Understanding Transcriptional Regulatory Redundancy by Learnable Global Subset Perturbations

Junhao Liu[1] Siwei Xu[1] Dylan Riffle[3] Ziheng Duan[1]
Martin Renqiang Min[2] Jing Zhang[1]

[1]University of California, Irvine [2]NEC Laboratories America [3]Cornell University

Introduction to GRIDS

Transcriptional regulation through cis-regulatory elements (CREs) is crucial for numerous biological functions, and its disruption can potentially lead to various diseases. It is well known that these CREs often exhibit redundancy, enabling them to compensate for one another in response to external disturbances. This underscores the need for methods to identify CRE sets that collaboratively regulate gene expression effectively. To address this:

- We present GRIDS, a computational method framing CRE dissection as a global feature explanation task.
- GRIDS first builds a differentiable surrogate function to approximate gene regulation and enable single-cell modality translation.
- It then uses learnable perturbations in a state transition framework to provide global explanations, efficiently exploring the feature landscape.

Preliminary of Single-Cell Data

The CRE is represented by the ATAC-seq binary vector $\mathbf{x} \in \{0,1\}^{d_a}$, where each dimension indicates a chromosome peak's state ("1" for open, "0" for closed). Typically, $d_a > 10^5$. Gene expression (RNA-seq) regulated by the CRE is denoted as $\mathbf{y} \in \mathbb{R}^{d_r}$, where d_a and d_r represent the number of peaks and genes, respectively. A single-cell multi-omics dataset consists of N cells $\mathcal{C} = \{\boldsymbol{c}^{(1)}, \boldsymbol{c}^{(2)}, \ldots, \boldsymbol{c}^{(N)}\}$, with each cell $\boldsymbol{c}^{(i)} = (\mathbf{x}^{(i)}, \mathbf{y}^{(i)})$ containing an ATAC-seq vector and its corresponding RNA-seq vector. Each cell also has a label $\ell^{(i)} \in \{1, \ldots, T\}$ indicating its type among T classes.

Regulatory Redundancy Problem

Gene expression is regulated by CREs through complex biological processes, modeled as $\mathbf{y} = \mathcal{F}(\mathbf{x})$, where $\mathcal{F}(\mathbf{x}) : \mathbb{R}^{d_a} \rightarrow \mathbb{R}^{d_r}$. Due to high experimental costs, frequent queries of the black-box function $\mathcal{F}$ are challenging. Regulatory redundancy dissection seeks a subset of L peak indices $\boldsymbol{r} = \{r_1, \ldots, r_L\}$ within the CRE (i.e., features in ATAC-seq $\mathbf{x}_{\boldsymbol{r}} \equiv \{\mathbf{x}_j | j \in \boldsymbol{r}\}$) that are critical for regulating gene expression across a cell population.

Global Feature Explanations for Regulatory Redundancy Dissection

We propose an in silico computational method by modeling it within a global feature explanation framework. Conventionally, global explanation is defined by how much a model's performance degrades over an observed population of samples when features are removed. In the context of regulatory redundancy, the global explanation objective can be expressed as

$$\boldsymbol{r}^* = \underset{\boldsymbol{r}}{\operatorname{argmin}}\, \mathbb{E}_{\mathbf{c}\sim\mathcal{C}}[\mathcal{L}(\mathcal{F}(\mathbf{x}_{\setminus \boldsymbol{r}}), \mathbf{y})]$$

where $\mathcal{L}$ is a loss measurement for expected gene expression degradation. $\mathbf{x}_{\setminus \boldsymbol{r}}$ denotes the perturbed CREs induced by $\boldsymbol{r}$, replacing the original feature $\mathbf{x}_{\boldsymbol{r}}$ with preset perturbation values $\mathbf{p} \in \mathbb{R}^{d_a}$.

Subset Transition Matrix for Gradient Estimation

- gradient w.r.t the input embedding

$$\mathbf{G} = \partial \mathbb{E}_{\mathbf{c}\sim\mathcal{C}}[\mathcal{L}(\hat{\mathcal{F}}(\mathbf{x}_{\setminus \boldsymbol{r}}), \mathbf{y})] / \partial \mathbf{W}^a_{\text{Emb}}(\mathbf{x}_{\setminus \boldsymbol{r}})$$

- construct the transition matrix $\mathbf{T} \in \mathbb{R}^{L \times d_a}$ via first-order approximation
- $\mathbf{T}_{i,j}$ represents the advantage of replacing index r_i with j

$$\mathbf{d}_j = \mathbf{G}_j \cdot (\mathbf{W}^a_{\text{Emb}}(\mathbf{p})_j - \mathbf{W}^a_{\text{Emb}}(\mathbf{x})_j)$$
$$\mathbf{T}_{i,j} = \mathbb{1}[j \notin \boldsymbol{r}]\mathbf{d}_j - \mathbb{1}[j \neq r_i]\mathbf{d}_{r_i}$$

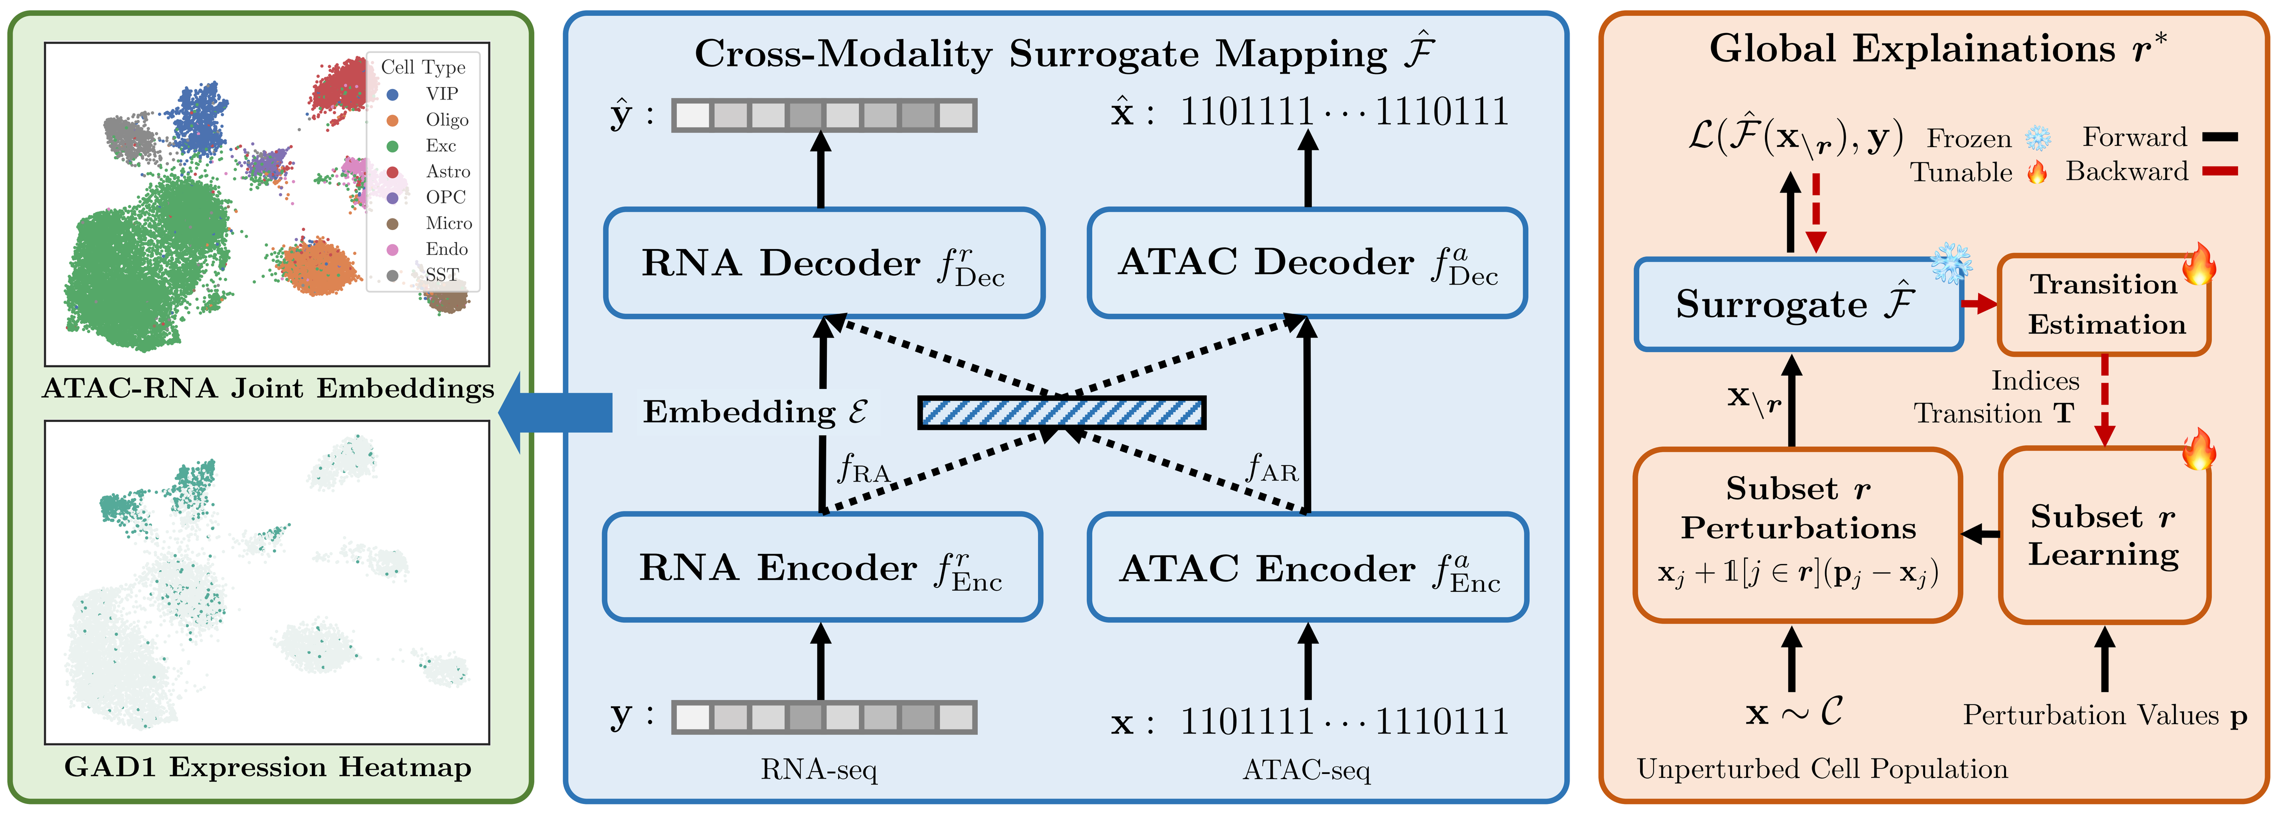

Figure 1. Overview of our proposed GRIDS method. It comprises two steps: training a cross-modality surrogate model and using a global explanation method to dissect regulatory redundancy.

Demo on MNIST

To evaluate GRIDS's global feature importance estimation, we tested its ability to identify key features in MNIST images. A binary classification model was trained to distinguish digits 8 and 3, achieving 97.9% accuracy on the test set. Various explanation methods were then used to identify the top $L = 64$ important pixels, masking them to zero ($\mathbf{p} = \mathbf{0}$). All methods produced subsets of L pixels based on their importance scores.

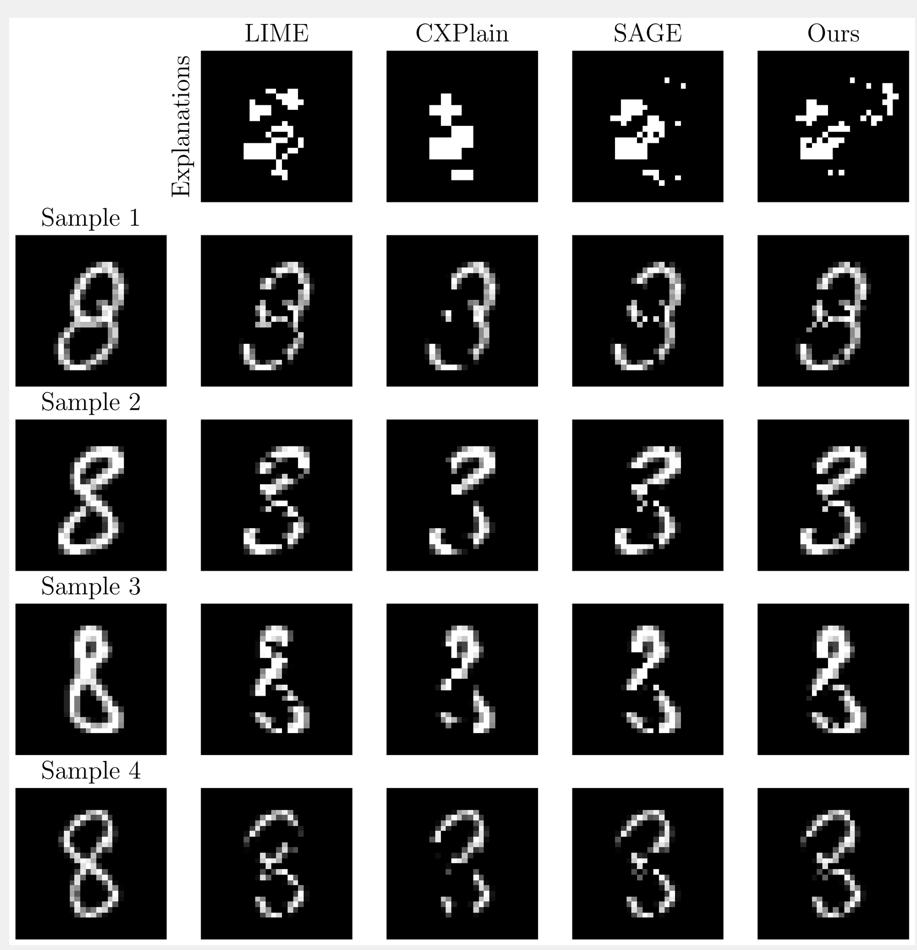

Figure 2. An examination of the most significant $L = 64$ pixels identified by various methods. Our subset perturbation learning method can find a similar combinatorial pattern as SAGE.

Experiment on Brain Data

We curated a set of deeply-sequenced single-cell multimodal data from postmortem human. We then evaluated the performance of GRIDS to dissect multi-CRE-to-gene regulatory redundancy by generating global feature importance explanations in the high-throughput single-cell multi-omics data. The global explanations were learned in the training set and then evaluated its performance on the test set.

Cell Type	Random Avg. Δ	Random Rel. Δ(%)	Saliency Avg. Δ	Saliency Rel. Δ(%)	SmoothGrad Avg. Δ	SmoothGrad Rel. Δ(%)	FIMAP Avg. Δ	FIMAP Rel. Δ(%)	GRIDS Avg. Δ	GRIDS Rel. Δ(%)
Astro	−0.085	−0.015	−2.163	−0.601	−2.155	−0.621	−13.502	−4.254	**−16.696**	**−5.837**
Endo	−1.073	−0.138	−4.974	−0.372	−9.726	−0.995	−38.997	−9.303	**−57.477**	**−11.816**
Micro	−0.012	−0.026	−23.757	−1.545	−32.944	−2.083	−73.752	−6.248	**−90.607**	**−7.671**
OPC	+0.823	−0.087	−54.645	−2.338	−48.438	−2.067	−77.167	−6.260	**−96.661**	**−8.256**
Oligo	−0.058	+0.026	−0.558	−0.173	−0.939	−0.220	−10.917	−4.252	**−16.760**	**−6.896**
SST	+0.159	+0.080	−5.201	−2.006	−5.201	−2.006	−16.453	−5.660	**−17.677**	**−6.365**
VIP	+0.012	+0.001	−0.654	−1.189	−0.634	−1.160	−2.732	−3.797	**−6.804**	**−7.195**
Avg.	+0.016	−0.021	−12.988	−1.209	−13.519	−1.290	−30.268	−5.367	**−39.103**	**−7.300**
Astro	−1.793	−0.533	−15.511	−4.853	−18.505	−6.217	−82.565	−24.766	**−100.556**	**−34.633**
Endo	+2.554	+0.468	−46.160	−6.217	−52.383	−7.893	−252.338	−41.790	**−259.920**	**−44.601**
Micro	−9.091	−0.490	−131.512	−9.122	−145.561	−10.116	−451.210	−39.695	**−470.430**	**−44.114**
OPC	−1.848	−0.165	−193.739	−10.260	−186.235	−9.891	**−415.231**	−35.687	−392.326	**−36.380**
Oligo	−1.134	−0.211	−19.809	−6.382	−21.136	−7.630	−69.460	−28.175	**−93.518**	**−38.982**
SST	−1.681	−0.615	−33.589	−11.675	−32.275	−11.115	−86.191	−29.198	**−93.772**	**−33.708**
VIP	+0.071	+0.002	−4.014	−4.876	−3.872	−4.782	−13.054	−16.757	**−19.703**	**−27.221**
Avg.	−1.843	−0.237	−68.620	−7.618	−70.292	−8.212	−202.368	−30.787	**−209.583**	**−36.893**

Table 1. Gene-focused benchmark results by comparing expression drops of marker genes across all cell types (upper: $L = 10$, bottom: $L = 128$).

Experiment on Brain Data (Continued)

Type	L	Method	Avg. Δ	Rel. Δ (%)
VIP-100	10	Random	−0.448	−0.009
		Saliency	−18.822	−0.915
		SmoothGrad	−18.424	−0.927
		FIMAP	−56.469	−3.087
		GRIDS	**−64.016**	**−3.827**
Microglia-100	10	Random	−0.333	−0.008
		Saliency	−42.372	−1.941
		SmoothGrad	−44.125	−2.073
		FIMAP	−115.092	−5.863
		GRIDS	**−141.339**	**−7.466**

Table 2. Cell-type-focused benchmark results in VIP and Microglia by comparing expression degradation of highly expressed genes after masking CRE features in the global explanation subset $\boldsymbol{r}$.

Method	$L = 10$ HR ↑	$L = 10$ SHR ↑	$L = 128$ HR ↑	$L = 128$ SHR ↑
Saliency	0.00	0.00	6.25	18.75
SmoothGrad	0.00	0.00	0.00	18.75
FIMAP	12.50	12.50	18.75	56.25
GRIDS	**18.75**	**25.00**	**31.25**	**68.75**

Table 3. The hit ratio of direct CRE-to-gene interactions.

Highlights

- We extend feature explanation techniques to scientific discovery on single-cell data.
- Through comprehensive benchmarking, GRIDS demonstrates superior explanatory capabilities compared to other leading methods.
- Moreover, GRIDS's global explanations reveal intricate regulatory redundancies across cell types and states, underscoring its potential to advance our understanding of cellular regulation in biological research.

Links

Paper GitHub Website

Acknowledgments

This work was supported by the National Institutes of Health [R01HG012572, R01NS128523].