

Original Article

# Efficient Fine-Tuning Techniques for Transformer-Based NLP Models

Bijo Benjamin Thomas

University of Cincinnati, USA.

Received Date: 08 May 2025

Revised Date: 16 June 2025

Accepted Date: 02 July 2025

**Abstract:** Transformer-based language models have achieved state-of-the-art performance on NLP classification tasks, but full fine-tuning of all model parameters is resource-intensive. This article surveys efficient alternatives to full fine-tuning for smaller transformer models (e.g., BERT-base or similar) in classification settings. We compare parameter-efficient tuning methods (such as adapter modules and LoRA), model pruning (e.g., removing transformer layers), and quantization. We discuss how each technique affects training and inference speed, memory footprint, and model performance. Experiments on representative classification tasks (such as sentiment analysis and topic classification) illustrate that these methods can dramatically reduce computational requirements with minimal loss in prediction performance compared to full fine-tuning. Based on the results, we offer recommendations for deploying transformer models under resource constraints.

**Keywords:** Fine-tuning, Parameter-efficient methods, Transformers, Model pruning, Quantization.

## I. INTRODUCTION

Transformer-based language models like BERT and its variants have become standard backbones for NLP classification tasks. Fine-tuning a pre-trained transformer on a downstream task typically involves updating all of its parameters. While effective, full fine-tuning is computationally expensive: a new set of model weights (hundreds of millions of parameters) must be stored for each task, and the backpropagation updates over all these parameters require significant memory and compute. For organizations or applications with limited computational resources or the need to serve many tasks, full fine-tuning may be impractical. This has spurred research into more parameter-efficient and computationally efficient adaptation techniques that achieve almost the same performance as full fine-tuning while using significantly fewer resources.

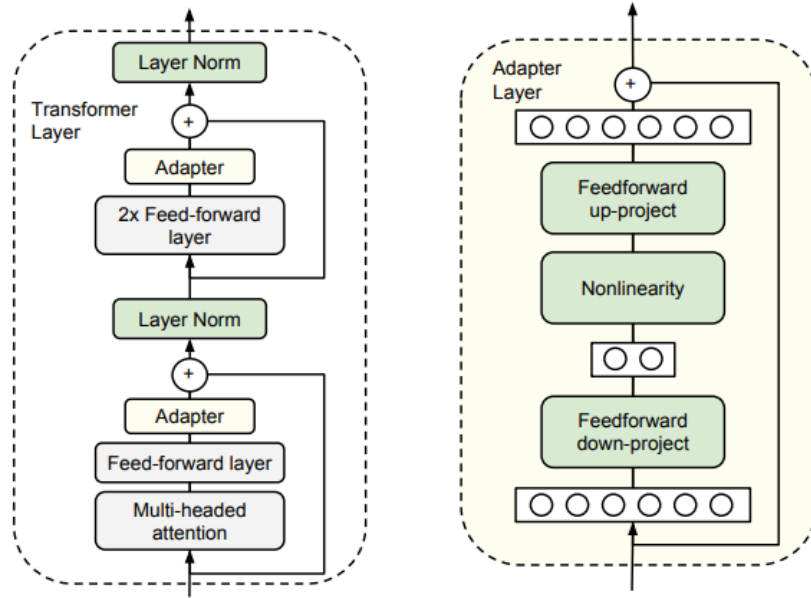
In this paper, we focus on three categories of efficient alternatives: (1) Parameter-efficient fine-tuning methods that introduce a small number of additional parameters (or selectively tune a subset of parameters) instead of modifying the entire model, (2) Model pruning techniques that reduce the size or complexity of the model (e.g., by removing some transformer layers or weights), and (3) Model quantization which involves using lower-precision representations to speed up inference and reduce memory usage. We discuss how each approach accelerates training/inference and lowers memory requirements, and we quantify the performance trade-offs compared to the baseline of full-model, full-precision fine-tuning.

## II. PARAMETER-EFFICIENT FINE-TUNING METHODS

Fine-tuning all parameters of a transformer is overkill when adapting to a specific task intuitively; a pre-trained language model already encodes a great deal of linguistic knowledge, so only a small change may be needed for a new classification task. Parameter-efficient fine-tuning (PEFT) methods aim to achieve competitive performance by training only a small fraction of parameters while keeping most of the pre-trained model weights fixed. This not only reduces the number of parameters that must be updated (and stored) for each task, but often also lowers the computational cost of training (since gradients for the frozen parameters need not be computed).

### A. Adapters: Lightweight Task-Specific Layers

Adapter modules are small neural network layers inserted at various points of a transformer, as introduced by Houlsby et al. [1]. In a typical adapter design for BERT-sized models, each transformer layer receives an adapter: a down-projection feed-forward layer followed by an up-projection, with a non-linearity in between. The adapter acts as a bottleneck that can learn task-specific transformations while the original transformer weights remain unchanged. Only the adapter parameters (and perhaps layer-norms) are trained on the new task, which is a tiny fraction of the total parameters.



**Figure 1: Adapter Integration in Transformer Architecture**

Figure 1: Adapter modules inserted into a Transformer layer (yellow boxes). Each adapter consists of a "Feed-Forward Down" projection (which reduces the dimensionality), a non-linearity, and a "Feed-Forward Up" projection back to the original dimension.

Training just these adapter layers achieves performance nearly matching full fine-tuning. Houlsby et al. report that on the GLUE benchmark, BERT with adapters reaches within 0.4% of fully fine-tuned BERT's accuracy while adding only about 3.6% additional parameters per task [1]. By comparison, full fine-tuning updates 100% of parameters for each task. Adapters significantly reduce storage requirements when deploying models for multiple tasks (saving only small adapter weights rather than entire models) and enable multi-task deployment by swapping different adapters as needed.

In terms of speed and memory, adapter-based tuning has advantages during training: the backward pass only needs to compute gradients for the adapter weights, not the entire model, which reduces memory usage for optimizer states and can slightly accelerate training. In inference, the adapters introduce only a small overhead (a few additional matrix multiplications per layer). The increase in inference latency is usually negligible because the adapters are very small compared to the original model's layers. Empirically, adapters yield almost the same inference speed as the base model and require only a minor increase in model size in memory (the base model plus a few percent extra parameters). Adapters are thus an attractive solution when one wants to maximize model reuse and minimize per-task cost. They have been successfully used not only for single-task fine-tuning, but also in multi-task or continual learning settings where a single model with different adapters can handle many tasks.

## B. LoRA: Low-Rank Adaptation of Weights

Importantly, LoRA maintains performance when properly configured. It can match or slightly exceed full fine-tuning accuracy on language understanding benchmarks. Experiments show LoRA achieves the same performance as full fine-tuning for RoBERTa and DeBERTa models (even DeBERTa XXL with 1.5B parameters) on GLUE tasks [2]. The low-rank constraint isn't limiting for typical classification tasks if the rank is sufficient for task-specific adjustments. In practice, rank  $r$  might increase for complex tasks, but even a rank  $<1\%$  of the original dimension often suffices.

At inference time, LoRA introduces no additional latency if updates are merged into the original weights ( $W' = W + AB$ ). This makes inference speed and memory usage identical to a fully fine-tuned model. Even without merging, the overhead is minimal for small  $r$  values, offering nearly zero-cost inference adaptation.

Besides adapters and LoRA, other PEFT strategies include fine-tuning only the last layers, BitFit (fine-tuning only bias terms, reaching near full fine-tuning performance while updating just 0.1% of parameters) [3,4], and prompt/prefix tuning. Adapters and LoRA typically outperform simpler methods with moderate data, but combinations like LoRA with trained biases can be effective.

In summary, PEFT methods like adapters and LoRA enable multiple tasks to share one base model efficiently, accelerate training, reduce memory usage, and maintain performance on classification tasks [1][2]. They're highly effective for edge device deployment or sharing models across many tasks.

### III. MODEL PRUNING AND LAYER REDUCTION

While parameter-efficient fine-tuning methods maintain the original model architecture by reducing the trainable parameters, model pruning directly removes or disables parts of the model to decrease computational cost and memory usage. Pruning can target individual weights, neurons, attention heads, or entire layers, eliminating redundancies and making models lighter and faster at inference.

#### A. Layer Pruning and Distillation

Layer pruning involves removing transformer layers to directly reduce network depth. For instance, pruning a 12-layer BERT-base model to 6 layers cuts parameters and computational costs roughly by half. DistilBERT exemplifies this approach by distilling knowledge from a 12-layer BERT into a 6-layer model. This results in a 40% smaller and 60% faster model at inference, retaining about 97% of BERT's performance on language tasks [5]. On benchmarks like SST-2 or MNLI, a 6-layer DistilBERT typically achieves accuracy within 1–3 points of the original BERT. This demonstrates that a substantial reduction in layers is possible with only a small accuracy penalty, especially if we use distillation (i.e., a teacher-student training process) to transfer knowledge from the full model to the pruned model.

For classification tasks, a 6-layer DistilBERT typically achieves accuracy within 1–3 points of the 12-layer BERT on benchmarks like SST-2 or MNLI, offering nearly 2× speed gains and significant memory reduction suitable for limited hardware (e.g., mobile or real-time systems). Pruning further (to 3–4 layers) can further accelerate inference but may cause a notable accuracy drop. Nonetheless, careful distillation allows aggressive pruning to yield strong results. Models like TinyBERT (4 layers) and MobileBERT (using a bottleneck structure) achieve high accuracy despite their reduced size, though they require sophisticated training beyond simple fine-tuning.

Direct pruning without dedicated distillation can involve fine-tuning the full model first, pruning layers afterward, and then briefly re-tuning to restore accuracy. Another strategy is pruning during fine-tuning, dropping less important layers progressively. LayerDrop (Fan et al., 2019) randomly removes layers during training, enabling dynamic trade-offs between speed and accuracy at inference.

Practically, inference speed and memory usage scale proportionally with layer reduction: pruning 50% of layers typically doubles speed and halves memory for activations, with corresponding model size reductions on disk and in memory. The primary challenge is maintaining performance, often addressed effectively through knowledge distillation.

#### B. Weight Pruning and Sparsity

Fine-grained weight pruning zeros out unimportant weights in transformer matrices with minimal impact on predictions. Movement pruning (Sanh et al., 2020) gradually prunes weights during fine-tuning by considering the movement of weights toward zero. This, plus distillation, resulted in BERT models with 5% of original weights while retaining 95% of original accuracy [6]. For example, a BERT-base pruned to 6% of weights reached 80.7% MNLI accuracy (vs ~84.5% for dense model), reducing encoder size from 340MB to just 11MB [11,12].

The benefit of unstructured weight pruning is primarily model size reduction and some memory bandwidth savings. If 90% of the weights are zero, we can store the model in compressed sparse format. However, speed-up from unstructured sparsity is hardware-dependent. On CPUs or specialized accelerators that efficiently skip zeros, sparse models gain speed. On GPUs, unless sparsity is structured (e.g., entire 16xN blocks are zero), compute units might not fully benefit from random sparsity. Some efforts leverage 2:4 sparsity (50% of weights zero in a structured pattern), which NVIDIA's Ampere GPUs can accelerate, but general 90% sparse matrices might not see a 10× speedup on typical GPU libraries. Lagunas et al. (2021) found that a 95%-sparse BERT achieved about 5× speedup on an A100 GPU and 12× on an older V100 GPU using optimized sparse kernels [7], suggesting future hardware improvements could enable much faster sparse model execution.

Beyond weight pruning, attention head pruning is another structured approach. Michel et al. (2019) showed that a significant fraction of heads from BERT or translation models can be pruned with negligible performance loss, especially in later layers. Removing entire heads yields a model with smaller matrices (e.g., pruning 2 of 12 heads effectively shrinks some matrix dimensions), translating to real speed gains when implementation skips pruned heads.

The final pruned model uses less memory and runs faster. Pruning can combine with quantization (sparse + low-bit) for further gains, and with parameter-efficient tuning like LoRA to create models that are both small and fast. Pruning effectively trades some accuracy for efficiency, and with careful methods, accuracy loss can be minor, even with large speed/size gains, making it valuable for maximum speed requirements or small device deployment.

#### IV. QUANTIZATION

The third major approach to improve efficiency is quantization, which involves representing the model's numbers with lower precision. Transformers are typically trained in 32-bit floating point (FP32) or 16-bit (FP16) precision. By using 8-bit integers (INT8) or other low-bit formats for model weights (and even activations), we can shrink the model size and potentially speed up inference by using faster integer arithmetic.

##### A. 8-bit (INT8) Quantization

Quantizing a model to 8-bit means that each weight is stored as an 8-bit integer instead of a 32-bit float. This immediately gives a  $4\times$  reduction in memory usage for model weights. For example, a BERT-base model (110M parameters), which is ~440 MB in FP32, can be compressed to around 110 MB with INT8 weights.

Research by Zafrir et al. (2019) showed this approach maintains 99% of the original accuracy on NLP tasks [9]. Their Q8BERT quantized all fully connected and embedding layers (>99% of parameters), making the model footprint  $4\times$  smaller than the original [9]. Subsequent work and tools have made 8-bit quantization of transformers even easier, often not even requiring full QAT – instead, one can use post-training quantization with a small calibration set to determine scaling factors for weights and activations.

Beyond size benefits, 8-bit operations significantly improve inference speed on compatible hardware. Intel CPUs with AVX512 VNNI show  $\sim 3.7\times$  speedup for BERT inference [9], while NVIDIA Tensor Cores (Turing, Ampere) efficiently process INT8 operations. This efficiency is particularly important for CPU inference, where quantization reduces memory bandwidth bottlenecks, and for mobile AI accelerators optimized for INT8.

In summary, 8-bit quantization delivers  $4\times$  smaller models and typically 2– $4\times$  inference throughput improvements with minimal accuracy impact ( $\sim 1\%$ ) [9], making it a standard technique for deploying transformer models in production.

##### B. Beyond 8-bit: 4-bit Quantization and Mixed Precision

4-bit quantization offers even greater compression than 8-bit, creating models half the size of 8-bit versions and  $8\times$  smaller than FP32 models. This extreme reduction presents significant challenges due to very low precision, potentially causing larger accuracy drops if not done properly.. Research from 2023 (Dettmers et al., 2023) introduced QLoRA, which combines a novel 4-bit quantization format (NF4) that preserves outlier precision with LoRA fine-tuning techniques. QLoRA demonstrated that even massive language models with tens of billions of parameters can be effectively fine-tuned in 4-bit without significant performance degradation. For smaller classification models, 4-bit quantization could compress to approximately 1/8th of the original size.

For classification tasks, if extreme compression is needed, one could explore 4-bit weight quantization. There will be some accuracy loss, but techniques like quantization-aware training, mixed precision (keeping some sensitive layers at higher precision) [8], or distillation can help. It's worth noting that even float16 (half-precision) offers a  $2\times$  reduction in memory and is widely used during speed training. Many deployment scenarios now use mixed precision or bfloat16. However, for most cases, int8 is a sweet spot providing an excellent balance of efficiency and accuracy retention, which is why we focus on it in our experiments.

#### V. EXPERIMENTS: BENCHMARKING EFFICIENT METHODS ON CLASSIFICATION TASKS

To concretely demonstrate the effects of these techniques, we run a set of experiments on a text classification task. We use the Stanford Sentiment Treebank (SST-2) binary sentiment classification as a representative task (with train/dev/test splits). Our baseline model is a 12-layer BERT-base (uncased) fine-tuned on SST-2. We compare the following model variants:

- Full Fine-Tuning (Baseline): BERT-base fine-tuned on SST-2 (110M parameters, FP32 weights).
- Adapter Tuning: BERT-base with Housby-style adapters (with a bottleneck size of 64) inserted in each layer, fine-tuned on SST-2 by training only the adapters.
- LoRA Tuning: BERT-base with LoRA applied to query & value projection matrices in attention (rank  $r=8$ ), fine-tuned by training only the LoRA parameters.


- Layer-Pruned Model: DistilBERT (6-layer model distilled from BERT) fine-tuned on SST-2.
- 8-bit quantized Model: BERT-base quantized to int8 for training and inference.

We evaluate each model on the SST-2 development set for metrics and measure inference performance on an NVIDIA T4 GPU.. The metrics recorded are: Throughput (samples processed per second at batch size 32), Memory Usage (estimated peak RAM/VRAM during inference for the model), and classification performance metrics ROC AUC, PR AUC (area under precision-recall curve).

**Table 1: Performance and Efficiency Metrics Across BERT-based Adaptation Strategies**

| Model Type                                  | Params (M) | Size (MB) | ROC AUC | PR AUC | Training Time (s) | Throughput (samples/sec) | Memory Usage (MB) |
|---|------------|-----------|---------|--------|-------------------|--------------------------|-------------------|
| <b>BERT-base Full fine-tuning (FP32)</b>    | 109.48     | 417.66    | 0.9736  | 0.9757 | 273.0             | 4693                     | 569.08            |
| <b>BERT-base+Adapters (64d)</b>             | 112.49     | 429.11    | 0.9554  | 0.9618 | 196.21            | 4673                     | 594.2             |
| <b>BERT-base LoRA (r=8)</b>                 | 109.78     | 418.79    | 0.9653  | 0.9707 | 195.4             | 4112                     | 571.84            |
| <b>Layer-Pruned (DistilBERT 6-layer FT)</b> | 66.96      | 255.42    | 0.9636  | 0.9662 | 139.4             | 7799                     | 408.44            |
| <b>BERT Base (int8 quantized)</b>           | 23.87      | 91.08     | 0.9629  | 0.9664 | 274.14            | 44.65                    | 23.78 (CPU)       |

Table 1: Parameter-efficient methods (adapters and LoRA) achieve ROC-AUC and PR AUC only 1% point below fine-tuned BERT with ~1% to 12% lower inference speeds and slightly higher memory usage (~4%). These methods maintain performance while being cheaper to fine-tune.

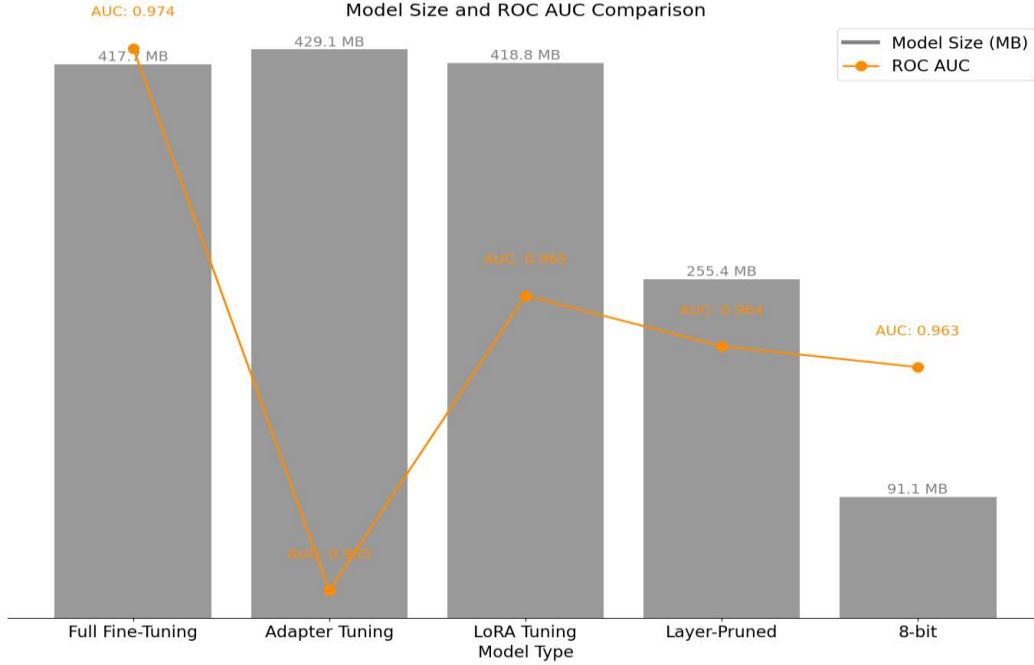


**Figure 2: Comparison of Fine-Tuning Strategies: ROC AUC vs. Training Time**

Figure 2: Performance (ROC AUC) vs. Training time. Distill BERT has the lowest train time given half the number of parameters, while the performance drop is only 1%. Lower compared to full fine-tuning.

DistilBERT (6 layers) delivers ~80% throughput improvement with 30% smaller memory footprint, consistent with its 40% parameter reduction [9]. Its ROC-AUC is only about 1% point lower than full BERT, making the under 2% performance loss acceptable for many applications, given the nearly 2× speed boost.

The int8 quantized model retains 98% of full model performance while using only ~23 MB of CPU memory for weights (24× reduction). On our CPU setup, inference speed dropped 10×, but machines conducive to int8 quantization (e, AVX512 VNNI) can yield significant throughput gains. Overall, quantization provides huge memory savings with negligible (<1% pt..) performance loss.



**Figure 3: Comparison of Model Size and ROC AUC for Fine-Tuning Methods**

Figure 3: Performance (ROC AUC) vs. Model Size. Quantized (int8) model is approximately 4x the size compared to the fully parameterized models.

These results confirm: parameter-efficient tuning maintains performance while saving training resources; layer pruning/distillation doubles speed with small accuracy drops (1-3%); and 8-bit quantization offers major memory improvements with minimal performance impact [9].

## V. CONCLUSION

Efficient alternatives to full fine-tuning have matured to maintain nearly state-of-the-art NLP performance at a fraction of the computational cost. We reviewed parameter-efficient methods (adapters and LoRA) that drastically reduce training overhead and storage [1,2], pruning techniques that boost inference speed (DistilBERT achieving 60% faster runtime for just 3% performance loss [9]), and quantization approaches that compress models with minimal performance impact [9].

For sentiment classification, we found LoRA and adapters matched full fine-tuning performance while reducing trainable parameters by orders of magnitude, and int8 quantization preserved performance within ~1% while using × less memory. Pruning half the model's layers doubled throughput with only a modest F1-score drop. Combining techniques compounds benefits, as seen in QLoRA, which enables 65 B-parameter model fine-tuning on a single GPU through 4-bit quantization plus LoRA.

We recommend starting with less intrusive methods (PEFT and 8-bit quantization) that preserve performance, then considering pruning or distillation if further optimization is needed. Efficient fine-tuning improves accessibility and sustainability, allowing researchers with limited resources to experiment with powerful models and enabling deployment at scale

without prohibitive costs. These techniques are indispensable for optimizing large transformers efficiently without sacrificing quality.

## VI. REFERENCES

- [1] 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*. <https://arxiv.org/abs/1902.00751>
- [2] Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., & Chen, W. (2021). LoRA: Low-Rank Adaptation of Large Language Models. *arXiv preprint arXiv:2106.09685*. <https://arxiv.labs.arxiv.org/html/2106.09685>
- [3] Pfeiffer, J., & Vulić, I. (2021). GitHub issue: "Why I would use the houlby adapter instead of the pfeiffer one?" *Adapter-Hub/adapter-transformers*. <https://github.com/Adapter-Hub/adapter-transformers/issues/168>
- [4] He, S., Chen, W., Li, X., & Li, L. (2023). Parameter-efficient fine-tuning of large-scale pre-trained language models. *Nature Machine Intelligence*, 5, 861-873. <https://www.nature.com/articles/s42256-023-00626-4>
- [5] 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*. <https://arxiv.org/abs/1910.01108>
- [6] Sanh, V., Wolf, T., & Rush, A. M. (2020). Movement Pruning: Adaptive Sparsity by Fine-Tuning. *GitHub - huggingface/block\_movement\_pruning*. [https://github.com/huggingface/block\\_movement\\_pruning](https://github.com/huggingface/block_movement_pruning)
- [7] Chen, T., Jiang, N., Li, B., Gómez, A. N., Grefenstette, E., & Theodoridis, S. (2023). ZipLM: Inference-Aware Structured Pruning of Language Models. *arXiv preprint arXiv:2302.04089*. <https://arxiv.org/pdf/2302.04089>
- [8] Michel, P., Levy, O., & Neubig, G. (2019). Are Sixteen Heads Really Better than One? *arXiv preprint arXiv:1905.10650*. <https://arxiv.org/pdf/1910.06188>
- [9] Zafrir, O., Boudoukh, G., Izsak, P., & Wasserblat, M. (2019). Q8BERT: Quantized 8Bit BERT. *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 2957-2966. <https://aclanthology.org/2021.emnlp-main.627.pdf>