

Multitask BERT Fine-Tuning and Generative Adversarial Learning for Auxiliary Classification

Stanford CS224N Default Project

Christopher Sun

Department of Computer Science
Stanford University
chrisun@stanford.edu

Abishek Satish

Department of Computer Science
Stanford University
absatish@stanford.edu

Abstract

In this study, we implement a novel BERT architecture for multitask fine-tuning on three downstream tasks: sentiment classification, paraphrase detection, and semantic textual similarity prediction. Our model, Multitask BERT, incorporates layer sharing and a triplet architecture, custom sentence pair tokenization, loss pairing, and gradient surgery. Such optimizations yield a 0.516 sentiment classification accuracy, 0.886 paraphrase detection accuracy, and 0.864 semantic textual similarity correlation on test data. We also apply generative adversarial learning to BERT, constructing a conditional generator model that maps from latent space to create fake embeddings in \mathbb{R}^{768} . These fake embeddings are concatenated with real BERT embeddings and passed into a discriminator model for auxiliary classification. Using this framework, which we refer to as AC-GAN-BERT, we conduct semi-supervised sensitivity analyses to investigate the effect of increasing amounts of unlabeled training data on AC-GAN-BERT’s test accuracy. Overall, aside from implementing a high-performing multitask classification system, our novelty lies in the application of adversarial learning to construct a generator that mimics BERT. We find that the conditional generator successfully produces rich embeddings with clear spatial correlation with class labels, demonstrating avoidance of mode collapse. Our findings validate the GAN-BERT approach and point to future directions of generator-aided knowledge distillation.

1 Key Information to include

TA Mentor: Neil Nie, External Collaborators: None, Sharing project: No, Equal author contributions

2 Introduction and Related Work

Bidirectional Encoder Representations from Transformers (BERT), a pretrained transformer-based model, has revolutionized the field of Natural Language Processing (NLP) with its ability to produce robust contextual word representations, which have in turn been used to fine-tune models and achieve state-of-the-art performance in various NLP tasks [1][2][3]. Because of BERT’s notable use of self-attention, the model can contextually learn and understand each word’s importance to a sentence.

In this work, we apply BERT embeddings and multitask learning to three downstream tasks: sentiment classification, paraphrase detection, and semantic textual similarity prediction. Multitask learning aims to enhance model performance by using shared representations to improve a classifier’s generalization ability and alleviate challenges of task interference and imbalanced learning priorities. Liu et al. (2019) describes the complications that follow with multitask learning; when constructing their Multi-Task Deep Neural Network (MT-DNN), the authors emphasize the importance of finding cross-task data and using regularization to fight probable overfitting [3]. Our multitask BERT model takes inspiration from Bi et al.’s summation of loss functions across tasks for gradient calculations

and parameter updates. However, such additive losses succumb to the issue of destructive gradients [4]. Yu et al. (2020) propose a solution called gradient surgery, a technique that projects the gradient of a specific task onto the normal plane of another task’s conflicting gradient. This method, called PCGrad, helps alleviate the destructive gradient problem [4].

The goal of this work is to propose a high-performing BERT-based architecture for the aforementioned NLP tasks. We first implement several optimizations for fine-tuning, including data augmentation, architectural modifications, hyperparameter tuning, and adjustments in loss and gradient calculation. We then apply generative adversarial learning to BERT, improving the GAN-BERT framework developed by Croce et al. (2020), in which a discriminator conducts a $k + 1$ classification problem using both real BERT embeddings and fake generator embeddings [5]. GAN-BERT has shown promise for the incorporation of largely-unlabeled training data in semi-supervised training [5]. Our GAN methodology comprises a significant part of our work going beyond the default final project.

3 Approach

3.1 Multitask BERT

We apply a suite of optimizations to improve BERT’s performance on downstream NLP tasks, including multitask fine-tuning, manipulation of loss and gradient computations, and external augmentation of training data. After implementing the base version of BERT (minBERT) and loading pretrained weights using the provided starter code, we devise a novel head architecture for the simultaneous completion of our three NLP tasks of interest. Figure 1 depicts the flow of information through our classification system: for each prediction head, we pass raw sentence(s) into BERT, extracting the special classification token $[CLS] \in \mathbb{R}^H$, where H is the the hidden dimension of 768. We use layer sharing as motivated by previous success of MT-DNN in Liu et al. (2019) and biomedical multitask learning in Peng et al. (2020) [3] [6]. In our *shared* block, layer normalization is applied in between dense blocks because of training time benefits and batch size independence compared to batch normalization [7].

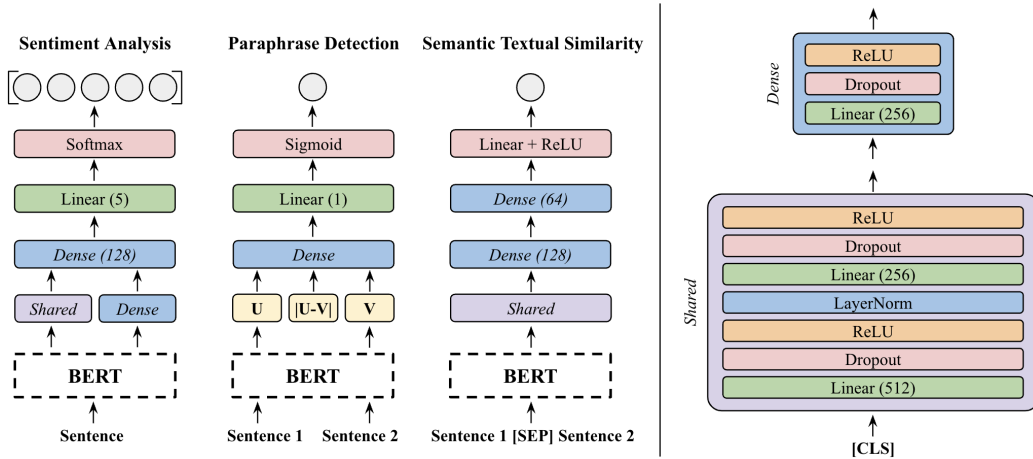


Figure 1: Multitask BERT Architecture

Sentiment Analysis To classify phrase sentiment, we gather the *shared* block output, simultaneously pass $[CLS]$ directly through a dense block, then concatenate these two representations and project them linearly. Softmax activation is applied to generate a probability distribution for multiclass classification.

Paraphrase Detection We apply a triplet network to learn thematic similarity between paired sentences, as motivated by Dor et al. (2018) [8]. In Figure 1, U and V are alternate notation for respective $[CLS]$ embeddings retrieved from BERT, and $|U - V|$ is the absolute difference in these embeddings. All three representations are concatenated and passed through dense layers, after which sigmoid activation is applied for binary classification.

Semantic Textual Similarity Though the STS task also involves sentence pairs, we opt for a different tokenization mechanism: connecting two sentences as one with the [SEP] token (we further discuss this decision in Section 4.3). This combined embedding passes through the same *shared block* as for sentiment analysis, as we reason logically that measures of semantic similarity align with sentiment relatedness.

Throughout the network, we apply dropout regularization to reduce overfitting [9].

3.2 Multitask Fine-tuning Optimizations

In multitask fine-tuning, a naive approach is to sum losses, then optimize parameters with respect to this summed loss, as in Bi et al. (2022) [10]. However, because of possible conflicting gradients between component losses for each prediction head, weights may be updated in suboptimal directions. To alleviate this issue, Yu et al. (2020) develops a gradient surgery technique called projecting conflicting gradients (PCGrad), which, simply put, involves modifying the gradient vectors before the parameter update step to ensure they are more aligned [4]. In our approach, we optimize Multitask BERT with “loss-pairing” and gradient surgery: in our training loop, we apply PCGrad* twice by pairing losses as follows, where *SST* corresponds to sentiment analysis, *P* corresponds to paraphrase detection, and *STS* corresponds to semantic similarity:

$$\mathcal{L}_1 = [\mathcal{L}_{SST}, \mathcal{L}_P], \mathcal{L}_2 = [\mathcal{L}_{STS}, \mathcal{L}_P].$$

The motivation for pairing losses as such is that we find paraphrase detection to be the highest-scoring, most smoothly-learning, most data-abundant task; taking this into account, we saw opportunity to use the gradient of \mathcal{L}_P as the “uniter” between $\nabla \mathcal{L}_{SST}$ and $\nabla \mathcal{L}_{STS}$. Because of PCGrad, we can rest assured that having one stable and reliable task in each loss pair \mathcal{L}_1 and \mathcal{L}_2 will prevent even an outlandish gradient of \mathcal{L}_{SST} or \mathcal{L}_{STS} from throwing off the optimal update trajectory. We apply PCGrad on \mathcal{L}_1 and \mathcal{L}_2 sequentially, allowing the more optimal paraphrase learning to positively affect the gradient updates of \mathcal{L}_{SST} and \mathcal{L}_{STS} independently. Hence, we propose a combination of PCGrad and paraphrase-paired loss as the optimal way to fine-tune Multitask BERT.

3.3 Summary of Semi-Supervised GAN Training Framework

We validate the GAN-BERT[†] framework and incorporate novel modifications, which we further assess in Section 5. We develop a conditional Generative Adversarial Network (cGAN) from minBERT that acts as an auxiliary classifier for sentiment analysis and paraphrase detection. We refer to this model as AC-GAN-BERT, illustrated in Figure 2.




Figure 2: Auxiliary Classifier GAN-BERT Architecture (number of hidden layers variable)

The mathematical formulations in this section are directly adapted from Croce et al. (2020) [5]. Let p_d denote the real data distribution, p_m denote model-predicted probabilities, and p_G denote the generated examples.

*We make use of Yu et al.’s source code linked in their paper.

†We inherit and make important modifications to Croce et al.’s source code.

Discriminator The Discriminator \mathcal{D} is a head for the BERT base, operating on both BERT [CLS] embeddings and Generator (below) outputs for multiclass classification. The Discriminator loss $\mathcal{L}_{\mathcal{D}}$ is a sum of supervised and unsupervised loss terms $\mathcal{L}_{\mathcal{D}_S}$ and $\mathcal{L}_{\mathcal{D}_U}$:

$$\begin{aligned}\mathcal{L}_{\mathcal{D}_S} &= -\mathbb{E}_{x,y \sim p_d} \log[p_m(\hat{y} = y|x, y \in (1, \dots, k))] \\ \mathcal{L}_{\mathcal{D}_U} &= -\mathbb{E}_{x \sim p_d} \log[1 - p_m(\hat{y} = y|x, y = k + 1)] - \mathbb{E}_{x \sim \mathcal{G}} \log[p_m(\hat{y} = y|x, y = k + 1)]\end{aligned}$$

In simple terms, $\mathcal{L}_{\mathcal{D}_S}$ represents the misclassification of a real labeled example belonging to one of k classes, and $\mathcal{L}_{\mathcal{D}_U}$ represents the misclassification of an unlabeled example as real or fake.

Conditional Generator The Conditional Generator \mathcal{G} conditions on a class label to project from latent space to a fake embedding, trying to fool the Discriminator into thinking the embedding is real. The Generator loss $\mathcal{L}_{\mathcal{G}}$ is a sum of “feature matching” and unsupervised loss terms $\mathcal{L}_{\mathcal{G}_{FM}}$ and $\mathcal{L}_{\mathcal{G}_U}$:

$$\begin{aligned}\mathcal{L}_{\mathcal{G}_{FM}} &= \|\mathbb{E}_{x \sim p_d} f(x) - \mathbb{E}_{x \sim \mathcal{G}} f(x)\|_2^2 \\ \mathcal{L}_{\mathcal{G}_U} &= -\mathbb{E}_{x \sim \mathcal{G}} \log[1 - p_m(\hat{y} = y|x, y = k + 1)]\end{aligned}$$

Since \mathcal{G} ideally generates embeddings similar to p_d , $\mathcal{L}_{\mathcal{G}_{FM}}$ examines the difference between an expected Generator example and the activation of an intermediate layer in \mathcal{D} (hence the $f(x)$) expressed as a squared L2 norm. Finally, $\mathcal{L}_{\mathcal{G}_U}$ considers the correct classification of fake (generated) embeddings by \mathcal{D} , which penalizes \mathcal{G} because of the inability of \mathcal{G} to fool \mathcal{D} .

4 Experiments

4.1 Data

Sentiment Analysis We use the Stanford Sentiment Treebank (SST) dataset for the sentiment classification task, which involves labeling phrases (215,154 in total) as negative, somewhat negative, neutral, somewhat positive, or positive [11]. During fine-tuning, we also bring in the Massive Text Embedding Benchmark (MTEB) dataset consisting of 31,015 additional rows of data [12]. However, because MTEB labels range from 0-2, we scale them by 2; though this creates class imbalance, we find this additional data to be beneficial (further discussion in Section 5).

Paraphrase Detection We use the Quora Question Pairs dataset to predict a binary label representing whether questions pairs (404,298 in total) are paraphrases of each other [13].

Semantic Textual Similarity We use the SemEval STS dataset for the regression task of predicting semantic similarity on a scale of 0 to 5 between two sentences (8,628 in total) [14]. During fine-tuning, we also bring in the Sentences Involving Compositional Knowledge (SICK) dataset consisting of 9,840 additional rows of data [15].

4.2 Experimental Details and Evaluation Methods

Before building the Multitask BERT architecture in Figure 1, we conduct a baseline sanity check with minBERT that simply projects from [CLS] embeddings to task predictions using one linear layer per output head, with a naive summing of losses and no gradient modifications.

Multitask BERT We train the sentiment analysis head with Categorical Cross Entropy loss, the paraphrase detection head with Binary Cross Entropy loss, and the semantic textual similarity head with Mean Square Error loss. As metrics, we track accuracy and Pearson correlation on train and dev data through each epoch. We run experiments according to the optimizations discussed in Section 3.2.

As hyperparameters, we decide to use a batch size of 112 (the maximum before our machines crash), learning rate of 1×10^{-5} , dropout probability of 0.5, and weight decay coefficient of 1×10^{-3} . We train Multitask BERT for 10-15 epochs.

We also introduce cyclic data loaders, essentially exhausting all the possible data for training. Due to large variance in dataset size across tasks, one zipped batch of data prevents data-abundant tasks from undergoing maximal training. In practice, using a cyclic data loader means that while iterating through the Quora Question Pairs data for paraphrase detection, we exhaust other tasks’ data and ensure SST and STS data are passed into training again.

AC-GAN-BERT We run a series of experiments to investigate AC-GAN-BERT, training the model as specified in Section 3.3:

- We test whether there is marginal benefit to using a conditional GAN (conditioned on class labels) vs. normal GAN as in Croce et al. (2020). We accomplish this both quantitatively through test accuracy and qualitatively through visual deciphering of embedding “quality.”
- We run experiments varying the depth of \mathcal{D} and \mathcal{G} , examining the relationship between the number of hidden layers of AC-GAN-BERT and model performance.
- We also conduct a sensitivity analysis on semi-supervised learning on the sentiment analysis task with AC-GAN-BERT, varying the proportion of labeled data to investigate how resistant the model is to unlabeled data. We accomplish this by using a random seed to mask out increasing numbers of phrases in the SST dataset; these masked out phrases become unlabeled data for the discriminator to classify as real or fake. We seek to quantify a point at which performance tapers off when using unlabeled data with respect to test accuracy.

We use a learning rate of 5×10^{-5} and a 100-length noise vector as generator latent space input. Because of computational resource limitations, we train AC-GAN-BERT for 5 epochs, and for the paraphrase detection task train data is truncated to 10,000 sentence pairs and test data to 2,000 sentence pairs. As evaluation methods, we quantitatively assess accuracy scores and qualitatively assess embedding representations.

4.3 Results

Multitask BERT We list performance metrics for our baseline, several iterations of Multitask BERT, and our final test leaderboard submission in Table 1. Some notable optimizations result in large performance boosts, such as tokenization of both sentences separated with [SEP] for semantic textual similarity as opposed to using the same triplet network approach as in paraphrase detection (Light Multitask BERT and beyond have this modification). Interestingly and unexpectedly, adding more data and larger, deeper layers did not positively affect sentiment classification accuracy.

Model Details	Sentiment	Paraphrase	Similarity	Overall
Baseline	0.410	0.532	0.102	0.498
Baseline with PCGrad	0.437	0.565	0.211	0.536
Pairwise Loss Addition with PCGrad	0.481	0.723	0.349	0.626
Double Triplet for Paraphrase and STS	0.513	0.756	0.361	0.650
Light Multitask BERT	0.496	0.870	0.825	0.760
Light Multitask BERT Data-Enhanced	0.502	0.878	0.822	0.764
Multitask BERT Data-Enhanced	0.501	0.888	0.872	0.775
Final Multitask BERT on Test Data	0.516	0.886	0.864	0.778

Table 1: Compilation of Multitask BERT Results

AC-GAN-BERT We list classification accuracies of GAN-BERT and AC-GAN-BERT in Table 2. There is no statistically significant boost in accuracy resulting from the inclusion of a conditional generator into the GAN-BERT framework. However, we observe qualitative improvements which are discussed in Section 5.

Model Details	Sentiment	Paraphrase
GAN-BERT	0.502	0.816
AC-GAN-BERT Depth 1	0.503	0.814
AC-GAN-BERT Depth 2	0.510	0.820
AC-GAN-BERT Depth 3	0.495	0.817

Table 2: GAN-BERT Performance - Varying Hidden Depth

Figure 3 and Table 3 display the results of the semi-supervised learning sensitivity analysis on the SST dataset, where λ represents the proportion of the SST dataset that is masked (where individual examples are selected with a random seed). The overall trend is that a larger fraction of unlabeled training data causes AC-GAN-BERT to perform with a lower accuracy. In particular, relative to the performance of AC-GAN-BERT on a fully labeled dataset ($\lambda = 0.0$), the performance drop becomes statistically significant ($p < 0.05$) when $\lambda = 0.4$, hence the color-coding of Figure 3. We calculate p -values using a one-tailed t-test for difference in sample means, where a sample consists of validation accuracies from all epochs. Table 4 contains all p -values in this analysis.




Figure 3: Unlabeling-Induced Accuracy Decrease

λ	Max Acc.	Min Acc.	Mean Acc.
0.0	0.503	0.480	0.491
0.1	0.495	0.476	0.485
0.2	0.505	0.437	0.473
0.3	0.498	0.465	0.485
0.4	0.490	0.474	0.481
0.5	0.486	0.446	0.469
0.6	0.481	0.454	0.465
0.7	0.460	0.425	0.447
0.8	0.441	0.411	0.422
0.9	0.451	0.410	0.432

Table 3: Sensitivity Analysis Results

5 Analysis

5.1 Multitask BERT

Discussion When devising our final architecture on limited compute and a low batch size of 4 or 8 examples, we observe that using a lower learning rate of 8×10^{-6} aided learning, as smaller updates are conducive to better convergence, especially given that fine-tuning is a learning rate sensitive task [16]. We also observe that doubling the number of hidden units in the *shared* and *dense* blocks in Figure 1 boosts performance. However, after obtaining more computational power, we increase batch size to 112, the maximum our GPU can handle; we also increase learning rate to 1×10^{-5} , as with more examples per batch, we reason that larger update steps can be taken. In general, though we apply dropout regularization, weight decay, and layer normalization, we still observe inevitable overfitting, with training data accuracies very quickly reaching nearly 100%.

Noteworthy Optimizations We observe large improvements in paraphrase detection accuracy with the triplet network architecture. Inspired by this, we originally attempt to use the same architecture for STS prediction, but realize that Pearson correlation plateaus at around 0.4. We then switch our STS tokenization mechanism, tokenizing sentence pairs simultaneously with the [SEP] token as a connector. This adjustment improves STS Pearson correlation significantly by around 0.3. Additionally, data augmentation through the SICK and MTEB datasets boosts performance on both paraphrase detection and semantic similarity prediction, but interestingly does not increase sentiment classification accuracy.

5.2 AC-GAN-BERT

Given that GANs are susceptible to mode collapse, we are especially interested in evaluating the *quality* of Generator output [17]. Recall that one of our main motivations for developing a generative adversarial framework is to assess the ability of the Generator to produce high-quality embeddings, in other words, fulfill the same role as BERT. However, “quality” can be defined many ways: we know that a good generator’s samples will mimic the true data manifold, and in computer vision applications, mere visual inspection can be used to evaluate quality [18][19]. In our work, we evaluate generator quality through visual inspection after mathematical manipulation. To qualitatively

assess embeddings, we apply t-distributed stochastic neighbor embedding (t-SNE) to reduce hidden dimensionality from 768 (the hidden dimension of [CLS]) to 2, allowing embeddings to be plotted on an xy-coordinate plane [20].

We extract both BERT and Generator embeddings for three models: **1)** Croce et al.’s unconditional GAN-BERT trained with 40% unlabeled data, **2)** AC-GAN-BERT trained with 40% unlabeled data, and **3)** AC-GAN-BERT trained with 80% unlabeled data. Comparison between the first and second allows us to isolate the marginal benefit in embedding quality of a conditional generator, and comparison between the second and third reveals the effect of highly-skewed datasets on embedding quality. Figure 4 illustrates t-SNE results of Model 1 embeddings, Figure 5 illustrates t-SNE results of Model 2 embeddings, and Figure 6 illustrates t-SNE results of Model 3 embeddings.

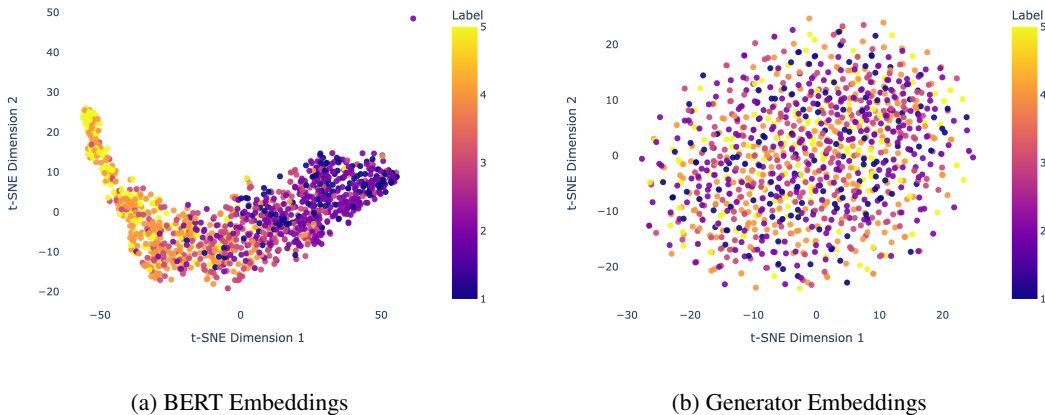


Figure 4: Unconditional GAN-BERT - 40% Unlabeled

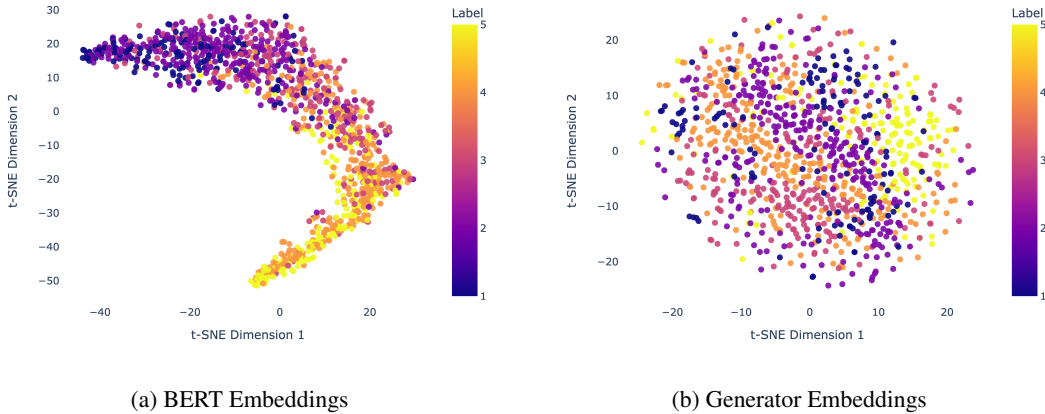


Figure 5: AC-GAN-BERT - 40% Unlabeled

Critically, we notice from Figure 4(b) that when using an unconditional Generator, there is *no spatial relationship* between embedding and class label. However, we notice from Figure 5(b) that embeddings seem to be clustered according to their class label, with equally-labeled embeddings residing in similar spatial regions. This observation is even more stark in Figure 6(b), which reveals that embedding regions are well-defined according to their class label. In all cases, BERT embeddings possess a spatial relationship with class label, which of course is a validation of BERT’s quality. But similar observations in conditional Generator embeddings serve as promising validation of their diversity and quality. In general, Generators can fail to map from latent space to their full output domain, instead creating outputs with very little variation. If this was the case in \mathbb{R}^{768} , then projecting

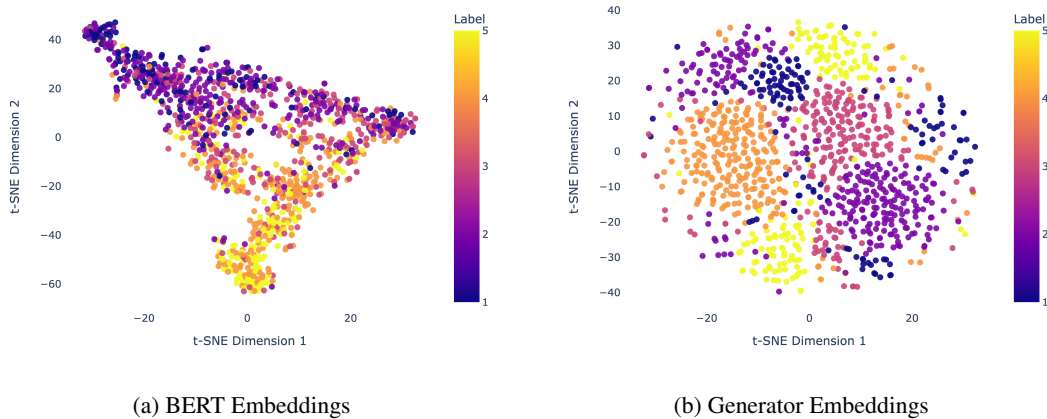


Figure 6: AC-GAN-BERT - 80% Unlabeled

embeddings down to \mathbb{R}^2 with t-SNE would surely cause mode collapse to manifest visually. But because we do not observe this, we can fairly confidently conclude an avoidance of mode collapse.

6 Conclusion

In this study, we present Multitask BERT, a novel architecture optimized for multitask fine-tuning on sentiment classification, paraphrase detection, and semantic textual similarity prediction. As a result of data augmentation, loss pairing, and custom tokenization, we achieve an overall performance score of 0.778 on test data. Additionally, we introduce AC-GAN-BERT, which incorporates generative adversarial learning to enhance BERT’s capabilities. The conditional generator within AC-GAN-BERT successfully produces class-conditioned embeddings, avoiding mode collapse and showing clear spatial correlation with class labels. Our results call for future work on more advanced loss functions and better regularization mechanisms to further maximize accuracy. Our conditional generator also offers potential as a student model for BERT knowledge distillation.

7 Ethics Statement

Classifiers like Multitask BERT and AC-GAN-BERT present the ethical challenge of propagation of the biases embedded in their training data. AC-GAN-BERT, in particular, poses a heightened risk of amplifying existing biases because of its adversarial nature, potentially reinforcing stereotypes or discrimination in downstream NLP tasks. To address these concerns, it is crucial to train these models on diverse and representative training data. Careful data curation can help with over- or under-representation of certain demographic groups or high-resource languages. Implementing a robust bias detection system during the data preprocessing stage can also help with filtering through biased data, thereby reducing bias propagation. In addition, semi-supervised training on AC-GAN-BERT means that both labeled and unlabeled data are necessary. This could raise concerns about the misrepresentation of sensitive information, such as false attribution of personal details, especially in domains like finance and healthcare. To mitigate this risk, training data should be anonymized and subject to differential privacy methods. Establishing such protocols will ensure the ethical use of models like Multitask BERT and AC-GAN-BERT.

References

- [1] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [2] Ian Tenney, Dipanjan Das, and Ellie Pavlick. Bert rediscovers the classical nlp pipeline. *arXiv preprint arXiv:1905.05950*, 2019.

- [3] Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. Multi-task deep neural networks for natural language understanding. *arXiv preprint arXiv:1901.11504*, 2019.
- [4] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. *Advances in Neural Information Processing Systems*, 33:5824–5836, 2020.
- [5] Danilo Croce, Giuseppe Castellucci, and Roberto Basili. Gan-bert: Generative adversarial learning for robust text classification with a bunch of labeled examples. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 2114–2119, 2020.
- [6] Yifan Peng, Qingyu Chen, and Zhiyong Lu. An empirical study of multi-task learning on bert for biomedical text mining. *arXiv preprint arXiv:2005.02799*, 2020.
- [7] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [8] Liat Ein Dor, Yosi Mass, Alon Halfon, Elad Venezian, Ilya Shnayderman, Ranit Aharonov, and Noam Slonim. Learning thematic similarity metric from article sections using triplet networks. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 49–54, 2018.
- [9] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- [10] Qiwei Bi, Jian Li, Lifeng Shang, Xin Jiang, Qun Liu, and Hanfang Yang. Mtrec: Multi-task learning over bert for news recommendation. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2663–2669, 2022.
- [11] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642, 2013.
- [12] Niklas Muennighoff, Nouamane Tazi, Loïc Magne, and Nils Reimers. Mteb: Massive text embedding benchmark. *arXiv preprint arXiv:2210.07316*, 2022.
- [13] Shankar Iyer, Nikhil Dandekar, and Kornel Csernai. First quora dataset release: Question pairs. Quora, 2017.
- [14] Eneko Agirre, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, and Weiwei Guo. * sem 2013 shared task: Semantic textual similarity. In *Second joint conference on lexical and computational semantics (*SEM), volume 1: proceedings of the Main conference and the shared task: semantic textual similarity*, pages 32–43, 2013.
- [15] Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli. A SICK cure for the evaluation of compositional distributional semantic models. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 216–223, 2014.
- [16] Hao Li, Pratik Chaudhari, Hao Yang, Michael Lam, Avinash Ravichandran, Rahul Bhotika, and Stefano Soatto. Rethinking the hyperparameters for fine-tuning. *arXiv preprint arXiv:2002.11770*, 2020.
- [17] David Bau, Jun-Yan Zhu, Jonas Wulff, William Peebles, Hendrik Strobelt, Bolei Zhou, and Antonio Torralba. Seeing what a gan cannot generate. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4502–4511, 2019.
- [18] Konstantin Shmelkov, Cordelia Schmid, and Karteek Alahari. How good is my gan? In *Proceedings of the European conference on computer vision (ECCV)*, pages 213–229, 2018.
- [19] Pegah Salehi, Abdolah Chalechale, and Maryam Taghizadeh. Generative adversarial networks (gans): An overview of theoretical model, evaluation metrics, and recent developments. *arXiv preprint arXiv:2005.13178*, 2020.

[20] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.

Appendix

λ	Columnwise p -values Relative to $\lambda = \frac{\text{column number}}{10}$									
0.0	0.500									
0.1	0.152	0.500								
0.2	0.085	0.170	0.500							
0.3	0.215	0.510	0.819	0.500						
0.4	0.033	0.190	0.738	0.250	0.500					
0.5	0.017	0.047	0.385	0.063	0.092	0.500				
0.6	<0.001	0.010	0.271	0.019	0.021	0.347	0.500			
0.7	<0.001	<0.001	0.044	<0.001	<0.001	0.035	0.044	0.500		
0.8	<0.001	<0.001	0.002	<0.001	<0.001	<0.001	<0.001	0.012	0.500	
0.9	<0.001	<0.001	0.007	<0.001	<0.001	<0.001	<0.001	0.075	0.849	0.500

Table 4: Summary of p -values Relative to Unlabeled Percentages

The way to interpret this table is as follows: in column i , the entry at row i will always equal 0.5, because the p -value resulting from a one-tailed t-test of difference in sample means, where both samples are the same, will always equal 0.5. In column i , the entry in row $i + j$, where j ranges from 1 to $9 - i$, is the p -value representing the significance of the decrease in validation accuracy between training AC-GAN-BERT with $\lambda = \frac{i}{10}$ and $\lambda = \frac{i+j}{10}$.