# PRESERVING DECISION-RELEVANT HISTORY: QUERY-RANK MEMORY FOR WRITE-BEFORE-QUERY PREDICTION

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

Before a query is known, a fixed-size writer must preserve the history distinctions that determine a later response. We formulate this write-before-query interface as compression of a history-indexed response field, and derive its linear response-operator geometry, capacity–allocation decomposition, and late-query timing gap. Query-Rank Memory (QRM) is the corresponding spectral writer with nested prefixes, estimated by weighted reduced-rank regression. Across six native history conditions, QRM lowers response MSE in five, while controlled, exchange, and known-effect tests connect the same response field to storage choices and decisions.

## 1 INTRODUCTION

A system may have to commit a compact state before it knows which prediction will be requested. This write-before-query interface appears whenever a history is observed first and a later horizon, object, sensor, or action query selects what matters. Figure 1 gives the canonical failure: History A and History B lead to the same current observation, yet the later query *retrieve the blue object* requires opposite actions. A writer that maps them to the same state forces the query-conditioned reader to act on an unresolved collision. The memory must therefore preserve distinctions that are invisible in the current observation but decisive for a later response.

Figure 2 turns this collision into a minimal analytic example. Let the history contain independent coordinates $h = (n, r_1, r_2)$ with variances $(100, 1, 1)$, where a later query asks for $r_1$ or $r_2$ with equal probability. The reconstruction score $S$ prefers the generic direction $n$, while the response

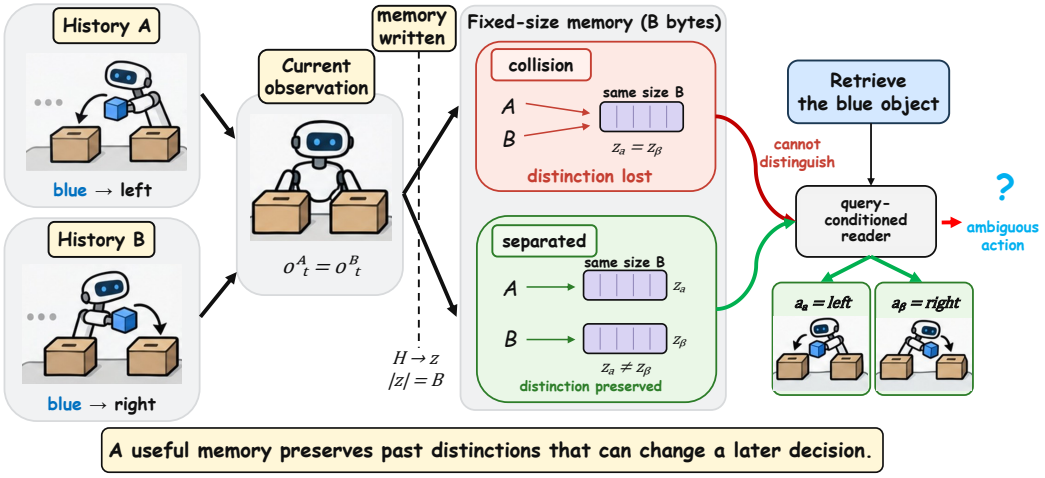



Figure 1: **A useful memory preserves past distinctions that can change a later decision.** History A and History B lead to the same current observation, while the later query *retrieve the blue object* requires opposite actions. A fixed-size writer can collide the two histories, $z_A = z_B$, leaving the query-conditioned reader with one ambiguous state; separated states preserve the distinction and support the two actions.



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.

score $C$ prefers $r_1$ and $r_2$. At fixed width $k = 2$, one valid PCA prefix is $P_{\mathrm{PCA}} = [n, r_1]$; the $r_1/r_2$ eigenvalue tie is immaterial. This prefix has response error $0.5$ even though its capacity loss is zero. Replacing $n$ with $r_2$ keeps the width fixed and removes the allocation loss, giving the response-optimal state $P^* = [r_1, r_2]$ with zero response error. The same memory budget is therefore ranked in opposite ways by history reconstruction and later response quality.

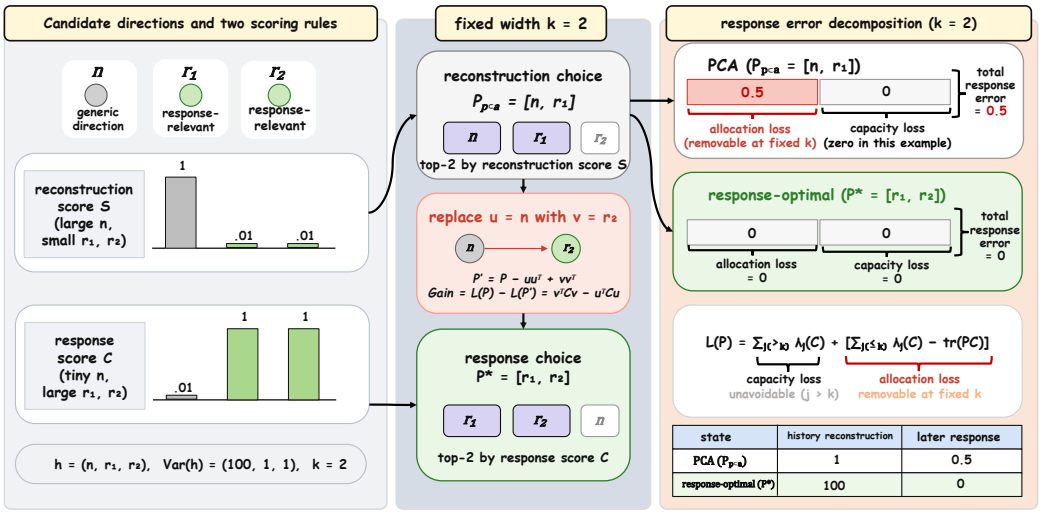



Figure 2: **Reconstruction and response score the same fixed width differently.** In the analytic width-two example, the reconstruction score $S$ favors the generic direction $n$, whereas the response score $C$ favors $r_1$ and $r_2$. One valid reconstruction choice is $P_{\mathrm{PCA}} = [n, r_1]$; the tied $r_1/r_2$ directions have identical reconstruction value. This state has allocation loss $0.5$ and capacity loss $0$; the response-optimal choice $P^* = [r_1, r_2]$ has zero response error. Replacing $n$ with $r_2$ changes only allocation at fixed width, as $L(P) - L(P') = v^\top C v - u^\top C u$.

Prior work provides tools for preserving future-relevant information, including supervised dimensionality reduction, reduced-rank regression, predictive-state representations, and predictive rate–distortion formulations (Barshan et al., 2011; Izenman, 1975; Littman et al., 2001; Boots & Gordon, 2010; Marzen & Crutchfield, 2016). We build on this foundation to study a committed state shared across a declared family of queries. The key quantity is response error at a given width: it reflects both the capacity available to represent the response family and the allocation of that capacity among history directions. Separating these effects quantifies the error required by width $k$, the loss associated with the chosen subspace, and the sharing available across query families.

The common object behind these examples is the complete response field $\Phi(H)$: a single history determines a response function over all declared queries. This view separates the information available at write time from the query used at read time. After whitening history, a Hilbert-space response operator maps the stored history features to this field. Its Gram operator yields an exact capacity–allocation decomposition, a late-query width law, and the exchange identity used in the experiments.

Query-Rank Memory (QRM) is the empirical spectral realization of this response operator. Weighted reduced-rank regression estimates the operator, while one query-independent writer provides nested prefixes for every width. The resulting write-before-query state contract makes the consequences of a storage choice directly measurable through a training-only exchange score and held-out responses.

The experiments examine this account at increasing levels of realism. Controlled rotations isolate allocation from capacity; recorded PH/MH histories test the response and reconstruction ranking under declared query families; held-out exchanges assess training-side gain predictions; and a known effect dynamical model measures the connection between response quality and action ranking. Across six native history task conditions, QRM lowers response MSE relative to PCA in five while increasing reconstruction MSE in all six. Together, these results examine response-aligned allocation as a state complexity principle for the specified offline write-before-query setting.

Our contributions are:

- We formulate write-before-query memory as compression of a history-indexed response field and characterize response-sufficient states through response equivalence.
- We derive a response-operator theory that unifies state width, capacity loss, allocation loss, query-family sharing, and the timing gap between early-query and late-query memory.
- We instantiate the theory with QRM, a query-independent spectral writer with nested prefixes, and evaluate its allocation and decision consequences with controlled, recorded-history, held-out exchange, and known-effect experiments.

## 2 Related Work

Weighted reduced-rank regression, PLS, and predictive-state identification provide established estimators for predictive low-rank subspaces (Izenman, 1975; Boots & Gordon, 2010). Under matched features, response weights, and a linear reader, QRM recovers the corresponding population subspace. Our contribution is the committed-state interface: the writer cannot observe the later query, so one finite prefix family must serve a declared query set. We make its cost explicit through $\Gamma_k(P)$, derive a fixed-width exchange identity, and evaluate a training-scored exchange on untouched episodes. Nested dropout supplies ordered prefixes with a different ordering objective (Rippel et al., 2014); QRM supplies the response-rank state contract and its testable allocation protocol.

Predictive-state representations use future tests; approximate information states and predictive rate–distortion provide broader views of useful compressed history, while decoder side-information coding studies a related separation between stored information and later side information (Littman et al., 2001; Singh et al., 2004; Subramanian et al., 2022; Marzen & Crutchfield, 2016; Wyner & Ziv, 1976). LoRANN studies query-weighted reduced-rank compression for approximate search, and task-aware answer preservation studies query-family risk and budget decisions for compressed audio (Jääsaari et al., 2024; Ivry, 2026). Robot memories such as MemoryVLA and MemER address history-dependent control; Compressor-VLA and Denoised MDPs learn task-relevant compressed states (Shi et al., 2025; Sridhar et al., 2025; Gao et al., 2025; Wang et al., 2022). Our focus is the response geometry of one state written before the query is available, and how that geometry accounts for capacity, allocation, and query-family sharing.

## 3 Problem Formulation

A history $H$ ends at time $\tau$. A writer stores $z_k = g_k(H) \in \mathbb{R}^k$. A later query $d \in \mathcal{D}$ selects an output through $\widehat{F}_H(d) = a(z_k, d)$. No history is passed to $a$, and reading does not rewrite the stored state. Queries select prediction horizons and response families in the recorded-history experiments; the synthetic dynamical experiment also ranks candidate actions under known effects. Each experiment specifies the prediction target and its units.

Let $\mu$ denote the history law and $\nu$ a declared probability measure over queries. A response is the conditional mean $F_H(d) = \mathbb{E}[r \mid H, d]$. Its complete response field and the memory risk are

$$\Phi(H) : d \mapsto F_H(d), \qquad \mathcal{Y} = L_2(\mathcal{D}, \nu; \mathbb{R}^m),$$
$$\mathcal{R}(g, a) = \mathbb{E}_{H\sim\mu} \|\Phi(H) - a(g(H), \cdot)\|^2_{\mathcal{Y}}. \tag{1}$$

We assume a separable $\mathcal{Y}$ and a measurable, square-integrable response field identified on the support of $\mu \otimes \nu$. Two histories are response-equivalent when $\Phi(H_1) = \Phi(H_2)$ in $\mathcal{Y}$, meaning equality for $\nu$-almost every query. A fixed metric integrates all declared horizons and output components; history-dependent query selection defines a different joint risk. History-independent responses enter a shared reader intercept. Output units and weights are fixed before fitting, and noisy targets add irreducible label noise.

If a new observation $e$ also enters a downstream predictor $p(e, \widehat{F})$, response error and final target error are separate metrics. We establish a target-risk relation for the specified Gaussian residual model and report both quantities.

Stored width counts per-history scalars. Storage precision, shared parameters, and writer/reader computation are separate resources.

# 4 RESPONSE-RANK THEORY

## 4.1 RESPONSE SUFFICIENCY AND LINEAR CAPACITY

**Theorem 1** (Response-sufficient state)**.** *For any square-integrable response field $\Phi$ and writer $g$, the best query-conditioned reader satisfies*

$$\inf_a \mathcal{R}(g, a) = \mathbb{E}_H \|\Phi(H) - \mathbb{E}[\Phi(H) \mid g(H)]\|_{\mathcal{Y}}^2.$$

*The infimum is zero if and only if $\Phi(H)$ is a measurable function of $g(H)$ up to $\mathbb{P}$-null sets.*

This conditional-projection characterization connects the state collision in Figure 1 to response equivalence. It covers nonlinear responses and continuous queries. To quantify width, let $\widetilde{h} = h - \mu_h$ be the centered history feature with covariance $S \succ 0$, and set $x = S^{-1/2}\widetilde{h}$. After removing the query-dependent intercept, assume $\Phi(H) = Ax$ for a bounded operator $A : \mathbb{R}^p \to \mathcal{Y}$. Its Gram matrix is $C = A^*A$. A descriptor representation with coefficient operator $T$ gives

$$(Ax)(d) = d^\top T x, \quad G = \mathbb{E}_{d\sim\nu}[dd^\top], \quad C = T^\top G T, \tag{2}$$

Vector responses stack output modes. The operator formulation covers continuous query families under the fixed product metric in Eq. (1); the actual history–query joint law is treated in Appendix E. Store $z = V^\top x$, where $V^\top V = I_k$ and $P = VV^\top$.

**Theorem 2** (Capacity and allocation loss of a width-$k$ state)**.** *For $\Phi(H) = Ax$, the minimum response error among readers linear in the stored coordinates is*

$$\begin{aligned}\mathcal{L}(P) &= \|A(I-P)\|_{\mathrm{HS}}^2 = \mathrm{tr}((I-P)C) \\ &= \underbrace{\sum_{j>k} \lambda_j(C)}_{\textit{capacity loss}} + \underbrace{\sum_{j\le k} \lambda_j(C) - \mathrm{tr}(PC)}_{\Gamma_k(P)\textit{: allocation loss}}.\end{aligned} \tag{3}$$

*The allocation loss is nonnegative and vanishes for a leading $k$-dimensional eigenspace of $C$. If $x$ is Gaussian, the same minimum holds over arbitrary measurable readers.*

Projection and the eigenvalue variational principle prove the identity (Appendix B). It separates the optimal width-$k$ error from the additional loss of a chosen subspace, as in Figure 2.

**Proposition 3** (Late-query timing gap)**.** *For scalar responses $(Ax)(d) = \langle a_d, x\rangle$, a writer that receives $d$ before writing attains zero error by storing that one scalar. A query-independent linear writer with an optimal linear reader has exact width $\mathrm{rank}(A)$ under Eq. (1). Its smallest width at tolerance $\epsilon \ge 0$ is*

$$k_{\mathrm{late}}(\epsilon) = \min\left\{ k : \sum_{j>k} \lambda_j(A^*A) \le \epsilon \right\}.$$

The result also permits arbitrary readers for Gaussian $x$. Exact risk concerns $\nu$-almost every query; finite query families with positive weights obtain coverage of every member. Continuous encoders require the separate support condition in Appendix C.

The best linear history reconstruction from the same state has error $\mathrm{tr}((I-P)S)$. PCA chooses directions using $S$, whereas response-optimal storage uses $C$. The comparison uses the same whitened coordinates and feature units so that $\Gamma_k$ measures a storage-choice penalty.

**Corollary 4** (Response overhead of reconstruction-ordered storage)**.** *For tolerance $\epsilon$, let $k_{\mathrm{resp}}(\epsilon)$ be the smallest width whose response spectral tail in Eq. (3) is at most $\epsilon$, and let $k_{\mathrm{PCA}}(\epsilon)$ be the smallest PCA width whose response loss is at most $\epsilon$. Then $k_{\mathrm{PCA}}(\epsilon) \ge k_{\mathrm{resp}}(\epsilon)$. The gap can grow with the number of high-variance nuisance directions: if $L$ nuisance directions precede $r$ response directions, exact response coverage requires $L+r$ PCA coordinates but only $r$ response coordinates.*

The corollary quantifies the cost of using input variance as a proxy for a later response field.

### 4.2 SHARED RESPONSE RANK

Let $\mathcal{V}_A$ and $\mathcal{V}_B$ be the spans of centered history-dependent responses for two query families. Their joint response rank obeys

$$r_{A\cup B} = r_A + r_B - \dim(\mathcal{V}_A \cap \mathcal{V}_B). \tag{4}$$

The intersection contains responses that both families can read from the same stored information. Separate caches may store these directions twice. Equation (4) therefore predicts when one shared state should be more efficient than two independent banks at the same total width, and why zero-overlap families provide a necessary negative control.

More generally, the response rank is the dimension of the span of all centered scalar response functions, including components of vector outputs. Storing a basis of these functions constructs a sufficient response state. A lower bound for arbitrary continuous encoders additionally requires a continuous section over an open response set; the precise condition and its failure mode are given in Appendix C.

### 4.3 IDENTIFIED-QUERY OMISSION

For an identified response operator $T$, a query $d$ requires the history direction $T^\top d$. Its omitted response energy is

$$\rho(d) = \|(I - P)T^\top d\|^2. \tag{5}$$

Thus a query can be unsupported by the stored state even when its descriptor is close to a previously seen query. Equation (5) applies to identified directions. When training queries span a strict subspace, the operator outside that subspace requires an additional assumption before an omission score is assigned. This separates query shift from ordinary allocation loss.

For a complete response table, the same projection identity also allows nonlinear target relationships. Let $Y \in \mathbb{R}^J$ collect its scalar targets, fix the history law $\mu$ and history-independent column weights $W \succeq 0$, and center and whiten history under $\mu$. Define $B = \mathbb{E}_\mu[x(Y - \mathbb{E}_\mu Y)^\top]$ and $C_{\mathrm{tab}} = BWB^\top$. The best linear reader then has target risk

$$L_{\mathrm{tab}}(P) = R_{\mathrm{lin}} + \sum_{j>k} \lambda_j(C_{\mathrm{tab}}) + \Gamma_k(P; C_{\mathrm{tab}}). \tag{6}$$

where $R_{\mathrm{lin}}$ is the best linear full-history risk and $\Gamma_k(P; C_{\mathrm{tab}})$ uses the definition in Eq. (3) with $C_{\mathrm{tab}}$. This separates noise and full-history approximation error from storage losses; Appendix B.6 gives the derivation.

**Proposition 5** (Fixed-width response exchange)**.** *Replacing a retained unit direction $u$ with an omitted unit direction $v$ preserves the stored width and changes only the allocation term. In the response-table setting, with $P' = P - uu^\top + vv^\top$, the ideal gain is*

$$L_{\mathrm{tab}}(P) - L_{\mathrm{tab}}(P') = v^\top C_{\mathrm{tab}} v - u^\top C_{\mathrm{tab}} u. \tag{7}$$

*The exchange identity also holds for Eq. (3) using $C$. Training scores predict this gain under the declared sample and column weights. Independently measured gains additionally reflect reader fitting and moment estimation; a fitted full-history baseline includes estimation error as well.*

**Lemma 6** (Decision preservation under a response margin)**.** *Let $\mathcal{A}$ be a finite action set and let $u(a, r)$ be an action utility that is $L$-Lipschitz in the response $r$. If the true response is $r_a$ and a stored state produces $\widehat{r}_a$ with $\|\widehat{r}_a - r_a\| \le \epsilon$ for every $a$, then any optimal action with margin*

$$m = \min_{a \neq a^\star} \{u(a^\star, r_{a^\star}) - u(a, r_a)\} > 2L\epsilon$$

*remains optimal when utilities are evaluated with $\widehat{r}_a$.*

The lemma gives a sufficient condition for preserving an action ordering under a uniform response bound and an action margin. The synthetic experiment instantiates the response-to-utility map in a known-effect model and reports response MSE, regret, and optimal-action rate as consequences of the same response state. The margin audit evaluates the uniform condition directly on the known synthetic model.

# 5 QRM METHOD

## 5.1 ONE ORDERED STATE FOR SEVERAL WIDTHS

Training provides history–query–target triples. We center and whiten history on the training split, fit a query-conditioned intercept and a ridge-regularized bilinear response operator $\widehat{T}$, and form

$$\widehat{K} = \widehat{G}^{1/2}\widehat{T}, \qquad \widehat{K} = U\Sigma V^\top, \qquad z_k(H) = V_k^\top \widehat{S}^{-1/2}(h(H) - \widehat{\mu}). \tag{8}$$

The query second moment $\widehat{G}$ specifies the operating weights used to measure response quality. It may differ from a broader training-query distribution used to identify the operator. Vector targets are handled by stacking query–output modes before the SVD. On a complete response table, the estimator is weighted multiresponse reduced-rank regression; with matching history whitening and future-response weights, predictive-state subspace identification yields the same right singular subspace (Izenman, 1975; Boots & Gordon, 2010). QRM uses this established spectral estimator as the computable realization of the response-state contract: its right singular vectors are ordered once, and each width is a prefix of one query-independent writer. Training-side exchange scores and held-out responses then evaluate the allocation implied by this ordering.

The writer receives history alone. After writing, a reader combines the stored prefix with $d$. All widths use prefixes of the same ordered writer. In matched comparisons, we refit a reader of the same functional form for every method and width. The nested property concerns the writer, while readers are refit at each width.

Training-side $\widehat{\Gamma}_k(P)$ estimates the cost of a proposed subspace, and held-out responses test its prediction. The fixed-coordinate perturbation bound is given in Appendix D. Width selection uses separate calibration episodes and a prespecified response-error failure criterion; Appendix F gives the simultaneous binomial bound and its assumptions.

In fixed common coordinates, an operator estimate with $\|\widehat{C} - C\|_{\mathrm{op}} \le \delta$ obeys

$$|\widehat{\Gamma}_k(P) - \Gamma_k(P)| \le 2k\delta, \qquad \Gamma_k(\widehat{P}_k) \le 2k\delta. \tag{9}$$

where $\widehat{P}_k$ is the leading estimated response subspace. If $\gamma_k = \lambda_k(C) - \lambda_{k+1}(C) > 0$, the same event gives the Davis–Kahan bound

$$\|\sin\Theta(\widehat{V}_k, V_k^\star)\|_{\mathrm{op}} \le \frac{2\delta}{\gamma_k}.$$

Thus any finite-sample concentration result of the form $\|\widehat{C} - C\|_{\mathrm{op}} \le \Delta_n(\eta)$ with probability $1-\eta$ immediately controls both allocation excess and subspace rotation. Whitening and fitted-reader uncertainty are handled separately.

## 5.2 WIDTH SELECTION WITH AN EXPLICIT FAILURE CRITERION

A response spectrum supplies a population allocation rule, while a finite deployment also needs evidence for a width decision. We freeze candidate writers and readers using training and validation data, then reserve independent calibration episodes. For width $k$ and query family $j$, an episode is a failure when its mean response loss exceeds a prespecified tolerance. With $b_{kj}$ failures in $n$ independent episodes, the one-sided exact binomial upper bound is

$$U_{kj} = \mathrm{Beta}^{-1}\left(1 - \frac{\alpha}{KJ}; b_{kj} + 1, n - b_{kj}\right). \tag{10}$$

with $U_{kj} = 1$ when every episode fails. We select the smallest width satisfying $U_{kj} \le \beta$ for every declared family. The union bound gives simultaneous coverage $1 - \alpha$ conditional on the frozen models and the stated episode sampling assumptions. A failed certificate records the protocol outcome and leaves adequacy of smaller states unresolved.

# 6 EXPERIMENTS

## 6.1 PROTOCOL AND MEASUREMENTS

The primary controlled generator has 24 history coordinates and rotates four response directions from high- to low-variance history directions while response energy and noise remain fixed. Five

| Memory | History SSE | Response MSE | Target MSE | Bayes risk |
|---|---|---|---|---|
| PCA | 59.9666 | 1.0218 | 0.6707 | 0.6667 |
| Random | 99.4744 | 0.9800 | 0.6661 | 0.6614 |
| Coefficient SVD | 186.1271 | $2.33 \times 10^{-5}$ | 0.4966 | 0.5000 |
| QRM | 186.1271 | $2.33 \times 10^{-5}$ | 0.4966 | 0.5000 |

Table 1: Equal-width controlled comparison: response rank four, stored width four, and orthogonal reconstruction/response directions. All values are measured means over five seeds except the separately labeled population Bayes-risk reference. History SSE measures feature reconstruction.

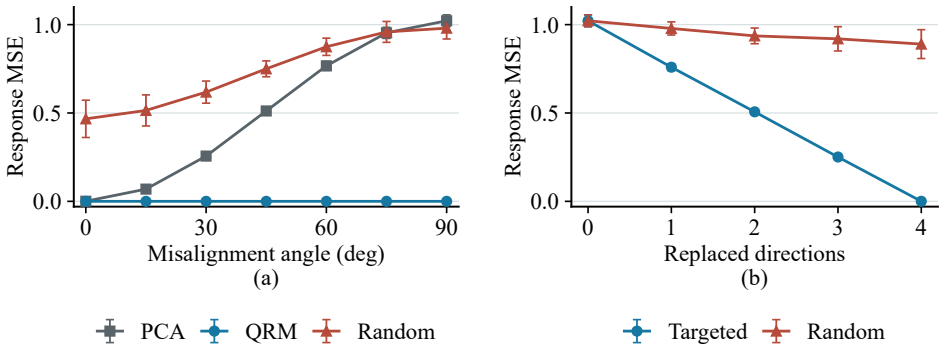



Figure 3: **Controlled rotation and direction exchange expose allocation loss.** The left panel rotates response demand away from high-variance history directions while response energy is fixed. The right panel exchanges response-ranked directions at fixed width four. Curves and intervals are computed from the same five generated splits used in Table 1.

independent seeds each supply 12,000 training, 2,000 validation, and 5,000 test triples. PCA, random projection, raw-coordinate coefficient SVD, and QRM use equal stored widths and the same query-conditioned linear reader. We measure reconstruction SSE, response MSE, and final-target MSE under the specified Gaussian residual model. Appendix H gives the generator, neural controls, and split details; Appendix G derives the final-target metric.

### 6.2 Controlled Response Allocation

Table 1 exhibits the ranking conflict under controlled rotation of response directions. PCA has lower reconstruction error but higher response error. QRM and raw-coordinate coefficient SVD both cover the full response at width four. The decisive test is the fixed-width intervention: as response demand rotates away from high-variance directions, PCA response loss rises, and replacing low-response-contribution directions with omitted high-contribution directions lowers held-out response error while increasing reconstruction error. A random replacement has a smaller effect, and restoring the original projection restores its original result. This is the experiment-level signature of $\Gamma_k(P)$ at fixed capacity. Figure 3 shows the angle scan and the fixed-width exchange with intervals across five generated splits.

The same response geometry predicts two further tests without changing the writer interface. Shared query families should benefit from a common state to the extent that their response spans intersect, while an identified query should incur the omitted energy in Eq. (5). In Figure 4, the learned response tail decreases with width and the observed query-angle curve follows the omitted-energy reference; the corresponding shared-family and repeated-read values are in the appendix.

Supporting controls examine shared query families, response-rank width selection, neural bottlenecks, storage precision, and history-dependent queries. These studies use the same declared response metric and ordered-state protocol; their full protocols and values are reported in Appendix H, Appendix F, and Appendix E.

### 6.3 Recorded-History Validation

We evaluate offline multi-horizon prediction from native low-dimensional robomimic demonstrations. The frozen proficient-human (PH) study covers Lift, Can, and Square, each with native and

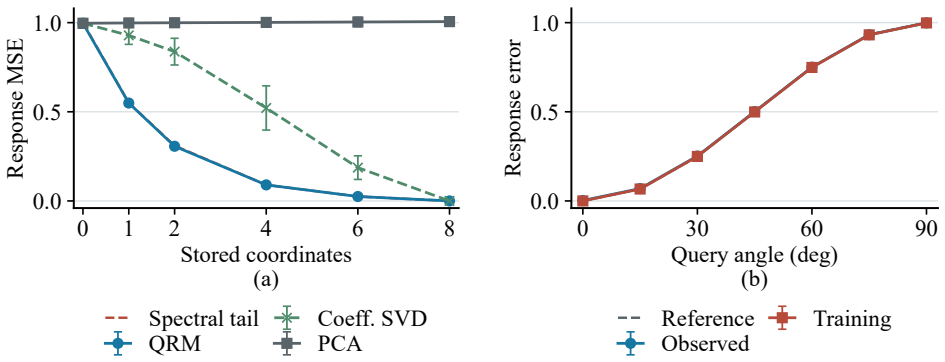



Figure 4: **Response spectra determine allocation under nonuniform demand.** The left panel shows the error tail as stored width grows; the right panel measures unresolved response as the query rotates away from the retained family. Curves use learned models and independent test responses.

| | | Response MSE | | | | | Recon. MSE |
|---|---|---|---|---|---|---|---|
| Data | Task | PCA | PLS | QRM | Current | GRU | PCA/QRM |
| PH | Can | 0.3198 | 0.2535 | 0.1479 | 0.1530 | 0.1443 | 0.2069/0.4306 |
| PH | Square | 0.2961 | 0.2253 | 0.1251 | 0.1316 | 0.1250 | 0.2467/0.4268 |
| PH | Lift | 0.4244 | 0.3335 | 0.7566 | 0.2873 | 0.3488 | 0.3727/0.7139 |
| MH | Can | 0.3499 | 0.2447 | 0.1130 | 0.1279 | 0.1016 | 0.2793/0.4789 |
| MH | Square | 0.3489 | 0.2389 | 0.0994 | 0.1160 | 0.0891 | 0.2851/0.4908 |
| MH | Lift | 0.6042 | 0.4017 | 0.3030 | 0.2655† | 0.2759 | 0.3555/0.7088 |

Table 2: Width-eight linear comparison with 64 B persistent states. PCA, PLS, QRM, and the compressed current-frame baseline use the same linear reader form and store eight float64 coordinates; GRU is a learned-reader reference with eight float32 coordinates (32 B). The audited QRM/PCA states also preserve their ranking after casting to 32 B; the matched persistent-state audit is shown in Figure 5a. Reconstruction is the affine-probe error. †The MH Lift current-frame row is a diagnostic; Appendix J.12 contains the complete frozen PH grid.

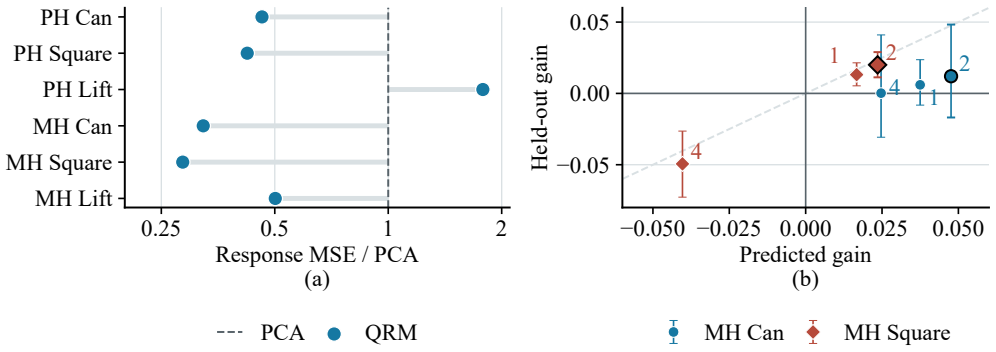



Figure 5: **Recorded histories connect response allocation to direction-exchange gains.** (a) Width-eight QRM/PCA response-MSE ratios at 32 B; ratios below one favor QRM. (b) Training-predicted versus held-out gains for fixed-width exchanges, with paired 95% bootstrap intervals over 45 episodes per task. Numbers denote exchanged directions; outlined points mark the two-direction exchanges selected using training data.

constructed recent-observation-dropout conditions. Independent multi-human (MH) Can and Square runs use six demonstrators per task; a separately fixed MH Lift run tests writer stability. In this protocol, each writer stores its state before the future offset and response family are selected. Response errors are standardized MSE with equal episode and query-family weights. Reconstruction uses the same affine probe for linear methods. These data test response-aligned allocation on recorded histories.

**Response allocation.** Table 2 reports the width-eight linear comparison at 64 B; Figure 5a reports QRM/PCA at a matched 32 B state budget. GRU is a learned-reader reference. Across six PH/MH conditions, QRM lowers response MSE in five and raises reconstruction MSE in all six, exhibiting the response–reconstruction trade-off quantified by the allocation decomposition. PH Lift marks an empirical boundary: at width eight, QRM response MSE is 0.7566 versus 0.4244 for PCA. Across the full frozen PH grid, the same point-estimate pattern appears in 23 of 30 correlated condition–width settings. On MH Can and Square, response and reconstruction rankings agree across all six observed demonstrators per task (Appendix J.9). Relative to a same-width current-frame state, QRM gains 0.0149 [0.0095, 0.0218] on Can and 0.0166 [0.0138, 0.0196] on Square. Appendix J.8 gives the full comparison and neural controls.

**Held-out direction exchange.** A training score selects two replacement directions for both tasks. All six candidate exchange estimates have matching training and held-out signs. The selected held-out gains are 0.0120 [-0.0169, 0.0483] on Can and 0.0200 [0.0113, 0.0290] on Square. Cluster sensitivity preserves the positive QRM/PCA gaps; the selected exchange remains uncertain on Can and positive on Square for the six observed demonstrators.

The complete PH intervention grid and MH stability analyses are reported in Appendices J.5–J.7; they provide the uncertainty measurements behind the exchange result in Eq. (7).

### 6.4 Decision validation with known action effects

Under known constant-acceleration dynamics, one width-six state supports future-position queries and ranking nine actions with known effects. QRM and coefficient-SVD coincide to displayed precision, isolating response-aligned storage from the solver: QRM yields MSE 0.0050 and action regret 0.0096, versus 4.5139 and 9.6009 for PCA; optimal-action rates are 0.9645 and 0.2527. An oracle-side margin audit finds the sufficient condition on 56.23% of QRM width-six test cases, with no action disagreements among those cases. This offline test evaluates the lemma's mechanism under the specified response-to-utility map. Appendix H.5 gives the width sweep and protocol.

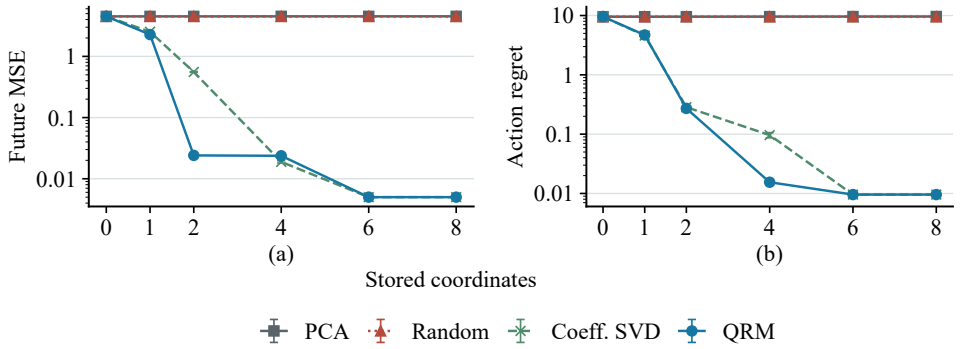



Figure 6: **Response-aligned memory improves one-step action ranking in a specified dynamical model.** The same width-six state supports future-position prediction and action ranking under known additive effects. Lower is better; coefficient-SVD and QRM coincide to displayed precision. The margin audit is an oracle-side finite-sample diagnostic.

## 7 Conclusion

We separate width-limited response loss from allocation loss and give a training-scored fixed-width exchange rule. Controlled and recorded histories show response gains alongside worse reconstruction, while PH Lift marks the empirical boundary set by the fitted response family. Held-out exchange estimates are positive on MH Can and MH Square, with a positive interval on Square and a wider interval on Can. In the known-effect model, the oracle-side margin condition holds on 56.23% of QRM width-six cases with no action disagreements among them. Nonlinear queries, query shift, finite precision, and closed-loop control remain open extensions.

## Reproducibility Statement

The accompanying code fixes the data generators, random seeds, data split rules, candidate memory widths, and evaluation metric definitions. Tables from the controlled experiments and the complete frozen PH experiment table are generated from CSV outputs. The MH experiment protocols, data-source hashes, model predictions, and paired audit results are saved with the new summary results. Appendix H describes all experiments that were executed, and Appendix J specifies the external-data experiment protocol.

## Statement on use of AI tools

Generative AI tools were used to assist with examining mathematical derivations, searching and checking references, and drafting and revising the manuscript. All AI-assisted outputs were reviewed and verified by the authors. The authors take responsibility for the final content, claims, results, and artifacts in this work.

## References


Elnaz Barshan, Ali Ghodsi, Zohreh Azimifar, and Mansoor Zolghadri Jahromi. Supervised principal component analysis: Visualization, classification and regression on subspaces and submanifolds. *Pattern Recognition*, 44(7):1357–1371, 2011. doi: 10.1016/j.patcog.2010.12.015.

Byron Boots and Geoffrey J. Gordon. Predictive state temporal difference learning. In *Advances in Neural Information Processing Systems*, volume 23, pp. 271–279, 2010. URL https://proceedings.neurips.cc/paper/2010/hash/3473decccb0509fb264818a7512a8b9b-Abstract.html.

Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11621–11631, 2020. doi: 10.1109/CVPR42600.2020.01164.

C. J. Clopper and E. S. Pearson. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, 26(4):404–413, 1934. doi: 10.1093/biomet/26.4.404.

Juntao Gao, Feiyang Ye, Jing Zhang, and Wenjing Qian. Compressor-VLA: Instruction-guided visual token compression for efficient robotic manipulation. *arXiv preprint arXiv:2511.18950*, 2025. URL https://arxiv.org/abs/2511.18950.

Amir Ivry. Task-aware answer preservation under audio compression for large audio language models. *arXiv preprint arXiv:2605.06631*, 2026. URL https://arxiv.org/abs/2605.06631.

Alan Julian Izenman. Reduced-rank regression for the multivariate linear model. *Journal of Multivariate Analysis*, 5(2):248–264, 1975. doi: 10.1016/0047-259X(75)90042-1.

Elias Jääsaari, Ville Hyvönen, and Teemu Roos. LoRANN: Low-rank matrix factorization for approximate nearest neighbor search. *arXiv preprint arXiv:2410.18926*, 2024. URL https://arxiv.org/abs/2410.18926.

Michael L. Littman, Richard S. Sutton, and Satinder Singh. Predictive representations of state. In *Advances in Neural Information Processing Systems*, volume 14, pp. 1555–1561, 2001. URL https://proceedings.neurips.cc/paper/2001/hash/1e4d36177d71bbb3558e43af9577d70e-Abstract.html.

Ajay Mandlekar, Danfei Xu, Josiah Wong, Soroush Nasiriany, Chen Wang, Rohun Kulkarni, Li Fei-Fei, Silvio Savarese, Yuke Zhu, and Roberto Martín-Martín. What matters in learning from offline human demonstrations for robot manipulation. In *Proceedings of the 5th Conference on Robot Learning*, volume 164, pp. 1678–1690, 2022. URL https://proceedings.mlr.press/v164/mandlekar22a.html.

Sarah E. Marzen and James P. Crutchfield. Predictive rate-distortion for infinite-order markov processes. *Journal of Statistical Physics*, 163:1312–1338, 2016. doi: 10.1007/s10955-016-1520-1.

James R. Munkres. *Topology*. Prentice Hall, 2 edition, 2000.

Oren Rippel, Michael A. Gelbart, and Ryan P. Adams. Learning ordered representations with nested dropout. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32, pp. 1746–1754, 2014. URL https://proceedings.mlr.press/v32/rippel14.html.

Hao Shi, Bin Xie, Yingfei Liu, Lin Sun, Fengrong Liu, Tiancai Wang, Erjin Zhou, Haoqiang Fan, Xiangyu Zhang, and Gao Huang. MemoryVLA: Perceptual-cognitive memory in vision-language-action models for robotic manipulation. *arXiv preprint arXiv:2508.19236*, 2025. URL https://arxiv.org/abs/2508.19236.

Satinder Singh, Michael R. James, and Matthew R. Rudary. Predictive state representations: A new theory for modeling dynamical systems. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, pp. 512–518. AUAI Press, 2004. doi: 10.5555/1036843.1036905. URL https://mlanthology.org/uai/2004/singh2004uai-predictive/.

Ajay Sridhar, Jennifer Pan, Satvik Sharma, and Chelsea Finn. MemER: Scaling up memory for robot control via experience retrieval. *arXiv preprint arXiv:2510.20328*, 2025. URL https://arxiv.org/abs/2510.20328.

Jayakumar Subramanian, Amit Sinha, Raihan Seraj, and Aditya Mahajan. Approximate information state for approximate planning and reinforcement learning in partially observed systems. *Journal of Machine Learning Research*, 23(12):1–83, 2022. URL https://www.jmlr.org/papers/v23/20-1165.html.

Tongzhou Wang, Simon S. Du, Antonio Torralba, Phillip Isola, Amy Zhang, and Yuandong Tian. Denoised MDPs: Learning world models better than the world itself. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162, pp. 22591–22612. PMLR, 2022. URL https://proceedings.mlr.press/v162/wang22c.html.

Aaron D. Wyner and Jacob Ziv. The rate-distortion function for source coding with side information at the decoder. *IEEE Transactions on Information Theory*, 22(1):1–10, 1976. doi: 10.1109/TIT.1976.1055508.

# A Notation and information access

$H$ is an observed history, $h(H) \in \mathbb{R}^p$ is the feature supplied to the writer, and $z_k(H) \in \mathbb{R}^k$ is the per-history stored state. At deployment, the reader receives a query $d$ after $z_k$ is fixed. The response $F_H(d)$ is the conditional mean of a declared target, as in Eq. (1). Its scalar components are indexed by both query and output component. A query-dependent intercept stores history-invariant effects and consumes no per-history coordinate.

The full learning process may observe query labels. The temporal restriction governs each deployed write, while training may use the complete set of query labels. Supervised comparison methods receive the same training targets and may learn a query-conditioned reader. A method supplied with the actual deployment query before writing is labeled an early-query reference because its information access differs.

All exact covariance statements use a common feature space and metric. $S$ is the history covariance; $x = S^{-1/2}(h - \mu)$ denotes whitening on its supported subspace. If $S$ is singular, the analysis first restricts to that support. $T$ maps whitened history to descriptor-indexed responses. $G$ is the query second-moment matrix, with centering determined by the declared risk. Query coordinates with positive probability define the weighted risk. Covering all declared queries additionally requires positive weighting, or an explicit coverage requirement, on all relevant descriptor directions.

A fixed response map and an optimal state can be approximated by a finite-data model with estimation error. The proofs distinguish the population response operator, the fitted operator, the writer subspace, and the fitted reader. The code reports population references for synthetic generators where these objects are known.

# B Projection error, allocation, and reconstruction

## B.1 Proof of the response-sufficient state theorem

Assume that the history and state spaces are standard Borel spaces and that $\mathcal{Y}$ is a separable Hilbert space. If $\Phi(H)$ is Bochner square-integrable, a regular conditional expectation $\mathbb{E}[\Phi(H) \mid g(H)]$ exists as a $\mathcal{Y}$-valued random element (defined up to a $\mathbb{P}$-null set). For a fixed writer $g$, the reader is an element of the Hilbert space $\mathcal{Y}$ for every stored state. The conditional expectation is the orthogonal projection of $\Phi(H)$ onto the closed subspace of $g(H)$-measurable square-integrable $\mathcal{Y}$-valued functions. Therefore the Pythagorean identity gives

$$\mathbb{E}\|\Phi - a(g(H), \cdot)\|_{\mathcal{Y}}^2 = \mathbb{E}\|\Phi - \mathbb{E}[\Phi \mid g(H)]\|_{\mathcal{Y}}^2 + \mathbb{E}\|\mathbb{E}[\Phi \mid g(H)] - a(g(H), \cdot)\|_{\mathcal{Y}}^2.$$

The first term is attained by the conditional-mean reader, with equality understood $\mathbb{P}$-almost surely as an element of $\mathcal{Y}$ (equivalently, $\nu$-almost everywhere in the query coordinate). It is zero exactly when $\Phi$ is $g(H)$-measurable up to $\mathbb{P}$-null sets, which proves Theorem 1. The separability assumption is automatic for the finite-query and finite-output experiments; it also covers the usual $L_2(\mathcal{D}, \nu; \mathbb{R}^m)$ response field.

## B.2 Proof of Theorem 2

The risk in Theorem 2 is the product-measure risk $\mathbb{E}_{H\sim\mu} \int_{\mathcal{D}} \|\cdot\|^2 \, \nu(\mathrm{d}d)$: $\nu$ is fixed before deployment and is independent of history. Let $P = VV^\top$ with $V^\top V = I_k$. The response decomposes into

$$d^\top T x = d^\top T P x + d^\top T (I - P) x. \tag{11}$$

The first term is computable from $z = V^\top x$. Since $x$ is white, $(I - P)x$ is uncorrelated with $V^\top x$. Pointwise in $d$, this makes $d^\top T V z$ the least-squares reader among functions linear in $z$; integrating over the fixed measure $\nu$ gives the same result for the product risk. Its residual loss is

$$\mathcal{L}(P) = \mathbb{E}\left[d^\top T (I - P) x x^\top (I - P) T^\top d\right] \tag{12}$$

$$= \mathrm{tr}\left(GT(I - P)T^\top\right) = \mathrm{tr}((I - P)C). \tag{13}$$

For Gaussian $x$, the retained and omitted projections are independent. The omitted conditional mean is then zero given $z$ and every fixed query $d$, so the same reader is optimal over all measurable readers

of $(z, d)$. With a general non-Gaussian $x$, whiteness gives uncorrelated projections, while independence requires the Gaussian condition; a nonlinear reader may infer part of an omitted projection. The general-distribution result therefore concerns the linear-reader class.

For a positive semidefinite $C$, the variational characterization of the sum of the largest eigenvalues gives

$$\max_{P=P^\top=P^2,\,\mathrm{tr}(P)=k} \mathrm{tr}(PC) = \sum_{j=1}^{k} \lambda_j(C). \tag{14}$$

Subtract and add this maximum to $\mathrm{tr}(C) - \mathrm{tr}(PC)$ to obtain Eq. (3). This also proves nonnegativity of $\Gamma_k(P)$. Tied eigenvalues can make the optimal subspace nonunique while preserving its optimal loss.

For a finite query set with descriptor vectors collected in a matrix $D$ and weights $w_j$ summing to one, let $G = \sum_j w_j d_j d_j^\top$. The map $A$ is represented in the weighted query coordinates by $G^{1/2}T$, so $A^*A = T^\top G T$. For a continuous query with square-integrable descriptors, the same identity is the operator equality obtained by integrating $T^\top dd^\top T$ against $\nu$. Thus the matrix used by QRM is the finite-coordinate Gram representation of the Hilbert-space response operator.

### B.3 Proof of the late-query timing gap

For an early query $d$, the writer can store $z_d = \langle a_d, x\rangle$, and the reader returns $z_d$; this is a one-coordinate representation with zero error for that query. For a late query, a linear state $z = V^\top x$ and a linear reader induce the projection $P = VV^\top$. The response operator factors through this state without error exactly when $A(I - P) = 0$, equivalently when the range of $A^*$ is contained in the range of $P$. The smallest such rank is $\dim \mathrm{range}(A^*) = \mathrm{rank}(A)$. For width $k$, Theorem 2 gives the minimum residual as the spectral tail of $A^*A$, so the smallest width below tolerance $\epsilon$ is the stated $k_{\mathrm{late}}(\epsilon)$. If arbitrary measurable readers are allowed, the same lower bound and tail formula hold under the Gaussian history assumption used in Theorem 2; without that assumption, a nonlinear reader can exploit distributional dependence between retained and omitted coordinates.

### B.4 Reconstruction and the introductory example

The best linear reconstruction is $\widehat{h} = S^{1/2}Px$. Its error is

$$\mathbb{E}\|h - \widehat{h}\|^2 = \mathrm{tr}\left(S^{1/2}(I-P)S^{1/2}\right) = \mathrm{tr}((I-P)S). \tag{15}$$

PCA therefore retains leading eigenvectors of $S$. It agrees with response-optimal allocation when the relevant leading subspaces of $S$ and $C$ agree; it can disagree otherwise.

In the example $S = \mathrm{diag}(100, 1, 1)$ and queries ask for the second or third coordinate with probability $1/2$. A two-coordinate reconstruction-optimal state retains the first coordinate and one direction from the last two. Its discarded history variance is one. Averaging over both response queries leaves error $1/2$, independent of the chosen direction inside the tied two-dimensional eigenspace. Retaining both response coordinates discards history variance 100 and incurs zero response error. Thus no uniqueness assumption about the PCA basis is needed.

### B.5 Extra capacity required by reconstruction

**Corollary 7** (Extra capacity required by reconstruction)**.** *For tolerance $\epsilon$, define $k_{\mathrm{resp}}(\epsilon)$ by the smallest spectral tail in Eq. (3) at most $\epsilon$, and define $k_{\mathrm{PCA}}(\epsilon)$ by the smallest PCA width whose response loss is at most $\epsilon$. Then $k_{\mathrm{PCA}}(\epsilon) \geq k_{\mathrm{resp}}(\epsilon)$. The ratio can grow with history dimension: with $L$ higher-variance nuisance coordinates preceding $r$ response coordinates, exact coverage requires $L + r$ PCA coordinates but only $r$ response coordinates.*

The corollary applies to reconstruction-ordered states; arbitrary trained writers require separate analysis.

**Proof.** At every width, the leading response subspace minimizes the specified response loss. It therefore reaches a fixed tolerance no later than a reconstruction-ordered prefix. To construct unbounded relative overhead, take $L$ independent nuisance coordinates with variances strictly greater than one and $r$ independent unit-variance response coordinates. Let queries read the latter with positive weights and ignore the nuisance coordinates. The response map has rank $r$, so a response state of width $r$ is sufficient. The first $L$ PCA coordinates contain no response information, and exact response coverage requires all remaining $r$ coordinates. Hence $k_{\mathrm{PCA}}(0)/k_{\mathrm{resp}}(0) = (L+r)/r$. Letting $L$ grow proves the claim. For nonzero tolerances, the exact retained-response count depends on the query weights. This construction establishes possibility; estimating how often such overhead occurs in natural data is an empirical question.

### B.6 Complete response tables and linear approximation

Fix a sampling measure $\mu$ for histories and complete target rows $Y \in \mathbb{R}^J$, with finite second moments, and a fixed column metric $W \succeq 0$. All means, whitening, regression moments, and risks below use this same history measure. A history-dependent sampling weight can be absorbed into $\mu$; a history-dependent column metric generally requires the joint weighted objective instead. Center $y = Y - \mu_Y$, where $\mu_Y = \mathbb{E}_\mu Y$, and whiten the supported history coordinates so that $\mathbb{E}_\mu x = 0$ and $\mathbb{E}_\mu[xx^\top] = I$. The conditional target mean may depend nonlinearly on $H$.

Let $B = \mathbb{E}_\mu[xy^\top]$, $\varepsilon = y - B^\top x$, and $R_{\mathrm{lin}} = \mathbb{E}_\mu\|W^{1/2}\varepsilon\|^2$. The normal equations give $\mathbb{E}_\mu[x\varepsilon^\top] = 0$. For $z = V^\top x$ and $P = VV^\top$, an optimal linear reader is $\mu_Y + B^\top P x$. Its target residual is $\varepsilon + B^\top(I-P)x$, whose terms are uncorrelated. Consequently,

$$\begin{aligned} L_{\mathrm{tab}}(P) &= \inf_A \mathbb{E}_\mu\|W^{1/2}(y - A^\top V^\top x)\|^2 \\ &= R_{\mathrm{lin}} + \mathrm{tr}((I-P)BWB^\top). \end{aligned} \tag{16}$$

The infimum is attained at $A = V^\top B$; a free intercept restores $\mu_Y$. The eigenvalue variational identity used in Theorem 2 now gives Eq. (6). Thus the table statement is a consequence of ordinary linear projection and reduced-rank prediction. It compares linear writers with optimal linear readers under the specified metric.

Writing $m(H) = \mathbb{E}_\mu[Y \mid H]$ further separates the full-history residual:

$$\begin{aligned} R_{\mathrm{lin}} &= \mathbb{E}_\mu\|W^{1/2}(Y - m(H))\|^2 \\ &\quad + \mathbb{E}_\mu\|W^{1/2}(m(H) - \mu_Y - B^\top x)\|^2. \end{aligned} \tag{17}$$

The first term is conditional target noise; the second is approximation error for the chosen history features and linear prediction class. Both remain even when all supported history coordinates are stored. A richer feature map or nonlinear reader can change the second term.

### B.7 Direction exchanges and fitted readers

For unit $u$ with $Pu = u$ and unit $v$ with $Pv = 0$, the matrix $P' = P - uu^\top + vv^\top$ is a rank-$k$ projector. Subtracting Eq. (16) at these projectors cancels $R_{\mathrm{lin}}$ and gives $\Delta_{\mathrm{ideal}} = \mathrm{tr}((P'-P)C_{\mathrm{tab}})$, which is Eq. (7). Several orthogonal exchanges contribute the sum of their gains. The reconstruction identity above similarly gives

$$L_{\mathrm{rec}}(P') - L_{\mathrm{rec}}(P) = u^\top S u - v^\top S v. \tag{18}$$

Here $S$ and $P$ refer to the same whitening and sampling measure. Positive response gain and positive reconstruction cost can therefore occur at a fixed width. Under the original independent-query linear response model, the identical trace argument uses $C = T^\top G T$.

For a fixed fitted affine reader $\widehat{f}_P(x)$ of $V^\top x$, projection orthogonality yields

$$\begin{aligned} R_{\mathrm{fitted}}(P) &= L_{\mathrm{tab}}(P) + E_{\mathrm{fit}}(P), \\ E_{\mathrm{fit}}(P) &= \mathbb{E}_\mu\|W^{1/2}(\widehat{f}_P(x) - \mu_Y - B^\top P x)\|^2 \geq 0. \end{aligned} \tag{19}$$

Thus $\Delta_{\mathrm{fitted}} = \Delta_{\mathrm{ideal}} + E_{\mathrm{fit}}(P) - E_{\mathrm{fit}}(P')$. A fitted full-history reader has risk $R_{\mathrm{lin}} + E_{\mathrm{fit}}(I)$, and a finite test mean adds sampling variation. Its measured error includes estimation and regularization effects, so the population approximation floor remains unidentified from that measurement alone.

### B.8 Proof of the decision-margin lemma

For every action $a$, Lipschitzness gives

$$|u(a, \widehat{r}_a) - u(a, r_a)| \le L\|\widehat{r}_a - r_a\| \le L\epsilon.$$

Hence the estimated utility of $a^\star$ is at least $u(a^\star, r_{a^\star}) - L\epsilon$, while that of any competitor $a$ is at most $u(a, r_a) + L\epsilon$. Their estimated gap is therefore at least $m - 2L\epsilon > 0$, which proves Lemma 6. The condition is uniform over actions; average MSE leaves it unverified without a separate concentration or bounded-error argument.

The same Lipschitz calculation gives a regret bound without a margin condition. Let $\widehat{a} \in \arg\max_a u(a, \widehat{r}_a)$ and $a^\star \in \arg\max_a u(a, r_a)$. Then

$$u(a^\star, r_{a^\star}) - u(\widehat{a}, r_{\widehat{a}}) \le 2L \max_{a \in \mathcal{A}} \|\widehat{r}_a - r_a\|. \tag{20}$$

Indeed, add and subtract the two estimated utilities and use their maximizing property. If the responses and errors are random, write $M$ for the true best-versus-second-best utility margin and $E = \max_a \|\widehat{r}_a - r_a\|$. For every $\epsilon > 0$,

$$\Pr(\widehat{a} \ne a^\star) \le \Pr(M \le 2L\epsilon) + \Pr(E > \epsilon). \tag{21}$$

The first term is a property of the deployment margin distribution and the second is a response-estimation tail probability. To see this, on the event $\{M > 2L\epsilon, E \le \epsilon\}$ the deterministic lemma applies, so disagreement can occur only in its complement. Equations (20)–(21) connect the response-space error and the action-ranking metrics used in the controlled experiment.

For an empirical table, the same identities hold with normalized row weights applied throughout centering, whitening, and regression. Equal-episode weighting assigns each row in episode $i$ weight $1/(N_{\mathrm{ep}} n_i)$, where $n_i$ is its row count and $N_{\mathrm{ep}}$ the number of episodes. Training moments and selected exchanges define a predicted gain before test evaluation. Matching the reader family and selecting penalties on validation permits a paired test of that prediction on untouched episodes. Test covariance can differ from the training whitening metric, and both moments and fitted readers are estimated; the training trace prediction is therefore assessed against the held-out gain. Computing both quantities from the same unregularized training fit would only verify the projection identity.

**Oracle-side decision-margin audit.** We audited the sufficient margin condition in the known-effect dynamical experiment using the true simulator utility. For each test instance, let $\delta = \max_{a \in \mathcal{A}} |u(a, \widehat{r}_a) - u(a, r_a)|$ be the maximum utility error across the nine candidate actions, and let $m$ be the gap between the best and second-best true utilities. The sufficient condition is $m > 2\delta$. For the QRM width-six state, the condition holds on $14{,}057$ of $25{,}000$ test instances ($56.23\%$), with per-seed coverage rates $56.76\%$, $56.72\%$, $56.20\%$, $55.60\%$, and $55.86\%$. None of the $14{,}057$ covered instances has an action disagreement. The audit uses the true action utility supplied by the generator and evaluates the implication of the margin condition under known effects. Response-space Lipschitz bounds and deployment certification require separate analysis. Instances outside the covered subset remain part of the empirical decision evaluation.

### B.9 When compressed histories are indistinguishable

Suppose two histories have the same state and the same query but target means $f_1$ and $f_2$. A deterministic reader outputs the same value $a$. Then

$$\frac{1}{2}\|a - f_1\|^2 + \frac{1}{2}\|a - f_2\|^2 = \left\|a - \frac{f_1 + f_2}{2}\right\|^2 + \frac{1}{4}\|f_1 - f_2\|^2. \tag{22}$$

Their average error is at least $\|f_1 - f_2\|^2/4$. This bound is independent of reader capacity. Approximate equality of states requires an additional reader-stability condition, and a numerical Lipschitz constant requires more than continuity alone. The executed paired-history test uses exact nullspace perturbations and reports both state distance and required response difference.

## C RANK AND MULTIPLE QUERY FAMILIES

Let $\{F_d\}$ be centered scalar functions of history. For vector responses, replace this set by all $F_{d,m}$ jointly. If their span has finite dimension $r$ and basis $b_1, \dots, b_r$, every response is a query-dependent linear combination of $b(H) = (b_1(H), \dots, b_r(H))$. Storing $b(H)$ is therefore sufficient. In a linear model with full support and a descriptor family spanning its declared space, this is the rank of the corresponding response operator.

**Proposition 8** (Conditional continuous-state lower bound)**.** *Assume the basis functions are continuous. Suppose there is a nonempty open $U \subset \mathbb{R}^r$ and a continuous map $s : U \to \mathcal{H}$ such that $b(s(u)) = u$. If a continuous writer $g : \mathcal{H} \to \mathbb{R}^k$ and a continuous reader recover every declared response, then $k \geq r$.*

*Proof.* Finite query/component evaluations span the basis, so a fixed linear combination of those responses recovers $b(H)$. Applying that combination to the continuous reader gives a continuous map $A$ with $A(g(s(u))) = u$. Hence $g \circ s$ is a continuous injection from $U$ into $\mathbb{R}^k$. If $k < r$, compose it with the canonical embedding of $\mathbb{R}^k$ into $\mathbb{R}^r$. Invariance of domain makes the image open in $\mathbb{R}^r$, although it is contained in a lower-dimensional coordinate plane, a contradiction (Munkres, 2000). □

The continuous section is a substantive assumption whose existence requires justification. For $H = t$, the responses $t$ and $t^2$ have a two-dimensional function span and both can be recovered from one stored scalar, illustrating why function independence alone is insufficient. The statement concerns continuous coordinates; bit-rate and training-sample lower bounds require separate arguments.

For two families with centered response spans $\mathcal{V}_A$ and $\mathcal{V}_B$, the vector-space dimension formula gives

$$r_{A\cup B} = r_A + r_B - \dim(\mathcal{V}_A \cap \mathcal{V}_B). \tag{23}$$

It extends to more families by taking their joint span. In the linear full-support setting, joint rank gives an exact common-state width, while independent banks may duplicate shared information. At a smaller budget, each family also has its own loss and weight; a mean metric can hide poor service for a rare family. The sharing experiment reports both mean and worst-family error.

## D FITTING, WEIGHTING, AND DIAGNOSTICS

### D.1 VECTOR OUTPUTS AND OUTPUT UNITS

For response components $F_m(H, d) = d^\top T_m x$ and positive semidefinite output weighting matrix $Q$, write $\widetilde{T}_i = \sum_m (Q^{1/2})_{im} T_m$. Stack $G^{1/2}\widetilde{T}_i$ vertically to form $K$. Then $C = K^\top K$ and the scalar proof applies unchanged. This formulation makes history scaling, query weights, and output units explicit. An unweighted sum of position, angle, and force errors has no stable physical meaning unless units are fixed deliberately.

The external response-table protocol normalizes each output using the training mean and standard deviation, then weights query families equally and coordinates equally within family. These weights define the evaluation metric for this study. A separate action-selection evaluation converts predicted utility coordinates back to original units before choosing an action.

### D.2 IMPLEMENTATION AND REDUCED-RANK REGRESSION

For triple data, the dense reader design is $d \otimes [1, z]$. Ridge regression estimates the query-conditioned intercept and response coefficients. An SVD of the fitted weighted response map defines the writer ordering. Fitted test responses use learned readers, while synthetic population references use the known operator.

For a complete $N \times J$ response table, let $\Omega$ contain normalized row weights on its diagonal, and let $X$ and $Y$ contain histories and targets centered under those weights. With $X^\top \Omega X = I$ on its support, multiresponse least squares gives $\widehat{B} = X^\top \Omega Y$. Uniform row weights recover $\widehat{B} = X^\top Y/N$. SVD truncation of $\widehat{B}W^{1/2}$ is exactly the classical rank-constrained prediction solution for

the weighted response metric. Isotropic ridge shrinks the fitted operator in these coordinates. Under matched coordinates and weights, QRM and reduced-rank regression share this estimator; results are organized around the QRM state contract and response metric, with this estimator correspondence stated explicitly.

Writer coordinates can be rescaled without changing the represented subspace. Reader fitting normalizes stored coordinates internally so that regularization strengths are comparable across methods, then maps coefficients back to the stored coordinates. The response-table reconstruction probe is fitted separately by least squares and uses no response-selected ridge penalty. Reconstruction quality therefore remains defined independently of response hyperparameter selection.

### D.3 Perturbation bound

For $\|\widehat{C} - C\|_{\mathrm{op}} \le \delta$ in fixed common coordinates, the allocation diagnostic obeys

$$|\widehat{\Gamma}_k(P) - \Gamma_k(P)| \le 2k\delta, \qquad \Gamma_k(\widehat{P}_k) \le 2k\delta. \tag{24}$$

Let $C$ and $\widehat{C}$ be in the same fixed coordinates, with operator-norm difference at most $\delta$. For a rank-$k$ projector, $|\operatorname{tr}(P(C - \widehat{C}))| \le k\delta$. The variational characterization also gives

$$\left| \sum_{j \le k} \lambda_j(C) - \sum_{j \le k} \lambda_j(\widehat{C}) \right| \le k\delta. \tag{25}$$

Adding the inequalities proves the first part of Eq. (24). If $\widehat{P}$ maximizes $\operatorname{tr}(P\widehat{C})$ and $P_\star$ maximizes $\operatorname{tr}(PC)$, then

$$\Gamma_k(\widehat{P}) = \operatorname{tr}(P_\star C) - \operatorname{tr}(\widehat{P}C) \tag{26}$$

$$\le \operatorname{tr}(P_\star(C - \widehat{C})) + \operatorname{tr}(\widehat{P}(\widehat{C} - C)) \le 2k\delta. \tag{27}$$

The bound uses an estimated $\delta$, while whitening contributes an additional error term. The following result gives a computable value for a frozen finite-dimensional response feature. It is stated in fixed common coordinates; estimation of the whitening map or regression operator contributes to the additional error term.

**Proposition 9** (Finite-sample spectral stability)**.** *Let $q_1, \dots, q_n \in \mathbb{R}^p$ be i.i.d., let $C = \mathbb{E}[q_1 q_1^\top]$, and assume $\|q_1\|_2 \le R$ almost surely. Suppose the fitted Gram matrix has the form*

$$\widehat{C} = \frac{1}{n} \sum_{i=1}^{n} q_i q_i^\top + E_n, \qquad \|E_n\|_{\mathrm{op}} \le b_n.$$

*For any $\eta \in (0, 1)$, define*

$$\Delta_n(\eta) = \frac{R^2}{\sqrt{n\eta}} + b_n.$$

*Then, with probability at least $1 - \eta$,*

$$\|\widehat{C} - C\|_{\mathrm{op}} \le \Delta_n(\eta), \qquad \Gamma_k(\widehat{P}_k) \le 2k\Delta_n(\eta).$$

*If $\gamma_k = \lambda_k(C) - \lambda_{k+1}(C) > 0$, the leading subspaces satisfy*

$$\|\sin\Theta(\widehat{V}_k, V_k^\star)\|_{\mathrm{op}} \le \min\left\{1, \frac{2\Delta_n(\eta)}{\gamma_k}\right\}.$$

*Here $\widehat{P}_k$ and $P_k^\star$ are the leading rank-$k$ projectors of $\widehat{C}$ and $C$.*

**Proof.** Set $Z_i = q_i q_i^\top - C$. Independence and centering give

$$\mathbb{E}\left\| \frac{1}{n} \sum_i Z_i \right\|_F^2 = \frac{1}{n} \mathbb{E}\|Z_1\|_F^2 = \frac{1}{n} \left( \mathbb{E}\|q_1 q_1^\top\|_F^2 - \|C\|_F^2 \right) \le \frac{R^4}{n}.$$

Markov's inequality and $\|M\|_{\mathrm{op}} \le \|M\|_F$ give the first bound for the empirical Gram matrix, and the triangle inequality adds $b_n$. Equation (24) gives the allocation bound. For completeness, the spectral-projector resolvent identity applied to the contour halfway between $\lambda_k(C)$ and $\lambda_{k+1}(C)$ gives

$$\|(I - P_k^\star)\widehat{P}_k\|_{\mathrm{op}} \le \frac{2\|\widehat{C} - C\|_{\mathrm{op}}}{\lambda_k(C) - \lambda_{k+1}(C)}$$

whenever the right side is below one; when it is at least one, the trivial bound $\|(I - P_k^\star)\widehat{P}_k\|_{\mathrm{op}} \le 1$ applies. This is the stated sin-$\Theta$ inequality. □

The proposition applies to a response feature such as $q = G^{1/2}Tx$ when the population whitening and response coordinates are frozen. When $\widehat{T}$, $S^{-1/2}$, or $G^{1/2}$ are estimated, their operator-norm errors enter $b_n$. The result therefore gives a direct sample-size and spectral-gap calculation for the fixed-coordinate component of QRM estimation; reader fitting and whitening uncertainty enter through the additional error term.

### D.4 Identified and unidentified query directions

For an identified response operator $T$, the omitted energy of query $d$ is

$$\rho(d) = \|(I - P)T^\top d\|^2. \tag{28}$$

If $T$ is known, Eq. (5) follows directly by projecting the demand vector $T^\top d$. A unit rotation $u\cos\theta + v\sin\theta$ with $Pu = u$, $Pv = 0$, and orthogonal $u, v$ has omitted squared norm $\sin^2\theta$.

If calibration queries lie in $U$, choose an operator difference $B$ satisfying $B^\top d = 0$ for $d \in U$. Then $T$ and $T + B$ are observationally equivalent on every available query. Outside $U$, the difference can be nonzero and arbitrarily large without additional bounds on $B$. A descriptor-support test therefore identifies when more calibration is needed; numerical unseen-response error requires bounds on $B$. The transfer experiment identifies all eight directions first, then restricts the operating memory to four. The separate unsupported-query test trains on four directions and returns an unknown-response status for the remaining directions.

## E Dependence between history and query

For arbitrary dependence, let $\phi = d \otimes [1, V^\top x]$ be the full bilinear-reader feature. With $A_V = \mathbb{E}[\phi\phi^\top]$ and $b_V = \mathbb{E}[\phi r]$, the best linear coefficient has minimum squared prediction error

$$\mathbb{E}[r^2] - b_V^\top A_V^\dagger b_V, \tag{29}$$

assuming finite second moments and the usual support condition. This expression depends on the joint law through $A_V$ and $b_V$; product moments from the marginal history and query distributions are insufficient in general.

We test this boundary with $x \sim \mathcal{N}(0, I_2)$ and response $r = d^\top x$. With probability $\eta$, the one-hot query selects the coordinate of larger absolute magnitude; otherwise it selects a coordinate uniformly at random. The same response function is used for every $\eta$. A width-one spectral state is compared with a search over 181 one-dimensional directions, each scored on training data using the true empirical joint response loss and a refitted reader. There are 9,000 training and 40,000 independent test examples per seed, for three seeds. The search serves as a two-dimensional diagnostic.

| Dependence fraction | Spectral QRM | Joint-loss search | Product prediction |
|---|---|---|---|
| 0.0 | 0.4991 | 0.4982 | 0.4890 |
| 0.5 | 0.6147 | 0.4499 | 0.4837 |
| 1.0 | 0.6454 | 0.2982 | 0.4850 |

Table 3: Means over three seeds. Dependence can invalidate the product-moment allocation rule and its error prediction. The independently evaluated direct joint-loss search exposes this boundary.

The independent case provides a control. At greater dependence, joint-loss optimization performs better and the product-moment prediction no longer reliably matches measured loss. Theorem 2 applies to the independent-query regime; this experiment probes its dependent-query complement. Real action or query selection based on history should use a joint-loss model or be evaluated empirically.

| Calibration episodes | Failures at width 4 | Upper bound | Selected width |
|---|---|---|---|
| 40 | 0 | 0.1128 | None |
| 100 | 0 | 0.0467 | 4 |
| 300 | 0 | 0.0158 | 4 |

Table 4: Independent capacity calibration for six frozen widths, one query family, confidence 0.95, and target failure probability 0.1. A failure is squared response error above 0.05. All five seeds give the shown selection.

## F Independent calibration for memory width

Candidate writers and readers are frozen using training and validation data before independent calibration episodes are examined. With $K$ widths, $J$ query families, $b_{kj}$ failures among $n$ episodes, and target simultaneous error $\alpha$, the one-sided exact bound is

$$U_{kj} = \mathrm{Beta}^{-1}\left(1 - \frac{\alpha}{KJ}; b_{kj} + 1, n - b_{kj}\right), \tag{30}$$

with $U_{kj} = 1$ if $b_{kj} = n$ (Clopper & Pearson, 1934). The smallest width satisfying $U_{kj} \le \beta$ for every family is selected; if none qualifies, the procedure returns no certified width.

Let the models be fixed before calibration. For independent episodes $i = 1, \dots, n$, candidate width $k$ and family $j$, define

$$B_{ikj} = \mathbb{1}\{\ell_{ij}(k) > \tau_j\}, \qquad p_{kj} = \Pr(B_{ikj} = 1). \tag{31}$$

Each $\ell_{ij}$ is a prespecified aggregation of potentially dependent reads within one episode. The episode event, tolerance, and family weights are fixed before examining calibration outcomes.

For each $(k, j)$, the exact one-sided binomial interval has failure probability at most $\alpha/(KJ)$ (Clopper & Pearson, 1934). Applying a union bound gives

$$\Pr\{p_{kj} \le U_{kj} \text{ for all } k, j\} \ge 1 - \alpha. \tag{32}$$

No independence between widths or families is needed. On this simultaneous event, any width selected using the upper bounds has every family failure probability at most $\beta$. Returning no width is permitted and required when no candidate passes. The guarantee is conditional on the frozen training/validation procedure, thresholds, query families, and deployment distribution.

With zero failures, the upper bound is $1 - (\alpha/(KJ))^{1/n}$. For the experiment $K = 6$, $J = 1$, $\alpha = .05$, and $\beta = .1$, at least

$$n \ge \left\lceil \frac{\log(\alpha/(KJ))}{\log(1 - \beta)} \right\rceil = 46 \tag{33}$$

episodes are needed even in the zero-failure case. Forty episodes therefore provide insufficient calibration evidence for a width. This sample requirement concerns calibration confidence; the number of demonstrations needed to learn the policy or operator is a separate quantity.

With 40 calibration episodes, no width is certified. With 100 or 300 episodes, width four qualifies in all five seeds; each selected model is evaluated on 20,000 fresh examples. These observations check the implementation in this generator, while the simultaneous guarantee follows from the sampling assumptions above.

For the external three-family, five-width protocol, the corresponding zero-failure requirement is 55 independent episodes. A split with 30 calibration demonstrations provides insufficient evidence for this particular 0.1 failure-probability claim even if all predictions pass the tolerance. The runner reports an unselected width when no candidate passes and leaves the bound unchanged after data inspection. Investigators may predeclare a different tolerance or confidence level before collecting an independent calibration set, or report descriptive errors without certification.

Episodes from the same physical recording session or repeated seed may remain dependent. Such data require grouping at the session or seed level. The software checks group disjointness and prevents frame-level trials; numerical arrays alone leave independence unverified. Confidence intervals over test episodes are descriptive paired-bootstrap intervals and are labeled separately from these exact calibration bounds.

## G PREDICTION RISK AND METRIC VALIDATION

Suppose $e = y + d^\top c + \varepsilon$, where $c \sim \mathcal{N}(0, I)$, $y \sim \mathcal{N}(0, A)$, and $\varepsilon \sim \mathcal{N}(0, N)$ are independent. A stored linear state reveals $Pc$ and leaves response variance $\rho(d) = d^\top(I - P)d$. Gaussian conditioning yields

$$\widehat{y}_\star = \frac{A}{A + N + \rho(d)}(e - d^\top Pc), \qquad R_\star(d) = \frac{A(N + \rho(d))}{A + N + \rho(d)}. \tag{34}$$

The function $f(\rho) = A(N + \rho)/(A + N + \rho)$ is increasing and concave. For a random residual sample $u$, $u^2$ need not equal its conditional variance, and in general $\mathbb{E}[f(u^2)] \neq f(\mathbb{E}[u^2])$. With $A = N = 1$ and omitted variance one, the no-memory Bayes risk is $2/3$, whereas the transformed-residual score is approximately $0.621$.

The revised evaluation constructs actual predictions and measures $(\widehat{y} - y)^2$. It separately records a population Bayes-risk reference and a measured oracle-predictor error. The learned prediction head estimates residual covariance from validation data; controlled validation responses are known from the generator. A finite test mean of $0.4966$ is consistent with the population optimum of $0.5$, which sets the attainable benchmark for this generator.

Minimizing mean omitted response energy need not minimize every nonlinear downstream risk. Concavity makes this especially relevant when comparing distributions of $\rho(d)$. The allocation theorem optimizes squared response error, and the final-target experiment measures the resulting consequence under its stated residual model. A universal success-rate conversion requires an additional model.

## H EXECUTED CONTROLLED PROTOCOLS

### H.1 PRIMARY GENERATOR

History covariance is diagonal with eigenvalues

$$(40, 35, 30, 25, 20, 15, 10, 5, 1, 1, 1, 1, 1, 1, 1, 1, .25, .25, .25, .25, .25, .25, .25, .25).$$

A standard normal $x$ generates $h = S^{1/2}x$. For rank $r$, response row $j$ is $\cos\theta\, e_j^\top + \sin\theta\, e_{8+j}^\top$. Queries are independent standard Gaussian vectors normalized to unit length. Response variance and rank are fixed as $\theta$ varies. The default scan uses $r = 4$ and angles $0, 15, 30, 45, 60, 75, 90$ degrees; the rank scan uses $r = 1, 2, 4, 6$ at 90 degrees. Reusing the random source across angles holds the histories fixed.

Widths are $0, 1, 2, 4, 6, 8$ and seeds are 5101, 5107, 5113, 5129, 5147. Each condition uses 12,000/2,000/5,000 train/validation/test triples. Training labels add standard deviation $0.05$ noise; the test response mean is noiseless. Ridge penalty is $10^{-5}$ under the sample-averaged normal-equation convention. All dense linear calculations use float64. Memory bytes are reported separately for the actual dtype and nominal float32 deployment.

### H.2 INTERVENTIONS

For the rank-four, 90-degree, width-four PCA state, the algorithm scores retained and omitted directions using the training-fitted response matrix. It replaces zero to four low-response retained directions with high-response omitted directions. Random replacements and restoring the original projection are paired controls. The reader is refitted using the same training triples after each replacement. Response and reconstruction errors are evaluated on the unchanged test set.

For paired histories, the evaluator projects a true query demand into the learned writer's nullspace. The pair $x_\pm = \pm v$ has equal states and a measurable response difference. Because writer nullspaces differ, each method induces its own pair; the QRM response difference quantifies its omitted component. Pair values and the lower bound are reported with the experiment.

| Shared directions | Union rank | Shared state MSE | Separate banks MSE |
|---|---|---|---|
| 0 | 8 | 0.4994 | 0.4983 |
| 2 | 6 | 0.2492 | 0.4929 |
| 4 | 4 | $1.22 \times 10^{-5}$ | 0.5070 |

Table 5: Two rank-four query families over the same histories. Both implementations store four scalars in total. Shared training uses 20,000 triples; separate banks use the same total split by family. Zero overlap is a negative control.

### H.3 Spectra and query transfer

The spectrum study uses eight response directions, 10,000 training and 10,000 test histories, and the same five seeds. Query-coordinate variances are respectively

$$(1/8, 1/8, 1/8, 1/8, 1/8, 1/8, 1/8, 1/8),$$
$$(.45, .24, .14, .08, .04, .025, .015, .01),$$
$$(.49, .29, .19, .01, .007, .006, .004, .003).$$

The reported spectral tail is a known population reference. Method curves are test errors of learned models. For transfer, all eight query directions are identified during training, but the design measure gives the first four 100 times the weight of the remaining four. A four-coordinate state is fixed, and a test query rotates from the first to the fifth direction. A separate experiment removes the last four directions from identification data and verifies that the support diagnostic returns unknown for those queries.

### H.4 Shared families

Two families each use four of eight response coordinates. Their intersection dimension is zero, two, or four. There are 10,000 distinct training histories and two queries per history, for 20,000 total triples. Both joint and separate implementations use that same total supervision. Independent test histories number 5,000. Total widths are two, four, six, and eight; the symmetric separate-bank comparison assigns half to each family, which is optimal under this constructed symmetry. The external implementation instead chooses bank allocation by validation to avoid relying on symmetry in real datasets.

At zero overlap, shared and separate states perform similarly. With full overlap, a joint four-coordinate state covers both families while two two-coordinate banks each omit part of the same information.

### H.5 Multiple reads of a dynamical history

Two-dimensional position, velocity, and acceleration at the write time form a six-dimensional latent variable. Three noisy position observations at times $-2, -1, 0$ are concatenated with 18 independent nuisance features. Targets are position $p + tv + \frac{1}{2}t^2a$. Training queries use horizons $0, .5, 1, 2$ and a polynomial descriptor $(1, t, t^2/2)$. Evaluation reads current position, horizons 1.5 and 3, and the horizon-two prediction used to rank nine actions. Each action adds a known $2a$ displacement; the goal is independently sampled. Mean realized cost minus the best candidate cost defines action regret.

There are 10,000 training and 5,000 test histories per seed. Reads share an unchanged stored state. The known action effect makes this a test of memory reuse and decision sensitivity. Future extrapolation uses the generator's polynomial structure.

## I Neural controls, units, and precision

The executed neural comparison uses 4,096 training and 5,000 test examples with seeds 5101, 5107, and 5113. The encoder has two 64-unit tanh layers and a linear output. A single global input scale preserves relative feature variances. Models train for 800 Adam steps at learning rate $0.002$ with batches of 256. Reconstruction models minimize history error. Supervised models minimize

| Memory | State MSE | Future MSE | Action regret | Optimal action rate |
|---|---|---|---|---|
| PCA | 1.0145 | 4.5139 | 9.6009 | 0.2527 |
| Random | 0.9644 | 4.4971 | 9.5338 | 0.2563 |
| Coefficient SVD | 0.0001 | 0.0050 | 0.0096 | 0.9645 |
| QRM | 0.0001 | 0.0050 | 0.0096 | 0.9645 |

Table 6: Width-six dynamical experiment, five seeds. Action regret is one-step candidate regret under known action effects.

response error. The nested model cycles through prefixes one, two, four, and eight of a width-eight writer. Separately trained baselines use widths two, four, and eight.

After fitting each writer, a common bilinear response probe and a linear reconstruction probe are refitted on training data. The table therefore evaluates stored content through common probes. Each architecture receives fixed compute, while the external runner provides early stopping and separate checkpoint selection.

| Memory | Reconstruction | Response MSE |
|---|---|---|
| Reconstruction AE | 60.6323 | 0.9998 |
| Supervised bottleneck | 186.4396 | 0.0069 |
| Nested supervised | 186.2277 | 0.0179 |

Table 7: Width-four neural controls, three seeds. The independent supervised model is stronger than the nested model in this run.

The robustness experiment changes input units, applies a random orthogonal rotation, standardizes individual features using training statistics, and varies training counts $256, 1024, 4096, 12000$. Each transformed comparison refits the same method from the transformed training data. Population prediction information is unchanged by invertible transforms; PCA's variance objective need not be invariant to rescaling. Numerical covariance truncation and ridge can still affect finite implementations. The complete results are supplied, including standardized conditions where PCA's error changes.

The precision experiment casts committed states to float64, float32, or float16 and then restores float64 for reader computation. It uses widths two, four, and eight, three seeds, and reports both response MSE and the squared difference from float64 predictions. These measurements characterize stored-value precision; half-precision training, hardware throughput, and robustness to arbitrary quantization remain separate evaluation questions.

## J External experiment specification

The recorded-history protocol below has been executed on three tasks, with native observations and a separately labeled artificial-dropout condition. Physical action branches remain an integration requirement. Schema fixtures provide programmatic file-handling tests, while the external evidence comes from recorded task data.

### J.1 Recorded robot histories: three manipulation tasks

The main external pipeline uses proficient-human Lift, Can, and Square data from the official robomimic v1.5 low-dimensional demonstrations (Mandlekar et al., 2022). Entire demonstrations are split before constructing windows. The task is offline multi-horizon prediction, with metrics distinct from the native robomimic policy-success benchmark.

Each history contains eight observations with stride three, ending at the write time. Query columns select future offsets 1, 5, 10, or 20 and end-effector position, the first three object-observation components, or gripper opening. The object coordinates are named by their array semantics until their physical meaning is checked for each task. Future actions are not supplied to the writer; the target is the demonstrator's observed continuation. Training/validation/calibration/test fractions are $.45/.10/.30/.15$ by episode. All windows and query labels from one episode remain in the same split.

Baseline writers are PCA, random projection, raw-coordinate coefficient SVD, weighted reduced-rank QRM, supervised PLS, the latest-frame control, and separately trained family banks. The coefficient-SVD comparison changes both column weights and direction coordinates relative to QRM; the same-coordinate ablation isolates column weighting. Bank allocation and reader penalties are selected on validation at the same total stored width. A neural runner adds a reconstruction autoencoder, an independent supervised MLP, a nested supervised MLP, and a GRU. The deployment writer receives history alone. Last-frame controls assess the added value of temporal history.

Training moments and standardization assign equal total weight to each source episode. The complete current-frame baseline has an independently validation-tuned reader. Neural comparisons use five seeds, widths 2, 4, 8, 16 and 32, hidden size 128, AdamW at $3 \times 10^{-4}$, at most 200 epochs, and validation patience 20. A common train-fitted affine reconstruction decoder measures retained history information; the autoencoder native decoder is evaluated separately.

Metrics are equal-episode, equal-family standardized response MSE; feature reconstruction error; prediction intervals across source episodes; actual state bytes; write/read latency; and calibration outcomes. The low-dimensional observation experiment characterizes state-feature reconstruction. Image reconstruction requires a separate visual-feature experiment using rendered observations. A controlled missing-observation variant masks recent object entries with a training-only mean and an explicit visibility indicator. This constructed dropout is reported separately from native observations.

### J.2 Physical candidate-action responses

An unexecuted collector is prepared to restore recorded robosuite states and execute a fixed set of candidate action sequences from each state. Candidate identity, output type, and horizon would define the later query. The planned labels are observed end-effector and object components, cumulative shaped reward, and task-success indicators. A repeated branch would verify restoration before labels are accepted. The planned writer receives the recorded history alone, while the privileged simulator state is used only to generate outcomes.

This protocol tests actual action-conditioned physical transitions, complementing the known additive action effect of the controlled generator. A candidate-selection evaluator measures regret against all replayed outcomes and reports whether action utilities differ at all. Offline macro-action selection rates are reported separately from full closed-loop task success. Simulator version compatibility and replay checks must pass on the machine running the experiment; this integration remains unexecuted in the present environment.

### J.3 Natural motion histories

The nuScenes adapter specifies multi-horizon object-displacement queries from annotation tracks (Caesar et al., 2020), with whole-scene splits, past-only positions, sizes, orientations, and observed time differences. Future annotation offsets define targets and a past-only constant-velocity reference. A diagnostic rerun reproduced 72 aggregate conditions within $3.9 \times 10^{-9}$; its label-permutation controls approach its better scores, so the rerun remains an audit. Full-domain and calibration estimates require evidence beyond the ten-scene annotation-track pilot.

### J.4 Data and reporting contract

Every prepared table contains one row per unique history, all declared response columns, a source-episode or scene group, a split label, and explicit query names and families. Train-only statistics define input and output scaling. The test file is evaluated after tuning and model freezing. The analysis record specifies source kind, split sizes, versions, and metrics. Synthetic schema-fixture results require an explicit flag and remain separate from external manuscript imports.

The reporting protocol retains every method and records unsuccessful calibration as no selection. Conclusions remain tied to each table's declared estimand. Analyses specified after test inspection are identified by timing; confirmation requires a new held-out split under a predeclared criterion. Physical data access, robot assets, and pretrained weight downloads remain subject to their original licenses.

### J.5 Recorded-history direction interventions

The intervention runner uses training-fitted response-energy scores to replace retained covariance directions at fixed width, refits the same reader, and evaluates unchanged test episodes. Original and restored rows are paired controls. The frozen result contains every declared condition, width, replacement strength, targeted direction, and random seed; no row is filtered by effect sign.

| Condition | Positive | Zero | Negative | Total |
|---|---|---|---|---|
| masked_can | 9 | 10 | 75 | 94 |
| masked_lift | 6 | 10 | 78 | 94 |
| masked_square | 5 | 10 | 79 | 94 |
| recorded_can | 9 | 10 | 75 | 94 |
| recorded_lift | 7 | 10 | 77 | 94 |
| recorded_square | 11 | 10 | 73 | 94 |

Table 8: Counts of response effects in the frozen intervention table. Positive means lower response MSE than the unperturbed same-width row; zero includes original/restored controls.

Across all 564 rows there are 47 positive, 60 zero, and 457 negative point estimates. Among targeted/random rows, 47 are positive and 457 are negative. The paired mean, percentile interval, and episode count are reported for every row. These measurements characterize offline response changes under fixed-width direction replacements.

| Condition | Method | Width | Replaced | Gain | 95% lower | 95% upper |
|---|---|---|---|---|---|---|
| recorded_can | targeted | 8 | 1 | 0.0147 | 0.0053 | 0.0244 |
| recorded_can | random | 8 | 1 | -0.0232 | -0.0349 | -0.0119 |
| recorded_lift | targeted | 8 | 1 | -0.0027 | -0.0404 | 0.0398 |
| recorded_square | random | 8 | 4 | -0.4201 | -0.4731 | -0.3706 |

Table 9: Representative paired intervention rows; the full targeted and random grid is frozen in the CSV.

### J.6 Measured table-kernel costs

The cost runner measures the fitted linear table writer and reader on the recorded test feature array. It uses 20 warmups and 300 repetitions per model and batch. Measurements are NumPy CPU timings. The two batches are one history and the largest batch allowed by the test split (min(64, test windows)); Lift therefore uses 52 for its second batch and Can/Square use 64. History extraction, RGB backbones, storage I/O, simulator execution, temporary BLAS buffers, peak process memory, and neural inference are outside this measurement.

| Condition | Batch | Model | State B | Param B | Writer median ($\mu$s) | Reader median ($\mu$s) |
|---|---|---|---|---|---|---|
| masked_can | 1 | pca_k8 | 64 | 21856 | 14.691 | 9.334 |
| masked_can | 1 | qrm_k8 | 64 | 21856 | 15.032 | 9.441 |
| masked_can | 64 | pca_k8 | 64 | 21856 | 156.323 | 23.163 |
| masked_can | 64 | qrm_k8 | 64 | 21856 | 158.241 | 23.578 |
| masked_lift | 1 | pca_k8 | 64 | 19296 | 13.966 | 8.973 |
| masked_lift | 1 | qrm_k8 | 64 | 19296 | 13.927 | 8.992 |
| masked_lift | 52 | pca_k8 | 64 | 19296 | 119.697 | 24.271 |
| masked_lift | 52 | qrm_k8 | 64 | 19296 | 119.542 | 23.958 |
| masked_square | 1 | pca_k8 | 64 | 21856 | 14.566 | 9.128 |
| masked_square | 1 | qrm_k8 | 64 | 21856 | 14.473 | 9.272 |
| masked_square | 64 | pca_k8 | 64 | 21856 | 169.673 | 26.154 |
| masked_square | 64 | qrm_k8 | 64 | 21856 | 143.937 | 26.296 |
| recorded_can | 1 | pca_k8 | 64 | 21216 | 15.129 | 10.175 |
| recorded_can | 1 | qrm_k8 | 64 | 21216 | 15.515 | 10.117 |
| recorded_can | 64 | pca_k8 | 64 | 21216 | 170.122 | 27.843 |
| recorded_can | 64 | qrm_k8 | 64 | 21216 | 170.264 | 27.880 |
| recorded_lift | 1 | pca_k8 | 64 | 18656 | 14.457 | 9.700 |
| recorded_lift | 1 | qrm_k8 | 64 | 18656 | 14.440 | 9.825 |
| recorded_lift | 52 | pca_k8 | 64 | 18656 | 124.943 | 23.241 |
| recorded_lift | 52 | qrm_k8 | 64 | 18656 | 108.764 | 25.782 |
| recorded_square | 1 | pca_k8 | 64 | 21216 | 6.789 | 4.086 |
| recorded_square | 1 | qrm_k8 | 64 | 21216 | 6.046 | 3.722 |
| recorded_square | 64 | pca_k8 | 64 | 21216 | 121.836 | 16.778 |
| recorded_square | 64 | qrm_k8 | 64 | 21216 | 119.372 | 16.135 |

Table 10: Width-eight CPU NumPy measurements for QRM and PCA. The complete cost table reports writer and reader p95 latency, input bytes, model conditions, and batch sizes.

At a fixed condition and width, QRM and PCA have the same persistent state bytes and runtime parameter bytes in the frozen measurements. Timing varies with condition and batch. These measurements characterize the NumPy writer and reader path summarized in the table.

### J.7 Calibration certificates: no selected external width

The frozen calibration rule uses $\alpha = 0.05$, maximum failure probability $0.1$, three query families, and the five predeclared widths. No width is selected for the external conditions. Lift has 51 valid calibration episodes in each mode, while Can and Square have 60 episodes; their joint maximum-over-family Clopper–Pearson upper bound remains above $0.1$ at every width. The table reports both scalar and joint upper bounds.

| Condition | Episodes | Min scalar CP UB | Min joint CP UB | Widths | Selected |
|---|---|---|---|---|---|
| masked_can | 60 | 0.152704 | 1.000000 | 2,4,8,16,32 | – |
| masked_lift | 51 | 0.888221 | 0.994180 | 2,4,8,16,32 | – |
| masked_square | 60 | 0.124183 | 0.983274 | 2,4,8,16,32 | – |
| recorded_can | 60 | 0.152704 | 0.983274 | 2,4,8,16,32 | – |
| recorded_lift | 51 | 0.901691 | 0.999935 | 2,4,8,16,32 | – |
| recorded_square | 60 | 0.124183 | 0.928225 | 2,4,8,16,32 | – |

Table 11: Frozen external calibration certificate summary. The minimum joint Clopper–Pearson upper bound is the minimum, over candidate widths, of the maximum upper bound across the three query families.

The simultaneous coverage argument uses episode-level independence. The table reports the resulting calibration status for the observed episode sets.

### J.8 Scope and limitations of the external evidence

The external results measure offline multi-horizon responses on recorded robomimic histories. Native low-dimensional observations and constructed recent-observation dropout are separate conditions. The PH Lift width-eight result, MH Can replacement interval, MH Lift ridge comparison, and current-only neural controls define the observed range across tasks and readers. The MH evidence covers three tasks and six demonstrators per task; trajectory-level intervals are conditional on this demonstrator mixture.

The frozen aggregate table reports state bytes for linear rows; the neural runs use 32 B float32 states at width eight. Original linear states use 64 B float64, and the fixed-reader storage cast yields 32 B linear states with negligible score change across the six native-history conditions. Equal persistent-state bytes leave model parameters, writer computation, reader computation, and training resources separate.

Direction interventions replace stored directions and refit an offline reader. The frozen PH grid contains positive, zero, and negative rows. The fixed MH Can/Square selection has a positive interval on Square and an interval containing zero on Can. These are response diagnostics. CPU NumPy table-kernel timings cover the stated warmups, repetitions, batches, and linear writer/reader path; GPU throughput, peak memory, preprocessing, storage I/O, neural inference, simulator execution, and end-to-end control cost are outside this measurement. Timings vary by condition and batch.

The frozen PH calibration rule selects no width under the stated sampling assumptions. The robosuite candidate-action branch remains an integration step; the current evidence concerns offline response and action-ranking metrics.

The PH Lift covariance diagnostics and the legacy nuScenes handoff reproduction use protocols distinct from the frozen primary comparisons. The MH Can/Square neural comparisons, MH Lift current-only GRU rerun, and PH current-only neural reruns are reader-specific controls. The MH Lift covariance-ridge test follows its separate fixed protocol, while the full-history GRU serves as a prelisted reference.

### J.9 Independent multi-human replication and current-frame checks

The MH Can and Square runs use the official robomimic v1.5 low-dimensional multi-human demonstrations from six demonstrators per task. Each task has 300 demonstrations. Entire demonstrations

are split by seed 202709 into 135 training, 30 validation, 90 calibration, and 45 test episodes, giving 10,208 Can and 13,807 Square valid windows in total. This data source is distinct from the frozen PH demonstrations. The observed continuation is the target; no physical action is executed by the experiment.

The protocol fixes eight observations at stride three; response offsets one, five, ten, and twenty; three equal-weight query families; and width eight. All standardization and direction scores are fit on training data, with reader ridge selected on validation episodes. Query and reconstruction differences are paired within each test episode. The 95% intervals use 5,000 paired bootstrap resamples of the 45 test episodes. Each six-demonstrator analysis uses 5–11 test episodes per person, and the episode-bootstrap interval is conditional on this demonstrator mixture.

A predeclared selection rule considers targeted one-, two-, and four-direction replacements, chooses the largest training-predicted query gain when positive, and otherwise keeps the original PCA state. The rule chooses two directions in both tasks. Table 12 reports the held-out results. Square meets the positive-gain criterion; the Can interval spans zero. The complete candidate comparison is included in the supplementary material.

| Task | Train prediction | Test gain vs. PCA [95% CI] | Gain vs. matched random [95% CI] |
|---|---|---|---|
| Can | 0.0476 | 0.0120 [-0.0169, 0.0483] | 0.1613 [0.1291, 0.2031] |
| Square | 0.0236 | 0.0200 [0.0113, 0.0290] | 0.2104 [0.1949, 0.2265] |

Table 12: Fixed-rule MH direction replacement at width eight. A gain is a reduction in standardized response MSE. The random comparator is the mean of five same-size random replacements.

For the full-history versus current-frame linear references, the MH Can and Square gains are 0.0204 [0.0136, 0.0289] and 0.0201 [0.0168, 0.0237]. All six demonstrators in each task have positive mean gains for this reference comparison, and also exhibit the QRM–PCA response and reconstruction ranking conflict. The current-frame and full-history references use different state sizes; the comparison therefore isolates the value of temporal history.

Table 13 compares separately trained supervised bottlenecks and GRUs with full histories and current-only inputs. Current-only training repeats the latest frame across all eight input slots. Each model pair uses the same network dimensions, encoder/reader parameter counts, five seeds, train/validation/test split, early-stopping rule, and 32 B float32 stored state. The reported interval is computed after averaging the five seed-specific per-window losses and then pairing by test episode; seed variability is retained in the accompanying CSV. PH current-only and MH neural runs were added as separate reader-specific checks. Neither entered the predeclared direction-swap selection or success criterion.

| Data | Task | Model | Full history | Current only | History gain [95% CI] |
|---|---|---|---|---|---|
| PH | Can | MLP | 0.1598 | 0.1222 | -0.0376 [-0.0510, -0.0257] |
| PH | Can | GRU | 0.1443 | 0.1218 | -0.0225 [-0.0322, -0.0132] |
| PH | Square | MLP | 0.1315 | 0.1141 | -0.0174 [-0.0234, -0.0120] |
| PH | Square | GRU | 0.1250 | 0.1111 | -0.0139 [-0.0199, -0.0083] |
| PH | Lift | MLP | 0.3803 | 0.3903 | 0.0100 [-0.0281, 0.0377] |
| PH | Lift | GRU | 0.3488 | 0.3364 | -0.0124 [-0.0296, 0.0028] |
| MH | Can | MLP | 0.1128 | 0.1047 | -0.0081 [-0.0156, -0.0021] |
| MH | Can | GRU | 0.1016 | 0.1003 | -0.0013 [-0.0090, 0.0053] |
| MH | Square | MLP | 0.1064 | 0.0951 | -0.0113 [-0.0196, -0.0049] |
| MH | Square | GRU | 0.0891 | 0.0941 | 0.0049 [0.0014, 0.0088] |

Table 13: Five-seed mean standardized response MSE at width eight. History gain equals current-only minus full-history MSE, so positive values favor access to older frames. Intervals describe variation across test episodes for the seed-averaged loss.

The stored-state precision audit casts each width-eight linear state from float64 to float32 and restores float64 only for the fixed reader computation. Across the six PH/MH native-history task conditions and both QRM and PCA, the largest absolute MSE change is $7.13 \times 10^{-10}$; the paired QRM–PCA rankings are unchanged. Neural and cast-linear states each occupy 32 B per history at width eight. Their writers, parameter counts, optimization, and compute costs differ, so stored-state bytes are the single matched resource in this descriptive comparison.

### J.10 PRE-FIXED MH LIFT WRITER-STABILITY TEST

The separate MH Lift test examines the adverse frozen PH Lift result. Its protocol and the covariance-ridge $\lambda = 1$ writer were fixed from the earlier PH diagnostic before reading MH Lift test outcomes. The official MH Lift collection has 300 demonstrations from six demonstrators. The same seed and proportions assign 135/30/90/45 complete demonstrations to train, validation, calibration, and test splits. One 41-frame test demonstration cannot supply both the eight-frame history and a 20-frame future, leaving 44 effective test episodes and 531 test windows. This duration rule is shared by all methods and uses no outcome scores.

The covariance-ridge writer forms the response-weighted cross operator after whitening the training history with $(S + I)^{-1/2}$, keeps its top eight directions, and fits the existing reader with validation-selected ridge. The original QRM, PCA, PLS, and complete current-frame reference use the unchanged benchmark runner. Table 14 gives all fixed comparisons and a same-width current-frame diagnostic outside the fixed Lift improvement criteria. All test scores use training-standardized response targets, equal query-family weights, and equal episode weights.

| Method | State B | Validation MSE | Test MSE |
|---|---|---|---|
| Covariance ridge $\lambda = 1$ | 64 | 0.2386 | 0.2965 |
| Original QRM | 64 | 0.2410 | 0.3030 |
| PCA | 64 | 0.5075 | 0.6042 |
| PLS | 64 | 0.3206 | 0.4017 |
| Complete current frame | 208 | 0.2203 | 0.2672 |
| Current frame, width eight | 64 | 0.2199 | 0.2655 |
| Full-history GRU, five seeds | 32 | – | 0.2759 |

Table 14: MH Lift response MSE. The current-frame width-eight row was added after the fixed ridge test. The GRU entry is the mean of five seeds (SD 0.0043) and uses a different writer and float32 state.

The fixed ridge reduces original QRM loss by 0.0065 [-0.0492, 0.0884], and the complete current frame is better than the ridge writer by 0.0293 [0.0045, 0.0573]. Intervals are paired 95% bootstrap intervals over the 44 effective test episodes, using 5,000 resamples. The same-width current-frame state added after the fixed test is better than the ridge writer by 0.0309 [0.0040, 0.0613]. These intervals describe the observed MH Lift comparison and its uncertainty.

Ordinary QRM improves on PCA response MSE by 0.3012 [0.1917, 0.4132], while its history reconstruction MSE is 0.7088 versus PCA's 0.3555. The MH Lift data therefore exhibit the same response–reconstruction ranking reversal as MH Can and Square, alongside the adverse frozen PH Lift width-eight row. Training uses 1,900 windows from 135 episodes for 208 history features; a 100-resample training-episode bootstrap raises median writer-subspace overlap from 0.789 to 0.888 under ridge one. This stability diagnostic is separate from the held-out endpoint.

The five-seed full-history GRU mean is 0.2759. A current-only GRU rerun conducted after the fixed ridge test has mean 0.2800; its history gain is 0.0041 [-0.0139, 0.0255]. Linear and neural persistent state bytes are compared after the separate float32 linear storage audit; parameter count and training remain different.

### J.11 SAME-COORDINATE QUERY-COLUMN WEIGHTING ABLATION

The unweighted baseline takes an SVD after mapping the fitted operator into raw history coordinates. Its contrast with QRM changes both response-column weighting and the coordinates in which directions are selected. To isolate weighting, the train-fitted history whitener, response operator, compiler ridge, width, splits, and linear reader are held fixed. Only the writer SVD changes, from equal total weight per query family to equal weight per scalar query column. Reader ridge is chosen on validation with the original family-equal evaluation metric.

All three intervals contain zero. The weighted spectral rule retains its defined population objective, while this ablation estimates the incremental effect of family weighting on the recorded tasks. The pre-fixed Can/Square swap and Lift covariance-ridge outcomes remain unchanged.

| MH task | Equal-column MSE | QRM MSE | Equal-column minus QRM [95% CI] |
|---|---|---|---|
| Can | 0.1127 | 0.1130 | -0.0003 [-0.0022, 0.0020] |
| Square | 0.0984 | 0.0994 | -0.0010 [-0.0025, 0.0002] |
| Lift | 0.3264 | 0.3030 | 0.0234 [-0.0066, 0.0734] |

Table 15: Same-coordinate writer-weight ablation at width eight. Positive differences favor family-equal QRM. Intervals are trajectory-paired bootstrap intervals over 45, 45, and 44 effective test episodes, respectively; the analysis was specified after earlier test results.

### J.12 Complete external method–width table

The following long table retains every row of the 342-row frozen main-results table, including control methods, all declared linear widths, and all neural widths. Response and reconstruction columns are test means; the state-byte column records persistent state size, with a dash for methods without a recorded value, and seeds are shown as successful/expected. The intervention and cost summaries report the corresponding row-level estimates and uncertainty intervals.

Table 16: All frozen robomimic main-result rows.

| Condition | Runner | Method | Width | State B | Response | Recon. | Seeds |
|---|---|---|---|---|---|---|---|
| masked_can | linear | full_history | 248 | 1984 | 0.1805 | 0.0000 | 1/1 |
| masked_can | linear | last_frame | 2 | 16 | 0.4226 | 0.4940 | 1/1 |
| masked_can | linear | last_frame | 4 | 32 | 0.3141 | 0.3326 | 1/1 |
| masked_can | linear | last_frame | 8 | 64 | 0.2850 | 0.2662 | 1/1 |
| masked_can | linear | last_frame | 16 | 128 | 0.2829 | 0.2301 | 1/1 |
| masked_can | linear | last_frame_full | 31 | 248 | 0.2829 | 0.2301 | 1/1 |
| masked_can | linear | no_memory | 0 | 0 | 0.9890 | 0.7322 | 1/1 |
| masked_can | linear | pca | 2 | 16 | 0.5205 | 0.3953 | 1/1 |
| masked_can | linear | pca | 4 | 32 | 0.4109 | 0.2441 | 1/1 |
| masked_can | linear | pca | 8 | 64 | 0.3426 | 0.1283 | 1/1 |
| masked_can | linear | pca | 16 | 128 | 0.2937 | 0.0524 | 1/1 |
| masked_can | linear | pca | 32 | 256 | 0.2363 | 0.0126 | 1/1 |
| masked_can | linear | pls | 2 | 16 | 0.4610 | 0.3975 | 1/1 |
| masked_can | linear | pls | 4 | 32 | 0.3761 | 0.2518 | 1/1 |
| masked_can | linear | pls | 8 | 64 | 0.2856 | 0.1567 | 1/1 |
| masked_can | linear | pls | 16 | 128 | 0.2377 | 0.0708 | 1/1 |
| masked_can | linear | pls | 32 | 256 | 0.1995 | 0.0161 | 1/1 |
| masked_can | linear | qrm | 2 | 16 | 0.3892 | 0.4939 | 1/1 |
| masked_can | linear | qrm | 4 | 32 | 0.2695 | 0.3610 | 1/1 |
| masked_can | linear | qrm | 8 | 64 | 0.1850 | 0.2849 | 1/1 |
| masked_can | linear | qrm | 16 | 128 | 0.1804 | 0.2470 | 1/1 |
| masked_can | linear | qrm | 32 | 256 | 0.1805 | 0.2008 | 1/1 |
| masked_can | linear | random | 2 | 16 | 0.6890 | 0.5724 | 1/1 |
| masked_can | linear | random | 4 | 32 | 0.5982 | 0.4105 | 1/1 |
| masked_can | linear | random | 8 | 64 | 0.4153 | 0.2510 | 1/1 |
| masked_can | linear | random | 16 | 128 | 0.3214 | 0.1261 | 1/1 |
| masked_can | linear | random | 32 | 256 | 0.2556 | 0.0437 | 1/1 |
| masked_can | linear | separate_banks | 2 | 16 | 0.5400 | – | 1/1 |
| masked_can | linear | separate_banks | 4 | 32 | 0.3473 | – | 1/1 |
| masked_can | linear | separate_banks | 8 | 64 | 0.2098 | – | 1/1 |
| masked_can | linear | separate_banks | 16 | 128 | 0.1770 | – | 1/1 |
| masked_can | linear | separate_banks | 32 | 256 | 0.1767 | – | 1/1 |
| masked_can | linear | unweighted | 2 | 16 | 0.6737 | 0.6340 | 1/1 |
| masked_can | linear | unweighted | 4 | 32 | 0.4577 | 0.5498 | 1/1 |
| masked_can | linear | unweighted | 8 | 64 | 0.2789 | 0.3995 | 1/1 |
| masked_can | linear | unweighted | 16 | 128 | 0.1863 | 0.2797 | 1/1 |
| masked_can | linear | unweighted | 32 | 256 | 0.1729 | 0.1441 | 1/1 |
| masked_can | neural | autoencoder | 2 | – | 0.4369 | 0.4481 | 5/5 |
| masked_can | neural | autoencoder | 4 | – | 0.2985 | 0.2947 | 5/5 |
| masked_can | neural | autoencoder | 8 | – | 0.2607 | 0.1766 | 5/5 |
| masked_can | neural | autoencoder | 16 | – | 0.2266 | 0.0954 | 5/5 |
| masked_can | neural | autoencoder | 32 | – | 0.2157 | 0.0351 | 5/5 |
| masked_can | neural | gru | 2 | – | 0.2524 | 0.4864 | 5/5 |
| masked_can | neural | gru | 4 | – | 0.2067 | 0.3585 | 5/5 |
| masked_can | neural | gru | 8 | – | 0.1847 | 0.2732 | 5/5 |
| masked_can | neural | gru | 16 | – | 0.1744 | 0.1998 | 5/5 |
| masked_can | neural | gru | 32 | – | 0.1714 | 0.1358 | 5/5 |
| masked_can | neural | nested | 2 | – | 0.2644 | 0.4795 | 5/5 |
| masked_can | neural | nested | 4 | – | 0.2031 | 0.3734 | 5/5 |
| masked_can | neural | nested | 8 | – | 0.1834 | 0.2811 | 5/5 |
| masked_can | neural | nested | 16 | – | 0.1818 | 0.2134 | 5/5 |
| masked_can | neural | nested | 32 | – | 0.1836 | 0.1468 | 5/5 |
| masked_can | neural | supervised | 2 | – | 0.2615 | 0.5109 | 5/5 |
| masked_can | neural | supervised | 4 | – | 0.1965 | 0.3695 | 5/5 |
| masked_can | neural | supervised | 8 | – | 0.1759 | 0.2789 | 5/5 |
| masked_can | neural | supervised | 16 | – | 0.1783 | 0.2078 | 5/5 |
| masked_can | neural | supervised | 32 | – | 0.1763 | 0.1350 | 5/5 |
| recorded_can | linear | full_history | 240 | 1920 | 0.1428 | 0.0000 | 1/1 |
| recorded_can | linear | last_frame | 2 | 16 | 0.3801 | 0.6793 | 1/1 |
| recorded_can | linear | last_frame | 4 | 32 | 0.2547 | 0.5195 | 1/1 |
| recorded_can | linear | last_frame | 8 | 64 | 0.1530 | 0.3998 | 1/1 |
| recorded_can | linear | last_frame | 16 | 128 | 0.1477 | 0.3392 | 1/1 |
| recorded_can | linear | last_frame_full | 30 | 240 | 0.1477 | 0.1509 | 1/1 |

| Condition | Runner | Method | Width | State B | Response | Recon. | Seeds |
|---|---|---|---|---|---|---|---|
| recorded_can | linear | no_memory | 0 | 0 | 0.9890 | 0.9899 | 1/1 |
| recorded_can | linear | pca | 2 | 16 | 0.5016 | 0.5790 | 1/1 |
| recorded_can | linear | pca | 4 | 32 | 0.4168 | 0.3923 | 1/1 |
| recorded_can | linear | pca | 8 | 64 | 0.3198 | 0.2069 | 1/1 |
| recorded_can | linear | pca | 16 | 128 | 0.2635 | 0.0928 | 1/1 |
| recorded_can | linear | pca | 32 | 256 | 0.1966 | 0.0292 | 1/1 |
| recorded_can | linear | pls | 2 | 16 | 0.4517 | 0.5826 | 1/1 |
| recorded_can | linear | pls | 4 | 32 | 0.3425 | 0.4115 | 1/1 |
| recorded_can | linear | pls | 8 | 64 | 0.2535 | 0.2501 | 1/1 |
| recorded_can | linear | pls | 16 | 128 | 0.1920 | 0.1219 | 1/1 |
| recorded_can | linear | pls | 32 | 256 | 0.1601 | 0.0362 | 1/1 |
| recorded_can | linear | qrm | 2 | 16 | 0.3800 | 0.6858 | 1/1 |
| recorded_can | linear | qrm | 4 | 32 | 0.2517 | 0.5322 | 1/1 |
| recorded_can | linear | qrm | 8 | 64 | 0.1479 | 0.4306 | 1/1 |
| recorded_can | linear | qrm | 16 | 128 | 0.1424 | 0.3856 | 1/1 |
| recorded_can | linear | qrm | 32 | 256 | 0.1428 | 0.3316 | 1/1 |
| recorded_can | linear | random | 2 | 16 | 0.8432 | 0.7390 | 1/1 |
| recorded_can | linear | random | 4 | 32 | 0.6673 | 0.5876 | 1/1 |
| recorded_can | linear | random | 8 | 64 | 0.4320 | 0.4124 | 1/1 |
| recorded_can | linear | random | 16 | 128 | 0.3335 | 0.2067 | 1/1 |
| recorded_can | linear | random | 32 | 256 | 0.2414 | 0.0882 | 1/1 |
| recorded_can | linear | separate_banks | 2 | 16 | 0.5392 | – | 1/1 |
| recorded_can | linear | separate_banks | 4 | 32 | 0.3298 | – | 1/1 |
| recorded_can | linear | separate_banks | 8 | 64 | 0.1715 | – | 1/1 |
| recorded_can | linear | separate_banks | 16 | 128 | 0.1409 | – | 1/1 |
| recorded_can | linear | separate_banks | 32 | 256 | 0.1404 | – | 1/1 |
| recorded_can | linear | unweighted | 2 | 16 | 0.7882 | 0.8741 | 1/1 |
| recorded_can | linear | unweighted | 4 | 32 | 0.3768 | 0.7036 | 1/1 |
| recorded_can | linear | unweighted | 8 | 64 | 0.2385 | 0.5913 | 1/1 |
| recorded_can | linear | unweighted | 16 | 128 | 0.1466 | 0.4157 | 1/1 |
| recorded_can | linear | unweighted | 32 | 256 | 0.1363 | 0.2327 | 1/1 |
| recorded_can | neural | autoencoder | 2 | – | 0.4007 | 0.6266 | 5/5 |
| recorded_can | neural | autoencoder | 4 | – | 0.3067 | 0.4703 | 5/5 |
| recorded_can | neural | autoencoder | 8 | – | 0.2472 | 0.2751 | 5/5 |
| recorded_can | neural | autoencoder | 16 | – | 0.2059 | 0.1481 | 5/5 |
| recorded_can | neural | autoencoder | 32 | – | 0.1880 | 0.0608 | 5/5 |
| recorded_can | neural | gru | 2 | – | 0.2477 | 0.6993 | 5/5 |
| recorded_can | neural | gru | 4 | – | 0.1719 | 0.5290 | 5/5 |
| recorded_can | neural | gru | 8 | – | 0.1443 | 0.4070 | 5/5 |
| recorded_can | neural | gru | 16 | – | 0.1378 | 0.2714 | 5/5 |
| recorded_can | neural | gru | 32 | – | 0.1335 | 0.1741 | 5/5 |
| recorded_can | neural | nested | 2 | – | 0.2661 | 0.6924 | 5/5 |
| recorded_can | neural | nested | 4 | – | 0.1908 | 0.5497 | 5/5 |
| recorded_can | neural | nested | 8 | – | 0.1652 | 0.4357 | 5/5 |
| recorded_can | neural | nested | 16 | – | 0.1604 | 0.3397 | 5/5 |
| recorded_can | neural | nested | 32 | – | 0.1633 | 0.2338 | 5/5 |
| recorded_can | neural | supervised | 2 | – | 0.2562 | 0.7066 | 5/5 |
| recorded_can | neural | supervised | 4 | – | 0.1892 | 0.5400 | 5/5 |
| recorded_can | neural | supervised | 8 | – | 0.1598 | 0.4184 | 5/5 |
| recorded_can | neural | supervised | 16 | – | 0.1581 | 0.3219 | 5/5 |
| recorded_can | neural | supervised | 32 | – | 0.1529 | 0.2113 | 5/5 |
| masked_lift | linear | full_history | 216 | 1728 | 1.0907 | 0.0000 | 1/1 |
| masked_lift | linear | last_frame | 2 | 16 | 0.6082 | 0.7110 | 1/1 |
| masked_lift | linear | last_frame | 4 | 32 | 0.4531 | 0.5815 | 1/1 |
| masked_lift | linear | last_frame | 8 | 64 | 0.4975 | 0.4995 | 1/1 |
| masked_lift | linear | last_frame | 16 | 128 | 0.5013 | 0.3405 | 1/1 |
| masked_lift | linear | last_frame_full | 27 | 216 | 0.5013 | 0.3405 | 1/1 |
| masked_lift | linear | no_memory | 0 | 0 | 0.9382 | 0.9121 | 1/1 |
| masked_lift | linear | pca | 2 | 16 | 0.7658 | 0.4839 | 1/1 |
| masked_lift | linear | pca | 4 | 32 | 0.5893 | 0.4048 | 1/1 |
| masked_lift | linear | pca | 8 | 64 | 0.4410 | 0.2682 | 1/1 |
| masked_lift | linear | pca | 16 | 128 | 0.3263 | 0.1128 | 1/1 |
| masked_lift | linear | pca | 32 | 256 | 0.3739 | 0.0098 | 1/1 |
| masked_lift | linear | pls | 2 | 16 | 0.6404 | 0.6023 | 1/1 |
| masked_lift | linear | pls | 4 | 32 | 0.4454 | 0.4175 | 1/1 |
| masked_lift | linear | pls | 8 | 64 | 0.3224 | 0.3089 | 1/1 |
| masked_lift | linear | pls | 16 | 128 | 0.4339 | 0.1399 | 1/1 |
| masked_lift | linear | pls | 32 | 256 | 0.3572 | 0.0125 | 1/1 |
| masked_lift | linear | qrm | 2 | 16 | 0.6034 | 0.8194 | 1/1 |
| masked_lift | linear | qrm | 4 | 32 | 0.5637 | 0.7729 | 1/1 |
| masked_lift | linear | qrm | 8 | 64 | 0.7478 | 0.6530 | 1/1 |
| masked_lift | linear | qrm | 16 | 128 | 1.0522 | 0.6364 | 1/1 |
| masked_lift | linear | qrm | 32 | 256 | 1.0624 | 0.8993 | 1/1 |
| masked_lift | linear | random | 2 | 16 | 0.8458 | 0.7174 | 1/1 |
| masked_lift | linear | random | 4 | 32 | 0.7360 | 0.6584 | 1/1 |
| masked_lift | linear | random | 8 | 64 | 0.5447 | 0.4455 | 1/1 |
| masked_lift | linear | random | 16 | 128 | 0.4259 | 0.2547 | 1/1 |
| masked_lift | linear | random | 32 | 256 | 0.4218 | 0.0450 | 1/1 |
| masked_lift | linear | separate_banks | 2 | 16 | 0.7536 | – | 1/1 |
| masked_lift | linear | separate_banks | 4 | 32 | 0.5526 | – | 1/1 |
| masked_lift | linear | separate_banks | 8 | 64 | 0.4894 | – | 1/1 |
| masked_lift | linear | separate_banks | 16 | 128 | 0.5539 | – | 1/1 |
| masked_lift | linear | separate_banks | 32 | 256 | 0.5539 | – | 1/1 |
| masked_lift | linear | unweighted | 2 | 16 | 1.5001 | 1.0731 | 1/1 |
| masked_lift | linear | unweighted | 4 | 32 | 1.6114 | 1.0971 | 1/1 |
| masked_lift | linear | unweighted | 8 | 64 | 1.3013 | 0.9857 | 1/1 |
| masked_lift | linear | unweighted | 16 | 128 | 0.9643 | 0.8171 | 1/1 |
| masked_lift | linear | unweighted | 32 | 256 | 0.8593 | 0.6026 | 1/1 |
| masked_lift | neural | autoencoder | 2 | – | 0.7559 | 0.5004 | 5/5 |
| masked_lift | neural | autoencoder | 4 | – | 0.6018 | 0.4167 | 5/5 |

| Condition | Runner | Method | Width | State B | Response | Recon. | Seeds |
|---|---|---|---|---|---|---|---|
| masked_lift | neural | autoencoder | 8 | – | 0.5471 | 0.3177 | 5/5 |
| masked_lift | neural | autoencoder | 16 | – | 0.5016 | 0.2377 | 5/5 |
| masked_lift | neural | autoencoder | 32 | – | 0.4974 | 0.2170 | 5/5 |
| masked_lift | neural | gru | 2 | – | 0.5875 | 0.7009 | 5/5 |
| masked_lift | neural | gru | 4 | – | 0.4324 | 0.5709 | 5/5 |
| masked_lift | neural | gru | 8 | – | 0.4245 | 0.3875 | 5/5 |
| masked_lift | neural | gru | 16 | – | 0.4174 | 0.3181 | 5/5 |
| masked_lift | neural | gru | 32 | – | 0.4211 | 0.2779 | 5/5 |
| masked_lift | neural | nested | 2 | – | 0.6084 | 0.7417 | 5/5 |
| masked_lift | neural | nested | 4 | – | 0.4474 | 0.6027 | 5/5 |
| masked_lift | neural | nested | 8 | – | 0.4332 | 0.4925 | 5/5 |
| masked_lift | neural | nested | 16 | – | 0.4223 | 0.3540 | 5/5 |
| masked_lift | neural | nested | 32 | – | 0.4208 | 0.2850 | 5/5 |
| masked_lift | neural | supervised | 2 | – | 0.5986 | 0.7125 | 5/5 |
| masked_lift | neural | supervised | 4 | – | 0.4446 | 0.5861 | 5/5 |
| masked_lift | neural | supervised | 8 | – | 0.4302 | 0.4368 | 5/5 |
| masked_lift | neural | supervised | 16 | – | 0.4335 | 0.3393 | 5/5 |
| masked_lift | neural | supervised | 32 | – | 0.4111 | 0.2877 | 5/5 |
| recorded_lift | linear | full_history | 208 | 1664 | 1.0736 | 0.0000 | 1/1 |
| recorded_lift | linear | last_frame | 2 | 16 | 0.5624 | 0.9537 | 1/1 |
| recorded_lift | linear | last_frame | 4 | 32 | 0.3651 | 0.7349 | 1/1 |
| recorded_lift | linear | last_frame | 8 | 64 | 0.2873 | 0.6178 | 1/1 |
| recorded_lift | linear | last_frame | 16 | 128 | 0.2872 | 0.4363 | 1/1 |
| recorded_lift | linear | last_frame_full | 26 | 208 | 0.2840 | 0.2919 | 1/1 |
| recorded_lift | linear | no_memory | 0 | 0 | 0.9382 | 1.1019 | 1/1 |
| recorded_lift | linear | pca | 2 | 16 | 0.7454 | 0.6416 | 1/1 |
| recorded_lift | linear | pca | 4 | 32 | 0.5036 | 0.5188 | 1/1 |
| recorded_lift | linear | pca | 8 | 64 | 0.4244 | 0.3727 | 1/1 |
| recorded_lift | linear | pca | 16 | 128 | 0.3010 | 0.1467 | 1/1 |
| recorded_lift | linear | pca | 32 | 256 | 0.3410 | 0.0145 | 1/1 |
| recorded_lift | linear | pls | 2 | 16 | 0.6145 | 0.7484 | 1/1 |
| recorded_lift | linear | pls | 4 | 32 | 0.4283 | 0.5303 | 1/1 |
| recorded_lift | linear | pls | 8 | 64 | 0.3335 | 0.3926 | 1/1 |
| recorded_lift | linear | pls | 16 | 128 | 0.3220 | 0.1725 | 1/1 |
| recorded_lift | linear | pls | 32 | 256 | 0.3119 | 0.0185 | 1/1 |
| recorded_lift | linear | qrm | 2 | 16 | 0.6100 | 0.9686 | 1/1 |
| recorded_lift | linear | qrm | 4 | 32 | 0.5701 | 0.9201 | 1/1 |
| recorded_lift | linear | qrm | 8 | 64 | 0.7566 | 0.7139 | 1/1 |
| recorded_lift | linear | qrm | 16 | 128 | 1.0289 | 0.7242 | 1/1 |
| recorded_lift | linear | qrm | 32 | 256 | 1.0388 | 0.8432 | 1/1 |
| recorded_lift | linear | random | 2 | 16 | 0.7671 | 0.8467 | 1/1 |
| recorded_lift | linear | random | 4 | 32 | 0.7422 | 0.7687 | 1/1 |
| recorded_lift | linear | random | 8 | 64 | 0.5907 | 0.5798 | 1/1 |
| recorded_lift | linear | random | 16 | 128 | 0.4902 | 0.3962 | 1/1 |
| recorded_lift | linear | random | 32 | 256 | 0.3894 | 0.0777 | 1/1 |
| recorded_lift | linear | separate_banks | 2 | 16 | 0.7557 | – | 1/1 |
| recorded_lift | linear | separate_banks | 4 | 32 | 0.5530 | – | 1/1 |
| recorded_lift | linear | separate_banks | 8 | 64 | 0.5245 | – | 1/1 |
| recorded_lift | linear | separate_banks | 16 | 128 | 0.5501 | – | 1/1 |
| recorded_lift | linear | separate_banks | 32 | 256 | 0.5501 | – | 1/1 |
| recorded_lift | linear | unweighted | 2 | 16 | 1.3293 | 1.2222 | 1/1 |
| recorded_lift | linear | unweighted | 4 | 32 | 1.3709 | 1.2079 | 1/1 |
| recorded_lift | linear | unweighted | 8 | 64 | 1.3352 | 1.1303 | 1/1 |
| recorded_lift | linear | unweighted | 16 | 128 | 1.0124 | 1.0116 | 1/1 |
| recorded_lift | linear | unweighted | 32 | 256 | 0.8691 | 0.7458 | 1/1 |
| recorded_lift | neural | autoencoder | 2 | – | 0.6953 | 0.6946 | 5/5 |
| recorded_lift | neural | autoencoder | 4 | – | 0.5588 | 0.5361 | 5/5 |
| recorded_lift | neural | autoencoder | 8 | – | 0.4837 | 0.4129 | 5/5 |
| recorded_lift | neural | autoencoder | 16 | – | 0.4653 | 0.3079 | 5/5 |
| recorded_lift | neural | autoencoder | 32 | – | 0.4303 | 0.2533 | 5/5 |
| recorded_lift | neural | gru | 2 | – | 0.5880 | 0.8835 | 5/5 |
| recorded_lift | neural | gru | 4 | – | 0.3843 | 0.7148 | 5/5 |
| recorded_lift | neural | gru | 8 | – | 0.3488 | 0.5119 | 5/5 |
| recorded_lift | neural | gru | 16 | – | 0.3371 | 0.3756 | 5/5 |
| recorded_lift | neural | gru | 32 | – | 0.3244 | 0.3043 | 5/5 |
| recorded_lift | neural | nested | 2 | – | 0.6093 | 0.9004 | 5/5 |
| recorded_lift | neural | nested | 4 | – | 0.4337 | 0.7337 | 5/5 |
| recorded_lift | neural | nested | 8 | – | 0.4027 | 0.6056 | 5/5 |
| recorded_lift | neural | nested | 16 | – | 0.3963 | 0.4939 | 5/5 |
| recorded_lift | neural | nested | 32 | – | 0.3795 | 0.3882 | 5/5 |
| recorded_lift | neural | supervised | 2 | – | 0.5875 | 0.8474 | 5/5 |
| recorded_lift | neural | supervised | 4 | – | 0.4193 | 0.7199 | 5/5 |
| recorded_lift | neural | supervised | 8 | – | 0.3803 | 0.5691 | 5/5 |
| recorded_lift | neural | supervised | 16 | – | 0.3511 | 0.4534 | 5/5 |
| recorded_lift | neural | supervised | 32 | – | 0.3578 | 0.3708 | 5/5 |
| masked_square | linear | full_history | 248 | 1984 | 0.1516 | 0.0000 | 1/1 |
| masked_square | linear | last_frame | 2 | 16 | 0.4512 | 0.4826 | 1/1 |
| masked_square | linear | last_frame | 4 | 32 | 0.3277 | 0.3894 | 1/1 |
| masked_square | linear | last_frame | 8 | 64 | 0.3019 | 0.2892 | 1/1 |
| masked_square | linear | last_frame | 16 | 128 | 0.3009 | 0.1840 | 1/1 |
| masked_square | linear | last_frame_full | 31 | 248 | 0.3009 | 0.1840 | 1/1 |
| masked_square | linear | no_memory | 0 | 0 | 0.9661 | 0.6827 | 1/1 |
| masked_square | linear | pca | 2 | 16 | 0.6385 | 0.3692 | 1/1 |
| masked_square | linear | pca | 4 | 32 | 0.4719 | 0.2543 | 1/1 |
| masked_square | linear | pca | 8 | 64 | 0.3060 | 0.1493 | 1/1 |
| masked_square | linear | pca | 16 | 128 | 0.2459 | 0.0614 | 1/1 |
| masked_square | linear | pca | 32 | 256 | 0.2012 | 0.0164 | 1/1 |
| masked_square | linear | pls | 2 | 16 | 0.5331 | 0.3879 | 1/1 |
| masked_square | linear | pls | 4 | 32 | 0.3566 | 0.2973 | 1/1 |
| masked_square | linear | pls | 8 | 64 | 0.2593 | 0.1518 | 1/1 |

| Condition | Runner | Method | Width | State B | Response | Recon. | Seeds |
|---|---|---|---|---|---|---|---|
| masked_square | linear | pls | 16 | 128 | 0.2055 | 0.0739 | 1/1 |
| masked_square | linear | pls | 32 | 256 | 0.1780 | 0.0215 | 1/1 |
| masked_square | linear | qrm | 2 | 16 | 0.4302 | 0.5014 | 1/1 |
| masked_square | linear | qrm | 4 | 32 | 0.2349 | 0.4221 | 1/1 |
| masked_square | linear | qrm | 8 | 64 | 0.1584 | 0.3169 | 1/1 |
| masked_square | linear | qrm | 16 | 128 | 0.1517 | 0.2810 | 1/1 |
| masked_square | linear | qrm | 32 | 256 | 0.1516 | 0.2525 | 1/1 |
| masked_square | linear | random | 2 | 16 | 0.8305 | 0.5219 | 1/1 |
| masked_square | linear | random | 4 | 32 | 0.6357 | 0.3978 | 1/1 |
| masked_square | linear | random | 8 | 64 | 0.5393 | 0.3029 | 1/1 |
| masked_square | linear | random | 16 | 128 | 0.3274 | 0.1687 | 1/1 |
| masked_square | linear | random | 32 | 256 | 0.2351 | 0.0548 | 1/1 |
| masked_square | linear | separate_banks | 2 | 16 | 0.5083 | – | 1/1 |
| masked_square | linear | separate_banks | 4 | 32 | 0.3039 | – | 1/1 |
| masked_square | linear | separate_banks | 8 | 64 | 0.1767 | – | 1/1 |
| masked_square | linear | separate_banks | 16 | 128 | 0.1521 | – | 1/1 |
| masked_square | linear | separate_banks | 32 | 256 | 0.1516 | – | 1/1 |
| masked_square | linear | unweighted | 2 | 16 | 0.7772 | 0.6040 | 1/1 |
| masked_square | linear | unweighted | 4 | 32 | 0.4833 | 0.5377 | 1/1 |
| masked_square | linear | unweighted | 8 | 64 | 0.1929 | 0.3890 | 1/1 |
| masked_square | linear | unweighted | 16 | 128 | 0.1535 | 0.2857 | 1/1 |
| masked_square | linear | unweighted | 32 | 256 | 0.1511 | 0.2080 | 1/1 |
| masked_square | neural | autoencoder | 2 | – | 0.3552 | 0.4120 | 5/5 |
| masked_square | neural | autoencoder | 4 | – | 0.2632 | 0.3068 | 5/5 |
| masked_square | neural | autoencoder | 8 | – | 0.2024 | 0.1936 | 5/5 |
| masked_square | neural | autoencoder | 16 | – | 0.1846 | 0.1031 | 5/5 |
| masked_square | neural | autoencoder | 32 | – | 0.1754 | 0.0379 | 5/5 |
| masked_square | neural | gru | 2 | – | 0.2290 | 0.5054 | 5/5 |
| masked_square | neural | gru | 4 | – | 0.1647 | 0.3915 | 5/5 |
| masked_square | neural | gru | 8 | – | 0.1501 | 0.2462 | 5/5 |
| masked_square | neural | gru | 16 | – | 0.1472 | 0.1737 | 5/5 |
| masked_square | neural | gru | 32 | – | 0.1412 | 0.1158 | 5/5 |
| masked_square | neural | nested | 2 | – | 0.2233 | 0.5045 | 5/5 |
| masked_square | neural | nested | 4 | – | 0.1635 | 0.3821 | 5/5 |
| masked_square | neural | nested | 8 | – | 0.1494 | 0.2995 | 5/5 |
| masked_square | neural | nested | 16 | – | 0.1491 | 0.2093 | 5/5 |
| masked_square | neural | nested | 32 | – | 0.1487 | 0.1368 | 5/5 |
| masked_square | neural | supervised | 2 | – | 0.2284 | 0.5046 | 5/5 |
| masked_square | neural | supervised | 4 | – | 0.1621 | 0.3983 | 5/5 |
| masked_square | neural | supervised | 8 | – | 0.1475 | 0.2930 | 5/5 |
| masked_square | neural | supervised | 16 | – | 0.1481 | 0.1914 | 5/5 |
| masked_square | neural | supervised | 32 | – | 0.1444 | 0.1260 | 5/5 |
| recorded_square | linear | full_history | 240 | 1920 | 0.1186 | 0.0000 | 1/1 |
| recorded_square | linear | last_frame | 2 | 16 | 0.4267 | 0.6709 | 1/1 |
| recorded_square | linear | last_frame | 4 | 32 | 0.2225 | 0.5471 | 1/1 |
| recorded_square | linear | last_frame | 8 | 64 | 0.1316 | 0.4129 | 1/1 |
| recorded_square | linear | last_frame | 16 | 128 | 0.1274 | 0.3268 | 1/1 |
| recorded_square | linear | last_frame_full | 30 | 240 | 0.1273 | 0.1488 | 1/1 |
| recorded_square | linear | no_memory | 0 | 0 | 0.9661 | 0.9263 | 1/1 |
| recorded_square | linear | pca | 2 | 16 | 0.6015 | 0.5470 | 1/1 |
| recorded_square | linear | pca | 4 | 32 | 0.4453 | 0.3780 | 1/1 |
| recorded_square | linear | pca | 8 | 64 | 0.2961 | 0.2467 | 1/1 |
| recorded_square | linear | pca | 16 | 128 | 0.2361 | 0.1082 | 1/1 |
| recorded_square | linear | pca | 32 | 256 | 0.1713 | 0.0364 | 1/1 |
| recorded_square | linear | pls | 2 | 16 | 0.5288 | 0.5629 | 1/1 |
| recorded_square | linear | pls | 4 | 32 | 0.3387 | 0.4136 | 1/1 |
| recorded_square | linear | pls | 8 | 64 | 0.2253 | 0.2527 | 1/1 |
| recorded_square | linear | pls | 16 | 128 | 0.1721 | 0.1316 | 1/1 |
| recorded_square | linear | pls | 32 | 256 | 0.1407 | 0.0493 | 1/1 |
| recorded_square | linear | qrm | 2 | 16 | 0.4261 | 0.6737 | 1/1 |
| recorded_square | linear | qrm | 4 | 32 | 0.2195 | 0.5560 | 1/1 |
| recorded_square | linear | qrm | 8 | 64 | 0.1251 | 0.4268 | 1/1 |
| recorded_square | linear | qrm | 16 | 128 | 0.1187 | 0.3821 | 1/1 |
| recorded_square | linear | qrm | 32 | 256 | 0.1186 | 0.3576 | 1/1 |
| recorded_square | linear | random | 2 | 16 | 0.9286 | 0.7406 | 1/1 |
| recorded_square | linear | random | 4 | 32 | 0.6876 | 0.5834 | 1/1 |
| recorded_square | linear | random | 8 | 64 | 0.4842 | 0.3967 | 1/1 |
| recorded_square | linear | random | 16 | 128 | 0.3550 | 0.2599 | 1/1 |
| recorded_square | linear | random | 32 | 256 | 0.2512 | 0.1122 | 1/1 |
| recorded_square | linear | separate_banks | 2 | 16 | 0.5070 | – | 1/1 |
| recorded_square | linear | separate_banks | 4 | 32 | 0.2896 | – | 1/1 |
| recorded_square | linear | separate_banks | 8 | 64 | 0.1442 | – | 1/1 |
| recorded_square | linear | separate_banks | 16 | 128 | 0.1190 | – | 1/1 |
| recorded_square | linear | separate_banks | 32 | 256 | 0.1186 | – | 1/1 |
| recorded_square | linear | unweighted | 2 | 16 | 0.5703 | 0.7923 | 1/1 |
| recorded_square | linear | unweighted | 4 | 32 | 0.4000 | 0.7081 | 1/1 |
| recorded_square | linear | unweighted | 8 | 64 | 0.1693 | 0.4849 | 1/1 |
| recorded_square | linear | unweighted | 16 | 128 | 0.1200 | 0.3876 | 1/1 |
| recorded_square | linear | unweighted | 32 | 256 | 0.1178 | 0.3093 | 1/1 |
| recorded_square | neural | autoencoder | 2 | – | 0.3467 | 0.6015 | 5/5 |
| recorded_square | neural | autoencoder | 4 | – | 0.2533 | 0.4559 | 5/5 |
| recorded_square | neural | autoencoder | 8 | – | 0.1900 | 0.3029 | 5/5 |
| recorded_square | neural | autoencoder | 16 | – | 0.1685 | 0.1743 | 5/5 |
| recorded_square | neural | autoencoder | 32 | – | 0.1552 | 0.0704 | 5/5 |
| recorded_square | neural | gru | 2 | – | 0.2218 | 0.6736 | 5/5 |
| recorded_square | neural | gru | 4 | – | 0.1500 | 0.5331 | 5/5 |
| recorded_square | neural | gru | 8 | – | 0.1250 | 0.3748 | 5/5 |
| recorded_square | neural | gru | 16 | – | 0.1231 | 0.2514 | 5/5 |
| recorded_square | neural | gru | 32 | – | 0.1197 | 0.1633 | 5/5 |
| recorded_square | neural | nested | 2 | – | 0.2327 | 0.6795 | 5/5 |

| Condition | Runner | Method | Width | State B | Response | Recon. | Seeds |
|---|---|---|---|---|---|---|---|
| recorded_square | neural | nested | 4 | – | 0.1545 | 0.5191 | 5/5 |
| recorded_square | neural | nested | 8 | – | 0.1324 | 0.4000 | 5/5 |
| recorded_square | neural | nested | 16 | – | 0.1302 | 0.2973 | 5/5 |
| recorded_square | neural | nested | 32 | – | 0.1292 | 0.2133 | 5/5 |
| recorded_square | neural | supervised | 2 | – | 0.2240 | 0.6827 | 5/5 |
| recorded_square | neural | supervised | 4 | – | 0.1506 | 0.5313 | 5/5 |
| recorded_square | neural | supervised | 8 | – | 0.1315 | 0.4153 | 5/5 |
| recorded_square | neural | supervised | 16 | – | 0.1286 | 0.2708 | 5/5 |
| recorded_square | neural | supervised | 32 | – | 0.1289 | 0.1922 | 5/5 |

The frozen results also contain seed standard deviations, nominal float32 bytes, certificate fields, source-row identifiers, and QRM–PCA deltas. This printed table presents the corresponding rows in a readable form.

### J.13 Demonstrator-cluster sensitivity

The multi-human intervals above weight test episodes equally and condition on the observed six-demonstrator mixture. A sensitivity analysis resampled demonstrators as clusters while preserving the episode-weighted estimand. Across 5,000 cluster bootstrap replicates, the QRM–PCA response gain interval was $[0.1696, 0.3571]$ for Can and $[0.2173, 0.2871]$ for Square. The selected two-direction exchange interval was $[-0.0117, 0.0434]$ for Can and $[0.0145, 0.0272]$ for Square. Thus the QRM–PCA comparison remains positive for both tasks, the selected Square exchange remains positive, and the selected Can exchange remains compatible with zero. These intervals are conditional on the observed demonstrator mixture; transfer to new demonstrators requires additional data.