

Supplement S2. Does added molecular detail transport a predictive advantage?

Main manuscript: What a description must preserve to compare system responses

Elkin Antonio Navarro Quiroz

Faculty of Basic and Biomedical Sciences, Center for Research in Life Sciences (CICV), Universidad Simón Bolívar, Barranquilla, Colombia

Correspondence: elkin.navarro@unisimon.edu.co

ORCID: <https://orcid.org/0000-0001-7567-6409>

Supplement to *What a description must preserve to compare system responses*.

Elkin Antonio Navarro Quiroz. Analysis completed 8 October 2026. Exploratory secondary analysis of published murine EfferoSeq measurements; not a new experiment or independent biological validation. This supplement retains the completed result, including unfavorable findings. It does not include the subsequent historical-learning or open-system extensions of the local development notes.

S2.1 Question and scope

We asked whether within-cell cross-products of gene expression improve the ordering of a contemporaneous acidification-related fluorescence signal in a sequencing-label domain excluded from fitting, beyond a local additive quadratic predictor. This is a limited test of a specified representation and estimator. It does not test intercellular organization, causal memory, future efferocytic capacity or clinical utility.

The 30-gene panel and earlier descriptive channel results were known before this analysis. The new fitting design preceded adjustment, but this was not a confirmatory preregistration. Its recorded consistency prediction required joint-minus-additive Spearman correlation to be positive in all four evaluations. No biological relevance margin was justified for the uncalibrated signal. Failure of that prediction does not establish information sufficiency or equivalence of the rival predictor.

S2.2 Data, correspondence and selection

The source is the murine in-vitro component of [GEO GSE300757](#), associated with [Balaban et al., Science Immunology \(2026\)](#). GEO describes bone-marrow-derived macrophages cocultured with pHrodo-labelled apoptotic thymocytes at an offered ratio of 1:5. Detailed assay attribution here is to the GEO record; the publisher's full methods were not accessible in the historical retrieval. No tumor component was used in this analysis.

The deposited `GSE300757_invitro_qc.h5ad` has 8,715 cells and 24,337 genes. The raw instrumental archive was joined to the matrix using amplification-plate identity and normalized well position. The exact channel was *YG 586/15 - 570LP-A. Of 8,715 matrix cells, 7,099 have this channel; label 2 lacks it entirely. Header-unrecognized archive members are enumerated in the historical `results/execution.json`, rather than silently regarded as complete coverage.

We selected 3,481 phago cells with finite positive channel values from labels 3 and 4. Terminal 3 h and 24 h measurements do not track the same cells longitudinally. These labels coincide with sequencing batches/dates; independent animal/preparation identity is unresolved.

Evaluation	Terminal time	Training cells	Test cells
Label 3 → label 4	3 h	893	807
Label 4 → label 3	3 h	807	893
Label 3 → label 4	24 h	1,153	628
Label 4 → label 3	24 h	628	1,153

Table S2.1. Cell counts identify the computation, not the number of independent biological replicates. Reversed evaluations reuse the two labels; they are not four independent preparations. No population confidence intervals or biological p-values were derived from cell counts.

S2.3 Models and evaluation

Expression was transformed as $\log_{10}(\text{UMI}/\text{cell total} \times 10,000)$, using totals recomputed from the deposited sparse matrix. Quality covariates Q were $\log_{10}(\text{total UMI})$, $\log_{10}(\text{number of detected genes})$ and deposited mitochondrial percentage. The positive instrumental target was transformed by $\log_{10}$. Matrix integrity checks included integer nonnegative counts, agreement with deposited totals, unique gene identifiers, unique joining keys and disjoint training/test cell IDs.

The fixed panel contained the following 30 genes, grouped only to define six within-cell means:

Group label	Genes
Recognition/engulfment	Mertk, Axl, Timd4, Itgav, Itgb5, Rac1
Processing	Lamp1, Lamp2, Ctsb, Ctsd, Lipa, Tfeb
Lipid handling	Abca1, Abcg1, Nr1h3, Apoe, Soat1
Metabolic capacity	Ucp2, Dnm1l, Pfkfb2, Slc16a1, Slc2a1
Contextual output	Vegfa, Hif1a, Il10, Tgfb1
Interferon response	Stat1, Irf1, Cxcl10, Isg15

Table S2.2. Labels aid interpretation; these are not measurements of six biological activities. Neither means nor products describe neighbors or interactions between cells.

Five predictors were fitted separately in each evaluation: Q only; Q plus six gene means; Q plus all 30 genes linearly; Q plus genes and their 30 individual squares (the **additive quadratic rival**); and the additive model plus all 435 distinct pair-products (the **joint expansion**). Nominal non-intercept feature counts were 3, 9, 33, 63 and 498. Training-constant columns were removed, yielding respectively 3, 9, 32, 61 and 417-445 retained columns in the saved primary fits. Thirty genes were specified; fewer variable columns do not mean a retrospectively selected smaller gene panel.

Each model had an intercept. Columns were standardized using training means and standard deviations only. Ridge penalty was fixed at $\lambda = 0.1 \times \text{training } n$; $\lambda/n = 0.01$ and 1 were sensitivities, not candidates selected using test outcomes. No outcome clipping or post-result refitting was used. Own ridge predictions were compared with scikit-learn in all 20 primary fits, agreeing to at most 4.0×10^{-14} .

The primary metric was held-out Spearman rank correlation. Secondary metrics were mean squared error and R^2 on the $\log_{10}$ signal. R^2 used the held-out outcome mean as its reference, which is available for evaluation but not deployment. There were four separate evaluations rather than a random cell split or a pooling of terminal times.

Twenty diagnostic permutations of the training target within UMI deciles were generated for each evaluation with seed 261008, retaining expression, label and time. They preserve only discretized UMI dependence and are not biological interventions or valid population null distributions. Saved sensitivity and permutation tables contain all evaluations, not a selected favorable comparison.

S2.4 Completed results

Training → test	Time	Additive quadratic p	Joint p	Joint – additive
3 → 4	3 h	−0.0242	0.0005	+0.0247
4 → 3	3 h	0.0509	0.0044	−0.0465
3 → 4	24 h	0.1040	−0.0066	−0.1106
4 → 3	24 h	0.0591	0.0520	−0.0072

Table S2.3. The declared positive-gain prediction failed in three of four evaluations. The primary rival includes individual squares; it is not the linear predictor. Differences are from `results/contrasts.tsv`, with full unrounded scores retained in `results/scores.tsv`.




Figure S2.1. Spearman correlation in the label excluded from fitting, shown for all five primary models and all four evaluation domains. Test-cell counts are 807, 893, 628 and 1,153 in the displayed order. Labels are confounded with sequencing batches/dates and independent animal/preparation identity is unresolved. No biological uncertainty interval is inferred from these cell counts. This is the unchanged figure from the finished pilot, not a new execution of the portable code.

All primary correlations were below 0.163. Neither sensitivity penalty gave a positive joint-minus-additive difference in all four evaluations. We therefore do not retain a claim of consistent advantage for this joint expansion in this domain.

All 20 primary R^2 values were negative. A diagnostic added after seeing these results compared model MSE with a deployable constant equal to the training outcome mean. The joint expansion had greater MSE in all four domains; the other models had mixed outcomes. Joint predictions in label 3 → 4 at 3 h reached $\log_{10} p = 33.246$, whereas the largest observed value in that test domain was 11.122. This indicates severe extrapolation; predictions were not clipped or refitted afterward. Complete diagnostics appear in `results/post_result_diagnostics.tsv` and are identified as post-result rather than part of the original fitting specification.

Every observed joint correlation fell within the corresponding diagnostic permutation min-max. These controls do not perfectly preserve continuous quality/gating dependence, create independent biological replicates or yield valid population p-values.

S2.5 Interpretation and limits

The tested expansion failed to transport a consistent ranking advantage over the specified additive quadratic rival. Model capacity, sparse cross-products, domain shift and measurement processes remain competing explanations. A negative finite-model comparison neither refutes distributed information generally nor proves that the simpler description is sufficient for an optimal predictor. A nonlinear predictor of the full within-cell gene vector already retains associations; adding pair-products is an estimator choice, not new spatial information.

pHrodo is an acidification-related signal, not a count of physical ingestion events, completed digestion or future function. The cells were selected post-exposure; calibration, compensation and interbatch gating remain unresolved, and terminal RNA can reflect ingested material and cell size. Offered ratio does not measure effective exposure or retained intracellular cargo. No linked animal/preparation map, later uptake target, same-cell history, neighboring cells or tissue coordinates was available. More molecular features cannot supply these missing observables.

The prospective repeated-efferocytosis application in the main article therefore remains a separate proposal. This pilot does not validate it, repair a novelty claim about the conceptual synthesis or identify a natural organizational mechanism. It illustrates why the usefulness of added detail must be assessed for a specified target and estimator, while retaining unfavorable results.

S2.6 Reproducibility and provenance

This package includes own analysis code, the three original pre-fit design/observability documents, saved model predictions, all primary scores, contrasts, sensitivities, permutation controls, post-result diagnostics and the original figure. The original per-cell observed_channel column has been omitted from the delivered predictions_model_outputs.tsv; retained prediction values and cell/fold/model correspondence are unchanged. No original input matrix, archive, per-cell instrumental readings or original-study software is redistributed. Inputs must be acquired separately from GEO and match the two identities in input_sources.json before reproduction. The portable runner then recreates the complete joined prediction table needed by the verifier; the delivered model-only table is not itself sufficient to recalculate target-dependent metrics.

The historical execution checked source SHA256 values before and after fitting; 20 metrics were recalculated from saved predictions and five central tables were byte-identical across local reruns. A Windows cp1252 final console-print failure was repaired using UTF-8 without changing the fitting specification or outcomes. Those checks occurred in the existing scientific environment, not a clean installation or independent experimental replication.

The delivered code has now been adapted only for source/output directory arguments. Static parsing/compilation and analysis-function comparisons, byte-identical copying of twelve historical files, and unchanged value strings after filtering the prediction table were checked; **the portable code has not been executed to refit or generate new scientific results**. results/execution.json remains an explicitly labelled historical receipt: its input path strings were normalized, but its original code hash, specifications, dates and outcomes were retained. README instructions distinguish the archived result from a future reproduction.

For source acquisition, input hashes, dependencies, commands and licensing boundaries, see README.md and input_sources.json. Source availability does not assert a universal CC0/CC BY license for GEO data. Author responsibility and substantive AI assistance are declared in the main article; these computational checks are not human peer review or biological validation.