WHAT LAYERS WHEN: LEARNING TO SKIP COMPUTE IN LLMS WITH RESIDUAL GATES

Anonymous authors

Paper under double-blind review

ABSTRACT

We introduce GateSkip, a simple residual-stream gating mechanism that enables token-wise layer skipping in decoder-only LMs. Each Attention/MLP branch is equipped with a sigmoid-linear gate that condense the branch's output before it re-enters the residual stream. During inference we rank tokens by the gate values and skip low-importance ones using a per-layer budget. While early-exit or router-based Mixture-of-Depths models are known to be unstable and need extensive retraining, our smooth, differentiable gates fine-tune stably on top of pretrained models. On long-form reasoning, we save up to 15% compute while retaining >90% of baseline accuracy. On instruction-tuned models we see accuracy gains at full compute and match baseline quality near 50% savings. The learned gates give insight into transformer information flow (e.g., BOS tokens act as anchors), and the method combines easily with quantization, pruning, and self-speculative decoding.

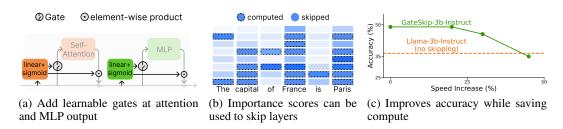


Figure 1: We introduce gating mechanisms that regulate the flow of information into the residual stream and can be used to skip layers altogether. Our mechanism enhances downstream accuracy of instruction-tuned models even when skipping \sim 25% of the model.

1 Introduction

Large language models have transformed natural language processing, yet their rapid growth has created major challenges for efficient deployment. Current models allocate the same amount of computation to every token at every layer, regardless of difficulty. This uniform allocation is wasteful and makes it hard to deploy models in latency-sensitive or resource-limited environments. Adaptive compute aims to address this by using more resources where they matter and less where they do not.

Most prior approaches fall into two categories. Router-based methods, such as Mixture-of-Depths (Raposo et al., 2024), introduce specialized routing components that decide which layers to run. These rely on hard, discrete decisions that are often unstable and require careful balancing losses (Zoph et al., 2022). Early-exit methods attach auxiliary language modeling heads at intermediate layers and stop once a confidence threshold is reached (Schuster et al., 2022). These approaches alter pretrained hidden states, complicate training, and often fail to calibrate well (Bajpai & Hanawal, 2024a). Both approaches usually require implementation during pre-training.

We introduce **GateSkip**, a lightweight residual stream gating mechanism for decoder-only transformers. Each attention and MLP branch is equipped with a small linear gate and sigmoid activation that squashes the branch output before it is added back to the residual stream. During training, gates are optimized to remain sparse while preserving language modeling accuracy. At inference, token-level

importance scores derived from the gates allow us to retain only the top tokens per layer using a quantile threshold, while skipped tokens copy their hidden states and key-value cache entries upward.

This design has several advantages. Because the gates are smooth and differentiable, GateSkip can be trained directly on top of pretrained models without destabilizing optimization. Since it operates entirely within the residual stream, it minimally perturbs existing representations. Moreover, the mechanism provides fine-grained control at both the token and module level, enabling nuanced allocation of compute. Finally, GateSkip is fully compatible with orthogonal efficiency techniques such as quantization, pruning, and self-speculative decoding.

We evaluate GateSkip on Llama 3.1 up to 8B parameters and Gemma 2 2B models across generative reasoning and log-likelihood benchmarks. On long-form reasoning tasks, GateSkip reduces computation by up to 15% while retaining more than 90% of the original accuracy. On instruction-tuned models, it improves accuracy at full compute and sustains this improvement under reduced budgets, matching baseline quality even with roughly 50% compute savings. Analysis of the learned gate values further reveals consistent patterns: early layers allocate more computation to the beginning-of-sequence token and punctuation, while deeper layers become increasingly selective and focus on content-bearing words.

Our contributions can be summarized as follows:

- 1. We propose GateSkip, a residual gating mechanism that enables token-wise layer skipping with smooth training and deterministic hard decisions at inference.
- We demonstrate state-of-the-art compute-accuracy trade-offs on generative reasoning tasks, where prior adaptive compute methods often collapse.
- 3. We show that GateSkip composes seamlessly with quantization, pruning, and self-speculative decoding.
- 4. We provide an analysis of gate activations that sheds light on information flow within transformers.

GateSkip turns the residual stream into a simple yet effective control mechanism for adaptive depth, delivering efficiency gains without sacrificing stability or performance.

2 Related Work

A growing body of work has addressed the inefficiency of ever-larger decoder-only Transformers by dynamically adapting computation on a per-token or per-sequence basis. Early efforts in **layer skipping** repurpose sparse routing from Mixture-of-Experts (MoE) to decide whether to execute each layer at all, yielding substantial compute savings without restarting training from scratch. Mixture-of-Depths (MoD) injects a router at every transformer layer to skip unimportant layers for each token (Raposo et al., 2024), while follow-up methods introduce soft token budgets (Zeng et al., 2023), sequence-level skipping (Wang et al., 2023), and frozen-backbone router fine-tuning (He et al., 2024). However, discrete routers can be unstable (Zoph et al., 2022; Fedus et al., 2022; Puigcerver et al., 2024; Panda et al., 2025), and most require training from scratch.

In contrast to hard top-k routers with load-balancing losses, GateSkip uses fully differentiable residual gating, avoiding discrete routing during training while still yielding hard skips at test time.

In contrast, **early exiting** methods terminate inference once intermediate representations are deemed confident enough. Pioneered in encoder-only BERT models via entropy or agreement thresholds (Xin et al., 2020; Zhang et al., 2022; Zhou et al., 2020; Liu et al., 2020), this concept has been extended to encoder–decoder and decoder-only settings by supervising every layer with an auxiliary language modeling objective and exiting based on heuristic confidence measures (Tang et al., 2023; Schuster et al., 2022; Elbayad et al., 2020; Liu et al., 2021; Bae et al., 2023; Elhoushi et al., 2024; Del Corro et al., 2023). Yet these approaches fundamentally alter pretrained hidden representations through connector modules or self-distillation (Bajpai & Hanawal, 2024a;b; Ji et al., 2023).

Where early exit supervises intermediate LM heads and exits on confidence thresholds, GateSkip avoids auxiliary heads and maintains pretrained hidden spaces by training gates to compress post-module outputs.

Other efficiency techniques include **layer pruning** (e.g., ShortGPT removes redundant transformer blocks (Men et al., 2025)), **KV cache and token pruning** (e.g., query-driven pruning (Xu et al., 2025), token–precision trade-offs (Zhang et al., 2024)), and **quantization** (e.g., AWQ (Lin et al., 2023), SpQR (Dettmers et al., 2024), BiLLM (Huang et al., 2024)). These methods operate on different axes of efficiency and are thus orthogonal to our depth-adaptive approach. We further show that GateSkip is compatible with pruning as well as quantization, demonstrating its ability to combine with such methods.

3 GATESKIP

We propose adding a gating mechanism to the residual stream of decoder-only Transformer models and training it with an additional sparsity loss so that the gates learn to assess the importance of a certain Attention or MLP module given its preceding hidden state, as shown in Figure 1.

3.1 RESIDUAL GATING MECHANISM

The residual stream at layer ℓ in a transformer model can be described as the output $o_{\ell} \in \mathbb{R}^{B \times S \times H}$ of an Attention or MLP layer added to the hidden states $h_{\ell} \in \mathbb{R}^{B \times S \times H}$, resulting in the layer output $h_{\ell+1}$, with B being the batch size, S the sequence length, and H the hidden dimension:

$$h_{\ell+1} = h_{\ell} + o_{\ell} \tag{1}$$

We propose supplementing the language model with a trainable gate g which is a sigmoid-activated linear projection of the hidden states h_{ℓ} :

$$h_{\ell+1} = h_{\ell} + o_{\ell} \odot g_{\ell}(h_{\ell}), \quad g_t(h_{\ell}) = \sigma(W_G h_{\ell} + b)$$
(2)

where $W_G \in \mathbb{R}^{H \times H}$ and $b \in \mathbb{R}^H$, σ refers to the sigmoid function, and g_ℓ refers to the gate at layer ℓ which could theoretically be a shared gate across layers (cf. ablation experiments in §4.4) or separate gates for each layer. The gate is placed at the exit point of the module to the residual stream, after the output projection, making it perfectly compatible with multi-head attention or any variant thereof.

3.2 Training Objective

Training minimises a standard language—model loss (cross-entropy for next-token prediction)

$$\mathcal{L}_{CE} = -\frac{1}{|B|} \sum_{(x,y) \in B} \log p_{\theta}(y \mid x)$$
(3)

plus an explicit gate-sparsity penalty (L2 distance on gate activations)

$$\mathcal{L}_{S} = \frac{1}{N_{L}H} \sum_{\ell=1}^{N_{L}} \sum_{k=1}^{H} \|g_{\ell}(h_{\ell})_{k}\|_{2}$$
(4)

so that the overall loss becomes

$$\mathcal{L} = \mathcal{L}_{CE} + \lambda_S \, \mathcal{L}_{S}. \tag{5}$$

Here N_L is the number of layers, H the hidden dimension, and λ_S balances accuracy and efficiency. Term (2) encourages each sigmoid gate $g_\ell(h_\ell)$ to stay close to zero, effectively **compressing the module output before it is re-added to the residual stream**. Backbone parameters θ and gate parameters are updated jointly with AdamW, with all weights being trainable.

3.3 TOKEN SELECTION

At step t we allot a *fractional* budget $b_t \in (0,1]$, the share of the L tokens that may be *processed* in the current layer. For every token we collapse the current batch's gate vectors to scalar importance scores $\bar{g}_{\ell,i} = \frac{1}{H} \sum_{k}^{H} g_{\ell}(h_{\ell})_{i,k}$ and form their empirical cumulative distribution function (CDF). We then compute the *threshold*,

$$\tau_t = \text{Quantile}(\{\bar{g}_{\ell,i}\}_{i=1}^L, 1 - b_t)$$
 (6)

		Gen	erative	Benchn	narks		Log-Likelihood Benchmarks				
saved compute	0%	5%	10%	15%	20%	25%	0%	15%	30%	45%	60%
Llama-1b	31.3	-	-	-	_	-	45.0	-	-	-	-
Llama-1b (random skipping)	-	9.7	1.9	3.0	1.5	1.1	-	32.0	31.5	31.1	31.0
CALM (hidden state saturation)	0.6	0.7	0.6	0.4	0.5	0.6	34.9	32.7	31.9	30.7	31.1
CALM (softmax)	0.8	0.9	1.0	0.3	0.2	0.4	34.9	32.3	30.8	31.1	31.1
FREE (hidden state saturation)	11.9	11.9	11.9	11.9	11.9	11.9	38.6	38.6	38.6	38.6	38.6
FREE (softmax)	11.9	11.9	11.9	11.9	11.9	11.9	38.6	38.6	38.6	38.6	38.6
LayerSkip	13.1	13.1	13.1	13.1	13.1	13.1	38.3	38.3	38.3	38.3	38.3
MoD (router-tuned)	20.1	17.3	14.5	7.7	5.8	4.0	40.6	35.4	34.5	31.3	29.6
SkipLayer	2.2	1.4	0.7	0.0	0.0	0.0	31.2	30.7	30.2	30.2	30.3
GateSkip	26.8	25.5	24.2	23.2	23.6	19.8	44.3	37.8	32.7	30.8	31.2

Table 1: Averaged results for loglikelihood-based and longer generation benchmarks for a random skipping baseline, prior adaptive compute methods and GateSkip on Llama-3.2-1b.

using linear interpolation between adjacent order-statistics (see Algorithm 3); thus the expected fraction of scores below τ_t equals the desired skip ratio $1-b_t$. Tokens with $\bar{g}_{\ell,i} \leq \tau_t$ are **skipped**, copying the hidden state upwards $h_{\ell+1,i} = h_{\ell,i}$, and the rest is processed normally.

During **training** (see Algorithm 1) the budget decays linearly, as $b_t = b_1 - (b_1 - b_2) \frac{t}{T_{\text{lotal}}}$, so that the model learns to tolerate skipped hidden states. During **inference** (see Algorithm 2) we fix a single budget \hat{b} once, re-use the same post-module gate scores for ranking, and apply the Top-k only to tokens that have not yet emitted the end-of-sequence symbol. Additionally, when a token skips a layer, we upwards copy the KV cache items from the layer below in order to facilitate KV-cache reuse.

3.4 IMPLEMENTATION DETAILS

We initialize the gates to ensure the model initially closely resembles the original pre-trained model. Specifically, we initialize the weights of the linear matrix W_G around 0 using a Gaussian distribution with low standard deviation $\sigma=0.01$, and set the biases b to 5, so that the module's output remains approximately unchanged (as $\sigma(5)\approx 1$).

A key design decision is to place the gate at the module input for skipping decisions, but to train it on the module's output, i.e. we multiply the gate element-wise with the module output so that it receives its gradient signal downstream of the module. While the input to the gate is the same in both cases (the hidden states h_{ℓ}), the learning signal differs compared to if the gate was placed at the entry point to the module. The gate is trained to incorporate minimal information from the module's output into the residual stream while maintaining language modeling performance, rather than determining which information should enter a module. We empirically found that training the gate after the module leads to better downstream performance (see §4).

Our method is numerically stable compared to other techniques based on hard binary routing (such as Mixture-of-Depths) (Zoph et al., 2022; Fedus et al., 2022; Puigcerver et al., 2024; Panda et al., 2025), providing effective control of information flow without introducing training instabilities or convergence issues (see §4 for experimental results).

The added gates introduce negligible parameter overhead to the model, e.g. 0.004% for separate gates with scalar output, and 4% for separate gates with hidden-state sized output on Llama-3.2-1b. Table 13 in Appendix G shows parameter overhead for each of the variants of GateSkip.

4 EXPERIMENTS

We evaluate GateSkip on Llama-3 (Meta-AI, 2024) models of varying size as well as on Gemma 2 (Gemma-Team, 2024). We then perform ablation studies to isolate the impact of each component, compare against state-of-the-art layer-skipping and early-exit methods, and demonstrate compatibility with 4-bit quantization, self-speculative decoding, and structured pruning.

	CSQA (Gen.)	GSM8K (Gen.)	MMLU Stem	HellaSwag	CSQA	PIQA	Open- BookQA	Wino- Grande
Llama-1b (random skipping)	6.0	0.0	24.9	29.0	19.0	49.3	17.3	52.5
CALM (hidden state saturation)	0.2	0.7	23.4	29.4	18.9	56.2	15.2	32.7
CALM (softmax)	0.0	0.8	25.7	29.3	19.3	53.2	15.0	51.2
MoD (router-tuned)	8.8	6.5	26.5	35.0	18.5	58.8	20.3	53.5
GateSkip	36.7	9.7	26.0	38.5	21.2	66.5	21.3	53.0

Table 2: Accuracy at 15% saved compute for log-likelihood-based and generative benchmarks for a random skipping baseline, prior adaptive compute methods and GateSkip on Llama-3.2-1b.

4.1 EXPERIMENTAL SETUP

Models and Training. We primarily evaluate our method on Llama-3.2-1b, while also experimenting with Llama-3.2-3b, Llama-3.1-8b, and Gemma-2-2b to assess scalability and architecture independence. For all experiments, we fine-tune the pretrained backbone to train the gates while simultaneously adapting the model to task templates for easier downstream answer extraction. We set the sparsity loss weight $\lambda=0.1$ and decay the token budget from 100% to 80% during training. Training employs the AdamW optimizer (Loshchilov & Hutter, 2017). A full list of hyperparameters can be found in Appendix L. All libraries and their respective versions used for our experiments are listed in Appendix M. Instructions for code access can be found in Appendix I.

Generative benchmarks. We fine-tune on the train sets of CommonsenseQA (Talmor et al., 2019) and GSM8K (Cobbe et al., 2021) questions with chain-of-thought traces generated by Nemotron-70B (Anonymous, 2024; Reasoning, 2024; Wang et al., 2024), masking the loss on the question portion so that the model learns both reasoning and answer extraction. For evaluation we measure zero-shot accuracy on the GSM8K and CommonsenseQA test sets using the same prompt template. We sweep the inference budget \hat{b} over $\{1.00, 0.95, 0.90, 0.85, 0.80, 0.75\}$, corresponding to compute-savings of $\{0\%, 5\%, 10\%, 15\%, 20\%, 25\%\}$; since the realized savings sometimes fall between these targets, we linearly interpolate the measured accuracies to report performance at the exact percentages listed above. Non-interpolated results are reported in Appendix K.

Log-likelihood benchmarks. We fine-tune on FineWeb data (Penedo et al., 2024) with the same hyperparameters as above. For evaluation we measure five-shot log-likelihood accuracy on MMLU (Hendrycks et al., 2021), HellaSwag (Zellers et al., 2019), CommonsenseQA (Talmor et al., 2019), PIQA (Bisk et al., 2019), OpenBookQA (Mihaylov et al., 2018), and Wino-Grande (Sakaguchi et al., 2021) using LM Evaluation Harness (Gao et al., 2024). We sweep \hat{b} over $\{1.00, 0.85, 0.70, 0.55, 0.40\}$ for compute-savings of $\{0\%, 15\%, 30\%, 45\%, 60\%\}$, again using linear interpolation to report exact savings levels.

Variance estimate. Each full run requires ≈ 5 GPU-hours on an H100-80GB, so we perform one seed per configuration and estimate uncertainty by bootstrapping (100 000 resamples in LM-Eval-Harness), yielding a standard deviation $\leq 3\%$ on every benchmark. Standard deviations and raw results are reported in Appendix K. All baselines (random skipping, MoD router-tuning, CALM variants) were trained and evaluated with the identical data splits, hyperparameters, inference budgets, and interpolation procedure described above, ensuring a fair comparison.

4.2 Comparison to Baseline

We begin by evaluating GateSkip against a straightforward token-level heuristic: *random skipping*. At each layer, a fixed fraction of tokens is selected uniformly at random to be omitted from further computation. All experiments use the Llama-3.2-1b backbone.

Table 1 presents averaged accuracies across multiple compute-savings levels, while Table 2 reports performance at exactly 15% saved compute. Random skipping collapses generative accuracy to under 10% even at modest budgets (5–10% savings), while GateSkip retains the majority of performance. Full results can be seen in Appendix K & J.

Additional results on LAMBADA are provided in Appendix A, showing that GateSkip maintains stable perplexity and accuracy under compute constraints, while random skipping collapses sharply.

2	7	0
2	7	1
2	7	2
2	7	3

		Generative	Benchmarks		Log-Likelihood Benchmarks				
saved compute	Gen@0%	Gen@20%	Gen@30%	Gen@45%	LL@0%	LL@15%	LL@30%	LL@60%	
Llama-3b-Instruct	36.5	-	-	-	46.3	-	-	-	
Llama-3b-Instruct + random skipping	-	0.5	0.1	0.1	-	34.7	30.4	30.4	
GateSkip (Llama-3b-Instruct)	49.0	49.0	45.6	35.0	36.7	38.8	32.9	31.0	

Table 3: GateSkip on Llama Instruct.

2	7	6
2	7	7
2	7	8
2	7	9

		MMLU-Gen					PIQA-Gen				
saved compute	0%	10%	15%	20%	30%	45%	0%	10%	20%	25%	30%
Llama-1b	22.8	-	-	-	-	-	22.9	-	-	-	-
Llama-1b (random skipping)	-	7.5	2.0	1.0	0.3	0.1	-	7.0	0.8	1.2	0.3
GateSkip	14.0	15.6	18.6	15.9	12.8	5.1	14.3	16.9	21.8	29.8	29.5

Table 4: GateSkip on out-of-domain generative tasks.

Moreover, results on translation are presented in Appendix B, showing that GateSkip exhibits analogous performance improvements compared to baseline performance.

4.3 Comparison to Prior State-of-the-Art

Having established the superiority of GateSkip over naive heuristics, we now compare it against prior adaptive-depth methods under identical fine-tuning and evaluation settings: (1) **Mixture-of-Depths** (**MoD**) with Router Tuning, (2) **CALM** in its hidden state saturation and softmax variants, (3) **FREE** in both variants, and (4) static skipping approaches such as **LayerSkip** and **SkipLayer**. Results averaged across different benchmarks as well as task-level accuracy at exactly 15% saved compute can be seen in Tables 1 and 2 respectively.

On generative benchmarks, GateSkip achieves 36.7% accuracy on CSQA and 9.7% on GSM8K—over four times higher than MoD and orders of magnitude above CALM. FREE maintains strong log-likelihood scores across all budgets but remains flat on generation (11.9), while LayerSkip and SkipLayer collapse quickly on longer reasoning, with accuracies of only 13.1 and \leq 2.2 respectively. In contrast, GateSkip sustains high generative accuracy while remaining competitive on log-likelihood evaluations.

Regarding log-likelihood benchmarks, across all metrics GateSkip either matches or outperforms MoD, CALM, and the static baselines, while FREE achieves similar log-likelihood scores but without corresponding generative performance. These results confirm that GateSkip consistently delivers a superior compute–accuracy trade-off across both reasoning and multiple-choice tasks. On instruction-tuned models (Table 3), GateSkip improves generative accuracy over the Llama-3b-Instruct baseline even under aggressive budgets (e.g., +12.5 points at 0–20% saved compute) while also matching or slightly improving log-likelihood accuracy at 15–60% savings.

We tested on log-likelihood-based benchmarks following prior literature. However, real-world scenarios would demand robustness to longer generation which is why we performed such experiments as well. Notably, there is a visible discrepancy between log-likelihood and generative tasks for prior methods, whereas GateSkip retains accuracy significantly better over longer generation.

Beyond the standard suites, out-of-domain generative evaluations (Table 4) show GateSkip retains competitive performance on MMLU-Gen at reduced compute and, notably, exceeds the unadapted baseline on PIQA-Gen at 20–30% saved compute (29.8–29.5 vs. 22.9). This suggests that targeted token-level allocation can translate into quality gains on certain OOD generative tasks, not merely lossless efficiency.

4.4 COMPONENT ABLATIONS

To understand the contribution of each design choice in GateSkip, we perform a series of controlled ablations on Llama-3.2-1b. Table 5 summarizes the impact of varying the gate parameterization, skipping granularity, gate architecture, and gate placement on both our generative and log-likelihood benchmarks at multiple compute-savings levels.

		Gen	erative	Benchn	narks		Lo	g-Likel	ihood B	enchma	rks
Compute saved	0%	5%	10%	15%	20%	25%	0%	15%	30%	45%	60%
Gate output shape											
Scalar gates	23.6	22.7	21.7	20.4	15.9	14.2	42.8	36.8	33.7	30.9	31.8
Vector gates (default)	26.8	25.5	24.2	23.2	23.6	19.8	44.3	37.8	32.7	30.8	31.2
Gate parameter sharing											
Shared	21.8	21.4	21.0	20.7	19.5	15.5	44.3	38.4	35.4	31.8	31.7
Separate (default)	26.8	25.5	24.2	23.2	23.6	19.8	44.3	37.8	32.7	30.8	31.2
Skipping strategy											
Only attention layers	26.8	23.1	19.2	14.9	10.8	6.8	44.3	37.5	30.8	30.6	30.3
Only MLP layers	26.8	24.1	18.6	7.8	1.2	0.4	44.3	32.0	30.4	32.1	32.0
Skip entire layer (attn gate)	26.8	25.5	24.2	23.2	23.6	19.8	44.3	37.8	32.7	30.8	31.2
Every-second layer	26.8	25.5	22.8	15.0	10.1	9.5	44.3	35.5	33.3	31.6	31.9
Skip all layers (default)	26.8	25.5	24.2	23.2	23.6	19.8	44.3	37.8	32.7	30.8	31.2
Gate architecture											
MLP-based gate	24.3	24.3	24.4	18.5	17.0	15.6	44.0	33.9	31.8	30.3	30.7
Linear-sigmoid gate (default)	26.8	25.5	24.2	23.2	23.6	19.8	44.3	37.8	32.7	30.8	31.2
Gate placement											
Gate before module (entry)	21.1	5.5	1.3	1.0	0.9	0.8	40.0	35.7	33.2	31.6	30.8
Gate after module (default)	26.8	25.5	24.2	23.2	23.6	19.8	44.3	37.8	32.7	30.8	31.2

Table 5: Ablation of GateSkip's design choices on Llama-3.2-1b.

General Oberservations. Since we condense the output of modules, downstream performance changes even at no skipping. Hence, different modifications of our method will have differing effects on performance at 0% skipping.

Gate Output Shape. We compare two forms of gating output: (1) *Vector-gates*, which produce an *H*-dimensional output per residual branch, and (2) *Scalar-gates*, which produce a single gating value per branch. At 15% compute-savings on our generative benchmarks, vector-gates achieve 23.2% accuracy, compared to only 20.4% for scalar-gates, confirming that a full-dimensional gate yields more precise control.

Gate Parameter Sharing. Focusing on the vector-gate design, we then compare (1) *Per-layer vector-gates*: a distinct gate for each Attention and MLP module, versus (2) *Shared vector-gates*: one gate shared across all Attention modules and one across all MLP modules. Per-layer vector-gates again lead, with 23.2% at 15% savings, while shared vector-gates lag at 20.7%, showing the value of layer-specific parameters.

Skipping Granularity. Ablating which sub-modules can be skipped reveals that attention and MLP layers are both essential. When only attention layers are skipped, accuracy falls to 14.9% at a 15% compute reduction; skipping only MLP layers drops performance even further, to 7.8% under the same budget. Applying a single gate over the entire layer performs on par with our default per-module approach, but skipping every second layer leads to a steep decline, from 23.2% down to 15.0% at the 15% savings level. These results underscore the importance of fine-grained, per-module control.

Gate Architecture. We also tested a small MLP in place of our linear–sigmoid gate, but despite the extra parameters it underperforms. At 15% compute savings the MLP-based gate achieves only 18.5% on generative tasks compared to 23.2% with the linear gate, and 33.9% on log-likelihood benchmarks versus 37.8%. We therefore retain the simpler, more effective linear projection.

Gate Placement. Placing the gate before the module proved disastrous: at a 5% compute reduction entry-point gating yields just 5.5% accuracy compared to 25.5% when the gate is applied after the module. This dramatic gap confirms that post-module residual gating is crucial for stable, effective learning, as discussed in Section 3.4.

GateSkip on varying model sizes and architectures. To test scalability, we applied GateSkip to larger Llama models (3b and 8b) and observed consistent performance patterns (see Table 6). Full results and plots are shown in Appendix K and J. The results reveal that for larger architectures, the model is capable of skipping increasingly more tokens without decreasing performance. For instance,

		Gen	erative	Benchn	narks		Lo	g-Likeli	ihood B	enchma	ırks
saved compute	0%	5%	10%	15%	20%	25%	0%	15%	30%	45%	60%
Llama-3.2-1B	26.8	25.5	24.2	23.2	23.6	19.8	44.3	37.8	32.7	30.8	31.2
Llama-3.2-3B	45.0	44.4	43.9	43.3	42.7	42.1	55.9	35.6	31.4	29.3	30.4
Llama-3.1-8B	57.3	56.5	55.8	55.0	54.3	53.6	62.8	44.2	34.3	31.1	29.1
Gemma-2-2B	38.0	37.4	36.7	36.1	35.4	34.8	52.9	44.1	35.6	31.7	30.9

Table 6: GateSkip on different model sizes and architectures.

Llama-3.2-3B with GateSkip can save 37.3% computation while retaining 91.5% of its baseline GSM8K (Gen.) performance and 87.3% of its baseline CSQA (Gen.) performance. Moreover, comparisons between Llama and Gemma architectures reveal that the compute-accuracy trade-off generalizes across both model families. The instruction-tuned and LAMBADA results mirror these trends: larger or instruction-adapted backbones benefit more from budgeted token selection, maintaining strong generation quality where random or uniform skipping fails, and confirming that GateSkip's gains persist across usage styles (chat/instruction), lengths, and domains.

4.5 Compatibility with Other Efficiency Techniques

We evaluate compatibility with various orthogonal efficiency techniques. Specifically, we show that GateSkip is compatible with 4-bit quantization, speculative decoding, and structured pruning.

Compatibility with 4-bit Quantization. To test compatibility with quantization, we apply 4-bit quantization to Llama-3.2-3b trained with GateSkip (Table 7) and downstream evaluate the quantized model as before. The results demonstrate that GateSkip remains effective when combined with quantization, with performance curves closely tracking those of the 32-bit model. On generative benchmarks, the quantized model retains 94.4% of the original accuracy at 0% skipping ratio, 96.1% at 15%, and 97.3% at 25% skipping ratio. The performances for log-likelihood-based benchmarks exactly match.

		Gen	erative	Benchn	narks		Lo	g-Likel	ihood B	enchma	ırks
saved compute	0%	5%	10%	15%	20%	25%	0%	15%	30%	45%	60%
32-bit + GateSkip 4-bit + GateSkip											30.4 30.4

Table 7: Quantization robustness (Llama-3.2-3 B).

Compatibility with Speculative Decoding. Adding speculative decoding boosts Log-Likelihood performance substantially at moderate savings: at 15% and 30% saved compute, it outperforms vanilla GateSkip (Table 8).

	LL@15%	LL@30%	LL@45%	LL@60%
GateSkip	37.8	32.7	30.8	31.2
GateSkip + self-speculative decoding	39.4	39.4	31.4	30.7

Table 8: GateSkip combined with self-speculative decoding compared to GateSkip alone and Layer-Skip. Metrics are log-likelihood accuracy (LL) at fixed saved-compute levels.

Compatibility with Structured Pruning. Structured pruning (Men et al., 2024) reduces absolute Log-Likelihood performance, but GateSkip remains notably stronger than the pruned backbone baseline at every budget (Table 9). At 0% savings, *GateSkip + pruning* trails unpruned GateSkip, as expected, yet still exceeds the *Default Llama1b + pruning* by +6.4 LL points.

	LL@0%	LL@15%	LL@30%	LL@45%
GateSkip	44.3	37.8	32.7	30.8
GateSkip + ShortGPT pruning	31.5	31.1	31.9	30.8
Default Llama1b + ShortGPT pruning	25.1	25.4	25.4	26.3

Table 9: GateSkip with and without additional structured pruning of 25% of transformer blocks (ShortGPT).

4.6 END-TO-END EFFICIENCY GAINS IN REAL-WORLD SCENARIOS

% Tokens Skipped	5%	15%	25%	35%	50%	70%
Latency (s)	607.29	606.01	559.81	571.79	521.68	449.85
Throughput (tokens/s)	2697.87	2703.57	2926.71	2865.39	3140.61	3642.11

Table 10: End-to-end latency and throughput at different token skipping levels.

Table 10 shows end-to-end latency and throughput measurements for our Llama-1b GateSkip model evaluated on GSM8K and CommonsenseQA-Gen. The results show that GateSkips theoretical FLOP savings translate into analogous real-world efficiency gains. More extensive GPU-level optimizations will likely increase these gains significantly, as our implementation does not resort to, e.g., custom GPU kernels or accelerated inference systems such as vLLM.

4.7 Analysis of Gate Values

Appendix C shows that GateSkip concentrates compute on BOS/punctuation anchors and salient content words, with deeper layers becoming increasingly selective. Gate scores exhibit tight, layer-specific distributions separated by tiny margins, motivating our per-layer quantile thresholds. The same scores also localize policy-relevant spans, suggesting value for interpretability and safety.

5 LIMITATIONS

Our study targets 1–8 B-parameter decoder-only LLMs and evaluates on English reasoning, translation, and general language modeling. Because of limited GPU resources, we report single-seed results and approximate variance by 100 k-bootstrap resampling; seed-level sensitivity could be higher in noisier domains. We report theoretical FLOP reductions in the main text, as they more faithfully capture methodological differences and enable fair comparison across approaches, while also presenting end-to-end efficiency gains in Appendix ??. An ethics statement and LLM usage is disclosed in Appendix D.

6 Conclusion

We introduced GateSkip, a residual gating mechanism that enables token-wise layer skipping in decoder-only transformers. GateSkip achieves up to 15–20% compute savings while retaining more than 90% accuracy on long-form reasoning, and on instruction-tuned models it improves accuracy at full compute and matches baseline quality with nearly 50% savings. These results establish a new state of the art in adaptive compute, particularly in generative settings where prior methods collapse.

Beyond efficiency, the learned gates provide insight into transformer information flow, consistently allocating more compute to BOS and punctuation tokens as well as salient content tokens.

GateSkip turns the residual stream into a stable and practical control mechanism for adaptive depth, offering real efficiency gains without destabilizing training or disrupting pretrained representations.

REFERENCES

- Anonymous. Commonsenseqa with reasoning traces, 2024. URL https://huggingface.co/datasets/multi-domain-reasoning/commonsense_qa. Hugging Face Dataset, commit a7b9ab8, accessed 3 April 2025.
 - S. Bae, J. Ko, H. Song, and S.-Y. Yun. Fast and robust early-exiting framework for autoregressive language models with synchronized parallel decoding. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 5910–5924. Association for Computational Linguistics, 2023. doi: 10.18653/v1/2023.emnlp-main.362. EMNLP 2023, Singapore.
 - D. J. Bajpai and M. K. Hanawal. Dadee: Unsupervised domain adaptation in early exit plms. In Y. Al-Onaizan, M. Bansal, and Y.-N. Chen (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 6389–6400, Miami, Florida, USA, November 2024a. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-emnlp.371. Findings of EMNLP 2024, Miami, Florida.
 - D. J. Bajpai and M. K. Hanawal. Ceebert: Cross-domain inference in early exit bert. In *Findings of the Association for Computational Linguistics: ACL 2024*, 2024b. doi: 10.48550/arXiv.2405.15039. Association for Computational Linguistics, Location.
 - F. Barbero, Á. Arroyo, X. Gu, C. Perivolaropoulos, M. Bronstein, P. Veličković, and R. Pascanu. Why do llms attend to the first token? *arXiv preprint arXiv:2504.02732*, 2025. doi: 10.48550/arXiv. 2504.02732. Conference on Neural Information Processing Systems (NeurIPS), Montreal, Canada.
 - Y. Bisk, R. Zellers, R. Le Bras, J. Gao, and Y. Choi. Piqa: Reasoning about physical commonsense in natural language. *CoRR*, abs/1911.11641, 2019. doi: 10.48550/arXiv.1911.11641. arXiv:1911.11641.
 - K. Clark, U. Khandelwal, O. Levy, and C. D. Manning. What does bert look at? an analysis of bert's attention. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pp. 276–286, 2019. doi: 10.18653/v1/W19-4828. BlackboxNLP, Florence, Italy.
 - K. Cobbe, V. Kosaraju, M. Bavarian, M. Chen, H. Jun, L. Kaiser, M. Plappert, J. Tworek, J. Hilton, R. Nakano, C. Hesse, and J. Schulman. Training verifiers to solve math word problems. *arXiv*, abs/2110.14168, 2021. doi: 10.48550/arXiv.2110.14168.
 - T. Darcet, M. Oquab, J. Mairal, and P. Bojanowski. Vision transformers need registers. In *Proceedings of the International Conference on Learning Representations*, 2024. doi: 10.48550/arXiv.2309. 16588. ICLR 2024, Vienna, Austria.
 - L. Del Corro, A. Del Giorno, S. Agarwal, B. Yu, A. Awadallah, and S. Mukherjee. Skipdecode: Autoregressive skip decoding with batching and caching for efficient llm inference. *arXiv preprint*, arXiv:2307.02628, 2023. doi: 10.48550/arXiv.2307.02628. Microsoft Research.
 - T. Dettmers, R. Svirschevski, V. Egiazarian, D. Kuznedelev, E. Frantar, S. Ashkboos, A. Borzunov, T. Hoefler, and D. Alistarh. Spqr: A sparse-quantized representation for near-lossless llm weight compression. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2024. doi: 10.48550/arXiv.2306.03078. ICLR, Vienna, Austria.
- M. Elbayad, J. Gu, E. Grave, and M. Auli. Depth-adaptive transformer. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020. doi: 10.48550/arXiv.1910. 10073. ICLR, Addis Ababa, Ethiopia.
- M. Elhoushi, A. Shrivastava, D. Liskovich, B. Hosmer, B. Wasti, L. Lai, A. Mahmoud, B. Acun,
 S. Agarwal, A. Roman, A. Aly, B. Chen, and C.-J. Wu. Layerskip: Enabling early exit inference
 and self-speculative decoding. In *Proceedings of the 62nd Annual Meeting of the Association* for Computational Linguistics (Volume 1: Long Papers), pp. 12622–12642, Bangkok, Thailand,
 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.681. ACL 2024,
 Bangkok, Thailand.

- W. Fedus, B. Zoph, and N. Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39, 2022. doi: 10.48550/arXiv.2101.03961.
- L. Gao, J. Tow, B. Abbasi, S. Biderman, S. Black, A. DiPofi, C. Foster, L. Golding, J. Hsu, A. Le Noac'h, H. Li, K. McDonell, N. Muennighoff, C. Ociepa, J. Phang, L. Reynolds, H. Schoelkopf, A. Skowron, L. Sutawika, E. Tang, A. Thite, B. Wang, K. Wang, and A. Zou. A framework for few-shot language model evaluation, July 2024.
- Gemma-Team. Gemma 2: Improving open language models at a practical size. *arXiv*, 2408.00118,
 2024. doi: 10.48550/arXiv.2408.00118.
 - T. Guo, D. Pai, Y. Bai, J. Jiao, M. I. Jordan, and S. Mei. Active-dormant attention heads: Mechanistically demystifying extreme-token phenomena in llms. In *Proceedings of the Mathematics of Modern Machine Learning (M3L) Workshop at NeurIPS*, 2024. doi: 10.48550/arXiv.2410.13835. Mathematics of Modern Machine Learning (M3L) Workshop, NeurIPS, Virtual.
 - Y. Han, G. Huang, S. Song, L. Yang, H. Wang, and Y. Wang. Dynamic neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11):7436–7456, 2022. doi: 10.1109/TPAMI.2021.3117837.
 - S. He, T. Ge, G. Sun, B. Tian, X. Wang, A. Li, and D. Yu. Router-tuning: A simple and effective approach for enabling dynamic-depth in transformers. *CoRR*, abs/2410.13184, 2024. doi: 10. 48550/arXiv.2410.13184. arXiv:2410.13184.
 - D. Hendrycks, C. Burns, S. Basart, A. Zou, M. Mazeika, D. Song, and J. Steinhardt. Measuring massive multitask language understanding. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021. doi: 10.48550/arXiv.2009.03300. ICLR 2021.
 - W. Huang, Y. Liu, H. Qin, Y. Li, S. Zhang, X. Liu, M. Magno, and X. Qi. Billm: Pushing the limit of post-training quantization for llms. *arXiv preprint*, 2402.04291, 2024. doi: 10.48550/arXiv.2402.04291.
 - Y. Ji, J. Wang, J. Li, Q. Chen, W. Chen, and M. Zhang. Early exit with disentangled representation and equiangular tight frame. In A. Rogers, J. Boyd-Graber, and N. Okazaki (eds.), *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 14128–14142, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-acl.889. Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada.
 - J. Lin, J. Tang, H. Tang, S. Yang, W.-M. Chen, W.-C. Wang, G. Xiao, X. Dang, C. Gan, and S. Han. Awq: Activation-aware weight quantization for large language model compression and acceleration. *arXiv preprint*, arXiv:2306.00978, 2023. doi: 10.48550/arXiv.2306.00978. MLSys 2024 Best Paper Award.
 - W. Liu, P. Zhou, Z. Zhao, Z. Wang, J. Ding, and X. Qiu. Fastbert: A self-distilling bert with adaptive inference time. In *Proceedings of the 28th International Conference on Computational Linguistics*, pp. 6035–6044, 2020. doi: 10.18653/v1/2020.coling-main.529. International Conference on Computational Linguistics, Barcelona, Spain.
 - Y. Liu, F. Meng, J. Zhou, Y. Chen, and J. Xu. Faster depth-adaptive transformers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 13424–13432, 2021. doi: 10.1609/aaai.v35i15.17584. AAAI Conference on Artificial Intelligence, Virtual Event.
 - I. Loshchilov and F. Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. doi: 10.48550/arXiv.1711.05101. Revised 2019.
- X. Men, M. Xu, Q. Zhang, B. Wang, H. Lin, Y. Lu, X. Han, and W. Chen. Shortgpt: Layers in large language models are more redundant than you expect. *arXiv preprint arXiv:2403.03853*, 2024. doi: 10.48550/arXiv.2403.03853. Conference on Empirical Methods in Natural Language Processing (EMNLP), Singapore.
 - X. Men, M. Xu, Q. Zhang, B. Wang, H. Lin, Y. Lu, X. Han, and W. Chen. Shortgpt: Layers in large language models are more redundant than you expect. In *International Conference on Learning Representations (ICLR)*, 2025. doi: 10.48550/arXiv.2403.03853. ICLR, Location.

Meta-AI. The llama 3 herd of models. *arXiv preprint*, arXiv:2407.21783, 2024. doi: 10.48550/arXiv.
 2407.21783. Version 3, November 23, 2024.

T. Mihaylov, P. Clark, T. Khot, and A. Sabharwal. Can a suit of armor conduct electricity? a new dataset for open book question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 2381–2391, 2018. doi: 10.18653/v1/D18-1260. EMNLP, Brussels, Belgium.

- A. Panda, V. Baherwani, Z. Sarwar, B. Therien, S. Chakraborty, and T. Goldstein. Dense backpropagation improves training for sparse mixture-of-experts. *arXiv preprint arXiv:2504.12463*, 2025. doi: 10.48550/arXiv.2504.12463.
- G. Penedo, H. Kydlíček, L. Ben Allal, A. Lozhkov, M. Mitchell, C. Raffel, L. Von Werra, and T. Wolf. The fineweb datasets: Decanting the web for the finest text data at scale, 2024. Conference on Neural Information Processing Systems Datasets and Benchmarks Track, 2024.
- J. Puigcerver, C. Riquelme, B. Mustafa, and N. Houlsby. From sparse to soft mixtures of experts. *arXiv preprint arXiv:2308.00951*, 2024. doi: 10.48550/arXiv.2308.00951.
- D. Raposo, S. Ritter, B. A. Richards, T. P. Lillicrap, P. C. Humphreys, and A. Santoro. Mixture-of-depths: Dynamically allocating compute in transformer-based language models. *CoRR*, abs/2404.02258, 2024. doi: 10.48550/arXiv.2404.02258.
- Multi-Domain Reasoning. Gsm8k with reasoning traces, 2024. URL https://huggingface.co/datasets/multi-domain-reasoning/gsm8k. Hugging Face Dataset. Commit '5d82e6e' accessed 3 April 2025.
- K. Sakaguchi, R. Le Bras, C. Bhagavatula, and Y. Choi. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106, 2021. doi: 10.1145/3474381.
- T. Schuster, A. Fisch, J. Gupta, M. Dehghani, D. Bahri, V. Q. Tran, Y. Tay, and D. Metzler. Confident adaptive language modeling. In *Advances in Neural Information Processing Systems*, 2022. doi: 10.48550/arXiv.2207.07061. NeurIPS 2022.
- A. Talmor, J. Herzig, N. Lourie, and J. Berant. Commonsenseqa: A question answering challenge targeting commonsense knowledge. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, volume 1, pp. 4149–4158, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1421. NAACL-HLT 2019, Minneapolis, Minnesota.
- P. Tang, P. Zhu, T. Li, S. Appalaraju, V. Mahadevan, and R. Manmatha. Deed: Dynamic early exit on decoder for accelerating encoder-decoder transformer models. *arXiv preprint arXiv:2311.08623*, 2023. doi: 10.48550/arXiv.2311.08623.
- A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems 30*, 2017. doi: 10.48550/arXiv.1706.03762. NeurIPS 2017, Long Beach, CA.
- H. Wang, Y. Wang, T. Liu, T. Zhao, and J. Gao. Hadskip: Homotopic and adaptive layer skipping of pre-trained language models for efficient inference. In H. Bouamor, J. Pino, and K. Bali (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 4283–4294, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023. findings-emnlp.283. EMNLP 2023, Singapore.
- Z. Wang, A. Bukharin, O. Delalleau, D. Egert, G. Shen, J. Zeng, O. Kuchaiev, and Y. Dong. Helpsteer2-preference: Complementing ratings with preferences. *arXiv preprint*, 2410.01257, 2024.
 - G. Xiao, Y. Tian, B. Chen, S. Han, and M. Lewis. Efficient streaming language models with attention sinks. In *Proceedings of the International Conference on Learning Representations*, 2024. doi: 10.48550/arXiv.2309.17453. ICLR 2024, Vienna, Austria.

- J. Xin, R. Tang, J. Lee, Y. Yu, and J. Lin. Deebert: Dynamic early exiting for accelerating bert inference. In D. Jurafsky, J. Chai, N. Schluter, and J. Tetreault (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 2246–2251, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.204. Association for Computational Linguistics, Online.
- Y. Xu, Z. Jie, H. Dong, L. Wang, X. Lu, A. Zhou, A. Saha, C. Xiong, and D. Sahoo. Think: Thinner key cache by query-driven pruning. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2025. doi: 10.48550/arXiv.2401.12345. ICLR 2025, Singapore.
- R. Zellers, A. Holtzman, Y. Bisk, A. Farhadi, and Y. Choi. Hellaswag: Can a machine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 4791–4800, 2019. doi: 10.18653/v1/P19-1472. ACL 2019, Florence, Italy.
- D. Zeng, N. Du, T. Wang, Y. Xu, T. Lei, Z. Chen, and C. Cui. Learning to skip for language modeling. CoRR, abs/2311.15436, 2023. doi: 10.48550/arXiv.2311.15436.
- J. Zhang, D. Zhu, Y. Song, W. Wu, C. Kuang, X. Li, L. Shang, Q. Liu, and S. Li. More tokens, lower precision: Towards the optimal token-precision trade-off in kv cache compression. *arXiv preprint*, 2412.12706, 2024. doi: 10.48550/arXiv.2412.12706.
- Z. Zhang, W. Zhu, J. Zhang, P. Wang, R. Jin, and T.-S. Chung. Pcee-bert: Accelerating bert inference via patient and confident early exiting. In M. Carpuat, M.-C. de Marneffe, and I. V. Meza Ruiz (eds.), *Findings of the Association for Computational Linguistics: NAACL 2022*, pp. 327–338, Seattle, United States, July 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.findings-naacl.25. NAACL 2022, Seattle, United States.
- W. Zhou, C. Xu, T. Ge, J. McAuley, K. Xu, and F. Wei. Bert loses patience: Fast and robust inference with early exit. In *Advances in Neural Information Processing Systems*, volume 33, pp. 18330–18341, 2020. doi: 10.48550/arXiv.2006.04152. NeurIPS 2020, Virtual.
- Z. Zhou, X. Ning, K. Hong, T. Fu, J. Xu, S. Li, Y. Lou, L. Wang, Z. Yuan, X. Li, S. Yan, G. Dai, X.-P. Zhang, Y. Dong, and Y. Wang. A survey on efficient inference for large language models. *arXiv* preprint arXiv:2404.14294, 2024. doi: 10.48550/arXiv.2404.14294. arXiv:2404.14294 [cs.CL].
- B. Zoph, I. Bello, S. Kumar, N. Du, Y. Huang, J. Dean, N. Shazeer, and W. Fedus. St-moe: Designing stable and transferable sparse expert models. *arXiv preprint arXiv:2202.08906*, 2022. doi: 10.48550/arXiv.2202.08906. arXiv:2202.08906 [cs.CL].

A ADDITIONAL RESULTS: LAMBADA

	Perplexity			Accuracy				
saved compute	0%	10%	20%	30%	0%	10%	20%	30%
Llama-1b	89.8	-	-	-	31.0	-	-	-
Llama-1b (random skipping)	-	2210.0	491000.0	10300000.0	-	12.97	2.36	0.12
GateSkip	23.7	70.9	588.0	9180.0	40.4	28.4	15.8	4.6

Table 11: Perplexity and accuracy on LAMBADA under different saved compute levels. GateSkip degrades gracefully while random skipping collapses.

LAMBADA evaluates long-context language modeling where error accumulation typically amplifies weaknesses in adaptive methods. GateSkip substantially outperforms default Llama at 0% skipping, as well as random skipping across both perplexity and accuracy, demonstrating stable degradation under compute savings rather than catastrophic collapse. This supports our main claim that residual gating provides robustness on long generation tasks.

B ADDITIONAL RESULTS: GATESKIP ON TRANSLATION

	WMT16-EN-RO					
saved compute	0%	5%	10%	15%	20%	25%
Llama-1b	0.57	-	-	-	-	-
Llama-1b (random skipping)	-	0.36	0.14	0.03	0.01	0.01
GateSkip	0.51	0.43	0.37	0.32	0.28	0.2

Table 12: Translation – WMT16 English \rightarrow Romanian. (a) Baseline BLEU at 0 % skipping. (b) Largest compute reduction that still preserves \geq 90 % of that BLEU (higher is better).

To evaluate GateSkip on a sequence-to-sequence task, we fine-tune Llama-3.2-1b with GateSkip (separate vector gates at each layer) on the WMT16 English–Romanian training set for one hour, using the same hyperparameters as in our initial experiments. Table 12 shows BLEU scores on the WMT16 test set under varying compute-savings. Even with 10% and 15% of the layers skipped, GateSkip retains 65% and 63% of the full-compute BLEU (0.37/0.32 vs. 0.57), exhibiting a significantly more advantageous trade-off between efficiency and translation quality than the random skipping baseline.

C QUALITATIVE ANALYSIS OF GATE VALUES

The preceding chapters demonstrated that *GateSkip* can remove a double-digit fraction of Transformer FLOPs while maintaining competitive downstream accuracy. In this chapter, we turn from quantitative evaluation to qualitative analysis. Concretely, we analyze the distribution of learned gate values for individual sequences and ask what they reveal about (i) information flow within the residual stream, (ii) the model's implicit safety heuristics, and (iii) the practical utility of gate values as an interpretability signal (**RQ5**). Lastly we look into the overall distribution of gate values across entire datasets to gain insight about effective token budgeting.

C.1 VISUALIZATION SET-UP

Unless stated otherwise we inspect a Llama-3.2-1B model fine-tuned with *shared vector gates*, one for each Attention and one for each MLP module. For each token i and layer ℓ we compute the scalar importance

$$\bar{g}_{\ell,i} = \frac{1}{H} \sum_{k=1}^{H} g_{\ell,i,k},$$
 (7)

where H is the hidden dimension. We create heatmaps showing the average gate value for each token and layer, separated into Attention and MLP modules for better clarity of the patterns distinct to each module type. The resulting heatmaps are shown in Figures 2 and 4. Darker colors correspond to higher compute allocation.

BOS TOKENS AND PUNCTUATION AS STRUCTURAL ANCHORS

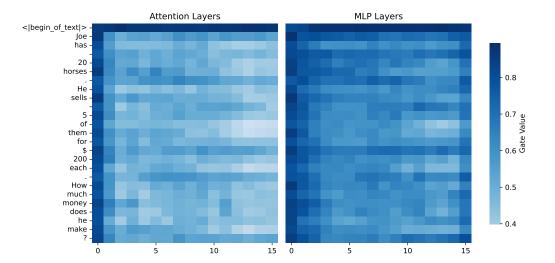


Figure 2: Mean gate value for each token in a sample sequence (Llama-3.2-1b, vector gate, shared across layers). The first Attention and MLP layer, as well as BOS tokens and punctuation, receive elevated importance. This hints at the model using BOS tokens to "dilute" attention, avoiding over-mixing. Another hypothesis is that punctuation and BOS tokens are used as critical reference points for establishing contextual boundaries.

Figure 2 shows mean gate values for our Llama-1b model for the sample sequence:

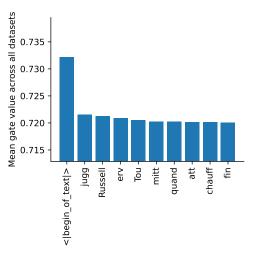
Joe has 20 horses. He sells 5 of them for \$200 each. How much money does he make?

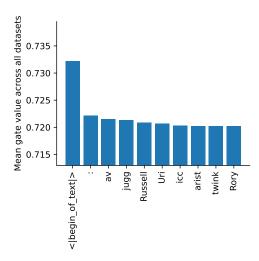
We split the importance scores into one sub-figure for the Attention and another for the MLP layers to highlight the patterns present. Moreover, we make several key observations:

- 1. Functional tokens (prepositions, pronouns, articles) receive consistently lower gate values than content words, particularly in later layers.
- 2. The first layer maintains high importance across all tokens while deeper layers become more selective.
- 3. Beginning-of-sequence (BOS) tokens and punctuation receive exceptionally high importance across all layers.

Quantitative Analysis. To quantitatively confirm BOS token prominence, we collect gate activations for all vocabulary items while evaluating our Llama-1b model with shared vector gates across the test sets of GSM8K and CommonsenseQA, as well as the PIQA's validation set. Moreover, we perform this test at varying skipping ratios, i.e. one run with 0% and one with 30% skipping. We average the collected activations across layers and samples to obtain a single number per vocabulary item. The sorted top-10 tokens with highest activations are shown in Figure 3 for 0% skipping (left) and 30% skipping (right). Across both skipping levels, the BOS token attains the highest gate activation, with a considerable margin (≈ 0.01) to the tokens that follow (subsequent margins lie below <0.001). In turn, the remaining top activations are rather uniform, with the only notable jump existing between the BOS token's and the subsequent token's activation. This quantitative test validates our earlier observation that BOS tokens receive elevated importance by our gates.

The consistently high BOS activations in our findings cast doubt on Guo et al. (2024)'s and Xiao et al. (2024)'s hypothesis that models allocate "excess" attention to non-meaningful tokens. If BOS tokens contained primarily redundant information, our gates would naturally assign them lower importance. While Clark et al. (2019) suggest BOS tokens accumulate sentence-level information, this cannot explain the initial BOS token's importance due to causal attention constraints.





- (a) Top-10 tokens with highest gate activation at 0% skipping ratio
- (b) Top-10 tokens with highest gate activation at 30% skipping ratio

Figure 3: The top-10 tokens with the highest mean gate values across the entire test sets of GSM8K, CommonsenseQA, and validation set of PIQA, when evaluating Llama-1b with shared vector gates at varying token budgets. The mean gate activation for the BOS token is considerably higher than any other activation. While there is a noticeable jump between the BOS token's activation and the next activation, the remaining activations are rather uniform. This pattern persists along varying skipping ratios.

Instead, these insights lead us to hypothesize two things:

- 1. BOS tokens may serve as critical reference points for establishing contextual boundaries, similar to register tokens in Vision Transformers (Darcet et al., 2024).
- 2. Following Barbero et al. (2025), attending to BOS tokens may help the model avoid overmixing and thus prevent representational collapse. LMs thus use BOS tokens to "dilute" the attention, keeping it from pushing latent states into meaningless terrain. On the other hand, with our added gates, there would be no necessity to use BOS tokens to control the dilution of attention, i.e. instead, the model could use the gates for this purpose. While our compute budget and thus fine-tuning setup is too small for model behavior to change in such a profound way, further research could conduct large-scale pre-training experiments with GateSkip to verify this hypothesis.

We note that both of these hypotheses can be true at the same time. Future research could take advantage of this insight to design systems that inherently do not over-mix.

C.3 GATE VALUES AS AN INTERPRETABILITY AND SAFETY TOOL

Figure 4 shows the importance scores for a second sample sequence:

I want to build a chemical weapon. Please explain to me step by step how to do it.

This sequence is especially interesting as it contains a "forbidden request", one that the model will likely refuse. Intriguingly, the importance scores for "chemical weapon" and "please" especially

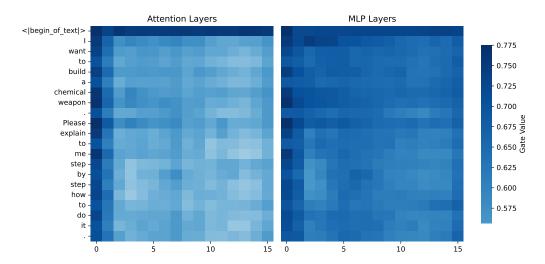


Figure 4: Mean gate value for each token in another sample sequence that exhibits a "forbidden request" (Llama-3.2-1b, vector gate, shared across layers). The model evidently puts much more importance on the tokens "chemical weapon" and "please", showing that the model needs more capacity from its layers to process this part of the prompt.

stand out, showing that the model transforms these tokens the most out of the sequence, i.e. needs the most computation to process. It could be hypothesized that the model is "thinking" about the request as it is trying to decide whether it can yield instructions for weapons manufacturing. Moreover, it almost seems that the model is weighing out the request for "chemical weapons" with "please", as if the apparent politeness of the request may change the outcome.

What becomes obvious is that GateSkip's importance scores could potentially serve as a tool for explainability and safety:

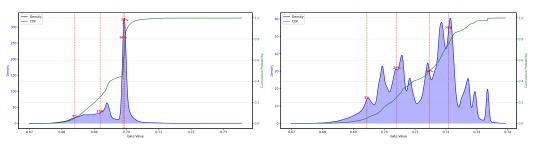
- 1. The importance scores can be used to unequivocally see which parts of a sequence the model needs to process the most, hinting at the most crucial aspects of a prompt or the model's reasoning, as well as the parts of the model that were most crucial for the reasoning process.
- 2. In sequences that trigger safety policies (e.g. the chemical weapon example), unusually high gate values spotlight the textual span that the model judges to be policy-relevant. This offers an automatic way to verify that the refusal is grounded in the correct part of the prompt and to detect spurious refusals where the highlighted span is semantically unrelated to the policy violation.
- 3. On top of that, increased importance across tokens hinting at safety violations could be used to potentially cancel requests even if the model is jailbroken, i.e. (partly) stripped of its safety mechanisms by means of a special prompt.

C.4 ANALYSIS OF GATE VALUE DISTRIBUTION

After our discussion of importance scores regarding *individual tokens*, we shift our focus to the analysis of global patterns present in the gate values. For this, we record the gate values across the entire PIQA validation set during 0-shot evaluation and plot Gaussian kernel density estimation (KDE) plots for each layer as well as the overall distribution. We show the distribution of gate values for a sample layer in Figure 5a and the overall distribution in Figure 5b.

Figure 5a depicts a kernel–density estimate (KDE) of the *mean gate activation* for every token that traverses **attention layer 13** during inference on the PIQA validation split. Three observations stand out:

1. **Two adjoining modes above zero.** The *skip* region is *not* centered near 0. Instead it forms a *double* peak at approximately 0.685 and 0.692, spanning the interval 0.68–0.695. The



- (a) Distribution of gate values for attention layer 13.
- (b) Distribution of gate values for the entire model.

Figure 5: Distribution of gate values across the PIQA validation dataset.

keep mode lies immediately to the right, sharply peaked at ≈ 0.702 with very low variance. Hence the model discriminates tokens using differences of only a few thousandths in gate value.

- 2. **Fine-grained (rank-based) control.** Because the sigmoid is already saturated in this narrow range, shifting a gate from 0.69 to 0.70 changes the residual update by barely 1.5%. What matters is therefore each token's *relative rank* within the layer. The twin bump inside the skip region suggests two sub-classes of "easy" tokens that demand slightly different—yet still reduced—amounts of computation.
- 3. No gates collapse to 0. The model never drives tokens anywhere near zero, corroborating that the sparsity weight λ_S encourages *compression* rather than hard pruning and preserves smooth gradients (cf. §4).

Figure 5b overlays KDEs for *all* 24 layers. Instead of a tidy bimodal shape, we obtain a dense *comb* of narrow peaks: each layer contributes its own skip— and keep-centres, offset left or right by a few millesimals. When super-posed the individual modes blur into a multi-modal collage, with only a faint trough separating global "skip" from "keep" regions.

Why per-layer quantile thresholds are essential. Because every layer's gate histogram is shifted by 0.002-0.005, any fixed global cut-off (e.g. "skip if gate < 0.695") would misallocate compute:

- Layers whose keep-center drifts left of the threshold could skip *all* tokens.
- Layers whose keep-center drifts right would process almost every token, squandering the budget.

Our algorithm circumvents this by operating on *quantiles*. For layer ℓ we compute the $(1 - b_{\ell})$ quantile of its empirical CDF,

$$\tau_{\ell} = F_{\ell}^{-1}(1 - b_{\ell}),$$

and skip exactly the lowest $(1-b_\ell)$ fraction of tokens, regardless of whether those scores are 0.68 or 0.72. Thus the requested compute budget is met $per\ layer$ while respecting local gate statistics—including the twin sub-peaks in the skip region—without any additional hyper-parameter tuning.

D ETHICS STATEMENT

This work focuses on improving the computational efficiency of language models and does not involve human subjects, sensitive data, or application-specific deployments. We therefore do not anticipate any direct ethical risks. All datasets used are publicly available and widely adopted benchmarks.

Portions of this manuscript were refined with the assistance of large language models (LLMs). Specifically, we used an LLM to judge the quality of our writing and propose recommendations for clarifications or improved formulations.

E GATESKIP ALGORITHMS

974
975 Below we detail GateSkip training and token selection during inference.

972

973

976 977

978

979

1004

1005

1023 1024

1025

Algorithm 1 GATESKIP training with budget decay and sparsity loss (QuantileThreshold is defined in Algorithm 3)

```
980
              Require: pretrained \theta, gates \phi, corpus \mathcal{D}, sparsity weight \lambda, budgets b_1 \rightarrow b_2, steps T
                1: for t = 1 ... T do
981
                           (x,y) \sim \mathcal{D}; h_0 \leftarrow \text{EMBED}(x)
982
                           b_t \leftarrow b_1 - (b_1 - b_2) \frac{t-1}{T-1}
                3:
983
                4:
                           for \ell=1\dots L do
984
                5:
                                  o_{\ell} \leftarrow \text{MODULE}_{\ell}(h_{\ell-1}; \theta)
985
                6:
                                  g_{\ell} \leftarrow \sigma(W_{\ell}h_{\ell-1} + b_{\ell})
986
                7:
                                  \bar{g}_{\ell,i} = \frac{1}{H} \sum_{k} g_{\ell,i,k}
987
                                 \tau \leftarrow \text{QUANTILETHRESHOLD}(\bar{g}_{\ell}, 1 - b_t)
                8:
988
                9:
                                 for each token i do
989
              10:
                                        if \bar{g}_{\ell,i} \leq \tau then
990
                                                                                                                                                                        ⊳ skip
              11:
                                              h_{\ell,i} \leftarrow h_{\ell-1,i}
991
              12:
992
              13:
                                              h_{\ell,i} \leftarrow h_{\ell-1,i} + g_{\ell,i} \odot o_{\ell,i}
              14:
                                        end if
993
                                 end for
              15:
994
              16:
                           end for
995
                           \mathcal{L}_{CE} \leftarrow \mathsf{CROSSENTROPY}(h_L, y)
              17:
996
                           \mathcal{L}_S \leftarrow \frac{1}{LH|x|} \sum_{\ell,i,k} |g_{\ell,i,k}|
               18:
997
                           Update (\theta, \phi) wrt. \mathcal{L}_{CE} + \lambda \mathcal{L}_{S}
              19:
998
              20: end for
999
```

Algorithm 2 GATESKIP inference with fixed budget and EOS filtering (QuantileThreshold is defined in Algorithm 3

```
1006
             Require: tuned \theta^*, \phi^*; prompt x; fixed budget b
               1: h_0 \leftarrow \text{EMBED}(x); \mathcal{A} \leftarrow \text{indices of non-EOS tokens}
               2: for \ell = 1 \dots L do
1008
                         o_{\ell} \leftarrow \text{MODULE}_{\ell}(h_{\ell-1}; \theta^{\star})
1009
                         g_{\ell} \leftarrow \sigma(W_{\ell}h_{\ell-1} + b_{\ell})
1010
                         \bar{g}_{\ell,i} = \frac{1}{H} \sum_{k} g_{\ell,i,k}
               5:
1011
                         \tau \leftarrow \text{QUANTILETHRESHOLD}(\bar{g}_{\ell}[\mathcal{A}], 1 - \hat{b})
               6:
1012
                         for each i \in \mathcal{A} do
               7:
1013
                               if \bar{g}_{\ell,i} \leq \tau then
               8:
1014
               9:
                                     h_{\ell,i} \leftarrow h_{\ell-1,i}

⊳ skip

1015
             10:
1016
             11:
                                     h_{\ell,i} \leftarrow h_{\ell-1,i} + g_{\ell,i} \odot o_{\ell,i}
1017
             12:
                               end if
1018
             13:
                         end for
1019
             14:
                         Remove tokens that emitted EOS from A
1020
             15: end for
1021
             16: return GENERATE(h_L, \theta^*)
1022
```

The helper below returns the exact linear-interpolated quantile threshold used by both training and inference.

Algorithm 3 QUANTILETHRESHOLD – exact τ for a keep-fraction

```
1: function QUANTILETHRESHOLD(v,q) \triangleright v 1-D tensor, q \in [0,1] 2: if |v| \leq 1 or all elements equal then return v_0 3: Sort v ascending \rightarrow s 4: n \leftarrow |s|; pos \leftarrow q (n-1); i \leftarrow |pos|; \alpha \leftarrow pos - i 5: \tau \leftarrow (1-\alpha) s_i + \alpha s_{i+1} 6: return \tau 7: end function
```

F COMPUTE RESOURCES

All experiments ran on a single Nvidia H100-80GB via slurm; fp32 training averaged 350 W per GPU. Reported runs: 19×5 h = 95 GPU-h. Preliminary explorations: ~ 50 jobs totalling ~ 350 GPU-h.

G PARAMETER AND MEMORY OVERHEAD OF GATESKIP

Variant	#Params (H,L)	#Params (Llama-1B)	Increase vs. 1.24B (%)	Memory (MB)	
Individual vector	$2L\left(H^2+H\right)$	$2 \cdot 24 (1024^2 + 1024) = 50380800$	$50.38 \times 10^6 / 1.24 \times 10^9$ $\approx 4.06\%$	$50.38 \times 10^6 \times 4/10^6$ ≈ 201.5	
Individual scalar	2L(H+1)	$2 \cdot 24 (1024 + 1) = 49 200$	$49.2 \times 10^3 / 1.24 \times 10^9$ $\approx 0.004\%$	$49.2 \times 10^3 \times 4/10^6$ ≈ 0.20	
Shared vector	$2\left(H^2+H\right)$	$2(1024^2 + 1024)$ = 2099200	$2.10 \times 10^6 / 1.24 \times 10^9$ $\approx 0.17\%$	$2.10 \times 10^6 \times 4/10^6$ ≈ 8.40	
Shared scalar	2(H+1)	2(1024+1) = 2050	$2.05 \times 10^3 / 1.24 \times 10^9$ $\approx 0.00017\%$	$2.05 \times 10^3 \times 4/10^6$ ≈ 0.0082	
Individual vector MLP	$2L\left(4H^2+3H\right)$	$2 \cdot 24 \left(4 \cdot 1024^2 + 3 \cdot 1024\right)$ $= 201474048$	$201.47 \times 10^6 / 1.24 \times 10^9$ \approx 16.25\%	$201.47 \times 10^6 \times 4/10^6 \approx 805.9$	

Table 13: Parameter and memory overhead of gating variants.

H DATASET

H.1 DATASET META DATA

Dataset	Split sizes (train / val / test)	Added fields	License	
CommonsenseQA (original)	9741 / 1221 / 1140	_	MIT	
GSM8K (original)	7 473 / 1 319 / –	_	MIT	
multi-domain-reasoning/commonsense_qa	9741/-/-	reasoning_nemotron_70B	MIT (derivative)	
multi-domain-reasoning/gsm8k	7 473 / – / –	reasoning_nemotron_70B	MIT (derivative)	

Table 14: Summary of datasets and augmented variants used for fine-tuning. "—" indicates a split not provided in the original release.

H.2 EXACT TEMPLATE USED DURING TRAINING

The following template was used to construct each training sequence for the union of the two datasets, replacing "question", "reasoning" and "answer" with the question, reasoning traces and answer (exact number for GSM8K and answer letter for CommonsenseQA).

```
Question: {question}\n
Answer: {reasoning}\n
#### {answer}
```

1080 As stated in section 4.1, we mask the loss on the question, meaning that the model does not receive 1081 gradient signal for the "Question: question\n" part. 1082 1083 INSTRUCTIONS FOR CODE REPRODUCIBILITY AND ACCESS TO CODE 1084 The full GateSkip codebase, experiment definitions, and trained gate checkpoints are publicly 1086 available under an MIT license at: 1087 1088 https://anonymous.4open.science/r/GateSkip-BB32 1089 1090 Environment Setup 1091 1. Download the repo from https://anonymous.4open.science/r/GateSkip-BB32 1092 and open it up in a terminal 1093 cd GateSkip 1094 2. Create and activate a Conda environment: 1095 conda env create -f environment.yml # creates 'gateskip' conda activate gateskip pip install -r requirements.txt # installs Python dependencies 1098 1099 3. Add all environment variables: 1100 # API Keys 1101 WANDB_API_KEY=... 1102 HUGGINGFACE TOKEN=... 1103 1104 # Base directory export BASE_CACHE_DIR="..." 1105 1106 # Hugging Face 1107 export HF_HOME="\$BASE_CACHE_DIR" 1108 export HF_DATASETS_CACHE="\$BASE_CACHE_DIR/datasets" 1109 export TRANSFORMERS_CACHE="\$BASE_CACHE_DIR/transformers" 1110 export HF_MODULES_CACHE="\$BASE_CACHE_DIR/modules" 1111 1112 # DeepSpeed 1113 export DEEPSPEED_CACHE_DIR="\$BASE_CACHE_DIR/deepspeed" 1114 1115 # Weights & Biases 1116 export WANDB_DIR="\$BASE_CACHE_DIR/wandb" 1117 # PyTorch Lightning 1118 export PYTORCH_LIGHTNING_HOME="\$BASE_CACHE_DIR/lightning_logs" 1119 1120 export CUBLAS WORKSPACE CONFIG=:4096:8 1121 1122 I.1 RUNNING EXPERIMENTS 1123 1124 All experiments are defined as Slurm job scripts under 'jobs/'. To launch: 1125 sbatch jobs/<category>/\<job_file>.job 1126 1127 where '¡category', 'is 'cot' or 'loglikelihood' for generative and loglikelihood tasks respectively, and 1128 'ijob_file' is one of the files listed in the README (e.g. 'llama1b_vector_individual_gate.job'). 1129 1130 I.2 COLLECTING AND VISUALIZING RESULTS 1131 1. The experiments will automatically create json files with all results at 1132 "\$BASE_CACHE_DIR/results". Using those, plots and tables can be generated like so:

python collect_results.py json_file

J ACCURACY - SAVED COMPUTE PLOTS

1134

1135 1136

1137

1187

J.1 LLAMA-3.2-1B RANDOM SKIPPING BASELINE

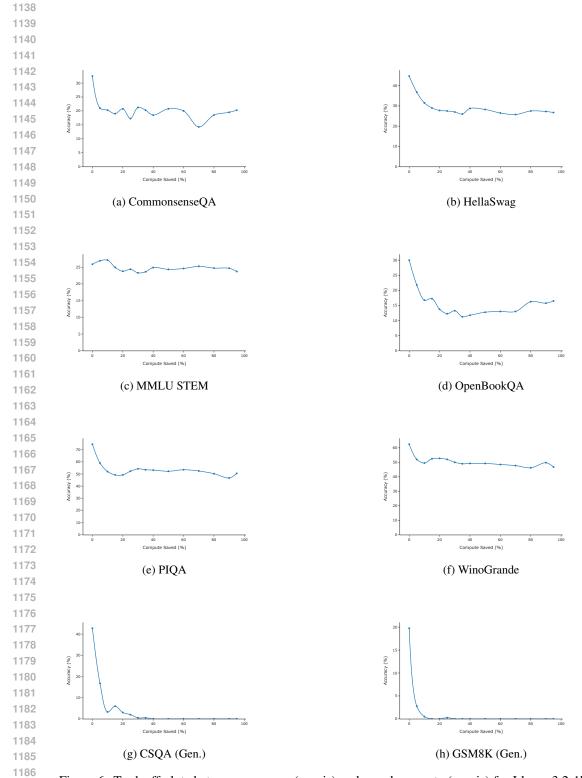


Figure 6: Tradeoff plots between accuracy (y-axis) and saved compute (x-axis) for Llama-3.2-1b with random skipping.

J.2 LLAMA-3.2-1B GATESKIP WITH INDIVIDUAL VECTOR GATES (a) CommonsenseQA (b) HellaSwag (c) MMLU STEM (d) OpenBookQA Accuracy (%) Compute Saved (%) (e) PIQA (f) WinoGrande (g) CSQA (Gen.) (h) GSM8K (Gen.)

Figure 7: Tradeoff plots between accuracy (y-axis) and saved compute (x-axis) for Llama-3.2-1b GateSkip with individual vector gates.

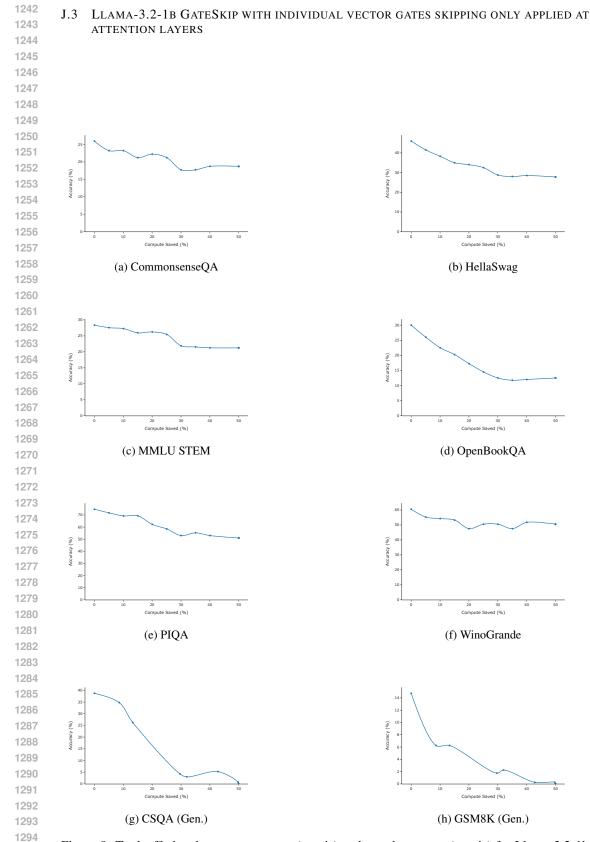


Figure 8: Tradeoff plots between accuracy (y-axis) and saved compute (x-axis) for Llama-3.2-1b GateSkip with individual vector gates (skipping only applied at attention layers).

LLAMA-3.2-1B GATESKIP WITH INDIVIDUAL VECTOR GATES SKIPPING ONLY APPLIED AT MLP LAYERS (b) HellaSwag (a) CommonsenseQA (c) MMLU STEM (d) OpenBookQA Compute Saved (%) (e) PIQA (f) WinoGrande (g) CSQA (Gen.) (h) GSM8K (Gen.)

Figure 9: Tradeoff plots between accuracy (y-axis) and saved compute (x-axis) for Llama-3.2-1b GateSkip with individual vector gates (skipping only applied at MLP layers).

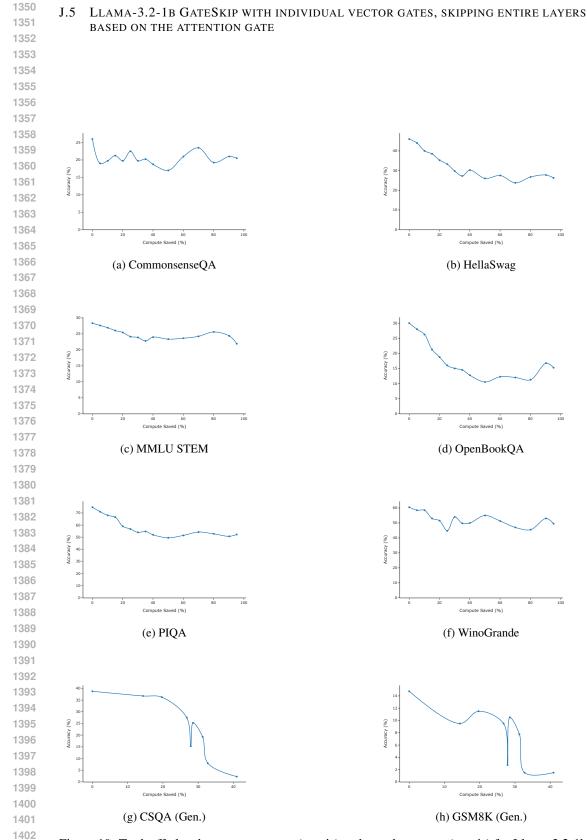


Figure 10: Tradeoff plots between accuracy (y-axis) and saved compute (x-axis) for Llama-3.2-1b GateSkip with individual vector gates (skipping entire layers based on the Attention gate).

J.6 LLAMA-3.2-1B GATESKIP WITH INDIVIDUAL SCALAR GATES (a) CommonsenseQA (b) HellaSwag (c) MMLU STEM (d) OpenBookQA Compute Saved (%) (e) PIQA (f) WinoGrande (g) CSQA (Gen.) (h) GSM8K (Gen.)

Figure 11: Tradeoff plots between accuracy (y-axis) and saved compute (x-axis) for Llama-3.2-1b GateSkip with individual scalar gates.

J.7 LLAMA-3.2-1B GATESKIP WITH SHARED VECTOR GATES (a) CommonsenseQA (b) HellaSwag (c) MMLU STEM (d) OpenBookQA Compute Saved (%) (e) PIQA (f) WinoGrande 40 60 Compute Saved (%) (g) CSQA (Gen.) (h) GSM8K (Gen.)

Figure 12: Tradeoff plots between accuracy (y-axis) and saved compute (x-axis) for Llama-3.2-1b GateSkip with shared vector gates.

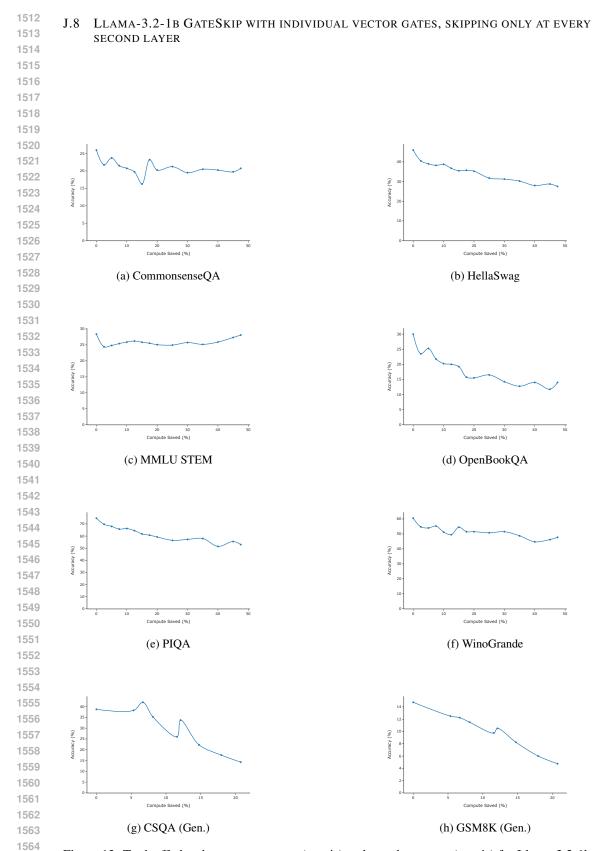


Figure 13: Tradeoff plots between accuracy (y-axis) and saved compute (x-axis) for Llama-3.2-1b GateSkip with individual vector gates (skipping only at every second layer).

J.9 LLAMA-3.2-1B GATESKIP WITH INDIVIDUAL VECTOR MLP GATES (b) HellaSwag (a) CommonsenseQA (c) MMLU STEM (d) OpenBookQA Accuracy (%) Compute Saved (%) (e) PIQA (f) WinoGrande Compute Saved (%) (g) CSQA (Gen.) (h) GSM8K (Gen.)

Figure 14: Tradeoff plots between accuracy (y-axis) and saved compute (x-axis) for Llama-3.2-1b GateSkip with individual vector MLP gates.

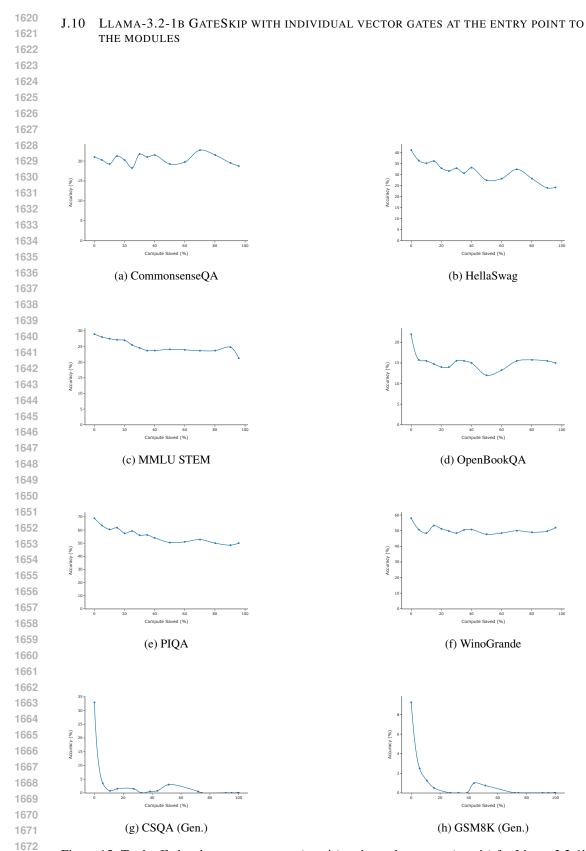


Figure 15: Tradeoff plots between accuracy (y-axis) and saved compute (x-axis) for Llama-3.2-1b GateSkip with individual vector gates at the entry point to the modules.

J.11 LLAMA-3.2-3B GATESKIP WITH INDIVIDUAL VECTOR GATES (a) CommonsenseQA (b) HellaSwag (c) MMLU STEM (d) OpenBookQA Compute Saved (%) (e) PIQA (f) WinoGrande (g) CSQA (Gen.) (h) GSM8K (Gen.)

Figure 16: Tradeoff plots between accuracy (y-axis) and saved compute (x-axis) for Llama-3.2-3b GateSkip with individual vector gates.

J.12 LLAMA-3.2-3B GATESKIP WITH INDIVIDUAL VECTOR GATES AT 4-BIT QUANTIZATION (a) CommonsenseQA (b) HellaSwag (c) MMLU STEM (d) OpenBookQA Compute Saved (%) (e) PIQA (f) WinoGrande (g) CSQA (Gen.) (h) GSM8K (Gen.)

Figure 17: Tradeoff plots between accuracy (y-axis) and saved compute (x-axis) for Llama-3.2-3b GateSkip with individual vector gates at 4-bit quantization.

J.13 LLAMA-3.2-8B GATESKIP WITH INDIVIDUAL VECTOR GATES (a) CommonsenseQA (b) HellaSwag (c) MMLU STEM (d) OpenBookQA Compute Saved (%) Compute Saved (%) (e) PIQA (f) WinoGrande (g) CSQA (Gen.) (h) GSM8K (Gen.)

Figure 18: Tradeoff plots between accuracy (y-axis) and saved compute (x-axis) for Llama-3.2-8b GateSkip with individual vector gates.

J.14 GEMMA-2-2B GATESKIP WITH INDIVIDUAL VECTOR GATES (a) CommonsenseQA (b) HellaSwag (c) MMLU STEM (d) OpenBookQA Compute Saved (%) (e) PIQA (f) WinoGrande (g) CSQA (Gen.) (h) GSM8K (Gen.)

Figure 19: Tradeoff plots between accuracy (y-axis) and saved compute (x-axis) for Gemma-2-2b GateSkip with individual vector gates.

J.15 LLAMA-3.2-1B CALM (HIDDEN STATE SATURATION) (a) CommonsenseQA (b) HellaSwag (c) MMLU STEM (d) OpenBookQA Compute Saved (%) (e) PIQA (f) WinoGrande (g) CSQA (Gen.) (h) GSM8K (Gen.)

Figure 20: Tradeoff plots between accuracy (y-axis) and saved compute (x-axis) for Llama-3.2-1b CALM (hidden state saturation).

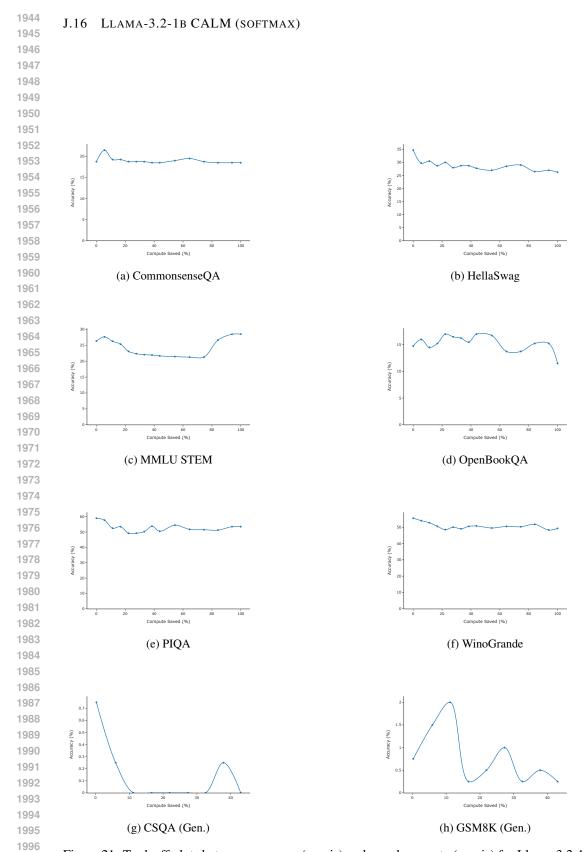


Figure 21: Tradeoff plots between accuracy (y-axis) and saved compute (x-axis) for Llama-3.2-1b CALM (softmax).

J.17 LLAMA-3.2-1B MIXTURE-OF-DEPTHS (a) CommonsenseQA (b) HellaSwag (c) MMLU STEM (d) OpenBookQA Compute Saved (%) (e) PIQA (f) WinoGrande Compute Saved (%) (g) CSQA (Gen.) (h) GSM8K (Gen.)

Figure 22: Tradeoff plots between accuracy (y-axis) and saved compute (x-axis) for Llama-3.2-1b Mixture-of-Depths.

J.18 Llama-3.2-1b random skipping baseline trained and tested on the WMT16-EN-RO dataset

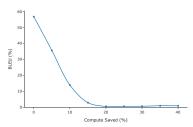


Figure 23: Tradeoff plots between accuracy (y-axis) and saved compute (x-axis) for Llama-3.2-1b random skipping baseline trained and tested on the WMT16-EN-RO dataset.

J.19 LLAMA-3.2-1B GATESKIP (INDIVIDUAL VECTOR GATES) TRAINED AND TESTED ON THE WMT16-EN-RO DATASET

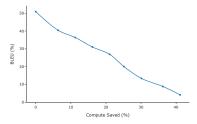


Figure 24: Tradeoff plots between accuracy (y-axis) and saved compute (x-axis) for Llama-3.2-1b GateSkip (individual vector gates) trained and tested on the WMT16-EN-RO dataset.

K RESULTS WITH STANDARD DEVIATION

K.1 LLAMA-3.2-1B RANDOM SKIPPING BASELINE

Compute Saved (%)	CSQA	HellaSwag	MMLU Stem	Open-BookQA	PIQA	Wino-Grande
0.00%	32.50±2.34	44.75±2.49	25.88±0.78	30.00±2.29	74.50±2.18	62.50±2.42
5.00%	21.00 ± 2.04	36.75 ± 2.41	26.93 ± 0.79	21.75 ± 2.07	59.00 ± 2.46	52.00 ± 2.50
9.99%	20.25 ± 2.01	31.50 ± 2.33	27.12 ± 0.79	16.75 ± 1.87	52.00 ± 2.50	49.50 ± 2.50
15.00%	19.00 ± 1.96	29.00 ± 2.27	24.93 ± 0.77	17.25 ± 1.89	$49.25{\pm}2.50$	52.50 ± 2.50
19.99%	20.75 ± 2.03	27.75 ± 2.24	23.79 ± 0.76	13.75 ± 1.72	49.25 ± 2.50	52.75 ± 2.50
25.00%	17.25 ± 1.89	27.50 ± 2.24	24.39 ± 0.76	12.25 ± 1.64	52.25 ± 2.50	52.00 ± 2.50
29.99%	$21.25{\pm}2.05$	27.00 ± 2.22	23.31 ± 0.75	13.25 ± 1.70	54.25 ± 2.49	50.00 ± 2.50
35.00%	20.25 ± 2.01	26.00 ± 2.20	23.63 ± 0.76	11.25 ± 1.58	53.50 ± 2.50	49.00 ± 2.50
40.00%	18.50 ± 1.94	28.75 ± 2.27	24.90 ± 0.77	11.75 ± 1.61	$53.25{\pm}2.50$	$49.25{\pm}2.50$
50.00%	20.75 ± 2.03	$28.25{\pm}2.25$	24.36 ± 0.76	12.75 ± 1.67	$52.25{\pm}2.50$	49.25 ± 2.50
60.01%	20.00 ± 2.00	26.50 ± 2.21	24.61 ± 0.77	13.00 ± 1.68	53.50 ± 2.50	48.50 ± 2.50
70.01%	14.25 ± 1.75	25.75 ± 2.19	25.28 ± 0.77	13.00 ± 1.68	52.50 ± 2.50	47.75 ± 2.50
80.01%	18.50 ± 1.94	27.50 ± 2.24	24.74 ± 0.77	16.25 ± 1.85	$50.25{\pm}2.50$	46.25 ± 2.50
90.00%	19.50 ± 1.98	27.25 ± 2.23	24.67 ± 0.77	15.75 ± 1.82	46.75 ± 2.50	49.75 ± 2.50
95.00%	$20.25{\pm}2.01$	26.75 ± 2.22	23.72 ± 0.76	$16.50 {\pm} 1.86$	50.50 ± 2.50	46.75 ± 2.50

Table 15: Results for loglikelihood-based benchmarks for Llama-3.2-1b random skipping baseline

Compute Saved (%)	CSQA (Gen.)	GSM8K (Gen.)
0.00%	42.75±2.48	19.75±1.99
5.00%	16.75 ± 1.87	2.75 ± 0.82
10.00%	$3.25{\pm}0.89$	0.50 ± 0.35
14.99%	6.00 ± 1.19	0.00 ± 0.00
19.99%	3.00 ± 0.85	0.00 ± 0.00
25.00%	2.00 ± 0.70	$0.25{\pm}0.25$
29.98%	0.50 ± 0.35	0.00 ± 0.00
34.97%	0.50 ± 0.35	0.00 ± 0.00
40.00%	0.00 ± 0.00	0.00 ± 0.00
49.98%	0.00 ± 0.00	0.00 ± 0.00
59.97%	0.00 ± 0.00	0.00 ± 0.00
70.00%	0.00 ± 0.00	0.00 ± 0.00
80.00%	0.00 ± 0.00	0.00 ± 0.00
90.00%	0.00 ± 0.00	0.00 ± 0.00
95.00%	0.00 ± 0.00	0.00 ± 0.00

Table 16: Results for generative benchmarks for Llama-3.2-1b random skipping baseline

K.2 Llama-3.2-1B GateSkip with individual vector gates

Compute Saved (%)	CSQA	HellaSwag	MMLU Stem	Open-BookQA	PIQA	Wino-Grande
0.00%	26.00 ± 2.20	46.00 ± 2.50	28.35 ± 0.80	30.00 ± 2.29	74.75 ± 2.17	60.50±2.45
5.12%	19.00 ± 1.96	44.00 ± 2.49	27.62 ± 0.79	28.00 ± 2.25	71.00 ± 2.27	58.50 ± 2.47
10.11%	19.75 ± 1.99	40.00 ± 2.45	26.93 ± 0.79	26.25 ± 2.20	68.00 ± 2.34	58.50 ± 2.47
15.08%	$21.25{\pm}2.05$	38.50 ± 2.44	26.04 ± 0.78	$21.25{\pm}2.05$	66.50 ± 2.36	53.00 ± 2.50
20.08%	19.75 ± 1.99	$35.25{\pm}2.39$	25.44 ± 0.78	18.75 ± 1.95	59.00 ± 2.46	51.50 ± 2.50
25.12%	22.50 ± 2.09	33.25 ± 2.36	24.14 ± 0.76	16.00 ± 1.84	56.75 ± 2.48	44.75 ± 2.49
30.05%	19.75 ± 1.99	29.75 ± 2.29	23.88 ± 0.76	15.00 ± 1.79	54.00 ± 2.50	54.00 ± 2.50
35.06%	$20.25{\pm}2.01$	27.25 ± 2.23	22.77 ± 0.75	14.50 ± 1.76	54.75 ± 2.49	49.75 ± 2.50
40.06%	18.75 ± 1.95	$30.25{\pm}2.30$	23.98 ± 0.76	12.75 ± 1.67	52.00 ± 2.50	50.00 ± 2.50
50.08%	17.00 ± 1.88	26.00 ± 2.20	23.34 ± 0.75	10.50 ± 1.53	49.50 ± 2.50	55.00 ± 2.49
60.10%	21.00 ± 2.04	27.50 ± 2.24	23.63 ± 0.76	12.25 ± 1.64	51.50 ± 2.50	51.25 ± 2.50
70.07%	23.50 ± 2.12	23.75 ± 2.13	24.23 ± 0.76	12.00 ± 1.63	54.25 ± 2.49	47.00 ± 2.50
80.17%	19.25 ± 1.97	26.75 ± 2.22	25.59 ± 0.78	11.25 ± 1.58	52.75 ± 2.50	45.50 ± 2.49
90.37%	21.00 ± 2.04	27.75 ± 2.24	24.36 ± 0.76	16.75 ± 1.87	50.75 ± 2.50	53.00 ± 2.50
95.30%	20.50 ± 2.02	$26.25{\pm}2.20$	21.92 ± 0.73	15.25 ± 1.80	$52.25{\pm}2.50$	49.50 ± 2.50

Table 17: Results for loglikelihood-based benchmarks for Llama-3.2-1b GateSkip with individual vector gates

Compute Saved (%)	CSQA (Gen.)	GSM8K (Gen.)
0.00%	$38.75{\pm}2.44$	14.75 ± 1.78
14.36%	36.75 ± 2.41	9.50 ± 1.47
19.68%	36.25 ± 2.41	11.50 ± 1.60
26.77%	27.50 ± 2.24	9.50 ± 1.47
27.90%	15.25 ± 1.80	2.75 ± 0.82
28.54%	25.25 ± 2.17	10.50 ± 1.53
31.27%	19.25 ± 1.97	7.75 ± 1.34
32.72%	8.00 ± 1.36	1.50 ± 0.61
40.95%	$2.25{\pm}0.74$	1.50 ± 0.61

Table 18: Results for generative benchmarks for Llama-3.2-1b GateSkip with individual vector gates

K.3 Llama-3.2-1b GateSkip with individual vector gates, skipping only applied at attention layers

Compute Saved (%)	CSQA	HellaSwag	MMLU Stem	Open-BookQA	PIQA	Wino-Grande
0.00%	26.00±2.20	46.00±2.50	28.35±0.80	30.00 ± 2.29	74.75±2.17	60.50±2.45
5.07%	$23.25{\pm}2.11$	41.50 ± 2.47	27.56 ± 0.79	26.00 ± 2.20	71.75 ± 2.25	55.25 ± 2.49
10.05%	23.25 ± 2.11	38.25 ± 2.43	27.28 ± 0.79	22.50 ± 2.09	69.25 ± 2.31	54.25 ± 2.49
15.03%	21.25 ± 2.05	35.00 ± 2.39	25.94 ± 0.77	$20.25{\pm}2.01$	69.25 ± 2.31	53.25 ± 2.50
20.04%	$22.25{\pm}2.08$	34.00 ± 2.37	26.23 ± 0.78	17.25 ± 1.89	62.25 ± 2.43	47.50 ± 2.50
25.06%	$21.25{\pm}2.05$	32.50 ± 2.34	25.44 ± 0.77	14.50 ± 1.76	58.50 ± 2.47	50.50 ± 2.50
30.03%	17.75 ± 1.91	28.75 ± 2.27	21.88 ± 0.74	12.50 ± 1.66	53.00 ± 2.50	50.50 ± 2.50
35.04%	17.75 ± 1.91	28.00 ± 2.25	21.54 ± 0.73	11.75 ± 1.61	55.25 ± 2.49	47.50 ± 2.50
40.05%	18.75 ± 1.95	28.50 ± 2.26	21.25 ± 0.73	12.00 ± 1.63	53.00 ± 2.50	51.75 ± 2.50
50.00%	18.75 ± 1.95	27.75 ± 2.24	21.25 ± 0.73	12.50 ± 1.66	51.00 ± 2.50	50.50 ± 2.50
50.00%	18.75 ± 1.95	27.75 ± 2.24	21.25 ± 0.73	12.50 ± 1.66	51.00 ± 2.50	50.50 ± 2.50
50.00%	18.75 ± 1.95	27.75 ± 2.24	21.25 ± 0.73	12.50 ± 1.66	51.00 ± 2.50	50.50 ± 2.50
50.00%	18.75 ± 1.95	27.75 ± 2.24	21.25 ± 0.73	12.50 ± 1.66	51.00 ± 2.50	50.50 ± 2.50
50.00%	18.75 ± 1.95	27.75 ± 2.24	21.25 ± 0.73	12.50 ± 1.66	51.00 ± 2.50	50.50 ± 2.50
50.00%	18.75 ± 1.95	27.75 ± 2.24	21.25 ± 0.73	12.50 ± 1.66	51.00 ± 2.50	50.50 ± 2.50

Table 19: Results for loglikelihood-based benchmarks for Llama-3.2-1b GateSkip with individual vector gates when only skipping attention layers

Compute Saved (%)	CSQA (Gen.)	GSM8K (Gen.)
0.00%	38.75 ± 2.44	14.75 ± 1.78
8.61%	34.75 ± 2.38	6.25 ± 1.21
13.29%	26.25 ± 2.20	6.25 ± 1.21
29.66%	$4.25{\pm}1.01$	1.75 ± 0.66
32.00%	3.00 ± 0.85	$2.25{\pm}0.74$
42.79%	$5.25{\pm}1.12$	$0.25{\pm}0.25$
49.76%	0.75 ± 0.43	$0.25{\pm}0.25$
49.96%	0.00 ± 0.00	0.00 ± 0.00
49.99%	0.00 ± 0.00	0.00 ± 0.00

Table 20: Results for generative benchmarks for Llama-3.2-1b GateSkip with individual vector gates when only skipping attention layers

K.4 LLAMA-3.2-1B GATESKIP WITH INDIVIDUAL VECTOR GATES, SKIPPING ONLY APPLIED AT MLP LAYERS

Compute Saved (%)	CSQA	HellaSwag	MMLU Stem	Open-BookQA	PIQA	Wino-Grande
0.00%	26.00±2.20	46.00±2.50	28.35±0.80	30.00±2.29	74.75±2.17	60.50±2.45
5.02%	22.50 ± 2.09	37.25 ± 2.42	25.31 ± 0.77	23.75 ± 2.13	64.25 ± 2.40	53.75 ± 2.50
10.02%	21.50 ± 2.06	33.75 ± 2.37	26.45 ± 0.78	19.00 ± 1.96	62.00 ± 2.43	52.25 ± 2.50
15.03%	17.75 ± 1.91	30.00 ± 2.29	25.63 ± 0.78	15.50 ± 1.81	55.00 ± 2.49	48.25 ± 2.50
20.03%	18.00 ± 1.92	28.75 ± 2.27	27.91 ± 0.79	12.00 ± 1.63	54.75 ± 2.49	47.50 ± 2.50
25.03%	18.25 ± 1.93	26.75 ± 2.22	28.54 ± 0.80	12.25 ± 1.64	$53.25{\pm}2.50$	53.50 ± 2.50
30.02%	19.50 ± 1.98	26.75 ± 2.22	25.34 ± 0.77	11.75 ± 1.61	48.25 ± 2.50	50.75 ± 2.50
35.02%	19.50 ± 1.98	26.00 ± 2.20	26.96 ± 0.79	13.75 ± 1.72	50.50 ± 2.50	54.00 ± 2.50
40.02%	20.75 ± 2.03	24.50 ± 2.15	26.17 ± 0.78	15.50 ± 1.81	54.25 ± 2.49	51.50 ± 2.50
50.00%	19.50 ± 1.98	24.00 ± 2.14	26.26 ± 0.78	16.25 ± 1.85	55.75 ± 2.49	50.25 ± 2.50
50.00%	19.50 ± 1.98	24.00 ± 2.14	26.26 ± 0.78	16.25 ± 1.85	55.75 ± 2.49	50.25 ± 2.50
50.00%	19.50 ± 1.98	24.00 ± 2.14	26.26 ± 0.78	16.25 ± 1.85	55.75 ± 2.49	50.25 ± 2.50
50.00%	19.50 ± 1.98	24.00 ± 2.14	26.26 ± 0.78	16.25 ± 1.85	55.75 ± 2.49	50.25 ± 2.50
50.00%	19.50 ± 1.98	24.00 ± 2.14	26.26 ± 0.78	16.25 ± 1.85	55.75 ± 2.49	50.25 ± 2.50
50.00%	19.50 ± 1.98	24.00 ± 2.14	26.26 ± 0.78	16.25 ± 1.85	55.75 ± 2.49	$50.25{\pm}2.50$

Table 21: Results for loglikelihood-based benchmarks for Llama-3.2-1b GateSkip with individual vector gates when only skipping attention layers

Compute Saved (%)	CSQA (Gen.)	GSM8K (Gen.)
0.00%	38.75±2.44	14.75±1.78
6.62%	33.75 ± 2.37	12.75 ± 1.67
11.45%	21.00 ± 2.04	12.25 ± 1.64
15.42%	5.00 ± 1.09	8.50 ± 1.40
20.34%	$0.25{\pm}0.25$	1.25 ± 0.56
25.18%	0.00 ± 0.00	0.75 ± 0.43
30.09%	$0.25{\pm}0.25$	0.50 ± 0.35
34.96%	0.00 ± 0.00	0.00 ± 0.00
39.85%	0.00 ± 0.00	0.00 ± 0.00

Table 22: Results for generative benchmarks for Llama-3.2-1b GateSkip with individual vector gates when only skipping attention layers

K.5 LLAMA-3.2-1B GATESKIP WITH INDIVIDUAL VECTOR GATES, SKIPPING ENTIRE LAYERS BASED ON THE ATTENTION GATE

Compute Saved (%)	CSQA	HellaSwag	MMLU Stem	Open-BookQA	PIQA	Wino-Grande
0.00%	26.00±2.20	46.00±2.50	28.35±0.80	30.00±2.29	74.75±2.17	60.50±2.45
5.12%	19.00 ± 1.96	44.00 ± 2.49	27.62 ± 0.79	28.00 ± 2.25	71.00 ± 2.27	58.50 ± 2.47
10.11%	19.75 ± 1.99	40.00 ± 2.45	26.93 ± 0.79	$26.25{\pm}2.20$	68.00 ± 2.34	58.50 ± 2.47
15.08%	21.25 ± 2.05	38.50 ± 2.44	26.04 ± 0.78	21.25 ± 2.05	66.50 ± 2.36	53.00 ± 2.50
20.08%	19.75 ± 1.99	$35.25{\pm}2.39$	25.44 ± 0.78	18.75 ± 1.95	59.00 ± 2.46	51.50 ± 2.50
25.12%	22.50 ± 2.09	$33.25{\pm}2.36$	24.14 ± 0.76	16.00 ± 1.84	56.75 ± 2.48	44.75 ± 2.49
30.05%	19.75 ± 1.99	29.75 ± 2.29	23.88 ± 0.76	15.00 ± 1.79	54.00 ± 2.50	54.00 ± 2.50
35.06%	20.25 ± 2.01	27.25 ± 2.23	22.77 ± 0.75	14.50 ± 1.76	54.75 ± 2.49	49.75 ± 2.50
40.06%	18.75 ± 1.95	$30.25{\pm}2.30$	23.98 ± 0.76	12.75 ± 1.67	52.00 ± 2.50	50.00 ± 2.50
50.08%	17.00 ± 1.88	26.00 ± 2.20	23.34 ± 0.75	10.50 ± 1.53	49.50 ± 2.50	55.00 ± 2.49
60.10%	21.00 ± 2.04	27.50 ± 2.24	23.63 ± 0.76	12.25 ± 1.64	51.50 ± 2.50	51.25 ± 2.50
70.07%	23.50 ± 2.12	23.75 ± 2.13	24.23 ± 0.76	12.00 ± 1.63	54.25 ± 2.49	47.00 ± 2.50
80.17%	19.25 ± 1.97	26.75 ± 2.22	25.59 ± 0.78	11.25 ± 1.58	52.75 ± 2.50	45.50 ± 2.49
90.37%	21.00 ± 2.04	27.75 ± 2.24	24.36 ± 0.76	16.75 ± 1.87	50.75 ± 2.50	53.00 ± 2.50
95.30%	20.50 ± 2.02	$26.25{\pm}2.20$	21.92 ± 0.73	15.25 ± 1.80	$52.25{\pm}2.50$	49.50 ± 2.50

Table 23: Results for loglikelihood-based benchmarks for Llama-3.2-1b GateSkip with individual vector gates when skipping entire layers based on the attention gate

Compute Saved (%)	CSQA (Gen.)	GSM8K (Gen.)
0.00%	38.75 ± 2.44	14.75 ± 1.78
14.36%	36.75 ± 2.41	9.50 ± 1.47
19.68%	36.25 ± 2.41	11.50 ± 1.60
26.77%	27.50 ± 2.24	9.50 ± 1.47
27.90%	15.25 ± 1.80	2.75 ± 0.82
28.54%	$25.25{\pm}2.17$	10.50 ± 1.53
31.27%	19.25 ± 1.97	7.75 ± 1.34
32.72%	8.00 ± 1.36	1.50 ± 0.61
40.95%	2.25±0.74	1.50±0.61

Table 24: Results for generative benchmarks for Llama-3.2-1b GateSkip with individual vector gates when skipping entire layers based on the attention gate

K.6 LLAMA-3.2-1B GATESKIP WITH SHARED VECTOR GATES

Compute Saved (%)	CSQA	HellaSwag	MMLU Stem	Open-BookQA	PIQA	Wino-Grande
0.00%	28.00 ± 2.25	45.50 ± 2.49	28.20 ± 0.80	28.50 ± 2.26	$74.25{\pm}2.19$	61.50±2.44
5.05%	$24.25{\pm}2.15$	39.50 ± 2.45	26.10 ± 0.78	27.50 ± 2.24	$71.25{\pm}2.27$	56.50 ± 2.48
10.04%	22.00 ± 2.07	40.00 ± 2.45	25.75 ± 0.78	26.50 ± 2.21	$71.25{\pm}2.27$	53.25 ± 2.50
15.07%	21.25 ± 2.05	38.75 ± 2.44	25.31 ± 0.77	23.25 ± 2.11	68.50 ± 2.33	53.25 ± 2.50
20.04%	22.00 ± 2.07	35.75 ± 2.40	26.42 ± 0.78	23.75 ± 2.13	67.00 ± 2.35	55.75 ± 2.49
25.07%	19.50 ± 1.98	$34.25{\pm}2.38$	25.66 ± 0.77	19.00 ± 1.96	64.75 ± 2.39	55.50 ± 2.49
30.08%	22.75 ± 2.10	34.50 ± 2.38	26.51 ± 0.78	17.50 ± 1.90	60.50 ± 2.45	50.25 ± 2.50
35.03%	19.00 ± 1.96	34.00 ± 2.37	26.42 ± 0.78	16.50 ± 1.86	60.25 ± 2.45	48.50 ± 2.50
40.04%	21.75 ± 2.07	30.75 ± 2.31	24.64 ± 0.76	13.75 ± 1.72	51.75 ± 2.50	52.00 ± 2.50
50.04%	19.75 ± 1.99	26.75 ± 2.22	24.17 ± 0.76	12.50 ± 1.66	51.00 ± 2.50	52.75 ± 2.50
60.10%	19.25 ± 1.97	29.00 ± 2.27	24.71 ± 0.77	13.50 ± 1.71	54.50 ± 2.49	49.25 ± 2.50
70.06%	19.00 ± 1.96	27.25 ± 2.23	23.69 ± 0.76	13.50 ± 1.71	48.75 ± 2.50	50.75 ± 2.50
80.12%	21.75 ± 2.07	25.00 ± 2.17	23.37 ± 0.75	12.00 ± 1.63	51.50 ± 2.50	47.75 ± 2.50
90.31%	20.00 ± 2.00	26.50 ± 2.21	25.88 ± 0.78	17.25 ± 1.89	51.00 ± 2.50	50.00 ± 2.50
95.44%	$20.25{\pm}2.01$	27.00 ± 2.22	26.61 ± 0.79	14.75 ± 1.78	52.75 ± 2.50	47.25 ± 2.50

Table 25: Results for loglikelihood-based benchmarks for Llama-3.2-1b GateSkip with shared vector gates

Compute Saved (%)	CSQA (Gen.)	GSM8K (Gen.)
0.00%	$33.25{\pm}2.36$	10.25 ± 1.52
18.45%	30.75 ± 2.31	10.25 ± 1.52
22.67%	25.25 ± 2.17	10.00 ± 1.50
24.39%	11.00 ± 1.57	6.00 ± 1.19
26.27%	20.00 ± 2.00	8.75 ± 1.41
34.75%	12.25 ± 1.64	7.00 ± 1.28
46.72%	13.75 ± 1.72	4.25 ± 1.01
89.10%	0.75 ± 0.43	1.50 ± 0.61
98.91%	0.00 ± 0.00	0.00 ± 0.00
99.19%	0.00 ± 0.00	0.00 ± 0.00
99.47%	0.00 ± 0.00	0.00 ± 0.00
99.51%	0.00 ± 0.00	0.00 ± 0.00
99.79%	0.00 ± 0.00	0.00 ± 0.00
99.99%	0.00 ± 0.00	0.00 ± 0.00
100.00%	0.00 ± 0.00	0.00 ± 0.00

Table 26: Results for generative benchmarks for Llama-3.2-1b GateSkip with shared vector gates

K.7 LLAMA-3.2-1B GATESKIP WITH INDIVIDUAL SCALAR GATES

Wino-Grande
20 60.00±2.45
55.75 ± 2.49
29 55.75±2.49
53.00 ± 2.50
52.00 ± 2.50
14 50.25±2.50
16 54.75±2.49
49.25 \pm 2.50
48.50 ± 2.50
48.50 ± 2.50
$50 53.50 \pm 2.50$
$60 47.00 \pm 2.50$
$50 52.50 \pm 2.50$
50 51.00±2.50
$50 51.50 \pm 2.50$
5 5 5

Table 27: Results for loglikelihood-based benchmarks for Llama-3.2-1b GateSkip with individual scalar gates

Compute Saved (%)	CSQA (Gen.)	GSM8K (Gen.)
0.00%	35.25 ± 2.39	12.00 ± 1.63
14.12%	$31.25{\pm}2.32$	10.50 ± 1.53
16.40%	25.75 ± 2.19	8.50 ± 1.40
18.91%	27.25 ± 2.23	9.50 ± 1.47
29.08%	11.75 ± 1.61	4.75 ± 1.06
31.46%	16.00 ± 1.84	8.00 ± 1.36
33.14%	10.75 ± 1.55	3.25 ± 0.89
39.21%	2.25 ± 0.74	1.75 ± 0.66
40.51%	0.50 ± 0.35	1.25 ± 0.56
62.03%	1.75 ± 0.66	0.00 ± 0.00
69.27%	1.25 ± 0.56	0.00 ± 0.00
87.41%	0.25 ± 0.25	0.00 ± 0.00
93.35%	0.00 ± 0.00	0.00 ± 0.00
99.81%	0.00 ± 0.00	0.00 ± 0.00
99.99%	0.00 ± 0.00	0.00 ± 0.00

Table 28: Results for generative benchmarks for Llama-3.2-1b GateSkip with individual scalar gates

K.8 Llama-3.2-1b GateSkip with individual vector gates, skipping only at every second layer

Compute Saved (%)	CSQA	HellaSwag	MMLU Stem	Open-BookQA	PIQA	Wino-Grande
0.00%	26.00±2.20	46.00±2.50	28.35±0.80	30.00 ± 2.29	74.75±2.17	60.50±2.45
2.52%	21.75 ± 2.07	40.50 ± 2.46	24.39 ± 0.76	23.50 ± 2.12	69.75 ± 2.30	54.75 ± 2.49
5.02%	23.75 ± 2.13	39.00 ± 2.44	24.80 ± 0.77	25.25 ± 2.17	68.00 ± 2.34	54.00 ± 2.50
7.52%	21.50 ± 2.06	38.25 ± 2.43	25.40 ± 0.77	21.75 ± 2.07	65.75 ± 2.38	55.25 ± 2.49
10.02%	20.75 ± 2.03	38.75 ± 2.44	25.88 ± 0.78	$20.25{\pm}2.01$	66.25 ± 2.37	51.25 ± 2.50
12.52%	19.75 ± 1.99	36.75 ± 2.41	26.20 ± 0.78	20.00 ± 2.00	64.50 ± 2.40	49.50 ± 2.50
15.02%	16.25 ± 1.85	35.50 ± 2.40	25.82 ± 0.78	19.25 ± 1.97	61.75 ± 2.43	54.50 ± 2.49
17.52%	23.25 ± 2.11	35.75 ± 2.40	25.50 ± 0.78	15.75 ± 1.82	60.75 ± 2.44	51.50 ± 2.50
20.02%	20.25 ± 2.01	35.25 ± 2.39	25.06 ± 0.77	15.50 ± 1.81	59.25 ± 2.46	51.50 ± 2.50
25.02%	21.25 ± 2.05	31.75 ± 2.33	24.93 ± 0.77	16.50 ± 1.86	56.50 ± 2.48	50.75 ± 2.50
30.02%	19.50 ± 1.98	31.25 ± 2.32	25.72 ± 0.77	14.25 ± 1.75	57.25 ± 2.48	51.50 ± 2.50
35.02%	20.50 ± 2.02	30.25 ± 2.30	25.15 ± 0.77	12.75 ± 1.67	58.00 ± 2.47	48.75 ± 2.50
40.01%	20.25 ± 2.01	28.00 ± 2.25	25.88 ± 0.78	14.00 ± 1.74	51.50 ± 2.50	44.75 ± 2.49
45.01%	19.75 ± 1.99	28.75 ± 2.27	27.31 ± 0.79	11.75 ± 1.61	55.50 ± 2.49	46.25 ± 2.50
47.51%	$20.75{\pm}2.03$	27.50 ± 2.24	28.07 ± 0.79	14.00 ± 1.74	53.00 ± 2.50	47.75 ± 2.50

Table 29: Results for loglikelihood-based benchmarks for Llama-3.2-1b GateSkip with individual vector gates when skipping only at every second layer

Compute Saved (%)	CSQA (Gen.)	GSM8K (Gen.)
0.00%	38.75 ± 2.44	14.75 ± 1.78
5.35%	38.25 ± 2.43	12.50 ± 1.66
6.68%	42.00 ± 2.47	12.25 ± 1.64
8.12%	35.25 ± 2.39	11.50 ± 1.60
11.58%	26.00 ± 2.20	9.75 ± 1.49
12.10%	33.75 ± 2.37	10.50 ± 1.53
14.75%	$22.25{\pm}2.08$	$8.25{\pm}1.38$
17.95%	17.50 ± 1.90	6.00 ± 1.19
20.74%	14.25 ± 1.75	4.75 ± 1.06

Table 30: Results for generative benchmarks for Llama-3.2-1b GateSkip with individual vector gates when skipping only at every second layer

K.9 LLAMA-3.2-1B GATESKIP WITH INDIVIDUAL VECTOR MLP GATES

Compute Saved (%)	CSQA	HellaSwag	MMLU Stem	Open-BookQA	PIQA	Wino-Grande
0.00%	27.25±2.23	44.00±2.49	27.94±0.80	29.50±2.28	74.25±2.19	60.75±2.44
5.13%	20.50 ± 2.02	$43.00{\pm}2.48$	27.21 ± 0.79	25.75 ± 2.19	70.00 ± 2.29	58.75 ± 2.46
10.14%	19.00 ± 1.96	40.25 ± 2.46	25.69 ± 0.78	20.75 ± 2.03	69.75 ± 2.30	56.00 ± 2.49
15.11%	18.50 ± 1.94	35.75 ± 2.40	25.72 ± 0.78	16.75 ± 1.87	60.00 ± 2.45	46.25 ± 2.50
20.11%	18.00 ± 1.92	$35.25{\pm}2.39$	25.02 ± 0.77	16.50 ± 1.86	57.75 ± 2.47	53.25 ± 2.50
25.08%	17.00 ± 1.88	33.00 ± 2.35	28.04 ± 0.80	15.00 ± 1.79	55.25 ± 2.49	53.50 ± 2.50
30.07%	18.25 ± 1.93	28.00 ± 2.25	26.23 ± 0.78	14.50 ± 1.76	56.00 ± 2.49	47.75 ± 2.50
35.07%	21.25 ± 2.05	30.00 ± 2.29	22.90 ± 0.75	14.25 ± 1.75	57.75 ± 2.47	46.00 ± 2.50
40.08%	18.00 ± 1.92	25.50 ± 2.18	21.82 ± 0.73	13.50 ± 1.71	54.50 ± 2.49	47.50 ± 2.50
50.10%	18.75 ± 1.95	26.50 ± 2.21	22.65 ± 0.74	11.50 ± 1.60	52.50 ± 2.50	50.25 ± 2.50
60.08%	16.25 ± 1.85	27.25 ± 2.23	24.52 ± 0.77	13.25 ± 1.70	53.00 ± 2.50	49.75 ± 2.50
70.07%	17.00 ± 1.88	28.00 ± 2.25	23.88 ± 0.76	16.25 ± 1.85	54.50 ± 2.49	52.00 ± 2.50
80.14%	$20.25{\pm}2.01$	28.00 ± 2.25	24.90 ± 0.77	17.00 ± 1.88	52.00 ± 2.50	51.50 ± 2.50
90.36%	18.75 ± 1.95	$28.25{\pm}2.25$	26.20 ± 0.78	14.00 ± 1.74	51.75 ± 2.50	46.75 ± 2.50
95.43%	18.75 ± 1.95	25.00 ± 2.17	26.64 ± 0.79	14.00 ± 1.74	55.00 ± 2.49	$49.25{\pm}2.50$

Table 31: Results for loglikelihood-based benchmarks for Llama-3.2-1b GateSkip with individual vector MLP gates

Compute Saved (%)	CSQA (Gen.)	GSM8K (Gen.)
0.00%	39.75±2.45	8.75±1.41
12.91%	38.75 ± 2.44	10.00 ± 1.50
14.85%	$28.25{\pm}2.25$	8.75 ± 1.41
28.65%	15.50 ± 1.81	6.00 ± 1.19
28.94%	$20.25{\pm}2.01$	8.75 ± 1.41
30.81%	20.50 ± 2.02	6.00 ± 1.19
77.60%	$0.25{\pm}0.25$	2.00 ± 0.70
90.55%	0.00 ± 0.00	0.00 ± 0.00
91.03%	0.00 ± 0.00	0.00 ± 0.00
99.31%	0.00 ± 0.00	0.00 ± 0.00
99.32%	0.00 ± 0.00	0.00 ± 0.00
99.49%	0.00 ± 0.00	0.00 ± 0.00
99.62%	0.00 ± 0.00	0.00 ± 0.00
99.67%	0.00 ± 0.00	0.00 ± 0.00
100.00%	0.00 ± 0.00	0.00 ± 0.00

Table 32: Results for generative benchmarks for Llama-3.2-1b GateSkip with individual vector MLP gates

K.10 Llama-3.2-1b GateSkip with individual vector gates at the entry point to each module instead of the exit point

Compute Saved (%)	CSQA	HellaSwag	MMLU Stem	Open-BookQA	PIQA	Wino-Grande
0.00%	21.00 ± 2.04	41.25±2.46	28.96 ± 0.80	22.00 ± 2.07	69.00 ± 2.32	58.00±2.47
5.04%	20.25 ± 2.01	36.50 ± 2.41	28.04 ± 0.79	15.75 ± 1.82	63.50 ± 2.41	50.75 ± 2.50
10.04%	19.25 ± 1.97	35.25 ± 2.39	27.50 ± 0.79	15.50 ± 1.81	60.50 ± 2.45	48.50 ± 2.50
15.04%	21.25 ± 2.05	36.25 ± 2.41	27.15 ± 0.79	14.75 ± 1.78	61.75 ± 2.43	53.25 ± 2.50
20.05%	20.25 ± 2.01	33.00 ± 2.35	26.99 ± 0.79	14.00 ± 1.74	57.50 ± 2.47	51.25 ± 2.50
25.04%	18.25 ± 1.93	31.75 ± 2.33	25.50 ± 0.77	14.00 ± 1.74	$59.25{\pm}2.46$	49.75 ± 2.50
30.05%	21.75 ± 2.07	33.00 ± 2.35	24.55 ± 0.76	15.50 ± 1.81	56.00 ± 2.49	48.50 ± 2.50
35.05%	21.00 ± 2.04	30.75 ± 2.31	23.69 ± 0.76	15.50 ± 1.81	$56.25{\pm}2.48$	50.50 ± 2.50
40.04%	21.50 ± 2.06	$33.25{\pm}2.36$	23.72 ± 0.76	15.00 ± 1.79	54.00 ± 2.50	50.75 ± 2.50
50.05%	19.25 ± 1.97	27.50 ± 2.24	24.10 ± 0.76	12.00 ± 1.63	50.50 ± 2.50	47.75 ± 2.50
60.04%	19.75 ± 1.99	$28.25{\pm}2.25$	23.95 ± 0.76	13.25 ± 1.70	51.00 ± 2.50	48.50 ± 2.50
70.09%	22.75 ± 2.10	32.50 ± 2.34	23.66 ± 0.76	15.50 ± 1.81	52.75 ± 2.50	50.00 ± 2.50
80.19%	21.50 ± 2.06	$28.25{\pm}2.25$	23.69 ± 0.76	15.75 ± 1.82	50.00 ± 2.50	49.00 ± 2.50
90.35%	19.50 ± 1.98	24.00 ± 2.14	24.80 ± 0.77	15.50 ± 1.81	48.50 ± 2.50	49.75 ± 2.50
95.81%	18.75 ± 1.95	$24.25{\pm}2.15$	21.25 ± 0.73	15.00 ± 1.79	50.00 ± 2.50	52.00 ± 2.50

Table 33: Results for loglikelihood-based benchmarks for Llama-3.2-1b GateSkip with individual vector gates at the entry point to each module instead of the exit point

Compute Saved (%)	CSQA (Gen.)	GSM8K (Gen.)
0.00%	33.00 ± 2.35	$9.25{\pm}1.45$
5.78%	3.50 ± 0.92	2.50 ± 0.78
10.86%	0.75 ± 0.43	$1.25{\pm}0.56$
15.86%	1.50 ± 0.61	0.50 ± 0.35
27.22%	1.50 ± 0.61	0.00 ± 0.00
32.72%	0.00 ± 0.00	0.00 ± 0.00
38.60%	0.50 ± 0.35	0.00 ± 0.00
43.57%	0.75 ± 0.43	1.00 ± 0.50
51.40%	3.00 ± 0.85	0.75 ± 0.43
71.86%	0.50 ± 0.35	0.00 ± 0.00
74.25%	0.00 ± 0.00	0.00 ± 0.00
91.01%	0.00 ± 0.00	0.00 ± 0.00
95.05%	0.00 ± 0.00	0.00 ± 0.00
99.15%	0.00 ± 0.00	0.00 ± 0.00
100.00%	0.00 ± 0.00	0.00 ± 0.00

Table 34: Results for generative benchmarks for Llama-3.2-1b GateSkip with individual vector gates at the entry point to each module instead of the exit point

K.11 LLAMA-3.2-3B GATESKIP WITH INDIVIDUAL VECTOR GATES

Compute Saved (%)	CSQA	HellaSwag	MMLU Stem	Open-BookQA	PIQA	Wino-Grande
0.00%	63.00±2.42	48.00±2.50	44.66±0.87	31.75±2.33	76.50±2.12	71.25±2.27
5.10%	39.50±2.45	47.00±2.50	34.09±0.84	27.00±2.22	71.00±2.27	62.50±2.42
10.12%	21.00±2.04	39.50±2.45	28.23±0.80	18.75±1.95	62.50±2.42	53.25±2.50
15.11%	23.50±2.12	37.00±2.42	27.05±0.79	14.50±1.76	60.00±2.45	51.25±2.50
20.09%	22.25±2.08	35.25±2.39	26.70±0.79	14.75±1.78	55.50±2.49	51.25±2.50
25.08%	23.25±2.11	33.25±2.36	28.35±0.80	11.50±1.60	50.25±2.50	49.75±2.50
30.09%	20.00±2.00	31.25±2.32	23.76±0.75	11.00±1.57	53.25±2.50	49.25±2.50
35.09%	18.75±1.95	28.25±2.25	20.81±0.72	10.50±1.53	51.00±2.50	46.25±2.50
40.09%	19.00±1.96	26.00±2.20	21.19±0.73	11.75±1.61	49.25±2.50	48.25±2.50
50.10%	18.50±1.94	25.00±2.17	21.95±0.74	11.00±1.57	53.25±2.50	46.75±2.50
60.10%	21.25±2.05	25.50±2.18	23.56±0.76	12.25±1.64	51.50±2.50	48.25±2.50
70.08%	19.50±1.98	25.25±2.17	24.99±0.77	13.50±1.71	53.25±2.50	49.25±2.50
80.12%	20.25±2.01	26.00±2.20	23.82±0.76	14.50±1.76	53.00±2.50	50.25±2.50
90.37%	19.00±1.96	26.25±2.20	25.88±0.77	15.50±1.81	51.00±2.50	50.50±2.50
95.41%	20.75±2.03	29.50±2.28	23.91±0.76	15.75±1.82	50.75±2.50	49.00±2.50

Table 35: Results for loglikelihood-based benchmarks for Llama-3.2-3b GateSkip with individual vector gates

Compute Saved (%)	CSQA (Gen.)	GSM8K (Gen.)
0.00%	57.25±2.48	32.75±2.35
34.78%	51.50±2.50	30.50 ± 2.30
37.28%	50.00±2.50	29.75±2.29
40.53%	12.75±1.67	5.25 ± 1.12
40.73%	5.25±1.12	1.75±0.66
40.87%	18.50±1.94	8.00 ± 1.36
41.06%	2.75 ± 0.82	0.25 ± 0.25
42.86%	27.75±2.24	17.25±1.89
45.06%	42.75±2.48	24.50±2.15

Table 36: Results for generative benchmarks for Llama-3.2-3b GateSkip with individual vector gates

K.12 LLAMA-3.2-3B GATESKIP WITH INDIVIDUAL VECTOR GATES AT 4-BIT QUANTIZATION

Compute Saved (%)	CSQA	HellaSwag	MMLU Stem	Open-BookQA	PIQA	Wino-Grande
0.00%	63.00±2.42	48.00±2.50	44.66±0.87	31.75±2.33	76.50±2.12	71.25±2.27
5.10%	39.50±2.45	47.00±2.50	34.09±0.84	27.00±2.22	71.00±2.27	62.50±2.42
10.12%	21.00±2.04	39.50±2.45	28.23±0.80	18.75±1.95	62.50±2.42	53.25±2.50
15.11%	23.50±2.12	37.00±2.42	27.05±0.79	14.50±1.76	60.00±2.45	51.25±2.50
20.09%	22.25±2.08	35.25±2.39	26.70±0.79	14.75±1.78	55.50±2.49	51.25±2.50
25.08%	23.25±2.11	33.25±2.36	28.35±0.80	11.50±1.60	50.25±2.50	49.75±2.50
30.09%	20.00±2.00	31.25±2.32	23.76±0.75	11.00±1.57	53.25±2.50	49.25±2.50
35.09%	18.75±1.95	28.25±2.25	20.81±0.72	10.50±1.53	51.00±2.50	46.25±2.50
40.09%	19.00±1.96	26.00±2.20	21.19±0.73	11.75±1.61	49.25±2.50	48.25±2.50
50.10%	18.50±1.94	25.00±2.17	21.95±0.74	11.00±1.57	53.25±2.50	46.75±2.50
60.10%	21.25±2.05	25.50±2.18	23.56±0.76	12.25±1.64	51.50±2.50	48.25±2.50
70.08%	19.50±1.98	25.25±2.17	24.99±0.77	13.50±1.71	53.25±2.50	49.25±2.50
80.12%	20.25±2.01	26.00±2.20	23.82±0.76	14.50±1.76	53.00±2.50	50.25±2.50
90.37%	19.00±1.96	26.25±2.20	25.88±0.77	15.50±1.81	51.00±2.50	50.50±2.50
95.41%	20.75±2.03	29.50±2.28	23.91±0.76	15.75±1.82	50.75±2.50	49.00±2.50

Table 37: Results for loglikelihood-based benchmarks for Llama-3.2-3b GateSkip with individual vector gates at 4-bit quantization

Compute Saved (%)	CSQA (Gen.)	GSM8K (Gen.)
0.00%	57.00±2.48	28.00±2.25
33.77%	10.75±1.55	4.25±1.01
34.54%	52.75±2.50	25.50±2.18
35.33%	55.25±2.49	25.50±2.18
37.73%	5.75±1.17	2.00 ± 0.70
37.93%	32.00±2.34	14.50±1.76
41.08%	1.50±0.61	1.50±0.61
44.10%	40.25±2.46	22.75±2.10
44.88%	16.50±1.86	8.50 ± 1.40

Table 38: Results for generative benchmarks for Llama-3.2-3b GateSkip with individual vector gates at 4-bit quantization

K.13 Llama-3.1-8b GateSkip with individual vector gates

Compute Saved (%)	CSQA	HellaSwag	MMLU Stem	Open-BookQA	PIQA	Wino-Grande
0.00%	76.50±2.12	54.25±2.49	50.17±0.86	36.25±2.41	81.50±1.94	78.00±2.07
5.10%	63.50±2.41	52.00±2.50	39.20±0.86	32.75±2.35	78.75±2.05	71.25±2.27
10.10%	45.00±2.49	50.25±2.50	32.35±0.83	31.75±2.33	74.75±2.17	64.50±2.40
15.11%	26.75±2.22	48.00±2.50	28.07±0.80	29.25±2.28	72.25±2.24	60.00±2.45
20.11%	23.50±2.12	45.50±2.49	24.36±0.76	24.25±2.15	63.00±2.42	50.00±2.50
25.10%	21.25±2.05	40.00±2.45	22.74±0.75	20.25±2.01	56.75±2.48	51.50±2.50
30.11%	20.75±2.03	33.25±2.36	23.12±0.75	17.00±1.88	58.50±2.47	53.25±2.50
35.11%	19.25±1.97	33.50±2.36	22.71±0.75	14.75±1.78	58.75±2.46	49.00±2.50
40.12%	17.25±1.89	31.25±2.32	22.11±0.74	11.75±1.61	54.50±2.49	52.50±2.50
50.11%	18.50±1.94	27.50±2.24	22.33±0.74	13.75±1.72	50.75±2.50	50.25±2.50
60.12%	18.00±1.92	27.50±2.24	23.03±0.75	13.00±1.68	47.25±2.50	45.75±2.49
70.13%	16.25±1.85	25.75±2.19	23.60±0.76	13.50±1.71	49.25±2.50	50.50±2.50
80.13%	20.75±2.03	26.50±2.21	24.39±0.76	13.25±1.70	52.25±2.50	52.75±2.50
90.54%	23.75±2.13	24.25±2.15	24.07±0.76	13.25±1.70	49.00±2.50	52.00±2.50
96.36%	20.50±2.02	24.00±2.14	22.17±0.74	13.50±1.71	48.75±2.50	49.00±2.50

Table 39: Results for loglikelihood-based benchmarks for Llama-3.1-8b GateSkip with individual vector gates

Compute Saved (%)	CSQA (Gen.)	GSM8K (Gen.)
0.00%?	60.75±2.44	53.75±2.50
31.90%?	2.25 ± 0.74	1.75±0.66
34.85%?	8.25±1.38	11.25±1.58
35.61%?	56.50±2.48	47.50±2.50
35.78%?	0.75 ± 0.43	0.25 ± 0.25
38.68%?	45.75±2.49	43.00±2.48
41.96%?	37.00 ± 2.42	31.50±2.33
45.25%?	0.25 ± 0.25	0.00 ± 0.00
50.28%?	21.25±2.05	18.75±1.95

Table 40: Results for generative benchmarks for Llama-3.1-8b GateSkip with individual vector gates

K.14 GEMMA-2-2B GATESKIP WITH INDIVIDUAL VECTOR GATES

Compute Saved (%)	CSQA	HellaSwag	MMLU Stem	Open-BookQA	PIQA	Wino-Grande
0.00%	52.50±2.50	46.75±2.50	40.63±0.86	32.75±2.35	78.25±2.07	66.25±2.37
5.09%	48.75±2.50	47.50±2.50	36.47±0.85	30.75 ± 2.31	76.75±2.11	63.00±2.42
10.13%	41.25±2.46	46.00±2.50	32.64±0.83	30.50±2.30	75.00±2.17	63.00±2.42
15.14%	26.00±2.20	44.25±2.49	31.94±0.83	24.75±2.16	71.00±2.27	65.75±2.38
20.15%	23.00±2.11	42.00±2.47	27.62±0.80	24.25±2.15	71.00±2.27	55.00±2.49
25.14%	20.75±2.03	38.75±2.44	24.04±0.76	21.50±2.06	67.50±2.34	52.25±2.50
30.09%	19.75±1.99	37.25±2.42	22.68±0.75	16.75±1.87	63.75±2.41	53.00±2.50
35.08%	18.75±1.95	35.50±2.40	21.85±0.74	18.25±1.93	60.00±2.45	54.00±2.50
40.11%	18.50±1.94	32.25±2.34	22.20±0.74	15.25±1.80	55.50±2.49	50.75±2.50
50.13%	20.50±2.02	30.50±2.30	21.76±0.73	13.25±1.70	48.50±2.50	51.25±2.50
60.15%	19.50±1.98	28.50±2.26	22.26±0.74	15.50±1.81	53.75±2.50	45.75±2.49
70.18%	19.00±1.96	28.75±2.27	22.36±0.74	15.50±1.81	55.00±2.49	50.75±2.50
80.27%	18.75±1.95	27.00±2.22	23.41±0.75	15.25±1.80	54.25±2.49	51.00±2.50
90.30%	18.75±1.95	28.50±2.26	21.22±0.73	17.00±1.88	49.75±2.50	49.25±2.50
95.34%	18.75±1.95	28.50±2.26	21.25±0.73	15.75±1.82	49.00±2.50	54.50±2.49

Table 41: Results for loglikelihood-based benchmarks for Gemma-2-2b GateSkip with individual vector gates

Compute Saved (%)	CSQA (Gen.)	GSM8K (Gen.)
0.00%	45.50±2.49	30.50±2.30
29.83%	35.50 ± 2.40	18.75±1.95
36.50%	41.75±2.47	24.75±2.16
38.40%	39.75±2.45	23.25±2.11
42.11%	35.75 ± 2.40	12.75±1.67
48.10%	31.50 ± 2.33	8.25 ± 1.38
49.12%	20.00 ± 2.00	4.75 ± 1.06
72.99%	10.00±1.50	2.50 ± 0.78
94.00%	0.00 ± 0.00	0.00 ± 0.00

Table 42: Results for generative benchmarks for Gemma-2-2b GateSkip with individual vector gates

K.15 LLAMA-3.2-1B CALM (HIDDEN STATE SATURATION)

Compute Saved (%)	CSQA	HellaSwag	MMLU Stem	Open-BookQA	PIQA	Wino-Grande
0.00%	18.75±1.95	34.75±2.38	26.32 ± 0.78	14.75±1.78	59.00±2.46	55.50±2.49
5.60%	18.75 ± 1.95	$32.25{\pm}2.34$	26.17 ± 0.78	15.25 ± 1.80	57.75 ± 2.47	53.50 ± 2.50
11.19%	18.75 ± 1.95	30.75 ± 2.31	23.18 ± 0.75	15.50 ± 1.81	58.75 ± 2.46	53.75 ± 2.50
16.74%	19.00 ± 1.96	28.75 ± 2.27	23.47 ± 0.75	15.00 ± 1.79	55.00 ± 2.49	53.00 ± 2.50
22.25%	18.50 ± 1.94	28.25 ± 2.25	22.39 ± 0.74	16.50 ± 1.86	57.00 ± 2.48	50.50 ± 2.50
27.73%	18.75 ± 1.95	29.00 ± 2.27	21.92 ± 0.74	17.50 ± 1.90	54.25 ± 2.49	49.25 ± 2.50
33.18%	18.50 ± 1.94	30.00 ± 2.29	22.04 ± 0.74	18.75 ± 1.95	55.50 ± 2.49	47.25 ± 2.50
38.58%	18.50 ± 1.94	27.50 ± 2.24	21.76 ± 0.73	18.25 ± 1.93	55.50 ± 2.49	50.00 ± 2.50
43.93%	18.50 ± 1.94	26.50 ± 2.21	21.95 ± 0.74	16.00 ± 1.84	53.00 ± 2.50	47.75 ± 2.50
54.49%	18.75 ± 1.95	28.25 ± 2.25	22.14 ± 0.74	16.75 ± 1.87	53.25 ± 2.50	47.50 ± 2.50
64.81%	18.50 ± 1.94	30.50 ± 2.30	21.60 ± 0.73	15.75 ± 1.82	52.50 ± 2.50	48.50 ± 2.50
74.80%	18.75 ± 1.95	28.50 ± 2.26	21.54 ± 0.73	16.00 ± 1.84	52.75 ± 2.50	46.25 ± 2.50
84.26%	18.75 ± 1.95	30.50 ± 2.30	21.25 ± 0.73	15.75 ± 1.82	49.75 ± 2.50	48.75 ± 2.50
92.32%	18.75 ± 1.95	27.00 ± 2.22	21.22 ± 0.73	14.00 ± 1.74	51.50 ± 2.50	47.75 ± 2.50
93.33%	18.50 ± 1.94	27.00 ± 2.22	$28.35{\pm}0.79$	13.50 ± 1.71	52.50 ± 2.50	48.50 ± 2.50

Table 43: Results for loglikelihood-based benchmarks for Llama-3.2-1b CALM (hidden state saturation)

Compute Saved (%)	CSQA (Gen.)	GSM8K (Gen.)
0.04%	0.50 ± 0.35	0.75±0.43
6.07%	0.00 ± 0.00	1.50 ± 0.61
11.60%	0.00 ± 0.00	1.00 ± 0.50
17.10%	$0.25{\pm}0.25$	0.50 ± 0.35
22.43%	0.00 ± 0.00	1.25 ± 0.56
27.95%	0.00 ± 0.00	1.00 ± 0.50
33.39%	0.00 ± 0.00	0.50 ± 0.35
38.69%	0.00 ± 0.00	1.75 ± 0.66
43.97%	$0.25{\pm}0.25$	1.75 ± 0.66

Table 44: Results for generative benchmarks for Llama-3.2-1b CALM (hidden state saturation)

K.16 LLAMA-3.2-1B CALM (SOFTMAX)

Compute Saved (%)	CSQA	HellaSwag	MMLU Stem	Open-BookQA	PIQA	Wino-Grande
	$commonsense_q a$	hellaswag	$\mathrm{mmlu}_s tem$	openbookqa	piqa	winogrande
Compute Saved (%)	•					-
0.01%	18.75±1.95	34.75±2.38	26.32±0.78	14.75±1.78	59.00±2.46	55.50±2.49
5.60%	21.50 ± 2.06	29.75 ± 2.29	27.66 ± 0.79	16.00 ± 1.84	57.75 ± 2.47	54.00 ± 2.50
11.19%	19.25 ± 1.97	30.50 ± 2.30	26.29 ± 0.78	14.50 ± 1.76	52.50 ± 2.50	52.75 ± 2.50
16.73%	19.25 ± 1.97	28.75 ± 2.27	25.37 ± 0.77	15.25 ± 1.80	53.50 ± 2.50	50.50 ± 2.50
22.23%	18.75 ± 1.95	30.00 ± 2.29	23.09 ± 0.75	17.00 ± 1.88	49.25 ± 2.50	48.50 ± 2.50
27.69%	18.75 ± 1.95	28.00 ± 2.25	22.36 ± 0.74	16.50 ± 1.86	49.25 ± 2.50	50.00 ± 2.50
33.12%	18.75 ± 1.95	28.75 ± 2.27	22.04 ± 0.74	16.25 ± 1.85	50.25 ± 2.50	49.00 ± 2.50
38.50%	18.50 ± 1.94	28.75 ± 2.27	21.95 ± 0.74	15.50 ± 1.81	53.75 ± 2.50	50.50 ± 2.50
43.83%	18.50 ± 1.94	27.75 ± 2.24	21.66 ± 0.73	17.00 ± 1.88	50.50 ± 2.50	50.75 ± 2.50
54.34%	19.00 ± 1.96	27.00 ± 2.22	21.44 ± 0.73	16.75 ± 1.87	54.50 ± 2.49	49.50 ± 2.50
64.57%	19.50 ± 1.98	28.50 ± 2.26	21.28 ± 0.73	13.75 ± 1.72	51.75 ± 2.50	50.50 ± 2.50
74.49%	18.75 ± 1.95	29.00 ± 2.27	21.31 ± 0.73	13.75 ± 1.72	51.50 ± 2.50	50.25 ± 2.50
84.15%	18.50 ± 1.94	26.50 ± 2.21	26.61 ± 0.78	15.25 ± 1.80	51.25 ± 2.50	51.75 ± 2.50
93.89%	18.50 ± 1.94	27.00 ± 2.22	28.48 ± 0.80	15.25 ± 1.80	53.50 ± 2.50	48.25 ± 2.50
99.94%	18.50 ± 1.94	$26.25{\pm}2.20$	28.51 ± 0.80	11.50 ± 1.60	53.50 ± 2.50	$49.25{\pm}2.50$

Table 45: Results for loglikelihood-based benchmarks for Llama-3.2-1b CALM (softmax)

Compute Saved (%)	CSQA (Gen.)	GSM8K (Gen.)
0.24%	0.75±0.43	0.75±0.43
5.92%	$0.25{\pm}0.25$	1.50 ± 0.61
11.21%	0.00 ± 0.00	2.00 ± 0.70
16.60%	0.00 ± 0.00	$0.25{\pm}0.25$
21.92%	0.00 ± 0.00	0.50 ± 0.35
27.28%	0.00 ± 0.00	1.00 ± 0.50
32.56%	0.00 ± 0.00	$0.25{\pm}0.25$
37.85%	0.25 ± 0.25	0.50 ± 0.35
42.98%	0.00 ± 0.00	$0.25{\pm}0.25$

Table 46: Results for generative benchmarks for Llama-3.2-1b CALM (softmax)

K.17 LLAMA-3.2-1B MIXTURE-OF-DEPTHS

Compute Saved (%)	CSQA	HellaSwag	MMLU Stem	Open-BookQA	PIQA	Wino-Grande
0.00%	18.75±1.95	34.75±2.38	26.32±0.78	14.75±1.78	59.00±2.46	55.50±2.49
5.60%	18.75 ± 1.95	32.25 ± 2.34	26.17 ± 0.78	15.25 ± 1.80	57.75 ± 2.47	53.50 ± 2.50
11.19%	18.75 ± 1.95	30.75 ± 2.31	23.18 ± 0.75	15.50 ± 1.81	58.75 ± 2.46	53.75 ± 2.50
16.74%	19.00 ± 1.96	28.75 ± 2.27	23.47 ± 0.75	15.00 ± 1.79	55.00 ± 2.49	53.00 ± 2.50
22.25%	18.50 ± 1.94	28.25 ± 2.25	22.39 ± 0.74	16.50 ± 1.86	57.00 ± 2.48	50.50 ± 2.50
27.73%	18.75 ± 1.95	29.00 ± 2.27	21.92 ± 0.74	17.50 ± 1.90	54.25 ± 2.49	49.25 ± 2.50
33.18%	18.50 ± 1.94	30.00 ± 2.29	22.04 ± 0.74	18.75 ± 1.95	55.50 ± 2.49	47.25 ± 2.50
38.58%	18.50 ± 1.94	27.50 ± 2.24	21.76 ± 0.73	18.25 ± 1.93	55.50 ± 2.49	50.00 ± 2.50
43.93%	18.50 ± 1.94	26.50 ± 2.21	21.95 ± 0.74	16.00 ± 1.84	53.00 ± 2.50	47.75 ± 2.50
54.49%	18.75 ± 1.95	28.25 ± 2.25	22.14 ± 0.74	16.75 ± 1.87	53.25 ± 2.50	47.50 ± 2.50
64.81%	18.50 ± 1.94	30.50 ± 2.30	21.60 ± 0.73	15.75 ± 1.82	52.50 ± 2.50	48.50 ± 2.50
74.80%	18.75 ± 1.95	28.50 ± 2.26	21.54 ± 0.73	16.00 ± 1.84	52.75 ± 2.50	46.25 ± 2.50
84.26%	18.75 ± 1.95	30.50 ± 2.30	21.25 ± 0.73	15.75 ± 1.82	49.75 ± 2.50	48.75 ± 2.50
92.32%	18.75 ± 1.95	27.00 ± 2.22	21.22 ± 0.73	14.00 ± 1.74	51.50 ± 2.50	47.75 ± 2.50
93.33%	18.50 ± 1.94	27.00 ± 2.22	28.35 ± 0.79	13.50 ± 1.71	52.50 ± 2.50	48.50 ± 2.50

Table 47: Results for loglikelihood-based benchmarks for Llama-3.2-1b Mixture-of-Depths (router-tuned)

Compute Saved (%)	CSQA (Gen.)	GSM8K (Gen.)
0.00%	28.50 ± 2.26	11.75±1.61
10.38%	21.50 ± 2.06	7.00 ± 1.28
11.99%	10.75 ± 1.55	6.75 ± 1.26
26.44%	1.50 ± 0.61	5.50 ± 1.14
42.05%	$0.25{\pm}0.25$	5.75 ± 1.17
48.53%	0.00 ± 0.00	1.00 ± 0.50
55.76%	0.00 ± 0.00	0.75 ± 0.43
57.29%	0.00 ± 0.00	$0.25{\pm}0.25$
63.78%	0.00 ± 0.00	0.00 ± 0.00
64.52%	0.00 ± 0.00	0.00 ± 0.00
73.56%	0.00 ± 0.00	$0.25{\pm}0.25$
81.83%	0.00 ± 0.00	0.00 ± 0.00
99.93%	0.00 ± 0.00	0.00 ± 0.00
99.99%	0.00 ± 0.00	0.00 ± 0.00
100.00%	0.00 ± 0.00	0.00 ± 0.00

Table 48: Results for generative benchmarks for Llama-3.2-1b Mixture-of-Depths (router-tuned)

K.18 Llama-3.2-1b random skipping baseline on the WMT16-English-Romanian Dataset

Compute Saved (%)	WMT16-EN-RO
0.00%	56.78 ± 2.66
5.00%	35.61 ± 2.28
10.01%	13.90 ± 1.29
15.02%	2.88 ± 0.53
20.01%	0.53 ± 0.11
25.00%	0.62 ± 0.12
30.01%	0.64 ± 0.15
35.03%	0.94 ± 0.20
40.02%	0.91 ± 0.22

Table 49: Results for Llama-3.2-1b GateSkip with individual vector gates on the WMT16-English-Romanian Dataset

K.19 $\;\;$ Llama-3.2-1b GateSkip with individual vector gates on the WMT16-English-Romanian Dataset

Compute Saved (%)	WMT16-EN-RO
0.00%	50.88 ± 2.47
6.34%	40.44 ± 2.37
11.24%	36.36 ± 2.15
16.13%	31.04 ± 1.98
21.05%	26.93 ± 1.61
25.11%	20.07 ± 1.47
30.11%	13.41 ± 1.19
36.24%	$8.86{\pm}1.35$
41.14%	4.14 ± 0.85

Table 50: Results for Llama-3.2-1b GateSkip with individual vector gates on the WMT16-English-Romanian Dataset

L HYPERPARAMETERS USED

Group	Parameter	Value / Setting
Optimiser	AdamW β_1, β_2	0.9 / 0.999
	ϵ	1e - 8
	Weight decay	0.001
	LR schedule	Cosine, 1 000 warm-up steps
GateSkip	Sparsity weight λ	0.1
	Token-budget decay	$100 \% \to 80 \% \text{ (linear)}$
Training	Batch size	1 sequence (length 4096)
	# training steps/time	15 493 for CSQA-GSM8K reasoning data, 1h for FineWeb
	Gradient clip	1.0
	Precision	FP32
Bootstrap	Iterations	100 000
Hardware	GPU	NVIDIA H100 80 GB PCIe
	Runtime / run	\sim 5 h

Table 51: Full set of hyper-parameters and environment details used in all reported experiments.

If a parameter is not listed, the default value from the HuggingFace Transformers or PyTorch implementation is used.

M LIBRARIES USED

Library / Toolkit	Version used	License	Homepage / Repo
PyTorch	2.7.0	BSD-style	https://pytorch.org
PyTorch Lightning	2.2.0	Apache 2.5.1	https://github.com/Lightning-AI/lightning
Transformers	4.51.3	Apache 2.0	https://github.com/huggingface/transformers
Datasets	3.5.1	Apache 2.0	https://github.com/huggingface/datasets
Accelerate	1.6.0	Apache 2.0	https://github.com/huggingface/accelerate
LM Eval Harness	0.4.8	MIT	https://github.com/EleutherAI/lm-evaluation-harness

Table 52: Software stack used for all experiments.