Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Evaluating porcine antisense oligonucleotide RNA-seq for cross- species virtual-knockout benchmarking

Xiaotong Zhao
Yunnan Agricultural University
Hua Chang
Yunnan Agricultural University
Zhuoyu Zhao
Yunnan Agricultural University
Xun Xiang
237915221@qq.com

Yunnan Agricultural University

Research Article

Keywords: virtual knockout, benchmark design, pig, RNA-seq, transcript mapping, response-scale assessment, gene ranking, cross-species genomics

Posted Date: September 10th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10564610/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Background

Human virtual-knockout models could extend perturbation-informed gene prioritisation to livestock, but their evaluation depends on experimental response labels that remain stable across quantification methods and samples. We assessed whether two porcine antisense oligonucleotide RNA-seq studies support whole-transcriptome scoring, leading-gene ranking or neither.

Results

We analysed two porcine satellite-cell antisense oligonucleotide studies (three libraries per treatment-time group). In SRP186451, STAR genome alignment was high (94.03-97.02%), whereas Salmon transcript mapping ranged from 28.57% to 65.29% and tracked PC1 (r = 0.965). Transcript- and genome-based quantification yielded concordant treatment effects. Stringent DESeq2 responses comprised 7,835 of 17,053 evaluable genes at 24 h and 7,859 at 36 h; edgeR TMM/TMMwsp returned comparable fractions (46.3-47.5%). Treatment and the STAR-Salmon mapping gap were nearly inseparable (maximum variance inflation factor = 522.88), precluding whole-response attribution. Leading ranks persisted across time. Among 13,719 one-to-one human-pig orthologues, direct ASO24-to-ASO36 Wald-rank transfer yielded a Top100 overlap of 0.570 and AUROC of 0.880. A chromosome-blocked source-only ridge model yielded 0.560 and 0.880; adding target-time-excluded features increased them to 0.590 and 0.893. In the locked PRJNA575967 stress test, mapping was balanced, time dominated PC1 and six of nine criteria passed. Whole-response scoring remained unsupported because global effect ordering varied across estimators. Leading-rank scoring was supported by leading-gene identity, sample-omission stability and recovery beyond random-gene and negative-control pseudo-contrasts (Top100 overlap = 0.28; AUROC = 0.629), although label-partition support was metric-dependent. Across both studies, leading-rank scoring was the broadest supported use. Contractile/myogenic and proliferation modules provided biological context, whereas QTL proximity remained within chromosome-stratified expectation.

Conclusions

The appropriate benchmark target depends on the response scale supported by estimator agreement, sample stability and null-calibrated rank persistence. In these studies, leading-gene ranks were more reproducible than whole-transcriptome effects and therefore provided the defensible target for future comparisons of human-derived models.

Background

Perturbation models are moving gene prioritisation from association towards intervention. CRISPR screens and Perturb-seq atlases now map transcriptional responses to targeted interventions [1–7], while models trained on these compendia extend prediction to unseen single and combinatorial perturbations [8–11]. Virtual-cell and foundation-model programmes seek to extend this logic across cell

types and molecular contexts [12–14]. Extending this capability to livestock could narrow the search space for functional studies of health, production and breeding traits, where experimental validation is costly and perturbation datasets are limited.

Benchmark quality has not advanced at the same pace as model complexity. A prediction score reflects both model error and the fidelity of the measured target. Recent evaluations have shown that simple baselines, systematic variation, task definition and metric choice can alter apparent model rankings [15–17]. A benchmark can therefore become misleading when a response vector changes across quantification routes, depends strongly on a small number of libraries or contains numerically unstable leading ranks. Transcriptome-wide analyses of perturbation atlases further show that the definition of the measured response is itself part of the benchmarking problem [18].

This issue is acute in livestock genomics. Pig reference genomes, expression atlases, single-cell resources and trait annotations have expanded rapidly [19–26]. The controlled porcine perturbation series assembled for this analysis comprised two satellite-cell ASO transcriptomes [26–30]. When such a small evidence base supplies benchmark targets for downstream models, read-level quality control alone does not determine whether an experiment supports genome-wide effect scoring, only a leading-gene ranking, or neither. These are distinct evaluation targets with different requirements for estimator agreement and reproducibility.

The available controlled porcine perturbation datasets are antisense oligonucleotide transcript-depletion studies rather than genetic-knockout experiments [27–28, 31]. They nevertheless provide measured intervention responses with which to test whether a model-evaluation target is reproducible; no equivalence between transient depletion and genetic knockout is assumed.

We treated construction of the experimental benchmark target as part of model evaluation. SRP186451 served as the primary study for assessing mapping structure, estimator agreement, sample influence, source-time persistence and label-based controls; the same criteria were applied under a prespecified plan to PRJNA575967. We asked whether each experiment supported whole-response scoring, leading-rank scoring or neither. We also tested whether target-time-excluded features improved on source-time persistence and whether response-stable genes formed coherent biological modules.

Methods

Datasets, references and quantification

We analysed two public *Sus scrofa* satellite-cell antisense oligonucleotide (ASO) RNA-seq studies. SRP186451/PRJNA523512 comprised vehicle and TCONS_00323213 ASO conditions at 24 and 36 h, and PRJNA575967 comprised negative-control and MSTRG.59589 ASO conditions at 30 and 40 h [27–30]. Each treatment–time group contained three paired-end libraries. All archived runs assigned to these groups were included. The sequencing library was the analysis unit, and ASO was contrasted with its matched control at each time point. Donor identity, cell-isolation batch, pairing and randomisation were

unavailable, so no blocking term was fitted. SRP186451 was the primary study; PRJNA575967 was analysed as a prespecified external stress test (Table 1).

Table 1
Datasets, gene universes and analysis roles.

Component	Dataset or resource	Scale	Use
Primary study	SRP186451 / TCONS_00323213 ASO	12 paired-end libraries; vehicle/ASO at 24 and 36 h	Estimator sensitivity and cross-time rank assessment
External stress test	PRJNA575967 / MSTRG.59589 ASO	12 paired-end libraries; control/ASO at 30 and 40 h	Prespecified application of response-scale criteria
ASO target evidence	Source-study qRT-PCR	TCONS_00323213 and MSTRG.59589; absent from Ensembl release 115	Target-engagement evidence
Reference and direct-rank universe	Sscrofa11.1 and Ensembl release 115	17055 coordinate-mapped genes shared across SRP186451 time contrasts	Quantification and direct-rank sensitivity
Cross-species-compatible universe	Ensembl human–pig orthology	13719 high-confidence one-to-one orthologues	Chromosome-blocked recovery and future model comparison
Functional resources	GO, Reactome, Bgee and STRING	Pathways, tissue expression and protein associations	Target-time-excluded features and post-ranking interpretation
Trait resource	Animal QTLdb release 59	Pig QTL and association records	Chromosome-calibrated trait context
Recovery models	Direct transfer, ridge, random forest and gradient boosting	Source-time and target-time-excluded features	Cross-time ranking; source-only performance defines the comparison floor

FASTQ byte sizes and compressed-stream integrity were checked against archive manifests, and read quality was assessed with FastQC and MultiQC [32–33]. Reads were analysed without trimming against the Sscrofa11.1 genome and Ensembl release 115 annotation [34–35]. target-engagement evidence was taken from source-study qRT-PCR [27–28]. Reference manifests and sample-level quality-control measures are provided in Additional file 1.

Transcript abundance was quantified with Salmon and imported through tximport as length-scaled gene counts [36–37]. SRP186451 was processed with both the original cDNA-only index and a genome-decoy index; PRJNA575967 used the decoy-aware index. Selective alignment applied mapping validation and sequence- and GC-bias correction [38]. In parallel, STAR aligned reads to the genome and generated stranded gene counts [39]. Salmon-inferred library orientation determined the STAR count column. Sample summaries included fragment counts, mapping rates, the STAR–Salmon mapping gap and

fragment-length estimates. Principal-component analysis used variance-stabilised counts, and mapping associations were evaluated by correlation and within-time treatment differences.

Differential expression and response stability

DESeq2 used the factorial design ~ time + treatment + time:treatment [40]. Vehicle 24 h and negative-control 30 h were the respective reference groups. Genes with at least 10 counts in three or more libraries were retained. Contrasts comprised ASO versus matched control at each time point and the treatment-by-time interaction. Stringent differential expression required a Benjamini–Hochberg-adjusted $P < 0.05$ and |log2 fold change| ≥ 1 [41].

For SRP186451, edgeR quasi-likelihood models with TMM and TMMwsp normalisation used the same design and filter [42–44]. Agreement with DESeq2 was assessed from effect correlations, stringent-gene fractions and Top100 overlap. Mapping-axis sensitivity was examined by adding one standardised mapping variable at a time to the factorial model. The tested variables were Salmon mapping, STAR total mapping and the STAR–Salmon gap for SRP186451, and the STAR–Salmon gap for PRJNA575967. These models were used only to assess coefficient sensitivity. Design rank, condition number and variance inflation factors (VIFs) were calculated; models with a maximum VIF ≥ 10 were classified as non-evaluable [45].

Latent structure in SRP186451 was assessed with num.sv and svaseq, followed by a secondary DESeq2 model containing the inferred surrogate variable [46–50]. Sample influence was evaluated by omitting each library and refitting every quantification and differential-expression branch. Full-data and leave-one-library-out estimates were compared by log2 fold-change correlation, Top100 identity overlap and cross-time ranking performance.

Cross-time rank evaluation

Direct rank persistence used coordinate-mapped genes shared by the two SRP186451 contrasts; FDR-based summaries additionally required defined *P* and adjusted *P* values. Cross-species model analyses used 13,719 high-confidence one-to-one human–pig orthologues obtained from Ensembl release 115 through biomaRt [34–35, 51–52]. Genes were ranked primarily by the absolute DESeq2 Wald statistic. A signed FDR score, sign(log2 fold change) × −log10(adjusted $P + 10^{-300}$), was analysed secondarily; when a tie crossed the Top100 boundary, the admissible overlap range was reported. Cross-time comparisons were ASO24 to ASO36 in SRP186451 and ASO30 to ASO40 in PRJNA575967.

Direct transfer and a source-only ridge model defined the temporal baseline. Ridge regression, random forest and gradient boosting then combined source-time response with target-time-excluded STRING, Bgee, GO and Reactome features [53–56]. Target-time and interaction-derived variables were excluded. Numeric missing values were set to zero after table merging. Ridge used an L2 penalty of 1.0 after within-fold standardisation; random forest and gradient boosting each used 180 estimators, with a

minimum leaf size of eight and maximum depth of three, respectively. No hyperparameter search was performed. Full feature definitions, seeds and settings are provided in Additional file 1.

Chromosome-group five-fold cross-validation was primary, with random five-fold cross-validation secondary [57]. Top100 overlap and Top10% AUROC were the primary metrics; AUPR, Spearman correlation, mean squared error and alternative top-*k* overlaps were secondary. Controls comprised annotation-only models, 2,000 random-gene rankings, control-library pseudo-contrasts and all 3-versus-3 label partitions within each time point [58]. Stability was assessed with 2,000 paired gene resamples, leave-one-canonical-chromosome-out analyses across 21 chromosomes and leave-one-library-out refitting. Exceedance fractions used the plus-one correction.

External response-scale assessment and biological context

Nine criteria and their thresholds were fixed before PRJNA575967 quantification. They covered genome alignment, treatment-associated mapping differences, PC1 mapping dominance, cross-quantifier agreement, mapping-axis sensitivity, stringent-gene retention, source-reported direction checks and random-rank calibration. Eighteen gene–time direction checks were derived from PALLD and source-reported qRT-PCR or myogenic panels. Whole-response scoring required agreement across quantifiers and mapping-sensitivity analyses in an evaluable design. Leading-rank scoring required concordant leading-gene identity and direction, stability to library omission and recovery beyond random-gene and negative-control pseudo-contrast distributions. The same evidence domains were summarised retrospectively for SRP186451. The criteria and thresholds are listed in Additional file 1; their application is reported in Table 3.

Table 3
Predefined evaluation criteria for PRJNA575967.

Criterion	Threshold	Observed	Outcome
Genome alignment	Mean STAR unique mapping ≥ 90%	94.70%	Pass
Treatment–mapping difference	Maximum within-time difference < 5 percentage points	0.861 percentage points	Pass
PC1 mapping dominance	Maximum absolute PC1–mapping correlation < 0.80	0.496	Pass (point-estimate rule)
Cross-quantifier rank agreement	Minimum Salmon–STAR Spearman ≥ 0.90	0.843	Fail
Mapping-sensitivity rank agreement	Minimum baseline–gap Spearman ≥ 0.90	0.651	Fail
Stringent DEG retention	Minimum retention ≥ 0.50	0.4972 (89/179)	Fail
Direction-check coverage	At least 14 of 18 checks evaluable	Salmon 18/18; STAR 17/18	Pass
Direction concordance	Minimum concordance ≥ 0.70	Salmon 15/18 (0.833); STAR 16/17 (0.941)	Pass
Random-rank calibration	Top100 exceedance fraction ≤ 0.05	0.00050	Pass

Biological interpretation used Gene Ontology, Reactome, Bgee, STRING, Ensembl orthology and Animal QTLdb [34–35, 51–52, 59–65]. Genes with concordant cross-time direction were grouped by functional annotation, tissue expression and STRING connectivity; network edges required a combined score ≥ 400. QTL support comprised interval overlap or proximity within 1 Mb. Observed counts were compared with 1,000 chromosome-stratified random gene sets.

Statistical analysis and reproducibility

Technical associations used Pearson or Spearman correlation. Pearson correlations were accompanied by Fisher *z* 95% confidence intervals. Differential-expression *P* values were adjusted within contrast by the Benjamini–Hochberg method. Empirical tail fractions were calculated for random-gene, label-partition and chromosome-stratified controls; tests were two-sided unless a directional tail was specified. Analyses used R 4.2.1, DESeq2 1.38.3, tximport 1.26.1, FastQC 0.12.1, MultiQC 1.10.1, Salmon 1.4.0, STAR 2.7.8a and scikit-learn 1.8.0. The normalisation analysis used R 4.6.0 and edgeR 4.9.9. Complete software versions, parameters, reference checksums and score definitions are provided in Additional file 1.

Results

SRP186451 contains a broad response coupled to transcript mapping

STAR aligned 94.03–97.02% of reads from all 12 libraries to the pig genome. Salmon transcript mapping was more variable (28.57–65.29%) and averaged 26.85 percentage points lower in ASO than vehicle libraries, whereas the corresponding STAR difference was 1.85 points. Salmon mapping closely tracked PC1 (r = 0.965; 95% CI, 0.878–0.990; $P = 3.73 \times 10^{-7}$) (Fig. 1a, b).

Treatment effects were reproducible across quantification and normalisation routes. Salmon–STAR log2 fold-change correlations were 0.958 at 24 h and 0.964 at 36 h; original and decoy-aware Salmon effects correlated by 0.980 at both time points (Fig. 2c). Adding genome decoys changed mean Salmon mapping from 50.38% to 48.58% but retained a 28.02-point ASO–vehicle difference (Fig. 2d). Salmon mapping correlated weakly with read depth and GC content but strongly with the deduplicated fraction (r = − 0.942). Of 17,055 coordinate-mapped genes, two lacked defined DESeq2 P values and were excluded only from FDR-based denominators. DESeq2 identified 7,835 of 17,053 evaluable genes at 24 h and 7,859 at 36 h; edgeR returned comparable stringent fractions (46.27–47.49%) and nearly identical effect estimates (Spearman ρ = 0.995–1.000) (Fig. 2a, b; Additional file 1).

Including the STAR–Salmon gap reduced stringent-gene counts across all three quantification branches, from baseline ranges of 6,402–13,431 genes to 15–378 genes (Table 2; Fig. 1c, d; Additional file 2: Fig. S2). Treatment and the gap were nearly collinear (maximum VIF = 522.88), and the Salmon-mapping model was similarly ill conditioned (maximum VIF = 681.13). The leading surrogate variable also aligned with treatment and mapping, but its inclusion retained thousands of stringent genes. The broad treatment-associated response was therefore reproducible, whereas its whole-transcriptome attribution was not separable from the mapping axis in this design.

Table 2
Baseline differential expression and mapping-gap projection in SRP186451.

Quantifier	Contrast	FDR-evaluable genes	Baseline stringent genes, n (%)	Gap-projection stringent genes, n (%)	Baseline–projection log2FC Spearman ρ	Design diagnostics (max VIF; condition no.)
Original Salmon	ASO24 vs vehicle 24 h	17053	7835 (45.95)	15 (0.09)	0.305	522.88; 108.17
Original Salmon	ASO36 vs vehicle 36 h	17053	7859 (46.09)	309 (1.81)	0.528	522.88; 108.17
STAR GeneCounts	ASO24 vs vehicle 24 h	23701	13431 (56.67)	16 (0.07)	0.370	522.88; 108.17
STAR GeneCounts	ASO36 vs vehicle 36 h	23701	12664 (53.43)	378 (1.59)	0.549	522.88; 108.17
Decoy-aware Salmon	ASO24 vs vehicle 24 h	16035	6402 (39.92)	27 (0.17)	0.284	531.93; 107.87
Decoy-aware Salmon	ASO36 vs vehicle 36 h	16035	6694 (41.75)	311 (1.94)	0.503	531.93; 107.87

Leading ranks persist across time in SRP186451

Across 17,055 coordinate-mapped genes, direct ASO24-to-ASO36 transfer produced a rank correlation of 0.833, Top100 overlap of 0.49, AUROC of 0.884 and AUPR of 0.652. Within the 13,719 one-to-one orthologues, direct transfer yielded a Top100 overlap of 0.570 and AUROC of 0.880.

The source-only ridge model reached a Top100 overlap of 0.560 and AUROC of 0.880 under chromosome-group validation (Fig. 3a, b). Adding target-time-excluded features increased these metrics to 0.590 and 0.893. The AUROC increment was 0.0125 across paired gene resamples (2.5th–97.5th percentiles, 0.00694–0.01858) and 0.0103–0.0139 across chromosome deletions (Fig. 3c). Random forest gave the highest Top100 overlap (0.600) but an AUROC of 0.885; annotation-only models remained near chance (Top100 overlap = 0.160; AUROC ≈ 0.502). Most recoverable structure was therefore already present at the source time point.

The FDR-score sensitivity analysis gave similar AUROC increments but created 103 ties at the Top100 boundary; admissible overlaps were 0.54–0.57 for source-only ridge and 0.57–0.60 for extended ridge (Additional file 2: Fig. S3). The observed Wald-rank metrics exceeded all 2,000 random-gene rankings (exceedance fraction = 0.00050) (Fig. 3d). Vehicle abundance ranks were highly stable across time, whereas vehicle pseudo-contrasts had a mean Top100 response overlap of 0.132. Leave-one-library-out refits retained log2 fold-change correlations ≥ 0.991, Top100 overlaps of 0.42–0.49 and AUROCs of 0.843–0.881 (Fig. 5a, b; Additional file 2: Fig. S1). Label-partition support varied by metric, but the combined controls supported leading-rank evaluation with source-time transfer as the comparison floor (Table 4).

Table 4
Evidence-to-decision mapping for virtual-knockout benchmark use.

Dataset	Response geometry	Whole-response decision	Leading-rank decision	Recommended benchmark use
SRP186451	Three shared criteria passed, two failed and four were not evaluable; treatment was nearly collinear with transcript mapping (maximum VIF 522.88)	Non-evaluable under the archived design	Supported within the 13719-gene one-to-one universe	Evaluate leading ranks; require source-time transfer as the performance floor
PRJNA575967	Six of nine fixed criteria passed in a well-conditioned design; weaker global effects reordered across estimators	Unsupported by the fixed global criteria	Supported by identity, direction, omission and random-gene/pseudo-contrast controls	Use as an external leading-rank stress test, not as a whole-response target

External stress testing in PRJNA575967

PRJNA575967 showed a different response geometry. Mean STAR unique mapping was 94.70%, and the largest within-time ASO–control mapping difference was 0.861 percentage points. PC1 explained 67.9% of expression variance and tracked time (r = 0.992; 95% CI, 0.971–0.998), whereas its largest absolute correlation with mapping was 0.496 (Fig. 4b; Additional file 2: Fig. S4).

Global effect ordering was less stable than leading-gene identity. Salmon–STAR log2 fold-change correlations were 0.843 at 30 h and 0.884 at 40 h, but Top100 overlaps were 0.89 and 0.84 with complete directional agreement among shared genes (Fig. 4c). Mapping-gap projection retained 89 of 179 stringent genes at 30 h and remained estimable (maximum VIF = 3.47). Source-study direction concordance was 15/18 for Salmon and 16/17 for STAR. Across 48 leave-one-library-out refits, log2 fold-change correlations exceeded 0.900 and Top100 overlap remained ≥ 0.82 (Fig. 5c, d).

ASO30-to-ASO40 effects had little genome-wide rank correlation (ρ = 0.034) but retained a Top100 overlap of 0.28, directional concordance of 27/28 among overlapping genes and an AUROC of 0.629 (Fig. 4c, d). Both ranking metrics exceeded all 2,000 random-gene draws and the maxima of nine negative-control pseudo-contrasts. AUROC ranged from 0.612 to 0.646 across gene resamples and from 0.626 to 0.631 across chromosome deletions. Six of nine prespecified criteria passed (Fig. 4a; Table 3). Failures involved global rank agreement and mapping-axis sensitivity, whereas leading-gene identity, direction and sample-omission stability were retained. PRJNA575967 therefore supported leading-rank evaluation but not whole-response scoring (Table 4).

Biological modules and trait context

Cross-time direction, network connectivity and functional annotation identified two SRP186451 modules (Fig. 6a–c). TNNT3, MYH3, TPM1 and MEF2C formed a concordantly downregulated contractile/myogenic module [66–73], whereas CDCA8, CDC20, UBE2C and ESPL1 formed a proliferation-related module with stronger induction at 36 h [59–62]. Nineteen Top100 genes overlapped priority-trait QTL intervals, close to the chromosome-stratified null mean of 17.35 (tail fraction = 0.354); proximity within 1 Mb was also within expectation (tail fraction = 0.843) (Additional file 2: Fig. S5). The annotations organised stable ranks into biological hypotheses without determining the ranking.

Discussion

Perturbation benchmarks commonly score models against a fixed experimental label, although apparent performance can change with the baseline, task definition, systematic variation and metric choice [15–17]. Our analysis places one decision upstream: whether the measured response supports whole-transcriptome comparison or only leading-gene ranking. This distinction is especially consequential in livestock, where a small number of public perturbation experiments may serve as reference targets for many subsequent models.

SRP186451 illustrates why response breadth and response estimability must be considered separately. Its broad treatment response was reproduced by DESeq2, edgeR, transcript-based quantification, genome-based counting and a decoy-aware index. Such breadth is biologically plausible when an ASO perturbs a differentiation-associated transcript [27, 70–73]. Treatment and transcript mapping, however, occupied almost the same sample direction. The gap-projection collapse is therefore informative as an estimability diagnostic: near-collinearity prevents attribution of the shared axis to technical or biological origins. Accordingly, the baseline response is reproducible, whereas whole-response attribution remains non-evaluable under the archived design.

PRJNA575967 provided an independent test of the prespecified criteria. Treatment-associated mapping differences were small, PC1 was dominated by time, and conditional models were well conditioned, yet weaker global effects were reordered between transcript- and genome-based estimators. Because the studies targeted different transcripts and used different time grids, PRJNA575967 tests the portability of the response-scale assessment across an independent experimental context. It did not meet the fixed

criteria for whole-response scoring, but leading-gene identity, direction, omission stability and null-calibrated cross-time persistence supported leading-rank use. The two datasets therefore exposed different constraints: non-evaluable attribution in SRP186451 and estimator-sensitive whole-response ordering in PRJNA575967.

This response-level assessment complements conventional RNA-seq quality control. FastQC, MultiQC and alignment summaries diagnose read quality and mapping performance [32–33, 36, 39], but they do not establish whether a perturbation vector is stable across estimators, whether a small number of libraries controls the leading ranks or whether a model merely reproduces source-time ordering. Linking these questions to a choice of evaluation scale turns quality control from a library-level screen into a benchmark-design decision.

The resulting hierarchy also changes how virtual-knockout models should be compared. Whole-response correlation, genome-wide ranking and leading-gene overlap are not interchangeable endpoints. In SRP186451, a one-feature source-only ridge model already achieved a Top100 overlap of 0.560 and AUROC of 0.880; extended models increased these values only to 0.600 and 0.893. Accordingly, complex architectures must exceed source-time persistence rather than only a random predictor, and their incremental value should be reported separately at genome-wide and leading-rank scales [15–17].

The primary Wald analysis further identified a practical rank-definition problem. FDR-derived scores can saturate when adjusted *P* values approach numerical zero, creating many genes tied at an arbitrary cap. Deterministic tie breaking can then alter Top100 overlap without changing the underlying evidence. Continuous Wald statistics removed the target-boundary tie in the primary analysis, while tie-aware ranges showed that the FDR-based conclusion remained qualitatively stable. Explicit reporting of the ranking variable, evaluated gene universe and boundary-tie sensitivity provides a practical safeguard for perturbation benchmarks.

External knowledge was most informative after the response scale had been established. Annotation-only models were near chance, indicating that annotation density did not reconstruct the perturbation ranking. STRING, pathway, tissue-expression and QTL resources instead organised stable genes into testable hypotheses. In SRP186451, concordant reduction of the TNNT3–MYH3–TPM1–MEF2C contractile/myogenic module is consistent with reduced contractile maturation [66–73]. The annotation-defined CDCA8–CDC20–UBE2C–ESPL1 module suggests retention of a proliferation-related programme later in the response [59–62]. These SRP186451 modules define focused follow-up hypotheses. QTL proximity remained within chromosome-matched expectation, showing that trait annotations provided context without generating the priorities.

The findings yield practical requirements for future livestock perturbation studies. Treatment and control libraries should be balanced across donor, cell preparation, extraction batch, library preparation and sequencing lane. RNA integrity, target-depletion efficiency and transcript composition should be recorded alongside conventional sequencing metrics. Benchmark gene universes, quantification branches, source-only baselines and primary rank definitions should be specified before model

comparison. These choices would widen the response scale that can be interpreted and make cross-study model rankings more comparable.

The analysis remains bounded by the archived experiments. Each treatment–time cell contained three libraries, donor identity and batch allocation were incompletely reported, and both quantification routes used the same genome build and annotation. The targeted transcripts were absent from Ensembl release 115, so RNA-seq did not directly measure ASO depletion. The one-to-one orthologue universe excludes duplicated and lineage-specific genes, and transient ASO depletion does not reproduce every feature of genetic knockout. Additional species, tissues and perturbation classes will be needed to calibrate general decision thresholds and to score human-derived virtual-knockout models directly.

Conclusions

Across two porcine ASO studies, whole-transcriptome response labels were either non-evaluable or sensitive to the expression estimator, whereas leading-gene ranks were reproducible across time and controls. Source-time transfer in SRP186451 also defined the performance floor that more complex models must exceed. Matching model evaluation to the response scale supported by the experiment provides a practical basis for future direct human-to-livestock virtual-knockout comparisons.

Abbreviations

ASO
antisense oligonucleotide
AUPR
area under the precision–recall curve
AUROC
area under the receiver operating characteristic curve
DEG
differentially expressed gene
FDR
false discovery rate
GO
Gene Ontology
NC
negative or vehicle control
PCA
principal-component analysis
qRT-PCR
quantitative reverse-transcription polymerase chain reaction
QTL
quantitative trait locus

VIF
variance inflation factor
VST
variance-stabilising transformation.

Declarations

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Availability of data and materials

All data generated or analysed during this study are included in this article and its additional files.

Competing interests

The authors declare that they have no competing interests.

Funding

This work was supported by the Yunnan Major Science and Technology Special Project: Breeding of High-Yield Dairy Goat Hybrid Lines and Research & Application of Efficient Breeding Technologies (Project No. 202602AE090078).

Authors' contributions

XZ conceived the study, developed the analysis design, performed the formal analyses and visualisation, and drafted the manuscript. HC contributed to study design, validation, interpretation and critical revision of the manuscript. ZZ contributed to data curation, reproducibility checks and critical revision of the manuscript. XX supervised the study, administered the project, acquired funding and critically revised the manuscript. All authors read and approved the final manuscript.

Acknowledgements

Not applicable.

References

1. Adamson B, Norman TM, Jost M, Cho MY, Nuñez JK, Chen Y, et al. A Multiplexed Single-Cell CRISPR Screening Platform Enables Systematic Dissection of the Unfolded Protein Response. Cell.

2016;167(7):1867–e188221. 10.1016/j.cell.2016.11.048.

2. Dixit A, Parnas O, Li B, Chen J, Fulco CP, Jerby-Arnon L, et al. Perturb-Seq: Dissecting Molecular Circuits with Scalable Single-Cell RNA Profiling of Pooled Genetic Screens. Cell. 2016;167(7):1853–e186617. 10.1016/j.cell.2016.11.038.
3. Datlinger P, Rendeiro AF, Schmidl C, Krausgruber T, Traxler P, Klughammer J, et al. Pooled CRISPR screening with single-cell transcriptome readout. Nat Methods. 2017;14(3):297–301. 10.1038/nmeth.4177.
4. Norman TM, Horlbeck MA, Replogle JM, Ge AY, Xu A, Jost M, et al. Exploring genetic interaction manifolds constructed from rich single-cell phenotypes. Science. 2019;365(6455):786–93. 10.1126/science.aax4438.
5. Replogle JM, Saunders RA, Pogson AN, Hussmann JA, Lenail A, Guna A, et al. Mapping information-rich genotype-phenotype landscapes with genome-scale Perturb-seq. Cell. 2022;185(14):2559–e257528. 10.1016/j.cell.2022.05.013.
6. Peidli S, Green TD, Shen C, Gross T, Min J, Garda S, et al. scPerturb: harmonized single-cell perturbation data. Nat Methods. 2024;21(3):531–40. 10.1038/s41592-023-02144-y.
7. Santinha AJ, Klingler E, Kuhn M, Farouni R, Lagler S, Kalamakis G, et al. Transcriptional linkage analysis with in vivo AAV-Perturb-seq. Nature. 2023;622(7982):367–75. 10.1038/s41586-023-06570-y.
8. Roohani Y, Huang K, Leskovec J. Predicting transcriptional outcomes of novel multigene perturbations with GEARS. Nat Biotechnol. 2024;42(6):927–35. 10.1038/s41587-023-01905-6.
9. Lotfollahi M, Wolf FA, Theis FJ. scGen predicts single-cell perturbation responses. Nat Methods. 2019;16(8):715–21. 10.1038/s41592-019-0494-8.
10. Lotfollahi M, Klimovskaia Susmelj A, De Donno C, Hetzel L, Ji Y, Ibarra IL, et al. Predicting cellular responses to complex perturbations in high-throughput screens. Mol Syst Biol. 2023;19(6):e11517. 10.15252/msb.202211517.
11. Wenkel F, Tu W, Masschelein C, Shirzad H, Hodgson L, Bendidi I, et al. TxPert: using multiple knowledge graphs for prediction of transcriptomic perturbation effects. Nat Biotechnol. 2026. 10.1038/s41587-026-03113-4.
12. Bunne C, Roohani Y, Rosen Y, Gupta A, Zhang X, Roed M, et al. How to build the virtual cell with artificial intelligence: Priorities and opportunities. Cell. 2024;187(25):7045–63. 10.1016/j.cell.2024.11.015.
13. Rood JE, Hupalowska A, Regev A. Toward a foundation model of causal cell and tissue biology with a Perturbation Cell and Tissue Atlas. Cell. 2024;187(17):4520–45. 10.1016/j.cell.2024.07.035.
14. Szałata A, Hrovatin K, Becker S, Tejada-Lapuerta A, Cui H, Wang B, et al. Transformers in single-cell omics: a review and new perspectives. Nat Methods. 2024;21(8):1430–43. 10.1038/s41592-024-02353-z.
15. Ahlmann-Eltze C, Huber W, Anders S. Deep-learning-based gene perturbation effect prediction does not yet outperform simple linear baselines. Nat Methods. 2025;22(8):1657–61. 10.1038/s41592-

025-02772-6.

16. Viñas Torné R, Wiatrak M, Piran Z, Fan S, Jiang L, Teichmann SA, et al. Systema: a framework for evaluating genetic perturbation response prediction beyond systematic variation. Nat Biotechnol. 2026;44(6):1050–9. 10.1038/s41587-025-02777-8.
17. Wei Z, Wang Y, Gao Y, Wang S, Li P, Si D, et al. Benchmarking algorithms for generalizable single-cell perturbation response prediction. Nat Methods. 2026;23(2):451–64. 10.1038/s41592-025-02980-0.
18. Nadig A, Replogle JM, Pogson AN, Murthy M, McCarroll SA, Weissman JS, et al. Transcriptome-wide analysis of differential expression in perturbation atlases. Nat Genet. 2025;57(5):1228–37. 10.1038/s41588-025-02169-3.
19. Groenen MAM, Archibald AL, Uenishi H, Tuggle CK, Takeuchi Y, Rothschild MF, et al. Analyses of pig genomes provide insight into porcine demography and evolution. Nature. 2012;491(7424):393–8. 10.1038/nature11622.
20. Warr A, Affara N, Aken B, Beiki H, Bickhart DM, Billis K, et al. An improved pig reference genome sequence to enable pig genetics and genomics research. Gigascience. 2020;9(6):giaa051. 10.1093/gigascience/giaa051.
21. Ma L, Meng Y, An Y, Han P, Zhang C, Yue Y, et al. Single-cell RNA-seq reveals novel interaction between muscle satellite cells and fibro-adipogenic progenitors mediated with FGF7 signalling. J Cachexia Sarcopenia Muscle. 2024;15(4):1388–403. 10.1002/jcsm.13484.
22. Cai S, Hu B, Wang X, Liu T, Lin Z, Tong X, et al. Integrative single-cell RNA-seq and ATAC-seq analysis of myogenic differentiation in pig. BMC Biol. 2023;21(1):19. 10.1186/s12915-023-01519-z.
23. Guo L, Han M, Xu J, Zhou W, Shi H, Chen S, et al. snRNA-Seq and Spatial Transcriptome Reveal Cell–Cell Crosstalk Mediated Metabolic Regulation in Porcine Skeletal Muscle. J Cachexia Sarcopenia Muscle. 2025;16(2):e13752. 10.1002/jcsm.13752.
24. Xu Z, Lin Q, Cai X, Zhong Z, Teng J, Li B, et al. Integrating large-scale meta-GWAS and PigGTEx resources to decipher the genetic basis of 232 complex traits in pigs. Natl Sci Rev. 2025;12(5):nwaf048. 10.1093/nsr/nwaf048.
25. Jahuey-Martínez FJ, Martínez-Quintana JA, Rodríguez-Almeida FA, Parra-Bracamonte GM. Exploration and Enrichment Analysis of the QTLome for Important Traits in Livestock Species. Genes (Basel). 2024;15(12):1513. 10.3390/genes15121513.
26. Liu J, Lang K, Tan S, Jie W, Zhu Y, Huang S, et al. A web-based database server using 43,710 public RNA-seq samples for the analysis of gene expression and alternative splicing in livestock animals. BMC Genomics. 2022;23:706. 10.1186/s12864-022-08881-2.
27. Li M, Liu Q, Xie S, Fu C, Li J, Tian C, et al. LncRNA TCONS_00323213 Promotes Myogenic Differentiation by Interacting with PKNOX2 to Upregulate MyoG in Porcine Satellite Cells. Int J Mol Sci. 2023;24(7):6773. 10.3390/ijms24076773.
28. Li L, Cheng X, Chen L, Li J, Luo W, Li C. Long Noncoding Ribonucleic Acid MSTRG.59589 Promotes Porcine Skeletal Muscle Satellite Cells Differentiation by Enhancing the Function of PALLD. Front Genet. 2019;10:1220. 10.3389/fgene.2019.01220.

29. National Center for Biotechnology Information. Sequence Read Archive study SRP186451. https://www.ncbi.nlm.nih.gov/Traces/study/?acc=SRP186451. Accessed 18 July 2026.
30. National Center for Biotechnology Information. BioProject PRJNA575967. https://www.ncbi.nlm.nih.gov/bioproject/PRJNA575967. Accessed 18 July 2026.
31. Crooke ST, Baker BF, Crooke RM, Liang XH. Antisense technology: an overview and prospectus. Nat Rev Drug Discov. 2021;20(6):427–53. 10.1038/s41573-021-00162-z.
32. Ewels P, Magnusson M, Lundin S, Käller M. MultiQC: summarize analysis results for multiple tools and samples in a single report. Bioinformatics. 2016;32(19):3047–8. 10.1093/bioinformatics/btw354.
33. Andrews S. FastQC: a quality control tool for high throughput sequence data. 2010. https://www.bioinformatics.babraham.ac.uk/projects/fastqc/. Accessed 18 July 2026.
34. Cunningham F, Allen JE, Allen J, Alvarez-Jarreta J, Amode MR, Armean IM, et al. Ensembl 2022. Nucleic Acids Res. 2022;50(D1):D988–95. 10.1093/nar/gkab1049.
35. Ensembl. Ensembl release 115 and Ensembl Genomes 62 released. https://www.ebi.ac.uk/about/news/updates-from-data-resources/ensembl-115-and-ensembl-genomes-62-released/. Accessed 24 July 2026.
36. Patro R, Duggal G, Love MI, Irizarry RA, Kingsford C. Salmon provides fast and bias-aware quantification of transcript expression. Nat Methods. 2017;14(4):417–9. 10.1038/nmeth.4197.
37. Soneson C, Love MI, Robinson MD. Differential analyses for RNA-seq: transcript-level estimates improve gene-level inferences. F1000Res. 2016;4:1521. 10.12688/f1000research.7563.2.
38. Srivastava A, Malik L, Sarkar H, Zakeri M, Almodaresi F, Soneson C, et al. Alignment and mapping methodology influence transcript abundance estimation. Genome Biol. 2020;21(1):239. 10.1186/s13059-020-02151-8.
39. Dobin A, Davis CA, Schlesinger F, Drenkow J, Zaleski C, Jha S, et al. STAR: ultrafast universal RNA-seq aligner. Bioinformatics. 2013;29(1):15–21. 10.1093/bioinformatics/bts635.
40. Love MI, Huber W, Anders S. Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2. Genome Biol. 2014;15(12):550. 10.1186/s13059-014-0550-8.
41. Benjamini Y, Hochberg Y. Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. J R Stat Soc Ser B Stat Methodol. 1995;57(1):289–300. 10.1111/j.2517-6161.1995.tb02031.x.
42. Robinson MD, McCarthy DJ, Smyth GK. edgeR: a Bioconductor package for differential expression analysis of digital gene expression data. Bioinformatics. 2010;26(1):139–40. 10.1093/bioinformatics/btp616.
43. Robinson MD, Oshlack A. A scaling normalization method for differential expression analysis of RNA-seq data. Genome Biol. 2010;11(3):R25. 10.1186/gb-2010-11-3-r25.
44. Chen Y, Lun ATL, Smyth GK. From reads to genes to pathways: differential expression analysis of RNA-seq experiments using Rsubread and the edgeR quasi-likelihood pipeline. F1000Res.

2016;5:1438. 10.12688/f1000research.8987.2.

45. Nygaard V, Rødland EA, Hovig E. Methods that remove batch effects while retaining group differences may lead to exaggerated confidence in downstream analyses. Biostatistics. 2016;17(1):29–39. 10.1093/biostatistics/kxv027.
46. Leek JT, Storey JD. Capturing Heterogeneity in Gene Expression Studies by Surrogate Variable Analysis. PLoS Genet. 2007;3(9):e161. 10.1371/journal.pgen.0030161.
47. Leek JT. svaseq: removing batch effects and other unwanted noise from sequencing data. Nucleic Acids Res. 2014;42(21):e161. 10.1093/nar/gku864.
48. Johnson WE, Li C, Rabinovic A. Adjusting batch effects in microarray expression data using empirical Bayes methods. Biostatistics. 2007;8(1):118–27. 10.1093/biostatistics/kxj037.
49. Risso D, Ngai J, Speed TP, Dudoit S. Normalization of RNA-seq data using factor analysis of control genes or samples. Nat Biotechnol. 2014;32(9):896–902. 10.1038/nbt.2931.
50. Stegle O, Parts L, Piipari M, Winn J, Durbin R. Using probabilistic estimation of expression residuals (PEER) to obtain increased power and interpretability of gene expression analyses. Nat Protoc. 2012;7(3):500–7. 10.1038/nprot.2011.457.
51. Vilella AJ, Severin J, Ureta-Vidal A, Heng L, Durbin R, Birney E. EnsemblCompara GeneTrees: Complete, duplication-aware phylogenetic trees in vertebrates. Genome Res. 2009;19(2):327–35. 10.1101/gr.073585.107.
52. Durinck S, Spellman PT, Birney E, Huber W. Mapping identifiers for the integration of genomic datasets with the R/Bioconductor package biomaRt. Nat Protoc. 2009;4(8):1184–91. 10.1038/nprot.2009.97.
53. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O et al. Scikit-learn: machine learning in Python. J Mach Learn Res. 2011;12:2825–2830. https://www.jmlr.org/papers/v12/pedregosa11a.html. Accessed 24 July 2026.
54. Breiman L. Random Forests. Mach Learn. 2001;45(1):5–32. 10.1023/a:1010933404324.
55. Friedman JH. Greedy function approximation: a gradient boosting machine. Ann Stat. 2001;29(5):1189–232. 10.1214/aos/1013203451.
56. Hoerl AE, Kennard RW. Ridge Regression: Biased Estimation for Nonorthogonal Problems. Technometrics. 1970;12(1):55–67. 10.1080/00401706.1970.10488634.
57. Roberts DR, Bahn V, Ciuti S, Boyce MS, Elith J, Guillera-Arroita G, et al. Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. Ecography. 2017;40(8):913–29. 10.1111/ecog.02881.
58. Phipson B, Smyth GK. Permutation P-values should never be zero: calculating exact P-values when permutations are randomly drawn. Stat Appl Genet Mol Biol. 2010;9:Article 39. 10.2202/1544-6115.1585
59. Ashburner M, Ball CA, Blake JA, Botstein D, Butler H, Cherry JM, et al. Gene Ontology: tool for the unification of biology. Nat Genet. 2000;25(1):25–9. 10.1038/75556.

60. Gene Ontology Consortium. The Gene Ontology knowledgebase in 2023. Genetics. 2023;224(1):iyad031. 10.1093/genetics/iyad031.
61. Szklarczyk D, Kirsch R, Koutrouli M, Nastou K, Mehryary F, Hachilif R, et al. The STRING database in 2023: protein–protein association networks and functional enrichment analyses for any sequenced genome of interest. Nucleic Acids Res. 2023;51(D1):D638–46. 10.1093/nar/gkac1000.
62. Milacic M, Beavers D, Conley P, Gong C, Gillespie M, Griss J, et al. The Reactome Pathway Knowledgebase 2024. Nucleic Acids Res. 2024;52(D1):D672–8. 10.1093/nar/gkad1025.
63. Bastian FB, Roux J, Niknejad A, Comte A, Fonseca Costa SS, de Farias TM, et al. The Bgee suite: integrated curated expression atlas and comparative transcriptomics in animals. Nucleic Acids Res. 2021;49(D1):D831–47. 10.1093/nar/gkaa793.
64. Hu ZL, Fritz ER, Reecy JM. AnimalQTLdb: a livestock QTL database tool set for positional QTL information mining and beyond. Nucleic Acids Res. 2007;35(Database issue):D604–9. 10.1093/nar/gkl946.
65. Animal QTLdb. Animal QTLdb release 59. https://www.animalgenome.org/QTLdb/release/59/. Accessed 24 July 2026.
66. Stefancsik R, Randall JD, Mao C, Sarkar S. Structure and sequence of the human fast skeletal troponin T (TNNT3) gene: insight into the evolution of the gene and the origin of the developmentally regulated isoforms. Comp Funct Genomics. 2003;4(6):609–25. 10.1002/cfg.343.
67. Gunning PW, Hardeman EC, Lappalainen P, Mulvihill DP. Tropomyosin: master regulator of actin filament function in the cytoskeleton. J Cell Sci. 2015;128(16):2965–74. 10.1242/jcs.172502.
68. Potthoff MJ, Olson EN. MEF2: a central regulator of diverse developmental programs. Development. 2007;134(23):4131–40. 10.1242/dev.008367.
69. Schiaffino S, Reggiani C. Fiber Types in Mammalian Skeletal Muscles. Physiol Rev. 2011;91(4):1447–531. 10.1152/physrev.00031.2010.
70. Bentzinger CF, Wang YX, Rudnicki MA. Building Muscle: Molecular Regulation of Myogenesis. Cold Spring Harb Perspect Biol. 2012;4(2):a008342. 10.1101/cshperspect.a008342.
71. Yin H, Price F, Rudnicki MA. Satellite Cells and the Muscle Stem Cell Niche. Physiol Rev. 2013;93(1):23–67. 10.1152/physrev.00043.2011.
72. Buckingham M, Rigby PWJ. Gene Regulatory Networks and Transcriptional Mechanisms that Control Myogenesis. Dev Cell. 2014;28(3):225–38. 10.1016/j.devcel.2013.12.020.
73. Mohammadabadi M, Bordbar F, Jensen J, Du M, Guo W. Key Genes Regulating Skeletal Muscle Development and Growth in Farm Animals. Anim (Basel). 2021;11(3):835. 10.3390/ani11030835.

Figures

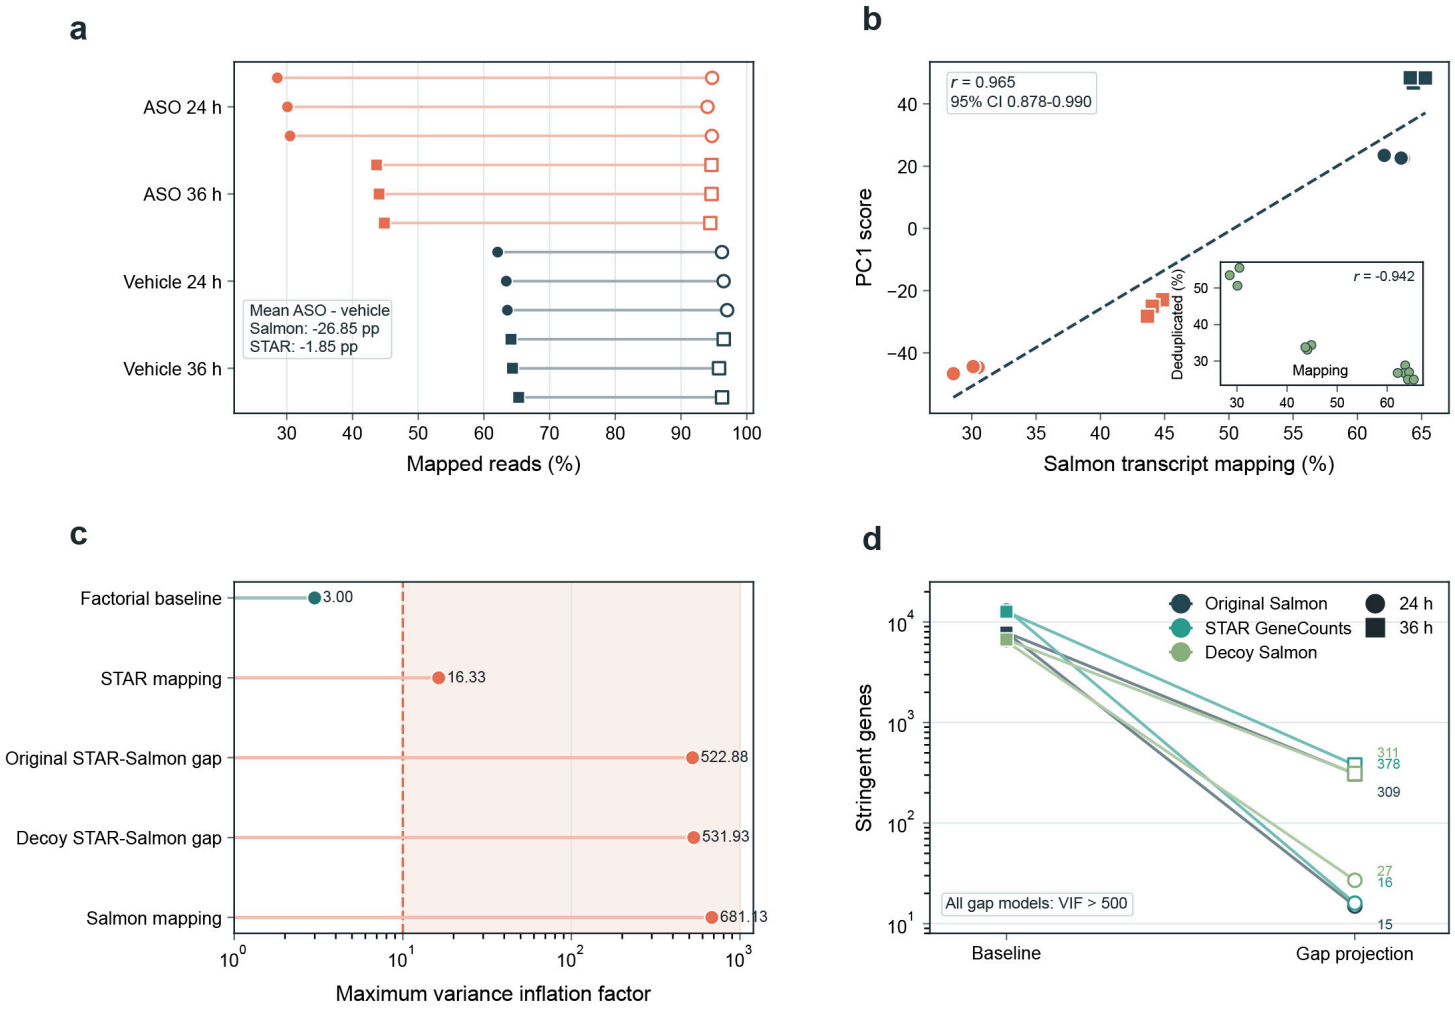


Figure 1

Treatment and transcript mapping are inseparable in SRP186451. a, Library-level dumbbells connect Salmon transcript-mapping rates (filled symbols) with STAR total genome-mapping rates (open symbols) for all 12 libraries. Brick red and deep blue denote ASO and vehicle, respectively; circles and squares denote 24 and 36 h. b, PC1 scores are plotted against Salmon mapping, with a fitted line and Pearson correlation. The inset relates Salmon mapping to the FastQC deduplicated percentage. c, Maximum variance inflation factors are shown on a logarithmic scale for the factorial baseline and four mapping-axis models. Deep green denotes a model below the study-defined VIF threshold of 10; brick red denotes models above the threshold. d, Stringent-gene counts are paired between baseline and STAR–Salmon-gap projection models for original Salmon, STAR GeneCounts and decoy-aware Salmon. Filled and open symbols denote baseline and projected values; circles and squares denote 24 and 36 h. Counts are plotted on a logarithmic scale. Because every gap model had a maximum VIF above 500, panel d is an estimability diagnostic rather than a corrected treatment estimate.

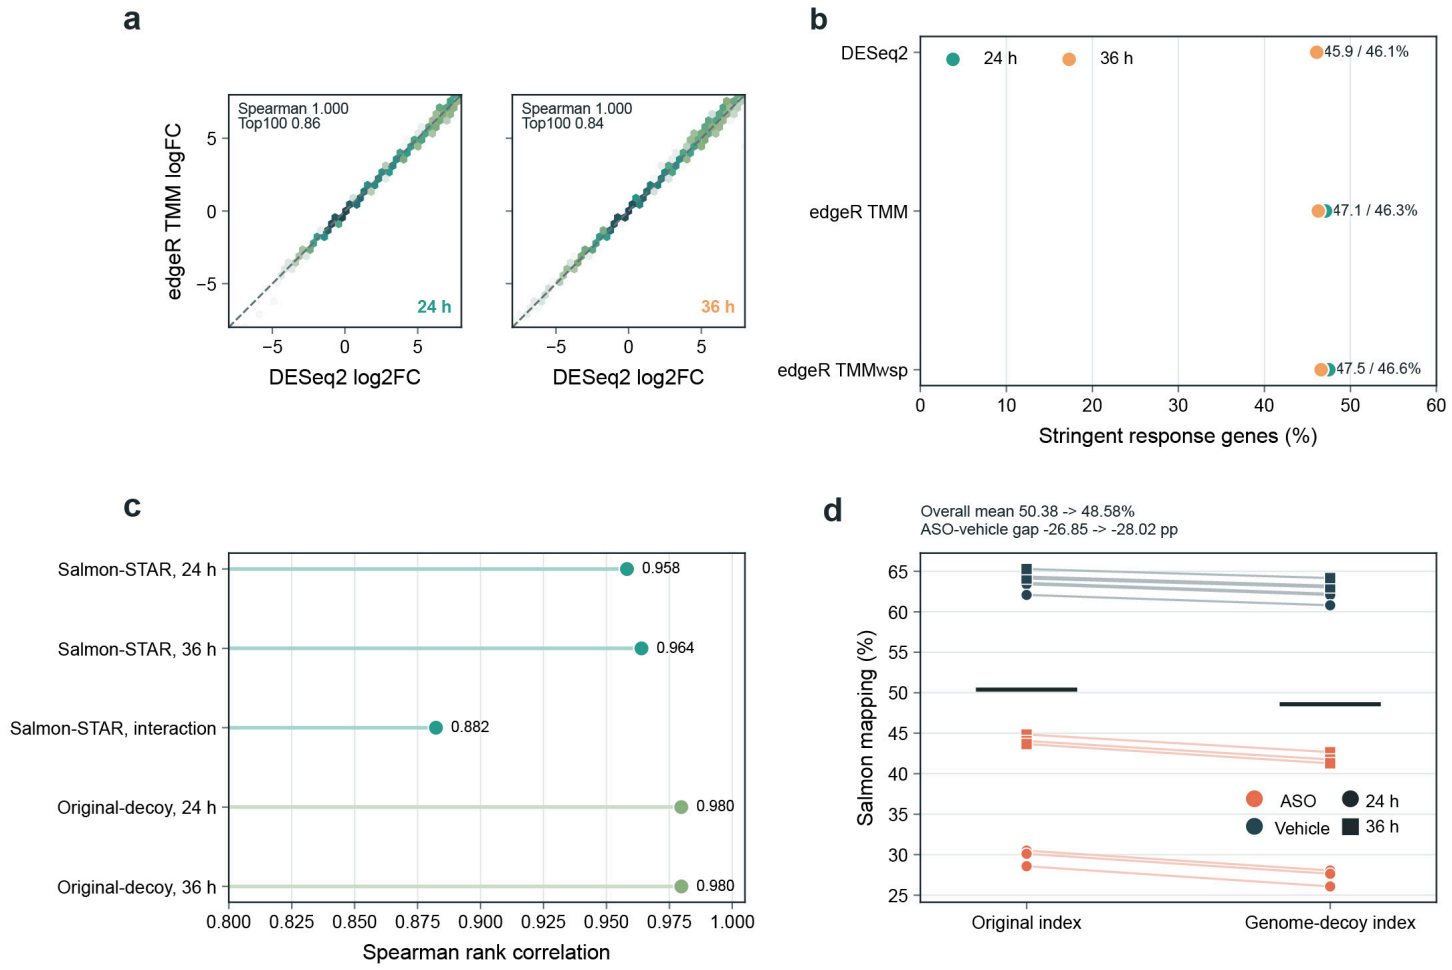


Figure 2

Response breadth is reproducible across estimators and normalisations in SRP186451. a, Hexagonal density plots compare DESeq2 and edgeR TMM log2 fold changes at 24 and 36 h; dashed lines mark identity, and annotations report Spearman correlation and Top100 overlap. b, Paired points show the percentage of stringent response genes at 24 h (teal) and 36 h (light orange) under DESeq2, edgeR TMM and edgeR TMMwsp. The horizontal axis begins at zero. c, Effect-rank agreement is shown across transcript-versus-genome and original-versus-decoy quantification routes. Teal denotes Salmon–STAR comparisons and sage denotes original–decoy comparisons. d, Paired library-level mapping rates compare the original and genome-decoy Salmon indices. Brick red and deep blue denote ASO and vehicle; circles and squares denote 24 and 36 h. Black horizontal segments show study-wide means.

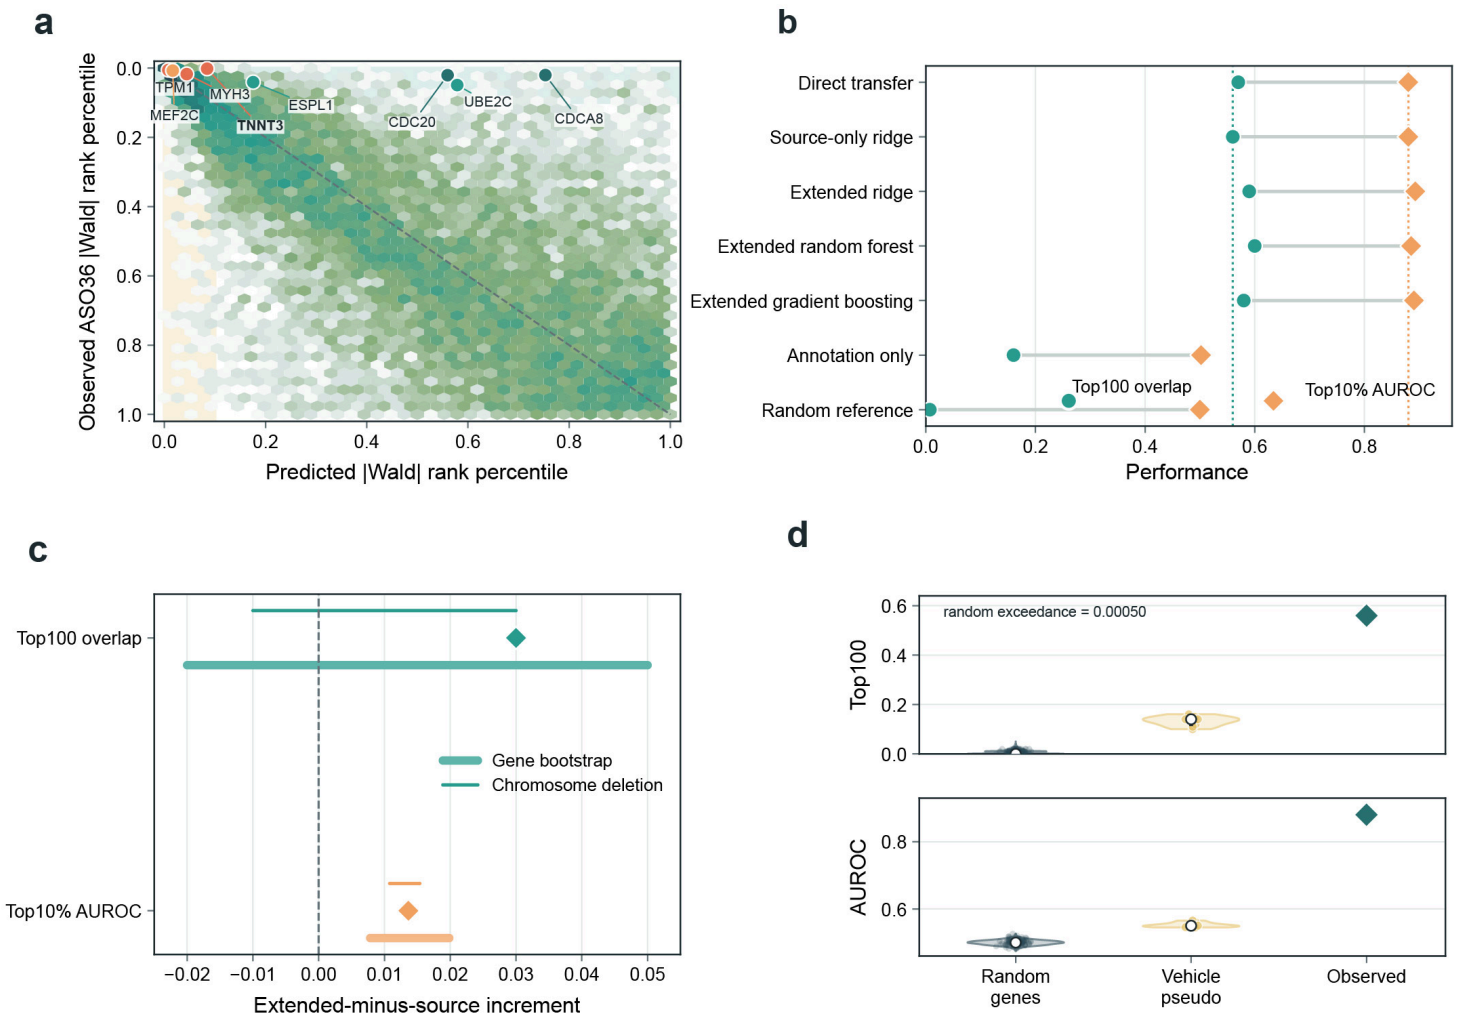


Figure 3

Source-time persistence defines the performance floor for cross-time recovery. a, Source-only ridge predictions are compared with observed ASO36 absolute Wald-rank percentiles under chromosome-group five-fold validation. The eight labelled genes comprise the contractile/myogenic and proliferation modules; shaded bands mark the predicted and observed top deciles. b, Wald-rank Top100 overlap (teal circles) and Top10% AUROC (light-orange diamonds) are compared across direct transfer, fitted models, an annotation-only baseline and a random reference. Dotted vertical lines mark source-only ridge performance. c, Diamonds show the full-data extended-minus-source increment. Thick pale segments are the 2.5th–97.5th percentiles from 2,000 paired gene resamples, and thin segments span 21 leave-one-chromosome-out estimates. Gene-resampling intervals do not represent biological-sample confidence intervals. d, Observed Top100 overlap and AUROC (deep-green diamonds) are compared with 2,000 random-gene rankings and nine vehicle pseudo-contrasts; violins contain the underlying null observations, vertical bars span the interquartile range and open circles mark medians.

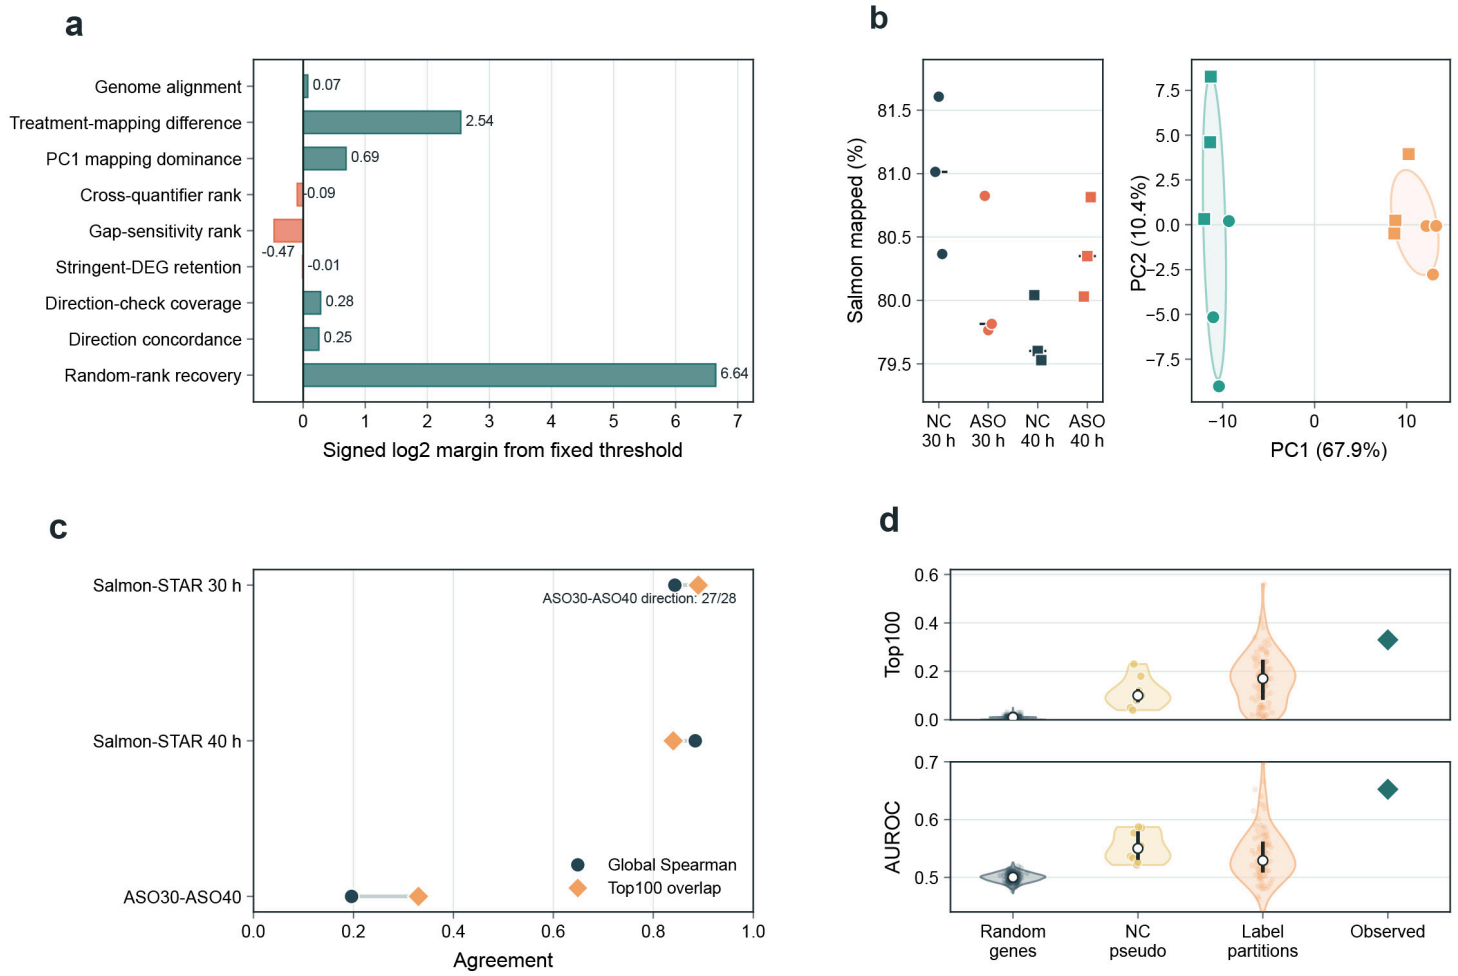


Figure 4

PRJNA575967 retains leading ranks despite estimator-sensitive global ordering. a, Each locked evaluation criterion is represented by its signed log2 margin from the fixed threshold. Positive margins indicate passage; deep green and brick red denote passed and failed criteria. b, The left panel shows all three mapping observations in each treatment–time cell with median segments; deep blue and brick red denote negative control and ASO, and circles and squares denote 30 and 40 h. The right panel shows VST-expression PCA, with teal and light orange denoting 30 and 40 h, circles and squares denoting negative control and ASO, and translucent covariance ellipses summarising each time point. c, Global Spearman correlation (deep-blue circles) and Top100 identity overlap (light-orange diamonds) are compared for the two Salmon–STAR contrasts and ASO30-to-ASO40 transfer. d, Observed cross-time Top100 overlap and AUROC (deep-green diamonds) are compared with 2,000 random-gene rankings, nine negative-control pseudo-contrasts and 100 unordered label-partition pairs. Violins show the underlying controls, vertical bars span the interquartile range and open circles mark medians.

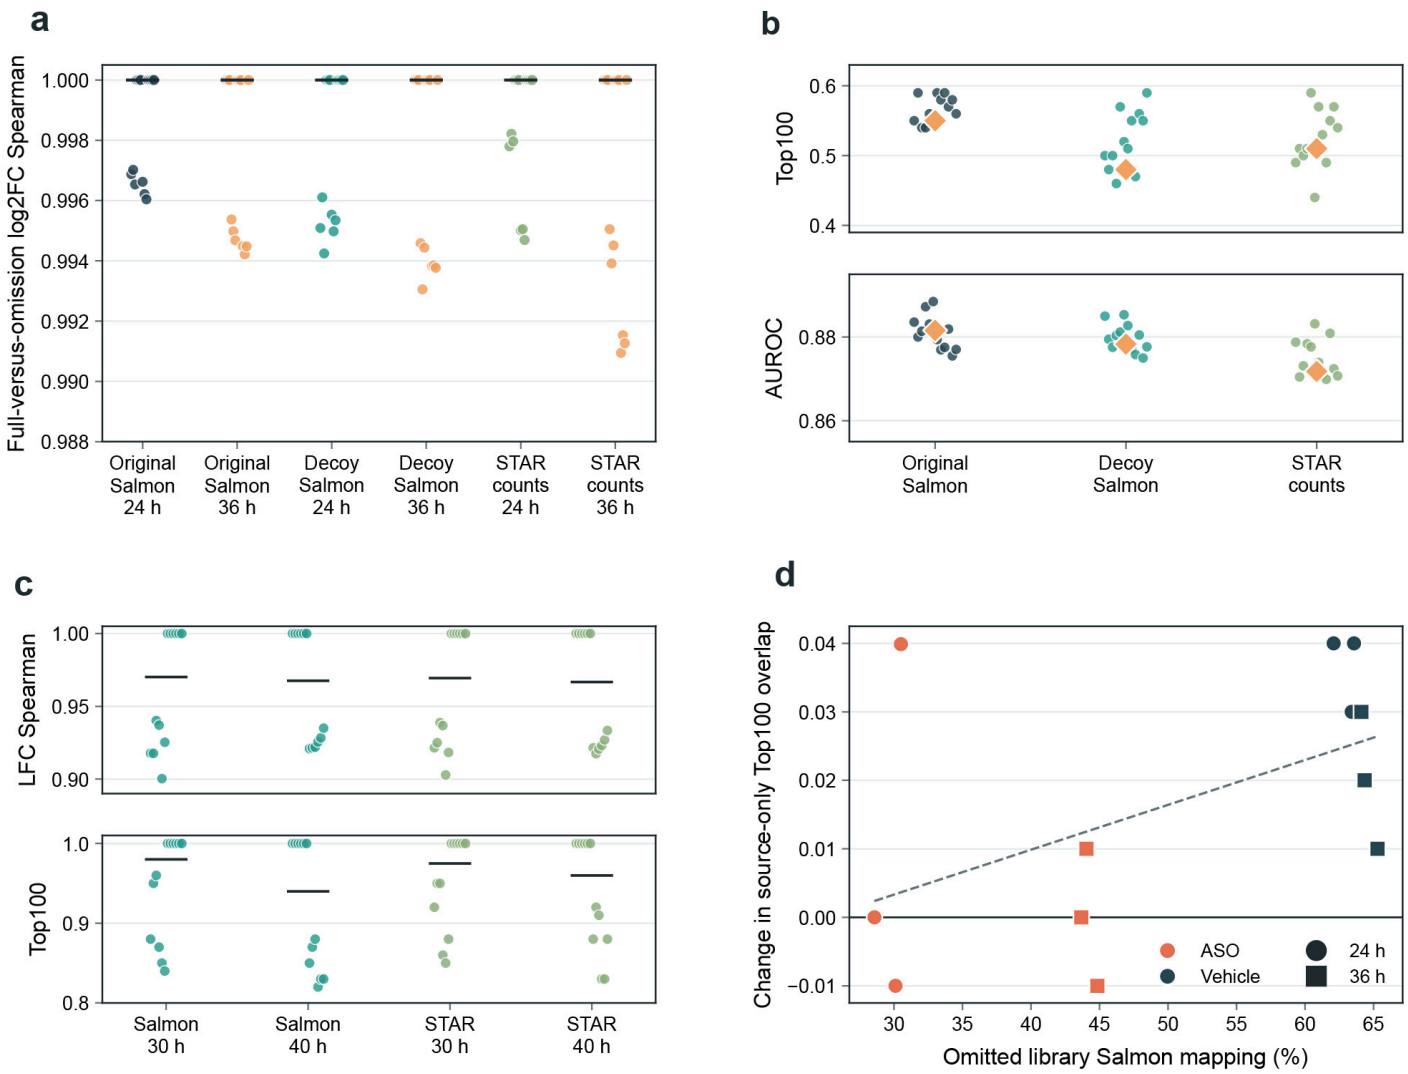


Figure 5

Leading ranks remain stable after single-library omission in both studies. a, Each point is a leave-one-library-out refit of the SRP186451 24- or 36-h treatment effect for original Salmon, decoy-aware Salmon or STAR GeneCounts. Horizontal segments mark medians; all 12 omissions are shown for each branch and time point. b, Source-only Top100 overlap and AUROC are shown after each SRP186451 omission; light-orange diamonds denote the corresponding full-data estimates. c, Full-versus-omission log2 fold-change correlation and Top100 identity overlap are shown for 48 PRJNA575967 refits across Salmon and STAR at 30 and 40 h. d, The omitted library's Salmon mapping rate is plotted against the resulting change in source-only Top100 overlap in SRP186451. Brick red and deep blue denote ASO and vehicle; circles and squares denote 24 and 36 h. The dashed line is a descriptive linear fit.

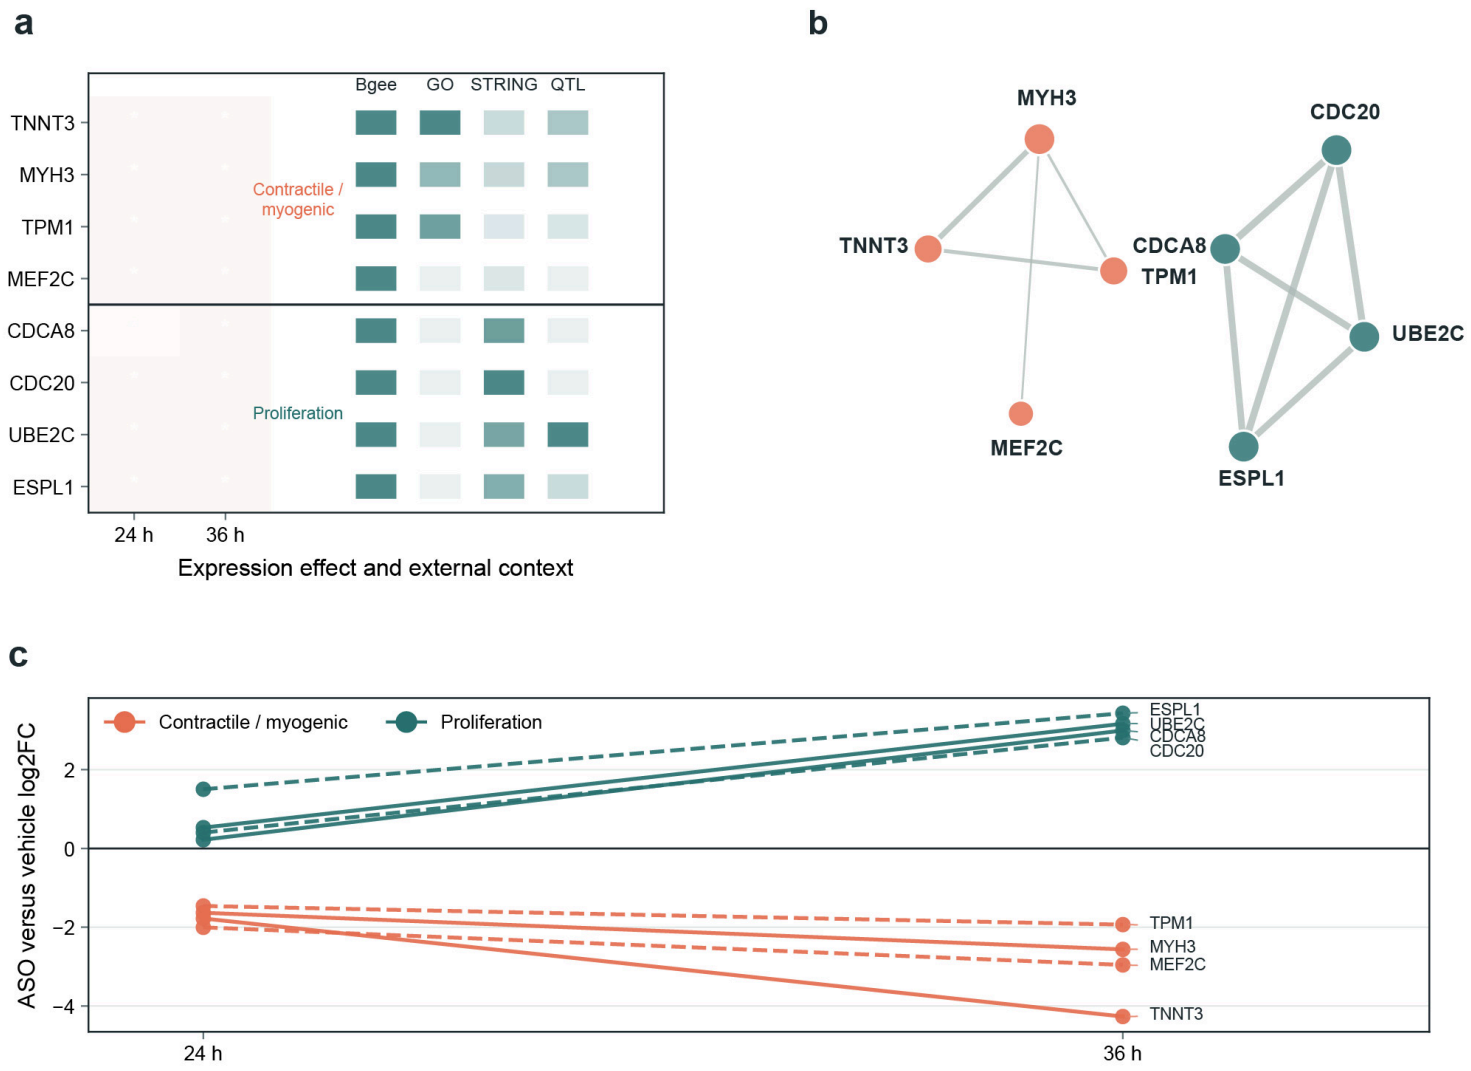


Figure 6

Stable leading genes form contractile/myogenic and proliferation modules. a, The expression heatmap shows ASO-versus-vehicle log2 fold changes at 24 and 36 h for eight leading genes; asterisks indicate Benjamini–Hochberg-adjusted $P < 0.05$. Adjacent tiles summarise Bgee muscle expression, muscle-related GO terms, STRING degree and priority-QTL context after column scaling. b, STRING associations among the eight genes are shown for combined scores ≥ 400. Brick-red nodes denote the contractile/myogenic module and deep-green nodes denote the proliferation module; node size scales with degree and edge width with the STRING combined score. c, Cross-time log2-fold-change trajectories show concordant suppression of the contractile/myogenic module and stronger late induction of the proliferation module. Solid and dashed lines distinguish individual genes within each module.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- Additionalfile1SupplementaryTables.xlsx

- Additionalfile2SupplementaryFigures.pdf