sc-GRIP: a Graph Convolutional Approach to Infer Gene Interaction Polarity from Single-cell Data

1st Carolina E. Atria
Faculty of Computer Science
UniVie Doctoral School Computer Science
Department of Neuroscience and Developmental Biology
University of Vienna
Vienna, Austria
carolina.elisabeth.atria@univie.ac.at

2nd Yitao Cai
Faculty of Computer Science
UniVie Doctoral School Computer Science
Institute for Theoretical Biochemistry
University of Vienna
Vienna, Austria
yitao.cai@univie.ac.at

3rd Pascal Weber
Faculty of Computer Science
UniVie Doctoral School Computer Science
University of Vienna
Vienna, Austria
pascal.weber@univie.ac.at

4th Anna Beer
Faculty of Computer Science
University of Vienna
Vienna, Austria
anna.beer@univie.ac.at

5th Nils Kriege
Faculty of Computer Science
ds:UniVie
University of Vienna
Vienna, Austria
nils.kriege@univie.ac.at

6th Christian Böhm
Faculty of Computer Science
University of Vienna
Vienna, Austria
christian.boehm@univie.ac.at

7th Roger Revilla-i-Domingo
Department of Neuroscience and Developmental Biology
Research Platform Single Cell Regulation of Stem Cells
University of Vienna
Vienna, Austria
roger.revilla-i-domingo@univie.ac.at

8th Claudia Plant
Faculty of Computer Science
ds:UniVie
University of Vienna
Vienna, Austria
claudia.plant@univie.ac.at

***Abstract*—Understanding how genes are regulated is fundamental to many biological research questions. While experimental and computational methods allow us to identify which genes can interact with each other, finding the polarity of these interactions (whether activation or repression) is a non-trivial problem. We introduce sc-GRIP (Single-cell Gene Regulation Interaction Polarity), a graph convolutional framework that infers the directionality of transcription factor gene interactions directly from single-cell RNA data. By combining gene expression profiles with gene interaction graphs, sc-GRIP learns latent representations of genes and predicts regulatory polarity using a bilinear decoder. Our method enables scalable, cell-type-specific inference without relying on prior species-specific annotations or extensive biological validation. We demonstrate sc-GRIP's effectiveness on curated datasets from human and mouse, and show the advantages over other methods. sc-GRIP particularly excels in the usage on non-model organisms, and we provide a case study on a morphologically simple animal, *Suberites domuncula*, where sc-GRIP manages to expand our understanding of gene regulations. sc-GRIP offers a novel computational approach for reconstructing biologically interpretable regulatory networks, not only in well-studied organisms but especially in emerging organisms, where large-scale experimental setups are often unfeasible.**

***Index Terms*—Graph Convolutional Network, Edge classification, Interaction polarity, Single-cell transcriptomics, Gene regulatory networks**

I. Introduction

Every cell in a living organism performs specific roles that are critical to the development and survival of the organism. These specific roles depend largely on the set of proteins -large, structurally diverse molecules- that is expressed in the cell at any given time. Animal genomes, which store the complete genetic material in the form of DNA, contain thousands of genes, each encoding a specific protein. Which genes are active or inactive in a cell at a given time is controlled by gene regulation. Central to this process are transcription factors (TFs), specialized proteins that bind to specific DNA sequences and determine whether the target gene is switched ON (expressed) or OFF (not expressed). When a gene is ON, its expression occurs in two main stages. First, transcription produces messenger RNA (mRNA), an intermediate that carries the information from DNA. Second, during translation, the mRNA serves as a template to generate the protein encoded by the gene.

Unraveling the complex networks of transcription factor (TF) interactions that control gene activity across cells is a key challenge in the life sciences. Recent advances in single-cell technologies have produced datasets that combine gene expression information (mRNA abundance for all genes) and maps of DNA regions accessible to TF binding, measured

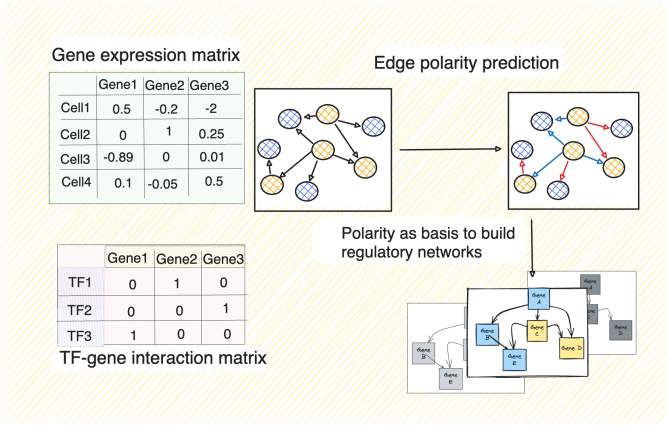

Fig. 1. Predicting edge polarity from an input adjacency matrix of known TF target gene interactions and a matrix of gene expression sets the stage for generating complex regulatory networks.

in thousands to millions of individual cells. Inferring regulatory networks from these data poses a major computational challenge, as the measurements are high-dimensional, sparse, and noisy. The tool presented here addresses a core aspect of this challenge: constructing the network of candidate TF–gene interactions and predicting the *polarity* of each interaction, i.e., whether it contributes to turning a gene ON or OFF.

Potential interactions between TFs and their target genes can be captured in regulatory graphs. While the graph structure can be directly determined from the data, inferring the weights and signs of the edges is challenging, as they are context-dependent: they can vary across cell or conditions.

Figure 1 gives an overview of our problem setting. Our input are high-dimensional, large gene expression matrices (top left) and the interaction matrices between TFs and their target genes (bottom left). As the latter lack information on signs and weights of the edges, our goal is to infer the edge polarity using the gene expression matrices. This builds the basis to the larger goal of constructing regulatory networks (bottom right) that are able to capture the molecular underlying processes of cell differentiation trajectories.

In contrast to traditional graph inference tasks, the signals in this context are entangled within high-dimensional manifolds and further obscured by variations in cell types. Furthermore, supervised labels for regulatory polarity are rare, incomplete, and biased toward well-studied genes. Thus, research in this field requires self-supervised or weakly supervised learning approaches that can generalize across unseen transcription factors and cell types.

In summary, to infer the gene interaction polarity, methods need to solve the following challenges:

1) model the latent regulatory structure,
2) infer edge labels, and
3) remain robust to biological noise and heterogeneity.

Prior work in regulatory inference has largely focused on binary interaction detection or correlation-based methods,

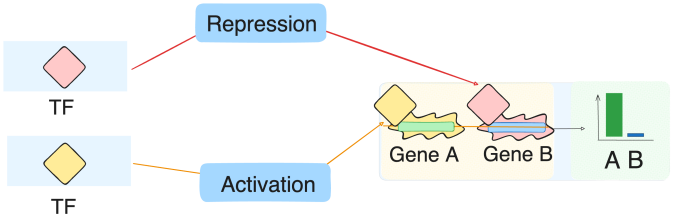

Fig. 2. Schematic overview of gene regulation via a transcription factors, showing their **activating** and **repressing** effects. Gene A is activated, while the expression of Gene B is repressed.

which lack the expressiveness needed to capture regulatory directionality at scale.

To tackle these challenges, we introduce **sc-GRIP (Single-cell Gene Regulation Interaction Polarity)**, a graph-based graph convolution framework for inferring the directionality (activation or repression) of transcription factor gene interactions from single-cell RNA sequencing data. To the best of our knowledge, our method is the first one to reconstruct gene regulatory relationships from gene expression features. SC-GRIP uses a graph-based autoencoder framework designed to learn cell-specific embeddings of transcription factor (TF)–gene interactions. It employs a two-layer Graph Convolutional Network (GCN) as the encoder and a bilinear decoder to classify interactions as activating or repressing. During training, the model leverages correlation patterns between TF and target gene expression across observations from single-cell data to generate weak supervision signals. sc-GRIP produces robust edge-level polarity scores reflecting the regulation types between TF and target gene.

Our experiments in Section IV show that sc-GRIP outperforms traditional polarity prediction approaches measured by ROC/AUC and precision on model-species where we can compare against the ground truth. As sc-GRIP works without prior knowledge that competitive approaches like WGCNA [1] and simiC [2] require, we can apply it to discover novel knowledge about non-model species, which we demonstrate on single-cell data of the sponge *Suberites domuncula*. The implementation of sc-GRIP is available at: *https://github.com/caroatria/sc-grip*

II. Related Work

Our work combines multiple research areas: biology, data mining, and machine learning. We give some biological background knowledge on transcription factors (TFs) in Section II-A, followed by advances in bioinformatics in Section II-B and related work in data mining and machine learning in Section II-C. Table I provides a brief overview of some relevant biological terms and a short explanation for ease of understanding.

A. Transcription factors

Transcription factors (TFs) can bind to specific sequences in the DNA and regulate the expression of target genes as Figure 2 illustrates. According to its interaction polarity, sc-GRIP

TABLE I. BIOLOGICAL TERMS

Term	Meaning
Genome	DNA molecules that constitute the complete set of genetic material of the organism and the instructions that determine the expression of each gene.
mRNA	Messenger molecule that carries sequence information of a specific gene and serves as template for the production of the corresponding protein.
Protein	Large, structurally diverse molecules that determine the capabilities, functions, and identity of a cell.
TF	Transcription factor: a protein that can regulate the expression of genes by binding to specific DNA sequences in the genome.
Gene expression	Process by which the protein encoded by a gene is expressed.
Single-cell mRNA-sequencing	Experimental method where mRNA abundances for all genes are measured in individual cells, resulting in a gene-by-cell matrix.
Cell type	Subset of cells with a shared function and set of active genes (e.g., stem cells, skin cells).
Stem cells	Cell type with the capacity to: (1) proliferate, giving rise to more cells of the same type; and (2) differentiate into other cell types.
Differentiation	Process by which a stem cell adopts the identity of another cell type.
Terminal cell type	Cells with a well-defined identity and function, that under normal conditions of the organism will not differentiate into other cell types.

models the function of each TF to be activating (indicating a net contribution in increasing the expression of the target gene) or repressive (indicating a net contribution in decreasing the expression of the target gene. Generally, genes are regulated by a combination of multiple TFs. Figure 2 illustrates the difference in gene regulation with activating (yellow) and repressing (red) TF-interactions.

B. Inferring TF activity in bioinformatics

Inferring gene regulatory and interaction networks from transcriptomic data is a long-standing challenge in biology and is rooted in a constant struggle between exhaustive laboratory experiments and computational predictions. Traditional bioinformatical approaches have leveraged correlation-based methods such as Weighted Gene Co-expression Network Analysis (WGCNA) [1], which cluster genes based on co-expression patterns but lack directionality and fine-grained prediction of regulatory influence. While tools like this are able to infer co-expression, they cannot predict features of gene regulatory interactions. Furthermore, many publications focus on predicting interaction polarity based on prior results from their research organism [3, 4]. While these methods have shown great potential, they are harshly limited in their application to a small list of model species [5, 6].

However, most of the single-cell RNA-sequencing analysis models are designed for use on well-established model organisms, such as humans, mice, and fruit flies. In less-studied non-model species, the single-cell RNA-sequencing datasets are left understudied [7].

In addition, assessing the quality of computational predictions often falls to laboratory experiments, which are time-consuming and associated with high costs per experiment.

In the following, we discuss four of the most commonly used methods for gene regulatory network inference: WGCNA [1], SCENIC [8], SimiC [2], and DoRothEA [4].

1) WGCNA: WGCNA [1] was one of the first approaches to construct gene regulatory interactions computationally, originally designed for bulk analyses (the precursor method of single-cell analyses, where a tissue sample is analysed as a whole). While the WGCNAs approach has definitive advantages, its heavy reliance on correlation between the gene expression data often fails to capture more complex, intricate interactions. Additionally, there is a large number of parameters involved in a WGCNA analysis, which require extensive user knowledge of the organism in question.

2) SCENIC: SCENIC [8] employs a three-step pipeline to infer Gene Regulatory Networks (gene regulatory networks). Initially, a random forest model predicts gene expression levels using a set of driver genes, deriving transcription factor weights from the trees based on their importance across predictions. This approach effectively captures non-linear relationships within the data. In the second step, overrepresented DNA motifs near the beginnings of genes are identified, mapping them to potential transcription factors using a motif-TF association database. Finally, the activity scores are computed at the single-cell level. While SCENIC has been shown to be an efficient tool, it cannot be run without extensive knowledge of the organisms' genes and regulatory landscape. Additionally, SCENIC is not able to produce weighted interactions, since it is primarily employed for model organisms, where this information has already been uncovered.

3) SimiC: In contrast to global approaches, SimiC [2] clusters cells into phenotypes and constructs individual gene regulatory networks for each phenotype. Each group's gene regulatory network is generated by a linear regression model using cell-specific driver genes to predict target gene expressions. Unlike SINCERITIES, which models temporal changes, SimiC focuses on phenotype-specific network inference. Lasso Regularization manages data dimensionality, while an additional regularization term ensures consistency in model weights across different cell phenotypes.

4) DoRothEA: The DoRothEA [4] approach was originally developed for the specific application in human data, and only later expanded to include mouse. Similarly to SCENIC, DoRothEA aims to compute gene regulatory networks with signed interactions between TF and target genes. Since this tool is only suitable for use with a mouse and humans, the approach depends on curated databases. Additionally, DoRothEA requires information on accessibility as well as gene expression. The accessibility of DNA can be used to measure which parts of the DNA sequence are physically reachable to be expressed.

Table II provides an overview of the presented related methods and their features concerning the prediction of gene interaction polarity. As, to the best of our knowledge, no published approach is aimed at the same input data as our

novel method sc-GRIP and the most relevant methods, that we described above, are not able to directly compute the polarity, we will instead compare our approach with variations of their presented related works, namely Pearson correlation and lasso regression edge weight prediction. This comparison also aims to show why a more nuanced machine learning approach is necessary to capture the complex gene regulatory interactions. The results of this comparison are shown in Figure 5 in Section IV.

TABLE II. FEATURES OF METHODS FOR GENE REGULATORY NETWORK INFERENCE.

Method	non-model species	no prior knowledge	edge polarities	single-cell samples
WGCNA [1]	♦	✓	✗	✗
SCENIC [8]	✗	✗	✗	✓
simiC [2]	♦	✓	✓	♦
DoRothEA [4]	✗	✗	✓	✓
sc-GRIP	✓	✓	✓	✓

✓ yes ✗ no ♦ potentially

C. Graph machine learning

With the rise of single-cell RNA sequencing, deep learning-based models have gained traction due to their ability to extract hierarchical and non-linear features. Autoencoders and variational Autoencoders (VAEs) have been applied to denoise expression data or extract latent representations [9], but few extend this to interaction prediction.

Graph-based neural networks, especially Graph Convolutional Networks (GCNs) [10, 11], provide a natural framework to model gene interactions. Graph Autoencoders (GAEs) have been used in various domain applications to learn node embeddings and infer edge presence [12]. In the context of gene networks, applications of GAEs are still emerging, but have been shown to be promising models for expression dynamics [13].

Recent work by Minici et al. addresses the problem of link polarity prediction in signed graphs under conditions of sparsity and noise. Their method leverages multiscale notions of social balance theory to enhance robustness, combining local and global structures through a principled optimization framework. Although developed in the context of social networks, their formulation offers valuable insight into handling uncertain or incomplete edge annotations—challenges that are similarly prevalent in gene regulatory inference, where true interactions are only partially observed and context-dependent. Their approach underscores the importance of structure-aware learning in graphs with noisy or indirect supervision.

Our approach builds upon this direction by integrating a bilinear GAE with correlation-derived labels, trained on single-cell expression graphs to predict probabilistic gene-gene interaction scores. This architecture allows for cell-level interaction inference, edge-specific confidence scores, and cluster-aware consensus prediction, offering a scalable and biologically interpretable framework for regulatory network reconstruction.

TABLE III. NOTATION AND ABBREVIATIONS

Notation	Meaning
$d \in \mathbb{N}$	dimensionality of the original feature space
$m \in \mathbb{N}$	dimensionality of the embedded space
$k \in \mathbb{N}$	number of clusters
$n \in \mathbb{N}$	number of points in the data set
$\mathscr{D} \subseteq \mathbb{R}^d$	set of all objects
$\mathscr{C}$	set of cluster labels
G	a gene regulatory graph
E	edge set
$A \in \mathbb{R}^{N \times N}$	adjacency matrix
$F \in \mathbb{R}$	feature dimension (in: input, out: output, h: hidden)
X	node feature space
H	hidden space
Z	latent space
$u \in \mathbb{R}$	source node(transcription factor)
$v \in \mathbb{R}$	target node(regulated gene)
$y \in \mathbb{R}$	label of a directed edge
$\tau \in \mathbb{R}$	threshold for binarizing correlation

III. METHODOLOGY

In this section, we describe the sc-GRIP approach in more detail. For this, we explain our objective in Section III-A and under which hypothesis we solved this objective in Section III-B. In Section III-C, we introduce our model architecture, and in Section III-D, we cover basic pre-processing steps. Finally, in Section III-E, we cover sc-GRIP-specific data preparation steps. The used terms are briefly explained in Table III.

A. Objective

The aim of the presented sc-GRIP model is to create an approach with the capability to predict the polarity of TF-gene regulation with limited input datasets. These datasets are a gene expression matrix (such as from a single-cell gene expression analysis) and an adjacency matrix that maps regulating interactions. Not relying on any prior knowledge from databases or experimental perturbations, sc-GRIP emerges as a polarity prediction approach that is especially suitable for non-model organisms.

B. Hypothesis

The hypothesis at the heart of our approach relies on the idea that a regulating TF must first be produced in order to affect (activate or repress) any target genes. Additionally, this regulation is limited by the physical production of the regulating TF. Additionally, in an activating relationship, the (normalised) gene levels of the regulating TF and its target gene would be at similar levels. Conversely, a regulating TF that strongly represses its targets would have a high TF level, while its target gene expression would be low.

Therefore, we hypothesise that the gene levels of a regulating TF and its target gene have a significant relationship

that must be taken into account when predicting the polarity of any such interaction. However, a simpler model, such as a straightforward correlation analysis, is not successful, due to the sparsity, noise, and complexity of gene regulation (see Section IV).

C. Architecture

To infer the polarity (activation or repression) of TF–gene interactions using single-cell RNA sequencing data, we suggest a graph-based autoencoder framework by learning a cell-specific representation of gene relationships. The sc-GRIP architecture is based on a two-layer Graph Convolutional Network (GCN) encoder and a bilinear decoder (gene interaction classifier). Figure 3 provides a schematic outline of sc-GRIP's approach.

Let $G = (V, E)$ be a directed, interaction graph, with V genes and E the potential regulatory edges, and a single-cell expression matrix $X \in \mathbb{R}^{|V|}$, with observations of each gene level in each cell. From the adjacency matrix, we can already infer the directionality of the interactions, as this matrix is an asymmetrical TF x target genes matrix. The TFs and target genes are both subsets of the full gene list, and interaction can also occur between different TFs. We represent the node features for cell c as $X_c \in \mathbb{R}^{N \times F_{in}}$, where $F_{in} = 1$ and $X_c[i]$ is the expression of gene i in cell c. Our goal is to predict a label $y_{uv} \in \{\text{activation, repression}\}$ for each directed edge $(u, v) \in E$. We construct the following elements: As sc-GRIP is designed for non-model species, where often no prior information is available, a supervised approach cannot be taken. Instead, we use proxy labels of interactions (y_{uv}), which are computed as the correlation of the expression levels. In this way, we want to make use of the gene expression levels as provided by our single-cell observations, while benefitting from the graph convolutional network that allows us to deal with more complex interactions.

Then the proxy label is:

$$y_{c,uv} = \mathbb{I}(cor_{c,uv} > \tau_{cor}) \tag{1}$$

where $\mathbb{I}(\cdot)$ is the indicator function and τ_{cor} is a correlation threshold (e.g., 0.5 in our case, or find an optimal threshold).

$y_{c,uv} = 1$ suggests that genes u and v have "correlated" expression behavior in cell c, which we take as a proxy for an activation interaction. $y_{c,uv} = 0$ suggests not correlated behavior, which we take as a proxy for a repression relationship.

The model aims to learn latent representations of genes (Z_c) that are predictive of these proxy labels.

Encoder (f_{enc}): Maps the up-projection embedding of input features $\mathbf{X}_e \in \mathbb{R}^{F_e}$ and graph structure $\mathbf{A}$ (adjacency matrix) to latent embeddings $\mathbf{Z}_c \in \mathbb{R}^{N \times F_{latent}}$.

$$\mathbf{X}_e = \mathbf{W}_e \mathbf{X}_c, \text{ where, } \mathbf{W}_e \in \mathbb{R}^{F_e \times F_{in}} \tag{2}$$

$$\mathbf{H}_c^{(1)} = \text{ReLU}(\text{GCNConv}(\mathbf{X}_c, \mathbf{A}; \mathbf{W}^{(1)})) \tag{3}$$

$$\mathbf{Z}_c = \text{GCNConv}(\mathbf{H}_c^{(1)}, \mathbf{A}; \mathbf{W}^{(2)}), \tag{4}$$

where $\mathbf{W}^{(1)} \in \mathbb{R}^{F_{in} \times F_h}$ and $\mathbf{W}^{(2)} \in \mathbb{R}^{F_h \times F_{out}}$ are learnable weight matrices. We assume local subgraphs have a more important influence on the relation with each other, so we adopted 2 GCN layers for the encoder.

A GCN layer operation [10] is:

$$\text{GCNConv}(\mathbf{H}, \mathbf{A}; \mathbf{W}) = \hat{\mathbf{D}}^{-1/2} \hat{\mathbf{A}} \hat{\mathbf{D}}^{-1/2} \mathbf{H}\mathbf{W} \tag{5}$$

with $\hat{\mathbf{A}} = \mathbf{A} + \mathbf{I}_N$, where $\mathbf{A}$ is the adjacency matrix, $\mathbf{I}_N$ the identity matrix and $\mathbf{D}$ the degree matrix.

Decoder (f_{dec}): Reconstructs edge interaction probabilities from the latent embeddings. For an edge $(u, v) \in E$, let $\mathbf{z}_{c,u}$ and $\mathbf{z}_{c,v}$ be the latent vectors of $\mathbf{Z}_c$ corresponding to genes u and v. The decoder predicts a probability of correlation.

$$s_{c,uv} = \mathbf{z}_{c,u}^T \mathbf{W} \mathbf{z}_{c,v} \tag{6}$$

$$p_{c,uv} = \sigma(s_{c,uv}) = \frac{1}{1 + e^{-s_{c,uv}}} \tag{7}$$

where $W \in \mathbb{R}^{F_{out} \times F_{latent}}$ is learnable parameters of the bilinear layer and σ is the sigmoid function. The model parameters are learned by minimizing the expected Binary Cross-Entropy (BCE) loss between the predicted probabilities $p_{c,uv}$ and the proxy labels $y_{c,uv}$ across all cells or clusters and all edges in E.

$$L(\theta) = \mathbb{E}_{c \sim \text{Dataset}} [L_c(\theta)] \tag{8}$$

where the loss for a single cell or cluster c is:

$$L_c(\theta) = -\frac{1}{|E|} \sum_{(u,v) \in E} \Big[y_{c,uv} \log(p_{c,uv}(\theta)) + (1 - y_{c,uv}) \log(1 - p_{c,uv}(\theta)) \Big]. \tag{9}$$

Minimizing $L(\theta)$ drives the model to learn parameters θ such that: If $y_{c,uv} = 1$ (high expression correlation), then $p_{c,uv}(\theta) \to 1$. If $y_{c,uv} = 0$ low expression correlation, then $p_{c,uv}(\theta) \to 0$. By achieving this, the learned embeddings $\mathbf{Z}_c$ capture information about gene expression correlations within the graph context, which is our proxy for interactions. The GCN layers are crucial because they allow information to propagate across the graph, enabling $\mathbf{Z}_c$ to reflect not just individual gene expression but also the expression patterns in its local neighbourhood, conditioned by E.

- **Encoder:** The encoder maps the input expression features X into a latent representation Z through two graph convolutional layers with ReLU activation:

$$Z = \text{GCN}_2(\text{ReLU}(\text{GCN}_1(X, E)), E)$$

- **Decoder:** The decoder reconstructs edge interactions by computing a bilinear function between node embeddings:

$$\hat{y}_{(i,j)} = \sigma\left(Z_i^\top W Z_j\right)$$

where σ is the sigmoid activation function, and W is a trainable weight matrix. The model parameters aim to predict whether a TF–gene interaction is activating or repressing based on the expression correlation across single cells of a tissue cell type.

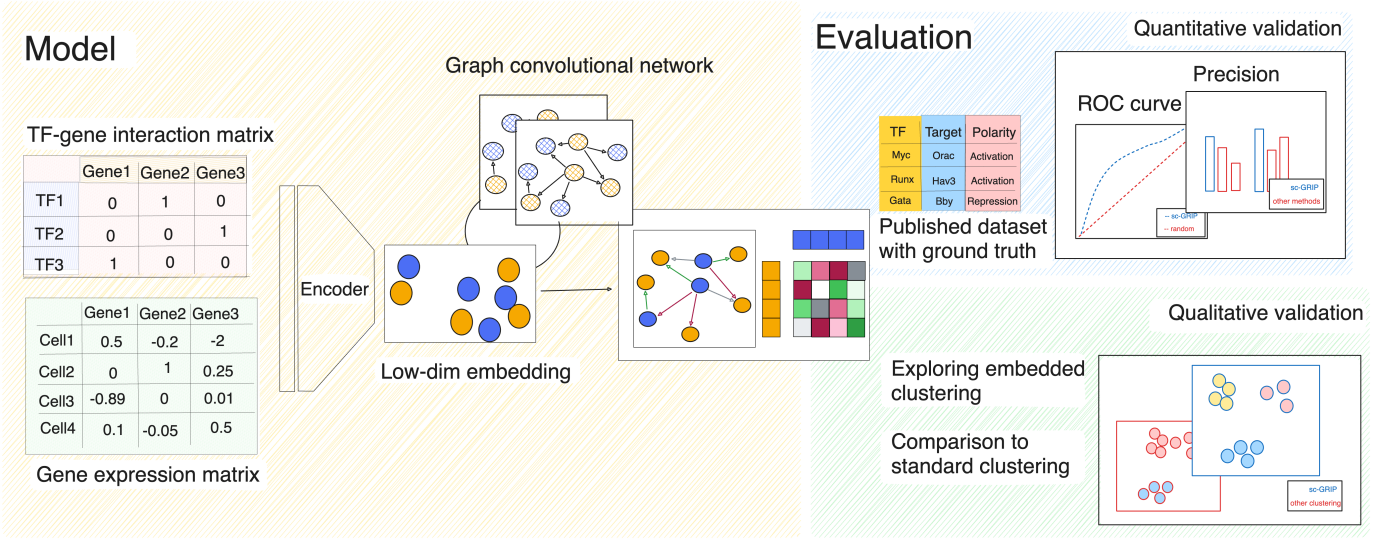

Fig. 3. Schematic overview of our method. The adjacency matrix and single-cell observations are both taken as input data, and the extracted graph is compressed to a lower-dimensional embedding. Following, a GCN-based encoder–decoder learns per-edge interaction probabilities, which are aggregated across cells to compute confidence scores.

D. Data Pre-processing

To apply sc-GRIP, users must provide two input matrices: (1) a gene expression matrix derived from single-cell RNA sequencing and (2) a TF–gene interaction adjacency matrix derived from motif scanning.

The gene expression matrix should be pre-processed using standard workflows such as those provided by Scanpy or Seurat [14, 15]. This includes filtering low-quality cells and genes, normalizing total counts per cell, applying a log-transformation, selecting highly variable genes, and reducing dimensionality with PCA. Cells should be clustered (e.g., using the Louvain or Leiden algorithm [16]) and annotated into cell types or phenotypes based on marker gene expression. Gene identifiers should be converted to standardized gene symbols (matching the adjacency matrix), and the final expression matrix, including cell and gene metadata, must be exported in the `.h5ad` format.

The adjacency matrix can be constructed either experimentally or by scanning gene promoter regions for known transcription factor binding motifs. Motif PWMs can be sourced from databases such as JASPAR or HOCOMOCO. Promoter sequences – typically defined as 2,000 base pairs upstream and 500 base pairs downstream of the transcription start site – should be scanned using FIMO (Find Individual Motif Occurrences) from the MEME suite [17]. For each TF, if a motif match is found in a gene's promoter region with a significance below a predefined threshold (e.g., $p < 1 \times 10^{-4}$), a regulatory edge is assigned. This results in a binary adjacency matrix where rows correspond to TFs, columns to target genes, and matrix entries indicate the presence or absence of a predicted regulatory interaction. These two inputs, (i) the processed expression data, and (ii) the motif-derived adjacency matrix, serve as the inputs for sc-GRIP, which then infers context-specific gene regulatory interactions across single cells.

E. sc-GRIP Data Preparation

To perform the analysis, two datasets must be available for processing. From the single-cell gene expression data, a matrix containing the gene expression levels is extracted. The adjacency matrix that records which transcription factors are able to influence which target genes is created by scanning the species' genome for known transcription factor binding sites. We process single-cell expression data and the interaction graph as follows:

- **Expression Matrix:** Each cell is represented by its expression levels across all genes and transcription factors.
- **Interaction Graph:** The edge list is constructed from motif-based predictions, indicating potential regulatory links.

For each single cell, we generate an individual graph data object with the fixed structure E but with cell-specific node features X. Algorithm 1 provides an overview of the sc-GRIP's model structure and workflow.

IV. Experiments

While the initial aim of the study was to implement a tool for polarity inference in non-model organisms, we can use model species to validate the approach. Therefore, we tested sc-GRIP on two large datasets for mouse and human, which have been collated from more than 10,000 publications into a manually curated network [18]. The networks include 828 and 800 transcription factors for the mouse and human datasets, respectively, and include information on the interaction polarity if available.

Additionally, to confirm the practicality of sc-GRIP for non-model species, one in-house dataset was tested, from an early-branching animal, the marine sponge (*Suberites domuncula*).

Algorithm 1 sc-GRIP with Correlation-Based Training

Input:
Cell expression dataset $\mathscr{D} = \{(X_i, E)\}_{i=1}^{N}$,
where X_i is the gene expression vector for cell i,
E is the fixed TF-target edge index
Correlation matrix $\mathscr{C}$ with values $\mathscr{C}_{tf,target}$
Number of epochs T, number of runs R, learning rate η
Correlation threshold $\theta = 0.5$

Initialize: GAE model parameters θ_{GAE}
for run $= 1$ to R **do**
  for epoch $= 1$ to T **do**
    for each cell graph (X_i, E) in $\mathscr{D}$ **do**
      Compute latent emb. $Z_i = \text{Encoder}(X_i, E)$
      Predict edge activations $\hat{Y}_i = \text{Decoder}(Z_i, E)$
      for each edge $e = (tf, target)$ **do**
        Retrieve label $Y_i[e] = \mathbb{1}\{\mathscr{C}_{tf,target} > \theta\}$
      Compute loss $\mathscr{L} = \text{BCELoss}(\hat{Y}_i, Y_i)$
      Update θ_{GAE} via backpropagation
  Store predicted edge activations $\hat{Y}^{(run)}$
Compute final activation scores as median over runs: $\hat{Y}_{\text{final}} = \text{median}(\{\hat{Y}^{(r)}\}_{r=1}^{R})$
Output: Final edge activation scores $\hat{Y}_{\text{final}}$ with mean and std over runs

To validate the polarity prediction, we perform sc-GRIP on one human single-cell dataset and observe the distribution of the predicted edge weights in comparison to the validated interaction types. These interaction types are ternary, each edge is classified as activation, repression or unknown interaction.

A. Validation datasets

Taking the experimentally validated interaction properties from the TRRUST database, we were able to quantitatively validate the polarity predictions. To evaluate the discriminative ability of sc-GRIP independent of a specific decision threshold, we computed the area under the ROC curve (AUC). On the real-world human blood cell dataset, the model achieved an AUC of 0.77, see Figure 4. Due to the small sample size, the ROC curve exhibits a stepwise shape, which is expected under these conditions.

To compare how sc-GRIP performs on datasets in comparison to the methods used in simiC, SCENIC, and WGCNA, the three mouse datasets were compared and analysed in Figure 5.

As can be seen in this comparison, sc-GRIP is able to outperform the LASSO and pearson correlation approaches, that make up the foundation of the related published models [1, 2, 8].

Table IV reflects the predictive performance with a focus on prediction precision against the published ground-truth labels. In all cases, sc-GRIP consistently outperforms the competing methods, demonstrating higher precision across datasets.

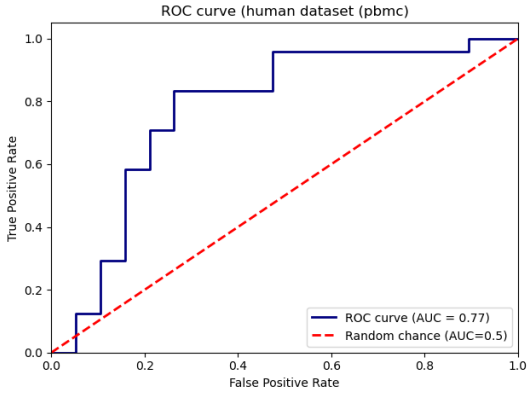

Fig. 4. Receiver operator curve for classification model on human blood cell dataset. sc-GRIP achieves an AUC of 0.77, showing its ability to distinguish between the two classes.

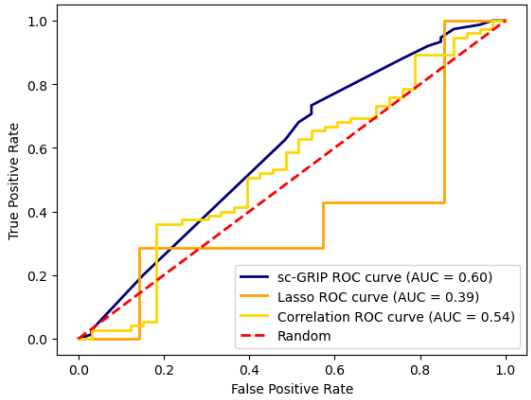

(a) Method comparison on mouse muscle dataset

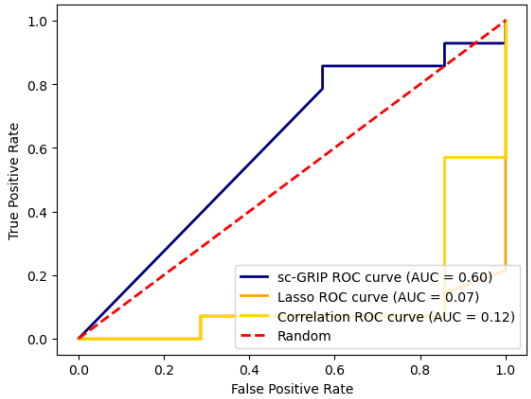

(b) Method comparison on human germ line dataset

Fig. 5. ROC/ AUC results for comparison of sc-GRIP to traditional polarity prediction approaches

B. Non-model species

To demonstrate the applicability in a non-model species, where experimental validations cannot be performed on a large scale, we used an in-house dataset (details omitted for double-blind review). The non-model organism in question is the marine sponge *Suberites domuncula*, which has been indicated as a novel species of interest to investigate the origins of

TABLE IV. PRECISION COMPARISON ACROSS DATASETS

Model	Mouse liver	Mouse muscle	Human germ line
sc-GRIP	1.00	0.65	0.65
LASSO	0.70	0.65	0.20
Correlation	0.70	0.60	0.17

animal stem cells [19]. Here we review a few examples of sc-GRIP predictions that illustrate the performance of our tool in this non-model species.

We first focused on the predicted targets of the stem cell-enriched TF Myc. While the function of Myc in sponge stem cells has not yet been tested, the function of this TF in mouse stem cells has been well studied. sc-GRIP predicts a positive regulation of a number of target genes in the sponge (Figure 6). A subset of these interactions has been corroborated by experiments at the collaborating lab. Additionally, comparison to experimentally-validated activating interactions of the mouse Myc TF in stem cells [20] shows a statistically very significant overlap (p-value $< 10e$-6). A group of genes that were previously manually analysed and predicted to be activated by Myc are also corroborated by the sc-GRIP's output. Figure 6 shows the network of Myc's interactions with target genes of interest.

sc-GRIP also predicts repressive interactions by the Myc TF. Most interesting among these is the gene encoding for the Smad4 TF. While the function of this TF has not been tested in the sponge, Smad4 has been shown to promote cell differentiation in the mouse [21]. This suggests a role of sponge Myc in inhibiting cell differentiation, consistent with its role in mouse stem cells [22].

In addition, a number of TFs known to play key roles in promoting cell differentiation in mouse stem cells (Pou1, Fox, Stat, Cebp), are predicted to have repressive interactions on *myc* (Pou and Fox; see Figure 6) or on some of the proliferation genes activated by Myc (e.g. *chek2*, *rps10*, *rpl27rt*, *rps25*, *rps27a*; Figure 6). This suggests a role of these TFs in antagonizing the function of Myc in sponge stem cells, inhibiting proliferation and promoting differentiation.

Consistently, the expression profile of these genes can be seen in the single-cell clustering of the *Suberites domuncula* dataset 7a. The colourmap shown in Subfigures 7b-7e corresponds to the gene expression levels across the clustered cells.

The expression profiles show that the genes are ON in cells that are differentiating (*stat*; Figure 7c) or fully differentiated (*pou1*, *fox*; Figure 7d), and they are OFF in stem cells (Figure 7e)).

These predictions can guide future experimental analyses and greatly facilitate the comparison of regulatory programs in different species.

V. APPLICATION OF SC-GRIP FOR NON-MODEL ORGANISMS

Most approaches in single-cell bioinformatics are developed for model organisms. These are species that have been used

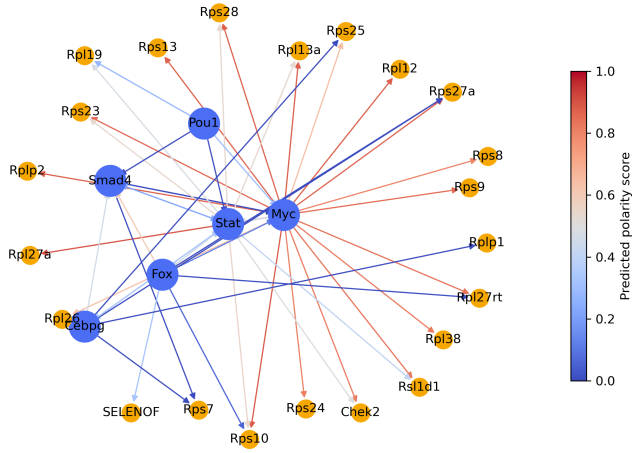

Fig. 6. Subset of predicted interactions by sc-GRIP in sponge dataset. TFs are depicted in blue, target genes in yellow, their interactions are coloured according to sc-GRIP's predicted polarity scores – ranging from 0 (repression) to 1 (activation).

by the scientific community over decades and have been studied in great detail. Examples of model organism include *Homo sapiens* (human), *Mus musculus* (house mouse), but also include organisms from other groups, such as *Drosophila melanogaster* (fruit fly), *Saccharomyces cerevisiae* (baker's yeast) or *Escherichia coli*.

While the scientific findings associated with model organisms have greatly impacted our understanding of biology and propelled research forward, these organisms represent less than 0.1% of all species known to exist. Still, a vast majority (reported around 75%) of publications use model organisms as their research focus, furthering the underrepresentation of all non-model species. To investigate broader topics that concern large phylogenetic groups, such as the origins of animals and the emergence of animal stem cells, require to investigate other organisms.

Therefore, focusing on less established species promises to be of great value for future research endeavours. In the analysis presented, the experiments were carried out on a sponge (Phylum *Porifera*) species. Sponges are likely to provide new and fundamental biological insights, as they represent one of the earliest-branching animal lineages, are structurally very simple, and contain fewer cell types than most other animals.

VI. CONCLUSION

In this publication, we present sc-GRIP, a polarity inference workflow that harnesses the information of single-cell gene expression matrices. By creating a latent space embedding of interacting nodes, sc-GRIP is able to predict edge labels for a polarity, which can be clustered into activating and repressing relationships. We have evaluated sc-GRIP on well-researched model species, where we can use published ground truth to validate our predictions and compare them to standard

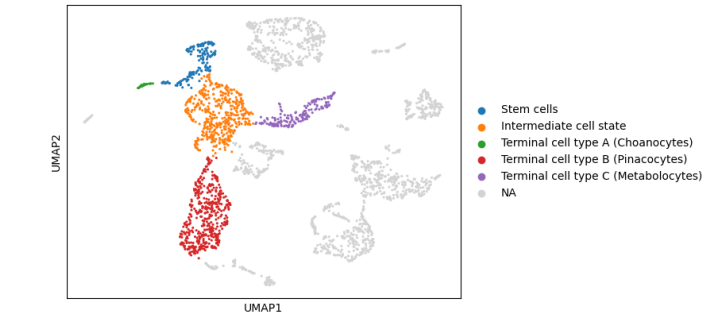

(a) Single-cell clustering

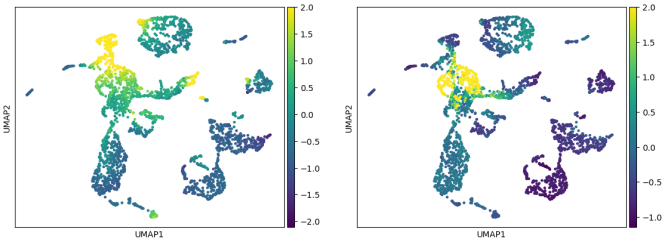

(b) Myc expression profile

(c) Stat expression profile

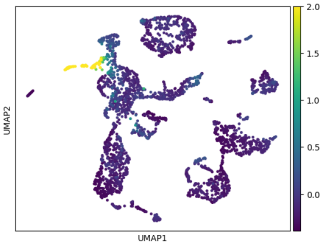

(d) Pou1 expression profile

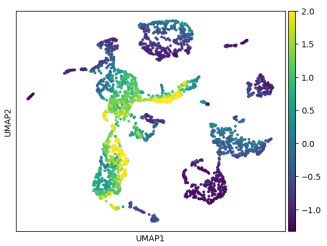

(e) Fox expression profile

Fig. 7. Gene expressions of four transcription factors (TFs) with annotated single-cell clusterings of *Suberites domuncula*: Myc, Stat, Pou1, and Fox. Visualized with UMAP.

methods in the field. Additionally, we are able to show the particular application of sc-GRIP on a non-model species. sc-GRIP makes novel and meaningful predictions (only a few examples of which are reviewed here) about the regulatory interactions that control stem cell functions in the sponge. These predictions are highly informative for guiding future experiments that will deepen our understanding of the complex network of TF interactions that control stem cell differentiation in this unique animal species.

Acknowledgment

This work has been partially funded by the Vienna Science and Technology Fund (WWTF) and the City of Vienna [Grant ID: 10.47379/ICT22059]. This research was funded in part by the Austrian Science Fund (FWF) [Grant DOI: 10.55776/P34955].

References

[1] P. Langfelder and S. Horvath, "Wgcna: an r package for weighted correlation network analysis," *BMC bioinformatics*, vol. 9, no. 1, pp. 1–13, 2008.

[2] J. Peng, G. Serrano, I. M. Traniello, M. E. Calleja-Cervantes, U. V. Chembazhi, S. Bangru, T. Ezponda, J. R. Rodriguez-Madoz, A. Kalsotra, F. Prosper *et al.*, "Simic enables the inference of complex gene regulatory dynamics across cell phenotypes," *Communications biology*, vol. 5, no. 1, p. 351, 2022.

[3] N. Papili Gao, S. M. Ud-Dean, O. Gandrillon, and R. Gunawan, "Sincerities: inferring gene regulatory networks from time-stamped single cell transcriptional expression profiles," *Bioinformatics*, vol. 34, no. 2, pp. 258–266, 2018.

[4] L. Garcia-Alonso, C. H. Holland, M. M. Ibrahim, D. Turei, and J. Saez-Rodriguez, "Benchmark and integration of resources for the estimation of human transcription factor activities," *Genome research*, vol. 29, no. 8, pp. 1363–1375, 2019.

[5] D. Hecker, M. Lauber, F. Behjati Ardakani, S. Ashrafiyan, Q. Manz, J. Kersting, M. Hoffmann, M. H. Schulz, and M. List, "Computational tools for inferring transcription factor activity," *Proteomics*, vol. 23, no. 23-24, p. 2200462, 2023.

[6] C. Z. Ma and M. R. Brent, "Inferring tf activities and activity regulators from gene expression data with constraints from tf perturbation data," *Bioinformatics*, vol. 37, no. 9, pp. 1234–1245, 2021.

[7] Z. Wong, Y. Chen, and L. Alvarez, "Exploring non-model organisms in single-cell transcriptomics," *Trends in Genetics*, 2024, in press.

[8] S. Aibar, C. B. González-Blas, T. Moerman, V. A. Huynh-Thu, H. Imrichova, G. Hulselmans, F. Rambow, J.-C. Marine, P. Geurts, J. Aerts *et al.*, "Scenic: single-cell regulatory network inference and clustering," *Nature methods*, vol. 14, no. 11, pp. 1083–1086, 2017.

[9] G. Eraslan, Ž. Avsec, J. Gagneur, and F. J. Theis, "Deep learning: new computational modelling techniques for genomics," *Nature Reviews Genetics*, vol. 20, no. 7, pp. 389–403, 2019.

[10] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," *arXiv preprint arXiv:1609.02907*, 2016.

[11] R. V. D. Berg, T. N. Kipf, and M. Welling, "Graph convolutional matrix completion," in *KDD Deep Learning Day*, 2017.

[12] T. N. Kipf and M. Welling, "Variational graph auto-encoders," *arXiv preprint arXiv:1611.07308*, 2016.

[13] R. Hasibi and T. Michoel, "A graph feature auto-encoder for the prediction of unobserved node features on biological networks," *BMC bioinformatics*, vol. 22, pp. 1–17, 2021.

[14] F. A. Wolf, P. Angerer, and F. J. Theis, "Scanpy: large-scale single-cell gene expression data analysis," *Genome*

biology, vol. 19, pp. 1–5, 2018.

[15] A. Gribov, M. Sill, S. Lück, F. Rücker, K. Döhner, L. Bullinger, A. Benner, and A. Unwin, "Seurat: visual analytics for the integrated analysis of microarray data," *BMC medical genomics*, vol. 3, pp. 1–6, 2010.

[16] V. A. Traag, L. Waltman, and N. J. Van Eck, "From louvain to leiden: guaranteeing well-connected communities," *Scientific reports*, vol. 9, no. 1, p. 5233, 2019.

[17] C. E. Grant, T. L. Bailey, and W. S. Noble, "Fimo: scanning for occurrences of a given motif," *Bioinformatics*, vol. 27, no. 7, pp. 1017–1018, 2011.

[18] H. Han, J.-W. Cho, S. Lee, A. Yun, H. Kim, D. Bae, S. Yang, C. Y. Kim, M. Lee, E. Kim *et al.*, "Trrust v2: an expanded reference database of human and mouse transcriptional regulatory interactions," *Nucleic acids research*, vol. 46, no. D1, pp. D380–D386, 2018.

[19] R. Revilla-i Domingo, C. Schmidt, C. Zifko, and F. Raible, "Establishment of transgenesis in the demosponge suberites domuncula," *Genetics*, vol. 210, no. 2, pp. 435–443, 2018.

[20] X. Chen, H. Xu, P. Yuan, F. Fang, M. Huss, V. B. Vega, E. Wong, Y. L. Orlov, W. Zhang, J. Jiang *et al.*, "Integration of external signaling pathways with the core transcriptional network in embryonic stem cells," *Cell*, vol. 133, no. 6, pp. 1106–1117, 2008.

[21] J. Xu, P. J. Gruber, and K. R. Chien, "Smad4 is essential for human cardiac mesodermal precursor cell formation," *Stem Cells*, vol. 37, no. 2, pp. 216–225, 2019.

[22] K. N. Smith, A. M. Singh, and S. Dalton, "Myc represses primitive endoderm differentiation in pluripotent stem cells," *Cell stem cell*, vol. 7, no. 3, pp. 343–354, 2010.

[23] Y. Hao, S. Hao, E. Andersen-Nissen, W. M. Mauck, S. Zheng, A. Butler, M. J. Lee, A. J. Wilk, C. Darby, M. Zager *et al.*, "Integrated analysis of multimodal single-cell data," *Cell*, vol. 184, no. 13, pp. 3573–3587, 2021.

[24] "A single-cell transcriptomic atlas characterizes ageing tissues in the mouse," *Nature*, vol. 583, no. 7817, pp. 590–595, 2020.

VII. Appendix

A. Data availability

The analyzed single-cell gene expression data, composed of 3,000 human blood cells, was provided from [23], and is a standard benchmark dataset for single-cell computational pipelines. The analysis files are available from the 10X Genomics website https://support.10xgenomics.com/single-cell-gene-expression/datasets/1.1.0/pbmc3k. The mouse single-cell transcriptomics data for the lung tissue were taken from [24] and are available from figshare https://figshare.com/articles/dataset/Processed_files_to_use_with_scanpy_/8273102/2

B. Analysed datasets

In this publication, we show the application of sc-GRIP on three different species: mouse, human, and a sponge representative. For the two model species, mouse and human, different published datasets were used in our analyses, with their respective features listed in Table V.

TABLE V. Dataset features

Dataset	#Nodes	#Observations	#Edges
Mouse liver	388	2334	108
Mouse muscle	444	1679	126
Human germ line	237	11	21
Human blood cells	235	2638	43
Sponge tissue	5650	2671	82068

C. Suberites domuncula dataset

For this internal dataset, the a section of the sponge *Suberites domuncula* was cut and left to regenerate for 14 days. Sponges are a group of animals with extraordinarily high regeneration potential, with their stem cells being able to heal all cell types necessary for survival. After 14 days of regeneration, the single-cell experiment takes place. The newly grown tissue, blastema, now contains diverse cell types, ranging from naive stem cells (prior to cell differentiation decisions) to terminally differentiated cell types.

D. Biological validation

If experimental validation is a possibility, sc-GRIP results can be confirmed either via chromatin immunoprecipitation-based methods (ChIP-seq), reporter assays, or using perturbation experiments, such as CRISPR knockout set-ups to test functional regulation of transcription factors. Using the sc-GRIP output as a guide for experimental analyses can increase the efficiency of the analyses and prevent the need for exhaustive testing.