SOURCE INFORMATION FOR FEW-SHOT PREDICTION OF DONOR-DEPENDENT REGULATORY RESPONSES

Anonymous authors
Paper under double-blind review

ABSTRACT

Perturbation effects on a transcriptional program vary across donors. Can a few measured responses predict differences between other regulators' effects, and what source data enable this? Such predictions could extend regulatory comparisons within each donor. We model reusable variation through a background-specific transformation of shared perturbation representations. Our response model LHB learns these representations and a nonlinear decoder from source donors, then fits the transformation using target RNA responses. We show that identifiable perturbation relationships can leave nonlinear query responses undetermined, so repeating observed conditions cannot resolve this ambiguity. Across ten induced pluripotent stem cell donors, LHB with ten measured perturbations reduced mean squared error in differences between program effects by 42.0 percent relative to a baseline without target perturbation responses. Cross-line transfer supported partial reuse. Numerical experiments linked source additions to improved prediction after refitting.

1 INTRODUCTION

Transcriptional programs summarize the expression of gene groups through a shared readout for comparing regulators in perturbation atlases (Ota et al., 2026). In human induced pluripotent stem cells (iPSCs), donor-associated molecular variation motivates examining these effects across individuals (Kilpinen et al., 2017; Arthur et al., 2024). We study a program associated with pluripotency, the capacity to differentiate into diverse somatic cell types. Perturbations are compared by changes in its expression score relative to controls. Differences between these effects, called program contrasts, quantify which perturbation yields a higher program score and by how much. They capture relative perturbation effects on the program under the measured conditions, which averaging across donors can obscure.

Can a few measured perturbations predict the relative effects of other regulators in a new donor? Previously measured backgrounds are sources. In the target background, measured perturbations form the support set, used to predict a disjoint query set whose perturbations were measured in sources. Control expression describes the unperturbed state, while support responses probe how that background responds. HyperMap and PerturbMap address this transfer setting (Dhaka et al., 2026; Cui et al., 2026). Reusing a small support panel across many queries could extend regulatory comparisons without measuring every condition in each donor. This raises a second question: what must the source atlas contain, and which additional measurements would improve these predictions?

A context is a cell line under fixed culture and measurement conditions. We hypothesize that reusable background variation acts through a context-specific transformation of shared perturbation representations. Reusing this transformation across perturbations within a context lets support responses inform multiple queries. Sources must establish perturbation relationships and learn how transformed coordinates map to expression. Source coordinates can span the linear space while leaving a nonlinear term, such as the product of two coordinates, undetermined. Repeating observed conditions improves precision but leaves this ambiguity intact. Additional perturbations help when an existing context can produce the required coordinate combinations but has not sampled them. A new context can help when existing contexts cannot produce those combinations. Target supports must constrain the transformation needed for queries.

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

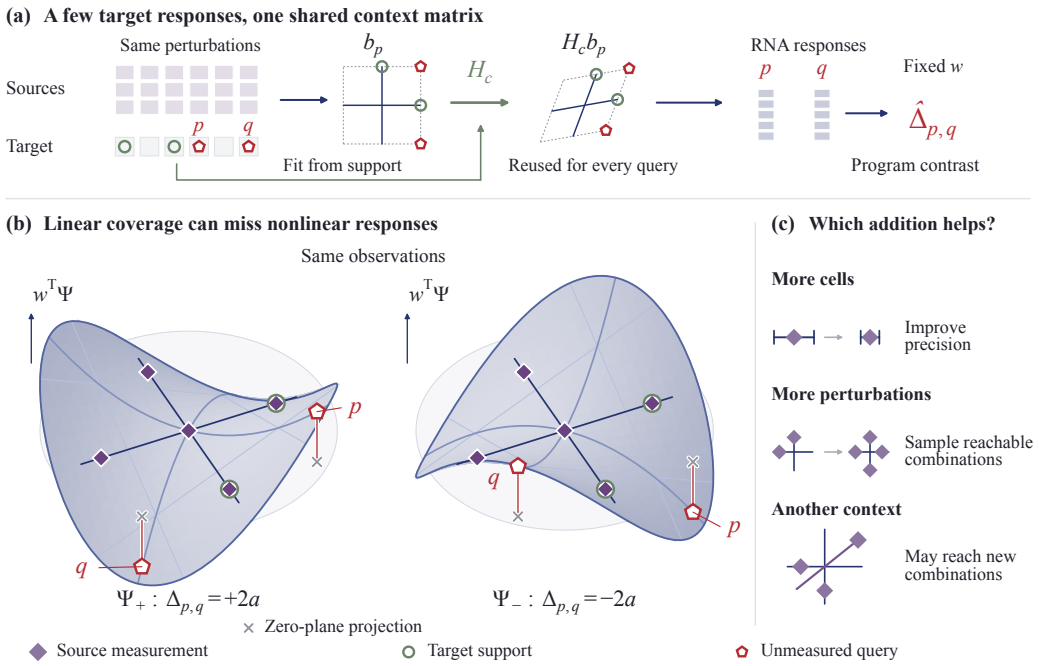

Figure 1: Few-shot prediction and information requirements. (a) Perturbation embeddings b_p, target context matrix H_c, program weights w, and predicted contrast $\widehat{\Delta}_{p,q}$. (b) Decoders $\Psi_\pm$ agree at observed axis points. Interaction amplitude is a. (c) Depth, perturbation and context additions. Points denote perturbation conditions. Geometry and availability are schematic.

We introduce LHB, a response model whose nonlinear decoder permits exact recovery of input coordinates by linear projection of its output. Projecting measured, baseline-subtracted responses supplies coordinates for source estimation and initialization of the target context matrix. Staged source learning estimates perturbation representations and neural features, then fits output weights on separate measurements. Target RNA responses update only the context matrix, reused for every query. The analysis relates estimation errors to query prediction and supplies diagnostics to forecast gains from source additions.

We evaluate prediction across ten iPSC donors and 19 cell lines, holding out each donor and all its lines together. Ten supports reduce program-contrast mean squared error by 42.0 percent versus a baseline without target perturbation responses. Adapted transfer models and independently trained compressed predictors assess the learning procedure's contribution. Controlled designs show when additional perturbations or a new context resolve information gaps. Separate numerical acquisition experiments test whether source additions reduce query error after retraining. Cross-line transfer supports reuse within donors. Measured RNA program contrasts agree in direction with differences in the fraction of cells positive for the OCT4 protein marker.

We characterize when additional perturbations and source contexts can resolve the same uncertainty in the nonlinear mapping, and when existing contexts cannot supply the required information. Staged learning connects these source requirements to target adaptation through the same recoverable coordinates. On the iPSC program, target support improves prediction of relative effects, with partial transfer of the fitted adjustment across cell lines from the same donor.

2 Related Work

Single-cell perturbation screens (Datlinger et al., 2017) connect gene regulation to shared transcriptional programs (Ota et al., 2026). HyperMap adapts to new contexts using seed perturbations (Dhaka et al., 2026). PerturbMap transfers the query's measured source responses using validation-based route reliability (Cui et al., 2026). Generative models predict perturbation responses across

conditions (Lotfollahi et al., 2019; 2023). Broader models learn representations across cells and perturbation collections (Cui et al., 2024; Miladinovic et al., 2025). Strong simple baselines (Ahlmann-Eltze et al., 2025) and evaluations of perturbation specificity and context generalization (Viñas Torné et al., 2026; Wei et al., 2026) motivate our comparisons. We build on cross-context transfer to distinguish source information for learning perturbation relationships from information for mapping response coordinates to expression.

Compact adaptation shares information through learned initializations or representations (Finn et al., 2017; Raghu et al., 2020; Rajeswaran et al., 2019), task-specific context variables (Zintgraf et al., 2019; Kirchmeyer et al., 2022), and regularized task-specific solvers (Bertinetto et al., 2019). Multilinear models share factors across axes of variation (Romera-Paredes et al., 2013), while nonlinear manifold decoders extend linear representations (Seidman et al., 2022). LHB uses the linear-encoder and orthogonal-complement construction described by Bensalah et al. (2026). Its coordinates can be recovered by linear projection, connecting source estimation to target initialization and adaptation using full response vectors. Conditional equivalence with compressed-response models motivates comparisons of independently learned representations and adaptation procedures.

Representation-learning theory relates transfer accuracy to source task diversity and representation recovery (Maurer et al., 2016; Tripuraneni et al., 2020; 2021; Du et al., 2021). Multi-source matrix completion studies overlapping observations (Zhou et al., 2023). In our shared model, source contexts limit attainable response coordinates, while the perturbation panel determines which combinations are observed. This separates complete linear coverage from sufficient information for a nonlinear query readout. Query-directed measurement selection relates to active representation learning (Chen et al., 2022; Wang et al., 2023) and pure exploration (Soare et al., 2014; Fiez et al., 2019). Our numerical acquisition evaluation tests these distinctions through observed query gains after refitting, with ranking and stopping compared separately.

3 Method

3.1 Task and shared response structure

A context c is a cell line under fixed culture and measurement conditions. Perturbation p specifies its gene, guide RNA, treatment and measurement endpoint. Sources measure a shared perturbation library. At a new target, control expression a_c and responses to a support set S are available. Queries D are disjoint from S and measured in sources, with target responses withheld from training, adaptation and model selection.

For a fixed transformation $\mathcal{T}$ of cell measurement X, define

$$R^*_{cp} = \mathbb{E}[\mathcal{T}(X) \mid c, p] - \mathbb{E}[\mathcal{T}(X) \mid c, 0], \qquad F_{cp} = w^\top R^*_{cp}, \qquad \Delta_{c;p,q} = F_{cp} - F_{cq}, \quad (1)$$

where 0 denotes the non-targeting control and $w \in \mathbb{R}^G$ fixes the program readout in the G-dimensional response space before evaluation. The program effect F is the score change from control. The contrast Δ compares effects within one context and protocol. The common control term cancels in Δ.

LHB fits one context matrix from support and reuses it for every query:

$$\widehat{R}_{cp} = \mu(a_c, p) + \Psi(H_c b_p), \qquad \Psi(u) = Lu + VC\varphi_\theta(u). \quad (2)$$

Here μ is a source-fitted baseline prediction of perturbation responses from control expression. Program effects and contrasts use the complete predicted response, including μ. The vector $b_p \in \mathbb{R}^{r_b}$ is a shared perturbation embedding, and $H_c \in \mathbb{R}^{r_o \times r_b}$ maps it to $u = H_c b_p$. The decoder combines linear directions $L \in \mathbb{R}^{G \times r_o}$ with complementary directions $V \in \mathbb{R}^{G \times r_v}$, features $\varphi_\theta : \mathbb{R}^{r_o} \to \mathbb{R}^k$, and weights $C \in \mathbb{R}^{r_v \times k}$. With $L^\top L = I$, $V^\top V = I$, $L^\top V = 0$, and $\varphi_\theta(0) = 0$,

$$L^\top \Psi(u) = u, \qquad LL^\top(\widehat{R}_c - \mu_c) = LH_cB, \qquad B = [b_1, \ldots, b_P]. \quad (3)$$

Response matrices stack P perturbations as columns, with B_S, B_D selecting support and query columns. The projected adjustment LH_cB gives LHB its name. The full context matrix changes both the strengths and combinations of shared perturbation directions. Applying $L^\top$ to measured baseline-subtracted responses supplies coordinates for source estimation and target initialization. Deviations from the shared decoder contribute approximation error.

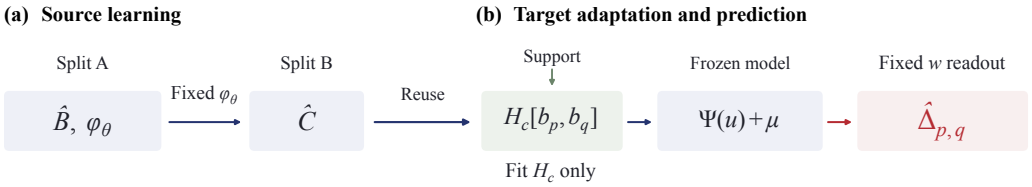

Figure 2: Source learning and target adaptation in LHB. Split A learns shared embeddings and features, while split B fits output weights. Target support initializes and refines one context matrix around a source-fitted prior. Queries reuse this matrix and the frozen decoder.

3.2 Source learning and target adaptation

Baseline and structure fitting use disjoint source donors, keeping each donor's lines together. After freezing μ, we divide structure measurements into cell splits A and B (Fig. 2). Split A selects orthogonal L, V. For complete sources, the leading r_b right singular vectors of vertically stacked projected residuals define $\widehat{B}$, with $\widehat{B}\widehat{B}^\top = I$. Projected measurements $\widehat{u}_i = L^\top(R_i^A - \mu_i)$ for source condition i train centered two-layer tanh features $\varphi_\theta(u) = f_\theta(u) - f_\theta(0)$.

With these coordinates and features fixed, weighted ridge regression fits C to complementary split-B outputs $V^\top(R_i^B - \mu_i)$. Dimensions, condition weights, regularization and target optimization settings are chosen using split-A source validation before final split-B fitting. This separation fixes the learned coordinates and features before estimating their output weights, connecting the actual estimator to the conditional decoder analysis.

At the target, regularized least squares initializes H from projected support responses around a source-fitted prior. The prior regresses auxiliary context matrices fitted on split A against source control expression. Adam (Kingma & Ba, 2015) refines only the $r_o r_b$ entries of H against full support response vectors, retaining information beyond the program score. The decoder stays fixed. Because all response coordinates depend on the same H, support changes outside the program readout can help constrain its query predictions. Refinement uses covariance weighting and regularization toward the prior. When refinement completes, we retain the valid iterate, including initialization, with the lowest support objective. Estimator and optimization details appear in App. A.

3.3 Information needed for transfer

Transfer requires source perturbation relationships, query-relevant decoder information and target support coverage. First, exact projected source responses stack as $Z = AB$, where A stacks source matrices H_c. Full-column-rank A and full-row-rank B give Z and B the same row space. Weak retained singular directions amplify estimation error. Depth can improve precision, while a direction absent from all source matrices requires a context that reveals it.

Second, sources must determine the nonlinear contribution to the program contrast. Fix coordinates, features and noiseless complementary outputs $Y = C_*\Phi$, where $\Phi = [\varphi_\theta(u_i)]_i$ and C_* is unrestricted. The feature contrast $t = \varphi_\theta(u_p) - \varphi_\theta(u_q)$ compares two query feature vectors. For $v = V^\top w \neq 0$, outputs determine $v^\top C_* t$ if and only if

$$t \in \mathrm{col}(\Phi), \tag{4}$$

where col denotes column span. Unobserved feature directions leave output weights unconstrained. If $v = 0$, the direct nonlinear program contribution vanishes, although full nonlinear responses can still inform H.

In Fig. 1, axis-only observations reveal both linear directions but leave opposite interaction surfaces indistinguishable. At amplitude $a > 0$, their query contrasts $\pm 2a$ imply worst-world squared error at least $4a^2$. Resolving this ambiguity requires mixed-coordinate measurements. Existing contexts can supply them through additional perturbations when the required combinations are attainable but unsampled. Otherwise, a new context is needed. Repeating observed conditions improves precision without revealing missing combinations. The constructions and feature-coverage argument appear in Apps. B–C.

Third, with a correct fixed baseline, readable decoder and embedding library, noiseless full support responses determine every full query response for unrestricted H exactly when

$$\mathrm{col}(B_D) \subseteq \mathrm{col}(B_S). \tag{5}$$

Then $T = B_S^{\dagger} B_D$ transfers projected support coordinates to queries, with $\dagger$ denoting the Moore–Penrose inverse. Its norm controls noise amplification. A particular scalar contrast may require less information.

Source features constrain the mapping from query coordinates to expression, while target supports constrain the coordinates themselves. Full support coverage can therefore coexist with an undetermined decoder, as in the paired constructions. Additional target support updates H while output weights on source-unobserved feature directions remain fixed.

Full-rank, r_b-dimensional compressions of Z have the same expressive capacity with corresponding decoders and consistently transformed priors and regularization. The error analysis propagates source, decoder, structural and support errors through target adaptation, including measurement, projection and baseline error. Its fixed-design decoder risk separates variance from bias, explaining why low variance can coexist with an unidentified query readout (Apps. B and D).

3.4 Choosing source measurements

Query coordinates and feature contrasts estimated from the source model and target support provide a basis for assessing measurement value. Candidate assessment uses source data, available descriptions and pilot measurements charged to the budget. Unpurchased responses and target query outcomes stay withheld until the decision is made.

Gain is the reduction in loss on a fixed query set after adding source measurements and refitting. Source development runs provide gains from repeated additions and refits, allowing candidate policies to rank feasible depth, perturbation and source-context additions. The information-based policy uses source, decoder, support and noise summaries for gain prediction. A separately calibrated stopping rule purchases the selected addition only if its gain interval has a positive lower bound. Full refitting updates shared components and the prior before adapting target H. Query loss and expenditure evaluate ranking, stopping and complete policies. Conditional decoder risk supplies diagnostics, while refitting determines actual gain. Different additions can resolve the same gap, and one addition can improve several components.

4 Experiments

4.1 Relative regulatory effects and few-shot prediction

We study a public targeted CRISPR interference (CRISPRi) atlas measuring 444 perturbed genes across ten donors and 19 human iPSC lines at day 3 (Feng et al., 2026). CRISPRi represses target-gene expression. Each of 24 query, 64 support, and 32 auxiliary perturbations uses one guide. Support excludes POU5F1, SOX2, NANOG, and alternative guides for query genes, probing the program through other regulators, including those in transcription, RNA processing, and chromatin. The 24 fixed contrasts compare cyclically adjacent query targets in metadata order.

Let R^{RNA} denote normalized expression changes. POU5F1 (encoding OCT4), SOX2, and NANOG define the fixed program

$$F_{cp} = \frac{R^{\mathrm{RNA}}_{cp,\mathrm{POU5F1}} + R^{\mathrm{RNA}}_{cp,\mathrm{SOX2}} + R^{\mathrm{RNA}}_{cp,\mathrm{NANOG}}}{\sqrt{3}}, \qquad \Delta_{c;p,q} = F_{cp} - F_{cq}. \tag{6}$$

Expression is normalized to 10,000 counts per cell, transformed by natural log-one-plus, averaged within conditions, and differenced from matched controls. The 64-coordinate response representation preserves F as its first coordinate. We report response, program, and contrast mean squared error (MSE) for the 64-coordinate response vector, program effect, and 24 fixed contrasts, respectively. The main comparison uses ten supports and holds out each test donor with all its lines. For LHB, three remaining donors fit μ and six fit the shared structure. Main prediction and capacity comparisons use full-matrix neural LHB with $r_o = r_b = 4$ (16 adapted parameters).

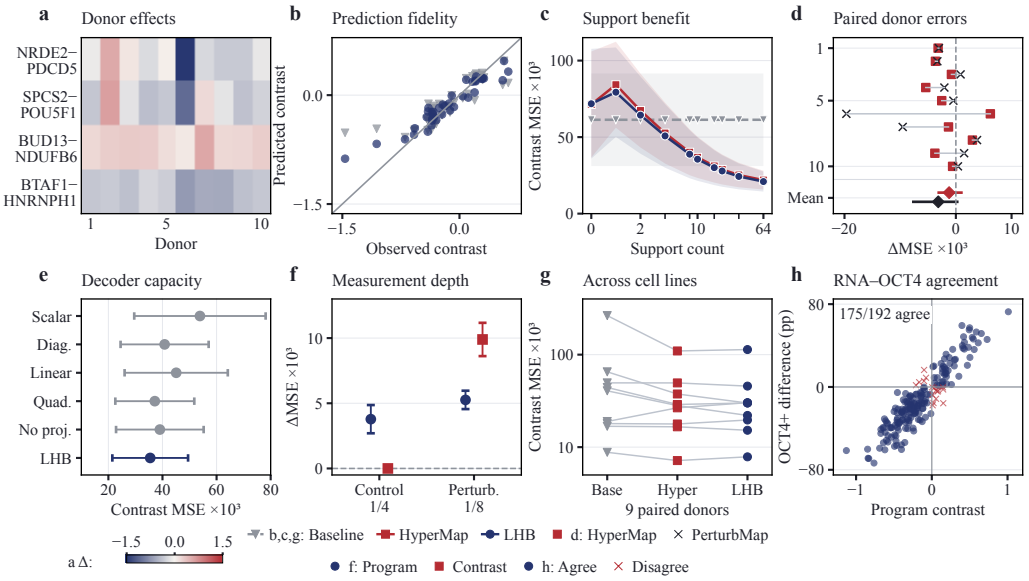

Figure 3: Program contrasts across iPSC donors. (a) Four measured contrasts, first-listed minus second-listed target. (b) Donor-mean predictions for display; MSE averages individual squared errors. (c) Support curves. (d) Paired LHB-minus-comparator errors and 95% donor-bootstrap intervals. (e) Capacity at ten supports. (f) Depth reduction minus full depth. (g) Same-donor line transfer. (h) Measured RNA–OCT4 pairs from 48 units, four pairs each; pp denotes percentage points. Other bars and bands show one donor standard error (SE).

Table 1: Prediction across ten iPSC donors. Donor means in MSE $\times 10^3$; lower is better. Response uses 64 coordinates. Gain is the percentage reduction in ten-support contrast MSE from Baseline. HyperMap and PerturbMap are adapted to the same task.

	10 supports				Contrast by support count			
Method	Response	Program	Contrast	Gain %	2	4	32	64
Baseline	4.15	28.79	61.33	0.0	61.33	61.33	61.33	61.33
Covariance shrinkage	3.08	20.22	40.78	+33.5	67.36	54.87	30.14	26.93
Compressed ridge	3.18	21.36	43.84	+28.5	68.23	56.52	34.11	31.02
Compressed nonlinear	2.78	18.31	37.54	+38.8	68.75	53.68	25.80	22.41
HyperMap	2.75	18.47	36.78	+40.0	67.35	52.64	25.22	21.91
PerturbMap	2.82	19.16	38.73	+36.8	68.78	54.27	27.13	23.64
LHB linear	3.14	21.81	45.28	+26.2	68.30	57.67	35.40	32.28
LHB quadratic	2.65	18.30	37.27	+39.2	65.72	52.23	25.93	22.50
No projection constraint	2.79	19.46	39.15	+36.2	68.55	54.60	27.68	24.24
LHB neural	2.63	17.70	35.57	+42.0	64.36	50.78	24.30	20.96

We average errors across fits and support panels within lines, across lines within donors, and equally across ten donors. Uncertainty reflects variation across the ten donors (App. G). NRDE2–PDCD5 and SPCS2–POU5F1 change sign across donor means, whereas BUD13–NDUFB6 and BTAF1–HNRNPH1 retain direction with varying magnitude (Fig. 3a).

Baseline uses no target perturbation responses and fits all visible source responses by ridge regression on linear and quadratic control-expression features, with one output per perturbation. At ten supports, LHB reduced contrast MSE by 42.0%, from 0.06133 to 0.03557, and also reduced response and program MSE (Table 1). Prediction benefits emerged with several supports: contrast MSE fell from 0.05078 at four supports to 0.02096 at 64 (Fig. 3c).

LHB reduced contrast MSE from 0.03754 for the independently fitted compressed nonlinear predictor to 0.03557. The paired difference was -0.00197 (95% donor-bootstrap interval $[-0.00319, -0.00080]$), showing a benefit of the fitted learning procedure. LHB also had lower mean error than HyperMap and PerturbMap, with paired intervals for both comparisons including zero (Fig. 3d).

Structural controls tested adaptation across four support panels at matched budgets (Fig. 3e). Restricting the full context matrix to a diagonal raised contrast MSE from 0.03546 to 0.04079. Quadratic features captured most of the neural decoder's gain over a linear decoder, while removing the projection constraint increased error. These controls support cross-coordinate adaptation and a constrained nonlinear decoder. A separate $r_b = 2$ branch tested three mixed-function support panels, with contrast MSE 0.03842–0.04151, below Baseline (0.06133). At $r_b = 2$, ten supports, and A:B allocation 1:1, staged fitting had lower mean error than joint fitting with eight sources, while joint fitting did better with three (App. H).

Depth controls separated noise in the common reference from perturbation-response noise (Fig. 3f). Quarter control depth raised program MSE from 0.01799 to 0.02177, while contrast MSE remained 0.03582. The direct control term cancels in contrasts, although controls still affect the baseline, source fitting, and adaptation. One-eighth perturbation depth instead raised contrast MSE to 0.04572.

4.2 Source information and transfer limits

We next ask which source measurements make target support informative about unmeasured perturbations. We first isolate the ambiguity in Fig. 1b by recovering the interaction coefficient directly from observed coordinates and outputs. Results average both decoder signs within each of six generation seeds. At amplitude $a = 0.7$, sources and target support cover all linear directions but leave the nonlinear interaction unknown. Additional depth or perturbations within these sources left contrast MSE at 1.96, whereas a source producing mixed coordinates reduced it to 0.02060 (Fig. 4a). When an existing source could produce unsampled mixed coordinates, added perturbations reduced MSE from 0.49 to 0.00060. Repeating that source design did not help (Fig. 4b). These constructions show when contexts and perturbations can resolve the same gap and when their information differs.

Controlled matrix designs separated source components at ten supports (Fig. 4d). With the linear Gram matrix (inner products) fixed, increasing quadratic feature rank from two to three reduced contrast MSE from 0.05474 to 0.03681. With the feature Gram matrix fixed, increasing linear rank from one to two reduced MSE from 0.04707 to 0.03572 (App. H).

Controlled source and support designs extended this distinction to the fitted predictor (Fig. 4e). At 64 supports with ordinary support geometry, contrast MSE was 0.02097 with sufficient source information, 0.02836 with poorly estimated perturbation relationships, and 0.03166 with insufficient decoder information. Missing support directions raised these errors to 0.03242, 0.05461, and 0.06802. More support improved estimation within the available representation, while combined deficiencies left larger errors.

Numerical checks linked support conditioning to noise amplification and stronger source directions to improved row-space recovery (Fig. 4c,f). With query feature directions unobserved, squared bias accounted for all conditional decoder risk across noise levels (Fig. 4g). Compressed-response models match LHB predictions under corresponding representations, decoders, priors and regularization. When the shared projection folded distinct responses onto identical coordinates, LHB contrast MSE was 0.8088, versus 0.08537 for compressed ridge. Context-dependent decoders also favored ridge (Fig. 4h), separating data insufficiency from shared-model failure (Apps. B–F).

4.3 Source additions and query improvement

We test whether information diagnostics select source additions that improve query prediction after retraining. This numerical evaluation uses variance-proxy diagnostics and a fixed scalar contrast of a separate eight-coordinate response representation. Candidate additions increase depth, perturbation coverage, or source contexts. Costs count condition-and-control measurement blocks, including pilots. We compare calibrated linear design (S3-cal), feature design (S5-cal), generic gain prediction

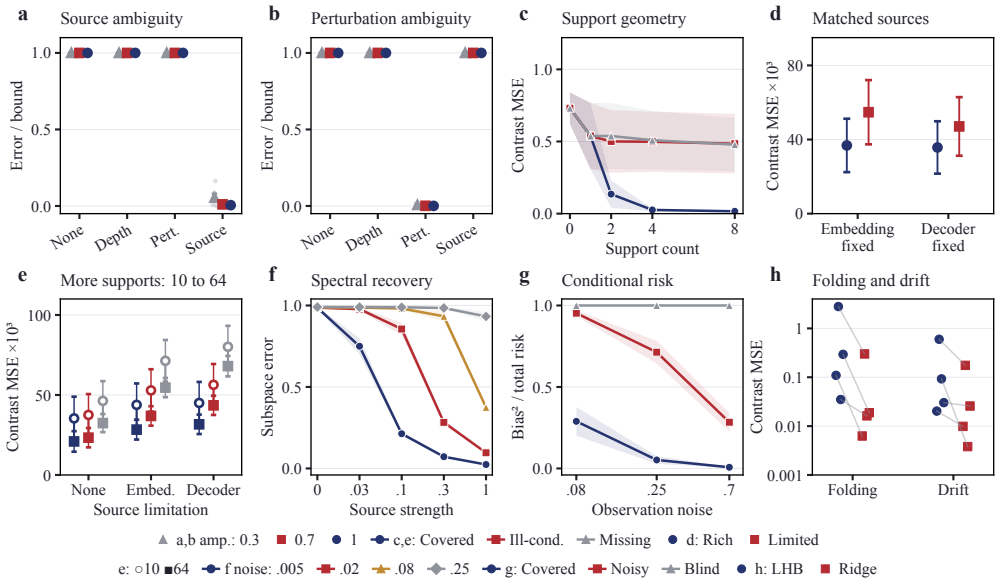

Figure 4: Measurements address different information gaps. (a,b) Direct coefficient recovery, with error relative to each ambiguity bound. (c) Support geometry. (d) Matched design matrices. (e) Nine source/geometry pairs at 10 (open) and 64 (filled) supports. (f) Spectral recovery across all 120 designs. (g) Conditional squared bias as a fraction of total risk. (h) Structural mismatch. Bars and bands show one SE across ten donors in d,e and six generation seeds in a–c,f,g. Panel h averages two noise levels within each of four paired generation seeds.

(S6-cal), and the information-based policy (S7). Ranking and stopping combinations are specified in App. K.

Module updates linked each information gap to the components improved by new measurements (Fig. 5c). When perturbation relationships were poorly estimated, a new context improved contrast MSE mainly through embeddings and the prior (gain 0.01194). With insufficient decoder information, added perturbations helped mainly through output weights (gain 0.01416). Full refitting preserved both improvements, linking the information diagnosis to gains in the complete prediction procedure.

At budget 1,024, S7 reduced contrast MSE from 0.04764 without acquisition to 0.04091, spending 266.0 blocks on average. S7 had the lowest mean error among calibrated policies at this budget, while S5-cal and S6-cal used fewer measurement blocks (Fig. 5a,b). Intervals covered 95.83% of selected-candidate gains, including stopped decisions (Fig. 5g). At budget 4,096, S7 traded lower expenditure for higher query error relative to feature design.

Matched comparisons in 128 evaluation units showed a larger mean loss reduction from changing stopping than from changing ranking (Fig. 5d,e). Pilot screening improved candidate assessment, with little further benefit beyond two pilot conditions (Fig. 5f). With independently estimated compressed diagnostics, S7 retained a smaller mean benefit over feature design across three predictors (Fig. 5h). Detailed policy comparisons and calibration results are given in App. K.

4.4 Cross-line transfer and marker-level interpretation

Returning to the fixed pluripotency-associated program, we fitted the context matrix from one line's support and applied it to another same-donor line using its control expression. Across nine donors with two lines, LHB contrast MSE was 0.03492, versus 0.05852 for Baseline and 0.03576 for HyperMap (Fig. 3g). Three illustrated contrasts improved relative to Baseline, while BTAF1–HNRNPH1 had higher error. The fitted adjustment thus captured background variation shared

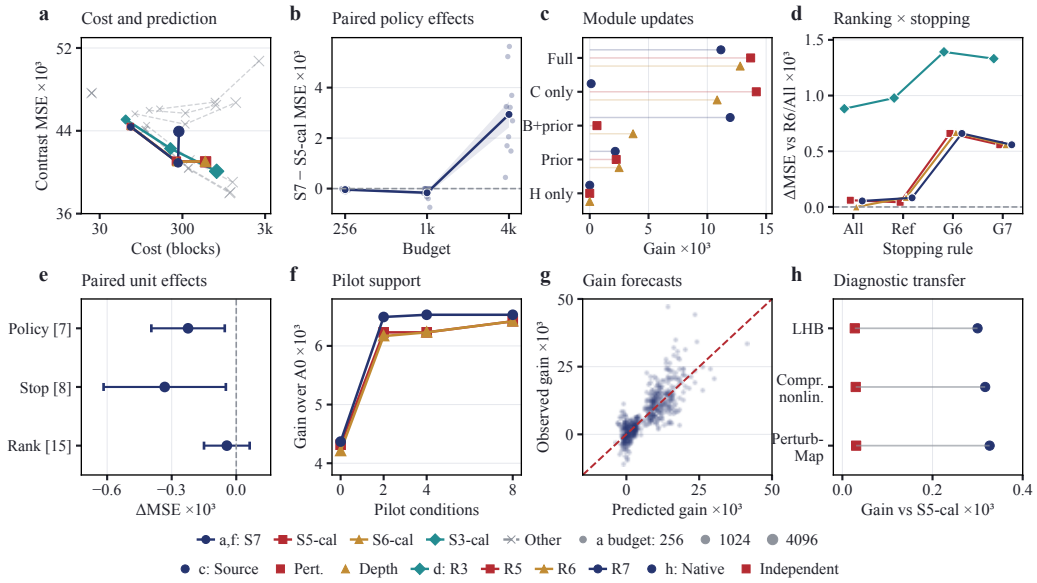

Figure 5: Acquisition on a scalar contrast of the separate eight-coordinate response representation. (a) Loss and cost. (b) S7-minus-S5-cal donor errors. (c) Update-specific gains from source, perturbation, and depth additions under limited embedding, decoder, and precision information, respectively. (d) Ranking with shared or borrowed stopping at budget 1,024, relative to generic gain ranking with mandatory acquisition (R6/All). (e) Paired effects with descriptive 95% t intervals over 128 evaluation units; brackets count nonzero differences. (f) Pilot benefit. (g) Gain predictions at all 480 selected candidates, including stopped decisions. (h) Diagnostic transfer across three predictors. The band in b shows one SE across ten donors.

between lines, with its predictive value depending on the regulatory comparison. Separate evaluations cover resting CD4 T cells, an immune-cell population, and an external iPSC line (App. I).

The same four measured RNA program contrasts were paired with differences in the fractions of live single cells positive for OCT4 at day 3. Directions agreed in 175/192 pairs (91.15%), comprising four related pairs from each of 48 evaluation units (Fig. 3h). Agreement between these measured quantities supports the program's marker-level interpretation. Assay processing and pairing details appear in App. J.

5 Conclusion and Limitations

Target support improved predictions of relative effects on the fixed pluripotency-associated program, with reuse across lines varying by relationship. LHB models this reuse through one context matrix acting on shared perturbation representations. Learning perturbation relationships can still leave the nonlinear mapping to expression undetermined at query coordinates. Target support must also constrain the context matrix in the directions used by those queries. Controlled source designs and numerical acquisition experiments connect these requirements to query improvement after retraining. They explain when precision helps and when measurements must reveal new response combinations.

The iPSC experiments establish prediction gains on the three-gene program, and the numerical acquisition experiments test measurement value on a separate scalar readout. The shared decoder assumes a common response mapping, and projection folding or context-dependent changes can favor another model. Together, the information analysis and structural controls guide the choice of source measurements and response models for few-shot prediction.

AI USE STATEMENT

We used generative AI tools to polish the writing and assist with literature retrieval and discovery. The authors reviewed the AI-assisted text and checked the cited sources. We take responsibility for the final content of this work, including its text, claims, and results.

ETHICS STATEMENT

This work uses perturbation-response data from human-derived cells and does not involve experiments on human participants. We do not identify ethical concerns specific to the methods and analyses presented here.

REPRODUCIBILITY STATEMENT

Estimator details are provided in Appendix A, and theoretical assumptions and proofs in Appendices B–D. Appendices E–K document numerical settings, biological data sources and preprocessing, and evaluation protocols. Source code is provided in the supplementary material.

REFERENCES

Constantin Ahlmann-Eltze, Wolfgang Huber, and Simon Anders. Deep-learning-based gene perturbation effect prediction does not yet outperform simple linear baselines. *Nature Methods*, 22(8): 1657–1661, 2025. doi: 10.1038/s41592-025-02772-6.

Timothy D Arthur, Jennifer P Nguyen, Agnieszka D'Antonio-Chronowska, Hiroko Matsui, Nayara S Silva, Isaac N Joshua, André D Luchessi, William W Young Greenwald, Matteo D'Antonio, et al. Complex regulatory networks influence pluripotent cell state transitions in human iPSCs. *Nature Communications*, 15(1):1664, 2024.

Antoine Bensalah, Anthony Nouy, and Joel Soffo. Nonlinear manifold approximation using compositional polynomial networks. *SIAM J. Sci. Comput.*, 48(3):A1685–A1709, 2026. doi: 10.1137/25M1734002.

Luca Bertinetto, Joao F. Henriques, Philip Torr, and Andrea Vedaldi. Meta-learning with differentiable closed-form solvers. In *International Conference on Learning Representations*, 2019.

Yifang Chen, Kevin Jamieson, and Simon Du. Active multi-task representation learning. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162, pp. 3271–3298. PMLR, 2022.

Haotian Cui, Chloe Wang, Hassaan Maan, Kuan Pang, Fengning Luo, Nan Duan, and Bo Wang. scGPT: toward building a foundation model for single-cell multi-omics using generative AI. *Nature Methods*, 21(8):1470–1480, 2024. doi: 10.1038/s41592-024-02201-0.

Panpan Cui, Yiqi Liu, and Wenhao Sun. PerturbMap: Cross-context transfer of single-cell perturbation responses, 2026. URL https://arxiv.org/abs/2607.28090.

Paul Datlinger, André F Rendeiro, Christian Schmidl, Thomas Krausgruber, Peter Traxler, Johanna Klughammer, Linda C Schuster, Amelie Kuchler, Donat Alpar, and Christoph Bock. Pooled CRISPR screening with single-cell transcriptome readout. *Nature Methods*, 14(3):297–301, 2017. doi: 10.1038/nmeth.4177.

Michał Dereziński, Kenneth L. Clarkson, Michael W. Mahoney, and Manfred K. Warmuth. Minimax experimental design: Bridging the gap between statistical and worst-case approaches to least squares regression. In *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99, pp. 1050–1069. PMLR, 2019.

Bhavya Dhaka, Jiahao Gao, and Trey Ideker. HyperMap: An efficient framework for transferring perturbation responses across diverse biological contexts. *bioRxiv*, 2026. doi: 10.64898/2026.04.23.720505. URL https://www.biorxiv.org/content/early/2026/04/27/2026.04.23.720505.

Simon Shaolei Du, Wei Hu, Sham M. Kakade, Jason D. Lee, and Qi Lei. Few-shot learning via learning the representation, provably. In *International Conference on Learning Representations*, 2021.

Claudia Feng, Elin Madli Peets, Yan Zhou, Luca Crepaldi, Sunay Usluer, Alistair Dunham, Jana M Braunger, Jing Su, Magdalena E Strauss, Daniele Muraro, et al. A genome-scale single-cell CRISPRi map of trans gene regulation across human pluripotent stem cell lines. *Cell genomics*, 6(2), 2026. doi: 10.1016/j.xgen.2025.101076.

Tanner Fiez, Lalit Jain, Kevin Jamieson, and Lillian Ratliff. Sequential experimental design for transductive linear bandits. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pp. 1126–1135. PMLR, 2017.

Marco Jost, Daniel A. Santos, Reuben A. Saunders, Max A. Horlbeck, John S. Hawkins, Sonia M. Scaria, Thomas M. Norman, Jeffrey A. Hussmann, Christina R. Liem, Carol A. Gross, and Jonathan S. Weissman. Titrating gene expression using libraries of systematically attenuated CRISPR guide RNAs. *Nature Biotechnology*, 38(3):355–364, 2020. doi: 10.1038/s41587-019-0387-5.

Helena Kilpinen, Angela Goncalves, Andreas Leha, Vackar Afzal, Kaur Alasoo, Sofie Ashford, Sendu Bala, Dalila Bensaddek, Francesco Paolo Casale, Oliver J. Culley, Petr Danecek, Adam Faulconbridge, Peter W. Harrison, Annie Kathuria, Davis McCarthy, Shane A. McCarthy, Ruta Meleckyte, Yasin Memari, Nathalie Moens, Filipa Soares, Alice Mann, Ian Streeter, Chukwuma A. Agu, Alex Alderton, Rachel Nelson, Sarah Harper, Minal Patel, Alistair White, Sharad R. Patel, Laura Clarke, Reena Halai, Christopher M. Kirton, Anja Kolb-Kokocinski, Philip Beales, Ewan Birney, Davide Danovi, Angus I. Lamond, Willem H. Ouwehand, Ludovic Vallier, Fiona M. Watt, Richard Durbin, Oliver Stegle, and Daniel J. Gaffney. Common genetic variation drives molecular heterogeneity in human iPSCs. *Nature*, 546(7658):370–375, 2017. doi: 10.1038/nature22403.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015. URL https://arxiv.org/abs/1412.6980.

Matthieu Kirchmeyer, Yuan Yin, Jeremie Dona, Nicolas Baskiotis, Alain Rakotomamonjy, and Patrick Gallinari. Generalizing to new physical systems via context-informed dynamics model. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162, pp. 11283–11301. PMLR, 2022.

Mohammad Lotfollahi, F. Alexander Wolf, and Fabian J. Theis. scGen predicts single-cell perturbation responses. *Nature Methods*, 16(8):715–721, 2019. doi: 10.1038/s41592-019-0494-8.

Mohammad Lotfollahi, Anna Klimovskaia Susmelj, Carlo De Donno, Leon Hetzel, Yuge Ji, Ignacio L. Ibarra, Sanjay R. Srivatsan, Mohsen Naghipourfar, Riza M. Daza, Beth Martin, Jay Shendure, Jose L. McFaline-Figueroa, Pierre Boyeau, F. Alexander Wolf, Nafissa Yakubova, Stephan Günnemann, Cole Trapnell, David Lopez-Paz, and Fabian J. Theis. Predicting cellular responses to complex perturbations in high-throughput screens. *Molecular Systems Biology*, 19(6), 2023. doi: 10.15252/msb.202211517.

Andreas Maurer, Massimiliano Pontil, and Bernardino Romera-Paredes. The benefit of multitask representation learning. *Journal of Machine Learning Research*, 17(81):1–32, 2016.

Djordje Miladinovic, Tobias Höppe, Mathieu Chevalley, Andreas Georgiou, Lachlan Stuart, Arash Mehrjou, Marcus Bantscheff, Bernhard Schölkopf, and Patrick Schwab. In silico biological discovery with large perturbation models. *Nature Computational Science*, 5(11):1029–1040, 2025. doi: 10.1038/s43588-025-00870-1.

Sami Nourreddine, Yesh Doctor, Amir Dailamy, Yi-Hung Lee, Jan N. Hansen, Rebecca Chinn, Antoine Forget, Benjamin Polacco, Monita Muralidharan, Alina Sigaeva, Sushant Sunder, Emily Pan, Jiahao Gao, Jake Y. Chen, Timothy Clark, Jillian Parker, Kirsten Obernier, Christian Metallo, Trey Ideker, Emma Lundberg, Nevan Krogan, and Prashant Mali. A genome-scale CRISPRi perturbation atlas of human induced pluripotent stem cells. *Nature Biotechnology*, 2026. doi: 10.1038/s41587-026-03199-w.

Mineto Ota, Jeffrey P. Spence, Tony Zeng, Emma Dann, Nikhil Milind, Alexander Marson, and Jonathan K. Pritchard. Causal modelling of gene effects from regulators to programs to traits. *Nature*, 650(8101):399–408, 2026. doi: 10.1038/s41586-025-09866-3.

Aniruddh Raghu, Maithra Raghu, Samy Bengio, and Oriol Vinyals. Rapid learning or feature reuse? towards understanding the effectiveness of MAML. In *International Conference on Learning Representations*, 2020.

Aravind Rajeswaran, Chelsea Finn, Sham M Kakade, and Sergey Levine. Meta-learning with implicit gradients. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

Bernardino Romera-Paredes, Hane Aung, Nadia Bianchi-Berthouze, and Massimiliano Pontil. Multilinear multitask learning. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28, pp. 1444–1452. PMLR, 2013.

Jacob Seidman, Georgios Kissas, Paris Perdikaris, and George J. Pappas. NOMAD: Nonlinear manifold decoders for operator learning. In *Advances in Neural Information Processing Systems*, volume 35, pp. 5601–5613. Curran Associates, Inc., 2022. doi: 10.52202/068431-0405.

Marta Soare, Alessandro Lazaric, and Rémi Munos. Best-arm identification in linear bandits. In *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.

Jordan W. Squair, Matthieu Gautier, Claudia Kathe, Mark A. Anderson, Nicholas D. James, Thomas H. Hutson, Rémi Hudelle, Taha Qaiser, Kaya J. E. Matson, Quentin Barraud, Ariel J. Levine, Gioele La Manno, Michael A. Skinnider, and Grégoire Courtine. Confronting false discoveries in single-cell differential expression. *Nature Communications*, 12(1):5692, 2021. doi: 10.1038/s41467-021-25960-2.

Nilesh Tripuraneni, Michael Jordan, and Chi Jin. On the theory of transfer learning: The importance of task diversity. In *Advances in Neural Information Processing Systems*, volume 33, pp. 7852–7862. Curran Associates, Inc., 2020.

Nilesh Tripuraneni, Chi Jin, and Michael Jordan. Provable meta-learning of linear representations. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139, pp. 10434–10443. PMLR, 2021.

Ramon Viñas Torné, Maciej Wiatrak, Zoe Piran, Shuyang Fan, Liangze Jiang, Sarah A. Teichmann, Mor Nitzan, and Maria Brbić. Systema: a framework for evaluating genetic perturbation response prediction beyond systematic variation. *Nature Biotechnology*, 44(6):1050–1059, 2026. doi: 10.1038/s41587-025-02777-8.

Yiping Wang, Yifang Chen, Kevin Jamieson, and Simon Shaolei Du. Improved active multi-task representation learning via lasso. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202, pp. 35548–35578. PMLR, 2023.

Per-Åke Wedin. Perturbation bounds in connection with singular value decomposition. *BIT Numerical Mathematics*, 12(1):99–111, 1972. doi: 10.1007/BF01932678.

Zhiting Wei, Yiheng Wang, Yicheng Gao, Shuguang Wang, Ping Li, Duanmiao Si, Yuli Gao, Siqi Wu, Danlu Li, Kejing Dong, Xingbo Yang, Chen Tang, Shaliu Fu, Xiaohan Chen, Wannian Li, Yuzhou You, Chen Zhang, Aibin Liang, Guohui Chuai, and Qi Liu. Benchmarking algorithms for generalizable single-cell perturbation response prediction. *Nature Methods*, 23(2):451–464, 2026. doi: 10.1038/s41592-025-02980-0.

Yi Yu, Tengyao Wang, and Richard J Samworth. A useful variant of the Davis–Kahan theorem for statisticians. *Biometrika*, 102(2):315–323, 2015.

Doudou Zhou, Tianxi Cai, and Junwei Lu. Multi-source learning via completion of block-wise overlapping noisy matrices. *Journal of Machine Learning Research*, 24(221):1–43, 2023.

Luisa Zintgraf, Kyriacos Shiarli, Vitaly Kurin, Katja Hofmann, and Shimon Whiteson. Fast context adaptation via meta-learning. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pp. 7693–7702. PMLR, 2019.

A Estimator details

The implementation stacks split-A residual vectors as rows of an uncentered matrix. Its leading r_o right singular vectors define L. To form $W = [L, V]$ with requested dimension $m = r_o + r_v$, it orthogonalizes L, the program vector w, and remaining singular directions, in that order. For the selected dimension m, the retained columns contain w in $\mathrm{col}(W)$, allowing the program readout to combine linear and complementary decoder components.

For source condition i, projected split-A coordinates $\widehat{u}_i = L^\top(R_i^A - \mu_i)$ train the features against $V^\top(R_i^A - \mu_i)$ through a temporary linear head with regularized squared error. The head is then discarded. With coordinates and features fixed, weighted ridge regression fits C to complementary split-B measurements:

$$\widehat{C} = Y\Gamma\widehat{\Phi}^\top(\widehat{\Phi}\Gamma\widehat{\Phi}^\top + \lambda_C I)^{-1}, \quad \widehat{\Phi} = [\varphi_\theta(\widehat{u}_i)]_i, \quad Y = [V^\top(R_i^B - \mu_i)]_i. \tag{7}$$

The condition-weight matrix $\Gamma \succ 0$ and ridge penalty $\lambda_C > 0$ are selected on split A before final split-B fitting. App. D states the conditional noise assumptions.

For target support, write $r_S = \mathrm{vec}(R^{\mathrm{obs}}_{c,S} - \mu_{c,S})$, where vec stacks matrix columns and $\|x\|_Q^2 = x^\top Q x$ for a positive-definite weight matrix Q. The full-response objective is

$$J_S(H) = \tfrac{1}{2}\|r_S - \mathrm{vec}\,\Psi(H\widehat{B}_S)\|^2_{\Omega_S^{-1}} + \tfrac{\lambda_H}{2}\|\mathrm{vec}\,H - h_0\|^2_{Q_H}. \tag{8}$$

Target adaptation uses the separable covariance

$$\Omega_S = T_S \otimes \mathrm{diag}(g), \qquad T_S = D_S + \kappa \mathbf{1}\mathbf{1}^\top, \tag{9}$$

where $g \in \mathbb{R}^G$ contains positive output-coordinate variance factors and D_S is a diagonal or dense condition matrix. Scalar and vector inputs give $D_S = dI$ and $D_S = \mathrm{diag}(d)$, respectively. The weights and shared-control coefficient κ satisfy $T_S \succ 0$. Standard scoring uses $g = \mathbf{1}$, $D_S = I$, and $\kappa = 0$.

The prior $h_0 = \mathrm{vec}\,H_0(a_c)$ regresses auxiliary split-A source context matrices against control expression. Full-matrix adaptation uses precision $Q_H = \alpha I$ with $\alpha > 0$ and regularization $\lambda_H > 0$.

For the target initialization, write $h = \mathrm{vec}(H)$, $P_S = I_{|S|} \otimes L^\top$, $z_S = P_S r_S$, $F_S = \widehat{B}_S^\top \otimes I_{r_o}$, and $\Omega_Z = P_S \Omega_S P_S^\top$. With $\Omega_S \succ 0$, $Q_H \succ 0$, and $\lambda_H > 0$, the regularized least-squares solution is

$$h^{\mathrm{init}} = (F_S^\top \Omega_Z^{-1} F_S + \lambda_H Q_H)^{-1}(F_S^\top \Omega_Z^{-1} z_S + \lambda_H Q_H h_0). \tag{10}$$

When refinement completes, the returned iterate minimizes the support objective over valid iterates, including initialization. Thus $J_S(\widehat{H}) \le J_S(H^{\mathrm{init}})$ when the initialization objective is finite. Refinement failure returns initialization, and initialization failure returns the prior. Eq. (25) bounds query error under its stated assumptions.

B Identification and error analysis

These results concern the shared-response model. Vector norms are Euclidean, and matrix subscripts 2 and F denote spectral and Frobenius norms. The superscript † denotes the Moore–Penrose inverse.

Lemma 1 (Projection and decoder stability)**.** *Suppose L, V have orthonormal columns and $L^\top V = 0$. Then $L^\top \Psi(u) = u$ and*

$$\|\Psi(u) - \Psi(u')\|^2 = \|u - u'\|^2 + \|C\varphi_\theta(u) - C\varphi_\theta(u')\|^2. \tag{11}$$

If the nonlinear component $C\varphi_\theta$ is $K_\perp$-Lipschitz on the coordinate domain, then

$$\|u - u'\| \le \|\Psi(u) - \Psi(u')\| \le \sqrt{1 + K_\perp^2}\,\|u - u'\|. \tag{12}$$

Proof. Projection removes the complementary component. Orthogonality eliminates the cross term in $\|L(u-u') + VC(\varphi_\theta(u) - \varphi_\theta(u'))\|^2$, giving Eq. (11) and both bounds. Quadratic features are Lipschitz on bounded domains, generally not on all of $\mathbb{R}^{r_o}$. □

Stacking source matrices H_c vertically gives A and exact projected responses $Z = AB$. Full-column-rank A gives $B = A^\dagger Z$, so Z and B share a row space. With full-row-rank B, their rank is r_b. Singular-subspace perturbation arguments then bound recovery error (Wedin, 1972; Yu et al., 2015).

Proposition 1 (Perturbation embedding recovery under noise)**.** *Suppose the completely observed matrix $\widehat{\mathbb{Z}} = AB_* + E$ has $B_* B_*^\top = I$, full-column-rank A, and $\|E\|_2 \le \eta$. The error E includes measurement, baseline, projection, and structural errors. Let $\widehat{s} = \sigma_{r_b}(\widehat{\mathbb{Z}}) > \eta$, where σ_j is the jth largest singular value. If the rows of $\widehat{B}$ are its leading r_b right singular vectors, truncated singular value decomposition (SVD) gives*

$$\|\widehat{B}^\top \widehat{B} - B_*^\top B_*\|_2 \le \min\{1, 2\eta/(\widehat{s} - \eta)\}. \tag{13}$$

Proof. Let $\widehat{P} = \widehat{B}^\top \widehat{B}$. Because AB_* has rank r_b, truncated SVD gives $\|\widehat{\mathbb{Z}}(I - \widehat{P})\|_2 \le \eta$. It follows that $\|AB_*(I - \widehat{P})\|_2 \le 2\eta$. Left multiplication by $A^\dagger$ bounds $\|B_*(I - \widehat{P})\|_2$ by $2\eta/\sigma_{r_b}(A)$. The nonzero singular values of AB_* equal those of A, since $B_* B_*^\top = I$. Singular-value perturbation yields $\sigma_{r_b}(A) \ge \widehat{s} - \eta$. Equal-rank orthogonal projectors have distance equal to the largest principal-angle sine, and this distance is at most one. Combining these facts proves the claim. □

The deterministic bound permits correlated noise and estimated-projection error covered by E.

Proposition 2 (Identification of the required decoder contrast)**.** *Fix coordinates, features, and noiseless source outputs $Y = C_*\Phi$. Let $t = \varphi_\theta(u_p) - \varphi_\theta(u_q)$ and $v = V^\top w \neq 0$. For unrestricted C_*, the contrast $v^\top C_* t$ is identified if and only if $t \in \mathrm{col}(\Phi)$.*

Proof. If $t = \Phi a$, then $v^\top C_* t = v^\top Y a$. If inclusion fails, take z with $z^\top \Phi = 0$ and $z^\top t \neq 0$. Decoder matrices C_* and $C_* + vz^\top$ yield the same source output, while their task values differ by $\|v\|^2 z^\top t$. For $v = 0$, the nonlinear contribution is zero. An estimator that refits C on target support adds those feature columns to the information set used in this argument. □

Proposition 3 (Polynomial feature coverage)**.** *For an integer $s \ge 1$, let $\Phi^{(s)}$ contain homogeneous degree-s features of the projected source coordinates in $\mathbb{R}^{r_o}$. Let $U_c = \mathrm{range}(H_c)$ and $\rho_c = \mathrm{rank}(H_c)$, and let $\mathrm{Sym}^s(U_c)$ be its degree-s symmetric tensor space. Then*

$$\mathrm{rank}\,\Phi^{(s)} \le \dim\left(\sum_c \mathrm{Sym}^s(U_c)\right) \le \min\left\{\binom{r_o + s - 1}{s}, \sum_c \binom{\rho_c + s - 1}{s}\right\}. \tag{14}$$

Proof. Every coordinate from source c lies in U_c. Its homogeneous polynomial lift lies in the symmetric tensor space $\mathrm{Sym}^s(U_c)$, of dimension $\binom{\rho_c + s - 1}{s}$. Observed columns span a subspace of the sum of these spaces. The sum's dimension is bounded by both the summed dimensions and the ambient feature dimension; attaining this bound also requires adequate sampling. □

Axis-restricted sources never observe the interaction $u_1 u_2$. A full-rank source context matrix permits mixed coordinates, but their observation still requires suitable perturbations. Neural feature coverage refers to the fitted feature matrix with parameters fixed.

Proposition 4 (Target support identification)**.** *Fix a correct baseline, a decoder satisfying $L^\top \Psi(u) = u$, and a perturbation embedding library. Noiseless full support responses identify all full query responses for every unrestricted context matrix H if and only if $\mathrm{col}(B_D) \subseteq \mathrm{col}(B_S)$.*

Proof. Under inclusion, $B_D = B_S T$ for $T = B_S^\dagger B_D$. Projection recovers HB_S from support, hence $HB_D = (HB_S)T$; decoding gives the query responses. If inclusion fails, a vector z satisfies $z^\top B_S = 0$ and $z^\top B_D \neq 0$. For any nonzero $a \in \mathbb{R}^{r_o}$, context matrices H and $H + az^\top$ give

identical support coordinates and different query coordinates. The projection identity makes the decoder injective, so the query responses differ. A fixed scalar projection can annihilate this difference; necessity here concerns full responses. □

Proposition 5 (Compressed source-response equivalence)**.** *If exact stacked projected source coordinates satisfy $Z = AB$ and A has full column rank, then $HB = HA^{\dagger}Z$. Let $\widetilde{Z} = QB$ be any full-row-rank compression of Z to r_b rows, with invertible Q. Then the map $K = HQ^{-1}$ satisfies $K\widetilde{Z} = HB$.*

Proof. The identities follow from $A^{\dagger}A = I$ and substitution. Both parameterizations have $r_o r_b$ free parameters and are bijective. Substitute $H = KQ$ into the full penalty $\|\operatorname{vec}(H) - h_0\|^2_{Q_H}$, including its prior. An unchanged Euclidean penalty on K generally changes the estimator despite identical expressivity. □

C Two indistinguishable constructions

Both constructions use a zero baseline and the decoder's final coordinate as the program readout. First, set $H_1 = \operatorname{diag}(1, 0)$, $H_2 = \operatorname{diag}(0, 1)$, and $H_t = I$. The source library is $e_1, e_2, -e_1, -e_2, e_1 + e_2, e_1 - e_2$, where e_j are standard basis vectors. Target support is e_1, e_2, and queries $b_p = e_1 + e_2$, $b_q = e_1 - e_2$ are compared as p minus q. With $w = e_3$, consider

$$\Psi_{\pm}(u) = (u_1, u_2, \pm a u_1 u_2)^{\top}, \qquad a > 0. \tag{15}$$

Source and target-support coordinates lie on the axes, where both decoders agree. Their query contrasts are $2a$ and $-2a$ (Fig. 1b), so any estimator returning z has average squared error

$$\tfrac{1}{2}\{(z - 2a)^2 + (z + 2a)^2\} = z^2 + 4a^2 \geq 4a^2. \tag{16}$$

The bound survives averaging over estimator randomness and lower-bounds worst-world risk. A source with context matrix I reveals the interaction on mixed perturbations, whereas the original two sources cannot.

For the second construction, use three embedding coordinates, source context matrix I, program readout $w = e_4$, and

$$H_t = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}, \qquad \Psi_{\pm}(u) = (u_1, u_2, u_3, \pm a u_1 u_2)^{\top}. \tag{17}$$

The measured library is $e_1, e_2, e_3 - e_2, e_3, e_1 - e_3$. The first three vectors form target support and span the embedding space; the last two are queries. The interaction vanishes on both the measured source library and target support. Query coordinates $(1, 1, 1)^{\top}$ and $(0, -1, -1)^{\top}$ give contrasts a and $-a$, hence worst-world squared error at least a^2. Adding $e_1 + e_2$ in the existing source reveals the sign. Repeating the old panel in another source with context matrix I leaves the ambiguity, whereas a different context matrix can resolve it.

The numerical checks recover the interaction coefficient directly, separating the information supplied by measurements from predictor retraining.

D Conditional decoder risk and complete prediction error

Experimental design distinguishes statistical noise models from worst-case fixed responses (Dereziński et al., 2019). Here the feature design is fixed, and risk concerns final-output noise. Right-multiply source output, feature, and error matrices by $\Gamma^{1/2}$ after feature evaluation, retaining their symbols. Coordinates, query vectors, and C_* are unchanged, and noise assumptions apply to the weighted errors. Set $M = \widehat{\Phi}\widehat{\Phi}^{\top} + \lambda_C I$ and $Y = C_*\widehat{\Phi} + \Xi + E$, where Ξ contains feature-input and decoder approximation errors and E is output noise. For fixed t and $v = V^{\top}w$, define $b(t) = v^{\top}(-\lambda_C C_* + \Xi\widehat{\Phi}^{\top})M^{-1}t$. Eq. (7) gives

$$\widehat{C} - C_* = (-\lambda_C C_* + \Xi\widehat{\Phi}^{\top} + E\widehat{\Phi}^{\top})M^{-1}. \tag{18}$$

Substitute $Y = C_*\widehat{\Phi}+\Xi+E$ and use $\widehat{\Phi}\widehat{\Phi}^\top = M-\lambda_C I$. Conditional on split A, the frozen baseline, and underlying means, $b(t)$ is deterministic. If $v^\top E$ has conditional mean zero and conditional covariance Σ_v, then

$$\mathbb{E}\big[(v^\top(\widehat{C}-C_*)t)^2\big] = b(t)^2 + t^\top M^{-1}\widehat{\Phi}\Sigma_v\widehat{\Phi}^\top M^{-1}t. \tag{19}$$

For independent zero-mean entries of $v^\top E$ with common variance τ^2, this becomes

$$\mathbb{E}\big[(v^\top(\widehat{C}-C_*)t)^2\big] = b(t)^2 + \tau^2 t^\top M^{-1}\widehat{\Phi}\widehat{\Phi}^\top M^{-1}t. \tag{20}$$

An unobserved feature direction can have zero variance while retaining squared bias $b(t)^2$. Shared controls across splits or feature fitting on final-output noise can violate conditional zero-mean noise. Here $\Xi = A_{\mathrm{dec}} + C_*(\Phi - \widehat{\Phi})$, with source decoder approximation error A_{dec}.

For prediction, write $\widehat{\Psi}(u) = Lu + V\widehat{C}\varphi_\theta(u)$ for the frozen decoder in J_S. Fix a reference context matrix H^{ref} and write

$$R_p^* = \mu_p + \widehat{\Psi}(H^{\mathrm{ref}}\widehat{b}_p) + \delta_p. \tag{21}$$

The residual δ_p collects baseline, projection, feature, embedding, and decoder errors. Let $B_* = [b_p^*]_p$ be a reference library satisfying $B_*B_*^\top = I$, and let H_* be a reference context matrix. Align coordinates by

$$O \in \operatorname*{arg\,min}_{O^\top O = I} \|\widehat{B} - OB_*\|_F, \qquad b_p^{\mathrm{ref}} = Ob_p^*, \qquad H^{\mathrm{ref}} = H_*O^\top. \tag{22}$$

Then $u_p^* = H_*b_p^* = H^{\mathrm{ref}}b_p^{\mathrm{ref}}$ and $e_{b,p} = \|\widehat{b}_p - b_p^{\mathrm{ref}}\|$. Write

$$R_p^* = \mu_p^* + Lu_p^* + VC_*\varphi_\theta(u_p^*) + a_p, \qquad e_{\mu,p} = \|\mu_p^* - \mu_p\|, \qquad e_{\mathrm{struct},p} = \|a_p\|. \tag{23}$$

Here μ_p^* is the reference baseline, and a_p collects approximation error under the learned projection and features. Alignment preserves H_*B_*. Column errors $e_{b,p}$ are reference quantities, distinct from the row-space bound. For the learned decoder, the projection lemma permits $K_\Psi = \sqrt{1+K_\perp^2}$ when $\widehat{C}\varphi_\theta$ is $K_\perp$-Lipschitz on the coordinate domain. Adding and subtracting $\widehat{\Psi}(u_p^*)$ gives

$$\|\delta_p\| \le e_{\mu,p} + e_{\mathrm{struct},p} + K_\Psi\|H^{\mathrm{ref}}\|_2 e_{b,p} + \|\widehat{C} - C_*\|_2\|\varphi_\theta(u_p^*)\|. \tag{24}$$

Proposition 6 (Error after context adaptation)**.** *Suppose $\widehat{B}_D = \widehat{B}_S T$, $\Omega_S \succ 0$, and $J_S(\widehat{H}) \le J_S(H^{\mathrm{init}})$. Let $R_S^{\mathrm{obs}} = R_S^* + \varepsilon_S$ and $\delta_S = [\delta_p]_{p\in S}$, with δ_D defined analogously. If $\widehat{\Psi}$ is K_Ψ-Lipschitz on all coordinates involved, then*

$$\|\widehat{R}_D - R_D^*\|_F \le K_\Psi\|T\|_2\left[\sqrt{2\lambda_{\max}(\Omega_S)J_S(H^{\mathrm{init}})} + \|\varepsilon_S + \delta_S\|_F\right] + \|\delta_D\|_F. \tag{25}$$

Proof. The nonnegative prior penalty and return condition bound the support residual by $\sqrt{2\lambda_{\max}(\Omega_S)J_S(H^{\mathrm{init}})}$. Adding $\varepsilon_S + \delta_S$ bounds decoded-support differences between $\widehat{H}$ and H^{ref}. Projection bounds $\|(\widehat{H} - H^{\mathrm{ref}})\widehat{B}_S\|_F$ by the same quantity. Multiplication by T, the decoder Lipschitz bound, and δ_D complete the result. □

For incomplete target coverage, choose $T = \widehat{B}_S^\dagger\widehat{B}_D$ and add $K_\Psi\|\widehat{H} - H^{\mathrm{ref}}\|_2\|\widehat{B}_D - \widehat{B}_S T\|_F$. Applying the contrast operator directly to response error preserves shared-noise cancellation.

E Numerical estimation settings

Generation seeds vary the numerical design, while noise draws repeat measurements within a fixed design (Fig. S1).

Compression checks use $r_o = r_b = 2$, six seeds, four coordinate condition numbers, three ridge values, and three feature families. Across 216 settings, the maximum relative prediction discrepancy was 5.84×10^{-14}, and the maximum absolute prior-penalty discrepancy was 1.08×10^{-13}. The penalty uses $\mathrm{vec}(H) = (Q^\top \otimes I)\,\mathrm{vec}(K)$ for $H = KQ$.

Embedding checks cross five source strengths, four noise scales, and six seeds. Using the known injected error norm, 42 of 120 settings meet the spectral-bound assumptions, and 18 give a bound below one. These designs use a fixed, correct projection.

Support checks cross three geometries, counts $0, 1, 2, 4, 8$, noise standard deviations $0.03, 0.15, 0.5$, ridge penalties $0.01, 0.1, 1$, and six seeds. All 810 settings use $H_0 = I$. Fig. S1f summarizes all geometry/noise/ridge combinations at eight versus two supports. Ill-conditioned coverage amplifies noise, whereas missing directions leave an uncovered residual. With ill-conditioned support and ridge 0.01, eight supports increased mean MSE at noise 0.15 and 0.5.

Fixed-design checks use ridge 0.2, covered, blind, and noisy feature regimes, noise standard deviations $0.08, 0.25, 0.70$, six generation seeds, and 4,096 noise draws per setting. For 32 source conditions, $\Sigma_{\text{cond}} = \sigma^2(0.65I + 0.35\mathbf{1}\mathbf{1}^\top)$ and $\Sigma_v = \|v\|^2\Sigma_{\text{cond}}$, with all-ones vector $\mathbf{1}$. Correlation 0.35 requires Eq. (19). Table S1 reports six-design means.

Table S1: Conditional risk on the linear contrast scale. Means over six designs with 4,096 noise draws each. Noise is the injected standard deviation. Gap is $100(\text{Empirical}/\text{Theory} - 1)$.

Regime	Noise	Bias^2	Variance	Theory	Empirical	Gap (%)
Blind	0.08	5.2446	0.0000	5.2446	5.2446	0.00
Blind	0.25	5.2446	0.0000	5.2446	5.2446	0.00
Blind	0.70	5.2446	0.0000	5.2446	5.2446	0.00
Covered	0.08	0.0008	0.0016	0.0023	0.0024	+0.97
Covered	0.25	0.0008	0.0154	0.0161	0.0163	+1.19
Covered	0.70	0.0008	0.1205	0.1212	0.1227	+1.18
Noisy features	0.08	0.0567	0.0014	0.0582	0.0581	−0.14
Noisy features	0.25	0.0567	0.0138	0.0706	0.0702	−0.49
Noisy features	0.70	0.0567	0.1084	0.1652	0.1636	−0.98

The paired constructions use amplitudes $0.3, 0.7, 1.0$ and six generation seeds, averaging both decoder signs within each seed. They compare no addition, increased depth, more source perturbations, and another context. Fig. 4a,b displays seed means and one SE.

F Learning under controlled structural mismatch

The structural study crosses 14 base worlds and four additional branches, noise standard deviations $0.06, 0.2$, four generation seeds, and up to eight predictors (Table S2). The branches isolate drift, baseline or input behavior, and complete baseline identities. Models share visible source responses and target support within each world. In Fig. S2, Mean, Baseline, Gain, and Ridge denote source mean, baseline, support-gain, and compressed ridge predictors. Linear, Quad., and Neural are LHB decoder families. "No proj." denotes the unconstrained decoder with its baseline retained.

Source fitting uses 45 epochs and target adaptation 15 steps where applicable. Ordinary worlds have 10 sources, 24 perturbations, two supports, four queries, and 12 response coordinates. One source fits the baseline and nine fit the structure. The axis-source world has six source perturbations and two queries. The three-dimensional construction has five source perturbations, three supports, and two queries. Neural models use 12 features, 24 hidden units, and fitting seed zero.

Table S2: Base worlds in the structural study.

World	Setting	World	Setting
G01	Linear response	G08	Missing support direction
G02	Quadratic decoder	G09	Ill-conditioned support
G03	Smooth neural decoder	G10	Projection folding
G04	Missing perturbation-embedding direction	G11	Low-dimensional baseline or gain
G05	Axis-restricted sources	G12	Incomplete source library
G06	Sparse source perturbations	G13	Selection and control dependence
G07	Coupled source restrictions	G14	Nonlinear structural equations

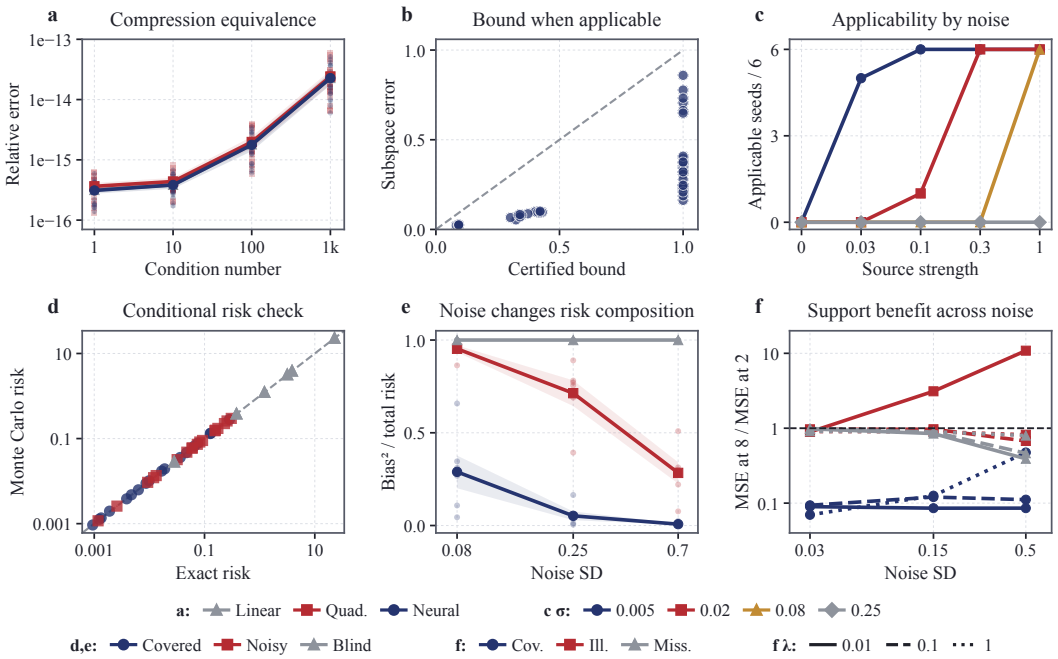

Figure S1: Numerical estimation checks. (a) Compression equivalence. (b,c) Subspace-error bounds and applicable designs. (d) Exact and empirical conditional risk, with Monte Carlo SE over 4,096 noise draws. (e) Squared-bias fractions. (f) Mean MSE at eight supports divided by that at two, for all geometry/noise/ridge combinations. Curves in a average ridge values within seeds; a,e show six-design means ± one SE.

The study reports 1,128 prediction evaluations. Another 24 LHB fits are unavailable in the incomplete-library world because the baseline sources lack responses for the requested perturbations. Fig. S2 compares methods within worlds with different response scales. Beyond folding and drift, ridge had lower contrast MSE in the shared neural-decoder world at noise 0.2 (0.04810 versus 0.09236). The world with coupled source restrictions (G07) favored neural LHB (0.08600 versus 0.30152). Model preference therefore depends on structural fit and source information.

Oracle replacements use known true components to diagnose estimation error across six worlds, four seeds, and eight component-replacement configurations. Of 192 combinations, 156 are evaluable and 36 structurally inapplicable. Folded responses need not admit a shared decoder, and the nonlinear structural-equation world has no true LHB parameterization. In the axis-source world, supplying true embeddings changed latent contrast MSE from 1.43205 to 1.46046. True geometry alone gave 1.45448; adding the true decoder reduced it to 0.05204, identifying decoder estimation as the main limitation. Joint replacements include interactions among component errors.

Query contrast error increased from initialization in 49 of 136 neural fits despite a non-increasing support objective, showing that improved support fit can worsen held-out prediction. The oracle panel uses latent contrast MSE, whereas trained-predictor matrices use errors against measured responses.

G Biological inputs, preprocessing, and evaluation units

The primary experiment uses the targeted human iPSC CRISPRi resource of Feng et al. (2026), Figshare dataset 27989294, with metadata file 51051809 and RNA file 51072902. Metadata contain 635,022 assigned cells: 626,781 targeting and 8,241 non-targeting controls, across ten donors, 19 lines, and 444 targets. Measurements are at day 3. Guide strength affects knockdown magnitude (Jost et al., 2020), so effects concern the specified guides and endpoint.

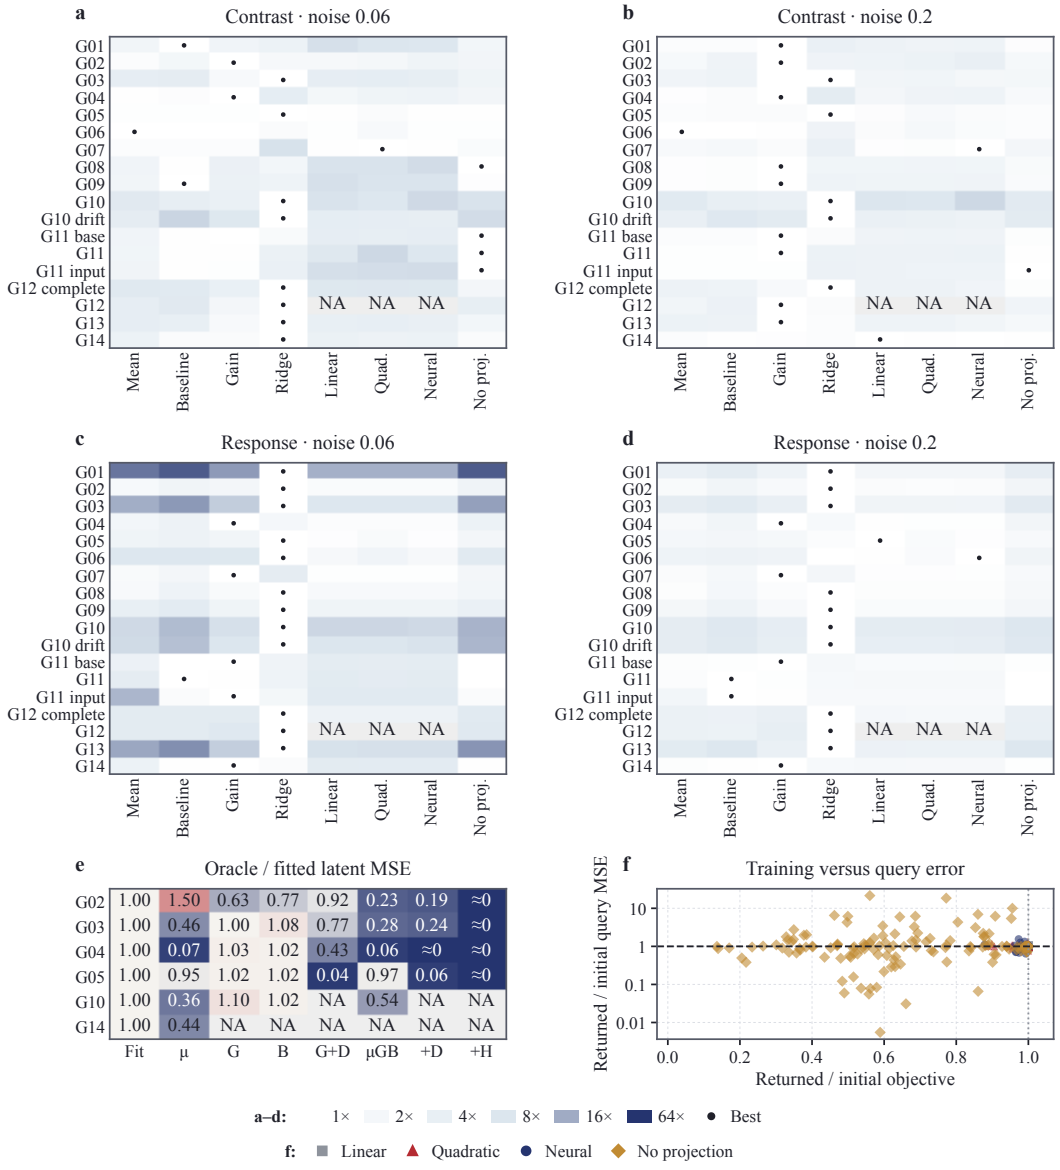

Figure S2: Controlled worlds and estimation limits. (a–d) Four-seed mean MSE relative to the best method within each world, marked by dots; colors use a logarithmic scale. (e) Latent contrast MSE relative to fitted neural LHB. G and D denote geometry and decoder; the final columns cumulatively add D and H to μGB. Values below 10^{-12} are marked ≈ 0; NA denotes structural inapplicability. (f) Support-objective and query-error changes after adaptation.

The first 120 distinct non-control targets in metadata order form 24 queries, 64 supports, and 32 auxiliary targets. Each uses the first guide represented in all 19 lines, excluding alternative query-target guides from support. Before branch-availability filtering, this pool contains 76,332 targeting cells and 8,241 controls. The 2,280 line–guide combinations have median 33 targeting cells (range 1–122), with 74 below four and 24 containing one.

Processing retains the source quality criteria: at least 2,000 detected genes, mitochondrial RNA fraction at most 10%, inlet-specific minimum unique molecular identifier (UMI) counts of 1,926–39,260, valid single-guide assignments, and removal of source-flagged low-quality cells. No additional highly variable gene selection is applied. Counts are normalized over the complete gene axis to 10,000 per cell and transformed by natural $\log(1 + x)$ before averaging. Responses subtract matched non-targeting-control means. Technical inlets are pooled within line and library pool. Cell splitting orders cells by a deterministic SHA256 hash of `0:cell_id`. Ordinary perturbations are divided equally between A/B, requiring two cells per branch. Controls use six disjoint pools for A, B, baseline, support, and two evaluation measurements. Branches with insufficient cells are excluded. Two of 66 control-containing line–library groups have fewer than 12 controls.

Each outer fold holds out a donor and all its lines. For LHB, the remaining donors are shuffled with seed zero: three fit μ and six fit shared structure. Source information determines transforms, geometry, features, and hyperparameters; target support adapts only H, and queries remain withheld. The 24 cyclic query contrasts have equal weight $1/24$. Metadata adjacency fixes evaluation pairs independently of regulatory annotation.

The native response loading is $P_{\mathrm{RNA}} = [w_{\mathrm{RNA}}, P_{\perp}]$, where w_{RNA} assigns POU5F1, SOX2, and NANOG weights $1/\sqrt{3}$ and other genes zero. The 63 columns of $P_{\perp}$ are source-derived loading vectors orthogonal to w_{RNA}. Thus $R = P_{\mathrm{RNA}}^{\top} R^{\mathrm{RNA}}$ preserves the program effect as its first coordinate, with $w = e_1$ in the 64-coordinate representation and no further gene-wise rescaling. Reported contrasts retain contributions from directly targeted genes.

None of the 64 supports is a program gene. Fig. S3 summarizes input coverage and retrospective functional annotation. The three technical panels contain 22, 21, and 21 targets and mix these functions, testing sensitivity to support-panel composition.

H PREDICTION COMPARISONS AND INFORMATION CONTROLS

Squared errors average over fits and panels within lines, lines within donors, and equally over donors. Donors define the biological replication level (Squair et al., 2021). Descriptive 95% donor-bootstrap intervals use 20,000 resamples and seed zero. The main comparison uses eight support panels and three computational repetitions, sharing deterministic source fits across repetitions. Its ten-support results come from the support-budget series, displayed linearly through two supports and logarithmically thereafter. Table S3 reports all 22 methods.

One or two supports gave higher contrast MSE than Baseline in this series. At ten supports, the paired LHB-minus-HyperMap difference was -0.00121, with a 95% donor-bootstrap interval of $[-0.00308, 0.00099]$. The interval for PerturbMap also included zero (Fig. 3d).

HyperMap, PerturbMap, and the large perturbation model (LPM) are adapted to the same source/support/query access and scoring interface. Their architectures and parameter counts differ from the originals and from each other. The HyperMap adapter uses a pointwise context gate and updates its response network from target support. PerturbMap uses source-response descriptors, affine source-to-target maps, and leave-one-out support validation for route reliability with at least three supports. LPM adapts a context vector with its shared multiplicative decoder fixed.

Main LHB uses $r_o = r_b = 4$ and $m = r_o + r_v = 16$: 16 adapted parameters and a 16-dimensional decoder adjustment within the 64-coordinate representation.

The capacity branch uses four support panels and three fits per line (228 scores per model); the depth branch uses two matched panels (114 scores per condition). Each depth reference uses the same panel and fit identifiers as its reduced-depth counterpart. Table S4 reports both branches. Fig. S4 extends capacity, depth, and adaptation controls across support budgets. Negative transfer compares query error with initialization within each run before donor aggregation, even when the support objective decreases.

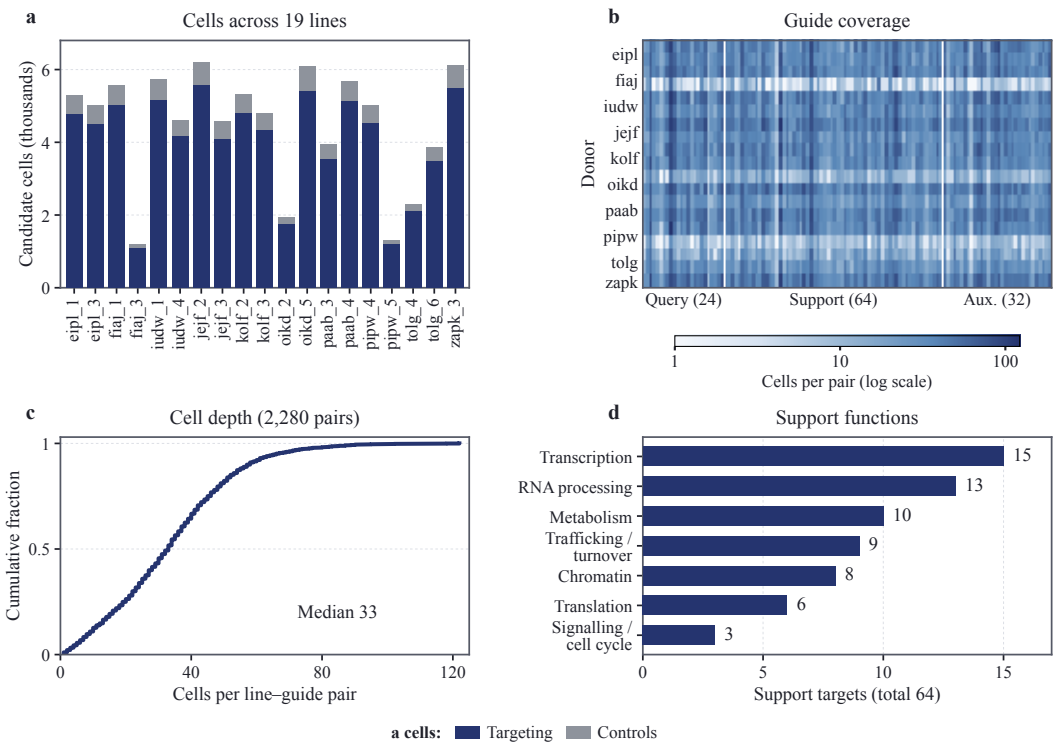

Figure S3: Input coverage. (a) Selected targeting and control cells across 19 lines. (b) Coverage of 24 query, 64 support, and 32 auxiliary guides. (c) Empirical cumulative distribution of targeting-cell counts across all 2,280 line–guide combinations. (d) Retrospective functional annotation of the support targets. Input coverage precedes disjoint-branch availability filtering.

Table S3: All main prediction methods at ten supports. Mean ± donor standard deviation, MSE $\times 10^3$, over ten donors. Response uses 64 coordinates. Gain is the percentage reduction in contrast MSE from Baseline.

Method	Response	Program	Contrast	Gain %
Baseline	4.15 ± 2.58	28.79 ± 40.12	61.33 ± 95.64	0.0
Source mean	6.19 ± 4.27	43.09 ± 64.93	94.59 ± 157.06	−54.2
Zero	13.30 ± 6.99	78.68 ± 101.54	170.43 ± 231.29	−177.9
Nearest baseline	5.21 ± 3.43	35.78 ± 50.82	77.70 ± 121.77	−26.7
Support shift	4.15 ± 2.58	28.62 ± 40.30	61.33 ± 95.64	0.0
Support gain	3.48 ± 1.61	22.00 ± 22.97	45.11 ± 57.21	+26.4
Covariance shrinkage	3.08 ± 1.38	20.22 ± 21.12	40.78 ± 48.93	+33.5
Compressed ridge	3.18 ± 1.68	21.36 ± 24.36	43.84 ± 58.25	+28.5
Generic context	3.09 ± 1.73	21.62 ± 26.82	44.62 ± 62.24	+27.2
Tucker	3.30 ± 1.79	22.85 ± 27.74	47.53 ± 65.90	+22.5
Compressed nonlinear	2.78 ± 1.20	18.31 ± 18.68	37.54 ± 44.61	+38.8
HyperMap adapted	2.75 ± 1.28	18.47 ± 18.24	36.78 ± 42.56	+40.0
PerturbMap adapted	2.82 ± 1.31	19.16 ± 21.28	38.73 ± 50.73	+36.8
LPM adapted	2.94 ± 1.38	19.48 ± 18.65	39.70 ± 44.62	+35.3
LHB scalar	3.82 ± 2.32	25.76 ± 32.85	53.92 ± 77.29	+12.1
LHB diagonal	2.90 ± 1.40	19.31 ± 21.29	40.91 ± 51.99	+33.3
LHB linear	3.14 ± 1.67	21.81 ± 25.76	45.28 ± 60.78	+26.2
LHB quadratic	2.65 ± 1.27	18.30 ± 19.87	37.27 ± 46.55	+39.2
No projection constraint	2.79 ± 1.37	19.46 ± 22.30	39.15 ± 51.68	+36.2
Support/query bypass	2.84 ± 1.49	19.32 ± 21.99	40.51 ± 52.81	+33.9
Joint training	2.71 ± 1.32	18.45 ± 20.96	38.19 ± 49.91	+37.7
LHB neural	2.63 ± 1.23	17.70 ± 18.96	35.57 ± 44.71	+42.0

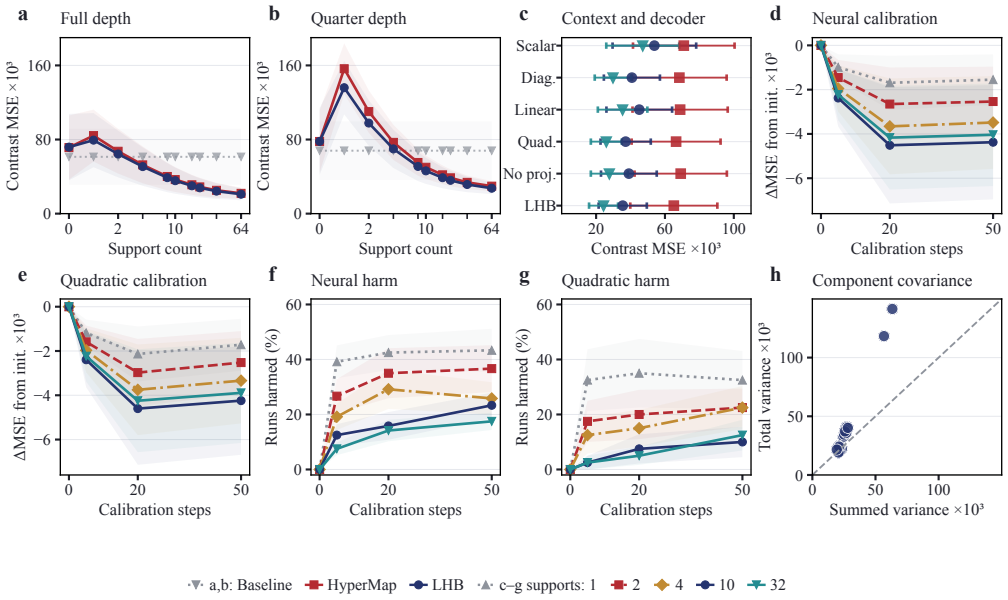

Figure S4: Model and measurement controls. (a,b) All support budgets at full and quarter perturbation depth. (c) Six context/decoder variants at 2, 10 and 32 supports. (d,e) Paired changes from zero-step initialization for neural and quadratic decoders. (f,g) Runs with query error above initialization, averaged within line and donor. (h) Total versus summed component variances across 19 lines; the diagonal denotes zero net covariance. Panels a–g show ten-donor means ± one SE, using paired differences in d,e. MSE and variance use $\times 10^3$.

Stage exchanges substitute geometry and/or embeddings between independently fitted LHB and compressed predictors, then refit downstream modules (Fig. S5a). Seven configurations use six sources, support budgets 4, 10 and 32, and 20 adaptation steps, except initialization-only controls. Refitting output weights with features fixed reproduces the reference. At ten supports and A:B allocation 1:1, staged versus joint contrast MSE is 0.03943 versus 0.03813 with three sources, and 0.03617 versus 0.03748 with eight.

The program prediction decomposes in the original score units as

$$\widehat{F}_{cp} = w^\top \mu(a_c, p) + w^\top L H_c b_p + w^\top V C \varphi_\theta(H_c b_p). \tag{26}$$

Across 24 queries in 19 lines, the fitted nonlinear program component has root mean square 0.03864. Mean within-line query variance is 0.04077, versus 0.02777 for summed component variances, so covariance affects component attribution. Decoder ablations capture the combined effects of direct nonlinear program output and estimation of H through full support responses.

Separate stage, source, and support controls use $r_b = 2$. They vary technical panels, source count, auxiliary coverage, depth, and support geometry. Matched designs fix either the perturbation-information or feature-design Gram matrix while varying the other component; Fig. S6a–c extends this comparison across budgets. Other controls vary measurement counts at condition, control, read, culture, line, and donor levels as sensitivity analyses of existing measurements.

For incomplete source matrices, masked alternating least squares fits observed entries, initialized from perturbations shared by all sources with rank at least r_b. The complete-matrix bound in Eq. (13) does not cover this estimator. The masking study yields predictions for 21 configurations, with nine disconnected designs remaining unevaluable (Fig. S7).

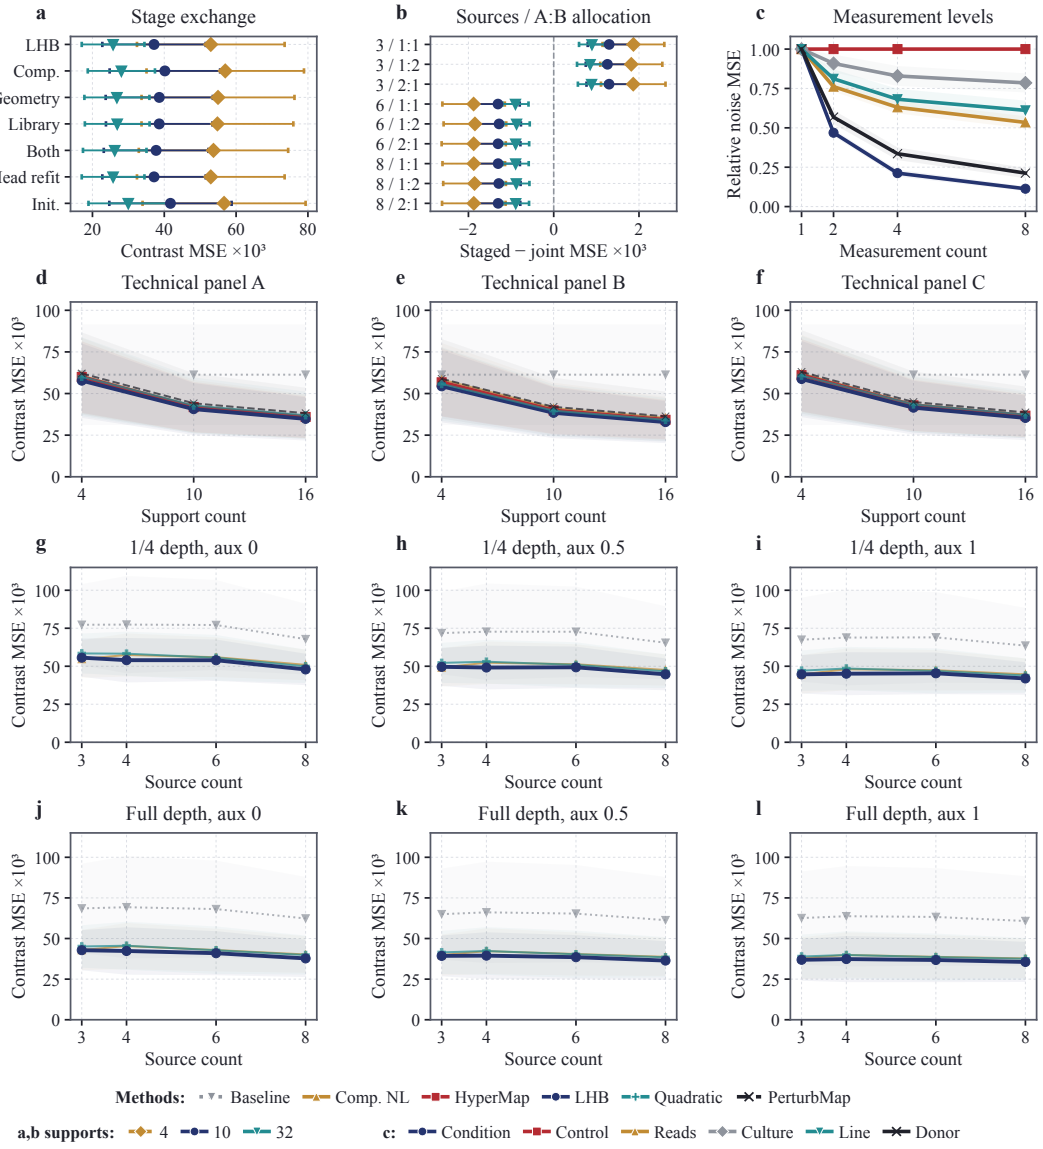

Figure S5: Pipeline and source-budget controls. (a) Stage exchanges. (b) Paired staged-minus-joint differences. (c) Relation-noise MSE relative to one measurement, with paired delta-method SE. (d–f) Mixed-function technical panels A/B/C. (g–i) Source count and auxiliary fraction at quarter depth. (j–l) The same design at full depth. Panels a,b,d–l show ten-donor means ± one SE. Contrast MSE and its differences use $\times 10^3$.

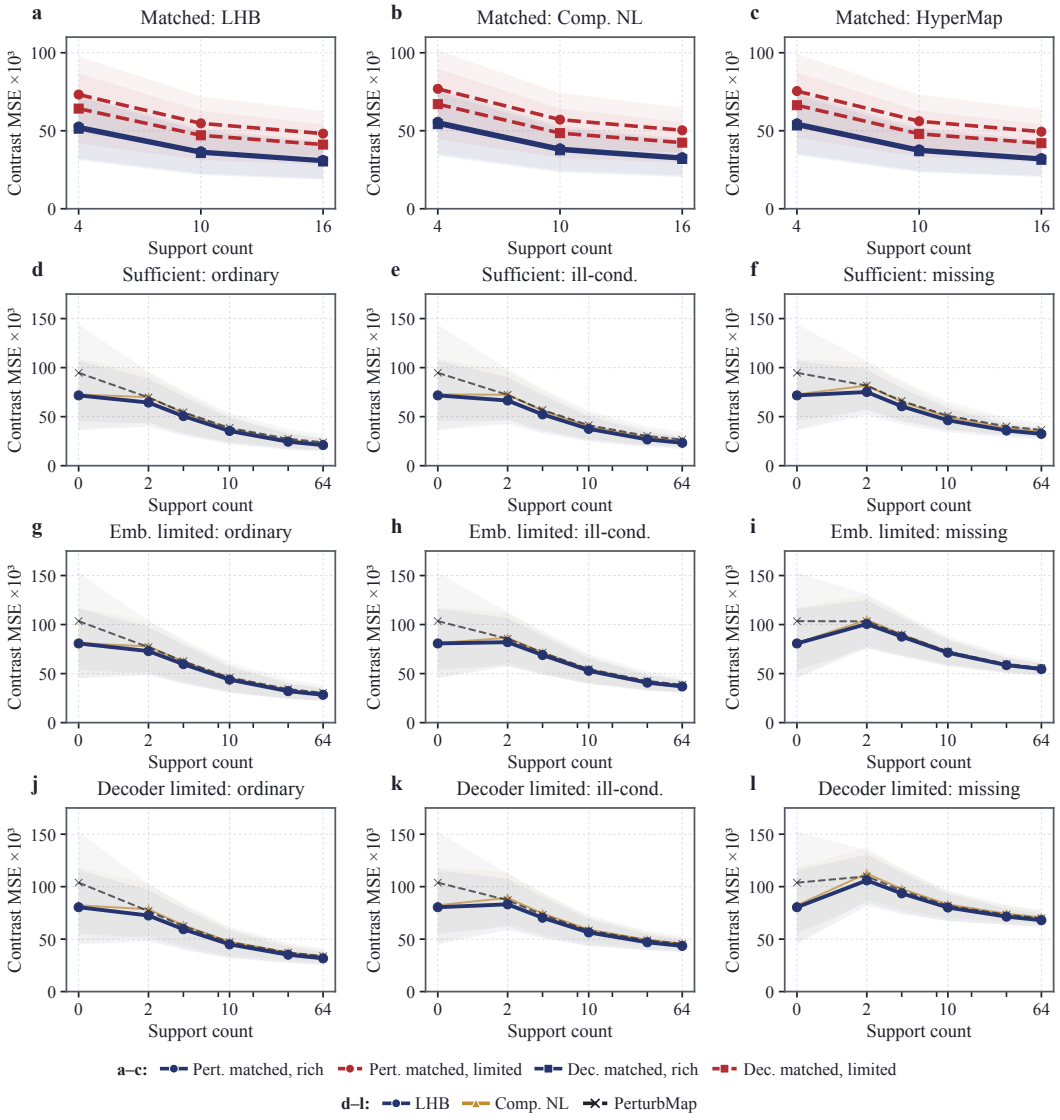

Figure S6: Source information and support geometry. (a–c) Rich versus limited information, matching either perturbation-information matrices (circles) or feature-design matrices (squares). Rich/limited refers to the unmatched component. (d–f) Three support geometries with sufficient source information. (g–i) Limited perturbation information. (j–l) Limited decoder information. Curves show ten-donor means ± one SE across six support budgets. Contrast MSE uses $\times 10^3$, with shared scales within a–c and d–l.

Table S4: Model and measurement controls at ten supports. Means ± standard error across ten donors. MSE and Diff. use $\times 10^3$. Diff. is contrast MSE minus LHB neural (A) or Full depth (B).

Condition	Response	Program	Contrast	Diff.
A. Model structure, four support panels				
Scalar context	3.82 ± 0.73	25.81 ± 10.36	53.85 ± 24.29	$+18.38$
Diagonal context	2.90 ± 0.44	19.37 ± 6.70	40.79 ± 16.28	$+5.33$
Linear decoder	3.14 ± 0.53	21.84 ± 8.12	45.07 ± 19.08	$+9.61$
Quadratic decoder	2.65 ± 0.40	18.35 ± 6.27	37.16 ± 14.60	$+1.70$
No projection constraint	2.79 ± 0.43	19.52 ± 7.03	39.03 ± 16.19	$+3.57$
LHB neural	2.63 ± 0.39	17.76 ± 5.97	35.46 ± 14.00	0.00
B. Measurement depth, two matched support panels				
Full depth	2.63 ± 0.39	17.99 ± 6.03	35.82 ± 14.10	0.00
Control depth $1/4$	2.69 ± 0.39	21.77 ± 5.83	35.82 ± 14.10	0.00
Perturbation depth $1/8$	2.72 ± 0.39	23.25 ± 6.22	45.72 ± 14.95	$+9.89$

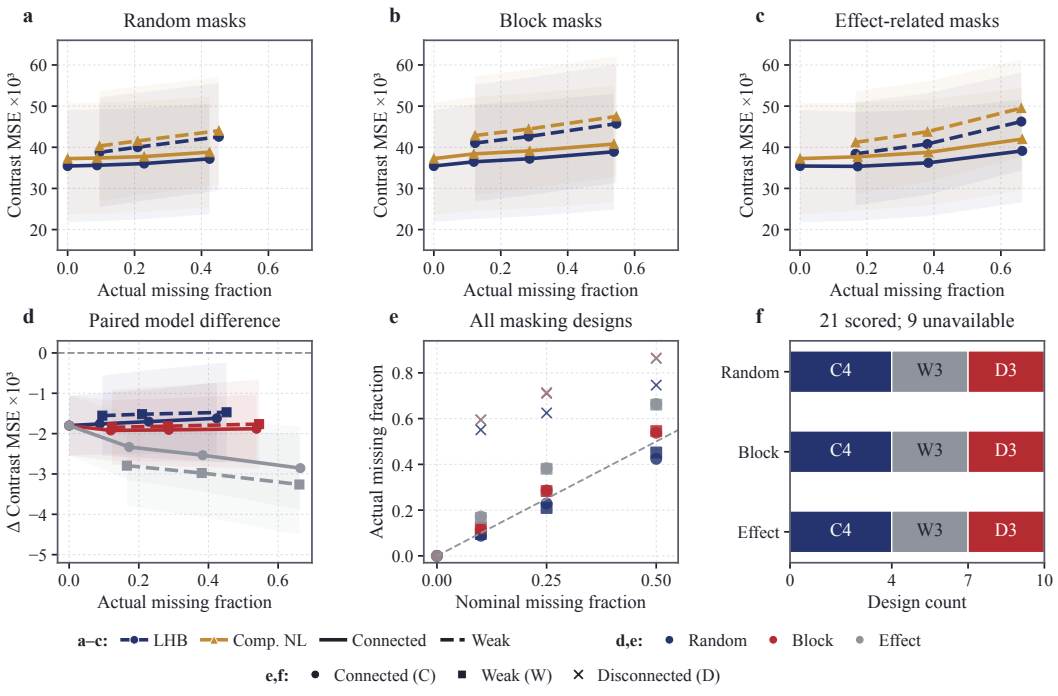

Figure S7: Missing source measurements. (a–c) Prediction under connected and weak overlap. (d) Paired LHB-minus-compressed-nonlinear errors; circles and squares denote connected and weak overlap. (e) Nominal versus actual missing fractions; crosses denote disconnected designs. (f) Design counts by missingness type. Nine disconnected configurations are unavailable. Panels a–d show ten-donor means ± one SE, paired in d.

I CROSS-LINE, EXTERNAL, AND ADDITIONAL EVALUATIONS

Cross-line transfer fits H from line A support and predicts line B in nine eligible donors, using the target-line baseline. Program weights, query identities, and cyclic contrasts remain fixed. Errors average over panels and fits before donors. The fitted matrix can reflect genetic and culture-associated differences. Table S5 gives cross-line errors for the four contrasts in Fig. 3a, with gains over Baseline for three contrasts and higher error for BTAF1–HNRNPH1.

Resting CD4 evaluation uses four donors from GSE314342/SRP643211, with a separate response representation and scalar endpoint. At 20 supports, LHB contrast MSE is 0.10918, versus 0.15485 for compressed ridge and 0.14682 for adapted PerturbMap.

The external iPSC development evaluation uses the single-line KOLF2.1J atlas of Nourreddine et al. (2026) (PRJNA1173491; Figshare 27261219), with a separate 64-coordinate representation of its public z-score summary. LHB contrast MSE falls from 0.06972 at zero supports to 0.03341 at ten supports. The control-expression Baseline has MSE 0.06629.

Table S5 and Fig. S8 also include 48 additional units and changed mixed-function supports on those same units. The changed-support series tests reuse across support-panel compositions.

Table S5: Additional evaluations at ten supports. Contrast MSE and Diff. use $\times 10^3$. Diff. is LHB minus comparator (Comp.). Gain is reduction from Baseline (%). Comp. is HyperMap for line transfers and additional units, PerturbMap for CD4, and compressed ridge for the external line. Here n counts donors, one external line, or the same 48 additional units. Response scales differ.

Series	n	Baseline	LHB	Comp.	Diff.	Gain %
Cross-line	9	58.52	34.92	35.76	−0.84	+40.3
CD4 series	4	271.55	142.07	174.47	−32.40	+47.7
External iPSC line	1	66.29	33.41	48.80	−15.38	+49.6
Additional units	48	65.15	33.74	34.97	−1.23	+48.2
Changed supports	48	65.15	36.09	37.38	−1.29	+44.6
Individual contrasts in the same nine line transfers						
NRDE2–PDCD5	9	100.17	40.68	47.84	−7.16	+59.4
SPCS2–POU5F1	9	90.05	53.18	60.13	−6.96	+40.9
BUD13–NDUFB6	9	17.60	10.45	12.33	−1.88	+40.6
BTAF1–HNRNPH1	9	7.88	15.05	12.10	+2.95	−90.9

J OCT4 ASSAY AND INTERPRETATION

RNA–OCT4 pairing uses donor, line, culture/preparation, target/guide, day 3 endpoint, batch, and matched control. The flow-cytometry protocol uses a BD LSRFortessa and BD Pharmingen PE anti-Oct3/4, clone 40/Oct-3, catalogue 560186, at 20 μl per test. Viability staining precedes fixation/permeabilization and intracellular staining. The protocol specifies fixed singlet, viability, and OCT4-positive gates, staining controls, appropriate compensation, and at least 10,000 valid singlets per condition as the acquisition target.

For pair (p, q), the assay contrast is $f_{\mathrm{OCT4+},p} - f_{\mathrm{OCT4+},q}$, with fractions in $[0, 1]$ and contrasts in $[-1, 1]$; multiplying by 100 gives percentage points. RNA combines log-transformed responses, so comparison concerns direction without a fitted quantitative mapping. With zero sign tolerance, agreement across 48 evaluation units is 43/48 for NRDE2–PDCD5 and 44/48 for each other illustrated contrast. The paired measurements compare RNA and protein-marker effects at the same endpoint.

Fig. 3h displays the paired contrasts and their directional agreement.

K MEASUREMENT SELECTION AND POLICY CONTROLS

Acquisition evaluates 24 queries through a signed unit-norm projection of a separate eight-coordinate response representation. Each budget comparison contains 480 decision starts across

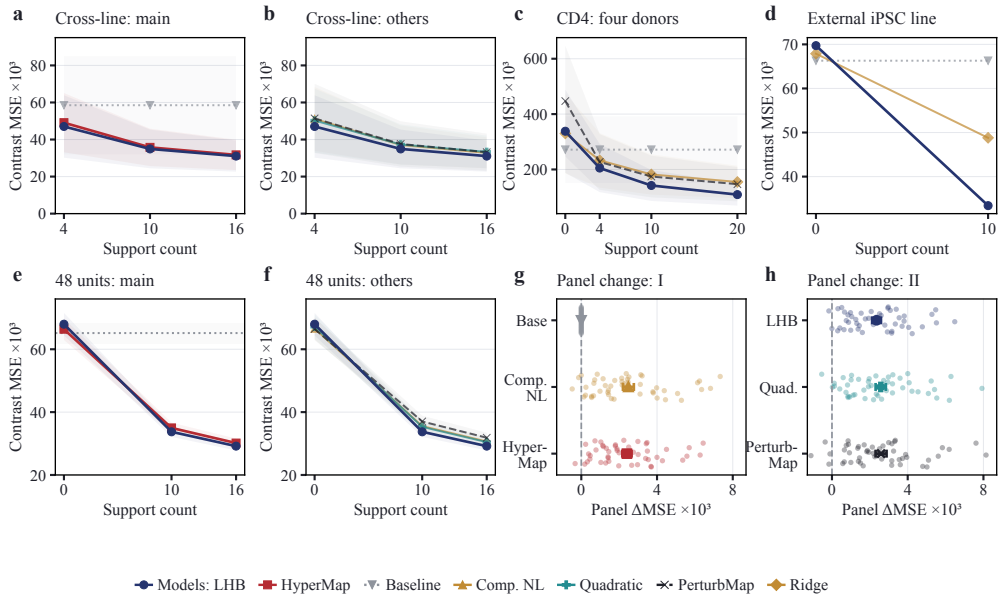

Figure S8: Extended prediction evaluations. (a,b) Nine same-donor line transfers. (c) Four CD4 donors. (d) One external iPSC line. (e,f) The same 48 additional units, with LHB repeated. (g,h) Changed-minus-original support-panel errors; points are units and large markers are means ± one unit SE. Curves in a–c,e,f show means ± one donor/unit SE; e,g use zero-support baselines. Response scales differ.

ten donors and three computational refits per start. The global menu has no addition (A0) and five intensities each for depth, source perturbations, and source contexts. Main comparisons sample three intensities per type, giving ten choices before budget filtering. Budgets 256, 1,024, and 4,096 cap condition-and-control measurement blocks, including a 24-block pilot even for A0.

S0 makes no addition, S1 chooses randomly, S2 adds contexts, and S4 adds depth. S3 uses linear task design, S5 feature design, S6 generic gain prediction, and S7 information diagnostics. R3/R5/R6/R7 denote their rankers. Calibrated variants add stopping; G_6, G_7, and G_{ref} denote generic, information-based, and common-reference rules. The rule labelled "All" always acquires, and Ref applies G_{ref} (Fig. 5d). Complete-policy and ranking comparisons use S5-cal and R6, respectively, selected on 96 development units.

Model-based diagnostics. The numerical policy evaluation uses variance-proxy summaries of perturbation, decoder, support, noise, and mismatch information. A separate model-based implementation computes diagnostics directly from fitted representations. The diagnostic, candidate-feature, and calibration formulas below specify this implementation. Here A, B, H_c denote fitted quantities, and Φ is the unweighted feature matrix. Let $G_A = A^\top A = \sum_c H_c^\top H_c$, $G_S = B_S B_S^\top$, $G_\varphi = \Phi\Phi^\top$ and $M_{\mathrm{diag}} = G_\varphi + \lambda_{\mathrm{diag}} I$. For each of J query contrasts, target support supplies estimated coordinates and hence a feature difference t_j. Write $\gamma = \|V^\top w\|$ and let P_Φ project onto the left singular vectors of Φ whose singular values exceed $\max\{e_\Phi, 10^{-8}\sigma_1(\Phi)\}$, where e_Φ is a feature-noise threshold. The default is $e_\Phi = 0$. The principal diagnostics are

$$d_A = \max\{0, \lambda_{\min}(G_A)\}, \qquad d_U = \frac{\gamma^2}{J}\sum_{j=1}^{J} \|(I - P_\Phi)t_j\|^2, \tag{27}$$

$$d_V = \frac{\gamma^2\widehat{\sigma}^2}{J}\sum_{j=1}^{J} t_j^\top M_{\mathrm{diag}}^{-1} G_\varphi M_{\mathrm{diag}}^{-1} t_j, \qquad d_S = \frac{\widehat{\sigma}^2}{|D|}\sum_{p\in D} b_p^\top (G_S + \lambda_{\mathrm{diag}} I)^{-1} b_p. \tag{28}$$

These proxies summarize source directions, uncovered feature contrasts, conditional output variance, and support conditioning. The diagnostic penalty is $\lambda_{\text{diag}} = \lambda_C$ in the diagnostic calculation; target adaptation uses its separate penalty λ_H. For gain prediction, this implementation uses $\widehat{\sigma}^2 = 0.05\ \mathrm{Var}(R^A) + 10^{-10}$ as an empirical scale proxy, with variance computed over available entries.

Candidate features and ranking. Candidate coordinates from current measurements or budgeted pilots give feature columns X_a that augment the design to $[\Phi, X_a]$. The candidate features include the resulting reductions in d_U and d_V, and $\ell_a^{\varphi} = \log\det(M_{\text{diag}} + X_a X_a^\top) - \log\det(M_{\text{diag}})$. An available candidate context estimate $\widehat{H}_a$ gives $G_{A,a} = G_A + \widehat{H}_a^\top \widehat{H}_a$. Its linear task proxy is

$$\ell_a^{\text{lin}} = \frac{\widehat{\sigma}^2}{|D|} \sum_{p \in D} b_p^\top \left[(G_A + \lambda_{\text{diag}} I)^{-1} - (G_{A,a} + \lambda_{\text{diag}} I)^{-1} \right] b_p.$$

For a depth addition with predicted remaining variance fraction ρ_a, the corresponding reductions are $(1-\rho_a)d_V$ and $\ell_a^{\text{lin}} = (1-\rho_a)d_S$. A binary indicator records whether candidate geometry is available. The input summaries also include response energy, source/support counts, noise scale, action type, cost, and numbers of added conditions and measurement blocks. The generic predictor includes candidate coordinate energy; the information predictor uses the source, readout, support and candidate-design diagnostics above.

Let x_a collect these candidate features. For a feasible positive-cost candidate with cost c_a, the rankers use

$$s_3(a) = \ell_a^{\text{lin}}/c_a, \qquad s_5(a) = \ell_a^{\varphi}/c_a,$$
$$s_6(a) = \widehat{g}_{\text{generic}}(x_a), \qquad s_7(a) = \widehat{g}_{\text{info}}(x_a).$$

Each learned gain model standardizes inputs using its training data, projects them onto at most four principal directions of those data, and fits ridge regression to the observed loss reduction after refitting, retaining negative gains. At budget $\mathcal{B}$, it chooses the highest score among candidates satisfying $c_a + c_{\text{pilot}} \le \mathcal{B}$; ties prefer lower cost and then the fixed action identifier.

Selected-action calibration and stopping. In this implementation, calibration uses out-of-fold selected-action paths for each ranker. Development splits exclude the held-out donor from all source, candidate, and query roles used to train the gain models. Within each outer training split, inner held-out selections train a separate gain regressor for the chosen action. Applying this regressor to outer held-out selections yields residuals $r_i = g_i - \widehat{g}_i$. Empirical residual quantiles give

$$I(a) = \left[\widehat{g}(a) + q_{0.025}(r), \widehat{g}(a) + q_{0.975}(r) \right].$$

The selected positive-cost action is executed only when its interval exists and its lower endpoint exceeds the frozen gain margin, zero in the default rule. Otherwise the policy returns A0 while retaining pilot cost. S3-cal, S5-cal and S6-cal use generic selected-action regressors; S7 uses an information regressor on the R7 selection path. The common reference rule uses only action-independent initial summaries, so swapping rankers leaves that rule's decision fixed. In the numerical evaluation, ranking uses proxy-based candidate forecasts and stopping uses their selected-action intervals. Gains and interval coverage are evaluated against the corresponding query losses after refitting.

Table S6 defines harm from fit-averaged decision gains. Missed improvement is zero for executed decisions. At budget 1,024, S7 coverage is 460/480 (95.83%) for selected candidates including stops and 223/238 (93.70%) for executed actions. For S7, harm occurs in 1.25% of single-fit runs but no fit-averaged decision.

Ranker comparisons hold reference stopping fixed; stopping comparisons fix the ranker and calibrate each rule along its selection path (Fig. S9d). Omitting perturbation, decoder, or noise diagnostics worsens mean loss. Omitting support diagnostics gives 0.04088 versus 0.04091 for the full policy, so component benefits depend on the evaluated mixture.

Module updates refit all source components, output weights only, embeddings and prior, prior only, or target H only. The last retains the source model and support, giving zero gain. Module, diagnostic-estimator, and pilot controls each use 120 starts across ten donors and three fits per start. Module gains average three intensities over 20 starts per information regime with equal donor

Table S6: Acquisition policies. Numerical contrast MSE, Diff. and Missed use $\times 10^3$. Gain is reduction from S0 (%). Diff. is MSE minus S5-cal. Cost is mean measurement blocks including pilots. Act is executed-addition rate. Harm is the percentage of starts with negative gain after averaging refits within decisions. Missed is best feasible gain forgone by stopping. Act, Harm and Missed average over all starts.

Policy	MSE	Gain %	Diff.	Cost	Act %	Harm %	Missed
Budget 256							
S0	47.64	0.0	+3.23	24.0	0.0	0.0	4.48
S1	45.97	+3.5	+1.56	118.9	100.0	24.6	0.00
S2	46.11	+3.2	+1.70	159.0	100.0	24.8	0.00
S4	45.64	+4.2	+1.23	80.0	100.0	21.2	0.00
S3-cal	45.09	+5.4	+0.68	61.8	41.2	0.4	1.85
S5-cal	44.41	+6.8	0.00	70.3	50.6	0.0	1.22
S6-cal	44.43	+6.7	+0.02	70.3	50.2	0.0	1.23
S7	44.37	+6.9	−0.04	71.4	52.5	0.0	1.17
S3	44.53	+6.5	+0.11	110.7	100.0	18.8	0.00
S5	44.17	+7.3	−0.25	106.7	100.0	17.9	0.00
S6	44.13	+7.4	−0.28	104.6	100.0	18.3	0.00
Budget 1,024							
S0	47.64	0.0	+6.57	24.0	0.0	0.0	8.89
S1	45.70	+4.1	+4.63	325.7	100.0	26.0	0.00
S2	46.78	+1.8	+5.71	744.0	100.0	32.7	0.00
S4	44.63	+6.3	+3.56	324.0	100.0	26.5	0.00
S3-cal	42.29	+11.2	+1.22	215.4	39.6	0.0	3.49
S5-cal	41.07	+13.8	0.00	253.6	47.7	0.0	2.26
S6-cal	41.01	+13.9	−0.06	256.9	48.3	0.0	2.20
S7	40.91	+14.1	−0.16	266.0	49.6	0.0	2.09
S3	41.23	+13.5	+0.16	373.9	100.0	19.2	0.00
S5	40.41	+15.2	−0.66	355.4	100.0	19.0	0.00
S6	40.35	+15.3	−0.72	352.7	100.0	18.5	0.00
Budget 4,096							
S0	47.64	0.0	+6.64	24.0	0.0	0.0	11.72
S1	46.47	+2.5	+5.46	760.1	100.0	28.3	0.00
S2	50.74	−6.5	+9.73	2524.0	100.0	41.2	0.00
S4	46.73	+1.9	+5.72	1324.0	100.0	37.3	0.00
S3-cal	40.09	+15.9	−0.91	777.6	41.9	0.0	3.98
S5-cal	41.00	+13.9	0.00	572.2	34.2	0.0	4.96
S6-cal	41.05	+13.8	+0.05	566.8	34.4	0.0	4.98
S7	43.94	+7.8	+2.94	271.6	18.5	0.0	7.97
S3	39.04	+18.1	−1.96	1212.6	100.0	19.8	0.00
S5	38.07	+20.1	−2.93	1139.1	100.0	19.8	0.00
S6	37.99	+20.3	−3.01	1116.8	100.0	19.2	0.00

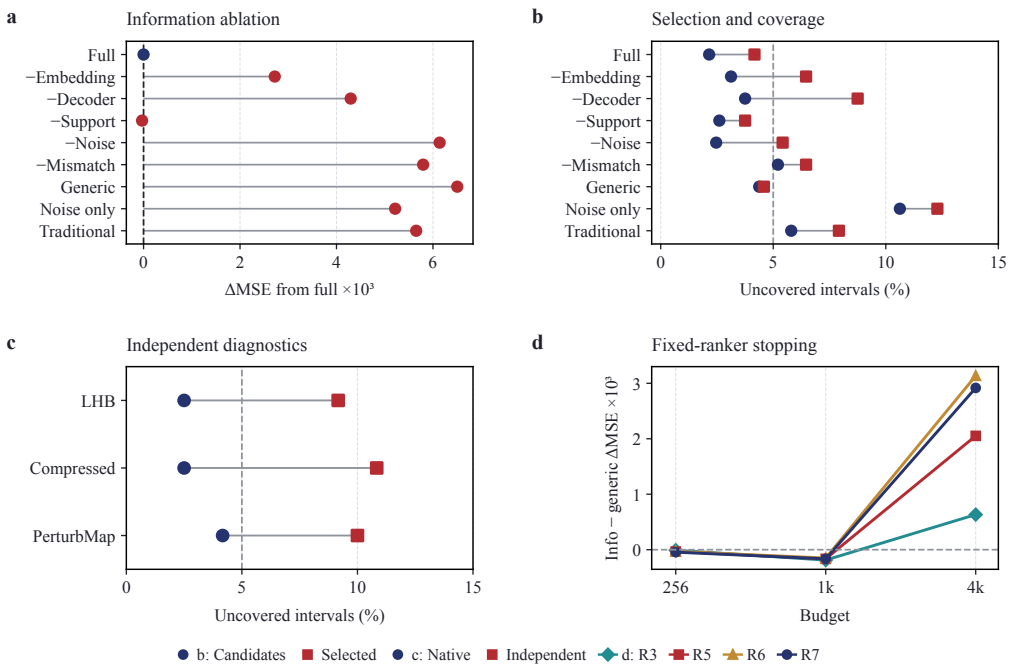

Figure S9: Diagnostic controls and calibration. (a) Nine input variants, relative to the full policy. Mismatch denotes an additional gain-predictor input; the panel shows the effect of omitting it. (b) Candidate and selected noncoverage, paired within each variant. (c) Selected noncoverage with native and independent diagnostics across three predictors, using 120 decision starts and ten donors. Dashed lines in b,c mark 5%. (d) Information-based minus generic stopping loss at each fixed ranker and budget, $\times 10^3$.

weights. With poorly estimated perturbation relationships, a new context gave gains of 0.01194 through embeddings and prior, 0.00012 through output weights, and 0.01115 through full refitting. With insufficient decoder information, added perturbations gave gains of 0.01416 through output weights, 0.00061 through embeddings and prior, and 0.01368 through full refitting.

Diagnostic-estimator comparisons use budget 1,024, three candidates per action type, four pilot conditions, and three predictors: LHB, compressed nonlinear, and adapted PerturbMap. Independent compressed diagnostics narrowed S7's mean gain over S5-cal from 0.000300–0.000327 to 0.000028–0.000030. Selected-candidate coverage fell from 95.83–97.50% to 89.17–90.83% (Figs. 5h and S9c).

Sensitivity checks cover six information regimes and 23 reweightings of existing observations. Pilot scans use 0, 2, 4, or 8 conditions (Fig. 5f). Menus with one, three, or five candidates per action type have four, ten, or 16 choices including A0; more choice does not uniformly improve performance. Across 120 evaluation paths per policy, at most one purchase occurs before stopping. Terminal S7 MSE is 0.04699 at cost 78.17, versus 0.04559 at cost 167.33 for S5-cal and S6-cal: lower expenditure accompanies higher loss.

The separate 128-unit comparison uses two panels and three fits per unit. S7-minus-S5-cal mean loss difference is -0.0002240 (95% interval $[-0.0003946, -0.0000534]$; 121 ties). R7-minus-R6 under common stopping gives -0.0000437 ($[-0.0001498, 0.0000624]$). Information-based minus generic stopping at fixed R6 gives -0.0003324 ($[-0.0006166, -0.0000483]$; 120 ties). These descriptive paired Student-t intervals summarize variation across the 128 evaluation units.