

000 001 002 003 004 005 GLOWQ: GROUP-SHARED LOW-RANK APPROXIMA- 006 TION FOR QUANTIZED LLMs 007 008 009

010 **Anonymous authors**
011

012 Paper under double-blind review
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029

ABSTRACT

030 Quantization techniques such as BitsAndBytes (Dettmers et al., 2022), AWQ (Lin
031 et al., 2024), and GPTQ (Frantar et al., 2023) are widely used as a standard method
032 in deploying large language models but often degrades accuracy when using low-
033 bit representations, e.g., 4 bits. Low-rank correction methods (e.g., LQER (Zhang
034 et al., 2024a), QERA (Zhang et al., 2024b), ASER (Zhao et al., 2025)) has been
035 proposed to mitigate this issue, however, they restore all layers and insert error-
036 correction modules into every decoder block, which increases latency and memory
037 overhead. To address this limitation, we propose GlowQ, a group-shared low-rank
038 approximation for quantized LLMs that caches a single shared right factor per
039 input-sharing group and restores only the groups or layers that yield the highest
040 accuracy benefit. GlowQ computes the high-precision projection once per input-
041 sharing group and reuses it across its modules, reducing parameter and memory
042 overhead, and retaining the expressivity of layer-specific corrections. We also
043 propose a selective variant, GlowQ-S, that applies the cached shared module only
044 where it provides the largest benefit. Compared with strong baselines, our
045 approach reduces TTFB by 5.6% and increases throughput by 9.6% on average,
046 while reducing perplexity on WikiText-2 by 0.17% and increasing downstream
047 accuracy by 0.42 percentage points. The selective model GlowQ-S further re-
048 duces latency, cutting TTFB by 23.4% and increasing throughput by 37.4%, while
049 maintaining accuracy within 0.2 percentage points on average.
050
051

1 INTRODUCTION

052 As large language models (LLMs) grow in width and depth, the cost of serving and adapting
053 them becomes a primary bottleneck for real-world use. Compression by post-training quanti-
054 zation (PTQ) alleviates memory and bandwidth pressure without altering the model architec-
055 ture, and has matured through methods such as GPTQ Frantar et al. (2023), AWQ Lin et al.
056 (2024), BITSANDBYTES Dettmers et al. (2022). A complementary thread augments quantized
057 weights with a small high-precision, low-rank term so that $\tilde{W} \approx W_q + AB$ and the inference output
058 is corrected by adding $A(BX)$ Zhang et al. (2024a;c); Zhao et al. (2025); Zhang et al. (2024b).
059 These lines enable competitive quality at quantized weights across modern transformer stacks.
060

061 Most low-rank compensation pipelines attach an independent (A, B) module to each layer or pro-
062 jection and evaluate the high-precision projection BX repeatedly along the network. This design
063 (i) duplicates the same expensive computation for modules that ingest the same input tensor, (ii)
064 increases memory traffic by materializing multiple BX values, and (iii) selects subspaces with ob-
065 jectives that often ignore the strong anisotropy of real activations Ethayarajh (2019); Godey et al.
066 (2024), misallocating limited rank to rarely used directions. As a result, the accuracy-efficiency
067 trade-off is weaker than necessary, especially under strict latency budgets.
068

069 We propose GlowQ, Group-Shared Low-Rank Approximation for Quantized LLMs. As illustrated
070 in Fig. 1, GlowQ treats modules that share the same input as a group Vaswani et al. (2017), learns
071 a single shared right factor B_{shared} for that group, and keeps module-specific left factors $\{A_i\}$. At
072 inference, it computes $R := B_{\text{shared}}X$ once per group and reuses it via A_iR , turning many large BX
073 multiplications into several cheap matrix-vector updates. To align limited rank with how inputs are
074 actually used, we adopt a covariance-aligned objective that emphasizes frequently visited directions.
075

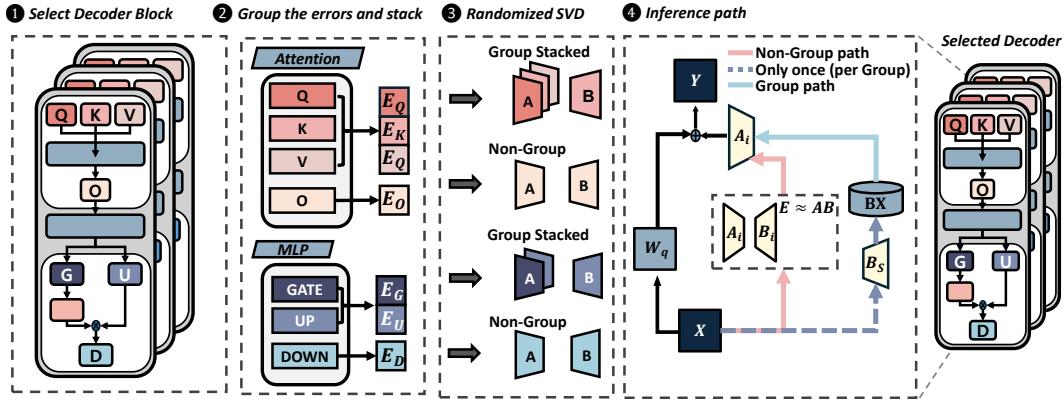


Figure 1: GlowQ Overview

Finally, a Selective Restore policy enables only high-payoff groups or layers under a deployment budget. Molchanov et al. (2019)

When modules share the input dimension, the joint least-squares problem with a single right factor is equivalent to approximating the vertical stack of module-wise error matrices; its minimizers are characterized by the right singular structure of the stacked matrix (“stacked SVD”). We then connect a usage-weighted risk $\min_{\mathbf{A}, \mathbf{B}} \|\mathbf{E}_{\text{cat}} - \mathbf{AB}\|_F^2$ to a right-weighted Frobenius norm, which yields a covariance-aligned objective whose global solution is governed by the SVD of the whitened error, where the whitened errors are those rescaled by the input covariance. This provides both a rationale for a shared \mathbf{B} and a principled way to steer it toward data-preferred axes.

To avoid forming tall whitened matrices, we introduce a QR-reduced randomized SVD routine: a thin QR compresses the stacked error into a $d \times d$ core; randomized SVD with oversampling and power iterations extracts the dominant right subspace; balanced recovery returns (A^*, B^*) with improved numerical stability. The solver drops into our grouping and caching runtime with no extra architectural changes.

- **Group-level shared- \mathbf{B} .** We formalize input-sharing groups and show that one shared right factor per group suffices for the joint least-squares objective, enabling one-shot \mathbf{BX} and multi-module reuse (Sec. 3).
- **Data-aware alignment.** We derive a covariance-aligned objective by bridging usage-weighted risk and a right-weighted Frobenius criterion; its global minimizer aligns the shared right subspace with data-preferred directions (Sec. 3.1).
- **QR-reduced RSVD.** We present a practical pipeline that performs QR reduction to a small core and applies randomized SVD with balanced factor recovery, avoiding tall whitened matrices while preserving accuracy (Sec. 3.2).
- **Caching & Selective Restore.** We implement a deployment path that caches $R = B_{\text{shared}} X$ once per group and activates only important groups/layers, translating algorithmic savings into latency/throughput gains (Sec. 3.3).
- **Empirical gains over strong baselines.** Across the evaluated model families and benchmarks, GlowQ consistently improves both efficiency and accuracy: it reduces time-to-first-byte (TTFB) by 5.6% and increases throughput by 9.6% on average, while reducing WikiText-2 perplexity by 0.17% and increasing downstream accuracy by 0.42 percentage points. The selective variant, GlowQ-S, further lowers latency, cutting TTFB by 23.4% and increasing throughput by 37.4%, while maintaining accuracy within 0.2 percentage points of full GlowQ.

2 RELATED WORK

Post-training quantization (PTQ). Today’s PTQ methods span a variety of designs that recover accuracy at the quantization stage without changing the model structure or runtime path.

108 GPTQ Frantar et al. (2023) uses second-order information to directly fit quantized weights and pre-
 109 serve layer outputs; AWQ Lin et al. (2024) protects important channels based on activation statistics
 110 via rescaling. In the BITSANDBYTES family, LLM.INT8() Dettmers et al. (2022) uses vector-wise
 111 quantization with an outlier-aware mixed-precision path where most channels run in INT8. These
 112 methods constitute standard baselines for LLM lightweighting.

113 **Quantization error correction via low-rank compensation.** Prior work shows that post-
 114 quantization errors can be effectively reduced by adding a low-rank term to the quantized weights
 115 or outputs. LQER approximates the per-layer quantization error as $E \approx AB$ and adds a high-
 116 precision correction without changing the inference graph (Zhang et al., 2024a). ZEROQUANT-V2
 117 systematizes low-rank compensation (LoRC) within PTQ pipelines and demonstrates that a small-
 118 rank correction can recover accuracy at low bit-widths (et al., 2023). QERA derives a closed-form,
 119 output-error-centric formulation that clarifies when low-rank correction benefits PTQ/PEFT (Zhang
 120 et al., 2024b). ASER combines a whitened-SVD-style low-rank corrector with activation smoothing
 121 to stabilize low-bit regimes (Zhao et al., 2025). While these works justify the AB correction prin-
 122 ciple, most deploy independent (A_ℓ, B_ℓ) at every layer and recompute the high-precision product
 123 $A_\ell(B_\ell X)$ for all layers and tokens, which increases latency and memory traffic; moreover, attaching
 124 a low-rank module to every layer inflates GPU memory usage.

125 **Stacked/collective SVD for a shared right subspace.** The idea of factorizing multiple matrices
 126 with a shared latent factor is established in collective/joint matrix factorization: when several
 127 matrices share the same input dimension, one can vertically concatenate their blocks and fit a single
 128 right subspace while allowing matrix-specific left factors (Singh & Gordon, 2008). Recent analyses
 129 also study the optimal recovery of shared singular subspaces across matrices (Ma & Ma, 2024).
 130 We adopt this principle for input-sharing modules in LLMs: we stack group-wise error blocks into
 131 E_{cat} and learn one B_{shared} per group. At inference, we compute the right projection once per group,
 132 $R = B_{\text{shared}}X$, cache it, and let each module apply only the lightweight left multiplication $A_i R$.
 133 This reduces high-precision matmuls and the number of resident correction parameters compared to
 134 layer-wise independent AB .

135 **Covariance-aligned selective restoration.** Because inputs are anisotropic, a plain stacked ob-
 136 jective may learn right subspaces misaligned with data-preferred directions. We therefore adopt a
 137 covariance-aligned (whitened) formulation, measuring residual error in the input-covariance met-
 138 ric so that the shared subspace is guided toward meaningful axes (Golub & Loan, 2013; Srebro &
 139 Jaakkola, 2003). Not all layers require restoration; following pruning-inspired saliency, we activate
 140 only the most beneficial groups under a budget, using (i) an SVD energy-capture score ($\|A\|_F^2$ per
 141 group), (ii) a normalized error ratio $\|E_g\|_F/\|W_g\|_F$ (Nagel et al., 2020; Banner et al., 2018), and (iii)
 142 a layer-order fallback when signals are weak. Coupled with the shared-right-subspace design and
 143 cached $R = B_{\text{shared}}X$, this selective restore achieves stronger accuracy-latency-memory trade-offs
 144 than per-layer low-rank baselines at the same cost.

3 METHOD: GLOWQ

148 In this section, we introduce our method, Group-Shared Low-Rank Approximation For Quantized
 149 LLMs (GlowQ). Prior low-rank restoration often (i) restores all layers and (ii) multiplies a per-
 150 layer low-rank module with activations, causing heavy overhead. We address both by (a) learning
 151 a **shared** right subspace for modules that share the same input and (b) **caching** the input projection
 152 once per group for reuse. We approximate each error matrix and its vertical concatenation by a rank-
 153 r factorization: $E_i \approx A_i \mathbf{B}$ and $E_{\text{cat}} \approx \mathbf{AB}$, where $\mathbf{A} = [\mathbf{A}_1; \dots; \mathbf{A}_m]$ and \mathbf{B} is shared within a
 154 group. At inference, the correction for each module i takes the form $\mathbf{A}_i(\mathbf{BX})$, where the projection
 155 \mathbf{BX} is computed once for the entire group.

3.1 GROUPING QUANTIZATION-ERROR CORRECTION MODULES

157 We aim to find the optimal shared low-rank correction module, in particular a shared right factor \mathbf{B} .
 158 To this end, we first formalize the problem via an unweighted baseline (Sec. 3.1.1), and propose a
 159 data-aware objective that incorporates covariance alignment to overcome the limitation induced by
 160 input anisotropy (Sec. 3.1.2).

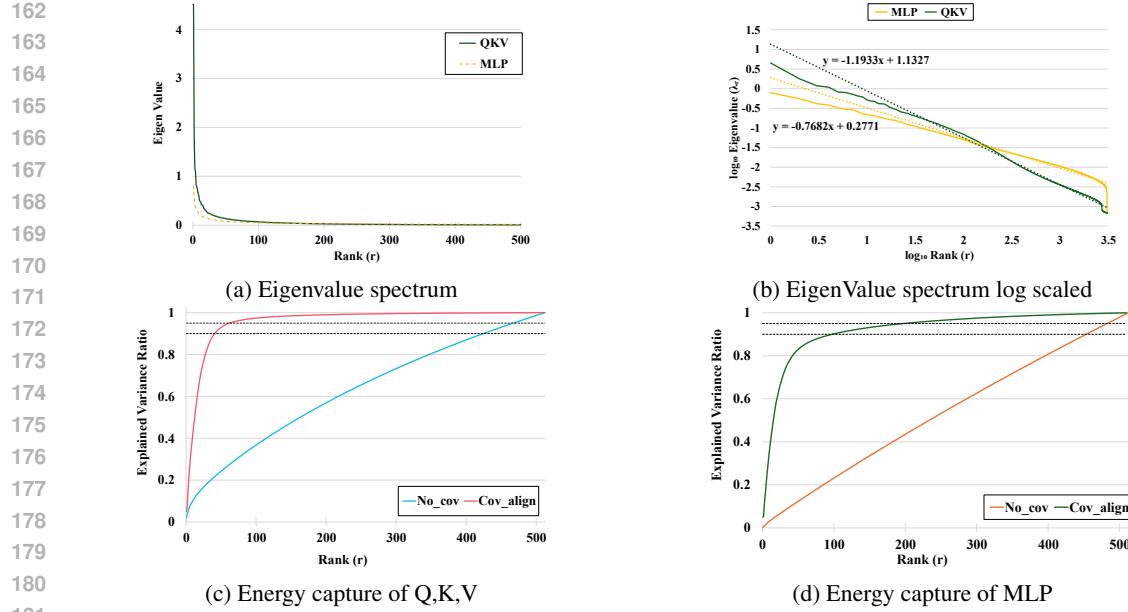


Figure 2: Input spectrum and energy-capture measurements. (a) We stream calibration samples through the model, collect the input activations at the target layer, and plot the eigenvalue spectrum of the empirical input covariance for the QKV and MLP groups, revealing a heavy-tailed profile. (b) The same spectra plotted in $\log_{10} \lambda_r$ - $\log_{10} r$ coordinates; dotted lines show least-squares fits over the approximately linear tail region, indicating power-law decay $\lambda_r \propto r^{-\alpha}$ with exponents $\alpha_{\text{MLP}} \approx 0.77$ and $\alpha_{\text{QKV}} \approx 1.19$. (c-d) For each group, we vertically stack the quantization-error matrices and plot the cumulative fraction of Frobenius energy recovered by the best rank- r approximation. We show both the unweighted baseline (No cov) and the covariance-aligned variant that weights errors by the observed inputs (Cov align). Horizontal dashed lines mark 90% and 95% energy capture.

3.1.1 UNWEIGHTED BASELINE: STACKED SVD

Let modules $i = 1, \dots, m$ share the same input dimension d . For error matrices $E_i \in \mathbb{R}^{O_i \times d}$, define the vertical concatenation

$$\mathbf{E}_{\text{cat}} := [E_1^T \ \dots \ E_m^T] \in \mathbb{R}^{d \times (\sum_i O_i)}, \quad \mathbf{A} := [\mathbf{A}_1^T \ \dots \ \mathbf{A}_m^T] \in \mathbb{R}^{r \times (\sum_i O_i)}. \quad (1)$$

We seek a shared right factor $\mathbf{B} \in \mathbb{R}^{r \times d}$ and blocks $\mathbf{A}_i \in \mathbb{R}^{O_i \times r}$ that minimize

$$\min_{\mathbf{A}, \mathbf{B}} \|\mathbf{E}_{\text{cat}} - \mathbf{AB}\|_F^2. \quad (2)$$

Proposition 1 (Shared-B is optimal). For modules that share the same input, jointly fitting with a single right factor \mathbf{B} is equivalent to one low-rank fit of the stacked matrix \mathbf{E}_{cat} . By Eckart-Young-Mirsky, an optimal \mathbf{B} spans the top- r right-singular subspace of \mathbf{E}_{cat} ; allowing per-module \mathbf{B}_i adds no extra expressivity because any differences can be absorbed into invertible reparameterizations of \mathbf{A}_i . Hence, a single shared \mathbf{B} is sufficient and optimal for the group. (Proof and identifiability details are deferred to Appendix A.1.)

Real inputs are anisotropic, which can be diagnosed by the eigenvalue spectrum of the covariance Σ_x . Fig. 2a exhibits a heavy-tailed profile, with an abrupt initial drop followed by a long tail, indicating that the usage of the representation space is strongly concentrated in a small number of axes. Under such a distribution, the relative importance between frequently used directions and the remaining ones diverges markedly. To quantify this behavior, Fig. 2b plots the eigenvalue spectra of the empirical input covariance for the MLP and QKV groups in $\log_{10} \lambda_r$ - $\log_{10} r$ scale: for each group we sort the eigenvalues $\{\lambda_r\}$ in descending order and plot $\log_{10} \lambda_r$ versus $\log_{10} r$. The dotted lines show least-squares linear fits over the approximately linear tail region, revealing power-law decay $\lambda_r \propto r^{-\alpha}$ with exponents $\alpha_{\text{MLP}} \approx 0.77$ and $\alpha_{\text{QKV}} \approx 1.19$, which quantitatively confirms the heavy-tailed, anisotropic input statistics that motivate our covariance-aligned objective.

216 However, the unweighted cluster SVD selects the shared right subspace purely from the geometry
 217 (variance structure) of the error matrices. This can misalign the selected subspace with the axes
 218 preferred by the data; at a fixed rank, such a misalignment reduces energy capture and weakens
 219 consistency within the group. Alignment arises naturally only in restrictive cases, such as isotropic
 220 input or near-simultaneous diagonalization.

221 Therefore, to treat anisotropy fairly, we should evaluate errors in a coordinate system where all
 222 directions carry equal usage. In such a space, frequently used directions are not under-weighted,
 223 and rarely used directions are not over-weighted, so the learned shared right subspace aligns better
 224 with the data-preferred axes.

226 3.1.2 DATA-AWARE COVARIANCE ALIGNMENT

228 The evidence in Fig. 2a shows strong input anisotropy; hence reconstruction should account not only
 229 for the geometry of error matrices but also for how inputs are actually used. For any factors (\mathbf{A}, \mathbf{B})
 230 and residual $\mathbf{M} := \mathbf{E}_{\text{cat}} - \mathbf{AB}$, the expected loss under the usage distribution is

$$232 \mathbb{E} \|\mathbf{M} \mathbf{x}\|_2^2 = \text{tr}(\mathbf{M} \Sigma_{\mathbf{x}} \mathbf{M}^{\top}) = \|\mathbf{M} \Sigma_{\mathbf{x}}^{1/2}\|_F^2. \quad (3)$$

234 which follows from the standard quadratic-form identity together with the Frobenius-trace iden-
 235 tity (Petersen & Pedersen, 2012). To balance direction-wise usage, we whiten by $\Sigma_{\mathbf{x}}^{1/2}$ so that the
 236 selected shared right subspace is steered toward axes preferred by the data.

238 Using the definitions from Sec. 3.1.1, we adopt the right-weighted objective

$$240 \min_{\mathbf{A}, \mathbf{B}} \|(\mathbf{E}_{\text{cat}} - \mathbf{AB}) \Sigma_{\mathbf{x}}^{1/2}\|_F^2 \equiv \min_{\mathbf{A}, \mathbf{B}} \|\tilde{\mathbf{E}} - \mathbf{AB}\|_F^2, \quad \tilde{\mathbf{E}} := \mathbf{E}_{\text{cat}} \Sigma_{\mathbf{x}}^{1/2}. \quad (4)$$

242 In the isotropic case ($\Sigma_{\mathbf{x}} \propto I$), Eq. 4 reduces to the unweighted baseline in Sec. 3.1.1. Empirically,
 243 Fig. 2c and 2d shows that, at a fixed rank, whitening yields substantially faster growth of the cumu-
 244 lative energy capture compared to the unweighted variant, indicating better alignment of the learned
 245 shared right subspace with data-preferred directions.

246 **Proposition 2 (Usage-weighted risk equals a right-weighted reconstruction error).** When in-
 247 puts are centered and have covariance $\Sigma_{\mathbf{x}}$, the model’s expected loss equals the residual energy aver-
 248 aged over draws from the input distribution. Equivalently, it is the residual measured after weighting
 249 columns according to how frequently and how strongly each input direction is used (as determined
 250 by $\Sigma_{\mathbf{x}}$). Therefore, minimizing the usage-weighted risk is exactly the same optimization as min-
 251 imizing the right-weighted reconstruction error in Eq. 4. A full derivation and the nonzero-mean
 252 case are deferred to Appendix A.2.

253 **Proposition 3 (Covariance-aligned minimizer)** The global minimizers $(\mathbf{A}^*, \mathbf{B}^*)$ of Eq. 4 are
 254 given by the rank- r SVD of the whitened error matrix $\tilde{\mathbf{E}} = \mathbf{E}_{\text{cat}} \Sigma_{\mathbf{x}}^{1/2}$. In particular, the opti-
 255 mal shared right subspace, $\text{row}(\mathbf{B}^*)$, is spanned by the top- r right singular vectors of $\tilde{\mathbf{E}}$; this is
 256 the standard Eckart-Young-Mirsky solution specialized to the whitened problem (Eckart & Young,
 257 1936) Golub & Loan, 2013.

259 3.2 SCALABLE IMPLEMENTATION VIA QR-REDUCED RANDOMIZED SVD

261 We present an implementation that solves the covariance-aligned objective

$$263 \min_{\mathbf{A}, \mathbf{B}} \|(\mathbf{E}_{\text{cat}} - \mathbf{AB}) \Sigma_{\mathbf{x}}^{1/2}\|_F^2 \quad (5)$$

265 without forming the tall whitened matrix. The method follows three steps: (i) **QR reduction** to
 266 compress the tall matrix into a $d \times d$ core, (ii) **Randomized SVD (RSVD)** on the core to capture
 267 the top- r right subspace, and (iii) **balanced recovery** to obtain $(\mathbf{A}^*, \mathbf{B}^*)$. This yields practical ad-
 268 vantages such as avoiding materialization of huge matrices, lower compute/memory cost, improved
 269 numerical stability via balanced factors, and direct compatibility with the caching/Selective-Restore
 pipeline in Sec. 3.3.

270 3.2.1 ALGORITHM & COMPLEXITY
271

272 Using a thin QR of \mathbf{E}_{cat} , we reduce the covariance-aligned objective to a $d \times d$ core (Alg. 1); a full
273 SVD on the core costs $\mathcal{O}(d^3)$, whereas randomized sketching recovers the leading right subspace in
274 $\mathcal{O}(d^2(r+p) + q d^2(r+p))$ time. Here, p denotes oversampling (extra sketch columns) and q denotes
275 the number of power iterations used to sharpen the subspace.

276 **Algorithm 1** Covariance-aligned QR reduction and randomized SVD on the core
277

278 **Require:** Stacked error $\mathbf{E}_{\text{cat}} \in \mathbb{R}^{m \times d}$, covariance $\Sigma_{\mathbf{x}} \succeq 0$, target rank r , oversampling p , power
279 iters q

280 **Ensure:** Low-rank factors $(\mathbf{A}^*, \mathbf{B}^*)$ for the covariance-aligned objective

- 281 1: **Thin QR of \mathbf{E}_{cat} :** compute $\mathbf{Q}_e \mathbf{R}_e = \mathbf{E}_{\text{cat}}$ with $\mathbf{Q}_e^\top \mathbf{Q}_e = \mathbf{I}_d$
- 282 2: **Core construction:** set $\mathbf{M} \leftarrow \mathbf{R}_e \Sigma_{\mathbf{x}}^{1/2} \in \mathbb{R}^{d \times d}$
- 283 3: **Random sketch / range finding:** draw $\Omega \sim \mathcal{N}(0, 1)^{d \times (r+p)}$, set $\mathbf{Y} \leftarrow \mathbf{M} \Omega$; optionally do q
284 power steps $\mathbf{Y} \leftarrow \mathbf{M}(\mathbf{M}^\top \mathbf{Y})$
- 285 4: **Orthonormalize:** $\mathbf{Q} \leftarrow \text{orth}(\mathbf{Y}) \in \mathbb{R}^{d \times (r+p)}$
- 286 5: **Compressed SVD:** $\mathbf{B}_{\text{small}} \leftarrow \mathbf{Q}^\top \mathbf{M}$; compute $\mathbf{B}_{\text{small}} = \tilde{\mathbf{U}} \Sigma \mathbf{V}^\top$
- 287 6: **Lift left factor:** $\mathbf{U} \leftarrow \mathbf{Q} \tilde{\mathbf{U}}$
- 288 7: **Truncate (top- r) & balance:** keep $(\mathbf{U}_r, \Sigma_r, \mathbf{V}_r)$ and set $\hat{\mathbf{A}}^* \leftarrow \mathbf{U}_r \Sigma_r^{1/2}$, $\hat{\mathbf{B}}^* \leftarrow \Sigma_r^{1/2} \mathbf{V}_r^\top$
- 289 8: **Lift to original variables:** $\mathbf{A}^* \leftarrow \mathbf{Q}_e \hat{\mathbf{A}}^*$, $\mathbf{B}^* \leftarrow \hat{\mathbf{B}}^* \Sigma_{\mathbf{x}}^{-1/2}$ \triangleright use a pseudoinverse if $\Sigma_{\mathbf{x}}$ is
290 singular

291
292 By left-orthogonal invariance of the Frobenius norm, the QR reduction collapses the tall- m problem
293 to a $d \times d$ core without loss for the covariance-aligned objective (formal proof in Appendix A.3;
294 (Golub & Loan, 2013)). Randomized sketching on the core provides an efficient and accurate es-
295 timate of the leading right subspace with controllable bias via (p, q) ; we summarize theoretical
296 guarantees in Appendix A.4 and present empirical runtime-accuracy trade-offs (vs. exact SVD) in
297 Sec. D.3.1 (Halko et al., 2011). Balanced recovery yields

$$298 \hat{\mathbf{A}}^* = \mathbf{U}_r \Sigma_r^{1/2}, \quad \hat{\mathbf{B}}^* = \Sigma_r^{1/2} \mathbf{V}_r^\top, \quad \|(\mathbf{E}_{\text{cat}} - \mathbf{A}^* \mathbf{B}^*) \Sigma_{\mathbf{x}}^{1/2}\|_F = \|\mathbf{M} - \mathbf{U}_r \Sigma_r \mathbf{V}_r^\top\|_F \quad (6)$$

300 and the resulting \mathbf{B}^* serves as the shared right factor used in Sec. 3.3 for once-per-group caching
301 ($R = \mathbf{B}_{\text{shared}} \mathbf{X}$) and Selective Restore.

303 3.3 CACHING AND SELECTIVE RESTORE
304

305 The group-shared factorization implies that modules within the same input-sharing group all rely on
306 the *same* right-side projection $\mathbf{X} \mathbf{B}_{\ell, \text{shared}}^\top$. Naively evaluating this projection for every module recre-
307 ates the primary inefficiency of layer-wise correction, i.e., multiple high-precision matrix-vector
308 multiplications along the critical path. To translate the theoretical shared structure into practical
309 inference gains, GlowQ introduces a caching mechanism that computes the right-sided projection
310 *once per group*, and a complementary selective-restore policy that activates correction only at groups
311 offering the largest accuracy benefit under a deployment budget.

312 For each layer group G_ℓ that shares the same input dimension, we compute a single intermediate

$$313 R_\ell := \mathbf{X} \mathbf{B}_{\ell, \text{shared}}^\top \in \mathbb{R}^{B \times T \times r} \quad (7)$$

314 once per group and reuse it across all modules in the group. Each module $i \in G_\ell$ then applies only
315 the small correction

$$316 y_i = \mathbf{W}_i^{(q)} \mathbf{X} + \mathbf{A}_{\ell, i} R_\ell, \quad (8)$$

317 where $\mathbf{A}_{\ell, i} \in \mathbb{R}^{O_i \times r}$ and $\mathbf{B}_{\ell, \text{shared}} \in \mathbb{R}^{r \times I}$. We adopt an anchor policy to materialize R_ℓ exactly
318 once and consume it a fixed number of times: in attention, q is the anchor and (k, v) are consumers;
319 in MLP, gate is the anchor and up is the consumer. Solo modules that do not share inputs (e.g.,
320 o_proj, down_proj) compute $R_i := \mathbf{X} \mathbf{B}_i^\top$ on the fly without reuse.

322 Given a latency or memory budget, we rank all candidate units (groups or solo layers) by an im-
323 portance score and activate only the top k . Importance is measured using two metrics: a GSVD-
324 based energy-capture score (Eq. 9) after covariance alignment (Paige & Saunders, 1981; Jolliffe &

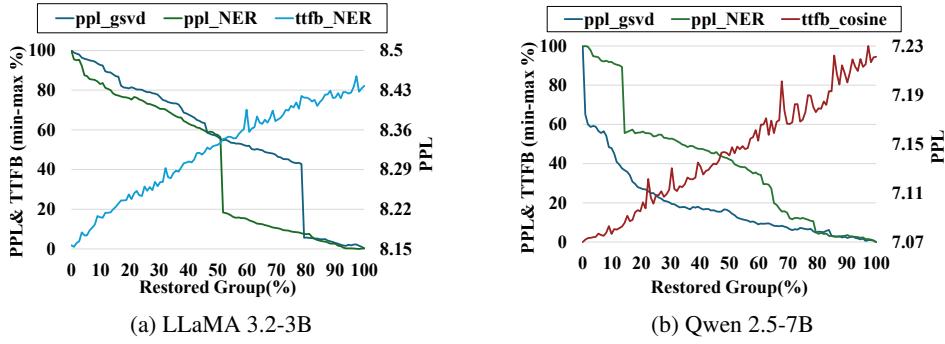


Figure 3: Perplexity (PPL) and time-to-first-byte (TTFB) versus the fraction of restored groups.

Cadima, 2016; Halko et al., 2011), and a normalized error ratio (Eq. 10) (Malinovskii et al., 2024; Dong et al., 2019). At runtime, we apply the cached low-rank correction only to the selected units, skipping inactive ones; because the cache is materialized only for active groups, selective restore naturally complements group-shared caching.

$$g_{\text{ec}}(u) = \frac{\sum_{j=1}^r \sigma_j(\mathbf{M}_u)^2}{\|\mathbf{M}_u\|_F^2} \quad (9) \quad g_{\text{ner}}(u) = \frac{\|\mathbf{E}_u\|_F^2}{\|\mathbf{W}_u\|_F^2} \quad (10)$$

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

We evaluate LLaMA 3 (3.2–3B, 3.1–8B) Dubey et al. (2024), LLaMA 2 (7B, 13B) Touvron et al. (2023), Qwen 2.5 (7B, 14B) Yang et al. (2024), Qwen 3 (8B, 14B) Yang et al. (2025), OPT (1.3B, 6.7B) Zhang et al. (2022), Mistral 7B Jiang et al. (2023), and Qwen1.5-MoE-A2.7B Bai et al. (2023) with Vicuna reported only in ablations.

All models use W4A16 (int4 weights, fp16 activations) with group size 128; the rank is fixed at 64. Calibration uses 64 sequences of length 2048 shared across methods, with no fine-tuning unless a baseline requires it. We compare GlowQ and GlowQ-S with various state-of-the-art baselines, including PTQ (BitsAndBytes Dettmers et al. (2022), AWQ Lin et al. (2024), GPTQ Frantar et al. (2023)) and error-correction methods in literature (L2QER Zhang et al. (2024a), ZeroQuant-V2 et al. (2023), QERA Zhang et al. (2024b)). All under the same protocol and recommended defaults.

We report perplexity on WikiText-2 Merity et al. (2016) and C4 Raffel et al. (2023), and zero-shot accuracy on ARC-E/ARC-C Clark et al. (2018), PIQA Bisk et al. (2019), HellaSwag Zellers et al. (2019), WinoGrande Sakaguchi et al. (2019), BoolQ Clark et al. (2019), and LAMBADA Paperno et al. (2016) via `lm-eval-harness` (defaults). We run the proposed method on A100 GPUs for covariance/SVD steps while inference is executed on an RTX 4090.

GlowQ-S Configuration GlowQ-S applies the cached correction only to a subset of groups, selected according to an importance score. Since different model families exhibit distinct restoration profiles, we adopt a model-specific scoring rule for GlowQ-S. We defer the full characterization of these curves and the selection policy to Section 4.6.

4.2 MAIN RESULTS: PERPLEXITY AND ZERO-SHOT ACCURACY

Perplexity. Table 1 reports test perplexity (lower is better) for WikiText-2 under a common protocol: W4A16 with int4 weight groups of 128 and a shared calibration set of 64 sequences at length 2048 for all methods. Overall, GlowQ achieves the best or tied-best perplexity on 9 of 11 model variants, including consistent gains on LLaMA 3 (3.2-3B/3.1-8B), Qwen 3 (8B/14B), Qwen 2.5-7B, and Mistral-7B. On Qwen 2.5-14B, GlowQ matches the strongest baselines. Exceptions occur on LLaMA 2-13B (where L2QER slightly leads) and OPT-1.3B (where QERA leads), while OPT 6.7B favors a pure PTQ path. These outcomes indicate that group-shared low-rank correction closes much of the int4 gap to FP16 across diverse architectures without task-specific tuning. Beyond the W4A16 setting, the lower block of Table 1 evaluates mixed-precision weight-activation quantization with W4A4 and W4A8. As expected, W4A4 increases perplexity for all methods, but GlowQ (and

378 Table 1: WikiText-2 test perplexity (lower is better). GlowQ-S restores 51% of layers for
 379 LLaMA 3.2-3B, while all other models use 50% restoration.

381 Method	Q config	382 LLaMA 2		383 LLaMA 3		384 Qwen 2.5		385 Qwen 3		Mistral	386 OPT	
		387 7B	388 13B	3.2-3B	3.1-8B	389 7B	390 14B	391 8B	392 14B	393 7B	394 1.3B	395 6.7B
384 FP16	-	5.48	4.90	7.81	6.24	6.86	5.29	9.73	8.64	5.32	14.62	10.85
385 BnB	NF4	5.64	4.97	8.29	6.66	7.10	5.64	9.97	8.88	5.51	15.16	10.94
386 AWQ	INT4, g128	5.61	4.97	8.24	6.64	7.11	6.17	10.19	9.00	5.51	15.22	11.23
387 GPTQ	INT4, g128	5.65	5.35	9.46	6.63	7.11	5.75	9.98	8.90	5.51	15.00	11.07
388 ZeroQuant-V2	INT4, g128	5.72	4.99	8.44	6.79	8.41	5.75	10.19	9.04	5.53	15.10	11.14
389 QERA	INT4, g128	5.61	4.98	8.22	6.64	8.09	5.64	10.07	8.85	5.48	14.85	11.00
390 L2QER	INT4, g128	5.68	4.94	8.30	6.75	8.14	5.66	10.07	8.85	5.46	15.30	11.16
391 GlowQ	INT4, g128	5.58	4.96	8.16	6.59	7.07	5.64	9.90	8.80	5.42	14.84	11.00
392 GlowQ-S	INT4, g128	5.60	4.96	8.22	6.62	7.09	5.68	9.97	8.89	5.45	15.00	11.00
393 L2QER	W4A4	5.90	5.18	9.42	7.65	9.11	6.52	10.76	9.36	5.73	27.40	11.32
394 L2QER	W4A8	5.69	4.95	8.31	6.76	8.15	5.67	10.11	8.86	5.47	14.90	11.00
395 GlowQ	W4A4	5.90	5.20	9.21	7.42	8.03	6.55	10.66	9.33	5.74	26.35	11.31
396 GlowQ-S	W4A4	5.92	5.20	9.25	7.45	8.05	6.61	10.72	9.37	5.79	27.42	11.33
397 GlowQ	W4A8	5.59	4.97	8.20	6.63	7.12	5.71	10.08	8.85	5.43	14.85	10.97
398 GlowQ-S	W4A8	5.60	4.97	8.24	6.64	7.13	5.77	10.10	8.92	5.48	14.99	10.99

401 GlowQ-S) remain competitive with or better than L2QER on most models, and the W4A8 configuration
 402 nearly recovers the W4A16 accuracy, indicating that our covariance-aware low-rank correction
 403 continues to be effective even under joint weight-activation quantization.

405 Table 2: Average accuracy (\uparrow) on seven downstream tasks and C4 perplexity (\downarrow).
 406

407 Method	Rank	408 LLaMA 3.2-3B		409 LLaMA 3.1-8B		410 Qwen 3-8B		411 Qwen 3-14B	
		412 Acc (\uparrow)	413 C4 (\downarrow)	414 Acc (\uparrow)	415 C4 (\downarrow)	416 Acc (\uparrow)	417 C4 (\downarrow)	418 Acc (\uparrow)	419 C4 (\downarrow)
416 FP16	-	67.14	10.30	73.29	9.00	71.48	14.52	74.10	13.08
417 ZeroQuant-V2		65.38	11.45	73.48	9.87	70.19	15.00	72.62	13.79
418 QERA		65.48	11.04	72.86	9.68	69.86	14.78	73.14	13.29
419 L2QER	64	66.19	11.04	72.43	9.63	69.52	14.82	73.24	13.80
420 GlowQ		66.90	10.98	73.33	9.59	70.71	14.60	73.84	13.26
421 GlowQ-S		66.33	11.07	72.62	9.78	70.29	14.77	73.24	13.48

422 **Overall quality.** Table 2 reports the zero-shot accuracy via *lm-eval-harness* along with the perplexity
 423 for the C4 dataset. Across four representative models (LLaMA 3.2-3B / 3.1-8B, Qwen 3
 424 8B / 14B), GlowQ attains the lowest C4 perplexity among quantized/error-corrected methods and de-
 425 liveries the strongest average zero-shot accuracy on LLaMA 3.2-3B and Qwen 3-8B/14B (ZeroQuant-
 426 V2 leads on LLaMA 3.1-8B). GlowQ improves over the best non-GlowQ baseline in the zero-shot
 427 accuracy by average +0.3%; in C4 perplexity, GlowQ improves by -0.2 ppl on average. Relative to
 428 FP16, the remaining C4 gap is +0.4 ppl on average, while average accuracy remains close to FP16
 429 across the board. The selective-restore variant (GlowQ-S) shows the expected efficiency trade-off:
 430 -0.55% on average accuracy and +0.15 ppl on average in C4 compared to GlowQ.

4.3 LATENCY AND THROUGHPUT BENEFITS FROM CACHING AND SELECTIVE RESTORE

431 **Latency on LLaMA 2 models.** Under a common generation protocol (3 prompts, batch=1,
 432 max_new_tokens=128, repeats=1, num_beams=1) and custom CUDA W4A16 kernels, we measure
 433 TTFB via a warm-start generate (max_new_tokens=1) and per-token decode latency using
 434 CUDA events (Table 3). We establish our baseline using a standard Layerwise method, which does
 435 not employ caching. This setup ensures a fair comparison, as both the Layerwise baseline and
 436 GlowQ utilize the identical custom CUDA W4A16 kernels compiled with the same optimization

Table 3: Latency comparison on LLaMA 2 models for Layerwise vs. GlowQ, GlowQ-S.

Models	Setting	TTFB(ms) ↓	tok/s ↑	Prefill(ms) ↓	Dec(ms/tok) ↓
LLaMA 2	7B	Layerwise	88.45	15.66	95.13
		GlowQ	82.66	17.12	92.23
		GlowQ-S	66.68	21.16	72.35
	13B	Layerwise	128.70	11.22	141.76
		GlowQ	122.78	12.33	136.53
		GlowQ-S	100.17	15.68	112.09
	Avg. ΔBX (%)		-5.57	+9.61	-3.37
	Avg. $\Delta R50$ (%)		-23.39	+37.44	-22.44
					-27.01

level, isolating the algorithmic impact of our caching strategy. Compared to this Layerwise baseline, GlowQ consistently reduces end-to-end latency across both sizes: on average TTFB drops by 5.51%, prefill time by 3.37%, and decode latency by 6.61%, yielding a 9.61% increase in throughput (tok/s).

Selective restore efficiency. The GlowQ-S, which are restoring about half of the units by an importance score, amplifies the gains: average TTFB, prefill, and decode fall by 23.39%, 22.44%, and 27.01%, respectively, and throughput increases by 37.44% over the Layerwise baseline.

4.4 MEMORY OVERHEAD AND EFFICIENCY ANALYSIS

On memory, GlowQ consistently uses less additional GPU memory than layer-wise restoration at the same rank r . This follows from maintaining a single shared right factor B_{shared} per input-sharing group and computing $R = B_{\text{shared}}X$ once per group for cache-and-reuse. Applying GlowQ-S further reduces overhead, yielding the flattest growth slope even at higher ranks. On accuracy, under an equal-memory budget in Fig. 4(b), GlowQ attains the lowest PPL, while GlowQ-S preserves PPL close to full GlowQ with substantially lower memory, consistently outperforming than layer-wise methods. Consequently, GlowQ is the preferred choice when maximizing performance within a fixed memory budget, whereas GlowQ-S offers a strong performance-efficiency compromise when memory constraints are tighter or latency minimization is prioritized.

4.5 COMPATIBILITY WITH PTQ METHODS AND GENERALIZATION TO MoE

We further examine the compatibility of GlowQ with diverse LLM configurations, focusing in particular on PTQ baselines and MoE architectures, as summarized in Table 4.

Layering GlowQ on top of PTQ baselines reduces perplexity by -0.59 ppl on average for GPTQ and -0.13 ppl on average for BnB. Improvements hold across both evaluated models in each setting, indicating consistent add-on gains independent of the underlying quantizer. GlowQ acts as an orthogonal, plug-and-play low-rank correction: it exchanges a small set of shared parameters for accuracy gains while remaining compatible with diverse PTQ pipelines.

On this MoE benchmark, GlowQ largely recovers the Wikitext-2 perplexity loss from 4-bit weight quantization and ends up only +0.02 PPL worse

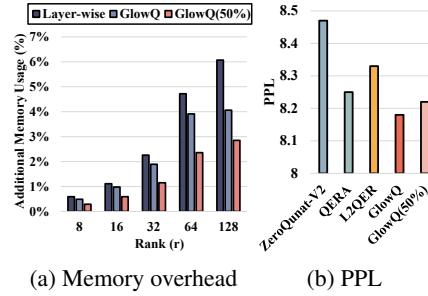


Figure 4: Comparison of memory and performance trade-off. (a) Memory overhead of different methods. (b) PPL at equal memory budget.

Table 4: Perplexity (\downarrow) on Wikitext-2 with and without GlowQ: dense models (top) and Qwen1.5-MoE-A2.7B (bottom).

Method	LLaMA 2-7B	LLaMA 3.2-3B
GPTQ	5.64	9.32
+GlowQ (on GPTQ)	5.60	8.19
BnB	5.64	8.29
+GlowQ (on BnB)	5.57	8.10
FP16 Quant only GlowQ Layerwise		
Qwen1.5-MoE-A2.7B	7.22	7.70
		7.41
		7.39

486 than the more expensive layer-wise low-rank baseline. The layer-wise variant attaches a separate
 487 error-correction module to every expert, whereas GlowQ uses a single shared right factor B_{shared} per
 488 group across experts and the shared MLP. The whitening-based alignment heatmaps in Fig. 7 and
 489 Fig. 8 show that expert-specific error subspaces are well aligned with this shared right subspace,
 490 explaining why the shared- B design can match layer-wise accuracy while reducing the memory
 491 footprint of the low-rank correction by about 63%. These results confirm that GlowQ remains ef-
 492 fective even on large MoE architectures.

493 Given the recent trend toward rotation-based saliency-aware PTQ and KV-cache compression,
 494 GlowQ can be viewed as a complementary low-rank correction layer that may be attached to strong
 495 PTQ baselines such as ROSAQ Yoon et al. (2025) and GuidedQuant Kim et al. (2025), and fur-
 496 ther extended to KV-cache compression frameworks like CommVQ Li et al. (2025); exploring such
 497 combinations remains an interesting direction for future work.

498

499 4.6 BEHAVIOR OF SELECTIVE RESTORATION ACROSS MODEL FAMILIES

500

501 Fig. 3 plots PPL and TTFB as a function of the restored fraction. On LLaMA 3.2-3B, PPL stays
 502 relatively flat and then exhibits an abrupt drop at an elbow point, after which marginal gains saturate
 503 quickly. In contrast, Qwen 2.5-7B shows a more gradual, near-monotone PPL decrease with increas-
 504 ing restoration, without a clear knee. Since TTFB generally grows with the restoration fraction,
 505 these shapes motivate different selective-restoration budgets. We verify that these family-specific
 506 tendencies persist across other sizes within each family in our ablation study (Sec. G).

507

508 Guided by the above curves, GlowQ-S restores (i) for LLaMA, the elbow (steep-drop) operating
 509 point to capture most PPL gains with limited overhead, and (ii) for Qwen, a fixed 50% of groups,
 510 which offers stable accuracy improvements with moderate TTFB growth. For unit ranking, we
 511 follow the importance metrics delineated in Sec. 3.3: covariance-aware error capture is adopted as
 512 the default criterion. For the model families, when two alternative metrics are available (covariance-
 513 aware error capture vs. normalized error ratio), we evaluate both on the validation split and, per
 514 model, adopt the metric that yields the stronger outcome; all reported results use this per-model best
 515 choice.

516

517 4.7 IMPACT OF COVARIANCE ALIGNMENT ON ACCURACY

518

519 On C4 (Table 5), the Σ_x -weighted Whitened SVD
 520 consistently outperforms the unweighted Stacked
 521 SVD across both layer-wise and group-shared vari-
 522 ants of Qwen 3-8B. Because the unweighted objec-
 523 tive ignores the input-usage distribution embodied
 524 in Σ_x , it tends to select right subspaces misaligned
 525 with the axes most exploited by the data, leading to
 526 a marked degradation in perplexity at a fixed rank;
 527 whitening, by evaluating errors in a data-aligned
 528 coordinate system, improves energy capture and yields lower PPL. Grouped restoration also domi-
 529 nates layer-wise under both weightings, and, taken together, these results identify White + Group as
 530 the preferred configuration.

531

532 5 CONCLUSIONS

533

534 We introduced GlowQ, a group-shared low-rank approximation for quantized LLMs that replaces
 535 per-layer correction with a single right subspace shared among input-sharing modules and a cache-
 536 and-reuse runtime. By connecting usage-weighted risk to a right-weighted reconstruction objective,
 537 our covariance-aligned (whitened) formulation steers the learned subspace toward data-preferred di-
 538 rections, and a QR-reduced randomized SVD provides an efficient, scalable solver. The deployment
 539 path computes one right-side projection per group and reuses it across modules, while a selective
 540 policy (GlowQ-S) activates only high-importance units under latency or memory budgets. Across
 541 modern model families and PTQ baselines, GlowQ consistently lowers perplexity, reduces time-
 542 to-first-byte, increases throughput, and decreases memory overhead relative to layer-wise correc-
 543 tion; whitening and grouping combine to yield the strongest results. The approach is architecture-
 544 agnostic, drop-in at inference, and complementary to existing PTQ pipelines.

545 Table 5: C4 Evaluation of Σ_x -weighted
 546 (Whitened SVD) vs. unweighted (Stacked SVD)
 547 on Qwen 3-8B; lower is better (\downarrow).

No-White		White	
Layer	Group	Layer	Group
14.97	14.60	13.85	13.40

540 REFERENCES
541

542 Eduardo Alvarez, Omri Almog, Eric Chung, Simon Layton, Dusan Stosic, Ronny
543 Krashinsky, and Kyle Aubrey. Introducing NVFP4 for efficient and accu-
544 rate low-precision inference. <https://developer.nvidia.com/blog/introducing-nvfp4-for-efficient-and-accurate-low-precision-inference/>,
545 June 2025. NVIDIA Technical Blog.

546

547 T. W. Anderson. *An Introduction to Multivariate Statistical Analysis*. Wiley, 3 edition, 2003.

548

549 Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge,
550 Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu,
551 Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi
552 Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng
553 Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi
554 Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang
555 Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. Qwen technical report, 2023. URL
556 <https://arxiv.org/abs/2309.16609>.

557

558 Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du,
559 Xiao Liu, Aohan Zeng, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. LongBench: A bilingual,
560 multitask benchmark for long context understanding. In *Proceedings of the 62nd Annual Meet-
561 ing of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 3119–3137,
562 Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/
2024.acl-long.172. URL <https://aclanthology.org/2024.acl-long.172>.

563

564 Ron Banner, Itay Nahshan, and Daniel Soudry. Post training 4-bit quantization of convolutional
565 networks for rapid-deployment. In *Proceedings of NeurIPS*, 2018.

566

567 Adi Ben-Israel and Thomas N. E. Greville. *Generalized Inverses: Theory and Applications*. CMS
Books in Mathematics. Springer, 2 edition, 2003.

568

569 Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006. URL <https://www.microsoft.com/en-us/research/wp-content/uploads/2006/01/Bishop-Pattern-Recognition-and-Machine-Learning-2006.pdf>.

570

571

572 Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. Piqa: Reasoning about
573 physical commonsense in natural language. In *AAAI Conference on Artificial Intelligence*, 2019.
574 doi: 10.1609/AAAI.V34I05.6239.

575

576 Åke Björck. *Numerical Methods for Least Squares Problems*. SIAM, 1996. doi: 10.1137/1.
577 9781611971484.

578

579 Jingfei Chang, Yang Lu, Ping Xue, Yiqun Xu, and Zhen Wei. Iterative clustering pruning for con-
580 volutional neural networks. *Knowledge-Based Systems*, 265:110386, 2023. ISSN 0950-7051.
581 doi: <https://doi.org/10.1016/j.knosys.2023.110386>. URL <https://www.sciencedirect.com/science/article/pii/S0950705123001363>.

582

583 Christopher Clark, Kenton Lee, Ming-Wei Chang, T. Kwiatkowski, Michael Collins, and Kristina
584 Toutanova. Boolq: Exploring the surprising difficulty of natural yes/no questions. In *North
585 American Chapter of the Association for Computational Linguistics*, 2019. doi: 10.18653/v1/
586 N19-1300.

587

588 Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and
589 Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge.
590 In *arXiv.org*, 2018.

591

592 Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. Llm.int8(): 8-bit matrix multi-
593 plication for transformers at scale. In *Proceedings of the 36th International Conference on Neural
Information Processing Systems*, NIPS '22, Red Hook, NY, USA, 2022. Curran Associates Inc.
ISBN 9781713871088.

594 Zhen Dong, Zhewei Yao, Amir Gholami, Michael W. Mahoney, and Kurt Keutzer. Hawq: Hessian
 595 aware quantization of neural networks with mixed-precision. In *Proceedings of the IEEE/CVF*
 596 *International Conference on Computer Vision (ICCV)*, 2019.

597

598 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, and et al.
 599 The llama 3 herd of models. In *arXiv preprint arXiv:2407.21783*, 2024. doi: 10.48550/arXiv.
 600 2407.21783.

601 C. Eckart and G. Young. The approximation of one matrix by another of lower rank. *Psychometrika*,
 602 1936.

603

604 Yao et al. Zeroquant-v2: Exploring post-training quantization in llms from comprehensive study
 605 to low rank compensation. *arXiv preprint arXiv:2303.08302*, 2023. URL <https://arxiv.org/abs/2303.08302>.

606

607 Kawin Ethayarajh. How contextual are contextualized word representations? Comparing the geom-
 608 etry of BERT, ELMo, and GPT-2 embeddings. In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiao-
 609 jun Wan (eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language*
 610 *Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-
 611 IJCNLP)*, pp. 55–65, Hong Kong, China, November 2019. Association for Computational Lin-
 612 guistics. doi: 10.18653/v1/D19-1006. URL <https://aclanthology.org/D19-1006/>.

613

614 Ky Fan. On a theorem of Weyl concerning eigenvalues of linear transformations. *Proceedings of*
 615 *the National Academy of Sciences of the United States of America*, 35(11):652–655, 1949. doi:
 10.1073/pnas.35.11.652.

616

617 Elias Frantar, Saleh Ashkboos, Torsten Hoefer, and Dan Alistarh. Gptq: Accurate post-training
 618 quantization for generative pre-trained transformers, 2023. URL <https://arxiv.org/abs/2210.17323>.

619

620 Amir Gholami, Sehoon Kim, Zhen Dong, Zhewei Yao, Michael W. Mahoney, and Kurt Keutzer. A
 621 survey of quantization methods for efficient neural network inference, 2021.

622

623 Nathan Godey, Éric de la Clergerie, and Benoît Sagot. Anisotropy is inherent to self-attention in
 624 transformers, 2024. URL <https://arxiv.org/abs/2401.12143>.

625

626 Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. Johns Hopkins Univer-
 627 sity Press, 4th edition, 2013. URL <https://www.ee.iitb.ac.in/~belur/upload/golub-van-loan-matrix-computations-2012-edition-4th.pdf>.

628

629 Nathan Halko, Per-Gunnar Martinsson, and Joel Tropp. Finding structure with randomness: Prob-
 630 abilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53(2):
 631 217–288, 2011.

632

633 Arthur E. Hoerl and Robert W. Kennard. Ridge regression: Biased estimation for nonorthogonal
 634 problems. *Technometrics*, 12(1):55–67, 1970. doi: 10.1080/00401706.1970.10488634.

635

636 Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, 2 edition,
 637 2012.

638

639 Albert Qiaochu Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh
 640 Chaplot, and et al. Mistral 7b. In *arXiv preprint arXiv:2310.06825*, 2023. doi: 10.48550/arXiv.
 2310.06825.

641

642 Ian T. Jolliffe and Jorge Cadima. Principal component analysis: a review and recent developments.
 643 *Philosophical Transactions of the Royal Society A*, 374(2065):20150202, 2016.

644

645 Jinuk Kim, Marwa El Halabi, Wonpyo Park, Clemens JS Schaefer, Deokjae Lee, Yeonhong Park,
 646 Jae W. Lee, and Hyun Oh Song. Guidedquant: Large language model quantization via exploiting
 647 end loss guidance. In *International Conference on Machine Learning (ICML)*, 2025.

648

649 Raghuraman Krishnamoorthi. Quantizing deep convolutional networks for efficient inference: A
 650 whitepaper, 2018. URL <https://arxiv.org/abs/1806.08342>.

648 Olivier Ledoit and Michael Wolf. A well-conditioned estimator for large-dimensional covariance
 649 matrices. *Journal of Multivariate Analysis*, 88(2):365–411, 2004. doi: 10.1016/S0047-259X(03)
 650 00096-4.

651

652 Jungi Lee, Junyoung Park, Soohyun Cha, Jaehoon Cho, and Jaewoong Sim. Mx+: Pushing the limits
 653 of microscaling formats for efficient large language model serving. MICRO '25, pp. 869–883,
 654 New York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400715730. doi:
 655 10.1145/3725843.3756118. URL <https://doi.org/10.1145/3725843.3756118>.

656

657 Junyan Li, Yang Zhang, Muhammad Yusuf Hassan, Talha Chafekar, Tianle Cai, Zhile Ren, Peng-
 658 sheng Guo, Binazir Karimzadeh, Colorado J Reed, Chong Wang, and Chuang Gan. Commvq:
 659 Commutative vector quantization for kv cache compression. In *Proceedings of the 42nd Interna-
 660 tional Conference on Machine Learning (ICML)*, 2025.

661

662 Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Wei-Ming Chen, Wei-Chen Wang, Guangxuan
 663 Xiao, Xingyu Dang, Chuang Gan, and Song Han. Awq: Activation-aware weight quantization for
 664 llm compression and acceleration. In *MLSys*, 2024.

665

666 Zhengchi Ma and Rong Ma. Optimal estimation of shared singular subspaces across multiple noisy
 667 matrices, 2024. URL <https://arxiv.org/abs/2411.17054>.

668

669 Vladimir Malinovskii, Andrei Panferov, Ivan Ilin, Han Guo, Peter Richtárik, and Dan Alistarh.
 Pushing the limits of large language model quantization via the linearity theorem. *arXiv preprint
 arXiv:2411.17525*, 2024.

670

671 Per-Gunnar Martinsson and Joel A. Tropp. Randomized numerical linear algebra: Foundations
 672 and algorithms. *Acta Numerica*, 29:403–572, 2020. doi: 10.1017/S0962492920000021. URL
 673 <https://doi.org/10.1017/S0962492920000021>. See also arXiv:2002.01387 for the
 674 preprint.

675

676 Gabryel Mason-Williams and Fredrik Dahlqvist. What makes a good prune? maximal unstructured
 677 pruning for maximal cosine similarity. In *The Twelfth International Conference on Learning
 678 Representations*, 2024. URL <https://openreview.net/forum?id=jsvvPVVzwf>.

679

680 Stephen Merity, Caiming Xiong, James Bradbury, and R. Socher. Pointer sentinel mixture models.
 In *International Conference on Learning Representations*, 2016.

681

682 Leon Mirsky. Symmetric gauge functions and unitarily invariant norms. *The Quarterly Journal of
 683 Mathematics*, 11(1):50–59, 1960. doi: 10.1093/qmath/11.1.50.

684

685 Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz. Importance estimation
 686 for neural network pruning. In *CVPR*, 2019.

687

688 Cameron Musco and Christopher Musco. Randomized block krylov methods for stronger and faster
 689 approximate singular value decomposition. In *Advances in Neural Information Processing Sys-
 690 tems (NeurIPS)*, 2015.

691

692 Markus Nagel, Rana Ali Amjad, Mart van Baalen, Tijmen Blankevoort, and Max Welling. Up or
 693 down? adaptive rounding for post-training quantization. In *Proceedings of ICML (Workshops) /
 694 arXiv preprint*, 2020.

695

696 Markus Nagel, Marios Fournarakis, Rana Ali Amjad, Yelysei Bondarenko, Mart van Baalen, and
 697 Tijmen Blankevoort. A white paper on neural network quantization, 2021. URL <https://arxiv.org/abs/2106.08295>.

698

699 Open Compute Project. Ocp microscaling formats (mx) specification, version 1.0. Technical re-
 700 port, Open Compute Project, September 2023. URL <https://www.opencompute.org/documents/ocp-microscaling-formats-mx-v1-0-spec-final-pdf>. Version
 701 1.0, Sept. 7, 2023.

702

703 Christopher C. Paige and Michael A. Saunders. Towards a generalized singular value decompositon.
 704 *SIAM Journal on Numerical Analysis*, 18(3):398–405, 1981. doi: 10.1137/0718026.

702 Denis Paperno, Germu00e1n Kruszewski, Angeliki Lazaridou, Q. N. Pham, R. Bernardi, Sandro
 703 Pezzelle, Marco Baroni, Gemma Boleda, and R. Fernu00e1ndez. The lambada dataset: Word
 704 prediction requiring a broad discourse context. In *Annual Meeting of the Association for Compu-*
 705 *tational Linguistics*, 2016. doi: 10.18653/v1/P16-1144.

706 Roger Penrose. A generalized inverse for matrices. *Mathematical Proceedings of the Cambridge*
 707 *Philosophical Society*, 51(3):406–413, 1955. doi: 10.1017/S0305004100030401.

709 Kaare Brandt Petersen and Michael Syskind Pedersen. The matrix cookbook. [https://www.](https://www.math.uwaterloo.ca/~hwolkowi/matrixcookbook.pdf)
 710 [math.uwaterloo.ca/~hwolkowi/matrixcookbook.pdf](https://www.math.uwaterloo.ca/~hwolkowi/matrixcookbook.pdf), 2012. Accessed: 2025-09-
 711 18.

712 Hadi Pouransari and Oncel Tuzel. Least squares binary quantization of neural networks. 2020
 713 *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*,
 714 pp. 2986–2996, 2020. URL <https://api.semanticscholar.org/CorpusID:210116550>.

717 Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi
 718 Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text
 719 transformer, 2023. URL <https://arxiv.org/abs/1910.10683>.

720 Benjamin Recht, Maryam Fazel, and Pablo A. Parrilo. Guaranteed minimum-rank solutions of
 721 linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010. doi:
 722 10.1137/070697835.

724 Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: An adver-
 725 sarial winograd schema challenge at scale, 2019. URL <https://arxiv.org/abs/1907.10641>.

727 Ajit P. Singh and Geoffrey J. Gordon. Relational learning via collective matrix factorization. In
 728 *Proceedings of KDD*, 2008.

729 Nathan Srebro and Tommi Jaakkola. Weighted low-rank approximations. In *Proceedings of ICML*,
 730 2003.

732 Hugo Touvron, Louis Martin, Kevin R. Stone, Peter Albert, Amjad Almahairi, and et al. Llama 2:
 733 Open foundation and fine-tuned chat models. In *arXiv preprint arXiv:2307.09288*, 2023.

734 Lloyd N. Trefethen and III Bau, David. *Numerical Linear Algebra*. SIAM, 1997. doi: 10.1137/1.
 735 9780898719574.

737 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez,
 738 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proceedings of the 31st Inter-*
 739 *national Conference on Neural Information Processing Systems*, NIPS’17, pp. 6000–6010, Red
 740 Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.

741 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, and et al. Qwen2.5 technical re-
 742 port. *CoRR*, abs/2412.15115, 2024. URL <https://doi.org/10.48550/arXiv.2412.15115>.

744 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, and et al. Qwen3 technical
 745 report, 2025. URL <https://arxiv.org/abs/2505.09388>.

747 Junho Yoon, Geom Lee, Donghyeon Jeon, Inho Kang, and Seung-Hoon Na. ROSAQ: rotation-based
 748 saliency-aware weight quantization for efficiently compressing large language models. *CoRR*,
 749 abs/2506.13472, 2025. doi: 10.48550/ARXIV.2506.13472. URL <https://doi.org/10.48550/arXiv.2506.13472>.

751 Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a ma-
 752 chine really finish your sentence?, 2019. URL <https://arxiv.org/abs/1905.07830>.

754 Cheng Zhang, Jianyi Cheng, George A. Constantinides, and Yiren Zhao. Lqer: low-rank quan-
 755 tization error reconstruction for llms. In *Proceedings of the 41st International Conference on*
Machine Learning, ICML’24. JMLR.org, 2024a.

756 Cheng Zhang, Jeffrey TH Wong, Can Xiao, George A Constantinides, and Yiren Zhao. Qera: an
 757 analytical framework for quantization error reconstruction. *arXiv preprint arXiv:2410.06040*,
 758 2024b.

759 Rongzhi Zhang, Kuang Wang, Liyuan Liu, Shuohang Wang, Hao Cheng, Chao Zhang, and Yelong
 760 Shen. Lorc: Low-rank compression for llms kv cache with a progressive compression strategy,
 761 2024c. URL <https://arxiv.org/abs/2410.03111>.

762 Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuhui Chen, Christopher
 763 Dewan, Mona T. Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer,
 764 Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer.
 765 Opt: Open pre-trained transformer language models. In *arXiv.org*, 2022.

766 Weibo Zhao, Yubin Shi, Xinyu Lyu, Wanchen Sui, Shen Li, and Yong Li. Aser: activation smoothing
 767 and error reconstruction for large language model quantization. In *Proceedings of the Thirty-Ninth
 768 AAAI Conference on Artificial Intelligence and Thirty-Seventh Conference on Innovative Appli-
 769 cations of Artificial Intelligence and Fifteenth Symposium on Educational Advances in Artificial
 770 Intelligence, AAAI'25/IAAI'25/EAAI'25*. AAAI Press, 2025. ISBN 978-1-57735-897-8. doi:
 771 10.1609/aaai.v39i21.34443. URL <https://doi.org/10.1609/aaai.v39i21.34443>.

774 A APPENDIX

775 A.0 NOTATION & SHAPES

776 Refer to Table 6

777 A.1 STACKED SVD: SHARED RIGHT SUBSPACE AND GLOBAL OPTIMUM (PROOF)

778 When multiple modules share the same input dimension, we vertically concatenate the module-wise
 779 error matrices $\mathbf{E}_i \in \mathbb{R}^{O_i \times d}$ into \mathbf{E}_{cat} . We then choose a shared right subspace (the row space of \mathbf{B})
 780 while allowing module-specific left factors \mathbf{A}_i , by solving

$$781 \min_{\mathbf{A}, \mathbf{B}} \|\mathbf{E}_{\text{cat}} - \mathbf{AB}\|_F^2.$$

782 This appendix shows that (i) the solution is well-defined, and (ii) the shared \mathbf{B} is also optimal in
 783 an energy/projection sense (Ky Fan; cf. Fan (1949); Golub & Loan (2013)). Consequently, a single
 784 shared \mathbf{B} serves as a strong representative of what one might otherwise try to learn as separate \mathbf{B}_i 's
 785 per module.

786 Problem (Unweighted Frobenius Approximation).

$$787 \min_{\mathbf{A}, \mathbf{B}} \|\mathbf{E}_{\text{cat}} - \mathbf{AB}\|_F^2. \quad (\text{A.1.1})$$

788 LEMMA A.1.1 - EQUIVALENCE OF SEARCH SETS: $\mathcal{M}_r = \mathcal{R}_r$.

789 Let

$$790 \mathcal{M}_r := \{\mathbf{AB} : \mathbf{A} \in \mathbb{R}^{m \times r}, \mathbf{B} \in \mathbb{R}^{r \times d}\}, \quad \mathcal{R}_r := \{\mathbf{X} \in \mathbb{R}^{m \times d} : \text{rank}(\mathbf{X}) \leq r\}.$$

791 Then $\mathcal{M}_r = \mathcal{R}_r$.

792 This is a standard consequence of rank-factorization and the SVD characterization of best rank- r
 793 approximants; see, e.g., Eckart & Young (1936); Mirsky (1960); Golub & Loan (2013); Horn &
 794 Johnson (2012). We omit the proof.

795 **Proof.** By Lemma A.1.1, the problem reduces to a rank- r approximation of \mathbf{E}_{cat} . By the Eckart-
 796 Young-Mirsky theorem Eckart & Young (1936); Mirsky (1960), the optimizer is the truncated SVD

$$797 \mathbf{X}^* = \mathbf{U}_r \mathbf{\Sigma}_r \mathbf{V}_r^\top,$$

798 so any global minimizer (\mathbf{A}, \mathbf{B}) must satisfy

$$799 \mathbf{AB} = \mathbf{X}^* = \mathbf{U}_r \mathbf{\Sigma}_r \mathbf{V}_r^\top. \quad (\text{A.1.2})$$

800 If $\sigma_r = \sigma_{r+1}$, the optimizer may be non-unique Golub & Loan (2013). \square

Table 6: Unified notation and shapes for stacked errors, low-rank factors, and covariance-weighted core.

Symbol	Meaning
$\mathbf{E}_i \in \mathbb{R}^{O_i \times d}$	Error matrix of module i with output dimension O_i and shared input dimension d .
$\mathbf{E}_{\text{cat}} := [\mathbf{E}_1; \dots; \mathbf{E}_m] \in \mathbb{R}^{m \times d}$	Row-stacked errors across modules; total rows $m := \sum_i O_i$.
$\mathbf{A} := [\mathbf{A}_1; \dots; \mathbf{A}_m] \in \mathbb{R}^{m \times r}$	Left factor formed by stacking per-module factors \mathbf{A}_i .
$\mathbf{B} \in \mathbb{R}^{r \times d}$	Shared right factor; target rank r .
$1 \leq r \leq \min\{m, d\}$	Admissible rank range.
$\mathbf{E}_{\text{cat}} = \mathbf{U}\Sigma\mathbf{V}^\top$	Thin SVD with $\mathbf{U} \in \mathbb{R}^{m \times d}$, $\Sigma \in \mathbb{R}^{d \times d}$, $\mathbf{V} \in \mathbb{R}^{d \times d}$ orthogonal.
$(\mathbf{U}_r, \Sigma_r, \mathbf{V}_r)$	Top- r SVD blocks: $\mathbf{U}_r \in \mathbb{R}^{m \times r}$, $\Sigma_r \in \mathbb{R}^{r \times r}$, $\mathbf{V}_r \in \mathbb{R}^{d \times r}$.
$\Sigma_x \succeq 0$	Input covariance; $\Sigma_x := \mathbb{E}[\mathbf{x}\mathbf{x}^\top]$ for centered inputs $\mathbb{E}[\mathbf{x}] = \mathbf{0}$.
$\Sigma_x^{1/2}$	(Pseudo-)square root of Σ_x .
$\mathbf{E}_{\text{cat}} = \mathbf{Q}_e \mathbf{R}_e$	Thin QR with $\mathbf{Q}_e \in \mathbb{R}^{m \times d}$, $\mathbf{Q}_e^\top \mathbf{Q}_e = \mathbf{I}_d$, $\mathbf{R}_e \in \mathbb{R}^{d \times d}$.
$\mathbf{M} := \mathbf{R}_e \Sigma_x^{1/2} \in \mathbb{R}^{d \times d}$	Covariance-weighted SVD core used for randomized SVD on the reduced space.
$\hat{\mathbf{A}} := \mathbf{Q}_e^\top \mathbf{A} \in \mathbb{R}^{d \times r}$	Variable change (reduced left factor).
$\hat{\mathbf{B}} := \mathbf{B} \Sigma_x^{1/2} \in \mathbb{R}^{r \times d}$	Variable change (covariance-weighted right factor).
Residual $(\mathbf{E}_{\text{cat}} - \mathbf{AB}) \in \mathbb{R}^{m \times d}$	Stacked error after factorization (no separate symbol reserved).

THEOREM A.1.2 - IDENTIFYING THE SHARED RIGHT SUBSPACE: $\text{row}(\mathbf{B}) = \text{span}(\mathbf{V}_r^\top)$.

We determine the optimal shared right subspace for the factorization $\min_{\mathbf{A}, \mathbf{B}} \|\mathbf{E}_{\text{cat}} - \mathbf{AB}\|_F^2$. Let $\mathbf{E}_{\text{cat}} = \mathbf{U}\Sigma\mathbf{V}^\top$ be a thin SVD, and let $r = \text{rank}(\mathbf{B})$. Denote $\mathbf{S} := \text{row}(\mathbf{B})$ and the orthogonal projector $\mathbf{P}_S := \mathbf{B}^\top(\mathbf{B}\mathbf{B}^\top)^{-1}\mathbf{B}$ (assume \mathbf{B} has full row rank; otherwise use the Moore-Penrose pseudoinverse).

Fixing \mathbf{B} , least-squares normal equations yield (see, e.g., (Golub & Loan, 2013, §5))

$$\mathbf{A}^* = \mathbf{E}_{\text{cat}} \mathbf{B}^\top (\mathbf{B}\mathbf{B}^\top)^{-1}, \quad (\text{A.1.3a})$$

$$\mathbf{A}^* \mathbf{B} = \mathbf{E}_{\text{cat}} \mathbf{P}_S. \quad (\text{A.1.3b})$$

Hence, with $\mathbf{G} := \mathbf{E}_{\text{cat}}^\top \mathbf{E}_{\text{cat}}$,

$$\|\mathbf{E}_{\text{cat}} - \mathbf{A}^* \mathbf{B}\|_F^2 = \|\mathbf{E}_{\text{cat}} (\mathbf{I} - \mathbf{P}_S)\|_F^2 = \|\mathbf{E}_{\text{cat}}\|_F^2 - \text{tr}(\mathbf{P}_S \mathbf{G}), \quad (\text{A.1.4})$$

where the last identity is the usual projection-trace formula (cf. Horn & Johnson (2012)).

864 Therefore, selecting \mathbf{S} of dimension r is equivalent to
 865

$$866 \quad \max_{\dim \mathbf{S}=r} \text{tr}(\mathbf{P}_{\mathbf{S}} \mathbf{G}). \quad (\text{A.1.5})$$

867 By Ky Fan's maximum principle Fan (1949), the maximizer \mathbf{S} is the span of the top- r eigenvectors
 868 of \mathbf{G} . Since $\mathbf{E}_{\text{cat}} = \mathbf{U} \Sigma \mathbf{V}^{\top}$ implies $\mathbf{G} = \mathbf{V} \Sigma^2 \mathbf{V}^{\top}$, its top- r eigenspace equals $\text{span}(\mathbf{V}_r)$. Thus
 869

$$870 \quad \text{row}(\mathbf{B}) = \text{span}(\mathbf{V}_r^{\top}). \quad (\text{A.1.6})$$

□

873 THEOREM A.1.3 - REPRESENTATIVENESS / ENERGY OPTIMALITY: SUM OF PROJECTION
 874 ENERGIES.

875 The shared right subspace $\mathbf{S} = \text{row}(\mathbf{B})$ of dimension r maximizes the total projection energy
 876 $\sum_i \|\mathbf{E}_i \mathbf{P}_{\mathbf{S}}\|_F^2$, where $\mathbf{P}_{\mathbf{S}}$ is the orthogonal projector onto \mathbf{S} (e.g., $\mathbf{P}_{\mathbf{S}} = \mathbf{Q} \mathbf{Q}^{\top}$ for any orthonormal
 877 basis \mathbf{Q} of \mathbf{S}).
 878

879 **Proof.** For each module \mathbf{E}_i ,

$$880 \quad \|\mathbf{E}_i \mathbf{P}_{\mathbf{S}}\|_F^2 = \text{tr}(\mathbf{P}_{\mathbf{S}} \mathbf{E}_i^{\top} \mathbf{E}_i), \quad (\text{A.1.7})$$

881 a standard identity using symmetry/idempotence of $\mathbf{P}_{\mathbf{S}}$ and trace cyclicity (see, e.g., Horn & John-
 882 son (2012); Golub & Loan (2013)). Summing over i yields
 883

$$884 \quad \max_{\dim \mathbf{S}=r} \sum_i \|\mathbf{E}_i \mathbf{P}_{\mathbf{S}}\|_F^2 = \max_{\dim \mathbf{S}=r} \text{tr}\left(\mathbf{P}_{\mathbf{S}} \sum_i \mathbf{E}_i^{\top} \mathbf{E}_i\right) = \max_{\dim \mathbf{S}=r} \text{tr}(\mathbf{P}_{\mathbf{S}} \mathbf{G}), \quad \mathbf{G} := \sum_i \mathbf{E}_i^{\top} \mathbf{E}_i = \mathbf{E}_{\text{cat}}^{\top} \mathbf{E}_{\text{cat}}. \quad (\text{A.1.8})$$

885 By Ky Fan's maximum principle Fan (1949) (cf. Eq. A.1.5), the maximizer is the span of the top- r
 886 eigenvectors of \mathbf{G} . Since $\mathbf{E}_{\text{cat}} = \mathbf{U} \Sigma \mathbf{V}^{\top}$ implies $\mathbf{G} = \mathbf{V} \Sigma^2 \mathbf{V}^{\top}$, it follows that
 887

$$888 \quad \mathbf{S}^* = \text{span}(\mathbf{V}_r) \iff \text{row}(\mathbf{B}) = \text{span}(\mathbf{V}_r^{\top}).$$

□

892 LEMMA A.1.4 - IDENTIFIABILITY AND "BALANCED" FACTORIZATION.
 893

894 Although the pair (\mathbf{A}, \mathbf{B}) is non-unique up to invertible reparameterizations, the right subspace
 895 $\text{row}(\mathbf{B})$ is identifiable; choosing the SVD half-split $\Sigma_r^{1/2}$ yields a numerically stable balanced fac-
 896 torization Golub & Loan (2013).
 897

898 **Non-uniqueness.** For any invertible $\mathbf{R} \in \mathbb{R}^{r \times r}$,

$$899 \quad (\mathbf{A}, \mathbf{B}) \mapsto (\mathbf{A}\mathbf{R}, \mathbf{R}^{-1}\mathbf{B}) \Rightarrow \mathbf{AB} \text{ invariant.}$$

900 Hence factors are not unique, while the projector onto $\text{row}(\mathbf{B})$ is unique (right singular subspace;
 901 cf. Theorem A.1.2 and Golub & Loan (2013)).
 902

903 **Balanced factorization.** Let $\mathbf{E}_{\text{cat}} = \mathbf{U} \Sigma \mathbf{V}^{\top}$ and denote by $\mathbf{U}_r, \Sigma_r, \mathbf{V}_r$ the top- r blocks. The
 904 half-split
 905

$$906 \quad \mathbf{A}^* = \mathbf{U}_r \Sigma_r^{1/2}, \quad (\text{A.1.9a})$$

$$907 \quad \mathbf{B}^* = \Sigma_r^{1/2} \mathbf{V}_r^{\top}, \quad (\text{A.1.9b})$$

$$908 \quad \mathbf{A}^* \mathbf{B}^* = \mathbf{U}_r \Sigma_r \mathbf{V}_r^{\top} \quad (\text{A.1.9c})$$

909 satisfies

$$910 \quad \mathbf{A}^{*\top} \mathbf{A}^* = \Sigma_r, \quad \mathbf{B}^* \mathbf{B}^{*\top} = \Sigma_r,$$

911 which avoids squaring condition numbers in normal equations and minimizes combined factor
 912 norms among reparameterizations:
 913

$$914 \quad \frac{1}{2} (\|\mathbf{A}\|_F^2 + \|\mathbf{B}\|_F^2) \geq \|\mathbf{U}_r \Sigma_r \mathbf{V}_r^{\top}\|_*,$$

915 with equality at $(\mathbf{A}^*, \mathbf{B}^*)$ Recht et al. (2010). (Standard facts; see Golub & Loan (2013); Recht
 916 et al. (2010).)
 917

918 **Block Recovery and the Pseudoinverse** Given the shared right factor \mathbf{B}^* , each module-specific
 919 left factor \mathbf{A}_i is obtained by a single least-squares solve. Using the Moore-Penrose pseudoinverse
 920 provides the minimum-norm solution and remains valid under rank deficiency Penrose (1955); Ben-
 921 Israel & Greville (2003); Golub & Loan (2013):
 922

$$\mathbf{A}_i^* = \mathbf{E}_i \mathbf{B}^{*\top} (\mathbf{B}^* \mathbf{B}^{*\top})^\dagger. \quad (\text{A.1.10})$$

923 It suggests that (i) when \mathbf{B}^* has full row rank, $(\cdot)^\dagger$ reduces to the inverse and Eq. A.1.10 coincides
 924 with the normal-equations solution; (ii) in general, $(\cdot)^\dagger$ yields the unique minimum-norm LS solution
 925 and is numerically stable under near-singularity Ben-Israel & Greville (2003); Golub & Loan
 926 (2013).
 927

928 **A.2 COVARIANCE-ALIGNED OBJECTIVE: BRIDGE EQUIVALENCE AND GLOBAL MINIMIZER
 929 (PROOF)**

930 Sec. 3.1.2 formulates the covariance-aligned objective
 931

$$\min_{\mathbf{A}, \mathbf{B}} \|(\mathbf{E}_{\text{cat}} - \mathbf{AB}) \Sigma_{\mathbf{x}}^{1/2}\|_F^2,$$

932 which weights errors by the input usage encoded in the covariance $\Sigma_{\mathbf{x}}$ Anderson (2003); Bishop
 933 (2006). This appendix provides a complete mathematical justification: (i) a bridge equivalence
 934 that converts $\mathbb{E}[\|(\mathbf{E}_{\text{cat}} - \mathbf{AB}) \mathbf{x}\|_2^2]$ into a Frobenius form via the trace identity $\mathbb{E}[\mathbf{x}^\top \mathbf{M} \mathbf{x}] =$
 935 $\text{tr}(\mathbf{M} \Sigma_{\mathbf{x}})$ Petersen & Pedersen (2012); (ii) a whitening reduction to a standard low-rank approx-
 936 imation by the change of variables $\tilde{\mathbf{B}} := \mathbf{B} \Sigma_{\mathbf{x}}^{1/2}$ (and $\mathbf{E}_{\text{cat}} \Sigma_{\mathbf{x}}^{1/2}$ on the right) Golub & Loan (2013);
 937 (iii) a closed-form global minimizer given by the truncated SVD of $\mathbf{E}_{\text{cat}} \Sigma_{\mathbf{x}}^{1/2}$ with balanced factors
 938 and the identity of the shared right subspace; and (iv) extensions to nonzero-mean inputs (centering)
 939 and singular $\Sigma_{\mathbf{x}}$ via pseudoinverse whitening Ben-Israel & Greville (2003); Penrose (1955).
 940

941 In our case, the (distribution-weighted) risk is the expected squared output error under the input law:
 942

$$\mathcal{R}(\mathbf{A}, \mathbf{B}) := \mathbb{E} \|\mathbf{M} \mathbf{x}\|_2^2.$$

943 Directions used more frequently or with larger magnitude (large variance) are weighted more heavily
 944 by $\Sigma_{\mathbf{x}}$, which motivates a right-weighted objective via $\Sigma_{\mathbf{x}}^{1/2}$ Bishop (2006); Anderson (2003). An
 945 empirical counterpart uses samples $\{\mathbf{x}_n\}_{n=1}^N$:

$$\widehat{\mathcal{R}}(\mathbf{A}, \mathbf{B}) := \frac{1}{N} \sum_{n=1}^N \|\mathbf{M} \mathbf{x}_n\|_2^2, \quad \widehat{\Sigma}_{\mathbf{x}} := \frac{1}{N} \sum_{n=1}^N \mathbf{x}_n \mathbf{x}_n^\top.$$

954 **THEOREM A.2.1 (BRIDGE EQUIVALENCE).**

955 In this subsection, we prove the bridge identity $\mathbb{E} \|\mathbf{M} \mathbf{x}\|_2^2 = \text{tr}(\mathbf{M} \Sigma_{\mathbf{x}} \mathbf{M}^\top) = \|\mathbf{M} \Sigma_{\mathbf{x}}^{1/2}\|_F^2$, which
 956 converts the distribution-weighted risk into a Frobenius norm amenable to SVD analysis (see the
 957 trace/expectation identities in Petersen & Pedersen (2012)).
 958

959 For zero-mean inputs with covariance $\Sigma_{\mathbf{x}} \succeq 0$,
 960

$$\mathbb{E} \|\mathbf{M} \mathbf{x}\|_2^2 = \text{tr}(\mathbf{M} \Sigma_{\mathbf{x}} \mathbf{M}^\top) = \|\mathbf{M} \Sigma_{\mathbf{x}}^{1/2}\|_F^2. \quad (\text{A.2.1})$$

961 **Proof.** (Vector norm \rightarrow trace). Since $\|\mathbf{y}\|_2^2 = \text{tr}(\mathbf{y} \mathbf{y}^\top)$ and trace is linear,
 962

$$\mathbb{E} \|\mathbf{M} \mathbf{x}\|_2^2 = \mathbb{E} \text{tr}(\mathbf{M} \mathbf{x} \mathbf{x}^\top \mathbf{M}^\top) = \text{tr}(\mathbf{M} \mathbb{E}[\mathbf{x} \mathbf{x}^\top] \mathbf{M}^\top) = \text{tr}(\mathbf{M} \Sigma_{\mathbf{x}} \mathbf{M}^\top).$$

963 (Trace \rightarrow Frobenius). Because $\|\mathbf{Z}\|_F^2 = \text{tr}(\mathbf{Z} \mathbf{Z}^\top)$ and $\Sigma_{\mathbf{x}}^{1/2} \Sigma_{\mathbf{x}}^{1/2} = \Sigma_{\mathbf{x}}$,
 964

$$\|\mathbf{M} \Sigma_{\mathbf{x}}^{1/2}\|_F^2 = \text{tr}((\mathbf{M} \Sigma_{\mathbf{x}}^{1/2})(\mathbf{M} \Sigma_{\mathbf{x}}^{1/2})^\top) = \text{tr}(\mathbf{M} \Sigma_{\mathbf{x}} \mathbf{M}^\top). \quad \square$$

965 Distribution-weighted risk equals the Frobenius norm of the right-whitened residual $\mathbf{M} \Sigma_{\mathbf{x}}^{1/2}$.
 966

972 LEMMA A.2.2 (NONZERO-MEAN INPUTS).
 973

974 In this subsection, we decompose the risk for $\mathbb{E}[\mathbf{x}] \neq \mathbf{0}$ into a covariance term and a deterministic
 975 mean term, showing $\mathbb{E}\|\mathbf{M}\mathbf{x}\|_2^2 = \text{tr}(\mathbf{M}\text{Cov}(\mathbf{x})\mathbf{M}^\top) + \|\mathbf{M}\boldsymbol{\mu}\|_2^2$ (cf. Anderson (2003); Bishop
 976 (2006)).

977 Let $\boldsymbol{\mu} := \mathbb{E}[\mathbf{x}]$ and $\text{Cov}(\mathbf{x}) := \mathbb{E}[(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})^\top]$. Then
 978

$$\mathbb{E}\|\mathbf{M}\mathbf{x}\|_2^2 = \text{tr}(\mathbf{M}\text{Cov}(\mathbf{x})\mathbf{M}^\top) + \|\mathbf{M}\boldsymbol{\mu}\|_2^2. \quad (\text{A.2.2})$$

980 **Proof.** Write $\mathbf{x} = (\mathbf{x} - \boldsymbol{\mu}) + \boldsymbol{\mu}$ and expand:
 981

$$\|\mathbf{M}\mathbf{x}\|_2^2 = \|\mathbf{M}(\mathbf{x} - \boldsymbol{\mu})\|_2^2 + 2\langle \mathbf{M}(\mathbf{x} - \boldsymbol{\mu}), \mathbf{M}\boldsymbol{\mu} \rangle + \|\mathbf{M}\boldsymbol{\mu}\|_2^2.$$

982 Taking expectations annihilates the cross term since $\mathbb{E}[\mathbf{x} - \boldsymbol{\mu}] = \mathbf{0}$, yielding the claim. \square
 983

984 Risk decomposes into a covariance term plus a mean-induced term.
 985

986 THEOREM A.2.3 (VARIABLE CHANGE AND WHITENING).
 987

988 In this subsection, we show that right-whitening reduces the covariance-aligned objective to a stan-
 989 dard Frobenius low-rank approximation by proving $\|(\mathbf{E}_{\text{cat}} - \mathbf{AB})\boldsymbol{\Sigma}_{\mathbf{x}}^{1/2}\|_F^2 = \|\tilde{\mathbf{E}} - \hat{\mathbf{A}}\hat{\mathbf{B}}\|_F^2$ with
 990 $\tilde{\mathbf{E}} = \mathbf{E}_{\text{cat}}\boldsymbol{\Sigma}_{\mathbf{x}}^{1/2}$, $\hat{\mathbf{B}} = \mathbf{B}\boldsymbol{\Sigma}_{\mathbf{x}}^{1/2}$ (standard whitening trick; cf. Golub & Loan (2013)).

991 Define

$$\tilde{\mathbf{E}} := \mathbf{E}_{\text{cat}}\boldsymbol{\Sigma}_{\mathbf{x}}^{1/2}, \quad \hat{\mathbf{A}} := \mathbf{A}, \quad \hat{\mathbf{B}} := \mathbf{B}\boldsymbol{\Sigma}_{\mathbf{x}}^{1/2}.$$

992 Then

$$\|(\mathbf{E}_{\text{cat}} - \mathbf{AB})\boldsymbol{\Sigma}_{\mathbf{x}}^{1/2}\|_F^2 = \|\tilde{\mathbf{E}} - \hat{\mathbf{A}}\hat{\mathbf{B}}\|_F^2. \quad (\text{A.2.3})$$

993 **Proof.** Direct substitution:
 994

$$\tilde{\mathbf{E}} - \hat{\mathbf{A}}\hat{\mathbf{B}} = \mathbf{E}_{\text{cat}}\boldsymbol{\Sigma}_{\mathbf{x}}^{1/2} - \mathbf{A}(\mathbf{B}\boldsymbol{\Sigma}_{\mathbf{x}}^{1/2}) = (\mathbf{E}_{\text{cat}} - \mathbf{AB})\boldsymbol{\Sigma}_{\mathbf{x}}^{1/2}.$$

995 Taking Frobenius norms yields the identity. \square

996 Whitening converts risk minimization into a plain Frobenius factorization.
 997

1000 LEMMA A.2.4 (WEIGHTED LEAST SQUARES FOR \mathbf{A} GIVEN \mathbf{B}).
 1001

1002 In this subsection, we derive the closed-form weighted least-squares minimizer $\mathbf{A}^* = \mathbf{E}\boldsymbol{\Sigma}_{\mathbf{x}}\mathbf{B}^\top(\mathbf{B}\boldsymbol{\Sigma}_{\mathbf{x}}\mathbf{B}^\top)^{-1}$ for fixed \mathbf{B} , and interpret the residual as a $\boldsymbol{\Sigma}_{\mathbf{x}}$ -weighted right projection
 1003 ((Golub & Loan, 2013, Ch. 5), Björck (1996); matrix derivatives in Petersen & Pedersen (2012)).
 1004

1005 Consider

$$f(\mathbf{A}) := \|(\mathbf{E}_{\text{cat}} - \mathbf{AB})\boldsymbol{\Sigma}_{\mathbf{x}}^{1/2}\|_F^2.$$

1006 Let $\mathbf{E} := \mathbf{E}_{\text{cat}}$, $\mathbf{E}_\sim := \mathbf{E}\boldsymbol{\Sigma}_{\mathbf{x}}^{1/2}$, and $\hat{\mathbf{B}} := \mathbf{B}\boldsymbol{\Sigma}_{\mathbf{x}}^{1/2}$. Then the unique least-squares minimizer is
 1007

$$\mathbf{A}^* = \mathbf{E}_\sim\hat{\mathbf{B}}^\top(\hat{\mathbf{B}}\hat{\mathbf{B}}^\top)^{-1} = \mathbf{E}\boldsymbol{\Sigma}_{\mathbf{x}}\mathbf{B}^\top(\mathbf{B}\boldsymbol{\Sigma}_{\mathbf{x}}\mathbf{B}^\top)^{-1}. \quad (\text{A.2.4})$$

1008 **Proof.** In whitened variables,
 1009

$$f(\mathbf{A}) = \|\mathbf{E}_\sim - \mathbf{A}\hat{\mathbf{B}}\|_F^2 = \text{tr}(\mathbf{E}_\sim\mathbf{E}_\sim^\top) - 2\text{tr}(\mathbf{A}\hat{\mathbf{B}}\mathbf{E}_\sim^\top) + \text{tr}(\mathbf{A}(\hat{\mathbf{B}}\hat{\mathbf{B}}^\top)\mathbf{A}^\top).$$

1010 Using $\frac{\partial}{\partial \mathbf{A}} \text{tr}(\mathbf{A}\mathbf{C}\mathbf{A}^\top) = 2\mathbf{A}\mathbf{C}$ for symmetric \mathbf{C} and $\frac{\partial}{\partial \mathbf{A}} \text{tr}(\mathbf{A}\mathbf{M}) = \mathbf{M}^\top$ Petersen & Pedersen
 1011 (2012),

$$\nabla_{\mathbf{A}} f(\mathbf{A}) = -2\mathbf{E}_\sim\hat{\mathbf{B}}^\top + 2\mathbf{A}(\hat{\mathbf{B}}\hat{\mathbf{B}}^\top) = 0 \Rightarrow \mathbf{A}^* = \mathbf{E}_\sim\hat{\mathbf{B}}^\top(\hat{\mathbf{B}}\hat{\mathbf{B}}^\top)^{-1}.$$

1012 Substituting $\mathbf{E}_\sim = \mathbf{E}\boldsymbol{\Sigma}_{\mathbf{x}}^{1/2}$ and $\hat{\mathbf{B}} = \mathbf{B}\boldsymbol{\Sigma}_{\mathbf{x}}^{1/2}$ gives the second form. \square
 1013

1014 In whitened variables, $\mathbf{E}_\sim - \mathbf{A}^*\hat{\mathbf{B}} = \mathbf{E}_\sim(\mathbf{I} - \mathbf{P}_{\hat{\mathbf{B}}})$ with $\mathbf{P}_{\hat{\mathbf{B}}} := \hat{\mathbf{B}}^\top(\hat{\mathbf{B}}\hat{\mathbf{B}}^\top)^{-1}\hat{\mathbf{B}}$, the orthogonal
 1015 projector onto row($\hat{\mathbf{B}}$) in the Euclidean metric. In original variables, $\mathbf{A}^*\mathbf{B} = \mathbf{E}\mathbf{P}_\Sigma$ with
 1016

$$\mathbf{P}_\Sigma := \boldsymbol{\Sigma}_{\mathbf{x}}\mathbf{B}^\top(\mathbf{B}\boldsymbol{\Sigma}_{\mathbf{x}}\mathbf{B}^\top)^{-1}\mathbf{B},$$

1017 the right projection under the $\boldsymbol{\Sigma}_{\mathbf{x}}$ -weighted inner product (a standard form of weighted/oblique
 1018 projection; cf. Golub & Loan (2013), Ben-Israel & Greville (2003); Björck (1996)).
 1019

1020 For fixed \mathbf{B} , the optimal \mathbf{A} is a weighted LS solution; the residual is a $\boldsymbol{\Sigma}_{\mathbf{x}}$ -weighted right projection.
 1021

1026 THEOREM A.2.5 (GLOBAL MINIMIZER; BALANCED FACTORS; RIGHT SUBSPACE).
 1027

1028 In this subsection, we obtain the global solution via the Eckart–Young–Mirsky theorem Eckart &
 1029 Young (1936); Mirsky (1960), choose balanced factors $\widehat{\mathbf{A}}^* = \mathbf{U}_r \Sigma_r^{1/2}$, $\widehat{\mathbf{B}}^* = \Sigma_r^{1/2} \mathbf{V}_r^\top$, and iden-
 1030 tify the optimal shared right subspace as $\text{row}(\mathbf{B}^*) = \text{row}(\mathbf{V}_r^\top \Sigma_x^{-1/2})$ (cf. Ky Fan’s principle and
 1031 the subspace discussion in Fan (1949); Golub & Loan (2013)).

1032 Let $\widetilde{\mathbf{E}} = \mathbf{U} \Sigma \mathbf{V}^\top$ be an SVD and $(\mathbf{U}_r, \Sigma_r, \mathbf{V}_r)$ the top- r blocks. Then
 1033

$$1034 \quad \widehat{\mathbf{A}}^* = \mathbf{U}_r \Sigma_r^{1/2}, \quad \widehat{\mathbf{B}}^* = \Sigma_r^{1/2} \mathbf{V}_r^\top$$

1035 achieve the global optimum of $\min_{\widehat{\mathbf{A}}, \widehat{\mathbf{B}}} \|\widetilde{\mathbf{E}} - \widehat{\mathbf{A}} \widehat{\mathbf{B}}\|_F^2$, with minimum value $\sum_{i>r} \sigma_i(\widetilde{\mathbf{E}})^2$ Eckart &
 1036 Young (1936); Mirsky (1960); Golub & Loan (2013). In original variables,
 1037

$$1038 \quad \mathbf{A}^* = \widehat{\mathbf{A}}^* = \mathbf{U}_r \Sigma_r^{1/2}, \quad \mathbf{B}^* = \widehat{\mathbf{B}}^* \Sigma_x^{-1/2} = \Sigma_r^{1/2} \mathbf{V}_r^\top \Sigma_x^{-1/2},$$

1039 and
 1040

$$1041 \quad \text{row}(\mathbf{B}^*) = \text{row}(\mathbf{V}_r^\top \Sigma_x^{-1/2}). \quad (\text{A.2.5})$$

1042 **Proof.** By Theorem A.2.3,

$$1044 \quad \min_{\mathbf{A}, \mathbf{B}} \|(\mathbf{E}_{\text{cat}} - \mathbf{AB}) \Sigma_x^{1/2}\|_F^2 = \min_{\widehat{\mathbf{A}}, \widehat{\mathbf{B}}} \|\widetilde{\mathbf{E}} - \widehat{\mathbf{A}} \widehat{\mathbf{B}}\|_F^2.$$

1046 Left/right orthogonal invariance of the Frobenius norm reduces the problem to $\min_{\text{rank}(\mathbf{Y}) \leq r} \|\Sigma -$
 1047 $\mathbf{Y}\|_F^2$, solved by the truncated SVD $\mathbf{Y}^* = \Sigma_r \oplus \mathbf{0}$; hence $\mathbf{X}^* = \mathbf{U}_r \Sigma_r \mathbf{V}_r^\top$ Eckart & Young (1936);
 1048 Mirsky (1960). Choosing $\widehat{\mathbf{A}}^* = \mathbf{U}_r \Sigma_r^{1/2}$ and $\widehat{\mathbf{B}}^* = \Sigma_r^{1/2} \mathbf{V}_r^\top$ produces $\mathbf{X}^* = \widehat{\mathbf{A}}^* \widehat{\mathbf{B}}^*$. Returning
 1049 to original variables gives the stated $(\mathbf{A}^*, \mathbf{B}^*)$ and the row-space identity (cf. Fan (1949); Golub &
 1050 Loan (2013)). \square

1051 In whitened variables: $(\widehat{\mathbf{A}}^*)^\top \widehat{\mathbf{A}}^* = \Sigma_r$ and $\widehat{\mathbf{B}}^* (\widehat{\mathbf{B}}^*)^\top = \Sigma_r$. In original variables: $(\mathbf{A}^*)^\top \mathbf{A}^* =$
 1052 Σ_r and $\mathbf{B}^* \Sigma_x (\mathbf{B}^*)^\top = \Sigma_r$ Golub & Loan (2013). For any orthogonal $\mathbf{R} \in \mathbb{R}^{r \times r}$, $(\mathbf{AR}, \mathbf{R}^\top \mathbf{B})$
 1053 attains the same objective value Golub & Loan (2013).

1055 The truncated SVD is globally optimal; the balanced factorization is well-conditioned, and the opti-
 1056 mal shared right subspace is $\text{row}(\mathbf{V}_r^\top \Sigma_x^{-1/2})$.
 1057

1058 **LEMMA A.2.6 (SINGULAR Σ_x AND PSEUDOINVERSE WHITENING).**

1059 In this subsection, we extend all results to rank-deficient Σ_x by showing the objective depends only
 1060 on $\text{Range}(\Sigma_x)$ and that pseudoinverse whitening preserves the conclusions on that subspace Ben-
 1061 Israel & Greville (2003); Penrose (1955).

1063 Let $\Sigma_x = \mathbf{Q} \Lambda \mathbf{Q}^\top$ with $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_{r_+}, 0, \dots, 0)$. Define
 1064

$$1065 \quad \Sigma_x^{1/2} = \mathbf{Q} \Lambda^{1/2} \mathbf{Q}^\top, \quad \Sigma_x^{-1/2} = \mathbf{Q} \Lambda^{\dagger/2} \mathbf{Q}^\top,$$

1067 where $\Lambda^{\dagger/2}$ applies $\lambda_i^{-1/2}$ to $\lambda_i > 0$ and 0 otherwise. Then the objective $\|(\mathbf{E}_{\text{cat}} - \mathbf{AB}) \Sigma_x^{1/2}\|_F^2$
 1068 depends only on $\text{Range}(\Sigma_x)$, and Theorems A.2.1–A.2.5 hold unchanged on that subspace.
 1069

1070 **Proof.** Let $\mathbf{Q} = [\mathbf{Q}_r \ \mathbf{Q}_0]$ with \mathbf{Q}_r spanning $\text{Range}(\Sigma_x)$ and $\Sigma_x^{1/2} = \mathbf{Q}_r \Lambda_r^{1/2} \mathbf{Q}_r^\top$. Then
 1071

$$1072 \quad \|(\mathbf{E} - \mathbf{AB}) \Sigma_x^{1/2}\|_F^2 = \|(\mathbf{E} \mathbf{Q}_r - \mathbf{A}(\mathbf{B} \mathbf{Q}_r)) \Lambda_r^{1/2}\|_F^2,$$

1073 which is the same Frobenius objective restricted to $\text{Range}(\Sigma_x)$. Components along \mathbf{Q}_0 vanish under
 1074 $\Sigma_x^{1/2}$ and contribute nothing. \square

1075 Pseudoinverse whitening discards the nullspace; all conclusions hold on $\text{Range}(\Sigma_x)$.
 1076

1077 In our implementation, to estimate and stabilize Σ_x , we perform ridge/shrinkage regularization
 1078 ($\widehat{\Sigma}_x \leftarrow \widehat{\Sigma}_x + \varepsilon \mathbf{I}$) while using diagonal approximations (cf. Bishop (2006); Anderson (2003); Ledoit
 1079 & Wolf (2004); Hoerl & Kennard (1970)) with mini-batch and sliding-window since computing full
 covariances are costly.

1080 A.3 QR REDUCTION: SMALL-CORE EQUIVALENCE AND GLOBAL SOLUTION (PROOF)
10811082 The covariance-aligned objective
1083

1084
$$\min_{\mathbf{A} \in \mathbb{R}^{m \times r}, \mathbf{B} \in \mathbb{R}^{r \times d}} \|(\mathbf{E}_{\text{cat}} - \mathbf{AB}) \Sigma_{\mathbf{x}}^{1/2}\|_F^2 \quad (\text{A.3.1})$$

1085

1086 can be solved without ever forming the tall whitened matrix $\tilde{\mathbf{E}} := \mathbf{E}_{\text{cat}} \Sigma_{\mathbf{x}}^{1/2} \in \mathbb{R}^{m \times d}$. A thin
1087 QR $\mathbf{E}_{\text{cat}} = \mathbf{Q}_e \mathbf{R}_e$ (with $\mathbf{Q}_e^\top \mathbf{Q}_e = \mathbf{I}_d$) collects all the information relevant to Eq. A.3.1 into the
1088 $d \times d$ core $\mathbf{M} := \mathbf{R}_e \Sigma_{\mathbf{x}}^{1/2}$ because $\tilde{\mathbf{E}} = \mathbf{Q}_e \mathbf{M}$ and the Frobenius norm is left-orthogonally invariant
1089 ($\|\mathbf{Q}\mathbf{Z}\|_F = \|\mathbf{Z}\|_F$ when $\mathbf{Q}^\top \mathbf{Q} = \mathbf{I}$) Golub & Loan (2013); Trefethen & Bau (1997). Thus we can
1090 reduce the large problem to an equivalent $d \times d$ problem, apply standard SVD/EYM analysis on
1091 the core, and lift the solution back (QR reduction to a core matrix; see also Halko et al. (2011);
1092 Martinsson & Tropp (2020) for randomized variants).
10931094 LEMMA A.3.1 (OPTIMAL \mathbf{A} LIES IN $\text{col}(\mathbf{Q}_e)$).
10951096 For any \mathbf{A} , decompose $\mathbf{A} = \mathbf{Q}_e \hat{\mathbf{A}} + \mathbf{A}_\perp$ with $\mathbf{Q}_e^\top \mathbf{A}_\perp = \mathbf{0}$ and set $\hat{\mathbf{B}} := \mathbf{B} \Sigma_{\mathbf{x}}^{1/2}$. Then
1097

1098
$$\|\tilde{\mathbf{E}} - \mathbf{A} \hat{\mathbf{B}}\|_F^2 = \|\mathbf{Q}_e(\mathbf{M} - \hat{\mathbf{A}} \hat{\mathbf{B}})\|_F^2 + \|\mathbf{A}_\perp \hat{\mathbf{B}}\|_F^2 \geq \|\mathbf{Q}_e(\mathbf{M} - \hat{\mathbf{A}} \hat{\mathbf{B}})\|_F^2,$$

1099

1100 where $\tilde{\mathbf{E}} = \mathbf{Q}_e \mathbf{M}$ and $\mathbf{M} := \mathbf{R}_e \Sigma_{\mathbf{x}}^{1/2}$. Hence any global minimizer satisfies $\mathbf{A}_\perp = \mathbf{0}$, i.e., $\mathbf{A}^* =$
1101 $\mathbf{Q}_e \hat{\mathbf{A}}^*$. It shrinks the search space for \mathbf{A} to the d -dimensional column space of \mathbf{Q}_e ; any component
1102 orthogonal to $\text{col}(\mathbf{Q}_e)$ only increases the loss. (Orthogonal decomposition/Pythagorean property of
1103 the Frobenius inner product; cf. Golub & Loan (2013); Trefethen & Bau (1997).)
11041105 **Proof.** Use $\tilde{\mathbf{E}} = \mathbf{Q}_e \mathbf{M}$ and orthogonality: $\mathbf{Q}_e^\top(\mathbf{Q}_e(\cdot)) = (\cdot)$ and $\mathbf{Q}_e^\top(\mathbf{A}_\perp \hat{\mathbf{B}}) = \mathbf{0}$, so the two
1106 terms are orthogonal in the Frobenius inner product and the squared norm splits. The minimum
1107 occurs at $\mathbf{A}_\perp = \mathbf{0}$. \square
11081109 THEOREM A.3.2 (CORE EQUIVALENCE).
11101111 By Lemma A.3.1 and left-orthogonal invariance of $\|\cdot\|_F$ (i.e., $\|\mathbf{Q}\mathbf{Z}\|_F = \|\mathbf{Z}\|_F$ for orthogonal \mathbf{Q} ;
1112 Golub & Loan (2013)),
1113

1114
$$\min_{\mathbf{A}, \mathbf{B}} \|(\mathbf{E}_{\text{cat}} - \mathbf{AB}) \Sigma_{\mathbf{x}}^{1/2}\|_F^2 = \min_{\hat{\mathbf{A}}, \hat{\mathbf{B}}} \|\mathbf{M} - \hat{\mathbf{A}} \hat{\mathbf{B}}\|_F^2, \quad \mathbf{M} = \mathbf{R}_e \Sigma_{\mathbf{x}}^{1/2}. \quad (\text{A.3.2})$$

1115

1116 Any minimizer $(\hat{\mathbf{A}}^*, \hat{\mathbf{B}}^*)$ lifts to a minimizer of the original problem via
1117

1118
$$\mathbf{A}^* = \mathbf{Q}_e \hat{\mathbf{A}}^*, \quad \mathbf{B}^* = \hat{\mathbf{B}}^* \Sigma_{\mathbf{x}}^{-1/2}, \quad (\text{A.3.3})$$

1119

1120 where $\Sigma_{\mathbf{x}}^{-1/2}$ denotes a (pseudo-)inverse square root when $\Sigma_{\mathbf{x}}$ is singular Ben-Israel & Greville
1121 (2003).
11221123 **Proof.** Restrict to $\mathbf{A} = \mathbf{Q}_e \hat{\mathbf{A}}$ (Eq. A.3.3). Then $\|(\mathbf{E}_{\text{cat}} - \mathbf{AB}) \Sigma_{\mathbf{x}}^{1/2}\|_F = \|\mathbf{Q}_e(\mathbf{M} - \hat{\mathbf{A}} \hat{\mathbf{B}})\|_F =$
1124 $\|\mathbf{M} - \hat{\mathbf{A}} \hat{\mathbf{B}}\|_F$. The lifting follows by inverting the change $\hat{\mathbf{B}} = \mathbf{B} \Sigma_{\mathbf{x}}^{1/2}$. \square
11251126 COROLLARY A.3.3 (PRESERVATION OF NONZERO SINGULAR VALUES AND RIGHT SINGULAR
1127 VECTORS).
11281129 Since $(\mathbf{Q}_e \mathbf{M})^\top (\mathbf{Q}_e \mathbf{M}) = \mathbf{M}^\top \mathbf{M}$, $\tilde{\mathbf{E}} = \mathbf{Q}_e \mathbf{M}$ and \mathbf{M} share the same nonzero singular values
1130 and the same right singular vectors. Hence the SVD of \mathbf{M} directly yields the optimal shared right
1131 subspace for the covariance-aligned objective (orthogonal invariance of SVD; e.g., Golub & Loan
1132 (2013); Trefethen & Bau (1997)).
1133**Proof.** Immediate from $\mathbf{Q}_e^\top \mathbf{Q}_e = \mathbf{I}_d$. \square

1134 THEOREM A.3.4 (BALANCED FACTORS, GLOBAL MINIMIZER, AND LIFTING).

1135 Let $\mathbf{M} = \mathbf{U}\Sigma\mathbf{V}^\top$ be an SVD and $(\mathbf{U}_r, \Sigma_r, \mathbf{V}_r)$ the top- r blocks. Then

$$1137 \quad \widehat{\mathbf{A}}^* = \mathbf{U}_r \Sigma_r^{1/2}, \quad \widehat{\mathbf{B}}^* = \Sigma_r^{1/2} \mathbf{V}_r^\top \quad (\text{A.3.4})$$

1139 achieve the global minimum of $\|\mathbf{M} - \widehat{\mathbf{A}}\widehat{\mathbf{B}}\|_F^2$ by the Eckart–Young–Mirsky theorem Eckart & Young
1140 (1936); Mirsky (1960); Golub & Loan (2013). Lifting to the original variables gives

$$1141 \quad \mathbf{A}^* = \mathbf{Q}_e \mathbf{U}_r \Sigma_r^{1/2}, \quad \mathbf{B}^* = \Sigma_r^{1/2} \mathbf{V}_r^\top \Sigma_x^{-1/2}. \quad (\text{A.3.5})$$

1142 The minimum value is $\|\mathbf{M} - \mathbf{U}_r \Sigma_r \mathbf{V}_r^\top\|_F^2$, and the shared right subspace is $\text{row}(\mathbf{B}^*) =$
1143 $\text{span}(\mathbf{V}_r^\top \Sigma_x^{-1/2})$ (cf. Ky Fan (1949)).

1144 It provides a closed-form global minimizer and a numerically well-conditioned (balanced) factorization.

1145 Truncated SVD is optimal; balancing ($\Sigma_r^{1/2}$) improves conditioning and scale regularity Golub &
1146 Loan (2013).

1147 **Proof.** Apply EYM to the core problem from Eq. A.3.2; choose balanced factors so that
1148 $\mathbf{U}_r \Sigma_r \mathbf{V}_r^\top = \widehat{\mathbf{A}}^* \widehat{\mathbf{B}}^*$. Use Eq. A.3.3 to obtain $(\mathbf{A}^*, \mathbf{B}^*)$. \square

1153 This process makes the thin QR cost $\mathcal{O}(md^2)$, while forming/using the core costs $\mathcal{O}(d^3)$ (or $\mathcal{O}(d^2)$)
1154 if $\Sigma_x^{1/2}$ is precomputed/structured). All subsequent optimization is on the $d \times d$ core Golub & Loan
1155 (2013); Trefethen & Bau (1997). After computing \mathbf{M} , we do not materialize \mathbf{M} ; instead we keep
1156 $\mathbf{z} \mapsto \mathbf{M}\mathbf{z} = \mathbf{R}_e(\Sigma_x^{1/2}\mathbf{z})$ and $\mathbf{y} \mapsto \mathbf{M}^\top \mathbf{y} = \Sigma_x^{1/2}(\mathbf{R}_e^\top \mathbf{y})$, and pass these to RSVD Halko et al.
1157 (2011); Martinsson & Tropp (2020).

1158 From Eq. A.3.5, $\text{row}(\mathbf{B}^*) = \text{span}(\mathbf{V}_r^\top \Sigma_x^{-1/2})$ defines the shared right subspace. In GLOWQ, this
1159 subspace is exactly the group-shared projection used to compute and cache $\mathbf{R} = \mathbf{B}_{\text{shared}} \mathbf{X}$ once
1160 per input-sharing group, thereby enabling efficient $\mathbf{A}_i \mathbf{R}$ reuse during inference while preserving
1161 expressivity via module-specific \mathbf{A}_i (cf. Sec. 3.3 and the Ky Fan view in Theorem A.1.3).

1163 A.4 RSVD ACCURACY GUARANTEES

1165 Let the core matrix be $\mathbf{M} := \mathbf{R}_e \Sigma_x^{1/2} \in \mathbb{R}^{d \times d}$ as defined by the QR reduction in Appendix A.3.
1166 We target rank $r \leq d$ with oversampling $p \geq 2$ and power iterations $q \geq 0$. By the core equivalence
1167 and preservation results, accuracy on \mathbf{M} transfers verbatim to the covariance-aligned objective.

1169 ALGORITHM A.4.1 - RSVD ON THE CORE \mathbf{M} .

1171 It computes the dominant right subspace (which defines the shared right factor) on the small $d \times$
1172 d core without ever materializing the tall whitened matrix (standard RSVD; (Halko et al., 2011;
1173 Martinsson & Tropp, 2020)).

1174 Procedure.

- 1176 (i) $\Omega \sim \mathcal{N}(0, 1)^{d \times (r+p)}$, $\mathbf{Y} \leftarrow \mathbf{M}\Omega$;
- 1177 (ii) Power iterations: repeat q times $\{\mathbf{Y} \leftarrow \mathbf{M}(\mathbf{M}^\top \mathbf{Y})\}$ with re-orthonormalization;
- 1178 (iii) $\mathbf{Q} \leftarrow \text{orth}(\mathbf{Y})$, $\mathbf{B} \leftarrow \mathbf{Q}^\top \mathbf{M}$;
- 1179 (iv) $\mathbf{B} = \widetilde{\mathbf{U}} \Sigma \mathbf{V}^\top$, $\mathbf{U} \leftarrow \mathbf{Q} \widetilde{\mathbf{U}}$; truncate to $(\mathbf{U}_r, \Sigma_r, \mathbf{V}_r)$;
- 1180 (v) Balanced core factors: $\widehat{\mathbf{A}}^* = \mathbf{U}_r \Sigma_r^{1/2}$, $\widehat{\mathbf{B}}^* = \Sigma_r^{1/2} \mathbf{V}_r^\top$.

1183 Find a good range \mathbf{Q} via randomized sketching (with optional power iterations), then refine by a
1184 small SVD on $\mathbf{Q}^\top \mathbf{M}$. *Justification.* Within the subspace $\mathcal{R}(\mathbf{Q})$, the best rank- r approximation is
1185 the truncated SVD of $\mathbf{Q}^\top \mathbf{M}$; lifting by \mathbf{Q} yields $\mathbf{U}_r \Sigma_r \mathbf{V}_r^\top$ as the optimal restricted approximation
1186 (Golub & Loan, 2013, Ch. 2). The randomized sketch ensures (in expectation or with high
1187 probability) that $\mathcal{R}(\mathbf{Q})$ captures the dominant right subspace of \mathbf{M} (Halko et al., 2011; Martinsson
1188 & Tropp, 2020). \square

1188 THEOREM A.4.2 (FROBENIUS ERROR, EXPECTATION).
 1189

1190 Let $\mathbf{M} = \mathbf{U}\Sigma\mathbf{V}^\top$ with singular values $\sigma_1 \geq \dots \geq \sigma_d$. For $p \geq 2$ and $q = 0$,

$$1192 \mathbb{E} \|\mathbf{M} - \mathbf{Q}\mathbf{Q}^\top \mathbf{M}\|_F \leq \left(1 + \frac{r}{p-1}\right)^{1/2} \left(\sum_{j>r} \sigma_j^2\right)^{1/2}. \quad (\text{A.4.1})$$

1194 (Halko–Martinsson–Tropp; e.g., (Halko et al., 2011, Thm. 10.5))
 1195

1196 It quantifies that RSVD matches the optimal tail energy up to a mild factor depending only on (r, p) .
 1197

1198 **Proof.** Write $\mathbf{M} = \mathbf{U} \begin{bmatrix} \Sigma_1 & \mathbf{0} \\ \mathbf{0} & \Sigma_2 \end{bmatrix} \mathbf{V}^\top$ with $\Sigma_1 \in \mathbb{R}^{r \times r}$ and Σ_2 the tail. Let $\mathbf{V}^\top \Omega = \begin{bmatrix} \Omega_1 \\ \Omega_2 \end{bmatrix}$ and
 1199 $\mathbf{Y} = \mathbf{M}\Omega$. Standard analysis of Gaussian sketches gives $\|(\mathbf{I} - \mathbf{P}_Q)\mathbf{M}\|_F \leq \|\Sigma_2\|_F \|\Omega_2 \Omega_1^\dagger\|_F$,
 1200 and $\mathbb{E} \|\Omega_2 \Omega_1^\dagger\|_F^2 \leq r/(p-1)$ for $p \geq 2$ (Halko et al., 2011). Taking square roots and expectations
 1201 yields Eq. A.4.1. \square
 1202

1204 THEOREM A.4.3 (SPECTRAL ERROR WITH q POWER ITERATIONS).
 1205

1206 For $q \geq 0$ and a modest constant $C_{r,p}$ (depending gently on r, p),

$$1207 \|\mathbf{M} - \mathbf{U}_r \Sigma_r \mathbf{V}_r^\top\|_2 \lesssim C_{r,p}^{1/(2q+1)} \sigma_{r+1}. \quad (\text{A.4.2})$$

1209 (Cf. (Halko et al., 2011; Martinsson & Tropp, 2020; Musco & Musco, 2015).)

1210 Power iterations shrink the subspace-angle gap geometrically toward the optimal σ_{r+1} bound.
 1211

1212 Each power iteration reduces the gap factor roughly by a $(\cdot)^{1/(2q+1)}$ exponent toward σ_{r+1} .
 1213

1214 **Proof.** After q power steps, $\mathbf{Y} = (\mathbf{M}\mathbf{M}^\top)^q \mathbf{M}\Omega = \mathbf{U} \Sigma^{2q+1} (\mathbf{V}^\top \Omega)$. Block-partitioning
 1215 $\mathbf{V}^\top \Omega = \begin{bmatrix} \Omega_1 \\ \Omega_2 \end{bmatrix}$ and analyzing principal angles between the exact and sketched right subspaces
 1216 gives
 1217

$$1218 \|\mathbf{I} - \mathbf{P}_Q\mathbf{M}\|_2 \leq \|\Sigma_2\|_2 \left\| \Sigma_2^{2q} \Omega_2 (\Omega_1)^\dagger \Sigma_1^{-2q} \right\|_2^{1/(2q+1)}.$$

1220 Bounding the Gaussian pseudo-inverse term by $C_{r,p}$ and using $\|\Sigma_2^{2q} \Sigma_1^{-2q}\|_2 = (\sigma_{r+1}/\sigma_r)^{2q}$
 1221 yields Eq. A.4.2. \square
 1222

1223 COROLLARY A.4.4 (TRANSFER TO THE COVARIANCE-ALIGNED OBJECTIVE).
 1224

1225 By Theorem A.3.2 and Corollary A.3.3.

$$1226 \|\mathbf{E}_{\text{cat}} - \mathbf{A}^* \mathbf{B}^*\|_F = \|\mathbf{M} - \mathbf{U}_r \Sigma_r \mathbf{V}_r^\top\|_F. \quad (\text{A.4.3})$$

1228 It links RSVD accuracy on the core directly to the original covariance-aligned objective.
 1229

1230 Core RSVD error bounds become the error bounds for the original problem, verbatim.
 1231

1232 **Proof.** We have $\tilde{\mathbf{E}} = \mathbf{E}_{\text{cat}} \Sigma_{\mathbf{x}}^{1/2} = \mathbf{Q}_e \mathbf{M}$ and Frobenius norms are left-orthogonally invariant; the
 1233 optimal truncated approximation on \mathbf{M} corresponds under lifting to the optimal approximation of
 1234 $(\mathbf{E}_{\text{cat}} - \mathbf{A}\mathbf{B})\Sigma_{\mathbf{x}}^{1/2}$, yielding Eq. A.4.3. \square
 1235

1236 PROPOSITION (Q-LESS LIFTING: BLOCKWISE RECOVERY OF \mathbf{A}_i^*).
 1237

1238 Write $\mathbf{E}_{\text{cat}} = [\mathbf{E}_1; \dots; \mathbf{E}_m]$ and $\mathbf{A}^* = [\mathbf{A}_1^*; \dots; \mathbf{A}_m^*]$ conformably. At fixed \mathbf{B}^* , each block admits
 1239 the closed form

$$1240 \mathbf{A}_i^* = \mathbf{E}_i (\mathbf{B}^*)^\top (\mathbf{B}^* (\mathbf{B}^*)^\top)^\dagger, \quad i = 1, \dots, m, \quad (\text{A.4.4})$$

1241 so the tall orthonormal factor \mathbf{Q}_e need not be stored (least-squares with pseudoinverse; cf. (Björck,
 1996; Ben-Israel & Greville, 2003; Golub & Loan, 2013)).

1242 It economizes memory: per-block factors are recovered directly from $(\mathbf{E}_i, \mathbf{B}^*)$ without retaining
 1243 \mathbf{Q}_e .

1244 Per-block least-squares with a pseudoinverse yields \mathbf{A}_i^* using only $(\mathbf{E}_i, \mathbf{B}^*)$.

1245 **Proof.** For each block, minimize $\|\mathbf{E}_i - \mathbf{A}_i \mathbf{B}^*\|_F^2$. The first-order optimality condition is
 1246 $\mathbf{A}_i^* \mathbf{B}^* (\mathbf{B}^*)^\top = \mathbf{E}_i (\mathbf{B}^*)^\top$. Multiplying on the right by the Moore–Penrose pseudoinverse gives
 1247 the minimal-norm solution $\mathbf{A}_i^* = \mathbf{E}_i (\mathbf{B}^*)^\top (\mathbf{B}^* (\mathbf{B}^*)^\top)^\dagger$, which is precisely Eq. A.4.4. \square

1250 B EFFECT OF RIGHT-WEIGHTED SHARED B

1251 In this section, we analyze the effect of the right-weighted shared-B on the GlowQ’s error correction
 1252 with the following procedure.

1253 **Procedure.** (1) Using calibration inputs $\{\mathbf{x}_n\}_{n=1}^N \subset \mathbb{R}^d$, estimate the layer input covariance

$$1254 \hat{\Sigma}_{\mathbf{x}} = \frac{1}{N} \sum_n \mathbf{x}_n \mathbf{x}_n^\top \quad (\text{optionally: } \hat{\Sigma}_{\mathbf{x}} \leftarrow \hat{\Sigma}_{\mathbf{x}} + \epsilon \mathbf{I}).$$

1255 (2) For each module $i \in \{q, k, v, \text{gate, up}\}$, form the quantization-error matrix $\mathbf{E}_i \in \mathbb{R}^{O_i \times d}$ and
 1256 the row-stack $\mathbf{E}_{\text{cat}} = [\mathbf{E}_1; \dots; \mathbf{E}_m] \in \mathbb{R}^{m \times d}$.

1257 (3) **Cov-aligned (whitened):** compute SVDs of $\tilde{\mathbf{E}}_i := \mathbf{E}_i \hat{\Sigma}_{\mathbf{x}}^{1/2}$ and $\tilde{\mathbf{E}}_{\text{cat}} := \mathbf{E}_{\text{cat}} \hat{\Sigma}_{\mathbf{x}}^{1/2}$, and take
 1258 the top- r right bases $\mathbf{V}_{i,r}$ and \mathbf{V}_r .

1259 (4) **Unweighted (no-cov):** repeat the same without whitening to obtain $\mathbf{V}_{i,r}^{(\text{no-cov})}$ and $\mathbf{V}_r^{(\text{no-cov})}$.

1260 (5) For each module, form the absolute cross-basis cosine matrix

$$1261 \mathbf{C}_i = |\mathbf{V}_r^\top \mathbf{V}_{i,r}| \in \mathbb{R}^{r \times r},$$

1262 Hungarian-reorder it to maximize the diagonal sum, and visualize as heatmaps.

1263 **Impact of optimization with the right weighted objective.** As illustrated in Fig. 5 and 6, the
 1264 whitened condition produces a bright near-diagonal across all groups (Q/K/V and MLP gate/up),
 1265 indicating a one-to-one alignment between the shared right subspace row (B_{shared}) and each module’s
 1266 right subspace (up to sign/permuation). The effect is strongest for Q/K, and slightly more diffuse for
 1267 V and for MLP (gate/up), but remains concentrated on the leading axes. In contrast, the unweighted
 1268 condition yields noise-like patterns with no diagonal structure.

1269 Right-side covariance weighting is crucial for estimating a shared B under anisotropic inputs: it
 1270 exposes a common right subspace across modules that ingest the same input tensor. This validates
 1271 the shared- B assumption and directly motivates our ABx caching strategy, i.e., computing $R =$
 1272 $B_{\text{shared}} X$ once per group and reusing $A_i R$ across modules. Unweighted stacked SVD fails to reveal
 1273 this alignment, weakening both the shared- B premise and the practical caching benefit.

1274 B.1 GROUP-CACHED (WEIGHTED STACKED RSVD) VS. LAYER-WISE (WEIGHTED RSVD)

1275 Table 7: Perplexity (lower is better) across model families on WikiText-2. Layer-wise applies layer-
 1276 wise SVD correction, whereas GLOWQ applies group-wise SVD with a shared right factor B ;
 1277 GLOWQ (Selective restore) denotes selective group restoration.

Method	LLaMA 2		LLaMA 3		Qwen 2.5		Qwen 3		OPT		Vicuna		Mistral
	7B	13B	3.2-3B	3.1-8B	7B	14B	8B	14B	1.3B	6.7B	7B	13B	7B
LAYERWISE	5.58	4.96	8.15	6.59	7.06	5.64	9.92	8.80	15.05	11.00	6.89	6.03	5.42
GlowQ	5.58	4.96	8.16	6.59	7.07	5.64	9.90	8.80	15.06	11.00	6.90	6.02	5.42
GlowQ-S	5.60	4.96	8.22	6.62	7.09	5.68	9.97	8.89	15.19	11.00	6.90	6.04	5.45

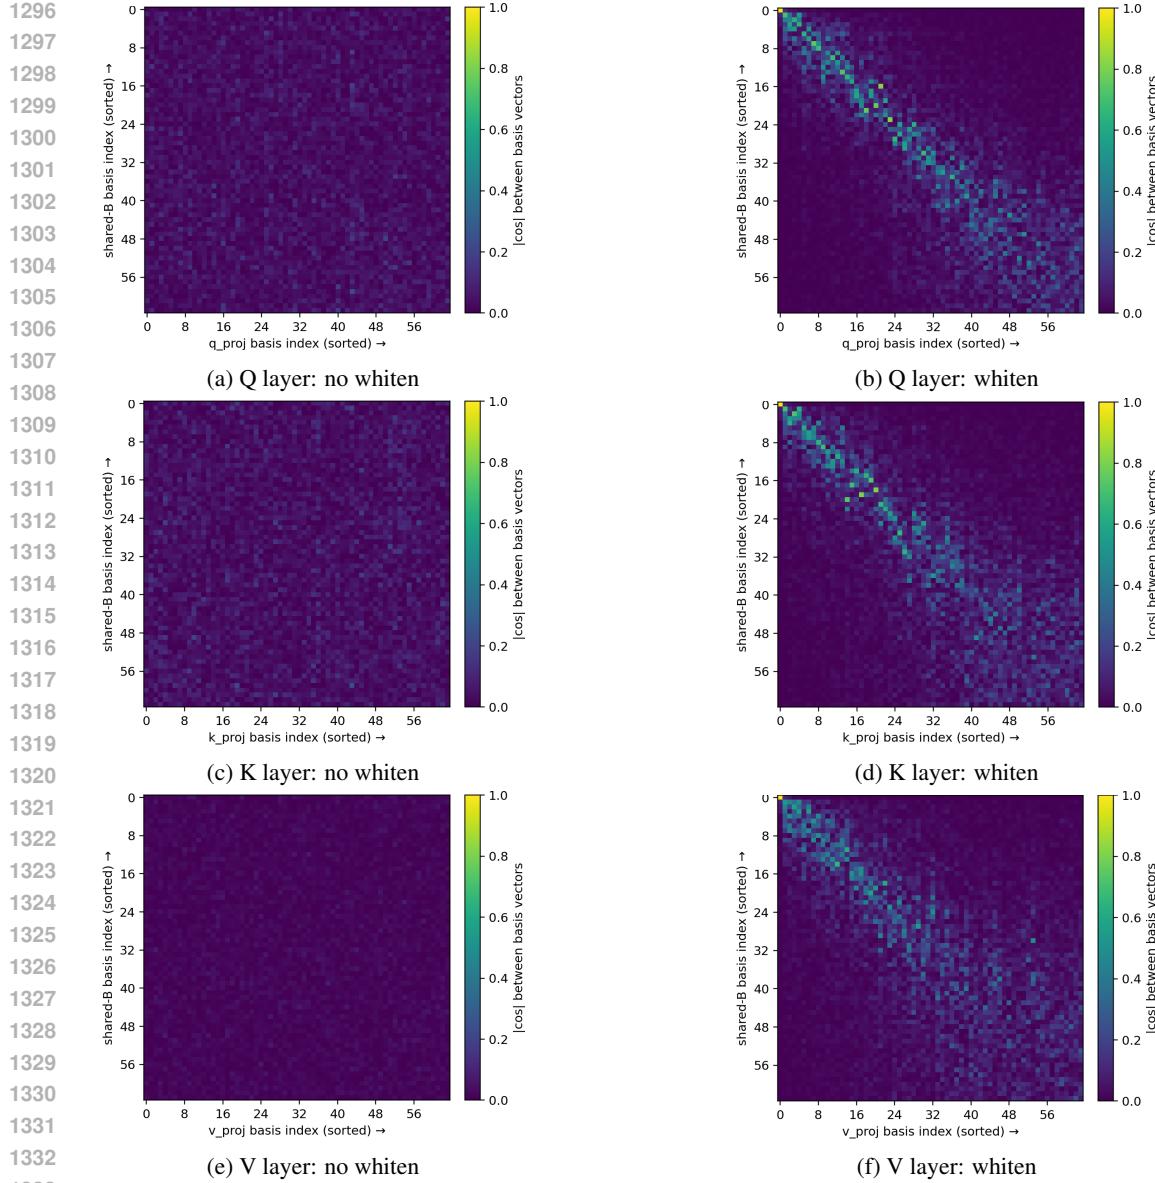


Figure 5: Whitening vs. non-whitening alignment matrices. For LLaMA 3.2-3B, we estimate a shared right basis B_{shared} from the stacked error either without covariance weighting (E_{cat} , left panels) or with covariance-aware whitening ($E_{\text{cat}} \Sigma_x^{1/2}$, right panels). Each heatmap shows the absolute basis alignment between row(B_{shared}) and the per-module right subspace for Q, K, V; brighter values denote larger absolute inner products. DiagScore and Affinity summaries are reported in the main text.

Results on Table 7. Across 13 model-size combinations, GlowQ and Layer-wise yield essentially identical perplexity: the mean gap is +0.001 ppl on average, with per-family fluctuations confined to ± 0.02 ppl. By design, GlowQ-S (Selective restore) trades a bit of accuracy for efficiency, trailing Layer-wise by +0.04 ppl on average. In short, the full shared- \mathbf{B} configuration matches layer-wise ($\mathbf{A}_i, \mathbf{B}_i$) on WikiText-2 without systematic degradation, while the selective variant incurs a small, consistent increase in ppl.

Observation on Fig. 5, 6. The covariance-aligned cross-basis heatmaps exhibit an almost perfectly diagonal structure after Hungarian matching, indicating a near one-to-one correspondence between the shared right subspace and each module's top- r directions. Whitening aligns input usage so that the shared \mathbf{B} spans (practically) the same right-singular space that the individual \mathbf{B}_i

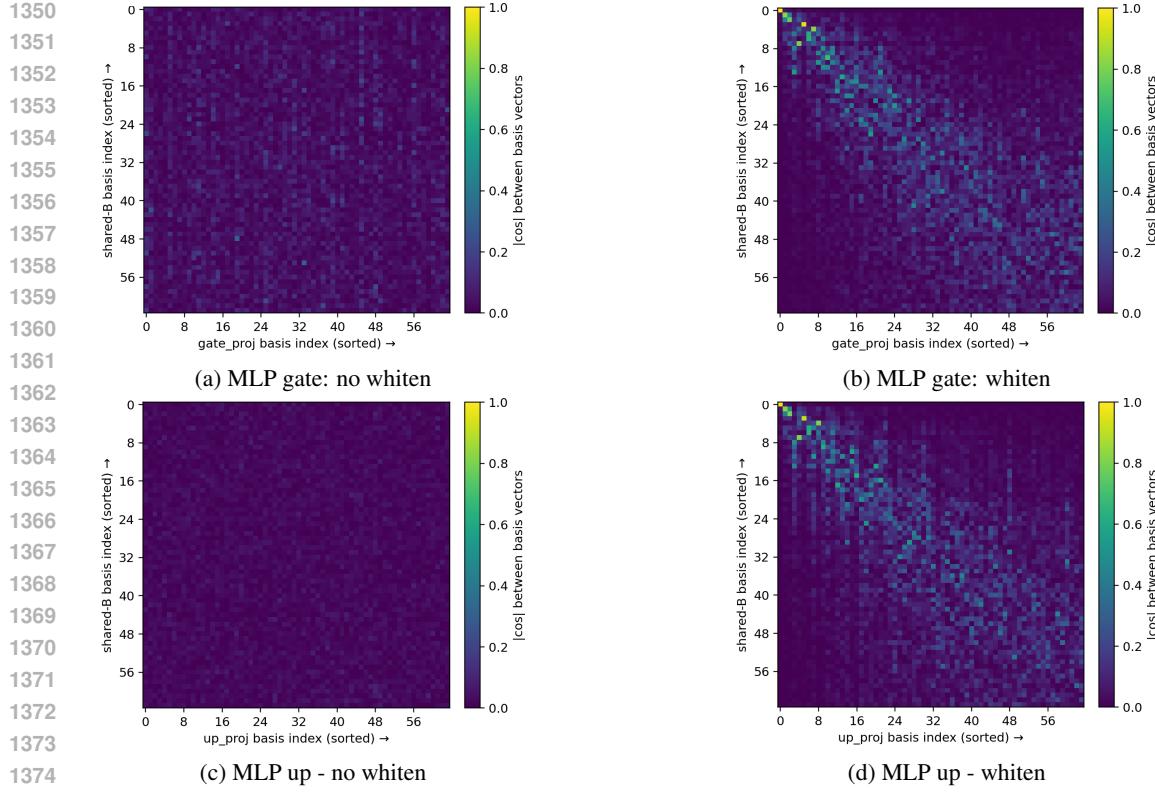


Figure 6: Whitening vs. non-whitening alignment matrices. MLP (up/gate).

would select, explaining why GlowQ’s perplexity tracks layerwise so closely, and why GlowQ-S, restoring only a subset, shows the small upward shift in ppl.

Observation on Fig. 7, 8, 9. For the MoE FFN of Qwen1.5-MoE-A2.7B, the covariance-aligned cross-basis heatmaps show the same qualitative behavior as in the dense models once whitening is enabled. Without whitening, all panels (expert gate/up, shared gate/up, and MoE attention) look almost uniformly dark, indicating that the shared right subspace and each expert’s local top- r directions are essentially uncorrelated. After whitening and Hungarian matching, the heatmaps become sharply diagonal for both the representative expert (e.g., `expert59_gate_proj` / `expert59_up_proj`) and the shared- B MLP/attention blocks, revealing a near one-to-one alignment between the shared basis and each expert’s own error subspace. This confirms that, once inputs are whitened, the grouped MoE FFNs and the shared MLP effectively live in the same right-singular space, so a single shared $\mathbf{B}_{\text{shared}}$ can serve all experts with only small residual mismatch. Consequently, GlowQ can compress all experts and the shared MLP with one shared right-hand matrix while closely tracking the layerwise baseline in perplexity, explaining the tiny +0.02 PPL gap we observe on Qwen1.5-MoE-A2.7B.

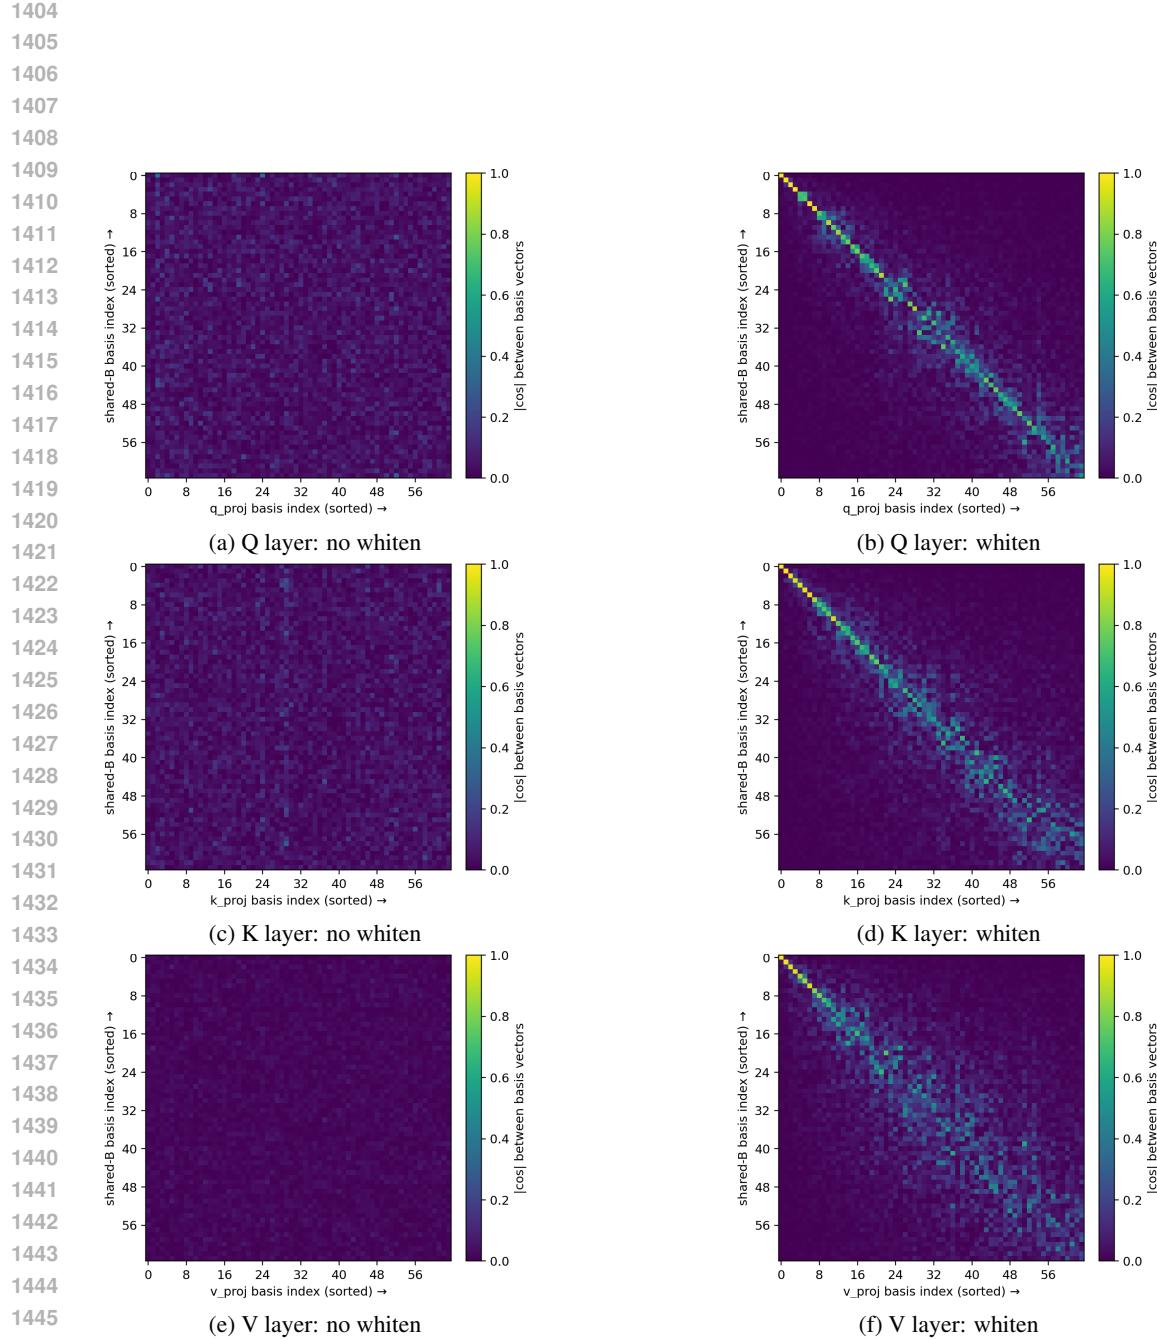


Figure 7: Whitening vs. non-whitening alignment matrices for Q/K/V in Qwen1.5-MoE-A2.7B. As in Fig. 5, we estimate a shared right basis B_{shared} from the stacked attention-projection error, either from the raw error (“no whiten”, left panels) or after covariance-aware whitening $E_{\text{cat}}\Sigma_x^{1/2}$ (“whiten”, right panels). Each heatmap shows the absolute basis alignment between row(B_{shared}) and the per-module right subspace for Q, K, and V; brighter values denote larger absolute inner products. Whitening again yields a sharply diagonally dominant structure, indicating that a single covariance-aligned basis captures the dominant error directions across Q/K/V.

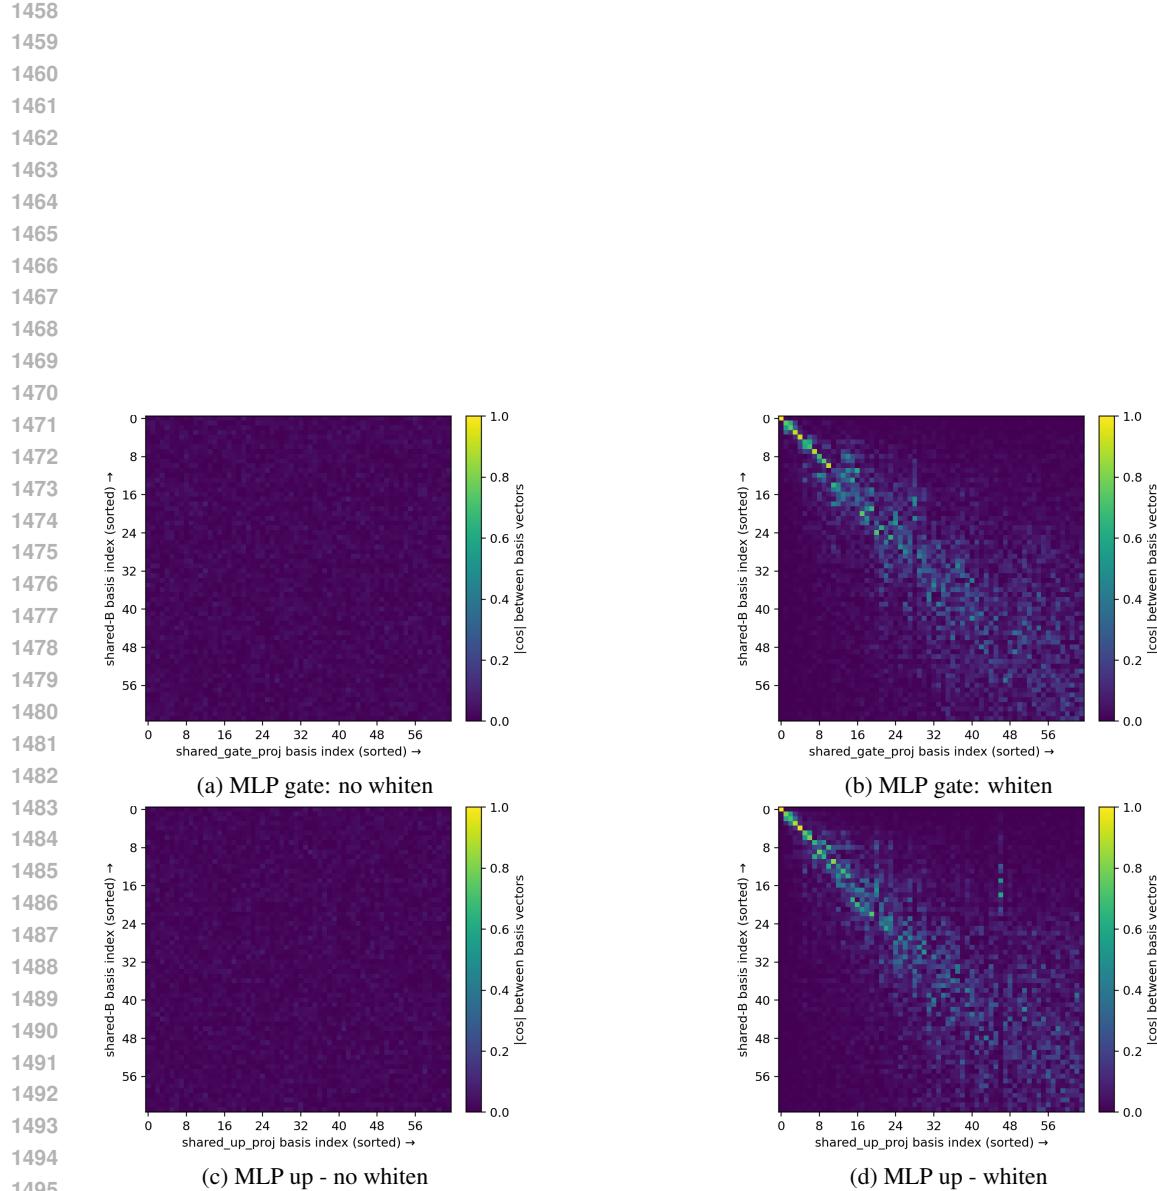


Figure 8: Whitening vs. non-whitening alignment matrices for MLP (gate and up) in Qwen1.5-MoE-A2.7B. The construction is identical to Fig. 7, but applied to the MLP gate and up projections aggregated over all experts. Whitening produces a diagonally dominant alignment, indicating that a shared covariance-aligned basis also captures the principal error directions of the MLP blocks.

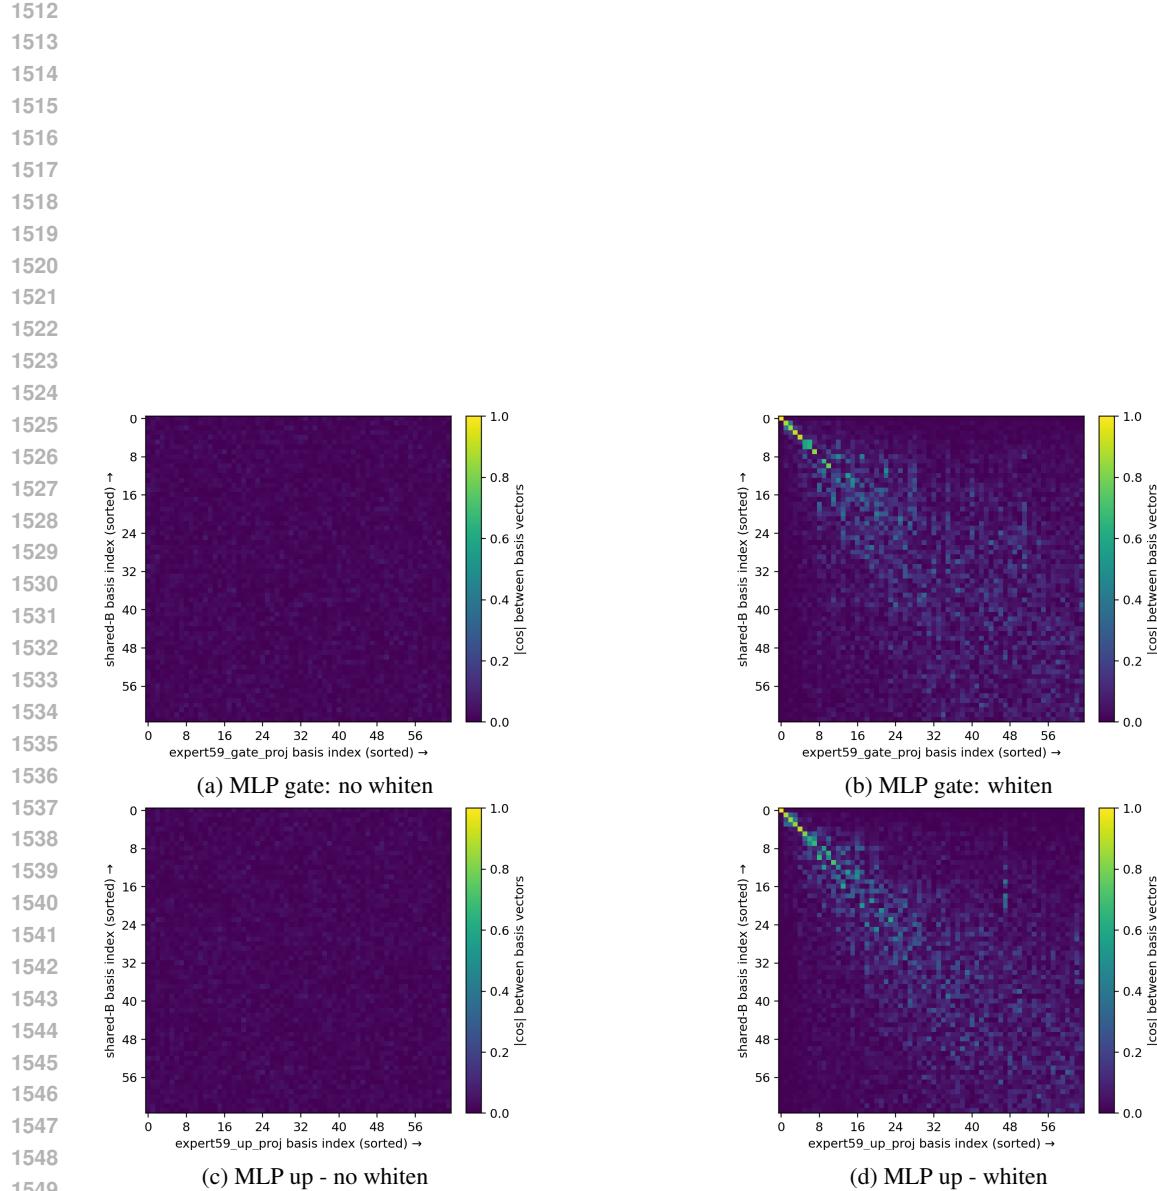


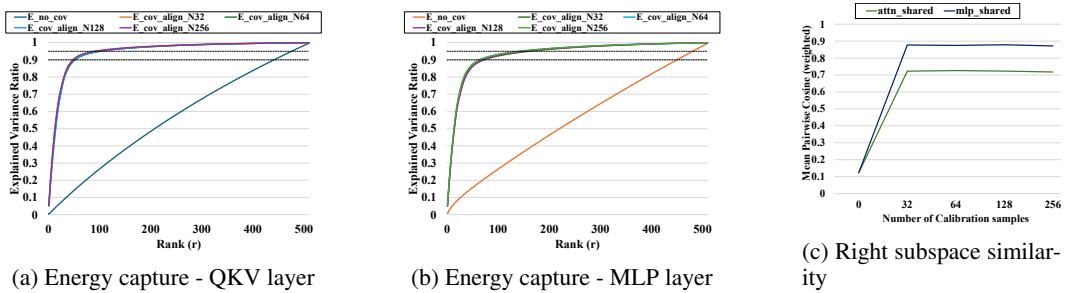
Figure 9: Whitening vs. non-whitening alignment matrices for the MLP (gate and up) of a single expert (Expert 59) in Qwen1.5-MoE-A2.7B. We apply the same construction as in Fig.8, but restrict the stacked error and shared basis B_{shared} to Expert 59 only. The diagonally dominant structure under whitening shows that the covariance-aligned basis remains meaningful even at the per-expert level.

1566 C TTFB & THROUGHPUT AROUND OTHER MODELS
15671568 Table 8: Latency comparison on LLaMA 3 models for Layerwise vs. GlowQ, GlowQ-S.
1569

Models	Setting	TTFB ↓ (ms)	tok/s ↑	Prefill ↓ (ms)	Dec ↓ (ms/tok)	
LLaMA 3	3.2-3B	Layerwise	70.83	17.46	71.07	
		GlowQ	64.92	18.94	66.20	
		GlowQ-S	53.17	21.37	60.69	
	3.1-8B	Layerwise	96.50	14.24	95.72	
		GlowQ	86.44	15.31	90.01	
		GlowQ-S	71.70	18.89	73.50	
Avg. ΔBX (%)		-9.38	+8.00	-6.41	-8.77	
Avg. $\Delta R50$ (%)		-25.32	+27.52	-18.91	-24.13	

1583
1584 **Results on Table 8.** Table 8 mirrors the LLaMA 2 evaluation under an identical runtime and mea-
1585 surement protocol. Two consistent trends emerge: (i) GlowQ reduces all latency components, with
1586 the largest relative gains on per-token decode; and (ii) GlowQ-S further amplifies these benefits.
1587 On LLaMA 3 (3.2-3B, 3.1-8B), GlowQ with BX caching improves serving latency over Layerwise:
1588 TTFB -9.38% , tok/s $+8.00\%$, Prefill -6.41% , and Dec -8.77% on average. GlowQ-S (selec-
1589 tive restore) amplifies these gains: TTFB -25.32% , tok/s $+27.52\%$, Prefill -18.91% , and Dec
1590 -24.13% on average. Improvements are consistent across both model sizes, with the largest reduc-
1591 tions appearing in the per-token Dec phase and end-to-end TTFB, reflecting reduced compute on the
1592 critical path. In practice, BX caching provides drop-in speedups without modifying weights, while
1593 the selective policy (GlowQ-S) offers a simple accuracy-latency knob by reducing the number of
1594 $\mathbf{A}_i\mathbf{R}$ applications (Sec. 3.3).

1595 **Observation on Table 8.** BX caching removes redundant right-projection work by reusing the
1596 shared subspace, so each decode step primarily executes lightweight $\mathbf{A}_i\mathbf{R}$ updates; this directly
1597 lowers Dec and TTFB. The selective-restore strategy further trims the executed paths across decoder
1598 blocks, yielding additional latency drops with a commensurate increase in throughput (tok/s). These
1599 mechanisms explain the near-linear percentage gains in the RSVD-driven core cost: caching reduces
1600 repeated right-side multiplies, while selective restoration shortens the active compute graph along
1601 the decoding trajectory.

1602 D HYPERPARAMETER CHANGE
16031604 D.1 CALIBRATION DIFFERENCE
16051614 Figure 10: Energy Capture and Cosine similarity of Rightspace over number of calibration samples
1615

1616

1617 Fig. 10a, 10b plot energy-capture curves versus rank for different numbers of calibration samples,
1618 while Fig. 10c reports the mean pairwise cosine (weighted) between the shared right
1619 subspaces as the calibration size varies.

1620 Varying the number of calibration samples $N \in \{32, 64, 128, 256\}$ leaves the energy-capture curves
 1621 in Fig. 10a, 10b nearly indistinguishable, especially for practical ranks $r \leq 128$. In Fig. 10c,
 1622 the weighted cosine similarity between the shared right subspace and layer-wise right subspaces is
 1623 already high at $N = 32$ and saturates for $N \geq 64$. These results indicate that a small calibration set
 1624 suffices to recover a stable, data-aligned right subspace, consistent with PCA stability under a clear
 1625 spectral gap (Jolliffe & Cadima, 2016; Horn & Johnson, 2012).

1626 We attribute the observed stability under a relatively small calibration set, e.g., $N = 32$ to the
 1627 following four reasons: (i) *Spectral-gap effect*: The input covariance Σ_x is heavy-tailed, so the
 1628 top directions are separated by a clear eigenvalue gap; the dominant r -dimensional right subspace
 1629 stabilizes quickly with modest N (Jolliffe & Cadima, 2016; Horn & Johnson, 2012). (ii) *Robust*
 1630 *weighted objective*. We optimize a right-weighted criterion,

$$\min_{\mathbf{A}, \mathbf{B}} \|(\mathbf{E}_{\text{cat}} - \mathbf{AB}) \Sigma_x^{1/2}\|_F^2,$$

1634 so small perturbations in the estimate $\widehat{\Sigma}_x$ have limited effect: large-eigenvalue axes dominate and
 1635 lead to the same top r -subspace (see also weighted low-rank formulations (Srebro & Jaakkola,
 1636 2003)). *Numerical regularization*. Shrinkage/normalization of $\widehat{\Sigma}_x$ reduces small-sample noise and
 1637 improves conditioning (Ledoit & Wolf, 2004; Hoerl & Kennard, 1970; Bishop, 2006). *Benefit of*
 1638 *group stacking*. Building the SVD core from vertically stacked errors increases the effective sample
 1639 support along rows, which smooths estimation of the shared right subspace (Paige & Saunders,
 1640 1981; Golub & Loan, 2013).

1641 To conclude, calibration sizes as small as $N \approx 32$ –64 already place the system in a saturated mode
 1642 since energy capture at a fixed r and the similarity between the shared and layer-wise right subspaces
 1643 change only marginally beyond this point. Thus, our covariance-aligned, group-shared \mathbf{B} achieves
 1644 stable performance with low calibration cost.

D.1.1 SHRINK ALPHA DIFFERENCE

1648 Table 9: Perplexity on WikiText-2 while sweeping calibration samples and shrink α (lower is better).

Calibration Samples	Shrink α	LLaMA 3		Qwen 3	
		3.2-3B	8B	3.1-8B	14B
32	0	8.16	6.59	9.89	8.82
	0.02	8.16	6.59	9.86	8.82
	0.05	8.16	6.59	9.88	8.82
64	0	8.15	6.59	9.90	8.81
	0.02	8.15	6.59	9.88	8.78
	0.05	8.16	6.59	9.87	8.79
128	0	8.16	6.59	9.92	8.80
	0.02	8.16	6.58	9.91	8.80
	0.05	8.15	6.58	9.90	8.81
256	0	8.16	6.58	9.93	8.81
	0.02	8.15	6.59	9.92	8.80
	0.05	8.16	6.58	9.92	8.82

1669 We apply a standard covariance shrinkage when forming the input statistic used for covariance-
 1670 aligned subspace estimation. Let $\widehat{\Sigma}_x$ be the sample covariance from N calibration sequences and d
 1671 the input dimension. We construct

$$\widehat{\Sigma}_x^{(\alpha)} = (1 - \alpha) \widehat{\Sigma}_x + \alpha \frac{\text{tr}(\widehat{\Sigma}_x)}{d} \mathbf{I}, \quad \alpha \in [0, 1],$$

i.e., a convex combination of the sample covariance and an isotropic target (scaled identity); small α reduces small-sample noise and improves conditioning without altering the dominant axes learned from data (Ledoit & Wolf, 2004; Bishop, 2006; Anderson, 2003).

Results on Table 9. Across calibration sizes $N \in \{32, 64, 128, 256\}$ and shrink $\alpha \in \{0, 0.02, 0.05\}$, perplexity remains essentially flat for LLaMA 3: for 3.2-3B and 8B, the sweep changes values by +0.01 ppl on average. Qwen 3 shows the same qualitative behavior, with a mild benefit from shrinkage: $\alpha \in [0.02, 0.05]$ yields -0.02 ppl on average for 3.1-8B and -0.01 ppl on average for 14B (relative to $\alpha=0$ at the same N). Aggregating all models, $\alpha=0.02$ improves by -0.01 ppl average, and increasing N beyond 64 produces only marginal changes ($\leq +0.01$ to $+0.02$ ppl on average depending on the family). In short, both the calibration size and a small shrink factor have only second-order effect on WikiText-2 perplexity, consistent with the stability suggested by the energy and cosine-similarity panels (Jolliffe & Cadima, 2016).

Observation on Table 9. The right subspace stabilizes quickly because (i) the input covariance exhibits a pronounced spectral gap, so the dominant r -dimensional space is identified with few samples (Jolliffe & Cadima, 2016; Horn & Johnson, 2012); (ii) the right-weighted objective emphasizes large-variance directions, making the solution insensitive to small perturbations in $\hat{\Sigma}_x$; (iii) mild shrinkage damps small-sample noise (Ledoit & Wolf, 2004; Bishop, 2006); and (iv) stacking modules to form the core increases effective sample support along rows (Paige & Saunders, 1981; Golub & Loan, 2013). Consequently, small calibration sets ($N \approx 32$ –64) already recover a data-aligned shared right subspace, explaining the near-constant perplexity across the sweep and the slight, consistent gains from $\alpha \in [0.02, 0.05]$ on Qwen 3.

D.1.2 MEMORY USAGE

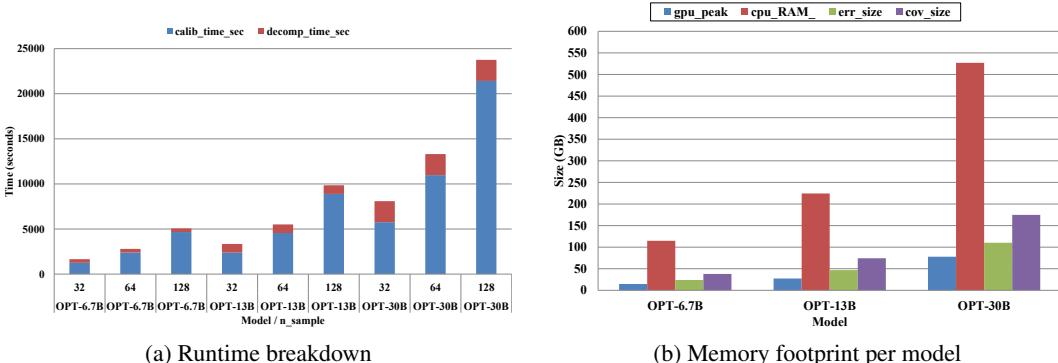


Figure 11: Calibration runtime and memory footprint as a function of model size and the number of calibration samples N . (a) Stacked bars show the runtime breakdown into calibration and decomposition for each (model, N) configuration; calibration dominates the total cost and grows nearly linearly with N , while decomposition time remains almost constant. (b) Memory footprint of the error tensor, covariance tensor, and peak GPU/CPU usage for each OPT model; the error and covariance tensors account for most of the memory and grow steeply with model size.

Results on Fig. 11. We profile calibration on a single A100 80GB GPU for three OPT models (6.7B, 13B, 30B) and calibration sizes $N \in \{32, 64, 128\}$, using SlimPajama-6B as the calibration corpus. Fig. 11a shows that the total wall-clock time is dominated by the calibration pass: for every model, the blue bars (forward passes used to estimate $\hat{\Sigma}_x$ and collect error tensors) account for most of the runtime and grow almost linearly with N , whereas the red bars (randomized GSVD / decomposition) contribute a relatively small and nearly constant overhead. Even for OPT-30B, increasing N from 32 to 128 scales the runtime by roughly the same factor, indicating that the cost is predictable and controlled by the choice of calibration size. Fig. 11b breaks down the memory footprint. Peak GPU memory (blue) grows moderately with model size and remains well below the CPU footprint, since we keep the model and activations on GPU but store error and covariance tensors on host memory. The green and purple bars show that these two tensors dominate the CPU

1728 usage and scale with model size: moving from OPT-6.7B to OPT-30B increases both `err_size`
 1729 and `cov_size` by several times, and the peak CPU RAM closely tracks their sum.
 1730

1731 **Observation on Fig. 11.** Overall, the results indicate that the main cost of our method comes from
 1732 a one-time, embarrassingly parallel calibration phase whose runtime scales linearly with N and
 1733 roughly with model size, while the decomposition step has almost fixed cost. Memory-wise, the
 1734 GPU footprint is modest and does not require larger-than-standard accelerators; the heavy objects
 1735 are the error and covariance tensors on CPU, which can be streamed, sharded, or discarded immedi-
 1736 ately after decomposition. Since Sections D.1–D.1.1 show that small calibration sets ($N \approx 32\text{--}64$)
 1737 already yield stable energy capture, right-subspace similarity, and perplexity, practitioners can op-
 1738 erate in this low- N regime. In practice, this keeps the calibration overhead to a few GPU hours
 1739 even for 30B models and confines the CPU memory requirement to a one-off offline preprocessing
 1740 step, directly addressing concerns about prohibitive calibration time and memory pressure for large
 1741 LLMs.
 1742

1742 D.2 RANK DIFFERENCE

1744 Table 10: Perplexity on WikiText-2 by rank and method, formatted like the calibration-sweep ta-
 1745 ble.(Lower is better.)
 1746

1748 Rank	1749 Method	1750 LLaMA 3		1752 Qwen 3	
		1753 3.2 3B	1754 3.1 8B	1755 8B	1756 14B
1750 8	GlowQ	8.22	6.64	9.95	8.84
	Layerwise	8.22	6.64	9.96	8.84
1753 16	GlowQ	8.20	6.63	9.95	8.81
	Layerwise	8.20	6.62	9.94	8.80
1755 32	GlowQ	8.18	6.61	9.91	8.81
	Layerwise	8.18	6.61	9.93	8.80
1758 64	GlowQ	8.16	6.59	9.87	8.80
	Layerwise	8.15	6.58	9.88	8.80
1760 128	GlowQ	8.12	6.56	9.83	8.79
	Layerwise	8.11	6.55	9.87	8.79

1763 **Results on Table 10.** Sweeping the rank r , GlowQ matches layer-wise restoration in perplexity:
 1764 the gap is +0.02 ppl average across models and ranks (never exceeding +0.04 ppl). Returns dimin-
 1765 ish beyond moderate ranks: from $r=8$ to $r=128$, the change is -0.09 ppl average across families.
 1766 Most of the gain is realized by $r \in \{32, 64\}$; increases beyond this window yield only marginal
 1767 improvements (e.g., $r=64 \rightarrow 128$ shifts by just a few hundredths of a ppl).
 1768

1769 **Observation on Table 10.** The rank-accuracy curve exhibits family-specific shapes: LLaMA
 1770 shows a knee around $r \approx 32\text{--}64$ (initially flat, then a brief drop), whereas Qwen decreases more
 1771 gradually without a sharp elbow. In practice, this suggests using $r=64$ for LLaMA and $r=32$ for
 1772 Qwen as strong defaults; GlowQ remains interchangeable with layer-wise restoration in accuracy at
 1773 fixed r , while retaining the runtime advantages established elsewhere.
 1774

1775 D.3 RANDOMIZED SVD PARAMETERS

1777 D.3.1 PROOF OF QR REDUCTION & RANDOMIZED SVD

1778 **Discussion.** Table 11 shows that Exact SVD on the $d \times d$ core \mathbf{M} takes 42.86 s in total (0.76 s
 1779 per layer on average), whereas Randomized SVD (RSVD) completes in 5.16–5.22 s (.09 s per
 1780 layer). This $\approx 8.2\text{--}8.3 \times$ wall-clock speedup is consistent with the complexity gap between $\mathcal{O}(d^3)$
 1781 and $\mathcal{O}((q+1)d^2(r+p) + d(r+p)^2)$ when $d \gg r+p$ (Golub & Loan, 2013; Halko et al., 2011;

1782 Table 11: SVD runtime (s) and perplexity on LLaMA 3.2-3B (WikiText-2). Exact =
 1783 `torch.linalg.svd` on the GSVD core M ; Randomized = Halko R-SVD with oversampling
 1784 p and power iterations q . *SVD-only* times factorization on M (CUDA-synced), excluding the core
 1785 QR used to build M . Total sums over layers; Layer(mean) averages across layers.

Method	q	p	SVD time (s) \downarrow		Perplexity \downarrow
			Total	Layer(mean)	
Exact SVD	—	—	42.86	0.76	8.16
Randomized SVD	0	0	5.16	0.09	8.22
		4	5.19	0.09	8.21
		8	5.20	0.09	8.21
Randomized SVD	1	16	5.21	0.09	8.21
		24	5.21	0.09	8.21
		0	5.17	0.09	8.17
Randomized SVD	2	4	5.19	0.09	8.16
		8	5.20	0.09	8.16
		16	5.21	0.09	8.16
Randomized SVD	2	24	5.21	0.09	8.16
		0	5.17	0.09	8.16
		4	5.20	0.09	8.16
Randomized SVD	2	8	5.20	0.09	8.15
		16	5.21	0.09	8.16
		24	5.22	0.09	8.16

1812 Martinsson & Tropp, 2020). Concretely, with $d=3072$, $r=64$, and $p \in \{0, \dots, 24\}$, we have
 1813 $(r+p)/d \leq 88/3072 \approx 2.9\%$, so the RSVD term $(q+1) d^2(r+p)$ scales roughly like a few percent
 1814 of d^3 up to constant factors, matching the observed order-of-magnitude reduction in runtime.

1815 **Effect of q and p .** Runtime varies only weakly across $p \in \{0, 4, 8, 16, 24\}$ and $q \in \{0, 1, 2\}$
 1816 (5.16 s \rightarrow 5.22 s). This is expected because the dominant RSVD cost is the matrix-block multiplies
 1817 $M\Omega$, $M^\top(\cdot)$; increasing p from 0 to 24 changes $(r+p)$ from 64 to 88 (only $\sim 38\%$), and the extra
 1818 q passes add a small multiple of the same GEMM cost. The lower-order term $d(r+p)^2$ is negligible
 1819 at this scale. In short, the linear dependence on $(r+p)$ and on $(q+1)$ predicted by

$$\mathcal{O}((q+1) d^2(r+p) + d(r+p)^2)$$

1821 manifests as a near-flat runtime curve because $d \gg r+p$ and GEMM kernels saturate the de-
 1822 vice (Halko et al., 2011; Martinsson & Tropp, 2020).

1823 **Accuracy.** Perplexity stays essentially unchanged: Exact = 8.16; RSVD is 8.22 at $(q=0, p=0)$ and
 1824 improves to 8.15–8.16 for $q \geq 1$ (with small p already sufficient). This aligns with randomized
 1825 SVD theory: even a single power iteration ($q=1$) sharpens separation between leading and trailing
 1826 singular directions and yields a right subspace that is effectively indistinguishable (for a rank- r
 1827 objective) from Exact SVD in downstream perplexity (Halko et al., 2011; Musco & Musco, 2015;
 1828 Martinsson & Tropp, 2020).

1829 The empirical results agree with the stated complexity: Exact SVD on M incurs $\mathcal{O}(d^3)$ time, while
 1830 RSVD retrieves the leading right subspace in $\mathcal{O}((q+1) d^2(r+p))$ time (plus a minor $d(r+p)^2$
 1831 term) (Golub & Loan, 2013; Halko et al., 2011; Martinsson & Tropp, 2020). In practice, $q=1$ with
 1832 a modest p (e.g., $p \in [4, 16]$) delivers near-Exact perplexity at $\sim 8 \times$ lower wall time, and increasing
 1833 p further yields diminishing returns (Halko et al., 2011; Martinsson & Tropp, 2020).

1836 D.3.2 POWER ITERATION & OVERSAMPLING DIFFERENCE
18371838 Table 12: Randomized SVD hyperparameters on WikiText-2, measured on LLaMA-3.2-3B. We
1839 sweep (a) oversampling p (fixed $q = 2$) and (b) power iterations q (fixed $p = 16$) and report perplex-
1840 ity (lower is better).

(a) Oversampling p sweep (fixed $q = 2$).			(b) Power iterations q sweep (fixed $p = 16$).		
Method	p	PPL \downarrow	Method	q	PPL \downarrow
LLaMA 3.2-3B	10	8.16	Llama 3.2-3B	0	8.21
	12	8.16		1	8.16
	16	8.16		2	8.16
	24	8.16			
LLaMA 3.1-8B	10	6.59	Llama 3.1-8B	0	6.63
	12	6.59		1	6.59
	16	6.59		2	6.59
	24	6.58			
Qwen 3-8B	10	9.90	Qwen 3-8B	0	9.97
	12	9.89		1	9.87
	16	9.88		2	9.88
	24	9.89			
Qwen 3-14B	10	8.81	Qwen 3-14B	0	8.79
	12	8.80		1	8.81
	16	8.81		2	8.81
	24	8.81			

1862 Table 12 contrasts oversampling p (with $q=2$ fixed; subtable 12a) and power iterations q (with $p=16$
1863 fixed; subtable 12b). Empirically, increasing p from 10 to 24 leaves PPL essentially unchanged
1864 across models (differences of ≤ 0.01), whereas raising q from 0 to 1 yields small but consistent
1865 gains (most visibly on QWEN3-8B), after which improvements saturate by $q=2$.1866 This pattern aligns with the standard analysis of randomized SVD (RSVD). Oversampling enlarges
1867 the sketch dimension to $\ell = r + p$, which reduces the probability of missing near-rank- r directions
1868 but ultimately does not change the target truncation rank r . Once r already captures the dominant
1869 subspace and the spectral gap is reasonable, the marginal benefit of additional p is small; theory
1870 predicts only a mild reduction of the residual as p grows (e.g., with expected error bounds that
1871 degrade roughly as $\sqrt{r/(p-1)}$), so practical guidance typically recommends $p \approx 5-10$ (Halko
1872 et al., 2011; Martinsson & Tropp, 2020).1873 By contrast, q directly amplifies spectral separation via the power scheme. Forming $\mathbf{Y} =
1874 (\mathbf{A}\mathbf{A}^\top)^q \mathbf{A}\Omega$ effectively reweights singular values as σ_i^{2q+1} , which boosts the ratio between σ_r
1875 and the tail $\{\sigma_{j>r}\}$ and thereby reduces leakage beyond rank r . As a result, the sampled subspace
1876 aligns better with the true top- r subspace, often yielding noticeable gains from $q=0$ to $q=1$, with
1877 diminishing returns thereafter; $q \in \{1, 2\}$ is commonly recommended in practice (Halko et al.,
1878 2011; Ma & Ma, 2024; Martinsson & Tropp, 2020).1879 When r already captures the dominant energy, increasing p beyond a modest buffer offers little
1880 accuracy benefit, while a single power iteration ($q=1$) can materially improve approximation for
1881 matrices with slowly decaying spectra. In our experiments, this theoretical expectation manifests as
1882 flat PPL curves across p and consistent but saturating improvements across q .
18831884 E COMPATIBILITY ACROSS QUANTIZATION DATATYPES
18851886 We apply weight-only quantization to Mistral-7B and evaluate on the WikiText-2 test set across both
1887 integer and floating-point-like datatypes (Table 13). For the integer settings (INT2/INT3/INT4),
1888 we use uniform weight-only quantization with shared scales within each weight group. For the
1889 floating-point-like settings (MXFP4, MXFP6, NVFP4), we adopt microscaling-style formats in

Table 13: WikiText-2 test perplexity (\downarrow) for different datatypes.

Method	FP16	INT			Floating-point-like		
		INT2	INT3	INT4	MXFP4	MXFP6	NVFP4
Quant only	5.32	1015.39	6.16	5.51	8.05	5.36	6.09
Quant + GlowQ		24.23	5.84	5.41	6.10	5.32	5.63

which weights are first normalized within a small block and then encoded using low-bit floating-point codes. Concretely, MXFP4 and MXFP6 follow the block-wise microscaling design of MX+ and the OCP MX specification, using a shared scale per block and 4-bit or 6-bit element codes, respectively Lee et al. (2025); Open Compute Project (2023). NVFP4 follows NVIDIA’s reference design with a microscaled FP4 representation for weights, as described in their low-precision inference guidelines Alvarez et al. (2025). These configurations allow us to test GlowQ not only on conventional integer quantization, but also on recent microscaling-based floating-point-like formats.

Layering GlowQ on top of the quant-only baselines reduces perplexity by -991.16 on INT2, -0.32 on INT3, -0.10 on INT4, -1.95 on MXFP4, -0.04 on MXFP6, and -0.46 on NVFP4, relative to the corresponding quant-only settings. Improvements hold across all six evaluated datatypes, indicating that GlowQ behaves as an orthogonal, plug-and-play low-rank correction rather than a mechanism tied to a single integer format or precision; in particular, it remains compatible with recent floating-point-like microscaling formats while providing consistent accuracy gains.

F LONGBENCH RESULTS

Table 14: The results of Llama-3.1-8B-Instruct on LongBench. The model is evaluated on the 15 English subsets using the official LongBench evaluation protocol, with up to 4K input tokens as context.

Method	NarrativeQA	Qasper	MultiFieldQA	HotpotQA	MuSiQue	2WikiMQA	GovReport	QMSum
Baseline	18.26	12.01	25.96	13.76	7.87	14.95	32.79	21.43
W4A4+GlowQ	14.68	10.80	24.95	14.21	8.39	14.20	32.01	22.01
W4A8+GlowQ	15.56	11.77	23.71	14.39	8.41	14.92	32.00	21.19
W4A16+GlowQ	15.46	11.82	23.68	14.39	7.77	14.53	32.32	21.20
	MultiNews	LCC	RepoBench-P	TriviaQA	SAMSum	TRec	PR	Avg
Baseline	26.95	51.93	47.00	87.76	44.72	70.00	37.50	34.19
W4A4+GlowQ	26.43	47.50	37.51	85.54	42.05	69.00	36.36	32.38
W4A8+GlowQ	27.03	51.50	35.97	84.10	42.62	68.50	37.08	32.58
W4A16+GlowQ	26.86	50.46	35.59	84.30	42.67	68.50	37.17	32.45

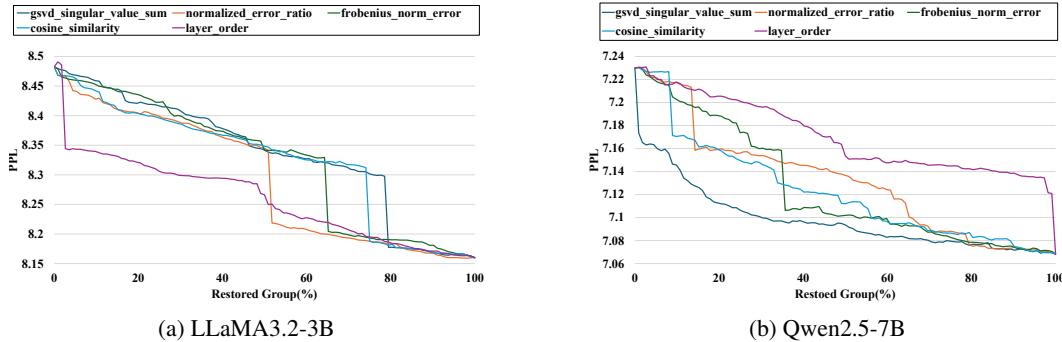
Table 14, 15 shows that across both 4K and 8K context settings on the English LongBench benchmark (Bai et al., 2024), applying W4 weight quantization with GlowQ (W4A4/8/16+GlowQ) leads to only small differences from the original LLaMA-3.1-8B-Instruct on the 15 English LongBench tasks. On most tasks, the scores remain within a few points of the baseline, and the relative difficulty and ranking among tasks are largely preserved. This indicates that, even under aggressive quantization of both weights and activations, the low-rank correction in GlowQ keeps the overall performance stable.

When we extend the context length from 4K to 8K, both the baseline and the GlowQ models improve their average scores by a similar margin. In other words, in scenarios that benefit from longer context, the GlowQ models track the same performance trends as the full-precision model, without a collapse in reasoning ability in the long-context regime. Overall, GlowQ enables 4-bit quantization while preserving LLaMA-3.1-8B-Instruct’s performance not only in standard contexts but also in long-context settings.

1944
1945 Table 15: The results of Llama-3.1-8B-Instruct on LongBench. The model is evaluated on the 15
1946 English subsets using the official LongBench evaluation protocol, with up to 8K input tokens as
1947 context.

Method	NarrativeQA	Qasper	MultiFieldQA	HotpotQA	MuSiQue	2WikiMQA	GovReport	QMSum
Baseline	23.50	13.54	27.87	16.83	10.94	16.44	34.27	22.87
W4A4+GlowQ	23.45	12.20	27.41	15.34	9.21	16.15	33.87	22.78
W4A8+GlowQ	25.38	12.61	25.71	15.37	9.93	15.30	34.07	22.67
W4A16+GlowQ	25.36	12.61	25.62	15.13	9.82	15.20	34.00	22.59
	MultiNews	LCC	RepoBench-P	TriviaQA	SAMSum	TRec	PR	Avg
Baseline	26.87	52.81	48.04	90.77	43.94	71.00	73.13	38.19
W4A4+GlowQ	26.39	48.73	38.83	88.78	42.43	70.50	70.52	36.44
W4A8+GlowQ	27.14	52.06	38.55	88.49	43.60	71.00	72.73	36.97
W4A16+GlowQ	26.96	51.12	38.84	88.67	43.44	71.00	73.50	36.92

G SELECTIVE RESTORATION ACROSS MODEL FAMILY



1973 Figure 12: Perplexity versus fraction of restored groups for different restoration metrics. For each
1974 metric, we sort the groups according to its score (GSVD singular-value sum, normalized error ratio,
1975 Frobenius-norm error, cosine similarity, or simple layer order), progressively restore groups back to
1976 full precision, and record the resulting perplexity.

1977
1978
1979 **Importance metric selection.** The performance of GlowQ-S depends on the policy used to rank
1980 groups for restoration. In Fig. 12, we evaluate five saliency metrics from quantization and pruning
1981 literature. These include: (1) gsvd singular value sum, our g_{ec} score (Eq. 9), which measures the
1982 captured error “energy” ($\|A\|_F^2$) in the low-rank factors and follows the standard practice of using
1983 singular-value energy to summarize PCA components (Jolliffe & Cadima, 2016; Halko et al., 2011);
1984 (2) normalized error ratio, our g_{ner} score (Eq. 10), a widely used PTQ-style proxy based on relative
1985 weight error $\|E_g\|_F/\|W_g\|_F$ (Nagel et al., 2021; Gholami et al., 2021; Krishnamoorthi, 2018); (3)
1986 frobenius norm error, the absolute error $\|E_g\|_F$ (Nagel et al., 2021; Pouransari & Tuzel, 2020; Zhao
1987 et al., 2025); (4) cosine similarity, measuring angular deviation between pre- and post-quantization
1988 weights or activations, which has been shown to be a strong pruning/quantization proxy (Mason-
1989 Williams & Dahlqvist, 2024; Chang et al., 2023); and (5) layer order as a simple baseline. The
1990 results show that gsvd singular value sum and normalized error ratio are consistently the most
1991 effective, yielding the steepest perplexity reduction. However, as noted in Sec. 3.3, no single metric is
1992 universally optimal. Therefore, our final policy (Sec. 4.6) pragmatically evaluates both g_{ec} and g_{ner}
1993 for a given model and selects the one that performs best, providing a robust, data-driven approach.

1994
1995 **Results on Fig. 13.** Across the four panels, LLaMA models exhibit a clear knee: perplexity drops
1996 steeply once a relatively small fraction of groups is restored, then plateaus. In contrast, Qwen and
1997 OPT show a gradual, near-linear descent as the restored fraction increases. The two evaluation
1998 curves in each subplot (ppl_gsvd vs. ppl_NER) track each other closely, differing mainly in the
1999 sharpness of the early descent.

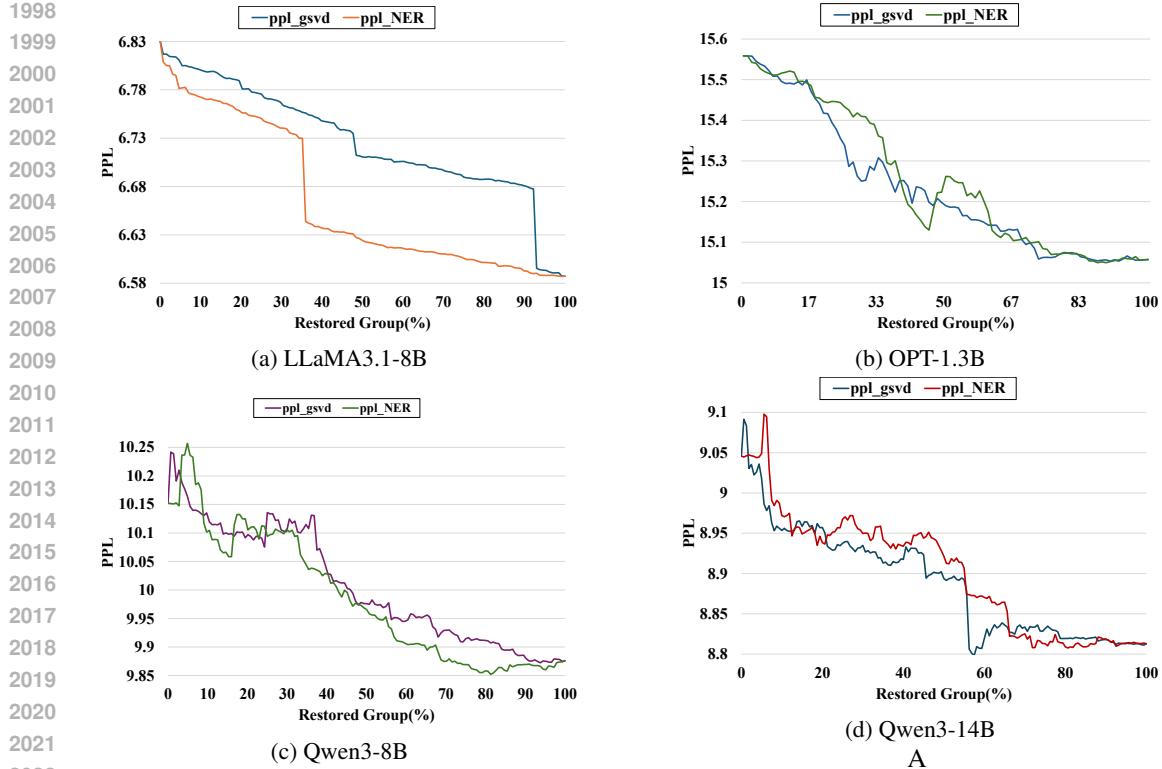


Figure 13: Perplexity as a function of restored group percentage for different model families (LLaMA 3.1-8B, Qwen 3-8B, Qwen 3-14B, OPT-1.3B). We compare GSVD-based restoration (ppl_gsvd) against NER-based restoration (ppl_NER).

These curves suggest selecting the error-recovery metric per model family: outlier/energy ranking with small budgets for knee-shaped profiles, and Hessian-/loss-weighted ranking with broader budgets for diffuse profiles. This family-aware policy aligns with known outlier, anisotropy, and curvature phenomena in modern LLMs.

2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051

2052

2053

2054

2055

2056

2057

2058

2059

2060

2061

2062

2063

2064

2065

2066

2067

2068

Table 16: Zero-shot results on LLaMA 3.2-3B.

Method	Rank	PIQA	ARC-C	ARC-E	HellaS	WinoG	BoolQ	LAMBADA	C4	AVG
		Acc ↑	Acc ↑	Acc ↑	Acc-norm ↑	Acc ↑	Acc ↑	Acc ↑	word PPL ↓	Acc ↑
FP16	-	72.33	39.33	72.33	63.67	70.33	77.00	71.00	10.30	67.14
ZeroQuant-V2		75.33	39.33	73.33	60.67	68.67	74.67	65.67	11.45	65.38
QERA		76.67	38.67	72.33	61.67	68.67	72.33	64.67	11.04	65.48
L2QER	64	75.33	40.00	71.67	64.00	68.33	73.33	68.33	11.04	66.19
GlowQ		77.67	39.67	72.00	64.00	70.33	74.33	70.33	10.98	66.90
GlowQ-S		77.33	39.67	71.67	64.00	69.67	71.67	70.33	11.07	66.33

2064

2065

2066

2067

Table 17: Zero-shot results on LLaMA 3.1-8B.

Method	Rank	PIQA	ARC-C	ARC-E	HellaS	WinoG	BoolQ	LAMBADA	C4	AVG
		Acc ↑	Acc ↑	Acc ↑	Acc-norm ↑	Acc ↑	Acc ↑	Acc ↑	word PPL ↓	Acc ↑
FP16	-	78.67	51.67	80.67	67.67	74.67	80.67	79.00	9.00	73.29
ZeroQuant-V2		78.00	51.33	81.67	68.67	76.00	84.33	74.33	9.87	73.48
QERA		77.00	51.33	80.33	69.00	74.33	82.67	75.33	9.68	72.86
L2QER	64	79.67	49.33	80.67	66.67	74.33	80.33	76.00	9.63	72.43
GlowQ		79.67	51.00	81.33	66.00	74.33	82.00	79.00	9.59	73.33
GlowQ-S		79.00	50.33	81.67	66.33	72.00	82.00	77.00	9.78	72.62

2078

2079

2080

2081

Table 18: Zero-shot results on Qwen 3-8B.

Method	Rank	PIQA	ARC-C	ARC-E	HellaS	WinoG	BoolQ	LAMBADA	C4	AVG
		Acc ↑	Acc ↑	Acc ↑	Acc-norm ↑	Acc ↑	Acc ↑	Acc ↑	word PPL ↓	Acc ↑
FP16	-	77.33	53.00	83.00	63.67	68.67	87.00	67.67	14.52	71.48
ZeroQuant-V2		75.67	52.33	80.33	63.00	71.00	85.33	63.67	15.00	70.19
QERA		76.33	51.33	79.00	62.33	69.67	85.67	64.67	14.78	69.86
L2QER	64	75.67	51.33	79.33	62.67	67.67	85.33	64.67	14.82	69.52
GlowQ		76.67	52.33	80.33	64.67	71.00	86.33	63.67	14.60	70.71
GlowQ-S		76.33	50.67	80.67	63.33	70.67	85.00	65.33	14.77	70.29

2092

2093

2094

2095

Table 19: Zero-shot results on Qwen 3-14B.

Method	Rank	PIQA	ARC-C	ARC-E	HellaS	WinoG	BoolQ	LAMBADA	C4	AVG
		Acc ↑	Acc ↑	Acc ↑	Acc-norm ↑	Acc ↑	Acc ↑	Acc ↑	word PPL ↓	Acc ↑
FP16	-	78.33	59.33	80.33	66.67	75.67	92.00	66.33	13.08	74.10
ZeroQuant-V2		78.33	59.33	78.00	65.67	73.00	92.00	62.00	13.79	72.62
QERA		76.98	57.67	79.33	67.00	74.00	92.00	65.00	13.29	73.14
L2QER	64	78.33	56.33	79.67	66.67	75.33	91.67	64.67	13.80	73.24
GlowQ		77.67	56.67	80.00	68.87	75.67	91.33	66.67	13.26	73.84
GlowQ-S		77.67	57.00	79.33	67.67	74.33	91.33	65.33	13.48	73.24

2105

Table 20: Zero-shot results on Vicuna-7B.

Method	Rank	PIQA	ARC-C	ARC-E	HellaS	WinoG	BoolQ	LAMBADA	C4	AVG
		Acc ↑	Acc ↑	Acc ↑	Acc-norm ↑	Acc ↑	Acc ↑	Acc ↑	word PPL ↓	Acc ↑
FP16	-	76.00	41.33	70.33	66.00	68.00	80.33	72.33	8.70	67.76
ZeroQuant-V2		75.67	43.00	71.00	65.00	65.33	80.00	68.33	9.07	66.90
QERA		76.00	42.67	70.67	67.00	67.33	81.00	69.67	8.91	67.76
L2QER	64	76.33	42.33	70.00	66.00	67.00	80.67	68.00	8.93	67.19
GlowQ		75.67	41.67	70.00	66.67	67.00	80.33	69.67	8.87	67.29
GlowQ-S		76.00	43.67	69.33	66.00	66.67	82.00	70.00	8.99	67.67

Table 21: Zero-shot results on Vicuna-13B.

Method	Rank	PIQA	ARC-C	ARC-E	HellaS	WinoG	BoolQ	LAMBADA	C4	AVG
		Acc ↑	Acc ↑	Acc ↑	Acc-norm ↑	Acc ↑	Acc ↑	Acc ↑	word PPL ↓	Acc ↑
FP16	-	77.33	49.33	73.67	67.33	74.33	86.00	74.33	7.76	71.76
ZeroQuant-V2		77.67	46.67	76.00	66.33	75.00	86.00	74.00	7.86	71.67
QERA		78.00	47.33	76.33	67.33	75.33	86.00	73.00	7.88	71.90
L2QER	64	77.67	48.67	75.67	67.00	74.00	84.67	74.00	7.79	71.67
GlowQ		78.33	47.67	75.67	67.33	74.67	85.33	74.00	7.85	71.86
GlowQ-S		77.67	57.00	79.33	67.67	74.33	91.33	65.33	7.86	73.24

2160 **LLM USAGE DISCLOSURE**

2161

2162 **WRITING POLISH**

2163

2164 After completing the full draft, we used a large language model (LLM) purely to aid proofreading
2165 and light copy-editing. Specifically, the LLM suggested fixes for grammar, spelling, punctuation,
2166 typographical errors, and minor wording for clarity and consistency.

2167

2168 **RETRIEVAL AND DISCOVERY**

2169

2170 We also used an LLM as a literature discovery assistant to broaden our search beyond papers we had
2171 already identified. The LLM helped generate alternative keywords and surface potentially relevant
2172 works.

2173

2174

2175

2176

2177

2178

2179

2180

2181

2182

2183

2184

2185

2186

2187

2188

2189

2190

2191

2192

2193

2194

2195

2196

2197

2198

2199

2200

2201

2202

2203

2204

2205

2206

2207

2208

2209

2210

2211

2212

2213