

scBench-Long: Verifiable Benchmarking of Long-Horizon Single-Cell Biology

Evaluating whether AI agents can recover complex scientific conclusions from raw single-cell biology data

Ian Diks^{1,*}, Zhen Yang^{1,*}, Arjun Banerjee¹, Tim Proctor¹, Kenny Workman¹

¹LatchBio, San Francisco, CA, USA

*Equal first authors

Correspondence: kenny@latch.bio

ABSTRACT

Single-cell studies require analysts to convert raw measurements into specific biological claims through multi-step workflows and integration of metadata, assay context, and auxiliary evidence. Existing AI-biology benchmarks largely measure broad knowledge, executable workflows, or local analysis steps. We introduce scBench-Long, a benchmark for long-horizon single-cell biology in which agents must recover scientific conclusions from raw or near-raw data without prescribed methods. The benchmark contains 21 evaluations spanning melanoma CD8 T-cell reactivity, CD8 RNA+ATAC regulatory inference, human-monkey chimera development, KRAS-driven lung tumor aging, and lethal COVID-19 lung pathology. Tasks cover paired scRNA/TCR sequencing, RNA and chromatin profiling, cross-species transcriptomics, combinatorial scRNA-seq, single-nucleus RNA-seq, immune repertoires, ortholog maps, ligand-receptor resources, and validation evidence. Candidate claims are reproduced, reviewed, and converted into controlled answer vocabularies with deterministic grading and trajectory rubrics. Across 1,068 completed trajectories, the strongest model-harness pair passes 16/63 runs (25.4%). scBench-Long evaluates whether agents can move beyond local analysis steps and make complex scientific claims that are supported by single-cell data.

Topline Benchmark Performance

We ran the current evaluation set across frontier model families and agent harnesses. Passing requires exact recovery of the graded structured answer for an evaluation attempt. Present systems show low but nonzero success rates: the strongest model-harness pair passes 16/63 completed runs across a 1,068-trajectory benchmark matrix, while most pairs solve only a small minority of tasks. This is consistent with the benchmark targeting long-horizon scientific workflow control rather than isolated tool execution.

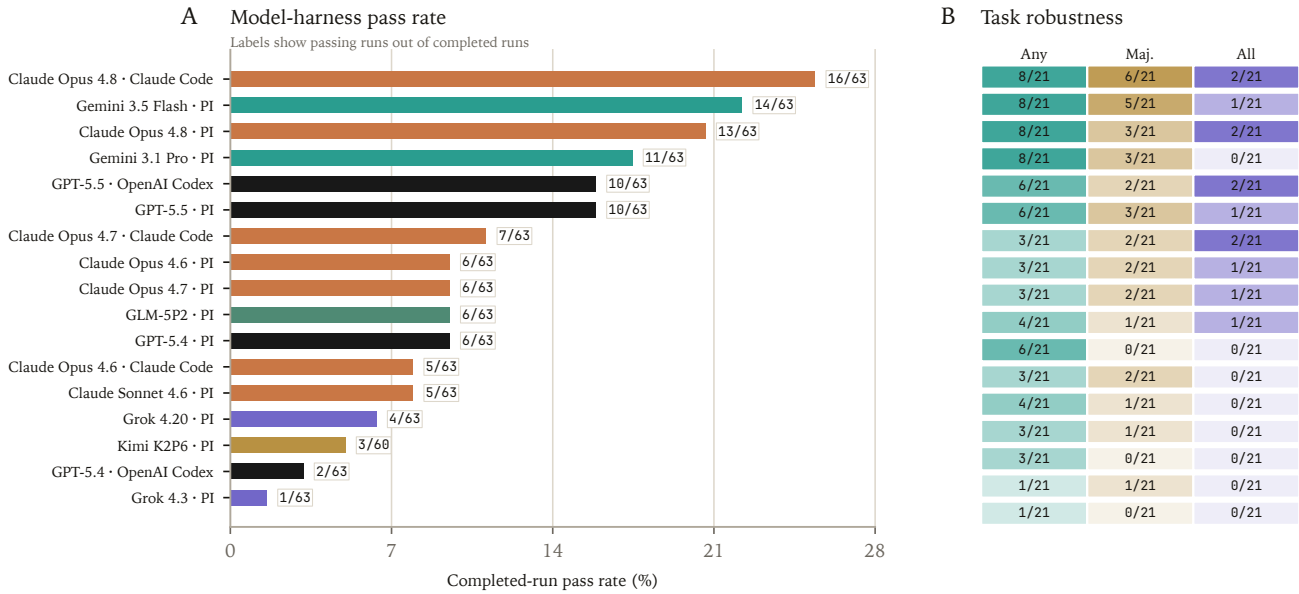


Figure 1: Topline benchmark performance. Completed-run pass rates and task-level robustness counts across 21 final evaluations and 17 model-harness pairs. Wilson confidence intervals are reported in Table 2 rather than plotted because runs are grouped within a shared evaluation set.

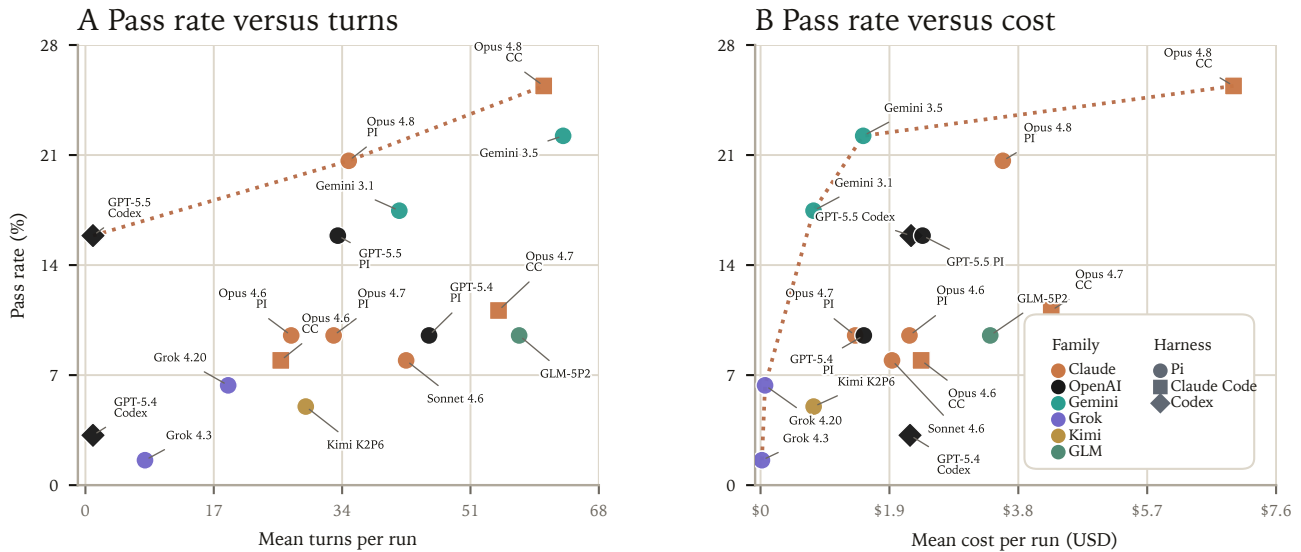


Figure 2: Efficiency frontiers for completed trajectories. Pass rate versus mean turns and mean cost per completed run. Points are model-harness pairs; dotted lines mark Pareto frontiers. Turn metadata were complete; cost metadata were available for 1,050/1,068 trajectories.

Introduction

Scientists use single-cell data to construct biological claims about cell states, perturbations, development, immunity, and disease. Raw single-cell measurements do not directly encode these claims. They must be processed through multi-step workflows, interpreted against experimental design, and integrated with assay-specific context, metadata, prior literature, and often additional assay data such as immune repertoires, chromatin accessibility, or functional validation experiments [17, 18, 19].

AI agents are beginning to show utility in biological data analysis, but current benchmarks largely test broad biology knowledge without deep treatment of the diverse task types specific to single-cell biology. Existing single-cell benchmarks focus on local analysis steps and do not evaluate realistic long-horizon science [34, 35, 36, 37, 38, 27, 28].

Verifiable ground truth is difficult to define in long-horizon single-cell benchmarking. The same dataset can support multiple valid claims, and published results do not always independently reproduce. A trustworthy evaluation must therefore

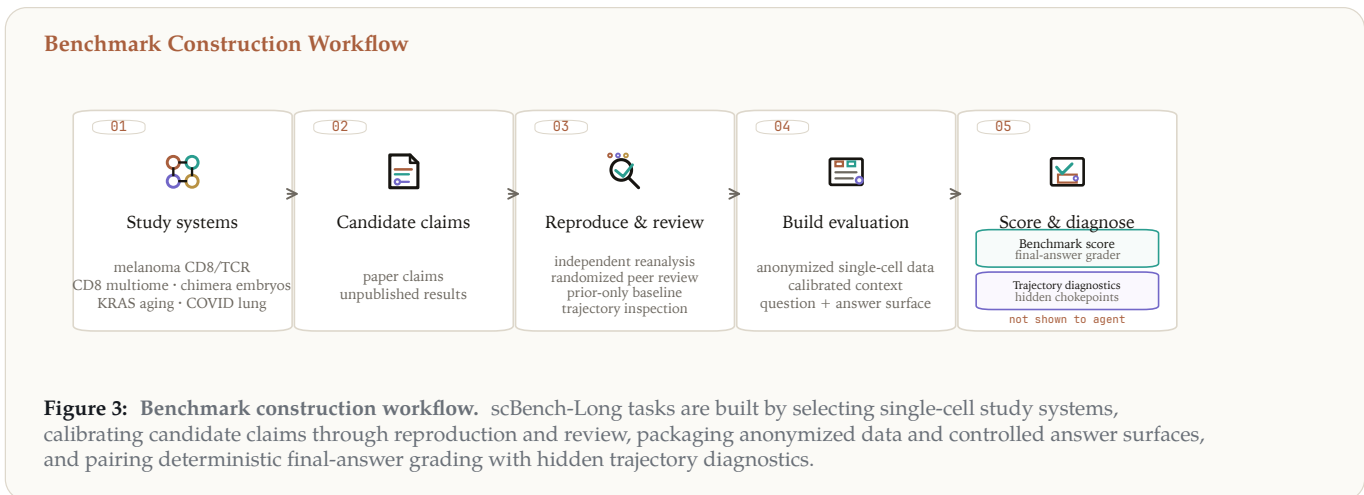
carefully constrain the set of admissible answers without penalizing valid and potentially unanticipated analysis paths [39, 40, 41, 43, 28, 29, 38].

We introduce scBench-Long, a benchmark for long-horizon single-cell biology. scBench-Long contains 21 evaluations across human melanoma CD8 T-cell reactivity, RNA+ATAC regulatory inference, human-monkey chimera development, KRAS-driven lung tumor aging, and lethal COVID-19 lung pathology [44, 45, 46, 47, 48]. Tasks provide raw or near-raw data, calibrated experimental context, auxiliary resources when needed, and compact scientific questions. Agents must recover structured biological conclusions that are graded deterministically over controlled vocabularies and symbols, with companion rubrics used to diagnose progress through key analysis chokepoints [27, 28, 29].

scBench-Long tests whether agents can move beyond accurately performing local analysis steps to reasoning about the kinds of realistic scientific tasks found in published studies or used in drug programs. We hope it serves as a useful measuring device for evaluating frontier agent systems as they improve in capability and are integrated into practical research and engineering.

Benchmark Design

Agents are evaluated on final scientific conclusions from 21 long-horizon single-cell evaluations spanning five study systems: human melanoma tumor-reactive CD8 T cells, RNA+ATAC regulation of CD8 tissue residency, human-monkey chimeric embryo development, age-dependent KRAS-driven lung tumorigenesis, and lethal COVID-19 lung pathology. The evaluations combine single-cell and single-nucleus transcriptomes, paired TCR or BCR repertoires, chromatin accessibility, cross-species ortholog mapping, ligand-receptor databases, developmental references, and functional validation data [44, 45, 46, 47, 48].



scBench-Long is structured around biological studies rather than independent datasets. Each study contributes multiple evaluations that revisit the same experimental system from different scientific perspectives, requiring agents to answer distinct questions about cell identity, state, regulation, composition, and mechanism. This mirrors how single-cell datasets are analyzed in practice: the same experiment is repeatedly

interrogated to support different biological claims [17, 18].

The benchmark builds on prior claim-conditioned benchmark construction rather than introducing a new general grading framework. Each evaluation includes raw or near-raw data, calibrated experimental context, scrubbed labels where needed, a compact scientific question, a controlled final-answer surface,

reproduction notes, and trajectory rubrics. Paper claims are used as sources for candidate evaluations, not as automatic ground truth. Candidate tasks are retained only when the target conclusion can be reproduced from the staged data and hardened through review, distractor design, and inspection of trajectories from multiple model families [28, 27, 36, 29].

The single-cell-specific contribution is the choice of study systems, evidence layers, answer surfaces, and failure probes. The 21 evaluations span human melanoma tumor-reactive CD8 T cell conjugates, RNA+ATAC regulation of CD8 tissue-resident-like programs, human-monkey chimeric embryo development, age-dependent KRAS-driven lung tumorigenesis, and lethal COVID-19 lung pathology. The underlying data include 10x Chromium 5' single-cell RNA sequencing with paired TCR sequencing, 10x Multiome RNA and chromatin profiling, Smart-seq2 cross-species transcriptomics, Parse Biosciences combinatorial single-cell RNA sequencing, and 10x Chromium 3' single-nucleus RNA sequencing. Several studies additionally require integrating T or B cell receptor repertoires, chromatin accessibility, ortholog mappings, ligand-receptor databases, developmental references, and immunohistochemistry [20, 21, 22, 23, 19, 24, 25].

Controlled answer surfaces are built around scientific objects that single-cell analysts reason with: cell-state identities, donor compositions, regulator and direction labels, clonotype or repertoire claims, ligand-receptor interpretations, and mechanism-caution conclusions [28, 29]. Distractors are chosen to separate data-supported conclusions from common shortcuts, including canonical marker priors, raw-abundance rankings, single-modality answers, and causal interpretations not supported by the data.

We use a small set of terms consistently throughout the paper. A *final evaluation* is one graded task with a controlled answer

surface. A *replicate* is an independent attempt at that evaluation. A *trajectory* is the completed agent run and analysis trace for one replicate. A *model-harness pair* is a base model evaluated inside a specific execution harness. The result set contains 17 model-harness pairs with three replicates per final evaluation where runs completed, for 1,068 completed trajectories.

Verifiable grading paired with rubric diagnostics

Benchmark scores use verifiable pass/fail grading on final scientific outcomes, following the endpoint-grading strategy used in prior verifiable biology-agent benchmarks [28, 27, 36]. Tasks require composing challenging local analysis steps: cell annotation, donor statistics, repertoire analysis, chromatin inspection, reference mapping, or literature-aware interpretation [17, 18, 19].

Verifiable endpoint grading is stable, but sparse. A model can fail the benchmark while solving many subproblems correctly, and deterministic grading necessarily penalizes answers outside the pre-specified target surface, including some valid claims not anticipated by the benchmark authors [30, 29]. We therefore retain rubric-based trajectory judging as a companion diagnostic rather than as the benchmark score. The general motivation for this design is shared with prior long-horizon biological benchmarks; the scBench-Long rubrics focus on single-cell chokepoints.

Evaluation authors define chokepoints after independent reproduction, review, and inspection of trajectories from multiple model families [28, 29]. Chokepoint rubrics are then used by LLM judges to score model trajectories, treating rubric scores as diagnostic evidence rather than replacement benchmark scores [29, 31, 32, 33].

Evaluation Inventory

scBench-Long is organized around study systems rather than independent datasets. Each system contributes a cluster of evaluations that asks different scientific questions with similar experimental context.

Across the five study systems, the benchmark contains 21 long-horizon evaluations. The melanoma CD8 tasks require agents to link gene-expression states with T-cell receptor clonotypes and interpret cell populations isolated by marker-based sorting. The CD8 multiome tasks ask whether proposed regulatory claims are supported by both RNA expression and chromatin accessibility. The chimeric embryo tasks require comparison of developmental states across species and evaluation of signaling between cell populations. The KRAS lung-tumor tasks test whether agents can separate age-associated effects from tumor-cell state and changes in the surrounding microenvironment. The COVID lung tasks require integration of single-nucleus expression profiles with immune-receptor, cell subset, and orthogonal validation evidence [44, 45, 46, 47, 48].

Study/system	Evals	Assays / data types	Main task themes
Melanoma tumor-reactive CD8 T cells	5	10x 5' scRNA-seq, paired scTCR-seq, sort-gate metadata	Link gene-expression states with T-cell receptor clonotypes; interpret marker-sorted cell populations
CD8 RNA+ATAC regulatory inference	3	10x Multiome snRNA-seq and scATAC-seq	Evaluate whether proposed regulatory claims are supported by both RNA expression and chromatin accessibility
Human-monkey chimeric embryos	4	Smart-seq2 scRNA-seq, ortholog maps, developmental references, ligand-receptor resources	Compare developmental states across species; evaluate signaling between cell populations
KRAS-driven lung tumor aging	4	Parse Biosciences scRNA-seq, mouse lung tumor cohorts, external aging signatures	Separate age-associated effects from tumor-cell state and surrounding microenvironment changes
Lethal COVID-19 lung pathology	5	Single-nucleus RNA-seq, BCR reconstruction, subset-level DE, immunohistochemistry	Integrate single-nucleus expression profiles with immune-receptor, cell subset, and orthogonal validation evidence

Table 1: Evaluation inventory. scBench-Long groups evaluations by study system while varying the claim, workflow, assay evidence, and controlled answer surface within each system.

Results

Topline Benchmark Performance

We aggregated 1,068 completed trajectories across 21 final evaluations and 17 model-harness pairs. Scores use deterministic pass/fail grading of the final answer. Claude Opus 4.8 with Claude Code has the highest pass rate, with 16/63 passing trajectories (25.4%; Wilson 95% CI, 16.3–37.3) followed by Gemini 3.5 Flash and PI at 14/63 (22.2%; 13.7–33.9).

Even the top agent systems solve only a minority of tasks consistently. Claude Opus 4.8 with Claude Code passes at least one replicate for 8/21 evaluations, at least two replicates for 6/21 evaluations, and all three replicates for 2/21 evaluations. Gemini 3.5 Flash with PI also reaches any-pass performance on 8/21 evaluations, but reaches majority-pass performance

on 5/21 and all-replicate performance on 1/21. Run-level pass rates therefore do not fully describe reliability and replicate-level outcomes remain necessary for interpreting long-horizon success.

Performance varies sharply across evaluations. Five evaluations have no successful trajectories, while the highest-pass evaluation, a pseudotime-delay task, has 26/51 passing trajectories (51.0%). Most evaluations sit in low-pass bins. Later results sections break these aggregate patterns into specific model behaviors and failure modes.

Harness choice also changes performance. Claude Opus 4.8 is strongest with Claude Code. GPT-5.4 is higher with PI than OpenAI Codex. GPT-5.5 has the same run-level pass rate under PI and OpenAI Codex, but different majority-pass and all-replicate counts.

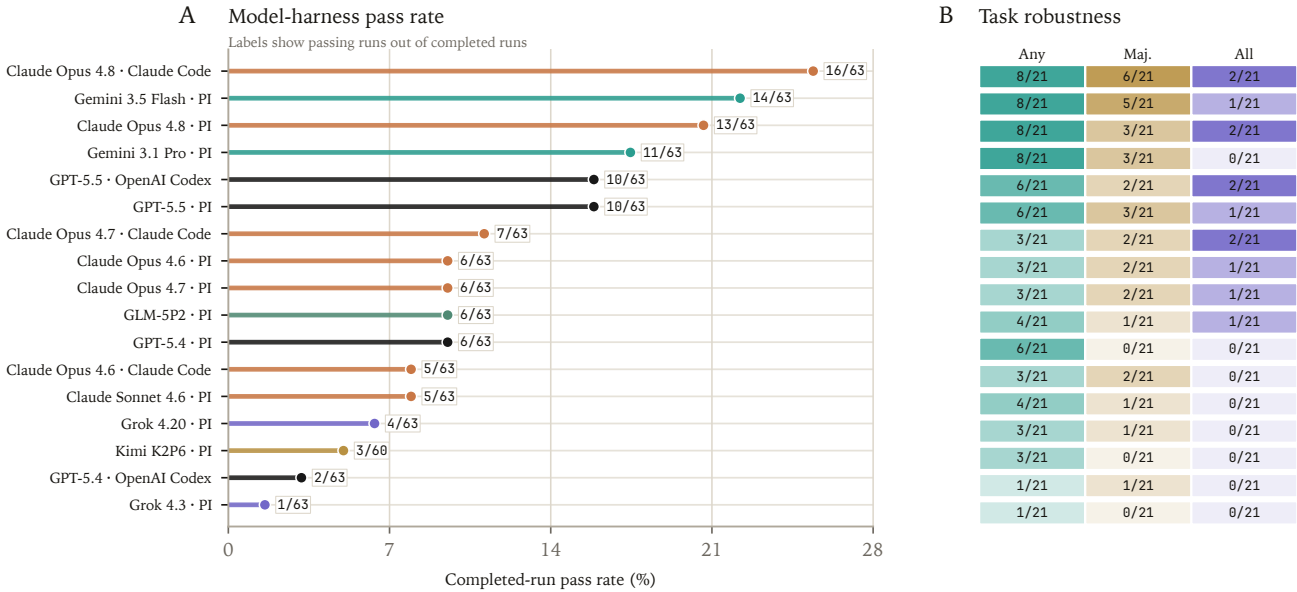


Figure 4: Main results summary. Completed-run pass rates and task-level robustness counts across 21 final evaluations and 17 model-harness pairs. Panel A reports passing runs out of completed runs. Panel B reports the number of evaluations with any passing replicate, a majority of passing replicates, or all replicates passing. Wilson confidence intervals are reported in Table 2.

Model-harness	Pass runs	Pass %	Wilson 95% CI	Any	Majority	All
Claude Opus 4.8 / Claude Code	16/63	25.40	16.28–37.34	8/21	6/21	2/21
Gemini 3.5 Flash / PI	14/63	22.22	13.72–33.91	8/21	5/21	1/21
Claude Opus 4.8 / PI	13/63	20.63	12.48–32.17	8/21	3/21	2/21
Gemini 3.1 Pro / PI	11/63	17.46	10.04–28.62	8/21	3/21	0/21
GPT-5.5 / OpenAI Codex	10/63	15.87	8.86–26.81	6/21	2/21	2/21
GPT-5.5 / PI	10/63	15.87	8.86–26.81	6/21	3/21	1/21
Claude Opus 4.7 / Claude Code	7/63	11.11	5.49–21.20	3/21	2/21	2/21
Claude Opus 4.6 / PI	6/63	9.52	4.44–19.26	3/21	2/21	1/21
Claude Opus 4.7 / PI	6/63	9.52	4.44–19.26	3/21	2/21	1/21
GLM-5P2 / PI	6/63	9.52	4.44–19.26	4/21	1/21	1/21
GPT-5.4 / PI	6/63	9.52	4.44–19.26	6/21	0/21	0/21
Claude Opus 4.6 / Claude Code	5/63	7.94	3.44–17.27	3/21	2/21	0/21
Claude Sonnet 4.6 / PI	5/63	7.94	3.44–17.27	4/21	1/21	0/21
Grok 4.20 / PI	4/63	6.35	2.50–15.22	3/21	1/21	0/21
Kimi K2P6 / PI	3/60	5.00	1.72–13.70	3/21	0/21	0/21
GPT-5.4 / OpenAI Codex	2/63	3.17	0.88–10.86	1/21	1/21	0/21
Grok 4.3 / PI	1/63	1.59	0.28–8.46	1/21	0/21	0/21

Table 2: Main verifiable results. Model-harness pass rates across the 21 final evaluations. Wilson intervals summarize run-level uncertainty. Any, Majority, and All report the number of evaluations with at least one passing replicate, a majority of passing replicates, or all replicates passing, respectively.

Rubric Judges Provide Dense but Imperfect Trajectory Diagnostics

Endpoint pass/fail grading, the deterministic final-answer score, remains the benchmark score, but it is too sparse to describe partial scientific progress. We therefore scored trajectories with four rubric judges using task-specific chokepoints. In the GLM-inclusive judge matrix, 1,067/1,068 trajectories had at least one judge row, 1,061 had at least three judge scores, 855 had all four judge scores, and 4,050/4,272 expected judge-score cells were present. Endpoint labels below come from the deterministic main results manifest; two stale fourth-replicate judge rows were excluded.

4.2.1 Judge scores are reproducible

Judge scores were consistent enough for diagnostic use. Mean pairwise judge correlation was 0.90, the minimum pairwise correlation was 0.87, and the mean absolute pairwise score difference was 7.6 percentage points. Same-trajectory judge variability was similar to variation across independent attempts on the same evaluation and model-harness group: judge SD averaged 5.5 points per trajectory, while replicate SD averaged 6.3 points.

4.2.2 Rubric scores enrich for success

Endpoint success remained sparse, with 125/1,068 passing trajectories (11.7%). By contrast, mean rubric score was nonzero for 1,066/1,068 trajectories. Among trajectories with at least three judge scores, endpoint-passing runs had higher mean rubric scores than endpoint-failing runs (75.9% vs. 58.7%). The association was useful but incomplete: Pearson $r = 0.29$, Spearman $\rho = 0.30$, and AUC = 0.77. The highest rubric decile contained 47/106 endpoint passes, so high rubric scores enrich for success but do not replace deterministic grading.

4.2.3 Rubric scores capture trajectory style

Rubric scores also varied by source model-harness. Opus 4.8 and GPT-5.5 trajectories received high rubric scores, GLM was intermediate, and Grok runs were lower. This ordering was not identical to endpoint pass rate: for example, Opus 4.8 with PI had a slightly higher mean rubric score than Opus 4.8 with Claude Code, while Claude Code had the higher endpoint pass rate. We therefore treat rubric scores as trajectory-style annotations and partial-progress diagnostics, not as calibrated scientific correctness scores.

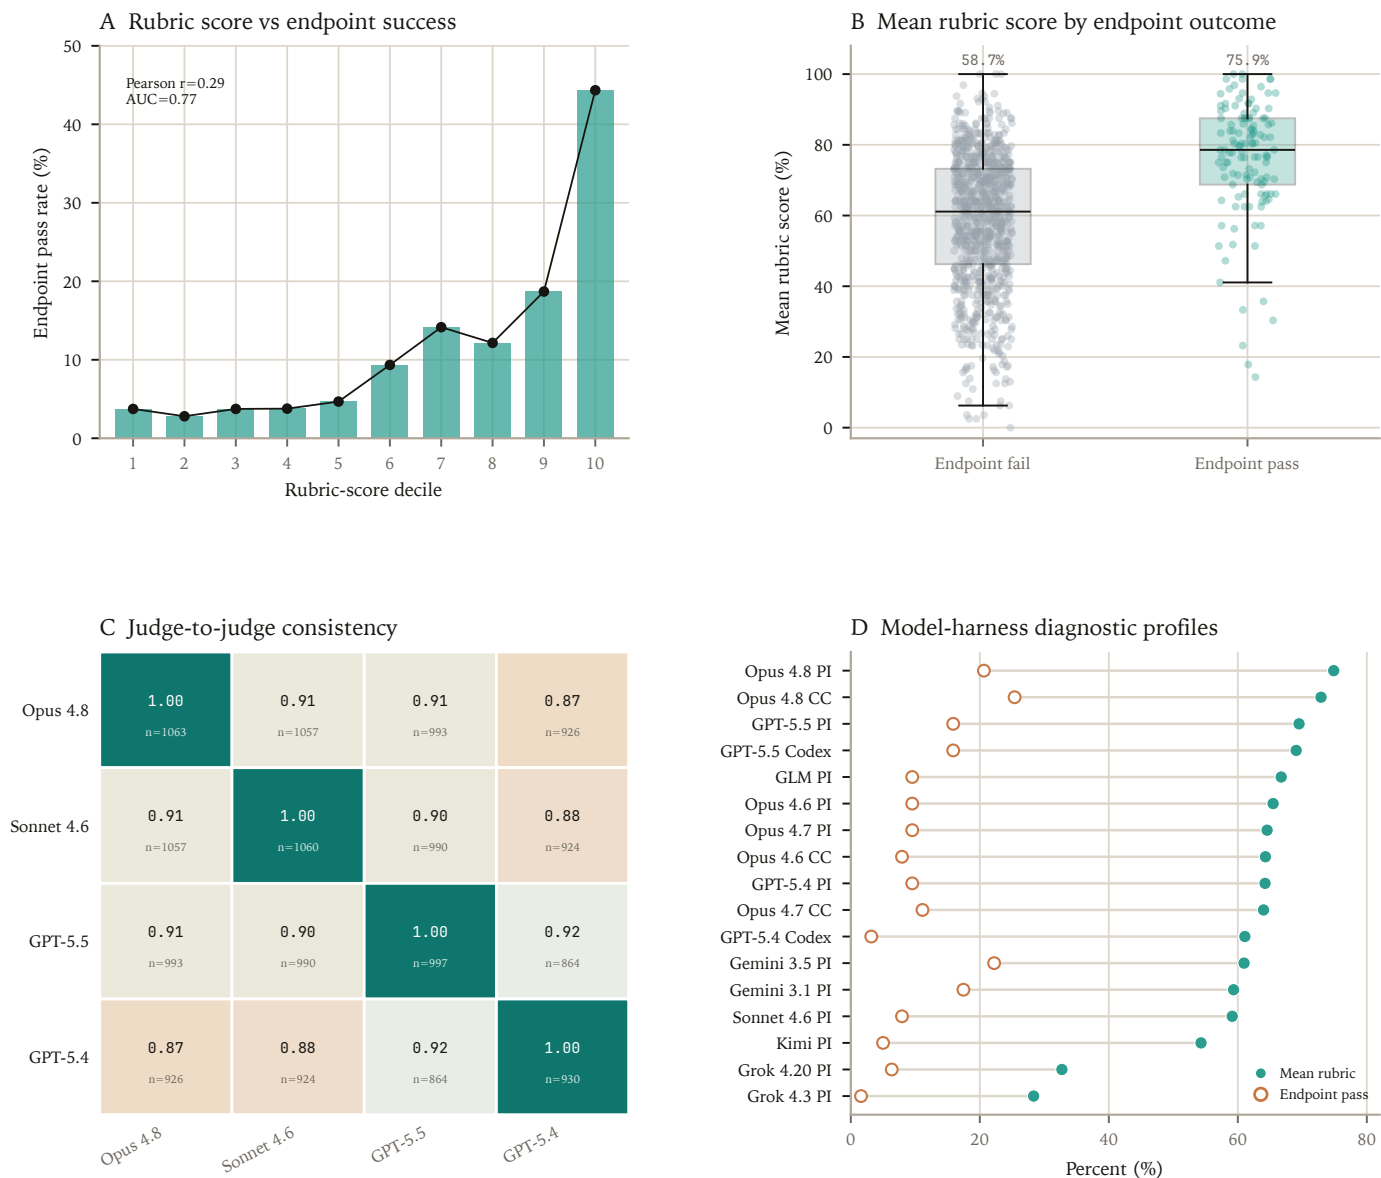


Figure 5: Rubric judges provide dense but imperfect diagnostics. Panel A shows endpoint pass rate by rubric-score decile, with endpoint correlation and AUC reported as diagnostics. Panel B shows mean rubric score by deterministic endpoint outcome. Panel C reports pairwise judge correlations on matched scored trajectories. Panel D compares endpoint pass rate and mean rubric score within each model-harness group. Endpoint labels use the GLM-inclusive deterministic results manifest; the single missing judge row is counted in the endpoint denominator as a fail and omitted from rubric-score means.

Failure Modes in Long-Horizon Biological Reasoning

Across the benchmark, models often recovered individual components of an analysis without assembling them into the correct biological conclusion. Claude, GPT, Gemini, GLM, Grok, and Kimi frequently identified relevant cell populations, recovered plausible ligand-receptor interactions, reconstructed immune receptor repertoires, and inferred broad regulatory programs. These intermediate results did not reliably translate into correct scientific claims. The recurring failure modes were not simple tool-use failures. Models prioritized abundant signals over specific evidence, substituted canonical biological expectations for measurements, drew causal conclusions from cross-sectional data, and simplified regulatory relationships

that required integrating multiple modalities.

4.3.1 Familiar Biological Priors Can Mask Task-Specific Evidence

Models will sometimes submit an answer based on a familiar biological prior, often some plausible conclusion from existing literature that shares some of the decision context, even after completing many correct analysis steps.

For instance, in melanoma CD8 T-cell tasks, models leaned on the familiar association between CD39, PD-1, TOX, LAG3, HAVCR2, CXCL13, and exhausted, antigen-experienced tumor-infiltrating CD8 T cells. This led them to treat CD39-high or PD-1-high singlets as the most tumor-reactive population, even

though the task required combining physical tumor/APC engagement with paired TCR clonal expansion.

Another task modified the input data by swapping two meta-data labels, `singlet` and `tumor_cluster`, while leaving the expression matrix and TCR data unchanged. This was to test whether an agent investigated the structure of the data empirically without drawing quick decisions from labels. Many models either relabeled the result back to the expected biology or partially followed the data and then overgeneralized.

4.3.2 Raw abundance can be mistaken for biological importance

Models often overattributed scientific importance to large quantitative values in several tasks.

In a human–monkey chimera task, agents analyzed ligand–receptor communication between hEPSC-derived human epiblast cells and host monkey epiblast cells from the same embryos [46, 24, 25]. The task required agents to decide whether significant ligand–receptor pairs supported a single dominant signaling family after accounting for directionality and database composition. Agents often treated the largest raw interaction class as the answer. However, the data contained many protein ligand–receptor pairs in both directions, with no single family dominating robustly in aggregate.

Only 2/50 trajectories passed, and both GPT-5.5 and Opus 4.8 were 0/6 across their paired harnesses.

4.3.3 Association Can Be Mistaken for Mechanism

In the melanoma tumor–conjugate task, agents compared malignant cells recovered from tumor:T-cell conjugates with malignant singlets. The task was to identify which tumor–cell programs were enriched in the conjugate population and to interpret those enrichments cautiously. Because the cells were sampled after physical pairing, the data show that certain tumor programs co-occur with T-cell contact. They do not, by themselves, show whether those programs were present before contact, induced by contact, or partly contaminated by transcripts from the co-sorted T cell.

The data supported two enriched malignant-cell programs with different levels of interpretability. A hypoxia/HIF program was robust across patients and controls, consistent with a tumor-intrinsic state enriched among conjugate-associated malignant cells. By contrast, the immune-response program

was more sensitive to decontamination and could not be assigned a direction from these data alone. It could reflect tumor cells predisposed to engage T cells, tumor cells responding to contact, selection of particular tumor states into conjugates, or residual RNA from the paired lymphocyte.

No model–harness pair passed this evaluation. Models often recovered plausible tumor states, interferon ligands, or antigen-presentation signals, but treated enrichment in the conjugate-associated population as evidence of a directional mechanism.

4.3.4 Failing to integrate multiple modalities

CD8 multiome tasks asked agents to infer regulatory control of the CD103+ Trm-like TIL program from paired snRNA-seq and scATAC-seq. These tasks required recognizing molecular patterns across RNA and chromatin accessibility [45, 19].

Model failures often only considered a single modality. In the KLF2 tasks, models often followed RNA differential expression and nominated genes such as *S1pr5* or *Cd160*, even though the relevant enhancer data did not support them as regulators of the Trm-like program. In another case, models recovered that KLF2 behaved as a repressor, but inferred its activity from *Klf2* RNA rather than from the coordinated behavior of its targets.

Harness effects in paired model comparisons

Harness choice changed endpoint outcomes for the same underlying model. We restricted this analysis to models evaluated under two harnesses: Claude Opus 4.6–4.8 under Claude Code and PI, and GPT-5.4–5.5 under OpenAI Codex and PI. Aggregate pass rate shifted for four of the five paired models. Claude Opus 4.8 was higher with Claude Code than PI (16/63 vs. 13/63), Claude Opus 4.7 was slightly higher with Claude Code (7/63 vs. 6/63), Claude Opus 4.6 was slightly higher with PI (6/63 vs. 5/63), and GPT-5.4 was higher with PI (6/63 vs. 2/63). GPT-5.5 had the same aggregate pass rate under OpenAI Codex and PI (10/63 for both).

Aggregate pass rate understated the harness effect. For each paired model, 3–6 of 21 evaluations changed pass count by harness, and GPT-5.5 solved different evaluations despite equal overall pass rate. These shifts show that long-horizon single-cell performance depends on the model–harness system, not only on the base model. We therefore report results at the model–harness level and interpret harnesses as part of the evaluated agent.

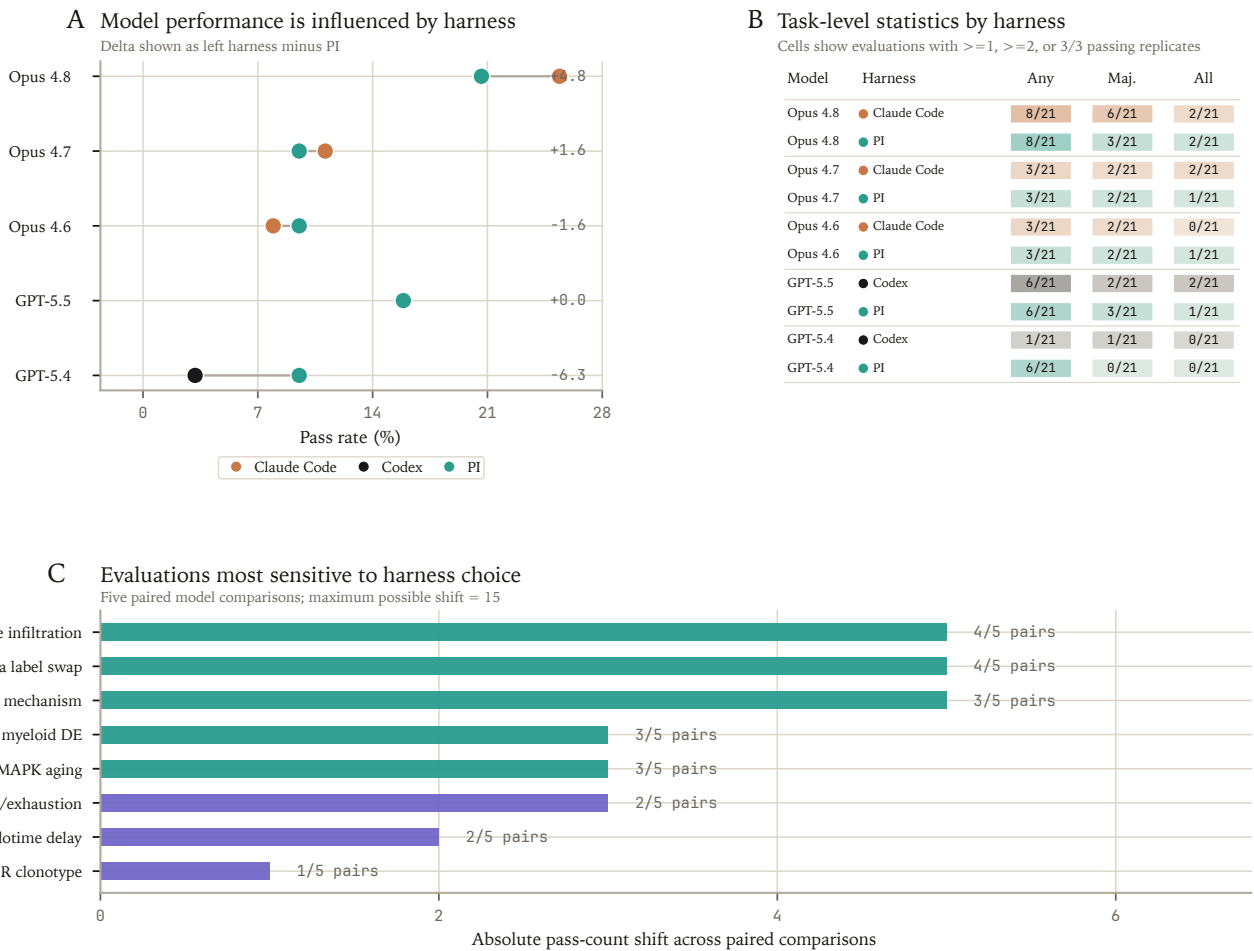


Figure 6: Harness effects in paired model comparisons. Panel A compares completed-run pass rates for identical models evaluated under two harnesses. Delta labels report the non-PI harness minus PI. Panel B reports task-level statistics by model-harness group: cells show evaluations with at least one, at least two, or all three passing replicates. Panel C ranks evaluations by the total absolute shift in passed replicates across paired harness comparisons.

Discussion

scBench-Long evaluates whether AI agents can recover claim-conditioned scientific conclusions from raw or near-raw single-cell data. The benchmark stages the data and context needed to test a specific biological claim, then grades the final structured answer deterministically. This follows prior verifiable long-horizon benchmark design while adapting it to single-cell tasks that require donor-aware reasoning, immune-repertoire analysis, chromatin integration, cross-species mapping, and orthogonal validation evidence [28, 27, 29].

Current agents show limited but measurable capability. The strongest model-harness pair, Claude Opus 4.8 with Claude Code, passes 16/63 completed runs (25.4%), while Gemini 3.5 Flash with PI passes 14/63 (22.2%). Success is not yet reliable: the top pair passes at least one replicate for 8/21 evaluations and all three replicates for only 2/21.

Models often execute plausible single-cell workflows but fail to reach the correct biological conclusion. We observed several recurring failure modes: substituting prior literature expectations for empirical evidence, conflating large numerical values with scientific importance, inferring mechanisms from molecular associations, and failing to integrate evidence across modalities. scBench-Long therefore evaluates agents beyond their ability to run local analysis steps, testing whether they can complete multi-step analyses and decide which final claims are supported by the data.

We hope scBench-Long serves both as a measurement tool and a diagnostic lens for developing agents that analyze single-cell data faithfully, transparently, and reproducibly. It is a focused contribution within a broader benchmark family spanning major biological data classes and work categories, including spatial biology, epigenomics, and therapeutics. More broadly, we view benchmarks as evolving specifications of computational biology workflows, supporting test-driven development of agent systems whose behavior can improve through both model training and harness engineering.

Methods

Evaluation construction

scBench-Long evaluations were constructed from study systems in which a specific single-cell claim could be tested from staged data and calibrated experimental context. We selected candidate claims from source studies and internal reproduction notes, then retained tasks only when the target conclusion could be independently reproduced from the data available to the agent. Source papers were therefore used as claim generators rather than as automatic ground truth.

Each evaluation was packaged as a task specification containing a compact scientific question, raw or near-raw data nodes, any required auxiliary resources, a controlled answer surface, metadata describing the biological dependency stack, and a deterministic grader. Data nodes included single-cell or single-nucleus expression matrices, immune-receptor tables, chromatin accessibility or multiome objects, cross-species mapping resources, ligand-receptor resources, developmental references, and vocabulary files as needed by the task. Task prompts were written to provide the experimental context needed for analysis without prescribing a solution path.

Constructing ground truths

Final answers were structured JSON objects over controlled vocabularies and symbols. Answer surfaces were chosen to match the scientific objects used by single-cell analysts, including cell-state labels, donor-level composition calls, regulator identities, direction labels, clonotype claims, ligand-receptor interpretations, and explicit mechanism-caution labels. Distractors and hard-fail conditions were added during review to separate data-supported conclusions from common shortcuts such as canonical marker recall, abundance-only ranking, single-modality reasoning, or causal interpretation of cross-sectional associations.

Benchmark scores use deterministic pass/fail grading of the final answer. Each grader combines typed checks over the submitted JSON answer, including list matching, dictionary matching, field-level predicates, required keys, multi-accepted synonyms or resolution-sensitive labels, and hard-fail conditions for biologically incompatible answers. Grading is intentionally endpoint-based: intermediate analyses are not required to follow a prescribed workflow, and agents can pass through any valid analysis path that reaches the controlled target surface.

Model-harness runs and result aggregation

The analyzed result set contains 21 final evaluations and 17 model-harness pairs. Each model-harness pair was run for three replicates per evaluation where runs completed, yielding 1,068 completed trajectories. The evaluated harnesses were

PI, Claude Code, and OpenAI Codex. The evaluated model families included Claude, Gemini, OpenAI GPT, Grok, Kimi, and GLM systems.

Run-level results were aggregated from the GLM-inclusive results manifest `results/scbench-long-results-w-glm.csv`. Each row records the model, harness, evaluation identifier, replicate, remote paths to the evaluation and result artifacts, and deterministic final-pass status. Model-harness summaries report completed runs, passing runs, Wilson confidence intervals, and task-level robustness counts: evaluations with at least one passing replicate, a majority of passing replicates, or all completed replicates passing. Missing attempts were excluded from run denominators and tracked separately during aggregation.

Cost and turn metadata were extracted from per-run result JSON artifacts when available. Turn metadata were complete for the analyzed trajectories; cost metadata were available for 1,050/1,068 trajectories. Efficiency summaries use mean cost and mean turn count over completed runs and are interpreted as harness- and logging-dependent diagnostics rather than benchmark scores.

Rubric-based trajectory diagnostics

Endpoint grading is stable but sparse, so each evaluation also includes trajectory rubrics for diagnostic analysis. Rubrics were written after reproduction and review to capture single-cell chokepoints expected to matter across plausible solution paths, such as identifying the correct biological comparison, preserving donor structure, using receptor evidence appropriately, checking chromatin support for RNA-level claims, mapping cross-species states, or avoiding causal claims unsupported by the data.

Rubric judges scored model trajectories independently of the deterministic benchmark score. The GLM-inclusive judge matrix contains one row per judged trajectory and per-judge fractional rubric scores from four judge models when available. Rubric summaries use matched deterministic endpoint labels from the main results manifest. Because rubric scores are prompt- and judge-dependent, they are treated as companion diagnostics for partial progress and failure-mode analysis, not as replacement benchmark scores.

References

- [1] Moses, L. & Pachter, L. Museum of spatial transcriptomics. *Nature Methods* **19**, 534–546 (2022). <https://doi.org/10.1038/s41592-022-01409-2>.
- [2] Williams, C. G., Lee, H. J., Asatsuma, T., Vento-Tormo, R. & Haque, A. An introduction to spatial transcriptomics for biomedical research. *Genome Medicine* **14**, 68 (2022). <https://doi.org/10.1186/s13073-022-01075-1>.
- [3] Dries, R., Chen, J., Del Rossi, N., Khan, M. M., Sistig, A. & Yuan, G. C. Advances in spatial transcriptomic data analysis. *Genome Research* **31**, 1706–1718 (2021). <https://doi.org/10.1101/gr.275224.121>.
- [4] Lyubetskaya, A., Rabe, B., Kavran, A., Bai, Y. et al. *In situ* multi-modal characterization of pancreatic cancer reveals tumor cell identity as a defining factor of the surrounding microenvironment. *Cell Reports* **45**, 116827 (2026). <https://doi.org/10.1016/j.celrep.2025.116827>.
- [5] Ishahak, M., Han, R. H., Annamalai, D., Woodiwiss, T., McCornack, C., Cleary, R. T., DeSouza, P. A., Qu, X., Dahiya, S., Kim, A. H. & Millman, J. R. Genetically engineered brain organoids recapitulate spatial and developmental states of glioblastoma progression. *Advanced Science* **12**, 2410110 (2025). <https://doi.org/10.1002/advs.202410110>.
- [6] Jones, M. G., Sun, D., Min, K. H. J., Colgan, W. N., Wang, H., Torok, T., Cardoso, E. C., Tian, L., Weir, J. A., Chen, V. Z., Koblan, L. W., Yost, K. E., Mathey-Andrews, N., D'Souza, E., Russell, A. J. C., Stickels, R. R., Balderrama, K. S., Rideout, W. M., Dai, M., Marrero, G., Kumar, V., Saqi, A., Chen, F., Weissman, J. S., Yosef, N. & Yang, D. Spatiotemporal lineage tracing reveals the dynamic spatial architecture of tumor growth and metastasis. *bioRxiv* (2024). <https://doi.org/10.1101/2024.10.21.619529>.
- [7] Yang, D., Jones, M. G., Naranjo, S., Rideout, W. M., Min, K. H. J., Ho, R., Wu, W., Replogle, J. M., Page, J. L., Quinn, J. J., Horns, F., Qiu, X., Chen, M. Z., Freed-Pastor, W. A., McGinnis, C. S., Patterson, D. M., Gartner, Z. J., Chow, E. D., Bivona, T. G., Chan, M. M., Yosef, N., Jacks, T. & Weissman, J. S. Lineage tracing reveals the phylogenetics, plasticity, and paths of tumor evolution. *Cell* **185**, 1905–1923.e25 (2022). <https://doi.org/10.1016/j.cell.2022.04.015>.
- [8] DuPage, M., Dooley, A. L. & Jacks, T. Conditional mouse lung cancer models using adenoviral or lentiviral delivery of Cre recombinase. *Nature Protocols* **4**, 1064–1072 (2009). <https://doi.org/10.1038/nprot.2009.95>.
- [9] Rodrigues, S. G., Stickels, R. R., Goeva, A., Martin, C. A., Murray, E., Vanderburg, C. R., Welch, J., Chen, L. M., Chen, F. & Macosko, E. Z. Slide-seq: A scalable technology for measuring genome-wide expression at high spatial resolution. *Science* **363**, 1463–1467 (2019). <https://doi.org/10.1126/science.aaw1219>.
- [10] Russell, A. J. C., Weir, J. A., Nadaf, N. M., Shabet, M., Kumar, V., Kambhampati, S., Raichur, R., Marrero, G. J., Liu, S., Balderrama, K. S., Vanderburg, C. R., Shanmugam, V., Tian, L., Iorgulescu, J. B., Yoon, C. H., Wu, C. J., Macosko, E. Z. & Chen, F. Slide-tags enables single-nucleus barcoding for multimodal spatial genomics. *Nature* **625**, 101–109 (2024). <https://doi.org/10.1038/s41586-023-06837-4>.

- [11] Groh, J., Feng, R., Yuan, X., Liu, L., Klein, D. et al. Microglia activation orchestrates CXCL10-mediated CD8+ T cell recruitment to promote aging-related white matter degeneration. *Nature Neuroscience* **28**, 1160–1173 (2025). <https://doi.org/10.1038/s41593-025-01955-w>.
- [12] Singhal, V., Chou, N., Lee, J., Yue, Y., Liu, J., Chock, W. K., Lin, L., Chang, Y.-C., Teo, E. M. L., Aow, J., Lee, H. K., Chen, K. H. & Prabhakar, S. BANKSY unifies cell typing and tissue domain segmentation for scalable spatial omics data analysis. *Nature Genetics* **56**, 431–441 (2024). <https://doi.org/10.1038/s41588-024-01664-3>.
- [13] Varrone, M., Tavernari, D., Santamaria-Martinez, A., Walsh, L. A. et al. CellCharter reveals spatial cell niches associated with tissue remodeling and cell plasticity. *Nature Genetics* **56**, 74–84 (2024). <https://doi.org/10.1038/s41588-023-01588-4>.
- [14] Qin, F., Luo, X., Lu, Q., Cai, B., Xiao, F. & Cai, G. Spatial pattern and differential expression analysis with spatial transcriptomic data. *Nucleic Acids Research* **52**, e101 (2024). <https://doi.org/10.1093/nar/gkae962>.
- [15] Kleshchevnikov, V., Shmatko, A., Dann, E., Aivazidis, A., King, H. W., Li, T., Elmentaite, R., Lomakin, A., Kedlian, V., Gayoso, A., Jain, M. S., Park, J. S., Ramona, L., Tuck, E., Arutyunyan, A., Vento-Tormo, R., Gerstung, M., James, L., Stegle, O. & Bayraktar, O. A. Cell2location maps fine-grained cell types in spatial transcriptomics. *Nature Biotechnology* **40**, 661–671 (2022). <https://doi.org/10.1038/s41587-021-01139-4>.
- [16] Biancalani, T., Scalia, G., Buffoni, L., Avasthi, R., Lu, Z., Sanger, A., Tokcan, N., Vanderburg, C. R., Segerstolpe, A., Zhang, M., Avraham-Davidi, I. & Regev, A. Deep learning and alignment of spatially resolved single-cell transcriptomes with Tangram. *Nature Methods* **18**, 1352–1362 (2021). <https://doi.org/10.1038/s41592-021-01264-7>.
- [17] Lähnemann, D., Köster, J., Szczurek, E., McCarthy, D. J., Hicks, S. C. et al. Eleven grand challenges in single-cell data science. *Genome Biology* **21**, 31 (2020). <https://doi.org/10.1186/s13059-020-1926-6>.
- [18] Heumos, L., Schaar, A. C., Lance, C., Litinetskaya, A., Drost, F. et al. Best practices for single-cell analysis across modalities. *Nature Reviews Genetics* **24**, 550–572 (2023). <https://doi.org/10.1038/s41576-023-00586-w>.
- [19] Hao, Y., Hao, S., Andersen-Nissen, E., Mauck, W. M., Zheng, S. et al. Integrated analysis of multimodal single-cell data. *Cell* **184**, 3573–3587.e29 (2021). <https://doi.org/10.1016/j.cell.2021.04.048>.
- [20] Zheng, G. X. Y., Terry, J. M., Belgrader, P., Ryvkin, P., Bent, Z. W. et al. Massively parallel digital transcriptional profiling of single cells. *Nature Communications* **8**, 14049 (2017). <https://doi.org/10.1038/ncomms14049>.
- [21] Han, A., Glanville, J., Hansmann, L. & Davis, M. M. Linking T-cell receptor sequence to functional phenotype at the single-cell level. *Nature Biotechnology* **32**, 684–692 (2014). <https://doi.org/10.1038/nbt.2938>.
- [22] Picelli, S., Faridani, O. R., Björklund, Å. K., Winberg, G., Sagasser, S. & Sandberg, R. Full-length RNA-seq from single cells using Smart-seq2. *Nature Protocols* **9**, 171–181 (2014). <https://doi.org/10.1038/nprot.2014.006>.
- [23] Rosenberg, A. B., Roco, C. M., Muscat, R. A., Kuchina, A., Sample, P. et al. Single-cell profiling of the developing mouse brain and spinal cord with split-pool barcoding. *Science* **360**, 176–182 (2018). <https://doi.org/10.1126/science.aam8999>.
- [24] Efremova, M., Vento-Tormo, M., Teichmann, S. A. & Vento-Tormo, R. CellPhoneDB: inferring cell-cell communication from combined expression of multi-subunit ligand-receptor complexes. *Nature Protocols* **15**, 1484–1506 (2020). <https://doi.org/10.1038/s41596-020-0292-x>.
- [25] Jin, S., Guerrero-Juarez, C. F., Zhang, L., Chang, I., Ramos, R. et al. Inference and analysis of cell-cell communication using CellChat. *Nature Communications* **12**, 1088 (2021). <https://doi.org/10.1038/s41467-021-21246-9>.
- [26] Workman, K., Yang, Z., Muralidharan, H. & Le, H. SpatialBench: Can agents analyze real-world spatial biology data? *arXiv arXiv:2512.21907* (2025). <https://doi.org/10.48550/arXiv.2512.21907>.
- [27] Workman, K., Yang, Z., Muralidharan, H., Abdulali, A. & Le, H. scBench: Evaluating AI agents on single-cell RNA-seq analysis. *arXiv arXiv:2602.09063* (2026). <https://doi.org/10.48550/arXiv.2602.09063>.
- [28] Diks, I., Muralidharan, H., Proctor, T. & Workman, K. Verifiable benchmarking of long-horizon spatial biology. *arXiv arXiv:2605.28065* (2026). <https://doi.org/10.48550/arXiv.2605.28065>.
- [29] Qu, Y., Lu, Y., Tu, X., Zhang, S., She, T., Shaw, A. G., Shih, J.-H., Zhao, B. et al. BiomniBench: Process-level evaluation of LLM agents for real-world biomedical research. *bioRxiv* (2026). <https://doi.org/10.64898/2026.05.12.724604>.
- [30] Lightman, H., Kosaraju, V., Burda, Y., Edwards, H., Baker, B., Lee, T., Leike, J., Schulman, J., Sutskever, I. & Cobbe, K. Let’s verify step by step. *arXiv arXiv:2305.20050* (2023). <https://doi.org/10.48550/arXiv.2305.20050>.
- [31] Zheng, L., Chiang, W.-L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E. P., Zhang, H., Gonzalez, J. E. & Stoica, I. Judging LLM-as-a-judge with MT-Bench and Chatbot Arena. *arXiv arXiv:2306.05685* (2023). <https://doi.org/10.48550/arXiv.2306.05685>.

- [32] Wang, P., Li, L., Chen, L., Cai, Z., Zhu, D., Lin, B., Cao, Y., Kong, L., Liu, Q., Liu, T. & Sui, Z. Large language models are not fair evaluators. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 9440–9450 (2024). <https://doi.org/10.18653/v1/2024.acl-long.511>.
- [33] Liu, Y., Iyer, D., Xu, Y., Wang, S., Xu, R. & Zhu, C. G-Eval: NLG evaluation using GPT-4 with better human alignment. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 2511–2522 (2023). <https://doi.org/10.18653/v1/2023.emnlp-main.153>.
- [34] Laurent, J. M., Janizek, J. D., Ruzo, M., Hinks, M. M., Hammerling, M. J., Narayanan, S., Ponnampati, M., White, A. D. & Rodrigues, S. G. LAB-Bench: Measuring capabilities of language models for biology research. *arXiv arXiv:2407.10362* (2024). <https://doi.org/10.48550/arXiv.2407.10362>.
- [35] Mitchener, L., Laurent, J. M., Andonian, A., Tenmann, B., Narayanan, S., Wellawatte, G. P., White, A., Sani, L. & Rodrigues, S. G. BixBench: a comprehensive benchmark for LLM-based agents in computational biology. *arXiv arXiv:2503.00096* (2025). <https://doi.org/10.48550/arXiv.2503.00096>.
- [36] Nair, S., Gunsalus, L., Orcutt-Jahns, B., Rossen, J., Lal, A., De Donno, C., Celik, M. H., Fletez-Brant, K., Xie, X., Corrada Bravo, H. & Eraslan, G. Agentic systems are adept at solving well-scoped, verifiable problems in computational biology. *bioRxiv* (2026). <https://doi.org/10.64898/2026.04.06.716850>.
- [37] Li, J. & Ho, A. GeneBench: Assessing AI agents for multi-stage inference problems in genomics and quantitative biology. *bioRxiv* (2026). <https://doi.org/10.64898/2026.04.22.720113>.
- [38] Anthropic. Evaluating Claude’s bioinformatics research capabilities with BioMysteryBench. *Anthropic Research* (2026). anthropic.com/research/BioMysteryBench.
- [39] Ioannidis, J. P. A. Why most published research findings are false. *PLOS Medicine* **2**, e124 (2005). <https://doi.org/10.1371/journal.pmed.0020124>.
- [40] Prinz, F., Schlange, T. & Asadullah, K. Believe it or not: how much can we rely on published data on potential drug targets? *Nature Reviews Drug Discovery* **10**, 712 (2011). <https://doi.org/10.1038/nrd3439-c1>.
- [41] Begley, C. G. & Ellis, L. M. Raise standards for pre-clinical cancer research. *Nature* **483**, 531–533 (2012). <https://doi.org/10.1038/483531a>.
- [42] Errington, T. M., Denis, A., Perfito, N., Iorns, E. & Nosek, B. A. Reproducibility in Cancer Biology: Challenges for assessing replicability in preclinical cancer biology. *eLife* **10**, e67995 (2021). <https://doi.org/10.7554/eLife.67995>.
- [43] Errington, T. M., Mathur, M., Soderberg, C. K., Denis, A., Perfito, N., Iorns, E. & Nosek, B. A. Investigating the replicability of preclinical cancer biology. *eLife* **10**, e71601 (2021). <https://doi.org/10.7554/eLife.71601>.
- [44] Ibáñez-Molero, S., Veldman, J., Simon Nieto, J., Traets, J. J., George, A., Hoefakker, K. et al. Tumour-reactive heterotypic CD8 T cell clusters from clinical samples. *Nature* **649**, 467–476 (2026). <https://doi.org/10.1038/s41586-025-09754-w>.
- [45] Green, W. D., Gomez, A. G., Plotkin, A. L., Pratt, B. M., Merritt, E. F. et al. Enhancer-driven gene regulatory networks reveal transcription factors governing CD8 T cell adaptation and differentiation in the tumor microenvironment. *Immunity* **58**, 1725–1741 (2025). <https://doi.org/10.1016/j.immuni.2025.04.030>.
- [46] Tan, T., Wu, J., Si, C., Dai, S., Zhang, Y. et al. Chimeric contribution of human extended pluripotent stem cells to monkey embryos ex vivo. *Cell* **184**, 2020–2032.e14 (2021). <https://doi.org/10.1016/j.cell.2021.03.020>.
- [47] Shuldiner, E. G., Karmakar, S., Tsai, M. K., Hebert, J. D., Tang, Y. J. et al. Aging represses oncogenic KRAS-driven lung tumorigenesis and alters tumor suppression. *Nature Aging* **5**, 2263–2278 (2025). <https://doi.org/10.1038/s43587-025-00986-z>.
- [48] Melms, J. C., Biermann, J., Huang, H., Wang, Y., Nair, A. et al. A molecular single-cell lung atlas of lethal COVID-19. *Nature* **595**, 114–119 (2021). <https://doi.org/10.1038/s41586-021-03569-1>.