

Are we ready for causal discovery in biological systems using deep learning?

Hock Chuan Yeo¹ and Kumar Selvarajoo^{1,2,3,4,*}

¹Bioinformatics Institute (BII), Agency for Science, Technology and Research (A*STAR), 30 Biopolis Street, Matrix #07-01, Singapore 138761, Singapore

²Synthetic Biology Translational Research Program, Yong Loo Lin School of Medicine, National University of Singapore (NUS), 10 Medical Drive, Singapore 117597, Singapore

³School of Biological Sciences, Nanyang Technological University (NTU), 60 Nanyang Drive, SBS-01s-45, Singapore 637551, Singapore

⁴Synthetic Biology for Clinical and Technological Innovation (SynCTI), National University of Singapore (NUS), 28 Medical Drive, Centre for Life Sciences #02-07, Singapore 117456, Singapore

*Corresponding author. Bioinformatics Institute (BII), Agency for Science, Technology and Research (A*STAR), 30 Biopolis Street, Matrix #07-01, Singapore 138761, Singapore.
E-mail: kumar_selvarajoo@bii.a-star.edu.sg

Abstract

The field of causal discovery has advanced considerably over the past three decades, in terms of perspectives, computational methods, and foundational concepts. Nevertheless, their application to biological systems that are commonly found in nature (i.e. large-scale, self-regulating), continues to face significant challenges. In this regard, we highlight emerging approaches that go beyond the traditional assumption of global acyclicity, instead leveraging efficient and scalable neural methods to infer pairwise causal relationships, directly from the data. Nonetheless, there remains five key technological hurdles, which must be overcome, to realize the deeper understanding and stronger inference biological causal networks promise.

Keywords systems biology, causal discovery, Bayesian networks, directed acyclicity, transformer, neural network

Advancements in causal discovery for biological systems

Causality remains a foundational area of inquiry across diverse scientific domains [1, 2]. Nonetheless, a series of advances in the last three decades, as to elemental concepts, perspective, and computation, have increasingly permitted large-scale causal discovery in biological systems. Emerging research is now pivoting toward powerful, efficient, and scalable neural approaches that can map individual causal edges between variables from input data, in lieu of the hard task of discerning entire candidate networks. Such capability, central to higher-order intelligence (Supplementary Perspectives 1 and 2), however remains an open question at the heart of biology, one best interpreted through the lens of developmental trends thus far. To this end, the techniques proposed to date may be categorized as follows, while acknowledging the field's inherent complexity and disorderly evolution:

1. Those fundamentally requiring Bayesian statistical formalism for causal inference
 - These methods [3–24] represent diverse, yet data and compute-hungry methods for searching plausible networks

in entirety, and as such, face the intractable issues of a super-exponentially growing search space (2^{n^2}) (as the number of variables n increases), coupled with the challenge of score-equivalent but graphically distinct solutions.

- Such classical techniques are subsequently superseded by innovatively embedding the Bayesian acyclicity constraint into the continuous objective search function [25], thereby enabling efficient exploration by powerful optimization algorithms. By notably allowing for stochastic gradient descent algorithms, the new formulation further unlocks neural approaches with novel capabilities [26–40]. Put together, the development enhances potential scalability to larger systems.
- Many of these improvised methods also circumvent the fundamental limitation of using observations to infer causality, by leveraging on intervention data (Supplementary Perspective 2). Although some even evolved to use supervised learning to achieve more efficacious discovery (e.g. [33, 34]), the vast majority hold onto the assumption of *global directed acyclicity* (absence of feedback loops), an Achilles' heel for causal discovery in self-regulating biological systems. Only

Received: September 12, 2025. **Revised:** December 29, 2025. **Accepted:** February 26, 2026

© The Author(s) 2026. Published by Oxford University Press.

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial License (<https://creativecommons.org/licenses/by-nc/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited. For commercial re-use, please contact reprints@oup.com for reprints and translation rights for reprints. All other permissions can be obtained through our RightsLink service via the Permissions link on the article page on our site—for further information please contact journals.permissions@oup.com.

Efficient Neural Causal DiscOvery (ENCO) [29] could discard the assumption under certain mild conditions.

2. Others rephrasing causal discovery as a simpler problem on binary classification

- Given the criticality of molecular feedback loops for cellular functions, such as timing and decision-making processes (Supplementary Perspective 2), some later lineages escape the clutch of *global directed acyclicity* by abandoning the Bayesian notion at its core. Rather, they focus on using supervised learning to discriminate the presence or absence of directed edge for all pairs of variables [2, 41–44]. These techniques as such need not explore a super-exponentially growing search space (as systems get larger) and yet can provide more reliable inference due to supervision.
- To exemplify this development, a recently-developed method, *deep discriminative causal learning* (D²CL) [2], has been reported to be *scalable up to 50 000 variables* for artificial systems, comparing favorably to that of classical state-of-the-art methods, e.g. 48 variables for *Structural Discovery from Interventions* (SDI) [30], 100 for *Differentiable Causal Discovery with Interventions* (DCDI) [27], and 1000 for *Differentiable Causal Discovery of Factor Graphs* (DCD-FG) [28]. Furthermore, D²CL *requires as little as 10³ samples* (both observation and interventional) to do so, in contrast to DCD-FG requiring at least 5×10^4 such samples. Along the same line, DCDI needs at least a magnitude more samples (10⁶) and interventional experiments on all nodes; ENCO also has sample size requirements and system scalability within the range of its peers [2].
- Inference-wise, D²CL surpasses ENCO, DCD-FG, and *Scalable Causal Learning* (SCL) [42], over a broad range of signal-to-noise ratio, as to both direct and ancestral causal effects in a nonlinear 1500-variables system. D²CL likewise outperforms several first-generation classical methods.
- To be clear, D²CL bests other techniques in the regime that it has been designed for (i.e. where biomedical applications lay):
 - one in which data dimensionality far exceeds sample size
 - but with existing knowledge of a subset of causal relationships
- D²CL aptly leverages the latter for supervised learning to mitigate the curse of the former, using powerful and highly scalable neural approaches. As a trade-off, D²CL does not seek to infer the specific data-generating model, unlike Bayesian methods that require many more samples than data dimensions to do so. Put differently, the latter are designed for a different use case: where system dimensions are markedly less than the number of available samples, and little is known about its cause and effect.
- D²CL however still uses assumptions that may be inconsistent with the behavior of biological systems. It assumes that each causal factor operates independently, and that by varying it, a response may be observed in its targets. Furthermore, the factors are presumed to act similarly to enable the learning of unknown causal edges from the knowledge of others. While it remains unclear as to the prediction

cost of these assumptions, the authors have demonstrated good performance across both yeast gene deletion datasets (ROCAUC 0.79–0.85 across scenarios) and human CRISPR-based intervention datasets (ROCAUC 0.65–0.73). Among the compared methods, only SCL (of the same category as D²CL) achieves similar results on the yeast dataset (ROCAUC 0.75–0.77).

- While D²CL focuses on inferring causal relationships in a system by using data and causal knowledge collected on, another emergent algorithm, *Causal Structure Induction via Attention* (CSiVA) [44], leverages transformer to supervise learning from diverse synthetic systems to unseen or even natural ones (i.e. meta-learning). Like D²CL, CSiVA has been shown to outperform DCDI and ENCO in inference (e.g. for up to 80 variables, different graph densities, and linear and nonlinear datasets) [44].
- Still, the algorithm relies on intervention data and assumes the full network is identifiable upon sufficiently rich interventions on each causal factor. This again implies that the causal factors are to act independently to some extent, so that their manipulation generates an observable change in their targets. Yet, even when the synthetic data in case studies are designed as such, CSiVA requires tremendous amount of data: for a system of only 80 variables, CSiVA requires a diversity of 40 000 networks and 1500 samples for each of them. In addition, the topology of the network-to-be-inferred must be sufficiently characterized to generate training data with matching distribution. On this, the authors reported instances of difficulties in generalizing learning across different types of networks. Larger and denser graphs, alongside some data-generating functions, are also found to be harder to learn from.
- Overall, CSiVA still represents a novel approach by harnessing attention mechanism [45] to pick up variations among samples and among factors that inform specific causation. In this regard, there is potential to pick up more causal information, such as among timepoints and among data modality with valid mechanisms (Fig. 1b).
- Compute-wise, although CSiVA and D²CL may take longer to train than ENCO takes to infer for small systems (<1000 variables), the time spent can be amortized over multiple usages while inference only takes minutes [44]. (DCDI takes many more hours to compute).

Moving forward

The development of CSiVA and D²CL signifies that technological breakthroughs may be on the horizon, as deep learning classifiers are expected to enable causal discovery in many more biological systems, that are commonly found in nature (i.e. those that are of a much larger scale and self-regulating), than is possible previously. However, there is currently a dearth of rigorous work on the theoretical limitations of emerging deep learning techniques, the boundary conditions of their application, or potential failure modes, as computer scientists are only beginning to explore their applicability. Furthermore, there are at least five technological hurdles that must be overcome before humanity can realize the deeper understanding and stronger inference (Supplementary Perspectives 1–6) that causality promises (Fig. 1a):




Figure 1 Advances necessary to enable meaningful, widespread applications of causal discovery tools for biological systems. (a) Key technological hurdles that must be overcome to realize the deeper understanding and stronger inference biological causal networks promise. These include two with significant recent progress (i.e. regarding feedback loops and scalability) and another five going forward. (b) Variations that may inform causal inference: Among factors within single modality (Type I), among factors of different data modality with regulatory relationship consistent with the Central Dogma (Type II), individual factors changes across samples (Type III), and individual factors changes across timepoints (Type IV). An example of each type is depicted in the diagram. Currently, only Type I and III have been implemented [44]. Here is the grand challenge: To develop a general architecture that can flexibly customize to the number and nature of interaction among data modalities.

- There is a need to distinguish between inhibiting and activating interactions to support adjacent analysis (e.g. control theory, dynamic modeling), like in many other scientific domains. This will allow for the derivation of sharper and more biological insights, as articulated in [Supplementary Perspective 2](#). Notably, such a large-scale approach is non-existent, perhaps due to the additional complexity of the problem.
- In addition, a focus on identifying direct causation via elemental interactions (e.g. physico-chemical [46]) would prepare the ground for inferring the mechanisms involved (e.g. reactions) and their governing equations. See Supplementary on the importance of using first-principles for developing a *true* cellular digital twin [47] ([Supplementary Perspective 2](#)), and an example of how mechanism informs the right corrective action ([Supplementary Perspective 5](#)). In this regard, an ambitious effort is ongoing in DeepMind to predict the ‘interactions of all of life’s molecules’. However, it is unclear what kind of approach or data would be universally applicable across the different classes of biomolecules, and if such data is already available or could be generated in sufficient quantities, given the high contextuality and dynamicity of the interactions.
- For cellular systems, the grand challenge is to use multi-omics data to infer the physico-chemical regulatory mechanisms, in line with the Central Dogma and the order of events ([Fig. 1b](#)). As computer scientists are still exploring the theoretical and practical limits of causal inference using single-modality data, very limited effort has been made to extend these methods to multi-omics data. Such an approach will however allow for the construction of a virtual cell, whose implications are elegantly described by Demis Hassabis in one of his LinkedIn post:
‘Imagine running “in silico” experiments orders of magnitude faster than in a wet lab. Scientists could rapidly test hypotheses, model complex pathways and see how a drug affects a cell. It would be an incredible boon not only for fundamental biology but also for medicine.’
 Undoubtedly, the challenge is enormous. However, if we are to turn the problem on its head, embedding known regulatory structures as first-principle constraints [48] may effectively ease issues of sample size requirement, model identifiability, noise robustness, and overfitting. The core problem may however be perennial, as larger, more complex, and heterogeneous biological networks will still demand even more and better data (and compute), putting a strain on infrastructure, and demanding further innovations such as in quantum computing to arrive at the solution efficiently. (We note somewhat related methods have been proposed to *explore* generic mediators of causation, which can use but however does not distinguish between the different types of omics data, e.g. [49, 50]).
- Synthetic data should strive for contextual realism by closely emulating fundamental laws and real noises (e.g. heteroscedastic noise for biological data) to foster confidence [51]. Also, as methods developed to date are based on data generated from artificial models as the ground truth, little is known about their potential applicability and generalizability to real-world problems. Compounding the issue, many natural laws are still poorly understood, creating a potential chicken-and-egg problem.
- Current sample requirements for efficacious causal prediction (> 1000 [2, 44]) far exceed what most biological studies can provide [48], except in the context of some single-cell omics. Yet even these are largely observational rather than interventional, aside from a few niche cases [52]. Coupled with the field’s infancy,

this means that no clinical applications or actionable biological hypotheses have yet emerged. To address the issue, advancement in the generation of omics datasets must be complemented with better rationalization and optimization of data requirement, such as by applying the *Design-Build-Test-Learn* cycle ([Fig. 1a](#)) and related recommendations given by Radivojević *et al.* [53].

We believe the ultimate solution (and hurdle) would have to effectively address several of these challenges at once.

Key Points

- Causal discovery methods split between classical Bayesian approaches and scalable neural classifiers more suited for biology (e.g. causal structure induction via attention, deep discriminative causal learning).
- Five key hurdles for cellular systems: distinguishing activating and inhibiting interactions, inferring direct physico-chemical causation, multi-modal integration, synthetic data, and noise realism, as well as sample size optimization.
- Uniting systems biology expertise with machine learning will drive the next leap in the field.

Acknowledgments

HCY & KS opined on the subject matter. HCY wrote the article. KS supervised & edited the article.

Supplementary material

[Supplementary material](#) is available at *Briefings in Bioinformatics* online.

Conflict of interest

None declared.

Funding

This work was supported by the core research budget of Bioinformatics Institute, Agency for Science, Technology and Research (A*STAR).

Data availability

There is no new data associated with this article.

References

1. Peters J, Janzing D, Schölkopf B. In: Bach F (ed.), *Elements of Causal Inference: Foundations and Learning Algorithms*. Cambridge, MA: MIT Press, 2017.
2. Lagemann K, Lagemann C, Taschler B *et al.* Deep learning of causal structures in high dimensions under data limitations. *Nature Mach Intell* 2023;**5**:1306–16.
3. Chickering DM. Optimal structure identification with greedy search. *J Mach Learn Res* 2003;**3**:507–54.
4. Cooper GF, Yoo C. Causal discovery from a mixture of experimental and observational data. In: *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*. Laskey KB, Prade H (eds.),

- pp. 116–25. Stockholm, Sweden: Morgan Kaufmann Publishers Inc., 1999.
5. Heckerman D, Geiger D, Chickering DM. Learning Bayesian networks: The combination of knowledge and statistical data. *Mach Learn* 1995;**20**:197–243.
 6. Tsamardinos I, Brown LE, Aliferis CF. The max-min hill-climbing Bayesian network structure learning algorithm. *Mach Learn* 2006;**65**:31–78.
 7. Huang B, Zhang K, Lin Y *et al*. Generalized score functions for causal discovery. In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. Guo Y, Farooq F (eds.), pp. 1551–60. London, United Kingdom: Association for Computing Machinery, 2018.
 8. Shimizu S, Hoyer PO, Hyvärinen A *et al*. A linear non-gaussian acyclic model for causal discovery. *J Mach Learn Res* 2006;**7**:2003–30.
 9. Spirtes P, Glymour C, Scheines R. In: Bach F (ed.), *Causation, Prediction, and Search*. Cambridge, MA: MIT Press, 2001.
 10. Maathuis MH, Kalisch M, Bühlmann P. Estimating high-dimensional intervention effects from observational data. *Annals Stat* 2009;**37**:3133–64.
 11. Hauser A, Bühlmann P. Characterization and greedy learning of interventional Markov equivalence classes of directed acyclic graphs. *J Mach Learn Res* 2012;**13**:2409–64.
 12. Colombo D, Maathuis MH, Kalisch M *et al*. Learning high-dimensional directed acyclic graphs with latent and selection variables. *Annals Stat* 2012;**40**:294–321.
 13. Hoyer PO, Janzing D, Mooij J *et al*. Nonlinear causal discovery with additive noise models. In: *Proceedings of the 22nd International Conference on Neural Information Processing Systems*. Koller D, Schuurmans D, Bengio Y *et al*. (eds.), pp. 689–96. Vancouver, BC, Canada: Curran Associates Inc., 2008.
 14. Peters J, Mooij JM, Janzing D *et al*. Identifiability of causal graphs using functional models. In: *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence*. Cozman F, Pfeffer A (eds.), pp. 589–98. Barcelona, Spain: AUAI Press, 2011.
 15. Daniušis P, Janzing D, Mooij J *et al*. Inferring deterministic causal relations. In: *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence*. Grünwald P, Spirtes P (eds.), pp. 143–50. Catalina Island, CA: AUAI Press, 2010.
 16. Monti RP, Zhang K, Hyvärinen A. Causal discovery with general non-linear relationships using non-linear ICA. In: Adams RP, Gogate V (eds.), *Uncertainty in Artificial Intelligence*, pp. 186–95. Virtual: PMLR, 2020.
 17. Sun X, Janzing D, Schölkopf B *et al*. A kernel-based causal learning algorithm. In: *Proceedings of the 24th International Conference on Machine Learning*. Ghahramani Z (ed.), pp. 855–62. Corvallis, OR, USA: Association for Computing Machinery, 2007.
 18. Zhang K, Peters J, Janzing D *et al*. Kernel-based conditional independence test and application in causal discovery. In: *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence*. Cozman F, Pfeffer A (eds.), pp. 804–13. Barcelona, Spain: AUAI Press, 2011.
 19. Eaton D, Murphy K. Bayesian structure learning using dynamic programming and MCMC. In: *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence*. Parr R, van der Gaag L (eds.), pp. 101–8. Vancouver, BC, Canada: AUAI Press, 2007.
 20. Rojas-Carulla M, Schölkopf B, Turner R *et al*. Invariant models for causal transfer learning. *J Mach Learn Res* 2018;**19**:1–34.
 21. Ghassami AE, Salehkaleybar S, Kiyavash N *et al*. Learning causal structures using regression invariance. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*. Guyon I, Von Luxburg U, Bengio S *et al*. (eds.), pp. 3015–25. Long Beach, CA, USA: Curran Associates Inc., 2017.
 22. Peters J, Bühlmann P, Meinshausen N. Causal inference by using invariant prediction: Identification and confidence intervals. *J R Stat Soc Series B Stat Methodology* 2016;**78**:947–1012.
 23. Budhathoki K, Vreeken J. Causal inference by stochastic complexity. *Knowl Inf Syst* 2018;**56**:657–91.
 24. Mitrovic J, Sejdinovic D, Teh YW. Causal inference via kernel deviance measures. In: *Proceedings of the 32nd International Conference on Neural Information Processing Systems*. Bengio S, Wallach H, Larochelle H *et al*. (eds.), pp. 6986–94. Montréal, Canada: Curran Associates Inc., 2018.
 25. Zheng X, Aragam B, Ravikumar P *et al*. DAGs with NO TEARS: Continuous optimization for structure learning. In: *Proceedings of the 32nd International Conference on Neural Information Processing Systems*. Bengio S, Wallach H, Larochelle H *et al*. (eds.), pp. 9492–503. Montréal, Canada: Curran Associates Inc., 2018.
 26. Yu Y, Chen J, Gao T *et al*. DAG-GNN: DAG structure learning with graph neural networks. In: Chaudhuri K, Salakhutdinov R (eds.), *International Conference on Machine Learning*, pp. 7154–63. PMLR: Long Beach, CA, 2019.
 27. Brouillard P, Lachapelle S, Lacoste A *et al*. Differentiable causal discovery from interventional data. In: *Proceedings of the 34th International Conference on Neural Information Processing Systems*. Larochelle H, Ranzato M, Hadsell R (eds.), et al. *Vancouver*, p. 2020. BC, Canada: Curran Associates Inc., 1834.
 28. Lopez R, Hütter J-C, Pritchard JK *et al*. Large-scale differentiable causal discovery of factor graphs. In: *Proceedings of the 36th International Conference on Neural Information Processing Systems*. Koyejo S, Mohamed S, Agarwal A (eds.), et al., p. 1402. New Orleans, LA, USA: Curran Associates Inc., 2022.
 29. Lippe P, Cohen T, Gavves E. Efficient neural causal discovery without acyclicity constraints. In: Ranzato M, Beygelzimer A, Dauphin Y *et al*. (eds.), *International Conference on Learning Representations*. Virtual: OpenReview.net, 2022.
 30. Ke NR, Bilaniuk O, Goyal A *et al*. Learning neural causal models from unknown interventions. In: *Proceedings of the Eighth International Conference on Learning Representations (ICLR 2020)*. Virtual, 2020.
 31. Zhu S, Ng I, Chen Z. Causal discovery with reinforcement learning. In: *Proceedings of the Eighth International Conference on Learning Representations (ICLR 2020)*. Virtual, 2020.
 32. Goudet O, Kalainathan D, Caillou P *et al*. Learning functional causal models with generative neural networks. In: Escalante HJ, Guyon I, Escalante-B S (eds.), et al. *Explainable and Interpretable Models in Computer Vision and Machine Learning*, pp. 39–80. Cham, Switzerland: Springer, 2018.
 33. Li H, Xiao Q, Tian J. Supervised whole dag causal discovery. arXiv preprint arXiv:2006.04697. 2020.
 34. Dai H, Ding R, Jiang Y *et al*. M4c: Seeing causality through latent vicinity. In: *Proceedings of the 2023 SIAM International Conference on Data Mining (SDM)*. Shekhar S, Zhou ZH, Chiang YY (eds.), et al., pp. 226–34. Minneapolis, MN: SIAM, 2023.
 35. Löwe S, Madras D, Zemel R *et al*. Amortized causal discovery: Learning to infer causal graphs from time-series data. In: *Conference on Causal Learning and Reasoning*. Schölkopf B, Uhler C, Zhang K (eds.), pp. 509–25. Virtual: PMLR, 2022.
 36. Wang X, Du Y, Zhu S *et al*. Ordering-based causal discovery with reinforcement learning. In: *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI-21)*, 2021.

37. Lachapelle S, Brouillard P, Deleu T *et al*. Gradient-based neural dag learning. In: *Proceedings of the Eighth International Conference on Learning Representations (ICLR 2020)*, p. 2020. Virtual.
38. Ke NR, Wang J, Mitrovic J *et al*. Amortized learning of neural causal representations. arXiv preprint arXiv:2008.09301. 2020.
39. Kalainathan D, Goudet O, Guyon I *et al*. Structural agnostic modeling: Adversarial learning of causal graphs. *J Mach Learn Res* 2022; **23**:1–62.
40. Scherrer N, Bilaniuk O, Annadani Y *et al*. Learning neural causal models with active interventions. In: *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS 2021)*, 2021.
41. Lopez-Paz D, Muandet K, Schölkopf B *et al*. Towards a learning theory of cause-effect inference. In: *International Conference on Machine Learning*. Bach F, Blei D (eds.), pp. 1452–61. Lille, France: PMLR, 2015.
42. Noè U, Taschler B, Täger J *et al*. Ancestral causal learning in high dimensions with a human genome-wide application. arXiv preprint arXiv:1905.11506. 2019.
43. Hill SM, Oates CJ, Blythe DA *et al*. Causal learning via manifold regularization. *J Mach Learn Res* 2019;**20**:127.
44. Ke NR, Chiappa S, Wang J *et al*. Learning to induce causal structure. In: *Proceedings of the 39th International Conference on Machine Learning*, p. 162. PMLR, 2022.
45. Vaswani A, Shazeer N, Parmar N *et al*. Attention is all you need. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*. Guyon I, Von Luxburg U, Bengio S (eds.), et al., pp. 6000–10. Long Beach, CA, USA: Curran Associates Inc., 2017.
46. Li Y-C, You Z-H, Yu C-Q *et al*. DeepCMI: A graph-based model for accurate prediction of circRNA–miRNA interactions with multiple information. *Brief Funct Genomics* 2023;**23**:276–85.
47. Wright L, Davidson S. How to tell the difference between a model and a digital twin. *Adv Model Simulation Eng Sci* 2020;**7**:13.
48. Yeo HC, Selvarajoo K. Machine learning alternative to systems biology should not solely depend on data. *Brief Bioinform* 2022;**23**:bbac436.
49. Wang X, Liu J, SeS H *et al*. HILAMA: High-dimensional multi-omics mediation analysis with latent confounding. *BMC Med Res Methodol* 2025;**25**:239.
50. Cai Q, Fu Y, Lyu C *et al*. A new framework for exploratory network mediator analysis in omics data. *Genome Res* 2024;**34**: 642–54.
51. Yeo HC, Vijay V, Selvarajoo K. Identifying effective evolutionary strategies-based protocol for uncovering reaction kinetic parameters under the effect of measurement noises. *BMC Biol* 2024;**22**:235.
52. Replogle JM, Saunders RA, Pogson AN *et al*. Mapping information-rich genotype-phenotype landscapes with genome-scale perturb-seq. *Cell* 2022;**185**:2559–2575.e28.
53. Radivojević T, Costello Z, Workman K *et al*. A machine learning automated recommendation tool for synthetic biology. *Nat Commun* 2020;**11**:4879.