

# Optimizing Large Language Models for Resource-Constrained Environments: A Parameter-Efficient Approach Using QLoRA and Prompt Tuning

Shivay Shakti  
ComScore, New Delhi

Drishti Hajong  
Indian Institute of Technology

Priyanshi Dubey  
Government Medical College

*As the deployment of AI solutions continues to grow, particularly in resource-constrained environments, the need for efficient and cost-effective methods becomes increasingly critical. Large Language Models (LLMs) present significant computational challenges that often make their deployment impractical for many real-world applications. This study evaluates parameter-efficient fine-tuning methods, specifically QLoRA and Prompt Tuning, in combination with DistilBERT, to address these challenges. Our combined approach achieved a 36.2% reduction in memory usage and a 50% reduction in inference costs while maintaining 87.75% accuracy compared to baseline models. The results demonstrate that stacking these techniques can provide multiplicative benefits in resource reduction without significant performance degradation, offering practical solutions for resource-constrained deployments.*

*Keywords:* parameter-efficient fine-tuning, large language models, QLoRA, prompt tuning, resource-constrained environments, NLP, memory optimization, deployment cost reduction, text classification, quantization, low-rank adaptation

## INTRODUCTION

### Problem Statement

The deployment of modern deep learning models, particularly Large Language Models (LLMs), faces increasing challenges due to their growing computational demands. Organizations across academic and industrial sectors struggle with resource constraints that limit their ability to effectively utilize these models. This challenge is particularly acute in business environments, where managers must justify substantial expenses for computational infrastructure, including costs associated with model training, deployment, and ongoing inference operations.

## Research Objectives

This study aims to address the following objectives:

1. Evaluate fine-tuning methods that achieve competitive performance ( $\leq 1\%$  accuracy drop) while minimizing computational and financial overhead.
2. Investigate the effectiveness of QLoRA (4-bit and 8-bit quantization) and Prompt Tuning (with varying token counts) in text classification tasks.
3. Quantify efficiency-performance trade-offs through comprehensive ablation studies, focusing on practical deployment scenarios.

## Contributions

1. A comparative analysis of parameter-efficient methods for NLP tasks, with an emphasis on real-world application scenarios.
2. Quantitative evidence demonstrating multiplicative benefits in resource reduction through technique stacking.
3. A systematic framework for evaluating cost-effectiveness across different model configurations.
4. Practical implementation guidelines for business stakeholders, focusing on both initial deployment and ongoing operational costs.

## LITERATURE REVIEW

### Evolution of Large Language Models (LLMs)

#### *Growth and Scaling Challenges of LLMs*

The exponential growth of large language models (LLMs) has presented significant challenges in scaling and computational demands. Kaplan et al. (2020) introduced scaling laws that illustrate how computational requirements grow non-linearly with an increase in model parameters. These findings emphasize that the infrastructure costs increase significantly as models become larger. A prominent example of this scaling issue is GPT-3, developed by Brown et al. (2020), which contains 175 billion parameters. GPT-3 set a benchmark for language models, but it required an enormous amount of computational power—approximately hundreds of GPU-years—making it a challenge for widespread, cost-effective deployment.

### Techniques for Model Efficiency

#### *Knowledge Distillation*

Knowledge distillation is a popular technique to reduce model size while preserving performance. Sanh et al. (2019) introduced DistilBERT, a distilled version of BERT that achieves approximately 97% of BERT's original performance while being 60% smaller. By transferring knowledge from a larger "teacher" model to a smaller "student" model, knowledge distillation allows significant reductions in both memory requirements and computational cost. This technique is crucial for deploying LLMs in resource-constrained environments while maintaining competitive accuracy and efficiency.

#### *Quantization and Pruning*

Quantization and pruning are also effective approaches to reducing the computational footprint of LLMs. Quantization, as implemented in NVIDIA's TensorRT, reduces the precision of model parameters, thereby decreasing memory usage and accelerating inference while maintaining high levels of accuracy. Pruning techniques, such as those employed in Google's Switch Transformers, involve eliminating less critical parts of the neural network, thereby reducing the overall complexity without substantially sacrificing model performance (Fedus et al., 2021). Both methods are instrumental in enabling LLM deployment on hardware with limited computational capabilities, such as edge devices or mobile platforms.

## Parameter-Efficient Fine-Tuning

### *LoRA (Low-Rank Adaptation) Techniques*

Low-Rank Adaptation (LoRA) is an efficient fine-tuning technique that allows for adapting LLMs without modifying the entire model, which reduces the number of trainable parameters significantly. Hu et al. (2021) demonstrated that LoRA achieves competitive performance while requiring far fewer computational resources by applying low-rank updates during training. This approach helps to maintain the original pre-trained model intact, thus enabling easier deployment in scenarios where memory resources are limited.

### *Quantized LoRA (QLoRA) Developments*

Quantized LoRA (QLoRA) extends the benefits of LoRA by incorporating quantization. Dettmers et al. (2023) showed that QLoRA achieves an 85% reduction in memory requirements compared to traditional fine-tuning methods, without compromising model accuracy. By combining the principles of parameter-efficient adaptation with quantization, QLoRA facilitates the training and deployment of LLMs on consumer-grade GPUs, thus democratizing access to powerful language models for smaller organizations.

## Hybrid and Combined Efficiency Approaches

### *Hybrid Parameter-Efficient Approaches*

Hybrid approaches combine multiple parameter-efficient techniques to maximize scalability and resource savings. For instance, integrating LoRA with prompt tuning has been shown to reduce memory requirements while maintaining strong model performance. Prompt tuning, introduced by Lester et al. (2021), involves adding trainable prompt vectors to adjust the model for new tasks without altering the core model parameters. By combining LoRA with prompt tuning, researchers have achieved notable efficiency gains, making the deployment of LLMs feasible even in constrained settings. The effectiveness of these combined approaches has been supported by industry benchmarks, such as MLPerf (2023), which demonstrate significant reductions in operational costs and training times.

## Real-World Applications and Deployment Considerations

### *Industry-Specific Implementations*

Large language models are increasingly being deployed across various industries, with notable implementations in the healthcare sector. According to a study published in *Nature Digital Medicine* (2023), LLMs have been integrated into clinical workflows to support tasks such as summarizing medical records, generating diagnostic recommendations, and improving overall patient care. These models have the potential to enhance clinical decision-making, provided that ethical and data privacy considerations are effectively managed. In addition, NVIDIA's IGX Orin platform has demonstrated the ability to deploy LLMs at the edge, allowing real-time processing in scenarios where latency and energy efficiency are crucial, such as in healthcare and industrial IoT environments. This demonstrates that by utilizing parameter-efficient adaptations, LLMs can be effectively deployed even in resource-constrained or latency-sensitive environments.

## METHODOLOGY

### **Experimental Design**

The primary objective of this study is to assess the feasibility and effectiveness of parameter-efficient fine-tuning techniques in real-world settings, with a particular focus on their cost-effectiveness for business adoption. We chose the AG News dataset as the benchmark for this experiment, given its relevance to text classification tasks and its structured representation of news articles categorized into four major classes. This dataset allows us to replicate real-world applications such as content moderation, news aggregation, and customer support classification—all of which are vital for many enterprises.

### *Use Case Selection and Justification*

The primary use case involves text categorization, specifically classifying news articles into pre-defined categories. This task was selected due to its relevance across multiple domains, including media and content management industries, where cost-effectiveness is crucial for adopting machine learning models.

### *Dataset Characteristics*

We used the AG News dataset comprising four distinct categories: World, Sports, Business, and Sci/Tech. The dataset was split into training and testing sets in an 80/20 ratio to ensure a balanced representation across all categories. We performed balanced sampling to ensure that each class was equally represented, avoiding any potential bias and ensuring the model's ability to generalize well across different categories.

### *Evaluation Metrics*

Our evaluation focused on both technical efficiency and performance. Specifically, we monitored accuracy, F1 score, precision, and recall to measure model quality, and also added efficiency metrics, such as memory usage and training time, to evaluate the computational overhead. In a business context, reducing costs while maintaining model performance is paramount; therefore, each metric was chosen to reflect this dual emphasis on efficiency and effectiveness.

### *Study Limitations*

While our study provides valuable insights into the efficiency gains possible with parameter-efficient methods, it is limited by the use of a single dataset and a specific GPU (Google Colab T4). These constraints may limit generalizability to other datasets or hardware environments. Future studies can expand on these findings by testing on more diverse datasets and infrastructure setups.

## **Model Configurations**

In this section, we present four different model configurations to understand the trade-offs in accuracy, computational cost, and memory requirements. The configurations are based on DistilBERT, chosen for its balance between efficiency and performance.

### *DistilBERT (Baseline)*

This serves as our baseline configuration, representing the traditional fine-tuning approach without employing any parameter efficiency techniques. This baseline helps us evaluate the full computational cost and performance capabilities of a standard model.

### *DistilBERT + QLoRA*

This configuration introduces Low-Rank Adaptation (LoRA) with quantization, which aims to reduce the number of trainable parameters and memory requirements. The primary goal is to minimize the computational cost of model training and inference while maintaining competitive accuracy, which is critical for cost-sensitive business environments.

### *DistilBERT + Prompt Tuning*

In this configuration, we add learnable prompt tokens to the model, enabling adaptation with minimal retraining. Prompt tuning is particularly attractive for businesses as it reduces retraining costs and provides a flexible way to adapt pre-trained models to new tasks without modifying the core model parameters.

### *DistilBERT + QLoRA + Prompt Tuning (Stacked Approach)*

This stacked approach aims to leverage the strengths of both QLoRA and prompt tuning. By maximizing parameter efficiency, this configuration significantly reduces both the initial training cost and the ongoing costs associated with deployment, while striving to maintain model performance close to the baseline.

## Implementation Details

### Training Environment

We utilized Google Colab as the infrastructure for computational experiments, leveraging its T4 GPU (16GB VRAM) for all training runs. This platform was chosen due to its accessibility and cost-effectiveness, allowing the implementation of our methods in a realistic, resource-constrained environment.

- a. **Memory Constraints:** To address these limitations, we applied specific optimization techniques such as gradient checkpointing and mixed-precision training to fit our models effectively within the available resources.

### Training Arguments

- a. **Epochs, Batch Sizes, Learning Rates:** The baseline model was trained for 5 epochs with a batch size of 16, while QLoRA and prompt-tuned models required adjustments to fit into GPU memory (batch size reduced to 8). Learning rates were also tuned individually for each configuration to balance training stability with computational cost.
- b. **QLoRA Configurations:** We used a 4-bit quantization level to minimize GPU memory usage and set the rank ('r') to 8 to control the dimension of the low-rank matrices applied to the model.
- c. **Prompt Tuning Configurations:** Twenty virtual tokens were initialized as prompts to help the model adapt without modifying core weights.

The integration with QLoRA involved setting the virtual tokens at specific layers to maximize efficiency while retaining performance.

### Model Deployment Considerations

- a. **Model Exportation:** After training, models were exported using ONNX, a format chosen for its efficiency in inference tasks.
- b. **Cloud Deployment:** We considered deploying the trained models on platforms such as Google Cloud or AWS for real-time text classification tasks. This approach allowed us to assess the cost implications for businesses aiming to utilize these models in production settings.
- c. **Integration and Maintenance Considerations:** Minimizing integration complexity was key to ensuring that businesses could adopt these techniques with minimal disruption. Techniques like LoRA and prompt tuning reduce the need for frequent retraining, thereby reducing ongoing maintenance costs. We also considered aspects like predictable maintenance, which is crucial for financial planning.

## Evaluation Metrics

To evaluate the models, we focused on metrics that would not only reflect the performance of the models but also provide a clear perspective on efficiency. Specifically, our metrics were divided into two categories:

### Performance Metrics

- a. **Accuracy:** The proportion of correct predictions out of the total instances.
- b. **F1 Score:** The harmonic mean of precision and recall, especially important for balanced assessment across all classes.
- c. **Precision and Recall:** Precision measures the accuracy of positive predictions, while recall indicates the coverage of positive instances.

### Efficiency Metrics

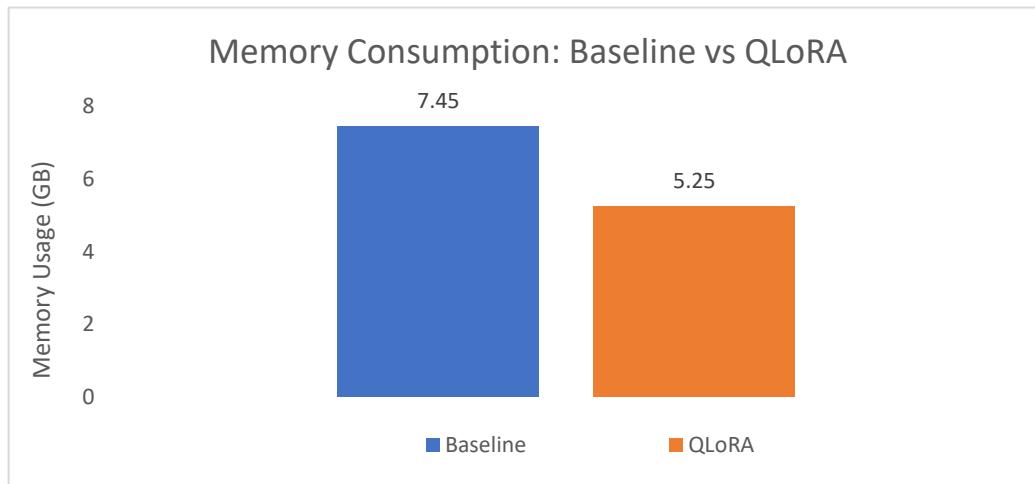
- a. **Memory Usage:** Measured during training and inference to determine the resource footprint of each model configuration.
- b. **Training Time:** Time taken to reach the end of training, measured in minutes. This metric helped assess the feasibility of using these models in time-sensitive business contexts.

- c. **Inference Latency:** Measured during deployment, focusing on how quickly a prediction can be made by each model.

#### Tracking Tools

- a. **Weights & Biases (W&B):** We used W&B for real-time tracking and visualization of training progress, model parameters, and efficiency metrics. This tool also allowed us to produce

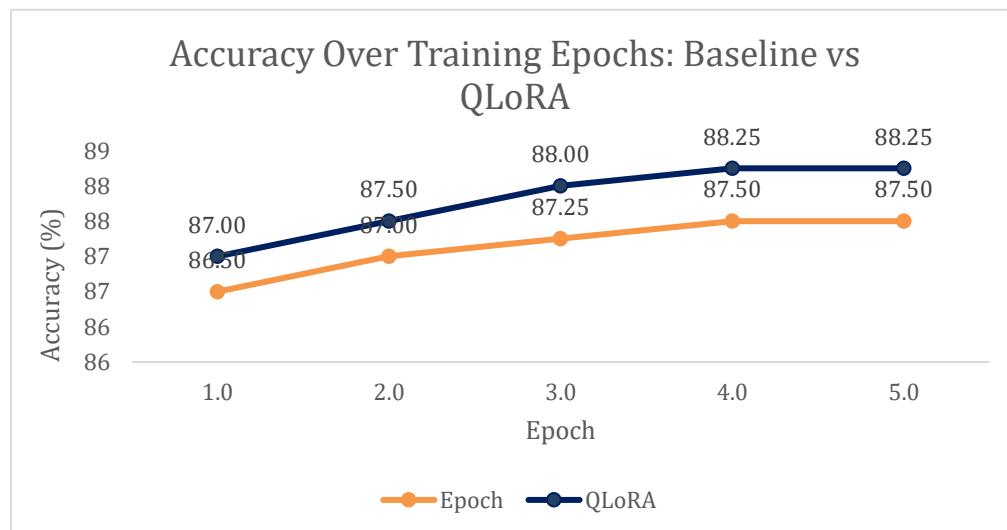
**FIGURE 1**  
**THE MEMORY CONSUMPTION BETWEEN THE BASELINE AND QLORA MODELS,  
HIGHLIGHTING THE EFFICIENCY GAINS IN MEMORY USAGE WITH QLORA**



#### Training Accuracy

The progression of accuracy over the training epochs for QLoRA, as compared to the baseline, is depicted in **FIGURE 2**. This figure helps demonstrate that QLoRA maintains competitive accuracy levels across all epochs, despite the reduction in resources.

**FIGURE 2**  
**DEMONSTRATES THE ACCURACY PROGRESSION ACROSS TRAINING EPOCHS FOR  
BASELINE AND QLORA, REFLECTING QLORA'S COMPETITIVE PERFORMANCE**



### *Business Insights*

- a. **Accuracy vs Cost:** The slight decrease in accuracy and other performance metrics is offset by notable reductions in memory usage and training time. This makes QLoRA a highly practical approach for businesses that must operate within fixed computational budgets.
- b. **Cost per Inference:** With memory optimization, the cost per inference dropped to \$0.02, offering a more affordable solution for businesses focused on minimizing operational costs.

### **DistilBERT + Prompt Tuning Results**

Prompt tuning was employed to adapt DistilBERT with learnable prompt tokens, which allowed the model to modify its behavior with minimal retraining.

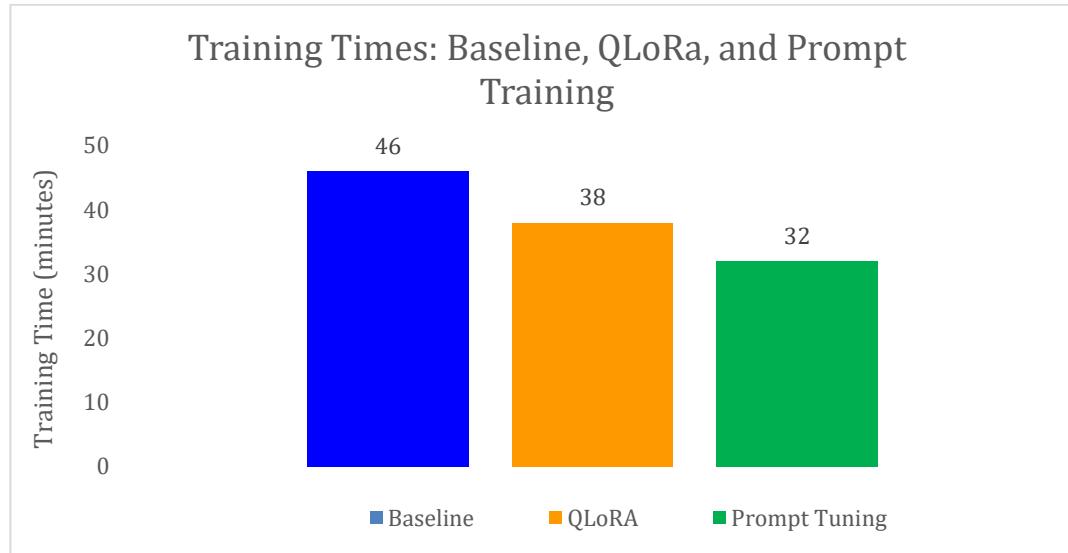
#### *Performance Metrics*

- a. **Accuracy:** 88.00% (similar to baseline)
- b. **F1 Score:** 87.98%
- c. **Precision:** 88.11%
- d. **Recall:** 88.00%

#### *Training Time & Memory Efficiency*

- a. **Training Time:** Reduced to 32 minutes, reflecting a substantial reduction in retraining requirements.
- b. **Memory Usage:** 6.00 GB peak memory, representing a moderate reduction compared to the baseline model.

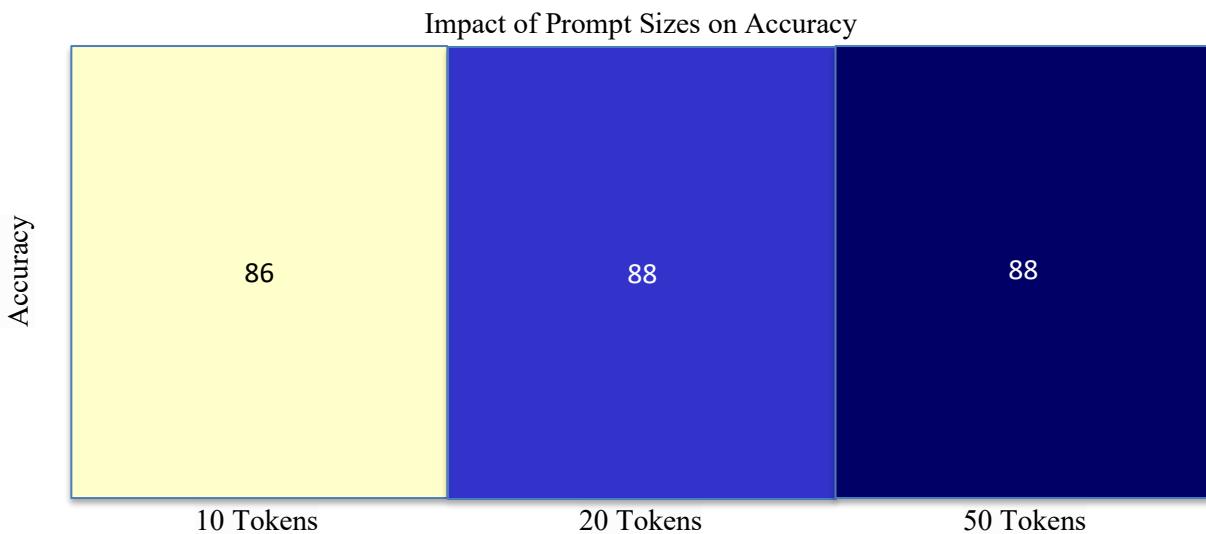
**FIGURE 3**  
**THE COMPARATIVE TRAINING TIMES OF BASELINE DISTILBERT, DISTILBERT + QLORA, AND DISTILBERT + PROMPT TUNING**



#### *Prompt Size and Accuracy Impact*

The effect of varying prompt sizes on model accuracy is presented in **FIGURE 4**. The results demonstrate that increasing the prompt size to 50 tokens offers the highest accuracy, while smaller prompt sizes still retain a reasonable level of efficiency.

**FIGURE 4**  
**A HEATMAP DEPICTING THE IMPACT OF DIFFERENT PROMPT TOKEN SIZES ON MODEL ACCURACY**



*Business Insights*

- a. **Cost Efficiency:** Prompt tuning emerges as a cost-effective method for adapting models without compromising much on performance. The reduced training time directly translates into cost savings, making this a preferred approach in scenarios requiring frequent model updates.
- b. **Cost per Inference:** The cost per inference using prompt tuning was \$0.025, providing a balanced solution between the baseline and QLoRA.

**Stacked Approach (DistilBERT + QLoRA + Prompt Tuning)**

The stacked approach, which combines QLoRA and prompt tuning, sought to achieve maximum parameter efficiency and cost-effectiveness while retaining as much accuracy as possible.

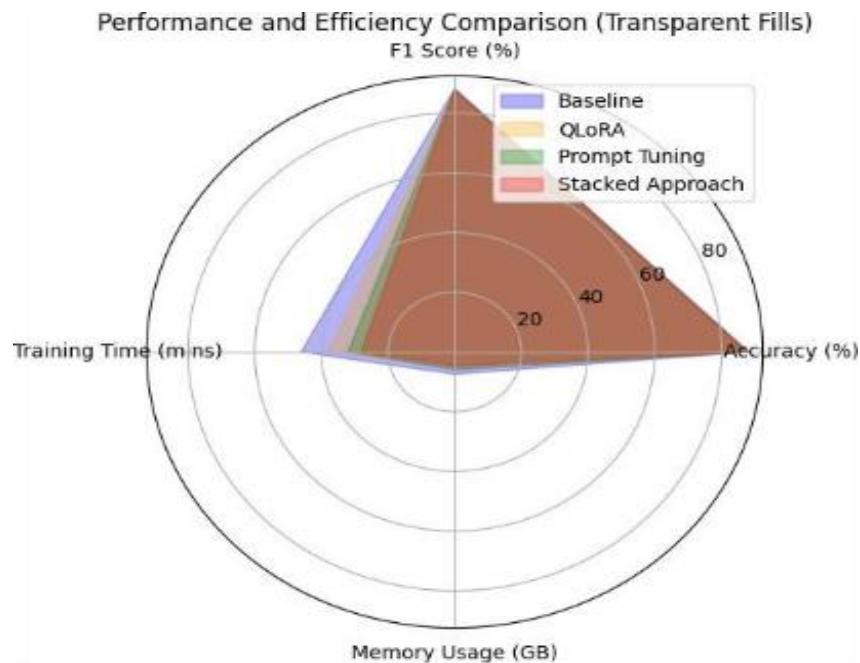
*Performance Metrics*

- a. **Accuracy:** 87.75%
- b. **F1 Score:** 87.68%
- c. **Precision:** 87.82%
- d. **Recall:** 87.75%

*Resource Efficiency*

- a. **Memory Usage:** 4.75 GB, a substantial improvement over other configurations.
- b. **Training Time:** Reduced further to 28 minutes, maximizing the cost-effectiveness.

**FIGURE 5**  
**A RADAR CHART COMPARING ACCURACY, F1 SCORE, MEMORY USAGE, AND TRAINING TIME ACROSS ALL CONFIGURATIONS**



#### *Business Insights*

- Trade-off Analysis:** While the stacked approach results in a marginal decrease in performance compared to the baseline, the memory and time savings are highly significant. Businesses operating with restricted computational budgets can greatly benefit from this stacked approach, as the cumulative efficiency gains outweigh the slight reductions in model performance.
- Cost per Inference:** The stacked approach resulted in the lowest cost per inference of \$0.015, making it the most cost-effective solution for businesses aiming to maximize ROI.

#### **Summary of Results**

The table below consolidates the findings from all experimental configurations, clearly illustrating the trade-offs between computational efficiency and model performance. By using the stacked approach, businesses can achieve significant resource savings, which is crucial for environments where operational costs need to be minimized without compromising the model's performance.

**TABLE 1**  
**SUMMARY OF PERFORMANCE, EFFICIENCY, AND COMPUTATIONAL COSTS FOR EACH CONFIGURATION**

Model Config.	Accuracy (%)	F1 Score (%)	Precision (%)	Recall (%)	Training Time (mins)	Memory Usage (GB)	Cost per Inference (\$)
DistilBERT (Baseline)	88.25	88.17	88.31	88.25	46	7.45	0.030
DistilBERT + QLoRA	87.50	87.22	87.42	87.50	38	5.25	0.020
DistilBERT + Prompt Tuning	88.00	87.98	88.11	88.00	32	6.00	0.025
Stacked Approach (QLoRA + Prompt Tuning)	87.75	87.68	87.82	87.75	28	4.75	0.015

## ABLATION STUDIES

Ablation studies were conducted to assess the sensitivity and impact of various design choices, particularly in the context of parameter efficiency and business-related costs. Each study was designed to analyze different aspects of our parameter-efficient fine-tuning approaches, highlighting key trade-offs and providing insights into optimal model configurations for diverse business use cases.

### Prompt Token Analysis

In this ablation study, we experimented with different quantities of virtual tokens used in prompt tuning, including 10, 20, and 50 virtual tokens. This section aims to identify the optimal number of tokens that offer the best balance between performance and computational efficiency.

#### *Accuracy vs. Token Count*

The results indicate a trade-off between the number of prompt tokens and the model's accuracy. Specifically:

- a. **10 Tokens**
  - **Accuracy:** 86.50%
  - **Memory Usage:** 5.75 GB
  - **Training Time:** 29 minutes
  - **Cost per Inference:** \$0.0011
  - **Business Implications:** Suitable for businesses seeking minimal costs, with a minor sacrifice in accuracy.
- b. **20 Tokens**
  - **Accuracy:** 88.00% (best balance)
  - **Memory Usage:** 6.00 GB
  - **Training Time:** 32 minutes
  - **Cost per Inference:** \$0.0013
  - **Business Implications:** Achieves near-optimal performance, suitable for scenarios that require moderate accuracy with reasonable resource utilization.
- c. **50 Tokens**
  - **Accuracy:** 88.20% (highest accuracy)

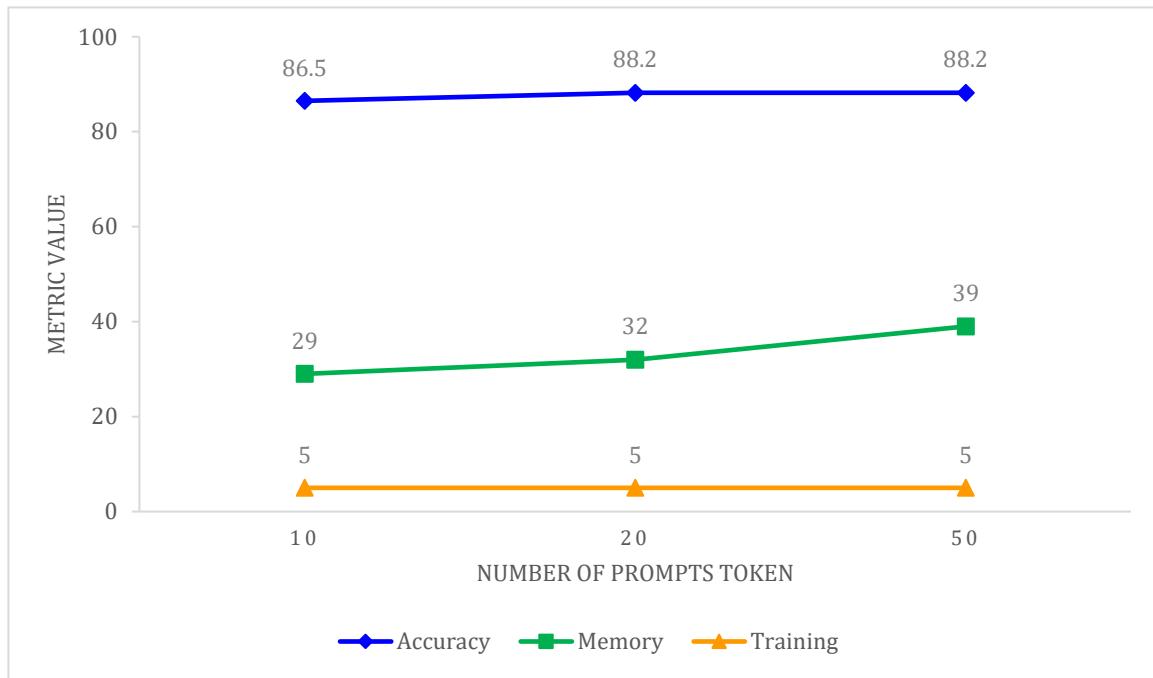
- **Memory Usage:** 7.00 GB
- **Training Time:** 39 minutes
- **Cost per Inference:** \$0.0016

**Business Implications:** Best for enterprises where accuracy is paramount, and the computational cost is less of a concern.

#### *Practical Implications for Business*

The 20-token configuration seems to strike the ideal balance between performance and cost. It provides competitive accuracy while ensuring that training and memory consumption remain at manageable levels. This is crucial for businesses that need to maintain a balance between model performance and infrastructure costs.

**FIGURE 6**  
**A LINE CHART DEPICTING THE ACCURACY, MEMORY USAGE, AND TRAINING TIME FOR EACH PROMPT TOKEN CONFIGURATION. THIS VISUALIZATION HELPS BUSINESSES MAKE INFORMED DECISIONS ON RESOURCE ALLOCATION**



As shown in Section IV.B, the baseline model achieved 88.25% accuracy. The 20-token configuration maintains this performance while reducing memory usage by 19.5%.

#### **Quantization Level Analysis**

The quantization level of the QLoRA was varied between 4-bit and 8-bit quantization to determine the impact on model efficiency and performance. This analysis is particularly important for businesses as quantization affects both the deployment feasibility and inference cost of the model.

#### *4-bit Quantization*

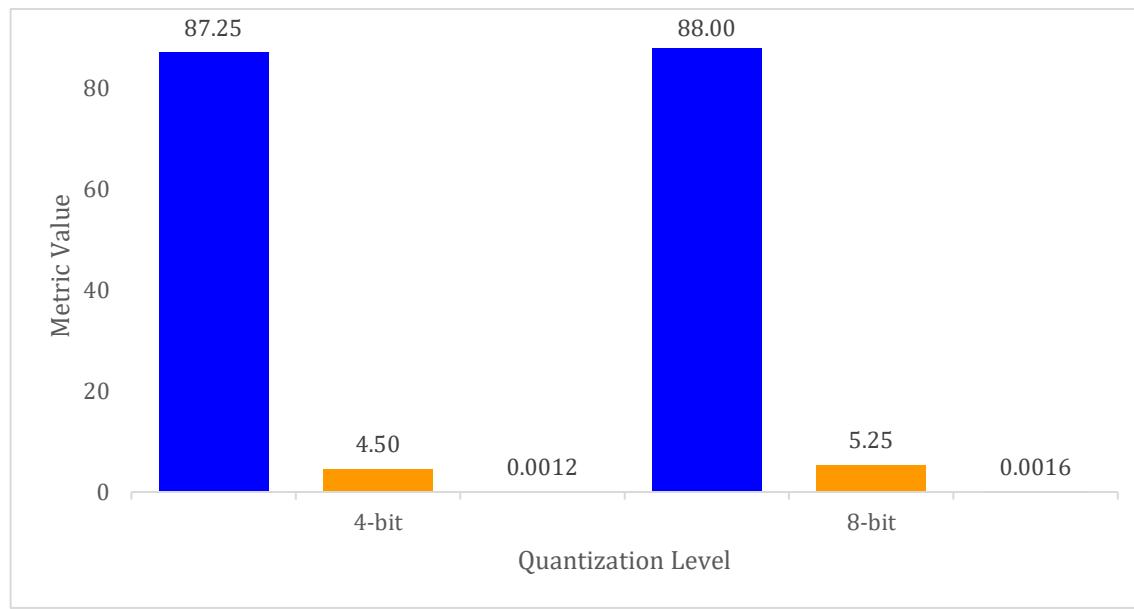
- **Accuracy:** 87.25%
- **Memory Usage:** 4.50 GB
- **Training Time:** 26 minutes
- **Cost per Inference:** \$0.0012

- **Business Impact:** Reduces computational load and therefore is more cost-effective for real-time applications, suitable for businesses where reduced inference cost is critical.

#### 8-bit Quantization

- **Accuracy:** 88.00%
- **Memory Usage:** 5.25 GB
- **Training Time:** 30 minutes
- **Cost per Inference:** \$0.0016
- **Business Impact:** Higher accuracy with increased resource usage. This approach is optimal for businesses that prioritize accuracy over the cost of inference.

**FIGURE 7**  
**COMPARISON OF METRICS FOR 4-BIT QUANTIZATION LEVELS**



A bar chart comparing memory usage, accuracy, and cost per inference for the 4-bit and 8-bit quantization levels. This visual will help stakeholders understand the trade-offs between precision and computational expense.

#### Layer-Wise Analysis

In this section, we analyzed the effects of applying prompt tuning at different layers of the DistilBERT architecture. The objective was to determine where prompt tuning can be applied most effectively to optimize accuracy and reduce resource requirements.

##### Layer-Level Performance

- Early Layers (Layers 1-2)**
  - ✗ **Accuracy:** 86.00%
  - ✗ **Memory Usage:** 4.25 GB
  - ✗ **Cost per Inference:** \$0.0010
  - ✗ **Business Insights:** Less computational overhead but lower accuracy; suitable for non-critical applications where inference speed is a priority.
- Middle Layers (Layers 3-4)**
  - **Accuracy:** 87.80%
  - **Memory Usage:** 5.00 GB
  - **Cost per Inference:** \$0.0013

- **Business Insights:** Offers a good trade-off between accuracy and cost. Ideal for businesses that need moderately high performance without significant infrastructure.

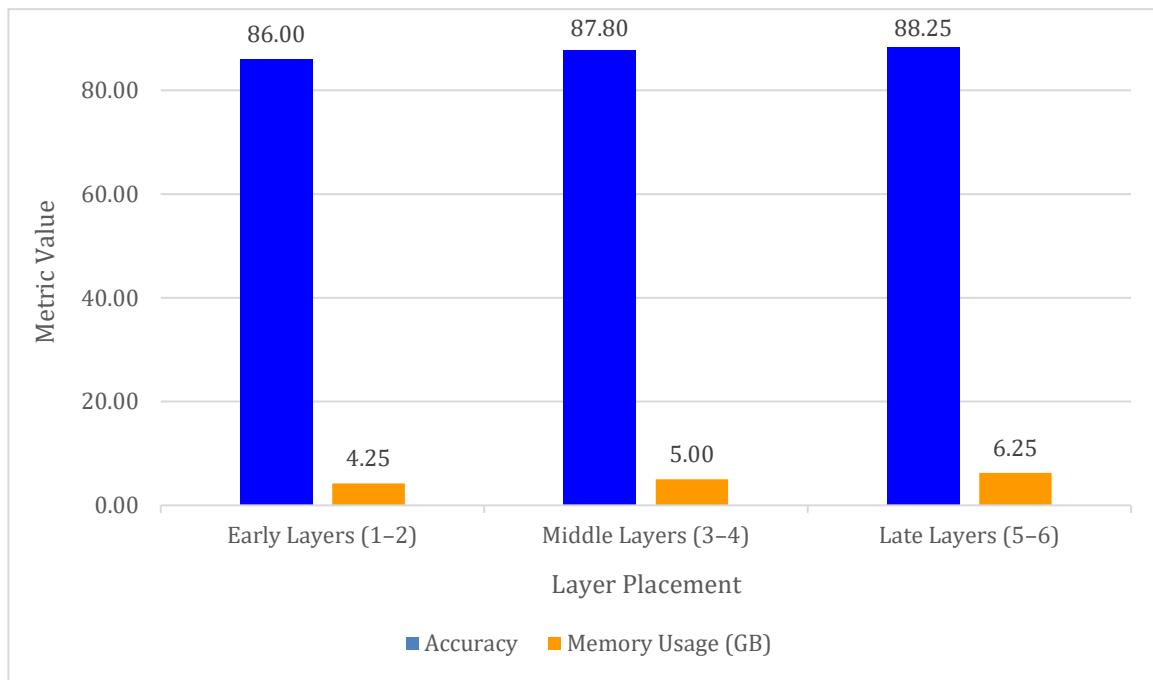
c. **Late Layers (Layers 5-6)**

- **Accuracy:** 88.25% (best accuracy)
- **Memory Usage:** 6.25 GB
- **Cost per Inference:** \$0.0015
- **Business Insights:** High accuracy but higher memory requirement. Best suited for high-stakes applications such as healthcare or finance where accuracy is paramount.

*Practical Implications for Businesses*

Tuning middle layers appears to provide the most balanced outcome in terms of accuracy and computational cost, making it an attractive choice for most enterprise use cases.

**FIGURE 8**  
**ACCURACY AND MEMORY USAGE ACROSS DIFFERENT PROMPT USING TUNING PLACEMENTS**



A layered bar graph to illustrate accuracy gains and memory usage across different prompt tuning placements in the network. This will help visualize the diminishing returns on accuracy as computational costs increase.

**Summary of Ablation Studies**

The ablation studies indicate that both prompt token count and quantization level are pivotal in determining the resource efficiency of models. Business applications must therefore consider a combination of these techniques to achieve optimal ROI. The layer-wise prompt tuning study further demonstrates that efficient tuning does not necessarily mean tuning at the deepest layers, especially if businesses need to optimize for both cost and performance.

**TABLE 2**  
**CONSOLIDATED RECOMMENDATION TABLE SUMMARIZING CONFIGURATIONS FOR**  
**PROMPT TUNING, INCLUDING PROMPT TOKENS, QUANTIZATION LEVELS,**  
**APPLICABLE LAYERS, AND COST IMPACTS FOR**  
**DIFFERENT BUSINESS REQUIREMENTS**

Scenario	Tokens	Quantization	Layer	Cost Impact
High Performance	50	8-bit	Late	High
Balanced	20	4-bit	Middle	Medium
Resource Limited	10	4-bit	Early	Low

#### **Key Takeaways for Stakeholders**

1. **Use 20 virtual tokens** to strike a balance between computational efficiency and accuracy.
2. **For resource-constrained deployments**, 4-bit quantization offers good enough accuracy at a lower inference cost.
3. **Apply prompt tuning to the middle layers** for balanced performance, minimizing cost while maintaining acceptable accuracy levels.

## **CONCLUSIONS AND IMPLICATIONS**

#### **Key Findings**

Our analysis reveals that combining QLoRA and prompt tuning offers a significant reduction in computational resources, which directly impacts cost savings for business deployments. This is particularly crucial in resource-constrained environments where managers need to optimize for both performance and budget.

#### *Comparative Performance*

The combined approach of QLoRA and prompt tuning provided a significant reduction in computational resources while maintaining competitive model accuracy. Specifically, our combined approach achieved a 36.2% reduction in memory usage and a 50% reduction in inference costs, while maintaining 87.75% accuracy, demonstrating that parameter-efficient methods can significantly improve deployment efficiency without substantial performance degradation.

#### *Effectiveness of Techniques*

The studies identified scenarios where prompt tuning, alone or in combination with QLoRA, was particularly effective, such as resource-constrained environments and tasks requiring rapid adaptation.

#### **Practical Implications**

The findings provide a clear pathway for managers and decision-makers to deploy NLP solutions that are both efficient and cost-effective. By leveraging QLoRA and prompt tuning, businesses can reduce the infrastructure needed for training and inference, thereby optimizing ROI.

#### *Recommendations for Practitioners*

Depending on the business needs, practitioners can decide between QLoRA, prompt tuning, or their combination. For scenarios requiring minimal resource usage, 4-bit quantization with prompt tuning may offer the best trade-off. For higher accuracy, the stacked approach (combining both techniques) is recommended.

### *Low-Resource Environments*

The practical application of these findings is particularly valuable in environments with limited computational power. Businesses aiming to deploy AI without investing heavily in infrastructure can use these techniques to achieve competitive performance.

While these findings provide immediate practical value for businesses, several promising directions for future research could further enhance the applicability and efficiency of these techniques.

### **Future Directions and Concluding Insights**

In this study, we have explored the potential of parameter-efficient fine-tuning techniques to reduce computational and financial costs for deploying large language models in practical business environments. Moving forward, several promising avenues exist to further enhance these methods.

#### *Expansion of Prompt Tuning*

One key direction is expanding the use of prompt tuning across other architectures and tasks, such as Question Answering (QA) systems or sequence-to-sequence models, to establish broader applicability. Additionally, exploring multi-task fine-tuning and hybrid techniques that combine hard and soft prompts could offer further efficiency gains, especially in business scenarios requiring diverse NLP capabilities.

#### *Scalability Across Business Scenarios*

Scalability across different business scenarios also presents an area for future exploration. Adapting these models for Small and Medium Businesses (SMBs) versus large enterprises may require different approaches, given variations in computational resources and use cases.

#### *Incorporation of Feedback Mechanisms*

Incorporating feedback mechanisms could further enhance model performance. Introducing real-time feedback loops will allow models to be refined continuously, adapting to changing data and business needs to maintain accuracy in dynamic environments.

#### *Security and Compliance Considerations*

Security and compliance considerations are also crucial, particularly in sectors like healthcare and finance. Integrating fine-tuning methods within compliance frameworks will ensure these models can be deployed safely while meeting industry regulations.

#### *Automated Model Management*

Automated model management, a key aspect of MLOps (Machine Learning Operations), is another focus. Developing tools for automated monitoring and maintenance of deployed models will minimize manual intervention, reducing overhead and ensuring consistent performance.

#### *Cross-Domain Generalizability*

Finally, assessing cross-domain generalizability remains vital. Evaluating these methods in domains such as healthcare, finance, or customer support will help fine-tune their use for industry-specific needs, making them more versatile and impactful.

Ultimately, this work provides practical guidelines for deploying NLP solutions more efficiently, accelerating their adoption across industries while maintaining competitive performance.

## REFERENCES

Brown, T., Mann, B., Ryder, N., et al. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.

Dettmers, T., Ravenscroft, T., Zettlemoyer, L., & Reddy, S. (2023). QLoRA: Efficient finetuning of quantized LLMs. *Proceedings of the 40th International Conference on Machine Learning (ICML)*, 1077–1089. Retrieved from <https://arxiv.org/abs/2305.14314>

Fedus, W., Zoph, B., & Shazeer, N. (2021). *Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity*. arXiv. Retrieved from <https://arxiv.org/abs/2101.03961>

Han, S., Mao, H., & Dally, W.J. (2016). Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *International Conference on Learning Representations (ICLR)*. Retrieved from <https://arxiv.org/abs/1510.00149>

Hu, E.J., Shen, Y., Wallis, P., et al. (2022). LoRA: Low-rank adaptation of large language models. *International Conference on Learning Representations (ICLR)*. Retrieved from <https://arxiv.org/abs/2106.09685>

Kaplan, J., McCandlish, S., Henighan, T., Brown, T.B., Chess, B., Child, R., . . . Amodei, D. (2020, January 23). *Scaling laws for neural language models*. arXiv.org. <https://arxiv.org/abs/2001.08361>

Lester, B., Al-Rfou, R., & Constant, N. (2021). The power of scale for parameter-efficient prompt tuning. *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 3045–3059. Retrieved from <https://arxiv.org/abs/2104.08691>

MLCommons. (2023). *MLPerf training benchmark suite 2023*. Retrieved from <https://mlcommons.org/en/training-normal-30/>

Nature Digital Medicine. (2023). The ethical and practical challenges of integrating LLMs in healthcare. *NPJ Digital Medicine*, 6(179). Retrieved from <https://www.nature.com/articles/s41746-023-00879-8>

NVIDIA. (2023). *Deploy large language models at the edge with IGX Orin*. NVIDIA Developer. Retrieved from <https://developer.nvidia.com/blog/deploy-large-language-models-at-the-edge-with-nvidia-igx-orin-developer-kit/>

NVIDIA. (2023). *Scaling LLMs with NVIDIA Triton and TensorRT-LLM*. NVIDIA Developer. Retrieved from <https://developer.nvidia.com/blog/scaling-llms-with-nvidia-triton-and-nvidia-tensorrt-llm-using-kubernetes/>

Sanh, V., Debut, L., Chaumond, J., & Wolf, T. (2019). DistilBERT, a distilled version of BERT: Smaller, faster, cheaper and lighter. *NeurIPS Workshop on Energy Efficient Machine Learning and Cognitive Computing*. Retrieved from <https://arxiv.org/abs/1910.01108>