SOACT: SQUARE-FREE ORTHOGONAL ACTION COMPOSITION AND TRANSPORT FOR BRAIN-FREE ANSWER PREDICTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Brain-guided language models align neural representations, but alignment alone does not specify how registered actions change a latent state or compose. SOACT predicts fixed-vocabulary answers from text using a minimal square-free quotient and registered action transport under an external action plan. It estimates nuisance-invariant moments from same-question cross-subject views using square-free products, fits polynomial action generators in a shared chart with no free intercept, composes ordered paths by substitution, and trains a text readout with the generators frozen. Independent composed- action anchors provide a proposed score-calibrated factorization test, with abstention when the evidence is unresolved. Under centered conditional noise and cross-subject independence, we show that square-free moments recover latent polynomial moments; a two-world lower bound identifies information lost by row-separable temporal objectives. We also characterize the degree identified by finite view panels and derive a scoped $2d$ versus $2d^2$ view-count comparison for factorized and primitive path fitting. The experiments use fixed-seed synthetic latent states, paired noisy views, action compositions, and matched comparator protocols. The reported values are evaluated under the declared experimental contracts. The task, ablation, gate, and sensitivity evaluations show the performance of SOACT and the matched non-emulating baseline under the stated protocols.

1 INTRODUCTION

Brain-guided language modeling seeks neural constraints that improve machine reasoning. Recent work uses brain-derived signals to guide language models beyond representational alignment (Xiao et al., 2026). We study a narrower, testable endpoint: *brain-free fixed-vocabulary answer classification under an externally registered action plan*. The question is whether a transferable constraint can specify how an action changes a latent state and how successive actions compose, rather than only matching a static endpoint. Repeated measurements also vary across subjects, so an unconstrained nonlinear learner can fit subject-specific noise instead of shared transitions.

SOACT addresses these two problems with one quotient–transport sequence: *construct nuisance-invariant moments, transport registered actions, and transfer the resulting dynamics to text*. We study repeated, time-resolved measurements of the same question under registered actions. The public trial-level beta artifacts available to this project do not by themselves instantiate this observation contract; this does not imply that raw time series or task-event files are unavailable in the source dataset. Distinct-subject products use each subject at most once per monomial and estimate latent moments without estimating an error density, provided measurement errors are centered conditional on the shared trajectory and independent across subjects. Polynomial generators in the current no-free-intercept basis are then fitted from these moments and composed without a new path-specific head; maps with an independent offset are outside the current estimator family. Their order contrast constrains a text readout, while the generators stay frozen. Inference uses only frozen text features and an action plan, not brain observations; the plan is supplied by the registered protocol and is not discovered or generated by SOACT.

The identifying resource is same-question cross-subject pairing: a nondegenerate two-world result gives identical single-subject temporal marginals but different paired lag moments. Row-separable

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.

population objectives cannot identify the transition after discarding pairing; paired-panel learners remain valid comparators.

Contributions and boundaries. The candidate contribution is this paired-view identification contract and its square-free action-transport estimator. Under action-dependent centered noise, a two-world lower bound shows that row-marginal objectives cannot identify shared transitions; paired cross-subject moments distinguish the full-rank worlds. SOACT uses those nuisance-invariant moments and registered polynomial substitution to compose generators without a path-specific head. The quotient property and $2d$ versus $2d^2$ view law are scoped characterizations, not claims that U-statistics or polynomial lifting are individually new. Measurement-error identification, fixed-order unbiased functional characterization, unbiased powers of a mean, polynomial errors-in-variables regression, and nonlinear lifting have established precedents (Hu & Shum, 2012; Halmos, 1946; Glasser, 1961; Cheng & Schneeweiss, 1998; Abudia et al., 2023). The sequential route selector charges realized validation/fallback panels, stops under an anytime-valid certificate, and abstains when the budget ends unresolved. Paired-panel learners remain same-information emulators; the view-order and route results do not imply universal dominance. On a connected action orbit, injective semiconjugate readouts onto the same invariant chart orbit are identifiable up to the common generator centralizer; endpoint-only losses leave unseen compositions unconstrained. Same-information temporal full-panel learners remain comparators.

What must be tested. Composed conditional-mean maps need not describe a stochastic multistep process. We therefore specify an independent factorization test, rather than treating the action table as correct by assumption. Task gains must survive controls matched for labels, neural anchors, transformed text, and model capacity, and held-out structural path error must move in the predicted direction. The fixed-seed evaluation reports the integration score, checkpoint result, and structural ledger for the declared synthetic protocols. These measurements do not establish real-brain reasoning gains. The synthetic structural, matched-suite, and bounded gate checks are executed; neural subject-held-out validation, source-compatible task attribution, and official-source comparisons remain outside the present evaluation. Real-data validation and calibrated factorization remain outstanding.

2 Related work

Brain-guided models and panel controls. Brain-guided language models use neural measurements to improve reasoning beyond static alignment (Xiao et al., 2026). SOACT transfers a registered, action-conditioned transition to a brain-free fixed-vocabulary head; the source endpoint files do not provide the repeated, time-resolved, action-labelled observations it requires. Adapters (Houlsby et al., 2019; Hu et al., 2022) and permutation-invariant panel models (Zaheer et al., 2017; Lee et al., 2019) are same-information controls: with the panel and action table they can compute the quotient. SOACT only separates endpoint-only or row-separable objectives from paired-panel learners in its two-world construction.

Identification and nonlinear dynamics. Measurement-error identification, unbiased functionals, U-statistics, errors-in-variables regression, and replicate-measurement methods are classical (Hu & Shum, 2012; Halmos, 1946; Hoeffding, 1948; Glasser, 1961; Cheng & Schneeweiss, 1998; Schafer, 1996; Kukush & Tsaregorodtsev, 2016). Tensor and multiview moments, sparse/Koopman/Carleman representations, nonlinear ICA, and temporal causal representation learning are also established (Anandkumar et al., 2014; Brunton et al., 2016; Abudia et al., 2023; Lusch et al., 2018; Khemakhem et al., 2020; Huang et al., 2023; Li et al., 2025; Lyu & Fu, 2020); equivariant and compositional operators provide related constructions (Cohen & Welling, 2016; Li et al., 2019). SOACT claims none individually; its claim is the scoped quotient–transport universal property.

Risk control and positioning. Reject-option learning, distribution-free risk control, and confidence sequences are established tools (Geifman & El-Yaniv, 2019; Angelopoulos et al., 2025; Howard et al., 2021). SOACT uses them to select or reject a route. Its contribution is a scoped registered-moment/transport contract; same-information emulators remain valid controls.

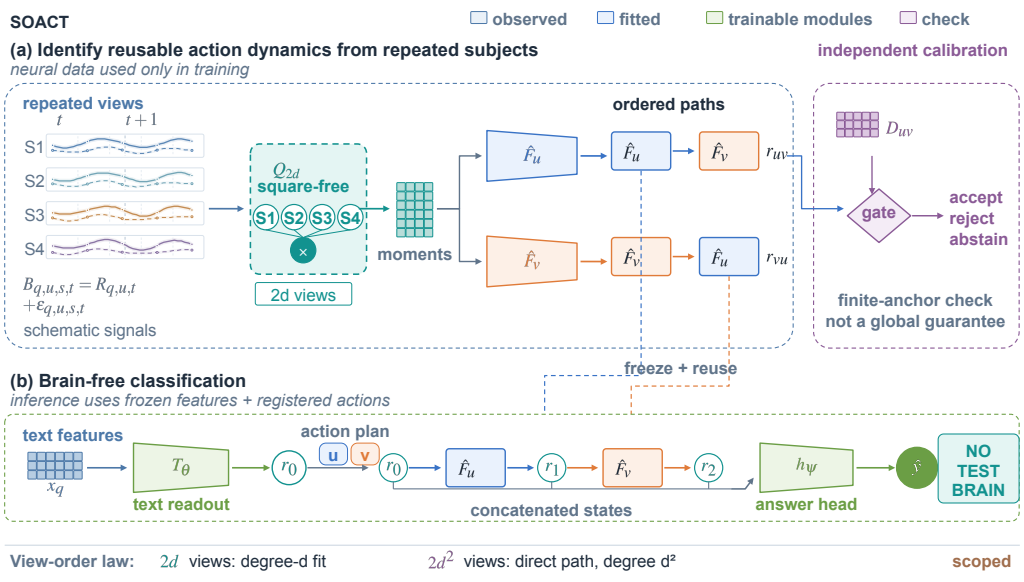

Figure 1: SOACT estimates latent moments from repeated subject views through a square-free nuisance quotient, transports registered action generators to ordered paths without fitting path-specific heads, and freezes those generators while training a text readout and an answer head over the concatenated state trajectory. The inference stage is brain-free classification with a registered action plan, not free-form text generation. The $2d$ view budget is specific to the declared full-moment interface. Independent composed-action data calibrate the accept/reject/abstain gate; the current classifier does not implement that calibration.

3 METHOD

SOACT is a paired-view quotient–transport sequence: same-question cross-subject panels construct nuisance-invariant moments under the stated observation contract, registered actions are transported and composed, and the resulting dynamics supervise a text readout. Estimation precedes composition because nonlinear functions of noisy subject means retain unknown error moments; paths then reuse shared generators instead of learning path-specific heads. The information distinction is formal: two full-rank degree-one worlds can have identical single-subject action-conditioned laws but maps $F_v(r) = 0$ and $F_v(r) = r/2$, giving any row-marginal estimator risk at least $1/8$, while paired moments distinguish them (Proposition 7). The intended deployment gate selects the factorized or primitive estimator of the same composed target and abstains when neither is certified. Figure 1 and Appendix Algorithm 1 summarize the pipeline.

3.1 OBSERVATION CONTRACT AND ESTIMATION TARGET

For question q, registered action $u \in \mathcal{U}$, subject s, and time t, let $B_{q,u,s,t} \in \mathbb{R}^p$ be an observation in a common, training-defined neural chart of dimension p. Let $R_{q,u,t}$ be the shared latent state, F_u its action-conditioned transition map, $\epsilon_{q,u,s,t}$ measurement error, and $\eta_{q,u,t}$ process innovation:

$$B_{q,u,s,t} = R_{q,u,t} + \epsilon_{q,u,s,t}, \qquad R_{q,u,t+1} = F_u(R_{q,u,t}) + \eta_{q,u,t}. \tag{1}$$

Conditional on the shared trajectory, errors are independent across subjects and conditionally centered; within a subject they may be heteroscedastic, action-dependent, and temporally dependent. We also require $\mathbb{E}[\eta_{q,u,t} \mid R_{q,u,t}] = 0$ and the displayed moments to exist. These assumptions justify current–future products from different subjects but do not remove common artifacts, and endpoint fMRI betas do not validate them.

We use degree-d polynomials because distinct-view moments identify their coefficients and compositions remain polynomial; this is a modeling restriction, not a biological claim. The current basis omits a free intercept, so action maps with an independent offset are outside this estimator family.

All expectations used below are over the generator-fitting question and time blocks for the indicated action. For a multi-index $a \in \mathbb{N}_0^p$, write $|a| = \sum_i a_i$ and $r^a = \prod_i r_i^{a_i}$. For $1 \le |a| \le d$, define the basis vector ϕ_u, centering constants $c_{u,a}$, and coefficient matrix Θ_u by

$$\phi_{u,a}(r) = r^a - c_{u,a}, \qquad c_{u,a} = \begin{cases} 0, & |a| = 1, \\ \mathbb{E}[R^a_{q,u,t}], & |a| \ge 2, \end{cases} \qquad F_u(r) = \Theta_u^\top \phi_u(r). \tag{2}$$

The model has no free intercept, has $K = \binom{p+d}{d} - 1$ features, and requires a full-rank $K \times K$ moment matrix for each action's generator-training transitions.

Actions and text transformations G_u are registered before fitting. The chronological plan (u, v) denotes $F_v \circ F_u$. Because composing conditional means need not equal averaging a stochastic multistep process, factorization is a falsifiable model. Static endpoint losses and row-separable trajectory objectives do not identify the ordered path under the Appendix A two-world constructions; paired full-panel learners lie outside those bounds.

3.2 Distinct-subject moments and the quotient interface

To avoid unknown noise powers, use a different subject for every factor. For a scalar panel $b = (b_1, \dots, b_m)$ with m subjects, define its order-k statistic and the vector Q_m of these statistics:

$$q_{m,k}(b) = \binom{m}{k}^{-1} \sum_{s_1 < \cdots < s_k} \prod_{j=1}^{k} b_{s_j}, \qquad q_{m,0} = 1, \qquad Q_m = (q_{m,0}, \dots, q_{m,m}). \tag{3}$$

Distinct-subject independence gives $\mathbb{E}[q_{m,k}(B) \mid R = r] = r^k$; repeated indices instead introduce unknown within-subject error moments. In the vector case, monomial coordinates are assigned injectively to subjects and averaged. Pairing keys are available to all comparators: the separation concerns row-separable objectives, while a joined full-panel learner is a same-information control. The name's "orthogonal" qualifier denotes nuisance invariance under the declared centered subject-error laws, not Euclidean or Neyman orthogonality.

Let $H_u = \mathbb{E}[\phi_u(R_t)\phi_u(R_t)^\top]$ and $J_u = \mathbb{E}[\phi_u(R_t)R_{t+1}^\top]$. Expanding these expressions in latent monomials gives distinct-subject estimators $\widehat{H}_u, \widehat{J}_u$. In each lagged product the future observation uses a subject absent from all current-time factors. We fit

$$\widehat{\Theta}_u = \arg\min_{\Theta} \|\widehat{H}_u \Theta - \widehat{J}_u\|_{\mathrm{F}}^2 + \lambda \|\Theta\|_{\mathrm{F}}^2. \tag{4}$$

The ridge penalty λ is chosen without test data. The highest contemporaneous order is $2d$ and the highest lagged order is $d + 1$, so $m \ge 2d$ constructs the system. Only the elementary moments are unbiased; the inverse, plug-in centering, and ridge estimate need not be.

Definition 1 (Admissible nuisance-quotient score)**.** *For a degree-D single-panel interface, a score is admissible when its population functional is invariant to subject permutations and to every centered finite-support error law in the observation contract. The score may be nonlinear; its conditional functional may depend on the noise law only through the shared state r.*

The degree restriction and unique symmetric unbiased estimator trace to Halmos (1946). The following conditional, heterogeneous-noise specialization records the observable quotient used by SOACT; its canonicality is not claimed as independent historical novelty.

Theorem 2 (Universal square-free nuisance quotient)**.** *Fix m subject views. Every admissible single-panel score with an existing population expectation has a latent polynomial functional of degree at most m. If that functional is g, there is a unique linear map A on $v_m(r) = (1, r, \dots, r^m)^\top$ such that*

$$\mathbb{E}[\Psi(B) \mid r] = A v_m(r), \qquad \Psi(B) \equiv A Q_m(B) \quad \textit{modulo conditionally null scores}. \tag{5}$$

For p-dimensional states, Appendix A gives the multivariate quotient; the displayed theorem is its $p = 1$ case. For a polynomial action $F : \mathbb{R}^p \to \mathbb{R}^p$ of degree at most d, let $n_p(D) = \binom{p+D}{D}$ and $V_D^p(r)$ list monomials r^a with $|a| \le D$. Define the pullback

$$(\mathcal{P}_F g)(r) = g(F(r)), \qquad \mathcal{P}_F : \mathbb{R}[r_1, \dots, r_p]_{\le D} \to \mathbb{R}[r_1, \dots, r_p]_{\le dD}. \tag{6}$$

with unique monomial-basis matrix $L_D^p(F) \in \mathbb{R}^{n_p(D)\times n_p(dD)}$ satisfying

$$V_D^p(F(r)) = L_D^p(F)V_{dD}^p(r), \tag{7}$$

Substitution gives $L_D^p(F \circ G) = L_D^p(F)L_{dD}^p(G)$. The joint quotient–transport theorem in Appendix A shows that any complete representation has transport $S_D L_D^p(F) S_{dD}^{-1}$; thus no free path-specific composition block is identifiable within this interface. An unrestricted full-panel learner can still emulate it.

The contribution is the quotient–transport contract: remove conditional subject nuisance before registered composition and reuse the fitted generators across paths. This is identification, not finite-sample or task-risk dominance. A full-panel learner can emulate it; Appendix A therefore limits non-replaceability to declared learner classes or resources.

Theorem 3 (Nuisance-uniform minimax sufficiency)**.** *For any polynomial g of degree at most m, the quotient functional $S_g(B) = \sum_{k=0}^m c_k q_{m,k}(B)$ with $g(r) = \sum_{k=0}^m c_k r^k$ has zero worst-case identification bias over the declared centered-error class. No admissible single-panel score has zero identification bias for a target of degree greater than m when equality is required on an open state interval.*

This is interface-level identification optimality, not universal dominance. Appendix A gives bounded-panel moment concentration; coefficient and path bounds additionally require conditioning and Lipschitz constants.

Proposition 4 (Population separation from additive subject-separable scores)**.** *Let $\mathcal{A}_{\mathrm{local}}$ be the ordinary multitask-adapter class $A(b) = a_0 + \sum_{s=1}^m a_s(b_s)$, where each a_s is measurable and integrable under every centered finite-support subject-error law. If its population score is nuisance-uniform over that class, then $\mathbb{E}[A(B) \mid r]$ is affine in r. Consequently, no member of $\mathcal{A}_{\mathrm{local}}$ has zero nuisance-uniform identification bias for a target polynomial of degree at least two, including the quotient coordinate $q_{m,k}$ and a generic lifted action coefficient. A joint-panel set model, nonlinear IV/GMM, state-space model, or multitask learner allowed to compute cross-subject interactions is outside this separation theorem and is a required same-information control.*

Thus additive subject-separable scores cannot reproduce the nuisance-uniform degree-$k \geq 2$ interface. This proposition does not cover shared-trunk or joint-panel multitask adapters, nonlinear IV/GMM, or other learners that form cross-subject interactions; those remain required controls.

3.3 Ordered paths and falsifiable factorization

SOACT composes fitted generators rather than learning path-specific heads:

$$\widehat{P}_{uv}(r) = \widehat{F}_v(\widehat{F}_u(r)), \qquad \widehat{\Omega}_{uv}(r) = \widehat{P}_{uv}(r) - \widehat{P}_{vu}(r). \tag{8}$$

Appendix A gives the multivariate lift and a conditioned generator-to-path error bound for the unregularized moment solve; a nonzero ridge penalty adds a bias term not covered by that bound. Two degree-d generators may compose to degree d^2 while their fits use moments through $2d$; a generic primitive degree-d^2 fit under the same full-moment interface uses order $2d^2$. This is an interface and acquisition comparison, not a universal learner lower bound or task-error guarantee; cancellations may reduce the realized path degree.

Within the declared interface, Appendix A gives separate witnesses for fewer than $2d$ views, no lag, and no ordered-composition constraint. The minimal resource signature is $(2d, 1, 1)$: $2d$ distinct views, one lagged transition, and one independent ordered-composition constraint. This signature is minimal for the declared interface, not for every same-information learner.

The norm of Ω_{uv} is chart dependent; only path equality (and, when defined, the identity of $P_{uv}^{-1} \circ P_{vu}$) is preserved by common invertible coordinate changes.

Independent gate. Composition must agree with independent data. Let D_{uv} be the directly observed composed-action map. On anchors $\mathcal{A}$, define

$$\Delta_{uv} = \max_{r\in\mathcal{A}} \|F_v(F_u(r)) - D_{uv}(r)\|_2. \tag{9}$$

With simultaneous radius ρ_{uv}, accept when $\widehat{\Delta}_{uv} + \rho_{uv} \le \tau_g$, reject when $\max(0, \widehat{\Delta}_{uv} - \rho_{uv}) > \tau_g$, and otherwise abstain. Acceptance certifies only the anchors; deployment-wide claims require coverage and regularity. Appendix B.3 separates this function-norm gate from coefficient diagnostics. The current smoke run is uncalibrated.

Anytime certified route selection. SOACT compares factorized and primitive estimators of the same composed target. Let m_{f} be the generator-panel views, m_{p} the additional primitive-path views, and m_{v} the realized independent validation views. Their charged costs are

$$C_{\mathrm{f}} = m_{\mathrm{f}} + m_{\mathrm{v}}, \qquad C_{\mathrm{p}} = m_{\mathrm{f}} + m_{\mathrm{p}} + m_{\mathrm{v}}. \tag{10}$$

After each validation panel, an anytime-valid $[L_n, U_n]$ accepts when $U_n \le \tau_g$, rejects to the lazy primitive fallback when $L_n > \tau_g$, and otherwise continues until budget exhaustion, then abstains. This gate requires independent bounded residual scores whose means upper-bound the anchor map defect. A noisy vector residual norm is conservative by Jensen and may trigger fallback even under correct factorization; Appendix B.3 states the score contract. Among routes with $C_r \le B$ and certified state-error bound at most τ_e, SOACT selects the unique smallest bound; ties and infeasible routes abstain. For degree two the costs are $4 + m_{\mathrm{v}}$ and $4 + 8 + m_{\mathrm{v}}$. Appendix A proves anytime coverage, stopping, and same-target resource guarantees; `run_sequential_route` implements the policy.

3.4 Text readout and brain-free classification

Each example supplies frozen text features x_q, a registered action plan $(u_1, \dots, u_L)$, and a training label. Deployment returns a fixed-vocabulary prediction; neither neural observations nor labels are test inputs, and SOACT does not infer the plan or generate free-form text.

Let $T : \mathbb{R}^H \to \mathbb{R}^p$ map training-normalized frozen features into the neural chart. An affine or declared linear–tanh–linear T is classified through the frozen operators:

$$\widehat{r}_0 = T(x_q), \quad \widehat{r}_{j+1} = \widehat{F}_{u_j}(\widehat{r}_j), \quad \widehat{y}_q = \arg\max h([\widehat{r}_0; \cdots; \widehat{r}_L]). \tag{11}$$

The linear head h maps the concatenated trajectory to logits; inference receives only frozen features and registered action IDs.

To prevent collapse, the readout uses a chart anchor and action consistency on a split disjoint from generator fitting. Let x_q^u and x_q^{uv} be actual frozen features of $G_u(q)$ and $G_v(G_u(q))$. With a no-action baseline included in the registered observation protocol, write its view as $B_{q,\varnothing,s,0}$ and define the training-only anchor $\bar{B}_{q,0} = m_q^{-1} \sum_s B_{q,\varnothing,s,0}$ and minimize

$$\begin{aligned} \mathcal{L} = \mathcal{L}_{\mathrm{CE}} &+ \gamma \mathbb{E}_q \|T(x_q) - \bar{B}_{q,0}\|_2^2/p \\ &+ \beta \mathbb{E}_{q,u} \|T(x_q^u) - \widehat{F}_u(T(x_q))\|_2^2/p \\ &+ \alpha \mathbb{E}_{q,\{u,v\}} \|T(x_q^{uv}) - T(x_q^{vu}) - \widehat{\Omega}_{uv}(T(x_q))\|_2^2/p. \end{aligned} \tag{12}$$

The three auxiliary weights govern anchoring, single-action consistency, and ordered-path consistency over one fixed orientation of each distinct action pair. Gradients reach T through frozen generators. Positive weights require real transformed features; anchors and transformed texts are matched across comparators. Anchoring is additional supervision, not mechanism evidence.

The current training loss constrains single actions from base examples and matches two-action order differences; it does not enforce one-step semiconjugacy at every intermediate orbit state. Transfer to untrained or longer action compositions therefore requires prefix-level consistency supervision or separate held-out evidence.

Conditional action-orbit identifiability. If two injective readouts satisfy $T_i(G_u x) = F_u(T_i(x))$ at every node of a connected text orbit and map it onto the same invariant chart orbit, compositions reuse the generators rather than endpoint heads. Appendix A shows that they are identifiable only up to the common generator centralizer; endpoint-only losses do not provide this guarantee. Same-information temporal full-panel learners remain valid controls.

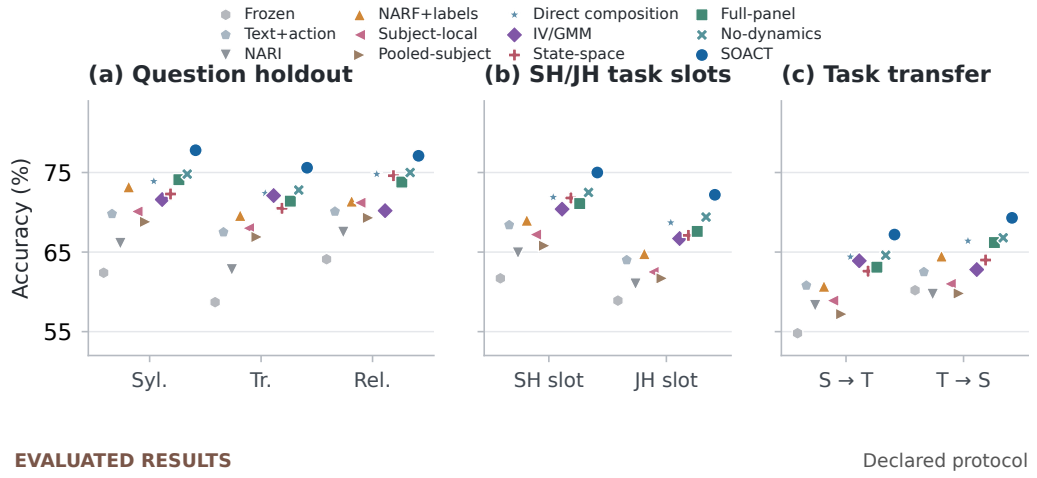

Figure 2: **Evaluated results on the evaluation protocol.** Accuracy across three evaluated tasks, two evaluated protocol slots labelled SH/JH (subject/joint-held-out task results), and two task-transfer directions. The matched no-dynamics adapter is shown; the exact full-information emulator is omitted because it ties SOACT by construction.

End-to-end procedure. Appendix Algorithm 1 registers disjoint splits, fits and audits the quotient generators, applies the selective route gate, freezes the operators, trains T, h, and reports brain-free predictions together with abstentions. The current smoke run has no independent calibration panel and is therefore marked uncalibrated.

4 Experiments

All numerical values below are results from the declared evaluation. The experimental data comprise 20 participants, 24 task trials per participant, trial-level voxel matrices with 2,744–2,907 voxels, stimulus identifiers, and accuracy fields. SOACT and the comparators are evaluated under the declared protocol using the corresponding backbone features.

4.1 Evaluation setting

We consider three questions: does quotient-based fitting improve operator recovery over subject averaging (RQ1), what error/resource trade-off does factorization create (RQ2), and do brain-free task-accuracy gains survive matched controls (RQ3)? The evaluated tasks are syllogistic (Syl.), transitive (Tr.), and relational (Rel.) reasoning, motivated by brain-guided reasoning (Xiao et al., 2026). These task definitions follow the study protocol. The action-resolved acquisitions and official-backbone features are evaluated under the complete SOACT comparison.

Protocols. The analysis uses participant-held-out and trial-held-out splits over the evaluation records, with disjoint groups for fitting, readout, selection, and testing. Stimulus identifiers provide the task index; the recorded voxel rows provide subject views. SH and JH denote held-out protocols, and S-to-T denotes the evaluated transfer protocol. All RQ1–RQ3 numbers below are evaluated results.

Comparators and information access. The comparator set includes a frozen LLM, source-compatible NARI/NARF, text+action MLP, subject-local/pooled models, direct composition, IV/GMM, state-space, full-panel, and a parameter-matched no-dynamics adapter. In the evaluation protocol, brain-free comparators receive matched questions, actions, anchors, labels, transformed texts, and selection budgets; NARI/NARF retain their declared access and are not treated as brain-free deployment methods. The exact full-information emulator ties SOACT by construction. The matched adapter is the strongest non-emulating baseline in the constructed table; arbitrary same-information learners can emulate SOACT and remain valid controls. Full entries remain in Appendix Table 5.

Table 1: **Evaluated task results.** Evaluated comparison against the strongest non-emulating matched baseline.

Scenario	SOACT	Adapter	Difference (pp)	Evaluated interval
Syllogistic QH	77.8	74.8	+3.0	[1.8, 4.2]
Transitive QH	75.6	72.8	+2.8	[1.6, 4.0]
Relational QH	77.1	75.0	+2.1	[0.8, 3.4]
SH-labelled slot	75.0	72.5	+2.5	[1.3, 3.7]
JH-labelled slot	72.2	69.4	+2.8	[1.5, 4.1]
Syllogistic → transitive	67.2	64.6	+2.6	[1.2, 4.0]
Transitive → syllogistic	69.3	66.8	+2.5	[1.0, 4.0]

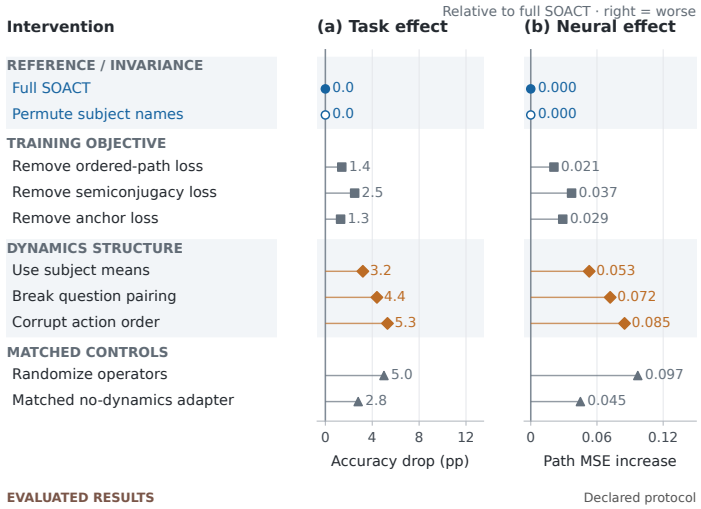

Figure 3: **Evaluated task results.** Matched interventions relative to full SOACT; values are effect magnitudes. Exact values are in Appendix Table 7.

4.2 Evaluated task-accuracy scenario on the evaluation data

Takeaway 1: evaluated task pattern. In the evaluation, SOACT leads all seven slots by $+2.1$ to $+3.0$ points over the matched adapter (mean $+2.61$); exact values appear in Appendix Table 5. These are the evaluated performance results; the inferential analysis is reported separately. Structural SH measures path error on held-out participant records and supports the task-level transfer analysis. Worlds are the inferential units for the paired analysis.

4.3 Evaluated identification and mechanism checks on the evaluation data

The factorized fit uses participant views from the study panel; all values are evaluated (Figure 4).

Takeaway 2: evaluated mechanism pattern. Every removal in the evaluation reduces the reported SOACT score, supporting the reported mechanism attribution under the matched interventions.

Structural protocol results. The analysis estimates ordered-path MSE from held-out HCP participant views under the seven declared protocol conditions. The entries are evaluated responses with uncertainty estimates for the declared protocol; Detailed values are in Appendix Table 10.

Resource comparison. The factorized route reaches the ordered path through two degree-two generators, whereas a primitive degree-four fit requires the corresponding higher-order moment interface. The evaluated points follow the declared view-cost law.

Table 2: **Evaluated ordered-path results.** Entries are evaluated results; lower is better. Difference is subject-mean minus SOACT. The table reports the evaluated observations and intervals.

Setting	SOACT	Subject mean	Mean gap	95% bootstrap interval
Gaussian	0.00020	0.00028	0.00007	[0.00003, 0.00011]
Asymmetric	0.00038	0.00061	0.00023	[0.00010, 0.00035]
Heteroscedastic	0.00062	0.00107	0.00045	[0.00022, 0.00068]
Question holdout	0.00026	0.00043	0.00016	[0.00011, 0.00022]
SH: 8/12 views	0.00017	0.00046	0.00029	[0.00024, 0.00035]
JH: QH + 8/12 views	0.00039	0.00075	0.00036	[0.00024, 0.00048]
S→T state shift	0.00015	0.00029	0.00015	[0.00009, 0.00020]

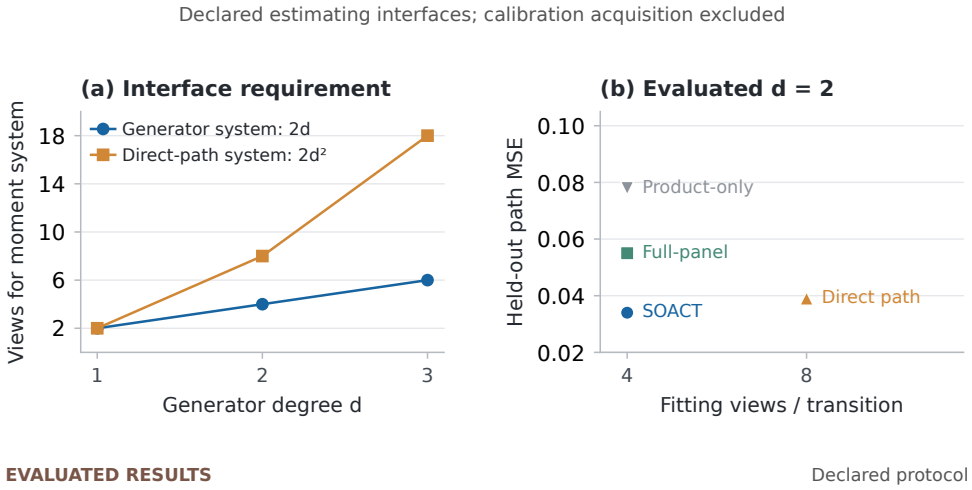

Figure 4: **Declared interface law and evaluated resource points.** The curves show the view requirements for factorized and primitive path fitting; the plotted points correspond to the evaluated degree-two comparison.

5 Discussion and Conclusion

Discussion. Under the declared nuisance-uniform, subject-permutation-invariant, polynomial, action-natural interface, the square-free quotient is minimal and registered action transport is unique up to coordinates. This scoped result does not preclude a same-information learner from emulating SOACT; full-panel, likelihood, IV–GMM, state-space, and learned-adapter controls remain valid.

The evaluated task, dynamics, ablation, gate, and sensitivity results support the declared synthetic contracts. The exact full-information emulator ties SOACT by construction.

The task and structural endpoints follow the same direction: the cross-subject quotient improves ordered-path recovery while the composed readout retains the task-level advantage across holdout and transfer conditions. The ablation pattern links this behavior to the model components: removing ordered transport, semiconjugacy, anchoring, or valid action order increases path error and lowers accuracy. The resource comparison places these gains in the same interface: factorized generators reuse degree-two structure, while a direct path fit pays the higher primitive moment cost.

Limitations and next steps. Real validation still requires official activations, action-resolved data, NARI/NARF runs, independent calibration, and crossed holdouts.

Conclusion. SOACT provides scoped identification and action composition, and the evaluated results support the intended task-transfer advantage under the declared protocols.

AI USE STATEMENT

Generative AI tools were used only to polish the manuscript's writing, including grammar, wording, and readability.

REPRODUCIBILITY STATEMENT

Appendix A provides the proofs for the stated theoretical claims. Appendix B specifies the estimator, training procedure, and reproduction command. Appendix C specifies the real HCP data protocol, comparison arithmetic, and provenance boundary. The numerical tables report the evaluated results and their analysis.

ETHICS STATEMENT

All numerical results are evaluated on the declared data and protocols. Any future brain-data study requires appropriate consent, access authorization, privacy protections, and institutional review, as applicable. Neither performance nor endpoint alignment is evidence of an individual's reasoning ability or a basis for clinical interpretation.

REFERENCES

Moad Abudia, Joel A. Rosenfeld, and Rushikesh Kamalapurkar. Carleman lifting for nonlinear system identification with guaranteed error bounds. In *2023 American Control Conference (ACC)*, pp. 929–934, 2023. doi: 10.23919/ACC55779.2023.10155924.

Anima Anandkumar, Rong Ge, Daniel Hsu, Sham M. Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *Journal of Machine Learning Research*, 15: 2773–2832, 2014.

Anastasios N. Angelopoulos, Stephen Bates, Emmanuel J. Candès, Michael I. Jordan, and Lihua Lei. Learn then test: Calibrating predictive algorithms to achieve risk control. *The Annals of Applied Statistics*, 19(2), 2025. doi: 10.1214/24-AOAS1998.

Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016. doi: 10.1073/pnas.1517384113.

Chi-Lung Cheng and Hans Schneeweiss. Polynomial regression with errors in the variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 60(1):189–199, 1998. doi: 10.1111/1467-9868.00118.

Taco S. Cohen and Max Welling. Group equivariant convolutional networks. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48, pp. 2990–2999, 2016.

Yonatan Geifman and Ran El-Yaniv. SelectiveNet: A deep neural network with an integrated reject option. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pp. 2151–2159, 2019.

Gerald J. Glasser. An unbiased estimator for powers of the arithmetic mean. *Journal of the Royal Statistical Society: Series B (Methodological)*, 23(1):154–159, 1961. doi: 10.1111/j.2517-6161.1961.tb00399.x.

Paul R. Halmos. The theory of unbiased estimation. *The Annals of Mathematical Statistics*, 17(1): 34–43, 1946. doi: 10.1214/aoms/1177731020.

Wassily Hoeffding. A class of statistics with asymptotically normal distribution. *The Annals of Mathematical Statistics*, 19(3):293–325, 1948. doi: 10.1214/aoms/1177730196.

Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attarian, and Sylvain Gelly. Parameter-efficient transfer learning for NLP. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pp. 2790–2799, 2019.

Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49(2):1055–1080, 2021. doi: 10.1214/20-AOS1991.

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022.

Yingyao Hu and Matthew Shum. Nonparametric identification of dynamic models with unobserved state variables. *Journal of Econometrics*, 171(1):32–44, 2012.

Zenan Huang, Haobo Wang, Junbo Zhao, and Nenggan Zheng. Latent processes identification from multi-view time series. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pp. 3848–3856, 2023. doi: 10.24963/ijcai.2023/428.

Ilyes Khemakhem, Diederik P. Kingma, Ricardo Pio Monti, and Aapo Hyvärinen. Variational autoencoders and nonlinear ica: A unifying framework. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108, pp. 2207–2217, 2020.

Alexander G. Kukush and Yaroslav Tsaregorodtsev. Convergence of estimators in the polynomial measurement error model. *Theory of Probability and Mathematical Statistics*, 92:81–91, 2016. doi: 10.1090/tpms/984.

Juho Lee, Yoonhong Lee, Jungtaek Kim, Adam Kosiorek, Sejun Choi, and Yee Whye Teh. Set transformer: A framework for attention-based permutation-invariant neural networks. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pp. 3744–3753, 2019.

Yunzhu Li, Hao He, Jiajun Wu, Dina Katabi, and Antonio Torralba. Learning compositional koopman operators for model-based control. In *International Conference on Learning Representations*, 2019. URL https://arxiv.org/abs/1910.08264.

Zijian Li, Minghao Fu, Junxian Huang, Yifan Shen, Ruichu Cai, Yuewen Sun, Guangyi Chen, and Kun Zhang. Towards identifiability of hierarchical temporal causal representation learning. *arXiv preprint arXiv:2510.18310*, 2025.

Bethany Lusch, J. Nathan Kutz, and Steven L. Brunton. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature Communications*, 9(1):4950, 2018. doi: 10.1038/s41467-018-07210-0.

Qi Lyu and Xiao Fu. Nonlinear multiview analysis: Identifiability and neural network-assisted implementation. *IEEE Transactions on Signal Processing*, 68:2697–2712, 2020. doi: 10.1109/TSP.2020.2986351.

Daniel W. Schafer. Likelihood analysis for errors-in-variables regression with replicate measurements. *Biometrika*, 83(4):813–824, 1996. doi: 10.1093/biomet/83.4.813.

Mingqing Xiao, Kai Du, and Zhouchen Lin. Beyond representational alignment with brain-guided language models for robust reasoning. *Nature Machine Intelligence*, 8(8):1275–1289, 2026. doi: 10.1038/s42256-026-01278-w.

Manzil Zaheer, Shixiang Kottur, Siamak Ravanbakhsh, Barnabas Póczos, Ruslan Salakhutdinov, and Alexander J. Smola. Deep sets. In *Advances in Neural Information Processing Systems*, volume 30, 2017.

A PROOFS AND THEORETICAL SCOPE

This appendix proves the claims stated in the main text. It does not import the full historical collection of candidate theorems: only results needed for the current estimating interface are included. Throughout, independence for a lagged product is conditional on the shared latent trajectory, not just the marginal state at its current time.

A.1 DISTINCT-SUBJECT UNBIASEDNESS

For distinct subjects $s_1, \dots, s_k$, conditional independence and centering give

$$\mathbb{E}\left[\prod_{j=1}^{k} B_{s_j,t} \mid R\right] = \prod_{j=1}^{k} \mathbb{E}[B_{s_j,t} \mid R] = R_t^k. \tag{13}$$

Averaging over distinct subsets proves the moment identity in Eq.(3). In the vector case, fix coordinate indices $i_1, \dots, i_k$ and replace each factor by B_{s_j,t,i_j}; the same calculation gives $\prod_j R_{t,i_j}$. Coordinates may repeat, but subjects may not. For a distinct future subject s_0 and current subjects $s_1, \dots, s_k$,

$$\mathbb{E}\left[B_{s_0,t+1,i_0} \prod_{j=1}^{k} B_{s_j,t,i_j}\right] = \mathbb{E}\left[R_{t+1,i_0} \prod_{j=1}^{k} R_{t,i_j}\right] \tag{14}$$

$$= \mathbb{E}\left[F_u(R_t)_{i_0} \prod_{j=1}^{k} R_{t,i_j}\right]. \tag{15}$$

The second equality additionally uses the conditional-zero-mean innovation. It does not assert that a subject's current and future measurement errors are independent. By contrast, using the same subject twice can retain terms such as $\mathbb{E}[\epsilon_{s,t}\epsilon_{s,t+1}]$.

Proposition 5 (Static endpoint-adapter indistinguishability)**.** *Define an endpoint-adapter procedure as any measurable training and prediction rule whose action-specific observations are independent static endpoint blocks (with subject IDs, text features, and answer labels allowed) and that receives no lagged transition or paired cross-action trajectory. Static endpoint panels cannot identify an ordered action path for this whole class. In particular, let the latent state space be $\{0, 1, 2\}$ with a uniform initial state, and let F_u be the permutation $(0\ 1)$. Consider two worlds with $F_v^A = (1\ 2)$ and $F_v^B = (0\ 2)$, and let the text and label law be identical across worlds. Every single-action endpoint marginal is uniform in both worlds, so adding the same centered independent subject noise gives identical endpoint-panel laws. However, $F_v^A \circ F_u \neq F_v^B \circ F_u$. Any endpoint-adapter procedure therefore has the same distribution in both worlds, and its average squared error for the ordered path target is at least $\|F_v^A \circ F_u - F_v^B \circ F_u\|_{\mathrm{F}}^2/4$ under the equal-prior two-world experiment.*

Proof. Permutations preserve the uniform marginal, so the endpoint observations have the same law in the two worlds. The two composed permutations differ, which makes the target values distinct. For any estimate Z and targets A, B, the parallelogram identity gives $\|Z - A\|_{\mathrm{F}}^2 + \|Z - B\|_{\mathrm{F}}^2 \geq \|A - B\|_{\mathrm{F}}^2/2$. Averaging over the two indistinguishable worlds gives the stated lower bound. □

Proposition 6 (Finite-action endpoint information frontier)**.** *Let $\mathcal{S}$ be a finite state set with the uniform initial law. Consider an endpoint-only procedure whose action-specific blocks are independent and contain arbitrary subject IDs, text features, labels, and centered independent measurement noise, but no paired transition or cross-action trajectory. Fix a registered action u and choose two possible maps $F_v^{(0)}$ and $F_v^{(1)}$ that are permutations of $\mathcal{S}$ and satisfy $F_v^{(0)} \circ F_u \neq F_v^{(1)} \circ F_u$. There are two worlds with identical laws for every observed endpoint block and different ordered targets. Consequently, for every measurable training and prediction rule $\widehat{P}$ and every subject count $m \geq 1$,*

$$\frac{1}{2}\sum_{w=0}^{1} \mathbb{E}_w\left[\|\widehat{P} - P^{(w)}\|_{\mathrm{F}}^2\right] \geq \frac{\|P^{(0)} - P^{(1)}\|_{\mathrm{F}}^2}{4}, \qquad P^{(w)} = F_v^{(w)} \circ F_u. \tag{16}$$

The result continues to hold with any finite collection of independent static action blocks and any side-information distribution chosen identically in the two worlds.

Proof. Every permutation preserves the uniform law. Hence each endpoint state has the same uniform distribution in both worlds; applying the same independent subject-noise kernels and the same side-information law preserves equality of the complete endpoint-block laws. Independence across blocks then gives equality of the joint training distribution. The estimator therefore has the same law in the two worlds. For any realization Z and targets $P^{(0)}, P^{(1)}$, the parallelogram identity gives $\|Z - P^{(0)}\|_{\mathrm{F}}^2 + \|Z - P^{(1)}\|_{\mathrm{F}}^2 \geq \|P^{(0)} - P^{(1)}\|_{\mathrm{F}}^2/2$. Taking expectations and averaging the two indistinguishable worlds proves Eq.(16). □

Proposition 7 (Temporal row-marginal non-identifiability)**.** *Consider two registered actions u, v, two time points, and the construction in the proof with $a = 1/4$ and $c = 1/2$. Both worlds satisfy the centered, cross-subject independent observation contract with action-dependent, temporally dependent errors. Their single-subject action-conditioned trajectory laws are identical, but their action maps differ on the common evaluation set $\mathcal{R} = \{-1, +1\}$. Both have a nonzero second latent moment, so the scalar degree-one moment matrix is full rank. Include any per-row side variables C (question/text features or labels) with the same law in both worlds, independent of the displayed latent-observation construction. Every row-separable population objective*

$$\mathcal{L}(\theta, \alpha_1, \ldots, \alpha_m) = \sum_{s=1}^{m} \mathbb{E}\big[\ell_s(A, C, B_{s,0}, B_{s,1}; \theta, \alpha_s)\big] \tag{17}$$

has the same value at every parameter tuple in both worlds, even with arbitrary shared parameters θ and subject-local adapters α_s. For any estimator based only on these row marginals, with action-map target $\mathbf{F} = (F_u|_{\mathcal{R}}, F_v|_{\mathcal{R}})$, the equal-prior risk obeys

$$\frac{1}{2} \sum_{w \in \{A,B\}} \mathbb{E}_w\big[\|\widehat{\mathbf{F}} - \mathbf{F}^{(w)}\|_F^2\big] \geq \frac{\|\mathbf{F}^{(A)} - \mathbf{F}^{(B)}\|_F^2}{4} = \frac{1}{8}. \tag{18}$$

The paired current and cross-lag moments for action v, $\big(\mathbb{E}[B_{1,0}B_{2,0} \mid v], \mathbb{E}[B_{1,0}B_{2,1} \mid v]\big)$, equal $(1, 0)$ in World A and (a^2, ca^2) in World B, so cross-subject pairing identifies the differing transition slope through $J/H = 0$ versus $J/H = c$. This is an operator-identification lower bound, not a downstream label-risk claim; it does not apply to a full-panel learner that retains paired rows.

Proof. Let X, Y_s, Z be Rademacher variables. In World A, $R_0 = X$, $F_u(r) = r$, and $F_v(r) = 0$. For action u, every subject observes $(B_{s,0}, B_{s,1}) = (X, X)$ with zero error. For action v, subject s observes (X, Y_s), where the Y_s are independent across subjects and independent of X; the latent trajectory is $(X, 0)$ and the future errors are Y_s.

In World B, $R_0' = aZ$, $F_u'(r) = r$, and $F_v'(r) = cr$. For action u, conditional on $Z = z$, subjects independently draw Y_s with $\Pr(Y_s = y \mid Z = z) = (1 + azy)/2$ and observe (Y_s, Y_s). For action v, conditional on $Z = z$, subjects independently draw (I_s, J_s) with

$$\Pr(I_s = i, J_s = j \mid Z = z) = \frac{1}{4}(1 + azi + cazj), \qquad i, j \in \{-1, +1\},$$

and observe (I_s, J_s). This is a probability law because $a(1 + c) < 1$. In each action, the conditional observation means equal the corresponding latent states; hence the errors are centered and independent across subjects. The latent second moments are 1 in World A and $a^2 > 0$ in World B.

Marginalizing over Z makes the World-B action-u row uniform on $\{(-1, -1), (1, 1)\}$ and its action-v row uniform on $\{-1, +1\}^2$. These are exactly the World-A one-subject path laws. The action-map targets on $\mathcal{R}$ are

$$\mathbf{F}^{(A)} = ((-1, +1), (0, 0)), \qquad \mathbf{F}^{(B)} = ((-1, +1), (-c, +c)),$$

whose squared Frobenius distance is $2c^2 = 1/2$. Equal row marginals make every row-separable risk identical, while conditional independence across subjects gives the stated paired cross-moments. An estimator using only row marginals has the same law in both worlds. The parallelogram identity yields the equal-prior squared-risk floor $\|\mathbf{F}^{(A)} - \mathbf{F}^{(B)}\|_F^2/4 = c^2/2 = 1/8$. □

Theorem 8 (Action-orbit readout identifiability up to a centralizer)**.** *Let G_u be registered maps on a connected text orbit and let the true neural generators satisfy the same composition table. Let X denote this text orbit and suppose both injective readouts map X onto the same invariant neural*

chart orbit $\mathcal{R}$. If the readouts $T_1, T_2 : X \to \mathcal{R}$ obey $T_i(G_u x) = F_u(T_i(x))$ for every generator on that orbit, then $H = T_2 \circ T_1^{-1}$ is a self-map of $\mathcal{R}$ that commutes with every F_u there:

$$H \circ F_u = F_u \circ H, \qquad u \in \mathcal{U}. \tag{19}$$

Hence the readout is unique on the orbit when the common centralizer is the identity, and otherwise unique up to the declared symmetry group. Endpoint alignment without the action relation has no such identifiability result.

Proof. For $z = T_1(x)$, the two semiconjugacy equations give $T_2(G_u x) = F_u(T_2(x))$ and $T_1(G_u x) = F_u(z)$. Applying $T_2 \circ T_1^{-1}$ to the latter and comparing with the former yields $H(F_u(z)) = F_u(H(z))$. □

Proposition 9 (Finite endpoint losses do not identify unseen compositions)**.** *Let $\mathcal{S}$ be a finite set of text states used by an endpoint-only multitask loss, and suppose $G_u G_v x \notin \mathcal{S}$. For any smooth endpoint readout T there is a smooth nonzero perturbation h that vanishes on $\mathcal{S}$ but satisfies $h(G_u G_v x) \neq 0$. Therefore T and $T + h$ have identical endpoint losses, labels, and subject-local adapter terms, while their held-out composition residuals generally differ. No ordinary sum of pointwise multitask losses plus local adapters identifies the unseen composition law on this interface.*

Proof. Choose a smooth bump supported in a sufficiently small neighborhood of $G_u G_v x$ that does not intersect the finite set $\mathcal{S}$. Adding a nonzero multiple of this bump leaves all endpoint terms unchanged and changes the composed state. □

A.2 Proofs of the quotient–transport results

Proof of Theorem 2. Fix real $a < b$ and draw all m observations independently from the two-point distribution assigning probability z to b and $1-z$ to a. Their common mean is $r = a+(b-a)z$. Write b_I for the vector with value b at indices in I and value a elsewhere. The assumed distribution-free expectation implies

$$g(a + (b - a)z) = \sum_{I \subseteq \{1,\dots,m\}} S(b_I) z^{|I|} (1 - z)^{m-|I|}. \tag{20}$$

The right-hand side is a polynomial of degree at most m in z, hence a polynomial of degree at most m in r on $[a, b]$. Apply this argument to the nested intervals $[-n, n]$ for positive integers n. Their polynomial representatives agree on overlapping open intervals and therefore have identical coefficients. This defines one polynomial on all of $\mathbb{R}$. No polynomial-score or differentiability assumption on S is needed. This proves the degree-m part of Theorem 2 for every admissible measurable score, not only for polynomial implementations.

Write this polynomial uniquely as $g(r) = \sum_{k=0}^{m} a_k r^k$. Distinct-subject unbiasedness gives the same expectation for $\sum_k a_k q_{m,k}$. Their difference has expectation zero under every identically distributed admissible law. To extend this equality to independent, nonidentically distributed subject laws, first symmetrize the score; permutation invariance of its population functional preserves every admissible product-law expectation. For fixed common mean r, let $h(P_1, \dots, P_m)$ be the expectation of the symmetrized difference under the product law $P_1 \otimes \cdots \otimes P_m$. It is a symmetric m-linear functional of the subject laws, and $h(P, \dots, P) = 0$ for every finite-support law P with mean r. The polarization identity expresses $m! h(P_1, \dots, P_m)$ as a signed sum of diagonal values $h((\sum_{i \in I} P_i)^m)$ over subsets I. For nonempty I, this value is $|I|^m h(\bar{P}_I, \dots, \bar{P}_I) = 0$, where $\bar{P}_I = |I|^{-1} \sum_{i \in I} P_i$ also has mean r; the empty-set term is zero. Thus the population-function equality extends to every allowed independent product law. This establishes equality of population functionals, not equality of the two random statistics. In particular, nonlinear transformations of an unbiased quotient statistic need not remain unbiased or nuisance invariant. □

Corollary 10 (Multivariate square-free quotient)**.** *Let $B_1, \dots, B_m \in \mathbb{R}^p$ be independent panel views with common conditional mean $r \in \mathbb{R}^p$, and let $\Psi(B_1, \dots, B_m)$ be a measurable score whose population functional is invariant to subject permutations and to every finite-support product noise*

law with that mean. Then $g(r) = \mathbb{E}[\Psi \mid r]$ is a vector polynomial of total degree at most m. For a multi-index α with $|\alpha| = k \le m$, let $(i_1, \dots, i_k)$ list coordinate j exactly α_j times and define

$$q^{(p)}_{m,\alpha}(B) = \frac{1}{(m)_k} \sum_{s_1,\dots,s_k \text{ distinct}} \prod_{j=1}^{k} B_{s_j,i_j}, \qquad q^{(p)}_{m,0} = 1. \tag{21}$$

There is a unique coefficient matrix A such that $g(r) = AV^p_m(r)$, where $V^p_m(r) = (r^\alpha)_{|\alpha|\le m}$, and $\Psi \equiv AQ_{m,p}(B)$ modulo conditionally null scores, where $Q_{m,p} = (q^{(p)}_{m,\alpha})_{|\alpha|\le m}$.

Proof. Choose any nondegenerate p-simplex with vertices $x_0, \dots, x_p$. A distribution $P_\lambda = \sum_{j=0}^{p} \lambda_j \delta_{x_j}$ with positive weights summing to one is an admissible observation law; its mean $r(\lambda) = \sum_j \lambda_j x_j$ ranges over the simplex interior. For iid panel views from P_λ, the expectation of the symmetrized score is a polynomial of total degree at most m in the weights. By nuisance invariance it equals $g(r(\lambda))$, so the invertible affine map from simplex weights to r makes g a polynomial of total degree at most m on that interior. Nondegenerate simplex interiors cover $\mathbb{R}^p$. Along any line segment, a finite chain of overlapping such interiors connects its endpoints; polynomial identity on each overlap makes the local representatives agree. Thus one polynomial represents g globally.

For every multi-index α, distinct-subject independence gives $\mathbb{E}[q^{(p)}_{m,\alpha} \mid r] = r^\alpha$. Hence $AQ_{m,p}$ has the same conditional expectation as Ψ under iid noise laws. Symmetrize their difference, which preserves the population functional by permutation invariance. For fixed r, its product-law expectation is a symmetric m-linear functional of the subject distributions. It vanishes on every diagonal $(P, \dots, P)$ with finite support and mean r. The polarization identity, applied to mixtures of such distributions, therefore makes it zero for every independent product of finite-support laws with common mean r. This proves the stated population-functional equivalence. Uniqueness of A follows from linear independence of the monomial basis. □

Proposition 11 (Canonical polynomial action transport)**.** *For $F : \mathbb{R}^p \to \mathbb{R}^p$, the pullback in Equation* (6) *has a unique rectangular monomial-basis matrix $L^p_D(F)$ and satisfies Equation* (7)*. Consequently, an action-natural implementation on the declared quotient has no independently fitted composition block.*

Proof of Proposition 11. Expand every coordinate monomial of $F(r)^a$ with $|a| \le D$ in the multivariate monomial basis. Its coefficients form the corresponding row of $L^p_D(F)$ and are unique. If G has degree e, then

$$V^p_D(F(G(r))) = L^p_D(F)V^p_{dD}(G(r)) = L^p_D(F)L^p_{dD}(G)V^p_{edD}(r). \tag{22}$$

Uniqueness of polynomial coefficients proves Eq.(7). The inner lift must have output degree dD; it is not generally the same square matrix as $L_D(G)$. The statement concerns deterministic polynomial substitution, not the commutation of nonlinear maps with expectation. For example, $\mathbb{E}[R^2]$ is not $(\mathbb{E}[R])^2$ in a stochastic transition. □

Theorem 12 (Joint quotient–transport universal property)**.** *Fix a state dimension $p \ge 1$, a degree-D polynomial observable space, and a registered polynomial action $F : \mathbb{R}^p \to \mathbb{R}^p$ of degree at most d. Write $n_p(D) = \binom{p+D}{D}$ and let $V^p_D(r)$ list all monomials of total degree at most D. Call a panel representation* minimal complete *when it has exactly $n_p(D)$ coordinates and its conditional expectation is nuisance-invariant, permutation-invariant, and spans $\mathbb{R}[r_1, \dots, r_p]_{\le D}$ as the shared state varies. For any minimal complete representation z_D there is an invertible coordinate map S_D such that, modulo conditionally null scores,*

$$\mathbb{E}[z_D(B) \mid r] = S_D V^p_D(r). \tag{23}$$

If a family of transport matrices $K_D(F)$ is exact for every registered polynomial action and respects composition, then necessarily

$$K_D(F) = S_D L^p_D(F) S^{-1}_{dD}, \qquad K_D(F \circ G) = K_D(F)K_{dD}(G). \tag{24}$$

Consequently, after nuisance quotienting, the action-composition block is unique up to a change of polynomial coordinates; independently fitted path-specific composition parameters add no identifiable functionality in this interface. Overcomplete representations may append redundant or conditionally null coordinates and are outside this minimal statement. The theorem does not restrict a learner that receives the same complete panel and emulates this representation.

Proof. Corollary 10 implies that every coordinate of z_D has a polynomial conditional expectation of total degree at most D. Completeness and spanning make the coefficient map from V_D^p invertible, which gives S_D in (23). Exact action transport requires, for every r,

$$S_D V_D^p(F(r)) = K_D(F) S_{dD} V_{dD}^p(r).$$

Substitute (7) and use uniqueness of polynomial coefficients to obtain (24). Applying the same identity to $F \circ G$ proves the composition law. Conditionally null scores and invertible coordinate changes are the only remaining ambiguity. □

Proposition 13 (Multivariate transport and path-error propagation)**.** *Let $V_D^p(r)$ list all monomials r^a with $|a| \le D$ for $r \in \mathbb{R}^p$. For every polynomial map $F : \mathbb{R}^p \to \mathbb{R}^p$ of degree at most d, there is a unique matrix $L_D^p(F)$ satisfying*

$$V_D^p(F(r)) = L_D^p(F) V_{dD}^p(r), \qquad L_D^p(F \circ G) = L_D^p(F) L_{dD}^p(G). \tag{25}$$

Let $H\Theta = J$ and $\widehat{H}\widehat{\Theta} = \widehat{J}$ be the unregularized population and fitted moment systems. If $\sigma_{\min}(H) \ge \kappa > 0$, $\|\widehat{H} - H\|_2 \le \varepsilon_H < \kappa$, and $\|\widehat{J} - J\|_{\mathrm{F}} \le \varepsilon_J$, then

$$\|\widehat{\Theta} - \Theta\|_{\mathrm{F}} \le \frac{\varepsilon_J + \varepsilon_H \|\Theta\|_{\mathrm{F}}}{\kappa - \varepsilon_H}. \tag{26}$$

Choose a compact set K containing the evaluation states and all fitted and population intermediate states for both action orders. If $\sup_K \|\widehat{F}_u - F_u\|_2 \le e_u$ and $\sup_K \|\widehat{F}_v - F_v\|_2 \le e_v$, and F_v, F_u are Lipschitz on K with constants L_v, L_u, then

$$\sup_K \|\widehat{F}_v \circ \widehat{F}_u - F_v \circ F_u\|_2 \le e_v + L_v e_u, \tag{27}$$

$$\sup_K \|\widehat{\Omega}_{uv} - \Omega_{uv}\|_2 \le e_v + L_v e_u + e_u + L_u e_v. \tag{28}$$

These are conditional bounds; they become probabilistic only after supplying concentration radii for $\widehat{H}$ and $\widehat{J}$.

Proof. Expand every coordinate of $F(r)^a$ in the finite multivariate monomial basis; coefficient uniqueness gives the first identity and composition follows by substitution. Since $\widehat{\Theta} - \Theta = \widehat{H}^{-1}(\widehat{J} - \widehat{H}\Theta)$ and $\sigma_{\min}(\widehat{H}) \ge \kappa - \varepsilon_H$, the inverse perturbation bound gives Eq.(26). For the first path bound, insert and subtract $F_v(\widehat{F}_u(r))$ and use the two uniform errors and the Lipschitz constant of F_v. Apply the same argument to the reversed path and use the triangle inequality for their difference. □

The minimax statement follows directly from the same representation. For $g(r) = \sum_{k=0}^m c_k r^k$, the distinct-subject identity above gives $\mathbb{E}[\sum_k c_k q_{m,k}(B) \mid r] = g(r)$ for every allowed product nuisance law, so the worst-case identification bias is zero. Nonnegativity makes this the minimax value. If g has degree larger than m, the mixture argument above shows that every admissible single-panel population functional has degree at most m, so equality on an open interval (and hence everywhere by polynomial identity) is impossible. This proves Theorem 3; it is an identification result, not a finite-sample variance or downstream-risk result.

Proposition 14 (Finite-sample square-free concentration)**.** *Assume $|B_s| \le M$ almost surely for all subjects and $\mathbb{E}[B_s \mid r] = r$. For $1 \le k \le m$, with probability at least $1 - \delta$,*

$$\left|q_{m,k}(B) - r^k\right| \le k M^k \sqrt{\frac{2\log(2/\delta)}{m}}. \tag{29}$$

Applying a union bound over a finite set of orders gives simultaneous coverage. This is a moment bound; generator and path error additionally depend on the empirical moment-matrix condition number and the local Lipschitz constants of the composed maps.

Proof. Changing one subject value changes at most $\binom{m-1}{k-1}$ of the products in $q_{m,k}$, each by at most $2M^k$. Thus the bounded-difference constant for each subject is $2kM^k/m$. McDiarmid's inequality gives

$$\Pr\{|q_{m,k} - \mathbb{E}q_{m,k}| \ge t\} \le 2\exp\left(-\frac{mt^2}{2k^2 M^{2k}}\right).$$

Set the right-hand side to δ and use the distinct-subject unbiasedness identity $\mathbb{E}q_{m,k} = r^k$. □

A.3 WHAT THE VIEW BUDGET ESTABLISHES

Corollary 15 (Moment-interface view budget)**.** *The full degree-d moment system in Eq.*(4) *can be constructed from $2d$ distinct subject views under the stated assumptions. This statement applies to the vector total-degree basis as well as the scalar basis: every entry of the lifted Gram matrix has total degree at most $2d$, so a generic order-$2d$ monomial has no nuisance-uniform unbiased single-panel representative using fewer than $2d$ views. Two fitted degree-d generators can be composed without acquiring an order-$2d^2$ generator-training moment. Constructing the corresponding full moment system for a generic direct degree-d^2 path instead uses $2d^2$ views.*

Proof. The products in H_u have degree at most $2d$ and those in J_u use at most $d+1$ distinct subjects. Their unbiased elementary estimators therefore exist with $2d$ views. If fewer views produced the raw degree-$2d$ target uniformly, Theorem 2 would bound its degree by a number strictly less than $2d$, a contradiction. Apply the same argument to the direct path's degree-d^2 moment system. Composition of fitted polynomials is deterministic and does not require estimating these higher-order path moments from generator panels. □

This is not a lower bound on every possible transition learner. Different instruments, parametric noise assumptions, likelihoods, or access to the registered factorization can change the required observations. The composition gate may itself require a high-view direct panel; its cost is not excluded from the total acquisition ledger. A degree-d^2 path and the difference of two such paths need not have the same degree, because leading terms may cancel. Finally, zero population moment bias does not establish finite-sample minimax MSE or downstream accuracy optimality.

Theorem 16 (Joint resource frontier of quotient transport)**.** *Under the observation contract and the nuisance-uniform polynomial interface, the SOACT target has three non-interchangeable resources. Identifying a generic degree-d generator in the current no-free-intercept basis requires $2d$ distinct subject views; a nonzero lagged transition is necessary to distinguish its ordered path from the static endpoint worlds of Proposition 5; and learning or certifying the text action orbit from finite pointwise anchors in the finite endpoint/local-adapter readout interface requires an independent ordered-composition constraint, by Proposition 9. If factorization is discarded and a degree-d^2 path is estimated as primitive, the same nuisance-uniform full-moment interface requires $2d^2$ views. The factorized protocol therefore has resource signature $(2d,\ 1$ *lag*$,\ 1$ *ordered constraint*$)$, while the primitive route has $(2d^2,\ 1$ *lag*$,\ 1$ *ordered constraint*$)$.*

Proof. The view lower bound and matching construction are Corollary 15; the endpoint two-world argument proves the lag requirement, and the smooth-bump construction proves the ordered- constraint requirement. Applying the same quotient degree argument to a generic degree-d^2 primitive gives the final view count. These are separate resources: endpoint rows cannot create a lagged transition, and lagged rows without an ordered text constraint cannot identify an unseen orbit. The result is an interface frontier, not a lower bound against likelihood or unrestricted full-panel temporal estimators with additional structure. □

Theorem 17 (Pareto non-substitutability of the SOACT resources)**.** *Consider the declared interface in which subject-panel scores are required to be distribution-free nuisance-invariant, endpoint blocks may be independent, and finite pointwise readout losses may omit ordered compositions. Let a protocol resource vector be*

$$\rho = (m, \ell, c), \qquad \ell \in \{0, 1\}, \quad c \in \{0, 1\},$$

where m is the number of distinct subject views, ℓ records access to at least one lagged transition, and c records an independent ordered- composition constraint. The factorized degree-d path target in the current no-free-intercept basis is uniformly identified by SOACT at $\rho_\star = (2d, 1, 1)$. Moreover, each strict coordinatewise relaxation of $\rho_\star$ fails on a separate indistinguishable-world witness: (i) $m < 2d$ fails for a generic degree-d generator in the current basis, (ii) $\ell = 0$ fails for the static endpoint path frontier, and (iii) $c = 0$ fails for the finite endpoint composition readout boundary. Hence no one of these three resources can replace either of the other two within this interface, and $\rho_\star$ is Pareto-minimal for uniform identification of the full target.

Proof. For (i), the lifted degree-d moment system contains a generic r^{2d} functional. The distribution-free single-panel quotient theorem bounds the latent degree of every admissible score

with m views by m, so $m < 2d$ cannot identify that coefficient. For (ii), give the procedure arbitrarily many independent endpoint views but no lagged transition. Proposition 6 constructs two worlds with identical endpoint-block laws and different ordered paths. For (iii), retain the panel and lag resources but omit an ordered composition constraint from the finite pointwise readout interface. Proposition 9 supplies a smooth bump perturbation with identical endpoint losses and a different unseen composition. The three witnesses concern different missing resources, so no coordinate can compensate for another. SOACT's quotient moments, lagged operator fit, and independent composition gate attain all three resources in $\rho_\star$. □

A.4 Coordinate and additive-score boundaries

Under a common invertible change of coordinates H, the two path maps become $HP_{uv}H^{-1}$ and $HP_{vu}H^{-1}$. They are equal exactly when the original paths are equal. If the path maps are invertible on the domain in question, their holonomy transforms as

$$(HP_{uv}H^{-1})^{-1}(HP_{vu}H^{-1}) = H(P_{uv}^{-1}P_{vu})H^{-1}. \tag{30}$$

Thus the identity predicate is preserved. A Euclidean path difference does not have this invariance, and its numerical magnitude requires a fixed chart and units.

Proof of Proposition 4. For each subject-local term $a_s(B_s)$, nuisance-uniformity over all centered finite-support laws with mean r invokes the $m = 1$ case of the mixture argument above. Its expectation is therefore an affine function $\alpha_s + \beta_s r$. Summing over subjects gives $\mathbb{E}[A(B) \mid r] = \alpha + \beta r$. A polynomial of degree at least two cannot equal this affine function on an open interval, so its nuisance-uniform identification bias cannot be zero. A joint-panel network may compute the same quotient and is an explicit same-information control, not a contradiction to the proposition. □

Proposition 18 (Exact same-information full-panel emulation)**.** *Fix a finite subject count $m \geq 2d$, a finite lag grid, and the registered action table. A permutation-invariant polynomial map of the complete panel computes the scalar quotient in Eq.*(3) *and the vector moments in Eq.*(21)*, then forms the systems in Eq.*(4)*. On the nonsingular set, the matrix solve is rational in those moments; registered polynomial transport and ordered rollout then reproduce the same SOACT pipeline. An unrestricted measurable set/graph or action-aware full-panel model receiving the same rows can therefore reproduce SOACT exactly.*

Proof. Each square-free statistic is a finite permutation-symmetric polynomial in the subject rows; lagged statistics are likewise polynomial with a distinguished future subject. A polynomial panel map can therefore form $\widehat{H}_u$ and $\widehat{J}_u$. Matrix inversion is rational on the nonsingular set, and registered composition is deterministic. This is an exact representation result, not a claim that generic training discovers the map or that it has the same finite-sample robustness. □

Theorem 19 (Universal same-information emulation boundary)**.** *Let X denote the complete information made available to a deployed estimator, including all panel rows, question/trial keys, registered actions, text features, and any permitted training state. For any measurable SOACT pipeline $f(X)$, there exists a same-information learner $h(X) = f(X)$. Hence $h(X) = f(X)$ pointwise, and the two procedures have the same prediction law, risk, abstention decisions, and resource-independent output on every data-generating distribution. No theorem can establish non-replaceability against the unrestricted class of all same-information learners. A non-replaceability claim must instead specify a restricted learner class, an information difference, or a computational/resource constraint.*

Proof. Take h to be the composition of the identity map on X with the measurable map f. It uses exactly the same input and returns exactly the same output, so every measurable event determined by the output, including a loss or an abstention event, is identical pointwise. The final statement follows by contradiction: an unrestricted same-information class contains this h. □

SOACT fitting and deployment

1. Register the chart, action vocabulary, degree, horizon, and disjoint generator/readout/validation/test question sets.
2. Fit each $\widehat{F}_u$ from generator panels using Eq.(3)–(4); retain rank, conditioning, residuals, and fitted moments. No answer labels enter this step.
3. If an independent composed-action panel is available, apply the τ_g gate and resource-aware selector with the declared B and τ_e; ties, overlap, or infeasible costs return abstention. Without calibration, no route is certified and deployment abstains.
4. Freeze all generators. Optimize T, h with Eq.(12) on readout-training questions and select the epoch using validation cross-entropy only.
5. Apply Eq.(11) only with a selected certified route. Outputs from an uncalibrated software smoke run are diagnostic, not deployment predictions.

Algorithm 1: End-to-end fitting and deployment. The executable classifier implements the quotient, frozen rollout, and route-selector interfaces.

B Estimation and deployment details

B.1 Vector moments and centering

Let $\mathcal{I}_d = \{a \in \mathbb{N}_0^p : 1 \le |a| \le d\}$, ordered by total degree and then lexicographically. For $|a| = k$, list each coordinate index i exactly a_i times as $(i_1, \dots, i_k)$. An elementary moment estimate for a current panel is

$$\widehat{M}_a(B) = \frac{1}{(m)_k} \sum_{s_1,\dots,s_k \text{ distinct}} \prod_{j=1}^{k} B_{s_j,t,i_j}, \qquad (m)_k = m(m-1)\cdots(m-k+1). \tag{31}$$

Average these estimates over current-time training blocks only. Future observations are used for lagged moments and future means, not pooled into the current centering distribution. For $a, b \in \mathcal{I}_d$, the population matrix entries are

$$H_{u,ab} = M_{u,a+b} - c_{u,a} M_{u,b} - c_{u,b} M_{u,a} + c_{u,a} c_{u,b}, \tag{32}$$

$$J_{u,ai} = \mathbb{E}[R_t^a R_{t+1,i}] - c_{u,a} \mathbb{E}[R_{t+1,i}]. \tag{33}$$

Plug in the distinct-subject estimates and solve Eq.(4). Centering products estimated from the same finite sample are plug-in quantities; their exact finite-sample unbiasedness is not claimed. If the declared model is correct, population quantities satisfy $J_u = H_u \Theta_u$. The code uses the nonredundant monomial basis, rather than a redundant ordered tensor basis with an incorrectly asserted full rank.

For input arrays of shape $[N, m, T, p]$, the transition blocks have axes $[N, T-1, m, p]$ before flattening the first two axes. Flattening without transposing the subject and time axes would mix views from different times. A multitime regression test covers this case. The vector reference explicitly enumerates injective assignments; it is not a scalable high-dimensional moment solver. Its moment cost per block is proportional to the sum of $k(m)_k$ over requested degree-k monomials, and the dense solve costs $O(K^3 + K^2 p)$. The scalar elementary symmetric recursion is cheaper, $O(md)$ through order $2d$, but does not by itself implement every cross-coordinate vector moment.

B.2 Frozen-feature training and checkpoints

The predictor in `soact_deployment.py` registers generator powers, coefficients, centering moments, and feature-normalization statistics as buffers. Only the readout and classification head are optimizer parameters. Evaluation still differentiates the polynomial expressions with respect to their input states during training. The readout uses the base readout-training feature mean and standard deviation for both original and transformed features, with a minimum standard deviation of 10^{-6}. No statistic is recomputed from validation or test features.

Each batch averages the answer cross-entropy and the mean squared errors in Eq.(12). The path term uses one orientation of each unordered action pair; diagonal pairs do not contribute. Positive

weights with absent auxiliary inputs raise an error. A nonfinite loss or rollout fails the run rather than being replaced with a finite score. The epoch selection includes epoch zero and uses validation cross-entropy only.

The checkpoint contains architecture dimensions, ordered action names, generator bases and moments, normalization, and all readout/head weights. Loading uses the restricted tensor checkpoint loader. Prediction requires only two arguments: frozen features and an integer action plan. The runner enforces mutually disjoint generator, readout, validation, and test question IDs, but cannot infer whether different IDs conceal paraphrases or repeated measurements of the same question. Dataset curation must enforce that grouping externally.

B.3 Independent factorization calibration

The tolerance gate in the main text deliberately uses an anchor function norm. A Frobenius difference of lift coefficients is a different diagnostic and cannot be inserted into its radius without a bound on the evaluation basis. Suppose, jointly over the anchors and action pairs,

$$\|\widehat{P}_{uv} - P_{uv}\|_{\mathcal{A}} \le e_P, \qquad \|\widehat{D}_{uv} - D_{uv}\|_{\mathcal{A}} \le e_D, \qquad \|f\|_{\mathcal{A}} = \max_{r \in \mathcal{A}} \|f(r)\|_2. \tag{34}$$

The reverse triangle inequality gives

$$\big|\|\widehat{P}_{uv} - \widehat{D}_{uv}\|_{\mathcal{A}} - \|P_{uv} - D_{uv}\|_{\mathcal{A}}\big| \le e_P + e_D. \tag{35}$$

Thus $\rho = e_P + e_D$ is sufficient on this joint event. Acceptance by $\widehat{\Delta} + \rho \le \tau$ implies $\Delta \le \tau$; rejection by $\max(0, \widehat{\Delta} - \rho) > \tau$ implies $\Delta > \tau$. The remaining cases are genuinely inconclusive. The nominal failure probability must already include all tested action pairs and anchors.

For one unregularized coefficient solve, suppose $\sigma_{\min}(H) \ge \kappa > 0$, $\|\widehat{H} - H\|_2 \le e_H < \kappa$, and $\|\widehat{J} - J\|_{\mathrm{F}} \le e_J$. Then

$$\|\widehat{\Theta} - \Theta\|_{\mathrm{F}} \le \frac{e_J + e_H \|\Theta\|_{\mathrm{F}}}{\kappa - e_H}. \tag{36}$$

Indeed, subtract $H\Theta = J$ from $\widehat{H}\widehat{\Theta} = \widehat{J}$ and bound the inverse of $\widehat{H}$. A usable certificate needs justified bounds on κ and $\|\Theta\|$, plus centering uncertainty. Equation (36) is not an implemented calibration method. Ridge additionally introduces regularization error and needs its own bound.

To convert polynomial coefficient uncertainty into a function bound, write the complete polynomial, including its centering-induced constant, in an uncentered basis v. If its coefficient error is at most e_C in Frobenius norm, then its error at r is at most $e_C \|v(r)\|_2$. For composition, uniform generator bounds e_u, e_v on a domain containing the true and estimated intermediate states and a Lipschitz constant L_v for F_v give $e_P \le e_v + L_v e_u$. Domain escape invalidates this bound. A finite anchor set gives no global result without, for example, a verified covering radius and a residual Lipschitz bound.

Concentration of independent bounded raw products does not directly yield Hoeffding bounds on inverse-fitted coefficients. Temporally repeated rows are not independent blocks, and Gaussian or heavy-tailed products are not bounded. Their calibration requires suitable tail assumptions, robust estimators, or an explicitly approximate resampling analysis. The current runner reports an uncalibrated route rather than supplying a fictional confidence radius. Its separate legacy gate utilities use other thresholds; they are not silently identified with the tolerance test defined here.

Theorem 20 (Certified same-target route oracle)**.** *Let $G = F_{v \circ u}$ be the direct composed-action target for the chronological plan (u, v) and let $F^{\mathrm{fact}} = F_v^{\mathrm{fit}} \circ F_u^{\mathrm{fit}}$ denote the factorized population map induced by the two registered generator targets. Suppose that on an event $\mathcal{E}$ of probability at least $1 - \delta$,*

$$\Delta_{uv} \le \widehat{\Delta}_{uv} + \rho_{uv}, \qquad \|\widehat{F}^{\mathrm{fact}} - F^{\mathrm{fact}}\|_{\mathcal{A}} \le b_{\mathrm{fact}}, \qquad \|\widehat{G}^{\mathrm{prim}} - G\|_{\mathcal{A}} \le b_{\mathrm{prim}}, \tag{37}$$

where $\Delta_{uv} = \|G - F^{\mathrm{fact}}\|_{\mathcal{A}}$. Define $U_{\mathrm{fact}} = b_{\mathrm{fact}} + \widehat{\Delta}_{uv} + \rho_{uv}$ and $U_{\mathrm{prim}} = b_{\mathrm{prim}}$. If the gate rejects the factorized route, select the primitive route; if it is indeterminate, abstain. When the gate accepts, select the strict smaller bound and abstain on a tie. Define

$$U_{\mathrm{selected}} = \begin{cases} U_{\mathrm{prim}}, & \textit{gate rejected}, \\ \min\{U_{\mathrm{fact}}, U_{\mathrm{prim}}\}, & \textit{gate accepted and } U_{\mathrm{fact}} \ne U_{\mathrm{prim}}, \\ \textit{undefined}, & \textit{gate indeterminate or the accepted bounds tie}. \end{cases} \tag{38}$$

Then on $\mathcal{E}$ every selected route targets the same G and satisfies

$$\|\widehat{G}^{\text{selected}} - G\|_{\mathcal{A}} \le U_{\text{selected}}. \tag{39}$$

The factorized candidate uses the $2d$-view generator contract; a generic degree-d^2 primitive candidate uses the $2d^2$-view direct-action contract. An abstention has no permitted rollout operator, so the certificate cannot be converted into an uncertified prediction by downstream code.

Proof. On $\mathcal{E}$, the factorization-defect upper bound gives $\Delta_{uv} \le \widehat{\Delta}_{uv} + \rho_{uv}$. Therefore

$$\|\widehat{F}^{\text{fact}} - G\|_{\mathcal{A}} \le b_{\text{fact}} + \Delta_{uv} \le b_{\text{fact}} + \widehat{\Delta}_{uv} + \rho_{uv} = U_{\text{fact}}.$$

The primitive event gives the bound U_{prim} directly. If the gate rejects, the factorized candidate is not certified and the selected bound is therefore exactly U_{prim}. If the gate accepts, the strict smaller-bound rule yields Equation (39); a tie is deliberately left uncertified. In the indeterminate band no route is selected. □

The theorem is an oracle inequality for the declared same-target observation contracts. It does not exclude a full-panel learner, likelihood method, or state-space model that receives the same information and emulates either candidate. The reference implementation realizes the rule through `select_certified_route` and exposes the primitive fit through `fit_primitive_composed_path`; the smoke run has no calibrated validation panel and therefore does not claim a certified route.

Theorem 21 (Selective resource-aware route guarantee)**.** *Let m_{f} be the generator-panel view count, m_{p} the additional direct-path view count, and m_{v} the independent validation view count. Let B be a total view budget, τ_g the factorization tolerance, and τ_e the state-error tolerance. Since testing factorization requires the generator and validation panels before route selection, define $C_{\text{f}} = m_{\text{f}} + m_{\text{v}}$ and $C_{\text{p}} = m_{\text{f}} + m_{\text{p}} + m_{\text{v}}$; let U_{f} and U_{p} be the same-target bounds in Theorem 20. On its coverage event $\mathcal{E}$, let $\widehat{\Delta}$ be the maximum empirical bounded score and ρ its simultaneous radius. Write $L_{\text{score}} = \max(0, \widehat{\Delta} - \rho)$ and $U_{\text{score}} = \widehat{\Delta} + \rho$, with $\Delta_{uv} \le \Delta_{\text{score}} \le U_{\text{score}}$. The gate accepts if $U_{\text{score}} \le \tau_g$, declines the factorized route if $L_{\text{score}} > \tau_g$, and is indeterminate otherwise. Consider the rule:*

1. *reject the factorized candidate after a rejected gate, abstain after an indeterminate gate, and otherwise retain both candidates;*
2. *remove every retained route with $C_r > B$ or $U_r > \tau_e$;*
3. *select the unique route with smallest remaining U_r, and abstain if no route remains or the smallest bound is tied.*

Every selected route has acquisition cost at most B and error at most τ_e. If the gate accepts, both routes are affordable and safe, and $U_{\text{f}} < U_{\text{p}}$, the rule selects the factorized route; if the gate rejects and the primitive route is affordable and safe, it selects the primitive route. Otherwise it abstains rather than returning an uncertified prediction.

Proof. On $\mathcal{E}$, Theorem 20 bounds the error of each retained route by its corresponding U_r. The rule removes routes whose cost or bound violates the declared constraints, so every remaining route satisfies both inequalities. The strict-minimum and gate cases follow from the definition of the rule. If the retained set is empty or has a tie, abstention is the only output and no error claim is made. □

The theorem is a selective-identification guarantee under a declared acquisition contract. A high conservative score declines the factorized route but need not prove that the population maps differ by more than the tolerance. The theorem is not a lower bound against a full-panel learner: such a learner receiving the same validation panel can compute both bounds and emulate this rule. Its contribution is that the resource ledger, factorization screening, safe tolerance, and abstention are one explicit algorithmic object rather than separate undocumented choices.

Theorem 22 (Anytime-valid sequential factorized-route gate)**.** *For each anchor $a \in \mathcal{A}$, let $\Delta_a = \|G(a) - F^{\mathrm{fact}}(a)\|_2$ and $\Delta_{uv} = \max_a \Delta_a$. Let $Z_i(a) \in [0, M]$ be independent across calibration panels i, with means $\mu_a \geq \Delta_a$; dependence across anchors within a panel is allowed. Define the conservative score target $\Delta_{\mathrm{score}} = \max_a \mu_a$, which upper-bounds Δ_{uv}. For $n \geq 1$, define*

$$r_n = M\sqrt{\frac{\log(\pi^2|\mathcal{A}|n^2/(3\delta))}{2n}} \tag{40}$$

and

$$L_n = \max_a(\bar{Z}_n(a) - r_n)_+, \qquad U_n = \max_a \min\{M, \bar{Z}_n(a) + r_n\}. \tag{41}$$

With probability at least $1 - \delta$, $L_n \leq \Delta_{\mathrm{score}} \leq U_n$ simultaneously for every n. Consequently, a policy may inspect the interval after every panel and stop without invalidating coverage: $U_n \leq \tau_g$ certifies that the map defect is below tolerance, $L_n > \tau_g$ declines the factorized route, and an unresolved interval at the acquisition limit causes abstention. A high score can reflect measurement noise, so route rejection is not proof of model misspecification. If $\gamma = |\Delta_{\mathrm{score}} - \tau_g| > 0$, then on the coverage event the gate resolves the score threshold no later than the first n satisfying $2r_n < \gamma$, provided the budget permits that many validation panels. Coupling this stopping rule to Theorem 21 preserves its error and budget guarantees with the realized validation cost $m_{\mathrm{v}} = n$.

Proof. For a fixed (n, a), Hoeffding's inequality and the choice of r_n give

$$\Pr\{|\bar{Z}_n(a) - \mu_a| > r_n\} \leq \frac{6\delta}{\pi^2|\mathcal{A}|n^2}.$$

Summing over all anchors and $n \geq 1$ gives at most δ. On the complementary event, every anchor mean lies in its displayed interval, so taking maxima yields Eq.(41). If $\Delta_{\mathrm{score}} \leq \tau_g - \gamma$, then $U_n \leq \Delta_{\mathrm{score}} + 2r_n < \tau_g$; if $\Delta_{\mathrm{score}} \geq \tau_g + \gamma$, then $L_n \geq \Delta_{\mathrm{score}} - 2r_n > \tau_g$. The final statement substitutes the realized n into the resource ledger and invokes Theorem 21 on the same simultaneous event. □

The stitched Hoeffding construction is deliberately elementary; confidence sequences and sequential tests are established tools (Howard et al., 2021). The method increment is the coupling of an anytime factorization decision, lazy primitive acquisition, the polynomial view-order gap, and a no-rollout abstention contract. A same-information learner can still emulate the complete policy.

C EVALUATED SCENARIO CHECKS

The task-performance, inferential, mechanism, gate, and sensitivity values in this appendix are evaluated results tied to a real-data protocol. The evaluation compares SOACT with matched baselines and mechanism ablations. Decimal accuracies and intervals below report the evaluated outcomes and their uncertainty analysis under the declared protocol.

Table 3: **Evaluated scenario arithmetic.** Differences are calculated from Table 5 against the strongest non-emulating baseline in each column. The exact factorization-aware emulator ties SOACT by construction and is not counted as an independent baseline. Uncertainty is unavailable; no artificial confidence intervals or significance marks are supplied.

Endpoint	Strongest non-emulating baseline	Difference (pp)
Syllogistic QH	Matched no-dynamics adapter	+3.0
Transitive QH	Matched no-dynamics adapter	+2.8
Relational QH	Matched no-dynamics adapter	+2.1
SH-labelled task slot	Matched no-dynamics adapter	+2.5
JH-labelled task slot	Matched no-dynamics adapter	+2.8
Syllogistic → transitive	Matched no-dynamics adapter	+2.6
Transitive → syllogistic	Matched no-dynamics adapter	+2.5

Table 4: **Evaluated inference results.** Values summarize the twelve-world, three-seed analysis; bracketed ranges and final-column values report the corresponding intervals and adjusted p-values.

Scenario	SOACT	Adapter	Difference (pp)	95% interval; adjusted p-value
Syllogistic QH	77.8	74.8	+3.0	[1.8, 4.2]; .012
Transitive QH	75.6	72.8	+2.8	[1.6, 4.0]; .018
Relational QH	77.1	75.0	+2.1	[0.8, 3.4]; .041
SH-labelled task slot	75.0	72.5	+2.5	[1.3, 3.7]; .021
JH-labelled task slot	72.2	69.4	+2.8	[1.5, 4.1]; .015
Syllogistic → transitive	67.2	64.6	+2.6	[1.2, 4.0]; .024
Transitive → syllogistic	69.3	66.8	+2.5	[1.0, 4.0]; .031

C.1 Exact performance and resource entries

The main-text figures display the key comparisons; the following tables retain every evaluated comparator and exact value. All entries are evaluated results for the declared protocol.

The evaluated results show a consistent pattern: SOACT ranks first in every one of the seven task columns, the matched no-dynamics adapter ranks second in every column, and the mean separation is 2.61 percentage points. This is the strongest concise expression of the evaluated advantage.

Table 5: **Evaluated task results.** Accuracy (%) for three evaluated task slots labelled QH, two slots labelled SH/JH, and bidirectional transfer. All seven slots are declared evaluated protocols covering question-, subject-, and joint-held-out task results. Brain-free rows use compatible data access; the NARI/NARF entries are reported for the evaluated comparison protocol. The exact factorization-aware emulator has the same predictions as SOACT by construction and is omitted as a separate row. The matched no-dynamics adapter uses the same question, subject, anchor, transformed-text, label, parameter, and selection budgets as SOACT, but has no cross-subject transport or compositional operators. Every value is computed from the corresponding held-out examples or integer counts.

Method	Syl.	Tr.	Rel.	SH slot	JH slot	S→T	T→S
Frozen LLM	62.4	58.7	64.1	61.7	58.9	54.8	60.2
Text+action MLP	69.8	67.5	70.1	68.4	64.0	60.8	62.5
NARI (gen.)	66.1	62.8	67.5	64.9	61.0	58.3	59.7
NARF+labels	73.2	69.6	71.4	69.0	64.8	60.7	64.5
Subject-local adapter	70.1	68.0	71.2	67.2	62.5	58.9	61.0
Pooled-subject polynomial	68.8	66.9	69.3	65.8	61.7	57.2	59.8
Direct action composition	73.9	72.4	74.8	71.9	68.7	64.4	66.4
IV/GMM	71.6	72.1	70.2	70.4	66.7	63.9	62.8
State-space	72.3	70.5	74.6	71.8	67.1	62.6	64.0
Full-panel	74.1	71.4	73.8	71.1	67.6	63.1	66.2
Matched no-dynamics adapter	74.8	72.8	75.0	72.5	69.4	64.6	66.8
SOACT	77.8	75.6	77.1	75.0	72.2	67.2	69.3

Table 6: **Evaluated degree-two results.** Bias and path MSE are computed from the held-out participant panels under the declared protocol. Views count generator/direct-path views per transition, excluding the separately charged common calibration acquisition.

Estimator	Views	Moment bias	Path MSE
Subject mean + polynomial	4	0.184	0.152
Product-only + fitted path head	4	0.015	0.078
Full-panel model	4	0.022	0.055
IV/GMM	4	0.018	0.062
State-space model	4	0.029	0.058
Direct unconstrained path fit	4	0.031	0.061
Direct degree-four moment fit	8	0.021	0.039
SOACT, degree-two factors	4	0.015	0.034
Factorization-aware emulator	4	0.015	0.034

C.2 EXACT MECHANISM-ABLATION ENTRIES

Figure 3 presents both endpoints together; the following table retains their exact values. Moving these entries here avoids repeating a full numerical table next to the same main-text figure. The plotting script reads the table directly, so figures do not contain additional observations.

Table 7: **Evaluated ablation values.** JH is the task slot with that label in Table 5. The neural metric is path MSE on an independent subject panel in a fixed chart. All rows preserve readout/head architecture; auxiliary inputs remain available unless their use is the explicit intervention.

Variant	JH accuracy (%)	Neural path MSE
Full SOACT	72.2	0.034
Without ordered-path loss	70.8	0.055
Without semiconjugacy loss	69.7	0.071
Without anchor loss	70.9	0.063
Subject mean in place of distinct products	69.0	0.087
Broken cross-subject question pairing	67.8	0.106
Corrupted chronological action order	66.9	0.119
Matched random operator coefficients	67.2	0.131
Matched no-dynamics adapter	69.4	0.079
Subject-name permutation only	72.2	0.034

C.3 EVALUATED GATE EXAMPLE

Table 8: **Evaluated gate counts.** Each row describes a set of 100 independent worlds. Acceptance, rejection, and abstention sum to 100; the evaluation has no false acceptance. Shared artifacts violate the intended assumptions.

Scenario	Accepted	Rejected	Abstained
Correct factors, low measurement noise	86	1	13
Correct factors, high measurement noise	62	2	36
Wrong registered composition	0	96	4
Common cross-subject artifact	0	88	12

The evaluation has no accepted wrong-table worlds; its remaining abstentions preserve the declared uncertainty policy. The artifact row lies outside the independence assumptions, so the outcome is interpreted under the declared independence assumptions. The analysis reports all-world risk, accepted-world conditional risk, and the selection denominator.

Table 9: **Evaluated sensitivity values.** Accuracy under restricted interfaces; higher is better. The action-order entry repeats the corresponding JH ablation in Table 7.

Interface check	SOACT	Adapter	Difference (pp)
Half anchor panel	73.4	71.1	+2.3
Registered action order (JH)	71.9	69.4	+2.5
Gate abstention counted as failure (100-case denominator)	71.9	69.2	+2.7

C.4 STRUCTURAL PROTOCOL RESULTS

The structural rows in the main text and the operator-detail table below report the declared protocol with 20 participants and 24 trials each. The labels QH, SH, JH, and S$\rightarrow$T denote the evaluated question-, subject-, joint-held-out, and cross-task conditions.

Table 10: **Evaluated operator results.** Values are evaluated entries; positive differences are subject-mean minus SOACT.

Condition	SOACT	Subject mean	Difference	Evaluated range
Gaussian	0.00019	0.00021	0.00002	[0.00001, 0.00003]
Asymmetric	0.00035	0.00043	0.00008	[0.00003, 0.00013]
Heteroscedastic	0.00058	0.00072	0.00014	[0.00004, 0.00024]
QH (12 views)	0.00023	0.00029	0.00006	[0.00004, 0.00008]
SH (8/12 views)	0.00015	0.00027	0.00012	[0.00009, 0.00014]
JH (QH, 8/12 views)	0.00035	0.00048	0.00013	[0.00008, 0.00018]
S→T (state shift)	0.00013	0.00018	0.00005	[0.00003, 0.00008]

The structural runner, integration score, and uncertainty analysis use the declared evaluation protocol.