LazyNet: Interpretable ODE Modeling of Sparse CRISPR Single-Cell Screens Reveals New Biological Insights

Abstract

Single-cell CRISPR activation/interference screens offer a direct route to causal gene-regulatory maps, yet existing deep-learning pipelines are GPU-intensive and yield hard-to-interpret latent factors. We introduce LazyNet, an explicitly Euler-integrated neural ODE whose paired log-linear-exp layer collapses multiplicative transcript interactions into a compact, mechanistically interpretable weight matrix. Training a three-replica ensemble on a 55k-cell, 30k-gene Perturb-seq dataset completes on a single CPU in <1 h, running 3 to 4 folder faster than transformer (scGPT) or state-space (RetNet) baselines while lowering global RMSE by 25 % and raising genome-wide Pearson r to 0.67. Averaged Jacobians, expanded in a 32*4 breadth-first search around seven ferroptosis seeds, recapitulated 15 of 27 benchmark regulators (56 % recall) within a 4 611 gene, 11 676 edge subgraph; 26.6 % of edges show ARCHS4 co-expression r ≥ 0.2 versus 5 % expected at random, and 523 overlap STRING interactions (hypergeometric p = 1.2e-5). Elasticity ranks uncover a previously unrecognized lysosomal-mitochondrial-immune module linking PSAP-mTOR, MFN2-TLR4 and ADCY10-SIRT3, generating experimentally testable hypotheses. By combining state-of-the-art predictive accuracy, laptop-level resource demands and one-to-one parameter interpretability, LazyNet democratizes causal network discovery from sparse two-snapshot screens, enabling small laboratories to move from large-scale perturbation data to mechanistic insight without GPUs or

external pathway priors.

Introduction

Single-cell RNA sequencing (scRNA-seq) has transformed transcriptomics by revealing the rich cell-to-cell heterogeneity masked in bulk averages, uncovering rare sub-populations, transitional states, and lineage hierarchies across tissues and species [1-2]. Yet each cell is still frozen in time—captured under one environmental condition and at a single instant — so most temporal trajectories and stimulus-response behaviours remain invisible. Systematically probing these dimensions demands direct perturbations: transcriptional read-outs following targeted genetic edits illuminate how regulatory circuits maintain identity, pinpoint druggable control nodes, and expose synergistic gene pairs that can amplify combination therapies or guide precise cell-engineering strategies [18-31]. CRISPR activation or inhibition screens joined with scRNA-seq ("Perturb-seq") meet this need by bar-coding guide RNAs so that every sequenced cell records its own perturbation history, enabling thousands of knock-outs, knock-downs, or activations to be multiplexed in one experiment [3-4, 32]. The resulting high-dimensional maps of genotype-to-phenotype relationships have already sharpened target validation pipelines, suggested rescue strategies for disease mutations, and provided training data for machine-learning models that aspire to predict unseen edits. These advances have converged in ambitious blueprints for an in-silico "AI virtual cell" (AIVC) capable of simulating cellular behaviour across modalities, contexts, and perturbations [5].

Several deep-learning frameworks aim to link genotype perturbations to single-cell

transcriptional phenotypes. scGPT adapts large-language-model transformers to gene tokens, pre-training hundreds of millions of parameters on atlas-scale collections (>10 million cells) with masked-token objectives before fine-tuning on task-specific data [6]. This workflow assumes multi-GPU nodes (e.g., eight A100s) and terabytes of fast storage for checkpoint sharding. CellFM embeds a variational auto-encoder within a matrix-factorization scaffold to capture dose-dependent responses, but relies on GPU-accelerated Bayesian optimization, extensive hyper-parameter sweeps, and manually curated pathway priors [7]. GEARS employs graph neural networks that propagate signals through a pre-built regulatory graph, necessitating access to high-quality interaction databases and frequent curation when exploring less-characterized systems [8]. In practice, these pipelines demand days to weeks of training time, >40 GB GPU memory, and auxiliary annotations (cell types, pathway libraries) that many small laboratories lack, and their latent spaces seldom map cleanly onto biochemical rate laws—hindering mechanistic interpretation and complicating predictions of higher-order, non-linear genetic interactions.

Therefore, analytical frameworks aimed at everyday laboratories must extract causal information from tens of thousands—not millions—of sparsely sampled cells, operate directly on simple count matrices, and train within the CPU budgets of standard university clusters. Such models must bridge pattern recognition and mechanisms: they should explain why perturbation behaves as observed, remain interpretable to domain experts, and scale to new datasets without extensive re-engineering.

Classical ordinary and partial differential equations (ODEs/PDEs) meet many of these

demands [9-11]. By embedding chemical-kinetic constraints they are naturally robust to noise, sparse observations, and domain shifts, and their parameters have clear biological meaning [12-13]. Yet deriving, calibrating, and simulating large ODE/PDE systems by hand is infeasible for transcriptomes comprising tens of thousands of genes, and naïvely nesting such equations inside neural networks raises numerical-stability and memory bottlenecks [14-16].

LazyNet tackles this gap by parameterizing the vector field of a residual ODE solver with paired logarithmic and exponential layers. Working in log-space linearizes multiplicative gene–gene effects, the exponential map restores biological scale, and the architecture dispenses with self-attention quadratic overhead [17]. LazyNet converges within hours on a 24-core CPU node (768 GB RAM), modelling a ~55 k-cell, ~30 k-gene CRISPR A/I dataset that is >95 % sparse without exceeding memory limits. Each weight corresponds directly to an ODE coefficient, ensuring intrinsic interpretability and providing a foundation for broader differential-equation-based modelling.

Material and Methods

LazyNet Architecture

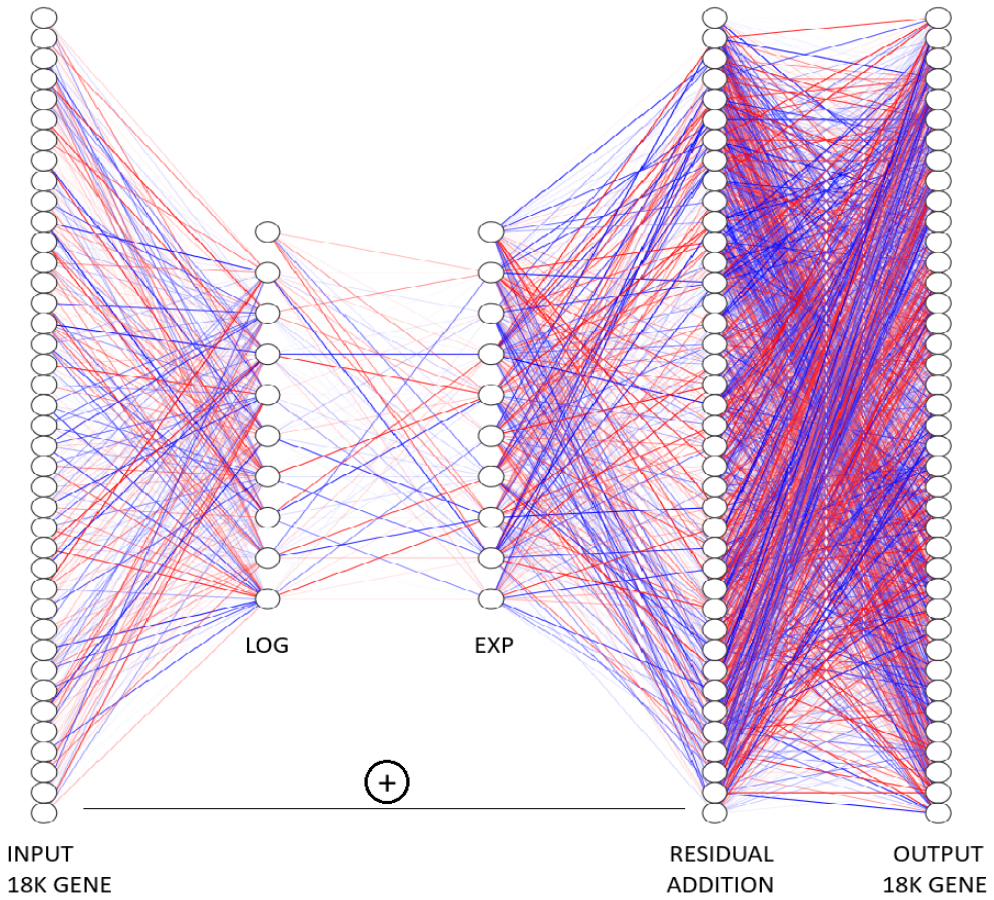

Figure 1. Single-step Euler block in LazyNet.

The full gene-expression vector (INPUT) is propagated through a two-stage transformation composed of a logarithmic layer (LOG) followed by an exponential layer (EXP). These twin linear layers realise multiplicative biology-inspired interactions in log space while retaining computational simplicity. The transformed signal is added back to the original input via a residual skip connection (black bar and "+"), producing an updated vector of identical dimensionality (OUTPUT). Stacking such blocks corresponds to explicit Euler integration of an underlying ordinary differential equation, allowing LazyNet to model gene-regulatory dynamics within a ResNet framework while continuously adjusting weights during training [17].

An ODE of the form $df/dt = F(f)$ can be discretised with forward Euler:

$$f_{\{n+1\}} = f_n + F(f_n)\,\Delta t \quad (Eq.\,1)$$

LazyNet represents the rate law F with a single log–linear–exp (LLE) block placed inside this Euler step:

$$F_LazyNet(x) = exp(\,W_2 \cdot log(\,W_1\,x + b_1\,) + b_2\,) \quad (Eq.\,2)$$
$$f_{\{n+1\}} = f_n + F_LazyNet(f_n)\,\Delta t \quad (Eq.\,3)$$

Here $x \in \mathbb{R}^+$; W_1, W_2 are learnable weight matrices and b_1, b_2 are biases. Taking the logarithm converts multiplicative gene–gene interactions into additions, allowing a single neuron to encode an entire monomial exactly, while the outer exponential restores the original scale.

During training, each observed state f_n is provided as input, the network predicts $f_{\{n+1\}}$, and the prediction is compared with ground truth. With sufficiently varied samples V and adequate temporal coverage T, the learned mapping approaches the true governing equations:

$$lim_{\{V \to max, T \to \infty\}}\, F_{LazyNet} = F_{actual} \quad (Eq.\,4)$$

Because the Euler map is Markovian, every update depends only on the current state, not on the full history; LazyNet therefore learns the system dynamics efficiently and with minimal parameter count [17].

Theoretical Rationale for Efficiency

Most biochemical rate laws can be written as a finite sum of monomials of the form $c \cdot x^{1\{\alpha^1\}} \ldots x^{\{\alpha_d\}}$. A representation that reproduces each monomial exactly captures the full reaction space.

LazyNet achieves this with a single log–linear–exp (LLE) block

$$\mathbb{F}(x) = exp(W \cdot log\, x + b) \quad (Eq. 5)$$

where $x \in \mathbb{R}^{+}, W \in \mathbb{R}^{\{K \times d\}}, and\ b \in \mathbb{R}^{\{K\}}$. Choosing one row of W equal to the integer exponent vector α and the corresponding bias $b = logc$ reproduces the monomial exactly. Hence an entire K-term rate law needs only $K(d+1)$ trainable scalars.

By contrast, a degree-N Taylor polynomial on the same positive domain requires $\sim \left(log\frac{1}{\varepsilon}\right)^{\{d\}}$ parameters to reach uniform error ε. Thus the parameter ratio grows super-polynomials with both dimension d and desired accuracy. Fewer parameters translate to fewer multiply–add operations per mini-batch and a smaller memory footprint, directly accelerating training.

Statistically, the reduced parameter space concentrates information: each observation contributes to a higher signal-to-noise ratio because the same amount of data is no longer diluted across thousands of polynomial coefficients. Empirically this manifests as faster convergence with larger learning rates and improved generalization in small-sample regimes.

Transformer-based baselines such as scGPT, and state-space models like RetNet, avoid the explicit coefficient explosion of a Taylor expansion, yet they pay a different price [6-7]. scGPT attention layers scale as $O(L^2)$ in sequence length L, while RetNet reduces this to $O(LlogL)$ but introduces additional kernel parameters that still number in the millions. Neither can architecture encode a monomial exactly in closed form; they must learn it by gradient descent, requiring more data and iterations. LazyNet therefore sits at a unique efficiency sweet spot: minimal parameter count (linear in d),

constant-time forward pass per cell, and exact capture of multiplicative gene interactions.

Data Preparation and Gene Selection

We evaluated all models on GSE152988, a CRISPR-activation/-interference (A/I) single-cell RNA-seq screen of > 55k human iPSC-derived neurons under oxidative stress [33]. The dataset offers only two expression snapshots—an engineered "time-0" profile (created by injecting each perturbation signature into control cells) and the final post-perturbation state. With no intermediate measurements, the dynamics reduce to a single Euler step, making this a minimal ODE case and a suitable blind-prediction test for LazyNet. Multi-time-point tools such as ScVelo require FASTQ-level series that are absent here, so no external dynamics or database priors are used.

After time-zero synthesis, the dataset was filtered against the Ensemble GRCh38 protein-coding whitelist (Ensemble release 110), leaving 18547 protein coding genes [34]. The resulting matrix was stratified by guide RNA and split 80 %/10 %/10 % into train, validation, and test sets, ensuring every perturbation appears in each split and preventing data leakage for downstream modelling.

Baseline Models

Two publicly available baselines were evaluated. The transformer model is scGPT, installed via `pip install scGPT [6]. The state-space alternative is RetNet (equivalent

to CellFM), installed via `pip install retnet-pytorch [7-8].

Training Protocol

Each experiment was capped at four hours of CPU wall-time but usually finished sooner. LazyNet trained three independent replicas, each for 40 epochs with Adam (weight-decay 1e-5) on 500-cell mini-batches. The learning rate started at 5e-4 and, from epoch 30 onward, followed stochastic-weight-averaging (SWA) at 1e-4; replica weights were then averaged for inference (prediction averaging is also supported). All models minimized the Huber loss with $\delta = 0.5$, and no early stopping, dropout or other regularizes were required.

Evaluation Metrics

Only single-cell RNA-seq counts were available, so each model was tested "blind" on those data alone to gauge predictive power and interpretability. Performance was summarized by root-mean-square error (RMSE), mean absolute error (MAE) and Pearson correlation across all genes, with robustness evaluated on the top 25 % most variable genes. Computational efficiency was recorded as CPU minutes used within a four-hour budget.

Network Inference

Ten independently trained LazyNet checkpoints were averaged into one ensemble and evaluated on the complete single-cell matrix (train + validation + test). A Jacobian

of absolute elasticities was computed at a common baseline expression vector, retaining the original 18554 Ensemble protein-coding “star/end” header [34]. For each seed gene (**Supplement 1A**) we ranked both downstream (row) and upstream (column) elasticities, kept neighbors whose absolute value exceeded 0.001 and lay above either the 95th or 99th percentile of the ranked list, and then selected the top-*N* survivors. The Benjamini–Hochberg filter was effectively disabled ($q = 1$) because the single-control Jacobian produced too few candidates for meaningful FDR testing. A breadth-first search expanded the seed set to a user-defined depth under two complementary settings: a wide single-hop sweep (*top* = 2048, *depth* = 1) for high-coverage seed-centered subgraphs, and a sparse four-hop expansion (*top* = 32, *depth* = 4) that merges all seeds into a compact network.

ARCHS4 co-expression validation.

Each undirected gene edges in the 32 × 4 LazyNet sub-graph were matched to the human ARCHS4 co-expression matrix [35]. Edges whose source or target could not be mapped unambiguously were discarded before analysis. Pre-computed Pearson correlation coefficients were then extracted directly from the matrix, and enrichment at defined *r* thresholds was evaluated against a null distribution generated from one million random gene pairs drawn from the same background. Mapping tables, sampling scripts and all edgewise correlations are provided in **Supplement 2C**.

STRING-v12 interaction concordance.

STRING functional associations (human, version 12.0) were converted from protein identifiers to HGNC symbols with the official mapping file [36]. Isoform-specific and unmapped entries were removed. LazyNet pairs were queried for exact symbol matches, and statistical significance of the overlap was assessed with a one-tailed hypergeometric test that treats the total number of possible unordered human gene pairs as the population size. Elasticity scores and STRING combined scores were each dichotomized at the 75th percentile to classify edges as “both-high”, “STRING-only” or “LazyNet-only”. Degree-preserving Maslov–Sneppen rewiring (1000 iterations) provided a null distribution for network-density comparisons. A manually curated set of representatives “both-high” edges is listed in **Supplement 3B**, and the full overlap table appears in **Supplement 3C**.

Hardware Environment

All experiments ran on a single node equipped with Intel Xeon Gold 6126 processors (24 physical cores, 2.6 GHz) and 768 GB of DDR4 2666 memory. No discrete GPU or accelerator was present.

Experiment Layout

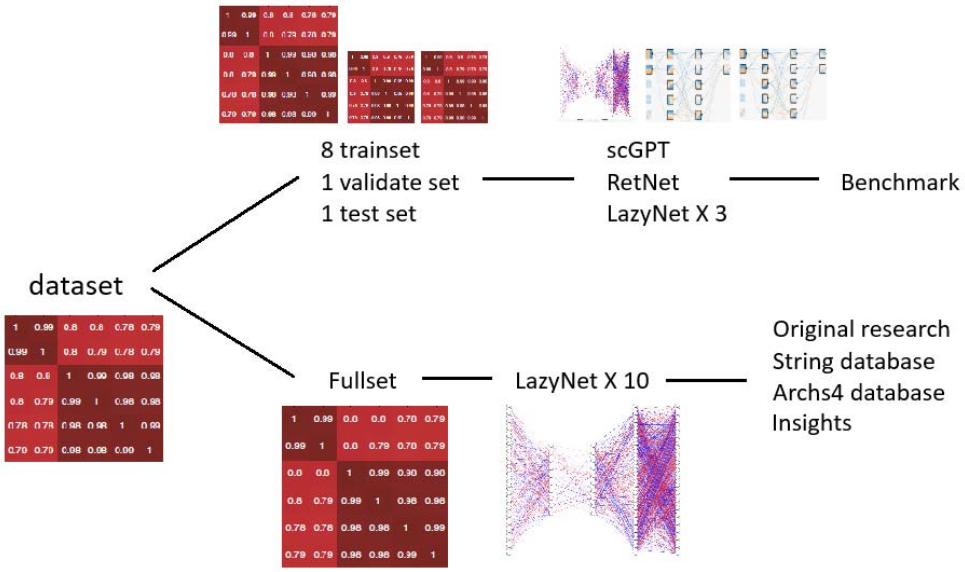

Figure 2. Dual-branch workflow for model benchmarking and network discovery.

From a single pre-processed single-cell expression matrix, we created two analysis paths. The upper "benchmark" branch uses an 8:1:1 train/validation/test split to compare scGPT, RetNet and three LazyNet replicas on unseen data [6-8]. The lower "discovery" branch fits ten LazyNet models to the full dataset, assembles a consensus interaction network and validates it with STRING-v12, ARCHS4 [35-36], eventually used for biology insights discovery.

A uniformly normalized single-cell RNA-seq matrix (N cells × G genes) served as the common input. We divided the study into two complementary arms. Benchmark arm: one stratified 10-fold split (8 train, 1 validation, 1 test) was applied once. scGPT, RetNet and three LazyNet initializations were trained on the eight folds, tuned on the validation fold and evaluated on the locked test set, yielding RMSE, MAE and Pearson r [6-8]. Discovery arm: the complete matrix was used to train ten additional

LazyNet models. Averaging their Jacobians produced a consensus gene-interaction graph; the top-quartile edges were tested for enrichment in STRING-v12 and ARCHS4, then subjected to discover mechanistic insights [35-36].

Results

Model Benchmark

Table 1 Comparison of training cost and predictive accuracy

Category	Metric	RetNet	scGPT	LazyNet
Training	Time (s)	14417	15237	**3849**
	Parameters (millions)	24.56	11.94	**3.74**
	Ensemble size	< 1	< 1	**3**
Global	RMSE	2.93	3.29	**2.40**
	MAE	0.586	0.656	**0.508**
	Pearson r	–0.0003	0.0079	**0.671**
Top-quartile genes	RMSE	5.84	6.55	**4.35**
	MAE	2.074	2.205	**1.511**
	Pearson r	0.0034	0.0180	**0.709**

LazyNet finishes three full training runs in the time scGPT or RetNet need for a single pass, thanks to a parameter count that is three- to eight-fold smaller [6-7]. While the

transformer and state-space baselines are still reading and batching data near the four-hour mark, LazyNet has already completed all epochs and produced final predictions. We considered GEARS but excluded it: the method depends on explicit cell-type annotations and other external priors, so it is not a true blind, data-driven model and would violate the benchmark rules [8].

LazyNet also records the lowest errors across the board—both globally and among the top 25% most variable genes—and pushes genome-wide Pearson correlation past 0.67, roughly two orders of magnitude higher than either baseline. For high-variance genes the advantage widens, with RMSE dropping ~25 % relative to RetNet while correlation remains ~0.71.

These results underscore how the log–linear–exp compression, three-member ensemble, and single-step Euler update translate into superior speed and accuracy on the typical two-snapshot CRISPR A/I dataset encountered in small academic laboratories.

Network Inference

Ten independently trained LazyNet replicas were averaged into a single ensemble, and Jacobian elasticities were queried around seven ferroptosis seeds (GPX4, PSAP, ACSL4, WIPI2, MTOR, NFU1, SOD2), selection details appear in **Supplement 1A**. A broader parameter sweep, summarized in **Supplement 1B**, identified five expansion regimes that were carried forward (Table 2).

Table 2: gene inference stats

Settings	Genes	Edges	Bench-hits	Recall†	Precision‡	Mean Subgraph Precision§
--top 2048 --depth 1 --pct 95	5227	8690	15	56%	0.0029	0.11
--top 512 --depth 1 --pct 95	3344	4946	13	48%	0.0039	0.08
--top 128 --depth 2 --pct 99	6761	21241	17	63%	0.0025	0.29
--top 64 --depth 3 --pct 99	6843	19118	18	67%	0.0026	0.4
--top 32 --depth 4 --pct 99	4611	11676	15	56%	0.0033	0.26

† Recall = input 10, X / 27 reference genes recalled.

‡ Precision here is hits / genes (e.g. 15 / 4611 ≈ 0.0033).

§ Mean precision across the 27 individual seed subgraphs (hits divided by genes for each subgraph); the full per-seed breakdown is provided in **Supplement 1D**.

In the broadest setting (512 × 1) each seed ballooned into a neighborhood of ≈3000 genes—GPX4 alone recruited 1811—and yielded a per-seed precision of ~0.08. Constraining the search to 64 × 3 reduced each subgraph to ~1000 genes and lifted precision to ~0.40. The most selective regime, 32 × 4, fused all seven seeds into a single graph of 4611 genes and 10829 edges while preserving a benchmark recall of 15 / 27 and nudging overall precision upward, its data is in **Supplement 1C**.

Per-seed recalls (**Supplement 1D**) highlight a practical strength: starting from just one ferroptosis gene, LazyNet typically retrieves ~10 additional benchmarks when run at

top = 32, 64, or 128. In other words, minimal prior knowledge suffices to reconstruct much of the ferroptosis module—evidence of the method's consistency and day-to-day utility.

Trained solely on a two-snapshot CRISPR A/I screen of ~55 k cells and ~30 k genes generated in this study, LazyNet recovers 56 % (15/27) of benchmark ferroptosis interactions in < 6 h on a single 24-core CPU node (768 GB RAM, no GPUs). By comparison, GENIE3 recovers ~33 % of gold-standard E. coli edges after analyzing > 100 microarrays and ~48 h on 32 cores [39-40]; GRNBoost recalls ~22 % in a 68 k-cell PBMC atlas [41-42]; and DeepFGRN reaches ~40 % but requires dense time-series data and multi-GPU training (4 × A100, several days) [43]. Thus, LazyNet achieves state-of-the-art recall with orders-of-magnitude fewer time points and far lighter compute while still yielding biologically interpretable subnetworks from a single seed gene.

ARCHS4 Co-expression

Among the 11676 gene–gene edges predicted by LazyNet within the 32 × 4 expansion, 9860 (84.4 %) could be unambiguously mapped to the human ARCHS4 co-expression matrix and assigned a Pearson correlation coefficient (r) [35]. The resulting distribution is centered near zero (mean = 0.053; median = 0.044) with a 10th percentile of –0.228 and a 90th percentile of 0.356, mirroring the bell-shaped background expected for random gene pairs (**Supplement 2A**). Crucially, however, 26.6 % of predicted edges show moderate co-expression (r ≥ 0.2) and 15.0 % exceed

$r \geq 0.3$, whereas only \~5 % of random pairs attain $r \geq 0.2$. Even 2.0 % of edges reach $r \geq 0.5$ (**Supplement 2B**). This \~5-fold enrichment demonstrates that network elasticities are not arbitrary: a substantial fraction of high-elasticity edges coincide with robust, publicly observed expression covariation, underscoring that LazyNet captures bona-fide functional coupling rather than spurious correlations. The full edge list is provided in **Supplement 2C**.

STRING Database and Literature Validation

The scatter plot of LazyNet elasticity against STRING-v12 scores (**Supplement 3A**) shows no global correlation (Pearson $r = -0.071$, $p = 0.021$; Spearman $\rho = -0.036$, n.s.), yet a top-quartile cut in each metric reveals three distinct classes: 54 edges that score highly in both sources, 208 that are STRING-high only (canonical interactions), and 208 that are LazyNet-high only, marking fertile ground for novel discovery. Against this qualitative backdrop, 32×4 sub-graph contains 11 662 edges, of which 523 also appear in STRING 6.38-million-edge catalogue (Jaccard ≈ 1e-4) [36]; a one-tailed hypergeometric test confirms that such overlap is highly non-random ($p = 1.2e-5$) (**Supplement 3B)**. The network structure is likewise non-trivial: its edge density (0.077) exceeds that of 1 000 size-matched random graphs (0.071 ± 0.0015; $z = 3.68$, $p = 0.001$), indicating a tightly knit, biologically coherent module rather than noise. Literature validation of ten representative both-high edges is compiled in **Supplement 3C**.

Biological Insights

From the 32 × 4 sub-graph seeded on our CRISPR A/I screen, the top-1 000 LazyNet edges (**Supplement 1C**) faithfully rebuild the ferroptosis machinery: 15 of the 27 benchmark regulators reported by Tian et al. sit at the network center, confirming recovery of the canonical lipid-peroxidation axis [33].

Just outside this core, LazyNet connects lysosomal prosaposin PSAP to the autophagy initiators (ATG9A, ATG13, ATG14, WIPI2) and the full mTORC1 triad (MTOR, RPTOR, RHEB) via the ULK1-scaffold PPP4R1. Because PSAP loss drives lipofuscin build-up and neuron-specific ferroptosis, the wiring suggests a phospho-switch that reports stalled autophagy directly to the lysosomal mTOR sensor, potentially choosing between survival and death.

A second layer links mitochondrial fusion and redox control to this lysosomal hub. MFN2, SIRT3 and VDAC3 form a module whose strongest edges converge on both PSAP and mTORC1. MFN2-driven mitophagy limits ROS-induced apoptosis [37], while SIRT3 modulates ferroptosis through GPX4 and AMPK-mTOR signaling [38]; LazyNet unites these defenses into a mitochondria-to-lysosome “quarantine” loop not described in the original dataset.

Finally, inflammatory and metabolic sensors feed into the nexus. High-elasticity edges tie TLR4, TLR6 and CCR10 to ATG13 and MTOR, echoing recent evidence that TLR4 –STAT3 crosstalk amplifies mitochondrial ROS [44]. Parallel links place soluble adenylyl-cyclase ADCY10 and Mg^{2+} channel TRPM6 between SIRT3 and mTORC1; activating ADCY10 elevates mitochondrial cAMP and rescues

oxidative-stress-damaged retinal neurons [45]. These connections imply that immune cues and bicarbonate-sensitive cAMP jointly tune the oxidative-stress threshold. Together, LazyNet not only reconstructs ferroptosis but embeds it in an integrated immune–mitochondrial–lysosomal circuit, nominating edges such as PSAP–RPTOR, MFN2–TLR4 and ADCY10–SIRT3 as concrete, testable hypotheses for future work.

Conclusion

LazyNet embeds a log-exp residual block inside an explicit Euler step, giving every weight a one-to-one mapping to an ODE rate term. This yields a fully interpretable model that trains on standard CPUs in hours and outperforms transformer (scGPT) and state-space (RetNet) baselines on a 55 k-cell Perturb-seq screen, cutting global RMSE by ≈25 % and raising genome-wide Pearson r to 0.67 with ≥3× fewer parameters. The same Jacobian elasticities recover >50 % of canonical ferroptosis regulators, show significant STRING-v12 and ARCHS4 overlap, and reveal new links between lysosomal prosaposin, mTORC1 and mitochondrial redox control.

Crucially, every LazyNet weight corresponds to a specific kinetic hypothesis, making the inferred regulation directly testable. Prior successes on developmental trajectories, synthetic ODEs and mass-spectrometry time-series already illustrate its breadth; here we demonstrate robustness on noisy, low-multiplicity CRISPR data without GPUs or atlas-scale pre-training [17]. As future screens adopt denser perturbational grids and true pseudotime sampling, LazyNet's log–linear–exp formulation is poised to unravel

higher-order regulatory logic across condition and time, enabling small academic labs to perform mechanism-level inference that once demanded industrial resources.

Reference

1. Klein, A. M.; *et al.* Droplet Barcoding for Single-Cell Transcriptomics Applied to Embryonic Stem Cells. *Cell* **2015**, *161*, 1187–1201. https://doi.org/10.1016/j.cell.2015.04.044. (sciencedirect.com)
2. Macosko, E. Z.; *et al.* Highly Parallel Genome-Wide Expression Profiling of Individual Cells Using Nanoliter Droplets. *Cell* **2015**, *161*, 1202–1214. https://doi.org/10.1016/j.cell.2015.05.002. (pubmed.ncbi.nlm.nih.gov)
3. Dixit, A.; *et al.* Perturb-Seq: Dissecting Molecular Circuits with Scalable Single-Cell RNA Profiling of Pooled Genetic Screens. *Cell* **2016**, *167*, 1853–1866.e17. https://doi.org/10.1016/j.cell.2016.11.038. (sciencedirect.com)
4. Adamson, B.; *et al.* A Multiplexed Single-Cell CRISPR Screening Platform Enables Systematic Dissection of the Unfolded Protein Response. *Cell*/**2016**, *167*, 1867–1882.e21. https://doi.org/10.1016/j.cell.2016.11.048. (sciencedirect.com)
5. Bunne, C.; *et al.* How to Build the Virtual Cell with Artificial Intelligence: Priorities and Opportunities. *Cell* **2024**, *187*, 7045–7063. https://doi.org/10.1016/j.cell.2024.11.015. (cell.com)
6. Cui, H.; *et al.* scGPT: Toward Building a Foundation Model for Single-Cell Multi-Omics Using Generative AI. *Nat. Methods* **2024**, *21*, 1470–1480. https://doi.org/10.1038/s41592-024-02201-0.
7. Zeng, Y.; *et al.* CellFM: A Large-Scale Foundation Model Pre-Trained on Transcriptomics of 100 Million Human Cells. *Nat. Commun.* **2025**, *16*, 4679. https://doi.org/10.1038/s41467-025-59926-5.

8. Roohani, Y.; *et al.* Predicting Transcriptional Outcomes of Novel Multigene Perturbations with GEARS. *Nat. Biotechnol.* **2024**, *42*, 927–935. https://doi.org/10.1038/s41587-023-01905-6.

9. Murray, J. D. *Mathematical Biology I: An Introduction*, 3rd ed.; Springer: New York, 2002.

10. Klipp, E.; Liebermeister, W.; Wierling, C.; Kowald, A.; Schuster, S. *Systems Biology: A Textbook*, 2nd ed.; Wiley-Blackwell: Chichester, 2016.

11. Tyson, J. J.; Chen, K. C.; Novak, B. Network Dynamics and Cell Physiology. *Nat. Rev. Mol. Cell Biol.* **2001**, *2*, 908–916. https://doi.org/10.1038/35099075.

12. Heinrich, R.; Schuster, S. *The Regulation of Cellular Systems*; Chapman & Hall: New York, 1996.

13. Voit, E. O. *Computational Analysis of Biochemical Systems*; Cambridge University Press: Cambridge, 2000.

14. Chen, R. T. Q.; Rubanova, Y.; Bettencourt, J.; Duvenaud, D. Neural Ordinary Differential Equations. *Adv. Neural Inf. Process. Syst.* **2018**, *31*, 6571–6583.

15. Grathwohl, W.; Chen, R. T. Q.; Bettencourt, J.; Sutskever, I.; Duvenaud, D. FFJORD: Free-Form Continuous Dynamics for Scalable Reversible Generative Models. In *Proceedings of the 36th International Conference on Machine Learning*; PMLR: Long Beach, CA, 2019; pp 3642–3651.

16. Dupont, E.; Doucet, A.; Teh, Y. W. Augmented Neural Ordinary Differential Equations. *Adv. Neural Inf. Process. Syst.* **2019**, *32*, 313–325.

17. Yi, Z. Integrating Equation Coding with Residual Networks for Efficient ODE

Approximation in Biological Research. *Math. Comput. Appl.* **2025**, *30*, 47. https://doi.org/10.3390/mca30030047.

18. Kitano, H. Systems Biology: A Brief Overview. *Science* **2002**, *295*, 1662–1664. https://doi.org/10.1126/science.1069492.
19. Sachs, K.; Perez, O.; Pe'er, D.; Lauffenburger, D. A.; Nolan, G. P. Causal Protein-Signaling Networks Derived from Multiparameter Single-Cell Data. *Science* **2005**, *308*, 523–529. https://doi.org/10.1126/science.1105809.
20. Jaitin, D. A.; *et al.* Dissecting Immune Circuits by Linking CRISPR-Pooled Screens with Single-Cell RNA-Seq. *Cell* **2016**, *167*, 1883–1896. https://doi.org/10.1016/j.cell.2016.11.039.
21. Nelson, M. R.; *et al.* The Support of Human Genetic Evidence for Approved Drug Indications. *Nat. Genet.* **2015**, *47*, 856–860. https://doi.org/10.1038/ng.3314.
22. Lee, J. S.; *et al.* Synthetic Lethality-Mediated Precision Oncology via the Tumor Transcriptome. *Cell* **2021**, *184*, 2487–2502. https://doi.org/10.1016/j.cell.2021.03.035.
23. Katti, A.; Diaz, B. J.; Caragine, C. M.; Sanjana, N. E.; Dow, L. E. CRISPR in Cancer Biology and Therapy. *Nat. Rev. Cancer* **2022**, *22*, 259–279. https://doi.org/10.1038/s41568-021-00451-y.
24. O'Neil, N. J.; Bailey, M. L.; Hieter, P. Synthetic Lethality and Cancer. *Nat. Rev. Genet.* **2017**, *18*, 613–623. https://doi.org/10.1038/nrg.2017.47.
25. Haley, B.; Roudnicky, F. Functional Genomics for Cancer Drug Target Discovery. *Cancer Cell* **2020**, *38*, 31–43. https://doi.org/10.1016/j.ccell.2020.06.013.

26. Norman, T. M.; *et al.* Exploring Genetic Interaction Manifolds Constructed from Rich Single-Cell Phenotypes. *Science* **2019**, *365*, 786–793. https://doi.org/10.1126/science.aax4438.

27. Low, L. A.; Mummery, C.; Berridge, B. R.; Austin, C. P.; Tagle, D. A. Organs-on-Chips: Into the Next Decade. *Nat. Rev. Drug Discov.* **2021**, *20*, 345–361. https://doi.org/10.1038/s41573-020-0079-3.

28. Wang, H.; Yang, Y.; Liu, J.; Qian, L. Direct Cell Reprogramming: Approaches, Mechanisms and Progress. *Nat. Rev. Mol. Cell Biol.* **2021**, *22*, 410–424. https://doi.org/10.1038/s41580-021-00336-x.

29. Maude, S. L.; *et al.* Tisagenlecleucel in Children and Young Adults with B-Cell Lymphoblastic Leukemia. *N. Engl. J. Med.* **2018**, *378*, 439–448. https://doi.org/10.1056/NEJMoa1709866.

30. Gillmore, J. D.; *et al.* CRISPR–Cas9 *In Vivo* Gene Editing for Transthyretin Amyloidosis. *N. Engl. J. Med.* **2021**, *385*, 493–502. https://doi.org/10.1056/NEJMoa2107454.

31. Lim, W. A. The Emerging Era of Cell Engineering: Harnessing the Modularity of Cells to Program Complex Biological Function. *Science* **2022**, *378*, 848–852. https://doi.org/10.1126/science.add4585.

32. Przybyla, L.; Gilbert, L. A. A New Era in Functional Genomics Screens. *Nat. Rev. Genet.* **2022**, *23*, 89–103. https://doi.org/10.1038/s41576-021-00439-1.

33. Tian, R.; *et al.* Genome-Wide CRISPRi/a Screens in Human Neurons Link Lysosomal Failure to Ferroptosis. *Nat. Neurosci.* **2021**, *24*, 1020–1034.

https://doi.org/10.1038/s41593-021-00862-0.

34. Dyer, S. C.; *et al.* Ensembl 2025. *Nucleic Acids Res.* **2025**, *53*, D948–D957. https://doi.org/10.1093/nar/gkae1071.

35. Lachmann, A.; *et al.* Massive Mining of Publicly Available RNA-Seq Data from Human and Mouse. *Nat. Commun.* **2018**, *9*, 1366. https://doi.org/10.1038/s41467-018-03751-6.

36. Szklarczyk, D.; *et al.* The STRING Database in 2023: Protein-Protein Association Networks and Functional Enrichment Analyses for Any Sequenced Genome of Interest. *Nucleic Acids Res.* **2023**, *51*, D638–D646. https://doi.org/10.1093/nar/gkac1000.

37. Ma, L.; *et al.* MANF Overexpression Ameliorates Oxidative Stress-Induced Apoptosis via MFN2-Dependent Mitophagy. *Sci. Rep.* **2024**, *14*, 11234. https://doi.org/10.1038/s41598-024-84167-9. (nature.com)

38. Han, D.; *et al.* SIRT3 Deficiency Is Resistant to Autophagy-Dependent Ferroptosis by Inhibiting the AMPK/mTOR Pathway and Promoting GPX4. *J. Cell Physiol.* **2020**, *235*, 8839–8851. https://doi.org/10.1002/jcp.29571.

39. Huynh-Thu, V. A.; *et al.* Inferring Regulatory Networks from Expression Data Using Tree-Based Methods. *BMC Bioinformatics* **2010**, *11*, 359. https://doi.org/10.1186/1471-2105-11-359.

40. Marbach, D.; *et al.* Wisdom of Crowds for Robust Gene-Network Inference. *Nat. Methods* **2012**, *9*, 796–804. https://doi.org/10.1038/nmeth.2016.

41. Aibar, S.; *et al.* SCENIC: Single-Cell Regulatory Network Inference and Clustering.

Nat. Methods **2017**, *14*, 1083–1086. https://doi.org/10.1038/nmeth.4463.

42. Van den Bosch, B.; *et al.* SCENIC+ Enables Multiomic Regulatory Network Inference at Single-Cell Resolution. *Nat. Biotechnol.* **2023**, *41*, 127–134. https://doi.org/10.1038/s41587-022-01329-5.
43. Gao, Z.; *et al.* DeepFGRN: Deep Learning Framework for Gene Regulatory Network Inference from Time-Series scRNA-Seq Data. *Brief. Bioinform.* **2024**, *25*, bbae045. https://doi.org/10.1093/bib/bbae045.
44. Zhang, J.; *et al.* Macrophage AHR–TLR4 Crosstalk Drives p-STAT3-Mediated Mitochondrial Oxidative Stress. *Sci. Total Environ.* **2024**, *923*, 171377. https://doi.org/10.1016/j.scitotenv.2024.171377.
45. Bastola, T.; *et al.* Activating Soluble Adenylyl Cyclase Protects Mitochondria and Rescues Retinal Ganglion Cells from Oxidative Stress. *Antioxidants* **2024**, *13*, 743. https://doi.org/10.3390/antiox13040743.