

EFFICIENT FINE-TUNING TECHNIQUES FOR SLOVENIAN LANGUAGE MODELS

Camile LENDERING,¹ Manfred GONZÁLEZ,¹ Joaquín FIGUEIRA,¹

¹Faculty of Computer and Information Science, University of Ljubljana

This work adapts and evaluates several parameter-efficient fine-tuning techniques (PEFT) for two large language models: the Slovenian-specific *SloBERTa* and the cross-lingual BERT-Base-Multilingual-Uncased. The study examines methods based on Low-Ranking Adaptation (LoRA), bias update, Prefix Tuning, and Infused Adapter for Attention Approximation (IA3). Performance is assessed using benchmarks from SloBENCH, a Slovenian language evaluation framework, focusing on the following tasks: natural language understanding, named entity recognition, dependency parsing, and textual entailment recognition. The techniques and models are compared in terms of performance, memory consumption, and runtime efficiency, providing a comprehensive evaluation of PEFT methods in diverse linguistic contexts.

Keywords: LLM, PEFT, NLP, Slovenian, LoRA, BitFit, p-tuning, student paper

1 INTRODUCTION

Large language models and, by extension, language models (LM) in general are becoming ubiquitous in the current technological landscape. As such, the amount of tasks and applications that are being performed by, or with the help of these models is continuously growing. Until recently, LMs were built for specific tasks, but with the advent of Large Language Models (LLM) this has become both unnecessary, as models are complex enough now that they can perform multiple tasks well out-of-the-box, and unfeasible, as building these models takes an enormous amount of computational resources. However, if one is to achieve the optimal performance on a given task or expand the capabilities of a LM, some adaptations of the model are in order.

These adaptations usually come in the form of a process called fine-tuning, in which the parameters of a pre-trained language model (PLM) are fitted for the specific task we want to achieve. Nonetheless, given the vast number of parameters in current LLMs, full fine-tuning (FFT) of the weights of a model is usually too computationally prohibitive for most applications. As a result, in the recent years many techniques have been developed to efficiently tune LMs, focusing specifically in reducing the number of trainable parameters (refer to (L. Xu et al., 2023) for a comprehensive review of these techniques). These are usually referred to as parameter efficient fine-tuning techniques (PEFT) and are mainly based on the Lottery Ticket Hypothesis, which states that “large models are needed in pre-training only to induce (in high probability) the existence of sub-networks initialized with the correct inductive bias for learning.”. In other words, that PLMs already contain the capacity and linguistic “knowledge” to perform a wide variety of tasks, but these abilities need to be accentuated through the proper learning stimulus that induces the use of the correct sub-networks when performing a particular task.

The purpose of this work is precisely to adapt and evaluate several PEFT techniques with a focus on the Slovenian language. To cover a wide range of approaches the analyzed techniques will span different families: Low-Rank Adaptation (LoRA) (Hu et al., 2022) and its derivatives, bias update (Zaken et al., 2021) and soft prompt based techniques. As a base model for analysis the SloBERTa LM will be used, which is a BERT model optimized for Slovenian. For performance assessment, the techniques will be evaluated across multiple NLP tasks, including Natural Language Inference (NLI), Named Entity Recognition (NER), Dependency Parsing (DP), and Recognizing Textual Entailment (RTE), using the benchmarks provided by SloBENCH Slovenian evaluation framework (Žitnik & Dragar, 2021).

2 LITERATURE REVIEW

2.1 LoRA

The field of adapting large language models (LLMs) for specific tasks has seen significant advancements with the introduction of LoRA. Hu et al. (Hu et al., 2022) pioneered this approach by freezing the pre-trained model weights and

integrating trainable rank decomposition matrices into the Transformer architecture. This method not only substantially reduces the number of trainable parameters required for downstream tasks but also lessens the GPU memory requirements, thereby enabling comparable or superior model quality with notable efficiency. The provision of a package to facilitate the integration of LoRA with PyTorch models, including implementations for popular models like RoBERTa, DeBERTa, and GPT-2, marks a significant contribution to the field.

Further enhancing the parameter efficiency of fine-tuning PLMs, Zhang et al. (2023) introduced IncreLoRA, an incremental parameter allocation method that adds trainable parameters based on the importance scores of each module. This approach, distinguished from structured pruning methods, not only improves parameter efficiency but also incorporates early learning and restart warmup techniques to improve training effectiveness and stability. The method demonstrated superior parameter efficiency and model performance, particularly in low-resource settings, through experiments on the GLUE benchmark.

On the deployment front, Xu et al. (2023) proposed the Quantization-Aware Low-Rank Adaptation (QA-LoRA) algorithm, aimed at the efficient deployment of LLMs on edge devices. By introducing group-wise operators, QA-LoRA enhances the quantization flexibility while streamlining adaptation, enabling the integration of quantized LLM and auxiliary weights without compromising accuracy. This method stands out by allowing for low-bit inference directly, overcoming the limitations of previous approaches like QLoRA and enabling faster model deployment on resource-constrained devices.

Lastly, the scalability and efficiency of serving multiple LoRA adapters derived from a base model have been addressed by Sheng et al. (2023) through the introduction of S-LoRA. This system, designed for scalable serving, significantly improves throughput and the capacity to serve numerous task-specific fine-tuned models by employing a unified memory management approach and optimized computation strategies. Complementing these efforts, Gao et al. (2024) introduced MoE-LoRA with Layer-wise Expert Allocation (MoLA), which optimizes the allocation of LoRA experts across different layers of the Transformer model, thereby enhancing model efficiency and performance across various NLP tasks. These advancements collectively signify

a leap forward in the efficient adaptation, deployment, and serving of LLMs, paving the way for broader application and innovation in the domain.

2.2 Bias update

BitFit (Zaken et al., 2021) is the first recorded technique to have implemented a sparse fine tuning of LMs using only the bias parameters and it falls under the category of partial fine-tuning according to Xu et al.'s taxonomy (L. Xu et al., 2023). The main mechanism of the technique is simple: during fine-tuning an LM on a particular task update only the bias parameters of the model's encoder layers. These parameters account for only a small fraction of all the parameters of the model (0.1% in the case of BERT). Additionally, the authors found that fitting only a small subset of these bias parameters (mainly the bias of the query encoders of the attention heads and the biases of one of the layers of the MLP inside of the encoder layer) leads to almost no performance drop and modifies only 0.04% of the parameters. According to Xu et al. (2023) the technique achieves good results with only a fraction of the memory footprint of other PEFTs. These promising results, combined with the existence of multiple pre-packaged implementations of the technique, led us to choose it as one of the subjects of analysis of this work.

The findings of the initial BitFit paper were further expanded in the works of Lawton et al. (2023) using neural architecture search (NAS), more specifically, iterative network pruning. The core idea of the method is to iteratively fine-tune the model using BitFit¹ and then prune its bias parameters according to a criteria based on the first order approximation of the loss that results from eliminating certain parameters from the network. The authors found that the resulting network architectures could maintain good performance with a large portion of their bias parameters pruned, further solidifying the findings in the initial BitFit paper that only a relatively small number of bias parameters are responsible for the fine-tuned performance. Unfortunately, there is no code or implementation freely available to replicate the results of this technique in our work.

¹The authors also fine-tuned the models using LoRA, but that falls outside of the scope of this analysis.

2.3 IA3

Infused Adapter for Attention Approximation (IA3) (Liu et al., 2022) is a PEFT technique that focuses on modifying the attention mechanism of a language model by introducing lightweight, learnable adapters that adjust its attention scores. This approach allows for significant reductions in the number of trainable parameters, making fine-tuning more resource-efficient while maintaining or even improving performance on specific downstream tasks. IA3 achieves this by infusing additional parameters only in critical parts of the attention mechanism, thus preserving the core structure and pre-trained knowledge of the original model. The main advantage of IA3 with respect to other PEFT techniques is precisely this infusion mechanism of the learnable parameters, which means that no additional weights have to be added to the model after training and hence inference time remains unchanged.

2.4 Soft prompts

Soft prompts, represent various methods to efficiently adapt LLMs for downstream tasks without altering the underlying model architecture or weights. This technique involves appending a sequence of tunable tokens, or "soft prompts," to the input of the model. During fine-tuning, these tokens are optimized to guide the model towards generating the desired task-specific output. Two prominent methods for implementing soft prompts are prompt tuning and Prefix Tuning (P-tuning).

2.4.1 PROMPT TUNING

Prompt tuning introduces a set of ℓ learnable tokens (soft prompts), denoted as $P = \{P_1, P_2, \dots, P_\ell\}$, and concatenates these tokens to beginning of the input to the model $X \in \mathbb{R}^{n \times d}$ to form $\hat{X} \in \mathbb{R}^{(n+\ell) \times d}$. Throughout the fine-tuning process, only the parameters associated with the prompt tokens P are adjusted via gradient descent, while the pre-trained model parameters are kept fixed. Hence, the length of the prompt and the dimensionality of the token embeddings determine the parameter cost for fine-tuning (Lester et al., 2021).

2.4.2 PREFIX-TUNING

Prefix-tuning introduces the idea of appending a set of soft prompts $P = \{P_1, P_2, \dots, P_\ell\}$, not to the input layer but to the hidden states within the multi-head attention layers of the model. This is different from prompt tuning, which concatenates soft prompts directly to the input. To promote stability during training, a feed-forward network (FFN) is used to parameterize these soft prompts. During fine-tuning, two distinct sets of prefix vectors, \hat{P}_K and \hat{P}_V , are concatenated to the attention layer's original key (K) and value (V) vectors, respectively. Hence, the only parameters that require optimization are those of \hat{P}_K , \hat{P}_V , and the FFN. Once the model is fine-tuned, the FFN is no longer needed, and only the optimized key and value prefix vectors are kept for model inference (Li & Liang, 2021).

3 PROPOSED METHODOLOGY

This section outlines the methodology used in our experiments, covering data preparation, model fine-tuning, evaluation, and resource measurement. All the code used to perform these tasks can be found in the following Github repository: <https://github.com/UL-FRI-NLP-2023-2024/ul-fri-nlp-course-project-pirates-of-the-caribbean>.

3.1 Data Preparation

We began by selecting datasets from the SloBENCH framework, focusing on tasks like Natural Language Inference (NLI), Named Entity Recognition (NER), Dependency Parsing (DP), and Recognizing Textual Entailment (RTE).

For NLI, we used the SI-NLI (Slovene Natural Language Inference) dataset, containing 5,937 Slovene sentence pairs labeled as "entailment," "contradiction," or "neutral." These pairs were derived from the Slovenian reference corpus cckres (Logar et al., 2013) and balanced by annotators who created three variations for each pair. The dataset is split into training (4,392 pairs), validation (547 pairs), and test sets (998 pairs), with the splits designed to evenly distribute difficult and easy instances. The dataset is provided as part of the SloBENCH evaluation framework (Žitnik & Dragar, 2021).

For both NER and DP tasks, we used the SSJ500K corpus, containing around 500,000 tokens annotated for tokenization, sentence segmentation, morphosyntactic tagging, and lemmatization. Half of the corpus includes syntactic dependencies, named entities, and verbal multiword expressions, while a quarter is annotated with semantic role labels. The data was split randomly into 80% for training and 20% for testing, with 10% of the training data further used for validation (Krek et al., 2021).

For the RTE task, we used the Slovene translation of the SuperGLUE benchmark, which includes translations from both machine and human translators. Specifically, the RTE dataset in this benchmark was translated by Google Machine Translation, with some portions also translated by human translators. The dataset was randomly split into 233 examples for training, 30 for validation, and 30 for testing (Žagar et al., 2020).

3.2 Fine-Tuning Techniques

We applied four PEFT methods: LoRA, BitFit, IA3, and Prefix Tuning. These methods were selected to cover a range of approaches, from modifying attention mechanisms to introducing trainable tokens. Each technique offers a different strategy for reducing the number of trainable parameters, freezing different parts of the language models to gain resource efficiency.

3.3 Experimental Setup

All experiments were conducted on a Tesla V100 GPU on the Arnes HPC cluster. Training parameters such as batch size, number of epochs, and learning rates were standardized across all methods to ensure fair comparisons. For each PEFT technique, we fine-tuned the SloBERTa and BERT Base Multilingual Uncased (BBMU) models on the selected datasets. While BBMU is a multilingual model designed to handle multiple languages, our experiments were conducted exclusively in Slovene, rather than in a cross-lingual setting. Training was performed for a fixed number of epochs.

3.4 Evaluation Metrics

We assessed model performance using standard metrics: precision, recall, and F1-score. These metrics were chosen to provide a comprehensive view of model accuracy and balance between precision and recall. Additionally, memory consumption and runtime during training were recorded for each PEFT method to evaluate the trade-off between performance and resource efficiency. As for the memory consumption, we recorded the allocated Memory which is the amount of GPU memory actively used by the tensors at each step of the process. This metric is essential for understanding the memory footprint of the model and ensuring that the GPU resources are utilized efficiently without exceeding capacity.

4 RESULTS

This section outlines the bench marking results for various Slovene and multi-lingual LLMs applied to different NLP tasks, including Named Entity Recognition (NER), Dependency Parsing (DP), Natural Language Inference (NLI), and Recognizing Textual Entailment (RTE). We evaluated the performance using popular PEFT methods to determine their effectiveness across these different language processing tasks.

4.1 Natural Language Inference (NLI)

NLI involves determining whether a "premise" sentence logically supports, contradicts, or is neutral to a "hypothesis" sentence. Using the SI-NLI dataset, this study uses 5,937 Slovene sentence pairs from the Slovenian reference corpus ccKres, each manually annotated with labels of "entailment," "contradiction," or "neutral."

Table 1: Performance Comparison on SI-NLI This table shows the results of training SloBERTa and BBMU models using various Parameter-Efficient Fine-Tuning techniques and Full Fine-Tuning on the SI-NLI dataset. Performance is measured by mean Accuracy, F1 Score, and training time (seconds) in 5 executions of the experiment.

Model	PEFT Method	Accuracy	F1 Score	Time (s)
SloBERTa	LoRa	68.5%	68.0%	260.5
SloBERTa	Prefix Tuning	35.3%	28.9%	235.1
SloBERTa	IA3	54.9%	56.4%	245.1
SloBERTa	BitFit	33.8%	20.3%	106.1
SloBERTa	Full	49.5%	47.7%	323.3
BBMU	LoRa	51.7%	50.8%	297.8
BBMU	Prefix Tuning	43.8%	44.0%	237.7
BBMU	IA3	52.3%	51.6%	247.2
BBMU	BitFit	32.2%	22.8%	107.2
BBMU	Full	51.1%	46.6%	321.4

The above results indicate that SloBERTa with LoRa achieves the highest accuracy (68.5%) and F1 score (68.0%), while BitFit shows the lowest performance for both models, with accuracies of 33.8% and 32.2%. IA3 offers a moderate performance, and Full fine-tuning results in varied outcomes. Training times range from 235 to 323 seconds, highlighting efficiency differences among methods.

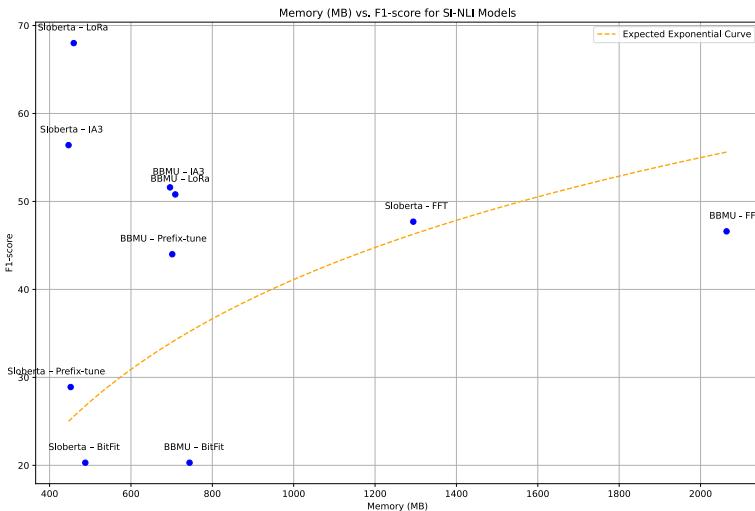


Figure 1: Memory (MB) vs. F1-score for various models and PEFT methods for the SI-NLI task.

The exponential curve in Figure 1 is included as a general reference to illustrate deviations from the expected trend, where increased memory usage typically correlates with improved performance. However, this relationship is not always straightforward, as demonstrated by our results, which show that some PEFT methods achieve higher efficiency, increasing performance without a proportional increase in memory usage.

Methods such as LoRa and IA3 for the SloBERTa model demonstrate high F1 scores with relatively lower memory usage compared to other methods. This suggests that these PEFT methods are more efficient, providing better performance without the need for many computational resources. However, BitFit has lower performance despite its lower memory requirements, indicating that not all PEFT methods offer the same benefits.

4.2 Name Entity Recognition (NER)

The objective of NER is to identify specific key elements in the words of a sentence. For instance, an NER model may be able to identify when a text mentions or refers to a person, a particular location, an organization, etc. In this

experiment we use the SSJ500K training dataset to benchmark several PEFT techniques in the same LM as in the previous task: the Slovene language specific *SloBERTa* and the multilingual BBMU. The corpus contains about 500.000 manually annotated tokens for token-level classification in the Slovene language. Furthermore, it contains several subsets for different tasks; we used the NER subset which contains roughly 9,500 manually annotated sentences using the entities of the CoNLL-2003 shared task.

Table 2: **Performance Comparison on SSJ500K’s NER subset** The table shows the result of training the *SloBERTa* and BBMU models using different PEFT techniques and FFT as a baseline in the NER subset of the SSJ500k corpus. All results were generated using three training epochs with batches of 16 samples. All the statistics obtained are averages from 5 execution of the experiments.

Model	Method	Accuracy	F1 score	Training Time (s)
SloBERTa	LoRA	85.7%	88.2%	85.5
SloBERTa	P-Tuning	80.8%	84.0%	68.3
SloBERTa	IA3	82.3%	85.3%	40.2
SloBERTa	BitFit	14.9%	16.9%	72.5
SloBERTa	FFT	85.8%	88.0%	107.5
BBMU	LoRA	80.0%	82.3%	115.2
BBMU	P-Tuning	75.1%	78.5%	89.8
BBMU	IA3	65.1%	68.5%	95.3
BBMU	BitFit	17.4%	18.5%	94.2
BBMU	FFT	80.9%	82.8%	153.2

The results show decent performance of almost all methods with the *SloBERTa* model, with LoRA even surpassing full fine-tuning with the same amount of training epochs and batch size. Furthermore, all methods reduced considerably the training time over full fine tuning. In fact, Prefix Tuning almost halved the training time of FFT for the *SloBERTa* model, achieving a training time of only 63% of that of FFT. Additionally, of all techniques LoRA performs by far the best. This is expected as LoRa has proven efficacy. Regarding the other methods, it’s surprising that such a simple method as Prefix Tuning can outperform IA3 by such a wide margin in one of the models, specially considering IA3 is intended to be an improvement on LoRA.

Regarding the performance comparison between the two models, it seems *SloBERTa* performs better overall in all experiments. This is to be expected given *SloBERTa* was specifically trained to perform well in Slovene. Additionally, it seems that it responds better to PEFT techniques, hinting at the fact that these techniques may not be as efficient when some cross-lingual transfer learning is required, such as the one required to fine-tune the BBMU model.

Finally, we can see a graph of the different models and techniques combinations performance (f1) vs. memory consumption in figure 2. We can observe similar results as in the previous NLI task, where the results don't seem to follow the expected logarithmic curve of diminishing returns. Rather, we see a much flatter curve where the same results can be achieved with an extensively reduced memory footprint. The only outlier seems to be BitFit as seen previously, which, while still using limited resources, achieves very poor performance.

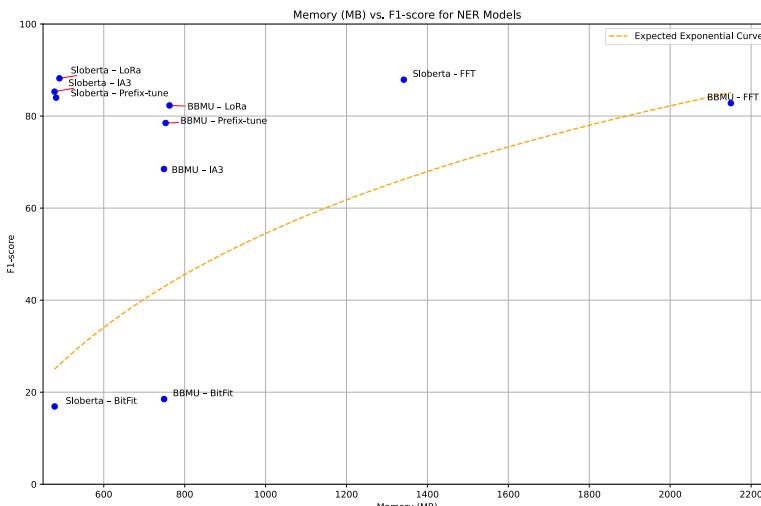


Figure 2: Memory (MB) vs. F1-score for various models and PEFT methods for the NER task.

4.3 Dependency Parsing

Using a different subset from the previous SSJ500k dataset we developed further experiments to evaluate the techniques' performance in Dependency

Parsing (DP). In particular, we wanted to predict for each word, the relation to its root. The subset we used contains about 11 400 synthetically annotated sentences with their corresponding dependency tree. Each word has assigned both a relation to its head or root and an index identifying the mentioned root. We used the same experimental setup as in the previous NER task. The results can be seen in table 3.

Table 3: Performance Comparison on SSJ500K’s DP subset The table shows the result of training the *SloBERTa* and BBMU models using different PEFT techniques and FFT as a baseline in the Dependency Parsing task subset of the SSJ500k corpus. All results were generated using three training epochs with batches of 32 samples, and all the statistics obtained are averages from 5 execution of the experiments.

Model	Method	Accuracy	F1 Score	Time (s)
SloBERTa	LoRA	94.0%	94.1%	45.6
SloBERTa	Prefix Tuning	91%	90.9%	37.2
SloBERTa	IA3	91.9%	91.9%	39.3
SloBERTa	BitFit	83.7%	82.1%	77.8
SloBERTa	FFT	93.1%	93.1%	56.0
BBMU	LoRA	90.4%	90.4%	58.3
BBMU	Prefix Tuning	85.6%	85.3%	50.8
BBMU	IA3	85.9%	85.5%	52.8
BBMU	BitFit	78.3%	77.0%	113.9
BBMU	FFT	90.3%	90.2%	78.7

In this task we observe similar results as in the previous one. LoRA achieves by far the best results in both models, surpassing FFT with significantly lower training times. However, it seems Prefix Tuning never performs better than IA3. The reason for this may be that the task is relatively more complex, requiring the model to discover deeper relationships between the tokens than in NER, so a more complex approach that focuses on enhancing the attention mechanism of the models may be a better strategy to tackle it. Overall, it seems that PEFT techniques perform quite well, specially with the *SloBERTa* model, obtaining similar or even superior results with significantly reduced training times.

Regarding the comparison between the models, similar results are observed as with the NER task: performance is better in the Slovene specific *SloBERTa* and the PEFT methods are less effective in the multilingual model. Finally, just as

in the previous experiments, we can see a comparison between memory consumption and performance for this task in figure 3. We see almost the same trend as before, with higher, but similar, performance and slightly lower memory consumption. The only difference is that BitFit performs much better in this task, indicating that BitFit seems to be very sensitive to its particular application: it may achieve good results with reduced memory usage on some very specific tasks, while completely failing in others. Finally, BitFit has by far the highest run time.

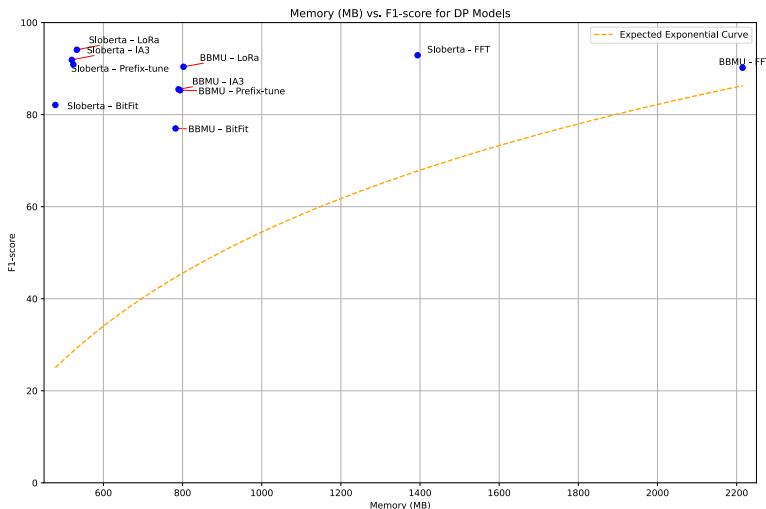


Figure 3: Memory (MB) vs. F1-score for various models and PEFT methods for the DP task.

4.4 Recognizing Textual Entailment (RTE)

In this section, we analyze how LoRA fine tuning influences over the F1 Scores over the human-translated SuperGLUE task RTE dataset from SloBENCH.

Table 4: Performance Comparison on RTE This table summarizes the mean F1 Score and training time (seconds) across five runs for SloBERTa and BBMU models fine-tuned with various Parameter-Efficient Fine-Tuning techniques and Full Fine-Tuning on the SuperGLUE Human-based RTE dataset from SloBENCH.

Model	PEFT Method	F1 Score	Training Time (s)
SloBERTa	LoRa	33.2%	17.53
SloBERTa	Prefix Tuning	54.9%	17.01
SloBERTa	IA3	32.4%	87.67
SloBERTa	BitFit	39.2%	97.94
SloBERTa	FFT	37.1%	10.58
BBMU	LoRa	65.5%	24.78
BBMU	Prefix Tuning	40.6%	26.90
BBMU	IA3	46.9%	84.36
BBMU	BitFit	47.8%	81.78
BBMU	FFT	38.1%	51.29

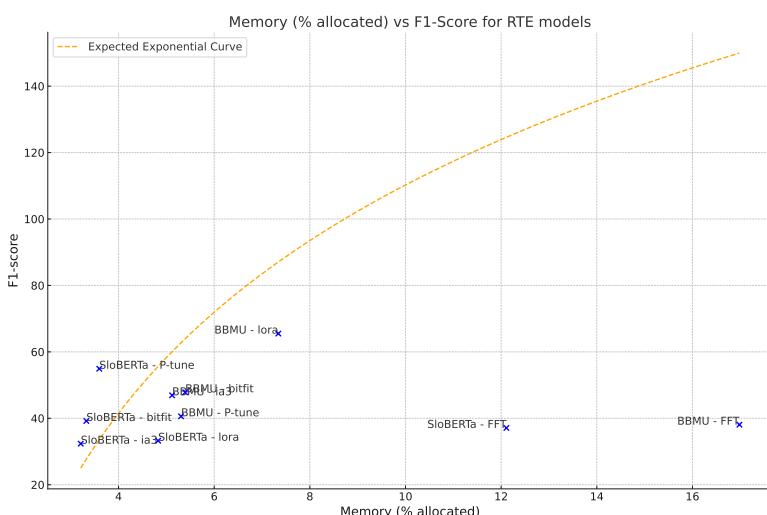


Figure 4: Memory (% allocated) vs. F1-score for various models and PEFT methods for the RTE task.

4.4.1 LORA HYPER-PARAMETER TUNING

In the table 4 we can observe that models like LoRA have bad performance with this small dataset. For this reason we want to explore a bit more this model and see opportunities of improvement by tweaking some of it's hyper parameters. In this experiment, the *SloBERTa* model is fine-tuned on the RTE dataset utilizing the LoRa approach to explore the impact of various hyper parameters on the model's performance, particularly focusing on the F1 score. The experiment iteratively tests combinations of three key hyper parameters: r , α , and *weight decay*.

- **r (rank):** This parameter determines the rank of the adaptation in LoRA (Hu et al., 2022), which affects the model's capacity to learn new patterns without significantly increasing complexity. Values explored were [8, 16, 32], enabling an investigation into how increasing the rank influences learning capacity and generalization.
- **α (scaling factor):** This controls the scaling of the updates in LoRA (Hu et al., 2022), impacting how aggressively the model adapts to the new dataset. Tested values were [16, 32, 64], providing insights into finding a balance between overly subtle and overly aggressive updates.
- **Weight Decay (regularization):** Used to prevent over fitting by penalizing larger weights, with values set at [0.01, 0.1, 0.2]. This parameter explores how varying levels of regularization affect over fitting and model performance on validation data.

The function used for this experiment orchestrates a grid search over these parameters, training the *SloBERTa* model for each configuration and logging the F1 score to identify the optimal setup. The configuration leading to the highest F1 score is noted as the best, reflecting the most effective balance of learning capacity, adaptability, and regularization for this task. Metrics and configurations are saved, and the best model setup is printed, showcasing the impact of tuning LoRA parameters on model efficacy in a sequence classification task on the RTE dataset. The results can be observed in the figures 5, 6, 7 where the best parameters to be used under this configuration are $r=16$, $\alpha=64$, weight decay=0.1. The best F1 Score achieved with *SloBERTa* was about 0.52 over this dataset, the future work will be focused on compare more models to see

which one performs better in this dataset that represents a challenge since it is small and in Slovene.

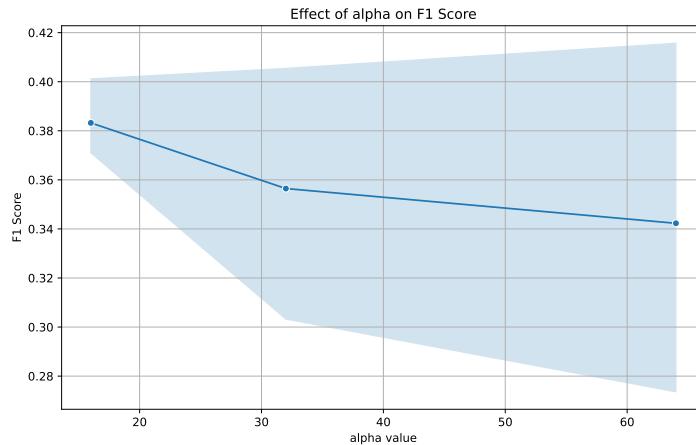


Figure 5: Influence of α on the F1-Score.

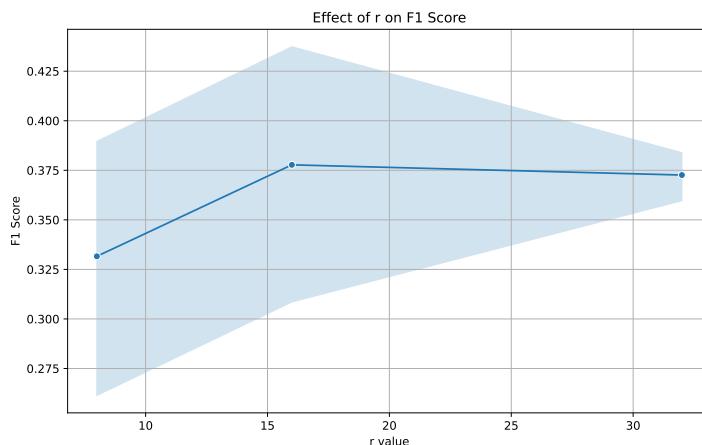


Figure 6: Influence of r on the F1-Score.

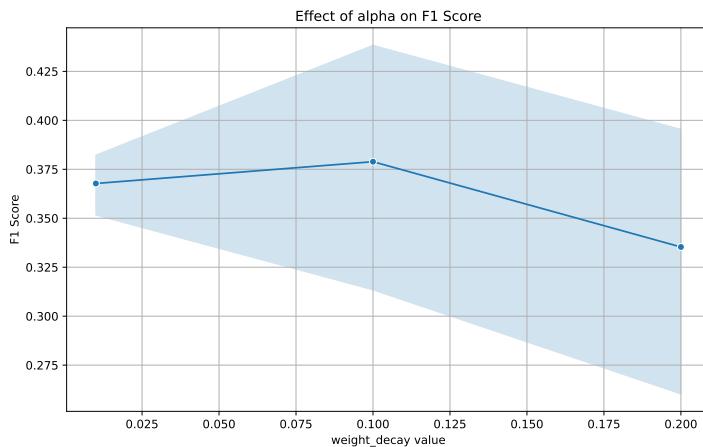


Figure 7: Influence of weight decay on the F1-Score.

4.4.2 RTE BEST MODELS

In this subsection we want to explore how early stopping can help by doing the correct parameter setting of these models to achieve better results than 4.

Method	Patience	Threshold	Eval F1	Epoch
BitFit	1	1×10^{-3}	53%	15
IA3	6	1×10^{-2}	77%	9
LoRA	1	1×10^{-2}	54%	3
Prefix Tuning	3	1×10^{-1}	71%	4

Table 5: Best evaluation F1 Scores for each approach with corresponding epoch, patience, and threshold parameters over SloBERTa model.

The aim in the experiments of table 5 was to identify the best combination of patience and threshold values for early stopping that maximizes the F1 score for each fine-tuning approach. We fine-tuned SloBERTa on a specific dataset using four different parameter-efficient methods: LoRA (Low-Rank Adaptation), Prefix Tuning, IA3 (Incremental Adaptation), and BitFit (Bias Fine-Tuning). For each method, we tested a range of patience and threshold values for early stopping. The patience values tested were [1, 3, 5, 6], and the

threshold values were $[1 \times 10^{-1}, 1 \times 10^{-2}, 1 \times 10^{-3}, 1 \times 10^{-4}]$. Epoch column in the table shows at which epoch the best F1 score was achieved. The patience parameter specifies the number of epochs to wait for an improvement in the validation loss

before stopping the training. Threshold parameter defines the minimum change in the monitored quantity to qualify as an improvement.

The results indicate that different early stopping parameters significantly affect the performance of the fine-tuning methods. Notably, the IA3 method achieved the highest F1 score of 0.77 with a patience of 6 and a threshold of 1×10^{-2} . Prefix Tuning also performed well, with an F1 score of 0.71, using a patience of 3 and a threshold of 1×10^{-1} . LoRA and BitFit methods showed moderate performance with the best F1 scores of 0.54 and 0.53, respectively. LoRA achieved its best performance with a patience of 1 and a threshold of 1×10^{-2} , while BitFit required a patience of 1 and a much lower threshold of 1×10^{-3} to achieve its best result.

These findings suggest that more aggressive early stopping (lower patience and higher threshold) may not always yield the best performance. Instead, a more balanced approach, with moderate patience and threshold values, appears to be more effective for certain fine-tuning methods. This experiment underscores the importance of tuning early stopping parameters to optimize model performance and convergence. These insights can guide future experiments and practical applications involving the fine-tuning of pre-trained models, especially in resource-constrained environments where parameter efficiency is crucial.

5 DISCUSSION

From the experiments and results we obtained, several conclusions can be drawn. Firstly, LoRA consistently emerged as the best performing method across most tasks, which aligns with its proven efficacy in the literature and its widespread use in the industry. The only exception to this general observation was the RTE benchmarks, in which it came out last. This last result highlights the fact that while LoRA is often a good choice when task specific context is missing, it's always best to evaluate tasks and datasets individually to determine the best selection of fine tuning techniques.

Additionally, both IA3 and Prefix Tuning demonstrated competitive performance with LoRA, offering slightly reduced memory footprints. In contrast, BitFit showed inconsistent performance overall, with memory consumption similar to other methods. In some instances, BitFit's performance was notably poor, occasionally dropping to single-digit F1 scores. This outcome is not surprising, as BitFit represents an early and relatively simplistic approach to PEFT.

Although it was pioneering at the time of its development, it has been surpassed by more advanced and sophisticated techniques in the rapidly evolving field of PEFT.

Furthermore, it is important to note that PEFT techniques outperformed FFT on multiple occasions during our benchmarks. This demonstrates that PEFT methods not only make large language models (LLMs) and advanced NLP techniques more accessible to lower-end hardware but also enhance performance on high-end systems. Our results show little to no correlation between memory consumption and performance across all tasks and benchmarks, with minor exceptions due to BitFit's inconsistent performance.

Given that this work heavily relied on Slovene benchmarks, a few remarks about multilingual learning are pertinent. Notably, SloBERTa outperformed BBMU in all tasks and across all PEFT techniques. This result is expected, as SloBERTa was specifically trained on the Slovene language. However, the significant performance improvements seen in BBMU when using PEFT methods highlight the effectiveness of these techniques even in cross-lingual settings.

Finally, we would like to address the sustainability aspect of our work. The widespread use of large language models (LLMs) has a substantial environmental impact (Jiang et al., 2024). Therefore, adopting PEFT techniques to reduce training time and computational resource requirements represents a significant advancement for the sustainability of the NLP field. However, a limitation of our study is that we have not yet tested these methods on larger language models with over 1 billion parameters, which would be an important direction for future research. The notable reductions in both time and memory usage observed across all models and tasks in our study strongly support this premise. By making NLP more efficient, PEFT techniques not only enhance accessibility but also contribute to more environmentally responsible NLP practices.

6 ACKNOWLEDGEMENT

We thank our colleagues from the Faculty of Computer and Information Science, University of Ljubljana, especially Boshko Koloski, for their insightful recommendations and feedback.

REFERENCES

Gao, C., Chen, K., Rao, J., Sun, B., Liu, R., Peng, D., Zhang, Y., Guo, X., Yang, J. & Subrahmanian, V. (2024). *Higher layers need more lora experts*.

Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., & Chen, W. (2022). LoRA: Low-rank adaptation of large language models. In *International conference on learning representations*. Retrieved from <https://openreview.net/forum?id=nZeVKeeFYf9>

Jiang, P., Sonne, C., Li, W., You, F., & You, S. (2024). Preventing the immense increase in the life-cycle energy and carbon footprints of llm-powered intelligent chatbots. *Engineering*. doi: 10.1016/j.eng.2024.04.002

Krek, S., Dobrovoljc, K., Erjavec, T., Može, S., Ledinek, N., Holz, N., Zupan, K., Gantar, P., Kuzman, T., Čibej, J., Arhar Holdt, Š., Kavčič, T., Škrjanec, I., Dafne, M., Jezeršek, L., & Zajc, A. (2021). *Training corpus ssj500k 2.3*. Retrieved from <http://hdl.handle.net/11356/1434> (Slovenian language resource repository CLARIN.SI)

Lawton, N., Kumar, A., Thattai, G., Galstyan, A., & Steeg, G. V. (2023). *Neural architecture search for parameter-efficient fine-tuning of large pre-trained language models*. arXiv. Retrieved from <https://arxiv.org/abs/2305.16597> doi: 10.48550/ARXIV.2305.16597

Lester, B., Al-Rfou, R., & Constant, N. (2021). *The power of scale for parameter-efficient prompt tuning*. arXiv. Retrieved from <https://arxiv.org/abs/2104.08691> doi: 10.48550/ARXIV.2104.08691

Li, X. L., & Liang, P. (2021). *Prefix-tuning: Optimizing continuous prompts for generation*. arXiv. Retrieved from <https://arxiv.org/abs/2101.00190> doi: 10.48550/ARXIV.2101.00190

Liu, H., Tam, D., Muqeeth, M., Mohta, J., Huang, T., Bansal, M., & Raffel, C. (2022). *Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning*. arXiv. Retrieved from <https://arxiv.org/abs/2205.05638> doi: 10.48550/ARXIV.2205.05638

Logar, N., Erjavec, T., Krek, S., Grčaj, M., & Holozan, P. (2013). *Written corpus ccKres 1.0*. Retrieved from <http://hdl.handle.net/11356/1034> (Slovenian language resource repository CLARIN.SI)

Sheng, Y., Cao, S., Li, D., Hooper, C., Lee, N., Yang, S., Chou, C., Zhu, B., Zheng, L., Keutzer, K., Gonzalez J.E., & Stoica, I. (2023). *S-lora: Serving thousands of concurrent lora adapters*.

Xu, L., Xie, H., Qin, S.-Z. J., Tao, X., & Wang, F. L. (2023). *Parameter-efficient fine-tuning methods for pretrained language models: A critical review and assessment*. arXiv. Retrieved from <https://arxiv.org/abs/2312.12148> doi: 10.48550/ARXIV.2312.12148

Xu, Y., Xie, L., Gu, X., Chen, X., Chang, H., Zhang, H., Chen, Z., Zhang, X., & Tian, Q. (2023). Qa-lora:

Quantization-aware low-rank adaptation of large language models.

Žagar, A., Robnik-Šikonja, M., Goli, T., & Arhar Holdt, Š. (2020). *Slovene translation of SuperGLUE*. Slovenian language resource repository CLARIN.SI. Retrieved from <http://hdl.handle.net/11356/1380>

Zaken, E. B., Ravfogel, S., & Goldberg, Y. (2021). *Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models*. arXiv. Retrieved from <https://arxiv.org/abs/2106.10199> doi: 10.48550/ARXIV.2106.10199

Zhang, F., Li, L., Chen, J., Jiang, Z., Wang, B., & Qian, Y. (2023). *Increlora: Incremental parameter allocation method for parameter-efficient fine-tuning*.

Žitnik, S., & Dragar, F. (2021). *SloBENCH evaluation framework*. Slovenian language resource repository CLARIN.SI. Retrieved from <http://hdl.handle.net/11356/1469>

UČINKOVITE TEHNIKE PRILAGAJANJA ZA SLOVENSKE JEZIKOVNE MODELE

To delo prilagaja in ocenjuje več tehnik učinkovitega prilagajanja parametrov (PEFT) za dva velika jezikovna modela: speci ično slovenski *SloBERTa* in večjezični BERT-Base-Multilingual-Uncased. Študija preučuje metode, ki temeljijo na nizko rangirani prilagoditvi (LoRA), posodobitvi pristranskosti, prilagajanju s predponami (Pre ix Tuning) in vstavljenem adapterju za približek pozornosti (IA3). Učinkovitost je ocenjena z uporabo merit SloBENCH, okvira za ocenjevanje slovenskega jezika, s poudarkom na naslednjih nalogah: razumevanje naravnega jezika, prepoznavanje poimenovanih entitet, odvisnostno razčlenjevanje in prepoznavanje besedilnih implikacij. Tehnike in modeli so primerjani glede na učinkovitost, porabo pomnilnika in čas izvajanja, kar zago-tavlja celovito oceno PEFT metod v različnih jezikovnih kontekstih.

Keywords: prilagajanje parametrov, PEFT, slovenski jezikovni modeli, SloBERTa, SloBENCH

To delo je ponujeno pod licenco Creative Commons: Priznanje avtorstva-Deljenje pod enakimi pogoji 4.0 Mednarodna.

This work is licensed under the Creative Commons Attribution-ShareAlike 4.0 International.

<https://creativecommons.org/licenses/by-sa/4.0/>

