J Biochem Technol (2026) 17(4): 23-29
https://doi.org/10.51847/v2hFoOjdoQ
ISSN: 0974-2328

Multiple Molecular Explanations for the Same Cellular Phenotype: An Identifiability Framework for Distinguishing Mechanistic Equivalence from Causal Specificity in Computational Biology

Daniel Okoro*, Grace Okafor

Received: 17 August 2026 / Received in revised form: 17 September 2026, Accepted: 20 September 2026, Published online: 10 November 2026

Abstract

A cellular phenotype rarely identifies its molecular cause uniquely. Regulatory redundancy, resource competition, compensatory interactions, alternative network configurations, and incomplete measurement can allow distinct molecular states to generate observations that are difficult or impossible to distinguish under a given assay. This creates an identifiability problem that is deeper than ordinary uncertainty in parameter estimation: the scientific question is whether competing molecular explanations are distinguishable at all from the available observations and, if not, what additional evidence could separate them. This article develops a non-empirical methodological framework for analyzing that problem in computational biology. It distinguishes mechanistic equivalence, defined relative to specified observables and experimental conditions, from structural and practical identifiability in mathematical network models, and then links these concepts to perturbation-based causal discrimination. The framework treats model fit, predictive accuracy, perturbational responsiveness, and causal specificity as different evidentiary achievements rather than interchangeable indicators of mechanistic truth. Evidence from single-cell perturbation studies, gene-regulatory-network inference, multimodal profiling, dynamic-model calibration, and identifiability analysis shows that additional measurements can narrow an explanation space without necessarily producing a unique causal account. The central implication is that causal specificity should be claimed only when competing explanations generate discriminably different predictions under explicitly defined observations or interventions and those differences survive relevant measurement, model, and biological-context limitations. Where such discrimination fails, mechanistic equivalence should be retained rather than resolved by model complexity or interpretive preference.

Daniel Okoro*
Department of Biochemistry, University of Malawi, Zomba, Malawi.

Grace Okafor
Department of Molecular Biology, University of Malawi, Zomba, Malawi.

***E-mail:** daniel.okoro@unima.ac.mw

Keywords: Mechanistic equivalence, Causal specificity, Structural identifiability, Practical identifiability, Gene regulatory networks, Perturbation biology

Introduction

Computational models increasingly connect molecular measurements to regulatory structure, pathway activity, and cellular phenotype, but successful reproduction of an observed state does not imply that the underlying molecular explanation has been uniquely determined. In dynamic biological models, even well-fitting parameterizations can remain weakly identifiable when observations are incomplete, noisy, or censored, so apparently precise model behavior may coexist with broad uncertainty about the parameter combinations that generated it (Porthiyas *et al.*, 2024). At the causal level, the problem is more demanding because an association or predictive relationship must be distinguished from a claim that a particular molecular route is responsible for the observed phenotype. Causal machine-learning approaches in single-cell genomics therefore depend on assumptions about interventions, confounding, measurement and transport across experimental conditions (Tejada-Lapuerta *et al.*, 2025).

Large-scale perturbation experiments partly change this evidentiary situation. Genome-scale Perturb-seq demonstrates that direct genetic interventions can produce information-rich genotype–phenotype maps at single-cell resolution, providing contrasts that are unavailable from passive observation alone (Replogle *et al.*, 2022; Tejada-Lapuerta *et al.*, 2025). Yet perturbation does not make interpretation automatic. A transcriptional response after gene disruption may reflect the targeted relation, downstream adaptation, feedback, compensatory regulation, or properties of the assay and cellular context. Intervention thus increases the possibility of discrimination without guaranteeing a uniquely identified molecular mechanism.

The same distinction applies to computational prediction. Network-inference methods can generate experimentally testable hypotheses about regulatory control, and in-silico perturbations have been used to nominate factors whose subsequent manipulation changes cell identity in the predicted direction (Kamimoto *et al.*, 2023). Nevertheless, predictive sophistication is not a substitute for discriminatory evidence. Recent benchmarking

© 2026 The Author(s). This is an **Open Access** article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0).
https://creativecommons.org/licenses/by/4.0/deed.en

has shown that deep-learning models for perturbation-effect prediction do not necessarily outperform simpler linear baselines under evaluated conditions (Ahlmann-Eltze *et al.*, 2025). Moreover, harmonized analysis of dozens of single-cell perturbation datasets reveals extensive differences in perturbation type, assay, cell context and effect structure across studies (Peidli *et al.*, 2024). A mechanism that appears discriminable in one dataset may therefore remain unresolved in another.

This article addresses the inferential space between observing a phenotype and assigning a specific molecular cause. Its central claim is conditional: distinct molecular explanations should be treated as mechanistically equivalent whenever they remain indistinguishable with respect to a defined set of observations, experimental conditions and model assumptions. Causal specificity is warranted only after additional evidence separates the competing explanations in a manner that cannot be accounted for by the relevant alternatives. The following sections distinguish biological phenotypic underdetermination from formal structural and practical identifiability, and then prepare the basis for evaluating how perturbations and model discrimination can reduce, or fail to reduce, an initially non-unique explanation space.

Materials and Methods

Mechanistic Equivalence and the Problem of Phenotypic Underdetermination

Phenotypic underdetermination arises when the mapping from molecular state to measured cellular outcome is many-to-one. The relevant molecular alternatives need not be minor parameter variations within one mechanism. Resource competition in synthetic biochemical circuits illustrates how interactions mediated through shared cellular resources can produce system-level phenotypes that differ from the behavior expected when circuit components are considered independently (Al-Radhawi *et al.*, 2023). Consequently, an observed output can be compatible with more than one internal allocation or regulatory state. Such cases motivate a restricted definition of mechanistic equivalence: two explanations are equivalent only with respect to the observables and conditions under which they cannot currently be separated. They need not be biochemically identical.

A related form of underdetermination appears in gene-regulatory-network reconstruction. Benchmarking of network-inference algorithms using single-cell transcriptomic data has shown that different computational approaches recover regulatory structure with substantially different performance and that incomplete ground truth complicates evaluation (Pratapa *et al.*, 2020). Transfer-learning approaches can nonetheless learn representations that contain biologically useful information about regulatory relationships and perturbational responses (Theodoris *et al.*, 2023). These results support a distinction between usefulness and uniqueness. A representation can predict a phenotype or regulatory response without demonstrating that the network encoded by the model is the only molecular organization capable of producing it.

This distinction becomes particularly important as models integrate increasingly rich data. Foundation models trained on large single-cell corpora can capture gene-context relationships that support perturbation and network-related tasks (Theodoris *et al.*, 2023; Cui *et al.*, 2024). Single-cell multiome approaches can further constrain regulatory inference by combining expression, chromatin accessibility and external atlas information (Yuan & Duren, 2025). In parallel, computational analyses show that network topology itself affects the distribution and informativeness of gene-perturbation effects (Aguirre *et al.*, 2025). Additional modalities and priors therefore reduce some ambiguities, but increased information should be described as constraining an explanation set rather than automatically collapsing it to one mechanism.

The resulting analytical distinction is summarized below.

Table 1 summarizes molecular and inferential configurations that generate or resolve mechanistic equivalence.

Table 1. Molecular and inferential configurations that generate or resolve mechanistic equivalence

Biological or inferential configuration	What remains observationally indistinguishable	Principal source of underdetermination	Evidence capable of increasing discrimination	Interpretation permitted
Resource-coupled network states	Different internal allocations producing a similar measured output	Shared limiting resources and indirect coupling	Orthogonal measurements of resource use, component activity or state-dependent response	Equivalent only at the measured output level
Alternative regulatory-network reconstructions	Competing edge structures compatible with the same expression snapshot	Partial observability and algorithm-dependent inference	Independent perturbation, temporal measurements or additional molecular modalities	Multiple candidate regulatory explanations remain
Learned latent representations	Different internal representations yielding comparable predictive performance	Prediction does not uniquely specify biological structure	Mechanistically targeted intervention or external network validation	Predictive equivalence without causal equivalence
Network-topology-dependent response states	Distinct structures generating overlapping perturbation phenotypes	Compensation, hierarchy, modularity and downstream convergence	Perturbations chosen to generate divergent	Conditional mechanistic discrimination

			candidate-specific predictions	
Genetic-interaction states	Different single or combinatorial perturbations entering related phenotype regions	Nonlinear interaction and convergence of downstream programs	Higher-dimensional molecular readouts and factorial perturbation	Shared phenotype does not establish shared mechanism
Multimodal but phenotypically convergent states	Similar global outcomes with distinct RNA, protein or pathway-level changes	Compression of several biological levels into one endpoint	Joint transcriptomic, proteomic and intervention evidence	Orthogonal mechanisms can be distinguished despite endpoint similarity
Matched-target chemical and genetic perturbations	Mechanistically related interventions that do not map cleanly by morphology	Off-target activity, assay sensitivity and incomplete phenotype capture	Target-proximal readouts, orthogonal assays and context-matched validation	Morphological similarity or dissimilarity alone is insufficient

High-dimensional perturbation experiments make the many-to-one structure of phenotype generation more visible. Combinatorial CRISPR screens have revealed organized manifolds of genetic interactions in which different perturbations produce related cellular states while retaining distinguishable interaction structure (Norman *et al.*, 2019). Multimodal Perturb-CITE-seq provides a complementary example: simultaneous perturbational, transcriptomic and protein measurements in cancer immune-evasion models separated mechanisms that could not be reduced to a single resistance pathway (Frangieh *et al.*, 2021). These studies show why phenotypic convergence should not be confused with mechanistic identity. Distinct molecular routes can enter nearby or apparently similar endpoint states while remaining separable when the experiment interrogates the appropriate intermediate levels.

The converse is also important: interventions believed to act on related targets do not always produce readily recognizable phenotypic matches. Large-scale paired chemical and genetic morphological profiling has demonstrated that known mechanistic relationships can be difficult to recover from image-derived phenotypes alone (Chandrasekaran *et al.*, 2024). Thus, failure to observe phenotypic equivalence does not prove mechanistic difference, just as apparent equivalence does not prove mechanistic sameness. Both conclusions depend on the measurement surface through which the biological system is observed.

This creates a useful representation of the problem: distinct molecular configurations may remain separated internally while converging at a limited observation surface, and additional measurements may either reveal their separation or leave them unresolved.

Figure 1 summarizes phenotypic convergence across molecularly distinct network regimes.

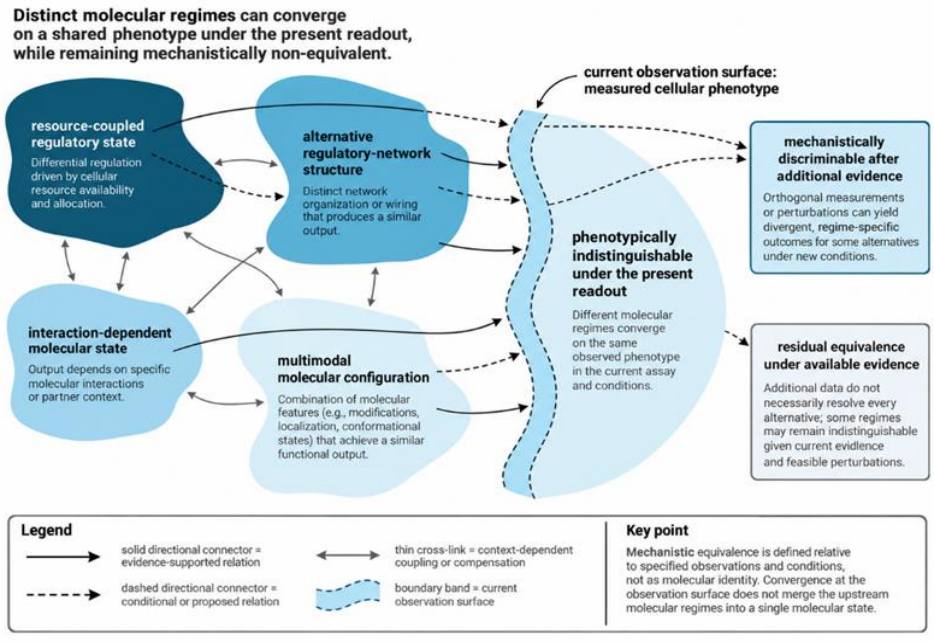


Figure 1. Phenotypic convergence across molecularly distinct network regimes

Phenotypic underdetermination therefore has at least two origins. Biology itself may contain redundancy, compensation, interaction and multiple routes to related endpoints; measurement may then compress those differences further. The methodological

consequence is that mechanistic equivalence cannot be adjudicated solely by asking whether two models reproduce the same phenotype. It requires asking whether their differences are observable in principle and, if they are, whether the available data are sufficiently informative to recover them. That question leads directly to structural and practical identifiability.

Structural and Practical Identifiability in Molecular Network Models

Structural and practical identifiability formalize different reasons why a model may fail to yield a unique explanation. Structural identifiability concerns whether parameters can, in principle, be uniquely determined from ideal, noise-free observations of specified model outputs. Practical identifiability concerns whether the available finite and noisy data actually constrain those parameters with useful precision (Wieland *et al.*, 2021). For the present argument, these concepts provide a disciplined way to separate two situations that may look similar empirically. One explanation may be non-unique because the model structure itself permits indistinguishable parameterizations; another may be distinguishable in principle but insufficiently resolved by the current experiment.

Structural analysis is consequently a property of the specified model–observation relationship rather than of goodness of fit. Scaling-based methods can identify parameter transformations that preserve model outputs and thereby reveal classes of structural non-identifiability (Castro & de Boer, 2020). Such tests, however, have defined domains of validity. Analyses of their limitations show that simple scaling procedures can miss non-identifiability produced by transformations outside the assumed class (Castro & de Boer, 2021). Passing a single diagnostic therefore should not be interpreted as proof that a biological mechanism is uniquely recoverable.

Practical identifiability introduces experimental design into the problem. Dynamic-model calibration protocols emphasize that optimization, uncertainty assessment and identifiability analysis must be considered together because a numerically satisfactory fit can coexist with poorly constrained parameters (Villaverde *et al.*, 2022). More recent computational work demonstrates how information-based analysis, regularization and experimental-design strategies can identify observations that improve practical recoverability in biological models (Wang & Hao, 2025). In the language of competing molecular explanations, this means that the evidentiary value of a measurement depends not simply on how much data are collected but on whether the measurement distinguishes predictions made by the candidates under consideration.

The computational infrastructure needed to perform these analyses has become increasingly mature. pyPESTO provides modular optimization and uncertainty-analysis methods for parameter-estimation problems (Schälte *et al.*, 2023), whereas PEtab formalizes interoperable definitions of models, parameters, observables and experimental conditions so that estimation problems can be reproduced across software environments (Schmiester *et al.*, 2021). AMICI provides scalable forward and adjoint sensitivity computation for large ordinary-differential-equation models (Fröhlich *et al.*, 2021). These developments improve reproducibility and computational tractability, but they do not change the logical status of the model. An optimization procedure can locate a best-supported parameter set within a specified model without demonstrating that the model is the only plausible molecular explanation of the data.

Model flexibility can make the distinction still sharper. Hybrid neural ordinary differential equations can represent incompletely specified biological dynamics, but increased functional flexibility also creates opportunities for unknown components to trade off against parameters intended to retain mechanistic meaning. Explicit identifiability analysis in such models has shown that apparently successful fitting can coexist with compromised recovery of mechanistic parameters (Giampiccolo *et al.*, 2024). At the opposite end of the scale, computational advances have made parameter estimation feasible for very large biochemical reaction networks (Fröhlich *et al.*, 2017). Scalability is scientifically valuable, but solving a larger optimization problem does not by itself establish uniqueness of the recovered biological explanation.

These considerations imply that identifiability should be assessed at more than one inferential layer. A parameter may be practically recoverable within a specified model while the model itself remains one member of a larger class of molecular explanations. Conversely, competing models may be structurally distinguishable in principle while the available experiment fails to generate observations with sufficient resolution to separate them. The methodological task is therefore not simply to label a model identifiable or non-identifiable. It is to specify what is being identified, relative to which observations, under which model assumptions, and at what biological scale.

For competing molecular explanations, the most consequential transition occurs when this diagnosis becomes prospective. Once an equivalence class has been identified, the next question is whether a new measurement, perturbation or experimental condition can force the candidates to make different observable predictions. Structural identifiability establishes whether separation is possible under the model and observation scheme; practical identifiability establishes whether the attainable data can support that separation. Neither alone demonstrates causal specificity. Causal specificity requires a further step in which competing explanations are confronted with interventions or observations that were chosen because their predicted consequences diverge.

Distinguishing Causal Specificity through Perturbation and Model Discrimination

Perturbation can convert an observational equivalence problem into a discrimination problem only when competing explanations predict different consequences. Models such as CPA, GEARS and PerturbNet show that cellular responses to previously unmeasured chemical or genetic perturbations can be predicted from learned structure (Lotfollahi *et al.*, 2023; Roohani *et al.*, 2024; Yu *et al.*, 2025). Their scientific value for identifiability lies in prospectively ranking interventions that should separate candidates. Predictive success alone, however, does not identify the causal mechanism.

Controlled simulation offers a complementary strategy. GRouNdGAN uses an explicit gene-regulatory structure to generate causal single-cell simulations and in-silico knockouts (Zinati *et al.*, 2024). Such systems can ask whether candidate mechanisms would remain observationally equivalent after a proposed intervention, but the encoded causal graph is an assumption of the simulation rather than an independently discovered truth. Experimental validation of network-derived perturbation predictions, as demonstrated in cell-identity studies, is therefore essential when causal interpretation is intended (Kamimoto *et al.*, 2023).

The evidentiary standard becomes stronger when intervention is combined with orthogonal measurement. Causal single-cell analysis requires assumptions about intervention, confounding and transport across contexts (Tejada-Lapuerta *et al.*, 2025). Multimodal Perturb-CITE-seq illustrates how RNA, protein and perturbational evidence can distinguish mechanistically different routes associated with a related endpoint (Frangieh *et al.*, 2021). Causal specificity is therefore best understood as a comparative claim: one explanation gains support because a discriminating intervention produces the outcome it predicts while credible rivals fail under the same conditions.

Model complexity cannot substitute for this comparison. Deep perturbation predictors may fail to outperform simple baselines under some benchmarks (Ahlmann-Eltze *et al.*, 2025), while morphology may fail to recover known chemical–genetic relationships even when perturbations have related targets (Chandrasekaran *et al.*, 2024). A useful discriminator must therefore produce candidate-specific differences at a measurement scale capable of resolving them.

Figure 2 summarizes experimental separability of competing molecular explanations across an identifiability boundary.

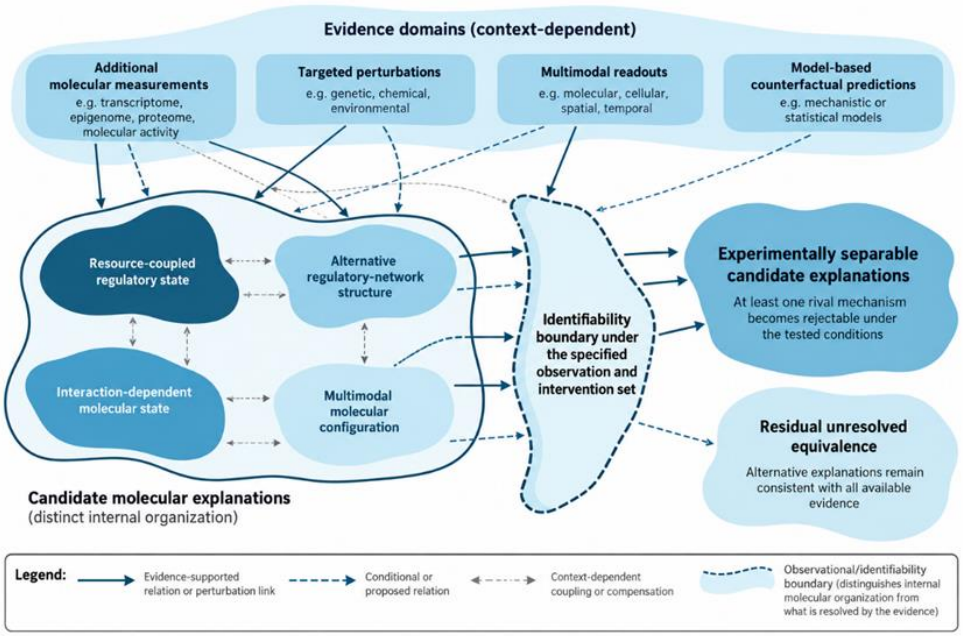


Figure 2. Experimental separability of competing molecular explanations across an identifiability boundary

Perturbation informativeness remains context dependent. Genome-scale Perturb-seq demonstrates how broad intervention can expose otherwise hidden genotype–phenotype structure (Replogle *et al.*, 2022). Cross-dataset analyses show that perturbation effects vary with assay and biological context (Peidli *et al.*, 2024). Network topology can also change which perturbations are informative (Aguirre *et al.*, 2025). The practical objective is therefore to select observations that maximize disagreement among plausible explanations while preserving the biological context in which the causal claim will be made.

An Identifiability Framework for Competing Molecular Explanations

The proposed framework begins by declaring the candidate explanations, the observable phenotype, and the level at which uniqueness is being claimed. Structural and practical identifiability should then be evaluated separately (Wieland *et al.*, 2021). Where experimental design can improve recoverability, the relevant question is whether an added observation reduces ambiguity among the specified candidates rather than merely tightening uncertainty within one model (Wang & Hao, 2025).

Mechanistic equivalence is retained when different explanations remain observationally indistinguishable under the current evidence. Resource-coupled circuits provide one example of how distinct internal interactions can yield similar system-level outputs (Al-Radhawi *et al.*, 2023), while network-structure analyses show that different architectures alter the distribution of perturbation effects (Aguirre *et al.*, 2025). Equivalence should therefore be stated relative to a defined phenotype, assay and condition.

A candidate measurement or intervention becomes discriminating only when rival explanations predict observably different outcomes. Explicit problem specification helps make those contrasts reproducible (Schmiester *et al.*, 2021), and predictive perturbation models can nominate prospective tests (Lotfollahi *et al.*, 2023). Such predictions do not count as causal confirmation until the resulting observations are obtained and compared with the alternatives.

Causal specificity is consequently bounded rather than absolute. Causal-learning theory emphasizes assumptions and unmeasured confounding (Tejada-Lapuerta *et al.*, 2025), while causal simulation shows how much discrimination depends on the structure supplied to the model (Zinati *et al.*, 2024). The strongest defensible conclusion is that one explanation is more specifically supported than named rivals under stated observations, interventions and biological conditions.

Results and Discussion

Interpretive Boundaries and Implications for Computational Biology

A central implication is that predictive sophistication and causal discrimination should be reported separately. Benchmark evidence shows that complex perturbation models do not automatically improve prediction over simpler alternatives (Ahlmann-Eltze *et al.*, 2025), and GRN inference performance remains dependent on assumptions, datasets and available ground truth (Pratapa *et al.*, 2020). Computational complexity is therefore not an evidentiary substitute for a discriminating experiment.

Transfer across contexts introduces another boundary. Harmonized perturbation datasets retain substantial biological and technical heterogeneity (Peidli *et al.*, 2024), while multiome-based network inference can depend on external atlas information and priors (Yuan & Duren, 2025). An observation that separates mechanisms in one cell type, assay or state may fail to do so elsewhere, so identifiability should be re-evaluated in the target context.

Model structure can also mimic or conceal practical uncertainty. Simple structural-identifiability tests have documented blind spots (Castro & de Boer, 2021), and flexible hybrid neural models can permit unknown components to trade off against mechanistic parameters (Giampiccolo *et al.*, 2024). Ambiguity should therefore not be assigned automatically to noise or insufficient sample size when the model class itself may permit alternative explanations.

Finally, generalization to unseen perturbations remains distinct from causal specificity. GEARS demonstrates predictive capacity for novel multigene perturbations (Roohani *et al.*, 2024). PerturbNet similarly predicts distributions of responses to previously unseen perturbations (Yu *et al.*, 2025). GRouNdGAN provides a different capability by generating controlled simulations conditioned on an explicit regulatory structure (Zinati *et al.*, 2024). These methods are most informative when used to design tests that expose disagreement among mechanisms. When plausible alternatives survive the available observations, computational biology should preserve that uncertainty rather than convert predictive adequacy into an unwarranted causal claim.

Conclusion

The same cellular phenotype can support multiple molecular explanations because biological systems contain redundancy, compensation and alternative network routes, while assays observe only part of that internal structure. The appropriate response is not to select the most elaborate or best-fitting explanation by default. Mechanistic equivalence should be retained until competing accounts generate separable predictions under relevant observations or perturbations. Structural identifiability establishes whether separation is possible in principle; practical identifiability asks whether attainable data can support it; causal specificity requires successful discrimination among credible rivals.

Failure to discriminate is itself interpretable: it identifies the current observation–intervention set as insufficient for stronger causal attribution. Progress therefore consists not merely in improving fit, but in designing evidence that changes which explanations remain scientifically defensible. The framework thus treats unresolved alternatives as an informative scientific state rather than a defect to be hidden by model complexity.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

Aguirre, M., Spence, J. P., Sella, G., & Pritchard, J. K. (2025). Gene regulatory network structure informs the distribution of perturbation effects. *PLOS Computational Biology, 21*(9), e1013387. doi:10.1371/journal.pcbi.1013387

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

Al-Radhawi, M. A., Del Vecchio, D., & Sontag, E. D. (2023). Identifying competition phenotypes in synthetic biochemical circuits. *IEEE Control Systems Letters, 7*, 211-216. doi:10.1109/LCSYS.2022.3186840

Castro, M., & de Boer, R. J. (2020). Testing structural identifiability by a simple scaling method. *PLOS Computational Biology, 16*(11), e1008248. doi:10.1371/journal.pcbi.1008248

Castro, M., & de Boer, R. J. (2021). The limitations, dangers, and benefits of simple methods for testing identifiability. *PLOS Computational Biology, 17*(10), e1009425. doi:10.1371/journal.pcbi.1009425

Chandrasekaran, S. N., Cimini, B. A., Goodale, A., Miller, L., Kost-Alimova, M., Jamali, N., Doench, J. G., Fritchman, B., Skepner, A., Melanson, M., et al. (2024). Three million images and morphological profiles of cells treated with

matched chemical and genetic perturbations. *Nature Methods*, *21*(6), 1114-1121. doi:10.1038/s41592-024-02241-6

Cui, H., Wang, C., Maan, H., Pang, K., Luo, F., Duan, N., & Wang, B. (2024). scGPT: Toward building a foundation model for single-cell multi-omics using generative AI. *Nature Methods*, *21*(8), 1470-1480. doi:10.1038/s41592-024-02201-0

Frangieh, C. J., Melms, J. C., Thakore, P. I., Geiger-Schuller, K. R., Ho, P., Luoma, A. M., Cleary, B., Jerby-Arnon, L., Malu, S., Cuoco, M. S., et al. (2021). Multimodal pooled Perturb-CITE-seq screens in patient models define mechanisms of cancer immune evasion. *Nature Genetics*, *53*(3), 332-341. doi:10.1038/s41588-021-00779-1

Fröhlich, F., Kaltenbacher, B., Theis, F. J., & Hasenauer, J. (2017). Scalable parameter estimation for genome-scale biochemical reaction networks. *PLOS Computational Biology*, *13*(1), e1005331. doi:10.1371/journal.pcbi.1005331

Fröhlich, F., Weindl, D., Schälte, Y., Pathirana, D., Paszkowski, Ł., Lines, G. T., Stapor, P., & Hasenauer, J. (2021). AMICI: High-performance sensitivity analysis for large ordinary differential equation models. *Bioinformatics*, *37*(20), 3676-3677. doi:10.1093/bioinformatics/btab227

Giampiccolo, S., Reali, F., Fochesato, A., Iacca, G., & Marchetti, L. (2024). Robust parameter estimation and identifiability analysis with hybrid neural ordinary differential equations in computational biology. *npj Systems Biology and Applications*, *10*(1), 139. doi:10.1038/s41540-024-00460-3

Kamimoto, K., Stringa, B., Hoffmann, C. M., Jindal, K., Solnica-Krezel, L., & Morris, S. A. (2023). Dissecting cell identity via network inference and in silico gene perturbation. *Nature*, *614*(7949), 742-751. doi:10.1038/s41586-022-05688-9

Lotfollahi, M., Klimovskaia Susmelj, A., De Donno, C., Hetzel, L., Ji, Y., Ibarra, I. L., Srivatsan, S. R., Naghipourfar, M., Daza, R. M., Martin, B., et al. (2023). Predicting cellular responses to complex perturbations in high-throughput screens. *Molecular Systems Biology*, *19*(6), e11517. doi:10.15252/msb.202211517

Norman, T. M., Horlbeck, M. A., Replogle, J. M., Ge, A. Y., Xu, A., Jost, M., Gilbert, L. A., & Weissman, J. S. (2019). Exploring genetic interaction manifolds constructed from rich single-cell phenotypes. *Science*, *365*(6455), 786-793. doi:10.1126/science.aax4438

Peidli, S., Green, T. D., Shen, C., Gross, T., Min, J., Garda, S., Yuan, B., Schumacher, L. J., Taylor-King, J. P., Marks, D. S., et al. (2024). scPerturb: Harmonized single-cell perturbation data. *Nature Methods*, *21*(3), 531-540. doi:10.1038/s41592-023-02144-y

Porthiyas, J., Nussey, D., Beauchemin, C. A. A., Warren, D. C., Quirouette, C., & Wilkie, K. P. (2024). Practical parameter identifiability and handling of censored data with Bayesian inference in mathematical tumour models. *npj Systems Biology and Applications*, *10*(1), 89. doi:10.1038/s41540-024-00409-6

Pratapa, A., Jalihal, A. P., Law, J. N., Bharadwaj, A., & Murali, T. M. (2020). Benchmarking algorithms for gene regulatory network inference from single-cell transcriptomic data. *Nature Methods*, *17*(2), 147-154. doi:10.1038/s41592-019-0690-6

Replogle, J. M., Saunders, R. A., Pogson, A. N., Hussmann, J. A., Lenail, A., Guna, A., Mascibroda, L., Wagner, E. J., Adelman, K., Lithwick-Yanai, G., et al. (2022). Mapping information-rich genotype-phenotype landscapes with genome-scale Perturb-seq. *Cell*, *185*(14), 2559-2575.e28. doi:10.1016/j.cell.2022.05.013

Roohani, Y., Huang, K., & Leskovec, J. (2024). Predicting transcriptional outcomes of novel multigene perturbations with GEARS. *Nature Biotechnology*, *42*(6), 927-935. doi:10.1038/s41587-023-01905-6

Schälte, Y., Fröhlich, F., Jost, P. J., Vanhoefer, J., Pathirana, D., Stapor, P., Lakrisenko, P., Wang, D., Raimúndez, E., Merkt, S., et al. (2023). pyPESTO: A modular and scalable tool for parameter estimation for dynamic models. *Bioinformatics*, *39*(11), btad711. doi:10.1093/bioinformatics/btad711

Schmiester, L., Schälte, Y., Bergmann, F. T., Camba, T., Dudkin, E., Egert, J., Fröhlich, F., Fuhrmann, L., Hauber, A. L., Kemmer, S., et al. (2021). PEtab—Interoperable specification of parameter estimation problems in systems biology. *PLOS Computational Biology*, *17*(1), e1008646. doi:10.1371/journal.pcbi.1008646

Tejada-Lapuerta, A., Bertin, P., Bauer, S., Aliee, H., Bengio, Y., & Theis, F. J. (2025). Causal machine learning for single-cell genomics. *Nature Genetics*, *57*(4), 797-808. doi:10.1038/s41588-025-02124-2

Theodoris, C. V., Xiao, L., Chopra, A., Chaffin, M. D., Al Sayed, Z. R., Hill, M. C., Mantineo, H., Brydon, E. M., Zeng, Z., Liu, X. S., et al. (2023). Transfer learning enables predictions in network biology. *Nature*, *618*(7965), 616-624. doi:10.1038/s41586-023-06139-9

Villaverde, A. F., Pathirana, D., Fröhlich, F., Hasenauer, J., & Banga, J. R. (2022). A protocol for dynamic model calibration. *Briefings in Bioinformatics*, *23*(1), bbab387. doi:10.1093/bib/bbab387

Wang, S., & Hao, W. (2025). A systematic computational framework for practical identifiability analysis in mathematical models arising from biology. *Advanced Science*, *12*(35), e04346. doi:10.1002/advs.202504346

Wieland, F.-G., Hauber, A. L., Rosenblatt, M., Tönsing, C., & Timmer, J. (2021). On structural and practical identifiability. *Current Opinion in Systems Biology*, *25*, 60-69. doi:10.1016/j.coisb.2021.03.005

Yu, H., Qian, W., Song, Y., & Welch, J. D. (2025). PerturbNet predicts single-cell responses to unseen chemical and genetic perturbations. *Molecular Systems Biology*, *21*(8), 960-982. doi:10.1038/s44320-025-00131-3

Yuan, Q., & Duren, Z. (2025). Inferring gene regulatory networks from single-cell multiome data using atlas-scale external data. *Nature Biotechnology*, *43*(2), 247-257. doi:10.1038/s41587-024-02182-7

Zinati, Y., Takiddeen, A., & Emad, A. (2024). GRouNdGAN: GRN-guided simulation of single-cell RNA-seq data using causal generative adversarial networks. *Nature Communications*, *15*(1), 4055. doi:10.1038/s41467-024-48516-6