Predicting Context-Aware Transcriptional Responses to Unseen Genetic Perturbation Subject to Interactome Distance Constraints

Feiyu Ma[1], **Yunfei Zhang**[2 *], **Hau-San Wong**[3] and **Si Wu**[1]
[1]School of Computer Science and Engineering, South China University of Technology
[2]School of Future Technology, South China University of Technology
[3]Department of Computer Science, City University of Hong Kong
mfy18043849871@gmail.com, cszhangyunfei@mail.scut.edu.cn, cshswong@cityu.edu.hk, cswusi@scut.edu.cn

Abstract

Predicting responses to genetic perturbation is pivotal for elucidating gene regulatory machinery. However, existing methods often rely on statistical perspectives to model differential expression, overlooking the constraints of the underlying molecular interactome, which renders predictions susceptible to spurious correlations. Based on the fact that a genetic perturbation is a molecular stimulus acting through functional connectivity to reconfigure cellular expression profiles, we propose PertDCR, a framework for context-aware Perturbation response prediction via Distance-Constrained Refinement. Specifically, PertDCR addresses the heterogeneity of perturbation effects by contextualizing the perturbation source with relevant gene programs inferred from the protein-protein interaction network as well as cell representation. To activate biologically valid responses, PertDCR predicts the distance between the perturbation source and highly responsive genes within the PPI network, and dynamically modulates the response intensity based on distance, thereby facilitating accurate perturbation prediction. Extensive experiments demonstrate that PertDCR achieves state-of-the-art performance across diverse unseen genetic perturbations, and the predictions are consistent with underlying biological mechanisms.

1 Introduction

Elucidating transcriptional responses to genetic perturbations is fundamental to dissecting disease pathologies and identifying therapeutic targets [Kitano, 2002; Sachs *et al.*, 2005]. While high-throughput screens like Perturb-seq [Dixit *et al.*, 2016] and CRISPR-screen [Datlinger *et al.*, 2017] enable the profiling of gene regulatory networks at scale, the sheer combinatorial complexity of gene-cell interactions precludes exhaustive experimental enumeration [Chen *et al.*, 2024; Lu *et al.*, 2024; Lee *et al.*, 2021; Chen *et al.*, 2026]. Consequently, developing computational methods capable of predicting reliable responses to unseen perturbations is an imperative strategy to navigate the vast search space of potential interventions.

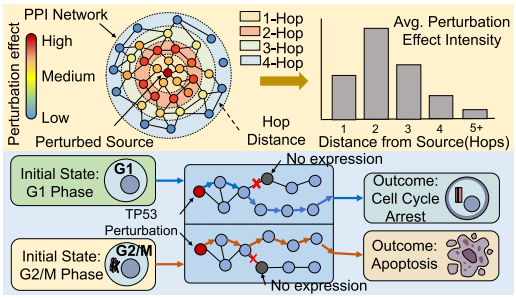

Figure 1: *Top*: Perturbation effects follow a distinct distribution constrained by interactome, peaking at specific hops rather than exhibiting simple linear decay. *Bottom*: TP53 activation illustrates context-aware heterogeneity, where the initial state (G1 vs. G2/M) drives distinct outcomes (cell cycle arrest vs. apoptosis).

To predict transcriptional responses to unseen perturbations, deep learning approaches typically focus on statistical distribution shifts induced by genetic interventions [Chen *et al.*, 2025]. scGen [Lotfollahi *et al.*, 2019] and CPA [Lotfollahi *et al.*, 2023] approximate the effects via latent space arithmetic, but reliance on simple algebraic operations limits their generalizability. To enhance representation robustness, scGPT [Cui *et al.*, 2024] and Geneformer [Theodoris *et al.*, 2023] leverage large-scale pretraining to extract co-expression patterns, yet the statistical focus remains susceptible to encoding spurious correlations. While recent frameworks like GEARS [Roohani *et al.*, 2024] and GRAPE [Chi *et al.*, 2025b] integrate biological priors to guide the feature learning, typically treating such knowledge merely as static inputs for feature aggregation, which offers limited performance gains. In summary, most existing approaches lack explicit modeling of interactome-based gene dependencies, which limits their ability to produce biologically valid responses.

Accurate response prediction necessitates moving beyond simple statistical correlations to incorporate the structural constraints of the underlying molecular interactome. A genetic perturbation triggers a signaling cascade that propagates through the protein-protein interaction (PPI) network to reconfigure downstream gene expression [Costanzo *et*

*Corresponding author.

al., 2016; Ideker and Krogan, 2012; Vidal *et al.*, 2011; Luck *et al.*, 2020]. As evidenced by the non-random propagation distribution in Figure 1, the network constrains the perturbation effects to specific neighborhoods. The distance-dependent propagation mechanism remains consistent even when generalizing to unseen perturbations [Kovács *et al.*, 2019; Cowen *et al.*, 2017]. Therefore, explicitly incorporating the interactome constraints provides a promising avenue to facilitate more precise and biologically plausible predictions.

In this work, we propose PertDCR, a framework for context-aware Perturbation response prediction via Distance-Constrained Refinement, enhancing the biological plausibility of the outcomes. Specifically, we adopt an Interactome-Embedded Graph Transformer (IEGT) to integrate the expression profile with prior information to learn the prior-informed gene and cell embeddings. To capture the available interaction network for the propagation of genetic effects, the Context-Aware Response Prediction module (CARP) employs task-relevant learnable queries to retrieve analogous gene programs, which are concatenated with cell and perturbation embeddings to form adaptive perturbation conditions. The resulting conditions are fed into a multi-task prediction module to derive a response tailored to the specific cellular and perturbation context. To further activate responses of the highly responsive genes, the preliminary response intensity is dynamically refined by the Distance-Constrained Refinement Module (DCRM) based on the predicted distance from the perturbation source within the PPI network, thereby yielding a final outcome with enhanced biological verifiability and prediction accuracy. The main contributions of this work are summarized as follows:

- Distinct from the existing methods focusing on statistical differential expression, PertDCR predicts responses subject to interactome constraints that align with the underlying biological mechanisms of genetic perturbation.
- We construct adaptive perturbation conditions by contextualizing the perturbation source with relevant gene programs, thereby yielding context-aware responses.
- Considering that genetic perturbations exert influence through functional connectivity, we refine the response intensity based on the distance from the perturbation source within the PPI network, enhancing the biological plausibility of the predictions and achieving state-of-the-art performance across diverse unseen perturbation.

2 Related Work

2.1 Deep Learning Approaches

Early computational methods for perturbation prediction primarily relied on linear assumptions [Kamimoto *et al.*, 2023]. Perturb-Seq modeled perturbation effects using regularized linear regression. However, the assumption of linear superposition limits its ability to capture complex non-additive gene interactions. Deep generative models have addressed these limitations and become a dominant paradigm. scGen pioneered this direction by using Variational Autoencoders (VAEs) and latent space arithmetic to predict responses. Subsequent frameworks extended this approach to handle more complex biological contexts. For example, CPA introduced adversarial disentanglement to predict combinatorial perturbations and continuous dosage effects. SAMS-VAE [Bereket and Karaletsos, 2023] incorporated sparse global latent variables to enhance the interpretability of the learned representations. Recent works have adopted advanced generative AI to better model unpaired single-cell data. scPPDM [Liang *et al.*, 2025] utilizes conditional diffusion models with dual-channel conditioning to generate high-fidelity perturbation profiles. Similarly, Departures [Chi *et al.*, 2025a] utilizes the Neural Schrödinger Bridge to achieve energy-optimal transport between control and perturbed distributions.

2.2 Single-Cell Foundation Models

Inspired by large language models, Transformer-based single-cell foundation models are reshaping perturbation prediction through large-scale self-supervised pretraining. scGPT pretrained a generative Transformer on over 33 million cells, treating genes as tokens to learn universal cell and gene embeddings for transfer learning. Geneformer [Theodoris *et al.*, 2023] introduced gene ranking-based language modeling on up to 100 million cells, capturing context-dependent gene interactions with strong cross-tissue transfer capabilities. To enhance zero-shot generalization, scGenePT [Istrate *et al.*, 2024] injected LLM-generated gene functional descriptions as prompts, demonstrating that natural language knowledge improves performance on unseen perturbations. RegFormer [Hu *et al.*, 2025] combines gene regulatory network (GRN) priors with Mamba architecture, using GRN-guided gene ordering and regulatory-aware pretraining objectives to model transcriptional hierarchies.

2.3 Knowledge-Enhanced Methods

Recognizing limitations of purely data-driven approaches [Zhou *et al.*, 2025; Wei *et al.*, 2025], recent work integrates biological prior knowledge for perturbation prediction. GEARS pioneered integrating deep learning with Gene Ontology-based knowledge graphs, enabling prediction for gene combinations never experimentally tested. Similarly, GRAPE introduced biotype information and constructed heterogeneous graphs with distinct node types for coding and non-coding genes to capture implicit gene relationships. Other methods focus on unifying different data types. PerturbNet [Yu *et al.*, 2025] unified multiple perturbation modalities to predict responses across small molecules, genetic perturbations, and missense mutations. Additionally, Biolord [Piran *et al.*, 2024] developed a disentanglement framework separating known attributes from unknown attributes via information constraints, enabling counterfactual generation of inaccessible cell states.

We summarize the key distinctions from existing methods as follows: (1) Unlike methods such as scGPT and Geneformer, which implicitly learn perturbation effects from data-driven statistical patterns, we construct adaptive perturbation conditions by retrieving analogous gene programs to derive context-aware responses; (2) In contrast to methods like GEARS, which incorporate prior information as static inputs for feature aggregation, we focus on directly refining the response intensity based on interactome distance constraints.

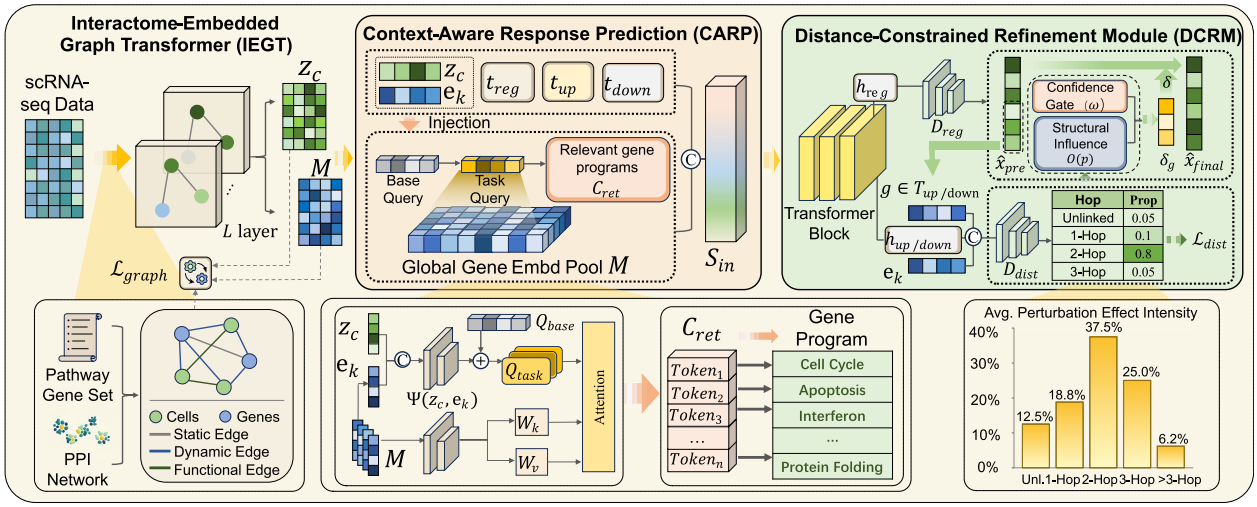

Figure 2: Overview of the proposed PertDCR. The Interactome-Embedded Graph Transformer (IEGT) integrates the control profile $\mathbf{x}_c$ with pathway and PPI networks, encoding prior-informed gene and cell embeddings. Subsequently, the Context-Aware Response Prediction (CARP) module employs task-relevant learnable queries to retrieve analogous gene programs, which are fused with the cell and perturbation embeddings to form adaptive perturbation conditions. These conditions are processed by a multi-task prediction module to derive preliminary responses. To activate biologically valid responses, the Distance-Constrained Refinement Module (DCRM) dynamically modulates the response intensity based on predicted topological distances, yielding the biologically plausible outcome $\hat{\mathbf{x}}_{final}$.

3 Methodology

3.1 Overview

We define the control population as $\mathcal{D}_{ctr} = \{\mathbf{x}_{ctr}^{(j)}\}_{j=1}^{M}$ and the perturbed population as $\mathcal{D}_{pert} = \{(\mathbf{x}_{pert}^{(i)}, k^{(i)})\}_{i=1}^{N}$. $\mathbf{x} \in \mathbb{R}^{|\mathcal{G}|}$ represents the gene expression profile, the perturbation source $k^{(i)}$ belongs to the set of genes $\mathcal{G}$. We ensure that $\mathcal{P}_{train} \cap \mathcal{P}_{test} = \emptyset$, where $\mathcal{P}_{train}$ and $\mathcal{P}_{test}$ be the sets of perturbed genes in training and testing, respectively. Our objective is to learn a mapping function $f_\theta : (\mathbf{x}_{ctr}, k) \rightarrow \hat{\mathbf{x}}_{final}$ that accurately predicts the transcriptional responses.

In this work, the Interactome-Embedded Graph Transformer (IEGT) processes the control profile $\mathbf{x}_{ctr}$, the pathway and the PPI network to encode the gene and cell embeddings. Subsequently, the Context-Aware Response Prediction (CARP) module employs task-relevant learnable queries to retrieve analogous gene programs. These retrieved tokens, fused with the cell and perturbation embedding, form the perturbation conditions. The resulting conditions are fed into a multi-task prediction module, which outputs a regression token for expression prediction and distinct activation/inhibition tokens for hops inference. Finally, the Distance-Constrained Refinement Module (DCRM) utilizes these tokens to predict the distances for highly responsive genes. The preliminary responses are then modulated based on the predicted distances to yield the final outcome $\hat{\mathbf{x}}_{final}$. The overall framework is shown in Figure 2.

3.2 Interactome-Embedded Graph Transformer

To integrate transcriptomic expression with prior information into cell and gene embedding, we construct a graph $G = (\mathcal{V}_c \cup \mathcal{V}_g, \mathcal{E}_{het})$ comprising both cell nodes $\mathcal{V}_c$ and gene nodes $\mathcal{V}_g$. The edges $\mathcal{E}_{het} = \mathcal{E}_{ppi} \cup \mathcal{E}_{exp} \cup \mathcal{E}_{path}$ include gene-gene edges derived from the PPI network, cell-gene edge weighted by expression levels, and cell-cell edge constructed via K-Nearest Neighbors on pathway activity scores.

Operating on the graph G, the IEGT E iteratively aggregates information from neighborhoods to update node representations. For a target gene node i at layer l, the updated feature $\mathbf{h}_i^{(l)}$ is calculated as follows:

$$\mathbf{h}_i^{(l)} = \psi\left(\sum_{j \in \mathcal{N}(i)} \alpha_{ij}^{(l)} \mathbf{W}_V^{graph} \mathbf{h}_j^{(l-1)}\right) + \mathbf{h}_i^{(l-1)}, \quad (1)$$

where $\mathcal{N}(i)$ represents the neighbor set of node i, $\mathbf{W}_V^{graph}$ denotes a learnable value projection matrix, and ψ denotes a feed-forward network. $\alpha_{ij}^{(l)}$ is computed via the interaction between the node i and the node j as follows:

$$\alpha_{ij}^{(l)} = \underset{j \in \mathcal{N}(i)}{\sigma_{soft}}\left(\frac{(\mathbf{W}_Q^{graph}\mathbf{h}_i^{(l-1)})^\top (\mathbf{W}_K^{graph}\mathbf{h}_j^{(l-1)})}{\sqrt{d}}\right), \quad (2)$$

where d is the scaling dimension, σ_{soft} denotes the softmax function and $\mathbf{W}_Q^{graph}, \mathbf{W}_K^{graph}$ are query and key projection matrices. By assigning varying α_{ij} to different edges, E adaptively modulates the influence of priors. After L layers of propagation, E yields the prior-informed cell embedding $\mathbf{z}_c$ and a global gene embedding pool $\mathbf{M} = \{\mathbf{m}_i\}_{i \in \mathcal{G}}$.

Furthermore, we impose a structural consistency regularization that encourages embeddings similarity between connected gene pairs in the PPI network while separating unconnected ones, defined as:

$$\mathcal{L}_{gene} = -\mathbb{E}_{(i,j)\sim\mathcal{E}_{ppi}}\left[\log \frac{\exp(\mathbf{m}_i^\top \mathbf{m}_j/\tau)}{\sum_{n \in \mathcal{G}} \exp(\mathbf{m}_i^\top \mathbf{m}_n/\tau)}\right], \quad (3)$$

where τ represents a temperature hyperparameter. Analogously, a structural constraint $\mathcal{L}_{cell}$ is applied to the cell-cell similarity graph. $\mathcal{L}_{graph} = \mathcal{L}_{gene} + \mathcal{L}_{cell}$ forces E to capture global structural dependencies, establishing a robust foundation for the subsequent perturbation prediction task.

3.3 Context-Aware Response Prediction

To account for the heterogeneity of perturbation effects across different cells, the Context-Aware Response Prediction module (CARP) constructs adaptive perturbation conditions to synthesize a preliminary response. To capture the available interaction network for the propagation of genetic effects, we instantiate a set of task-relevant queries $\mathbf{Q}_{task}$ by modulating a set of learnable base tokens $\mathbf{Q}_{base}$ with the specific condition vector derived from the cell embedding and perturbation source: $\mathbf{Q}_{task} = \mathbf{Q}_{base} + \Psi(\mathbf{z}_c, \mathbf{e}_k)$, where Ψ is a function implemented as concatenation followed by a linear projection. $\mathbf{Q}_{task}$ serve as probes to retrieve relevant gene programs from $\mathbf{M}$, generating retrieved contexts $\mathbf{C}_{ret}$ as follows:

$$\mathbf{C}_{ret} = \sigma_{soft}\left(\frac{\mathbf{W}_Q^{ret}\mathbf{Q}_{task}(\mathbf{M}\mathbf{W}_K^{ret})^\top}{\sqrt{d}}\right)\mathbf{M}\mathbf{W}_V^{ret}, \quad (4)$$

where $\mathbf{W}_Q^{ret}$, $\mathbf{W}_K^{ret}$, $\mathbf{W}_V^{ret}$ are learnable query, key and value matrices for the retrieval mechanism. Subsequently, CARP constructs the perturbation conditions by concatenating $\mathbf{C}_{ret}$ with the cell embedding, the perturbation source, and a set of specialized learnable tokens $\mathbf{T}_{special} = [\mathbf{t}_{reg}, \mathbf{t}_{up}, \mathbf{t}_{down}]$. These specialized tokens are designed to separately aggregate latent information for the expression regression task and the distance prediction of activation/ repression targets. The composite perturbation conditions $\mathbf{S}_{in} = [\mathbf{C}_{ret}, \mathbf{z}_c, \mathbf{e}_k, \mathbf{T}_{special}]$ are fed into the multi-task prediction module.

The updated regression token $\mathbf{h}_{reg}$ is projected via a regression decoder D_{reg} to generate the preliminary post-perturbation expression estimate: $\hat{\mathbf{x}}_{pre} = D_{reg}(\mathbf{h}_{reg})$. The preliminary response serves as context-aware baseline, and the updated activation token $\mathbf{h}_{up}$ and repression token $\mathbf{h}_{down}$ are reserved for the subsequent distance inference.

3.4 Distance-Constrained Refinement

To activate biologically valid responses while at the same time suppressing spurious correlations, the Distance-Constrained Refinement Module (DCRM) dynamically modulates the preliminary response $\hat{\mathbf{x}}_{pre}$ based on interactome distance. Given that genetic perturbations trigger sparse regulatory cascades in reality, DCRM focuses on the set of all highly responsive genes $\mathcal{T} = \mathcal{T}_{up} \cup \mathcal{T}_{down}$ selected by predicted shift magnitudes, where $\mathcal{T}_{up}$ and $\mathcal{T}_{down}$ comprise genes with the strongest predicted upregulation and downregulation, respectively.

For each candidate gene $g \in \mathcal{T}_{up}$, a distance predictor D_{dist} utilizes the context-aware token $\mathbf{h}_{up}$ to estimate the probability distribution of topological hops from source $\mathbf{e}_k$

$$\mathbf{p}_{k,g} = \sigma_{soft}\left(D_{dist}([\mathbf{h}_{up} \parallel \mathbf{e}_k \parallel \mathbf{m}_g])\right), \quad (5)$$

where $[\cdot \parallel \cdot]$ denotes concatenation operation. An identical process using the repression token $\mathbf{h}_{down}$ is applied to the repression set $\mathcal{T}_{down}$. To ensure that the predicted probabilities reflect authentic molecular interactome, we impose a distance constraint $\mathcal{L}_{dist}$ as follows:

$$\mathcal{L}_{dist} = -\sum_{g\in\mathcal{T}}\sum_{b=1}^{B}\mathbb{I}(b_{gt} = b)\log(p_{k,g,b}), \quad (6)$$

where b_{gt} denotes the ground-truth hop distance in the PPI network, $p_{k,g,b}$ is the b-th element of $\mathbf{p}_{k,g}$, B is the number of distance intervals, and $\mathbb{I}$ is the indicator function.

To flexibly refine the response intensity based on predicted distance, we calculate a scalar correction term δ_g for each target gene $g \in \mathcal{T}$ as follows:

$$\delta_g = (\mathbf{p}_{k,g}^\top\omega)\cdot O(\mathbf{p}_{k,g}), \quad (7)$$

where ω denotes learnable weights assigning importance to each hop interval. The first term $(\mathbf{p}_{k,g}^\top\omega)$ quantifies the expected structural influence of the perturbation on the target gene. The second term $O(\mathbf{p}_{k,g})$ serves as a confidence gate that modulates the correction magnitude based on the uncertainty of the distance prediction, and is formulated as follows:

$$O(\mathbf{p}) = \sigma_{sig}\left(\beta_0\cdot(1-\frac{H(\mathbf{p})}{H_{max}}) + \beta_1\cdot(p_{(1)} - p_{(2)}) + \beta_2\right), \quad (8)$$

where $H(\mathbf{p})$ is the entropy measuring prediction ambiguity, normalized by the maximum entropy H_{max}, $p_{(1)} - p_{(2)}$ represents the margin between the top two probabilities, $\beta_0, \beta_1, \beta_2$ are learnable parameters, and σ_{sig} denotes sigmoid function. The confidence gate ensures that refinements are applied only when the model is confident in the propagation process, preventing corrections driven by ambiguous structural inference. The final transcriptomic profile is updated as follows:

$$\hat{\mathbf{x}}_{final} = \hat{\mathbf{x}}_{pre} + \boldsymbol{\delta}, \quad (9)$$

where $\boldsymbol{\delta}$ is a sparse vector populated with δ_g at targeted indices ($g \in \mathcal{T}_{up} \cup \mathcal{T}_{down}$) and zero elsewhere. DCRM refines the response based on the distance between genes in the PPI network, enhancing the biological plausibility of final response.

3.5 Model Optimization

To align the final transcriptomic profile $\hat{\mathbf{x}}_{final}$ with the experimentally observed state, we enforce a regression objective that minimizes the divergence between the predicted responses and the ground truth:

$$\mathcal{L}_{reg} = \mathbb{E}_{\mathbf{x}_{ctr}\sim\mathcal{D}_{ctr},(\mathbf{x}_{pert},k)\sim\mathcal{D}_{pert}}\left[\|\hat{\mathbf{x}}_{final} - \mathbf{x}_{pert}\|_2^2\right] \quad (10)$$

where $\mathbf{x}_{pert}$ denotes the verified post-perturbation gene expression profile associated with target k. In summary, the total optimization objective is formulated as follows:

$$\mathcal{L}_{total} = \mathcal{L}_{reg} + \lambda\mathcal{L}_{dist} + \gamma\mathcal{L}_{graph} \quad (11)$$

where λ is the weight of $\mathcal{L}_{dist}$ and γ is the weight of $\mathcal{L}_{graph}$. By minimizing $\mathcal{L}_{total}$, PertDCR enhances the robust generalization to unseen perturbations by refining predictions based on the interactome distance.

4 Experiment

4.1 Experimental Settings

We begin by introducing the experimental settings. Subsequently, we benchmark PertDCR against state-of-the-art methods and conduct ablation studies to validate individual components. We further investigate the interpretability of the distance-constrained refinement and assess the model's generalization capabilities on cross-cell-line datasets.

Datasets. To comprehensively evaluate the performance of PertDCR, we utilize datasets spanning diverse cell lines, including Adamson, Norman, Datlinger, Papalexi, Replogle and Tian [Peidli *et al.*, 2022] . Notably, the Norman dataset includes combinatorial perturbation types. To further validate the capability of PertDCR in handling perturbations across different cell lines, we constructed a curated cross-cell-line dataset by integrating Datlinger, Papalexi, and Tian, referred to Mixture dataset. For data preprocessing, we follow the standard protocols established in GEARS [Roohani *et al.*, 2024] and scGPT [Cui *et al.*, 2024] and enforce a rigorous data splitting strategy where the perturbation types in the test set are strictly disjoint from those in the training set. The PPI network is constructed using data from the STRING database [Szklarczyk *et al.*, 2022]. The concise description of datasets is provided in Table 1.

Implementation details. All models are implemented in PyTorch on an NVIDIA A6000 GPU. We set the distance into $B = 5$ intervals representing 1-hop, 2-hop, 3-hop, >3-hop and unlinked. The hyperparameters λ and γ are initialized to 0.05 and 0.1, respectively. We optimize the model using the Adam optimizer [Kingma and Ba, 2015] with a learning rate of 1×10^{-4} and a batch size of 64.

Evaluation metrics. To comprehensively assess the performance of PertDCR, we employ five metrics. Pearson Correlation (Pearson) measures the linear relationship between the predicted and ground-truth expression profiles across all genes. Pearson Delta (Pearson$_\delta$) calculates the correlation between the predicted and actual expression changes (delta) over the entire genome. Pearson DE Delta (Pearson$_{\delta+\text{de}}$) restricts this correlation calculation to the top-20 Differentially Expressed Genes (DEGs) that are most significantly perturbed. Directionality Accuracy (Dir. Acc) verifies whether the model correctly identifies the sign of regulation (upregulation or downregulation) for genes with significant changes. Pathway NES Correlation (NES Cor) measures the alignment of gene set enrichment analysis scores [Tilford and Siemers, 2009] between prediction and ground truth.

Dataset	Cell Line	Perturbation	Scale (#cells × #genes)
Adamson [Adamson *et al.*, 2016]	K562	Single-gene	68,603 × 5,060
Norman [Norman *et al.*, 2019]	A549	Single&Double-gene	91,205 × 5,045
Datlinger [Datlinger *et al.*, 2017]	Jurkat	Single-gene	5,905 × 5,009
Papalexi [Papalexi *et al.*, 2021]	THP-1	Single-gene	20,729 × 5,016
Replogle [Replogle *et al.*, 2022]	K562	Single-gene	171,542 × 6,017
Tian [Tian *et al.*, 2021]	iNeuron	Single-gene	4,844 × 5,117
Mixture	/	Single-gene	15,899 × 5,125

Table 1: Dataset descriptions.

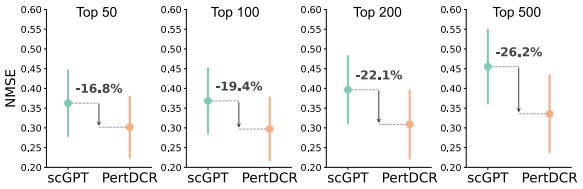

Figure 3: Comparison of NMSE between PertDCR and scGPT across varying ranges of DEGs on the Adamson dataset.

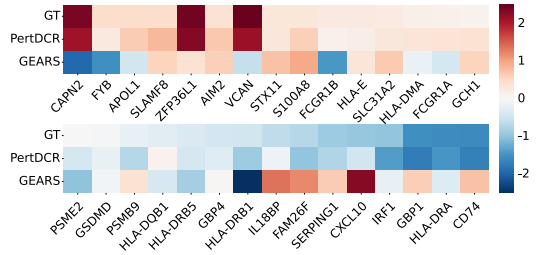

Figure 4: Heatmap comparing the Z-score normalized expression of activated (*top*) and inhibited (*bottom*) DEGs between PertDCR and GEARS under JAK2 perturbation on the Papalexi dataset.

4.2 Comparison to State-of-the-Arts

To demonstrate the robustness of PertDCR across diverse perturbation types, we compare PertDCR with the competing methods, including Linear Regression (LR), CPA, GEARS, scGPT, biolord and PerturbNet on different datasets. As detailed in Table 2, PertDCR consistently achieves superior performance across all metrics. On the Datlinger dataset, PertDCR achieves a Pearson$_\delta$ of 0.556 and a Pearson$_{\delta+\text{de}}$ of 0.754, outperforming scGPT which achieves 0.382 and 0.514 respectively. Similarly, on the Replogle dataset, PertDCR achieves a Pearson$_\delta$ 0.588 compared to 0.250 for scGPT and 0.154 for GEARS. These substantial margins indicate that integrating interactome constraints allows the model to effectively predict precise response. For the Norman dataset which contains combinatorial two-gene perturbations, PertDCR achieves 0.618 in terms of NES Cor, surpassing the GEARS score of 0.580, suggesting that DCRM in PertDCR effectively enhances the biological plausibility of the predictions even in complex multi-gene perturbation scenarios.

Additionally, we visualize the Normalized Mean Squared Error (NMSE) which quantifies the error normalized by the variance of the ground truth on the Adamson dataset across varying numbers of top DEGs. As illustrated in Figure 3, PertDCR reduces the NMSE by about 17 percentage points compared to scGPT on the top-50 DEGs and expands this advantage to a 26.2% reduction on the top-500 DEGs, demonstrating that PertDCR offers precise predictions for the targets that undergo the most significant transcriptional shifts. We also visualize the Z-score normalized expression profiles for the top-30 DEGs following JAK2 perturbation on the Papalexi dataset in Figure 4. PertDCR exhibits remarkable concordance with the ground truth, accurately reconstructing both the strongly upregulated cluster and the downregulated cluster in terms of direction and magnitude. In contrast, GEARS displays directional reversals for several key genes.

Method	Venue	Adamson					Norman					Datlinger				
		Pearson	Pearson_{δ}	$\text{Pearson}_{\delta+\text{de}}$	Dir. Acc	NES Cor	Pearson	Pearson_{δ}	$\text{Pearson}_{\delta+\text{de}}$	Dir. Acc	NES Cor	Pearson	Pearson_{δ}	$\text{Pearson}_{\delta+\text{de}}$	Dir. Acc	NES Cor
LR	-	0.9826	0.387	0.620	78.0%	0.351	0.9877	0.523	0.634	80.7%	0.453	0.9884	0.215	0.288	60.7%	0.178
CPA	Mol. Syst. Biol. (2023)	0.9280	0.121	0.494	61.5%	0.113	0.9817	0.177	0.297	55.5%	0.101	0.9593	0.022	0.441	56.1%	0.071
GEARS	Nat. Biotechnol. (2023)	0.9892	0.531	0.678	81.8%	0.574	0.9845	0.583	0.715	85.7%	0.580	0.9777	0.031	0.011	54.0%	0.057
scGPT	Nat. Methods (2024)	0.9906	0.615	0.769	84.8%	0.613	0.9852	0.561	0.732	85.2%	0.549	0.9761	0.382	0.514	77.4%	0.352
biolord	Nat. Biotechnol. (2024)	0.9914	0.285	0.725	55.6%	0.179	0.9854	0.528	0.659	81.8%	0.436	0.9895	0.203	0.241	60.8%	0.149
PerturbNet	Mol. Syst. Biol. (2025)	0.7309	0.146	0.189	72.8%	0.003	0.4840	0.066	0.062	43.5%	0.054	0.8421	0.010	0.351	48.9%	0.011
PertDCR	-	**0.9932**	**0.747**	**0.841**	**89.6%**	**0.645**	**0.9878**	**0.587**	**0.734**	**86.5%**	**0.618**	**0.9890**	**0.556**	**0.754**	**83.6%**	**0.473**
Method	Venue	Papalexi					Tian					Replogle				
		Pearson	Pearson_{δ}	$\text{Pearson}_{\delta+\text{de}}$	Dir. Acc	NES Cor	Pearson	Pearson_{δ}	$\text{Pearson}_{\delta+\text{de}}$	Dir. Acc	NES Cor	Pearson	Pearson_{δ}	$\text{Pearson}_{\delta+\text{de}}$	Dir. Acc	NES Cor
LR	-	0.9774	0.277	0.341	65.4%	0.264	0.9876	0.200	0.438	60.3%	0.134	0.9929	0.131	0.261	74.8%	0.340
CPA	Mol. Syst. Biol. (2023)	0.9671	0.130	0.263	61.0%	0.142	0.9800	0.004	0.307	53.7%	0.008	0.9769	0.017	0.023	47.4%	0.003
GEARS	Nat. Biotechnol. (2023)	0.9794	0.324	0.505	69.8%	0.304	0.9850	0.051	0.200	54.4%	0.017	0.9920	0.154	0.277	74.8%	0.347
scGPT	Nat. Methods (2024)	0.9841	0.296	0.421	72.8%	0.315	0.9842	0.632	0.842	83.4%	0.495	0.9918	0.250	0.451	75.1%	0.345
biolord	Nat. Biotechnol. (2024)	0.9886	0.362	0.484	75.8%	0.211	0.9816	0.143	0.150	58.8%	0.086	0.9926	0.238	0.427	69.1%	0.303
PerturbNet	Mol. Syst. Biol. (2025)	0.8913	0.124	0.267	52.9%	0.003	0.6947	0.040	0.011	45.4%	0.051	0.6344	0.027	0.069	50.9%	0.012
PertDCR	-	**0.9886**	**0.488**	**0.649**	**78.8%**	**0.401**	**0.9856**	**0.656**	**0.887**	**83.4%**	**0.510**	**0.9932**	**0.588**	**0.695**	**86.2%**	**0.421**

Table 2: Quantitative comparison on six datasets for unseen perturbation prediction.

Variants	Pearson	Pearson_{δ}	$\text{Pearson}_{\delta+\text{de}}$	Dir. Acc	NES Cor
PertDCR w/o PE	0.9910	0.694	0.775	86.8%	0.572
PertDCR w/o RM	0.9923	0.701	0.785	88.0%	0.585
PertDCR w/o DCRM	0.9926	0.717	0.811	88.3%	0.633
PertDCR w/o CE	0.9926	0.713	0.819	88.6%	0.608
PertDCR	**0.9932**	**0.747**	**0.841**	**89.6%**	**0.645**

Table 3: Quantitative comparison of PertDCR and different variants.

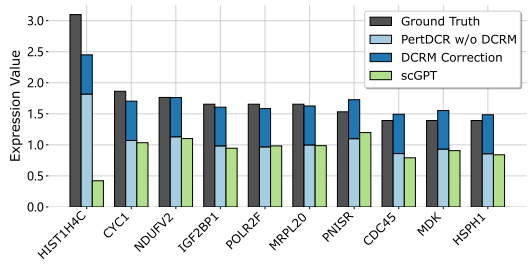

Figure 5: Performance comparison on the top-10 DEGs for PRMT5 perturbation. The total height of the stacked bar indicates the final prediction of PertDCR, composed of the preliminary prediction (light blue) and the learned DCRM Correction (dark blue).

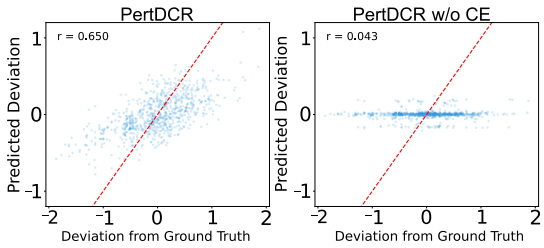

Figure 6: Visualization of expression deviations (PertDCR vs. Ground Truth) from the population mean for the CHERP perturbation. Each point represents a gene-cell pair. The higher Pearson correlation (r) and alignment with the identity line (red dashed) indicate that PertDCR better preserves cellular heterogeneity.

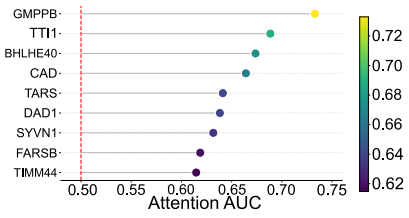

Figure 7: Ranking of the top-10 perturbation types based on attention AUC for ground truth downstream targets. The red dashed line represents the random baseline.

4.3 Analysis of the Main Components

To comprehensively demonstrate the roles of the main modules in our model, we conduct a series of experiments with quantitative results summarized in Table 3.

Is the distance-constrained refinement effective? To verify the effectiveness of the refinement module, we compare PertDCR with 'PertDCR w/o DCRM' which omits the final distance-based rectification. As presented in Table 3, the removal of DCRM leads to a clear decline in Pearson_{δ} from 0.747 to 0.717. We further visualize the top-10 differentially expressed genes for the PRMT5 perturbation in Figure 5. As shown by the stacked bars, the DCRM correction term (dark blue) dynamically compensates for the preliminary prediction (light blue), effectively bridging the gap to the Ground Truth. Consequently, the final response (total height) demonstrates superior alignment with the actual expression levels compared to scGPT, confirming that the module accurately scales refinements according to structural constraints.

Is the specific cell embedding necessary? We build a variant, 'PertDCR w/o CE', which disables the injection of specific Cell Embedding. The results in Table 3 show a decrease in $\text{Pearson}_{\delta+\text{de}}$ to 0.819, indicating a loss in cell-specific predictive capability. We further visualize the alignment between the predicted and ground truth deviations from the mean expression in Figure 6. PertDCR exhibits a positive correlation of 0.650 with data points distributed along the diagonal while the variant displays a collapsed distribution where the majority of predicted values cluster around zero deviation. This implies that without cell embeddings, the model fails to distinguish individual cellular states and degenerates into predicting the population average, thereby losing the biological heterogeneity intrinsic to single-cell data.

Does the retrieval mechanism capture meaningful gene programs? We analyze the attention distribution of the learnable queries using the Area Under the Curve metric as illus-

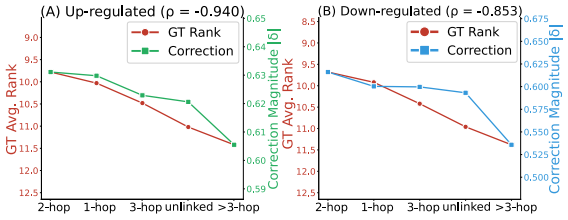

Figure 8: Concordance between ground truth gene ranking and learned correction magnitudes. The top-20 DEGs are stratified by topological distance (X-axis) and ordered by their ground truth perturbation intensity (Left Y-axis, lower rank indicates stronger variation). The parallel trend with the model's correction magnitude $|\delta|$ (Right Y-axis) and the strong negative correlation (ρ) demonstrate that the DCRM dynamically scales refinements according to the actual severity of the perturbation effect at each structural depth.

trated in Figure 7. PertDCR achieves a mean AUC of 0.659 in distinguishing the top-40 most responsive genes from other genes, consistently outperforming the random baseline across all perturbation types. The results demonstrate that the retrieval mechanism identifies task-relevant gene programs.

Do biological priors enhance functional understanding? To validate the contribution of prior-based edges, we compare PertDCR with 'PertDCR w/o PE' which retains only cell-gene expression edges. Table 3 reveals that this variant exhibits a notable drop in NES Cor from 0.645 to 0.572, suggesting that the integration of biological priors is crucial for capturing functional pathway activities and maintaining the structural integrity of the embedding.

4.4 Analysis of Distance-Constrained Refinement

To validate the interpretability of the distance constrained refinement module, we investigate whether the learned correction magnitudes align with the biological signal strength. Intuitively, genes structurally closer to the perturbation source or involved in key regulatory cascades should exhibit sharper expression shifts, necessitating larger refinement terms. As shown in Figure 8, we stratify the top-20 DEGs by their hop distance. We observe a strong concordance between the ground truth perturbation intensity (indicated by Ranking, red line) and the model's learned correction magnitude (blue/green line). Notably, the model assigns the largest corrections dynamically according to impact severity rather than rigid hop sequence. The genes at the 2-hop distance exhibit the strongest actual activation/repression (lowest rank) and correctly receive the highest correction magnitude from DCRM. The strong negative correlation between ground truth rankings and correction magnitudes ($\rho = -0.968$ for up-regulation) confirms that DCRM selectively amplifies the response signal for genes that undergo the most drastic changes.

4.5 Performance on Cross-Cell-Line Dataset

We evaluate the model's capability to generalize across cell lines using the Mixture dataset. As presented in Table 4, PertDCR consistently outperforms scGPT by significant margins across all metrics. Specifically, it improves prediction precision ($\text{Pearson}_{\delta+\text{de}}$) by about 50% (0.788 vs. 0.527) and accurately captures regulatory trends with a Directionality Accuracy of 90.3% compared to 73.6% for scGPT. Crucially, the superior Pathway NES score (0.786) confirms that PertDCR successfully reconstructs biologically plausible responses. This robustness is further evidenced in Figure 9, where PertDCR surpasses scGPT in 86.5% of individual perturbation tasks, demonstrating exceptional stability in handling complex cross-cell-line scenarios.

Method	Pearson	Pearson_{δ}	$\text{Pearson}_{\delta+\text{de}}$	Dir. Acc	NES Cor
biolord	0.735	0.538	0.544	76.8%	0.587
scGPT	0.808	0.578	0.527	73.6%	0.601
PertDCR	**0.898**	**0.790**	**0.788**	**90.3%**	**0.786**

Table 4: Quantitative comparison on the Mixture dataset.

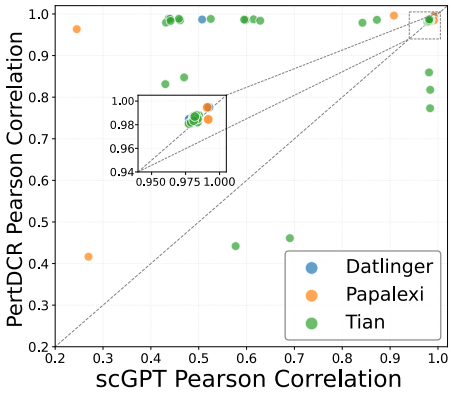

Figure 9: Comparison of the Pearson correlation between PertDCR and scGPT across distinct perturbation types on the Mixture dataset. The points above the diagonal indicate superior performance by PertDCR. PertDCR wins 86.5% of perturbations.

5 Conclusion

In this work, we proposed PertDCR, a framework that models genetic perturbations as topology-constrained molecular stimuli rather than simple statistical shifts. To address the heterogeneity of perturbation effects, PertDCR constructs adaptive perturbation conditions by integrating cell representations with interactome-inferred gene programs, deriving context-aware responses. Furthermore, the distance-constrained refinement module modulates response intensity based on the distance-dependent propagation mechanism within the PPI network. Experimental results demonstrate that PertDCR achieves state-of-the-art performance across diverse unseen scenarios, providing a robust tool for elucidating complex gene regulatory machineries.

Discussion. As high-throughput perturbational screens become integral to drug discovery and regenerative medicine, PertDCR serves as a valuable computational method by facilitating predicted responses adherence to underlying biological mechanisms within the active interactome. A unique strength of PertDCR is its ability to contextualize the perturbation source within the active cellular interactome, allowing for the precise mapping of how signals propagate through topological neighborhoods. This feature is promising for identifying cell-specific therapeutic targets where the same gene perturbation may yield divergent outcomes.

Acknowledgments

This work was supported in part by Scientific Research Innovation Capability Support Project for Young Faculty (Project No. ZYGXQNJSKYCXNLZCXM-H8), and in part by the GuangDong Basic and Applied Basic Research Foundation (Project No. 2024A1515011437).

References

[Adamson *et al.*, 2016] Britt Adamson, Thomas M Norman, Marco Jost, Min Y Cho, James K Nuñez, Yuwen Chen, Jacqueline E Villalta, Luke A Gilbert, Max A Horlbeck, Marco Y Hein, et al. A multiplexed single-cell crispr screening platform enables systematic dissection of the unfolded protein response. *Cell*, 167(7):1867–1882, 2016.

[Bereket and Karaletsos, 2023] Michael Bereket and Theofanis Karaletsos. Modelling cellular perturbations with the sparse additive mechanism shift variational autoencoder. *Advances in Neural Information Processing Systems*, 36:1–12, 2023.

[Chen *et al.*, 2024] Tianyi Chen, Xindian Wei, Lianxin Xie, Yunfei Zhang, Cheng Liu, Wenjun Shen, Si Wu, and Hau-San Wong. Self-former: multi-scale gene filtration transformer for single-cell spatial reconstruction. *Briefings in Bioinformatics*, 25(6):bbae523, 2024.

[Chen *et al.*, 2025] Tianyi Chen, Yunfei Zhang, Lianxin Xie, Wenjun Shen, Si Wu, and Hau-San Wong. Spotdiff: Spatial gene expression imputation diffusion with single-cell rna sequencing data integration. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 15848–15856, 2025.

[Chen *et al.*, 2026] Tianyi Chen, Wen Xue, Yunfei Zhang, Yongcan Luo, Cheng Liu, Wenjun Shen, Si Wu, and Hau-San Wong. S2potae: multimodal spatial spot autoencoder integrating image and transcriptomic features for deconvolution. *Briefings in Bioinformatics*, 27(1):bbag020, 2026.

[Chi *et al.*, 2025a] Changxi Chi, Yufei Huang, Jun Xia, Jiangbin Zheng, Yunfan Liu, Zelin Zang, and Stan Z Li. Departures: Distributional transport for single-cell perturbation prediction with neural schr\" odinger bridges. *arXiv preprint arXiv:2511.13124*, 2025.

[Chi *et al.*, 2025b] Changxi Chi, Jun Xia, Jingbo Zhou, Jiabei Cheng, Chang Yu, and Stan Z. Li. Grape: Heterogeneous graph representation learning for genetic perturbation with coding and non-coding biotype. In James Kwok, editor, *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, pages 4949–4957. International Joint Conferences on Artificial Intelligence Organization, 8 2025. Main Track.

[Costanzo *et al.*, 2016] Michael Costanzo, Benjamin VanderSluis, Elizabeth N Koch, Anastasia Baryshnikova, Carles Pons, Guihong Tan, Wen Wang, Matej Usaj, Julia Hanchard, Susan D Lee, et al. A global genetic interaction network maps a wiring diagram of cellular function. *Science*, 353(6306):aaf1420, 2016.

[Cowen *et al.*, 2017] Lenore Cowen, Trey Ideker, Benjamin J Raphael, and Roded Sharan. Network propagation: a universal amplifier of genetic associations. *Nature Reviews Genetics*, 18(9):551–562, 2017.

[Cui *et al.*, 2024] Haotian Cui, Chloe Wang, Hassaan Maan, Kuan Pang, Fengning Luo, Nan Duan, and Bo Wang. scgpt: toward building a foundation model for single-cell multi-omics using generative ai. *Nature methods*, 21(8):1470–1480, 2024.

[Datlinger *et al.*, 2017] Paul Datlinger, André F Rendeiro, Christian Schmidl, Thomas Krausgruber, Peter Traxler, Johanna Klughammer, Linda C Schuster, Amelie Kuchler, Donat Alpar, and Christoph Bock. Pooled crispr screening with single-cell transcriptome readout. *Nature methods*, 14(3):297–301, 2017.

[Dixit *et al.*, 2016] Atray Dixit, Oren Parnas, Biyu Li, Jenny Chen, Charles P Fulco, Livnat Jerby-Arnon, Nemanja D Marjanovic, Danielle Dionne, Tyler Burks, Raktima Raychowdhury, et al. Perturb-seq: dissecting molecular circuits with scalable single-cell rna profiling of pooled genetic screens. *cell*, 167(7):1853–1866, 2016.

[Hu *et al.*, 2025] Luni Hu, Hua Qin, Yilin Zhang, Yi Lu, Ping Qiu, Zhihan Guo, Lei Cao, Wenjian Jiang, Qianqian Chen, Yanbang Shang, et al. Regformer: A single-cell foundation model powered by gene regulatory hierarchies. *bioRxiv*, pages 2025–01, 2025.

[Ideker and Krogan, 2012] Trey Ideker and Nevan J Krogan. Differential network biology. *Molecular systems biology*, 8(1):565, 2012.

[Istrate *et al.*, 2024] Ana-Maria Istrate, Donghui Li, and Theofanis Karaletsos. scgenept: Is language all you need for modeling single-cell perturbations? *bioRxiv*, pages 2024–10, 2024.

[Kamimoto *et al.*, 2023] Kenji Kamimoto, Blerta Stringa, Christy M Hoffmann, Kunal Jindal, Lilianna Solnica-Krezel, and Samantha A Morris. Dissecting cell identity via network inference and in silico gene perturbation. *Nature*, 614(7949):742–751, 2023.

[Kingma and Ba, 2015] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.

[Kitano, 2002] Hiroaki Kitano. Systems biology: a brief overview. *science*, 295(5560):1662–1664, 2002.

[Kovács *et al.*, 2019] István A Kovács, Katja Luck, Kerstin Spirohn, Yang Wang, Carl Pollis, Sadie Schlabach, Wenting Bian, Dae-Kyum Kim, Nishka Kishore, Tong Hao, et al. Network-based prediction of protein interactions. *Nature communications*, 10(1):1240, 2019.

[Lee *et al.*, 2021] Joo Sang Lee, Nishanth Ulhas Nair, Gal Dinstag, Lesley Chapman, Youngmin Chung, Kun Wang, Sanju Sinha, Hongui Cha, Dasol Kim, Alexander V Schperberg, et al. Synthetic lethality-mediated precision on-

cology via the tumor transcriptome. *Cell*, 184(9):2487–2502, 2021.

[Liang *et al.*, 2025] Zhaokang Liang, Shuyang Zhuang, Xiaoran Jiao, Weian Mao, Hao Chen, and Chunhua Shen. scppdm: A diffusion model for single-cell drug-response prediction. *arXiv preprint arXiv:2510.11726*, 2025.

[Lotfollahi *et al.*, 2019] Mohammad Lotfollahi, F Alexander Wolf, and Fabian J Theis. scgen predicts single-cell perturbation responses. *Nature methods*, 16(8):715–721, 2019.

[Lotfollahi *et al.*, 2023] Mohammad Lotfollahi, Anna Klimovskaia Susmelj, Carlo De Donno, Leon Hetzel, Yuge Ji, Ignacio L Ibarra, Sanjay R Srivatsan, Mohsen Naghipourfar, Riza M Daza, Beth Martin, et al. Predicting cellular responses to complex perturbations in high-throughput screens. *Molecular systems biology*, 19(6):e11517, 2023.

[Lu *et al.*, 2024] L Lu, Wen Xue, Xindian Wei, et al. Sctrans: multi-scale scrna-seq sub-vector completion transformer for gene-selective cell type annotation. In *33rd International Joint Conference on Artificial Intelligence (IJCAI 2024)*, pages 5954–62. International Joint Conferences on Artificial Intelligence, 2024.

[Luck *et al.*, 2020] Katja Luck, Dae-Kyum Kim, Luke Lambourne, Kerstin Spirohn, Bridget E Begg, Wenting Bian, Ruth Brignall, Tiziana Cafarelli, Francisco J Campos-Laborie, Benoit Charloteaux, et al. A reference map of the human binary protein interactome. *Nature*, 580(7803):402–408, 2020.

[Norman *et al.*, 2019] Thomas M Norman, Max A Horlbeck, Joseph M Replogle, Alex Y Ge, Albert Xu, Marco Jost, Luke A Gilbert, and Jonathan S Weissman. Exploring genetic interaction manifolds constructed from rich single-cell phenotypes. *Science*, 365(6455):786–793, 2019.

[Papalexi *et al.*, 2021] Efthymia Papalexi, Eleni P Mimitou, Andrew W Butler, Samantha Foster, Bernadette Bracken, William M Mauck III, Hans-Hermann Wessels, Yuhan Hao, Bertrand Z Yeung, Peter Smibert, et al. Characterizing the molecular regulation of inhibitory immune checkpoints with multimodal single-cell screens. *Nature genetics*, 53(3):322–331, 2021.

[Peidli *et al.*, 2022] Stefan Peidli, Tessa D. Green, Ciyue Shen, Torsten Gross, Joseph Min, Jake P. Taylor-King, Debora S. Marks, Augustin Luna, Nils Blüthgen, and Chris Sander. scperturb single-cell perturbation data: Rna and protein h5ad files, August 2022.

[Piran *et al.*, 2024] Zoe Piran, Niv Cohen, Yedid Hoshen, and Mor Nitzan. Disentanglement of single-cell data with biolord. *Nature Biotechnology*, 42(11):1678–1683, 2024.

[Replogle *et al.*, 2022] Joseph M Replogle, Reuben A Saunders, Angela N Pogson, Jeffrey A Hussmann, Alexander Lenail, Alina Guna, Lauren Mascibroda, Eric J Wagner, Karen Adelman, Gila Lithwick-Yanai, et al. Mapping information-rich genotype-phenotype landscapes with genome-scale perturb-seq. *Cell*, 185(14):2559–2575, 2022.

[Roohani *et al.*, 2024] Yusuf Roohani, Kexin Huang, and Jure Leskovec. Predicting transcriptional outcomes of novel multigene perturbations with gears. *Nature Biotechnology*, 42(6):927–935, 2024.

[Sachs *et al.*, 2005] Karen Sachs, Omar Perez, Dana Pe'er, Douglas A Lauffenburger, and Garry P Nolan. Causal protein-signaling networks derived from multiparameter single-cell data. *Science*, 308(5721):523–529, 2005.

[Szklarczyk *et al.*, 2022] Damian Szklarczyk, Rebecca Kirsch, Mikaela Koutrouli, Katerina Nastou, Farrokh Mehryary, Radja Hachilif, Annika L Gable, Tao Fang, Nadezhda T Doncheva, Sampo Pyysalo, Peer Bork, Lars J Jensen, and Christian von Mering. The string database in 2023: protein–protein association networks and functional enrichment analyses for any sequenced genome of interest. *Nucleic Acids Research*, 51(D1):D638–D646, 11 2022.

[Theodoris *et al.*, 2023] Christina V Theodoris, Ling Xiao, Anant Chopra, Mark D Chaffin, Zeina R Al Sayed, Matthew C Hill, Helene Mantineo, Elizabeth M Brydon, Zexian Zeng, X Shirley Liu, et al. Transfer learning enables predictions in network biology. *Nature*, 618(7965):616–624, 2023.

[Tian *et al.*, 2021] Ruilin Tian, Anthony Abarientos, Jason Hong, Sayed Hadi Hashemi, Rui Yan, Nina Dräger, Kun Leng, Mike A. Nalls, Andrew B. Singleton, Ke Xu, Faraz Faghri, and Martin Kampmann. Genome-wide crispri/a screens in human neurons link lysosomal failure to ferroptosis. *Nature Neuroscience*, 24(7):1020–1034, Jul 2021.

[Tilford and Siemers, 2009] Charles A Tilford and Nathan O Siemers. Gene set enrichment analysis. In *Protein networks and pathway analysis*, pages 99–121. Springer, 2009.

[Vidal *et al.*, 2011] Marc Vidal, Michael E Cusick, and Albert-László Barabási. Interactome networks and human disease. *Cell*, 144(6):986–998, 2011.

[Wei *et al.*, 2025] Xindian Wei, Tianyi Chen, Xibiao Wang, Wenjun Shen, Cheng Liu, Si Wu, and Hau-San Wong. Come: contrastive mapping learning for spatial reconstruction of single-cell rna sequencing data. *Bioinformatics*, 41(3):btaf083, 2025.

[Yu *et al.*, 2025] Hengshi Yu, Weizhou Qian, Yuxuan Song, and Joshua D Welch. Perturbnet predicts single-cell responses to unseen chemical and genetic perturbations. *Molecular Systems Biology*, 21(8):960, 2025.

[Zhou *et al.*, 2025] Xianyong Zhou, Xindian Wei, Cheng Liu, Wenjun Shen, Ping Xuan, Si Wu, and Hau-San Wong. Robust subspace structure discovery for cell type identification in scrna-seq data. *BMC bioinformatics*, 26(1):295, 2025.