arXiv:2611.34667v1 [cs.LG] 25 Nov 2026

LEARNING TO BE TRUNCATED: PREPARING LLM CHECKPOINTS FOR SVD COMPRESSION

Chenyue Fang
EcoCloud, EPFL
chenyue.fang@epfl.ch

Danila Mishin
EcoCloud, EPFL
danila.mishin@epfl.ch

Babak Falsafi
EcoCloud, EPFL
babak.falsafi@epfl.ch

ABSTRACT

Compression based on singular value decomposition (SVD) often requires subsequent fine-tuning or weight correction to mitigate performance degradation. Budget-specific optimization must be repeated for different compression ratios. We introduce LLift, which prepares a pretrained model for SVD compression across ratios without further fine-tuning. LLift fine-tunes a low-rank approximation at a single ratio using the original model as a teacher. We merge the learned updates into the original weights while preserving the complementary components excluded from the training forward pass. We apply SVD to the prepared model's target matrices and reuse the decompositions at different truncation ranks. On LLaMA-7B, we demonstrate accuracy gains across three compression ratios using a single prepared model. The gains extend across model families, model sizes, and SVD compression operators. We analyze gradients and an early training update to provide local evidence that updates learned at one ratio can benefit compression at other ratios. These results show that LLift can improve compression across deployment budgets without fine-tuning each compressed model separately.

1 INTRODUCTION

Deploying large language models across hardware with different memory, latency, and compute budgets motivates compression methods that produce models of different sizes from a single pretrained checkpoint. Singular value decomposition (SVD) supports post hoc compression by replacing weight matrices with low-rank factors of adjustable rank. These models use standard dense matrix operations (Wang et al., 2025b). However, truncation can degrade the model's predictions. Existing methods reduce this degradation through activation-aware reconstruction, rank allocation, and error correction (Yuan et al., 2023; Wang et al., 2025a;b; Hu et al., 2026; Abbasi et al., 2026).

Figure 1 compares different methods for producing compressed models at multiple compression ratios. Post-compression fine-tuning and budget-specific reconstruction improve individual compressed models, but require repeating optimization for each target. Elastic methods such as FlexRank and ACIP amortize training across budgets using shared low-rank factors or a learned component ordering (Zaccone et al., 2026; Genzel et al., 2025). These methods instantiate deployment models from the learned representations. We study a complementary goal: updating the original model, then truncating shared SVDs of the prepared model's target matrices at different ranks.

Training with future compression in mind provides a foundation for this goal. Compression-aware training encourages low-rank matrices (Alvarez & Salzmann, 2017), while DecomposableNet jointly trains full- and low-rank networks (Yaguchi et al., 2021). More recently, Prehab prepares the original model by optimizing a task loss together with spectral regularization (Qin et al., 2025), encouraging weight spectra that are more favorable to truncation. Our approach instead uses prediction errors exposed by a low-rank model to guide updates to the original model.

We introduce LLift, which fine-tunes a low-rank approximation of the original model at one preparation budget. This approximation is the low-rank view: its forward pass uses only the retained components, while the original model provides supervision. After fine-tuning, we merge the learned update into the original weights and preserve the components excluded from the training forward pass. With fixed activation transforms, we apply SVD once per target matrix of the prepared model and select different truncation ranks from the shared decompositions. This separates the budget used to expose compression error during training from the budget selected for deployment.

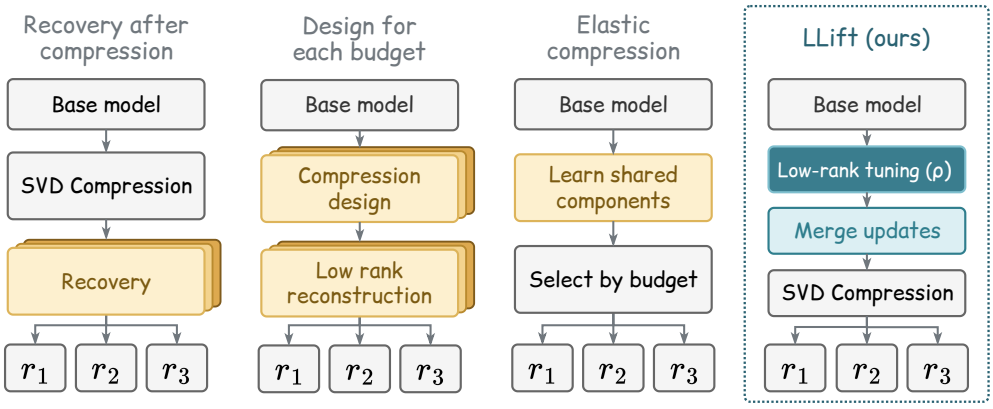


Figure 1: Compression workflows across budgets. Post-compression fine-tuning and budget-specific reconstruction repeat optimization for each budget, while elastic methods select from learned shared components. LLift fine-tunes a low-rank approximation at one preparation budget ρ and merges the learned updates into the original weights to prepare one model for compression across budgets. Stacked boxes indicate budget-specific work; r_1, r_2, r_3 denote deployment budgets.

The challenge is to make updates learned at one budget useful when we compress the prepared model at another. Exporting only the fine-tuned student would discard components absent from the training view, including information that larger deployment budgets could retain. LLift preserves these complementary components while incorporating the learned correction into the original weights. Recompressing the prepared model can then draw on information outside the preparation subspace.

We show that preparing each model at one budget improves compression quality at other evaluated budgets, with gains across model families, sizes, and SVD operators. We analyze gradients at initialization and identify shared descent directions as local evidence of cross-budget compatibility.

Our contributions are threefold:

- We introduce LLift, a model-preparation method that uses supervision at one compression budget. We truncate shared SVDs of the prepared model's target matrices to obtain models with competitive performance across deployment budgets without further fine-tuning.
- We show that updates learned by the low-rank student remain useful after merging the updates and applying SVD to the prepared model. Preserving the original complementary weights improves recompressed accuracy over exporting the fine-tuned student alone at deployment.
- We evaluate model preparation across model families, sizes, budgets, and SVD compression operators, and analyze local cross-budget compatibility through gradient alignment at initialization and recompression after the first nonzero AdamW update to the original model.

2 Related Work

SVD-based LLM compression. Post hoc SVD compression has evolved from weight-space decomposition toward activation-aware reconstruction, rank selection, and error correction. Existing methods use activation-derived scaling (Yuan et al., 2023), input-space whitening (Wang et al., 2025b), differentiable rank optimization (Wang et al., 2025a; Sundrani et al., 2025), cumulative-error suppression (Hu et al., 2026), loss-aware singular-component selection (Abbasi et al., 2026), activation-aware decomposition with dynamic rank allocation (Qi et al., 2026), Fisher-weighted decomposition with Kronecker-factored approximations (Chekalina et al., 2026), and token-wise Fisher statistics with rank allocation using a calibrated error model and integer linear programming (Thoma et al., 2026). Low-Rank Prehab encourages a more compressible weight spectrum before applying SVD (Qin et al., 2025), while NuMuon promotes low-rank structure during pretraining through constrained optimizer updates (Dolatabadi et al., 2026). PiSSA (Meng et al., 2024) adapts principal singular components for parameter-efficient fine-tuning, retaining a frozen residual in its forward pass. LLift prepares a pretrained model by supervising its end-to-end predictions

through an activation-aware low-rank view. Its objective targets compressed-model behavior without explicitly regularizing the weight spectrum, and we evaluate the prepared model with SVD-LLM, ZS-SVD, and KFAC-SVD with LEMS rank allocation at deployment.

Multi-budget and elastic low-rank deployment. Nested and elastic models support multiple deployment budgets through shared representations or subnetworks, with smaller configurations contained in larger ones (Kusupati et al., 2022; Devvrit et al., 2024; Cai et al., 2024; Taghibakhshi et al., 2025). For low-rank deployment, selecting subsets of shared components adjusts model size. FlexRank combines shared low-rank factors with budget-specific rank configurations and trains the resulting nested models across budgets (Zaccone et al., 2026). ACIP learns a global importance ordering over SVD components together with a shared low-rank correction, allowing deployment at different model sizes by truncating the ordering (Genzel et al., 2025). Both methods organize multi-budget deployment around a shared low-rank representation or component ordering and use these representations to instantiate deployment models directly at different budgets.

LLift instead prepares the original model through supervision of a low-rank view at a single budget and merges the learned updates into the original weights. At deployment, we obtain different model sizes by selecting truncation ranks from shared SVDs of the prepared model's target matrices. We test whether this preparation improves performance at other compression budgets.

3 PRELIMINARIES

Activation-aware low-rank approximation. For a linear layer with weight $W \in \mathbb{R}^{d_{\text{out}} \times d_{\text{in}}}$ and calibration activations $X \in \mathbb{R}^{d_{\text{in}} \times N}$, activation-aware compression seeks a rank-k weight approximation that best preserves the original layer outputs on these calibration activations:

$$\min_{\text{rank}(\widehat{W}) \le k} \left\| (W - \widehat{W})X \right\|_F^2 . \tag{1}$$

Following the activation-whitening formulation of SVD-LLM (Wang et al., 2025b), we use $XX^\top$ to summarize the calibration activations' second-order statistics. Adding ϵI, with $\epsilon > 0$, makes this matrix positive definite. We can therefore obtain an invertible factor T satisfying $TT^\top = XX^\top + \epsilon I$. We can compute T by Cholesky decomposition and obtain the following identity:

$$\left\| (W - \widehat{W})T \right\|_F^2 = \left\| (W - \widehat{W})X \right\|_F^2 + \epsilon \left\| W - \widehat{W} \right\|_F^2 . \tag{2}$$

Approximating WT minimizes the output reconstruction error together with a weight reconstruction penalty. We decompose $WT = U\Sigma V^\top$, with singular values in descending order. Since T^{-1} preserves rank, we solve the damped objective by truncating WT and multiplying by T^{-1}:

$$\widehat{W}^{(k)} = U_{:k}\Sigma_{:k}V_{:k}^\top T^{-1}, \tag{3}$$

Here $U_{:k}$ and $V_{:k}$ contain the first k columns of U and V, respectively. We retain the leading $k \times k$ diagonal block of Σ in $\Sigma_{:k}$, corresponding to the k largest singular values of WT.

Nested rank truncation. For a fixed checkpoint and activation transform, the components retained at rank k_1 are a subset of those retained at rank k_2 whenever $k_1 < k_2$. A single SVD can therefore produce multiple model sizes by retaining different numbers of leading components. This nesting permits reuse of the decomposition but does not by itself guarantee task performance at unseen budgets. LLift uses activation-aware SVD to initialize its low-rank training view. At deployment, we apply SVD to the prepared model's transformed target matrices and select different truncation ranks from these shared decompositions to obtain models at each deployment budget.

4 METHOD

LLift prepares a model for SVD compression across deployment budgets. We fine-tune a low-rank student at one preparation anchor, using fixed subspaces derived from the original weights and activation statistics. We merge these updates into the original weights, preserving the complementary components excluded from the student forward pass. For deployment, we truncate shared SVDs of the prepared model's transformed target matrices at the ranks specified by each budget (Figure 2).

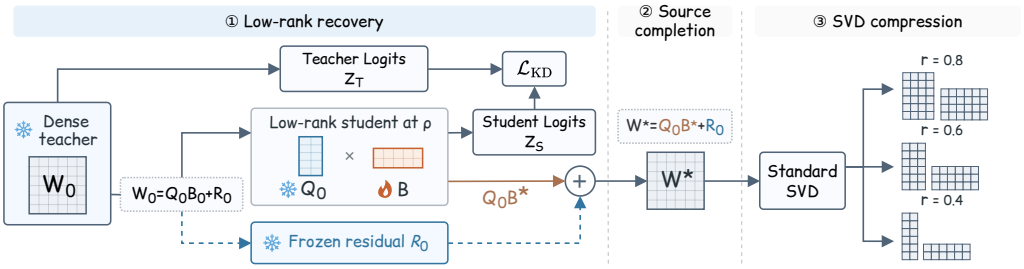


Figure 2: Fine-tuning a low-rank view and merging its updates into the original model. LLift uses the frozen original model to supervise a student whose target layer uses Q_0B at preparation anchor ρ, with Q_0 fixed and B trainable. Snowflakes and flames denote frozen and trainable quantities. After training, we merge the update $Q_0(B^\star - B_0)$ into W_0, yielding $W^\star = Q_0B^\star + R_0$, where $R_0 = W_0 - Q_0B_0$ is the frozen complementary component that bypasses the training forward pass. With a fixed activation transform T, we decompose $W^\star T$ once and truncate the shared decomposition to the ranks specified by r. The diagram shows one target matrix; the prepared model also retains the embeddings, output head, and normalization parameters adapted during student fine-tuning.

4.1 Learning through a Low-Rank View

Activation-aware SVD constructs each compressed layer from local weight and activation statistics, while the quality of the resulting model depends on the joint effect of compression across layers. LLift therefore exposes all target layers to their low-rank views simultaneously and uses the pretrained model to supervise the resulting end-to-end behavior. We use one preparation anchor to define the training view and select deployment ranks when compressing the prepared model.

Let $f(x;\Theta_0)$ denote the original model's logits for input x, with target matrices $\{W_{\ell,0}\}_{\ell\in\mathcal{S}}$ and remaining parameters Φ_0. We use ρ for preparation and r for deployment; both specify nominal retained-parameter fractions in the target matrices. We assign ρ uniformly across target matrices. A rank-k_ℓ factorization stores $k_\ell(d_{\text{out},\ell}+d_{\text{in},\ell})$ parameters instead of $d_{\text{out},\ell}d_{\text{in},\ell}$, so we choose

$$k_\ell(\rho) = \max\left\{1, \left\lfloor \frac{\rho\, d_{\text{out},\ell}d_{\text{in},\ell}}{d_{\text{out},\ell}+d_{\text{in},\ell}} \right\rfloor\right\}. \tag{4}$$

The rank $k_\ell(\rho)$ determines the dimensions of the two factors used in the training forward pass. Uniform parameter retention, subject to rank rounding, does not imply the same retained fraction of singular directions across target matrices with different input and output dimensions.

The initial view preserves the leading activation-aware components of the pretrained weights. We estimate T_ℓ^{train} from the original model on calibration text and decompose $W_{\ell,0}T_\ell^{\text{train}} = U_\ell\Sigma_\ell V_\ell^\top$ as in Eq. 3. We use the retained left singular vectors as a fixed basis and project the weights:

$$Q_{\ell,0} = U_{\ell,:k_\ell(\rho)}, \qquad B_{\ell,0} = Q_{\ell,0}^\top W_{\ell,0}. \tag{5}$$

Here $Q_{\ell,0}\in\mathbb{R}^{d_{\text{out},\ell}\times k_\ell(\rho)}$ and $B_{\ell,0}\in\mathbb{R}^{k_\ell(\rho)\times d_{\text{in},\ell}}$. The product $Q_{\ell,0}B_{\ell,0} = Q_{\ell,0}Q_{\ell,0}^\top W_{\ell,0}$ is the rank-$k_\ell(\rho)$ projection induced by the activation-aware decomposition in Eq. 3. We associate each coordinate matrix B_ℓ with a corresponding model weight parameterized as

$$W_\ell(B_\ell) = W_{\ell,0} + Q_{\ell,0}(B_\ell - B_{\ell,0}).$$

Since $Q_{\ell,0}^\top Q_{\ell,0} = I$, its projection onto the fixed column space of $Q_{\ell,0}$ satisfies $Q_{\ell,0}Q_{\ell,0}^\top W_\ell(B_\ell) = Q_{\ell,0}B_\ell$. The student evaluates only the projected map $h\mapsto Q_{\ell,0}(B_\ell h)$. The complementary component of the original weight remains unchanged and does not participate in this forward pass. Preparation therefore does not require optimizing full target matrices. We keep $Q_{\ell,0}$ fixed.

The original model serves as a frozen teacher. Let $z_t^{\text{T}} = f(x;\Theta_0)_t$ denote its logits at position t and $z_t^{\text{S}}(\mathcal{B},\Phi)$ those of the student with target weights $\{Q_{\ell,0}B_\ell\}_{\ell\in\mathcal{S}}$, where $\mathcal{B}=\{B_\ell\}_{\ell\in\mathcal{S}}$ and Φ collects other student parameters. We minimize the token-level student-to-teacher KL loss:

$$\mathcal{L}_{\text{LLift}}(\mathcal{B},\Phi) = \mathbb{E}_{x\sim\mathcal{D}_{\text{train}}}\left[\frac{1}{|\Omega(x)|}\sum_{t\in\Omega(x)}\text{KL}\big(p_\tau(z_t^{\text{S}}(\mathcal{B},\Phi))\,\big\|\,p_\tau(z_t^{\text{T}})\big)\right], \tag{6}$$

where $p_\tau(z) = \mathrm{softmax}(z/\tau)$ and $\Omega(x)$ excludes padded positions. Supervision at the final logits accounts for the combined effect of compressing all target layers. During training, we omit the original complementary components $W_{\ell,0} - Q_{\ell,0}B_{\ell,0}$ from the student forward pass. This exposes prediction errors caused by compression and provides a learning signal for adapting the retained coordinates. We initialize $(\mathcal{B}, \Phi) \leftarrow (\mathcal{B}_0, \Phi_0)$ and jointly optimize these parameters, including the embeddings, output head, and normalization parameters. We keep the teacher and bases $Q_{\ell,0}$ fixed.

4.2 Constructing the Prepared Model

We update the original weights by merging the coordinate change $B^\star_\ell - B_{\ell,0}$ into $W_{\ell,0}$ through $Q_{\ell,0}$, preserving the original components that the training forward pass excludes. We can equivalently combine the fine-tuned low-rank factors with the original truncation residual.

We retain the original truncation residual

$$R_{\ell,0} = W_{\ell,0} - Q_{\ell,0}B_{\ell,0}, \qquad Q_{\ell,0}^\top R_{\ell,0} = 0. \tag{7}$$

We write each target weight in the prepared model in two equivalent forms:

$$\begin{aligned} W^\star_\ell &= W_{\ell,0} + Q_{\ell,0}(B^\star_\ell - B_{\ell,0}) \\ &= Q_{\ell,0}B^\star_\ell + R_{\ell,0}. \end{aligned} \tag{8}$$

Reattaching $R_{\ell,0}$ yields the minimum-change completion of the fine-tuned coordinates:

$$W^\star_\ell = \underset{W:\, Q_{\ell,0}^\top W = B^\star_\ell}{\arg\min} \|W - W_{\ell,0}\|_F^2. \tag{9}$$

The constraint fixes the trained component, and the nearest completion preserves the original complement. We prove this minimum-distance property in Appendix A.1; it does not establish task-loss optimality. With no adaptation, Eq. 8 recovers $W_{\ell,0}$. After adaptation, projection onto the fixed basis still recovers the trained view: $Q_{\ell,0}Q_{\ell,0}^\top W^\star_\ell = Q_{\ell,0}B^\star_\ell$.

The fixed $R_{\ell,0}$ requires no gradients or optimizer states. We reattach it only when exporting the prepared model. The target matrices require optimizer states only for B_ℓ. The exported model is $\Theta^\star = (\{W^\star_\ell\}_{\ell\in\mathcal{S}}, \Phi^\star)$, including the adapted remaining parameters. In Appendix A, we derive when factorized and full-matrix parameterizations induce identical updates to the target weights.

4.3 Recompression beyond the Training Budget

The intended benefit of preparation is improved quality after recompressing the prepared model, including at budgets not seen during preparation. Deployment compression acts on the completed weights $W^\star_\ell = Q_{\ell,0}B^\star_\ell + R_{\ell,0}$, so both the learned update and the preserved complementary component can influence the singular directions retained by SVD at each deployment budget.

The orthogonality condition $Q_{\ell,0}^\top R_{\ell,0} = 0$ refers to the fixed preparation basis. For a deployment activation transform T, it does not imply $B^\star_\ell TT^\top R_{\ell,0}^\top = 0$. The cross terms in $(W^\star_\ell T)(W^\star_\ell T)^\top$ can therefore couple the fine-tuned component to the preserved complement. Applying SVD to $W^\star_\ell T$ can select left singular directions that differ from the preparation basis, allowing the complement to affect compression even when $r \le \rho$. At $r = \rho$, the recompressed model consequently need not reproduce the predictions of the low-rank student trained at that same preparation budget.

Restoring $R_{\ell,0}$ allows deployment SVD to retain components absent from $Q_{\ell,0}B^\star_\ell$. We show that this can reduce reconstruction error relative to the original teacher weight $W_{\ell,0}$ after compression at the same rank, and derive sufficient conditions for improvement at any deployment rank. Above the preparation rank, the restored residual supplies additional directions that the trained low-rank weight alone cannot provide. Appendix A.2 gives the derivations.

We hold the deployment activation transforms T_ℓ^{deploy} fixed across budgets and apply SVD once to each $W^\star_\ell T_\ell^{\mathrm{deploy}}$. For budget r, we retain the first $k_\ell(r)$ components of these shared decompositions and construct the layer factors as in Eq. 3. Let $\mathcal{C}_r(\Theta; \mathcal{D}_{\mathrm{cal}}^{\mathrm{deploy}})$ denote compression with these fixed transforms. Applying this operator to the prepared model yields the deployment models

$$\Theta_r = \mathcal{C}_r\left(\Theta^\star; \mathcal{D}_{\mathrm{cal}}^{\mathrm{deploy}}\right), \qquad r \in \mathcal{R}. \tag{10}$$

Each budget requires its own truncation ranks and factor materialization, while reusing the SVD of each target matrix without additional fine-tuning. Section 5.3 examines cross-budget transfer through the alignment of the ratio-specific fine-tuning gradients at initialization.

5 EXPERIMENTS

We test whether preparation at one anchor improves compression across budgets, examine residual restoration, and assess generality across SVD operators, model families, and scales.

5.1 EXPERIMENTAL SETUP

Preparation. LLaMA-7B is our reference model. Unless otherwise stated, we prepare each model once at $\rho = 0.6$ on the WikiText-2 (WT2) (Merity et al., 2017) training split using KL distillation from the original model. We construct the fixed preparation bases from original-model activations on WT2. We use student-to-teacher KL as the default objective and report CE and KL+CE controls in Appendix B.3. In particular, removing KL supervision and training with CE alone substantially increases post-compression perplexity, especially at smaller deployment budgets.

Deployment. We use SVD-LLM to obtain deployment models at $r \in \{0.8, 0.6, 0.4\}$ without further fine-tuning. For matched comparisons, we estimate deployment activation transforms from the original model on a separate WT2/Alpaca (Taori et al., 2023) calibration set and hold them fixed across models and budgets. We reuse each model's SVDs across all three budgets by varying the truncation ranks. Re-estimating these transforms from the prepared model changes accuracy only modestly and provides no consistent perplexity benefit (Appendix B.4). Following Eq. 4, r denotes the nominal ratio of factor parameters to original target-matrix parameters.

Evaluation. We report held-out perplexity on WT2, Penn Treebank (PTB) (Marcus et al., 1993), and C4 (Raffel et al., 2020), together with the equally weighted mean of plain zero-shot accuracy (Avg.) on seven tasks: OpenBookQA, ARC-Easy, ARC-Challenge, WinoGrande, HellaSwag, PIQA, and MathQA (Mihaylov et al., 2018; Clark et al., 2018; Sakaguchi et al., 2019; Zellers et al., 2019; Bisk et al., 2019; Amini et al., 2019).

5.2 QUALITY ACROSS DEPLOYMENT BUDGETS

Comparison with SVD compression methods. Table 1 compares LLift with ASVD (Yuan et al., 2023), SVD-LLM (Wang et al., 2025b), Dobi-SVD (Wang et al., 2025a), SAES-SVD (Hu et al., 2026), ZS-SVD (Abbasi et al., 2026), Swift-SVD* (Qi et al., 2026), and LEMS (Thoma et al., 2026), which denotes KFAC-SVD with LEMS rank allocation. LLift raises mean accuracy from the reported SVD-LLM scores of $0.44/0.37/0.31$ to $0.50/0.46/0.38$ at $r = 0.8/0.6/0.4$. At $r = 0.6/0.4$, it achieves the highest mean accuracy and lowest perplexity on all three corpora among the methods in the table. LEMS achieves $0.51/0.43/0.34$ accuracy at $r = 0.8/0.6/0.4$. It outperforms LLift at $r = 0.8$ on mean accuracy and all three perplexities, while LLift performs better on these metrics at $r = 0.6/0.4$.

Comparison with elastic methods. Table 2 compares LLift, FlexRank, and ACIP with and without its rank-32 correction Δ. ACIP without Δ obtains $0.49/0.40/0.33$ accuracy at $r = 0.8/0.6/0.4$, versus $0.50/0.46/0.38$ for LLift. FlexRank obtains accuracies of $0.49/0.46/0.41$ at these budgets, matching LLift in mean accuracy at $r = 0.6$ and outperforming it at $r = 0.4$. FlexRank's preparation takes 27.0 A100 GPU-hours, mostly for per-layer rank profiling, compared with 2.07 for LLift (Appendix C.1).

Budget-specific fine-tuning. We compare LLift with budget-specific SVD-LLM recovery using the same WT2 training data, KL objective, optimizer settings, and deployment calibration. Each recovery model receives the training exposure of one LLift preparation run. On LLaMA-7B and Qwen3-8B, LLift achieves broadly comparable seven-task mean accuracy at $r = 0.8/0.6$, while recovery yields lower WT2 and C4 perplexity at all three budgets. At $r = 0.4$, recovery also achieves higher mean accuracy: 0.40 versus 0.38 on LLaMA-7B and 0.39 versus 0.34 on Qwen3-8B

Table 1: SVD compression on LLaMA-7B. Baselines use official calibration settings. WT2, PTB, and C4 report perplexity (↓). Avg. is seven-task mean plain accuracy (↑). Bold and underlining mark the best and second-best in these columns per budget. †: budget-specific optimization. LLift uses one WT2 preparation run shared across deployment budgets. Appendix B.7 compares budget-specific recovery under matched training conditions.

r	Method	WT2↓	PTB↓	C4↓	OBQA↑	ARC-e↑	ARC-c↑	WinoG↑	HellaS↑	PIQA↑	MathQA↑	Avg.↑
Uncompressed		5.68	8.35	7.34	0.31	0.75	0.42	0.69	0.56	0.78	0.26	0.54
0.8	ASVD	11.14	16.55	15.93	0.25	0.53	0.27	0.64	0.41	0.68	0.24	0.43
	SVD-LLM	7.94	16.22	15.84	0.22	0.58	0.29	0.63	0.43	0.69	0.24	0.44
	Dobi-SVD†	8.54	14.83	10.01	0.26	0.59	0.31	0.66	0.44	0.70	0.23	0.46
	SAES-SVD†	7.17	15.16	13.77	0.29	0.68	0.36	0.65	0.45	0.75	0.25	0.49
	ZS-SVD	6.74	11.87	10.74	0.31	0.70	0.37	0.68	0.49	0.74	0.24	0.50
	ZS-SVD 5×†	6.43	11.17	10.42	0.32	0.72	0.39	0.67	0.50	0.75	0.25	**0.51**
	Swift-SVD*	11.64	14.74	11.76	0.24	0.63	0.31	0.64	0.44	0.71	0.23	0.46
	LEMS	**6.23**	**10.27**	**9.15**	0.29	0.73	0.38	0.67	0.52	0.76	0.24	**0.51**
	LLift	7.49	11.72	11.15	0.30	0.67	0.37	0.66	0.52	0.74	0.25	0.50
0.6	ASVD	1407	3292	1109	0.13	0.28	0.22	0.48	0.26	0.55	0.19	0.30
	SVD-LLM	13.11	63.75	49.83	0.19	0.42	0.25	0.58	0.33	0.60	0.21	0.37
	Dobi-SVD†	13.54	46.38	23.54	0.22	0.41	0.27	0.58	0.34	0.61	0.23	0.38
	SAES-SVD†	10.42	45.13	32.79	0.23	0.50	0.29	0.62	0.36	0.65	0.23	0.41
	ZS-SVD	11.44	43.19	34.13	0.23	0.52	0.26	0.62	0.35	0.64	0.22	0.41
	ZS-SVD 5×†	9.45	36.52	26.20	0.23	0.56	0.29	0.62	0.38	0.66	0.22	0.42
	Swift-SVD*	36.87	56.15	25.20	0.16	0.40	0.22	0.54	0.30	0.61	0.22	0.35
	LEMS	9.22	21.46	19.29	0.22	0.60	0.29	0.61	0.39	0.67	0.22	0.43
	LLift	**8.40**	**15.96**	**14.77**	0.25	0.63	0.31	0.62	0.45	0.71	0.25	**0.46**
0.4	ASVD	57057	45218	43036	0.12	0.26	0.21	0.49	0.26	0.53	0.18	0.29
	SVD-LLM	53.74	438.58	345.49	0.14	0.28	0.22	0.50	0.27	0.55	0.21	0.31
	Dobi-SVD†	46.18	238.91	190.62	0.15	0.31	0.20	0.52	0.28	0.54	0.22	0.32
	SAES-SVD†	22.01	116.83	93.97	0.16	0.33	0.25	0.52	0.30	0.54	0.23	0.34
	ZS-SVD	45.17	334.85	212.57	0.15	0.31	0.21	0.54	0.28	0.55	0.21	0.32
	ZS-SVD 5×†	20.41	163.58	88.56	0.17	0.37	0.22	0.55	0.30	0.57	0.21	0.34
	Swift-SVD*	257.70	400.56	143.85	0.13	0.27	0.21	0.49	0.26	0.55	0.21	0.30
	LEMS	21.93	104.05	83.94	0.16	0.37	0.22	0.55	0.30	0.59	0.22	0.34
	LLift	**13.48**	**53.47**	**32.24**	0.17	0.48	0.24	0.57	0.34	0.62	0.22	**0.38**

Table 2: Elastic methods. Seven-task average accuracy (↑). Δ is a rank-32 correction whose parameters are excluded from r.

Method	$r = 0.8$	$r = 0.6$	$r = 0.4$
ACIP	0.49	0.40	0.33
ACIP ($+\Delta$)	0.48	0.45	0.38
FlexRank	0.49	0.46	0.41
LLift	0.50	0.46	0.38

Table 3: Model preparation and residual restoration under matched `mix75` calibration. Seven-task average accuracy (↑).

Model	$r = 0.8$	$r = 0.6$	$r = 0.4$
Original	0.44	0.40	0.31
Spectral regularization	0.48	0.41	0.33
Fine-tuned $Q_0 B^\star$ only	0.44	0.44	0.37
LLift: $Q_0 B^\star + R_0$	0.50	0.46	0.38

(Table 16). For the three deployment budgets, LLift incurs a total pipeline cost of 3.13 A100 GPU-hours on LLaMA-7B, compared with 5.10 for separate recovery (Table 24). These results highlight the trade-off: LLift amortizes preparation across deployment budgets without additional fine-tuning, whereas budget-specific recovery can achieve higher per-model quality, particularly under stronger compression, at the cost of repeated optimization.

5.3 Model Preparation and Cross-Budget Transfer

Relation to Low-Rank Prehab. Low-Rank Prehab (Qin et al., 2025) prepares a model before SVD, optimizing a task loss on the original model with spectral rank regularization. LLift supervises a low-rank view at one anchor and restores the original complement. We compare our Prehab reimplementation with LLift at matched training exposure in Appendix B.11. However, the reimplementation's uncompressed seven-task mean accuracy already drops from 0.54 to 0.35. We therefore treat this comparison as a reproduction diagnostic. A separate control using truncated Schatten

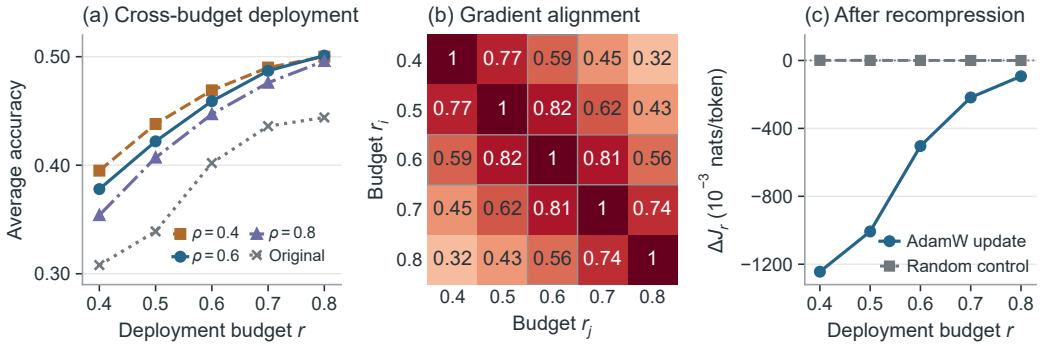


Figure 3: Cross-budget transfer and local compatibility on LLaMA-7B. (a) Seven-task mean accuracy across deployment budgets for models prepared at different anchors; we include the original model as a reference. (b) Cosine similarities of fixed-projector teacher-KL gradients at initialization. (c) Change in held-out student-to-teacher KL after updating the original model and recompressing it with SVD. We compare the first nonzero AdamW update ($s = 1$) with a norm-matched random update. Negative values indicate lower KL than when we compress the original model.

spectral regularization improves mean accuracy over Original but remains below LLift at all three budgets (Table 3). We detail the control's objective and training setup in Appendix B.8.

Comparison with Decomposable-Net. We evaluate a LLaMA-7B adaptation of Decomposable-Net (Yaguchi et al., 2021) under matched WT2 training exposure of approximately 11.5M token positions. With the same SVD-LLM deployment operator, Decomposable-Net achieves mean accuracies of $0.47/0.42/0.34$ at retention ratios of $0.8/0.6/0.4$, compared with $0.50/0.46/0.38$ for LLift. Appendix B.12 provides the adaptation details and results with both deployment operators.

Model preparation and residual restoration. In Table 3, LLift raises mean accuracy from $0.44/0.40/0.31$ to $0.50/0.46/0.38$ at $r = 0.8/0.6/0.4$. A single prepared model therefore benefits deployment budgets both above and below the preparation anchor.

To isolate residual restoration, we apply SVD-LLM to the fine-tuned student $Q_0 B^\star$ and the prepared model $Q_0 B^\star + R_0$ from the same preparation run, with identical adapted non-target parameters, deployment transforms, and ranks. Restoring R_0 improves seven-task mean accuracy from $0.44/0.44/0.37$ to $0.50/0.46/0.38$ at $r = 0.8/0.6/0.4$.

Our paired diagnostic shows that residual restoration reduces measured teacher reconstruction error for all 224 target matrices at each deployment budget. The sufficient condition in Appendix A.2 certifies improvement for $100.0\%/73.7\%/79.9\%$ of these matrices at $r = 0.8/0.6/0.4$. The remaining matrices also improve empirically, indicating that the sufficient condition is conservative. Appendix B.1 provides the detailed diagnostic results.

Role of non-target parameters. We freeze LLaMA-7B's embeddings, output head, and normalization parameters and train only the factor coordinates. Both variants' mean accuracies round to $0.50/0.46/0.38$, whereas freezing these parameters increases WT2 perplexity by 12–13% and C4 perplexity by approximately 4–5% across the three deployment budgets (Appendix B.2). We find that adapting these parameters primarily improves perplexity in this control.

Spectral energy retention. At the reported precision, the original and prepared models differ by at most 2×10^{-4} in retained whitened spectral energy fractions across the three budgets (Appendix B.5). This suggests that retained energy alone does not fully explain compression quality.

Preparation anchor. Preparing at $\rho = 0.4$ improves accuracy at $r = 0.6/0.4$ relative to the default configuration (Figure 3(a)). The default $\rho = 0.6$ yields the lowest WT2 perplexity at all three deployment budgets among the tested anchors (Appendix B.3), revealing a tradeoff between task accuracy and language-modeling quality. The anchor-matched points ($\rho = r$) in Figure 3(a)

provide budget-specific preparation references. The single default $\rho = 0.6$ prepared model remains competitive, with accuracy differences of $+0.50/0/-1.66$ percentage points at $r = 0.8/0.6/0.4$. Figure 3(a) further evaluates intermediate deployment budgets $r = 0.5$ and 0.7, which we do not use as preparation anchors, providing denser evidence of transfer away from the preparation budget.

Local cross-budget compatibility. We probe teacher-KL objectives at the original checkpoint while holding activation transforms and budget-specific projectors P_r fixed. Let $g_r = \nabla_W \widetilde{\mathcal{L}}_r(W_0; P_r)$, where W collects the sampled target matrices, and we keep all remaining parameters fixed. For the diagnostic projector P_ρ, an ordinary gradient step in the sampled fixed-basis factor coordinates induces $\Delta W = -\eta g_\rho$ in model weights. For sufficiently small η,

$$\Delta\widetilde{\mathcal{L}}_r = -\eta\langle g_r, g_\rho\rangle_F + O(\eta^2). \tag{11}$$

Positive alignment therefore identifies a shared local descent direction within the sampled parameters. The pairwise mean alignments in Figure 3(b) are positive. At $\rho = 0.6$, cosine similarity is 0.59 ± 0.02 for $r = 0.4$ and 0.56 ± 0.11 for $r = 0.8$, with standard deviations across five disjoint batches. This diagnostic tests local compatibility at initialization; the model comparisons evaluate transfer after AdamW preparation and SVD compression of the prepared model. Appendix B.6 provides the factor-coordinate derivation and probe details. This fixed-projector diagnostic does not test whether the same local benefit survives an AdamW update followed by SVD of the updated weights; the following probe evaluates the combined effect of these two operations.

Realized update after recompression. We scale the first nonzero AdamW update from the default $\rho = 0.6$ run by s and apply it to the original model, then recompute the deployment SVD factors from the updated weights. Let ΔJ_r denote the change in held-out WT2 student-to-teacher KL relative to compressing the original model at the same budget r. KL decreases at all five tested budgets for $s \in \{0.25, 0.5, 1\}$, with ΔJ_r ranging from -1.244 to -0.093 nats/token at $s = 1$ (Figure 3(c)). At $s = 1$, updating only the target matrices yields nearly identical KL changes, while the norm-matched random control stays near zero. These measurements provide local evidence of cross-budget transfer for the first nonzero update. Appendix B.6 reports the controls.

5.4 Generality Across Compression Operators and Models

Table 4: Effect of LLift preparation across deployment compressors on LLaMA-7B. All methods use `mix75` calibration (75% WT2 and 25% Alpaca). All LLift rows use the same prepared checkpoint. Bold values mark the best result within each Original/LLift pair.

		$r = 0.8$				$r = 0.6$				$r = 0.4$			
Compressor	Checkpoint	WT2	PTB	C4	Avg.	WT2	PTB	C4	Avg.	WT2	PTB	C4	Avg.
SVD-LLM	Original	8.03	14.03	13.89	0.44	14.27	38.34	35.03	0.40	63.25	300.26	218.52	0.31
	LLift	**7.49**	**11.72**	**11.15**	**0.50**	**8.40**	**15.96**	**14.77**	**0.46**	**13.48**	**53.47**	**32.24**	**0.38**
ZS-SVD	Original	**6.86**	11.42	10.46	**0.51**	12.41	34.31	28.04	0.42	56.71	291.02	184.77	0.33
	LLift	7.84	**11.39**	**10.11**	**0.51**	**9.86**	**19.22**	**15.38**	**0.45**	**16.88**	**88.31**	**35.33**	**0.38**
LEMS	Original	**6.25**	**10.12**	**8.93**	**0.52**	9.49	19.28	17.15	0.44	23.21	91.16	64.58	0.35
	LLift	7.76	11.22	9.40	**0.52**	**9.03**	**15.67**	**13.42**	**0.46**	**13.72**	**38.20**	**27.13**	**0.39**

Compression with different operators. Table 4 tests whether LLift preparation transfers across deployment operators. We compress the original LLaMA-7B model and the same default LLift-prepared checkpoint with SVD-LLM, one-shot ZS-SVD, and LEMS, using `mix75` calibration and matched evaluation settings within each pair. At $r = 0.6/0.4$, preparation improves mean accuracy and lowers all three perplexities for every compressor. With LEMS, mean accuracy rises from $0.44/0.35$ to $0.46/0.39$, while C4 perplexity decreases from $17.15/64.58$ to $13.42/27.13$.

At $r = 0.8$, preparation leaves LEMS mean accuracy nearly unchanged and increases all three perplexities. For ZS-SVD, WT2 perplexity increases, whereas PTB and C4 perplexity decrease. The gains therefore depend on both the compressor and deployment budget. These comparisons show that one checkpoint prepared with the SVD-LLM basis can also benefit a Fisher-based decomposition with nonuniform rank allocation.

Table 5: Evaluation across model families. Seven-task average accuracy (↑).

Model	Method	r=0.8	r=0.6	r=0.4
Qwen3-8B	SVD-LLM	0.49	0.35	0.31
	ZS-SVD	0.55	0.41	0.31
	LLift	0.55	0.47	0.34
Gemma-2-9B	SVD-LLM	0.52	0.40	0.32
	ZS-SVD	0.57	0.39	0.32
	LLift	0.57	0.51	0.40

Table 6: Scaling to larger LLaMA models. Seven-task average accuracy (↑).

Model	Method	r=0.8	r=0.6	r=0.4
LLaMA-13B	SVD-LLM	0.50	0.42	0.32
	ZS-SVD	0.54	0.39	0.34
	LLift	0.54	0.49	0.41
LLaMA-30B	SVD-LLM	0.55	0.45	0.33
	ZS-SVD	0.59	0.52	0.38
	LLift	0.57	0.53	0.45

Model families and scales. Tables 5 and 6 extend the evaluation to Qwen3-8B, Gemma-2-9B, LLaMA-13B, and LLaMA-30B. We prepare each model independently and use the same prepared model for all three deployment budgets. At $r = 0.6/0.4$, LLift achieves higher mean accuracy than the SVD-LLM and ZS-SVD references across all four models. At $r = 0.6$, accuracy on Qwen3-8B and Gemma-2-9B rises from $0.35/0.40$ with SVD-LLM to $0.47/0.51$ with LLift. At $r = 0.4$, accuracy reaches 0.41 on LLaMA-13B and 0.45 on LLaMA-30B, compared with 0.32 and 0.33 for SVD-LLM compression. At $r = 0.8$, LLift matches ZS-SVD's seven-task mean accuracy on Qwen3-8B, Gemma-2-9B, and LLaMA-13B, but trails it on LLaMA-30B (0.57 versus 0.59).

6 Conclusion

LLift prepares a reusable model for SVD compression across deployment budgets by fine-tuning a low-rank view at a single preparation budget. The frozen original model supervises the low-rank student. We merge the learned updates into the original weights while preserving components excluded from the student forward pass. We reuse SVDs of the prepared model's transformed target matrices at different truncation ranks to obtain deployment models without further fine-tuning.

On LLaMA-7B, a single prepared model raises seven-task mean accuracy from the reported SVD-LLM scores of $0.44/0.37/0.31$ to $0.50/0.46/0.38$ at $r = 0.8/0.6/0.4$. Controlled comparisons show that preserving the complementary components of the original weights is important for subsequent recompression, while the similar retained spectral energy before and after preparation indicates that spectral energy alone does not explain compression quality. The improvements further extend across model families, model sizes, and SVD-based compression operators. LLift thus shifts fine-tuning from individual compressed models to model preparation, allowing the same prepared model to support multiple deployment budgets without repeating compression-aware optimization.

Acknowledgments

We are especially grateful to Ananda (Theertha) Suresh for guidance throughout this project and for reviewing the manuscript. We thank Amir Yazdanbakhsh for enabling this project and providing valuable advice. We also thank Xin Lin and Henry Zheng for helpful discussions and suggestions.

REFERENCES

Ali Abbasi, Chayne Thrash, Haoran Qin, Shansita Sharma, Sepehr Seifi, and Soheil Kolouri. Zero sum SVD: Balancing loss sensitivity for low rank LLM compression, 2026. URL https://arxiv.org/abs/2602.02848.

Jose M. Alvarez and Mathieu Salzmann. Compression-aware training of deep networks. In *Advances in Neural Information Processing Systems*, volume 30, 2017. URL https://papers.nips.cc/paper_files/paper/2017/hash/db85e2590b6109813dafa101ceb2faeb-Abstract.html.

Aida Amini, Saadia Gabriel, Shanchuan Lin, Rik Koncel-Kedziorski, Yejin Choi, and Hannaneh Hajishirzi. MathQA: Towards interpretable math word problem solving with operation-based formalisms. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 2357–2367, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1245. URL https://aclanthology.org/N19-1245.

Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. PIQA: Reasoning about physical commonsense in natural language, 2019. URL https://arxiv.org/abs/1911.11641.

Ruisi Cai, Saurav Muralidharan, Greg Heinrich, Hongxu Yin, Zhangyang Wang, Jan Kautz, and Pavlo Molchanov. Flextron: Many-in-one flexible large language model, 2024. URL https://arxiv.org/abs/2406.10260.

Viktoriia A. Chekalina, Daniil Moskovskiy, Tatyana Matveeva, Andrey Kuznetsov, and Evgeny Frolov. Scalable Kronecker-factored Fisher approximation for neural network parameter sensitivity. In *Forty-third International Conference on Machine Learning*, 2026. URL https://proceedings.mlr.press/v306/chekalina26a.html.

Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try ARC, the AI2 reasoning challenge, 2018. URL https://arxiv.org/abs/1803.05457.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems, 2021. URL https://arxiv.org/abs/2110.14168.

Devvrit, Sneha Kudugunta, Aditya Kusupati, Tim Dettmers, Kaifeng Chen, Inderjit Dhillon, Yulia Tsvetkov, Hannaneh Hajishirzi, Sham Kakade, Ali Farhadi, and Prateek Jain. MatFormer: Nested transformer for elastic inference, 2024. URL https://arxiv.org/abs/2310.07707.

Hadi Mohaghegh Dolatabadi, Thalaiyasingam Ajanthan, Sameera Ramasinghe, Chamin P Hewa Koneputugodage, Shamane Siriwardhana, Violetta Shevchenko, Karol Pajak, James Snewin, Gil Avraham, and Alexander Long. NuMuon: Nuclear-norm-constrained Muon for compressible LLM training, 2026. URL https://arxiv.org/abs/2603.03597.

Martin Genzel, Patrick Putzky, Pengfei Zhao, Sebastian Schulze, Mattes Mollenhauer, Robert Seidel, Stefan Dietzel, and Thomas Wollmann. Choose your model size: Any compression of large language models without re-computation, 2025. URL https://arxiv.org/abs/2502.01717.

Xing Hu, Dawei Yang, Yuan Cheng, Zhixuan Chen, and Zukang Xu. SAES-SVD: Self-adaptive suppression of accumulated and local errors for SVD-based LLM compression, 2026. URL https://arxiv.org/abs/2602.03051.

Aditya Kusupati, Gantavya Bhatt, Aniket Rege, Matthew Wallingford, Aditya Sinha, Vivek Ramanujan, William Howard-Snyder, Kaifeng Chen, Sham Kakade, Prateek Jain, et al. Matryoshka representation learning. In *Advances in Neural Information Processing Systems*, December 2022.

Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics*, 19(2):313–330, 1993. URL https://aclanthology.org/J93-2004/.

Fanxu Meng, Zhaohui Wang, and Muhan Zhang. PiSSA: Principal singular values and singular vectors adaptation of large language models, 2024. URL https://arxiv.org/abs/2404.02948v3.

Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models. In *International Conference on Learning Representations*, 2017. URL https://openreview.net/forum?id=Byj72udxe.

Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct electricity? a new dataset for open book question answering. In *EMNLP*, 2018.

Ruoling Qi, Yirui Liu, Xuaner Wu, Xiangyu Wang, Ming Li, Chen Chen, Jian Chen, Yin Chen, and Qizhen Weng. Swift-SVD: Theoretical optimality meets practical efficiency in low-rank LLM compression. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=nAQ4h8FpdM.

Haoran Qin, Shansita Sharma, Ali Abbasi, Chayne Thrash, and Soheil Kolouri. Low-rank Prehab: Preparing neural networks for SVD compression, 2025. URL https://arxiv.org/abs/2512.01980.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020. URL http://jmlr.org/papers/v21/20-074.html.

Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. WinoGrande: An adversarial Winograd schema challenge at scale. *arXiv preprint arXiv:1907.10641*, 2019.

Sidhant Sundrani, Francesco Tudisco, and Pasquale Minervini. Low-rank compression of language models via differentiable rank selection, 2025. URL https://arxiv.org/abs/2512.13733.

Ali Taghibakhshi, Sharath Turuvekere Sreenivas, Saurav Muralidharan, Ruisi Cai, Marcin Chochowski, Ameya Sunil Mahabaleshwarkar, Yoshi Suhara, Oluwatobi Olabiyi, Daniel Korzekwa, Mostofa Patwary, Mohammad Shoeybi, Jan Kautz, Bryan Catanzaro, Ashwath Aithal, Nima Tajbakhsh, and Pavlo Molchanov. Nemotron Elastic: Towards efficient many-in-one reasoning LLMs, 2025. URL https://arxiv.org/abs/2511.16664.

Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca, 2023.

Moritz Thoma, Maximilian Groezinger, Maximilian Forstenhäusler, Emad Aghajanzadeh, Manoj Rohit Vemparala, Christos Anagnostopoulos, Pierpaolo Mori, Nael Fasfous, Alexander Frickenstein, Daniel Mueller-Gritschneder, and Ulf Schlichtmann. Advancing SVD-based LLM compression via layer-wise error model search. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=IjIgNPFuCt.

Qinsi Wang, Jinghan Ke, Masayoshi Tomizuka, Yiran Chen, Kurt Keutzer, and Chenfeng Xu. Dobi-SVD: Differentiable SVD for LLM compression and some new perspectives, 2025a. URL https://arxiv.org/abs/2502.02723.

Xin Wang, Yu Zheng, Zhongwei Wan, and Mi Zhang. SVD-LLM: Truncation-aware singular value decomposition for large language model compression. In *International Conference on Learning Representations (ICLR)*, 2025b. URL https://openreview.net/forum?id=LNYIUouhdt.

Atsushi Yaguchi, Taiji Suzuki, Shuhei Nitta, Yukinobu Sakata, and Akiyuki Tanizawa. Decomposable-Net: Scalable low-rank compression for neural networks. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, pp. 3249–3256, 2021. doi: 10.24963/ijcai.2021/447. URL https://www.ijcai.org/proceedings/2021/447.

Zhihang Yuan, Yuzhang Shang, Yue Song, Qiang Wu, Yan Yan, and Guangyu Sun. ASVD: Activation-aware singular value decomposition for compressing large language models, 2023.

Riccardo Zaccone, Stefanos Laskaridis, Marco Ciccone, and Samuel Horvath. FlexRank: Nested low-rank knowledge decomposition for adaptive model deployment. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=DK0kvnNelx.

Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. HellaSwag: Can a machine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019.

A Additional Derivations

We prove the minimum-change completion and analyze reconstruction after residual restoration. With fixed Q, we also show when training B through QB and training W through $QQ^\top W$ produce identical weight updates.

A.1 Minimum-Change Completion

We omit the layer index and let $Q_0^\top Q_0 = I$, $B_0 = Q_0^\top W_0$, and $R_0 = (I - Q_0 Q_0^\top) W_0$. Every feasible W satisfying $Q_0^\top W = B^\star$ has the unique decomposition $W = Q_0 B^\star + R$ with $Q_0^\top R = 0$. The two terms of $W - W_0 = Q_0(B^\star - B_0) + (R - R_0)$ are orthogonal in the Frobenius inner product. Consequently,

$$\|W - W_0\|_F^2 = \|B^\star - B_0\|_F^2 + \|R - R_0\|_F^2. \tag{12}$$

The learned coordinates fix the first term, and $R = R_0$ uniquely minimizes the second. This proves Eq. 9 for each target matrix. The minimization concerns target weights in Frobenius norm; we optimize the remaining model parameters with the training objective in Eq. 6.

A.2 Residual Restoration and Reconstruction after Compression

We compare compression with and without residual restoration using the same deployment transform T and rank m. Omitting the layer index, define

$$H = Q_0 B^\star T, \qquad R = (W_0 - Q_0 B_0)T, \qquad D = Q_0(B^\star - B_0)T.$$

The teacher matrix is $A_0 = W_0 T = H + R - D$. Let $[X]_m$ denote a best approximation to X of rank at most m in Frobenius norm. We compare

$$Y_- = [H]_m, \qquad Y_+ = [H + R]_m$$

through their teacher reconstruction errors:

$$\mathcal{E}(Y) = \|A_0 - Y\|_F^2, \qquad \Delta E_m = \mathcal{E}(Y_-) - \mathcal{E}(Y_+).$$

Define

$$J_m = \|H + R - Y_-\|_F^2 - \|H + R - Y_+\|_F^2,$$
$$I_m = 2\langle D, Y_+ - Y_- \rangle_F.$$

Since Y_+ optimally approximates $H + R$ and Y_- satisfies the same rank constraint, $J_m \geq 0$. Expanding $\mathcal{E}(Y) = \|(H + R - Y) - D\|_F^2$ gives

$$\Delta E_m = J_m - I_m. \tag{13}$$

Thus, restoration improves teacher reconstruction exactly when $J_m > I_m$.

Let P be the orthogonal projector onto $\mathrm{col}(Q_0)$. Since $PD = D$, Cauchy–Schwarz gives

$$\Delta E_m \geq J_m - 2\|D\|_F \|P(Y_+ - Y_-)\|_F =: M_m. \tag{14}$$

Therefore, $M_m > 0$ guarantees improvement at any deployment rank. This bound does not require $Q_0^\top R_0 = 0$.

For deployment rank $m = k + t > k$, where k is the preparation rank, we have $Y_- = H$. The matrix $H + [R]_t$ has rank at most m, so optimality of Y_+ implies

$$\|F\|_F \leq \|R - [R]_t\|_F, \qquad F = H + R - Y_+.$$

Because the two teacher reconstruction errors are $\|R - D\|_F^2$ and $\|F - D\|_F^2$,

$$\Delta E_m = \|R\|_F^2 - \|F\|_F^2 - 2\langle R, D \rangle_F + 2\langle F, D \rangle_F.$$

Applying Cauchy–Schwarz and the SVD identity $\|R\|_F^2 - \|R - [R]_t\|_F^2 = \|[R]_t\|_F^2$ yields

$$\Delta E_m \geq \|[R]_t\|_F^2 - 2\|D\|_F \|R - [R]_t\|_F - 2\langle R, D \rangle_F. \tag{15}$$

When $Q_0^\top R_0 = 0$, the final term vanishes. In this case,

$$J_m \geq \|[R]_t\|_F^2, \qquad \|P(Y_+ - Y_-)\|_F = \|PF\|_F \leq \|R - [R]_t\|_F,$$

so Eq. 14 is at least as tight as Eq. 15. The latter makes the contribution of additional rank explicit through the residual spectrum. These bounds characterize reconstruction of individual teacher weight matrices under the fixed deployment transform.

A.3 Fixed-Basis Updates under SGD and AdamW

For a fixed orthonormal basis Q and projector $P = QQ^\top$, write

$$W = QB + R, \qquad B = Q^\top W, \qquad R = (I - P)W. \tag{16}$$

The full-matrix forward PW equals the factorized forward QB. For upstream gradient $G = \partial\mathcal{L}/\partial(PW)$, their gradients satisfy $\nabla_W\mathcal{L} = QQ^\top G$ and $\nabla_B\mathcal{L} = Q^\top G$. Under plain SGD without momentum or weight decay, an update to the full matrix therefore gives

$$W_{t+1} = W_t - \eta_t QQ^\top G_t = Q\big(B_t - \eta_t Q^\top G_t\big) + R_0. \tag{17}$$

With matched initialization, minibatches, learning rates, and updates to other parameters, the two implementations produce the same student trajectory and the same prepared model with target weights $QB_* + R_0$. For this fixed-basis SGD setting, factor coordinates therefore suffice to represent the full-matrix training trajectory.

This equivalence does not generally extend to AdamW. Its elementwise second moment statistics and adaptive scaling depend on whether the optimizer acts in the coordinates of W or B. Even when the gradient with respect to W lies in the column space of Q, the adaptively scaled direction A_t need not do so. For an AdamW update to W with decay coefficient λ,

$$W_{t+1} = (1 - \eta_t\lambda)W_t - \eta_t A_t, \qquad R_{t+1} = (1 - \eta_t\lambda)R_t - \eta_t(I - P)A_t. \tag{18}$$

Thus full-matrix AdamW can update the complementary component, whereas factorized training followed by residual restoration preserves R_0. Refreshing the projector additionally changes the training subspace.

B Additional LLift Results and Evaluation Details

We examine preparation choices, compression operators, and cross-budget diagnostics. We report accuracy on a 0–1 scale and calculate differences before rounding scores to two decimal places.

B.1 Paired Diagnostic of Residual Restoration

We use the LLaMA-7B checkpoint prepared at $\rho = 0.6$ (Table 3). We compress $Q_0B^\star$ and $Q_0B^\star + R_0$ with identical adapted non-target parameters, frozen `mix75` deployment transforms, and ranks at $r = 0.8/0.6/0.4$.

For target matrix ℓ at budget r, we evaluate the quantities $J_{\ell,r}$, $I_{\ell,r}$, and $M_{\ell,r}$ from Appendix A.2 in FP64. We reconstruct $R_0 = W_0 - Q_0B_0$ in FP64 and obtain P by QR factorization. We separately measure reconstruction gains using the exported compression factors:

$$\widehat{\Delta E}_{\ell,r} = \|(W_{\ell,0} - \widehat{W}^-_{\ell,r})T_\ell\|_F^2 - \|(W_{\ell,0} - \widehat{W}^+_{\ell,r})T_\ell\|_F^2, \tag{19}$$

where $\widehat{W}^-_{\ell,r}$ and $\widehat{W}^+_{\ell,r}$ denote the compressed weights without and with restoration. We normalize each gain by $N_\ell = \|W_{\ell,0}T_\ell\|_F^2$. The maximum normalized discrepancy between the theoretical and exported reconstruction gains is 1.14×10^{-5}.

Table 7: Paired reconstruction diagnostic on LLaMA-7B. Positive $M_{\ell,r}$ certifies improvement for the exact truncated matrices defined in Appendix A.2. The measured gain $\widehat{\Delta E}_{\ell,r}$ uses exported compression factors. Counts, means, and maxima are taken over 224 target matrices.

r	$M_{\ell,r} > 0$	$\widehat{\Delta E}_{\ell,r} > 0$	Mean $\widehat{\Delta E}_{\ell,r}/N_\ell$	Max. $\lvert I_{\ell,r}\rvert/J_{\ell,r}$
0.8	224/224	224/224	3.882×10^{-2}	0.60%
0.6	165/224	224/224	8.055×10^{-3}	1.93%
0.4	179/224	224/224	5.948×10^{-3}	2.16%

The gain $J_{\ell,r}$ is strictly positive for every matrix at all three budgets, and $|I_{\ell,r}|/J_{\ell,r}$ never exceeds 2.2%. Thus, the interaction with the learned update does not offset the approximation gain in this

comparison. The measured teacher reconstruction error decreases for all 224 matrices at each budget. The sufficient condition certifies improvement for $224/165/179$ matrices at $r = 0.8/0.6/0.4$, while the remaining matrices also improve empirically. At $r = 0.8$, the bound expressed through the residual spectrum in Eq. 15 is positive for 216 matrices.

At $r = 0.6/0.8$, the deployment rank is at least the preparation rank for every matrix, so $[H]_m = H$ in exact arithmetic. Increasing the rank beyond k therefore retains no additional components of H, whereas residual restoration makes additional components available to SVD. Residual restoration also improves seven-task mean accuracy at all three budgets (Table 3).

B.2 Freezing the Remaining Parameters

To isolate the contribution of Φ, we freeze the token embeddings, output head, and all 65 RMSNorm modules in LLaMA-7B and train only the factor coordinates $\mathcal{B}$. Both LLift variants use KL at $\rho = 0.6$ and restore the original residuals before compression, with all other settings unchanged. Original receives no preparation.

Table 8: Freezing non-target parameters Φ during LLaMA-7B preparation. All checkpoints use SVD-LLM with the same fixed original-model `mix75` activation transforms. Avg. is seven-task mean plain accuracy.

r	Checkpoint	WT2 PPL ↓	C4 PPL ↓	Avg. ↑
0.8	Original	8.03	13.89	0.44
	LLift (frozen Φ)	8.40	11.58	0.50
	LLift (trainable Φ)	7.49	11.15	0.50
0.6	Original	14.27	35.03	0.40
	LLift (frozen Φ)	9.46	15.46	0.46
	LLift (trainable Φ)	8.40	14.77	0.46
0.4	Original	63.25	218.52	0.31
	LLift (frozen Φ)	15.12	33.65	0.38
	LLift (trainable Φ)	13.48	32.24	0.38

With Φ frozen, LLift reduces C4 perplexity relative to Original by $16.6/55.9/84.6\%$ and raises mean accuracy from $0.44/0.40/0.31$ to $0.50/0.46/0.38$ at $r = 0.8/0.6/0.4$ (Table 8). Compared with training Φ, freezing it increases WT2 perplexity by 12–13% and C4 perplexity by approximately 4–5% across these budgets, while changing unrounded mean accuracy by at most 0.0022.

B.3 Objective and Anchor Ablations

We vary the loss while keeping the remaining recipe fixed. KL+CE gives each term weight one; pure CE uses no teacher. Figure 4 shows that CE increases C4 perplexity at all three budgets, from 32.2 to 90.6 at $r = 0.4$. Its accuracy slightly exceeds KL at $r = 0.6$ but falls to 0.37 at $r = 0.4$, versus 0.38 for KL. KL+CE improves mean accuracy over KL by 0.10–0.83 percentage points while slightly increasing C4 perplexity at every budget. We use KL to balance task accuracy and language-modeling quality.

On LLaMA-7B, $\rho = 0.6$ yields the lowest WT2 perplexity at all three budgets among the tested anchors, while $\rho = 0.4$ gives higher accuracy at $r = 0.6/0.4$ (Table 9).

Anchor selection beyond LLaMA-7B. We evaluate Qwen3-8B at $\rho \in \{0.4, 0.6, 0.8\}$ using seed 42, with the checkpoint from the main experiments used for $\rho = 0.6$. We compress each checkpoint at $r \in \{0.8, 0.6, 0.4\}$ with SVD-LLM and fixed original-model `mix75` activation transforms (Table 10).

The default $\rho = 0.6$ improves mean accuracy over direct compression by $2.41/9.28/3.39$ percentage points at $r = 0.8/0.6/0.4$. Among the tested anchors, $\rho = 0.4$ gives the highest accuracy and lowest WT2/C4 perplexity at $r = 0.4$, whereas $\rho = 0.8$ gives the lowest perplexity at $r = 0.8$. The preferred anchor therefore depends on the deployment budget and metric.

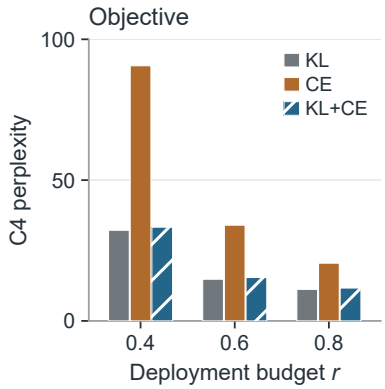


Figure 4: Training-objective ablation on LLaMA-7B. C4 perplexity for KL, CE, and equally weighted KL+CE at each deployment budget, with preparation anchor $\rho = 0.6$.

Table 9: Preparation-anchor ablation on LLaMA-7B under SVD-LLM. WT2 and C4 report perplexity (↓); Avg. is seven-task mean plain accuracy (↑).

	$r = 0.8$			$r = 0.6$			$r = 0.4$		
Checkpoint	WT2	C4	Avg.	WT2	C4	Avg.	WT2	C4	Avg.
Original	8.03	13.89	0.44	14.27	35.03	0.40	63.25	218.52	0.31
LLift ($\rho = 0.4$)	9.72	12.78	0.50	10.89	17.08	0.47	14.07	30.78	0.39
LLift ($\rho = 0.6$)	7.49	11.15	0.50	8.40	14.77	0.46	13.48	32.24	0.38
LLift ($\rho = 0.8$)	7.70	11.22	0.50	9.53	16.91	0.45	20.20	52.31	0.35

Table 10: Preparation-anchor ablation on Qwen3-8B under SVD-LLM. WT2 and C4 report perplexity (↓); Avg. is seven-task mean plain accuracy (↑).

	$r = 0.8$			$r = 0.6$			$r = 0.4$		
Checkpoint	WT2	C4	Avg.	WT2	C4	Avg.	WT2	C4	Avg.
Original	14.21	28.97	0.52	31.79	89.05	0.38	239.87	896.79	0.31
LLift ($\rho = 0.4$)	17.95	37.03	0.54	21.09	54.99	0.47	32.40	120.76	0.37
LLift ($\rho = 0.6$)	13.95	29.55	0.55	17.85	46.78	0.47	44.60	173.05	0.34
LLift ($\rho = 0.8$)	12.66	26.23	0.54	18.76	51.12	0.43	93.18	434.36	0.32

Across LLaMA-7B and Qwen3-8B, the tested anchors $\rho = 0.4$ and 0.6 generally provide a favorable balance of task accuracy and perplexity over the evaluated deployment budgets. These results support using 0.4–0.6 as a practical anchor range for cross-budget preparation.

B.4 Preparation Bases and Compression Operators

Table 11 groups the preparation-basis comparison by deployment operator. We prepare each LLift model once at $\rho = 0.6$ using either an SVD-LLM or a ZS-SVD basis, then compress it with each operator at three budgets. Both deployment protocols use `mix75` calibration. Each deployment block evaluates the two preparation bases under the same compression operator and evaluation settings.

Relative to the default SVD-LLM basis, the ZS-SVD basis changes accuracy by $-0.11/+1.29/-1.75$ percentage points under ZS-SVD deployment at $r = 0.8/0.6/0.4$, and by $-0.60/-0.82/-2.32$ points under SVD-LLM deployment. Matching the preparation and deployment operators therefore does not uniformly improve transfer.

In Table 4, SVD-LLM uses fixed original-model activation transforms. For LEMS, we compute KFAC-SVD statistics and select ranks for each checkpoint using the same `mix75` windows. We

Table 11: LLaMA-7B results across preparation bases and compression operators under `mix75` calibration. Parentheses identify the preparation basis; both LLift variants use KL at $\rho = 0.6$. WT2 and C4 report perplexity; task columns and Avg. report zero-shot accuracy. Each panel reports one deployment operator.

Preparation	r	WT2	C4	OBQA	ARC-e	ARC-c	WinoG	HellaS	PIQA	MathQA	Avg.
SVD-LLM deployment											
Original	0.8	8.03	13.89	0.23	0.62	0.31	0.63	0.39	0.69	0.23	0.44
	0.6	14.27	35.03	0.24	0.53	0.26	0.57	0.34	0.64	0.22	0.40
	0.4	63.25	218.52	0.15	0.28	0.21	0.49	0.27	0.55	0.22	0.31
LLift (SVD-LLM)	0.8	7.49	11.15	0.30	0.67	0.37	0.66	0.52	0.74	0.25	0.50
	0.6	8.40	14.77	0.25	0.63	0.31	0.62	0.45	0.71	0.25	0.46
	0.4	13.48	32.24	0.17	0.48	0.24	0.57	0.34	0.62	0.22	0.38
LLift (ZS-SVD)	0.8	7.74	11.61	0.27	0.67	0.34	0.66	0.51	0.74	0.26	0.50
	0.6	9.41	17.05	0.24	0.63	0.29	0.63	0.44	0.70	0.24	0.45
	0.4	17.96	52.73	0.17	0.42	0.21	0.55	0.32	0.60	0.22	0.36
ZS-SVD deployment											
Original	0.8	6.86	10.46	0.30	0.71	0.37	0.69	0.50	0.75	0.25	0.51
	0.6	12.41	28.04	0.23	0.59	0.29	0.62	0.37	0.67	0.22	0.42
	0.4	56.71	184.77	0.15	0.34	0.21	0.53	0.28	0.56	0.21	0.33
LLift (SVD-LLM)	0.8	7.84	10.11	0.30	0.70	0.39	0.66	0.54	0.76	0.25	0.51
	0.6	9.86	15.38	0.23	0.62	0.30	0.61	0.44	0.69	0.23	0.45
	0.4	16.88	35.33	0.17	0.46	0.25	0.58	0.34	0.61	0.23	0.38
LLift (ZS-SVD)	0.8	7.27	10.02	0.29	0.71	0.37	0.66	0.54	0.76	0.26	0.51
	0.6	8.58	14.34	0.25	0.65	0.31	0.62	0.45	0.70	0.23	0.46
	0.4	16.38	42.18	0.16	0.43	0.21	0.55	0.33	0.60	0.22	0.36

compare the original and prepared checkpoints within each compressor under matched evaluation settings.

Deployment recalibration. We re-estimate activation statistics from the LLift-prepared model, apply SVD-LLM to its weights using the resulting transforms, and truncate the shared decompositions at $r \in \{0.4, 0.5, 0.6, 0.7, 0.8\}$. Table 12 reports changes relative to compressing the same weights with original-model statistics.

Table 12: Deployment recalibration of LLaMA-7B LLift under SVD-LLM. Entries are paired differences (prepared-model statistics minus original-model statistics) for the same weights. Accuracy differences use the 0–1 scale; PPL differences are absolute.

Metric	$r = 0.4$	$r = 0.5$	$r = 0.6$	$r = 0.7$	$r = 0.8$
Δ WT2 PPL	$+1.17$	$+0.83$	$+0.61$	$+0.44$	$+0.36$
Δ C4 PPL	-0.06	$+0.31$	$+0.30$	$+0.26$	$+0.19$
Δ Avg. accuracy	-0.0021	$+0.0073$	$+0.0035$	$+0.0001$	$+0.0008$

Recalibration changes mean accuracy by -0.21 to $+0.73$ percentage points. WT2 perplexity increases by 0.36–1.17 at every budget; C4 perplexity changes by -0.06 to $+0.31$ and increases at four of the five budgets. We retain original-model statistics in the main comparisons because recalibration provides no consistent perplexity benefit.

B.5 Retained Spectral Energy

We apply SVD to each $W_\ell T_\ell$ using the same frozen original-model `mix75` transforms and select prefixes according to Eq. 4. We measure retained whitened spectral energy as the fraction of squared

singular values that remain after truncation. Table 13 shows that the reported retained fractions differ by at most 2×10^{-4} despite the improvements in compressed-model quality.

Table 13: Retained whitened spectral energy after truncation at each deployment budget. Both models use the same frozen `mix75` activation transforms.

Model	$r = 0.4$	$r = 0.6$	$r = 0.8$
Original	0.8818	0.9325	0.9633
LLift	0.8820	0.9327	0.9634

B.6 Gradient Alignment and Recompression Probes

We examine cross-budget transfer through gradient alignment at initialization and through the measured effect of an AdamW update followed by SVD recompression.

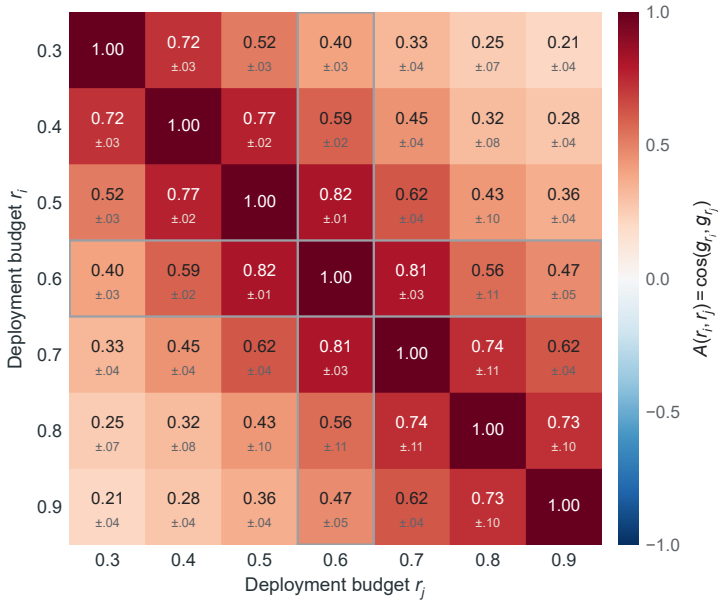


Figure 5: Cross-budget gradient alignment at initialization. Cells report mean cosine similarity and standard deviation across five disjoint WT2 batches. Figure 3(b) shows the 0.4–0.8 submatrix of means.

We use SVD-LLM projections and student-to-teacher KL from Eq. 6, with temperature one. We sample 56 of 224 target matrices across eight depths and use five disjoint WT2 batches of 32×1024 input positions. We evaluate all budgets at the original model and hold activation statistics, each budget's projector, and the remaining parameters fixed during differentiation. Let W collect the sampled target matrices and define

$$g_r = \nabla_W \widetilde{\mathcal{L}}_r(W; P_r)\Big|_{W=W_0}, \qquad A(\rho, r) = \frac{\langle g_\rho, g_r \rangle_F}{\|g_\rho\|_F \|g_r\|_F}. \tag{20}$$

Here $\widetilde{\mathcal{L}}_r$ is the fixed-projector KL surrogate, and inner products sum over the sampled matrices.

For an orthonormal basis Q_ρ of the diagnostic projector $P_\rho = Q_\rho Q_\rho^\top$ and upstream gradient G_ρ, the corresponding factor coordinates satisfy $\nabla_B \mathcal{L}_\rho = Q_\rho^\top G_\rho$. An ordinary gradient step confined to the sampled factor coordinates therefore induces the model-weight update $\Delta W = Q_\rho \Delta B = -\eta Q_\rho Q_\rho^\top G_\rho = -\eta g_\rho$. A first-order expansion of $\widetilde{\mathcal{L}}_r$ along this update gives Eq. 11 for sufficiently

small η. All off-diagonal means in Figure 5 are positive, identifying a shared local descent component at initialization. Standard deviations measure variation across probe batches at the same model checkpoint.

Recompression probe and controls. For the update probe in Figure 3(c), let $\Delta\Theta$ denote the first nonzero AdamW displacement from the default $\rho = 0.6$ run. We measure $J_r(\Theta) = \mathrm{KL}(p_{\mathrm{S},r}(\Theta)\|p_{\mathrm{T}})$ on 64 held-out WT2 validation sequences of length 1024, with no overlap with training batches, at temperature one. Here $p_{\mathrm{S},r}(\Theta)$ comes from compressing Θ at budget r and p_{T} comes from the frozen original teacher. For each scale s, we recompute the weighted SVDs of the updated target matrices using the frozen original-model `mix75` deployment transforms and report

$$\Delta J_r(s) = J_r(\Theta_0 + s\Delta\Theta) - J_r(\Theta_0).$$

Table 14 reports the full update at $s \in \{0.25, 0.5, 1\}$ and the target-only and norm-matched random controls at $s = 1$; negative values indicate improvement. To isolate the target-matrix contribution, we write $\Delta\Theta = (\Delta W, \Delta\Phi)$ and repeat the probe with $(\Delta W, 0)$, keeping the remaining parameters at Φ_0.

Table 14: Change in held-out student-to-teacher KL after scaling the first nonzero AdamW update and recomputing the deployment SVD factors. Entries report $\Delta J_r(s)$ in 10^{-3} nats/token; negative values indicate improvement over compressing the original model. Signed zeros reflect rounding.

Update	s	$r = 0.4$	$r = 0.5$	$r = 0.6$	$r = 0.7$	$r = 0.8$
Full $(\Delta W, \Delta\Phi)$	0.25	-570	-508	-283	-120	-53
	0.5	-932	-793	-412	-177	-79
	1	-1244	-1007	-504	-218	-93
Target only $(\Delta W, 0)$	1	-1223	-994	-497	-213	-91
Norm-matched random	1	$+0.8$	$+0.4$	$+0.2$	-0.0	$+0.0$

At $s = 1$, the target-only changes match the full-update changes within 2.3%, while the random control stays near zero. For this first nonzero update at $s \in \{0.25, 0.5, 1\}$, we also compare KL changes under fixed projectors with those after recomputing the deployment SVD factors. The changes have the same sign; the largest reported relative difference is about 37%, at $s = 0.25$ and $r = 0.4$.

B.7 Budget-Specific Fine-Tuning

We compress the original LLaMA-7B model with SVD-LLM using WT2 calibration, then fine-tune each compressed model independently on Alpaca. We first train LoRA adapters on the `*_u_proj` factors and merge them, then train and merge adapters on the `*_v_proj` factors. We use a fixed configuration with seed 42, three epochs per stage, and a learning rate of 10^{-4}. We retain the weights from the final epoch of each stage, without model selection. We evaluate mean plain accuracy over the same seven tasks.

Table 15: Budget-specific fine-tuning and LLift on LLaMA-7B. SVD-LLM + FT uses WT2 calibration followed by two-stage LoRA fine-tuning on Alpaca at each budget; Before FT reports the accuracy of its compressed initialization. LLift uses one model prepared at $\rho = 0.6$ and frozen `mix75` deployment transforms for all budgets. WT2 and C4 report perplexity; Avg. is seven-task mean plain accuracy.

r	Before FT Avg.	SVD-LLM + FT			LLift		
		WT2 ↓	C4 ↓	Avg. ↑	WT2 ↓	C4 ↓	Avg. ↑
0.8	0.38	7.62	12.01	0.50	7.49	11.15	0.50
0.6	0.35	9.54	15.78	0.46	8.40	14.77	0.46
0.4	0.31	15.19	26.79	0.39	13.48	32.24	0.38

Table 15 compares the resulting models with the single LLift model prepared at $\rho = 0.6$. At $r = 0.8/0.6/0.4$, LLift's accuracy is $0.15/0.22/1.13$ percentage points lower, while its WT2 perplexity is lower at all three budgets. LLift also has lower C4 perplexity at $r = 0.8/0.6$, while budget-specific fine-tuning has lower C4 perplexity at $r = 0.4$.

Matched recovery. We apply SVD-LLM to the original checkpoint at each budget using the same frozen WT2/Alpaca 75:25 calibration transforms and rank rule as LLift. We then recover each compressed model with LLift's WT2 KL training recipe, training the right factors and non-target parameters while freezing the left factors. Each recovery run receives the same training exposure as LLift: approximately 11.5M token positions for LLaMA-7B and 10.2M for Qwen3-8B. All runs use seed 42.

Table 16: Budget-specific SVD-LLM recovery and LLift. WT2 and C4 report perplexity (↓). Avg. is seven-task mean plain accuracy (↑).

	$r = 0.8$			$r = 0.6$			$r = 0.4$		
Method	WT2	C4	Avg.	WT2	C4	Avg.	WT2	C4	Avg.
LLaMA-7B									
LLift	7.49	11.15	0.50	8.40	14.77	0.46	13.48	32.24	0.38
Recovery	6.75	10.52	0.50	7.96	14.05	0.46	10.97	23.40	0.40
Qwen3-8B									
LLift	13.95	29.55	0.55	17.85	46.78	0.47	44.60	173.05	0.34
Recovery	12.08	28.45	0.55	15.78	38.82	0.48	26.26	79.88	0.39

Table 16 compares three independently recovered models with one LLift-prepared checkpoint. Mean accuracy is broadly comparable at $r = 0.8/0.6$, while recovery achieves higher accuracy at $r = 0.4$ and lower perplexity across all three budgets.

Table 17 reports Qwen3-8B GSM8K results under the evaluation protocol in Appendix D. Per-budget exposure assigns one LLift run's training exposure to each recovery model. Total exposure divides that exposure across $N = 3$ budgets ($r = 0.8/0.6/0.4$) or $N = 5$ budgets ($r = 0.8/0.7/0.6/0.5/0.4$). LLift uses the same prepared checkpoint throughout.

Table 17: Qwen3-8B GSM8K strict exact-match accuracy (%, ↑).

Method	$r = 0.8$	$r = 0.7$	$r = 0.6$	$r = 0.5$	$r = 0.4$
LLift	64.67	46.32	17.13	2.05	0.38
Recovery, per-budget ($N = 3$)	61.41	—	17.97	—	1.21
Recovery, total ($N = 3$)	60.65	—	13.65	—	1.44
Recovery, total ($N = 5$)	58.15	22.82	6.97	2.35	0.76

B.8 Preparation and Baseline Protocols

Target matrices and preparation. We factorize the attention projections (`q_proj`, `k_proj`, `v_proj`, `o_proj`) and MLP projections (`gate_proj`, `up_proj`, `down_proj`) in every Transformer block. LLaMA-7B has 224 such matrices across 32 blocks. We group the token embeddings, output head, and normalization parameters into Φ and train them without factorization. For the LLaMA-7B reference, we use seed 42, teacher KL with temperature one, and the WT2 training split. We run AdamW for four epochs (176 optimizer steps) with learning rate 10^{-4}, warmup fraction 0.05, weight decay 0.01, gradient clipping at norm 1, sequence length 1024, and global batch size 64. We construct the preparation bases from original-model activations on WT2 and freeze both the teacher and these bases.

Deployment calibration and scoring. We estimate activation transforms from the original model on 256 sequences of length 2048. This calibration set contains 75% WT2 and 25% Alpaca (`mix75`); we construct it separately from preparation calibration. We freeze these transforms across models

and budgets, decompose each $W_\ell T_\ell$ once, and select prefixes for each deployment budget. We evaluate WT2 perplexity on its held-out test split and C4 perplexity separately. We average unrounded plain accuracies equally over the seven tasks in Sec. 5. The paired residual-restoration control keeps the adapted non-target parameters and deployment transforms identical.

Compression and elastic baselines. We use the SVD baselines documented by Abbasi et al. (2026); Hu et al. (2026). For cross-model experiments, SVD-LLM uses 256 WT2 sequences of length 2048, and ZS-SVD uses 256 WT2 samples.

We run Swift-SVD* with the released dynamic-rank configuration ($\alpha = 0.6$) and C4 calibration, and LEMS with WT2 calibration. Both use 256 sequences of length 2048. LEMS uses 32 rank-search samples and the HiGHS solver in place of Gurobi. Our LLaMA-7B Swift-SVD reproduction falls below the paper's reported accuracies. Both methods select ranks for each target budget without budget-specific model-weight training.

For the elastic comparison, ACIP uses C4 and FlexRank uses FineWeb-Edu for preparation, while LLift uses WT2. ACIP omits the rank-32 correction Δ; ACIP ($+\Delta$) retains it. The reported budget for ACIP ($+\Delta$) excludes this correction's 79.95M parameters.

Spectral-regularization control. For the spectral control in Table 3, we train the dense LLaMA-7B student with student-to-teacher KL at temperature one, without CE or a low-rank forward pass. We match LLift's AdamW recipe and 11.5M-token training budget described above. For each of the 224 attention and MLP target matrices $W_\ell \in \mathbb{R}^{m_\ell \times n_\ell}$, we penalize the normalized truncated Schatten-1 tail

$$\mathcal{R}_\ell(W_\ell) = \frac{\sum_{i>k_\ell} \sigma_i(W_\ell T_\ell)}{\|W_\ell T_\ell\|_F}, \qquad k_\ell = \left\lfloor \frac{0.6\, m_\ell n_\ell}{m_\ell + n_\ell} \right\rfloor. \tag{21}$$

We compute and freeze the whitening factors T_ℓ from the original model on 256 WT2 sequences of length 1024. We compute spectra from FP32 Gram eigenvalues floored at 10^{-6}. At each optimizer step, we apply a separate Muon update for the spectral term with coefficient 10^{-3}, learning rate 10^{-5}, momentum 0.95, five Newton–Schulz iterations, and RMS scale 0.2. We compress the trained weights with SVD-LLM using frozen original-model `mix75` transforms and the shared evaluation protocol described above.

B.9 Additional Comparisons on Qwen3-8B

Table 18 compares SVD-LLM, one-shot ZS-SVD, the released Swift-SVD* configuration (Qi et al., 2026), and KFAC-SVD with LEMS rank allocation (Thoma et al., 2026), denoted LEMS, on Qwen3-8B. We use the shared evaluation protocol to evaluate SVD-LLM perplexity and ZS-SVD. LLift + SVD-LLM and LLift + LEMS apply SVD-LLM and KFAC-SVD with LEMS rank allocation, respectively, to the same checkpoint prepared at $\rho = 0.6$.

Table 18: Additional SVD compression results on Qwen3-8B. WT2 and C4 report perplexity (↓); Avg. is seven-task mean plain accuracy (↑).

	$r = 0.8$			$r = 0.6$			$r = 0.4$		
Method	WT2	C4	Avg.	WT2	C4	Avg.	WT2	C4	Avg.
SVD-LLM	14.06	37.63	0.49	28.96	136.92	0.35	186.53	1344.52	0.31
ZS-SVD	12.53	27.63	0.55	61.29	213.74	0.40	165.57	720.59	0.31
Swift-SVD*	23.57	21.97	0.53	66.82	41.62	0.43	301.21	130.99	0.33
LEMS	10.26	18.01	0.56	15.04	54.35	0.43	40.48	222.46	0.33
LLift + SVD-LLM	13.95	29.55	0.55	17.85	46.78	0.47	44.60	173.05	0.34
LLift + LEMS	12.71	21.44	0.59	15.24	29.75	0.51	26.68	71.90	0.38

LLift + SVD-LLM achieves higher mean accuracy than SVD-LLM and Swift-SVD* at all three budgets and than ZS-SVD and LEMS at $r = 0.6/0.4$, while LEMS performs better at $r = 0.8$. Deploying the same prepared checkpoint with LEMS yields accuracies of $0.59/0.51/0.38$ at $r = 0.8/0.6/0.4$ and lower WT2/C4 perplexity than LLift + SVD-LLM at every budget. This composition requires no further model training.

B.10 Deployment Calibration Sensitivity

We keep the LLaMA-7B LLift checkpoint fixed and compare deployment calibration on WT2, C4, and the default WT2/Alpaca 75:25 mixture (`mix75`), using 256 sequences of length 2048 per corpus. For each corpus, we apply SVD-LLM to both checkpoints with shared original-model activation transforms. Table 19 shows seven-task mean accuracy ranges of 0.49–0.50, 0.45–0.46, and 0.35–0.38 for LLift at $r = 0.8/0.6/0.4$, respectively. At $r = 0.4$, LLift achieves $0.35/0.36/0.38$ with WT2/C4/`mix75`, compared with $0.31/0.30/0.31$ for the original checkpoint under matched calibration. Preparation gains therefore persist with WT2 and C4, while `mix75` further improves LLift's accuracy.

Table 19: Deployment calibration sensitivity on LLaMA-7B under SVD-LLM. Original and LLift share original-model activation transforms within each calibration setting. WT2 and C4 report perplexity; Avg. is seven-task mean plain accuracy.

		WT2 PPL ↓		C4 PPL ↓		Avg. ↑	
r	Calibration	Original	LLift	Original	LLift	Original	LLift
0.8	WT2	7.90	7.14	15.85	11.46	0.38	0.50
	C4	12.01	12.02	12.01	10.66	0.46	0.49
	`mix75`	8.03	7.49	13.89	11.15	0.44	0.50
0.6	WT2	13.65	7.88	54.06	16.08	0.36	0.45
	C4	40.49	20.67	25.17	14.61	0.38	0.45
	`mix75`	14.27	8.40	35.03	14.77	0.40	0.46
0.4	WT2	63.95	12.19	353.48	41.36	0.31	0.35
	C4	306.60	67.50	149.04	34.87	0.30	0.36
	`mix75`	63.25	13.48	218.52	32.24	0.31	0.38

Perplexity is more domain-sensitive. At $r = 0.4$, switching LLift's calibration from WT2 to C4 increases WT2 perplexity from 12.19 to 67.50 while reducing C4 perplexity from 41.36 to 34.87. At this budget, LLift lowers both perplexities relative to the original checkpoint under all three matched calibration settings.

B.11 Low-Rank Prehab and Matched Training Exposure

We compare our Low-Rank Prehab (Qin et al., 2025) reimplementation with LLift on LLaMA-7B using the same first 256 WT2 training chunks of length 1024 for three epochs (0.79M token positions), with global batch size 16, 48 optimizer steps, and seed 42. Prehab minimizes CE plus $\lambda \sum_\ell \|W_\ell T_\ell\|_*^2 / \|W_\ell T_\ell\|_F^2$, using frozen original-model whitening factors, $\lambda = 10^{-3}$, and AdamW with learning rate 10^{-5}. LLift retains its default KL objective, $\rho = 0.6$ anchor, and 10^{-4} learning rate. We apply SVD-LLM to each checkpoint with shared frozen `mix75` transforms, identical deployment ranks, and no further fine-tuning. Table 20 includes a CE-only control ($\lambda = 0$) and default LLift (11.53M positions) as references.

At $r = 0.8/0.6/0.4$, LLift at 0.79M positions is $0.49/0.14/0.79$ percentage points below default LLift in mean accuracy (computed from unrounded scores), with higher WT2 and C4 perplexity. It exceeds our Prehab reimplementation's mean accuracy at all three budgets. However, the latter's uncompressed accuracy already falls from 0.54 to 0.35. We therefore treat this single-seed comparison as a reproduction diagnostic, without claiming superiority over a validated Prehab implementation.

B.12 Comparison with Decomposable-Net

We adapt Decomposable-Net (Yaguchi et al., 2021) to pretrained LLaMA-7B using next-token CE supervision. We retain its joint dense and low-rank objective,

$$\mathcal{L}_{\mathrm{DNet}} = 0.5\,\mathcal{L}_{\mathrm{CE}}(W) + 0.5\,\mathcal{L}_{\mathrm{CE}}(\widetilde{W}), \tag{22}$$

where $\widetilde{W}$ is constructed by truncated SVD at each training step, with gradients propagated through the decomposition. We sample a parameter-retention ratio uniformly from $[0.4, 0.8]$ and allocate

Table 20: Training-exposure comparison on LLaMA-7B with SVD-LLM and frozen `mix75` calibration. Exposure counts training token positions in millions, including repeated epochs. Dense denotes the uncompressed checkpoint; Avg. is seven-task mean plain accuracy.

Budget	Model	Exposure (M)	WT2 PPL ↓	C4 PPL ↓	Avg. ↑
Dense	Original	—	5.68	7.34	0.54
	CE fine-tuning	0.79	5.27	7.72	0.54
	Prehab (reimpl.)	0.79	5.22	7.82	0.35
	LLift	0.79	8.04	8.13	0.54
	LLift	11.53	7.38	8.23	0.54
$r = 0.8$	Original	—	8.03	13.89	0.44
	CE fine-tuning	0.79	7.53	15.27	0.33
	Prehab (reimpl.)	0.79	8.04	17.74	0.33
	LLift	0.79	9.46	11.89	0.50
	LLift	11.53	7.49	11.15	0.50
$r = 0.6$	Original	—	14.27	35.03	0.40
	CE fine-tuning	0.79	13.44	42.18	0.39
	Prehab (reimpl.)	0.79	16.76	62.88	0.30
	LLift	0.79	11.73	16.63	0.46
	LLift	11.53	8.40	14.77	0.46
$r = 0.4$	Original	—	63.25	218.52	0.31
	CE fine-tuning	0.79	65.40	303.82	0.30
	Prehab (reimpl.)	0.79	94.64	459.60	0.30
	LLift	0.79	20.98	39.36	0.37
	LLift	11.53	13.48	32.24	0.38

Table 21: Comparison with Decomposable-Net on LLaMA-7B at matched training exposure. D-Net denotes native truncated SVD with global rank allocation. SVD-LLM denotes activation-whitened SVD with shared calibration and deployment ranks across checkpoints. WT2 and C4 report perplexity (↓), and Avg. is seven-task mean plain accuracy (↑).

Budget	Preparation	Deployment	WT2 PPL ↓	C4 PPL ↓	Avg. ↑
Dense	Original	—	5.68	7.34	0.54
	Decomposable-Net	—	9.44	23.14	0.51
	LLift	—	7.38	8.23	0.54
$r = 0.8$	Original	D-Net	9682.78	10689.00	0.30
	Original	SVD-LLM	8.03	13.89	0.44
	Decomposable-Net	D-Net	9.06	33.83	0.45
	Decomposable-Net	SVD-LLM	9.87	43.27	0.47
	LLift	SVD-LLM	7.49	11.15	0.50
$r = 0.6$	Original	D-Net	14968.86	29171.23	0.30
	Original	SVD-LLM	14.27	35.03	0.40
	Decomposable-Net	D-Net	11.83	60.26	0.39
	Decomposable-Net	SVD-LLM	12.68	85.71	0.42
	LLift	SVD-LLM	8.40	14.77	0.46
$r = 0.4$	Original	D-Net	31626.26	30290.67	0.30
	Original	SVD-LLM	63.25	218.52	0.31
	Decomposable-Net	D-Net	20.04	139.06	0.33
	Decomposable-Net	SVD-LLM	34.89	460.14	0.34
	LLift	SVD-LLM	13.48	32.24	0.38

ranks across the 224 target matrices by globally ordering their singular values. Relative to the original formulation, we use a pretrained language model, AdamW, and sampling over parameter-retention ratios. Training follows LLift's WT2 schedule with sequence length 1024, global batch size 64, four epochs (176 optimizer steps), and approximately 11.5M token positions. We use a learning rate of 10^{-4} and seed 42.

We evaluate the Decomposable-Net checkpoint using both its native truncated-SVD deployment with global rank allocation, denoted D-Net in Table 21, and SVD-LLM. The latter uses the same frozen original-model `mix75` calibration transforms and deployment ranks as LLift. Dense denotes each uncompressed checkpoint.

Table 21 shows that Decomposable-Net preparation improves compressed-model accuracy under its native deployment procedure. LLift achieves higher mean accuracy and lower WT2/C4 perplexity at all three retention ratios than either Decomposable-Net deployment variant.

C Computational Cost and Resource Usage

For deployment ratios $\mathcal{R}$, we account for preparation, calibration, decomposition, and per-budget factor materialization:

$$C_{\mathrm{LLift}}(\mathcal{R}) = C_{\mathrm{prep}} + C_{\mathrm{cal}} + C_{\mathrm{SVD}} + \sum_{r \in \mathcal{R}} C_{\mathrm{materialize}}(r). \tag{23}$$

With one prepared model and fixed activation transforms, we apply weighted SVD to each target matrix of the prepared checkpoint once and materialize factors at each budget. We compute GPU-hours as elapsed stage time times the number of devices held, counting shared work once and excluding quality evaluation.

C.1 Preparation and Training Resources

Table 22 reports preparation costs for all five models on A100-80GB GPUs, covering WT2 whitening, basis construction, optimization, and merging $W^\star = Q_0 B^\star + R_0$. We compute stage costs from recorded runtimes and log-based estimates, including loading and saving only where timed. Deployment costs are accounted for separately in Eq. 23.

Table 22: LLift preparation-stage costs in A100-80GB GPU-hours. Subtotals are computed before rounding.

Model	Whitening	Bases	Optimization	Merge	Subtotal
LLaMA-7B	0.40	0.09	1.56	0.02	2.07
Qwen3-8B	0.55	0.10	1.75	0.01	2.42
Gemma-2-9B	0.87	0.15	1.43	0.01	2.47
LLaMA-13B	3.97	0.15	4.97	0.02	9.11
LLaMA-30B	11.48	0.55	8.17	0.05	20.25

Optimization uses four GPUs, while basis construction and merging use one. Whitening uses one GPU for the 7B/8B/9B models and holds four for LLaMA-13B/30B, including during single-GPU transform derivation.

On four A100-80GB GPUs, FlexRank (Table 2) trains on about 12M FineWeb-Edu token positions and takes 27.0 GPU-hours, of which 24.0 are per-layer profiling for its rank search. LLift's preparation takes 2.07 GPU-hours with about 11.5M WT2 token positions. Because the corpora differ, these costs describe each method's own workflow rather than a matched comparison.

Table 23 reports optimization wall-clock time, throughput, and peak allocated GPU memory. All models use BF16, ZeRO-2, sequence length 1024, and global batch size 64 on four A100-80GB GPUs. The 7B/8B/9B models use microbatch size four and four accumulation steps per GPU. LLaMA-13B/30B use CPU optimizer offload, microbatch size one, and 16 accumulation steps. We estimate token positions per second per GPU as sample throughput $\times 1024/4$, including repeated epochs and padding. Peak memory covers rank 0 during optimization and excludes host memory.

Table 23: LLift optimization resources on four A100-80GB GPUs. Peak denotes allocated GPU memory on rank 0.

Model	Steps	Time (s)	Positions/s/GPU	Peak (GiB)	CPU offload
LLaMA-7B	176	1405.4	2045	32.18	No
Qwen3-8B	156	1573.2	1601	47.74	No
Gemma-2-9B	152	1289.1	1887	60.67	No
LLaMA-13B	176	4472.4	643	42.89	Yes
LLaMA-30B	176	7350.2	391	41.12	Yes

C.2 Budget-Specific Fine-Tuning and Deployment Timings

We follow SVD-LLM's sequential LoRA procedure, adapting U and then V on Alpaca for each budget (Appendix B.7). This takes approximately 5.4 GPU-hours per budget on average, totaling 16.1 GPU-hours for $r \in \{0.8, 0.6, 0.4\}$, excluding compression and deployment.

Table 24 reports estimated pipeline costs for the matched LLaMA-7B recovery comparison in Appendix B.7. Pipeline costs include preparation or recovery, deployment calibration, and SVD compression. LLift requires 3.13 A100 GPU-hours for three budgets, compared with 5.10 when each recovery run receives LLift's training exposure. Training positions are summed across runs. We sum stage-level GPU-hours, count shared work once, and exclude evaluation. The 2.07 GPU-hours in Table 22 cover preparation alone. LLift supports additional deployment budgets without further training.

Table 24: LLaMA-7B training and pipeline costs for three deployment budgets ($r = 0.8/0.6/0.4$), in A100 GPU-hours.

Method	Training runs	Positions (M)	Optimization	Pipeline
LLift	1	11.50	1.55	3.13
Recovery	3	34.50	4.00	5.10

LLift's optimization time comes from a profile with matching factor shapes and training steps. Both methods use the decomposition cost recorded for a four-budget batched SVD run.

On A100 GPUs, one-shot ZS-SVD and ZS-SVD $5\times$ take 0.25 and 2.11 GPU-hours, respectively, for one deployment budget, excluding LLift preparation.

C.3 Inference Resources

We benchmark dense and factorized shapes with random FP16 weights on one A100-80GB GPU using eager Hugging Face execution with SDPA and no compilation. The factorized shapes use the SVD-LLM rank rule and two-factor layout shared by LLift, with parameter counts matching all 20 model–budget configurations. This benchmark measures the corresponding structures, without timing the trained compressed checkpoints.

We measure prefill latency for 2048 tokens at batch size one. Decode generates 128 tokens from a four-token prompt with KV caching and greedy decoding at batch sizes 1, 16, and 64. Each setting uses three warmup iterations and ten timed iterations. Table 25 reports weight memory, peak allocated GPU memory during decoding at batch size 64, prefill latency, and aggregate decode throughput.

Dense-to-compressed weight-memory ratios range from 1.20 to 2.45. Decode throughput is 0.68–0.93 times the dense baseline, while prefill latency varies with model and budget. On the shared node, repeated LLaMA-7B runs differ by up to approximately 18% in decode throughput. Timing equivalence with trained compressed checkpoints remains unverified.

Table 25: Inference resources for dense and factorized shapes with random FP16 weights on one A100-80GB GPU. Peak denotes allocated GPU memory during decoding at batch size 64.

Model	Budget	Memory (GiB)		Prefill (ms)	Decode (tokens/s)		
		Weights	Peak		$B = 1$	$B = 16$	$B = 64$
LLaMA-7B	Dense	12.55	16.73	157	40.4	624.8	2505.5
	0.8	10.14	14.32	210	31.7	456.2	1831.6
	0.6	7.72	11.91	264	31.3	459.9	1728.4
	0.4	5.31	9.50	173	30.6	459.8	1860.1
Qwen3-8B	Dense	15.26	16.62	183	26.5	419.0	1689.2
	0.8	12.67	14.03	412	23.0	339.7	1332.2
	0.6	10.08	11.44	336	22.8	333.3	1332.0
	0.4	7.49	8.85	195	22.7	341.9	1353.1
Gemma-2-9B	Dense	17.21	20.25	235	22.1	353.4	1410.1
	0.8	14.11	17.14	513	18.3	278.4	1004.0
	0.6	11.01	14.04	227	18.3	271.8	1116.8
	0.4	7.91	10.94	216	18.8	275.8	1099.9
LLaMA-13B	Dense	24.24	30.75	274	27.7	407.2	1646.2
	0.8	19.51	26.02	331	21.0	312.5	1240.1
	0.6	14.79	21.29	427	22.1	336.5	1120.9
	0.4	10.06	16.57	203	21.1	325.7	1309.3
LLaMA-30B	Dense	60.59	73.19	665	17.8	261.0	836.2
	0.8	48.62	61.23	910	14.8	210.2	639.4
	0.6	36.67	49.27	558	14.7	224.4	777.7
	0.4	24.71	37.31	993	13.9	200.2	598.1

D GSM8K GENERATION

We evaluate Qwen3-8B on GSM8K (Cobbe et al., 2021) to assess multi-step generation. We use the `gsm8k` task in lm-eval 0.4.11 on all 1,319 test examples, with five-shot prompting, greedy decoding, and seed 42. All models use the same prompts without a chat template and run in FP16 with vLLM 0.17.1 and a maximum total sequence length of 2048 tokens. We reconstruct each compressed weight as the product of its two factors before evaluation. Table 26 reports strict exact-match accuracy. The original dense model achieves 88.63% accuracy. LLift outperforms SVD-LLM and ZS-SVD at $r = 0.8/0.6$. All methods approach zero at $r = 0.4$. These results show gains in compressed-model generation, although substantial degradation remains.

Table 26: Qwen3-8B GSM8K strict exact-match accuracy (%, ↑).

Method	$r = 0.8$	$r = 0.6$	$r = 0.4$
SVD-LLM	24.94	0.00	0.00
ZS-SVD	13.19	0.00	0.00
LLift	64.67	17.13	0.38