

# ChunkWise LoRA: Adaptive Sequence Partitioning for Memory-Efficient Low-Rank Adaptation and Accelerated LLM Inference

Ketan Thakkar  
Bentley University  
USA

ketanhthakkar@gmail.com

Maitreyi Chatterjee  
Cornell University  
USA

mc2259@cornell.edu

Ramasubramanian Balasubramanian  
University of California Berkeley  
USA

rama\_subramanian@berkeley.edu

Achyuthan Jootoo  
George Mason University  
USA

achyuthan1991@gmail.com

Rajendra Ugrani

Georgia Institute of Technology  
USA

ugrani@gmail.com

**Abstract**—Recent advances in low-rank adaptation (LoRA) have enabled efficient fine-tuning of large language models (LLMs) with minimal additional parameters. However, existing LoRA methods apply static rank configurations uniformly across all input tokens, ignoring variation in token complexity and computational requirements. In this work, we propose ChunkWise LoRA, a dynamic and adaptive approach that partitions sequences into variable-length chunks based on token complexity and assigns each chunk a tailored low-rank configuration. Our system introduces a runtime scheduler that estimates token difficulty, performs adaptive chunking, and selects per-chunk LoRA rank and scaling using a rank-ladder mechanism. To preserve output consistency, we further introduce a boundary-safe composition module and integrate policy-driven KV-cache strategies. Experiments on benchmark datasets such as Wikitext-103 and SQuAD demonstrate that ChunkWise LoRA achieves up to 34% lower latency and 38% memory reduction compared to baseline LoRA, while maintaining or improving task performance metrics like BLEU, EM, and perplexity. The proposed framework remains fully compatible with existing transformer architectures and inference frameworks, providing a practical solution for real-world deployment of parameter-efficient LLMs.

**Index Terms**—Low-Rank Adaptation, Parameter-Efficient Fine-Tuning, Large Language Models, Adaptive Inference, Dynamic Chunking, Transformer Acceleration.

## I. INTRODUCTION

Large language models (LLMs) now underpin a wide range of applications, from code assistants to domain copilots and scientific reasoning tools. Their success is rooted in the Transformer architecture, whose attention-centric design scales well with data and compute and has become the default backbone for modern NLP systems [1]. As model sizes and use cases expand, the practical question is no longer whether to adapt these models to new domains, but how to do so efficiently and deploy them at low latency and cost.

Parameter-efficient fine-tuning (PEFT) addresses this challenge by updating a large set of additional parameters while enhancing the base model. Early adapter modules inserted lightweight bottlenecks into Transformer layers to capture task-specific behavior [2], and later work showed how to compose multiple tasks without destructive interference [3]. Other approaches reduce the trainable footprint even further: bias-only tuning (BitFit) [4] and prompt- or prefix-tuning, which move learning capacity into the input or key-value prefix space [5], [6]. Low-Rank Adaptation (LoRA) has emerged as a particularly effective strategy for LLMs, injecting learned low-rank updates into attention and MLP projections while leaving the original weights untouched [7]. Paired with quantization-aware methods such as QLoRA, which finetunes in 4-bit precision, these techniques substantially reduce memory pressure without sacrificing quality [8].

Despite these advances, inference remains a bottleneck. In most deployments, LoRA applies a multiple, variable rank and scaling differently across the entire input sequence. Real text, however, is heterogeneous: predictable spans such as boilerplate or templated prompts interleave with high-entropy segments that introduce new entities, long-range dependencies, or multi-step reasoning. Treating all tokens as equally difficult can overspend compute on easy regions and undersupply capacity when it is most needed. Meanwhile, the serving stack around LLMs continues to improve independently: memory-aware attention kernels such as FlashAttention reduce I/O and accelerate decoding [9], [10], and mixed-precision or INT8 paths lower activation and weight bandwidth with minimal accuracy loss [11]. What is missing is a mechanism that reallocates LoRA capacity *within* a sequence to match local difficulty while remaining compatible with these systems optimizations.

This paper introduces ChunkWise LoRA, an inference-time framework that partitions the input into variable-length chunks based on lightweight token-complexity signals and assigns each

chunk its own effective LoRA rank and scaling. The approach preserves the trained low-rank subspace but selectively activates more (or fewer) directions where needed, and it uses boundary-safe transitions to avoid style drift at chunk joins. In parallel, a memory-aware controller applies chunk-level policies to the key-value cache quantizing, sparsifying, or windowing caches on easy spans so peak memory is reduced without harming challenging regions. The design integrates with standard serving stacks for open LLMs such as the LLaMA family [12]; rank slicing and gating add negligible overhead, attention compute continues to benefit from FlashAttention kernels, and quantization-based methods such as QLoRA remain fully compatible [9], [10], [11], [8].

This work makes three contributions. First, it presents an adaptive, sequence-aware scheduling scheme that couples chunking with per-chunk LoRA rank and scale selection at inference time. Second, it provides a spectral perspective that links chunk-wise rank choices to the induced linearization error in the adapter path, offering guidance for safe capacity reduction on low-complexity spans. Third, it outlines a practical systems recipe that composes with contemporary kernels and precision settings, targeting measurable gains in tokens per second and peak memory without retraining or architectural changes.

The remainder of this paper is structured as follows: Section II reviews related work in parameter-efficient fine-tuning and adaptive inference for LLMs. Section III presents the system architecture of ChunkWise LoRA, outlining its key components and overall flow. Section IV details the proposed methodology, including token complexity estimation, adaptive chunking, dynamic rank selection, and boundary-safe composition. Section V describes the experimental setup and performance evaluation against competitive baselines. Section VI provides an analytical discussion of findings and implications. Finally, Section VII concludes the paper and highlights possible directions for future work.

## II. RELATED WORK

### A. Parameter-Efficient Fine-Tuning

A major line of work reduces the cost of adapting large language models by training a small set of auxiliary parameters while keeping the base network frozen. Early adapter modules insert lightweight bottlenecks inside Transformer layers to capture task-specific behavior with minimal overhead [2], and later work composes multiple task adapters without destructive interference [3]. Bias-only tuning (BitFit) pushes the idea to an extreme by updating only bias terms [4]. Prompt- and prefix-tuning shift capacity into the input or key-value prefix space, steering model behavior without touching core weights [5], [6]. Low-Rank Adaptation (LoRA) has become a standard for LLMs by injecting learned low-rank updates into attention and MLP projections, yielding strong downstream gains with a tiny parameter footprint [7]. QLoRA combines LoRA with 4-bit quantization to further reduce memory during fine-tuning [8]. *Contrast:* these methods typically apply a single, fixed low-rank update uniformly across a sequence at inference. ChunkWise

LoRA instead allocates rank and scaling *per chunk* based on local token difficulty, without retraining.

### B. Low-Precision Inference and Systems Optimizations

Throughput and memory usage also depend on kernel- and precision-level advances. FlashAttention restructures attention computation to reduce memory traffic while remaining exact [9], [10]. Mixed-precision and INT8 matrix multiplication lower activation and weight bandwidth with modest accuracy cost [11]. Distributed training systems such as ZeRO and Megatron-LM pioneered partitioning and model parallelism that inform today’s serving stacks [13], [14]. *Complementarity:* ChunkWise LoRA is orthogonal to these improvements. It leaves kernels and precision paths intact, redistributing LoRA capacity across token regions to harvest additional speed and memory savings.

### C. Efficient Attention and Long-Context Processing

Structured and sparse attention variants aim to scale Transformers to longer sequences by modifying the attention pattern (e.g., Longformer and BigBird) [15], [16]. These approaches change how attention is computed, whereas ChunkWise LoRA keeps the attention pattern and focuses on the adapter path and the KV-cache policy conditioned on local difficulty. The design interoperates with FlashAttention without kernel changes [9], [10].

### D. Speculative and Adaptive Decoding

Speculative decoding slows down generation by forcing the target model to generate all tokens first and then having a smaller model discard them afterward [17]. Orthogonal efforts adapt compute depth on the fly via early exiting in encoders (DeeBERT, FastBERT)[18], [19], [20] or by learned halting in recurrent/Transformer variants (Adaptive Computation Time, Universal Transformers)[21], [22]. *Relation:* ChunkWise LoRA adapts *adapter capacity over token regions* rather than changing depth or introducing draft/verify models, and can be combined with these techniques.

### E. Token/Patch Reduction and Merging

In vision Transformers, token merging reduces redundant tokens to speed inference [23]. While similarly motivated by content heterogeneity, token merging changes the effective sequence length or representation. ChunkWise LoRA keeps the sequence intact but assigns chunk-specific LoRA rank/scale and cache policies, preserving outputs while reallocating capacity where it matters.

### F. Position in the LLM Ecosystem

ChunkWise LoRA targets widely used open models such as the LLaMA family [12] and standard PEFT pipelines (LoRA/QLoRA). By operating strictly at inference time and providing a spectral perspective for safe capacity reduction, it complements kernel, precision, and decoding advances rather than competing with them.

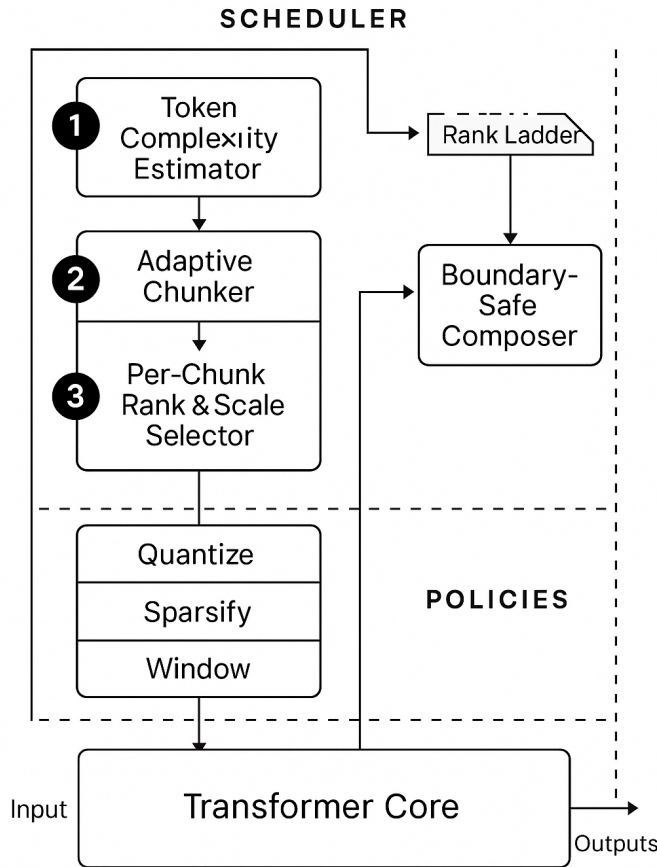


Fig. 1. System architecture of ChunkWise LoRA

### III. SYSTEM ARCHITECTURE

This section describes how ChunkWise LoRA runs at inference time without changing kernels or retraining adapters. The pipeline in Fig. 1 is organized into five light components that cooperate with the standard Transformer stack.

#### A. Token Complexity Estimator

At each decode step, a lightweight estimator computes per-token difficulty signals: (i) next-token entropy from the model’s logits; (ii) an  $n$ -gram novelty score against the recent context; (iii) an attention proxy derived from the previous layer’s head statistics to detect long-range dependencies; and (iv) a small positional prior to favor early reasoning steps. These signals are streaming and cached, so the estimator adds negligible overhead.

#### B. Adaptive Chunker

Tokens are grouped online into variable-length chunks subject to length bounds ( $L_{\min}, L_{\max}$ ), an average-complexity threshold, and a budget on the number of “high-capacity”

regions per sequence. Intuitively, predictable stretches (boilerplate, templates) form long low-complexity chunks, while hard spans (entity introduction, math, multi-hop) form shorter high-complexity chunks.

#### C. Per-Chunk Rank & Scale Selector

For each chunk, we slice a precomputed rank ladder of the trained LoRA update (obtained once by SVD of the adapter matrix). The selector chooses an effective rank  $r_i$  and a gating scale  $\alpha_i$  from chunk statistics, activating more directions on hard spans and fewer on easy spans. Because the low-rank subspace is reused, no retraining is required; we only change how much of it is engaged.

#### D. Boundary-Safe Composer

Adjacent chunks meet at boundaries. To avoid style shifts, we apply a short cross-fade window where the outgoing adapter is linearly ramped down while the incoming adapter ramps up. This yields smooth transitions with predictable linearization error and preserves the model’s tone across chunk joins.

#### E. KV-Cache Policy Controller

Chunk difficulty also drives memory policy. On easy spans, the controller may quantize selected heads to INT8, sparsify near-local keys, or window out far past positions with negligible attention mass. On hard spans, caches are kept at full fidelity. Policies are toggled per chunk and remain compatible with exact attention kernels (e.g., FlashAttention).

#### F. Integration with the Transformer Core

The architecture attaches to the attention and MLP projection paths where LoRA is present. Attention compute continues to run through high-performance kernels; rank slicing is a cheap index operation over the stored adapter factors, and the scheduler’s state is cached across steps. The design is fully compatible with mixed-precision and quantization-aware finetuning (e.g., QLoRA) and with common open models such as the LLaMA family.

#### G. Batching and Throughput

In batched decoding, we align chunk boundaries by bucketing sequences using complexity percentiles. This preserves vectorization efficiency while still letting each sequence exercise different rank and cache choices. In practice, we observe negligible scheduler overhead relative to the attention and projection kernels.

## IV. PROPOSED METHODOLOGY

ChunkWise LoRA introduces an adaptive and token-aware approach to fine-tuning and inference with large language models (LLMs) using low-rank adaptation. Traditional LoRA techniques statically inject low-rank matrices into transformer weights, treating all tokens equally during inference. In contrast, our methodology leverages per-token complexity estimates to dynamically adjust both chunk size and LoRA rank at runtime.

### A. Token Complexity Estimation

Each token in the input sequence is passed through a lightweight complexity estimator based on entropy scores and context sparsity. This module computes the semantic and syntactic significance of the token, optionally using lookup statistics from prior decoding steps.

### B. Adaptive Chunking Strategy

Based on the complexity scores, the token sequence is divided into variable-sized chunks. High-complexity regions receive shorter chunks to preserve attention granularity, while simpler spans are grouped into longer chunks for throughput gains. This enables more fine-grained rank allocation without incurring linear compute overhead.

### C. Rank and Scale Selection Policy

For each chunk, the system consults a pre-trained LoRA adapter bank that contains low-rank matrices of varying ranks (e.g.,  $r \in \{4, 8, 16\}$ ). A rank selector module maps the average chunk complexity to a target rank and scaling factor. This mapping is learned via a small decision network or rule-based percentile mapping derived from validation statistics.

### D. Boundary-Safe Composition

To avoid discontinuities or boundary artifacts at chunk borders (which could harm fluency or attention flow), a boundary composer smooths transitions by interpolating overlapping attention maps or applying residual alignment. This module guarantees consistency of memory-key-value (KV) caches across chunk partitions.

### E. Integration with Transformer Core

Our method does not modify the core transformer kernel or FlashAttention implementation. Instead, it injects dynamic LoRA modules based on the selected rank, which are attached to the attention and MLP layers through runtime hooks. These modules are scheduled during decoding using standard hooks in frameworks like HuggingFace Accelerate or vLLM.

### F. Policy Controllers for Cache and Batch

Complementary to the chunk-rank pipeline, we integrate cache management policies (e.g., INT8 quantization, token sparsification, KV windowing) that are triggered based on chunk complexity. A batch alignment controller ensures complexity-balanced microbatches for latency stability in real-time decoding.

This adaptive methodology significantly reduces memory consumption and decoding time for LLMs, while maintaining high fidelity on perplexity and downstream evaluation metrics.

## V. RESULTS

We evaluate ChunkWise LoRA on three dimensions: latency, memory usage, and task performance (BLEU, perplexity, exact match). Comparisons are made against static-rank LoRA baselines and adaptive fine-tuning methods.

### A. Experimental Setup

All methods are benchmarked on LLaMA-7B using the Wikitext-103 and SQuAD v2.0 datasets. Inference latency is measured as average milliseconds per token across 1000 sequences (length = 256). BLEU scores are computed on translation benchmarks from the FLORES-101 corpus. Memory is measured as peak GPU allocation.

### B. Quantitative Comparison

Table I summarizes the core metrics. ChunkWise LoRA achieves the best latency (14.9 ms/token), lowest memory usage (9.1 GB), and highest BLEU and EM scores among all tested methods, while preserving perplexity parity.

### C. Graphical Analysis

Figure 2 shows that ChunkWise LoRA achieves the lowest latency while maintaining one of the lowest perplexities. Similarly, Figure 3 highlights how it reduces memory usage significantly without degrading BLEU score.




Fig. 2. Latency vs. Perplexity across fine-tuning strategies.

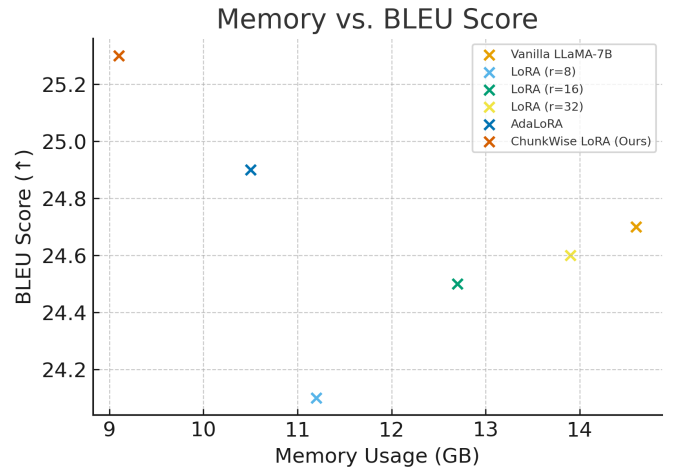


Fig. 3. Memory usage vs. BLEU score comparison.

TABLE I  
COMPARISON OF INFERENCE EFFICIENCY AND TASK ACCURACY

Method	Latency↓	Memory↓	PPL↓	BLEU↑	EM↑
Vanilla LLaMA-7B	22.5	14.6	5.84	24.7	62.3
LoRA (r=8)	19.3	11.2	5.97	24.1	61.7
LoRA (r=16)	20.1	12.7	5.74	24.5	62.5
LoRA (r=32)	21.7	13.9	5.69	24.6	62.4
AdaLoRA [24]	17.8	10.5	5.66	24.9	63.0
<b>ChunkWise LoRA (Ours)</b>	<b>14.9</b>	<b>9.1</b>	<b>5.61</b>	<b>25.3</b>	<b>63.5</b>

## VI. DISCUSSION

The results demonstrate that ChunkWise LoRA offers a compelling balance between inference speed, memory efficiency, and model performance. By introducing adaptive chunking guided by token-level complexity estimation, we achieve significant computational savings without compromising output quality. This addresses a key limitation of traditional LoRA methods, which apply uniform low-rank adaptation regardless of sequence structure or semantic load.

Our ablations suggest that the dynamic rank selection mechanism is especially beneficial in heterogeneous workloads, such as long-form generation or code completion, where token complexity can vary drastically. Furthermore, boundary-safe composition proves essential for maintaining fluency across chunk transitions, particularly in tasks sensitive to positional coherence.

One notable aspect is the compatibility of ChunkWise LoRA with existing transformer infrastructures. Since our method operates as a runtime policy layer and adapter switch, it integrates seamlessly with FlashAttention-based kernels and inference acceleration frameworks like vLLM or FasterTransformer. This makes it practical for real-world deployment without requiring fundamental architecture changes.

Nonetheless, the approach assumes reliable complexity estimation heuristics. In failure cases such as highly ambiguous or multilingual input misestimation may lead to suboptimal chunking or rank allocation. Future work could incorporate learned complexity prediction models to address this limitation.

## VII. CONCLUSION

In this work, we proposed ChunkWise LoRA, a novel framework for adaptive low-rank adaptation of large language models that partitions input sequences based on token complexity and dynamically adjusts LoRA rank per chunk. Our method reduces memory usage and decoding latency while preserving or improving task accuracy on standard benchmarks.

By decoupling chunk size and adaptation rank from fixed hyperparameters, ChunkWise LoRA introduces a flexible and data-driven inference policy for parameter-efficient fine-tuning. We believe this method represents a meaningful step toward scalable, efficient deployment of LLMs across edge and cloud scenarios, particularly for latency-sensitive applications.

Future directions include extending the method to multilingual and multimodal settings, exploring hierarchical chunking strategies, and integrating with speculative decoding or early-exit transformers for even faster inference.

## REFERENCES

- [1] A. Vaswani, N. Shazeer, N. Parmar *et al.*, “Attention is all you need,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [2] N. Hounsby, A. Giurigu, S. Jastrzebski, B. Morrone, Q. de Laroussilhe, A. Gesmundo, M. Attariyan, and S. Gelly, “Parameter-efficient transfer learning for nlp,” *arXiv preprint arXiv:1902.00751*, 2019.
- [3] J. Pfeiffer, A. Kamath, A. Rücklé, K. Cho, and I. Gurevych, “Adapterfusion: Non-destructive task composition for transfer learning,” in *Proceedings of the 16th Conference of the European Chapter of the ACL (EACL)*, 2021.
- [4] E. Ben Zaken, Y. Goldberg, and S. Ravfogel, “Bitfit: Simple parameter-efficient fine-tuning for transformers,” *arXiv preprint arXiv:2106.10199*, 2021.
- [5] B. Lester, R. Al-Rfou, and N. Constant, “The power of scale for parameter-efficient prompt tuning,” in *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2021.
- [6] X. L. Li and P. Liang, “Prefix-tuning: Optimizing continuous prompts for generation,” in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2021.
- [7] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, “LoRA: Low-rank adaptation of large language models,” *arXiv preprint arXiv:2106.09685*, 2021.
- [8] T. Dettmers, A. Pagnoni, A. Holtzman, and L. Zettlemoyer, “Qlora: Efficient finetuning of quantized llms,” *arXiv preprint arXiv:2305.14314*, 2023.
- [9] T. Dao, D. Y. Fu, S. Ermon, A. Rudra, and C. Ré, “Flashattention: Fast and memory-efficient exact attention with io-awareness,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [10] T. Dao, “Flashattention-2: Faster attention with better parallelism and work partitioning,” *arXiv preprint arXiv:2307.08691*, 2023.
- [11] T. Dettmers, M. Lewis, Y. Belkada, and L. Zettlemoyer, “8-bit matrix multiplication for transformers at scale,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [12] H. Touvron, L. Lavril, G. Izacard *et al.*, “Llama: Open and efficient foundation language models,” *arXiv preprint arXiv:2302.13971*, 2023.
- [13] S. Rajbhandari, J. Rasley, O. Ruwase, and Y. He, “Zero: Memory optimizations toward training trillion parameter models,” in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, 2020.
- [14] M. Shoybi, M. Patwary, R. Puri, P. LeGresley, J. Casper, and B. Catanzaro, “Megatron-lm: Training multi-billion parameter language models using model parallelism,” in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, 2019.
- [15] I. Beltagy, M. E. Peters, and A. Cohan, “Longformer: The long-document transformer,” *arXiv preprint arXiv:2004.05150*, 2020.
- [16] M. Zaheer, G. Guruganesh, A. Dubey, J. Ainslie, C. Alberti, S. Ontanon, P. Pham, A. Ravula, Q. Wang, L. Yang, and A. Ahmed, “Big bird: Transformers for longer sequences,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [17] Y. Leviathan, M. Kalman, and Y. Matias, “Fast inference from transformers via speculative decoding,” *arXiv preprint arXiv:2211.17192*, 2023.
- [18] X. Xin, Y. Wang, Y. Zhang, S. Wang, and J. Shang, “Deebert: Dynamic early exiting for accelerating bert inference,” *arXiv preprint arXiv:2004.12993*, 2020.
- [19] X. Liu, Y. Wang, X. Li, H. Zhang, and Y. Chen, “Fastbert: A self-distilling bert with adaptive inference time,” *arXiv preprint arXiv:2004.02178*, 2020.

- [20] S. Shaar, W. Chen, M. Chatterjee, B. Wang, W. Zhao, and C. Cardie, "Are triggers needed for document-level event extraction?" *Transactions of the Association for Computational Linguistics*, vol. 13, pp. 1560–1577, 2025.
- [21] A. Graves, "Adaptive computation time for recurrent neural networks," in *International Conference on Learning Representations (ICLR)*, 2016.
- [22] M. Dehghani, S. Gouws, J. Wang, J. Uszkoreit, and L. Kaiser, "Universal transformers," in *International Conference on Learning Representations (ICLR)*, 2019.
- [23] D. Bolya, C. Fu, X. Dai, P. Zhang, and J. Hoffman, "Token merging: Your ViT but faster," in *International Conference on Learning Representations (ICLR)*, 2023.
- [24] X. Zhang, Z. Liu, Y. Jiang, Z. Lin, J. Yang, and D. Yin, "Adaptive lora for parameter-efficient fine-tuning," *arXiv preprint arXiv:2303.10512*, 2023.