

Token Merging for Multilingual Speech Recognition: A Systematic Study Across Model Scale and Fine-Tuning

Dylan Luke Holyoak

Geffen Academy At UCLA

dholyoa40@geffenacademy.ucla.edu

Abstract

Leading multilingual speech recognition models such as Whisper transcribe diverse, low-resource languages without language-specific training but are computationally expensive to deploy. Token merging mitigates this inefficiency by dynamically combining redundant features, shortening the sequence length during inference without requiring retraining the language model. In this paper, we systematically evaluate token merging on the Whisper model family across sixteen diverse languages and three different model sizes. We also test how token merging interacts with fine-tuning (DoRA) on low-resource languages. Our findings show that merging tokens increases computational efficiency with almost no loss in transcription accuracy across most low-resource languages and model sizes, and it works even after the model has been fine-tuned. Our results demonstrate that token merging is a highly practical method for making multilingual speech recognition faster and cheaper to deploy. Code and frozen evaluation results are available at <https://github.com/chironk/token-merging-multilingual-asr>.

1 Introduction

Speech models have traditionally been trained on large amounts of labeled training data, which is only available for a small number of languages. More recently, multilingual automatic speech recognition (ASR) systems have shown that a single model can be trained to transcribe speech in many languages, rather than a separate model per language. Such systems are particularly useful for low-resource languages, where the amount of audio and transcripts available is inadequate to support a per-language model. Recent multilingual systems include weakly supervised models such as MMS (Pratap et al., 2024) and Whisper (Radford et al., 2023), and self-supervised models such as XLS-R (Conneau et al., 2021). These

models demonstrate that multilingual training can be achieved by learning powerful representations from speech audio alone within a single network. These multilingual models show competitive transcription performance for roughly one hundred to even over a thousand languages.

Among these multilingual ASR models, Whisper (Radford et al., 2023) is the most widely used open-source model, taking speech audio as input and producing transcriptions across roughly one hundred languages without fine-tuning for each language. Although Whisper achieves near-human performance in speech recognition, its audio encoder introduces a significant computational bottleneck. The model is explicitly designed to process fixed-length, 30-second audio inputs, which are encoded into 1,500 tokens and passed through transformer layers. In other words, whether processing a 1-second voice command or a full 30-second monologue, the audio encoder evaluates the 30-second padded sequence. This fixed cost dominates inference for short utterances, creating a major barrier to cost-effective deployment at scale. This inefficiency is especially pronounced in multilingual settings, where average utterance durations can vary significantly across languages.

Token merging mitigates this inefficiency by adaptively combining encoder tokens that share highly similar representations. This operation reduces the sequence length within the model while preserving critical information. Because no retraining is required, the compressed model still transcribes every language it was originally trained on. The technique of token merging originated in computer vision, where soft merging of similar tokens in vision transformers was introduced by ToMe (Bolya et al., 2023), alongside related approaches that prune tokens with learned gating networks (Rao et al., 2021), reorganize tokens by attention score (Liang et al., 2022), halt tokens adaptively per layer (Yin et al., 2022), or vary depth

and width per input (Meng et al., 2022).

By contrast, improving the efficiency of speech recognition systems via token merging remains relatively underexplored. Current literature is primarily limited to adjacent token merging for monolingual English RNN-T speech recognition models (Li et al., 2023) and analysis of adjacent-token redundancy in speech language models (Xiang et al., 2026). To our knowledge, the technique has not been systematically evaluated on a multilingual ASR model, and key questions remain unanswered. Does it generalize across encoder scales spanning roughly an order of magnitude in parameter count? Does it generalize across typologically diverse languages, including the tonal and low-resource cases for which capacity matters most? And does it remain free after adapter-based fine-tuning, when the decoder has been adapted but the encoder has not?

This paper provides such an evaluation. We apply adjacent token merging, because speech audio inputs differ from images in that the encoder sequence has a temporal order that the decoder cross-attends to monotonically. Accordingly, merging must be restricted to adjacent positions to preserve this alignment. The token merging module is based on non-overlapping pairing on similarity, and is placed intra-block between the attention residual and the feed-forward network. We examine the token merging method using Whisper variants with different model scales, and evaluate on sixteen languages, spanning tonal and non-tonal phonology, eleven language (sub)families, and four benchmark resource tiers (Conneau et al., 2023).

Contributions. The present paper makes three contributions:

1. **Cross-lingual evaluation.** We evaluate on sixteen languages covering tonal and non-tonal phonology, eleven language (sub)families, and low-through high-resource tiers, demonstrating that the technique generalizes substantially beyond the English-centric setting of prior work.
2. **Cross-scale evaluation.** We show that canonical adjacent token merging causes negligible word-error-rate (WER) drift at three Whisper encoder scales spanning approximately an order of magnitude in parameter count, indicating that the technique is not specific to one model size.

3. **Composition with fine-tuning.** We show that the technique composes with DoRA (Liu et al., 2024) decoder fine-tuning on all six low-resource target languages, and generalizes to ten held-out languages.

2 Related Work

Token reduction in transformers. The quadratic complexity of self-attention (Vaswani et al., 2017) with respect to sequence length has motivated extensive research into compressing transformer sequences. Early methods, such as DynamicViT (Rao et al., 2021), pruned tokens via learned gating networks. Alternatively, Bolya et al. (2023) demonstrated that bipartite soft *merging*—combining representations rather than discarding them—accelerates vision transformers with negligible accuracy loss and without the need for retraining. In the speech domain, A-ToMe (Li et al., 2023) adapted this idea to an English RNN-T recognizer using a cascade of evenly spaced merge layers to show that token merging also works for ASR models. We adopt their merging schedule and key-based similarity signal, but expand the method to the multilingual encoder-decoder Whisper architecture and explore its composition with parameter-efficient fine-tuning. Furthermore, concurrent research on English LLMs (Xiang et al., 2026) identifies significant adjacent-token redundancy in deep layers, independently corroborating our structural premise.

Efficient speech recognition. Several methods improve speech recognition efficiency by altering the model architecture or weights. For instance, Distil-Whisper (Gandhi et al., 2023) distills the base model into a $5.8\times$ faster student. LiteASR (Kamahori et al., 2025) employs low-rank factorization to halve the size of Whisper-large-v3 while surpassing Whisper-medium’s accuracy. BaldWhisper (Sy et al., 2025) combines embedding compression with layer merging for deployment. However, these approaches all produce new, retrained models, introducing significant overhead when a single model must serve many languages. Other strategies, such as self-supervised pretraining (e.g., wav2vec 2.0 (Baevski et al., 2020), XLS-R (Conneau et al., 2021)) or utilizing parameter-efficient architectures like the Conformer (Gulati et al., 2020), offer alternative methods to lighter encoders but still demand full retraining or per-language fine-tuning. In contrast, token merging fundamentally differs

by reducing the input sequence length rather than the parameter count. It operates entirely at inference time on existing Whisper models, requiring zero retraining and no per-language calibration. A complementary family of methods instead reduces encoder cost from the input side, by cropping the input before it reaches the encoder. WhisperX (Bain et al., 2023) pre-segments long-form audio with a voice-activity-detection front-end and merges the resulting speech regions back into 30-second chunks aligned on low-activity boundaries, reporting a twelve-fold transcription speedup via batched inference. These input-side savings target the padding cost of Whisper’s fixed 30-second window, whereas token merging targets the redundancy that remains inside whatever the front-end forwards.

Multilingual ASR. Modern multilingual ASR systems, such as MMS (Pratap et al., 2024) and XLS-R (Conneau et al., 2021), have scaled speech recognition to hundreds of languages; however, they often rely on self-supervised pretraining and require per-language fine-tuning to transcribe effectively. In contrast, Whisper (Radford et al., 2023) provides highly capable zero-shot transcription without language-specific adaptation. We select Whisper as our primary testbed because of its robust zero-shot capabilities and the public availability of its models across multiple encoder scales, an essential requirement for our cross-scale analysis. To rigorously assess cross-lingual performance, we utilize the FLEURS dataset (Conneau et al., 2023), a standardized benchmark spanning 102 languages. Furthermore, while recent analyses have investigated Whisper’s decoder behavior at a fine sub-token granularity across varying resource tiers (Liang et al., 2025), our work complements this research by targeting the computational efficiency of its encoder.

Parameter-efficient fine-tuning. Parameter-efficient fine-tuning mitigates the computational burden of full model retraining by updating only a minimal subset of parameters while the majority remain frozen. Within this paradigm, LoRA (Hu et al., 2022) introduces trainable low-rank update matrices alongside the original weights, whereas DoRA (Liu et al., 2024) builds upon this by decomposing the updates into distinct magnitude and directional components to enhance learning stability. Other specialized designs, such as depth-aware adaptation (Xiao et al., 2026), selec-

tively distribute adapters across encoder layers to optimize low-resource speech recognition. For our study, we apply DoRA exclusively to the Whisper decoder (leaving the encoder completely frozen) to cleanly isolate and investigate the interaction between adapter-based decoder fine-tuning and inference-time encoder token merging.

3 Method

3.1 Whisper Encoder Architecture

Whisper processes a 30-second log-Mel spectrogram through a two-layer convolutional stem (stride 1 then stride 2) that projects to a fixed-length sequence of $T_0=1500$ tokens. The sequence then passes through L transformer encoder layers, each applying multi-head self-attention followed by a position-wise feed-forward network. The relevant architectural constants are $(L, d, d_{ff}, H, n_{\text{mels}})=(12, 768, 3072, 12, 80)$ for Whisper-small, $(24, 1024, 4096, 16, 80)$ for Whisper-medium, and $(32, 1280, 5120, 20, 128)$ for Whisper-large-v3, where L is the model layers, d the model dimension, d_{ff} the feed-forward dimension, H the number of attention heads, and n_{mels} the number of mel bins. The encoder output is consumed by an autoregressive decoder through cross-attention.

3.2 Token Merging Procedure

We apply the adjacent-pair token merging procedure of Bolya et al. (2023) as adapted to speech recognition by Li et al. (2023), with four details specialized to the encoder-decoder Whisper setting.

Similarity signal. For each candidate adjacent pair $(t, t+1)$ at a merge layer, we compute cosine similarity on per-head-averaged key vectors

$$\bar{k}_t = \frac{1}{H} \sum_{h=1}^H k_t^{(h)}, \quad k_t^{(h)} \in \mathbb{R}^{d/H}, \quad (1)$$

where $k_t^{(h)}$ is the layer’s key projection for head h . We use keys rather than hidden states because they are the representation actually consulted by attention; per-head averaging reduces the per-pair similarity to a single scalar in $[-1, 1]$ while remaining faithful to the similarity structure attention sees.

Placement. Merging is applied intra-block, between the multi-head attention residual update and the layer-norm preceding the feed-forward network. This is the canonical placement of Bolya et al.

(2023) and yields bit-exact equivalence to the unmerged forward when no pairs are merged, which we verify empirically (see §4).

Greedy non-overlapping selection. Given a target number of merges m at the current layer, we sort candidate adjacent pairs by descending cosine similarity and greedily select the top pairs subject to the constraint that no token participates in more than one pair. The merge operation replaces each selected pair $(t, t+1)$ with the mean of their hidden representations and removes the second position, shortening the sequence by exactly m .

Adjacency constraint. We restrict merging to adjacent token pairs. This is essential in the speech setting: encoder tokens correspond to fixed time windows in the input spectrogram, and the decoder cross-attends in approximately monotonic temporal order. Merging non-adjacent tokens would fuse acoustically unrelated time windows and disrupt the alignment on which the decoder relies.

3.3 Merge-Layer Schedule

We adopt the cascade-of-evenly-spaced-layers schedule of Li et al. (2023). Concretely, we apply merging at the fixed subset of encoder layers determined by

$$\mathcal{M}_L = \{2, 5, 8, 11, \dots\} \cap \{1, 2, \dots, L-1\}, \quad (2)$$

which selects every third layer starting from layer 2 (one-indexed) and excludes the final encoder layer. This yields $|\mathcal{M}_L| = 4, 8, 10$ merge layers for $L = 12, 24, 32$ respectively. Excluding the final layer preserves a smoothing buffer before the decoder reads the encoder output. The per-layer reduction rate r is derived from a global token reduction ratio (TRR) by

$$r = 1 - (1 - \text{TRR})^{1/|\mathcal{M}_L|}, \quad (3)$$

so that the multiplicative effect of $|\mathcal{M}_L|$ merge layers compounds to the desired global sequence reduction. At each merge layer the number of merges is $m = \lfloor (T_{\text{pre}} - 1) \cdot r \rfloor$, where T_{pre} is the sequence length entering the layer. We do not tune the merge-layer schedule per language or per model scale.

4 Experimental Setup

Models. We evaluate three Whisper variants drawn from the official OpenAI releases: Whisper-small (244M parameters), Whisper-medium

(769M), and Whisper-large-v3 (1.55B). Our evaluation spans across small to large-scale ASR models. In addition, we fine-tune Whisper-medium using Weight-Decomposed Low-Rank Adaptation (DoRA) (Liu et al., 2024) across six low-resource languages to investigate the potential interplay between token merging and model fine-tuning. All checkpoints are pinned to a fixed HuggingFace revision for reproducibility.

Fine-tuning. We fine-tune the decoder of the Whisper-medium model only (encoder frozen) using DoRA (Liu et al., 2024) on standard ASR transcription — the same task as Whisper’s pre-training — using six mid/low-resource languages: Vietnamese, Hausa, Lingala, Tamil, Maltese, and Javanese. This set of languages is selected to be balanced and diverse: three tonal (Vietnamese, Hausa, Lingala) and three non-tonal (Tamil, Maltese, Javanese) languages, six distinct language families (Austroasiatic, Chadic, Bantu, Dravidian, Semitic, Austronesian), and baseline WER for these languages spanning 14% to 89%. We use the following DoRA hyperparameters: rank 32, $\alpha=64$, target modules $\{\text{q_proj}, \text{k_proj}, \text{v_proj}, \text{out_proj}\}$, learning rate 10^{-5} , 200-step warmup, effective batch size 32, 2000 steps. Training mixes the six languages with temperature ($T=0.5$) plus a 10% English anchor. We use DoRA to fine-tune the model with a training set including the six mid/low-resource languages as listed above. We then test the model using the hold-out test set for these languages. In addition, we evaluate the model on ten additional languages, including four high-resource languages (English, French, German, Spanish) and six untrained mid/low-resource languages (Thai, Swahili, Afrikaans, Icelandic, Welsh, Kazakh).

Data and languages. We evaluate on FLEURS (Conneau et al., 2023) using the Whisper models with different scales. Table 1 lists the sixteen evaluation languages covering tonal and non-tonal phonology, eleven language (sub)families, and four FLEURS resource tiers (high/medium/low/very low). The set comprises four high-resource languages and 12 languages of med/low/very low resources. We sample 264 utterances per language from FLEURS to compute base WER for the test set.

Merging conditions. We test token reduction ratios $\text{TRR} \in \{0, 0.05, 0.10, 0.20, 0.30, 0.40\}$, where $\text{TRR} = 0$ is the baseline without token merg-

Table 1: Evaluation languages from FLEURS (Conneau et al., 2023), ordered by resource tier (H=High, M=Medium, L=Low, VL=Very Low) and tonal status (T=tonal, NT=non-tonal). Base WER (%) is from the Whisper-medium model baseline on $n=264$ utterances per language.

Language	Family	Ton.	Res.	Base WER (%)
English	Germanic	NT	H	5.2
French	Romance	NT	H	8.8
German	Germanic	NT	H	6.6
Spanish	Romance	NT	H	3.5
Thai	Tai-Kadai	T	M	40.6
Vietnamese	Austroasiatic	T	M	14.3
Swahili	Bantu	NT	M	51.7
Tamil	Dravidian	NT	M	35.7
Hausa	Chadic	T	L	88.8
Lingala	Bantu	T	L	83.2
Afrikaans	Germanic	NT	L	45.1
Icelandic	Germanic	NT	L	49.8
Maltese	Semitic	NT	L	82.2
Welsh	Celtic	NT	L	36.9
Javanese	Austronesian	NT	VL	67.2
Kazakh	Turkic	NT	VL	53.4

ing. The merge-layer schedule is fixed across models per Equations 2 and 3.

Evaluation. We report word error rate (WER) on Whisper-normalized transcripts (lower-cased, punctuation stripped, with digits preserved) and report Δ WER in percentage points (pp) relative to the baseline of the same model without token merging on the same samples. We also compare the model efficiency by comparing the empirical measure of computation time before and after token merging.

5 Results

We organize the results around four empirical questions. Does merging generalize across typologically diverse languages (§5.1)? Does it compose with adapter-based decoder fine-tuning (§5.2)? Does it generalize across encoder scales (§5.3)? And what is the realized efficiency benefit (§5.4)? All Δ WER values are reported in percentage points relative to the unmerged baseline of the same model on the same samples. Since WER is an error metric, lower error indicates better performance of the ASR model. Hence, negative Δ WER indicates less error (better performance) of the model with token merging than the baseline model. Positive Δ WER suggests that token merging degrades performance (more errors).

5.1 Cross-Lingual Robustness

We evaluate adjacent token merging on the Whisper-medium model across sixteen FLEURS languages. Figure 1 reports Δ WER as a function of token reduction ratios TRR; per-language values are reported in Table 3 of Appendix B.

We first examine the twelve med-/low-resource languages that span tonal (Vietnamese, Thai, Hausa, Lingala) and non-tonal phonology, Latin-script (Afrikaans, Welsh, Maltese, Icelandic) and non-Latin-script (Tamil, Thai, Kazakh) writing systems. At $\text{TRR} = 0.40$, the cohort-mean Δ WER is -0.21 pp, with the largest single-language improvement reaching -1.41 pp (Lingala). The aggregate is essentially flat across all five TRRs, varying only between -0.07 and -0.28 pp. Notably, nine of the twelve languages exhibit a negative Δ WER under token merging, i.e., slightly lower error than the baseline model without token merging. Hence, for mid-/low-resource languages, token merging not only improves computational efficiency but also leaves accuracy intact and in most cases marginally improves it.




Figure 1: Δ WER (percentage points) as a function of token reduction ratio with the Whisper-medium model for the four high-resource languages (top) and the 12 main-sweep languages (bottom).

To verify that the result is not an artifact of med-/low-resource languages where the baseline without token merging often performs poorly, we additionally evaluate four high-resource anchor languages with Whisper-medium under the same protocol. As shown in the top panel of Figure 1, the mean Δ WER at TRR = 0.40 is only +0.10 pp (English -0.04, Spanish +0.06, German +0.24, French +0.15). Despite baseline WERs for these high-resource languages already near the floor (3.5–8.8%), token merging remains essentially cost-free in accuracy on these languages.

5.2 Composition with Fine-Tuning

Table 2 compares the Whisper-medium against the DoRA fine-tuned model trained using the six languages, with and without merging at TRR = 0.40. Fine-tuning yields substantial baseline gains on four of the six languages, reducing WER for fine-tuned models for Lingala (83.2 $\rightarrow$ 52.2%), Javanese (67.2 $\rightarrow$ 51.2%), Hausa (88.8 $\rightarrow$ 75.3%), and Maltese (82.2 $\rightarrow$ 75.9%). Vietnamese (14.3 $\rightarrow$ 14.1%) and Tamil (35.7 $\rightarrow$ 36.0%) show essentially no change.

Applying token merging on top of fine-tuning costs at most +0.59 pp on any of the six languages at TRR = 0.40 (Lingala), and improves WER on three of these languages. Fine-tuning the multilingual ASR models and token merging are therefore additive in practice: fine-tuning unlocks double-digit WER reductions on the weakest languages (Lingala -31, Javanese -16, Hausa -13 pp), while token merging remains approximately free on the resulting model. The mean Δ WER at TRR = 0.40 is -0.29 pp on the fine-tuned model versus -0.50 pp on the baseline model over the same six languages—i.e. merging still reduces WER, just by 0.21 pp less.

Beyond the trained six languages used in the fine-tuning, we evaluate the fine-tuned model on ten held-out languages: four high-resource anchor languages (English, French, German, Spanish) and six untrained mid-/low-resource languages (Thai, Swahili, Afrikaans, Icelandic, Welsh, Kazakh). Fine-tuning preserves accuracy on the high-resource languages (mean drift from baseline +0.02 pp; max +0.29 pp on German) and induces only modest drift on the six untrained languages (mean +1.77 pp; max +3.78 pp on Kazakh). Token merging at TRR = 0.40 remains approximately free across the untrained languages: mean Δ WER = +0.23 pp, with the largest single-

language excursion +0.97 pp (Icelandic). Per-language detail is given in Table 5 of Appendix B.

5.3 Cross-Scale Robustness

We next test whether the cross-lingual robustness of token merging reported above persists across model scales. Figure 2 plots Δ WER on the six languages at every TRR for three model variants, including Whisper-small (244M), Whisper-medium (769M), and Whisper-large-v3 (1.55B). The six languages, Vietnamese, Tamil, Javanese, Maltese, Lingala, and Hausa, were chosen as representative mid-/low-resource languages spanning a wide range of baseline WER (14% to 89% on Whisper-medium). At TRR = 0.40, the mean Δ WER across the six languages is +0.08 pp for Whisper-small, -0.50 pp for Whisper-medium, and +0.32 pp for Whisper-large-v3. The magnitude of the mean change stays below half a percentage point at every scale. Table 2 reports the results at TRR = 0.40 on the six languages; per-language and per-TRR detail behind the figure is given in Table 4 of Appendix B.

Several structural patterns emerge. Vietnamese, which has the lowest baseline WER of the six languages, changes little at any model scale ($|\Delta$ WER| < 0.5 pp at TRR = 0.40). Whisper-large-v3 carries the only consistent per-language costs above one pp at TRR = 0.40—Javanese (+1.65 pp) and Hausa (+1.07 pp), neither of which appears at the smaller scales.

Figure 3 reveals the structural basis for why token merging is effective for the multilingual ASR models. The per-layer redundancy profile follows the same characteristic shape in every language we test: high adjacent-token similarity in the early and middle encoder layers, tapering in the deepest layers. The curves cluster tightly across languages (cross-language standard deviation $\leq$ 0.03 at every layer). This shared, language-agnostic redundancy structure is precisely what token merging exploits.

5.4 Computational Efficiency

We measure the computational efficiency using the inference latency of Whisper-medium under token merging relative to a baseline model without token merging, evaluated at token-reduction ratios of 0.20, 0.30, and 0.40. All measurements use CUDA event timers bracketed by device synchronization rather than host-side wall-clock calls. Each condition is preceded by discarded warm-up iterations, and a global warm-up drives the GPU to its sustained clock before timing begins, so the

Table 2: Whisper-medium vs the DoRA-fine-tuned variant at TRR = 0.40 on the six trained target languages. For each model we report the unmerged baseline WER (%) and the merging effect $\Delta@0.40$ (pp, vs that model’s own baseline). Mean row (**bold**) averages all six languages.

Language	Whisper-medium		Whisper-medium DoRA	
	Base WER (%)	$\Delta@0.40$	Base WER (%)	$\Delta@0.40$
Vietnamese	14.3	−0.09	14.1	+0.17
Tamil	35.7	−0.20	36.0	−1.41
Javanese	67.2	−0.88	51.2	−0.41
Maltese	82.2	+0.02	75.9	+0.13
Lingala	83.2	−1.41	52.2	+0.59
Hausa	88.8	−0.40	75.3	−0.84
Mean (6)	61.9	−0.50	50.8	−0.29




Figure 2: Per-language ΔWER (pp) versus token reduction ratios (TRR) on the six languages, relative to each model’s baseline without token merging; the rightmost panel in the bottom row is the mean across the six languages. Four lines per panel: Whisper-small, Whisper-medium, Whisper-medium with DoRA fine-tuning, and Whisper-large-v3. The mean stays within ± 0.75 pp at every TRR; the largest single-language excursion is Javanese on Whisper-large-v3 ($\leq +2.2$ pp; $+1.65$ pp at TRR = 0.40).




Figure 3: Per-layer adjacent K-cosine similarity at the three encoder scales, with no merging applied ($n=264/\text{lang}$, 16 languages). Light-grey lines are per-language profiles, the bold coloured line is the mean across languages, and dashed verticals mark the encoder layers at which we merge.

baseline (always measured first) is not penalized by clock ramp-up. Baseline and token-merged con-

ditions for a given utterance are measured back-to-back so they share thermal and memory state,

and speedups are computed as per-utterance ratios that are then averaged across utterances, which cancels inter-sample thermal drift. Results are over six languages (Vietnamese, Tamil, Javanese, Maltese, Lingala, Hausa), eight utterances each, on an NVIDIA RTX 3080 Ti (PyTorch 2.6.0, Transformers 5.5.3, float32, SDPA attention).

We examine encoder speedups relative to the unmerged baseline, which is a stable ≈ 94 ms across all six languages. Token merging yields a monotonic encoder speedup that grows with the reduction ratio, reaching $1.27\times$ at $\text{TRR} = 0.40$ (encoder time $94 \rightarrow 74$ ms). This result is highly consistent across languages (per-language range $1.26\text{--}1.27\times$), reflecting that encoder cost is governed by the merge schedule rather than by linguistic content. Hence, token merging delivers a reliable $1.27\times$ encoder speedup at $\text{TRR} = 0.40$ with negligible cross-language variance. The encoder speedup with token merging reaches $1.18\times$ at $\text{TRR} = 0.30$ and $1.10\times$ at $\text{TRR} = 0.20$. Because the savings are encoder-localized, the largest end-to-end gains are expected in deployment settings dominated by encoder cost (e.g., large batch sizes or short transcripts).

6 Discussion

Fine-tuning unlocks accuracy and composes with merging. DoRA fine-tuning on our six target languages yields substantial baseline improvements for four of them: Lingala (-31 pp), Javanese (-16 pp), Hausa (-13 pp), and Maltese (-6 pp), while Vietnamese and Tamil remain essentially unchanged. Crucially, applying token merging on top of this fine-tuning induces a maximum WER degradation of just +0.59 pp across the trained languages at a token reduction ratio (TRR) of 0.40, and at most +0.97 pp across the ten held-out languages. These two efficiency techniques are therefore highly complementary: the DoRA adapter unlocks massive accuracy gains for languages that traditionally perform poorly in multilingual ASR models, while token merging reduces inference costs almost for free on the resulting model—generalizing seamlessly across both seen and unseen languages.

When to apply token merging. Our findings indicate that deployment recommendations are safely decoupled from the target language. The same fixed merging schedule and TRR setting hold across all sixteen evaluated languages, as well as across both the stock and fine-tuned check-

points, eliminating the need for language-specific or cohort-specific recalibration. Setting $\text{TRR} = 0.40$ maintains a strict 1.5 pp WER budget across every language in our main simulation and all high-resource anchor languages, delivering a reliable $1.27\times$ encoder speedup. For tighter accuracy constraints, lower TRR values trade compute savings for an additional safety margin. Consequently, practitioners can establish the compute-accuracy trade-off a single time at deployment based entirely on their global WER budget, rather than relying on granular, language-by-language fine-tuning.

7 Conclusion

In this work, we presented a systematic evaluation of adjacent token merging for multilingual Whisper models. Our study encompassed three encoder scales (spanning an order of magnitude in parameter count), sixteen typologically diverse languages, and decoder-side DoRA fine-tuning. We demonstrated that adjacent token merging induces negligible word error rate (WER) degradation across all tested model sizes and, in low-resource language scenarios, even marginally improves accuracy for some languages. Furthermore, the technique composes seamlessly with adapter-based fine-tuning across both target and held-out languages. At the most aggressive reduction setting, merging eliminates roughly a quarter of the theoretical encoder compute, which translates to a reliable empirical speedup. Consequently, token merging serves as a highly practical, training-free, and post-hoc efficiency method for deploying multilingual ASR systems.

Limitations

We acknowledge two primary limitations to our study. First, we employ a static merging schedule across all languages and model scales, so a single conservative rate has to cover the hardest clip in the batch even when most clips could tolerate more aggressive merging. A natural next step is to train a small neural network that takes the clip’s audio features as input and outputs a reduction rate for each merge layer, so every clip gets its own schedule instead of the shared one used here. This network can be trained by comparing the transcripts Whisper produces under its predicted rates against the transcripts the unmerged Whisper produces, rewarding rates that merge aggressively while leaving the transcript unchanged. Furthermore, our fine-tuning

composition experiments rely exclusively on the DoRA approach; it remains to be tested whether this additive performance result extends to other parameter-efficient adaptation methods, such as prefix tuning or alternative adapter architectures.

Second, while our held-out evaluation spans ten languages across diverse resource tiers, this subset represents only a fraction of Whisper’s complete pretraining corpus. Consequently, our findings should be interpreted as strong supporting evidence of broad generalization rather than a conclusive proof. Validating the observed trade-off between accuracy and computation speed on a wider array of typologically distant languages, particularly those severely underrepresented in the original pre-training data, would further support these claims.

References

- Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in neural information processing systems*, 33:12449–12460.
- Max Bain, Jaesung Huh, Tengda Han, and Andrew Senior. 2023. [WhisperX: Time-accurate speech transcription of long-form audio](#). In *Interspeech 2023*, pages 4489–4493.
- Daniel Bolya, Cheng-Yang Fu, Xiaoliang Dai, Peizhao Zhang, Christoph Feichtenhofer, and Judy Hoffman. 2023. Token merging: Your ViT but faster. In *International Conference on Learning Representations*.
- Alexis Conneau, Alexei Baevski, Ronan Collobert, Abdelrahman Mohamed, and Michael Auli. 2021. Un-supervised cross-lingual representation learning for speech recognition. In *Interspeech*, pages 2426–2430.
- Alexis Conneau, Min Ma, Simran Khanuja, Yu Zhang, Vera Axelrod, Siddharth Dalmia, Jason Riesa, Clara Rivera, and Ankur Bapna. 2023. Fleurs: Few-shot learning evaluation of universal representations of speech. In *2022 IEEE Spoken Language Technology Workshop (SLT)*, pages 798–805. IEEE.
- Sanchit Gandhi, Patrick von Platen, and Alexander M. Rush. 2023. [Distil-Whisper: Robust knowledge distillation via large-scale pseudo labelling](#). *Preprint*, arXiv:2311.00430.
- Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, and Ruoming Pang. 2020. Conformer: Convolution-augmented transformer for speech recognition. In *Interspeech*, pages 5036–5040.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*.
- Keisuke Kamahori, Jungo Kasai, Noriyuki Kojima, and Baris Kasikci. 2025. LiteASR: Efficient automatic speech recognition with low-rank approximation. In *Findings of the Association for Computational Linguistics: EMNLP 2025*. ArXiv:2502.20583.
- Yuang Li, Yu Wu, Jinyu Li, and Shujie Liu. 2023. Accelerating transducers through adjacent token merging. In *Interspeech*. ArXiv:2306.16009.
- Siyu Liang, Nicolas Ballier, Gina-Anne Levow, and Richard Wright. 2025. [Beyond WER: Probing Whisper’s sub-token decoder across diverse language resource levels](#). *Preprint*, arXiv:2509.25516.
- Youwei Liang, Chongjian Ge, Zhan Tong, Yibing Song, Jue Wang, and Pengtao Xie. 2022. Not all patches are what you need: Expediting vision transformers via token reorganizations. In *International Conference on Learning Representations*.
- Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-Ting Cheng, and Min-Hung Chen. 2024. DoRA: Weight-decomposed low-rank adaptation. In *Proceedings of the 41st International Conference on Machine Learning*. PMLR.
- Lingchen Meng, Hengduo Li, Bor-Chun Chen, Shiyi Lan, Zuxuan Wu, Yu-Gang Jiang, and Ser-Nam Lim. 2022. AdaViT: Adaptive vision transformers for efficient image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Vineel Pratap, Andros Tjandra, Bowen Shi, Paden Tomasello, Arun Babu, Sayani Kundu, Ali Elkahky, Zhaozheng Ni, Apoorv Vyas, Maryam Fazel-Zarandi, Alexei Baevski, Yossi Adi, Xiaohui Zhang, Wei-Ning Hsu, Alexis Conneau, and Michael Auli. 2024. Scaling speech technology to 1,000+ languages. *Journal of Machine Learning Research*, 25(97):1–52.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. Robust speech recognition via large-scale weak supervision. In *Proceedings of the 40th International Conference on Machine Learning*, pages 28492–28518. PMLR.
- Yongming Rao, Wenliang Zhao, Benlin Liu, Jiwen Lu, Jie Zhou, and Cho-Jui Hsieh. 2021. DynamicViT: Efficient vision transformers with dynamic token sparsification. In *Advances in Neural Information Processing Systems*, volume 34, pages 13937–13949.
- Yaya Sy, Christophe Cerisara, and Irina Illina. 2025. [BaldWhisper: Faster Whisper with head shearing and layer merging](#). *Preprint*, arXiv:2510.08599.

- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30.
- Bajian Xiang, Tingwei Guo, Xuan Chen, and Yang Han. 2026. [Do we need distinct representations for every speech token? unveiling and exploiting redundancy in large speech language models.](#) *Preprint*, arXiv:2604.06871.
- Yang Xiao, Eun-Jung Holden, and Ting Dang. 2026. [Adapting where it matters: Depth-aware adaptation for efficient multilingual speech recognition in low-resource languages.](#) *Preprint*, arXiv:2602.01008.
- Hongxu Yin, Arash Vahdat, Jose M. Alvarez, Arun Mallya, Jan Kautz, and Pavlo Molchanov. 2022. A-ViT: Adaptive tokens for efficient vision transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

A Theoretical FLOP Derivation

The encoder-FLOP figures reported in §5.4 (24.3–25.1% reduction at TRR=0.40 across model scales) are computed by simulating the integer per-layer arithmetic ($m = \lfloor (T_{\text{pre}} - 1) \cdot r \rfloor$) and summing, at each encoder layer, attention FLOPs at the pre-merge length and feed-forward FLOPs at the post-merge length:

$$\begin{aligned}\text{Attn}(T) &= 8Td^2 + 4T^2d, \\ \text{FFN}(T) &= 4Td \cdot d_{ff}.\end{aligned}\tag{4}$$

The fixed cost of the two-layer convolutional stem is included in both the merged and baseline totals so that the reported reduction is honest with respect to the encoder as a whole. Layer-norm contributions are below 0.1% of layer FLOPs and are omitted per the standard convention (Bolya et al., 2023).

The gap between the nominal sequence reduction (40%) and the realised FLOP reduction ($\sim 25\%$) at TRR=0.40 is structural. Only a subset of layers ($|\mathcal{M}_L|$ out of L) perform merging, layers before the first merge always operate at the full length T_0 , and the feed-forward term scales linearly rather than quadratically in T , dominating the per-layer cost at our d_{ff}/d ratio of 4.

B Per-Language Tables

Table 3: 12-language main sweep on Whisper-medium with merging at the schedule of Eq. 2. Columns 0.05–0.40 are ΔWER (pp) at each token reduction ratio; baseline WER (%) at TRR=0 shown alongside. The mean row (**bold**) averages all twelve languages.

Language	Base WER (%)	0.05	0.10	0.20	0.30	0.40
Vietnamese	14.3	+0.02	−0.11	−0.30	−0.09	−0.09
Tamil	35.7	+0.16	+0.50	−0.10	+0.13	−0.20
Welsh	36.9	−0.30	−0.43	−0.58	−0.56	−0.25
Thai	40.6	+0.34	−0.03	+0.08	+0.01	−0.27
Afrikaans	45.1	+0.05	+0.17	−0.02	−0.19	−0.24
Icelandic	49.8	−0.19	−0.48	−0.10	+0.10	−0.39
Swahili	51.7	+0.88	+0.37	+0.83	+1.34	+1.01
Kazakh	53.4	+1.09	+0.88	+0.73	+0.90	+0.54
Javanese	67.2	−1.82	−1.89	−1.37	−1.76	−0.88
Maltese	82.2	+0.49	+0.36	+0.33	+0.03	+0.02
Lingala	83.2	−1.31	−0.59	−2.78	−1.48	−1.41
Hausa	88.8	−0.23	−0.80	−0.06	−0.14	−0.40
Mean (12)	54.1	−0.07	−0.17	−0.28	−0.14	−0.21

Table 4: Per-language Δ WER (pp) on the six trained languages at each TRR, for each of the three stock model scales plus the DoRA fine-tuned medium variant. Baseline (TRR=0) WER (%) is shown in parentheses next to each language. The mean row (**bold**) averages over the six trained languages.

Model	Language	0.05	0.10	0.20	0.30	0.40
Whisper-small	Vietnamese (21.9)	+0.12	+0.25	-0.14	+0.01	-0.14
	Tamil (43.5)	-0.46	-0.16	-0.84	-0.76	-0.73
	Javanese (82.2)	+0.34	+1.16	+0.26	+0.41	0.00
	Maltese (89.3)	+0.02	+0.17	+0.08	+0.21	+0.19
	Lingala (90.4)	-0.24	-0.05	-0.07	+0.23	+0.82
	Hausa (88.4)	-0.10	+0.23	+0.30	+0.45	+0.36
	Mean	-0.06	+0.27	-0.07	+0.09	+0.08
Whisper-medium	Vietnamese (14.3)	+0.02	-0.11	-0.30	-0.09	-0.09
	Tamil (35.7)	+0.16	+0.50	-0.10	+0.13	-0.20
	Javanese (67.2)	-1.82	-1.89	-1.37	-1.76	-0.88
	Maltese (82.2)	+0.49	+0.36	+0.33	+0.03	+0.02
	Lingala (83.2)	-1.31	-0.59	-2.78	-1.48	-1.41
	Hausa (88.8)	-0.23	-0.80	-0.06	-0.14	-0.40
	Mean	-0.45	-0.42	-0.71	-0.55	-0.50
med. DoRA	Vietnamese (14.1)	+0.19	+0.12	+0.14	+0.14	+0.17
	Tamil (36.0)	-1.76	-1.55	-1.76	-1.32	-1.41
	Javanese (51.2)	0.00	-0.52	-0.36	-0.79	-0.41
	Maltese (75.9)	+0.38	+0.14	-0.08	-0.10	+0.13
	Lingala (52.2)	-0.21	+0.12	+0.02	+0.44	+0.59
	Hausa (75.3)	-0.43	-0.48	-0.82	-0.58	-0.84
	Mean	-0.31	-0.36	-0.48	-0.37	-0.29
large-v3	Vietnamese (10.8)	+0.41	+0.36	-0.16	-0.15	-0.19
	Tamil (32.1)	-0.46	-0.21	-0.60	-0.88	-0.18
	Javanese (60.1)	+1.87	+2.04	+2.19	+1.44	+1.65
	Maltese (68.1)	-0.25	-0.47	-0.74	-0.55	-0.40
	Lingala (70.1)	+0.37	+0.05	-0.05	-0.28	-0.02
	Hausa (83.5)	+0.68	+1.27	+1.08	+1.30	+1.07
	Mean	+0.44	+0.51	+0.29	+0.15	+0.32

Table 5: Held-out cohort merging sweep on the DoRA-fine-tuned Whisper-medium model ($n=264/\text{lang}$). Columns 0.05–0.40 are Δ WER (pp) at each TRR vs the DoRA-fine-tuned unmerged baseline; “Base WER (%)” is the DoRA unmerged WER at TRR=0. Four high-resource anchors above the rule; six untrained mid/low-resource languages below. Mean rows (**bold**) report each tier separately and the cohort overall.

Language	Base WER (%)	0.05	0.10	0.20	0.30	0.40
English	4.9	+0.05	+0.07	+0.11	+0.12	+0.05
French	9.0	+0.11	+0.06	+0.01	+0.01	+0.10
German	6.9	+0.15	+0.15	+0.15	+0.10	+0.17
Spanish	3.4	+0.10	+0.15	+0.13	+0.15	+0.10
Mean (4 anchors)	6.1	+0.10	+0.11	+0.10	+0.10	+0.11
Thai	42.4	-0.06	-0.26	+0.16	-0.37	-0.08
Welsh	38.3	+0.31	+0.43	+0.33	+0.55	+0.25
Afrikaans	44.8	+0.39	+0.38	+0.32	+0.12	+0.15
Icelandic	51.2	+1.06	+1.26	+1.06	+1.84	+0.97
Swahili	54.2	+0.57	+0.40	+0.59	+0.72	+0.46
Kazakh	57.2	+0.75	+0.41	+0.52	+0.30	+0.13
Mean (6 untrained)	48.0	+0.50	+0.44	+0.50	+0.53	+0.31
Mean (all 10)	31.2	+0.34	+0.31	+0.34	+0.35	+0.23