Additional file 1

Assay resolution governs the transfer of molecular dependence

Anonymous Authors

This supplement contains the identified-set and response proofs, the adaptation-only prediction protocol, and the conditional coexpression analyses. Earlier estimator comparisons and the complete historical study record remain in the original public benchmark [5].

1 Assay resolution and recipient coexpression

This section gives the theory and retrospective analyses added after the original fixed-margin benchmark. All resolution comparisons hold the source interaction fixed.

1.1 Sharp identified interval under arbitrary marginal constraints

Let x, y range over finite assay-profile supports, let $S(x, y)$ be finite, and consider positive joint distributions

$$q(x, y) = \exp\{S(x, y) + a(x) + b(y)\}. \tag{1}$$

Let $\phi(x)$ and $\psi(y)$ collect any finite sets of observed within-assay features, such as all individual marker indicators. For observed means μ_x, μ_y, define the marginal polytopes

$$\mathcal{R}(\mu_x) = \left\{ r : \sum_x r(x) = 1,\ r \geq 0,\ \sum_x r(x)\phi(x) = \mu_x \right\},$$
$$\mathcal{C}(\mu_y) = \left\{ c : \sum_y c(y) = 1,\ c \geq 0,\ \sum_y c(y)\psi(y) = \mu_y \right\}.$$

For each r, c, let $Q_S(r, c)$ be the unique minimizer over couplings with those marginals of

$$\Phi_S(q) = \sum_{x,y} q(x, y)\{\log q(x, y) - S(x, y)\}, \tag{2}$$

with $0 \log 0 = 0$. For positive marginals, the Lagrange equations give Eq. (1). We define admissible reconstructions directly as $Q_S(r, c)$, allowing zero-mass assay states. This includes summary constraints that exclude every strictly positive assay distribution.

Proposition S1 (Identified interval under observed assay summaries)**.** *Suppose $\mathcal{R}(\mu_x)$ and $\mathcal{C}(\mu_y)$ are nonempty. Let $U = f(x)$ and $V = g(y)$ be binary queries whose frequencies m_U, m_V are fixed by the observed features. Define*

$$\ell = \min_{r \in \mathcal{R}(\mu_x), c \in \mathcal{C}(\mu_y)} \sum_{x,y} f(x)g(y)Q_S(r, c)(x, y),$$
$$h = \max_{r \in \mathcal{R}(\mu_x), c \in \mathcal{C}(\mu_y)} \sum_{x,y} f(x)g(y)Q_S(r, c)(x, y).$$

Then the query tables obtained from admissible reconstructions are exactly $\{P(t) : \ell \leq t \leq h\}$, where

$$P(t) = \begin{pmatrix} 1 - m_U - m_V + t & m_V - t \\ m_U - t & t \end{pmatrix}.$$

Both endpoints are attained. If $\ell < h$, both query frequencies lie in $(0, 1)$, and the unique minimax prediction under forward KL loss, among all single query-table predictions, is $P_* = P(t(\theta_*))$, where

$$\theta_* = \frac{\mathcal{H}(P(\ell)) - \mathcal{H}(P(h))}{h - \ell}, \qquad \mathcal{H}(P) = -\sum P \log P,$$

and $t(\theta)$ is the cell probability with margins m_U, m_V and log odds θ. For $\ell = h$, the query is identified and $P_* = P(\ell)$.

Proof. The two marginal sets are compact convex polytopes. For fixed r, c, the transportation polytope is compact, and the strictly convex objective in Eq. (2) has a unique minimizer. In finite dimensions its feasible couplings vary continuously with their margins. To see this at a boundary, let $r_n \to r$ and $c_n \to c$. Given any coupling q of r, c, apply maximal-coupling kernels carrying r to r_n and c to c_n, independently conditional on its two coordinates. The resulting coupling q_n has the new margins and satisfies

$$\|q_n - q\|_1 \le \|r_n - r\|_1 + \|c_n - c\|_1.$$

Conversely, any limit of couplings with margins r_n, c_n has margins r, c. The objective is continuous at zero probabilities because $u \log u \to 0$ and S is finite. Compactness, these approximating couplings, and uniqueness of the minimizer imply $Q_S(r_n, c_n) \to Q_S(r, c)$. The queried cell probability is consequently a continuous real-valued function on the compact connected set $\mathcal{R}(\mu_x) \times \mathcal{C}(\mu_y)$. Its image is a compact connected subset of the real line, hence the interval $[\ell, h]$. The fixed query frequencies determine the other three cells, proving the identified-set claim.

For a fixed prediction, KL loss is convex in the true table, so its maximum over the interval occurs at an endpoint. Parameterize fixed-margin predictions by log odds θ. For an endpoint cell probability t_0, $\partial_\theta \operatorname{KL}(P(t_0)\|P(t(\theta))) = t(\theta) - t_0$. The two endpoint losses therefore vary in opposite directions between ℓ and h and have a unique equalizer. Direct subtraction gives the stated θ_*. To prove minimax optimality over all predictions, choose w such that $P_* = wP(\ell) + (1-w)P(h)$. A prediction assigning zero to an endpoint-positive cell has infinite worst-case loss. For any other prediction R,

$$\begin{aligned} &w \operatorname{KL}(P(\ell)\|R) + (1-w)\operatorname{KL}(P(h)\|R) \\ &\quad = R_* + \operatorname{KL}(P_*\|R), \end{aligned}$$

where $R_* = \operatorname{KL}(P(\ell)\|P_*) = \operatorname{KL}(P(h)\|P_*)$. The worst endpoint loss cannot be smaller than this weighted average, with equality uniquely at P_*. Populations throughout the identified interval have the same supplied summaries. Convexity in the prediction also makes randomization unhelpful. □

Taking ϕ and ψ to be all marker indicators gives the identified set for the means-only benchmark. Taking them to be indicators of complete assay profiles makes both marginal polytopes singletons and collapses the interval. Thus the proposition compares exactly the information supplied to the two experimental arms. The means-only information projection is a point in this interval; the interval allows every complete assay marginal compatible with the reported frequencies.

1.1.1 Closed form with only the queried frequencies

Suppose the only constraints are the frequencies $m_U, m_V \in (0, 1)$. For every rectangle crossing the four query fibers, define

$$d = S(x_0, y_0) + S(x_1, y_1) - S(x_0, y_1) - S(x_1, y_0),$$

where $f(x_0) = g(y_0) = 0$ and $f(x_1) = g(y_1) = 1$. Write $L = \min d$ and $H = \max d$. The closure of compatible coarse log odds is $[L, H]$, so Proposition S1 has $\ell = t(L)$ and $h = t(H)$.

Every fine distribution in Eq. (1) satisfies

$$q(x_0, y_0)q(x_1, y_1) = e^d q(x_0, y_1)q(x_1, y_0).$$

Summing over rectangles in the four coarse blocks expresses the coarse odds ratio as a positively weighted average of e^d. Its log therefore belongs to $[L, H]$. Concentrate the fine RNA masses $1 - m_U, m_U$ on an extremizing x_0, x_1 and protein masses $1 - m_V, m_V$ on y_0, y_1. Their two-by-two reconstruction attains the corresponding endpoint. Positive fine distributions approximate these concentrated distributions while preserving the coarse margins. The positive fine-marginal domain is connected, and its entropic reconstruction is continuous, so the closure of attainable coarse tables is exactly the stated interval. The equalizing prediction in Proposition S1 is the corresponding information-radius minimax rule [6].

For the three-marker example in the main text, take the RNA distribution $r(z_1, z_2) = (1 + \rho z_1 z_2)/4$, with $|\rho| < 1$. The joint distribution

$$q(z_1, z_2, w) = r(z_1, z_2)\frac{e^{J z_2 w}}{2\cosh J}$$

has all three individual means zero and the required fine interaction. Conditional expectation gives $\mathbb{E}(W \mid Z_1, Z_2) = Z_2 \tanh J$ and hence $\mathbb{E}(Z_1 W) = \rho \tanh J$. All probabilities are positive, so the example does not depend on structural zeros. With all three individual means fixed at zero, ρ ranges over $[-1, 1]$ in closure. The sharp all-marker interval for $\Pr(Z_1 = 1, W = 1)$ is therefore $[(1 - |\tanh J|)/4, (1 + |\tanh J|)/4]$.

1.2 Response to a transferred interaction

The expansion below specializes the fixed-marginal score-projection identities of Sei and Yano [4]. Fix positive marginal masses $r(x)$ and $c(y)$ on finite supports. Write $\bar{x} = x - \mathbb{E}_r X$, $\bar{y} = y - \mathbb{E}_c Y$, and $h = \bar{x}^\top B \bar{y}$. Centering changes the interaction only by additive assay functions. The reconstruction can therefore be written

$$q_t(x, y) = r(x)c(y) \exp\{th(x, y) + \alpha_t(x) + \beta_t(y)\}.$$

Let $s_t = \partial_t \log q_t$. Differentiating the two fixed marginal constraints gives

$$\mathbb{E}_{q_t}(s_t \mid X) = 0, \qquad \mathbb{E}_{q_t}(s_t \mid Y) = 0.$$

Because the assay means do not vary with t,

$$\frac{d}{dt} \operatorname{Cov}_{q_t}(X, Y) = \mathbb{E}_{q_t}\{\bar{X}\bar{Y}^\top s_t\}, \tag{3}$$

$$\operatorname{Cov}_{q_1}(X, Y) = \int_0^1 \mathbb{E}_{q_t}\{\bar{X}\bar{Y}^\top s_t\}\, dt. \tag{4}$$

These identities are exact. At $t = 0$, $q_0 = rc$ and the centered interaction has zero conditional mean in either argument, so $s_0 = h$.

At $t = 0$, the adjusting functions vanish. After fixing their additive gauge, the finite marginal equations have a nonsingular linearization at independence, so their solution is analytic near zero. Differentiating the marginal constraints once gives zero conditional means for $\partial_t \log q_t$. Centering h in both arguments therefore gives $\partial_t \log q_t|_{t=0} = h$. Differentiating twice gives

$$\left.\frac{\partial_t^2 q_t}{rc}\right|_{t=0} = h^2 - \mathbb{E}_c(h^2 \mid x) - \mathbb{E}_r(h^2 \mid y) + \mathbb{E}_{rc} h^2.$$

When multiplied by $\bar{x}_i \bar{y}_j$ and averaged under rc, the last three terms vanish. Define the third central moment tensors by $T^x_{iac} = \mathbb{E}_r(\bar{X}_i \bar{X}_a \bar{X}_c)$ and $T^y_{jbd} = \mathbb{E}_c(\bar{Y}_j \bar{Y}_b \bar{Y}_d)$. It follows that

$$\operatorname{Cov}_{q_t}(X_i, Y_j) = t(\Sigma_x B \Sigma_y)_{ij} + \frac{t^2}{2} \sum_{a,b,c,d} B_{ab} B_{cd} T^x_{iac} T^y_{jbd} + O(t^3).$$

For an explicit example, take three spins in each assay and let

$$r(x) = \frac{1 + \rho_x x_1 x_2 x_3}{8}, \qquad c(y) = \frac{1 + \rho_y y_1 y_2 y_3}{8},$$

with $x_i, y_i \in \{-1, 1\}$ and $|\rho_x|, |\rho_y| < 1$. Every individual mean is zero and each within-assay covariance matrix is the identity. For diagonal $B = \operatorname{diag}(b_1, b_2, b_3)$,

$$\mathbb{E}_{q_t}(X_1 Y_1) = tb_1 + t^2 \rho_x \rho_y b_2 b_3 + O(t^3).$$

The pairwise maximum-entropy replacement is uniform in each assay and loses the second term. Positive distributions with identical means and covariances can therefore give different RNA–protein associations under the same bilinear interaction.

1.2.1 Query discrepancy from covariance mismatch

Consider recipient and pooled reconstructions with the same binary marker frequencies $m, n \in (0, 1)$ and interaction $tx^\top By$. Let k_R and k_P be their respective queried entries of $\Sigma_x B \Sigma_y$. Both query tables equal the independent table P_0 at $t = 0$, and

$$P_R(t) - P_P(t) = t(k_R - k_P)V + O(t^2),$$

where $V = \left(\begin{smallmatrix} 1 & -1 \\ -1 & 1 \end{smallmatrix}\right)$. Taylor expansion of KL divergence about the positive table P_0 gives

$$\begin{aligned} &2\,\mathrm{KL}(P_R(t) \| P_P(t)) \\ &\quad = t^2(k_R - k_P)^2 \sum_{u,v} \frac{V_{uv}^2}{(P_0)_{uv}} + O(t^3) \\ &\quad = \frac{t^2(k_R - k_P)^2}{m(1-m)n(1-n)} + O(t^3). \end{aligned}$$

The coefficient is the response score used to rank queries. If the recipient model describes the true joint law, it is also the leading excess deviance from substituting pooled coexpression. The empirical comparison tests the ranking at the fitted interaction strength without imposing that equality.

1.3 Finite-strength response diagnostic

We scaled the frozen interaction over $t \in \{0.05, 0.10, 0.25, 0.50, 0.75, 1\}$ without changing assay marginals, recipient pools, or queries. At each value, exact matrix scaling reconstructed the recipient covariance model and leave-one-recipient-out cohort-pool model. Their query discrepancy was twice the forward KL divergence from the recipient table to the pool table. We compared this quantity with the local response coefficient fixed before the continuation.

For each cohort and value of t, we computed the pooled Spearman correlation, the mean within-recipient Spearman correlation, overlap between response-ranked and exact top-27 queries, and the ratio of summed exact discrepancy to its quadratic approximation. Recipient-bootstrap 95% intervals used 20,000 draws with seed 906095. At $t = 1$, independently reconstructed tables agreed with the frozen primary predictions to maximum absolute error 5.56×10^{-16}.

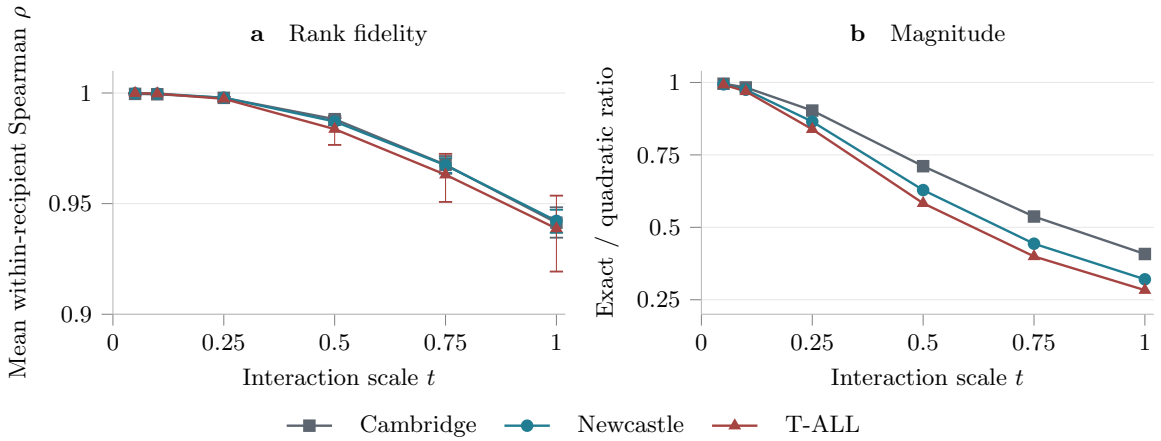

Figure 1: Mean within-recipient correlation between the local response score and exact discrepancy between recipient-covariance and cohort-pool models (left), and the ratio of exact discrepancy to its quadratic approximation (right). Bars are recipient-bootstrap 95% intervals. This diagnostic uses the conditional fixed-frequency protocol.

1.4 Adaptation-only prediction

The stricter analysis retains the source and recipient membership, nine-marker panel and deterministic cell selection. For each person, RNA is binary detection and protein is strict exceedance of the median count among the first 256 cells. Ties are low. Applying this threshold to the remaining 256 cells does not modify the threshold. Source donors use the same definition.

Four random within-donor matchings of the 512 source cells per donor, with seed 906096, form the pairing contrasts. The source fit uses penalty 10^{-4} without retuning. Each source-prior marker probability is $(n_1 + 1/2)/(6144 + 1)$, where n_1 counts the high state among the 12 source donors' 6144 cells. Their product-Bernoulli distribution supplies total pseudocount mass four to each recipient adaptation histogram, normalized together with its 256 observed cells. No scoring-cell summary is supplied.

The four resolution arms constrain means, one complete assay distribution, or both complete distributions, as in the conditional analysis below. An independence arm is the product of the same estimated marker frequencies. All predictions were saved before forming the empirical scoring tables and their losses. Each loss averages twice forward KL over all 81 marker pairs. Twenty thousand paired recipient-bootstrap samples, seed 906097, provide nominal 95% intervals and Bonferroni percentile intervals across the nine full-pattern contrasts with means, protein patterns, and independence.

Table 1: Adaptation-only prediction. Mean recipient deviance over all 81 pairs.

Recipient information	Cambridge	Newcastle	T-ALL
Individual frequencies	0.022292	0.032327	0.058471
RNA patterns, protein frequencies	0.020598	0.027637	0.045103
RNA frequencies, protein patterns	0.019638	0.026090	0.038081
Both assay patterns	0.018585	0.022889	0.025986
Independence at estimated frequencies	0.034503	0.043524	0.074591

An independent implementation recounted raw thresholds and scoring tables, reconstructed the 12,288 source contrasts, and reproduced all 380 model reconstructions by iterative proportional fitting plus 95 independence

tables. Maximum probability and loss differences were 7.68×10^{-11} and 5.79×10^{-12}. Altering scoring counts left thresholds and predictions unchanged; permuting adaptation profiles separately within each assay also left predictions unchanged. All nine bootstrap contrasts were reproduced.

1.5 Matched resolution and covariance analyses

The three recipient cohorts contain 24 Cambridge donors, 56 Newcastle donors, and 15 T-ALL patients. The same nine RNA and nine protein markers give 81 ordered queries. All analyses retain the original deterministic 512-cell sample, donor-wide protein ranks, 256 adaptation cells, 256 scoring cells, and 12-donor source interaction with pairing-contrast penalty 10^{-4}. The paired source fit is the only source of cross-assay interaction information. Recipient adaptation uses separate-assay patterns and the scoring sample's individual marker frequencies, not its pairings.

The means-only reconstruction constrains all 18 individual frequencies. Each one-pattern arm additionally constrains that assay's full binary-profile histogram. The both-pattern arm constrains both histograms. Histograms receive source-prior mass four and exponential tilts to the scoring-cell marker frequencies. Constant markers are restricted to their observed states. Matrix scaling fits complete pattern constraints, and moment matching fits individual-frequency constraints. These are reconstructions from the same interaction, not separately trained interaction models.

The covariance arm replaces each smoothed and tilted histogram with its maximum-entropy distribution subject to the same first and second moments. For binary profiles its log density contains unary and pairwise terms. This replacement can induce higher-order correlations; it fixes pairwise moments rather than forcing all higher-order moments to vanish. The retained-gain fraction is

$$\frac{L_{\text{means}} - L_{\text{covariance}}}{L_{\text{means}} - L_{\text{patterns}}},$$

where each L is the cohort's mean donor deviance. It is a descriptive loss decomposition, not an estimate of a mediated biological effect.

1.6 Recipient-specific and pooled coexpression

The source pool uses the 12 source donors' adaptation halves: 3,072 cells per assay. For each query person, the cohort pool uses every other recipient's adaptation half and excludes the query person. It contains 5,888 cells in Cambridge, 14,080 in Newcastle, and 3,584 in T-ALL. Both controls use four prior units of prior mass per contributing person, preserving the prior fraction, then undergo the same tilting and maximum-entropy reconstruction at the query person's marker frequencies. The cohort pool uses recipient-cohort data but no scoring-cell patterns. These controls vary the origin of covariance without changing the interaction, marker frequencies, or scoring tables.

1.7 Query ranking and uncertainty

For RNA marker i and protein marker j, define

$$s_{ij} = \frac{\{(\Sigma_x B \Sigma_y)_{ij}^{\text{recipient}} - (\Sigma_x B \Sigma_y)_{ij}^{\text{pool}}\}^2}{m_i(1-m_i)n_j(1-n_j)}.$$

The source-strength comparator is $B_{ij}^2 m_i(1-m_i)n_j(1-n_j)$. Scores are zero for constant-marker queries. Both rankings use the fixed marker order to break ties. Pair-specific gain is the cohort-pool deviance minus recipient-covariance deviance, using unchanged predictions. The top 27, bottom 27, and source-ranked top 27 were specified before computing the pair-specific losses. Cumulative gain curves sum gains over people and selected queries, then divide by the sum over all 81 queries; they do not average person-specific percentages.

All loss endpoints average twice KL over every marker pair, including constants, then equally over people. Each comparison uses 20,000 paired person-level bootstrap draws. The four analysis families use seeds 906091–906094 and contain, respectively, nine pattern-resolution contrasts, six covariance-replacement contrasts, six recipient-versus-pool contrasts, and six query-ranking contrasts. For a family of M comparisons, the Bonferroni percentile endpoints are $0.025/M$ and $1 - 0.025/M$. These intervals condition on the fixed source fit and fitted pools.

1.8 Relation to full-profile interaction fitting

The earlier T-ALL comparison used all 15 patients, the same permitted recipient information, and additional protein-abundance coordinates. A common recipient adapter compared the no-prior CHAMPOLLION inverse-transport objective [3] with donor-wise joint fitting, query-targeted fitting, and a multinomial comparator.

Table 2: Query-ranking comparisons. Values are mean deviance-gain differences per selected pair. Intervals resample people and correct across six contrasts.

Cohort	Comparison	Difference	Favorable people	Bonferroni interval
Cambridge	Response top vs bottom third	0.001711	15/24	−0.000126 to 0.003883
Newcastle	Response top vs bottom third	0.004983	44/56	0.002657 to 0.007787
T-ALL	Response top vs bottom third	0.024259	14/15	0.014180 to 0.037793
Cambridge	Response vs source-ranked top third	0.000038	11/24	−0.001043 to 0.001154
Newcastle	Response vs source-ranked top third	0.001642	38/56	0.000612 to 0.002920
T-ALL	Response vs source-ranked top third	0.013881	14/15	0.007619 to 0.021405

Mean losses were 0.008441, 0.008476, 0.008602, and 0.010630, respectively. Query-targeted fitting improved on the multinomial comparator by 19.07% (95% interval, 13.05–24.36%); its relative improvement over the CHAMPOLLION objective was −1.91% (−6.26 to 2.25%). This comparison did not meet its all-control superiority criterion. It neither benchmarks the complete CHAMPOLLION workflow nor establishes a better interaction learner. The subsequent binary resolution analyses hold one fitted interaction constant.

1.9 Computational verification

Separate implementations reconstructed all 380 initial resolution predictions, 190 covariance-preserving assay distributions and their 95 joint predictions, and 190 pooled-control predictions. Maximum absolute differences in predicted probabilities were 5.04×10^{-11}, 1.14×10^{-10}, and 1.27×10^{-10}, respectively. Verification recounted the scoring tables, checked exact marker margins, and checked source membership and exclusion of the query person from each cohort pool. Independently calculated query scores differed by at most 1.47×10^{-10}; pair-specific gain and bootstrap-summary differences were below 4×10^{-16}. Tests cover marginal constraints, reconstruction equivalence, covariance replacement, pool construction, and query scoring.

Analysis plans preceded these computations but followed examination of earlier outcomes; the analyses are retrospective. The original public v2.0.4 benchmark is a separate release. The accompanying analysis artifact contains the derived assay distributions, fitted interactions, predictions, and losses needed to reproduce the subsequent reconstructions and statistical summaries. It does not redistribute raw count matrices or reproduce their initial acquisition.

2 Independent colon-biopsy test

The colon-biopsy study [2] provides a separate tissue and study for the fixed blood-reference interaction. We used the public count layer of Figshare file 54027674 in article 21919356, version 3, associated with GSE250490. Only cells labeled as paired gene-expression and protein libraries were eligible. Biopsies were pooled within physical donors. Eleven of twelve donors met the fixed budget of 512 cells; the remaining donor had 433 eligible cells. Cell selection ordered barcodes by SHA-256 with the fixed salt `colon-generalization-v1` and physical-donor identifier. The selected matrix rows are released under anonymous donor labels.

The protocol, selection, source fit, reconstruction code, and evaluation code were fixed publicly in commit `b6e6cab` before count access [1]. Metadata inspection had established tissue, clinical groups, assay compatibility, and available cell numbers. The first 256 selected cells defined RNA detection, protein thresholds, and smoothed within-assay distributions exactly as in the adaptation-only analysis. The remaining 256 cells supplied the endpoint. Predictions were saved and hashed before their scoring counts were opened. The source interaction and prior were unchanged, and no recipient parameter was selected.

Table 3: Colon-biopsy mean deviance. Clinical-group columns are descriptive; the inferential unit is the physical donor and the primary analysis combines all eleven donors. Controls number four, aminosalicylate-treated patients four, and vedolizumab-treated patients three.

Recipient information	All	Controls	Aminosalicylates	Vedolizumab
Individual frequencies	0.02588	0.02316	0.02796	0.02672
RNA patterns	0.02367	0.02171	0.02482	0.02476
Protein patterns	0.02308	0.02094	0.02428	0.02432
Both patterns	0.02183	0.02036	0.02212	0.02339
Independence	0.03169	0.02793	0.03584	0.03118

The prespecified success criterion required at least 10% reduction against individual frequencies and positive adjusted lower bounds against individual frequencies, protein patterns, and independence. Paired donor-bootstrap

intervals used 20,000 draws with seed 9132026 and Bonferroni percentile endpoints $0.025/3$ and $1 - 0.025/3$. The adjusted intervals were 10.89–22.40%, 2.56–9.49%, and 24.85–38.43%, respectively. All five prediction arms, eleven donor losses, and 81 query-level mean losses are released. No clinical group or query was used to select a result or modify the model.

An independent implementation reproduced all 55 donor-by-arm predictions with maximum probability difference 6.33×10^{-11}. It separately recounted the 891 scoring tables and reproduced all three bootstrap comparisons. Independent permutations within the two adaptation assays left prediction unchanged. The separate colon artifact includes acquisition and checksum verification, sufficient marginal distributions, predictions, losses, and scripts that reproduce reconstruction and the frozen endpoint.

References

[1] Assay-resolution study contributors. Blood-to-colon dependence transfer: frozen analysis protocol and results. https://github.com/sushaan-k/coupling-fields-benchmark/tree/main/analysis/colon_generalization, 2026. Protocol frozen before count access in commit b6e6cab, 13 September 2026.

[2] Elvira Mennillo, Yang Joon Kim, Gyehyun Lee, Iulia Rusu, Ravi K. Patel, et al. Single-cell and spatial multi-omics highlight effects of anti-integrin therapy across cellular compartments in ulcerative colitis. *Nature Communications*, 15:1493, 2024. doi: 10.1038/s41467-024-45665-6.

[3] Jules Samaran, Gabriel Peyré, and Laura Cantini. CHAMPOLLION: Robust multi-omics integration via inverse optimal transport using paired cells. *bioRxiv*, 2026. doi: 10.64898/2026.04.28.721317.

[4] Tomonari Sei and Keisuke Yano. Minimum information dependence modeling. *Bernoulli*, 30:2623–2643, 2024. doi: 10.3150/23-BEJ1687.

[5] sushaan-k. Coupling Fields: public benchmark, version 2.0.4, 2026. URL https://github.com/sushaan-k/coupling-fields-benchmark/releases/tag/coupling-fields-v2.0.4-public-benchmark. Accessed 5 September 2026.

[6] Tim van Erven and Peter Harremoës. Rényi divergence and Kullback–Leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014. doi: 10.1109/TIT.2014.2320500.