that can fool the  $D$  model, while the  $D$  model tries to accurately distinguish between real and generated data like a police officer inspecting forgeries. However, achieving a balance between the two models is essential, since an imbalance can lead to one model not converging. Specifically, the  $G$  model needs to create realistic outputs that can fool the  $D$  model, while the  $D$  model needs to accurately distinguish between real and generated data. Despite the challenges, GANs have shown great potential in generating high-quality and diverse samples in a range of applications, such as image synthesis, text generation, and music composition.

#### 2.6. RLHF

Reinforcement-learning-based training methods have shown state-of-the-art success in recent years in natural language processing tasks. One of the most advanced training methods is the reinforcement learning with hybrid feedback (RLHF) approach. This method combines the advantages of both human feedback and self-supervised learning, enabling the model to learn from its own mistakes while also benefiting from human expertise. RLHF has been successfully applied in various tasks, including machine translation, text generation, and summarization.

This is due to the absence of human feedback in the approach proposed in this article. So we used the feedback provided by ChatGPT [33].

### 3. EvoText

In this section, the training process of EvoText is introduced, and the theoretical representation and algorithmic implementation are given.

#### 3.1. Prior Learning of Discriminator

Given an a priori dataset  $Data = \{(x_1, y_1), (x_2, y_2), \dots, (x_N, y_N)\}$ , which could be related to tasks, such as grammar judgment or semantic rationalization, we aim to fine-tune the discriminator model using the following objective function:

$$\min_{\theta} \mathbb{E}_{Data} \left[ \sum_{n=1}^N L(D_{\theta}(x_n), y_n) \right] \quad (2)$$

Here,  $D_{\theta}$  represents the pretrained natural language understanding model, and  $L$  represents the loss function.

### 3.2. Pre-Warm-Up Training for Generator

We need to freeze all Transformer blocks in the pretrained natural language generative model  $G$  and add new linear and softmax layers on top. We then train the newly added layers using the prior dataset through the following equation:

$$\min_{\theta} \mathbb{E}_{Data} \left[ \sum_{n=1}^N L \left( F_{\theta} \left( G_{\Theta}^{blocks} (x_n) \right), y_n \right) \right] \tag{3}$$

where  $G_{\Theta}^{blocks}$  refers to all the frozen Transformer blocks of the pretrained model  $G$  with parameters  $\Theta$ , and  $F_{\theta}$  represents the newly added linear layer and its trainable parameters. The objective is to minimize the expected value of the loss function  $L$  on the prior dataset, where  $y_n$  is the ground truth label for input example  $x_n$ .

### 3.3. Training Dataset

Give a new token  $Z_K^{in} = \{z_1, z_2, \dots, z_k\}$  representing the input generative model  $G$ , such as "Yesterday", the token is fed into the  $G$  model to obtain  $Z_K^{out} = \{z_1, z_2, \dots, z_{k+n}\}$ , e.g., "Yesterday, a man named Jack said he saw an alien, the skies near New Orleans". It can express as  $Z_K^{out} = G_{\Theta}(Z_K^{in})$ . The generated token  $Z_K^{out}$  is fed into the discriminator model  $D$  to obtain the label  $Y_K = D_{\theta}(Z_K^{out})$  for each token. Suppose we need to construct  $M$  samples, and finally, we obtain the posterior data as  $Z = \{(Z_1^{out}, Y_1), (Z_2^{out}, Y_2), \dots, (Z_M^{out}, Y_M)\}$ .

### 3.4. Supervised Fine-Tuning for Generators

Supervised training has a greater impact on the model than unsupervised training, even when using the same dataset size. To fine-tune the parameters of all Transformer blocks in the generator  $G_{\Theta}^{blocks}$ , we enable the gradient of all generator parameters. The optimization objective is defined as follows:

$$\min_{\Phi} \mathbb{E}_{(Z^{out}, Y) \sim Z} \left[ \sum_{n=1}^M L \left( F_{\theta} \left( G_{\Theta}^{blocks} (Z_n^{out}) \right), Y_n \right) \right] \tag{4}$$

where  $\Phi$  denotes the parameters of the model  $G_{\Theta}^{blocks}$ , and  $\theta$  represents the parameters of the linear layer.

### 3.5. Semi-Supervised Fine-Tuning for Generators

Assuming that grammatical sentences are labeled as  $Y_K = 1$  and ungrammatical sentences are labeled as  $Y_K = 0$ , we extract all tokens with  $Y_K = 1$  in the dataset  $Z$  as  $Z^{fit} = \{z_1, z_2, \dots, z_m\}$ . The generator model  $G$  is again unsupervised pretrained using corpus  $Z^{fit}$ . The generator model  $G$  is then unsupervised pretrained using the corpus  $Z^{fit}$ , where  $zk = (u_1, u_2, \dots, u_n)$  represents an unsupervised token. To accomplish this, we use autoregressive language modeling to maximize the following likelihood function:

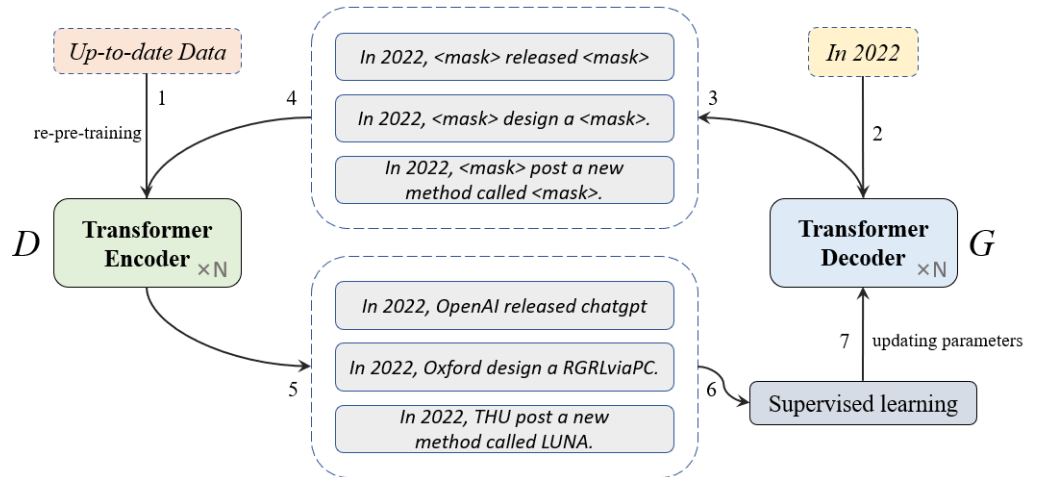
$$L_1 \left( Z^{fit} \right) = \sum_i \log G(u_i | u_{i-k}, \dots, u_{i-1}; \Theta) \tag{5}$$

where  $k$  is the size of the context window.

### 3.6. Self-Escalation

You can adjust the discriminator model  $D$  to make it perform better, and EvoText is even superior. You can also continue to pretrain the discriminator model  $D$  with new knowledge to bring it up-to-date. Certain words of text in the  $Z$  dataset are randomly masked, and the masked words are predicted using the  $D$  model, which is then input to  $G$  for supervised fine-tuning. As illustrated in Figure 2, the proposed method consists of several steps. First, the up-to-date data is retrained to update the discriminator. Second,

the generator is given a command containing a specific year (e.g., “In 2022”) to generate the data. Third, 15% of the words in the generated data is masked. Fourth, the discriminator performs word completion on the masked words. Finally, the completed data are subjected to a supervised fine-tuning of the generator, which is labeled one (by default, all statements are grammatically correct after the discriminator’s completion).



**Figure 2.** The process of learning up-to-date knowledge. The probability of being masked in this case is 15%, which is the same proportion as in the pretraining of BERT.

### 3.7. Algorithm Implementation

As shown in Algorithms 1 and 2, this is the full process of EvoText.

---

#### Algorithm 1 Priori Learning.

---

**Require:** Training samples  $Data = \{(x, y)\}$ , fine-tuning step  $N_{ep}$ , learning rate  $\tau_1$  for  $D$ , and learning rate  $\tau_2$  for  $G$ .

- 1: Initialize  $\theta$  and  $\vartheta$
  - 2: **for**  $t = 1 \cdots N_{ep}$  **do**
  - 3:   **for** minibatch  $B \subset Data$  **do**
  - 4:      $g_1 \leftarrow \mathbb{E}_B \left[ \sum_{n=1}^N \nabla_{\theta} L(D_{\theta}(x_n), y_n) \right]$
  - 5:      $\theta \leftarrow \theta - \tau_1 g_1$
  - 6:      $g_2 \leftarrow \mathbb{E}_B \left[ \sum_{n=1}^N \nabla_{\vartheta} L \left( F_{\vartheta} \left( G_{\vartheta}^{blocks}(x_n) \right), y_n \right) \right]$
  - 7:      $\vartheta \leftarrow \vartheta - \tau_2 g_2$
  - 8:   **end for**
  - 9: **end for**
-

**Algorithm 2** Training Process.

---

**Require:** Input token  $\mathbf{Z}_M^{in} = \{\mathbf{Z}_K\} = (z_k)$ , sequence maximum length  $l$ , termination character  $\alpha$ , minibatch size  $M$ , Supervised Fine-tuning learning rate  $\tau_3$ , Semi-supervised Fine-tuning learning rate  $\tau_4$ .

- 1: Initialize  $\Theta$  and  $\Phi$ .
- 2: **for**  $b = 1 \dots M$  **do**
- 3:   **for**  $t = 1 \dots l$  **do**
- 4:      $z_{k+t} \leftarrow G_{\Theta}(\mathbf{Z}_b)$
- 5:     **if**  $z_{k+t} = \alpha$  **then**
- 6:       **break**
- 7:     **end if**
- 8:      $K \leftarrow K + 1$
- 9:      $\mathbf{Z}_b \leftarrow (z_{k-t}, z_k)$
- 10:   **end for**
- 11: **end for**
- 12:  $\mathbf{Y} \leftarrow D_{\theta}(\mathbf{Z}_M^{in})$
- 13:  $\mathbf{Z}_M^{out} = \{(\mathbf{Z}_M^{in}, \mathbf{Y})\}$
- 14: **if** Supervised Fine-tuning **then**
- 15:   The cross-entropy loss function  $L$ .
- 16:    $g \leftarrow \mathbb{E}_B \left[ \sum_{n=1}^M \nabla_{\Phi} L \left( F_{\theta} \left( G_{\Theta}^{blocks}(\mathbf{Z}_n^{out}) \right), Y_n \right) \right]$
- 17:    $\Phi \leftarrow \Phi - \tau_3 g$
- 18: **end if**
- 19: **if** Semi-supervised Fine-tuning **then**
- 20:   Unsupervised loss function  $L1$ .
- 21:   **for**  $n = 1 \dots m$  **do**
- 22:     **if**  $Y_n = 1$  **then**
- 23:        $g \leftarrow \mathbb{E}_B \left[ \nabla_{\Theta} L1 \left( G_{\Theta}(\mathbf{Z}_n^{out}), \mathbf{Z}_n^{out} \right) \right]$
- 24:     **end if**
- 25:      $\Theta \leftarrow \Theta - \tau_4 g$
- 26:   **end for**
- 27: **end if**

---

**4. Experimental Setup**

In this section, we describe our experimental setup to demonstrate the performance of our algorithm.

*4.1. Experimental Environment*

We utilize a server configuration consisting of a 120-core Xeon(R) Platinum 8358P CPU @ 2.60 GHz and 8 NVIDIA A100 (80 GB) GPUs. In order to ensure optimal efficiency, we release the GPU when the model is not being trained after deploying it in a real-world scenario.

*4.2. Experimental Model*

To demonstrate the performance of EvoText, we adopted the GAN approach and selected 4 natural language understanding models and 8 natural language generation models. (These models are part of the PyTorch-Transformers library: <https://github.com/huggingface/pytorch-transformers> (accessed on 17 February 2023)).

*4.2.1. BERT*

The BERT model is widely recognized as one of the most outstanding models in recent years, having topped the GLUE tasks list when it was first released [34]. For this experiment, we exclusively utilize the large-cased version of BERT as the discriminator model. We apply the EvoText method to the fine-tuning of this model. Notably, the BERT *large-cased* model boasts 16 layers of Transformer encoders, 24 self-attentive heads,



and 330 million parameters, while the BERT<sub>base-cased</sub> model has 12 layers of Transformer encoders, 12 self-attentive heads, and 104 million parameters.

#### 4.2.2. RoBERTa

The RoBERTa model is an improved version of the BERT model that requires longer training time, a larger batch size, and more training data. Unlike BERT, RoBERTa uses dynamic masking and text encoding, moving away from BERT's NSP task. It modifies the key hyperparameters in BERT based on BERT's language masking strategy, resulting in better generalization to downstream tasks. Despite these modifications, the overall number of parameters in RoBERTa is consistent with BERT.

#### 4.2.3. GPT-2

The primary contribution of GPT-2 is its exploration of the performance of larger-scale models in ZERO-SHOT scenarios, where no fine-tuning is used. With only pretraining, hints, and predictions, GPT-2 achieved state-of-the-art results in 8 out of 9 tasks. Additionally, it is an exceptional model for natural language generation. The GPT-2<sub>small</sub> model and BERT<sub>medium</sub> have 24 and 12 layers of Transformer decoders, 24 and 12 self-attentive heads, and 335M and 124M parameters, respectively. Moreover, GPT-2 also offers larger models, such as GPT-2<sub>large</sub> with 774M parameters and GPT-2<sub>xl</sub> with 1.5B parameters.

#### 4.2.4. GPT-Neo

The GPT-Neo [35,36] 1.3B is a Transformer model trained on the Pile using cross-entropy loss. As an autoregressive language model, it learns to predict the next token in a given sequence of English text, thereby capturing internal representations of English. These representations can then be used to extract features that are useful for downstream tasks. Although language models have many applications beyond this, there are still many unknowns in this area of research.

#### 4.2.5. OPT

The OPT model [37] is primarily pretrained using a causal language model (CLM) target, which belongs to the family of GPT-3 models. The pretrained model can be used to evaluate prompts and generate text for downstream tasks. Additionally, the model can be fine-tuned on downstream tasks using CLM instances. The experiments in this paper use models with 125M and 350M parameters.

#### 4.2.6. Transformer-XL

The Transformer-XL model [38] introduces two innovations to the Vanilla Transformer, a recurrence mechanism and relative positional coding. An additional advantage of Transformer-XL over Vanilla Transformer is that it can be used for word-level and character-level language modeling. It achieves state-of-the-art language modeling results on several different datasets and combines a circularity mechanism with an attention mechanism that allows the model to learn long-term dependencies.

#### 4.2.7. Language Models with Pre-trained Word Embeddings and without Pre-Trained Word Embeddings

In addition to attention-based models, pretrained word embedding models such as Word2Vec [39] or Glove [40] can also yield good results when incorporated into the word embedding layer. Similarly, scratch-trained word embedding layers can be effective for specific tasks, such as hate detection or text toxicity detection [41–46]. In this paper, we evaluate and compare various pretrained and scratch-trained word embedding models as discriminators to assess their impact on the overall training system.

#### 4.3. Dataset

In the subsequent experiments, we chose an a priori dataset for discriminators.

#### 4.3.1. CoLA

The CoLA (corpus of linguistic acceptability) [47] consists of 10,657 sentences from 23 linguistic publications, professionally annotated for acceptability (grammaticality) by their original authors. The public version presented here contains 9594 sentences from the training and development sets, excluding 1063 sentences from the retention test set. The goal is to classify the sentences into either acceptable or unacceptable categories based on their grammaticality. Due to its carefully curated and annotated nature, CoLA is a valuable resource for evaluating the performance of various NLP models and techniques in the domain of language understanding.

#### 4.3.2. LAMBADA

The source of the corpus constructed by LAMBADA [48] is unpublished anthologies. The rationale is to minimize the influence of generic knowledge on the answers, i.e., it is difficult for the model to derive answers from generic knowledge. It consists of 5325 novels and 465 million words. LAMBADA has been widely used for language generation tasks and language understanding tasks, such as language modeling and text comprehension, where the goal is to predict the next word in a given sentence based on the preceding context.

#### 4.3.3. CBT

The children's book test (CBT) aims to directly measure the extent to which language models exploit the wider language environment. The CBT is built from books that are freely accessible. The CBT has been widely used for evaluating the performance of various NLP models and techniques in the domain of language understanding and generation.

#### 4.3.4. WikiText

The WikiText [49] dataset is a large-scale language modeling dataset that is widely used in natural language processing research. It is created by extracting articles from the English Wikipedia and is available in three versions: WikiText-2, WikiText-103, and WikiText-500k. The dataset includes articles covering a wide range of topics, providing a diverse range of text for training and evaluation. The WikiText dataset has been used in various language modeling tasks, including next word prediction, text generation, and text classification. It is a valuable resource for training and evaluating natural language processing models, and its use has contributed significantly to the development of language modeling research.

#### 4.3.5. PTB

The PTB (penn treebank dataset) [50] contains 42,000, 3000, and 3000 English sentences for the training set, validation set, and test set. "<eos>" is the start signal of each sentence, and "<eos>" is the end signal of each sentence. The dataset is annotated with part-of-speech tags, constituency parse trees, and semantic roles, providing rich linguistic annotations for various natural language processing tasks. The PTB has been used in a wide range of natural language processing tasks, including language modeling, part-of-speech tagging, named entity recognition, parsing, and machine translation.

#### 4.3.6. enwiki8 and text8

The text8 (the dataset is available for download at <https://huggingface.co/datasets/enwik8> (accessed on 17 February 2023)) comes from enwiki8 (the dataset is available for download at <http://mattmahoney.NET/dc/text8.zip> (accessed on 17 February 2023)), which was first used to conduct text compression. Simply put, enwiki8 is the first 100,000,000 characters picked up from Wikipedia; and text8 is the result of removing all kinds of strange symbols and non-English characters from these characters, then converting uppercase characters into lowercase characters and transforming numbers into the corresponding English words. This dataset aims to learn distributed representations of words that capture their semantic and

syntactic relationships, and it has been used in various natural language processing tasks, including language modeling, text generation, and word embeddings.

#### 4.3.7. 1BM

The 1BW (one billion word) [51] dataset is a large English language corpus used for pretraining language models. It contains one billion words and is freely available for research purposes. This benchmark dataset is widely used for evaluating the performance of statistical language models and is composed of various genres and topics, including news, technology, and novels. It was proposed by the Google Brain team and is considered a standard for measuring progress in the field of natural language processing. The 1BW dataset has been used for pretraining language models to improve their performance on downstream NLP tasks, such as text classification, sentiment analysis, and language generation.

#### 4.4. Model Evaluation Indicators

We use PPL (perplexity), ACC (accuracy), and BPC (bits-per-character) as performance metrics for our experiments. PPL measures the average number of choices available to the model when predicting the next word in a sentence and is calculated using the following formula:

$$\text{PPL}(S) = \sqrt[m]{\frac{1}{p(w_1, w_2, \dots, w_m)}} = \sqrt[m]{\prod_{i=1}^m \frac{1}{p(w_i | w_1, \dots, w_{i-1})}} \quad (6)$$

where  $S$  is the sentence being evaluated,  $m$  is the length of the sentence, and  $p(w_i | w_1, \dots, w_{i-1})$  is the probability of the  $i$ -th word given the preceding words in the sentence. A lower PPL value indicates better model performance.

ACC measures the percentage of correct judgments out of all judgment cases and is calculated using the following formula:

$$\text{ACC} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}} \quad (7)$$

where TP (true positive) is the number of cases correctly judged as positive, TN (true negative) is the number of cases correctly judged as negative, FP (false positive) is the number of cases incorrectly judged as positive, and FN (false negative) is the number of cases incorrectly judged as negative.

In this work, accuracy refers to the percentage of correctly predicted tokens in the test dataset. In other words, it measures how often the model predicted the correct next word given the previous words in the sentence. This metric is commonly used to evaluate the performance of language models.

BPC measures the number of bits required on average to encode each character in the text and is calculated using the following formula:

$$\text{BPC} = -\frac{1}{m} \sum_{i=1}^m \log_2 p(w_i | w_1, \dots, w_{i-1}) \quad (8)$$

where  $m$  is the length of the text, and  $\log_2 p(w_i | w_1, \dots, w_{i-1})$  is the number of bits required to encode the  $i$ -th character given the preceding characters in the text. A lower BPC value indicates better model performance.

Specifically, BPC measures the number of bits needed to encode each character in the text. Lower BPC values indicate better compression, which in turn indicates that the model has learned to better capture the patterns and structure of the text.

## 5. Experimental Procedure

This section shows the parameters that need to be tuned in the actual training and the comparison with other models.

### 5.1. Data Preprocessing

In this paper, we preprocessed the data using common techniques such as regular expression substitution and expanding English abbreviations. Table 1 shows the details of the preprocessing steps.

**Table 1.** Example of data preprocessing for the training set in downstream tasks.

Operation	Text
Original Text	<p>our friends <u>won't</u> buy this analysis, let alone the next one we propose.            one more pseudo generalization and <u>i'm</u> giving up .            one more pseudo generalization or <u>i'm</u> giving up .            i <u>ll</u> fix you a drink .  <u>we're</u> dancing the night away .            My little #ObsessedWith MyDog @Cafe Solstice Capitol Hill            More #tinyepic things #tinyepicwestern , this one is crazy @user I may be one of your . .            Last night @Omnia Night Club At Caesars Palace            friendship at its finest. ....#pixar #toystory #buzz #woody #friends #friendship #bff...            I L VE working for a <u>cause! Yesterday's</u> balloon decor for SNN 11th Annual Back 2 School Health ...</p>
After preprocessing	<p>our friends <u>will not</u> buy this analysis let alone the next one we propose            one more pseudo generalization and <u>i am</u> giving up            one more pseudo generalization or <u>i am</u> giving up            i <u>will</u> fix you a drink  <u>we are</u> dancing the night away            My little ObsessedWithMyDog <u>Cafe</u> Solstice Capitol Hill            More tinyepic things <u>tinyepicwestern</u> this one is crazy <u>user</u> I may be one of your            Last night <u>Omnia</u> Night Club At Caesars Palace            friendship at its finest <u>pixar toystory buzz woody friends friendship bff</u>            I L VE working for a <u>cause Yesterdays</u> balloon decor for SNN 11th Annual Back 2 School Health</p>

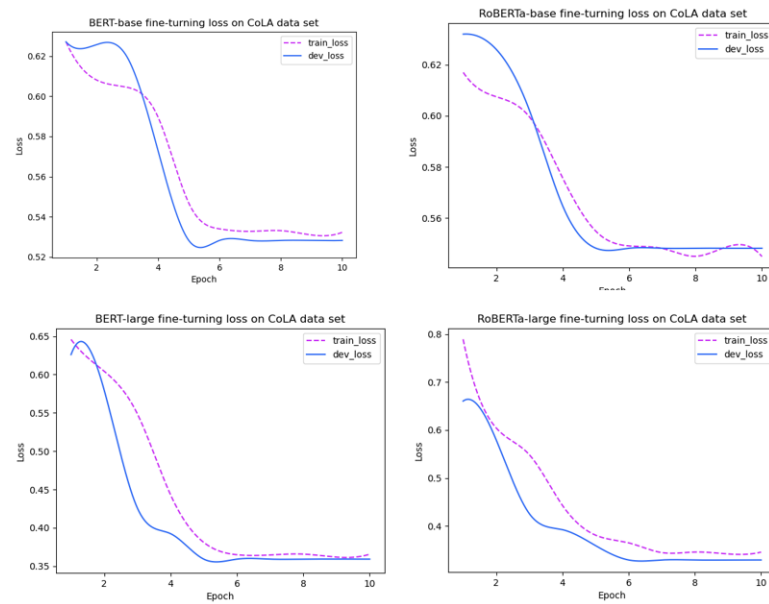
### 5.2. Fine-Tuning of Discriminators For Priori Datasets

In this paper, we employed BERT<sub>large</sub>, BERT<sub>base</sub>, RoBERTa<sub>large</sub>, or RoBERTa<sub>base</sub> as the discriminator model. Since these models are pretrained, they need to be fine-tuned to achieve optimal results in downstream tasks. During the fine-tuning process, it is recommended to use a lower learning rate and a smaller number of epochs to update the model. This is because using large learning rates and epochs may cause the model to fail to converge or overfit, which can negatively impact the model's performance in this task.

**Results** As illustrated in Figure 3, we observed that the large model performed better on the CoLA task, with RoBERTa exhibiting the lowest loss rate. During the pretraining process, we set the maximum length of the tokenizer to 45, enabled padding, and used a minibatch size of 512. We fine-tuned the model using several of the most commonly used parameter settings. As summarized in Table 2, we achieved the best results with a learning rate of  $3 \times 10^{-5}$  and 10 epochs. Following the fine-tuning process, the RoBERTa<sub>large</sub> model demonstrated the ability to make judgments about grammatical plausibility.

**Table 2.** Setting different parameters allows the RoBERTa<sub>large</sub> model to fine-tune the loss rate of the validation set in the CoLA task.

		Learning Rate		
		$2 \times 10^{-5}$	$3 \times 10^{-5}$	$4 \times 10^{-5}$
epoch	5	0.378204	0.392302	0.425639
	10	0.360829	<b>0.359063</b>	0.370926
	15	0.363742	0.362149	0.360642



**Figure 3.** BERT<sub>large</sub>, BERT<sub>base</sub>, RoBERTa<sub>large</sub>, and RoBERTa<sub>base</sub> model fine-tuned loss rate in CoLA dataset.

### 5.3. Prewarm-Up Training of Generator

During pretraining of GPT-2<sub>medium</sub>, we followed the same data preprocessing steps as before, with one exception: GPT’s tokenizer does not auto-pad sentences to the maximum length. Therefore, we used the special token “<|endoftext|>” to pad sentences that were not long enough. Unlike the input to the BERT<sub>large</sub> model, the generator model used in GPT-2<sub>medium</sub> is an autoregressive language model that requires mask attention to model the data. As such, we needed to provide masks in the input to the GPT-2<sub>medium</sub> model to ensure optimal performance.

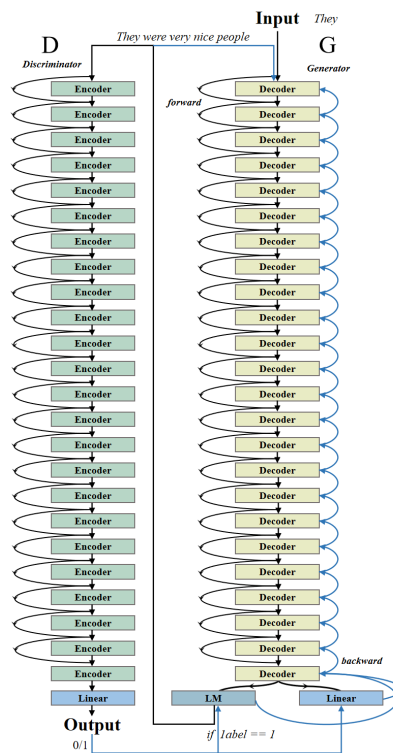
As shown in Table 3, the best performance was achieved with a learning rate of  $1 \times 10^{-2}$  and 10 epochs. During training, we only updated the parameters of the last linear layer, which allowed the model to be easily fine-tuned for supervised tasks and prevented extensive updates to the linear layer parameters during fine-tuning.

**Table 3.** With all Transformer block parameters of GPT-2 frozen, only the linear layer is trained to validate the results of the set on the CoLA task.

Model Input	Loss ↓								
	lr = $5 \times 10^{-3}$	lr = $5 \times 10^{-3}$	lr = $5 \times 10^{-3}$	lr = $1 \times 10^{-2}$	lr = $1 \times 10^{-2}$	lr = $1 \times 10^{-2}$	lr = $2 \times 10^{-2}$	lr = $2 \times 10^{-2}$	lr = $2 \times 10^{-2}$
	Epoch = 5	Epoch = 10	Epoch = 15	Epoch = 5	Epoch = 10	Epoch = 15	Epoch = 5	Epoch = 10	Epoch = 15
Token	2.3275	2.3275	2.3275	2.3275	2.3275	2.3275	2.3275	2.3275	2.3275
Token+mask attention	0.6910	0.6825	0.6722	0.6810	<b>0.6721</b>	0.6730	0.6754	0.6772	0.6772

### 5.4. Training Process

As illustrated in Figure 4, our framework involves training the discriminator model BERT<sub>large</sub> and the generator model GPT-2<sub>medium</sub> in a training loop. Firstly, we identify common text word-initial words and use GPT-2<sub>medium</sub> to complete the sentences. Subsequently, the completed sentences are fed into the fine-tuned BERT<sub>large</sub> model to evaluate their grammatical plausibility. This evaluation result is then utilized to conduct supervised fine-tuning of the GPT-2<sub>medium</sub> model. To avoid the discriminator model’s errors significantly affecting the generator model, we adopt a minimal learning rate and train the generator model only one round.




**Figure 4.** Training process of discriminator ( $BERT_{large}$ ) and generator ( $GPT-2_{medium}$ ) in EvoText. The black line indicates the forward propagation process, and the blue line indicates the reverse fine-tuning process after acquiring the label.

Table 4 presents some examples of text generated by  $GPT-2_{medium}$ . Subsequently, we fed this data to a discriminator model to assess their syntactic plausibility. The discriminator model’s output value of 1 or 0 indicates whether the modified sentence is grammatically valid or invalid, respectively. Next, we used these labeled data to perform supervised fine-tuning of the  $GPT-2_{medium}$  model. Figure 5 shows that we used a learning rate of  $1 \times 10^{-4}$  and a minibatch size of 64 for fine-tuning.

**Table 4.** We did not fine-tune the text generated by the previous  $GPT-2_{medium}$  model. Then, it is input to the discriminator for judgment. The red marker denotes the data generated by the generator, while a D Output of 0 indicates that the statement is not grammatically correct, and 1 indicates that it is grammatically correct.

G Input	G Output	D Output
That	That doesn’t have any significance, right?	0
It	It was a beautiful night of sunshine with some gorgeous light falling.	1
He	He said: ""[W]ith this being an attack of our religion I do feel the time will not have arrived.""	0
We	We have already started the implementation phase and will keep the project in mind throughout.	1
I	I’ve done all those jobs.	1



**Figure 5.** Supervised fine-tuning of the training loss of the  $GPT-2_{medium}$  model.

### 5.5. Semi-Supervised Fine-Tuning Generator Model

Based on the discriminator output presented in Table 4, the sentences labeled with a value of 1 are fed back into the generator model for further pretraining. This process aims to enhance the generator’s ability to produce grammatically correct text.

Table 5 illustrates the evaluation of 10 tasks using 4 different natural language understanding models. The results indicate that RoBERTa<sub>large</sub> outperforms the other models. Thus, we selected RoBERTa<sub>large</sub> as the discriminator for subsequent experiments. As demonstrated in Table 6, utilizing pretrained word embeddings as a replacement for the model’s word embedding layer or reinitializing the parameters to train from scratch resulted in inferior performance compared to the original RoBERTa<sub>large</sub> model.

**Table 5.** After 156 minibatch-sized EvoText sessions, we evaluated the performance of various natural language understanding models (D) and the same language generation model (GPT<sub>medium</sub>) on 7 natural language processing tasks ZERO-SHOT. Each model was tested five times on each task, and the results were averaged.

Model	LAMBADA		CBT		WikiText		PTB	enwiki8	text8	1BW
	PPL ↓	ACC ↑	ACC ↑ (CN)	ACC ↑ (NE)	PPL ↓ (WikiText2)	PPL ↓ (WikiText103)	PPL ↓	BPC ↓	BPC ↓	PPL ↓
<b>Baseline</b>										
GPT-2 <sub>medium</sub>	15.60	55.48	92.35	87.10	22.76	26.37	47.33	1.01	1.06	55.72
<b>Ours</b>										
GPT-2 <sub>medium</sub> (G)+ RoBERTa <sub>large</sub> (D)	<b>14.21</b>	<b>57.20</b>	<b>92.85</b>	87.00	<b>21.94</b>	<b>25.12</b>	<b>45.93</b>	<b>1.00</b>	<b>1.06</b>	<b>55.01</b>
RoBERTa <sub>base</sub> (D)	14.91	56.21	92.84	87.15	22.40	25.49	45.99	1.01	1.07	55.82
BERT <sub>large</sub> (D)	14.28	57.14	92.83	<b>87.32</b>	21.97	<b>25.12</b>	45.94	1.00	1.06	55.03
BERT <sub>base</sub> (D)	14.92	56.17	92.81	87.02	21.44	25.32	45.98	1.01	1.07	55.80

**Table 6.** After 156 minibatch-sized EvoText sessions, we evaluated the performance of various natural language understanding models (D) with or without pretrained word embeddings and the same language generation model (GPT<sub>medium</sub>) on 7 natural language processing tasks using ZERO-SHOT. Each model was tested five times on each task, and the results were averaged.

Model	LAMBADA		CBT		WikiText		PTB	enwiki8	text8	1BW
	PPL ↓	ACC ↑	ACC ↑ (CN)	ACC ↑ (NE)	PPL ↓ (WikiText2)	PPL ↓ (WikiText103)	PPL ↓	BPC ↓	BPC ↓	PPL ↓
<b>Baseline</b>										
GPT-2 <sub>medium</sub>	15.60	55.48	92.35	87.10	22.76	26.37	47.33	1.01	1.06	55.72
<b>Ours</b>										
GPT-2 <sub>medium</sub> (G)+ RoBERTa <sub>large</sub> (D)	<b>14.21</b>	<b>57.20</b>	<b>92.85</b>	87.00	21.94	<b>25.12</b>	<b>45.93</b>	<b>1.00</b>	<b>1.06</b>	<b>55.01</b>
RoBERTa <sub>large</sub> + Word2Vec	15.19	57.13	92.85	87.03	23.01	27.72	46.91	1.03	1.06	55.61
RoBERTa <sub>large</sub> + Glove	15.29	57.03	92.79	<b>87.13</b>	23.11	27.22	46.93	1.04	1.06	55.10
RoBERTa <sub>large</sub> + Scratch-trained	15.01	56.13	92.79	87.03	<b>21.81</b>	26.42	46.01	1.04	1.06	55.10

### 5.6. Experimental Results

After a training process consisting of 156 minibatch iterations, which represents the average of the LAMBADA, CBT, WikiText, PTB, enwiki8, text8, and 1BW dataset sizes, we evaluated the performance of six natural language generation models on various datasets, including LAMBADA, CBT, WikiText, PTB, enwiki8, text8, and 1BW.

Table 7 shows that EvoText can steadily improve the performance of eight natural language generation models. Notably, the training process with just 156 steps can produce better results without altering the model architecture or the initial pretraining method. It is surprising to see that EvoText improves the performance of the GPT<sub>small</sub> model to surpass that of the OPT<sub>125M</sub> model. These results indicate that EvoText can significantly enhance the performance of the model without requiring extensive modifications. Our approach demonstrates favorable performance compared to the current state-of-the-art

RLHF in terms of Chatgpt feedback on nearly every task. Based on the results presented in Table 8, it can be observed that the EvoText training approach is highly effective in rectifying a significant portion of the grammatical errors produced by the model, which is an impressive outcome.

**Table 7.** ZERO-SHOT performance of the model on 7 natural language processing tasks after 156 minibatch-sized EvoText sessions. Each model was run five times on each task and averaged.

Model	LAMBADA		CBT		WikiText		PTB	enwiki8	text8	1BW
	PPL ↓	ACC ↑	ACC ↑ (CN)	ACC ↑ (NE)	PPL ↓ (WikiText2)	PPL ↓ (WikiText103)	PPL ↓	BPC ↓	BPC ↓	PPL ↓
<b>Baseline</b>										
GPT-2 <sub>small</sub>	35.13	45.99	87.65	83.4	29.41	37.50	65.85	1.16	1.17	75.20
GPT-2 <sub>medium</sub>	15.60	55.48	92.35	87.1	22.76	26.37	47.33	1.01	1.06	55.72
GPT-2 <sub>large</sub>	10.87	60.12	93.45	88.0	19.93	22.05	40.31	0.97	1.02	44.575
GPT-2 <sub>xl</sub>	8.63	63.24	93.30	89.05	18.34	17.48	35.76	0.93	0.98	42.16
GPT-Neo1.3B	7.50	57.2	-	-	-	13.10	-	-	-	-
OPT <sub>125M</sub>	32.93	47.2	88.24	86.81	28.14	38.23	60.15	1.17	1.16	72.89
OPT <sub>350M</sub>	13.29	56.99	92.25	87.31	20.99	25.01	<b>46.23</b>	1.00	1.06	50.72
Transformer-XL	-	-	-	-	-	-	54.5	0.99	1.08	-
<b>RLHF by Chatgpt</b>										
GPT-2 <sub>small</sub>	30.43	45.79	88.36	84.44	29.20	37.48	64.91	1.16	1.16	74.99
GPT-2 <sub>medium</sub>	15.29	55.01	92.41	<b>87.80</b>	22.66	26.28	47.30	1.01	1.06	55.69
GPT-2 <sub>large</sub>	10.88	60.13	93.39	88.09	<b>19.23</b>	<b>22.05</b>	39.91	0.97	1.02	44.50
GPT-2 <sub>xl</sub>	8.57	<b>63.24</b>	93.87	89.07	18.21	<b>17.47</b>	35.75	0.93	0.98	42.11
GPT-Neo1.3B	7.52	57.18	-	-	15.37	13.00	-	-	-	-
OPT <sub>125M</sub>	30.71	47.29	88.17	<b>86.99</b>	28.12	37.03	60.24	1.16	1.16	70.99
OPT <sub>350M</sub>	13.01	57.09	92.55	<b>87.43</b>	<b>20.81</b>	25.01	46.28	1.00	1.06	50.54
Transformer-XL	-	-	-	-	-	-	54.25	0.99	1.08	-
<b>Ours(D is RoBERTa<sub>large</sub>)</b>										
EvoText GPT-2 <sub>small</sub>	<b>28.42</b>	<b>46.29</b>	<b>89.23</b>	<b>84.48</b>	<b>28.31</b>	<b>35.72</b>	<b>62.48</b>	<b>1.14</b>	<b>1.16</b>	<b>73.82</b>
EvoText GPT-2 <sub>medium</sub>	<b>14.21</b>	<b>57.20</b>	<b>92.85</b>	87.0	<b>21.94</b>	<b>25.12</b>	<b>45.93</b>	<b>1.00</b>	<b>1.06</b>	<b>55.01</b>
EvoText GPT-2 <sub>large</sub>	<b>10.57</b>	<b>60.14</b>	<b>93.80</b>	<b>88.32</b>	19.93	22.06	<b>38.20</b>	<b>0.96</b>	<b>1.02</b>	<b>44.41</b>
EvoText GPT-2 <sub>xl</sub>	<b>8.09</b>	63.23	<b>93.90</b>	<b>89.15</b>	17.92	17.50	<b>34.91</b>	<b>0.92</b>	<b>0.98</b>	<b>42.09</b>
EvoText GPT-Neo1.3B	<b>7.41</b>	<b>57.19</b>	-	-	15.29	<b>12.81</b>	-	-	-	-
EvoText OPT <sub>125M</sub>	<b>27.51</b>	<b>50.09</b>	<b>91.01</b>	85.24	<b>28.00</b>	<b>35.89</b>	<b>55.99</b>	<b>1.14</b>	<b>1.15</b>	<b>70.21</b>
EvoText OPT <sub>350M</sub>	<b>12.10</b>	<b>57.41</b>	<b>92.69</b>	<b>88.91</b>	21.99	<b>25.00</b>	46.33	<b>1.00</b>	<b>1.06</b>	<b>50.02</b>
EvoText	-	-	-	-	-	-	<b>53.21</b>	<b>0.99</b>	<b>1.07</b>	-
Transformer-XL	-	-	-	-	-	-	-	-	-	-

**Table 8.** We entered “Once upon a time” into Baseline’s GPT-2<sub>xl</sub> and EvoText GPT-2<sub>xl</sub>, respectively, for comparison.

<p><b>Generation by Baseline GPT-2<sub>xl</sub></b>  Once upon a time, girl name is Lisa.  Lisa is like to go on walk in park, but yesterday, she goes on walk and she lost.  She asks help from man which he see, but man doesn’t speak English.  She feels very scared and doesn’t know how to come back to her house.  Suddenly, she saw a police car and she run to them.  Police helped her and she come back to her house safely.</p>
<p><b>Generation by EvoText GPT-2<sub>xl</sub></b>  Once upon a time, <b>there was a girl named Lisa.</b>  Lisa <b>enjoyed going for walks in the park,</b> but yesterday, she got lost <b>while</b> on a walk.  She asked <b>for</b> help from a man she <b>saw,</b> but he <b>didn’t</b> speak English.  She <b>felt very scared</b> and didn’t know how to get back home.  Suddenly, she saw a police car and <b>ran towards them.</b>  <b>The</b> police officers <b>helped</b> her and she was able to return home safely.</p>

### 5.7. Up-to-Data Knowledge Update

We collected abstracts of preprints published on arXiv from June to September 2022 as the most up-to-date knowledge dataset. We partitioned the dataset into a training dataset, validation dataset, and testing dataset with a split ratio of 7:1:2. In the conventional methodology, the generator model is directly subjected to fine-tuning. In contrast, our methodology as shown in Section 3.6 entails solely fine-tuning the discriminator model, followed by EvoText training. To ensure equitable experimental outcomes, we employ the identical epoch across all trials.



Table 9 demonstrates that the up-to-date knowledge update generator model, implemented with EvoText’s approach, outperforms retraining the generator model while maintaining its performance on the ZERO-SHOT task. However, the scalability of EvoText on larger datasets and more complex natural language processing tasks is not discussed in this paper. Nevertheless, based on the results presented in Table 9, it can be observed that the model not only acquires new knowledge but also avoids catastrophic forgetting of the original knowledge, which is promising for future research on scalability.

**Table 9.** We evaluated the performance of both models on the arXiv dataset by retraining the discriminator at a learning rate of  $1 \times 10^{-4}$  and the generator at a learning rate of  $5 \times 10^{-5}$ .

Model	PPL↓ (arXiv)	PPL↓ (LAMBADA) (ZERO-SHOT)
<b>Baseline</b>		
GPT-2 <sub>small</sub>	37.14 (ZERO-SHOT)	35.13
OPT <sub>125M</sub>	36.80 (ZERO-SHOT)	32.93
<b>Traditional</b>		
GPT-2 <sub>small</sub>	20.42	<b>28.10</b>
OPT <sub>125M</sub>	20.13	28.01
<b>Ours</b>		
EvoText GPT-2 <sub>small</sub>	<b>19.81</b>	28.12
EvoText OPT <sub>125M</sub>	<b>18.94</b>	<b>28.00</b>

### 5.8. Ablation Experiments

To investigate the effect of each module in EvoText, we performed ablation experiments by removing them one by one. These modules include the fine-tuning discriminator model, the prewarm-up generator, and the supervised and semi-supervised fine-tuning generator models.

Based on Table 10, it can be observed that each module of EvoText has an impact on the final results. Removing any of these modules may cause some negative impact on the overall performance. It is worth noting that the supervised learning module is deemed necessary in our approach.

**Table 10.** We evaluated the impact of removing a particular module of EvoText GPT-2<sub>media</sub> on ZERO-SHOT performance in seven tasks.

Model	LAMBADA		CBT		WikiText		PTB	enwiki8	text8	1BW
	PPL ↓	ACC ↑	ACC ↑ (CN)	ACC ↑ (NE)	PPL ↓ (WikiText2)	PPL ↓ (WikiText103)	PPL ↓	BPC ↓	BPC ↓	PPL ↓
<b>Init</b>										
Full module EvoText	<b>14.21</b>	<b>57.20</b>	<b>92.85</b>	<b>87.0</b>	<b>21.94</b>	<b>25.12</b>	<b>45.93</b>	<b>1.00</b>	<b>1.06</b>	<b>55.01</b>
<b>Ablation Experiments</b>										
Remove D pretraining	14.80	55.94	92.95	86.53	22.46	27.54	46.99	1.01	1.06	55.72
Remove G	14.34	57.03	92.90	86.93	22.12	25.54	45.99	1.00	1.06	55.11
Pre-warm-up Training										
Remove supervised fine-tuning	15.12	55.62	92.89	87.23	22.70	27.64	47.90	1.01	1.06	56.11
Remove semi-supervised fine-tuning	14.42	57.04	92.90	86.23	22.08	25.69	46.70	1.00	1.06	55.88

## 6. Conclusions and Future Work

In this article, we introduced EvoText, a training process for two pretrained models aimed at addressing the challenges of insufficient sample data and computational resources, allowing models to continue learning after deployment. Through fine-tuning discriminators and prewarm-up training generators, we achieved better model performance with just 156 training steps, significantly improving performance without requiring additional training data. This approach steadily improves the performance of natural language understanding and generation tasks without changing the model structure, with the potential for even greater performance improvements over time. EvoText is an effective and scalable training model that holds great promise for low-resource NLP tasks. Our extensive

experiments demonstrate the potential for improving pretrained model performance and highlight the importance of supervised learning.

Future research directions may include exploring the potential for EvoText in other NLP tasks and applications, investigating the impact of different discriminator and generator architectures on model performance, and further exploring the potential for continued learning after deployment in other settings. Additionally, our study highlights the importance of supervised learning in NLP and suggests that future research should continue to focus on developing effective training processes for pretrained models in low-resource settings.

**Author Contributions:** Conceptualization, methodology, formal analysis, and software, Z.Y.; verification, Z.Y., C.Z. and H.X.; writing—original draft preparation, Z.Y., C.Z. and H.X.; writing—review and editing, Y.L.; visualization, C.Z. and H.X.; supervision, Y.L.; funding acquisition, Y.L. All authors have read and agreed to the published version of the manuscript.

**Funding:** The authors gratefully acknowledge the support of the AIMTEEL 202201 Open Fund for Intelligent Mining Technology and Equipment Engineering Laboratory in Anhui Province and the Anhui Provincial Department of Education Scientific Research Key Project (Grant No. 2022AH050995). The financial assistance provided by these projects was instrumental in carrying out the research presented in this paper. We would like to thank all the members of the laboratory for their valuable support and assistance. Without their help, this research would not have been possible. Finally, we would like to express our gratitude to the Anhui Polytechnic University for providing the necessary facilities and resources for this study.

**Institutional Review Board Statement:** The study did not require ethical approval.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** All data generated by the generators in this article are placed in: <https://github.com/DLYuanGod/Auto-learning/blob/main/Gen.csv> (accessed on 9 March 2023).

**Conflicts of Interest:** The authors declare no conflict of interest.

## References

1. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, L.; Polosukhin, I. Attention is all you need. In *Advances in Neural Information Processing Systems*; Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2017; Volume 30. Available online: <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf> (accessed on 17 February 2023).
2. Huang, S.; Renals, S. Hierarchical bayesian language models for conversational speech recognition. *IEEE Trans. Audio Speech Lang. Process.* **2020**, *18*, 1941–1954. [[CrossRef](#)]
3. Wang, C.; Dai, S.; Wang, Y.; Yang, F.; Qiu, M.; Chen, K.; Zhou, W.; Huang, J. Arobert: An asr robust pre-trained language model for spoken language understanding. *IEEE/ACM Trans. Audio Speech Lang. Process.* **2022**, *30*, 1207–1218. [[CrossRef](#)]
4. Yu, F.-H.; Chen, K.-Y.; Lu, K.-H. Non-autoregressive asr modeling using pre-trained language models for chinese speech recognition. *IEEE/ACM Trans. Audio Speech Lang. Process.* **2022**, *30*, 1474–1482. [[CrossRef](#)]
5. Peters, M.E.; Neumann, M.; Iyyer, M.; Gardner, M.; Clark, C.; Lee, K.; Zettlemoyer, L. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, New Orleans, LA, USA, 1–6 June 2018; pp. 2227–2237. Available online: <https://aclanthology.org/N18-1202> (accessed on 17 February 2023).
6. Radford, A.; Narasimhan, K. Improving Language Understanding by Generative Pre-Training. *J. Mach. Learn. Res.* **2018**, *23*, 1–18.
7. Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; Liu, P.J. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.* **2020**, *21*, 1–67. Available online: <http://jmlr.org/papers/v21/20-074.html> (accessed on 17 February 2023).
8. Devlin, J.; Chang, M.-W.; Lee, K.; Toutanova, K. 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)*, Minneapolis, MN, USA, 2–7 June 2019; pp. 4171–4186. Available online: <https://aclanthology.org/N19-1423> (accessed on 17 February 2023).
9. Zhuang, L.; Wayne, L.; Ya, S.; Jun, Z. A robustly optimized BERT pre-training approach with post-training. In *Proceedings of the 20th Chinese National Conference on Computational Linguistics, Huhhot, China, 13–15 August 2021*; pp. 1218–1227. Available online: <https://aclanthology.org/2021.ccl-1.108> (accessed on 17 February 2023).
10. Kipf, T.N.; Welling, M. Variational graph auto-encoders. In *Proceedings of the NIPS Workshop on Bayesian Deep Learning*, Barcelona, Spain, 9–10 December 2016.

11. Goodfellow, I.J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; Bengio, Y. Generative training nets. In Proceedings of the 27th International Conference on Neural Information Processing Systems, Montreal, QC, Canada, 8–13 December 2014; Volume 2; pp. 2672–2680.
12. Ho, J.; Jain, A.; Abbeel, P. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*; Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2020; Volume 33, pp. 6840–6851. Available online: <https://proceedings.neurips.cc/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf> (accessed on 17 February 2023).
13. Brown, T.B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language Models are Few-Shot Learners. In Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS), Virtual, 6–12 December 2020.
14. Christiano, P.F.; Shah, Z.; Mordatch, I.; Schneider, J.; Blackwell, T.; Tobin, J.; Abbeel, P.; Zaremba, W. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems*; Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2017; Volume 30, pp. 4302–4310.
15. Shannon, M.; Zen, H.; Byrne, W. Autoregressive models for statistical parametric speech synthesis. *IEEE Trans. Audio Speech Lang. Process.* **2013**, *21*, 587–597. [[CrossRef](#)]
16. Niedz, M.; Ciolek, M.; Cisowski, K. Elimination of impulsive disturbances from stereo audio recordings using vector autoregressive modeling and variable-order kalman filtering. *IEEE/ACM Trans. Audio Speech Lang. Process.* **2015**, *23*, 970–981.
17. Deng, F.; Bao, C.; Kleijn, W.B. Sparse hidden markov models for speech enhancement in non-stationary noise environments. *IEEE/ACM Trans. Audio Speech Lang. Process.* **2015**, *23*, 1973–1987. [[CrossRef](#)]
18. Özdemir, O.; Kerzel, M.; Weber, C.; Lee, J.H.; Wermter, S. Language model-based paired variational autoencoders for robotic language learning. *IEEE Trans. Cogn. Dev. Syst.* **2022**. (accessed on 17 February 2023). [[CrossRef](#)]
19. Zhang, Z.; Wu, Y.; Zhou, J.; Duan, S.; Zhao, H.; Wang, R. Sg-net: Syntax guided transformer for language representation. *IEEE Trans. Pattern Anal. Mach. Intell.* **2022**, *44*, 3285–3299. [[CrossRef](#)]
20. Choudhary, T.; Goyal, V.; Bansal, A. Wtasr: Wavelet transformer for automatic speech recognition of indian languages. *Big Data Min. Anal.* **2023**, *6*, 85–91. [[CrossRef](#)]
21. Aloysius, N.; Nedungadi, G.M.P. Incorporating relative position information in transformer-based sign language recognition and translation. *IEEE Access* **2021**, *9*, 145929–145942. [[CrossRef](#)]
22. Li, Z.; Li, Z.; Zhang, J.; Feng, Y.; Zhou, J. Bridging text and video: A universal multimodal transformer for audio-visual scene-aware dialog. *IEEE/ACM Trans. Audio, Speech, Lang. Process.* **2021**, *29*, 2476–2483. [[CrossRef](#)]
23. Qi, Q.; Lin, L.; Zhang, R.; Xue, C. Medt: Using multimodal encoding-decoding network as in transformer for multimodal sentiment analysis. *IEEE Access* **2022**, *10*, 28750–28759. [[CrossRef](#)]
24. Liu, P.J.; Saleh, M.; Pot, E.; Goodrich, B.; Sepassi, R.; Kaiser, L.; Shazeer, N. Generating wikipedia by summarizing long sequences. *Int. Conf. Learn. Represent.* **2018**. Available online: <https://openreview.net/forum?id=Hyg0vbWC-> (accessed on 17 February 2023).
25. Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I. Language models are unsupervised multitask learners. *OpenAI Blog* **2019**, *1*, 9.
26. Bai, Y.; Yi, J.; Tao, J.; Tian, Z.; Wen, Z.; Zhang, S. Fast end-to-end speech recognition via non-autoregressive models and cross-modal knowledge transferring from bert. *IEEE/ACM Trans. Audio Speech Lang. Process.* **2021**, *29*, 1897–1911. [[CrossRef](#)]
27. Lan, Z.; Chen, M.; Goodman, S.; Gimpel, K.; Sharma, P.; Soricut, R. Albert: A lite bert for self-supervised learning of language representations. *Int. Conf. Learn. Represent.* **2020**. Available online: <https://openreview.net/forum?id=H1eA7AEtV5> (accessed on 17 February 2023).
28. Yang, Z.; Dai, Z.; Yang, Y.; Carbonell, J.; Salakhutdinov, R.; Le, Q.V. *XLNet: Generalized Autoregressive Pretraining for Language Understanding*; Curran Associates Inc.: Red Hook, NY, USA, 2019.
29. Jiao, X.; Yin, Y.; Shang, L.; Jiang, X.; Chen, X.; Li, L.; Wang, F.; Liu, Q. TinyBERT: Distilling BERT for natural language understanding. In *Findings of the Association for Computational Linguistics: EMNLP 2020*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2020; pp. 4163–4174. Available online: <https://aclanthology.org/2020.findings-emnlp.372> (accessed on 17 February 2023).
30. Clark, K.; Luong, M.-T.; Le, Q.V.; Manning, C.D. Electra: Pre-training text encoders as discriminators rather than generators. *Int. Conf. Learn. Represent.* **2020**. Available online: <https://openreview.net/forum?id=r1xMH1BtvB> (accessed on 17 February 2023).
31. Salimans, T.; Goodfellow, I.; Zaremba, W.; Cheung, V.; Radford, A.; Chen, X. Improved techniques for training gans. In Proceedings of the 30th International Conference on Neural Information Processing Systems, Barcelona, Spain, 6–12 December 2016; pp. 2234–2242.
32. Isola, P.; Zhu, J.-Y.; Zhou, T.; Efros, A.A. Image-to-image translation with conditional training networks. In Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 21–26 July 2017; pp. 5967–5976.
33. Schulman, J.; Zoph, B.; Kim, C.; Hilton, J.; Menick, J.; Weng, J.; Uribe, J.F.C.; Fedus, L.; Metz, L.; Pokorny, M.; et al. ChatGPT: Optimizing Language Models for Dialogue. *OpenAI*, 30 November 2022. Available online: <https://openai.com/blog/chatgpt> (accessed on 9 March 2023).

34. Wang, A.; Singh, A.; Michael, J.; Hill, F.; Levy, O.; Bowman, S. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP, Brussels, Belgium, 1–2 November 2018; pp. 353–355. Available online: <https://aclanthology.org/W18-5446> (accessed on 9 March 2023).
35. Black, S.; Leo, G.; Wang, P.; Leahy, C.; Biderman, S. GPT-Neo: Large Scale Autoregressive Language Modeling with Mesh-Tensorflow. 2021. Available online: <https://www.semanticscholar.org/paper/GPT-Neo%3A-Large-Scale-Autoregressive-Language-with-Black-Gao/7e5008713c404445dd8786753526f1a45b93de12> (accessed on 9 March 2023).
36. Gao, L.; Biderman, S.; Black, S.; Golding, L.; Hoppe, T.; Foster, C.; Phang, J.; He, H.; Thite, A.; Nabeshima, N.; et al. The pile: An 800gb dataset of diverse text for language modeling. *arXiv* **2020**, arXiv:2101.00027.
37. Iyer, S.; Lin, X.V.; Pasunuru, R.; Mihaylov, T.; Simig, D.; Yu, P.; Shuster, K.; Wang, T.; Liu, Q.; Koura, P.S.; et al. Opt-impl: Scaling language model instruction meta learning through the lens of generalization. *arXiv* **2022**, arXiv:2212.12017.
38. Dai, Z.; Yang, Z.; Yang, Y.; Carbonell, J.; Le, Q.V.; Salakhutdinov, R. Transformer-xl: Attentive Language Models Beyond a Fixed-Length Context. 2019. Available online: <https://arxiv.org/abs/1901.02860> (accessed on 9 March 2023).
39. Mikolov, T.; Chen, K.; Corrado, G.; Dean, J. Efficient Estimation of Word Representations in Vector Space. In Proceedings of the International Conference on Learning Representations (ICLR), Scottsdale, AZ, USA, 2–4 May 2013.
40. Pennington, J.; Socher, R.; Manning, C.D. GloVe: Global Vectors for Word Representation. In Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP), Doha, Qatar, 25–29 October 2014.
41. Subba, B.; Kumari, S. A heterogeneous stacking ensemble based sentiment analysis framework using multiple word embeddings. *Comput. Intell.* **2022**, *38*, 530–559. [[CrossRef](#)]
42. Rodriguez, P.L.; Spirling, A. Word embeddings: What works, what doesn't, and how to tell the difference for applied research. *J. Politics* **2022**, *84*, 101–115. [[CrossRef](#)]
43. Siino, M.; Di Nuovo, E.; Tinnirello, I.; La Cascia, M. Fake News Spreaders Detection: Sometimes Attention Is Not All You Need. *Information* **2022**, *13*, 426. [[CrossRef](#)]
44. Saleh, H.; Alhothali, A.; Moria, K. Detection of Hate Speech using BERT and Hate Speech Word Embedding with Deep Model. *arXiv* **2021**, arXiv:2111.01515.
45. Marco, S.; Elisa, D.N.; Ilenia, T.; Marco, L.C. Detection of hate speech spreaders using convolutional neural networks. In *Proceedings of Conference and Labs of the Evaluation Forum (CLEF)*, CEUR Workshop Proceedings, Bucharest, Romania, 21–24 September 2021; pp. 2126–2136.
46. Incitti, F.; Urli, F.; Snidaro, L. Beyond word embeddings: A survey. *Inf. Fusion* **2023**, *89*, 418–436. [[CrossRef](#)]
47. Warstadt, A.; Singh, A.; Bowman, S.R. Neural Network Acceptability Judgments. *Trans. Assoc. Comput. Linguist.* **2019**, *7*, 625–641. Available online: <https://aclanthology.org/Q19-1040> (accessed on 9 March 2023). [[CrossRef](#)]
48. Paperno, D.; Kruszewski, G.; Lazaridou, A.; Pham, N.Q.; Bernardi, R.; Pezzelle, S.; Baroni, M.; Boleda, G.; Fernandez, R. The LAMBADA dataset: Word prediction requiring a broad discourse context. In Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Berlin, Germany, 7–12 August 2016; pp. 1525–1534. Available online: <http://www.aclweb.org/anthology/P16-1144> (accessed on 9 March 2023).
49. Merity, S.; Xiong, C.; Bradbury, J.; Socher, R. Pointer sentinel mixture models. In Proceedings of the International Conference on Learning Representations (ICLR), Toulon, France, 24–26 April 2017. Available online: <https://openreview.net/forum?id=Byj72udxe> (accessed on 10 April 2023).
50. Marcus, M.P.; Santorini, B.; Marcinkiewicz, M.A. Building a large annotated corpus of English: The Penn Treebank. *Comput. Linguist.* **1993**, *19*, 313–330. Available online: <https://www.aclweb.org/anthology/J93-2004> (accessed on 9 March 2023).
51. Chelba, C.; Mikolov, T.; Schuster, M.; Ge, Q.; Brants, T.; Koehn, P.; Robinson, T. One billion word benchmark for measuring progress in statistical language modeling. In Proceedings of the 2013 International Conference on Computational Linguistics, Potsdam, Germany, 19–22 March 2013; pp. 2634–2643.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.