

Optimizing Transformer Models for Resource-Constrained Environments: A Study on Model Compression Techniques

Ziqian Luo ¹, Hanrui Yan ² and Xueting Pan ¹

¹ Oracle Corporation, Seattle, WA 98101, USA; luoziqian98@gmail.com (Z.L.); xtpan8800@gmail.com (X.P.)

² ASCENDING Inc, Fairfax, VA 22031, USA; hanruiy137@gmail.com

Abstract: Recent progress in computer vision has been driven by transformer-based models, which consistently outperform traditional methods across various tasks. However, their high computational and memory demands limit their use in resource-constrained environments. This research addresses these challenges by investigating four key model compression techniques: quantization, low-rank approximation, knowledge distillation, and pruning. We thoroughly evaluate the effects of these techniques, both individually and in combination, on optimizing transformers for resource-limited settings. Our experimental findings show that these methods can successfully strike a balance between accuracy and efficiency, enhancing the feasibility of transformer models for edge computing.

Keywords: transformer model; computer vision; image processing; model compression

1. Introduction

Transformers, as introduced by Vaswani et al. [1], have revolutionized machine learning with their high capacity for modeling data, scalability, and exceptional aptitude for capturing long-range dependencies. Originating in the natural language processing (NLP) sphere with successes in machine translation and text summarization [2], Transformers now extend their utility to computer vision tasks such as image classification [3] and object detection [4]. More recently, their prowess has been harnessed in the field of multimodal sentiment analysis [5], where they integrate and interpret diverse data streams—textual, visual, and acoustic [6,7]—to assess emotional responses [8]. The application of Transformers has also expanded into the healthcare domain [9], facilitating advancements in areas such as disease detection [10], medical imaging analysis [11,12], and the processing of electronic health records. These technologies are pivotal in developing predictive models for patient outcomes, enhancing diagnostic accuracy, and personalizing treatment strategies, showcasing the potential of Transformers to impact patient care significantly. Unlike previous research in NLP that grapples with the quadratic complexity of softmax-attention in lengthy sequences as noted by Child et al. [13], Vision Transformers (ViTs) manage shorter, fixed-length input sequences. As a result, softmax-attention represents a reduced share of the total floating-point operations (FLOPs) in ViTs, presenting a distinct set of optimization challenges and underscoring the need for specialized model compression strategies.

The Vision Transformer (ViT) [14], an innovative architecture derived from NLP methodologies, has

demonstrated that transforming image patches into token sequences processed by transformer blocks can achieve, and sometimes surpass, the accuracy of traditional CNN architectures. This breakthrough has catalyzed a paradigm shift in image processing. However, the advanced performance of ViT is accompanied by a substantially large parameter set, numbering in the hundreds of millions, which results in considerable memory and computational overhead during inference. Such demands render ViTs less viable for devices with limited resources or power constraints. Consequently, the development and refinement of model compression techniques for ViTs have become critical for their adoption in industrial applications, where efficiency and resource management are paramount.

In contrast to well-explored deep learning models, the model compression landscape for Vision Transformers remains relatively untapped. This paper presents a systematic exploration of model compression strategies for ViTs, focusing on quantization, low-rank approximation, knowledge distillation, and pruning. Through a thorough comparative analysis, we assess the individual and collective impact of these techniques on model efficiency. Our investigation also delves into the potential synergistic effects of combining different methods, with the goal of enhancing performance. The extensive experimental findings confirm that our proposed approaches facilitate a favorable balance between maintaining accuracy and improving computational efficiency, which is essential for practical deployments.

2. Related Work

2.1. Quantization

Quantization has emerged as a cornerstone technique for promoting efficient inference in neural networks. This process involves converting a network into a low-bit representation, thereby reducing both computational demands and memory footprint, with minimal impact on model performance. A critical aspect of this technique is establishing an appropriate clipping range for the weights. Krishnamoorthi [15] suggests determining this range by evaluating all weights within a layer's convolutional filters, whereas Shen et al. [16] adopt a groupwise quantization approach for Transformers. To mitigate the potential loss in accuracy that quantization may introduce, Quantization-Aware Training (QAT) has been proposed. QAT involves conducting the standard forward and backward passes using the floating-point representation of the quantized model, followed by the re-quantization of model parameters after each gradient update, thus preserving accuracy while leveraging the benefits of quantization.

2.2. Low-Rank Approximation

The Vision Transformer (ViT) capitalizes on the self-attention mechanism, which inherently entails a quadratic computational complexity, a significant challenge for scalability. Chen et al. [17] highlight that the attention matrices within ViTs inherently possess low-rank properties [18], presenting an opportunity for complexity reduction [19]. Leveraging low-rank matrix approximation on these attention matrices emerges as a promising approach to mitigating computational costs. A variety of methodologies have been developed for this purpose, including Nyström-based methods [20, 21], Performer [22], and Linformer [23], each with its own unique implementation and suitability for integration with pre-trained ViT models during fine-tuning and validation phases. Moreover, the combination of low-rank approximation with sparse attention mechanisms, as suggested by Chen et al. [17], has been shown to yield even more refined approximations, enhancing the efficiency and effectiveness of ViTs.

2.3. Knowledge Distillation

Knowledge distillation is a refined model compression technique wherein a compact 'student' model is trained to emulate a more complex 'teacher' model by utilizing the teacher's soft labels. These soft labels are recognized for their rich informational content, often leading to superior performance in the student model as compared to training with hard labels [24]. The efficacy of soft labels in enhancing student learning has been corroborated by Yuan et al. [25] and Wei et al. [26]. A novel advancement in this domain is the introduction of a

distillation token in Vision Transformers [14] by Touvron et al. [27]. This token, analogous to the class token but dedicated to capturing the teacher's predictions, engages through the self-attention mechanism, optimizing the distillation process. Such bespoke approaches have demonstrated considerable gains over conventional distillation techniques, underscoring the potential for transformer-specific optimization strategies.

2.4. Pruning

Pruning represents a widely endorsed method to streamline the architecture of Vision Transformers by reducing their dimensional complexity [28]. Central to this technique is the assignment of an importance score to each model dimension, allowing for the selective elimination of dimensions deemed less critical, based on their scores. This targeted reduction aims to maintain a robust pruning ratio while preserving the model's accuracy. The strategy of dimensional redistribution, as proposed by Yang et al. [29], may be integrated into the pruning process, further refining the model's performance. Intriguingly, studies have shown that a model, post-pruning, can occasionally surpass the original in performance, indicating the potential of pruning to not only simplify but also to enhance the functionality of the model [30].

3. Methodology

3.1. Quantization

3.1.1. Basic Concept

The overarching objective of quantization is to reduce the precision of model parameters (θ) and intermediate activation maps to a lower precision format, such as 8-bit integers, while minimizing the impact on the model's generalization performance. The initial step in this process involves defining a quantization function capable of mapping weights and activations to a discrete set of values. A commonly utilized function for this purpose is delineated as follows:

$$Q(r) = \text{Int}(r/S) - Z, \quad (1)$$

where Q represents the quantization mapping function, r denotes a real-valued input (e.g, weights, activation), S is a scaling factor, and Z is an integer zero point. This mechanism, known as *uniform quantization*, ensures the equidistant spacing of resultant values. It's noteworthy that alternative *non-uniform quantization* strategies exist. Moreover, the original real value r can be approximated from its quantized counterpart $Q(r)$ through a process known as *dequantization*:

$$\tilde{r} = S(Q(r) + Z), \quad (2)$$

where the approximation \tilde{r} may differ from r due to rounding errors inherent in quantization.

A critical aspect of quantization is determining the optimal scaling factor S , which effectively partitions real values r into discrete segments:

$$Q(r) + Z, \quad (3)$$

with $[\alpha, \beta]$ representing the clipping range and b denoting the bit width of quantization. The selection of the clipping range $[\alpha, \beta]$, a process termed as *calibration*, is pivotal. A straightforward method involves employing the minimum and maximum of the inputs as the clipping range, i.e, $\alpha = r_{\min}$ and $\beta = r_{\max}$, corresponding to an *asymmetric quantization* scheme where $-\alpha = \beta$. Alternatively, a *symmetric quantization* approach, where $-\alpha = \beta = \max(|r_{\max}|, |r_{\min}|)$, can be employed. In such cases, the quantization function in Eq. 1 can be simplified by setting $Z = 0$.

3.1.2. Post Training Quantization

Post Training Quantization (PTQ) streamlines the quantization process by adjusting weights directly, without necessitating further fine-tuning. This efficiency, however, may lead to notable accuracy declines due to the inherent precision loss of quantization. Liu et al. [31] observed substantial accuracy reductions when applying quantization to LayerNorm and Softmax layers within Transformer architectures. Lin et al. [32] attributed these discrepancies to the polarized distribution of activation values in LayerNorm layers and attention map values. Specifically, significant inter-channel variability within LayerNorm layer inputs (as

illustrated on the left side of Figure 1) induces considerable quantization errors when employing layer-wise quantization approaches. Moreover, a predominance of small-value distributions in attention maps—with only sparse outliers approaching a value of 1—further exacerbates performance declines under uniform quantization strategies. Addressing these challenges, Lin et al. [32] introduced a novel quantization approach employing Powers-of-Two Scale for LayerNorm and Log-Int-Softmax for Softmax layers, aiming to mitigate the adverse effects of traditional quantization methods.

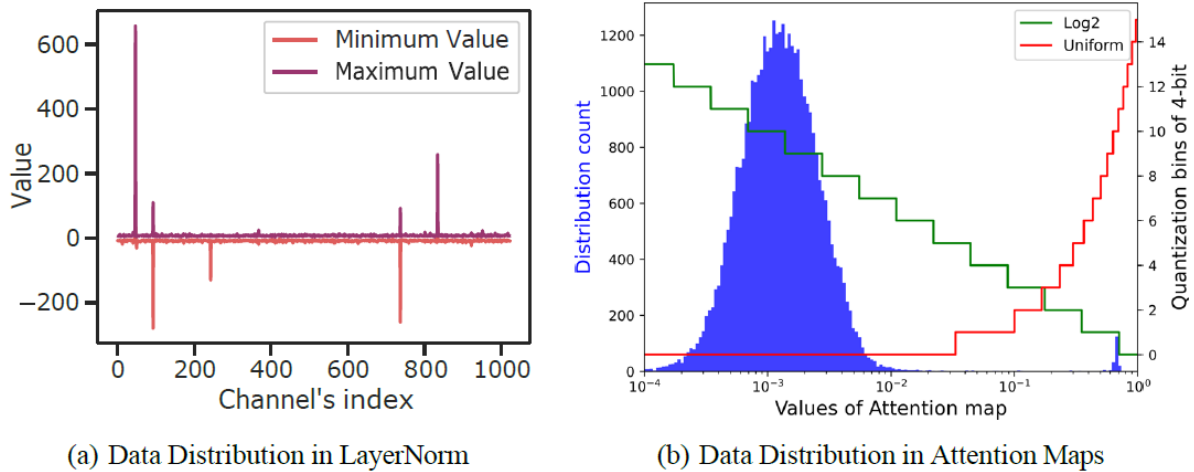


Figure 1. (a) Channel-wise minimum and maximum values of the last LayerNorm inputs in full precision Swin-B [32]. (b) Distribution of the attention map values from the first layer of ViT-L, and visualizing the quantized bins using uniform or Log2 quantization with 4-bit [32].

3.1.3. Quantization Aware Training

Applying quantization directly to a fully trained model can inadvertently perturb model parameters, leading to notable performance declines. An effective strategy to circumvent this issue involves re-training the model with quantized parameters, thereby guiding the model towards a more favorable loss landscape. Quantization Aware Training (QAT) stands out as a prominent technique for this purpose. In QAT, the standard forward and backward processes are executed on a model represented in floating-point, yet parameters are re-quantized following each gradient update, ensuring the model adapts to quantization-induced changes.

Learned Step Size Quantization (LSQ) [33], an advancement in this field, refines the quantizer configuration process and has set new benchmarks in quantization performance by optimizing the quantization intervals. Similarly, DIFFQ [34] introduces a differentiable model compression technique that eschews the need for gradient approximation methods such as the Straight Through Estimator (STE). By employing pseudo quantization noise, DIFFQ achieves an approximation of the quantization process during training that is fully differentiable, thereby facilitating more nuanced adjustments to both the weights and quantization bit-depth.

3.2. Knowledge Distillation

Knowledge distillation techniques, such as soft and hard distillation, facilitate the transfer of knowledge from a complex ‘teacher’ model to a simpler ‘student’ model. Soft distillation focuses on minimizing the Kullback-Leibler (KL) divergence between the softened logits (or outputs) of both the teacher and student models.

The DeiT [27] method introduces a novel approach specific to Transformers, incorporating a ‘distillation token’ into the architecture that functions analogously to the class token but focuses on mimicking the teacher’s predictions. This mechanism allows for a direct interaction between the distillation token and other components through the self-attention layers, demonstrating superior performance in distillation. Our experimental setup involves applying the DeiT framework for knowledge distillation on the CIFAR dataset, adjusting for computational resource constraints.

3.3. Pruning

Pruning in Vision Transformers focuses primarily on reducing the model's complexity by decreasing the number of parameters, specifically by adjusting the dimensions of weight kernels between hidden layers. The goal is to ensure the incremental loss incurred from this reduction remains below a predefined threshold δ , preserving the integrity of the model for subsequent tasks. Determining which dimensions to prune involves the use of importance scores, a concept learned either during pre-training or fine-tuning. Zhu et al. [28] and Yang et al. [29] derive these scores from the gradient magnitude of each weight, proposing the integration of a "soft gate" layer post-pruning which hardens to zero-out less critical dimensions during inference:

$$S_B(W) = \left(\sum_{b \in B} \frac{\partial S}{\partial \omega_b} \omega_b \right)^2 \quad (4)$$

Alternatively, Yu et al. [30] employ KL divergence to calculate importance scores, focusing on the divergence between model performances with and without specific modules across a dataset Ω . This method facilitates both within-layer and across-module pruning:

$$S_B(W) = \sum_{i \in \Omega} D_{KL}(p_i || q_i) \quad (5)$$

where q_i corresponds to the loss with the full model, and p_i to the loss sans the pruned weights. Recent innovations have introduced even more nuanced importance scoring systems. Tang et al. [35] devised a score reflecting the theoretical impact of each patch on the final error, enhancing patch slimming efficiency. Rao et al. [36] combined local and global features for a more holistic assessment of token significance. Similarly, Yi et al. [37] synthesized various scores into a unified loss function, further refining the pruning process.

3.4. Low-Rank Approximation

In Vision Transformers (ViTs), each self-attention block initially projects an input sequence X using weights W_Q , W_K , and W_V to obtain feature representations $Q = W_Q X$, $K = W_K X$, and $V = W_V X$. The self-attention mechanism, computed as $\text{soft}(QK^T / \sqrt{d_q})V$, introduces computational and spatial complexities of $O(n^2)$, where n is the sequence length [21].

Given the formal proof of self-attention's low-rank nature [23], leveraging this property for low-rank approximation emerges as a strategic choice to enhance computational efficiency. Such approximations aim to preserve accuracy while significantly reducing both time and space complexity to approximately $O(n)$, even when integrating with pre-existing or newly trained models [17,21,20].

Notably, low-rank approximation does not inherently reduce the model size, as the original weights W_Q , W_K , W_V are retained. However, it does offer substantial reductions in computational time and memory usage, particularly during fine-tuning or validation phases for pre-trained models. This is because the approximation calculations are performed subsequent to input reception.

Various methodologies for low-rank approximation exist, including Nyström-based approaches like Nyströmformer [21] and SOFT [20], which linearize self-attention through the Nyström method. Alternative linearization techniques, such as Linformer [23] and Performer [22], along with strategies integrating both low-rank and sparse attention mechanisms [17], further enhance approximation accuracy (as shown in Figure 2).

Our experimental focus is on the Nyströmformer-based ViT, adapting the softmax attention matrix calculation to utilize the Nyström method across all self-attention blocks. This allows for the use of pre-trained Vanilla ViT weights, facilitating direct comparisons in performance. The Nyströmformer technique employs landmark points for an efficient approximation, circumventing the need for full QK^T calculations. We evaluate the efficacy of this approach with landmark counts (m) set at 24, 32, and 64, assessing its impact on model performance.

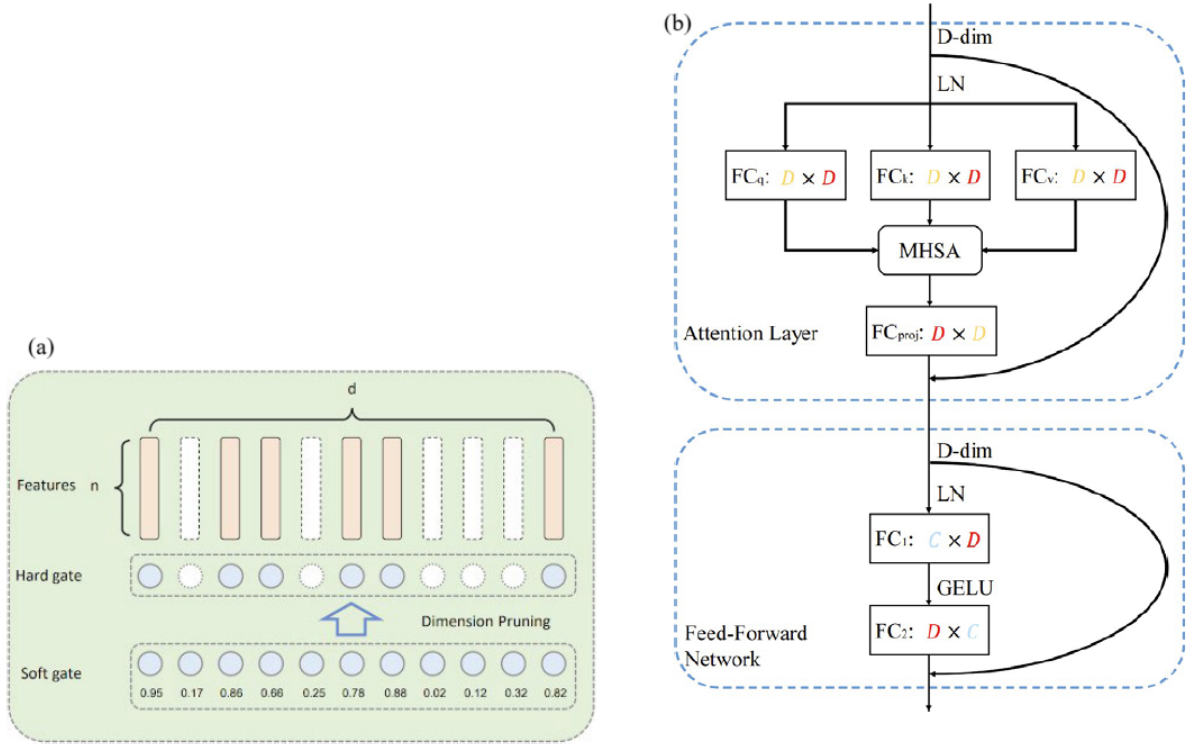


Figure 2. Different pruning models. (a) Soft-gate based model proposed by Zhu et. al. [28]. (b) Skip-dimension model proposed by Yu et. al. [30].

4. Experiments

This section delineates a thorough comparative analysis of various model compression techniques applied to Vision Transformers, including quantization, knowledge distillation, pruning, and low-rank approximation. Additionally, we investigate the synergistic potential of combining these methods to ascertain enhancements in performance metrics.

4.1. Experimental Settings

The experimental framework is established on a Tesla V100-SXM2 16GB GPU, with PyTorch serving as the primary platform for code implementation (It is pertinent to note that PyTorch's current support for quantization is limited to CPU-based operations, necessitating CPU-based inference speed tests for certain methodologies.). The scope of our dataset utilization is confined to CIFAR-10 and CIFAR-100, attributed to computational resource constraints. Our primary metrics of interest include model size and inference speed, acknowledging the inherent trade-off between accuracy and these efficiency parameters. An optimal compression technique is thus characterized by minimal impact on accuracy coupled with substantial reductions in model size and enhancements in inference speed. Results of the comparative analysis across CIFAR-10 and CIFAR-100 datasets are systematically presented in Tables 1 and 2.

Table 1. Evaluation results on CIFAR-10. The Speed values are iterations per second.

Model	Method	Accuracy GPU	Speed CPU	Speed Size (MB)
Vanilla ViT [14]	-	98.94 → 4.48	0.050	327
Dynamic Quantization	Quantization (PTQ)	98.73 → -	0.062	84
FQ-ViT [32]	Quantization (PTQ)	97.31 → -	-	-
DIFFQ with LSQ [33]	Quantization (QAT)	93.37 → 2.10	-	41

Table 1. Cont.

Model	Method	Accuracy	GPU Speed (MB)	CPU Speed	Size
DIFFQ with diffq [34]	Quantization (QAT)	60.29 → 12.20	-	-	2
DeiT base [27]	Knowledge Distillation	98.47 → 7.04	0.096	-	327
DeiT tiny [27]	Knowledge Distillation	95.43 → 16.78	-	-	21
ViT-Pruning (r = 0.1) [28]	Pruning	88.36 → 4.86	-	-	301
ViT-Pruning (r = 0.2) [28]	Pruning	80.56 → 5.54	-	-	254
ViT-Nyströmformer (m = 24) [21]	Low-rank Approximation	65.91 → 4.67	-	-	327
ViT-Nyströmformer (m = 32) [21]	Low-rank Approximation	75.94 → 4.57	-	-	327
ViT-Nyströmformer (m = 64) [21]	Low-rank Approximation	91.70 → 4.38	-	-	327
DeiT base + Dynamic Quantization	Knowledge Distillation + PTQ	96.75 → -	0.117	-	84

Table 2. Evaluation results on CIFAR-100. The Speed values are iterations per second.

Model	Method	Accuracy	GPU Speed (MB)	CPU Speed	Size
Vanilla ViT [14]	-	92.87 → 4.34	0.093	-	327
Dynamic Quantization	Quantization (PTQ)	90.87 → -	0.122	-	84
FQ-ViT [32]	Quantization (PTQ)	84.87 → -	-	-	-
DIFFQ with LSQ [33]	Quantization (QAT)	76.08 → 2.10	-	-	41
DIFFQ with diffq [34]	Quantization (QAT)	41.02 → 12.00	-	-	2
DeiT base [27]	Knowledge Distillation	87.35 → 6.97	0.149	-	327
DeiT tiny [27]	Knowledge Distillation	75.90 → 16.16	-	-	21
ViT-Pruning (r = 0.1) [28]	Pruning	74.46 → 4.69	-	-	302
ViT-Pruning (r = 0.2) [28]	Pruning	64.27 → 5.19	-	-	272
ViT-Nyströmformer (m = 24) [21]	Low-rank Approximation	38.51 → 4.77	-	-	327
ViT-Nyströmformer (m = 32) [21]	Low-rank Approximation	50.31 → 4.65	-	-	327
ViT-Nyströmformer (m = 64) [21]	Low-rank Approximation	74.01 → 4.46	-	-	327
DeiT base + Dynamic Quantization	Knowledge Distillation + PTQ	82.61 → -	0.196	-	84

4.2. Comparison of Different Model Compression Methods

In assessing the impact of model compression on Model Size, we find that quantization and pruning strategies offer substantial size reductions with minimal accuracy compromise. Notably, quantization techniques, particularly Dynamic Quantization (Dynamic Quantization was implemented using the PyTorch Quantization API: https://pytorch.org/tutorials/advanced/dynamic_quantization_tutorial.html), have demonstrated superior efficacy, reducing model size to 25. Contrarily, weight pruning, particularly with simplistic importance scoring, does not facilitate an optimal balance between model size and accuracy. A pruning rate of 0.1 (indicating 10% of parameters pruned) led to a significant accuracy reduction in both CIFAR-10 and CIFAR-100 datasets compared to the unpruned ViT. Further investigation, as depicted in Figure 3, reveals that a majority of parameters are deemed critically important (scores above 0.99), suggesting inherent limitations in simple form importance scoring for weight pruning. Enhancements could stem from the

integration of more sophisticated importance scores [30] or adopting strategies like input patch reduction or slimming, as opposed to direct weight pruning [35,36].

For Inference Speed, a spectrum of enhancements is observed across different model compression strategies, with methods centered around knowledge distillation particularly standing out for their efficiency gains. Notably, the DeiT base model, despite not undergoing significant size reduction, achieves an inference speed nearly double that of the standard Vision Transformer (ViT), all while preserving accuracy to a remarkable degree. An intriguing case is observed with the DeiT tiny configuration on the CIFAR-10 dataset, where it attains 95.43% accuracy—a figure closely aligned with the Vanilla ViT—yet delivers a quadruple increase in speed and is compressed to merely 6% of the original model’s size.



Figure 3. Number of parameters vs importance score. Blue: CIFAR-10. Red: CIFAR-100.

Furthermore, the application of Nyströmformer-based techniques to ViT illustrates a nuanced balance between accuracy and speed, particularly influenced by the selection of the number of landmarks (m). Opting for a larger m value enhances the precision of approximations at the expense of processing velocity. Additionally, Dynamic Quantization contributes to inference speed improvements in the range of 10–20% on CPU platforms, underscoring the practical benefits of model compression beyond just reductions in size.

4.3. Exploration of Mixed Methods

The examination of individual model compression techniques suggests that a hybrid approach, leveraging the strengths of both quantization and knowledge distillation, warrants further investigation. Particularly, when a slight decrease in accuracy is acceptable, such a combined strategy appears promising for optimizing both model compactness and processing efficiency. As demonstrated in Tables 1 and 2, employing a composite method—integrating the DeiT base model with Dynamic Quantization—significantly enhances inference speed, achieving more than a twofold increase, while concurrently reducing the model’s size to one-fourth of its original dimensions. This is achieved with a manageable trade-off in accuracy, highlighting the potential of mixed approaches in striking a balanced compromise between speed, size, and performance [38–55].

5. Conclusion

This study has been dedicated to an empirical investigation of model compression techniques aimed at enhancing the efficiency and deployment viability of Vision Transformers (ViTs). We meticulously examined four predominant compression methods—quantization, low-rank approximation, knowledge distillation, and pruning—complemented by a review of cutting-edge research in the field. Through comparative analyses conducted on the CIFAR-10 and CIFAR-100 datasets, our findings underscore the efficacy of post-training quantization and knowledge distillation as standout strategies. These methods not only significantly reduce model size but also expedite inference times, all while maintaining acceptable levels of performance degradation. Further exploration into the synergistic potential of combining quantization and knowledge distillation has revealed a compelling avenue for optimization. Particularly evident within the CIFAR-10 dataset, this hybrid approach markedly accelerated inference speeds—surpassing baseline speeds by more than a factor of two—while concurrently diminishing model size to merely a quarter of its initial footprint. The insights

garnered from this comprehensive examination advocate for a holistic, multi-faceted approach to model compression. Integrating diverse compression methodologies holds substantial promise for refining the operational efficiency of Vision Transformers, heralding a robust direction for future research in this domain. Moreover, this paper can be further integrated in the field of machine learning, transportation engineering, biological engineering, etc.

Funding

Not applicable.

Author Contributions

Conceptualization, Z. L., H. Y. and X. P.; writing—original draft preparation, Z. L., H. Y. and X. P.; writing—review and editing, Z.L., H.Y. and X.P.; All of the authors read and agreed to the published the final manuscript.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Data Availability Statement

Not applicable.

Conflicts of Interest

The authors declare no conflict of interest.

References

- 1 Vaswani A, Shazeer N, Parmar N, *et al.* Attention Is All You Need. *Advances in Neural Information Processing Systems* 2017; **30**: 1–15.
- 2 Luo Z. Knowledge-guided Aspect-based Summarization. In Proceedings of the 2023 International Conference on Communications, Computing and Artificial Intelligence (CCCAI), Shanghai, China, 23–25 June 2023; pp. 17–22. <https://doi.org/10.1109/CCCAI59026.2023.00012>.
- 3 Chen F, Chen N, Mao H, *et al.* Assessing Four Neural Networks on Handwritten Digit Recognition Dataset (MNIST). *arXiv* 2018, arXiv:1811.08278.
- 4 Luo Z, Zeng X, Bao Z, *et al.* Deep Learning-Based Strategy for Macromolecules Classification with Imbalanced Data from Cellular Electron Cryotomography. In Proceedings of the 2019 International Joint Conference on Neural Networks (IJCNN), Budapest, Hungary, 14–19 July 2019; pp. 1–8.
- 5 Chen F, Luo Z. Sentiment Analysis using Deep Robust Complementary Fusion of Multi-Features and Multi-Modalities. *arXiv* 2019, arXiv:1904.08138.
- 6 Luo Z, Xu H, Chen F. Utterance-Based Audio Sentiment Analysis Learned by a Parallel Combination of CNN and LSTM. *arXiv* 2018, arXiv:1811.08065.
- 7 Luo Z, Xu H, Chen, F. Audio Sentiment Analysis by Heterogeneous Signal Features Learned from Utterance-Based Parallel Neural Network. In Proceedings of the AffCon@ AAI, Honolulu, HI, USA, 27 January 2019; pp. 80–87.
- 8 Chen F, Luo Z. Learning Robust Heterogeneous Signal Features from Parallel Neural Network for Audio Sentiment Analysis. *arXiv* 2019, arXiv:1811.08065.
- 9 Zhang J, Du Y, Zhou P, *et al.* Predicting Unseen Antibodies' Neutralizability via Adaptive Graph Neural Networks. *Nature Machine Intelligence* 2022; **4**: 964–976.
- 10 Wu Y, Gao M, Zeng M, *et al.* BridgeDPI: A Novel Graph Neural Network for Predicting Drug–Protein

- Interactions. *Bioinformatics* 2022; **38**: 2571–2578.
- 11 Chen F, Jiang Y, Zeng X, *et al.* PUB-SalNet: A Pre-Trained Unsupervised Self-Aware Backpropagation Network for Biomedical Salient Segmentation. *Algorithms* 2020; **13**: 126.
 - 12 Liu S, Ban X, Zeng X, *et al.* A Unified Framework for Packing Deformable and Non-Deformable Subcellular Structures in Crowded Cryo-Electron Tomogram Simulation. *BMC Bioinformatics* 2020; **21**: 399. <https://doi.org/10.1186/s12859-020-03660-w>.
 - 13 Child R, Gray S, Radford A, *et al.* Generating Long Sequences with Sparse Transformers. *arXiv* 2019, arXiv:1904.10509.
 - 14 Alexey D. An Image is Worth 16×16 Words: Transformers for Image Recognition at Scale. *arXiv* 2021, arXiv:2010.11929.
 - 15 Krishnamoorthi R. Quantizing Deep Convolutional Networks for Efficient Inference: A Whitepaper. *arXiv* 2018, arXiv:1806.08342.
 - 16 Shen S, Dong Z, Ye J, *et al.* Q-Bert: Hessian Based Ultra Low Precision Quantization of Bert. In Proceedings of the AAAI Conference on Artificial Intelligence, New York, NY, USA, 7–12 February 2020; pp. 8815–8821.
 - 17 Chen B, Dao T, Winsor E, *et al.* Scatterbrain: Unifying Sparse and Low-rank Attention Approximation. *Advances in Neural Information Processing Systems* 2021; **34**: 17413–17426.
 - 18 Chen F, Chen N, Mao H, *et al.* The Application of Bipartite Matching in Assignment Problem. *arXiv* 2019, arXiv:1902.00256.
 - 19 Chen F, Chen N, Mao H, *et al.* An Efficient Sorting Algorithm-Ultimate Heapsort (UHS). *arXiv 2019 arXiv*: 1902.00257
 - 20 Lu J, Yao J, Zhang J, *et al.* Soft: Softmax-Free Transformer with Linear Complexity. *Advances in Neural Information Processing Systems* 2021; **34**: 21297–21309.
 - 21 Xiong Y, Zeng Z, Chakraborty R, *et al.* Nyströmformer: A Nystöm-Based Algorithm for Approximating Self-Attention. In Proceedings of the AAAI Conference on Artificial Intelligence, Virtually, 2–9 February 2021.
 - 22 Choromanski K, Likhoshesterov V, Dohan D, *et al.* Rethinking Attention with Performers. *arXiv* 2022, arXiv: 2009.14794. <https://doi.org/10.48550/arXiv.2009.14794>.
 - 23 Wang S, Li BZ, Khabsa M, *et al.* Linformer: Self-Attention with Linear Complexity. *arXiv* 2020, arXiv: 2006.04768.
 - 24 Hinton G, Vinyals O, Dean J. Distilling the Knowledge in a Neural Network. *arXiv* 2015, arXiv:1503.02531.
 - 25 Yuan L, Chen Y, Wang T, *et al.* Tokens-to-Token Vit: Training Vision Transformers from Scratch on Imagenet. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Montreal, BC, Canada, 11–17 October 2021; pp. 558–567.
 - 26 Wei L, Xiao A, Xie L, *et al.* Circumventing Outliers of Autoaugment with Knowledge Distillation. In *Computer Vision—ECCV*; Springer International Publishing: Cham, Switzerland, 2020; Volume 12348, pp. 608–625. https://doi.org/10.1007/978-3-030-58580-8_36.
 - 27 Touvron H, Cord M, Douze M, *et al.* Training Data-Efficient Image Transformers & Distillation through Attention. In Proceedings of the 38th International Conference on Machine Learning, PMLR, Virtual, 18–24 July 2021; pp. 10347–10357.
 - 28 Zhu M, Tang Y, Han K. Vision Transformer Pruning. *arXiv* 2021, arXiv:2104.08500. <https://doi.org/10.48550/arXiv.2104.08500>.
 - 29 Yu S, Chen T, Shen J, *et al.* Unified Visual Transformer Compression. *arXiv* 2022, arXiv:2203.08243. <https://doi.org/10.48550/arXiv.2203.08243>.
 - 30 Yang H, Yin H, Molchanov P, *et al.* Nvit: Vision Transformer Compression and Parameter Redistribution. 2021. Available online: <https://openreview.net/forum?id=LzBBxCg-xpa>(accessed on 29 May 2024).
 - 31 Yu H, Wu J. A Unified Pruning Framework for Vision Transformers. *Science China Information Sciences* 2023; **66**: 179101. <https://doi.org/10.1007/s11432-022-3646-6>.
 - 32 Liu Z, Wang Y, Han K, *et al.* Post-Training Quantization for Vision Transformer. *Advances in Neural*

- Information Processing Systems* 2021; **34**: 28092–28103.
- 33 Lin Y, Zhang T, Sun P, *et al.* Fq-Vit: Fully Quantized Vision Transformer without Retraining. *arXiv* 2021, arXiv:2111.13824.
- 34 Esser SK, McKinstry JL, Bablani, D, *et al.* Learned Step Size Quantization. *arXiv* 2020, arXiv:1902.08153. <https://doi.org/10.48550/arXiv.1902.08153>.
- 35 Défossez A, Adi Y, Synnaeve G. Differentiable Model Compression via Pseudo Quantization Noise. *arXiv* 2022.
- 36 Tang Y, Han K, Wang Y, *et al.* Patch Slimming for Efficient Vision Transformers. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, New Orleans, LA, USA, 18–24 June 2022; pp. 12165–12174.
- 37 Rao Y, Zhao W, Liu B, *et al.* Dynamicvit: Efficient Vision Transformers with Dynamic Token Sparsification. *Advances in Neural Information Processing Systems* 2021; **34**: 13937–13949.
- 38 Kong C, Li H, Zhang L, *et al.* Link Prediction on Dynamic Heterogeneous Information Networks. In *Computational Data and Social Networks*; Springer International Publishing: Cham, Switzerland, 2019; Volume 11917, pp. 339–350. https://doi.org/10.1007/978-3-030-34980-6_36.
- 39 Chen F, Luo Z, Xu Y, *et al.* Complementary Fusion of Multi-Features and Multi-Modalities in Sentiment Analysis. *arXiv* 2019, arXiv:1904.08138.
- 40 Chen F, Luo Z, Chen N, *et al.* Assessing Four Neural Networks on Handwritten Digit Recognition Dataset (MNIST). *Journal of Computer Science Research* 2024; **6(3)**:17–22.
- 41 Zhou L, Luo Z, Pan X. Machine Learning-Based System Reliability Analysis with Gaussian Process Regression. *arXiv* 2024, arXiv:2403.11125.
- 42 Chen F, Luo Z, Zhou L, *et al.* Comprehensive Survey of Model Compression and Speed Up for Vision Transformers. *arXiv* 2024, arXiv:2404.10407.
- 43 Pan X, Luo Z, Zhou L. Navigating the Landscape of Distributed File Systems: Architectures, Implementations, and Considerations. *arXiv* 2024, arXiv:2403.15701.
- 44 Pan X, Luo Z, Zhou L. Comprehensive Survey of State-of-the-Art Convolutional Neural Network Architectures and Their Applications in Image Classification. *Innovations in Applied Engineering and Technology* 2022; **1(1)**: 1–16.
- 45 Luo Z, Chen F, Chen X, *et al.* A Novel Framework for Text-Image Pair to Video Generation in Music Anime Douga (MAD) Production. *Artificial Intelligence Advances* 2024; **6(1)**: 25–33.
- 46 Luo Z, Pan X. Advancing Mathematical Reasoning in AI: A Comprehensive Approach with the MathQA Dataset. *HAL Open Science* 2024; hal-04667520.
- 47 Luo Z, Pan X. Defense Against Adversarial Attacks on Speech Systems. *HAL Open Science* 2024; hal-04667394.
- 48 Chen Z, Fu C, Wu R, *et al.* Wa with Heterogeneous RGCN for Medical ICD Coding Generation. In Proceedings of the 31st ACM International Conference on Multimedia, Ottawa ON Canada, 29 October–3 November 2023; pp. 5428–5435. <https://doi.org/10.1145/3581783.3612564>.
- 49 Chen Z, Fu C, Tang X. Multi-Domain Fake News Detection with Fuzzy Labels. In *Database Systems for Advanced Applications*. DASFAA 2023 International Workshops; Springer Nature: Cham, Switzerland, 2023; Volume 13922, pp. 331–343. https://doi.org/10.1007/978-3-031-35415-1_23.
- 50 Wang Y, Chen J, Wang M, *et al.* A Closer Look at Classifier in Adversarial Domain Generalization. In Proceedings of the 31st ACM International Conference on Multimedia, Ottawa ON Canada, 29 October–3 November 2023; pp. 280–289. <https://doi.org/10.1145/3581783.3611743>.
- 51 Gu Y, Yan D, Yan S, *et al.* Price Forecast with High-Frequency Finance Data: An Autoregressive Recurrent Neural Network Model with Technical Indicators. In Proceedings of the 29th ACM International Conference on Information & Knowledge Management, Virtual, 19–23 October 2020; pp. 2485–2492. <https://doi.org/10.1145/3340531.3412738>.
- 52 Gu Y, Chen K. GAN-Based Domain Inference Attack. In Proceedings of the AAAI Conference on Artificial Intelligence, Washington, DC, USA, 7–14 February 2023. <https://doi.org/10.1609/aaai.v37i12.26663>.

- 53 Gu Y, Sharma S, Chen K. Image Disguising for Scalable GPU-accelerated Confidential Deep Learning. In Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security, Copenhagen Denmark 26–30 November 2023; pp. 3679–3681. Available online: <https://dl.acm.org/doi/abs/10.1145/3576915.3624364> (accessed on 29 May 2024).
- 54 Du S, Chen Z, Wu H, *et al.* Image Recommendation Algorithm Combined with Deep Neural Network Designed for Social Networks. *Complexity* 2021; **2021**: 5196190. <https://doi.org/10.1155/2021/5196190>.
- 55 Wang Y, Chen Z, Fu C. Synergy Masks of Domain Attribute Model DaBERT: Emotional Tracking on Time-Varying Virtual Space Communication. *Sensors* 2022; **22**: 8450.

© The Author(s) 2023. Published by Global Science Publishing (GSP).



This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.