BEYOND SCALAR RECOVERY: RESIDUAL GEOMETRY IN QUANTIZED VISION-LANGUAGE-ACTION MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Post-quantization recovery in vision-language-action (VLA) models is typically summarized by scalar error or task-level metrics, which measure how much error remains but not how that error is organized across a multi-step action chunk. We characterize this missing structure through teacher-relative residual geometry in a controlled OpenVLA-OFT setting. Across 13 heterogeneous ternary quantization configurations, mean pairwise Spearman similarity increases from 0.037 before adaptation to 0.557 afterward, with a demonstration-bootstrap median gain of $+0.520$ (95% CI $[0.472, 0.568]$). This geometric reorganization is not explained by uniform error shrinkage, and scalar recovery magnitude can diverge substantially from final residual alignment. The shared-alignment trend persists across additional optimization seeds, matched ground-truth-only adaptation, and a distinct 3-bit quantizer. Matched full-precision controls further demonstrate that shared alignment does not require an initial quantization perturbation. Finally, a leave-one-out template formed directly from other configurations' adapted residuals strongly predicts the final residual geometry of a held-out configuration, reaching mean Spearman correlation 0.678 and substantially outperforming both raw geometry and a leave-one-out recovery-field predictor. Shared chunk-position and action-dimension effects explain much of this predictability, while the complete template retains additional predictive structure in reconstruction error.

1 INTRODUCTION

Vision-language-action (VLA) models and related generalist robot policies connect visual-language conditioning with direct robot control (Zitkovich et al., 2023; Ghosh et al., 2024; Black et al., 2025). OpenVLA provides an open VLA backbone (Kim et al., 2025b), while OpenVLA-OFT adapts it for downstream manipulation through action chunking (Kim et al., 2025a). The scale of these models makes low-bit inference attractive for resource-constrained deployment, motivating recent work on VLA quantization (Xu et al., 2026b; Zhang et al., 2026).

Quantization can substantially perturb action predictions, while additional training can recover part of this degradation (Park et al., 2024; 2025b). Such post-quantization adaptation is usually summarized by scalar quantities such as mean action error or task-level performance. These metrics measure *how much* error remains, but not *where* that error is concentrated within a structured action output. For an action-chunking policy with horizon H and D motion dimensions, teacher-relative error naturally forms an $H \times D$ residual map. We use the term *residual geometry* to denote the relative organization of teacher-relative residual magnitudes across action-chunk positions and action dimensions. We therefore ask: *how does this residual geometry change during post-quantization adaptation?*

In a controlled task-adapted OpenVLA-OFT setting, we show that adaptation substantially reorganizes teacher-relative residual geometry. Across 13 heterogeneous ternary quantization configurations, mean pairwise Full-map Spearman similarity increases from 0.037 before adaptation to 0.557 afterward. The shared-alignment increase persists after removing additive main effects and residualizing measured nuisance structure, and is not explained by uniform shrinkage or scalar recovery magnitude. It also persists across heterogeneous layer selections, additional optimization seeds, matched GT-only adaptation, and a distinct 3-bit quantizer; matched full-precision controls

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

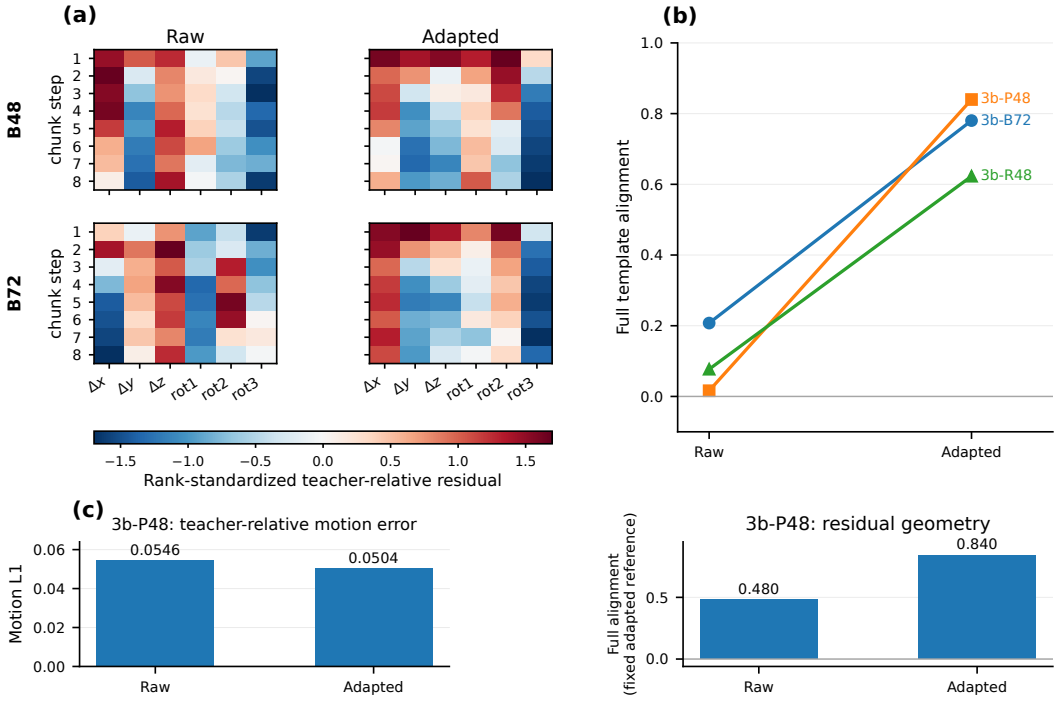


Figure 1: **Post-quantization adaptation reshapes teacher-relative action-chunk residual structure.** **(a)** Rank-standardized 8×6 residual maps for representative ternary configurations B48 and B72 before and after adaptation. **(b)** A distinct 3-bit quantizer also yields substantial increases in alignment to the shared ternary residual structure; across the combined set of 13 ternary and three 3-bit configurations, the demonstration-bootstrap median increase in mean pairwise Full-map alignment is $+0.517$ (95% CI $[0.476, 0.556]$). **(c)** Scalar recovery and geometry can diverge: 3b-P48 reduces teacher-relative motion error by only 7.6%, while Full alignment to a fixed adapted ternary reference increases from 0.480 to 0.840. Panel (b) uses the corresponding raw and adapted ternary reference templates for the two evaluation stages, whereas panel (c) evaluates both states against the same adapted ternary reference.

demonstrate an adaptation-associated shared structure without requiring an initial quantization perturbation.

Under configuration holdout, a leave-one-out adapted-endpoint template strongly predicts the final residual geometry of the held-out configuration ($\rho = 0.678$), compared with the raw shape (-0.042) and recovery-field predictor (0.339).

Our contributions are threefold:

- We show that, across heterogeneous quantization perturbations, post-quantization adaptation substantially reorganizes teacher-relative action-chunk residual geometry while strengthening a shared cross-configuration organization that scalar recovery magnitude and uniform error shrinkage do not capture.
- We establish that this shared-alignment trend persists across heterogeneous ternary configurations, additional optimization seeds, matched GT-only adaptation, independently sampled random target plans, and a distinct 3-bit quantizer. Matched full-precision controls further demonstrate that shared alignment does not require an initial quantization perturbation.
- We establish that a leave-one-out adapted-endpoint template strongly predicts held-out final residual geometry under configuration holdout. Shared chunk-position and action-dimension main effects account for much of this predictability, while the complete mean-final template retains additional predictive structure, most clearly in reconstruction error.

Together, these results establish residual geometry as a complementary view of post-quantization VLA adaptation: scalar metrics summarize how much error remains, while residual geometry reveals how that remaining error is organized.

2 Related Work

Quantization and efficient VLA deployment. Post-training quantization for large language models has developed second-order, scaling- and activation-aware, calibration-based, and rotation-based methods for preserving model quality at low precision (Frantar et al., 2023; Xiao et al., 2023; Lin et al., 2024; Shao et al., 2024; Ashkboos et al., 2024; Liu et al., 2025). Recent work extends these ideas to vision-language-action models through channel- or action-aware sensitivity, calibration, mixed-precision quantization, and output-aware compensation (Xu et al., 2026b; Zhang et al., 2026; Xu et al., 2026a; Akbari et al., 2026; Ranjan & Savakis, 2026), while other studies explore binary, ternary, and other low-bit VLA representations (Wang et al., 2025; Yan et al., 2026; Wang et al., 2026). These works primarily ask how to preserve VLA behavior under compression. We instead use heterogeneous quantization configurations as controlled perturbations and analyze how their structured action-output residuals change during downstream adaptation.

Adaptation and recovery after compression. Quantized robotic policies can recover degraded behavior through post-quantization training, expert supervision, or guidance from a full-precision policy (Park et al., 2024; 2025b). Related low-bit adaptation methods for language models combine quantized weights with low-rank fine-tuning or quantization-aware training (Dettmers et al., 2023; Xu et al., 2024; Li et al., 2024; Chen et al., 2025), and recent VLA studies investigate supervised or reinforcement-learning recovery after compression (Chen et al., 2026; Sun et al., 2026). Recoverability can also depend strongly on the parameter subspace and optimization setting (Xia & Pang, 2026). Our question is complementary to these recovery methods: rather than focusing only on how much performance or scalar error is recovered, we ask how adaptation reorganizes the remaining error across action-chunk coordinates. Matched GT-only and full-precision controls further test whether the shared residual structure persists without teacher distillation or without an initial quantization perturbation.

Structured quantization error and residual geometry. Quantization error is not fully characterized by a single local distortion statistic. Layer-local outliers need not predict network-level degradation, quantization error can propagate across layers, and quantization robustness can depend on training dynamics (Vlassis et al., 2026; Arai & Ichikawa, 2025; Catalan-Tatjer et al., 2026). Prior methods exploit such structure through adaptive rounding, block reconstruction, or training procedures designed to improve subsequent quantization robustness (Nagel et al., 2020; Li et al., 2021; Zhao et al., 2025; Park et al., 2025a).

Residual structure has also been studied in VLA fine-tuning. FiberTune analyzes teacher-structured residuals in intermediate visual representations and uses them as a training signal (Lin et al., 2026). Our analysis instead studies the teacher-relative $H \times D$ residual in the final action-chunk output space. We track how this geometry changes from raw quantization through adaptation, separate geometric alignment from scalar recovery, test matched full-precision adaptation without quantization, and compare a recovery-field construction with a direct shared adapted-endpoint model. Residual geometry therefore serves here as an analysis of post-quantization VLA behavior rather than as a new quantization or adaptation method.

3 Experimental Setting for Residual-Geometry Analysis

This section defines the controlled experimental setting used to characterize how teacher-relative residual geometry changes during post-quantization adaptation. We first specify the model, data, quantization conditions, and adaptation controls, and then define the residual-geometry evaluation used throughout the study.

3.1 Model, Task, Data, and Quantization

We study a task-adapted OpenVLA-OFT policy (Kim et al., 2025a), used as the frozen full-precision teacher throughout the analysis.

The task is simulated cube stacking with a MyCobot 280 M5 in Isaac Lab, using the instruction "stack the blue cube on top of the red cube" and environment `Isaac-Stack-Cube-MyCobot-IK-Rel-Visuomotor-Mimic-v0`. We analyze fixed action outputs rather than closed-loop task success.

The teacher dataset contains 385 demonstrations, split into 347 demonstrations for teacher optimization and 38 for validation. Our residual analysis uses exactly these 38 teacher-validation demonstrations, which are held out from the teacher's gradient-based fine-tuning updates. The frozen teacher is the fixed final checkpoint at 15,000 steps rather than a checkpoint selected by validation performance.

From each of the 38 demonstrations, we take eight fixed trajectory-phase samples, yielding 304 action chunks. A frozen demonstration-level split with seed 20260827 assigns 19 demonstrations to calibration and 19 to test, producing 152 samples per split. All final adaptation runs use the calibration split, while the test split is excluded from final adaptation, within-run validation, and checkpoint selection.

The B48/B72 depth-based and P24/P48 projection-based target sets were frozen from sensitivity rankings computed over the full 304-sample analysis pool before final post-quantization adaptation. In contrast, R12–R72 and S0–S4 use frozen random target plans and do not use these sensitivity scores for target selection. Detailed provenance is given in Appendix A.

The policy outputs action chunks with horizon $H = 8$ and seven action dimensions. Ground-truth supervision uses all seven dimensions, while teacher-relative geometry uses the six continuous motion dimensions, giving an 8×6 residual map.

Our primary study includes 13 heterogeneous ternary weight-quantization configurations spanning depth-based, nested-random, separately sampled random, and projection-based layer selections: B48/B72, R12/R24/R48/R72, S0–S4, and P24/P48. We use these heterogeneous configuration families as controlled perturbations for studying post-quantization residual geometry.

Ternary quantization uses codes $\{-1, 0, +1\}$ with a tensor-wise mean-absolute-weight scale, following a BitNet-style low-bit parameterization (Ma et al., 2024). Gradients through the quantized forward weights are propagated using a straight-through estimator (Bengio et al., 2013).

We additionally evaluate B72, R48, and P48 using a distinct per-output-channel symmetric-uniform 3-bit quantizer. Configuration construction and exact quantizer definitions are given in Appendix B.

3.2 Adaptation and Control Conditions

Each quantized student is adapted with ground-truth imitation and teacher-matching knowledge distillation (Hinton et al., 2015),

$$\mathcal{L}_{\text{adapt}} = \mathcal{L}_{\text{GT}} + \lambda \mathcal{L}_{\text{KD}}, \tag{1}$$

where $\mathcal{L}_{\text{GT}}$ is the mean absolute error over all $H \times 7$ action entries and $\mathcal{L}_{\text{KD}}$ is the mean absolute error to the frozen teacher over the $H \times 6$ continuous-motion entries.

The primary condition uses $\lambda = 0.5$, six adaptation epochs, learning rate 7×10^{-5}, and training seed 20260903. Only the latent full-precision weights underlying the selected quantized matrices are trainable; the action head and all other parameters remain frozen.

We use two matched controls. First, GT-only adaptation sets $\lambda = 0$ for six configurations while retaining the same calibration split, learning rate, six-epoch budget, and optimization seed. Second, matched full-precision controls for B48, R48, and P48 adapt the same parameter subsets from the frozen teacher under the same GT+KD protocol, but without quantization or an STE forward pass. These controls test, respectively, whether the shared residual structure persists without teacher distillation and without an initial quantization perturbation.

Additional optimization and implementation details are reported in Appendix A.

3.3 Residual Geometry and Evaluation

Let $\hat{a}^{(c)}_{n,h,d}$ denote the prediction of configuration c for sample n, chunk step h, and continuous motion dimension d, and let $a^{(T)}_{n,h,d}$ denote the corresponding frozen-teacher output. For final teacher-relative evaluation, we generate the teacher reference from the same frozen checkpoint using the same deterministic inference path used for the evaluated test predictions. All reported teacher-relative scalar metrics and residual-geometry analyses use this path-matched reference.

Over N evaluation samples, we define the teacher-relative residual map

$$E^{(c)}_{h,d} = \frac{1}{N}\sum_{n=1}^{N}\left|\hat{a}^{(c)}_{n,h,d} - a^{(T)}_{n,h,d}\right|, \qquad E^{(c)} \in \mathbb{R}^{8\times 6}. \tag{2}$$

All residuals are computed in the model's normalized action space. Our scalar diagnostic, T6, averages the same absolute teacher-relative error over samples, chunk steps, and the six motion dimensions. Thus, T6 measures overall residual magnitude, whereas $E^{(c)}$ preserves how that residual is distributed across the action chunk.

To compare residual organization independently of global magnitude, we flatten the 48 cells, replace their values by ranks, and standardize the resulting vector. The *Full* representation uses the residual map directly before this rank transformation. We additionally evaluate an *Interaction* representation that removes additive chunk-position and action-dimension main effects, and a *Nuisance* representation that linearly residualizes the map against measured cell-wise teacher variability, ground-truth variability, and teacher–ground-truth discrepancy. Exact definitions are provided in Appendix C.

For the 13 ternary configurations, we summarize shared structure using mean pairwise Spearman similarity and cosine alignment to a leave-one-out template. For leave-one-out alignment, each configuration is compared with a template formed from the other 12 ternary configurations at the same evaluation stage, so the evaluated configuration never contributes to its own reference. As a reference-stage control, we additionally evaluate both the raw and adapted states of each held-out configuration against the same leave-one-out template formed from the adapted residuals of the other 12 configurations (Appendix C).

For the 3-bit experiments, the primary analysis uses stage-specific templates formed from all 13 ternary configurations, while a separate *Full* control holds the adapted template fixed across states (Appendix H); the 3-bit models never contribute to either reference.

Finally, to assess how well final residual geometry is captured by a shared adapted endpoint, we compare four leave-one-out predictors of each configuration's final residual shape: its own raw shape, a shared recovery-field predictor estimated from the other configurations, an additive mean-final predictor retaining only shared chunk-position and action-dimension main effects, and a mean-final predictor formed directly from the other configurations' adapted shapes. Prediction quality is evaluated using Spearman correlation and RMSE. The centered-log construction and full predictor definitions are given in Appendix F.

Here, leave-one-out prediction holds out a configuration rather than test demonstrations: all predictors and targets are evaluated on the same fixed 19-demonstration test set, while the held-out configuration is excluded from construction of the cross-configuration predictors.

Unless stated otherwise, all residual-geometry results use the fixed 152-sample test split. Uncertainty analyses resample at the demonstration level rather than treating action chunks from the same demonstration as independent.

4 Residual Geometry Beyond Scalar Recovery

4.1 Adaptation Strengthens Shared Cross-Configuration Geometry

Immediately after quantization, different layer-selection configurations produce substantially different teacher-relative residual patterns despite sharing the same full-precision teacher, task, and evaluation set.

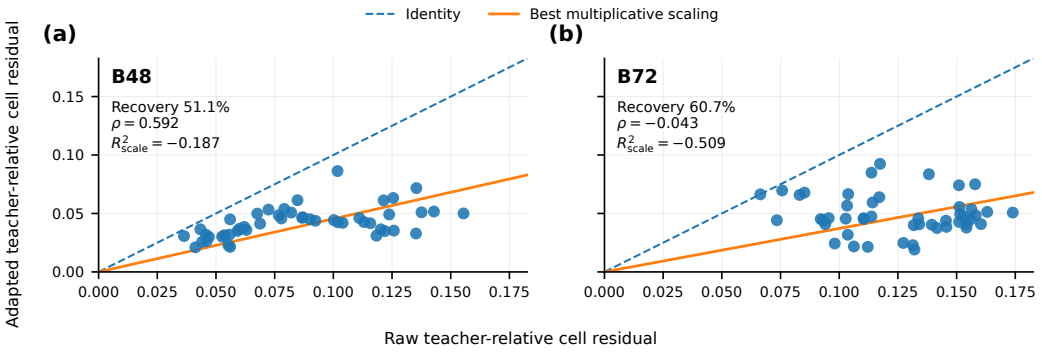


Figure 2: **Adaptation reorganizes residuals rather than uniformly rescaling them.** Each point represents one of the 48 motion cells in the 8×6 teacher-relative residual map for **(a)** B48 and **(b)** B72. Dashed lines show the no-change reference $y = x$, while solid lines show the best zero-intercept multiplicative scaling fit. Despite substantial scalar recovery, uniform rescaling poorly describes the adapted maps, especially for B72.

Across the 13 ternary configurations, mean pairwise Spearman similarity in the *Full* representation increases from 0.037 before adaptation to 0.557 afterward. The corresponding demonstration-bootstrap median gain is $+0.520$, with 95% CI $[0.472, 0.568]$.

Interaction similarity increases from 0.053 to 0.321, with bootstrap median gain $+0.286$ and 95% CI $[0.191, 0.374]$, after removing additive chunk-position and action-dimension main effects. *Nuisance* similarity increases from 0.033 to 0.570, with bootstrap median gain $+0.527$ and 95% CI $[0.466, 0.569]$, after residualizing measured cell-wise nuisance structure.

The positive *Interaction* gain demonstrates that the alignment increase persists after removing additive chunk-position and action-dimension main effects. Separately, the additive-template analysis in Section 5.2 demonstrates that these main effects account for much of the cross-configuration predictive signal.

The shared-alignment increase also persists when the reference stage is held fixed. Evaluating both raw and adapted residuals against the same leave-one-out adapted template gives mean leave-one-out alignment gains of $+0.770$ for *Full*, $+0.333$ for *Interaction*, and $+0.825$ for *Nuisance*. The corresponding demonstration-bootstrap medians are $+0.763$, $+0.356$, and $+0.819$, with 95% CIs $[0.731, 0.793]$, $[0.258, 0.463]$, and $[0.767, 0.860]$, respectively. Full-map alignment increases for all 13 configurations under this fixed-reference control (Appendix C).

The projection-based P48 configuration provides a representative leave-one-out example. Its *Full* alignment to a template constructed from the other 12 ternary configurations increases from -0.339 before adaptation to $+0.662$ afterward. Its *Interaction* alignment similarly increases from -0.331 to $+0.578$. Final alignments remain heterogeneous across configurations. Adaptation therefore strengthens a shared cross-configuration component without making the residual maps identical or eliminating configuration-specific structure.

4.2 Residual Reorganization Is Not Uniform Shrinkage

A simple explanation for lower post-adaptation error would be that each raw residual pattern is approximately preserved while its magnitude is uniformly reduced:

$$E_{\text{adapted}} \approx \alpha E_{\text{raw}}, \qquad 0 < \alpha < 1. \tag{3}$$

Under this model, cells with large raw residuals should remain relatively large after adaptation, and a single multiplicative scale factor should describe the adapted map.

Figure 2 shows that uniform scaling does not describe either representative configuration. B48 reduces teacher-relative motion error by 51.1%, yet its raw-to-adapted cell-rank correlation is only $\rho = 0.592$. The best zero-intercept multiplicative scaling model gives $R^2 = -0.187$. Here, R^2

is the standard centered coefficient of determination, so a negative value means that the fitted zero-intercept model performs worse than predicting the adapted-map mean.

For B72, scalar recovery reaches 60.7%, while the corresponding cell-rank correlation is $\rho = -0.043$ and the multiplicative scaling model gives $R^2 = -0.509$.

Substantial scalar recovery can therefore accompany a major change in which action-chunk coordinates remain most erroneous. The increase in cross-configuration agreement cannot be explained as a mechanical consequence of uniformly shrinking each configuration's original residual pattern.

4.3 Scalar Recovery and Geometric Alignment Can Decouple

Ternary R48 reduces teacher-relative motion error from 0.392 to 0.082 (approximately 79% scalar recovery), yet its final *Full* leave-one-out alignment is 0.483. Strong scalar recovery therefore does not imply correspondingly strong alignment with the shared residual structure.

The matched full-precision controls provide a complementary scale reference: after adaptation, their teacher-relative T6 ranges from 0.048 to 0.066, overlapping the final T6 of many quantized models. R12 and 3b-P48 already begin near this adaptation-associated residual scale. R12 changes only from 0.041 to 0.042 while reaching final *Full* leave-one-out alignment 0.683; 3b-P48 changes from 0.0546 to 0.0504 while its *Full* alignment to a fixed adapted ternary reference increases from 0.480 to 0.840.

These cases show why scalar recovery and residual geometry need not move together. Scalar recovery is inherently relative to the magnitude of the initial perturbation, whereas residual geometry captures how the remaining error is organized across the action chunk. The two measurements therefore remain complementary, while the matched full-precision controls identify adaptation itself as an important source of the final residual scale.

5 Dissecting the Shared Residual Structure

5.1 Matched Full-Precision Controls Show Shared Structure Without Quantization

Because all quantized configurations undergo the same downstream adaptation, the shared geometry can reflect adaptation itself rather than a quantization-specific recovery mechanism. We therefore construct matched full-precision controls for B48, R48, and P48. Each starts from the same full-precision teacher and adapts the same parameter subset as its quantized counterpart under the same GT+KD objective, learning rate, six-epoch budget, and optimization seed, but without quantization or an STE forward pass.

Under the path-matched evaluation reference, the unadapted full-precision controls coincide with the frozen teacher and therefore have zero teacher-relative T6. We therefore evaluate them as adapted-endpoint controls rather than as raw-to-adapted recovery experiments. For each control, the adapted ternary reference excludes the corresponding quantized configuration.

Under this counterpart-excluded reference, *Full*-map alignment is 0.811 for FP-B48, 0.421 for FP-R48, and 0.843 for FP-P48, with all three pointwise demonstration-level 95% bootstrap intervals above zero (Figure 3c). Across *Full*, *Interaction*, and *Nuisance* representations, all nine configuration–representation pointwise 95% bootstrap intervals remain above zero, although *Nuisance* alignment for FP-R48 is substantially weaker (Appendix E).

Matched unquantized adaptation therefore develops substantial alignment with the shared adapted geometry. These controls demonstrate that the shared alignment does not require an initial quantization perturbation. Quantization sets the initial perturbation, while configuration-specific residual variation persists after adaptation.

5.2 A Shared Adapted Endpoint Captures Much of the Final Geometry

We compare three leave-one-out cross-configuration predictors against each configuration's own raw shape: a shared *recovery field* estimated from the other 12 configurations' raw-to-adapted changes,

Table 1: Leave-one-out prediction of final residual geometry across 13 ternary configurations. The held-out configuration is excluded from predictor construction; all predictors and targets use the fixed 19-demonstration test set. Higher Spearman correlation and lower centered-log RMSE over the 48 cells are better.

Predictor	Mean Spearman $\rho \uparrow$	Mean RMSE $\downarrow$
Raw shape	-0.042	0.557
Recovery field	0.339	0.496
Additive mean-final	0.658	0.225
Mean-final	$\mathbf{0.678}$	$\mathbf{0.218}$

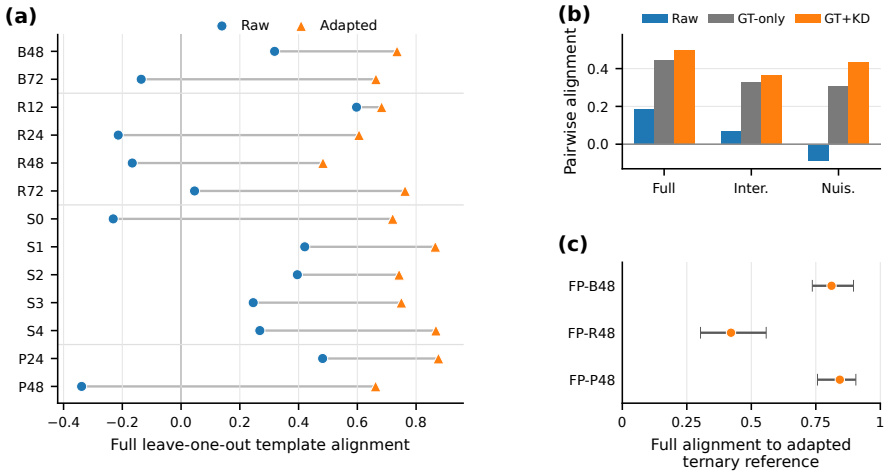


Figure 3: **Shared residual alignment across ternary configurations and controls.** **(a)** Full leave-one-out alignment for 13 ternary configurations. **(b)** GT-only and GT+KD adaptation. **(c)** Full-precision controls against counterpart-excluded adapted ternary references; error bars show demonstration-bootstrap 95% CIs.

an *additive mean-final* template retaining only chunk-position and action-dimension main effects, and the complete *mean-final* template formed directly from the other configurations' adapted shapes.

Table 1 shows that the recovery-field predictor improves over the raw baseline, increasing mean Spearman correlation from -0.042 to 0.339 and reducing mean RMSE from 0.557 to 0.496. The complete mean-final template is substantially stronger than both the raw and recovery-field baselines, reaching 0.678 correlation and 0.218 RMSE.

The additive mean-final baseline already reaches 0.658 correlation and 0.225 RMSE, demonstrating that shared chunk-position and action-dimension effects account for much of this predictive signal. Within this additive structure, the centered-log chunk-position effect is largest at step 1 ($+0.37$) and decreases to -0.11 by step 8, while rot3 has the smallest action-dimension effect (-0.49), with Δy ($+0.19$) and rot2 ($+0.17$) among the largest.

The complete template retains additional predictive structure: its mean correlation is higher by 0.019, although the paired demonstration-bootstrap interval includes zero (median $+0.019$, 95% CI $[-0.001,\ 0.036]$), while its RMSE advantage remains positive under resampling (median additive-minus-mean-final $+0.008$, 95% CI $[0.005,\ 0.011]$).

5.3 Robustness Across Objectives, Seeds, and Quantizers

The shared-alignment trend also persists across the five separately sampled Random-24 configurations S0–S4 (*Full*: $0.313 \rightarrow 0.664$).

In the optimization-seed analysis, alignment increases in 26 of 27 configuration–representation–seed comparisons across B48, R48, and P48 against fixed counterpart-excluded adapted references over the three tested seeds. The sole negative change is P48 Nuisance at seed 20260905 (-0.014). Its final Nuisance alignment is also seed-sensitive, taking values of 0.702, 0.082, and -0.227 for seeds 20260903, 20260904, and 20260905, respectively (Appendix C.1, C.2).

Within the matched six-configuration control, the shared-alignment increase is reproduced without teacher distillation. *Full* mean pairwise similarity increases from 0.186 before adaptation to 0.445 under GT-only training and 0.497 under GT+KD (Appendix G).

The distinct 3-bit quantizer reproduces the shared-alignment trend: *Full* alignment increases for B72, P48, and R48, and eight of nine configuration–representation point estimates increase under the stage-specific comparison, with 3b-R48 *Interaction* as the sole exception. Across the combined 13 ternary and three 3-bit configurations, the demonstration-bootstrap median increase in mean pairwise *Full*-map alignment is $+0.517$ (95% CI $[0.476, 0.556]$).

Together, these controls establish robustness across the tested quantized settings and, with the full-precision and adapted-endpoint analyses, a strong adaptation-associated shared structure alongside persistent configuration-specific variation.

6 Discussion

Scalar recovery measures how much teacher-relative error remains, whereas residual geometry captures where it remains across the action chunk. Their divergence demonstrates substantial reorganization, while matched full-precision controls support an adaptation-associated rather than quantization-specific interpretation.

Common task supervision, action-coordinate coupling, the shared frozen action head, and repeated optimization are plausible but not isolated sources of this structure. Because configuration also determines the parameter subset updated during adaptation, persistent configuration-specific variation cannot be attributed uniquely to quantization location. Uniform scaling fails to explain the reorganization (Section 4.2), which persists under a fixed adapted reference (Appendix C).

The shared adapted endpoint is a major source of cross-configuration predictability: its leave-one-out template substantially outperforms raw and recovery-field predictors. Additive chunk-position and action-dimension effects account for much of this signal, while the complete template retains additional structure, most clearly in its positive RMSE advantage under resampling.

Our analysis covers one task-adapted OpenVLA-OFT policy, one simulated task, selected quantization conditions, and fixed-input outputs; cross-task, cross-policy, and closed-loop generalization remain open.

7 Conclusion

We characterize post-quantization adaptation in a task-adapted OpenVLA-OFT policy through teacher-relative action-chunk residual geometry. Adaptation substantially reorganizes residual structure and strengthens shared cross-configuration organization beyond scalar recovery or uniform shrinkage.

This trend persists across heterogeneous ternary configurations, additional optimization seeds, matched GT-only adaptation, independently sampled random target plans, and a distinct 3-bit quantizer. Matched full-precision controls demonstrate that shared alignment does not require an initial quantization perturbation, while configuration-specific variation persists.

Under configuration holdout, leave-one-out adapted-endpoint templates strongly predict final residual geometry. Shared chunk-position and action-dimension effects explain much of this signal, while the complete template retains additional predictive structure in reconstruction error.

AI USE STATEMENT

Generative AI tools were used to provide feedback on research methodology and experimental design, assist with software development and debugging, support literature search and interpretation of experimental results, and assist with drafting and editing portions of the manuscript. We did not use generative AI tools to generate synthetic datasets or to formulate or prove mathematical claims.

All experiments were executed by the authors. AI-assisted code, analyses, and numerical claims were reviewed and verified against the corresponding experimental outputs before inclusion in the manuscript. The authors made the final decisions regarding research design, analysis, interpretation, and presentation, and take responsibility for the content of this work.

REPRODUCIBILITY STATEMENT

Section 3 specifies the model, simulated manipulation task, dataset and demonstration split, quantization conditions, adaptation protocol, and residual-geometry evaluation used in this study.

Appendix A provides detailed implementation and data-provenance information, including trajectory-phase sampling, action preprocessing and normalization, optimization settings, trainable parameter subsets, and the path-matched teacher-reference procedure used for final evaluation.

Appendix B gives the exact ternary and 3-bit quantizer definitions and configuration construction.

Appendices C–H report the complete configuration-level analyses, robustness experiments, matched full-precision and objective controls, adapted-endpoint analysis, and cross-quantizer results.

Appendix I specifies the demonstration-level bootstrap procedures, sampling units, replicate counts, random seeds, and confidence-interval construction.

The anonymous supplementary material provides the configuration files, selection plans, data-split manifests, analysis and statistical-audit scripts, and canonical prediction outputs needed to reproduce the reported analyses from frozen prediction outputs.

Together, these sections and the supplementary material document the experimental and analytical procedures used to obtain the reported results.

REFERENCES

Arash Akbari, Arman Akbari, Masih Eskandar, Qitao Tan, Yixiao Chen, Jingwu Luo, Bertha Pangaribuan, Liyun Zhang, Jennifer Dy, Geng Yuan, Xue Lin, Gaowen Liu, Stratis Ioannidis, and Yanzhi Wang. ActQuant: Sub-4-bit action-guided quantization for vision-language-action models. *arXiv preprint arXiv:2605.24011*, 2026. URL https://arxiv.org/abs/2605.24011.

Yamato Arai and Yuma Ichikawa. Quantization error propagation: Revisiting layer-wise post-training quantization. In *Advances in Neural Information Processing Systems*, volume 38, 2025. doi: 10.52202/085713-5080.

Saleh Ashkboos, Amirkeivan Mohtashami, Maximilian L. Croci, Bo Li, Pashmina Cameron, Martin Jaggi, Dan Alistarh, Torsten Hoefler, and James Hensman. QuaRot: Outlier-free 4-bit inference in rotated LLMs. In *Advances in Neural Information Processing Systems*, volume 37, 2024. doi: 10.52202/079017-3180.

Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013. URL https://arxiv.org/abs/1308.3432.

Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Robert Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, Laura Smith, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. π_0: A vision-language-action flow model for general robot control. In *Proceedings of Robotics: Science and Systems*, Los Angeles, CA, USA, June 2025. doi: 10.15607/RSS.2025.XXI.010.

Albert Catalan-Tatjer, Niccolò Ajroldi, and Jonas Geiping. Training dynamics impact post-training quantization robustness. In *International Conference on Learning Representations*, 2026.

Mengzhao Chen, Wenqi Shao, Peng Xu, Jiahao Wang, Peng Gao, Kaipeng Zhang, and Ping Luo. EfficientQAT: Efficient quantization-aware training for large language models. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 10081–10100, Vienna, Austria, 2025. Association for Computational Linguistics. doi: 10.18653/v1/2025.acl-long.498. URL https://aclanthology.org/2025.acl-long.498/.

Yuxuan Chen, Yixin Han, Yize Huang, and Xiao Li. RLRC: Reinforcement learning-based recovery for compressed vision-language-action models. *IEEE Robotics and Automation Letters*, 11(7): 8864–8871, 2026.

Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. QLoRA: Efficient finetuning of quantized LLMs. In *Advances in Neural Information Processing Systems*, volume 36, pp. 10088–10115, 2023. doi: 10.52202/075280-0441.

Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. OPTQ: Accurate post-training quantization for generative pre-trained transformers. In *International Conference on Learning Representations*, 2023.

Dibya Ghosh, Homer Rich Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Tobias Kreiman, Charles Xu, Jianlan Luo, You Liang Tan, Lawrence Yunliang Chen, Quan Vuong, Ted Xiao, Pannag R. Sanketi, Dorsa Sadigh, Chelsea Finn, and Sergey Levine. Octo: An open-source generalist robot policy. In *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, July 2024. doi: 10.15607/RSS.2024.XX.090.

Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. URL https://arxiv.org/abs/1503.02531.

Moo Jin Kim, Chelsea Finn, and Percy Liang. Fine-tuning vision-language-action models: Optimizing speed and success. In *Proceedings of Robotics: Science and Systems*, Los Angeles, CA, USA, June 2025a. doi: 10.15607/RSS.2025.XXI.017.

Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan P. Foster, Pannag R. Sanketi, Quan Vuong, Thomas Kollar, Benjamin Burchfiel, Russ Tedrake, Dorsa Sadigh, Sergey Levine, Percy Liang, and Chelsea Finn. OpenVLA: An open-source vision-language-action model. In *Proceedings of The 8th Conference on Robot Learning*, volume 270 of *Proceedings of Machine Learning Research*, pp. 2679–2713. PMLR, 2025b. URL https://proceedings.mlr.press/v270/kim25c.html.

Yixiao Li, Yifan Yu, Chen Liang, Nikos Karampatziakis, Pengcheng He, Weizhu Chen, and Tuo Zhao. LoftQ: LoRA-fine-tuning-aware quantization for large language models. In *International Conference on Learning Representations*, 2024.

Yuhang Li, Ruihao Gong, Xu Tan, Yang Yang, Peng Hu, Qi Zhang, Fengwei Yu, Wei Wang, and Shi Gu. BRECQ: Pushing the limit of post-training quantization by block reconstruction. In *International Conference on Learning Representations*, 2021.

Haihao Lin, Xiangsheng Huang, Xiao Yang, Weibang Zhou, Yiqi Zhang, Bo Yang, Simin Zeng, Jiawei Yang, Zhengyang Wang, and Jiahui Du. FiberTune: Preserving action-fiber visual residuals in vision-language-action fine-tuning. *arXiv preprint arXiv:2606.08653*, 2026. URL https://arxiv.org/abs/2606.08653. Accepted at CoRL 2026.

Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Wei-Ming Chen, Wei-Chen Wang, Guangxuan Xiao, Xingyu Dang, Chuang Gan, and Song Han. AWQ: Activation-aware weight quantization for on-device LLM compression and acceleration. In *Proceedings of Machine Learning and Systems*, volume 6, 2024.

Zechun Liu, Changsheng Zhao, Igor Fedorov, Bilge Soran, Dhruv Choudhary, Raghuraman Krishnamoorthi, Vikas Chandra, Yuandong Tian, and Tijmen Blankevoort. SpinQuant: LLM quantization with learned rotations. In *International Conference on Learning Representations*, 2025.

Shuming Ma, Hongyu Wang, Lingxiao Ma, Lei Wang, Wenhui Wang, Shaohan Huang, Li Dong, Ruiping Wang, Jilong Xue, and Furu Wei. The era of 1-bit LLMs: All large language models are in 1.58 bits. *arXiv preprint arXiv:2402.17764*, 2024. URL https://arxiv.org/abs/2402.17764.

Markus Nagel, Rana Ali Amjad, Mart Van Baalen, Christos Louizos, and Tijmen Blankevoort. Up or down? Adaptive rounding for post-training quantization. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 7197–7206. PMLR, 2020. URL https://proceedings.mlr.press/v119/nagel20a.html.

Jungwoo Park, Taewhoo Lee, Chanwoong Yoon, Hyeon Hwang, and Jaewoo Kang. Outlier-safe pre-training for robust 4-bit quantization of large language models. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 12582–12600, Vienna, Austria, 2025a. Association for Computational Linguistics. doi: 10.18653/v1/2025.acl-long.618. URL https://aclanthology.org/2025.acl-long.618/.

Seongmin Park, Hyungmin Kim, Wonseok Jeon, Juyoung Yang, Byeongwook Jeon, Yoonseon Oh, and Jungwook Choi. Quantization-aware imitation-learning for resource-efficient robotic control. *arXiv preprint arXiv:2412.01034*, 2024. URL https://arxiv.org/abs/2412.01034.

Seongmin Park, Hyungmin Kim, Sangwoo Kim, Wonseok Jeon, Juyoung Yang, Byeongwook Jeon, Yoonseon Oh, and Jungwook Choi. Saliency-aware quantized imitation learning for efficient robotic control. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 13140–13150, 2025b. doi: 10.1109/ICCV51701.2025.01221.

Navin Ranjan and Andreas Savakis. Mix-QVLA: Task-evidence-aware mixed-precision quantization of vision-language-action models. *arXiv preprint arXiv:2606.19565*, 2026. URL https://arxiv.org/abs/2606.19565.

Wenqi Shao, Mengzhao Chen, Zhaoyang Zhang, Peng Xu, Lirui Zhao, Zhiqian Li, Kaipeng Zhang, Peng Gao, Yu Qiao, and Ping Luo. OmniQuant: Omnidirectionally calibrated quantization for large language models. In *International Conference on Learning Representations*, 2024.

Guoheng Sun, Kaixi Feng, Shwai He, Xiaochuan Gong, Yexiao He, Ziyao Wang, Zheyu Shen, Wanghao Ye, Ramana Rao Kompella, Gaowen Liu, and Ang Li. Drop-then-recovery: How redundant are vision-language-action models? *arXiv preprint arXiv:2606.27755*, 2026. URL https://arxiv.org/abs/2606.27755.

Georgios Vlassis, Saleh Ashkboos, Alexandra Volkova, Torsten Hoefler, and Dan Alistarh. Beyond outliers: A study of optimizers under quantization. In *International Conference on Learning Representations*, 2026.

Hongyu Wang, Chuyan Xiong, Ruiping Wang, and Xilin Chen. BitVLA: 1-bit vision-language-action models for robotics manipulation. *arXiv preprint arXiv:2506.07530*, 2025. URL https://arxiv.org/abs/2506.07530.

Xinyu Wang, Mingze Li, Sicheng Lyu, Dongxiu Liu, Kaicheng Yang, Ziyu Zhao, Yufei Cui, Xiao-Wen Chang, and Peng Lu. HoloQ-VLA: Uniform W4A4 quantization of vision-language-action models. *arXiv preprint arXiv:2605.28803*, 2026. URL https://arxiv.org/abs/2605.28803.

Peng Xia and Junbiao Pang. SandwichQuant: Which parameters matter before and after quantization? *arXiv preprint arXiv:2608.24173*, 2026. URL https://arxiv.org/abs/2608.24173.

Guangxuan Xiao, Ji Lin, Mickael Seznec, Hao Wu, Julien Demouth, and Song Han. SmoothQuant: Accurate and efficient post-training quantization for large language models. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 38087–38099. PMLR, 2023. URL https://proceedings.mlr.press/v202/xiao23c.html.

Siyuan Xu, Tianshi Wang, Fengling Li, Lei Zhu, and Heng Tao Shen. DA-PTQ: Drift-aware post-training quantization for efficient vision-language-action models. *arXiv preprint arXiv:2604.11572*, 2026a. URL https://arxiv.org/abs/2604.11572.

Yuhao Xu, Yantai Yang, Zhenyang Fan, Yufan Liu, Yuming Li, Bing Li, and Zhipeng Zhang. QVLA: Not all channels are equal in vision-language-action model's quantization. In *International Conference on Learning Representations*, 2026b. URL https://arxiv.org/abs/2602.03782.

Yuhui Xu, Lingxi Xie, Xiaotao Gu, Xin Chen, Heng Chang, Hengheng Zhang, Zhengsu Chen, Xiaopeng Zhang, and Qi Tian. QA-LoRA: Quantization-aware low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2024.

Xin Yan, Zhenglin Wan, Feiyang Ye, Xingrui Yu, Hangyu Du, Yang You, and Ivor Tsang. HB-VLA: Pushing 1-bit post-training quantization for vision-language-action models. *arXiv preprint arXiv:2602.13710*, 2026. URL https://arxiv.org/abs/2602.13710.

Jingxuan Zhang, Yunta Hsieh, Zhongwei Wan, Haokun Lin, Xin Wang, Ziqi Wang, Yingtie Lei, and Mi Zhang. QuantVLA: Scale-calibrated post-training quantization for vision-language-action models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 39539–39549, 2026.

Weibo Zhao, Yubin Shi, Xinyu Lyu, Wanchen Sui, Shen Li, and Yong Li. ASER: Activation smoothing and error reconstruction for large language model quantization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 22822–22830, 2025. doi: 10.1609/aaai.v39i21.34443.

Brianna Zitkovich, Tianhe Yu, Sichun Xu, Peng Xu, Ted Xiao, Fei Xia, Jialin Wu, Paul Wohlhart, Stefan Welker, Ayzaan Wahid, Quan Vuong, Vincent Vanhoucke, Huong Tran, Radu Soricut, Anikait Singh, Jaspiar Singh, Pierre Sermanet, Pannag R. Sanketi, Grecia Salazar, Michael S. Ryoo, Krista Reymann, Kanishka Rao, Karl Pertsch, Igor Mordatch, Henryk Michalewski, Yao Lu, Sergey Levine, Lisa Lee, Tsang-Wei Edward Lee, Isabel Leal, Yuheng Kuang, Dmitry Kalashnikov, Ryan Julian, Nikhil J. Joshi, Alex Irpan, Brian Ichter, Jasmine Hsu, Alexander Herzog, Karol Hausman, Keerthana Gopalakrishnan, Chuyuan Fu, Pete Florence, Chelsea Finn, Kumar Avinava Dubey, Danny Driess, Tianli Ding, Krzysztof Marcin Choromanski, Xi Chen, Yevgen Chebotar, Justice Carbajal, Noah Brown, Anthony Brohan, Montserrat Gonzalez Arenas, and Kehang Han. RT-2: Vision-language-action models transfer web knowledge to robotic control. In *Proceedings of The 7th Conference on Robot Learning*, volume 229 of *Proceedings of Machine Learning Research*, pp. 2165–2183. PMLR, 2023. URL https://proceedings.mlr.press/v229/zitkovich23a.html.

A Experimental and Reproducibility Details

Task and teacher policy. All experiments use the same task-adapted OpenVLA-OFT policy, retained as the frozen full-precision teacher throughout the residual analysis. The policy is adapted to a simulated MyCobot 280 M5 cube-stacking task in Isaac Lab with the language instruction "stack the blue cube on top of the red cube." The corresponding environment is `Isaac-Stack-Cube-MyCobot-IK-Rel-Visuomotor-Mimic-v0`. The experiments in this paper analyze fixed action outputs rather than closed-loop task success.

Teacher-data provenance. The task-adaptation dataset contains 385 demonstrations. Using teacher split seed 42, 347 demonstrations are assigned to gradient-based teacher fine-tuning and 38 to teacher validation. The 38 demonstrations used in the residual analysis are exactly these teacher-validation demonstrations and therefore are not used for the teacher's gradient-based optimization. The teacher used throughout the present study is the fixed final 15,000-step checkpoint rather than a checkpoint selected according to performance on the 38 validation demonstrations.

MimicGen dataset provenance. The 385-demonstration corpus is an exact subset of a larger 3,000-demonstration MimicGen-generated pool. In the generation pipeline, source demonstrations are reselected separately for each of four subtasks using object-pose nearest-neighbor matching, with uniform random selection among the three nearest candidates ($k = 3$). Generated demonstrations are executed as new simulator rollouts with action noise and are not recursively added back to the source pool. The exported generated HDF5 does not retain the selected source-demonstration indices for each subtask, so exact source-segment ancestry of the retained demonstrations cannot be reconstructed from the final dataset.

Output space and action normalization. The policy predicts action chunks with horizon $H = 8$ and seven action dimensions. Ground-truth supervision is applied to all seven action dimensions, whereas teacher matching and residual-geometry analysis use the six continuous motion dimensions. The resulting teacher-relative residual geometry therefore contains $8 \times 6 = 48$ cells.

All reported teacher-relative residuals are computed in the model's normalized action space. For normalization-enabled action dimensions, dataset preprocessing uses the 1st and 99th percentile statistics, q_{01} and q_{99}:

$$a_{\text{norm}} = \text{clip}\left(2\frac{a - q_{01}}{q_{99} - q_{01} + 10^{-8}} - 1, -1, 1\right). \tag{4}$$

Residual-geometry analysis uses only the six continuous motion dimensions.

Frozen calibration and test split. From each of the 38 analysis demonstrations, we use eight fixed trajectory-phase targets uniformly spaced from 0.05 to 0.95. For a demonstration containing T frames, a phase target p is mapped to frame index

$$\text{round}(p(T-1)). \tag{5}$$

This procedure yields eight distinct sampled frames for every demonstration, producing 304 action-chunk samples in total. The sample identities are fixed by a frozen manifest.

Using downstream split seed 20260827, the 38 demonstrations are divided into 19 calibration demonstrations and 19 test demonstrations, yielding 152 samples in each split. The calibration and test splits are demonstration-disjoint. All final post-quantization adaptation runs use the complete 152-sample calibration split for optimization. The fixed 152-sample test split is not used for final adaptation training, within-run validation, or checkpoint selection. Final reported evaluation uses the fixed sixth-epoch model.

The depth-based B48/B72 and projection-based P24/P48 target sets were frozen from sensitivity rankings computed over the broader 304-sample analysis set before final post-quantization adaptation. The nested-random R12–R72 plan and the separately sampled S0–S4 plans use fixed random target selections rather than these sensitivity scores. Accordingly, the fixed test split is a holdout with respect to final adaptation, within-run validation, and checkpoint selection, but it is not independent of the earlier 304-sample sensitivity rankings used to define B48/B72 and P24/P48.

Path-matched teacher reference. Teacher targets used during adaptation and the reference used for final evaluation originate from the same frozen 15,000-step checkpoint. For all reported teacher-relative test metrics, we evaluate this checkpoint under the same deterministic inference path used for the stored test predictions and use those outputs as the path-matched teacher reference.

The matched inference path disables flash and memory-efficient SDP kernels, enables the math SDP path, disables cuDNN benchmarking, enables deterministic cuDNN behavior, and enables deterministic PyTorch algorithms. The raw FP-B48, FP-R48, and FP-P48 predictions are bitwise identical under this path and therefore have exactly zero teacher-relative T6 before adaptation.

Teacher outputs cached for adaptation were generated from the same frozen checkpoint under the training execution path. Relative to the path-matched final-evaluation reference, these cached outputs differ by mean absolute value 0.003106 over the six motion dimensions, with maximum absolute difference 0.017578. The checkpoint, samples, preprocessing, and action definition are otherwise unchanged. All final teacher-relative scalar and residual-geometry results reported in this manuscript use the path-matched reference.

Adaptation objective. For the primary GT+KD experiments, the adaptation objective is

$$\mathcal{L}_{\text{adapt}} = \mathcal{L}_{\text{GT}} + \lambda \mathcal{L}_{\text{KD}}, \tag{6}$$

with $\lambda = 0.5$. Here, $\mathcal{L}_{\text{GT}}$ is the mean absolute error between the student prediction and ground-truth action over all $H \times 7$ action entries, whereas $\mathcal{L}_{\text{KD}}$ is the mean absolute error between the student and frozen teacher over the $H \times 6$ continuous-motion entries. The primary GT+KD condition uses six epochs, learning rate 7×10^{-5}, and training seed 20260903.

Matched GT-only controls. The matched objective-control experiment uses $\lambda = 0$ while retaining the same calibration split, six-epoch adaptation budget, learning rate 7×10^{-5}, and training seed 20260903. The six configurations used for this comparison are B48, B72, R12, R24, R48, and R72.

Trainable parameters for quantized adaptation. For each quantization configuration, all policy parameters are initially frozen, after which only the latent full-precision weight tensors corresponding to the selected quantized linear matrices are made trainable. The action head and all unselected weights and biases remain frozen. Gradients through the quantized forward weights are propagated to the trainable latent weights using the straight-through estimator.

Matched full-precision controls. The FP-B48, FP-R48, and FP-P48 controls start from the same frozen full-precision teacher and make trainable the same parameter subset as their corresponding quantized configurations. All three matched controls use $\lambda = 0.5$, learning rate 7×10^{-5}, six epochs, training seed 20260903, `PagedAdamW8bit`, and zero weight decay. Each control contains 2,164,260,864 trainable parameters.

The selected weights use the ordinary full-precision forward pass: quantization is disabled and no STE forward patch is installed. The action head and all remaining parameters remain frozen. Thus, these controls preserve the downstream adaptation protocol while removing the initial quantization perturbation.

Optimization. The primary training loop processes one action-chunk sample per forward/backward pass and uses gradient accumulation of one, so each optimizer update corresponds to one sample. Forward and loss computation use BF16 autocast. Optimization uses `PagedAdamW8bit` with zero weight decay. The 8-bit optimizer-state representation is distinct from the ternary or 3-bit quantization applied to model weights.

For the optimization-seed robustness analysis, B48, R48, and P48 are additionally trained using seeds 20260904 and 20260905 while holding the data split, objective, learning rate, and training schedule fixed.

Hardware. All adaptation runs reported in this work were executed on the same single-GPU workstation equipped with an NVIDIA GeForce RTX 5090 with 32 GB of GPU memory.

Table 2: Primary settings used for the final adaptation experiments.

Setting	Value
Model	OpenVLA-OFT
Task	simulated MyCobot cube stacking
Teacher data	385 demonstrations
Teacher split seed	42
Teacher optimization split	347 demonstrations
Teacher validation / analysis set	38 demonstrations
Teacher checkpoint	fixed final step 15,000
Analysis samples	38 demos / 304 samples
Downstream split seed	20260827
Calibration set	19 demos / 152 samples
Test set	19 demos / 152 samples
Action chunk	8×7
Geometry space	8×6 motion cells
Epochs	6
Learning rate	7×10^{-5}
GT+KD coefficient	$\lambda = 0.5$
Matched GT-only coefficient	$\lambda = 0$
Optimizer	PagedAdamW8bit
Weight decay	0
Per-step batch size	1
Gradient accumulation	1
Forward precision	BF16 autocast
Primary training seed	20260903
Additional seeds	20260904 / 20260905
Final-run validation	none
Checkpoint selection	fixed final epoch
Hardware	NVIDIA GeForce RTX 5090 (32 GB)

B Quantization Configurations and Quantizer Definitions

Ternary quantization. For each selected weight tensor W, we define the scalar scale

$$s = \max\left(\frac{1}{|W|}\sum_{i,j}|w_{ij}|, 10^{-6}\right), \tag{7}$$

where $|W|$ denotes the number of entries in W.

Each weight is assigned a ternary code

$$q_{ij} = \begin{cases} +1, & w_{ij}/s > 0.5, \\ -1, & w_{ij}/s < -0.5, \\ 0, & \text{otherwise}, \end{cases} \tag{8}$$

and the quantized forward weight is

$$\hat{w}_{ij} = s\, q_{ij}. \tag{9}$$

Thus, the realized weight values are $\{-s, 0, +s\}$, with a tensor-wise mean-absolute-weight scale and thresholds at $\pm 0.5s$. Weights exactly at either threshold map to the zero code.

During post-quantization adaptation, the forward pass uses $\hat{W}$ while gradients are propagated to the latent full-precision weights using a straight-through estimator.

The 13 ternary target sets remain fixed throughout the final adaptation and evaluation runs. They are intentionally heterogeneous and provide distinct raw quantization perturbations rather than defining a single optimized layer-selection strategy.

Configuration selection. For the present residual-geometry analysis, B48/B72 and P24/P48 were frozen from sensitivity-based rankings computed over the full 304-sample analysis set before final

post-quantization adaptation. These plans are used here as fixed perturbation configurations rather than as claims about optimal layer selection. B48 and B72 use depth-based target rankings with budgets of 48 and 72 matrices, respectively, while P24 and P48 use projection-based target rankings with a fixed random depth tie-break and budgets of 24 and 48 targets, respectively. By contrast, the nested R12–R72 family is generated by a fixed-seed random ranking, and S0–S4 use five separately sampled frozen Random-24 plans. These random-family target identities therefore do not depend on the sensitivity rankings. Random selection operates over 24 depth–projection cells, each containing the same projection matrix in four consecutive transformer blocks; a Random-24 configuration therefore selects six such cells.

Depth-based configurations. B48 quantizes the `gate_proj`, `up_proj`, and `down_proj` matrices in transformer blocks 12–27, for a total of 48 linear matrices. B72 extends the same construction to blocks 4–27, for a total of 72 linear matrices.

Nested random configurations. R12, R24, R48, and R72 are loaded from a single frozen random selection plan containing the exact target-module identities for each budget. The target sets satisfy

$$R12 \subset R24 \subset R48 \subset R72, \tag{10}$$

with 12, 24, 48, and 72 unique target matrices, respectively. No target resampling is performed during post-quantization adaptation.

Separately sampled Random-24 configurations. S0–S4 use five separately sampled and frozen random selection plans generated with random seeds $0, 1, 2, 3, 4$, respectively. Each plan specifies a 24-matrix target set that is reused unchanged during post-quantization adaptation.

These independently sampled configurations provide distinct random target realizations in addition to the nested R12–R72 sequence and test whether the observed geometry depends on one particular random selection.

The exact module identities for the nested and independently sampled random configurations are fixed in the corresponding frozen selection-plan files used throughout the experiments.

Projection-based configurations. P24 and P48 are constructed from a frozen projection-sensitivity ranking with a fixed depth tie-break. P24 contains 24 `up_proj` matrices. P48 contains 32 `up_proj` matrices and 16 `gate_proj` matrices, for 48 targets in total. P24 is a subset of P48.

3-bit symmetric-uniform quantization. For the cross-quantizer robustness experiment, each output channel i uses scale

$$s_i = \frac{\max_j |w_{ij}|}{3}. \tag{11}$$

Weights are quantized according to

$$q_{ij} = \mathrm{clip}\left(\mathrm{round}\left(\frac{w_{ij}}{s_i}\right), -3, 3\right), \qquad \hat{w}_{ij} = s_i q_{ij}. \tag{12}$$

The forward pass uses $\hat{w}_{ij}$, while the backward pass applies an identity straight-through estimator to the latent full-precision weight.

We evaluate this quantizer on B72, P48, and R48. The 3-bit configurations never contribute to construction of the ternary reference templates.

C Complete Ternary Recovery Results

T6 is the mean absolute teacher-relative residual over the six continuous motion dimensions and all eight chunk steps. Scalar recovery is defined as

$$100\frac{\mathrm{T6}_{\mathrm{raw}} - \mathrm{T6}_{\mathrm{final}}}{\mathrm{T6}_{\mathrm{raw}}}. \tag{13}$$

Table 3: Reference construction for the residual-geometry alignment analyses. "Primary" denotes the primary training seed 20260903.

Analysis	Evaluated target	Reference configurations	Reference stage
Ternary LOO, raw	raw ternary configuration c	other 12 ternary configurations	raw
Ternary LOO, adapted	adapted ternary configuration c	other 12 ternary configurations	adapted
Fixed adapted ref., raw	raw ternary configuration c	other 12 ternary configurations	adapted
Fixed adapted ref., adapted	adapted ternary configuration c	other 12 ternary configurations	adapted
Additional optimization seed	adapted B48/R48/P48 at alternate seed	other 12 primary-seed ternary configurations	adapted
3-bit, raw	raw 3-bit configuration	all 13 ternary configurations	raw
3-bit, adapted	adapted 3-bit configuration	all 13 ternary configurations	adapted
Full-precision control	adapted FP-B48/FP-R48/FP-P48	adapted ternary set excluding matched counterpart	adapted

For the geometry analysis, let $E^{(c)} \in \mathbb{R}^{8\times 6}$ denote the teacher-relative residual map for configuration c. The *Full* representation uses $E^{(c)}$ directly. The *Interaction* representation removes chunk-position and action-dimension main effects:

$$E^{(c)}_{\mathrm{int},h,d} = E^{(c)}_{h,d} - \overline{E}^{(c)}_{h,\cdot} - \overline{E}^{(c)}_{\cdot,d} + \overline{E}^{(c)}. \tag{14}$$

For *Nuisance* residualization, each of the 48 cells is associated with three measured nuisance features computed across test samples: teacher-output standard deviation, ground-truth standard deviation, and mean absolute teacher–ground-truth discrepancy. These three features are standardized across cells, an intercept is added, and the flattened residual map is linearly residualized against them.

For all three representations, the resulting 48-cell vector is replaced by its ranks and standardized across cells. Let $z^{(c)}_{r,s} \in \mathbb{R}^{48}$ denote this rank-standardized vector for representation r and evaluation stage $s \in \{\mathrm{raw}, \mathrm{adapted}\}$.

Pairwise similarity and leave-one-out alignment. For a set of K configurations, mean pairwise similarity is

$$\overline{\rho}_{r,s} = \frac{2}{K(K-1)} \sum_{c<c'} \rho_{\mathrm{S}} \left(z^{(c)}_{r,s}, z^{(c')}_{r,s} \right), \tag{15}$$

where ρ_{S} denotes Spearman correlation. The primary ternary analysis uses $K = 13$.

For leave-one-out evaluation of configuration c, the same-stage reference template is

$$t^{(-c)}_{r,s} = \frac{1}{K-1} \sum_{c' \neq c} z^{(c')}_{r,s}, \tag{16}$$

and alignment is

$$A^{(c)}_{r,s} = \frac{z^{(c)}_{r,s} \cdot t^{(-c)}_{r,s}}{\left\| z^{(c)}_{r,s} \right\|_2 \left\| t^{(-c)}_{r,s} \right\|_2}. \tag{17}$$

Thus, the evaluated configuration never contributes to its own reference template. For two individually rank-standardized vectors, cosine similarity is equivalent to Spearman correlation; because the leave-one-out reference is an average of multiple standardized vectors, we retain the cosine definition explicitly for template alignment.

Table 3 summarizes the reference construction used in the primary and robustness analyses.

All 13 configurations increase their Full leave-one-out alignment after adaptation despite substantial heterogeneity in their raw states. All 13 also increase after nuisance residualization. The Interaction representation is less uniform: R12 and S2 show configuration-level decreases even though aggregate Interaction similarity increases substantially.

At the pairwise aggregate level, mean cross-configuration similarity changes from 0.037 to 0.557 for Full residual maps, from 0.053 to 0.321 for Interaction, and from 0.033 to 0.570 for Nuisance.

The complete scalar results further illustrate the distinction between scalar recovery and residual geometry. R48 recovers approximately 79% of its scalar teacher-relative error while ending with Full

Table 4: Complete scalar-recovery and leave-one-out residual-geometry results for all 13 ternary configurations. Each LOO template is constructed from the other 12 configurations at the same evaluation stage. The statement that all 13 configurations improve refers to the Full residual representation.

	Scalar			Full LOO		Interaction LOO		Nuisance LOO	
Config	Raw T6	Final T6	Rec. (%)	Raw	Final	Raw	Final	Raw	Final
B48	0.0865	0.0423	+51.1	+0.318	+0.735	+0.328	+0.528	+0.201	+0.762
B72	0.1245	0.0489	+60.7	-0.136	+0.663	-0.142	+0.685	-0.345	+0.630
R12	0.0406	0.0421	-3.7	+0.597	+0.683	+0.505	+0.495	+0.476	+0.680
R24	0.0841	0.0517	+38.6	-0.214	+0.606	+0.146	+0.400	-0.357	+0.558
R48	0.3916	0.0822	+79.0	-0.166	+0.483	+0.319	+0.576	-0.163	+0.374
R72	0.4308	0.0710	+83.5	+0.046	+0.762	+0.116	+0.765	-0.030	+0.844
S0	0.6213	0.0676	+89.1	-0.231	+0.720	-0.030	+0.632	-0.207	+0.791
S1	0.3190	0.0463	+85.5	+0.421	+0.865	+0.121	+0.326	+0.611	+0.891
S2	0.0609	0.0489	+19.6	+0.396	+0.742	+0.151	+0.102	+0.527	+0.759
S3	0.6772	0.0591	+91.3	+0.246	+0.750	+0.062	+0.598	+0.508	+0.799
S4	0.4288	0.0457	+89.4	+0.268	+0.868	+0.369	+0.526	+0.448	+0.883
P24	0.2016	0.0472	+76.6	+0.482	+0.876	+0.425	+0.624	+0.636	+0.878
P48	0.4281	0.0710	+83.4	-0.339	+0.662	-0.331	+0.578	-0.521	+0.702

Table 5: Fixed adapted-reference control across the 13 ternary configurations. Both raw and adapted states are evaluated against the same leave-one-out adapted reference. Bootstrap statistics report the adapted-minus-raw mean alignment gain under demonstration-level resampling.

Representation	Raw	Adapted	Point Δ	Bootstrap median Δ	95% CI
Full	-0.046	0.724	$+0.770$	$+0.763$	$[0.731, 0.793]$
Interaction	0.192	0.526	$+0.333$	$+0.356$	$[0.258, 0.463]$
Nuisance	-0.091	0.735	$+0.825$	$+0.819$	$[0.767, 0.860]$

leave-one-out alignment of 0.483. R12 provides a complementary low-error case: its T6 changes from 0.0406 to 0.0421, while its final Full leave-one-out alignment remains 0.683. The 3b-P48 result in Appendix H provides a stronger case of limited scalar recovery with strong geometric alignment.

Fixed adapted-reference control. The primary leave-one-out analysis uses a reference template matched to the evaluation stage. To test whether the observed alignment increase depends on changing that reference from raw to adapted, we repeat the analysis while holding the reference stage fixed. For each held-out configuration, both its raw and adapted residual vectors are evaluated against the same leave-one-out template formed from the adapted residual vectors of the other 12 ternary configurations. The held-out configuration never contributes to this reference.

The fixed-reference gain remains positive throughout the demonstration-bootstrap interval for all three residual representations. At the point estimate, alignment increases for 13 of 13 configurations in Full, 11 of 13 in Interaction, and 13 of 13 in Nuisance. For Full, every configuration also has a positive pointwise 95% bootstrap interval. Thus, the shared-alignment increase persists when raw and adapted residuals are evaluated against the same fixed adapted reference, showing that the aggregate trend does not depend on changing the reference stage.

C.1 Separately Sampled Random-24 Subset Analysis

Several configurations in the full ternary set are structurally related; in particular, the R12–R72 family is nested. We therefore repeat the pairwise residual-geometry analysis using only S0–S4, five separately sampled and frozen Random-24 configurations.

Mean pairwise similarity increases from 0.313 to 0.664 for Full, from 0.009 to 0.183 for Interaction, and from 0.211 to 0.686 for Nuisance. Thus, the aggregate increase remains after removing the nested R12–R72 family from this analysis.

The result is also stable under demonstration-level bootstrap resampling. The median adapted-minus-raw increase is $+0.343$ (95% CI $[0.287, 0.394]$) for Full, $+0.217$ ($[0.100, 0.320]$) for Interaction, and $+0.423$ ($[0.328, 0.507]$) for Nuisance.

Table 6: Residual-geometry similarity restricted to the five separately sampled Random-24 configurations (S0–S4). Bootstrap statistics report the adapted-minus-raw change in mean pairwise similarity under demonstration-level resampling.

Representation	Raw	Adapted	Δ	Bootstrap median Δ (95% CI)
Full	0.313	0.664	+0.351	+0.343 [0.287, 0.394]
Interaction	0.009	0.183	+0.174	+0.217 [0.100, 0.320]
Nuisance	0.211	0.686	+0.475	+0.423 [0.328, 0.507]

Table 7: Optimization-seed robustness for B48, R48, and P48. All rows are evaluated against the same fixed counterpart-excluded primary-seed adapted reference for each configuration and representation. "Raw" reports the raw configuration against this fixed adapted reference, and the three seed rows report the corresponding adapted checkpoints against the same reference.

Config	State / seed	Full	Interaction	Nuisance
B48	Raw	0.421	0.479	0.169
B48	20260903	0.735	0.528	0.762
B48	20260904	0.661	0.694	0.618
B48	20260905	0.853	0.713	0.834
R48	Raw	0.070	0.166	0.050
R48	20260903	0.483	0.576	0.374
R48	20260904	0.612	0.747	0.499
R48	20260905	0.632	0.723	0.420
P48	Raw	-0.237	-0.466	-0.213
P48	20260903	0.662	0.578	0.702
P48	20260904	0.618	0.745	0.082
P48	20260905	0.621	0.541	-0.227

Not every individual configuration pair increases in similarity. This result therefore establishes an aggregate cross-configuration trend rather than convergence of every pair toward an identical residual map.

C.2 Optimization-Seed Robustness

We evaluate optimization-seed sensitivity on three representative ternary configurations, B48, R48, and P48. In addition to the primary seed 20260903, we repeat the same six-epoch GT+KD adaptation protocol using seeds 20260904 and 20260905. The data split, quantization configuration, learning rate, objective, and evaluation set are unchanged.

For each alternate-seed run, the evaluated configuration is compared with the fixed adapted primary-seed template formed from the other 12 ternary configurations, as summarized in Table 3. This keeps the external reference fixed while varying the optimization seed of the evaluated configuration.

Across the nine configuration–seed runs in Table 7, adapted alignment exceeds the corresponding raw alignment in 26 of 27 representation-level comparisons against the same fixed adapted reference. The sole exception is P48 Nuisance at seed 20260905, with a small point change of approximately -0.014.

As an aggregate robustness check, we replace B48, R48, and P48 in the 13-configuration analysis by the corresponding alternate-seed runs while leaving the other ten primary-seed configurations fixed. With the primary seed, mean pairwise similarity is 0.557 for Full, 0.321 for Interaction, and 0.570 for Nuisance. Replacing the three evaluated configurations by their seed-20260904 runs gives 0.562, 0.368, and 0.506, respectively; using their seed-20260905 runs gives 0.583, 0.356, and 0.477. All remain substantially above the corresponding raw similarities 0.037, 0.053, and 0.033.

The exact final residual maps are not seed-invariant. Substantial seed-to-seed variation remains; the most pronounced observed case is nuisance-residualized P48, whose seed-20260903 and seed-20260905 maps have Spearman correlation approximately -0.002. Thus, the multiseed experiment

establishes robustness of the shared-alignment trend across the tested optimization seeds without implying convergence to an identical residual map across seeds.

D Complete Residual-Geometry Maps

Figure 4 provides a visual audit of all 13 ternary configurations. Each individual heatmap is constructed by averaging absolute teacher-relative residual magnitude over the fixed test samples and then rank-standardizing the resulting 48 cells. Colors therefore encode the relative ordering of residual magnitude within a configuration rather than absolute T6 magnitude.

The raw and adapted mean maps at the top are visualization aids obtained by averaging the already rank-standardized maps across all 13 configurations. These averaged maps are not re-standardized after averaging, so their color strength reflects the degree to which configuration-level patterns survive cross-configuration averaging. They are not used as leave-one-out reference templates or evaluation statistics.

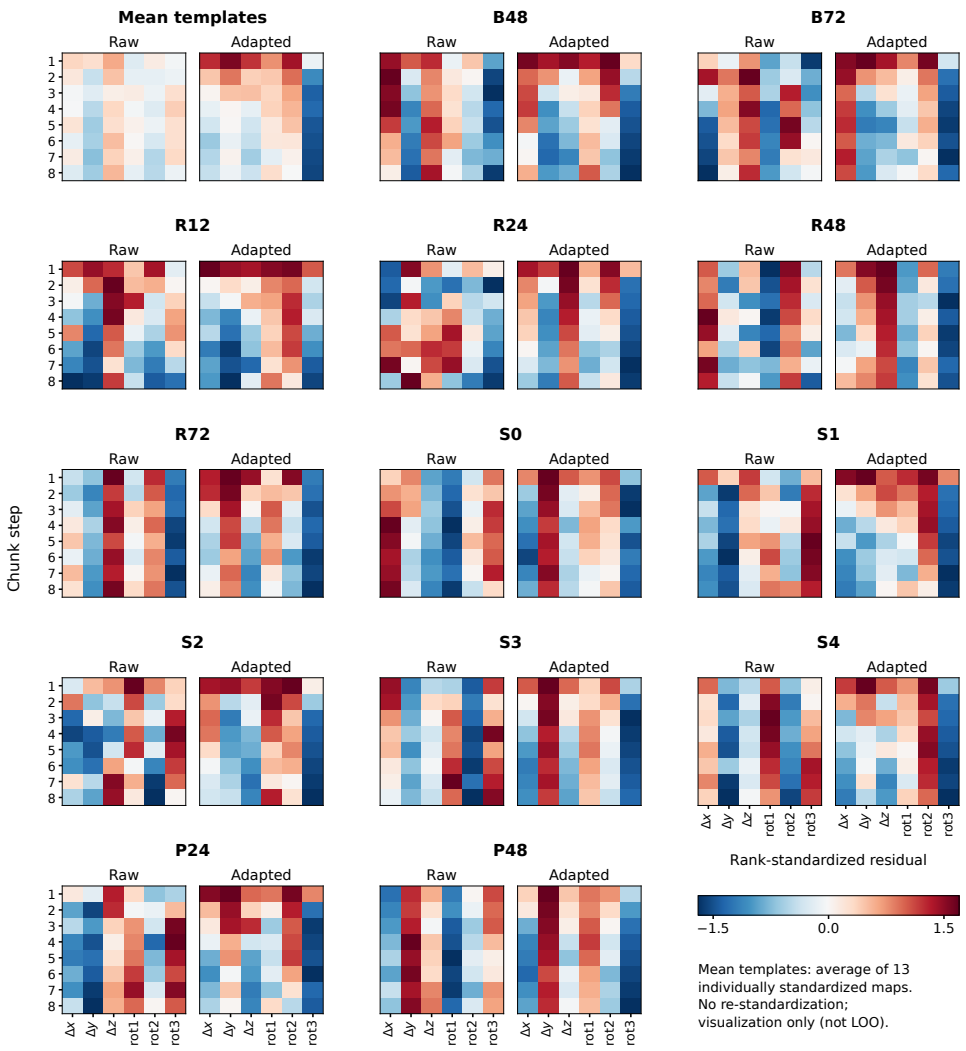


Figure 4: Complete teacher-relative residual geometry across the 13 ternary quantization configurations. The top row shows raw and adapted mean rank-standardized residual maps across all 13 configurations as visual references; these mean maps are not used as leave-one-out evaluation statistics. Individual raw and adapted 8×6 residual maps are shown below. The weak raw mean reflects heterogeneous initial residual geometries, whereas the adapted mean reveals a stronger shared component while substantial configuration-specific structure remains.

E Matched Full-Precision Adaptation Controls

The matched full-precision controls test whether alignment with the shared adapted geometry requires an initial quantization perturbation. FP-B48, FP-R48, and FP-P48 use the same parameter subsets, GT+KD objective, learning rate, six-epoch budget, primary training seed, optimizer, and frozen action head as their corresponding quantized experiments, but retain the selected weights in full precision and do not use the quantized STE forward pass.

Under the path-matched teacher reference, the unadapted full-precision controls coincide exactly with the frozen teacher and therefore have teacher-relative motion T6 equal to zero. After downstream adaptation, T6 is 0.053515 for FP-B48, 0.065867 for FP-R48, and 0.047759 for FP-P48.

Table 8: Matched full-precision control alignment to the adapted ternary reference structure. The corresponding quantized configuration is excluded from each reference template. Confidence intervals are obtained by demonstration-level bootstrap resampling over the 19 test demonstrations (2,000 replicates; seed 20260916).

	Full		Interaction		Nuisance	
Control	Point	95% CI	Point	95% CI	Point	95% CI
FP-B48	0.811	[0.736, 0.896]	0.543	[0.358, 0.684]	0.802	[0.732, 0.883]
FP-R48	0.421	[0.303, 0.558]	0.664	[0.518, 0.769]	0.183	[0.019, 0.319]
FP-P48	0.843	[0.757, 0.905]	0.489	[0.291, 0.736]	0.874	[0.778, 0.921]

We therefore interpret these experiments as adapted-endpoint controls rather than scalar-recovery experiments.

For each full-precision control, the adapted ternary reference template excludes the corresponding quantized configuration. Thus, B48 is excluded when evaluating FP-B48, R48 is excluded when evaluating FP-R48, and P48 is excluded when evaluating FP-P48.

All nine configuration–representation pointwise 95% bootstrap intervals lie above zero. The strength of the alignment nevertheless remains configuration-specific: FP-B48 and FP-P48 show strong Full and Nuisance alignment, whereas Nuisance alignment for FP-R48 is substantially weaker.

Together, these controls demonstrate that matched unquantized adaptation develops substantial alignment with the shared adapted residual structure. The shared alignment therefore does not require an initial quantization perturbation. Substantial configuration-specific residual variation nevertheless remains.

F Shared Adapted Endpoint and Recovery-Field Analysis

For a residual map E, we define its centered log-shape as

$$\ell(E)_{h,d} = \log(E_{h,d} + \epsilon) - \frac{1}{HD} \sum_{h',d'} \log(E_{h',d'} + \epsilon)\,, \tag{18}$$

where $\epsilon = 10^{-8}$. Mean-centering in log space removes a configuration-wide multiplicative scale, so this analysis focuses on residual shape rather than overall residual magnitude.

For each ternary configuration c, its centered raw-to-adapted recovery field is

$$G^{(c)} = \ell\Big(E^{(c)}_{\text{final}}\Big) - \ell\Big(E^{(c)}_{\text{raw}}\Big)\,. \tag{19}$$

For a held-out configuration c, the leave-one-out shared recovery field is

$$\bar{G}_{-c} = \frac{1}{12} \sum_{c' \neq c} G^{(c')}. \tag{20}$$

We compare four predictors of the held-out adapted shape. The raw-shape baseline is

$$\hat{\ell}^{(c)}_{\text{raw}} = \ell\Big(E^{(c)}_{\text{raw}}\Big)\,. \tag{21}$$

The shared recovery-field predictor is

$$\hat{\ell}^{(c)}_{\text{field}} = \ell\Big(E^{(c)}_{\text{raw}}\Big) + \bar{G}_{-c}. \tag{22}$$

The mean-final predictor directly averages the adapted centered log-shapes of the other 12 configurations:

$$\hat{\ell}^{(c)}_{\text{mean-final}} = \frac{1}{12} \sum_{c' \neq c} \ell\Big(E^{(c')}_{\text{final}}\Big)\,. \tag{23}$$

Table 9: Leave-one-out prediction of final teacher-relative residual geometry. Additive retains only chunk-position and action-dimension main effects from the leave-one-out mean-final template. Bold entries denote the best predictor in each row and metric.

	Spearman ρ				RMSE			
Config	Raw	Field	Additive	Mean-final	Raw	Field	Additive	Mean-final
B48	+0.592	**+0.774**	+0.561	+0.654	0.318	0.411	0.201	**0.192**
B72	-0.043	+0.283	+0.487	**+0.524**	0.448	0.462	0.232	**0.227**
R12	+0.440	**+0.739**	+0.554	+0.632	0.295	0.297	0.260	**0.254**
R24	+0.077	+0.080	+0.424	**+0.448**	0.384	0.474	0.258	**0.253**
R48	+0.186	+0.407	**+0.416**	+0.372	0.474	0.504	0.282	**0.277**
R72	+0.074	+0.216	**+0.788**	+0.786	0.517	0.642	0.133	**0.118**
S0	-0.442	+0.382	**+0.790**	+0.746	0.631	0.509	**0.303**	0.306
S1	-0.146	+0.415	+0.773	**+0.827**	0.662	0.415	0.160	**0.152**
S2	-0.073	+0.560	+0.601	**+0.653**	0.434	0.312	0.239	**0.235**
S3	-0.363	+0.121	**+0.785**	+0.776	0.857	0.628	0.282	**0.278**
S4	-0.449	-0.239	+0.835	**+0.837**	0.976	0.805	0.164	**0.156**
P24	-0.498	+0.283	+0.851	**+0.876**	0.583	0.373	0.106	**0.084**
P48	+0.101	+0.392	**+0.695**	+0.680	0.662	0.614	0.307	**0.304**
Mean	-0.042	+0.339	+0.658	**+0.678**	0.557	0.496	0.225	**0.218**

To determine how much of this cross-configuration prediction is captured by additive chunk-position and action-dimension structure, let

$$M^{(c)} = \hat{\ell}^{(c)}_{\text{mean-final}}. \tag{24}$$

We define the additive mean-final predictor as

$$\hat{\ell}^{(c)}_{\text{add},h,d} = \overline{M}^{(c)}_{h,\cdot} + \overline{M}^{(c)}_{\cdot,d} - \overline{M}^{(c)}_{\cdot,\cdot}, \tag{25}$$

where the three terms denote the row mean, column mean, and grand mean of the leave-one-out mean-final template, respectively. This predictor therefore retains only chunk-position and action-dimension main effects from the shared adapted endpoint. The held-out configuration's final residual map is not used in its construction.

The recovery-field predictor can equivalently be written as

$$\hat{\ell}^{(c)}_{\text{field}} = \hat{\ell}^{(c)}_{\text{mean-final}} + \left[\ell\Big(E^{(c)}_{\text{raw}}\Big) - \frac{1}{12} \sum_{c' \neq c} \ell\Big(E^{(c')}_{\text{raw}}\Big) \right]. \tag{26}$$

This identity makes the relationship between the recovery-field and mean-final predictors explicit. The mean-final predictor represents the shared adapted endpoint directly, whereas the recovery-field predictor additionally carries forward the held-out configuration's deviation from the mean raw residual shape of the other configurations.

All predictors are evaluated against $\ell(E^{(c)}_{\text{final}})$ over the 48 action-chunk cells using Spearman correlation and RMSE. The reported RMSE is therefore measured in centered-log shape space rather than in the original residual or T6 scale. This leave-one-out procedure holds out a configuration, not test demonstrations: all predictors and targets are evaluated on the same fixed 19-demonstration test set. The analysis is distinct from the rank-standardized cosine-template analysis in Appendix C; here the predictors operate directly on centered log-shapes.

The recovery-field predictor improves mean Spearman correlation relative to the raw-shape baseline from -0.042 to 0.339 and reduces mean RMSE from 0.557 to 0.496. The recovery-field construction is therefore informative relative to raw geometry.

The direct mean-final predictor is substantially stronger than both the raw and recovery-field predictors. Its mean Spearman correlation is 0.678, and its mean RMSE is 0.218. Mean-final has higher Spearman correlation than the raw baseline for all 13 configurations and higher correlation than the recovery-field predictor for 10 of 13. It also has lower RMSE than the recovery-field predictor for all 13 configurations.

The additive mean-final baseline reaches mean Spearman correlation 0.658 and mean RMSE 0.225. Thus, shared chunk-position and action-dimension main effects capture a large fraction of the cross-configuration predictive signal in the shared adapted endpoint. The complete mean-final template

Table 10: Demonstration-bootstrap comparisons of the mean-final predictor. Spearman contrasts use mean-final minus the comparator; the RMSE contrast uses additive minus mean-final, so positive values favor mean-final in every row.

Contrast	Point	Bootstrap median	95% CI
Spearman: mean-final − raw	+0.720	+0.722	[0.662, 0.788]
Spearman: mean-final − field	+0.338	+0.338	[0.297, 0.383]
Spearman: mean-final − additive	+0.019	+0.019	[-0.001, 0.036]
RMSE: additive − mean-final	+0.007	+0.008	[0.005, 0.011]

Table 11: Aggregate teacher-relative and GT-relative prediction under the raw-shape, recovery-field, and mean-final predictors. Higher Spearman correlation and lower RMSE are better.

	Mean Spearman $\rho \uparrow$			Mean RMSE $\downarrow$		
Reference	Raw	Field	Mean-final	Raw	Field	Mean-final
Teacher	-0.042	0.339	0.678	0.557	0.496	0.218
GT	0.499	0.492	0.963	0.178	0.195	0.0165

nevertheless retains additional predictive structure: it has higher point-estimate Spearman correlation than the additive baseline for 8 of 13 configurations and lower point-estimate RMSE for 12 of 13.

For reference, mean pairwise Spearman similarity among the adapted configurations is 0.557. The stronger mean-final prediction is consistent with averaging away part of the configuration-specific variation while retaining shared adapted structure. The additive comparison further shows that a substantial part of this shared endpoint is organized through common chunk-position and action-dimension effects, while non-additive cell structure contributes additional predictive information.

The mean-final Spearman advantages over the raw-shape and recovery-field predictors are positive in all demonstration-bootstrap resamples. Relative to the additive baseline, the mean-final Spearman advantage is smaller: its bootstrap median is $+0.019$ with 95% CI $[-0.001,\ 0.036]$. By contrast, the RMSE advantage of the complete mean-final template is consistent under resampling. The additive-minus-mean-final RMSE contrast has bootstrap median $+0.008$ with 95% CI $[0.005,\ 0.011]$, and is positive in all 2,000 bootstrap replicates.

At the observed point estimate, mean-final exceeds the additive predictor in Spearman correlation for 8 of 13 configurations and in RMSE for 12 of 13. Across bootstrap replicates, the median numbers of mean-final wins are 8 of 13 for Spearman and 12 of 13 for RMSE, with 95% bootstrap-quantile ranges of 6–11 and 10–13, respectively.

These results establish the shared adapted endpoint as a major source of cross-configuration predictability in the final residual geometry. The additive control shows that much of this predictability is carried by shared chunk-position and action-dimension effects. The complete residual template nevertheless retains additional predictive structure, most clearly in reconstruction error, where its aggregate advantage remains positive under paired demonstration-level resampling.

The recovery-field predictor improves over the raw-shape baseline, but it does not outperform the direct mean-final predictor on average. Thus, the recovery-field construction is informative relative to raw geometry, while these results do not establish additional predictive value beyond the shared adapted endpoint.

GT-Relative Recovery-Field Diagnostic. We repeat the recovery-field construction after replacing the frozen teacher with the demonstration target as the residual reference. Unlike the teacher-relative analysis, the GT-relative recovery-field predictor does not improve aggregate prediction relative to the raw-shape baseline.

For GT-relative residuals, the recovery-field predictor does not improve over the raw-shape baseline: mean Spearman correlation changes from 0.499 to 0.492, while mean RMSE changes from approximately 0.178 to 0.195. By contrast, the GT-relative mean-final template strongly predicts

final residual geometry, reaching mean Spearman correlation 0.963 and mean RMSE 0.0165. Thus, the shared adapted endpoint is highly predictive under the GT-relative reference as well, whereas the recovery-field gain observed in the teacher-relative frame is not reproduced. Because the teacher-relative and GT-relative analyses begin from substantially different baseline geometries, we do not interpret differences in their absolute predictive performance as evidence for causal teacher specificity.

G Matched Objective-Control Experiments

The matched objective-control analysis uses six ternary configurations: B48, B72, R12, R24, R48, and R72. It compares the same raw quantized states after matched GT-only adaptation and after the primary GT+KD adaptation.

Both conditions use the same calibration split, learning rate 7×10^{-5}, six-epoch training budget, training seed 20260903, optimizer, weight decay, trainable parameter subsets, and frozen action head. They differ in the teacher-distillation term: $\lambda = 0$ for GT-only and $\lambda = 0.5$ for GT+KD.

To verify that the GT-only condition performs effective downstream adaptation rather than changing residual geometry in isolation, we additionally evaluate ground-truth action error on the fixed test split. GT7 is the mean absolute error over all 8×7 preprocessed action entries on the fixed 152-sample test split.

Across the six configurations, mean GT7 decreases from 0.3654 before adaptation to 0.2848 under GT-only training and 0.2753 under GT+KD. Thus, the shared-geometry increase under GT-only adaptation is accompanied by a substantial aggregate reduction in ground-truth action error on the fixed test split. GT+KD attains the lowest mean GT7, while its incremental reduction relative to GT-only is modest compared with the overall improvement from the raw quantized states.

The same matched control shows that teacher distillation is not required for the shared-alignment increase to emerge. Table 13 reports the aggregate pairwise alignment for the three residual representations.

GT-only adaptation reproduces most of the aggregate increase observed under GT+KD in all three representations. The demonstration-bootstrap GT+KD-minus-GT-only contrast has median $+0.055$ for Full, with 95% CI $[-0.011,\ 0.142]$ and bootstrap-positive fraction 0.942. For Interaction, the median contrast is $+0.045$, with 95% CI $[-0.124,\ 0.237]$ and bootstrap-positive fraction 0.682. For Nuisance, the median contrast is $+0.133$, with 95% CI $[0.008,\ 0.265]$ and bootstrap-positive fraction 0.981.

Mean alignment is higher under GT+KD in this matched subset, but only the Nuisance contrast has a 95% bootstrap interval that excludes zero. Thus, explicit teacher matching is not necessary for the shared-alignment increase to emerge in this matched six-configuration setting, although the Nuisance representation shows an additional GT+KD advantage under the reported bootstrap analysis.

Table 12: Ground-truth action error on the fixed test split for the matched objective-control experiments. GT7 is mean absolute error over all 8×7 preprocessed action entries on the fixed 152-sample test split. Lower is better.

Configuration	Raw	GT-only	GT+KD
B48	0.3122	0.2776	0.2696
B72	0.3360	0.2723	0.2709
R12	0.2772	0.2809	0.2682
R24	0.2848	0.2831	0.2713
R48	0.4922	0.2931	0.2915
R72	0.4903	0.3018	0.2801
Mean	0.3654	0.2848	0.2753

Table 13: Aggregate residual-geometry alignment for the matched six-configuration objective control. The final column reports the demonstration-bootstrap GT+KD-minus-GT-only contrast.

Representation	Raw	GT-only	GT+KD	Bootstrap median Δ (95% CI)
Full	0.186	0.445	0.497	+0.055 [-0.011, 0.142]
Interaction	0.068	0.325	0.367	+0.045 [-0.124, 0.237]
Nuisance	-0.086	0.305	0.431	+0.133 [0.008, 0.265]

H Cross-Quantizer Robustness

To test whether the shared-alignment trend is specific to ternary quantization, we evaluate a distinct per-output-channel symmetric-uniform 3-bit quantizer on B72, P48, and R48. The exact 3-bit quantizer definition is given in Appendix B.

The three 3-bit configurations are evaluated against fixed ternary reference templates and never contribute to construction of those templates. Raw 3-bit models are evaluated against the corresponding raw ternary template, while adapted 3-bit models are evaluated against the corresponding adapted ternary template.

All three 3-bit configurations increase their Full alignment to the ternary reference after adaptation. Full alignment changes from 0.208 to 0.780 for 3b-B72, from 0.016 to 0.840 for 3b-P48, and from 0.077 to 0.623 for 3b-R48.

Because the stage-specific comparison changes the ternary reference between the raw and adapted states, we also evaluate both states against the same adapted ternary *Full* reference. Under this fixed reference, *Full* alignment changes from 0.486 to 0.780 for 3b-B72, from 0.480 to 0.840 for 3b-P48, and from 0.486 to 0.623 for 3b-R48. Thus, all three configurations retain positive alignment increases when the reference stage is held fixed. In particular, 3b-P48 reduces teacher-relative T6 from 0.0546 to 0.0504 (7.6% scalar recovery), while its fixed-reference *Full* alignment increases from 0.480 to 0.840.

Across the three configurations and three residual representations, eight of nine raw-to-adapted alignment point estimates increase. The sole exception is the Interaction representation of 3b-R48, which decreases from 0.335 to 0.059. Its point-estimate change is -0.276. Under demonstration-level bootstrap resampling, the corresponding gain has median -0.200 with 95% CI $[-0.479,\ 0.100]$, so the interval includes zero.

The other two 3b-R48 representations increase substantially: Full alignment changes from 0.077 to 0.623, and Nuisance alignment changes from 0.101 to 0.388. Thus, the distinct 3-bit quantizer reproduces the shared-alignment trend broadly across the tested configurations and representations, while the 3b-R48 Interaction result provides a clear representation-level boundary case.

Table 14: Complete cross-quantizer results for the three 3-bit configurations. Full, Interaction, and Nuisance columns report alignment to the corresponding stage-specific ternary reference template at the same evaluation stage. The 3-bit configurations do not contribute to construction of the reference templates.

	Scalar			Full		Interaction		Nuisance	
Config	Raw T6	Final T6	Rec. (%)	Raw	Final	Raw	Final	Raw	Final
3b-B72	0.0951	0.0493	48.1	+0.208	+0.780	+0.172	+0.753	-0.018	+0.876
3b-P48	0.0546	0.0504	7.6	+0.016	+0.840	-0.326	+0.695	-0.497	+0.888
3b-R48	0.0960	0.0729	24.1	+0.077	+0.623	+0.335	+0.059	+0.101	+0.388

I Statistical Audit

Demonstration-level bootstrap. Because multiple test samples originate from the same demonstration, all reported bootstrap uncertainty analyses resample demonstrations rather than treating individual action chunks as independent observations. As described in Appendix A, the retained demonstrations were generated by MimicGen from a shared source pool, and the exported HDF5 does not preserve subtask-level source-demonstration indices. The bootstrap therefore treats each retained generated demonstration as the resampling unit and does not additionally cluster demonstrations by shared source-segment ancestry.

Each bootstrap replicate samples the 19 test demonstrations with replacement. Whenever a demonstration is selected, all eight of its fixed trajectory-phase samples are kept together. Residual maps and all sample-dependent quantities required by the corresponding statistic are then recomputed from the resampled data, including the applicable representation transforms, reference templates, and alignment statistics. For paired contrasts, the same bootstrap demonstration draw is used for both conditions within each replicate.

We use 2,000 bootstrap replicates and report percentile 95% intervals defined by the 2.5th and 97.5th percentiles of the bootstrap distribution. Primary geometry and cross-quantizer analyses use bootstrap seed 20260912, while matched-control, adapted-endpoint, and fixed-reference analyses use seed 20260916.

The bootstrap resamples test demonstrations, not model configurations or optimization seeds. Accordingly, these intervals quantify test-demonstration sampling variability conditional on the frozen configuration plans and fixed trained checkpoints. Optimization-seed variability is evaluated separately in Appendix C.2.

For the primary 13-configuration ternary analysis, the adapted-minus-raw increase in mean pairwise alignment remains positive throughout the demonstration-bootstrap interval for all three residual representations. The median gains are $+0.520$ for *Full* (95% CI $[0.472, 0.568]$), $+0.286$ for *Interaction* ($[0.191, 0.374]$), and $+0.527$ for *Nuisance* ($[0.466, 0.569]$). All 2,000 bootstrap replicates yield a positive aggregate gain for each representation.

For the pooled 3-bit analysis, the bootstrap intervals are above zero for all three representations. For the combined 16-configuration analysis, the raw-to-adapted change in mean pairwise alignment is likewise positive throughout the reported 95% bootstrap intervals. In particular, the Full representation has median increase $+0.517$ with 95% CI $[0.476, 0.556]$, matching the cross-quantizer robustness result reported in the main text.

At the individual-configuration level, 3b-R48 provides the principal representation-level boundary case. Its Full alignment gain has bootstrap median $+0.544$ with 95% CI $[0.393, 0.701]$, while its Nuisance gain has median $+0.345$ with 95% CI $[0.125, 0.546]$. Both intervals remain above zero.

By contrast, the 3b-R48 Interaction gain has bootstrap median -0.200 with 95% CI $[-0.479, 0.100]$. The corresponding point estimate is negative, but the bootstrap interval includes zero. Thus, demonstration-level resampling supports the aggregate cross-quantizer alignment increase and the positive Full and Nuisance 3b-R48 gains, while it does not establish a reliably negative 3b-R48 Interaction effect at the 95% interval level.

Table 15: Demonstration-bootstrap audit of cross-quantizer and combined geometric alignment gains. "Pooled 3-bit gain" is the mean raw-to-adapted alignment change across 3b-B72, 3b-P48, and 3b-R48 relative to the corresponding stage-specific ternary reference templates. "16-config pairwise Δ" is the raw-to-adapted change in mean pairwise alignment over the combined 13 ternary and three 3-bit configurations.

	Pooled 3-bit gain		16-config pairwise Δ	
Representation	Median	95% CI	Median	95% CI
Full	0.643	[0.537, 0.753]	0.517	[0.476, 0.556]
Interaction	0.447	[0.266, 0.615]	0.291	[0.195, 0.384]
Nuisance	0.853	[0.665, 1.016]	0.535	[0.488, 0.574]