GREmLN: A Cellular Graph Structure Aware Transcriptomics Foundation Model

Anonymous authors
Paper under double-blind review

Abstract

The ever-increasing availability of large-scale single-cell profiles presents an opportunity to develop foundation models to capture cell properties and behavior. However, standard language models such as transformers benefit from sequentially structured data with well defined absolute or relative positional relationships, while single cell RNA data have orderless gene features. Molecular-interaction graphs, such as protein-protein interaction (PPI) networks or gene regulatory networks (GRN), offer graph structure-based models that effectively encode both non-local gene token dependencies, as well as potential causal relationships. We introduce GREmLN (**G**ene **R**egulatory **Em**bedding-based **L**arge **N**eural model), a foundation model that leverages graph signal processing to embed gene token graph structure directly within its attention mechanism, producing biologically informed single cell specific gene embeddings. Our model faithfully captures transcriptomic landscapes and achieves superior performance relative to state-of-the-art baselines on cell type annotation, graph structure understanding, and fine-tuned reverse perturbation prediction tasks. It offers a unified and interpretable framework for learning high-capacity foundational representations that capture complex long-range regulatory dependencies from high-dimensional single-cell transcriptomic data. The quality of network inference is not a focus of this study. Our framework applies to all graph-structured priors.

1 Introduction

Interaction networks have been broadly used to represent the molecular logic that mediates the effect of genetic and epigenetic events in determining cell states and whose dysregulation is responsible for a variety of human diseases, including cancer (Califano & Alvarez, 2017; Chen et al., 2014). We propose that advances in multi-omics single cell profiling, such as single cell RNA-seq (scRNA-seq), coupled with cellular graphs, provide a unique opportunity to create foundation models that improve our ability to elucidate the mechanisms that underlie tumorigenesis, tumor plasticity, and drug responses. Recent attempts to train large-scale transformer-based foundation models with scRNA data demonstrated potential in terms of both cell-level tasks (e.g.: cell type prediction, batch correction) and gene-level tasks (e.g.: gene regulatory network analysis, perturbation response prediction) (Cui et al., 2024; Hao et al., 2024; Theodoris et al., 2023). However, these models represent genes as discrete tokens and model cells as sequences, applying standard multi-head self-attention mechanisms, either in their original form (Theodoris et al., 2023), with attention biases informed by gene-gene relationships (Cui et al., 2024), or using randomized nonlinear kernel functions to encode structural priors (Hao et al., 2024). These formulations overlook the fact that genes products in the cell, as represented in scRNA-seq profiles, lack inherent sequential order or positional semantics, which is beneficial for inducing stable long-range attention decay patterns that facilitate the Transformer's capacity to model dependencies and generalize to sequences beyond those encountered during training (Kazemnejad et al., 2023; Su et al., 2021; Ke et al., 2020).

In this work, we introduce GREmLN (**G**ene **R**egulatory **Em**bedding-based **L**arge **N**eural model), a scRNA foundation model that leverages gene token graph structure to encode biologically meaningful relative position information and long-range dependency into single cell level gene embeddings. We first detail the formulation of our attention mechanism based on graph signal processing and numerical analysis. By applying a diffusion kernel to the graph Laplacian, we construct a kernel Gram matrix that can be used to transform the query embeddings, thereby enabling the self-attention mech-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

anism to be constrained and thus structured by the underlying graph. The diffusion process itself further helps capture long-range dependencies between gene nodes in the network. We also implement a Chebyshev polynomial based approximation of the kernel Gram matrix to scale to large graphs and long gene sequences. We then evaluate model performance with both cell level and gene level tasks such as cell type annotation with cell embeddings and graph structure understanding of gene embeddings to demonstrate superior performances over scRNA foundation model baselines.

2 Related Work

2.1 Graph Transformers

A growing body of work integrates graph structural priors into Transformer architectures to enhance relational reasoning and introduce domain-specific inductive biases (Min et al., 2022a). These methods generally fall into two categories: (1) incorporating graph structure into positional embeddings, and (2) modifying attention computations using graph information.

Popular positional encoding schemes include classical Laplacian eigenvectors (Dwivedi & Bresson., 2020) and their permutation-equivariant variants (Lim et al.; Huang et al.). Other approaches include SVD-derived embeddings (Hussain et al., 2021), centrality and degree features (Ying et al., 2021), and hop-based relative embeddings via auxiliary GNNs (Zhang et al., 2020). However, such embeddings often scale poorly to large or high-order graphs and risk introducing noisy or biased priors, limiting generalization. The work in Kanatsoulis et al. (2025) recently proposed learnable, GNN-generated encoding that act as drop-in, computationally efficient alternatives for Laplacian-based methods.

For attention modification, methods include designing equivariant attention mechanisms (Tholken & Fabritiis, 2022), using ego-centric subgraphs (Zhao et al., 2021), or appending GNN-based features with permutation-invariant Transformers (Wu et al., 2021). Others use attention masking (Mialon et al., 2021; Min et al., 2022b) or bias terms derived from shortest path distances or proximity metrics (Ying et al., 2021; Zhao et al., 2021). Yet, hard masking blocks all information flow and impairs gradient stability, while additive biases only shift attention scores without enabling structural message passing.

In contrast, we embed graph structure into a spectral kernel that filters noise and injects multi-hop, spectrum-aware inductive bias at low cost. This enhances convergence, preserves gradient propagation across regulatory edges, and enables each query to integrate contextual information from biologically related neighbors.

2.2 Biological Foundation Models

Similar to their use in natural language understanding, foundation models learn robust representations from biological data. Early work such as ESM (Rives et al., 2021) or ProGen (Madani et al., 2020) focused on proteins sequences but have since expanded to many other modalities such as RNA sequences (Zou et al., 2024), small molecules (Chithrananda et al., 2020), and most relevantly scRNA-seq (Cui et al., 2024; Theodoris et al., 2023). Foundation models for scRNA-seq frame the problem analogously to natural language, where each cell is a document and each gene is a word. All methods aim to develop a meaningful latent representation of RNA expression in a given cell that can subsequently be used for downstream tasks such as batch effect removal (Cui et al., 2024), cell type classification (Theodoris et al., 2023), drug sensitivity prediction (Hao et al., 2024), or the effects of genetic perturbations (Theodoris et al., 2023).

One key difference of scRNA-seq to other modalities is that scRNA-seq represents the expression value of each gene as a continous rather than discrete variable. Most models solve this by simply discretizing expression into bins and modeling expression using standard auto-regressive (Cui et al., 2024) or masked-token modeling (Theodoris et al., 2023) objectives. Newer models such as scFoundation model the read depth of each cell and employ a masked regression objective to better model the continuous nature of RNA expression (Hao et al., 2024). Another challenge with modeling scRNA-seq is the inherent lack of sequential structure. scRNA-seq expression vectors are unordered sets, and so standard sequential based positional encodings cannot be applied. Existing methods simply fix an arbitrary gene order; thus, this challenge remains unsolved.

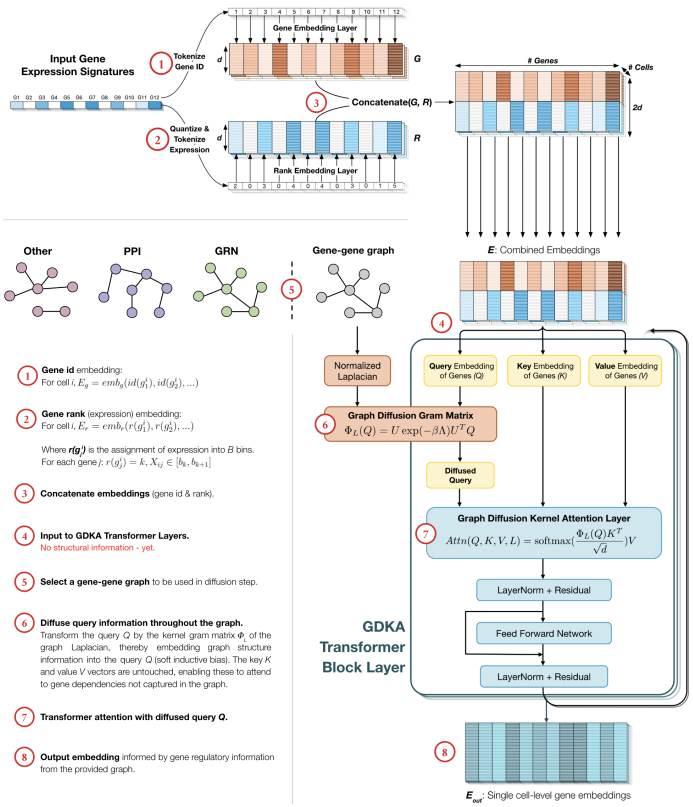

Figure 1: GREmLN Architecture: Input gene expression signatures are tokenized and embedded into gene ID (G) & gene expression (R) embeddings. The G & R embeddings are concatenated and fed to the Graph Diffusion Kernel Attention (GDKA) transformer. The GDKA transformer embeds graph structure information into the query Q, providing a soft inductive graph bias. The model can attend to gene relationships not captured in the graph via the key K and value V vectors.

3 Model Architecture

We propose a new model architecture that integrates gene token graph structure to facilitate the learning of single-cell-specific gene embeddings. By employing this approach, the model gains the ability to capture meaningful long-range dependencies among gene tokens, while also enriching the feature space with relevant positional and structural relationships between these tokens. Such a framework applies to arbitrary gene token interaction graphs, including gene regulatory networks, protein-protein interaction networks and protein subcellular localization networks.

3.1 Tokenization & Initial Embeddings

We can extract two types of information from a scRNA count matrix: (1) Gene identity, and (2) Gene expression value, which is a measurement of how expressed a gene is in a cell. An unprocessed scRNA count matrix (raw count) contains positive integers representing the number of RNA transcript copies of a gene. As illustrated in Figure 1, the inputs to the model are a gene identity embedding matrix (denoted as E^g), and a gene rank embedding matrix (denoted as E^r). We construct the gene identity embedding for cell i as:

$$E^g = emb_g(id(g_1^i),\ id(g_2^i),\ \ldots,\ id(g_G^i))$$

where emb_g is a learnable embedding layer. Gene IDs are never masked. To construct the gene rank embedding, we first compute the rank of gene j in cell i as:

$$r(g_j^i) = k, \hat{X}_{ij} \in [b_k, b_{k+1}]$$

where $\hat{X}$ is the normalized expression matrix, and b_k, b_{k+1} are lower and upper bounds of bin k. The normalized expression matrix is obtained by normalizing the total counts of all cells to sum up to a fixed number (in our case 10^6) and then log-plus-one transforming it. The <MASK> rank token is randomly assigned to 30% of the rank values. The gene rank embedding of cell i is then given by:

$$E^r = emb_r(r(g_1^i),\ r(g_2^i),\ \ldots,\ r(g_G^i))$$

where emb_r is a learnable embedding layer. The input embedding to the transformer is then given by the concatenation of two representations:

$$E = [E^g || E^r]$$

We also prepend a <CLS> token to the beginning of each cell sentence that attends to all genes. The <CLS> token is not part of the graph and is not involved in any graph operations.

3.2 Graph Diffusion Kernel Attention

Leveraging the fact that we observe existing token graphs, we incorporate graph-based structural information into the attention mechanism by transforming each query embedding with a diffusion kernel Gram matrix (KGM) derived from the token graph (Figure 1). In doing so, the query vectors become conditioned on graph topology - effectively biasing the attention process to reflect the gene-gene relationships encoded in the graph. Key and value embeddings, meanwhile, are not transformed to preserve the original information in token embeddings. We aim to develop general embeddings that inherently capture token graph structures. More specifically, assume the token graph adjacency matrix is A, we compute the normalized Laplacian matrix as:

$$L = I - D^{-1/2} A D^{-1/2},$$

where D is the diagonal degree matrix. Using the normalized Laplacian matrix ensures that the eigenvalues remain bounded when mapping the graph to the spectral domain. To that end, let $L \in \mathbb{R}^{G\times G}$ be a symmetric matrix with spectral decomposition

$$L \;=\; U\Lambda U^{\top},\; U = \big[u_1, \ldots, u_G\big],\; \Lambda = \operatorname{diag}\big(\lambda_1, \ldots, \lambda_G\big).$$

For a non–negative spectral filter

$$\Phi_L \;=\; \sum_{\ell=1}^{G} \kappa(\lambda_\ell)\, u_\ell u_\ell^{\top} \;=\; U\,\kappa(\Lambda)\,U^{\top} \tag{1}$$

with $\kappa(\Lambda) = \operatorname{diag}\big(\kappa(\lambda_1), \ldots, \kappa(\lambda_G)\big)$. Because $\kappa(\lambda_i) \geq 0$ for every i, the matrix Φ_L is symmetric positive semidefinite and therefore constitutes the matrix representation of an orthogonal projection operator onto the Reproducing Kernel Hilbert Space $\mathcal{H}_\kappa$ whose reproducing kernel satisfies

$$k_\kappa(i,j) \;=\; e_i^{\top}\Phi_L e_j, \qquad i,j \in \{1,\ldots,G\}.$$

where $\kappa(.,.)$ is any kernel function.

$$\Phi_L(Q) = U \exp(-\beta\Lambda) U^T Q$$

We then compute dot-product attention with the diffused query embedding as:

$$Attn(Q,K,V,L) = \operatorname{softmax}(\frac{\Phi_L(Q)K^T}{\sqrt{d}})V$$

where un-softmaxed attention score for each pair is given by the following similarity measure:

$$S_{ij} = q_i^{\top}\Phi_L^{1/2} k_j$$

This bilinear score incurs asymmetric smoothing which steers the attention mechanism toward interactions that respect the graph-implied low-frequency (long-range) decay, while still allowing keys to contribute high-frequency detail.

3.3 THE PASSTHROUGH METHOD

When only a subset of genes $S \subseteq \{1, \ldots, G\}$ is retained, due to a selection criterion or computational limitations, the conventional approach recomputes the kernel on the induced subgraph $G[S]$,

$$\Phi_{L_S} = U_S \exp(-\beta \Lambda_S) U_S^\top, \tag{2}$$

where $L_S = U_S \Lambda_S U_S^\top$ is the normalized Laplacian of $G[S]$. This discards every edge incident to $\{1, \ldots, G\} \setminus S$, sparsifying the network and disconnecting otherwise connected components.

We instead subset the *master* KGM Φ_L of Eq. equation 1, computed once on the full graph:

$$\Phi_L^{\text{pass}} = \Phi_L[S, S]. \tag{3}$$

Each entry of Φ_L is determined solely by the participation of its two nodes in the global eigenmodes,

$$(\Phi_L)_{ij} = \sum_{\ell=1}^{G} e^{-\beta \lambda_\ell} \, u_\ell(i) \, u_\ell(j), \tag{4}$$

and is therefore independent of every other pair. Subsetting Φ_L consequently preserves the full-graph scores exactly, whereas recomputation on a subgraph alters the eigenstructure:

$$\begin{aligned} (\Phi_L^{\text{pass}})_{ij} &= (\Phi_L)_{ij}, \\ (\Phi_{L_S})_{ij} &\neq (\Phi_L)_{ij} \quad \text{in general}, \quad i, j \in S. \end{aligned} \tag{5}$$

Because Φ_L is positive semidefinite and Φ_L^{pass} is a principal submatrix of it, Φ_L^{pass} too is positive semidefinite and thereby a valid kernel on S. The passthrough method thus conditions on the full graph for any node subset at no additional runtime cost beyond the one-time precomputation of Φ_L. Intuitively, this can be seen as a full graph with only the nodes in S emitting heat.

3.4 TRAINING OBJECTIVES

We formulate the pre-training objective as a masked modeling problem conditioned on graphs. Assume we have G genes, let $X^r = (X_1^r, X_2^r, .., X_G^r)$ be the bin values of genes in a cell, L^c be the normalized Laplacian matrix of the token graph specific to cell type c which the cell belongs to, $m = (m_1, m_2, ..., m_G)$ be a random binary mask vector where $m = 1$ represent genes being masked. The likelihood of all masked genes is given by

$$p_\theta(X_m^r | L^c, X_{\bar{m}}^r, E^g) = \prod_{i, \forall m_i = 1} p_\theta(x_i^r | L^c, X_{\bar{m}}^r, E^g)$$

where E^g is the gene identity embedding, X_m^r are the bin values of masked genes, θ are all learnable parameters, and $X_{\bar{m}}^r$ are the bin values of unmasked genes. Denote the transformer encoder with graph diffusion kernel attention as enc and the linear decoder network as dec, a masked gene i's conditional distribution can be parameterized as:

$$\begin{aligned} E_{\bar{m}}^r &= emb_r(X_{\bar{m}}^r) \\ H &= enc(L^c, E_{\bar{m}}^r, E^g) \\ p_\theta(x_i^r | L^c, X_{\bar{m}}^r, E^g) &= dec(H_i) = \sigma(W H_i + b) \end{aligned}$$

where H is the learned embedding by the transformer encoder. The pre-training objective is then given by maximizing the expected conditional log likelihood:

$$\mathcal{L} = -\mathbb{E}_{X^r \sim \mathcal{D}} \mathbb{E}_{m \sim p(m)} [\sum_{i, \forall m_i = 1} \log p_\theta(x_i^r | L^c, X_{\bar{m}}^r, E^g)]$$

where $\mathcal{D}$ is some ground truth data distribution and $p(m)$ is the sampling distribution for the mask.

3.5 APPROXIMATION OF LARGE KERNEL GRAM MATRICES

Due to the size of biological networks and length of cell sentences, performing operations such as matrix exponential and spectral decomposition every batch to transform query embeddings becomes

computationally expensive and difficult to scale. To alleviate this issue, we take advantage of Chebyshev polynomials to approximate the kernel Gram matrix. More specifically, if a token graph has G nodes, we approximate the transformed query embedding with Chebyshev polynomial truncated at term $K \ll G$ as:

$$\Phi_L(Q) = U \exp(-\beta\Lambda)U^T Q \approx \sum_{k=1}^{K} c_k T_k(L) Q$$

to efficiently compute the kernel gram matrix of large graphs. This approximation can be solved with Gauss-Chebyshev quadrature, which we detail in Appendix A.3. We also provide an analysis discussing the asymptotic bound of runtime for this operation in Appendix A.4.

4 Validation Experiments

In this section, we compare our model to state-of-the-art transformer-based single-cell transcriptomics foundation models: scGPT (Cui et al., 2024), Geneformer (Theodoris et al., 2023), and scFoundation (Hao et al., 2024). For graph input, we use gene regulatory networks and PPI networks as examples due to their broad availability, though other gene token graphs could be used (network generation details in Appendix A.1 and A.10). We also evaluate against minimal-capacity models - linear models and shallow MLPs - so that performance predominantly reflects the pre-trained embeddings rather than the downstream model's expressiveness. We first assess gene expression reconstruction and cell type classification to evaluate cell embedding quality, then evaluate how well the learned gene embeddings capture graph structural information, and lastly evaluate the model's ability to predict perturbation labels from single-target Perturb-Seq data, using both frozen and fine-tuned embeddings.

4.1 Evaluation Datasets

We select evaluation datasets that are non-overlapping with those used during pre-training to assess the generalizability of all models. In particular, we focus on a dataset consisting of healthy human immune cells, as we are interested in downstream applications of our model to human immune cell state engineering. Aside from healthy immune cells, we also test the model on diverse datasets derived from cancer-infiltrating immune cells and non-immune cells. For perturbation tasks, we use the dataset from Adamson et al. (2016) and Norman et al. (2019). Details and availability of the datasets can be found in Appendix A.8.

4.2 Transcriptomic Landscape Learning & Cell Type Annotation

We first validate of our model's effective learning of mRNA transcripts by training a simple linear regression model to reconstruct gene expression from frozen pre-trained embeddings. Our model is able to accurately reconstruct gene expression for both healthy and diseased human cells (Appendix A.9). We then assess the quality of our learned embeddings' cellular representations by measuring their ability to recapitulate cell type.

Since our model requires graphs as input, we only allow the model to see integrated networks that combines all cell-type specific GRNs to avoid inadvertently leaking cell-type labels.

Bayesian Graph Integration Let X_{train} be the training gene-expression matrix and let $X_{\text{test}}^{(i)}$ denote the profile of a single test cell i. For every cell type $c \in \mathcal{C}$ we compute the posterior edge-exclusion probability $\Pr\big(e \notin \mathcal{G}_c \mid X_{\text{train}}, C = c\big)$ and declare the edge present in the type-specific estimate $\widehat{\mathcal{G}}_c$ whenever this posterior does not exceed a user-defined false-edge level α,

$$e \in \widehat{\mathcal{G}}_c \iff \Pr\big(e \notin \mathcal{G}_c \mid X_{\text{train}}, C = c\big) \le \alpha. \tag{6}$$

A k-nearest-neighbor classifier fitted exclusively on X_{train} yields an empirical posterior $w_c^{(i)} = \Pr\big(C = c \mid X_{\text{test}}^{(i)}\big)$, with $\sum_{c \in \mathcal{C}} w_c^{(i)} = 1$. Using the law of total probability,

$$p_i(e) = \sum_{c \in \mathcal{C}} w_c^{(i)} \Pr\big(e \notin \mathcal{G}_c \mid X_{\text{train}}, C = c\big). \tag{7}$$

Table 1: Test cell type annotation performance across multiple tissue types. CI is computed across 3 random initializations of test splits. scGPT is excluded from evaluation on the Held-out Non-Immune Cell dataset, as it does not constitute a zero-shot setting. Leading zeros are omitted from confidence intervals.

Dataset	Human Immune Cells			Non-Immune Cells (zero-shot)		
Model	Precision	Recall	Macro F1	Precision	Recall	Macro F1
GREmLN	**0.941 ± .002**	**0.937 ± .003**	**0.939 ± .001**	**0.936 ± .002**	**0.936 ± .001**	**0.937 ± .001**
scGPT	0.928 ± .003	0.924 ± .001	0.924 ± .002	-	-	-
scFoundation	0.887 ± .013	0.879 ± .017	0.879 ± .015	0.913 ± .002	0.904 ± .003	0.903 ± .003
Geneformer	0.800 ± .012	0.797 ± .010	0.792 ± .010	0.774 ± .015	0.768 ± .010	0.768 ± .012

Table 2: Benchmarking performance for graph understanding. Left: Human Immune dataset. Right: Cancer Infiltrating Myeloid dataset. Metrics: area under the ROC curve (AUROC), with confidence intervals from cross-validation, and average precision (AP) from the precision-recall curve. Leading zeros are omitted from confidence intervals.

	Human Immune		Cancer Myeloid	
Model	AUROC	AP	AUROC	AP
GREmLN	**0.957 ± .005**	**0.947**	**0.976 ± .001**	**0.964**
Geneformer	0.944 ± .006	0.944	0.932 ± .002	0.916
scGPT	0.916 ± .009	0.922	0.897 ± .003	0.867
scFoundation	0.775 ± .015	0.783	0.965 ± .001	0.949
Chance	0.500	-	0.500	-

The regulatory graph supplied to the model for cell i is finally obtained via the same α-level threshold,

$$\widehat{\mathcal{G}}_i = \big\{\, e : \; p_i(e) \le \alpha \big\}. \tag{8}$$

Equations equation 6-equation 8 collectively recast the procedure as two successive Bayesian updates: first, learning type-specific edge posteriors on X_{train}; second, averaging these posteriors for each unseen cell under the k-NN-derived type posterior, thereby integrating multiple cell-type-specific graphs without ever exposing the true test-cell labels.

Training cell type classifier For each cell i, we begin with its embedding matrix $H_i \in \mathbb{R}^{L\times d}$, where L is the gene dimension and d is the embedding dimension. Next, we obtain a fixed-size cellular representation $h_i \in \mathbb{R}^d$ via mean pooling over the gene sequence dimension. Next, we fit a linear classifier to predict cell type given h_i and the integrated graph $\widehat{\mathcal{G}}_i$ by minimizing cross-entropy loss between the estimated class distribution and the true cell type label. More specifically, assuming we have cell type label y_i, we aim to learn the distribution of cell type labels conditioning on the cell representation and the input graph. For cell type c the conditional density is given by:

$$p(y_i = c | h_i, \widehat{\mathcal{G}}_i, f_\theta)$$

where f_θ are the classifiers being trained. We learn f_θ by optimizing cross-entropy loss as:

$$\mathcal{L}(\theta) = -\sum_{c\in C}\sum_{i=1}^{N} \log p(y_i = c | h_i, \widehat{\mathcal{G}}_i, f_\theta)$$

Our model achieved highest performances on all classification metrics on the human immune cells dataset (Table 1). Additionally, for held-out non-immune cells we test the model's ability to perform zero-shot cell type annotation. Since scGPT is trained with the full CELLxGENE dataset, we exclude it from this evaluation. We achieve superior zero-shot performances on the held-out non immune cell dataset (Table 1) for all metrics. We also point out that GREmLN contains only 14.7 million learnable parameters (less than half of all baselines' parameter sets and one sixth of scFoundation, see details in Appendix A.5), which highlights the predictive power of gene regulatory network guidance. Our key contribution lies in tailoring the attention mechanism to encode biological inductive bias, rather than increasing model size.

4.3 GRAPH STRUCTURE UNDERSTANDING

We assess the learned embeddings' capacity to capture and generalize graph structure by evaluating edge prediction on unseen gene regulatory networks from held-out cells. Specifically, we randomly mask a subset (15%) of edges and measure how well the model recovers them when conditioned on the learned embeddings and the incomplete graph.

Let H be learned gene embeddings from a model and A be the adjacency matrix of the graph, the objective is to learn the conditional probability of a masked edge existing, i.e.:

$$p_\theta(A_{i,j} = 1 | \tilde{A} = A \odot \mathcal{M}, H)$$

We use the learned embeddings H from pre-trained models and parameterize the conditional likelihood with a model f_ψ where ψ are learned model parameters. More specifically, for an edge (i, j):

$$p_\theta(A_{i,j} = 1 | \tilde{A} = A \odot \mathcal{M}, H) = f_\psi(i, j, \tilde{A}, H_i, H_j)$$

where H_i, H_j are learned embeddings of nodes i, j respectively. To learn the model f_ψ we treat an edge-label pair $\{(i, j), y_{ij}\}$ as i.i.d samples from some data generating distributions. The label $y_{ij} = 1$ if the edge (i, j) is a real masked edge and $y_{ij} = 0$ if the edge is a fake edge that does not exist in A. The training objective is then given by:

$$\begin{aligned}\mathcal{L}(\psi) = - \sum_{\{(i,j),y_{ij}\}\sim\mathcal{D}} & [y_{ij} \log f_\psi(i, j, \tilde{A}, H_i, H_j) \\ & + (1 - y_{ij}) \log(1 - f_\psi(i, j, \tilde{A}, H_i, H_j))]\end{aligned}$$

In practice, we freeze our encoder weights and incorporate the partially observed graph $\tilde{A}$ through the attention mechanism to condition our embeddings. For baseline models, we similarly condition on $\tilde{A}$ by passing the learned gene embeddings through a graph attention network using the same masked graph. We first benchmark GRN structure understanding on healthy human datasets, evaluating seen cell types but unseen data/graph. GREmLN achieves the best performance by a significant margin on both ROC and PRC, while Geneformer is also competitive due to its tokenization and pre-training objective, which focus on learning GRN-aware gene embeddings (Table 2).

To further evaluate structural generalization, we tested the model on graph structure understanding for previously unseen cell types. Since pathological conditions such as cancer often induce local rewiring of gene regulatory networks, we specifically evaluated cancer-infiltrating myeloid cells from distinct tissue origins, held out to simulate out-of-distribution generalization. GREmLN consistently outperformed all baseline models across all metrics (Table 2), demonstrating superior generalization to unseen regulatory graphs within complex, heterogeneous biological systems.

4.4 FINE-TUNING FOR REVERSE PERTURBATION PREDICTION

We lastly evaluate fine-tuning for reverse perturbation prediction on Perturb-Seq data: given a perturbed cell expression profile, the model needs to infer the corresponding perturbation label. Single-gene CRISPR perturbations induce only subtle transcriptomic shifts, yielding low class separability and making the task intrinsically challenging.

For a fair assessment, we first benchmark GREmLN with scFoundation and Geneformer, which report state-of-the-art performance, using frozen embeddings. We train GREmLN with both GRN and curated PPI network inputs and examine their respective performance from frozen embeddings. The PPI networks used are HINT (Das & Yu, 2012) and STRING (Szklarczyk et al., 2025); see Appendix A.10 for more PPI network details. We also include a linear model as a baseline.

We then fine-tune the best performing GREmLN model from the frozen benchmark, the GRN model, with classification entropy loss and report its fine-tuned performance. We didn't fine-tune the baseline models due to time and computational resource constraints.

Table 3 shows GREmLN achieves state-of-the-art performance with and without fine-tuning, with the GRN model outperforming all other frozen models - emphasizing the generalizability of our graph-structure based attention mechanism. These results also capture the benefit of the context-specific design of gene regulatory networks, which encode cell-type-specific circuitry. Notably,

Table 3: Reverse perturbation prediction across models. Top-10 AUC: for each held-out cell, we rank candidate perturbations by predicted score and compute top-k accuracy - the fraction of cells whose true perturbation falls within the k highest-ranked candidates - for $k = 1, \ldots, 10$, then report the mean of these values (the normalized area under the cumulative match characteristic curve up to rank 10). Frozen-embedding results cover baseline foundation models and GREmLN variants pre-trained on different graphs, followed by fine-tuned GREmLN (GRN) at increasing depth; GRN was chosen for fine-tuning based on its superior frozen-embedding performance.

Model	Parameters	Accuracy	F1 Score	Top-10 AUC
Frozen embeddings				
random guessing	-	0.05	0.05	0.50
scFoundation-frozen	100M	0.143	0.103	0.515
Geneformer-frozen	30M	0.162	0.134	0.585
Lasso logistic regression	-	0.281	0.127	0.608
GREmLN-HINT-frozen (3L)	14.7M	0.421	0.357	0.781
GREmLN-STRING-frozen (3L)	14.7M	0.427	0.360	0.783
GREmLN-GRN-frozen (3L)	**14.7M**	**0.435**	**0.372**	**0.795**
Fine-tuned embeddings				
GREmLN-GRN-finetuned (1L)	8.4M	0.461	0.405	0.809
GREmLN-GRN-finetuned (3L)	14.7M	0.572	0.532	0.872
GREmLN-GRN-finetuned (6L)	**24.2M**	**0.601**	**0.564**	**0.877**

GREmLN attains this performance with less than half the trainable parameters of the next smallest model, Geneformer. Finally, fine-tuning the GRN model yields a further 14.5% increase in predictive accuracy.

4.5 Scaling Behavior

We lastly examine whether GREmLN exhibits scaling behavior, using the GRN configuration since it performed best on the Adamson reverse perturbation task from frozen embeddings (Table 3). We pre-trained and fine-tuned three architectures of increasing depth (1, 3, and 6 Graph Diffusion Kernel Attention Transformer blocks), yielding models with 8.4M, 14.7M, and 24.2M parameters, respectively, all trained for 4 epochs with identical configuration, including dataset and weight decay. On the reverse perturbation prediction task (Section 4.4), these models showed monotonic improvement with increasing size (Table 3), demonstrating that GREmLN benefits systematically from scale.

5 Conclusion & Discussion

In this work, we introduce GREmLN, a graph-aware foundation model that integrates gene token graphs, learning biologically informed gene embeddings through graph-induced ordering. These embeddings can reconstruct gene expression via a linear regression model, demonstrating that transcriptomics information is well-captured. The model additionally achieves state-of-the-art performance in predicting held-out cell types and unseen regulatory structures, illustrating an ability to capture generalizable rules of gene regulation and cell state representation. The model can be pre-trained with both PPI and GRN inputs, and can be fine-tuned to distinguish cells with small transcriptomics differences from Perturb-seq data with different perturbation labels.

While the functional search space for gene interactions is intractable even for large models, incorporating validated regulatory interactions significantly improves performance across downstream tasks (Ideker & Krogan, 2012; Bar-Joseph et al., 2003; Alvarez et al., 2016). Our architecture supports arbitrary molecular interaction graphs and leverages well-established methods for network inference (Margolin & Nemenman, 2006) alongside diffusion-based graph kernels (Paull et al., 2013; Leiserson et al., 2015). Future directions include fine-tuning cell and gene embeddings for specific tasks - such as modeling combinatorial perturbations and identifying optimal interventions - while using attention-based interpretability to recover core regulatory modules.

AI USE STATEMENT

In this work, we used generative AI tools for: providing feedback on research methodology or experiments, implementing methods, assisting in the writing of proofs, supporting qualitative and thematic data analysis, interpreting results. We have not used generative AI tools for: helping develop theoretical models or conceptual frameworks, formulating mathematical claims, providing critical ingredients for proving mathematical claims, proposing or refining hypotheses, cleaning and reformatting datasets. Generating synthetic data sets and assisting with translation are not applicable to this work.

Additionally, we used generative AI tools for: creating or modifying scientific figures or images, suggesting experimental parameters, creating or editing software code, drafting parts of a research paper, summarizing or analyzing existing literature, discovering research topics or identifying gaps, brainstorming, sourcing/searching for information, editing a research paper to improve readability, identifying relevant literature, formatting references.

We have reviewed all AI-assisted work. LLM-generated code was verified and tested for correctness by two or more authors. Any other AI-assisted work in research methodology feedback, proof writing, data analysis, and interpreting results was scrutinously reviewed by the authors. The same precaution was applied to AI-assisted work in figure visualization code generation, text polishing recommendations, summarization of literature, and other aforementioned applications. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

ETHICS STATEMENT

The authors have no concerns pertaining to the Code of Ethics to report.

REPRODUCIBILITY STATEMENT

We describe the proposed methods conceptually in this paper (main body and Appendix). Experiment code is not included with the submission. All datasets drawn from existing literature are publicly available and cited in Appendix A.8, where we also explain why we chose them. We formally define the evaluation metrics and explain why we chose them in Section 4. We partially report the hyperparameter search ranges, the selection criteria, and the final hyperparameter values in Appendix A.5, and we partially describe the computing infrastructure in Appendix A.5 & A.6.

REFERENCES

Britt Adamson, Marco Jost Thomas M Norman, Min Y Cho, and et. al. A multiplexed single-cell crispr screening platform enables systematic dissection of the unfolded protein response. *Cell*, 2016.

Mariano J Alvarez, Yao Shen, Federico M Giorgi, Alexander Lachmann, B Belinda Ding, B Hilda Ye, and Andrea Califano. Network-based inference of protein activity helps functionalize the genetic landscape of cancer. *Nature Genetics*, 48(8):838, 2016.

Mariano J Alvarez, Yao Shen, Federico M Giorgi, Alexander Lachmann, B Belinda Ding, B Hilda Ye, and Andrea Califano. Functional characterization of somatic mutations in cancer using network-based inference of protein activity. *Nature Genetics*, 2019.

Ziv Bar-Joseph, Georg K Gerber, Tong Ihn Lee, Nicola J Rinaldi, Jane Y Yoo, François Robert, D Benjamin Gordon, Ernest Fraenkel, Tommi S Jaakkola, Richard A Young, et al. Computational discovery of gene modules and regulatory networks. *Nature biotechnology*, 21(11):1337–1342, 2003.

Andrea Califano and Mariano J. Alvarez. The recurrent architecture of tumour initiation, progression and drug sensitivity. *Nature Reviews Cancer*, 2017.

James C. Chen, Mariano J. Alvarez, and Flaminia Talos et al. Identification of causal genetic drivers of human disease through systems-level analysis of regulatory networks. *Cell*, 2014.

Seyone Chithrananda, Gabriel Grand, and Bharath Ramsundar. Chemberta: Large-scale self-supervised pretraining for molecular property prediction. *arxiv*, 2020.

Haotian Cui, Chloe Wang, and et. al. scgpt: Toward building a foundation model for single-cell multi-omics using generative ai. *Nature Methods*, 2024.

Tri Dao, Daniel Fu, Stefano Ermon, et al. Flashattention: Fast and memory-efficient exact attention with io-awareness. *arxiv*, 2022.

Jishnu Das and Haiyuan Yu. HINT: High-quality protein interactomes and their applications in understanding human disease. *BMC Systems Biology*, 6(1):92, 2012.

Vijay Prakash Dwivedi and Xavier Bresson. A generalization of transformer networks to graphs. *arxiv*, 2020.

Minsheng Hao, Jing Gong, and et. al. Large-scale foundation model on single-cell transcriptomics. *Nature Methods*, 2024.

Yinan Huang, William Lu, Joshua Robinson, Yu Yang, Muhan Zhang, Stefanie Jegelka, and Pan Li. On the stability of expressive positional encodings for graphs. In *The Twelfth International Conference on Learning Representations*.

Md Shamim Hussain, Mohammed J Zaki, and Dharmashankar Subramanian. Edge-augmented graph transformers: Global self-attention is enough for graphs. *arxiv*, 2021.

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

Charilaos I Kanatsoulis, Evelyn Choi, Stephanie Jegelka, Jure Leskovec, and Alejandro Ribeiro. Learning efficient positional encodings with graph neural networks. In *The Thirteenth International Conference on Learning Representations*, 2025.

Amirhossein Kazemnejad, Inkit Padhi, and et. al. The impact of positional encoding on length generalization in transformer models. *NeurIPS*, 2023.

Guolin Ke, Di He, and Tie-Yan Liu. Rethinking positional encoding in language pre-training. *ICLR*, 2020.

Mark DM Leiserson, Fabio Vandin, Hsin-Ta Wu, Jason R Dobson, Jonathan V Eldridge, Jacob L Thomas, Alexandra Papoutsaki, Younhun Kim, Beifang Niu, Michael McLellan, et al. Pan-cancer network analysis identifies combinations of rare somatic mutations across pathways and protein complexes. *Nature genetics*, 47(2):106–114, 2015.

Derek Lim, Joshua David Robinson, Lingxiao Zhao, Tess Smidt, Suvrit Sra, Haggai Maron, and Stefanie Jegelka. Sign and basis invariant networks for spectral graph representation learning. In *The Eleventh International Conference on Learning Representations*.

Malte D. Luecken, M. Büttner, and et. al. Benchmarking atlas-level data integration in single-cell genomics. *Nature Methods*, 2021.

Ali Madani et al. Progen: Language modeling for protein generation. *arxiv*, 2020.

Adam A Margolin and Ilya Nemenman. Aracne: An algorithm for the reconstruction of gene regulatory networks in a mammalian cellular context. *BMC Bioinformatics*, 2006.

Gregoire Mialon, Dexiong Chen, Margot Selosse, and Julien Mairal. Graphit: Encod- ing graph structure in transformers. *arxiv*, 2021.

Erxue Min, Runfa Chen, and Yatao Bian. Transformer for graphs: An overview from architecture perspective. *arxiv*, 2022a.

Erxue Min, Yu Rong, Tingyang Xu, and Yatao Bian. Masked transformer for neighhourhood-aware click-through rate prediction. *arxiv*, 2022b.

Prabhjot S. Mundi, Filemon S. Dela Cruz, A. Grunn, et al. A transcriptome-based precision oncology platform for patient-therapy alignment in a diverse set of treatment-resistant malignancies. *Cancer Discovery*, 2023.

Thomas M Norman, Max A Horlbeck, Joseph M Replogle, Alex Y Ge, and et. al. Exploring genetic interaction manifolds constructed from rich single-cell phenotypes. *Science*, 2019.

Adam Paszke, Sam Gross, Francisco Massa, et al. Pytorch: An imperative style, high-performance deep learning library. *arxiv*, 2019.

Evan O Paull, Daniel E Carlin, Mario Niepel, Peter K Sorger, David Haussler, and Joshua M Stuart. Discovering causal pathways linking genomic events to transcriptional states using tied diffusion through interacting events (tiedie). *Bioinformatics*, 29(21):2757–2764, 2013.

Evan O. Paull, Alvaro Aytes, Sunny J. Jones, et al. A modular master regulator landscape controls cancer transcriptional identity. *Cell*, 2021.

Alexander Rives, Joshua Meier, Tom Sercu, Siddharth Goyal, Zeming Lin, Jason Liu, Demi Guo, Myle Ott, C. Lawrence Zitnick, Jerry Ma, and Rob Fergus. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proceedings of the National Academy of Sciences*, 118(15):e2016239118, 2021. doi: 10.1073/pnas.2016239118.

Jianlin Su, Yu Lu, and et. al. Roformer: Enhanced transformer with rotary position embedding. *arxiv preprint*, 2021.

Damian Szklarczyk, Katerina Nastou, Mikaela Koutrouli, Rebecca Kirsch, Farrokh Mehryary, Radja Hachilif, Dewei Hu, Matteo E. Peluso, Qingyao Huang, Tao Fang, Nadezhda T. Doncheva, Sampo Pyysalo, Peer Bork, Lars J. Jensen, and Christian von Mering. The STRING database in 2025: Protein networks with directionality of regulation. *Nucleic Acids Research*, 53(D1):D730–D737, 2025.

Christina V. Theodoris, Ling Xiao, and et. al. Transfer learning enables predictions in network biology. *Nature*, 2023.

Philipp Tholken and Gianni De Fabritiis. Torchmd-net: Equivariant transformers for neural network based molecular potentials. *ICLR*, 2022.

Zhanghao Wu, Paras Jain, and et. al. Representing long-range context for graph neural networks with global attention. *NeurIPS*, 2021.

Chengxuan Ying, Tianle Cai, and et. al. Do transformers really perform bad for graph representation? *arxiv*, 2021.

Tizita Z Zeleke, Qingfei Pan, and et. al. Network-based assessment of hdac6 activity predicts preclinical and clinical responses to the hdac6 inhibitor ricolinostat in breast cancer. *Nature Cancer*, 2023.

Jiawei Zhang, Haopeng Zhang, and et. al. Graph-bert: Only attention is needed for learning graph representations. *arxiv*, 2020.

Jianan Zhao, Chaozhou Li, and et. al. Gophormer: Ego-graph transformer for node classification. *arxiv*, 2021.

Shuxian Zou, Tianhua Tao, and et. al. A large-scale foundation model for rna function and structure prediction. *bioarxiv*, 2024.

A Appendix

A.1 Gene Regulatory Network Construction

To construct gene regulatory networks, we use the ARACNe algorithm (Margolin & Nemenman, 2006) which infers the transcriptional targets of regulatory proteins (represented as edges in a network) by estimating the mutual information between their gene expression via bootstrapping and adaptive partitioning, followed by elimination of indirect targets via the data processing inequality theorem. ARACNe is a key component of several well-validated algorithms that have been effective in elucidating master regulator proteins representing mechanistic determinants of tumor cell states as well as in determining their sensitivity to small molecules, including use in clinical trials. Additionally, the networks from the algorithm are highly predictive and has a potentially higher false negative rate, but very low false positive rate as a result, since it produces the sparsest network that can still account for all the information transfer in the cell due to the pruning. Such constitution acts excellent first approximation of the underlying network, which is ideal for our setting where graph is used as an inductive bias. (Alvarez et al., 2016; 2019; Mundi et al., 2023; Paull et al., 2013; 2021; Zeleke et al., 2023).

Gene regulatory networks vary widely across cell types, so prior to model pre-training, we infer a unique GRN for each cell type in our training corpus. To reconstruct each network, we first run ARACNe within every Louvain-algorithm-inferred cluster. The resulting subnetworks are then consolidated into a single network by evaluating the likelihood of edge recurrence under a binomial null model, allowing the false discovery rate to be controlled at 5%. Forming a consensus network from cluster-specific subnetworks offers several advantages: (1) consensus networks are robust to potentially significant batch effects; (2) almost all cells of each cell type contribute to the final network, and (3) multiple possible cell states are faithfully represented in the single network, enabling embeddings to capture the full granularity of single-cell biological states.

While we use ARACNe in this work, we emphasize that the model is agnostic to the specific GRN inference procedure: any inferred GRN can serve as input, and we do not aim to benchmark or assess the relative quality of network construction methods here.

A.2 Preprocessing Pipeline and Implementation Details

The pre-training dataset consists of 11M scRNA-seq profiles spanning 19K genes from healthy human cells, sourced from the expansive CELLxGENE dataset (`https://cellxgene.cziscience.com/collections`). Our single cell dataset covers 379 cell-types from various tissues (heart, blood, brain, lung, kidney, intestine, pancreas, and others). The specificity of the cell-types is as indicated in the `obs.cell_type` attribute of the CELLxGENE corpus for corresponding cells.

Preprocessing Pipeline Single cell RNA sequencing data is inherently noisy, subject to batch effects, sequencing errors, and high dropout rates. In fact, the sparsity of the expression profiles in our dataset is approximately 80%. Extracting a meaningful signal from such profiles is challenging; thus, it is critical to pre-process this data appropriately before feeding it to the model. In this work we choose to discretize the cell sentences. Discretization is not only necessary for the masked token prediction pre-training task but also to attenuate batch effects in scRNA-seq data. Indeed, rank-based binning produces a dataset-invariant scaling. Since we optimize cross entropy loss during pre-training, the optimal choice of number of bins would be enough bins to produce a near-uniform occupancy distribution. In practice, the bin number is treated as a hyper-parameter and selected via grid search on validation data with an increment of 50 bins. Previous single cell transcriptomics foundation models such as scGPT and Geneformer also rely on similar pre-processing steps.

Our pre-processing pipeline consists of the following steps: quality control and normalization, quantization, regulatory network generation, and caching. (1) The cells are controlled for quality by filtering out any cells not meeting our UMI sequencing depth requirements. For a given cell, the UMI count of a gene is a measure of that gene's expression level. UMI count thresholds are set to a minimum and maximum of 1K and 100K reads, respectively, and - if needed - are dynamically updated to keep only the middle 90% of cells by total UMI count. Cells with insufficient total UMI counts are removed as they likely represent faulty or empty reads from the sequencing process, and cells

with exceedingly large total UMI counts are removed as they likely represent doublets and multiplets (the read contains information from more than one cell). The remaining expressions are log1p normalized to reduce the range of expression values and stabilize the expression variance. (2) After quality controlling, the cells are quantized into 100 expression bins. The zero-th bin is assigned to all zero-expression genes and the remaining expressed genes are assigned to their corresponding quantile for that cell; the bin number increasing with expression. This provides a discretized relative expression of the genes in each cell, preserving cell insights while reducing the computational burden of processing this data during training. (3) Gene Regulatory Networks (GRNs) are then generated for each cell-type using the ARACNe algorithm (Margolin & Nemenman, 2006). The use of GRNs is at the heart of the novelty of this model, framing expression signatures in a regulatory context. (4) In the final caching step, a GRN sub-network including only the non-zero expressed genes is created for each cell. It is also required that each single cell meets the minimum number of 500 expressed genes.

To evaluate the model's zero-shot capabilities, 15% of cell-types, along with their GRNs, are held out as the unseen-graph validation set. The cells of the remaining 85% are randomly split into the seen-graph validation set and the training set. The seen-graph validation set accounts for approximately 15% of total cells, leaving the last 70% for the training set.

Implementation Details The model is implemented using PyTorch (Paszke et al., 2019) and Flashattention (Dao et al., 2022). Optimization are performed with the AdamW optimizer in PyTorch for both pre-training and validation models. Validation-based early stopping are implemented for all downstream regressors/classifiers that takes in frozen embeddings. The architecture of downstream models for all validation experiments are identical across the different foundation models within the same task.

A.3 Chebyshev Approximation of Diffused Query Embeddings

We first compute the Chebyshev coefficients c_k. Since our Laplacian matrix is normalized, the range of eigen values is $[0, 2]$. Since Chebyshev polynomials are orthogonal on $[-1, 1]$, we first map to this domain by setting

$$x = \frac{2\lambda}{\lambda^*} - 1$$

where λ^* is the max eigenvalue, which means $\lambda^* = 2$. Solving for λ yields:

$$\lambda = \frac{\lambda^*}{2}(x+1) = (x+1)$$

We can then compute the Chebyshev coefficients numerically with Gauss-Chebyshev quadrature. The Chebyshev coefficients are given by:

$$c_k = \frac{2}{\pi}\int_{-1}^{1} \frac{f(\lambda)T_k(x)}{\sqrt{1-x^2}}dx \tag{9}$$

$$= \frac{2}{\pi}\int_{-1}^{1} \frac{\exp(-\beta\lambda)T_k(x)}{\sqrt{1-x^2}}dx \tag{10}$$

$$= \frac{2}{\pi}\int_{-1}^{1} \frac{\exp(-\beta(x+1))T_k(x)}{\sqrt{1-x^2}}dx \tag{11}$$

To compute this integral numerically, the non-angular quadrature points x_i can be converted to a function of angles θ as:

$$x_i = \cos\theta_i = \cos(\frac{\pi(i-\frac{1}{2})}{N})$$

where N are number of total quadrature points. Notice that $T_k(x) = T_k(\cos(\theta)) = \cos(k\theta)$. We differentiate both sides to obtain:

$$\frac{dx}{d\theta} = -\sin\theta \rightarrow dx = -\sin\theta d\theta$$

Denote $g(x) = \exp(-\beta(x+1))$, substitute this change-of-variables form into the original integral yields:

$$c_k = \frac{2}{\pi}\int_{-1}^{1} T_k(x)\frac{g(x)}{\sqrt{1-x^2}}dx \tag{12}$$

$$= \frac{2}{\pi}\int_{0}^{\pi} \cos(k\theta)\frac{g(\cos\theta)}{\sin\theta} * \sin\theta d\theta \tag{13}$$

$$= \frac{2}{\pi}\int_{0}^{\pi} \cos(k\theta)g(\cos\theta)d\theta \tag{14}$$

Notice that in (5) the integral range after variable change becomes $[\pi, 0]$ and we flip the range because of the negative sign from the change-of-variables form. The resulting integral after change-of-variables can be approximated as:

$$c_k \approx \frac{2}{N}\sum_{i=1}^{N} g(\cos\theta_i)\cos(k\theta_i)$$

which can be efficiently computed. The Chebyshev polynomial terms and the subsequent transform of the query embedding Q is given by:

$$T_0(L)Q = Q$$
$$T_1(L)Q = LQ$$
$$T_k(L)Q = 2LT_{k-1}(L)Q - T_{k-2}(L)Q$$

A.4 Runtime Analysis

Here we provide run time analysis of both our pre-processing steps and our attention computation.

GRN inference Network inference imposes a modest and effectively fixed cost, even when considering hundreds to thousands of samples. Importantly, since network generation is a post-clustering step and the ARACNe algorithm requires tissue specific consistency, there is no need to use more than a small subset of sample - typically 500 profiles - regardless of cluster or dataset size. As there are a limited number of stable and transitory biological states physiologically possible, the computational cost will scale far below linear when approaching billions of cells, and we will update the text to clarify this critical point. Additionally, network generation is a pre-processing step so it's performed once all together before model training.

Attention computation Our model's runtime is bounded by the runtime complexity of a vanilla attention model. Assume we have G tokens, hidden size d, N quadrature points, and K terms of Chebyshev recursion, the runtime for performing the operation of estimating Gram matrix and transforming the query is given by

$$\mathcal{O}(NK) + \mathcal{O}(K \cdot G\delta \cdot d)$$

where δ is the average degree. The first term is the runtime of computing the chebyshev coefficients(see appendix A.3), while the second part is the runtime of computing the sparse matrix multiplication of transforming query embeddings. Since biological networks are very sparse in general, we have $\delta, K \ll G$. Hence, our attention computation is asymptotically bounded by the runtime of a regular dot-product attention model with runtime $\mathcal{O}(G^2 d)$

A.5 Hyperparameter Selection and Training Details

Model hyperparameters are outlined in Table 4. The model was pre-trained for 4 full epoch on 2 Nvidia H100 GPUs in parallel. For the diffusion rate, larger betas lead to models capturing global patterns (longer diffusions), while smaller betas lead to local patterns (shorter diffusions). A good default value is $\frac{1}{\lambda_{min}}$ where λ_{min} is the minimum non-zero eigenvalue of the Laplacian matrix. For the truncation order K, a good default value is close to the rank of the Laplacian matrix. In our specific case, we selected parameter values by grid-search around the default values.

Table 4: All hyperparameters and their values used to train the model are listed below. The tested model contains 14.7 million parameters.

Hyperparameters	Value
Transformer model dimension	512
Feed forward dimension	2048
Number of Transformer blocks	3
Number of attention heads per block	8
Activation	GeLU
Learning rate	5e-4
Batch size	256
Chebyshev truncation index(K)	8
Diffusion rate(β)	0.1
Number of quadrature points(N)	100
Graph integration false edge threshold(α)	0.05

A.6 Computational Costs

Pretraining To streamline pretraining, the fixed kernel gram matrices are precomputed during data preprocessing. Thus, the kernel gram matrix ϕ_L is loaded per sample and is used to transform the learned query Q, $\phi_L(Q)$. This leads to significant pretraining efficiency gains. Ultimately, GREmLN finishes 1 epoch in 3 hours versus 1.5 hours for a same-sized un-diffused transformer model (excluding validation steps) on a single H100 GPU. Pre-computing the kernel gram matrices also allows for efficient scaling as the same gram matrix is used in every transformer layer. Additionally, given that gene-filtering takes place for every sample, the use of the passthrough method (i.e. subsetting the full-graph KGM rather than building a KGM on the subgraph) allows the same computational resources to be used while benefiting from diffusion across the full graph structure.

A.7 Baseline Model Checkpoints

Here we list the exact pre-trained checkpoint for the models used in the experiments:

scGPT We used the 33M parameter model pre-trained on normal human cells, which is available at `https://github.com/bowang-lab/scGPT`

Geneformer We used the 30M parameter model with version number GF-6L-30M-i2048, which is available at `https://huggingface.co/ctheodoris/Geneformer`

scFoundation We use the 100M parameter model available at `https://github.com/biomap-research/scFoundation`

A.8 Details of Validation Datasets

Human Immune Cells This dataset includes cells from both human bone marrow and peripheral blood, covering 16 distinct immune cell types. We use the version provided by scGPT (`https://figshare.com/ndownloader/files/25717328`), with original cell type annotations from Luecken et al. (2021). This dataset serves as our primary benchmark for evaluating the model's ability to capture gene regulatory graph structure and to learn meaningful cell embeddings. We apply the same preprocessing steps as those used for the CELLxGENE corpus for model pre-training.

Cancer Infiltrating Myeloid This dataset comprises tumor-infiltrating myeloid cells (TMIs) from nine different human cancer types. TMIs are key regulators of tumor progression, with properties and functions that vary across cancer types. Similar to scGPT, we designate TMIs from six cancer types as a reference set for training and use the remaining types for out-of-distribution inference (query set). Pre-processing steps are identical to those applied to the CELLxGENE corpus. The dataset is available at: `https://drive.google.com/drive/folders/1VbpApQufZq8efFGakW3y8QDDpY9MBoDS`

Held Out Non-Immune Cells This dataset consists of 10 non-immune cell types that our model does not see during pre-training and validation, sourced from the CELLxGENE corpus (available at: `https://cellxgene.cziscience.com/collections`) with the same pre-processing steps as the training dataset. The selected cell types are retinal rod, epithelial, perivascular, amacrine, respiratory basal, smooth muscle, neural progenitor, secretory, myofibroblast, and megakaryocyte, testing the model's ability to perform zero-shot annotation across distinct non-immune cell types.

Adamson et. al. Perturb-seq This dataset consists of cells from the K562 leukemia cell line perturbed by Perturb-seq. In total, there are 87 unique single-gene perturbations. Since many genetic perturbations have an almost undetectable impact on the transcriptional states of cells, we select the 20 perturbations that produce the greatest distributional shift in transcriptional states relative to a control cell population, as measured by centroid distance in PCA space. The resulting pared-down dataset used for reverse perturbation modeling thus contains 9,416 cells across 20 perturbation conditions. The dataset is publicly available at: `https://dataverse.harvard.edu/api/access/datafile/6154417`

Norman et. al. Perturb-seq This dataset consists of cells from the K562 leukemia cell line perturbed by CRISPRa Perturb-seq. In total, there are 105 unique single-gene perturbations and 131 unique dual-gene perturbations. Since many genetic perturbations have an almost undetectable impact on the transcriptional states of cells, we select the 20 perturbations that produce the greatest distributional shift in transcriptional states relative to a control cell population, as measured by centroid distance in PCA space. The resulting pared-down dataset used for reverse perturbation modeling thus contains 21,999 cells across 20 perturbation conditions. The dataset is publicly available at: `https://dataverse.harvard.edu/api/access/datafile/6154020`

A.9 Linear Probing of Gene Expression

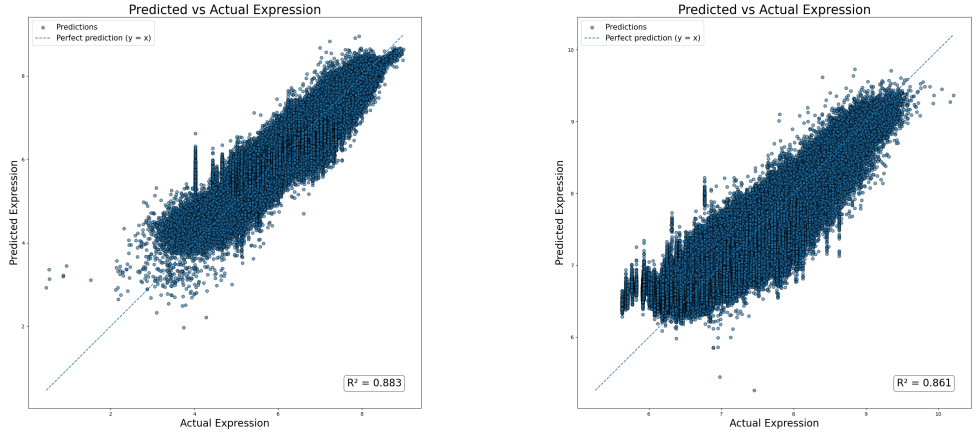

(a) **Human Immune Cells,** $R^2 = 0.883$ (b) **Cancer Infiltrating Myeloid,** $R^2 = 0.861$

Figure 2: **Reconstructing gene expression.** Shown are predicted and ground-truth gene expression from the learned GREmLN embedding. A linear regressor was trained on the GREmLN embedding to predict gene expression.

A.10 PPI Network

We obtained our PPI network from HINT(High Quality Interactomes), which is avaliable at: `https://hint.yulab.org/`. HINT is a curated aggregate of PPIs from eight sources (BioGRID, MINT, iRefWeb, DIP, IntAct, HPRD, MIPS, PDB). It applies automated and manual QC to remove low-quality/erroneous interactions, updates nightly, and supports both interactive queries and bulk downloads. We also tried pre-train with PPI networks from STRING-DB(`https:`

//string-db.org/) with high edge confidence cutoffs. However, we observe empirically that HINT networks significantly outperforms STRING-DB networks. It is outside the scope of this work to compare and investigate PPI network qualities.