Frontiers in Computing and Intelligent Systems
ISSN: 2832-6024 | Vol. 12, No. 3, 2025

Predicting Multigene Transcriptional Outcomes Based on Dual-Stream Mamba Neural Network

Hao Li [1, 2], Liyu Liu [1, 2, *]

[1] Postgraduate Training Base Alliance of Wenzhou Medical University, Wenzhou, Zhejiang 325000, China
[2] Wenzhou Key Laboratory of Biophysics, Wenzhou Institute, University of Chinese Academy of Sciences, Wenzhou, Zhejiang 325001, China

*** Corresponding author:** Liyu Liu

Abstract: The integration of gene perturbation technology with single-cell RNA sequencing and CRISPR editing has promoted the standardization of perturbation effect studies at the single-cell level. How- ever, the diversity of human cell types and computational complexity of perturbation combinations make exhaustive experimental exploration infeasible. In this context, transcriptional prediction models for multi-gene perturbations have undergone a paradigm shift from traditional statistical modeling to deep learning-driven approaches. Nevertheless, challenges persist, including difficulties in adapting to genome-scale applications and limitations to single-gene perturbation analysis. To address these issues, this study proposes a Mamba-based deep learning algorithm that constructs a dual-stream integration framework for multi-gene perturbation response prediction. The model establishes dual-stream architecture to capture both statistical and biological features of gene ex- pression data, circumventing the limitations of single data- or knowledge-driven approaches. The core Mamba component leverages its linear computational complexity and selective memory mechanism to achieve precise modeling of gene-gene interactions while maintaining computational efficiency. Its input-dependent parameterization strategy enables effective mining of complex data features. Analytical results demonstrate that the proposed model exhibits superior performance in gene perturbation prediction while providing biological interpretability, establishing it as an effective methodology.

Keywords: Genetic Perturbation Prediction; Genomics; Mamba Encoder; Variational Autoencoder.

1. Introduction

Despite centuries of biological exploration, fundamental challenges persist in understanding the operational mechanisms of cells—the fundamental units of life—within the human body. While modern biotechnology has enabled the static mapping of molecular compositions (e.g., genome, transcriptome, proteome, metabolome) and their spatial distributions, the inherent complexity of cells remains difficult to resolve. This complexity primarily arises from the highly dynamic interactions between molecules that drive cellular functional differentiation [1]. Research findings derived from genetic perturbation techniques, a core methodology for dissecting genetic interaction networks and their impact on cellular function, are progressively forming a crucial scientific foundation for modern biotechnological innovation and disease treatment theories [2].

Recent years have witnessed the breakthrough development of single-cell RNA sequencing (scRNA-seq) technology [3]. Its integration with high-throughput genetic perturbation screening [4] and CRISPR-mediated gene editing [5] has driven technological innovation and established standardized research paradigms within the field [6]. This technological convergence enables researchers to systematically analyze the downstream regulatory effects of genetic perturbations on cell states at single-cell resolution [7]. However, the vast diversity of human cell types and the exponential growth in computational complexity associated with potential perturbation combinations render exhaustive experimental characterization of cellular responses under all perturbation conditions temporally and resource-wise infeasible. Consequently, the research paradigm for predicting transcriptomic effects of multi-gene perturbations has undergone a fundamental shift: evolving from early traditional statistical models to deep learning frameworks integrating biological prior knowledge, achieving breakthroughs in key dimensions such as interpretability and computational efficiency. Crucially, single-cell multi-omics data provides an ideal substrate for deep learning techniques, facilitating efficient modeling in core tasks including cell segmentation [8, 9], subpopulation clustering [10], and drug response prediction [11].

During the early generative model and linear propagation phase, research primarily relied on generative models to capture gene expression distribution characteristics. In 2019, Lotfollahi et al. proposed scGen, the first to introduce Variational Autoencoders (VAEs) for single-cell perturbation prediction. By interpolating within a latent space to estimate gene expression changes, it achieved a Pearson correlation coefficient of 0.78 in cross-cell-type prediction [12]. However, this method was limited to single-gene perturbations and lacked mechanistic interpretability. Concurrently, Aibar et al. proposed a Gene Regulatory Network (GRN)-based method propagating perturbation effects linearly using gene co-expression matrices to construct regulatory relationships [13]. Yet, its prediction error increased by 42% for multi-gene combinations, and the linear assumption struggled to capture non-linear genetic interactions.

The exponential increase in computational complexity with the number of genes further spurred the emergence of knowledge-enhanced models. Post-2020, researchers began integrating biological pathway knowledge to enhance model generalizability. The GEARS model, proposed by Roohani et al., utilized a gene co-expression graph to encode transcriptional synergy, constructing a dual knowledge graph. It employed Gene Ontology (GO) pathway similarity graphs

to guide perturbation propagation [14] and aggregated neighborhood information via Graph Neural Networks (GNNs). In 131 dual-gene perturbation experiments, GEARS reduced the Mean Squared Error (MSE) by 37% and improved prediction accuracy for unseen gene combinations by 53%. Around the same time, Kamimoto et al. introduced the CellOracle model, which combined GRNs with causal inference to infer regulatory paths by backpropagating perturbation response gradients, albeit restricted to single-gene analysis [15]. Its strengths included incorporating pathway databases like Reactome to construct attention masks, developing a self-focusing loss function to enhance weights for differentially expressed genes, and using a direction-aware loss to correct expression trend biases. However, the message-passing mechanism of GNNs proved inefficient for hierarchical pathway structures and incapable of modeling long-range gene interactions.

Most of these methods, including scGen, PerturbNet [16, 17], and CPA [18], primarily predict transcriptional responses to chemical treatments rather than genetic perturbations. They often treat genetic perturbations as analogs of different doses, which inadequately captures the complex interactions between different genes and specific gene knockouts, resulting in limited performance for genetic perturbation prediction specifically.

The introduction of the Transformer architecture significantly enhanced the modeling capacity for cross-gene relationships, as its self-attention mechanism can capture long-range dependencies. In 2024, Bai et al. proposed AttentionPert, integrating attention mechanisms and GNNs into perturbation analysis, overcoming the limitations of traditional linear models or single-scale approaches [19]. Their proposed multi-scale modeling framework is applicable not only to genetic perturbation studies but can also be transferred to scenarios like drug combination prediction and environmental stimulus response, providing a general methodology for systems biology. That same year, Tan et al. introduced the dual-stream network BioDSNN [20]. It employs a Variational Autoencoder to construct a data-driven stream for reconstructing global gene expression features, while an attention mechanism processes Reactome pathways to form a stream governed by biological prior knowledge. This dual-stream design of BioDSNN avoids the limitations of relying solely on data or knowledge, leveraging both statistical patterns and biological logic to enhance model robustness. Compared to the previous phase, this stage saw substantial improvements across all gene prediction metrics. Models not only utilized hybrid architectures to integrate data-driven and knowledge-driven learning but also incorporated diffusion models to improve generation quality. Furthermore, uncertainty quantification metrics (e.g., Bayesian confidence in GEARS) were developed to enhance model robustness. However, the quadratic computational complexity of Transformers and their limitations in processing long sequences restricted their application at whole-genome scale [21].

The introduction of the Mamba state space model (SSM) brought innovation to long-sequence processing. Its selective memory mechanism, employing an input-dependent parameterization strategy, achieves 94% sequence information capture while maintaining linear computational complexity, enabling the extraction of more complex data features [22]. In this work, we propose a novel Mamba-based dual-stream framework model for predicting multi-gene perturbation responses, aiming to demonstrate its unique advantages for this task.

2. Model Algorithm

The dual-stream Mamba neural network constructed in this paper can be briefly divided into three layers, as shown in Figure 1, namely the input layer, the encoding layer and the output layer. Next, we will provide a detailed introduction to the three layers of this model.

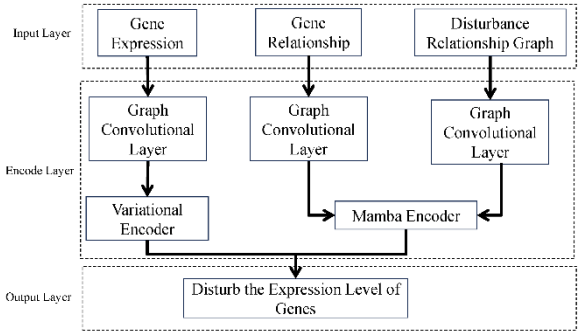

Figure 1. Diagram of the dual-stream mamba neural network model

(1) Input Layer

The input layer mainly consists of three parts of data, namely gene expression data, gene relationship map data, and perturbation relationship map data. These three parts of data constitute the input of the entire model. The data of gene expression levels and gene perturbation were derived from the Norman, RPE1, and K562 datasets. In order to predict the transcription results caused by perturbations, this paper selects these three publicly available perturbation sets. These three datasets are composed of a single gene or multiple genes. For single genes, this paper uses data from two different genetic interference screenings, consisting of 1543 (RPE1 cells) and 1092 (K562 cells) interferons respectively, with each interference measuring more than 170,000 cells [22]. Furthermore, the Norman dataset [23] contains 131 two-gene perturbations, so we can evaluate the performance at different perturbation levels. The summary of the dataset is shown in Table 1.

Table 1. Dataset information

DataSet	Cell Quantity	Gene number	Number of gene perturbations
K562	162,751	5000	1092
RPE1	162,733	5000	1543
Norman	91,205	5045	131

The gene relationship map data comes from the processing results of GO data in Gears. Because GO data is used to measure the functional similarity between genes, this paper is used to measure the functional relationship between genes and explore more perturbation information. For example, Nu is the Go pathway set of the gene u, and the Jaccard index between a pair of genes u and v is calculated as:

$$J_{(u,v)} = \left| \frac{N_u \cap N_v}{N_u \cup N_v} \right| \tag{1}$$

The Jaccard index is used to measure the proportion of shared paths among them. For each gene, this paper selects the gene with the highest Jaccard index to construct the

perturbation similarity map. Then, all the initialized possible gene perturbation embeddings will be fed into the graph convolutional network for expansion. In this way, when each disturbance is embedded, the information of adjacent disturbances can be integrated.

(2) Coding Layer

The encoding layer consists of two data streams, namely: the data stream of gene expression level data passing through the graph convolutional layer and the variational autoencoder, and the data stream of gene relationship graph data and perturbation relationship graph data passing through the graph convolutional layer and the Mamba encoding layer. Next, this article will elaborate in detail on the composition of each data stream. Both data streams selected graph convolutional networks as the data embedding layer. When capturing the characteristics of gene expression level data, variational encoders were used. Variational autoencoders have multiple advantages as generative networks: they can learn the latent distribution of data and generate diverse and continuously changing new samples; High-dimensional data can be compressed into a low-dimensional latent space to extract semantic features. The structure is flexible and easy to combine with other technologies; As a probabilistic model, it can not only explain the data generation process but also estimate the uncertainty of the generation results. Therefore, in this paper, variational encoders are adopted to capture the uncertainty and variability of gene expression data. Variational autoencoders can measure the deviation between the learning distribution of potential variables (given the original gene expression data) and the prior distribution (usually Gaussian prior distribution). Specifically, the variational autoencoder maximizes the lower bound of evidence, known as the evidence lower bound (ELBO), because the marginal likelihood is difficult to handle. The variational lower bound of the marginal likelihood of the input data is the objective function of the variational autoencoder. The marginal likelihood is obtained by summing the marginal likelihood of different data points, as shown below:

$$\log\left(p\left(x^{(1)},\cdots,x^{(N)}\right)\right)=\sum_{i=1}^{N}\log\left(p_\theta\left(x^i\right)\right) \tag{2}$$

In this paper, the Mamba encoder is adopted to capture the associations between the gene relationship map data and the perturbation map data. As a new and emerging architecture, Mamba focuses more on the correlations within the data, enabling the model to focus more on key information. Meanwhile, compared with the attention mechanism, Mamba has greater throughput and faster processing speed, and has better universality and adaptability. The Standard State Space Model (SSM) usually assumes that parameters remain constant, that is, it is believed that the state transition matrix, etc. do not change during sequence processing. The Mamba encoder introduces time-varying parameters, allowing the parameters to change along with the iteration step size. This design, while maintaining the degree of linear complexity in computation, brings greater flexibility and expressiveness to sequence modeling, which is one of the key innovations that distinguish the Mamba encoder from other encoders. The generation formula of time-varying parameters is shown in (3).

$$\begin{aligned}\Delta_t &= \tau_\Delta(W_\Delta x_t)\\ \mathrm{B}_t &= W_B x_t\\ C_t &= W_c x\end{aligned} \tag{3}$$

Then, the Mamba model discretizes the Continuous State Space model (SSM) through the Zero-Order Hold (ZOH) method. The zero-order preservation method discretizes the parameter generation that supports input perception, realizes the selective mechanism, and simultaneously achieves linear complexity and full parallel computing, successfully breaking through the bottleneck of long sequences. It is specifically shown as Formula (4).

$$\overline{A}_t = e^{\Delta_t A}, \overline{B}_t = (\Delta_t A)^{-1}(e^{\Delta_t A}-1)\Delta_t B_t \tag{4}$$

Then, Parallel state updates are carried out (as shown in Formula (5)), and the parallel Scan algorithm (Parallel Scan) is used to avoid recursive computations. The parallel scanning algorithm has disruptive value. While improving the training speed, it significantly saves memory. It is the only parallel scheme that supports dynamic parametric state transitions and provides a computational basis for the selectivity mechanism of Mamba.

$$h_t = \overline{A}_t h_{t-1} + \overline{B}_t x_t \tag{5}$$

$$y_t = C_t h_t \tag{6}$$

In this study, to determine the optimal value of the core parameter of the Mamba model - the state dimension (d state), three groups of values of 4, 16, and 128 were selected, and experiments were conducted on the REP1 and K652 datasets. The experimental results are shown in Figure 2: When the value of the state dimension is 16, the model performance reaches the optimum.

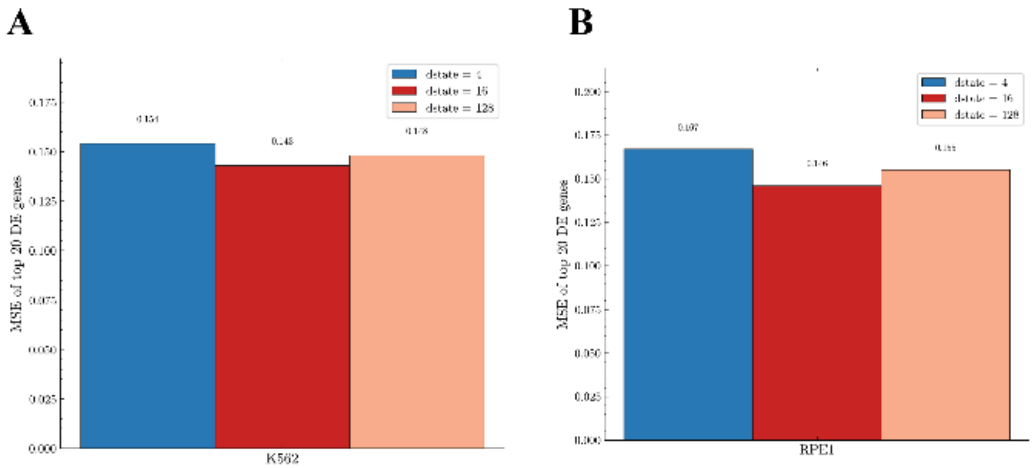

Figure 2. Comparison Chart of Results from Different State Dimension Settings. (A) K562 dataset (B) RPE1 dataset

(3) Parameter setting

During the model training process, the learning rate scheduling adopts the attenuation strategy based on the gamma coefficient (gamma=0.5). When the preset attenuation conditions (such as the stagnation of the validation set performance) are met, the learning rate will be adjusted by multiplying by the gamma factor: the initial learning rate is 1e-3, it drops to 5e-4 after the first attenuation, and further drops to 2.5e-4 after the second attenuation. This progressive attenuation mechanism enables the model to make more refined parameter adjustments when approaching the optimal solution, effectively enhancing the convergence stability and generalization ability. The main training parameters are configured as follows: The learning rate is 1e-3, the training rounds are 20, and the loss function adopts the combined form of KL divergence and mean square error to accurately quantify the difference between the prediction results and the real labels and guide the optimization direction. Aiming at the problem of intense gradient fluctuations in the early stage of training, a higher momentum parameter beta1=0.9 is set, which not only supposes the risk of parameter update oscillation but also ensures precise convergence in the later stage. To prevent overfitting, L2 regularization is introduced and the weight attenuation coefficient 5e-4 is set. By constraining the model complexity, overlearning of the training noise is avoided. As shown in Figure 3, the training process is stable and eventually converges.

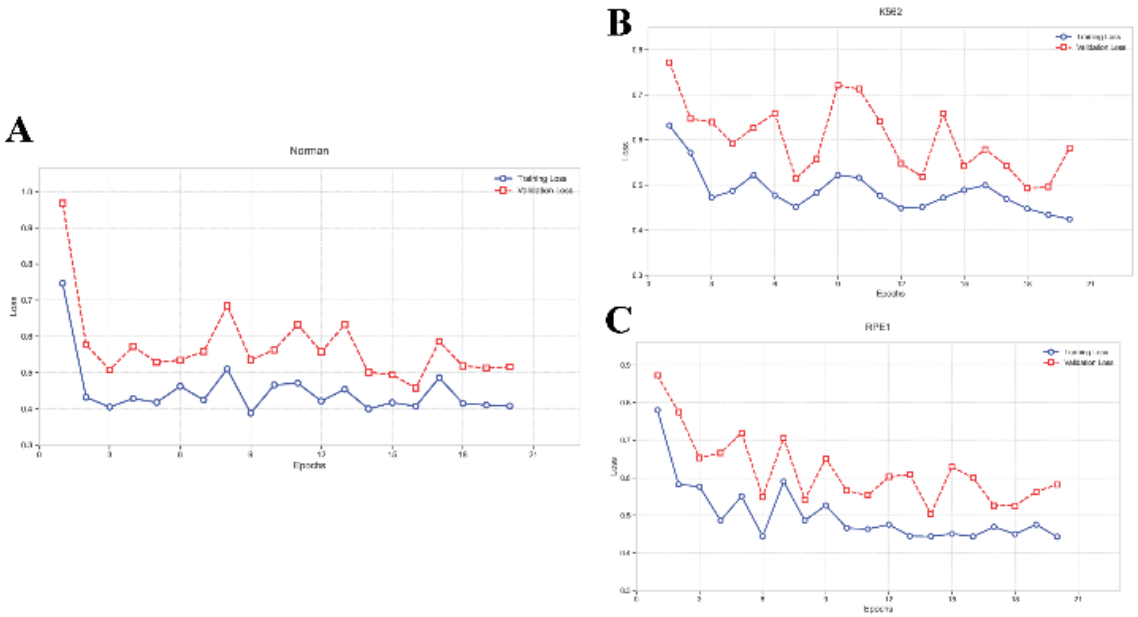

Figure 3. Training loss curves for different datasets. (A) Noman dataset (B) K562 dataset (C) RPE1 dataset

3. Analysis Result

As illustrated in Figure 4, to assess differences in single-cell gene expression before and after perturbation, we calculated the Mean Squared Error (MSE) for each gene between the control and experimental groups. Given that the vast majority of genes exhibit no substantial changes between unperturbed and perturbed states, only the top 20 differentially expressed genes (DEGs) were considered. In the figure, the groups labeled "0/2 unseen," "1/2 unseen," and "2/2 unseen" correspond to experimental conditions where zero, one, or both perturbation targets were observed in the training set, respectively. This grouping setting precisely simulates the challenge of data sparsity inherent in gene perturbation prediction tasks: fewer perturbation targets observed during training correspond to larger information gaps the model must fill, with prediction difficulty increasing exponentially. The experimental results in Figure 4 demonstrate that the proposed model exhibits significant advantages across all data groups, with prediction performance substantially surpassing all baseline models. Even under the extremely challenging "0/2 unseen" scenario, the model's prediction error remains lower than traditional methods, robustly validating its superior generalization capability for complex multi-gene perturbation prediction tasks.

The model's exceptional performance stems from two core innovative designs. Firstly, the dual-stream network architecture overcomes the limitations of single data input modalities. By constructing parallel information channels, it achieves deep integration of gene expression data with domain knowledge: one pathway focuses on extracting statistical features from vast transcriptomic data, while the other encodes structured biological knowledge—such as gene regulatory networks (GRNs) and functional annotations—into learnable embedding vectors. These streams dynamically interact via an attention mechanism, ensuring the model simultaneously incorporates data-driven pattern recognition and biologically logical constraints. Secondly, the novel Mamba encoder, leveraging its unique linear-complexity architecture and dynamic attention mechanism, effectively overcomes the efficiency bottleneck of traditional Transformers in long-sequence modeling. This encoder not only captures long-range dependencies within gene expression data but also adaptively adjusts feature weights according to different perturbation contexts. Consequently, it precisely identifies key regulatory nodes within complex gene regulatory networks, significantly enhancing the model's prediction accuracy and robustness.

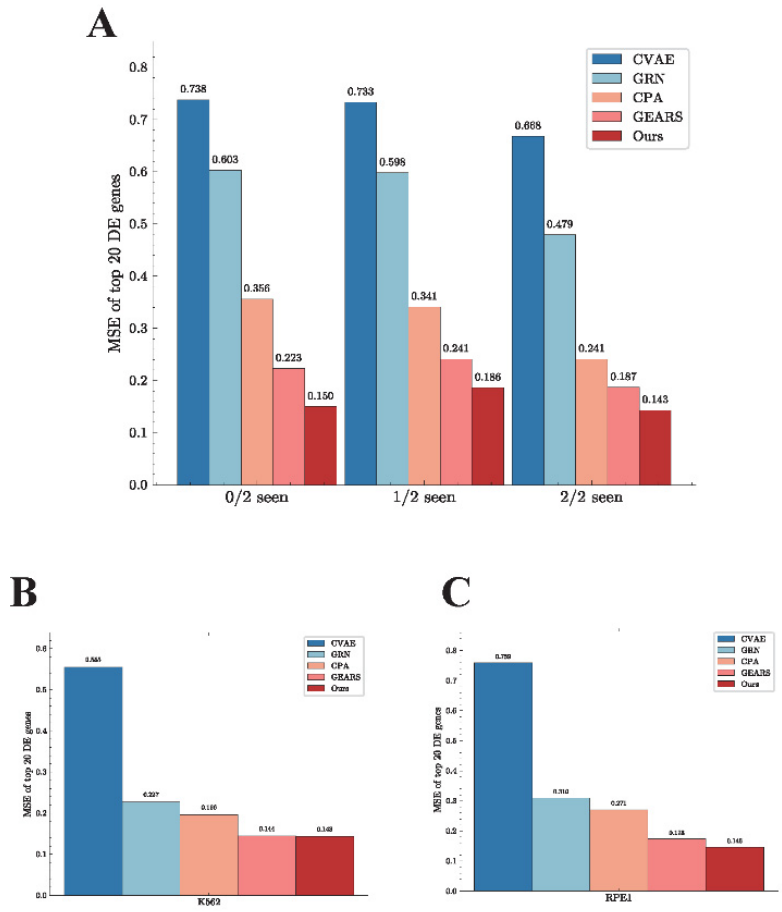

Figure 4. Distribution of Mean Squared Error (MSE) in Gene Perturbation Response Prediction. (A) Noman dataset (B) K562 dataset (C) RPE1 dataset

In this paper, an uncertainty score is adopted to measure the confidence level of the model prediction in new disturbances. The Gaussian likelihood function is adopted to model the gene expression values under perturbation. Meanwhile, an additional gene-specific layer is added to predict the logarithmic variance term (Formula (8)), and the uncertainty score of the model is modified through the modified Bayesian neural network loss (Formula (9)) [24].

$$S_u = \log \sigma_u^2 = w_u^{unc} h^{post-pert} + b_u^{unc} \tag{7}$$

$$L_{unc} = \frac{1}{T}\sum_{k=1}^{T}\frac{1}{T}\sum_{l=1}^{T_k}\frac{1}{G}\sum_{u=1}^{k}\exp(-s_u)(g_u - \hat{g}_u)^{(2+y)} \tag{8}$$

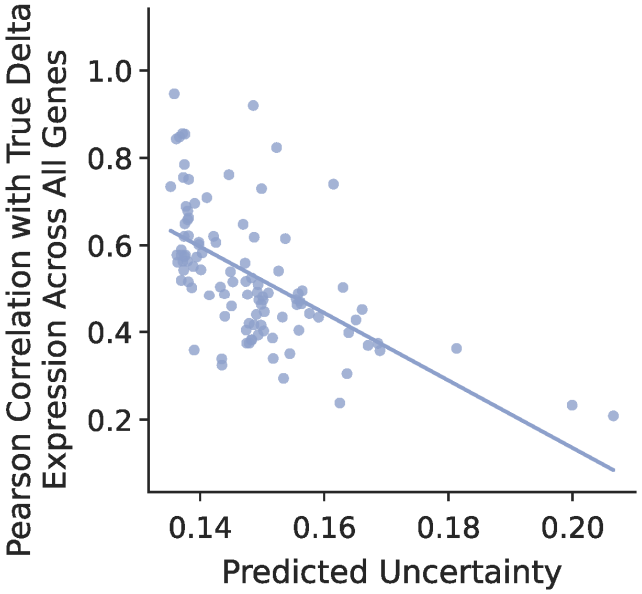

Figure 5. Distribution of Model Uncertainty Scores

The distribution of uncertainty scores for the model proposed in this paper is shown in Figure 5. Analyzing this figure shows that the model exhibits a low level of prediction uncertainty and a significant correlation between the predicted gene expression differences and the actual values. This result indicates that the model can not only reliably predict the expression changes after gene perturbation, but also effectively capture the complex dynamic features of gene interactions, providing an analytical tool with high confidence for the precise resolution of multi-gene perturbation effects.

In the field of disease research, dynamic changes in gene expression patterns contain rich information, which can not only provide valuable potential biomarkers for disease diagnosis and prognosis assessment, but also can be used to identify key targets for therapeutic interventions and drug discovery by screening for genes with significantly altered expression levels. In order to explore the response mechanism of individual genes under different perturbation scenarios, this study analyzed the prediction of perturbation results for three typical gene combinations, which correspond to different levels of visibility in the training set: double-visible gene combinations (both genes present in the training process), single-visible gene combinations (only one gene present in the training data), and zero-visible gene combinations (both genes absent in the training data). training data exist). Taking specific experiments as an example, the ETS2 and CEBPE gene combination belongs to the double-visible type, and the complete perturbation data of both genes can be obtained during the training period; the FOSB and CEBPB combination is of the single-visible type, in which CEBPB did not experience experimental perturbation during the training phase; and the FOXL2 and MEIS1 combination is in zero-visible status, and the data related to both genes are completely lacking during the training process [25][26].

As shown in Figure 6, the dual-stream model proposed in this study demonstrated outstanding performance in all three complex perturbation scenarios, achieving accurate and stable prediction of the expression level of genes after perturbation. This model not only performs excellently in the analysis of multi-gene interactions, but also the high correlation between its prediction results and the real data fully confirms the robustness and reliability of the model when capturing the complex mechanism of gene disturbances.

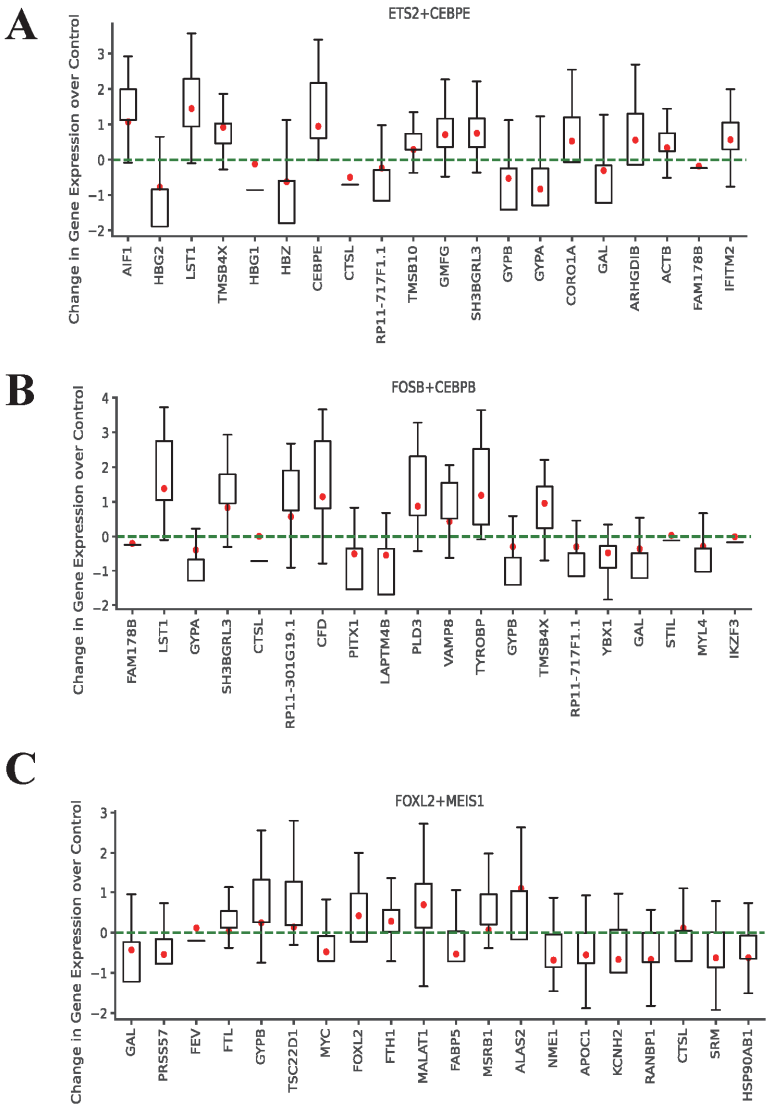

Figure 6. Distribution of gene expression levels under dual-gene perturbation. (A) ETS2 and CEBPE (B) POSB and CEBPB (C) FOXL2 and MEIS1

4. Conclusion

This study innovatively constructed a two-stream deep learning framework based on the Mamba state space model, achieving precise prediction of multi-gene perturbation responses. This architecture analyzes the statistical patterns of gene expression profiles through data-driven streams. Meanwhile, it integrates biological prior knowledge such as the Reactome pathway database and gene ontology (GO) functional annotations with the aid of knowledge-driven streams. Moreover, it adopts a dynamic cross-modal attention mechanism to achieve adaptive fusion of heterogeneous information, effectively overcoming the limitations of modeling based on a single data source or knowledge-driven approaches. The core Mamba module accurately models the long-range interaction relationship between genes while ensuring the operational efficiency through its linear computational complexity and selective memory mechanism. Its input-dependent parameterization strategy significantly enhances the extraction ability of complex biological characteristics. Although the dual-stream Mamba framework demonstrates superior performance in the prediction of multi-gene perturbations, it is still limited by the inherent complexity of biological systems when dealing with ultra-high-dimensional combined perturbations. Future improvements should focus on three strategic directions: (1) Systematically integrate multi-omics features such as chromatin accessibility and protein interaction networks to construct a panoramic gene regulation map; (2) Develop dynamic priority learning mechanisms for key biological pathways; (3) Establish an uncertainty quantification framework to enhance the interpretability of the model in clinical transformation scenarios. These technological breakthroughs will drive the next generation of algorithms to solve more complex gene regulation problems in precision medicine.

References

[1] Jaitin, D.A., Weiner, A., Yofe, I., Lara-Astiaso, D., Keren-Shaul, H., David, E., et al. (2016) Dissecting Immune Circuits by Linking Crispr-Pooled Screens with Single-Cell RNA-Seq. Cell, 167, 1883-1896.e15. [1]Jaitin, D.A., Weiner, A., Yofe, I., Lara-Astiaso, D., Keren-Shaul, H., David, E., et al. (2016) Dissecting Immune Circuits by Linking Crispr-Pooled Screens with Single-Cell RNA-Seq. Cell, 167, 1883-1896.e15.

[2] Katti, A., Diaz, B.J., Caragine, C.M., Sanjana, N.E. and Dow, L.E. (2022) CRISPR in Cancer Biology and Therapy.Nature Reviews Cancer, 22, 259-279.

[3] Adamson, B., Norman, T.M., Jost, M., Cho, M.Y., Nuñez, J.K., Chen, Y., et al. (2016) A Multiplexed Single-Cell CRISPR Screening Platform Enables Systematic Dissection of the Unfolded Protein Response. Cell, 167, 1867-1882.e21.

[4] Hanna, R.E. and Doench, J.G. (2020) Design and Analysis of CRISPR-Cas Experiments. Nature Biotechnology, 38, 813-823.

[5] Nakamura, M., Gao, Y., Dominguez, A.A. and Qi, L.S. (2021) CRISPR Technologies for Precise Epigenome Editing.Nature Cell Biology, 23, 11-22.

[6] Frangieh, C.J., Melms, J.C., Thakore, P.I., Geiger-Schuller, K.R., Ho, P., Luoma, A.M., et al. (2021) Multimodal Pooled Perturb-Cite-Seq Screens in Patient Models Define Mechanisms of Cancer Immune Evasion. Nature Genetics, 53, 332-341.

[7] Przybyla, L. and Gilbert, L.A. (2021) A New Era in Functional Genomics Screens. Nature Reviews Genetics, 23, 89-103.

[8] Stringer, C., Wang, T., Michaelos, M. and Pachitariu, M. (2020) Cellpose: A Generalist Algorithm for Cellular Segmen- tation. Nature Methods, 18, 100-106.

[9] Littman, R., Hemminger, Z., Foreman, R., Arneson, D., Zhang, G., Gómez-Pinilla, F., et al. (2021) Joint Cell Segmen- tation and Cell Type Annotation for Spatial Transcriptomics. Molecular Systems Biology, 17, e10108.

[10] Zeng, Y., Zhou, X., Rao, J., Lu, Y. and Yang, Y. (2020) Accurately Clustering Single-Cell RNA-Seq Data by Capturing Structural Relations between Cells through Graph Convolutional Network. 2020 IEEE International Conference on Bi- oinformatics and Biomedicine (BIBM), Seoul, 16-19 December 2020, 519-522.

[11] Fan, Z., Zhao, H., Zhou, J., Li, D., Fan, Y., Bi, Y., et al. (2024) A Versatile Attention-Based Neural Network for Chemical Perturbation Analysis and Its Potential to Aid Surgical Treatment: An Experimental Study. International Journal of Surgery, 110, 7671-7686.

[12] Lotfollahi, M., Wolf, F.A. and Theis, F.J. (2019) scGen Predicts Single-Cell Perturbation Responses. Nature Methods, 16, 715-721.

[13] Aibar, S., González-Blas, C.B., Moerman, T., Huynh-Thu, V.A., Imrichova, H., Hulselmans, G., et al. (2017) SCENIC: Single-Cell Regulatory Network Inference and Clustering. Nature Methods, 14, 1083-1086.

[14] Roohani, Y., Huang, K. and Leskovec, J. (2023) Predicting Transcriptional Outcomes of Novel Multigene Perturbations with Gears. Nature Biotechnology, 42, 927-935.

[15] Kamimoto, K., Stringa, B., Hoffmann, C.M. et al. (2023) Dissecting Cell Identity via Network Inference and in Silico Gene Perturbation. Nature, 614, 742-751.

[16] Yu, H. and Welch, J.D. (2024) Perturbnet Predicts Single-Cell Responses to Unseen Chemical and Genetic Perturbations.

[17] Lotfollahi, M., Klimovskaia Susmelj, A., De Donno, C., Hetzel, L., Ji, Y., Ibarra, I.L., et al. (2023) Predicting Cellular Responses to Complex Perturbations in High-Throughput Screens. Molecular Systems Biology, 19, e11517.

[18] Bai, D., Ellington, C.N., Mo, S., Song, L. and Xing, E.P. (2024) AttentionPert: Accurately Modeling Multiplexed Genetic Perturbations with Multi-Scale Effects. Bioinformatics, 40, i453-i461.

[19] Tan, Y., Xie, L., Yang, H., Zhang, Q., Luo, J. and Zhang, Y. (2024) BioDSNN: A Dual-Stream Neural Network with Hybrid Biological Knowledge Integration for Multi-Gene Perturbation Response Prediction. Briefings in Bioinformatics, 26, bbae617.

[20] Vaswani, A., Shazeer, N., Parmar, N. et al. (2017) Attention Is All You Need. arXiv: 1706.03762.

[21] Gu, A. and Dao, T. (2023) Mamba: Linear-Time Sequence Modeling with Selective State Spaces. arXiv: 2312.00752.

[22] Replogle, J.M., Saunders, R.A., Pogson, A.N., Hussmann, J.A., Lenail, A., Guna, A., et al. (2022) Mapping Information- Rich Genotype-Phenotype Landscapes with Genome-Scale Perturb-Seq. Cell, 185, 2559-2575.e28.

[23] Norman, T.M., Horlbeck, M.A., Replogle, J.M., et al. (2019) Exploring Genetic Inter-Action Manifolds Constructed from Rich Single-Cell Phenotypes. Science, 365, 786-793.

[24] Kendall, A. and Gal, Y. (2017) What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? Proceeding of 31st International Conference on Neural

Information Processing Systems, Long Beach, 4-9 December 2017, 5580-5590.

[25] Replogle, J.M., Norman, T.M., Xu, A., Hussmann, J.A., Chen, J., Cogan, J.Z., et al. (2020) Combinatorial Single-Cell CRISPR Screens by Direct Guide RNA Capture and Targeted Sequencing. Nature Biotechnology, 38, 954-961.

[26] Bock, C., Datlinger, P., Chardon, F., Coelho, M.A., Dong, M.B., Lawson, K.A., et al. (2022) High-Content CRISPR Screening. Nature Reviews Methods Primers, 2, Article No. 8.