

# Mechanistic inference of stochastic gene expression from structured single-cell data

Christopher E. Miles<sup>a</sup>

Single-cell gene expression measurements encode variability spanning molecular noise, cell-to-cell heterogeneity, and technical artifacts. Mechanistic stochastic models provide powerful approaches to disentangle these sources, yet inferring underlying dynamics from standard snapshot sequencing data faces fundamental identifiability limitations. This review focuses on how structured datasets with temporal, spatial, or multimodal features offer constraints to resolve these ambiguities, but they demand more sophisticated models and inference strategies, including machine-learning techniques with inherent tradeoffs. We highlight recent progress in the judicious integration of structured single-cell data, stochastic model development, and innovative inference strategies to extract predictive, gene-level insights. These advances lay the foundation for scaling mechanistic inference upward to regulatory networks and multicellular tissues.

## Address

Department of Mathematics, Center for Complex Biological Systems, University of California, Irvine, CA, USA

Corresponding author: Miles, Christopher E. ([chris.miles@uci.edu](mailto:chris.miles@uci.edu))

<sup>a</sup> Present address: Department of Mathematics, University of Utah, UT, USA

**Current Opinion in Systems Biology** 2025, **42**:100555

This review comes from a themed issue on Mathematical Modelling 2025

Edited by **Alejandro F. Villaverde** and **Matthew Simpson**

For a complete overview see the [Issue](#) and the [Editorial](#)

Available online 18 August 2025

<https://doi.org/10.1016/j.coisb.2025.100555>

2452-3100/© 2025 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

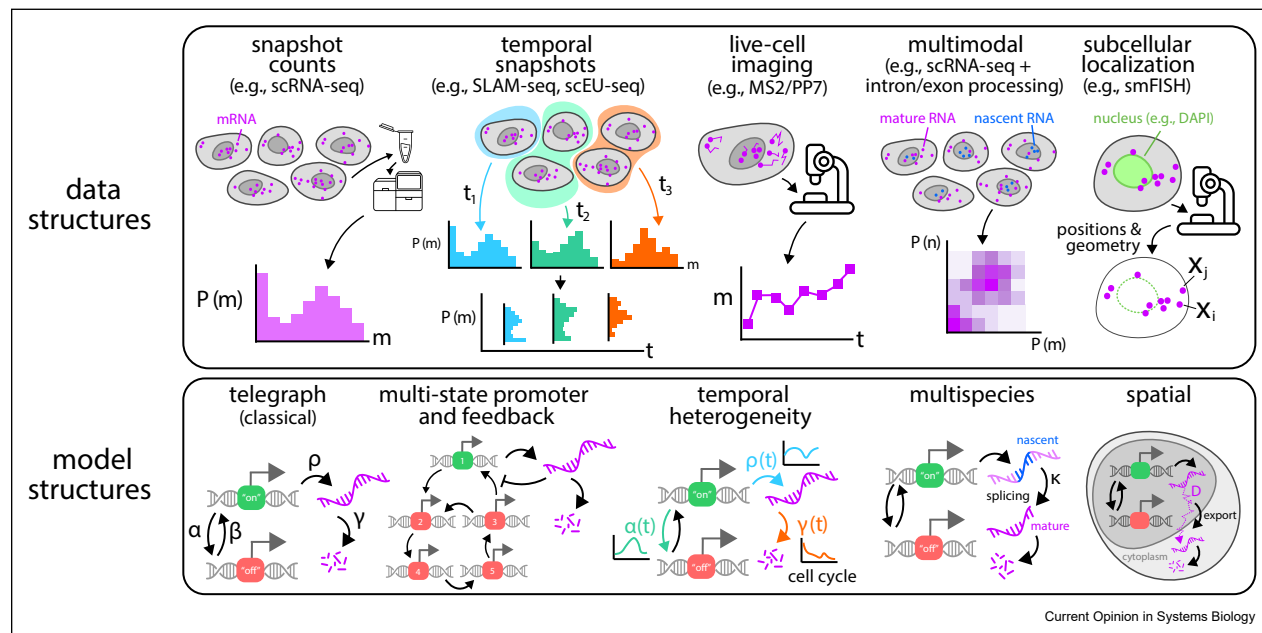
## Introduction

Single-cell gene expression measurements encode a tapestry of variability spanning intrinsic molecular noise, cellular heterogeneity, and confounding technical artifacts. Far from a mere nuisance, the biological portion of these fluctuations shapes cellular identities [1] and disease progression [2]. Disentangling the interwoven noise components is therefore essential for extracting

causal understanding of genetic processes from single-cell data. Beyond statistical pattern-finding, mechanistic stochastic models offer an explanatory framework for this task by linking specific molecular and experimental processes to observed distributions [3]. Application of such models to single-cell data has empowered numerous quantitative insights, from cataloging transcriptional dynamics across the human genome [4] to understanding how these kinetics are regulated [5–8] and vary by cell state [9,10]. More than just descriptive, these models provide mechanistic parameters that move beyond observing *whether* a system has changed to explaining *how* its underlying molecular processes are altered under perturbation, a distinction essential for therapeutic design and cellular engineering.

Despite successes in dissecting transcriptional dynamics, reliably inferring kinetics and distinguishing between hypotheses using typical high-throughput single-cell data remains challenging. Standard ‘snapshot’ measurements of transcript counts from conventional scRNA-seq data often lack sufficient information to uniquely constrain parameters or discriminate between models. This review examines how structured single-cell datasets, sophisticated stochastic models, and advanced inference strategies, together, can overcome these limitations to disentangle the dynamics. We focus on inference for individual genes, which, despite their apparent simplicity, pose surprising inferential obstacles and provide a valuable foundation for understanding regulatory and multicellular scales. We first illustrate fundamental challenges in mechanistic inference from snapshot data, then explore how richer experimental designs incorporating temporal dynamics (e.g. metabolic labeling), spatial organization, and multimodal measurements (e.g. nascent/mature RNA) provide critical constraints for interpretation (see [Figure 1](#)). As datasets and models grow in complexity, we evaluate the evolving toolkit of inference methods, from classical approaches to machine learning techniques, with particular attention to reliability challenges introduced by approximations and surrogate models. Throughout, we emphasize that the integration of appropriate data structures, models, and carefully validated inference is essential for extracting reliable biological insights from single-cell data.

Figure 1




Landscape of single-cell gene expression data and model structures considered throughout the review. Data with additional structure can resolve model ambiguities, but this necessitates developing new models and inference techniques.

## Promises and challenges of mechanistic inference with the telegraph model

A natural entry point into model-based inference with single-cell data is the classical telegraph model. This model minimally captures transcriptional bursting (the stochastic switching between transcriptionally active and silent states [11]) and has thus served as a cornerstone of quantitative single-cell analyses. By examining the model's promises and limitations, we identify broader challenges in mechanistic inference that motivate the need for richer data and more flexible models.

The telegraph model describes a gene promoter stochastically switching between inactive  $G$  ('off') and active  $G^*$  ('on') states with respective rates  $\alpha$  and  $\beta$ . While in the active state, transcription produces mRNA at a rate  $\rho$ , and mRNA is degraded at a rate  $\gamma$ . These processes are summarized by the scheme



Observations of the model are encoded by the *chemical master equations* for  $P_G(m, t)$  and  $P_{G^*}(m, t)$ , the probability distributions of mRNA counts in each promoter state with  $m = 0, 1, \dots$ ,

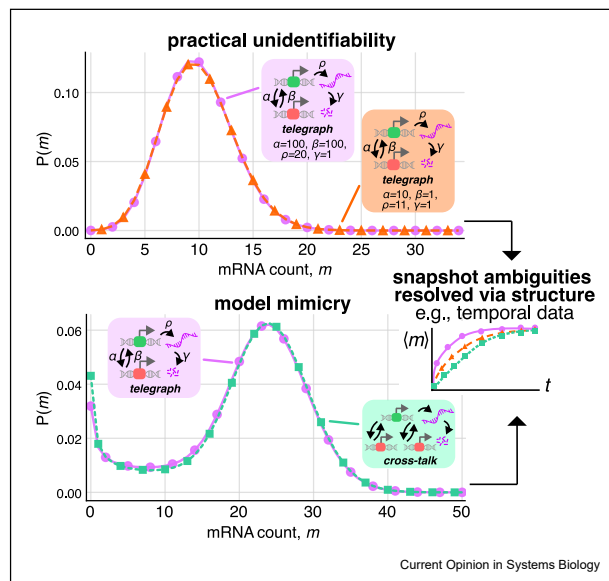
$$\begin{aligned} dP_G/dt &= \beta P_{G^*}(m, t) + \gamma(m+1)P_G(m+1, t) \\ &\quad - (\alpha + \gamma m)P_G(m, t) \end{aligned} \quad (1)$$

$$\begin{aligned} dP_{G^*}/dt &= \alpha P_G(m, t) + \rho P_{G^*}(m-1, t) \\ &\quad + \gamma(m+1)P_{G^*}(m+1, t) \\ &\quad - (\beta + \rho + \gamma m)P_{G^*}(m, t) \end{aligned} \quad (2)$$

The resulting steady-state distribution for  $P(m) = P_G(m) + P_{G^*}(m)$  is the key to both the model's utility and inferential challenges. It classically follows a Poisson-Beta distribution with three effective parameters: the relative promoter switching rates  $a = \alpha/\gamma$  and  $b = \beta/\gamma$ , and the relative synthesis rate  $r = \rho/\gamma$ . This Poisson-Beta distribution can reproduce the unimodal or bimodal observations of mRNA counts, and reduces to the widely observed negative binomial distribution in the limit of short, infrequent pulses of transcription ( $\alpha \ll \gamma, \alpha \ll \beta$ ) [12]. This analytical simplicity provides a valuable basis for extensions that address practical experimental complexities. For example, modeling measurement noise (e.g. mRNA capture efficiency) upon the telegraph framework improves inference quality [13–15].

However, the resulting Poisson-Beta distribution highlights a fundamental inference challenge: the original four kinetic parameters in the model cannot be resolved from steady-state snapshots alone. Such a limitation is an example of *structural unidentifiability*, where model parameters (or combinations thereof) can fundamentally

Figure 2



Inference challenges with snapshot observations of the telegraph model. **Practical identifiability:** starkly different kinetics can produce the same distribution, shown in Ref. [16]. **Model mimicry:** more complex promoter structures, such as a crosstalk pathway model, can produce a nearly identical distribution, shown in Ref. [17].

not be resolved from particular data [18]. Even these three effective parameters are challenging to infer reliably. The three-parameter distribution exhibits *practical unidentifiability*, where different parameter combinations yield nearly identical distributions (Figure 2). This parameter ambiguity is exacerbated by technical noise [16,19] and by extrinsic variability like cell-cycle progression [20,21]. One might assume these inferential obstacles stem from the model's simplicity. However, fundamentally more complex variants (e.g. involving multiple promoter states or feedback) can generate snapshot distributions statistically indistinguishable from the simple telegraph model, a phenomenon called *model mimicry* in Jiao et al. [17], as shown in Figure 2.

### Leveraging structured data to overcome ambiguity

These challenges highlight that snapshot counts alone are often insufficient for reliable mechanistic interpretation. Jiao et al. [17] point toward a resolution: while steady-state distributions may be indistinguishable, transient dynamics often differ distinctly. Resolving these inherent ambiguities thus requires richer experimental data to provide additional constraints. Information embedded in temporal dynamics, spatial organization, or the relationships between multiple molecular species measured simultaneously holds the

potential to distinguish between mechanisms that appear degenerate in static measurements. Leveraging these structured data types requires carefully constructing appropriate mechanistic models and robust methods to connect them with data and pinpoint what information can be extracted. Next, we will highlight recent case studies that exemplify successes in leveraging single-cell data structures for mechanistic insight.

### Temporal dynamics reveal absolute rates and pathway complexity

Observing systems changing over time is the most intuitive way to resolve ambiguities inherent to static, snapshot-based inference, where often only ratios of kinetic rates can be inferred. Temporal data provides the necessary information to overcome this degeneracy and pinpoint absolute timescales. Such data in single-cell gene expression take a variety of forms. The ideal dataset for any temporal inference is a genuine time series, where temporal correlations encode valuable information. While such systems exist (e.g. MS2/PP7 labeling [22]) and have been used with mechanistic modeling [23], these measurements are challenging at scale. More commonly, temporal insights are gleaned from a series of destructive snapshots taken over time, where different cells are sampled, giving the marginal distributions of a population, as shown in Figure 1. Even without connected time points, these temporal snapshots have seen recent success in distinguishing models and parameters, as discussed next.

Metabolic labeling techniques (e.g. using 4sU followed by sequencing protocols like scEU-seq or SLAM-seq [24,25]) are a powerful tool to generate temporal measurements. These methods allow the distinction between newly synthesized ("new" or "labeled") and pre-existing ("old" or "unlabeled") RNA within single cells, often measured across time courses or after pulse-chase experiments. By tracking the dynamics of both labeled and unlabeled populations, one can decouple and directly estimate the absolute kinetic rates of RNA synthesis  $\rho$  and degradation  $\gamma$ , parameters that are inherently intertwined in static measurements. Volterras et al. [26] demonstrate the power of this approach by integrating time-resolved metabolic labeling data to resolve absolute rates. Beyond the inference of rate parameters, the authors ask: how much temporal heterogeneity can be extracted from these measurements? Remarkably, the authors can identify how the kinetic rates, e.g.  $\rho(t)$ ,  $\gamma(t)$  vary temporally, averaging over technical complexities introduced by gene-doubling during replication. This genome-wide quantification of absolute, temporally-heterogeneous burst kinetics and mRNA degradation rates reveals complex regulatory patterns, such as distinct waves of burst frequency versus decay rate modulation across the cell cycle.

Beyond determining absolute rates, temporal data can also clarify the complexity of the underlying pathway, which we previously saw can be challenging with the telegraph model and snapshots (Figure 2, model mimicry). Nicoll et al. [27] provide an analytical calculation that shows that for a large class of transcription models, the initial increase in mean mRNA levels following induction (assuming negligible initial counts and a known delay) follows a general power law,  $\langle m(t) \rangle \propto t^n$ . The exponent  $n$  reflects the number of sequential, irreversible rate-limiting steps (promoter transitions, splicing, and export) for mature transcript production. Measuring this transient scaling behavior, potentially in the minutes following induction, thus provides a method, robust to certain types of extrinsic noise, to estimate a lower bound on pathway complexity, distinguishing between models (e.g. with different numbers of inactive promoter states) that might yield identical steady-state distributions.

Surprisingly, even without any direct temporal measurements, modeling temporal variations is a valuable pursuit. For instance, Sukys and Grima [28] derive a latent cell cycle phase  $\tau$ , computationally inferred from scRNA-seq snapshots, as a biological pseudo-time axis. By fitting mechanistic models that explicitly incorporate this latent cell-cycle progression,  $\rho(\tau)$ , they estimate phase-specific burst parameters for thousands of genes even without temporal measurements. Their findings reveal dosage-compensation mechanisms (e.g. burst frequency halving postreplication while burst size remains constant) and highlight that neglecting such inherent temporal structure can significantly bias kinetic parameter estimates derived from naive snapshot analysis. Popular RNA velocity methods also extract pseudotemporal information from snapshots by bridging the temporal component of this review with multimodal data (see Multimodal readouts decouple linked processes), particularly nascent and mature RNA. Recent variants employ fully mechanistic models to extract interpretable kinetics (e.g. Refs. [9,29,30]). These examples highlight how extracting temporal information is possible even without temporally connected time points. Understanding what information can (or cannot) be extracted from temporal snapshots (marginal distributions) remains an active area of fundamental theoretical research [31,32].

### Spatial context connects location to kinetics and function

Beyond temporal observation, the spatial organization of molecules within cells and tissues is increasingly accessible through high-resolution imaging, ranging from single-molecule FISH (smFISH) for quantifying a few specific gene transcripts to highly multiplexed methods like MERFISH, capable of mapping hundreds simultaneously [33]. These spatial patterns offer

another valuable data modality to dissect mechanisms obfuscated in other measurements. Various computational toolkits and statistical frameworks (e.g. Ref. [34]) have emerged to quantify mRNA enrichment in specific compartments (e.g. nucleus, cytoplasm, and protrusions) or proximity to cellular structures. These methods offer functional insights, such as cell-type-specific pattern variations [35], but often lack mechanistic details of how spatial organization arises. Some mechanistic models (e.g. Ref. [36]) consider well-mixed compartments, overlooking the fine-grained spatial information available from imaging.

Explicitly modeling molecular positions, however, unlocks the mechanistic link between these fine-grained patterns and the underlying kinetics. For instance, transcription occurs at a visible transcriptionally active site (TAS) modeled as a point source at a location  $z$ , within a nuclear domain  $\Omega$  delineated by a fluorescent stain (e.g. DAPI). The subsequent approximate diffusive motion of mRNA through the crowded nucleoplasm (with diffusivity  $D$ ) followed by processing, nuclear export, or degradation, encodes spatial mRNA patterns, although challenging to dissect due to the stochasticity and low molecular counts inherent to these systems. Miles [37] provides an example of quantifying these spatial gradients by coupling telegraph transcription kinetics with a spatial movement model and explicit nuclear export at a geometric boundary. While the resulting distribution of mRNA positions  $x_1, \dots, x_m$  is challenging to describe due to stochasticity in both the quantity and value, it is concisely encoded as a spatial Cox process [38].

$$\begin{aligned} x_1, \dots, x_m &\sim \text{Pois}(u(x, T)), \\ \partial_t u &= D \partial_{xx} u - \gamma u + \rho \Lambda(t) \delta(x - z), \\ \kappa u(x, t) &= -D \partial_x u(x, t), \quad \text{for } x \in \partial\Omega. \end{aligned} \quad (3)$$

where  $\kappa$  encodes nuclear export and  $\Lambda(t)$  is the telegraph process that switches at rates  $\alpha, \beta$ . Crucially, this spatial information can resolve parameter ambiguities inherent in total counts. For example, if the diffusion coefficient  $D$  is known (e.g. from live-cell tracking), the spatial patterns enable inference of production and degradation rates, which are otherwise confounded in nonspatial steady-state measurements, as we saw in Promises and challenges of mechanistic inference with the telegraph model. Cell-to-cell heterogeneity in spatial features (e.g. nuclear shape, gene location) can also sometimes aid parameter identifiability [39]. Altogether, harnessing spatial data offers unique opportunities to constrain transport kinetics and understand localization-dependent regulation.

### Multimodal readouts decouple linked processes

Measuring only a single molecular species, such as mature mRNA, necessarily integrates the effects of all



upstream processes. This aggregation often makes it impossible to distinguish, for example, whether low mRNA levels are due to slow transcription or rapid degradation, see Figure 2. Simultaneously measuring and modeling multiple linked molecular species provides intermediate readouts that can deconvolve connected processes, resolve ambiguities, and offer insight into unobservable aspects of the regulatory cascade [20].

A key example in transcriptomics is the joint measurement of nascent (unspliced) and mature (spliced) RNA, often achievable from standard scRNA-seq data [40]. The joint distribution of these counts  $P(n, m)$  contains information on intermediate rates like splicing, which is lost when only considering the marginal distribution of total mRNA. Fitting such multispecies models at scale, however, presents a significant computational challenge. To address this, Carilli et al. [41] developed biVI, which embeds a biophysical model within a variational autoencoder (VAE) framework. This hybrid VAE approach captures complex-cell state heterogeneity in low dimensions while inferring interpretable, cell-state-dependent kinetic parameters from joint nascent/mature counts.

Further theoretical foundations for multispecies models are rapidly advancing, from refined models of nascent/mature RNA dynamics [42–44] to compartment models that distinguish nuclear and cytoplasmic mRNA [36,45]. The latter have proven crucial for explaining observations of ultra-low, sub-Poissonian noise [46] and are becoming practical for large-scale inference through new

analytical solutions [36]. Together, these pursuits underscore the value of integrating tailored models with multimodal data.

## Inference in practice

The previous section highlights how structured data and richer models can, in principle, provide deeper mechanistic insights. However, increased complexity introduces challenges in connecting structured data with models. For instance, how does one infer cell-cycle-dependent forms like  $\rho(t)$  rather than a simple rate  $\rho$ ? Or how should one handle unwieldy multivariate distributions from multimodal or spatial data, especially with cellular heterogeneities? These newly introduced obstacles amplify the need for carefully chosen inference strategies that push the boundary of classical inference. This section reviews emerging inference strategies for stochastic models, highlighting tradeoffs, practical considerations, and successful applications in translating structured single-cell data into reliable biological insight, as summarized in Table 1.

## The limits of classical inference for complex models

Classical statistical inference hinges on evaluating the likelihood function,  $L(\theta) = P(\text{data}|\text{model}, \theta)$ . For molecular counts, often  $\log L(\theta) = \sum_m N(m) \log P(m|\theta)$ , where  $N(m)$  is the number of cells with RNA counts  $m$ , simple models like the telegraph model allow  $P(m)$  to be computed analytically. For more complex models, this approach becomes challenging but can still be fruitful, particularly through two distinct directions. The first is

**Table 1**

**Summary of inference techniques for stochastic models.**

Methods	Core principle	Tradeoffs (+strengths/- limitations)	Examples
Likelihood-Based	Evaluate $P(\text{Data} \text{Model}, \theta)$ using likelihood derived exactly (e.g. inverting GF) or approximately (e.g. FSP)	+ Strong statistical grounding + Enables standard UQ & ID analysis (FIM, profiles) - Requires tractable likelihood, limits model complexity - Can be computationally expensive	[3,36,37,47,48]
Moment-Based	Match model's statistical moments (or other summaries) to data estimates	+ Computationally scalable, avoids full likelihood calculation + Loses distributional information, prone to bias - Sensitive to moment choice - UQ/ID challenging (bounds sometimes possible [49])	[15,27]
Simulation-Based (SBI)	Compare data to simulations (via stats/dist/NNs). Uses e.g. ABC, NPE/NLE	+ Handles intractable likelihoods & complex models + Neural: Sim-efficient; learns features implicitly + ABC: Simpler setup - Computationally intensive (esp. ABC) - Setup sensitive (ABC: stats; Neural: Architecture/training)	[14,26,50]
Hybrid (ML-enhanced)	Use ML (NNs, GNNs, VAEs) to approx./accelerate/embed model components	- Validation essential (approx. Posteriors, reliability, ID) + Scalable and flexible; handles complex data/dynamics + Can embed mechanism for interpretability, e.g. biVI [41] - Surrogate errors can bias inference; "black box" risks - Challenging UQ and validation; can obscure mechanistic ID	[41,51,52]

UQ, uncertainty quantification; ID, identifiability; GF, generating functions; FSP, finite-state projection; NPE/NLE, neural posterior/likelihood estimation; ML, machine learning; NN, neural network; GNN, graph neural network; VAE, variational autoencoder; ABC, approximate Bayesian computation.

to compute  $P(m)$  using generating functions,  $G(z) = \sum_m z^m P(m)$ , which can be derived with lengthy calculations for surprisingly complex models [3,36] and then inverted numerically to recover  $P(m)$ . Second, finite state projection (FSP) methods, which truncate the state space  $m = 0, 1, \dots, M$  in a systematic way, provide approximate likelihoods for complex models [48,53,54]. If full likelihood computation is infeasible but some analytical progress is possible, moment-based methods match model-derived statistical moments (e.g. time-dependent means [27]) to data estimates, circumventing likelihood calculation. Moment-based approaches inherently lose distributional information, and uncertainty quantification can be challenging, although recent work shows promise in deriving parameter bounds [49]. At the frontier of addressing the computational burden of computing these quantities for complex models, efficient approximation methods like Holimap [55] aim to map complex network dynamics onto simpler, solvable systems.

#### Expanding the toolkit: simulation-based and ML-enhanced inference

When likelihood or moment calculations are infeasible for models with complex data structures, simulation-based inference (SBI) offers a practical alternative that uses forward simulations to infer parameters without an explicit likelihood. The most classical SBI approach, approximate Bayesian computation (ABC), compares simulations with data through the lens of user-chosen summary statistics. Parameters with simulations ‘close’ to the data are accepted to build an approximate posterior distribution. This compression into lower-dimensional statistics enables data-model comparison to be feasible, but the performance of ABC critically depends on the choice of context-specific summary statistics that are sufficiently informative. For example, to infer absolute rates and cell-cycle dependence from complex metabolic labeling data, Volteras et al. [26] use dynamic correlations between labeled/unlabeled RNA and moments (means, variances) of the populations to pin down dynamics and how they vary temporally.

While conceptually straightforward, ABC’s critical dependence on manually chosen summary statistics motivates the use of modern alternatives. Neural SBI approaches (e.g. neural posterior/likelihood estimation, NPE/NLE) employ deep learning to train neural networks on simulation outputs as surrogate approximations to the posterior distribution  $P(\theta|\text{data})$  or the likelihood function. In doing so, neural SBI methods implicitly learn informative features and show promise for superior computational efficiency over ABC [56–58]. However, enjoying these benefits is not entirely riskless. The challenge shifts to decisions in network architecture, training procedures, and the crucial step of validating the surrogate. Neural SBI

methods can yield overconfident posteriors when the underlying model is misspecified [59] and are complicated by the difficulty of validating the surrogate’s accuracy [58].

There is growing momentum in combining deep learning’s representational power with the mechanistic interpretability of classical approaches. Hybrid methods at this intersection have emerged that leverage neural networks to accelerate simulations, create surrogates for model components, or embed mechanistic parameters within flexible representations. For instance, biVI [41] uses a variational autoencoder (VAE) with analytically computed basis functions that describe an otherwise intractable joint distribution of nascent and mature mRNA counts. Sukys and Grima [28] demonstrate another clever hybrid by using deep learning to extract a latent cell-cycle time  $\tau$  for each observation, which is then used in a more traditional moment-approximated inference scheme to reveal how kinetic rates vary across the cell cycle, e.g.  $\rho(\tau)$ . These two examples highlight the merit of hybrid methods, where careful construction of the interface between mechanistic and machine learning components can yield powerful interpretability, computational tractability, and statistical expressiveness.

#### Tradeoffs and reliability

Beyond models of gene expression, deep-learning-based methods for general stochastic reaction models [51,60] and spatial stochastic models [58] show remarkable promise for scalability and flexibility. For instance, graph neural networks trained on simple geometries can extrapolate predictions of stochastic spatial reaction models to new complex geometries [52].

Despite their promise, these deep learning methods must be used with caution and undergo extensive validation to avoid the ‘black box’ issues that can confound their output [61]. This need for scrutiny arises because such methods often rely on neural network surrogates to approximate the true posterior or the likelihood, and the reliability of the final inference hinges on the fidelity of this approximation [58]. The choice of what these surrogates approximate, e.g. the posterior or the likelihood, determines whether the valuable toolkit of classical inference is forfeited or made approximate. Methods that directly learn an approximate posterior, such as neural posterior estimation (NPE), largely bypass the likelihood and its associated toolkit. In contrast, other strategies, including the hybrid approaches discussed previously, use neural networks to creatively compute an *approximate* likelihood. These methods can serve as standalone surrogates for the likelihood function itself (i.e. neural likelihood estimation, NLE) or be embedded within larger models to resolve specific bottlenecks, such as approximating an intractable

distribution [41] or learning unknown terms to render a master equation solvable [52]. In these likelihood-approximating cases, the classical toolkit remains accessible in principle, but its power for evaluating obstacles like practical identifiability and model mimicry is now contingent on the surrogate's accuracy. This toolkit spans profile likelihood analysis to evaluate parameter unidentifiability [62], likelihood-based information criteria for model selection [48], and information-theoretic quantities (e.g. the Fisher information) to direct optimal experimental design [47,63]. By comparison, the equivalent toolkit for modern, deep-learning-based inference is under active development but remains less refined, marking an active frontier for the field.

## Conclusion and outlook

Recent advances in mechanistic inference for stochastic gene expression signal a maturing field, moving beyond simply fitting models to critically evaluating what can be reliably learned from available data. The increasing availability of structured single-cell datasets (whether temporal, spatial, or multimodal) has created new opportunities to resolve regulatory mechanisms, while also sharpening the need for rigorous validation and principled model selection. The next frontier lies in integrating an even richer array of molecular measurements. This expansion, however, introduces formidable new challenges alongside immense promise.

Experimental techniques are pushing the field along several, often intertwined, axes of complexity. One axis pushes toward greater molecular detail for single genes. For instance, long-read sequencing [64] and imaging [65] methods can distinguish mRNA isoforms, or CRISPR-based live-cell RNA tracking [66] can encode post-transcriptional dynamics with unprecedented detail. Another axis expands the breadth of measurement. Genuinely multiomic datasets that pair transcriptomics with epigenomics (e.g. chromatin accessibility) or proteomics (e.g. surface protein levels) are now widely available. These modalities hold tremendous potential for creating more complete regulatory models but require synthesizing data types that differ fundamentally in their biological timescales and technical biases [67]. A third axis scales these models upward to larger systems like gene regulatory networks (GRNs) and multicellular tissues [68,69]. Applying these insights at such scales magnifies inferential challenges due to a combinatorial explosion of potential interactions, making it difficult to infer causal regulatory links from transcriptomic data alone [70].

Tackling these frontiers underscores this review's central thesis: the principled approaches developed for single-gene inference are not merely an inspiration but constitute the essential groundwork for this next era.

They provide the mechanistic building blocks for larger models and, critically, a conceptual framework for disentangling the contributions of intrinsic, extrinsic, and technical noise, a task that becomes paramount in more complex systems. Ultimately, as experimental techniques advance, the primary bottleneck shifts from data acquisition to robust, interpretable analysis. Extracting reliable insights requires more than rich data; it demands an inference ecosystem that integrates experimental design with rigorous diagnostics. While tools for identifiability analysis, sensitivity diagnostics, and principled model comparison are well-developed alongside classical inference, a critical gap remains in adapting them for the simulation-based and deep-learning methods essential for these complex systems [58,71]. Building generalizable, modular tools that integrate these diagnostics directly into the inference workflow will be the crucial step in ensuring our conclusions are truly meaningful.

## Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the author used Gemini 2.5 Pro (Google) to improve language and readability. After using this tool/service, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication.

## Declaration of competing interest

The author has nothing to declare.

## Acknowledgements

The work was partially supported by NSF CAREER DMS-2339241 and NSF/NIGMS DMS-2451263. The author thanks Fangyuan Ding, Alex Mogilner, and Elizabeth Read for helpful comments on an earlier draft, and the anonymous reviewers for their careful reading and thoughtful suggestions.

## Data availability

No data was used for the research described in the article.

## References

Papers of particular interest, published within the period of review, have been highlighted as:

- of special interest
  - of outstanding interest
1. Coomer MA, Ham L, Stumpf MP: **Noise distorts the epigenetic landscape and shapes cell-fate decisions.** *Cell Syst* 2022, **13**: 83–102.e6, <https://doi.org/10.1016/j.cels.2021.09.002>.
  2. Alachkar N, Norton D, Wolkensdorfer Z, Muldoon M, Paszek P: **Variability of the innate immune response is globally constrained by transcriptional bursting.** *Front Mol Biosci* 2023, **10**, 1176107, <https://doi.org/10.3389/fmolb.2023.1176107>.

3. Gorin G, Vastola JJ, Pachter L: **Studying stochastic systems biology of the cell with single-cell genomics data.** *Cell Syst* 2023, **14**:822–843.e22, <https://doi.org/10.1016/j.cels.2023.08.004>.
4. Dar RD, Razooky BS, Singh A, Trimeloni TV, McCollum JM, Cox CD, Simpson ML, Weinberger LS: **Transcriptional burst frequency and burst size are equally modulated across the human genome.** *Proc Natl Acad Sci* 2012, **109**:17454–17459, <https://doi.org/10.1073/pnas.1213530109>.
5. Larsson AJM, Johnsson P, Hagemann-Jensen M, Hartmanis L, Faridani OR, Reinius B, Segerstolpe Å, Rivera CM, Ren B, Sandberg R: **Genomic encoding of transcriptional burst kinetics.** *Nature* 2019, **565**:251–254, <https://doi.org/10.1038/s41586-018-0836-1>.
6. Luo S, Wang Z, Zhang Z, Zhou T, Zhang J: **Genome-wide inference reveals that feedback regulations constrain promoter-dependent transcriptional burst kinetics.** *Nucleic Acids Res* 2023, **51**:68–83, <https://doi.org/10.1093/nar/gkac1204>.
7. Mahat DB, Tipples ND, Martin-Rufino JD, Waterton SK, Fu J, Blatt SE, Sharp PA: **Single-cell nascent RNA sequencing unveils coordinated global transcription.** *Nature* 2024, **631**:216–223, <https://doi.org/10.1038/s41586-024-07517-7>.
8. Chen P-T, Levo M, Zoller B, Gregor T: **A conserved coupling of transcriptional ON and OFF periods underlies bursting dynamics.** *Nat Struct Mol Biol* 2025:1–13, <https://doi.org/10.1038/s41594-025-01615-4>.
9. Gao CF, Vaikuntanathan S, Riesenfeld SJ: **Dissection and integration of bursty transcriptional dynamics for complex systems.** *Proc Natl Acad Sci* 2024, **121**, e2306901121, <https://doi.org/10.1073/pnas.2306901121>.
10. Chari T, Gorin G, Pachter L: **Biophysically interpretable inference of cell types from multimodal sequencing data.** *Nature Computational Science* 2024, **4**:677–689, <https://doi.org/10.1038/s43588-024-00689-2>.
11. Tunnacliffe E, Chubb JR: **What is a transcriptional burst?** *Trends Genet* 2020, **36**:288–297, <https://doi.org/10.1016/j.tig.2020.01.003>.
12. Kim JK, Marioni JC: **Inferring the kinetics of stochastic gene expression from single-cell RNA-sequencing data.** *Genome Biol* 2013, **14**:R7, <https://doi.org/10.1186/gb-2013-14-1-r7>.
13. Gorin G, Pachter L: **Distinguishing biophysical stochasticity from technical noise in single-cell RNA sequencing using monod.** 2022, <https://doi.org/10.1101/2022.06.11.495771>.
14. Tang W, Jørgensen ACS, Marguerat S, Thomas P, Shahrezaei V: **Modelling capture efficiency of single-cell RNA-sequencing data improves inference of transcriptome-wide burst kinetics.** *Bioinformatics* 2023, **39**:btad395, <https://doi.org/10.1093/bioinformatics/btad395>.
15. Grima R, Esmenjaud P-M: **Quantifying and correcting bias in transcriptional parameter inference from single-cell data.** *Biophys J* 2024, **123**:4–30, <https://doi.org/10.1016/j.bpj.2023.10.021>.
16. Gu J, Laszik N, Miles CE, Allard J, Downing TL, Read EL: **Scalable inference and identifiability of kinetic parameters for transcriptional bursting from single cell data.** *bioRxiv preprint* 2025, <https://doi.org/10.1101/2025.05.13.653671>.
17. Jiao F, Li J, Liu T, Zhu Y, Che W, Bleris L, Jia C: **What can we learn when fitting a simple telegraph model to a complex gene expression model?** *PLoS Comput Biol* 2024, **20**, e1012118, <https://doi.org/10.1371/journal.pcbi.1012118>.  
Demonstrates that fitting the simple telegraph model to steady-state data from complex gene expression processes can yield misleading parameters, highlighting model mimicry.
18. Cinquemani E: **Identifiability and reconstruction of biochemical reaction networks from population snapshot data.** *Processes* 2018, **6**:136, <https://doi.org/10.3390/pr6090136>.
19. Wang Y, Shu Z, Cao Z, Grima R: **From noise to models to numbers: evaluating negative binomial models and parameter estimations in single-cell RNA-seq.** 2025, <https://doi.org/10.1101/2025.05.05.652189>.
20. Ham L, Jackson M, Stumpf MP: **Pathway dynamics can delineate the sources of transcriptional noise in gene expression.** *eLife* 2021, **10**, e69324, <https://doi.org/10.7554/eLife.69324>.
21. Fu X, Patel HP, Coppola S, Xu L, Cao Z, Lenstra TL, Grima R: **Quantifying how post-transcriptional noise and gene copy number variation bias transcriptional parameter inference from mRNA distributions.** *eLife* 2022, **11**, e82493, <https://doi.org/10.7554/eLife.82493>.
22. Tutucci E, Vera M, Biswas J, Garcia J, Parker R, Singer RH: **An improved MS2 system for accurate reporting of the mRNA life cycle.** *Nat Methods* 2018, **15**:81–89, <https://doi.org/10.1038/nmeth.4502>.
23. Burton J, Manning CS, Rattray M, Papalopulu N, Kursawe J: **Inferring kinetic parameters of oscillatory gene regulation from single cell time-series data.** *J R Soc Interface* 2021, **18**, 20210393, <https://doi.org/10.1098/rsif.2021.0393>.
24. Rodrigues SG, Chen LM, Liu S, Zhong ED, Scherrer JR, Boyden ES, Chen F: **RNA timestamps identify the age of single molecules in RNA sequencing.** *Nat Biotechnol* 2021, **39**:320–325, <https://doi.org/10.1038/s41587-020-0704-z>.
25. Ramsköld D, Hendriks G-J, Larsson AJM, Mayr JV, Ziegenhain C, Hagemann-Jensen M, Hartmanis L, Sandberg R: **Single-cell new RNA sequencing reveals principles of transcription at the resolution of individual bursts.** *Nat Cell Biol* 2024, **26**:1725–1733, <https://doi.org/10.1038/s41556-024-01486-9>.
26. Volterras D, Shahrezaei V, Thomas P: **Global transcription regulation revealed from dynamical correlations in time-resolved single-cell RNA sequencing.** *Cell Syst* 2024, **15**:694–708.e12, <https://doi.org/10.1016/j.cels.2024.07.002>.  
Quantifies genome-wide, cell-cycle-dependent burst kinetics and degradation rates using time-resolved metabolic labeling, stochastic modeling, and scalable ABC inference.
27. Nicoll AG, Szavits-Nossan J, Evans MR, Grima R: **Transient power-law behaviour following induction distinguishes between competing models of stochastic gene expression.** *Nat Commun* 2025, **16**:2833, <https://doi.org/10.1038/s41467-025-58127-4>.  
Shows theoretically that transient power-law scaling of mean mRNA following induction distinguishes mechanistic models indistinguishable at steady state.
28. Sukys A, Grima R: **Cell-cycle dependence of bursty gene expression: insights from fitting mechanistic models to single-cell RNA-seq data.** *Nucleic Acids Res* 2025, **53**:gkaf295, <https://doi.org/10.1093/nar/gkaf295>.  
Estimates cell-cycle phase-dependent burst parameters by fitting mechanistic models that account for extrinsic noise, using deep learning to infer cell age from scRNA-seq.
29. Gayoso A, Weiler P, Lotfollahi M, Klein D, Hong J, Streets A, Theis FJ, Yosef N: **Deep generative modeling of transcriptional dynamics for RNA velocity analysis in single cells.** *Nat Methods* 2024, **21**:50–59, <https://doi.org/10.1038/s41592-023-01994-w>.
30. Fang M, Gorin G, Pachter L: **Trajectory inference from single-cell genomics data with a process time model.** *PLoS Comput Biol* 2025, **21**, e1012752, <https://doi.org/10.1371/journal.pcbi.1012752>.
31. Maddu S, Chardès V, Shelley MJ: **Inferring biological processes with intrinsic noise from cross-sectional data.** preprint. 2024, <https://doi.org/10.48550/arXiv.2410.07501>. arXiv:2410.07501.
32. Guan V, Janssen J, Rahmani H, Warren A, Zhang S, Robeva E, Schiebinger G: **Identifying drift, diffusion, and causal structure from temporal snapshots.** 2025, <https://doi.org/10.48550/arXiv.2410.22729>. arXiv:2410.22729.
33. Chen KH, Boettiger AN, Moffitt JR, Wang S, Zhuang X: **Spatially resolved, highly multiplexed RNA profiling in single cells.** *Science* 2015, **348**:6090, <https://doi.org/10.1126/science.aaa6090>.
34. Wang JX, Zhou X: **ELLA: modeling subcellular spatial variation of gene expression within cells in high-resolution spatial transcriptomics.** *bioRxiv preprint* 2024, <https://doi.org/10.1101/2024.09.23.614515>.



35. Takei Y, Yang Y, White J, Goronzy IN, Yun J, Prasad M, Ombelets LJ, Schindler S, Bhat P, Guttman M, Cai L: **Spatial multi-omics reveals cell-type-specific nuclear compartments.** *Nature* 2025;1–11, <https://doi.org/10.1038/s41586-025-08838-x>.
36. Wang Y, Szavits-Nossan J, Cao Z, Grima R: **Joint distribution of nuclear and cytoplasmic mrna levels in stochastic models of gene expression: analytical results and parameter inference.** *Phys Rev Lett* 2025, **135**, 068401, <https://doi.org/10.1103/q5sd-tpms>.
37. Miles CE: **Incorporating spatial diffusion into models of bursty stochastic transcription.** *J R Soc Interface* 2025, **22**, 20240739, <https://doi.org/10.1098/rsif.2024.0739>.
38. Schnoerr D, Grima R, Sanguinetti G: **Cox process representation and inference for stochastic reaction–diffusion processes.** *Nat Commun* 2016, **7**, 11729, <https://doi.org/10.1038/ncomms11729>.
39. Miles CE, McKinley SA, Ding F, Lehoucq RB: **Inferring stochastic rates from heterogeneous snapshots of particle positions.** *Bull Math Biol* 2024, **86**:74, <https://doi.org/10.1007/s11538-024-01301-4>.
40. Sullivan DK, Hjørleifsson KE, Swarna NP, Oakes C, Holley G, Melsted P, Pachter L: **Accurate quantification of nascent and mature RNAs from single-cell and single-nucleus RNA-seq.** *Nucleic Acids Res* 2025, **53**:gkae1137, <https://doi.org/10.1093/nar/gkae1137>.
41. Carilli M, Gorin G, Choi Y, Chari T, Pachter L: **Biophysical modeling with variational autoencoders for bimodal, single-cell RNA sequencing data.** *Nat Methods* 2024, **21**:1466–1469, <https://doi.org/10.1038/s41592-024-02365-9>.  
Introduces biVI, a hybrid VAE framework embedding biophysical models to infer transcriptional kinetics from multimodal nascent/mature scRNA-seq data
42. Gorin G, Vastola JJ, Fang M, Pachter L: **Interpretable and tractable models of transcriptional noise for the rational design of single-molecule quantification experiments.** *Nat Commun* 2022, **13**:7620, <https://doi.org/10.1038/s41467-022-34857-7>.
43. Gorin G, Pachter L: **Modeling bursty transcription and splicing with the chemical master equation.** *Biophys J* 2022, **121**:1056–1069, <https://doi.org/10.1016/j.bpj.2022.02.004>.
44. Szavits-Nossan J, Grima R: **Steady-state distributions of nascent RNA for general initiation mechanisms.** *Phys Rev Res* 2023, **5**, 13064, <https://doi.org/10.1103/PhysRevResearch.5.013064>.
45. Ma M, Szavits-Nossan J, Singh A, Grima R: **Analysis of a detailed multi-stage model of stochastic gene expression using queueing theory and model reduction.** *Math Biosci* 2024, **373**, 109204, <https://doi.org/10.1016/j.mbs.2024.109204>.
46. Weidemann DE, Holehouse J, Singh A, Grima R, Hauf S: **The minimal intrinsic stochasticity of constitutively expressed eukaryotic genes is sub-poissonian.** *Sci Adv* 2023, **9**, eadh5138, <https://doi.org/10.1126/sciadv.adh5138>.
47. Vo HD, Forero-Quintero LS, Aguilera LU, Munsky B: **Analysis and design of single-cell experiments to harvest fluctuation information while rejecting measurement noise.** *Front Cell Dev Biol* 2023, **11**, <https://doi.org/10.3389/fcell.2023.1133994>.
48. Kilic Z, Schweiger M, Moyer C, Shepherd D, Pressé S: **Gene expression model inference from snapshot RNA data using Bayesian non-parametrics.** *Nature Computational Science* 2023, **3**:174–183, <https://doi.org/10.1038/s43588-022-00392-0>.
49. Li Z, Barahona M, Thomas P: **Moment-based parameter inference with error guarantees for stochastic reaction networks.** *J Chem Phys* 2025, **162**, 135105, <https://doi.org/10.1063/5.0251744>.
50. Sierra MAR, Sokolowski TR: **AI-powered simulation-based inference of a genuinely spatial-stochastic gene regulation model of early mouse embryogenesis.** *PLoS Comput Biol* 2024, **20**, e1012473, <https://doi.org/10.1371/journal.pcbi.1012473>.
51. Jiang Q, Fu X, Yan S, Li R, Du W, Cao Z, Qian F, Grima R: **Neural network aided approximation and parameter inference of Non-Markovian models of gene expression.** *Nat Commun* 2021, **12**:2618, <https://doi.org/10.1038/s41467-021-22919-1>.
52. Cao Z, Chen R, Xu L, Zhou X, Fu X, Zhong W, Grima R: **Efficient and scalable prediction of stochastic reaction–diffusion processes using graph neural networks.** *Math Biosci* 2024, **375**, 109248, <https://doi.org/10.1016/j.mbs.2024.109248>.
53. Vo HD, Fox Z, Baetica A, Munsky B: **Bayesian estimation for stochastic gene expression using multifidelity models.** *J Phys Chem B* 2019, **123**:2217–2234, <https://doi.org/10.1021/acs.jpcc.8b10946>.
54. Fox ZR, Munsky B: **The finite state projection based fisher information matrix approach to estimate information and optimize single-cell experiments.** *PLoS Comput Biol* 2019, **15**, e1006365, <https://doi.org/10.1371/journal.pcbi.1006365>.
55. Jia C, Grima R: **Holimap: an accurate and efficient method for solving stochastic gene network dynamics.** *Nat Commun* 2024, **15**:6557, <https://doi.org/10.1038/s41467-024-50716-z>.
56. Lueckmann J-M, Boelts J, Greenberg D, Goncalves P, Macke J: **Benchmarking simulation-based inference.** In *Proceedings of the 24th international conference on artificial intelligence and statistics*. PMLR; 2021:343–351. In <https://proceedings.mlr.press/v130/lueckmann21a.html>.
57. Frazier DT, Kelly R, Drovandi C, Warne DJ: **The statistical accuracy of neural posterior and likelihood estimation.** 2024, <https://doi.org/10.48550/arXiv.2411.12068>. arXiv:2411.12068.
58. Ramirez Sierra MA, Sokolowski TR: **Comparing AI versus optimization workflows for simulation-based inference of spatial-stochastic systems.** *Mach Learn: Sci Technol* 2025, **6**, 10502, <https://doi.org/10.1088/2632-2153/ada0a3>.  
Compares neural SBI (SNPE) with simulated annealing for parameter inference of spatial-stochastic models, assessing posterior quality and computational tradeoffs.
59. Kelly RP, Nott DJ, Frazier DT, Warne DJ, Drovandi C: **Misspecification-robust sequential neural likelihood for simulation-based inference.** *Trans Mach Learn Res* 2024. <https://openreview.net/forum?id=tbOYJwXhcY>; 2024.
60. Huang Z, Luo S, Wang Z, Zhang Z, Jiang B, Nie Q, Zhang J: **Deep learning linking mechanistic models to single-cell transcriptomics data reveals transcriptional bursting in response to DNA damage.** *bioRxiv preprint* 2024, <https://doi.org/10.1101/2024.07.10.602845>.
61. Ahlmann-Eltze C, Huber W, Anders S: **Deep learning-based predictions of gene perturbation effects do not yet outperform simple linear baselines.** *bioRxiv preprint* 2025, <https://doi.org/10.1101/2024.09.16.613342>.
62. Simpson MJ, Maclaren OJ: **Profile-wise analysis: a profile likelihood-based workflow for identifiability analysis, estimation, and prediction with mechanistic mathematical models.** *PLoS Comput Biol* 2023, **19**, e1011515, <https://doi.org/10.1371/journal.pcbi.1011515>.
63. Cook J, Ron E, Svetlov D, Aguilera LU, Munsky B: **Sequential design of single-cell experiments to identify discrete stochastic models for gene expression.** In *2024 IEEE 63rd conference on decision and control (CDC)*; 2024:7778–7785, <https://doi.org/10.1109/CDC56724.2024.10886397>.
64. Belchikov N, Hsu J, Li XJ, Jarroux J, Hu W, Joglekar A, Tilgner HU: **Understanding isoform expression by pairing long-read sequencing with single-cell and spatial transcriptomics.** *Genome Res* 2024, **34**:1735–1746, <https://doi.org/10.1101/gr.279640.124>.
65. Ding F, Elowitz MB: **Constitutive splicing and economies of scale in gene expression.** *Nat Struct Mol Biol* 2019, **26**: 424–432, <https://doi.org/10.1038/s41594-019-0226-x>.
66. Xia C, Colognori D, Jiang XS, Xu K, Doudna JA: **Single-molecule live-cell RNA imaging with CRISPR–csm.** *Nat Biotechnol* 2025:1–8, <https://doi.org/10.1038/s41587-024-02540-5>.

67. Baião AR, Cai Z, Poulos RC, Robinson PJ, Reddel RR, Zhong Q, Vinga S, Gonçalves E: **A technical review of multi-omics data integration methods: from classical statistical to deep generative approaches.** *Briefings Bioinf* 2025, **26**:bbaf355, <https://doi.org/10.1093/bib/bbaf355>.
68. Joly-Smith E, Wang ZJ, Hilfinger A: **Inferring gene regulation dynamics from static snapshots of gene expression variability.** *Phys Rev* 2021, **104**, 44406, <https://doi.org/10.1103/PhysRevE.104.044406>.
69. Herbach U: **Harissa: stochastic simulation and inference of gene regulatory networks based on transcriptional bursting.** In *Computational methods in systems biology*. Edited by Pang J, Niehren J, Cham: Springer Nature Switzerland; 2023:97–105, [https://doi.org/10.1007/978-3-031-42697-1\\_7](https://doi.org/10.1007/978-3-031-42697-1_7).
70. Kernfeld E, Keener R, Cahan P, Battle A: **Transcriptome data are insufficient to control false discoveries in regulatory network inference.** *Cell Syst* 2024, **15**:709–724.e13, <https://doi.org/10.1016/j.cels.2024.07.006>.
71. Noordijk B, Garcia Gomez ML, Ten Tusscher KHWJ, De Ridder D, Van Dijk ADJ, Smith RW: **The rise of scientific machine learning: a perspective on combining mechanistic modelling with machine learning for systems biology.** *Frontiers Syst Biol* 2024, **4**, 1407994, <https://doi.org/10.3389/fsysb.2024.1407994>.