Một tiếp cận học đa nhiệm dựa trên mô hình ngôn ngữ nhỏ cho một số bài toán xử lý ngôn ngữ tự nhiên

Lê Xuân Tuyên, Lê Quang Hùng[*]

Khoa Công nghệ thông tin, Trường Đại học Quy Nhơn, Việt Nam

TÓM TẮT

Trong lĩnh vực Xử lý ngôn ngữ tự nhiên (Natural Language Processing - NLP), các mô hình dựa trên Transformer đã mang lại những kết quả hấp dẫn về khả năng hiểu và sinh ngôn ngữ. Tuy nhiên, số lượng tham số cực lớn và chi phí tính toán cao của các mô hình này gây khó khăn cho việc triển khai trên các thiết bị hạn chế về tài nguyên. Điều này thúc đẩy việc nghiên cứu phát triển các phương pháp hiệu quả, mang lại sự cân bằng giữa kích thước mô hình và hiệu suất. Bài báo này đề xuất một tiếp cận học đa nhiệm mới, dựa trên mô hình ngôn ngữ nhỏ như TinyBERT cho một số bài toán (tác vụ) NLP, bao gồm: phân tích cảm xúc, phát hiện diễn giải và độ tương đồng ngữ nghĩa của văn bản. Chúng tôi thực hiện huấn luyện ban đầu thông qua cơ chế học đa nhiệm để nắm bắt các đặc trưng ngôn ngữ chung. Sau đó, mô hình huấn luyện trước này được tinh chỉnh cho từng tác vụ cụ thể. Chúng tôi tiến hành các thực nghiệm để đánh giá hiệu quả của phương pháp này trên bộ dữ liệu chuẩn GLUE. Kết quả chứng minh tính hiệu quả của phương pháp của chúng tôi trong thiết lập học đa nhiệm cho các tác vụ NLP khác nhau. Chúng tôi cung cấp cài đặt thực nghiệm của nghiên cứu này trên kho lưu trữ Github[1].

Từ khóa: *NLP, học đa nhiệm, tinh chỉnh, TinyBERT, phân tích cảm xúc, phát hiện diễn giải, độ tương đồng ngữ nghĩa.*

---

[*]*Tác giả liên hệ chính*
 Email: lequanghung@qnu.edu.vn
[1]*https://github.com/Tuyenle2/Project-code-2025.1143.93*

A Multi-Task Learning Approach Based on Small Language Model for Natural Language Processing Tasks

Xuan-Tuyen Le, Quang-Hung Le[*]

Department of Information Technology, Quy Nhon University, Vietnam

ABSTRACT

In the field of Natural Language Processing (NLP), Transformer-based models have yielded appealing results in language understanding and generation. However, these models' extremely high number of parameters and computationally expensive are challenging to deployment on resource-constrained devices. This motivates research in the development of efficient methods that offer an attractive balance between model size and performance. This paper proposes a novel multi-task learning approach, which relies on a small language model like TinyBERT for NLP tasks, including: sentiment analysis, paraphrase detection, and semantic textual similarity. We perform initial training through a Multi-Task Learning (MTL) mechanism to capture general language features. Subsequently, this pre-trained model is fine-tuned for each specific task. We conduct experiments to evaluate the effect of this approach on the GLUE benchmark dataset. The results demonstrate the effectiveness of our method within a multi-task learning setup for diverse NLP tasks. We provide an implementation of this work on Github repository[1].

Keywords: *NLP, multi-task learning, fine-tuning, TinyBERT, sentiment analysis, paraphrase detection, semantic similarity.*

1. INTRODUCTION

Large language models such as BERT[1], RoBERTa[2], T5[3], XLNet[4], ELECTRA[5], and GPT[6] have revolutionized the field of NLP, achieving state-of-the-art performance on a wide range of tasks. However, their massive size and high computational requirements pose significant barriers to deployment on resource-constrained devices. TinyBERT[7], a compact version of BERT created through knowledge distillation, has emerged as a potential solution, offering an attractive balance between model size and performance. To fully exploit the potential of these compact models, multi-task fine-tuning (MT-FT) has become an increasingly popular research direction in the field of NLP[8]. The primary goal of MTL is to improve the performance of multiple related tasks by leveraging useful information shared among them[10]. Instead of training separate models for each task, MTL allows a single model to learn simultaneously from multiple tasks, thereby promoting knowledge sharing and improving generalization capabilities[8,9]. Various multi-task architectures and strategies have been proposed[10], from shared encoder layers with task-specific output layers as in multi-task deep neural networks[8], to unified "text-to-text" approaches like T5[3,11], or efficient adapter-based techniques[13]. Although modern NLP models continually push the boundaries of performance, the effective application of MT-FT,

---

[*]*Corresponding author*
Email: lequanghung@qnu.edu.vn
[1]*https://github.com/Tuyenle2/Project-code-2025.1143.93*

especially on compact models like TinyBERT, still faces many challenges. Key challenges include selecting and balancing tasks to avoid negative transfer[14], designing optimal network architectures for efficient parameter sharing[10], dealing with catastrophic forgetting when learning multiple tasks[14], and particularly addressing the limited capacity of small models when handling the complexity of multiple data sources and learning objectives[13]. Building on these considerations, we pose three research questions (RQ) and present main contributions to address them.

- **RQ1:** How can TinyBERT be leveraged in a multi-task learning architecture to simultaneously learn these three tasks, exploiting their commonalities to improve generalization and computational efficiency?

- **RQ2:** After initial multi-task training, to what extent can individual fine-tuning for each task improve performance compared to traditional single-task training?

- **RQ3:** How does this approach balance performance and computational efficiency compared to larger models (like BERT-base) or other approaches in current research?

To answer these questions, we design a multi-task learning framework built upon Tiny-BERT[7] to tackle multiple NLP tasks including: sentiment analysis, paraphrase detection, and semantic textual similarity. Our main contributions are summarized as follows:

- First, we propose a novel multi-task learning approach, which relies on a small language model for NLP tasks, aiming to address the efficiency-performance trade-off under resource constraints.

- Second, we demonstrate that leveraging shared representations through MTL followed by individual task-specific fine-tuning significantly enhances the overall performance on sentiment analysis, paraphrase detection, and semantic textual similarity tasks.

- Third, we provide empirical evidence showing that despite its compact size, TinyBERT can achieve competitive performance with considerably lower computational cost, offering a practical alternative to larger models such as BERT-base in real-world applications.

The rest of this paper is organized as follows. Section 2 provides preliminaries on problem statement, multi task learning, fine tuning, and TinyBERT that form the basis for the development of the proposed method detailed in Section 3. Section 4 describes the experimental setup, results and comparing the result with BERT-base in Section 5. Section 6 briefly recalls some related work. Finally, Section 7 wraps up the paper with conclusions and future work.

2. BACKGROUND

2.1. Multi-Task Learning

Multi-task learning (MTL for short) is a machine learning approach in which a single model is trained to perform multiple related tasks simultaneously[15]. The core idea is that different tasks can share information or common representations, which helps to improve overall performance, enhance generalization capabilities, and reduce the risk of overfitting, especially when data for individual tasks is limited[16]. In NLP, MTL is often applied by sharing the encoder layers of the model across multiple tasks, while having separate output layers for each specific task[8].

Formally, the total loss in MTL can be defined as:

$$L_{\text{total}} = \sum_{i=1}^{T} L_i \qquad (1)$$

where L_i denotes the task-specific loss for the i-th task, and T is the number of tasks. To better balance task learning, dynamic weighting strategies have been proposed[14]. One such formulation is:

$$L_{\text{total}} = \sum_{i=1}^{T} w_i L_i \quad \text{with} \quad w_i = \frac{L_{\text{total}}}{T \times L_i} \quad (2)$$

This joint learning allows the model to capture fundamental language features useful for various objectives.

2.2. Problem Statement

Let $\mathcal{D}$ be the set of input texts (or text pairs, depending on the specific task), where each element $x \in \mathcal{D}$ is a sentence, a text passage, or a pair of sentences/words. For each text x (or text pair), there will be corresponding labels for each task:

- Sentiment analysis (SA): $y_{sa} \in (0;1)$ where $y_{sa} = 0$ if the text has negative sentiment, and $y_{sa} = 1$ if positive.

- Paraphrase detection (PD): For a pair of texts $(x_1;x_2)$, the label $y_{pd} \in (0;1)$, where $y_{pd} = 0$ if the two texts are not paraphrases of each other, and $y_{pd} = 1$ if they are paraphrases.

- Semantic textual similarity (STS): For a pair of words $(w_1;w_2)$, the label $y_{pd} \in (0;1)$, where $y_{sts} = 0$ if the two words are antonyms, and $y_{sts} = 1$ if they are synonyms (in some cases, this task might involve determining if a word in a specific context can be replaced by another word without significantly changing the sentence's meaning).

The MTL problem with TinyBERT can be modeled as a function f_{mtl}:

$$f_{mtl} : \mathcal{D} \to (y_{sa}, y_{pd}, y_{sts}) \qquad (3)$$

Here, $f_{mtl}(x)$ includes the model's predicted labels $(y_{sa}, y_{pd}, y_{sts})$ or the input x across all three tasks. $(y_{sa}, y_{pd}, y_{sts})$ represents the set of possible labels for each task.

2.3. Fine-Tuning

Fine-tuning (FT for short) is a popular transfer learning technique[17,18] where a model that has been pre-trained on a large amount of data (often self-supervised tasks like language modeling) is then adjusted or "fine-tuned" on a smaller, task-specific dataset for a target task[19]. The pre-training process helps the model learn rich and contextual language representations[20].

The fine-tuning step allows the model to adapt this learned knowledge to the nuances of the specific task, often leading to much better

performance than training a model from scratch on limited task data[1]. In classification tasks, the typical loss function used during fine-tuning is the cross-entropy loss:

$$L_{\text{CE}} = -\sum_{c=1}^{C} y_c \log(p_c) \qquad (4)$$

where C is the number of classes, y_c is the true label distribution, and p_c is the predicted probability for class c. In more recent setups, supervised contrastive learning[22,23] is also applied to encourage discriminative feature learning:

$$L_{\text{SCL}} = -\frac{1}{N} \sum_{i=1}^{N} \frac{1}{|P(i)|} \sum_{p \in P(i)} \log \frac{\exp(s_{i,p}/\tau)}{\sum_{a \in A(i)} \exp(s_{i,a}/\tau)} \qquad (5)$$

where:

- N: total samples in a batch.

- $P(i)$: set of positives (samples with the same label as i, excluding i).

- $|P(i)|$: number of positives for anchor i.

- $A(i)$: set of all samples in the batch excluding i (positives and negatives).

- $s_{i,p}$: cosine similarity between embeddings of i and p.

- τ: temperature parameter controlling sharpness of the distribution.

2.4. TinyBERT

While large models like BERT[1,12] achieve impressive performance, their size and computational cost are significant barriers to deployment on resource-constrained devices. To address this issue, more compact models have been developed. TinyBERT[7,21] utilizes *knowledge distillation* to reduce the size of BERT while retaining a large portion of its performance.

This distillation process involves transferring knowledge from a large *teacher* model (the original BERT) to a smaller *student* model (TinyBERT) during both the pre-training and task-specific fine-tuning stages.
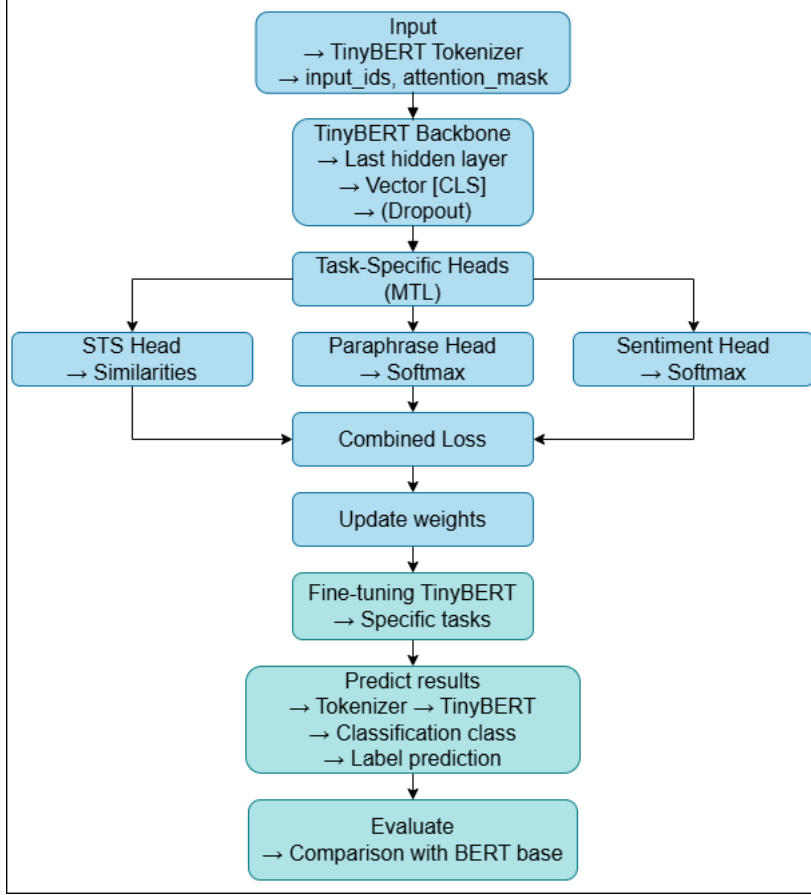
4

Figure 1. Overview of proposed method

The knowledge distillation loss in Tiny-BERT typically integrates several components:

$$L_{KD} = \alpha L_{emb} + \beta L_{hidden} + \gamma L_{logit} \quad (6)$$

where L_{emb}, L_{hidden}, and L_{logit} refer to the embedding-level, hidden-state-level, and output-level distillation losses, respectively.

3. METHODOLOGY

This section presents the architecture and training process of our proposed approach, which combines MTL and fine-tuning strategies using the TinyBERT model. First, we detail the architecture that supports shared learning across multiple tasks. Then, we describe how the TinyBERT backbone is trained in a multi-task setting. Finally, we explain the individual fine-tuning phase applied after MTL for task-specific optimization.

3.1. Model Architecture

Our architecture utilizes TinyBERT as a shared backbone to extract contextual representations across all tasks. As shown in Fig 1, raw text data is fed into the system as character strings. This text is pre-processed by the TinyBERT-Tokenizer to convert it into corresponding tokens. The encoded tokens from each task are then passed to the TinyBERT backbone. This is the shared architectural component, sharing weights across all tasks. TinyBERT processes these token sequences and generates contextual representations at the last hidden layer. Specifically, the embedding vector corresponding to the special [CLS] token from the last hidden layer is usually taken as a representation for the entire input sequence. This vector is expected to capture the composite information of the sentence. A dropout layer is applied after extracting the [CLS] vector to help the model prevent overfitting during the MTL training process. The [CLS] representation (after

the dropout layer) is then fed into separate heads, each designed for a specific task within the set of MTL tasks. The loss function from each task head is computed independently. These loss functions are then combined to form a joint loss function for the entire MTL model.

$$L_{total} = w_{sa}L_{sa} + w_{pd}L_{pd} + w_{si}L_{si} \quad (7)$$

where w_{sa}, w_{pd}, w_{si} are weights to balance the contribution of each task. The optimization process relies on this joint loss function to simultaneously update the weights of the TinyBERT backbone and the task-specific heads. This allows TinyBERT to learn robust and flexible language representations from various information sources. After TinyBERT has been pre-trained with MTL and possesses good language understanding capabilities, it is then fine-tuned for the specific task of interest. In here, we fine-tune for all three aforementioned tasks to compare the model's performance after fine-tuning with the model after multi-task training. After fine-tuning, the complete TinyBERT model (including the TinyBERT backbone and the classification layer for the target task) is used to make predictions. With a new input text, the system will process it through the tokenizer, feed it into Tiny-BERT, pass it through the classification layer, and finally predict the label based on the class with the highest probability. The training process comprises two main phases:

- **Phase 1: Multi-task Training.** The input text x (or its components, such as $x_1;x_2$ for PD, or words $w_1;w_2$ in context for SI) is fed through the TinyBERT model. TinyBERT generates contextual embeddings (or a set of embeddings). The key point of MTL is that most of TinyBERT's parameters (the Transformer layers) are shared across all tasks, allowing the model to learn general representations useful for multiple tasks.

- **Phase 2: Fine-tuning for Specific NLP Tasks.** From the shared representation ex obtained from TinyBERT separate output layers are used for each task. A classification layer is applied to the representation of the [CLS] token from ex to predict the sentiment label y_{sa}. The representation of the sentence pair from TinyBERT

(typically the [CLS] token's representation when the two sentences are concatenated and fed into the model) is passed through a classification layer to predict the label y_{pd}. Similarly, the representation of the word pair from TinyBERT is passed through a classification layer to predict the label y_{sts}.

This architecture enables TinyBERT to learn generalized features that are beneficial across tasks, while still allowing task-level specialization through individual fine-tuning.

3.2. Multi-task TinyBERT Training

Algorithm 1 describes the training procedure for the multi-task TinyBERT model. First, the TinyBERT model is initialized with pre-trained weights. Then, appropriate labels and classification layers are added for each task. Data is divided into batches, and each batch is tokenized using TinyBERT's tokenizer to convert text into a suitable tokenized format. Subsequently, the TinyBERT model extracts an embedding representing the entire sentence from the [CLS] token. Next, the [CLS] token (after the dropout layer) is fed into separate heads, each designed for a specific task within the set of MTL tasks. Finally, the loss for each task is calculated based on its type, and model parameters are updated through a gradient descent algorithm. This process is repeated over multiple epochs, and the AdamW optimizer[24] is used to optimize the model's. The TinyBERT backbone is shared among tasks, allowing the model to learn general representations useful for multiple tasks, while separate outputs enable the model to adjust predictions for each specific task. During training, the model receives input as input-ids and attention-mask, along with a task parameter specifying the current task. Based on the task, the model uses the corresponding classification or regression layer to produce output. With the provided labels, the model calculates the appropriate loss function (crossEntropyLoss for classification (3)(4), MSELoss for regression (5)) for use during training. This architecture allows the model to learn shared general representations through the shared TinyBERT layer and then specialize for each task through separate output layers.

Algorithm 1 Multi-task TinyBERT Training

Require: Multi-task dataset $\mathcal{D}$ for SST-2, QQP, and STS-B; pre-trained TinyBERT model; number of tasks N_C; task-specific weights $W = \{w_{\text{sa}}, w_{\text{pd}}, w_{\text{sts}}\}$, initially set to 1.0; hyperparameters: number of epochs E, batch size B, embedding size f

Ensure: Trained multi-task TinyBERT model with shared backbone and task-specific heads; optimized combined loss function $\mathcal{L}_{\text{total}}$

1: Load pre-trained TinyBERT as shared backbone

2: Add task-specific heads:
   - SST-2: Linear + Softmax for binary classification ($y_{\text{sa}} \in \{0, 1\}$)
   - QQP: Linear + Softmax for binary classification ($y_{\text{pd}} \in \{0, 1\}$)
   - STS-B: Linear layer for regression ($y_{\text{sts}} \in [0, 5]$)

3: Load and preprocess dataset $\mathcal{D}$

4: Tokenize text using TinyBERT-Tokenizer

5: **for** epoch $e = 1$ to E **do**

6:     Divide $\mathcal{D}$ into mini-batches of size B

7:     **for** each mini-batch **do**

8:        Pass input through TinyBERT to obtain [CLS] embeddings

9:        Apply dropout to [CLS] embeddings

10:       Compute predictions from task-specific heads
$$\mathcal{L}_{\text{sa}} = -\frac{1}{B}\sum_{i=1}^{B}\sum_{c=0}^{1} y_{i,c} \log(P_{i,c})$$
$$\mathcal{L}_{\text{pd}} = -\frac{1}{B}\sum_{i=1}^{B}\sum_{c=0}^{1} y_{i,c} \log(P_{i,c})$$
$$\mathcal{L}_{\text{sts}} = \frac{1}{B}\sum_{i=1}^{B}(\text{score}_{\text{sts},i} - y_{\text{sts},i})^2$$

11:       Combined Loss:

12:         $\mathcal{L}_{\text{total}} = w_{\text{sa}}\mathcal{L}_{\text{sa}} + w_{\text{pd}}\mathcal{L}_{\text{pd}} + w_{\text{sts}}\mathcal{L}_{\text{sts}}$

13:       Compute gradients of $\mathcal{L}_{\text{total}}$

14:       Update model parameters using optimizer

15:     **end for**

16: **end for**

17: **return** Trained TinyBERT model, $\mathcal{L}_{\text{total}}$, task-specific head parameters

Algorithm 2 Fine-tuning TinyBERT for Specific NLP Tasks

Require: Pre-trained multi-task TinyBERT model, task-specific dataset D_{task}, loss function $\mathcal{L}_{\text{task}}$, hyperparameters: number of epochs E, batch size B

Ensure: Fine-tuned TinyBERT model optimized for the specific task

1: **Initialize:** Load pre-trained TinyBERT (shared backbone + task-specific head)

2: **Prepare Dataset:** Load and preprocess D_{task} using TinyBERT tokenizer to get input IDs and attention masks

3: **for** epoch $e = 1$ to E **do**

4:     Divide D_{task} into mini-batches of size B

5:     **for** each mini-batch **do**

6:        Forward pass:
   - Pass tokenized inputs through TinyBERT to get [CLS] embeddings
   - Feed embeddings to task-specific head to get predictions

7:        Compute loss $\mathcal{L}_{\text{task}}$:
   -SST-2 or QQP (classification):
   $\mathcal{L}_{\text{task}} = -\frac{1}{B}\sum_{i=1}^{B}\sum_{c=0}^{1} y_{i,c}\log(P_{i,c})$
   -STS-B (regression):
   $\mathcal{L}_{\text{task}} = \frac{1}{B}\sum_{i=1}^{B}(\text{score}_{\text{pred},i} - y_{\text{true},i})^2$

8:        Backward pass: Compute gradients of $\mathcal{L}_{\text{task}}$

9:        Update model parameters using optimizer

10:     **end for**

11: **end for**

12: **return** Fine-tuned TinyBERT model

3.3. Fine-tuning TinyBERT for Specific NLP Tasks

Algorithm 2 describes the process of fine-tuning TinyBERT for specific NLP tasks. After being pre-trained using MTL acquiring strong language understanding capabilities, TinyBERT is further

fine-tuned on the target task of interest. Here, we fine-tune the model for all three aforementioned tasks to compare its performance after fine-tuning with its performance immediately after multi-task training. We use the multi-task trained TinyBERT model as the starting point and further train it on data specific to each individual task, utilizing the corresponding loss function. The model's weights are updated to optimize performance for that task. This fine-tuning process allows the model to adapt the representations learned during the MTL phase to better fit the unique characteristics and requirements of each task, thereby enhancing overall accuracy and performance.

4. EXPERIMENTS

In this section, we present the experimental implementation based on the method described in Section 3. Our aim is to explore the effectiveness of using a small language model like TinyBERT in a multi-task learning setup for diverse NLP tasks.

4.1. Dataset

The GLUE benchmark (General Language Understanding Evaluation)[1] is a widely recognized collection of diverse NLP tasks designed to assess the language understanding capabilities of machine learning models. This benchmark comprises various tasks that test different aspects of natural language understanding. The study utilizes three datasets from GLUE benchmark:

- SST-2 (Stanford Sentiment Treebank)[2]: A sentiment classification dataset comprising sentences from movie reviews, labeled as either positive or negative. It includes approximately 67,000 sentences for training and about 872 sentences for development (validation) and 1,821 sentences for testing.

- QQP (Quora Question Pairs)[3]: A paraphrase detection dataset containing pairs of questions from the Quora website, labeled to indicate whether the two questions are semantically equivalent. This is a large dataset with over 363,780 question pairs for training and around 40,431 pairs for development and 390,965 pairs for testing. Question pairs can have different phrasing but still convey similar meanings, or vice versa. It uses binary labels: duplicate or non-duplicate.

- STS-B (Semantic Textual Similarity Benchmark)[4]: A semantic textual similarity dataset containing sentence pairs from various sources, labeled with scores from 1 to 5 indicating their degree of semantic similarity. It includes approximately 7,000 sentence pairs for training, 1,500 pairs for development, and 1,400 pairs for testing. Unlike SST-2 and QQP with discrete labels, STS-B uses a continuous scale, requiring the model to understand a more detailed level of semantic correlation.

The dataset is visualized in the following Fig 2.

---

[1] https://h2o.ai/wiki/glue/

[2] https://paperswithcode.com/sota/text-classification-on-glue-sst2

[3] https://paperswithcode.com/dataset/glue

[4] https://medium.com/nlplanet/two-minutes-nlp-glue-tasks-and-2022-leaderboard-517baedfa597
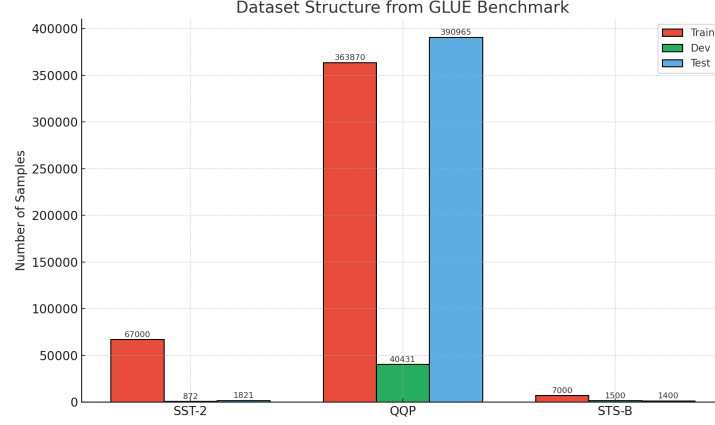
Figure 2. Dataset visualization.

4.2. Evaluation

We evaluate the model's performance on the validation sets for each task. For sentiment analysis and paraphrase detection, we use evaluate-classification:

$$\text{Accuracy} = \frac{\text{number of correct predictions}}{\text{total number of samples}} \quad (8)$$

And for semantic textual similarity, we use evaluate-STS, which calculates the Pearson correlation coefficient:

$$r = \frac{\sum(x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum(x_i - \bar{x})^2 \sum(y_i - \bar{y})^2}} \quad (9)$$

- x_i: Predicted score ($score_{si}$).

- y_i: Actual score (y_{si}).

- $\bar{x}, \bar{y}$: Mean values of predicted and actual scores.

4.3. Experimental Setup

The experiments were conducted using PyTorch and the Hugging Face Transformers li-brary on a single GPU. The pre-trained Tiny-BERT model "huawei-noah/TinyBERT-General-4L-312D"[7] was used as the backbone. The specific version TinyBERT-General-4L-312D refers to a TinyBERT model with a general configuration ("General"), comprising 4 Transformer layers and a hidden dimension size of 312 [7]. Combining MTL with fine-tuning of a compact model like TinyBERT-General-4L-312D promises to leverage the advantages of both approaches: the ability to learn rich representations from multiple tasks and the deployment efficiency of a small model. The batch size was set to 16. The AdamW optimizer[24] was employed with a learning rate of 1e-5 and a dropout rate of 0.01. The number of training epochs for both the initial multi-task training and individual fine-tuning for each task was 5. Early stopping was triggered if the validation score did not improve for 2 consecutive epochs. Mixed precision training was utilized to accelerate training and reduce memory consumption.

5. RESULTS AND DISCUSSION

The model's validation performance at different training stages is presented in Table 1 and visualized in the following Fig 3.

Table 1. Validation performance at different training stages

Task	Measure	After MTL	After FT
Sentiment analysis	Accuracy	0.8704	0.8922
Paraphrase detection	Accuracy	0.8737	0.8822
Semantic textual similarity	Pearson	0.8624	0.8690

Figure 3. Performance comparison across tasks.

The results show that initial multi-task training enables the model to achieve significant performance across all three tasks. After individual fine-tuning, the accuracy for sentiment and paraphrase tasks improved, while the Pearson correlation for the semantic textual similarity task also slightly increased.

The high performance achieved after the initial multi-task training phase indicates that the TinyBERT model can effectively learn shared representations beneficial for all three tasks: sentiment analysis, paraphrase detection, and semantic textual similarity. This suggests that some fundamental linguistic and semantic knowledge can be transferred among these tasks. Thus, the answer to **RQ1** is that TinyBERT can learn effectively from three tasks simultaneously, achieving good performance with a compact model.

The improvement in performance after individual fine-tuning for each task demonstrates that continuing to train the model on task-specific data allows it to adjust its weights to better capture the specific nuances of each task, leading to improved performance. From this, we derive the answer to **RQ2**: Individual fine-tuning after MTL is clearly effective, helping to optimize for each task and narrowing the gap with larger models.

The initial MTL phase appears to have facilitated some positive knowledge transfer among tasks, as evidenced by the reasonable performance achieved across all three tasks with a single shared model. The further improvements observed during individual fine-tuning indicate that additional specialization on each task is beneficial for maximizing performance.

Comparing performance across tasks, the paraphrase detection task achieved the highest accuracy in both the initial multi-task training phase and after individual fine-tuning. This could be due to the clear binary nature of the task (duplicate or non-duplicate) and the relatively large size of the QQP dataset. The semantic textual similarity task, being a regression task predicting a continuous score, had a slightly lower Pearson correlation compared to the accuracy achieved in the classification tasks.

Table 2. Comparing TinyBERT and BERT-base

Criteria	TinyBERT	BERT-base
Parameters	~14.5M	~110M
Transformer Layers	4	12
Hidden Dimensions	312	768
Inference Time	~4–6x faster	Standard
Memory Required	~1/7	High

Table 2 compares TinyBERT and BERT-base models, TinyBERT is significantly more compact than BERT-base, featuring 7.6 times fewer parameters. This leads to faster inference speeds and lower memory requirements.

We compare our method and results with related work on multi-task learning and TinyBERT on GLUE tasks: Table 3 compares performance.

Table 3. Comparing TinyBERT and BERT-base performance

Tasks	TinyBERT (MTL)	TinyBERT (FT)	BERT-base (Single)	BERT-base (MTL)
SST-2	0.8704	0.8922	0.927	0.913
QQP	0.8737	0.8822	0.913	0.904
STS-B	0.8624	0.8690	0.894	0.885

- BERT-base (single task): Devlin et al.[1], fine-tuned individually on each GLUE task.

- BERT-base (MTL): Liu et al.[8] multi-task trained on similar tasks.

TinyBERT achieves quite good performance (87-88%) even before fine-tuning, indicating that MTL helps leverage common knowledge among tasks. However, it is still about 3-4% lower than BERT-base MTL (90-91%), reflecting the difference in representational capacity due to model scale. This suggests that while TinyBERT can learn common features, it's not as robust as BERT-base due to its limited number of layers and parameters. After fine-tuning, TinyBERT improves by 1-2% across all three tasks, narrowing the gap with BERT-base. However, it does not surpass BERT-base.

Compared to BERT-base, multitask TinyBERT achieves 2-4% lower performance across all three tasks (SST-2, QQP, STS-B), even after individual fine-tuning. However, it excels in computational efficiency, being 7.6 times smaller and 4-6 times faster in inference. This makes TinyBERT a reasonable choice in scenarios where efficiency is prioritized over maximum performance, such as deployment on edge devices or in real-time applications. The combined MTL and FT approach of this study also shows potential in leveraging a single model for multiple tasks, reducing training costs compared to individually training separate BERT-base models. This directly answers **RQ3**: Our method achieves a reasonable balance between performance and computational efficiency, which is particularly useful in practical applications requiring resource savings.

6. RELATED WORK

Recent advances in NLP have explored both large-scale pre-trained models and compact architectures optimized for efficiency. In particular, MTL and model compression have emerged as two key strategies to improve performance and reduce computational costs.

Stickland and Murray introduced projected attention layers to adapt BERT for MTL in[8]. By adding lightweight task-specific adapters, they achieved strong performance on the GLUE benchmark while maintaining efficiency. Their results confirmed that parameter-efficient adaptation strategies can make large models more scalable across multiple tasks.

Jiao et al. proposed TinyBERT in[7], which applied a two-stage knowledge distillation approach first at the pre-training level and then at the task-specific fine-tuning level to compress BERT-base into a smaller model while retaining competitive performance. TinyBERT-4L, in particular, achieved over 96% of BERT-base's performance on GLUE with only about 13% of its parameters.

While TinyBERT was initially designed for single-task fine-tuning, later studies have extended its application to MTL. A notable example is the work by Yu et al.[25], which explored the effect of multi-task fine-tuning on small models such as Phi-3-Mini in the financial domain. Their study revealed that combining related tasks significantly improved performance even outperforming large models like GPT-4-o in domain-specific benchmarks.

In addition, Houlsby et al.[13] proposed adapter layers for parameter-efficient fine-tuning, enabling multitask training with minimal parameter overhead. The adapters can be selectively

trained for new tasks, making them particularly useful in low-resource and multi-domain settings.

Furthermore, a growing body of literature focuses on dynamic task weighting to improve MTL stability and mitigate negative transfer. For instance, Kendall et al.[27] proposed weighting losses based on task uncertainty, while Lakkapragada et al.[14] used exponential moving average loss strategies to maintain balance during training.

Some recent surveys, for example those by Zhang and Yu[10], Crawshaw[26], and Yu et al.[25], Liu et al.[28] have analyzed the training strategies and task relatedness in MTL. These works emphasize the importance of sharing knowledge across tasks, selecting compatible task combinations, and using dynamic weighting schemes to mitigate negative transfer.

7. CONCLUSION

In this study, we proposed a novel multi-task learning approach, which relies on a small language model like TinyBERT for NLP tasks, including: sentiment analysis, paraphrase detection, and semantic textual similarity. We first trained a shared TinyBERT model using a multi-task learning framework, then fine-tuned it individually for each task. The experimental results showed that the initial MTL phase enabled the model to capture shared representations that achieved strong performance across all tasks. Subsequent task-specific fine-tuning further improved accuracy and correlation scores, demonstrating the value of combining MTL with fine-tuning. Specifically, TinyBERT achieved competitive performance—while remaining significantly smaller and faster than BERT-base making it suitable for resource-constrained environments. These findings highlighted the potential of small language models like TinyBERT to serve as efficient and effective solutions for multi-task NLP scenarios. This approach also offers practical benefits in deployment settings, such as reduced memory usage and training costs, without a significant compromise in task performance.

In the future, we will explore different task balancing strategies in the initial MTL phase to potentially improve learned shared representations. Investigating other fine-tuning techniques

and hyper parameters for each task could also lead to better individually fine-tuned models, which could then benefit the aggregation process. Experimenting with more sophisticated methods to determine task similarity, possibly considering uncertainty or variance in individual model performance, could also be valuable. Evaluating the performance of all models on the held-out GLUE test sets would provide a more comprehensive assessment of their generalization capabilities.

REFERENCES

1. Devlin, J., Chang, M.-W., Lee, K., Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, 4171–4186. https://aclanthology.org/N19-1423/

2. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V. (2019). RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv preprint arXiv:1907.11692*. https://arxiv.org/abs/1907.11692

3. Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., Liu, P.J. (2020). Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research*, 21(140), 1–67. https://jmlr.org/papers/v21/20-074.html

4. Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R., Le, Q.V. (2019). XL-Net: Generalized Autoregressive Pretraining for Language Understanding. *Advances*

in Neural Information Processing Systems (NeurIPS).

5. Clark, K., Luong, M.-T., Le, Q.V., Manning, C.D. (2020). ELECTRA: Pre-training Text Encoders as Discriminators Rather Than Generators. *International Conference on Learning Representations (ICLR).*

6. Brown, T., Mann, B., Ryder, N., Subbiah, M., et al. (2020). Language Models are Few-Shot Learners. *Advances in Neural Information Processing Systems (NeurIPS).*

7. Jiao, X., Yin, Y., Shang, L., Jiang, X., Chen, X., Li, L., Wang, F., Liu, Q. (2020). TinyBERT: Distilling BERT for Natural Language Understanding. *Findings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 4163–4174. https://aclanthology.org/2020.findings-emnlp.372/

8. Liu, P., Qiu, X., Huang, X. (2019). Multitask Deep Neural Networks for Natural Language Understanding. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*, 4487–4496. https://aclanthology.org/P19-1441/

9. Arivazhagan, N., et al. (2019). Massively Multilingual Neural Machine Translation in the Wild: Findings and Challenges. *arXiv preprint arXiv:1903.00089.*

10. Zhang, Y., Yu, Z. (2023). A Survey of Multi-Task Learning in Natural Language Processing: Regarding Task Relatedness and Training Methods. *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*, 901–913. https://aclanthology.org/2023.eacl-main.66/

11. Raffel, C., et al. (2020). Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research*

12. Clark, K., Khandelwal, U., Levy, O., Manning, C.D. (2019). What Does BERT Look At? An Analysis of BERT's Attention. *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL).*

13. Houlsby, N., Giurgiu, A., Jastrzebski, S., Morrone, B., De Laroussilhe, Q., Gesmundo, A., Attariyan, M., Gelly, S. (2019). Parameter-Efficient Transfer Learning for NLP. *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2790–2799. https://proceedings.mlr.press/v97/houlsby19a.html

14. Lakkapragada, A., Gupta, V., Rajkumar, R. (2022). Mitigating Negative Transfer in Multi-Task Learning with Exponential Moving Average Loss Weighting Strategies. *arXiv preprint arXiv:2211.12999.* https://arxiv.org/abs/2211.12999

15. Ruder, S. (2017). An Overview of Multi-Task Learning in Deep Neural Networks. *arXiv preprint arXiv:1706.05098.* https://arxiv.org/abs/1706.05098

16. Caruana, R. (1997). Multitask Learning. *Machine Learning*, 28(1), 41–75. https://link.springer.com/article/10.1023/A:1007379606734

17. Ruder, S. (2019). Neural Transfer Learning for Natural Language Processing. PhD Thesis.

18. Sun, C., Qiu, X., Xu, Y., Huang, X. (2019). How to Fine-Tune BERT for Text Classification?. *Chinese Computational Linguistics (CCL).*

19. Howard, J., Ruder, S. (2018). Universal Language Model Fine-tuning for Text Classification. *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*, 328–339. https://aclanthology.org/P18-1031/

20. McCann, B., Bradbury, J., Xiong, C., Socher, R. (2017). Learned in Translation: Contextualized Word Vectors. *Advances in Neural Information Processing Systems (NeurIPS).*

21. Sanh, V., Debut, L., Chaumond, J., Wolf, T. (2019). DistilBERT, a Distilled Version of BERT: Smaller, Faster, Cheaper and Lighter. *arXiv preprint arXiv:1910.01108.*

22. Gunel, B., Du, J., Conneau, A., Stoyanov, V. (2021). Supervised Contrastive Learning for Pre-trained Language Model Fine-tuning. *International Conference on Learning Representations (ICLR)*. `https://arxiv.org/abs/2011.01403`

23. Khosla, P., et al. (2020). Supervised Contrastive Learning. *Advances in Neural Information Processing Systems (NeurIPS)*. `https://arxiv.org/abs/2004.11362`

24. Zhuang, Z., Liu, M., Cutkosky, A., Orabona, F. (2022). Understanding AdamW through Proximal Methods and Scale-Freeness. *arXiv preprint arXiv:2202.00089*. `https://arxiv.org/abs/2202.00089`

25. Yu, J., Dai, Y., Liu, X., Huang, J., Shen, Y., Zhang, K., et al. (2024). Unleashing the Power of Multi-Task Learning: A Comprehensive Survey Spanning Traditional, Deep, and Pretrained Foundation Model Eras. *arXiv preprint arXiv:2404.18961*. `https://arxiv.org/abs/2404.18961`

26. Crawshaw, M. (2020). Multi-Task Learning with Deep Neural Networks: A Survey. *arXiv preprint arXiv:2009.09796*. `https://arxiv.org/abs/2009.09796`

27. Kendall, A., Gal, Y., Cipolla, R. (2018). Multi-task Learning Using Uncertainty to Weigh Losses for Scene Geometry and Semantics. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 7482–7491. `https://openaccess.thecvf.com/content_cvpr_2018/papers/Kendall_Multi-Task_Learning_Using_CVPR_2018_paper.pdf`

28. Liu, H., Shen, Y., Jin, D., et al. (2021). Multi-Task Learning for Natural Language Processing: A Survey. *ACM Computing Surveys*.

29. Aghajanyan, A., et al. (2021). Intrinsic Dimensionality Explains the Effectiveness of Language Model Fine-Tuning. *International Conference on Learning Representations (ICLR)*.