

Revealing Regulatory Network Organization
through Single-Cell Perturbation Profiling
and Maximum Entropy Models

Thesis by
Jialong Jiang

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2024
Defended 30 May 2024

© 2024

Jialong Jiang
ORCID: 0000-0001-8560-8397

All rights reserved

ACKNOWLEDGEMENTS

I would like to take the opportunity to express my heartfelt gratitude to my advisor, Matt Thomson. The scene is still fresh in my memory when I met Matt six years ago during my interview at Caltech. We discussed the possibility and challenge of inferring the properties of biological systems using the language of statistical physics and agreed that perturbations would be essential for the inference. With the six years of journey, we transformed this concept into the modeling framework and computational pipeline in this thesis, which is truly enabled by the mentoring and support of Matt. Matt has been an extraordinary advisor, and my every interaction with him has been profoundly beneficial. His scientific enthusiasm, and insights into theory, computation, and biology would enlighten me in our every discussion. I received meticulous guidance on how to drive computation to biological results, and to organize and present the results to achieve scientific impacts. His precious advice was a leading light for my entire job search process. I am especially grateful for his unwavering support during the pandemic when I was stuck in China due to visa issues.

I also want to thank my committee members, Michael Elowitz, Rob Phillips, Lior Pachter, and Erik Winfree, for their invaluable insights and suggestions. Michael Elowitz provided tremendous help in applying the program discovery methods to other modularity design questions in system biology. I also substantially benefited from his advice on organizing my presentations and talks. Rob Phillips provided constructive feedback on the maximum entropy framework and suggested the use of pedagogical examples for the intuition behind perturbation-based inference. Lior Pachter gave important advice on the analysis of single-cell data and network inference methods. Erik Winfree contributed helpful discussions on the landscape of spin network models. Additionally, I am grateful to other Caltech faculty for their helpful discussions: Andrew Stuart and Venkat Chandrasekaran for their input on the development of inference and optimization theory for the model and perturbation design; Long Cai for his guidance on transferring the modeling framework to spatial transcriptomics; Kai Zinn, Woj Wojtowicz, and Barbara Wold for their advice on applications of the model. I also would like to express special thanks to Charles Trimble, for supporting my research with the Trimble fellowship.

Specifically for the D-SPIN project, I am grateful for all the collaborators who made the work possible. Sisi Chen was the director of the Caltech Single-Cell Profiling

and Engineering Center and also led the single-cell team in the Thomson lab. I learned a great amount about single-cell technologies and immune cell biology from her. It was a great pleasure to work with her on the design and analysis of the drug profiling experiments. David Sivak provided helpful discussions on the development of the theory, and we also benefited a lot from his careful revision of the manuscript. Zev Gartner helped to build the multiplexing experiments and provided valuable advice on the drug action mechanisms. Inna-Marie Strazhnik is our scientific illustrator, and I feel very lucky to work with her and learn the visual presentation of scientific results. She provided enormous help in organizing the manuscript figures, posters, and presentations. Yingying Gong was a SURF student I worked with who helped to convert the computational pipeline into a Python package. Ariane Helou and Guy Riddihough provided great help with the writing and revision of the manuscript. Tiffany Tsou and Jeff Park performed a major part of the drug profiling experiments. I would also like to thank Chris McGinnis for helping with the multiplexing experiments, Qin Zhu for collecting the drug-target databases, Eric Chow for sequencing the drug profiling experiments, John Hanna for discussion of the proteasome complexes, Tami Khazaei for preliminary drug data analysis, and Jost Vielmetter for the use of liquid-handling robots.

I also want to thank other members of the Thomson lab for helpful discussions and interactions, which have greatly enriched my research experience, including Jerry Wang, Zijie Qu, David Brown, Meera Prasad, Meng Wang, Xiaoqiao Chen, Michael Zellinger, Shichen Liu, Paul Rivaud, Manu Flores Bautista, Alec Lourenco, Guru Raghavan, Arjuna Subramanian, James Gornet and Brian Williams. During my time at Caltech, I also received various help from collaborators and friends at Caltech and beyond, including Michal Polonsky, Yujing Yang, Lynn Fang, Bo Gu, Yunyi Shen, Cassie Zhang, Jingjing Sun, Meichen Fang, Vickie Yang, Konnor von Emster, Yifan Chen, Fangzhou Xiao, Xinying Ren, Peng He, Xun Wang, Weiwei Yang, Ziyun Zhang, Shuo Cao, Tianzhe Zheng, Sheng Wang, Yitong Ma, Ronghui Zhu, Xiaozhe Ding, Xinhong Chen, Yuelin Shi, Xiang Fu, Bin Xu Wang, Hao Sun, Bin Wang and Yusen Zhu.

The warm assistance from the staff at Caltech has also been indispensable for my life here, including BBE administrator Liz Ayala, Thomson lab coordinators Rebecca Smith and Jessica Silva, Elowitz lab coordinators Jo Leonardo and Anthony Vasquez, Phillips lab coordinator Rui Malinowski, and Pachter lab coordinator Charlene Kim. I would also like to especially thank the Caltech International Student Program team

Laura Flower Kim and Daniel Yoder for their prompt and professional handling of all the international student documents.

Lastly, I want to express my deepest gratitude to my parents, Shuxian Jiang and Wenming Wu. Their unconditioned love and accompany have always been my source of inner peace, and keep me strong to face all challenges.

ABSTRACT

Gene regulatory networks within cells modulate the expression of the genome in response to signals and changing environmental conditions. Reconstructions of gene regulatory networks can reveal the information processing and control principles used by cells to maintain homeostasis and execute cell-state transitions. In this thesis, we introduce a computational framework, D-SPIN, that generates quantitative models of gene regulatory networks from single-cell mRNA-seq datasets collected across thousands of distinct perturbation conditions. D-SPIN models the cell as a collection of interacting gene-expression programs, and constructs a probabilistic model to infer regulatory interactions between gene-expression programs and external perturbations. Using large Perturb-seq and drug-response datasets, we demonstrate that D-SPIN models reveal the organization of cellular pathways, sub-functions of macromolecular complexes, and the logic of cellular regulation of transcription, translation, metabolism, and protein degradation in response to gene knockdown perturbations. D-SPIN can also be applied to dissect drug response mechanisms in heterogeneous cell populations, elucidating how combinations of immunomodulatory drugs can induce novel cell states through additive recruitment of gene expression programs. D-SPIN provides a computational framework for constructing interpretable models of gene-regulatory networks to reveal principles of cellular information processing and physiological control.

PUBLISHED CONTENT AND CONTRIBUTIONS

Jiang, Jialong et al. (2023). “D-SPIN constructs gene regulatory network models from multiplexed scRNA-seq data revealing organizing principles of cellular perturbation response”. In: *bioRxiv*. doi: 10.1101/2023.04.19.537364. J.J. participated in the conception of the project, developed the computational methods, analyzed the data, wrote the software, and participated in the writing, visualization, and revision of the manuscript.

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	vi
Published Content and Contributions	vii
Table of Contents	vii
List of Illustrations	x
Chapter I: Introduction	1
Chapter II: Inferring a unified regulatory network model from single-cell perturbation data	5
2.1 Introduction	5
2.2 Overview of D-SPIN framework	7
2.3 Mathematical formulation of D-SPIN framework	9
2.4 Computational methods for network inference in D-SPIN	11
2.5 Supplementary information	17
Chapter III: Benchmarking and analyzing the inferred regulatory networks	32
3.1 Introduction	32
3.2 State-of-the-art performance in inferring an HSC regulatory network model	35
3.3 D-SPIN constructs an interpretable, generative model of HSC perturbation response	38
3.4 D-SPIN achieves superior accuracy in large-scale regulatory network inference	42
3.5 Supplementary information	47
Chapter IV: Dissecting homeostasis maintenance strategies from genome-wide Perturb-seq	53
4.1 Introduction	53
4.2 Constructing a regulatory network model from genome-wide Perturb-seq data	54
4.3 D-SPIN reconstructs the architecture of core cellular pathways	61
4.4 Generating hypotheses for macromolecular complex sub-functions	62
4.5 Coarse-grained D-SPIN models provide insight into global perturbation response strategies	66
4.6 Supplementary information	72
Chapter V: Analyzing immunomodulatory drug and combination action on human immune cell population	80
5.1 Introduction	80
5.2 Modeling immunomodulatory drug responses in primary human immune cells	82
5.3 Classification of small molecules using D-SPIN	90
5.4 Drug combinations that generate hyper-suppression	99

5.5 Drug combinations that activate the innate immune system while repressing T-cell activation	105
5.6 Supplementary information	107
Chapter VI: Conclusions and discussions	117
Bibliography	122

LIST OF ILLUSTRATIONS

<i>Number</i>	<i>Page</i>
2.1 D-SPIN constructs unified gene regulatory network models from single-cell mRNA-seq data collected across perturbation conditions	6
2.2 D-SPIN yields a probabilistic model that can be interpreted as a gene regulatory network	9
3.1 HSC network model generates various cell-state distributions under perturbations.	34
3.2 D-SPIN achieves state-of-the-art network inference accuracy and reveals network-level mechanisms of cell-fate modulation by transcription factor perturbations.	36
3.3 D-SPIN identifies perturbation responses and reproduces perturbed cell-state distributions.	39
3.4 Edge sensitivity analysis reveals network-level mechanisms of cell state modulations.	41
3.5 D-SPIN achieves superior accuracy on large modular networks with up to a thousand nodes.	44
3.6 D-SPIN achieves superior accuracy on large random Erdős-Rényi networks with up to a thousand nodes.	46
4.1 D-SPIN generates a regulatory network model from genome-wide Perturb-seq data that reveals the organization of cellular pathways.	55
4.2 Evaluating single-gene network; D-SPIN reconstructs cell-cycle states.	58
4.3 The choice of gene-program number; Uncertainty of the inferred network.	60
4.4 D-SPIN reveals potential function partitions for molecular complex subunits.	63
4.5 Extended examples of molecular complex partitions identified by D-SPIN.	67
4.6 Coarse-graining of D-SPIN network model identifies global perturbation response strategies in K562 cells for distinct classes of gene knockdowns.	68
4.7 Response strategies exhibit as coherent up and downregulation of genes of relevant functions.	70

5.1	Experimental system of drug profiling in T-cell-induced immune activation.	82
5.2	Immunostaining validates CD3/CD28 antibody does not activate monocyte alone.	83
5.3	Timecourse of T-cell mediated immune activation shows various dynamics of inflammatory and signaling genes.	84
5.4	Large-scale drug profiling revealed a diversity of immune cell states under the modulation of 502 small molecule drugs.	85
5.5	Expression of cell-typing markers in each major cell type.	86
5.6	The choice of gene-program number; Uncertainty of the inferred network.	87
5.7	D-SPIN derives a drug-response network model from human immunomodulatory drug-response single-cell mRNA-seq profiling experiments.	89
5.8	Examples of gene program discretization and gene program visualization on immune cell population UMAP.	91
5.9	D-SPIN generates cell state distributions that are highly similar to experimental data under various metrics.	93
5.10	Drug response vectors identified by D-SPIN group drug into categories of phenotypic class.	94
5.11	Drug classification derived from D-SPIN aligns with known drug targets and mechanisms of action.	96
5.12	D-SPIN identifies phenotypical classes of drug categories and strength of immune inhibitors.	97
5.13	Drug combination profiling discovers novel cell states.	100
5.14	Two drug combinations generate novel cell states distinct from single drugs.	102
5.15	The combination of Dasatinib and Halcinonide produces novel cell states with coherent superposition.	103
5.16	Drug dosage combination profiling reveals phase diagrams of drug combination effect.	104
5.17	The combination of Tacrolimus and Motolimod produces novel cell states with incoherent dominance.	106

Chap ter 1

INTRODUCTION

The human genome encodes more than 30,000 genes, but a typical human cell expresses only 5,000 genes at any given time (Consortium* et al., 2022). Gene regulatory networks within cells modulate gene expression based upon environmental cues and the cell’s internal state (Bray, 1995; Regev and Shapiro, 2002, 2004; Davidson, 2010b). In gene regulatory networks, transcription factors and associated regulatory proteins interact with one another and gene promoters to activate or repress gene transcription (Bintu, Buchler, Garcia, Gerland, Hwa, Kondev, and Phillips, 2005; Davidson and Erwin, 2006; Davidson, 2010a). Gene regulatory networks play a central role in cellular decision-making processes (Kueh et al., 2013; Yosef et al., 2013; Sivak and Thomson, 2014; Sokolik et al., 2015). However, we have a limited understanding of the global logic encoded within gene regulatory networks and the underlying principles of cellular information processing.

Global reconstruction of gene regulatory networks in *E. coli*, yeast, and sea urchin embryos have revealed features of biological information processing, including network modularity, recurring network motifs, and combinatorial logic through the assembly of transcription factor complexes at gene promoters (Monod, 1949; Pardee, Jacob, and Monod, 1959; T. I. Lee et al., 2002; Milo et al., 2002; Shen-Orr et al., 2002; G. P. Wagner, Pavlicev, and Cheverud, 2007; Hnisz et al., 2013). However, our understanding of the information-processing principles of gene regulatory networks is based on a limited number of network reconstructions performed on model organisms over a limited range of physiological conditions. Even in *E. coli*, a large fraction of genes have unknown regulators (Ireland et al., 2020). In metazoans, gene network models have primarily focused on sub-circuits involved in specific processes like T-cell activation, T-cell fate selection, and embryonic stem cell differentiation (Kueh et al., 2013; Yosef et al., 2013; Sokolik et al., 2015). The comparative analysis of gene regulatory network architectures across cell types, developmental stages of an organism, and species therefore remains nascent, and there are very few quantitative models of regulatory networks that can predict the response of a cell to signals, genetic perturbations, or therapeutics. Cells simultaneously regulate tens of different processes including transcription, metabolism, protein degradation, and differentiation. How gene regulatory networks globally modulate core cellu-

lar processes in parallel, in response to environmental conditions, remains poorly understood.

Historically, gene regulatory network reconstruction and modeling have been constrained by the number of biochemical or genetic measurements required to reconstruct networks with hundreds to thousands of interacting protein components. Classical biochemical approaches, for example, perform bottom-up gene network reconstructions through pairwise binding measurements requiring order- M^2 biochemical assays for a network with M components (Harbison et al., 2004; Maerkl and Quake, 2007). Genetic approaches use a top-down strategy that infers gene regulatory network architecture from perturbation-response experiments where genes are knocked down or protein activity is inhibited with small molecules. Perturbation of one network component alters gene expression for subsets of genes; through the top-down association of perturbations with responses, perturbation approaches enable the identification of gene regulators, the organization of regulators into pathways, and the construction of integrated network models (Ferguson, Sternberg, and Horvitz, 1987; L. S. Huang and Sternberg, 1995; Yuh, Bolouri, and Davidson, 1998; Z. Hu, Killion, and Iyer, 2007). For global analysis of gene regulatory networks, however, perturbation approaches require knocking out hundreds to thousands of genes while monitoring transcription across thousands of genes. Historically, experimental realization of genome-wide network reconstruction through genetic perturbation has, like biochemistry, been limited by experimental scale when genes were knocked out in bulk assays.

Developments in single-cell genomics and perturbation barcoding circumvent some of the conventional limitations of perturbation-driven network reconstruction. Perturbation barcoding approaches, including Perturb-seq, click-tags, and MULTI-seq, allow the transcriptional state of each cell in a cell population to be measured across thousands of different genetic, signaling, and small molecule conditions (Dixit et al., 2016; McGinnis et al., 2019; Gehring et al., 2020; Schraivogel et al., 2020; Replogle et al., 2022). These experimental approaches identify the perturbation delivered to every single cell in a population while also measuring changes in the entire transcriptome through single-cell mRNA-seq. Given that single-cell perturbation measurements allow interrogation of a cell population's response to thousands of genetic, signaling, or small molecule perturbations, a core challenge is developing computational methods to integrate data from thousands of such experiments into a gene regulatory network model that can classify perturbations, map the flow

of information across a regulatory network, and predict cellular response to novel perturbations.

In this thesis, I will present a mathematical modeling and network inference framework, D-SPIN (Dimension-reduced Single-cell Perturbation Integration Network), that constructs gene regulatory network models directly from single-cell perturbation-response data. D-SPIN models the cell as a collection of interacting gene-expression programs and constructs a probabilistic model, known as a Markov random field or spin network, to infer regulatory interactions between gene-expression programs and applied perturbations.

In Chapter 2, we develop the framework of D-SPIN from both the theoretical perspectives and computational perspectives. Spin network models originated in statistical physics and have been generalized and applied to a wide range of systems and applications through work on maximum entropy methods in physics and on associative memory of Hopfield networks in machine learning (Jaynes, 1957; Hopfield, 1982). To enable information integration in large datasets with thousands of conditions and millions of single cells, D-SPIN exploits a natural factoring within the mathematical structure of Markov random fields inference to separate the learning problem into inference of a unified regulatory network and perturbation responses. D-SPIN can operate on single genes or gene-expression programs, and the program-level description enables D-SPIN to derive a coarse-grained description of a cellular regulatory network. D-SPIN can accommodate a wide range of different perturbation types including genetic perturbations, small molecules, signaling conditions, and even physiological states of health and disease.

In Chapter 3, we demonstrate the unique advantages of perturbation-based network inference by benchmarking the network reconstruction accuracies of D-SPIN on synthetic datasets. In synthetic data, gene expression profiles are simulated from known regulatory network models with differential equations, and the inferred networks are compared with the ground truth network architecture. On hematopoietic stem cell development networks taken from biological literature (Krumseck et al., 2011; Pratapa et al., 2020), D-SPIN achieves state-of-the-art performance on network inference accuracy. Moreover, through formal reasoning on the probabilistic graphical model generated by D-SPIN, we identify how each network edge contributes to the observed perturbation responses. On large synthetic networks with up to a thousand nodes with random or modular edge connections, D-SPIN surpasses the inference accuracy of existing methods by a large margin. Conceptually,

we demonstrate that the significant accuracy improvement is enabled by D-SPIN’s intrinsic ability to incorporate perturbations into a unified network model.

In Chapters 4 and 5, we apply D-SPIN to construct gene regulatory network models from experimental datasets with thousands of perturbations and millions of cells, including two of the largest single-cell mRNA-seq datasets in existence: a genome-wide Perturb-seq experiment on the K562 Chronic Myelogenous Leukemia cell line (Replogle et al., 2022) and a new human immune cell drug-response experiment collected in the Thomson lab. The integrated regulatory network models reveal global organizing principles of cellular regulation and perturbation responses, including the organizing principles of cellular pathways and networks of pathways used to maintain homeostasis in K562 cells in response to perturbation of core cellular processes such as respiration, transcription, and translation. Applied to profile the response of human immune cells to combinations of immunomodulatory drugs, D-SPIN models demonstrate that drug combinations can generate novel cell states through additive recruitment or superposition of drug-specific gene expression programs.

In the final chapter, we summarize and discuss insights derived from our research in regulatory network models of single-cell perturbation profiling. Broadly, D-SPIN provides a computational framework for large-scale modeling and comparative analysis of gene regulatory networks across cell types, physiological conditions, and organisms. D-SPIN can become a discovery engine for converting the increasing single-cell atlas datasets into mechanistic models of cellular regulation and perturbation responses to advance our knowledge of principles of cellular function and guide our control of biological systems.

Chapter 2

INFERRING A UNIFIED REGULATORY NETWORK MODEL FROM SINGLE-CELL PERTURBATION DATA

2.1 Introduction

In this chapter, we develop mathematical and computational strategies to integrate data across many different single-cell perturbation experiments into a gene regulatory network model that provides qualitative and quantitative insights into the architecture of cellular pathways and regulatory strategies. D-SPIN solves the gene regulatory network inference problem by constructing a probabilistic model that encodes the distribution of transcriptional states within a cell population across a set of perturbed conditions (Figure 2.1). D-SPIN can be applied to model cell populations perturbed by inputs including gene knockdown/activation, altered signaling conditions, and disease states, as well as the treatment of small molecules or other therapeutics.

Mathematically, D-SPIN builds a spin network model or Markov random field, which can be interpreted as an interaction network. Spin-network models were initially introduced in physics to study magnetic materials known as spin glasses, whose properties emerge through positive and negative interactions between localized magnetic moments or spins (Jaynes, 1957; Hopfield, 1982; Castellana and Bialek, 2014). In physics and machine learning, spin network models have been generalized and applied to many different systems of interacting elements, including neural networks, bird flocks, and proteins (Hopfield, 1982; Schneidman et al., 2006; Bialek et al., 2012; Santhanam and Wainwright, 2012; Lang et al., 2014; Teschendorff and Feinberg, 2021), and also to study the storage of memories in spin networks known as Hopfield networks.

However, the application of spin network models in the single-cell context remains limited due to the absence of an effective statistical inference framework that can scale to large datasets, incorporate perturbations, and capture the information contained in the entire cell state distribution measured in single-cell experiments. D-SPIN exploits a factoring of the spin-network Hamiltonian to develop an efficient computational procedure for inferring spin-network models from large single-cell datasets with millions of cells and thousands of perturbation conditions.

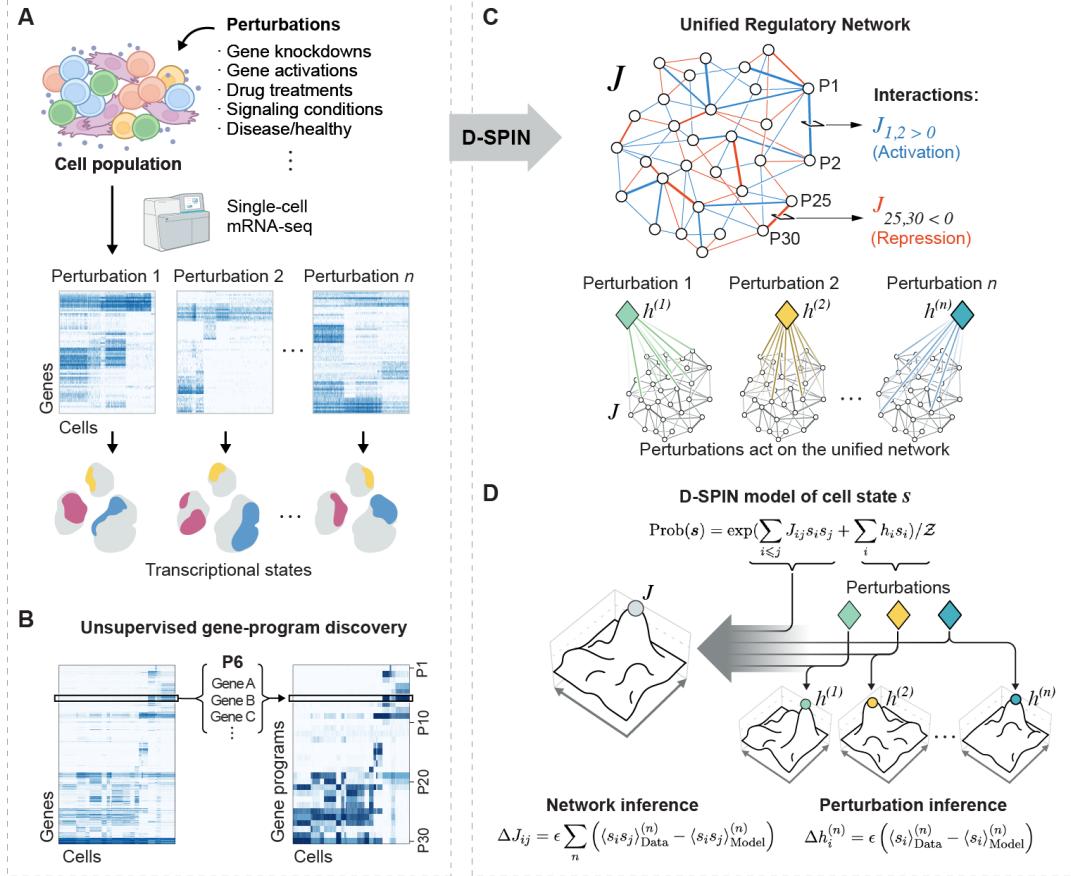


Figure 2.1: D-SPIN constructs unified gene regulatory network models from single-cell mRNA-seq data collected across perturbation conditions. (A) D-SPIN accepts as input single-cell mRNA-seq data from a cell population profiled across a series of different perturbed experimental conditions such as genetic perturbations, drug treatments, signaling conditions, or disease/healthy conditions. (B) To reduce the dimensionality of the network-inference task and enhance interpretability, D-SPIN extracts gene-expression programs from the data using orthogonal non-negative matrix factorization. D-SPIN is also capable of constructing gene-level networks with thousands of nodes using pseudolikelihood-based inference. (C) D-SPIN then constructs a unified regulatory network J , whose edges represent inferred interactions between gene-expression programs. D-SPIN also infers interactions between the network and applied perturbations. (D) Mathematically, D-SPIN constructs the regulatory network model by inferring a probabilistic model that encodes the probability of each transcriptional state in terms of an interaction matrix J and sample-specific perturbation response vectors $h^{(n)}$. D-SPIN scales to large datasets with millions of cells through a parallelized, maximum-likelihood procedure. At each step of inference, D-SPIN generates model samples for each perturbation condition in parallel, and then uses the difference between model and data cross-correlation and mean to update the network weights J and each perturbation response vector $h^{(n)}$.

2.2 Overview of D-SPIN framework

In general, spin network models consider a network of M interacting elements s_i . The goal of the framework is to construct a model that encodes the probability $P(s_1, \dots, s_M)$ of each potential configuration of the network of elements based upon inferred interactions between elements, as quantified in a coupling matrix \mathbf{J} . In D-SPIN, s_i is chosen to be three-state discrete variables in $[-1, 0, 1]$ instead of two-state variable in $[-1, 1]$ as is common in statistical physics. The three-state representation is chosen to model intermediate levels of gene expression, which can be functionally relevant for gene regulation and for improved quantitative agreement between model and experimental cell state distributions (Section 2.3). Similar discretization methods have been deployed to reduce noise and accommodate nonlinear dependencies from the pioneering works in modeling gene expression networks using microarray measurements (Friedman et al., 2000; Pe'er et al., 2001; Zhou et al., 2019).

D-SPIN can be applied to model networks with thousands of individual genes. We adopt an optimized inference algorithm for building gene-level networks where we scale inference to networks that contain thousands of genes using the concept of pseudolikelihood (Besag, 1974; Ravikumar, Wainwright, and Lafferty, 2010; Aurell and Ekeberg, 2012; Nguyen, Zecchina, and Berg, 2017). However, gene-level networks with thousands of nodes are difficult to interpret. Therefore, in addition to providing a scalable method for inferring D-SPIN models on individual genes, we also design D-SPIN to construct reduced- or low-dimensional spin network models by modeling interactions between gene programs, i.e., co-regulated groups of genes (Segal, Shapira, et al., 2003; Roguev et al., 2008; X. Chen, S. Chen, and Thomson, 2022). Cells regulate their transcriptional state by modulating transcription factors that impact sets of genes related to differentiation states, cellular pathways, and core physiological processes. By focusing on gene expression programs, rather than single genes, D-SPIN generates low-dimensional models of cellular regulation that are interpretable and can be inferred from data through a computationally efficient inference procedure. Building upon previously published work, D-SPIN contains an automatic procedure for identifying co-regulated groups of genes (gene programs) from single-cell data, based on unsupervised orthogonal nonnegative matrix factorization (oNMF) and phenotypic gene-program annotation (Figure 2.1 B, Section 2.5)(Choi, 2008; Vavasis, 2010; Heimberg et al., 2016; S. Chen et al., 2020), but can also accommodate pre-defined gene sets from prior biological knowledge (Section 2.5).

Following the extraction of gene-expression programs, D-SPIN applies a maximum likelihood statistical inference procedure to infer a network of regulatory interactions between gene programs and between programs and applied perturbations (Figure 2.1 C). Regulatory interactions between gene programs are represented in an interaction matrix \mathbf{J} where each entry J_{ij} quantifies the inferred interaction between program i and j . Interactions between the regulatory network and perturbations are encoded in a set of perturbation response vectors $\mathbf{h}^{(n)}$ where $h_i^{(n)}$ quantifies the interaction of the perturbation n with gene program i .

Following inference of \mathbf{J} and \mathbf{h} , D-SPIN models can be analyzed as network models (undirected graphical models) that reveal regulatory connections between cellular processes and identify interactions between applied perturbations and cellular gene-expression programs (Figure 2.2). D-SPIN models can also be used to generate the distribution of transcriptional states in a cell population under normal and perturbed conditions and to explain how interactions within a gene network lead to a specific distribution of cell states in an experimental population. The models provide a computational strategy that can bridge geometric, cell clustering approaches like Uniform Manifold Approximation and Projection (UMAP) with mechanistic modeling approaches that seek to interpret changes in cell population structure in terms of underlying regulatory interactions. Geometric embedding approaches like UMAP, or even principal components analysis (PCA), can help identify shifts in gene expression distributions within a cell population across conditions; however, interpretation of such shifts in terms of genes and gene interactions remains a major challenge (Chari and Pachter, 2023). D-SPIN provides a quantitative framework that can reveal how perturbations to specific nodes of a network alter the distribution of transcriptional states in a cell population and, thus, alter the rendering of a cell population in linear or non-linear embedding methods like PCA and UMAP, respectively.

Compared with existing methods of regulatory network inference, D-SPIN has three major advantages. First, D-SPIN integrates information from multiple perturbations into a unified regulatory network model, significantly improving the network reconstruction accuracy. Second, D-SPIN is generative and defines a complete probabilistic distribution of all transcriptional states with a set of minimal parameters of gene (program) interactions and perturbation bias, so that we can perform reasoning on the network model. Third, D-SPIN is highly scalable with the number of cells and conditions as the computation time scales linearly with cell number, and its

algebraic structure enables parallelized inference.

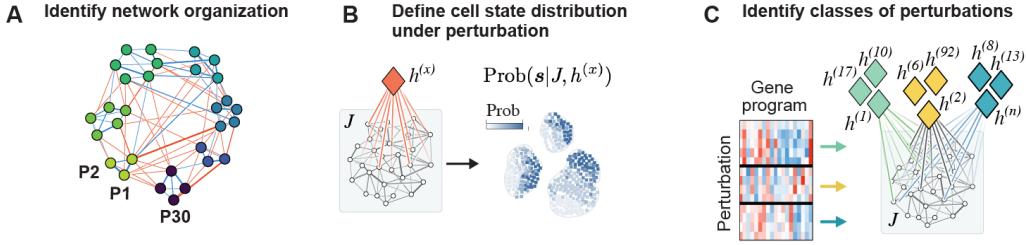


Figure 2.2: **D-SPIN yields a probabilistic model that can be interpreted as a gene regulatory network.** The D-SPIN model can be applied to (A) estimate the distribution of transcriptional states defined by the regulatory network model under a specific condition, (B) reveal principles of gene regulatory network organization, including modularity and sub-network architecture, and (C) classify perturbations into groups that reflect biological function.

2.3 Mathematical formulation of D-SPIN framework

The probabilistic graphical model we used in the formulation of D-SPIN originates from the Ising model in statistical physics, also called a spin glass, spin network, or Markov random field in different disciplines. Spin network models were initially introduced to analyze phase transitions in magnetic materials known as spin glasses. The spins interact with one another through a set of symmetric couplings and also with an applied magnetic field which biases spins to point up or down. Following the development of spin glass models, the maximum entropy formulation of statistical mechanics revealed that spin network models provide the most general description of a set of M objects that are interacting through pairwise interactions (Jaynes, 1957; Nguyen, Zecchina, and Berg, 2017).

Specifically, the model defines an energy function $E(\mathbf{s})$ computes the effective energy of a given transcriptional state \mathbf{s} by balancing the influence of the regulatory network (encoded in \mathbf{J}) and the perturbation-specific inputs (encoded in each $\mathbf{h}^{(n)}$). The discretized state vector $\mathbf{s} = [s_1, s_2, \dots, s_M]$, $s_i \in \{-1, 0, 1\}$, where M is the number of genes or gene programs. The resulting probability of a given transcriptional state \mathbf{s} is

$$P(\mathbf{s}) = \frac{1}{Z} \exp[-E(\mathbf{s}; \mathbf{J}, \mathbf{h}^{(n)})] \quad (2.1)$$

$$E(\mathbf{s}; \mathbf{J}, \mathbf{h}^{(n)}) = - \sum_{ij} J_{ij} s_i s_j - \sum_i h_i^{(n)} s_i \quad (2.2)$$

$$Z = \sum_{\mathbf{s}} \exp[-E(\mathbf{s}; \mathbf{J}, \mathbf{h}^{(n)})], \quad (2.3)$$

where E is the energy function and Z is a normalizing constant called partition function. States with lower energies have higher probabilities; therefore, positive interactions $J_{ij} > 0$ promote s_i and s_j being simultaneously on or off, and negative interactions promote the two variables being at different states. The bias vectors h_i increase probability of the variable s_i being on ($h_i > 0$) or off ($h_i < 0$).

D-SPIN assumes that across perturbation conditions, the interaction network \mathbf{J} stays the same and the bias vector $\mathbf{h}^{(n)}$ is modulated by each perturbation condition n to activate different modes in the network and to produce different cellular transcription state distributions. We train the model from data by gradient ascent of the log-likelihood function.

$$\log \mathcal{L}(\mathbf{J}, \mathbf{h}^{(n)}) = \sum_n^N \sum_{\mathbf{s}_{\text{Data}}^{(n)}} \log P(\mathbf{s}_{\text{Data}}^{(n)} | \mathbf{J}, \mathbf{h}^{(n)}) \quad (2.4)$$

We compute the gradient of the objective function as

$$\begin{aligned} \frac{\partial \log \mathcal{L}}{\partial J_{ij}} &= \sum_n^N \sum_{\mathbf{s}_{\text{Data}}^{(n)}} \frac{\partial}{\partial J_{ij}} \left[-E(\mathbf{s}_{\text{Data}}^{(n)}) - \log Z \right] \\ &= \sum_n^N \sum_{\mathbf{s}_{\text{Data}}^{(n)}} \left\{ s_{i\text{Data}}^{(n)} s_{j\text{Data}}^{(n)} - \frac{1}{Z} \sum_{\mathbf{s}} s_i s_j \exp[-E(\mathbf{s})] \right\} \end{aligned} \quad (2.5)$$

The first term is the sum $s_i s_j$ over all experimental samples, and the second term is a constant, which is the expectation of $s_i s_j$ of the current model, defined by the parameters \mathbf{J}, \mathbf{h} . We normalize the gradient by sample number to improve numerical stability under a given step size. With similar derivation for $\partial \log \mathcal{L} / \partial h_i$, the gradients of the objective function have the following form.

$$\frac{\partial \log \mathcal{L}}{\partial \mathbf{J}} = \frac{1}{N} \sum_n^N \left(\langle s_i s_j \rangle_{\text{Data}}^{(n)} - \langle s_i s_j \rangle_{\text{Model}}^{(n)} \right) \quad (2.6)$$

$$\frac{\partial \log \mathcal{L}}{\partial \mathbf{h}^{(n)}} = \langle s_i \rangle_{\text{Data}}^{(n)} - \langle s_i \rangle_{\text{Model}}^{(n)} \quad (2.7)$$

where $\langle \cdot \rangle_{\text{Data}}^{(n)}$ is the average over the data of n -th conditions, and $\langle \cdot \rangle_{\text{Model}}^{(n)}$ is the expectation on the distribution defined by current model parameters $\mathbf{J}, \mathbf{h}^{(n)}$. Therefore, the modeling learning is essentially matching the cross-correlation $\langle s_i s_j \rangle$ and mean $\langle s_i \rangle$ between the model and data.

The model has some unique advantages:

1. From a statistical perspective, the spin network is the maximum entropy model given the mean and pairwise cross-correlation of data (Nguyen, Zecchina, and

Berg, 2017). Entropy describes how a distribution is “spread out” over all possible states, and the principle of maximum entropy states that the best model describing a system is the model that generates data distributions with the maximum entropy while agreeing with relevant statistics of the data. Such a model has no extra assumptions about the structure of the system apart from the measured statistics.

2. From a computational perspective, the inference of the spin network, i.e., maximization of the log-likelihood function, is a concave problem where the only local maximum is the global maximum (Section 2.5). Therefore, optimization techniques to avoid traps of local minimums are not necessary. However, it is worth noting that the optimization is still a difficult NP-hard problem in the field of computational complexity (Ravikumar, Wainwright, and Lafferty, 2010), primarily due to the computational complexity of accurate gradient estimation and large condition numbers from model identifiability issues (Santhanam and Wainwright, 2012). Inferring the model across perturbation conditions can mitigate these issues, and we computationally evaluated the curvature of the inference problem in our dataset (Section 2.5).
3. From a physics perspective, the spin network is a Boltzmann distribution defined by the energy of each state. The model can be connected to the distribution of microscopic state in thermodynamics and dynamics of a microscopic system under thermal noise (Section 2.5).

2.4 Computational methods for network inference in D-SPIN

Given single-cell gene-expression data collected from a series of perturbation conditions, D-SPIN performs an optimization procedure through gradient ascent to estimate model parameters \mathbf{J} and $\mathbf{h}^{(n)}$ (Figure 2.1 D). For spin network models, the optimization procedure yields a closed-form optimization gradient as in Equation 2.6 2.7. At each step of model inference, the procedure samples the gradient from the model for each perturbation condition and computes the update of the regulatory network \mathbf{J} and each perturbation vector $\mathbf{h}^{(n)}$.

In general, spin network model inference for a single condition is a convex optimization problem, and we show in the supplementary information that the multi-condition optimization problem solved by D-SPIN remains convex, so that the D-SPIN optimization has a single unique solution (Section 2.5). As a statistical model, Fisher information can be applied to quantify the uncertainty of edges in the D-SPIN

model to guide interpretation and experimental design (Section 2.5) (Jiang, Sivak, and Thomson, 2019). The D-SPIN inference procedure solves the full maximum likelihood inference problem associated with spin-network inference. Therefore, D-SPIN inference is distinct from the conventional Hebbian learning rule developed to encode memories in Hopfield networks, as applied in a pioneering study to perform regulatory network inference from bulk microarray data (Hopfield, 1982; Lang et al., 2014).

Numerically, various computational methods have been proposed to solve the spin network inference problem (Nguyen, Zecchina, and Berg, 2017), and we adapt three of the most accurate methods to the context of D-SPIN. Each method has its advantages and specific niches depending on the problem setting, such as network size, number of cells, and number of perturbation conditions.

Exact maximum likelihood inference For small networks, the probability distribution $P(\mathbf{s}|\mathbf{J}, \mathbf{h})$ can be explicitly computed, thus the exact gradient of \mathbf{J}, \mathbf{h} can be computed using Equation 2.6 2.7. With the exact gradient, the optimization problem can be solved using gradient ascent or other optimizers such as Momentum or Adam (Kingma and Ba, 2014).

As the gradient estimation requires enumerating all possible states, the computational complexity scales exponentially with the number of nodes. Though exact, the method is only applicable to small networks of around 10 nodes ($3^{10} \approx 6 \times 10^4$ states).

Markov Chain Monte Carlo (MCMC) maximum likelihood inference As the gradient only requires computing the mean and cross-correlation of the samples, we can approximate the complete distribution $P(\mathbf{s}|\mathbf{J}, \mathbf{h})$ by sampling an empirical distribution. Without evaluating the exact distribution, we can construct a Markov Chain between states whose stationary distribution is $P(\mathbf{s}|\mathbf{J}, \mathbf{h})$. Specifically, we utilize the Gibbs sampling scheme. Starting from a random initial state $\mathbf{s} = [s_1, s_2, \dots, s_M]^\top$, at each step we randomly take an index k in the M nodes, and update the value of s_k by its conditional distribution given all other nodes $P(s_k|s_1, \dots, s_{k-1}, s_{k+1}, \dots, s_M, \mathbf{J}, \mathbf{h})$. After a burn-in period of steps to allow the Markov Chain to equilibrate, the sequence of samples is an empirical distribution of $P(\mathbf{s}|\mathbf{J}, \mathbf{h})$ and can be used to estimate the gradient Equation 2.6 2.7.

In each sampling step, computing the marginal distribution is of complexity $\mathcal{O}(M)$.

For accurate cross-correlation estimation, the required number of samples scales as $\mathcal{O}(M^2)$, without considering the burn-in period of the MCMC process. Therefore the overall computational complexity is at least $\mathcal{O}(M^3)$. The MCMC method applies to medium-sized networks up to $30 \sim 50$ nodes.

Pseudolikelihood method The major challenge involved in scaling inference to large networks is the partition function Z in the distribution $P(\mathbf{s}|\mathbf{J}, \mathbf{h})$ which involves exponentially many terms with the network size M . An alternative approximation method called pseudolikelihood was developed originally for spatial statistics and adapted to spin network problems (Besag, 1974; Aurell and Ekeberg, 2012; Nguyen, Zecchina, and Berg, 2017). Rather than performing inference using the global likelihood function constrained by pairwise correlations, pseudolikelihood methods consider the conditional probability distribution, $P(s_k|s_1, \dots, s_{k-1}, s_{k+1}, \dots, s_M)$ for a single network component s_k given data about all other network components excluding s_k . Parameter inference is performed by maximizing the logarithm of a pseudolikelihood function, constructed as the product of M single spin conditional likelihood functions, yielding a regression-like problem. As an approximation to the likelihood function, pseudolikelihood methods are computationally efficient and scalable. Model parameters estimated by pseudolikelihood converge to the ground truth, with error decreasing exponentially with sample size in theoretical analysis and numerical tests on classic spin network models from physics (Ravikumar, Wainwright, and Lafferty, 2010; Aurell and Ekeberg, 2012). Additionally, the framework achieves high accuracy in cases where networks have strong interactions with local fields, and the number of data points required to estimate network parameters scales logarithmically with the number of network nodes (Ravikumar, Wainwright, and Lafferty, 2010). Pseudolikelihood methods can also be applied to estimate directed interactions (Aurell and Ekeberg, 2012; Nguyen, Zecchina, and Berg, 2017). Pseudolikelihood-based inference methods enable scaling of D-SPIN to infer gene-level models in gene regulatory networks with thousands of nodes.

Specifically, we denote $\mathbf{s}_{\setminus k} = [s_1, \dots, s_{k-1}, s_{k+1}, \dots, s_M]^T$ as the state \mathbf{s} except s_k .

The pseudolikelihood function for the inference problem of a single condition is

$$\begin{aligned} \text{Pseudo}P(\mathbf{s}|\mathbf{J}, \mathbf{h}) &= \prod_k^M P(s_k|\mathbf{s}_{\setminus k}, \mathbf{J}, \mathbf{h}) \\ &= \prod_k^M \frac{\exp E(s_k, \mathbf{s}_{\setminus k}, \mathbf{J}, \mathbf{h})}{\sum_{r \in \{-1, 0, 1\}} \exp E(s_k = r, \mathbf{s}_{\setminus k}, \mathbf{J}, \mathbf{h})} \\ &= \prod_k^M \frac{\exp(s_k \Theta_k + s_k^2 J_{kk})}{\exp(\Theta_k + J_{kk}) + 1 + \exp(-\Theta_k + J_{kk})}, \end{aligned} \quad (2.8)$$

where $\Theta_k = h_k + \sum_{j \neq k} J_{jk} s_j$ is the effective field conditioned on all other nodes $\mathbf{s}_{\setminus k}$. The pseudolikelihood function decouples the mutual dependence between nodes, thus removing the exponentially complex partition function Z . As a cost, the pseudolikelihood function in general does not sum to 1, and thus is not a distribution; this is where the “pseudo” name comes from. The gradient for the log pseudolikelihood objective function is

$$\begin{aligned} \frac{\partial \mathcal{L}_{\text{Pseudo}}}{\partial J_{ij, i \neq j}} &= \frac{\partial}{\partial J_{ij}} \sum_k^M \left(s_k \Theta_k + s_k^2 J_{kk} - \log \sum_{r \in \{-1, 0, 1\}} \exp(r \Theta_k + r^2 J_{kk}) \right) \\ &= 2s_i s_j - \sum_{k \in \{i, j\}} \frac{\exp(\Theta_k + J_{kk}) - \exp(-\Theta_k + J_{kk})}{\exp(\Theta_k + J_{kk}) + 1 + \exp(-\Theta_k + J_{kk})} \end{aligned} \quad (2.9)$$

$$\frac{\partial \mathcal{L}_{\text{Pseudo}}}{\partial J_{ii}} = s_i^2 - \sum_{k=i}^M \frac{\exp(\Theta_k + J_{kk}) + \exp(-\Theta_k + J_{kk})}{\exp(\Theta_k + J_{kk}) + 1 + \exp(-\Theta_k + J_{kk})} \quad (2.10)$$

$$\frac{\partial \mathcal{L}_{\text{Pseudo}}}{\partial h_i} = s_i - \sum_{k=i}^M \frac{\exp(\Theta_k + J_{kk}) - \exp(-\Theta_k + J_{kk})}{\exp(\Theta_k + J_{kk}) + 1 + \exp(-\Theta_k + J_{kk})}. \quad (2.11)$$

The computational complexity of the gradient computation scales with $\mathcal{O}(M^2)$. Also, note that the effective field Θ_k depends on each cell state \mathbf{s} , so the gradient computational scales with the total number of cells. Practically, the pseudolikelihood method is highly scalable and even applies to networks of thousands of nodes. The approximation of pseudolikelihood is more accurate when the number of samples is high. However, typically the number of observed samples is far lower than the total number of possible states (3^M). Therefore exact maximum likelihood and MCMC maximum likelihood are preferred when they are computationally feasible.

Furthermore, the form of pseudolikelihood is closely related to regression models, enabling assigning directionality to the inferred network. The distribution of a single node conditioned on all other nodes

$$P(s_k|\mathbf{s}_{\setminus k}, \mathbf{J}, \mathbf{h}) = \frac{\exp(s_k \Theta_k + s_k^2 J_{kk})}{\exp(\Theta_k + J_{kk}) + 1 + \exp(-\Theta_k + J_{kk})} \quad (2.12)$$

can be interpreted as a regression problem, where we predict the state of the dependent variable s_k with other variables are predictors. The interactions J_{ij} are the coefficients of the regression problem. If gene A predicts gene B better than B predicts A, this suggests a regulation direction of A to B. To compute the directional network, the gradient estimation Equation 2.9 can be simply replaced with an asymmetric version of

$$\frac{\partial \mathcal{L}_{\text{Pseudo}}}{\partial J_{ij, i \neq j}} = s_i s_j - \sum_{k=i} \frac{\exp(\Theta_k + J_{kk}) - \exp(-\Theta_k + J_{kk})}{\exp(\Theta_k + J_{kk}) + 1 + \exp(-\Theta_k + J_{kk})} \quad (2.13)$$

Even though D-SPIN is capable of inferring a directed network, we focus our analysis on undirected networks because directed networks cannot define a stationary distribution on the data when time information is not supplied. In the context of probabilistic graphical models, directed models are always constrained to be acyclic, i.e., with no cycle. Such a constraint is reasonable in the field of causal inference, where the circularity of causal relations is rare. However, in cellular regulatory networks, feedback loops are prevalent to maintain homeostasis or signal amplification. For example, regulatory network models of hematopoietic stem cell differentiation contain several loops, such as Pu1-Gata1-Gata2-Pu1 (Krum siek et al., 2011).

The acyclic constraint of directed probabilistic graphical models is fundamental, as cycles in the conditional dependence between variables will produce inconsistent distributions. As mentioned in the seminal work of Judea Pearl on causal inference (Pearl, 1987, 2022), even three cyclic dependent variables "will normally lead to inconsistencies", and the acyclic constraint can ensure consistency of the distributions. Constructing distributions on even two mutually dependent variables requires nontrivial constraints (Arnold and Press, 1989). Another approach to include directionality in regulatory network models is including explicit time dependence, which also has limited application due to the lack of experimental approaches to measure dynamical evolution of transcriptome profiles. Therefore, in the context of probabilistic models of regulatory networks, undirected networks are a more appropriate choice.

Parallelization of the inference

The form of the spin network model learning rules enables us to develop a highly efficient and scalable training algorithm for D-SPIN by exploiting intrinsic parallelization. Training of general graphical models requires estimating the gradient

function for every data point separately, which is typically computationally extensive. In contrast, in D-SPIN the gradient is pooled for each experimental condition, and only depends on the mean $\langle s_i \rangle$ and cross-correlation $\langle s_i s_j \rangle$ between gene programs in each condition. Consequently, the training of the network can be deployed in parallel, with each computational thread computing the gradient based on the mean and cross-correlation of a single experimental condition. This approach eliminates the need to estimate gradients on every data point separately (i.e., every cell's transcriptional state) and minimizes data communication cost by only requiring the mean and cross-correlation to be exchanged during parallelization. Therefore, we routinely use hundreds of CPU cores for the optimization, enabling efficient network inference over large datasets. Practically, to avoid overfitting on a large number of similar conditions, we estimate the network using a subset of experimental conditions with balanced samples from different phenomenological condition groups and utilize the inferred network to identify response vectors of each condition independently.

Regularization of the inference

In statistical inference, it is common to leverage prior knowledge about the potential form of the solution, a process known as regularization. Regularization nudges the solution towards a preferred direction, which can also be interpreted as assigning a prior distribution of the model parameters in the framework of Bayesian inference. For example, ℓ_1 (Lasso) regularization promotes sparsity of the solution, while ℓ_2 (Ridge) regularization promotes solutions with smaller magnitudes.

Specifically in the D-SPIN framework, we sometimes have prior knowledge of perturbation action from the experimental design. For example, in single-cell profiling of gene knockdown or activation, the response vector relative to control should have a strong inhibition/activation at the target gene. Suppose we have an estimation of the relative action of all the perturbations as $\mathbf{h}_r^{(n)}$. Then we can infer shared unperturbed single gene (program) activity \mathbf{h}_0 and penalize the difference between $\mathbf{h}^{(n)} - \mathbf{h}_0$ and $\mathbf{h}_r^{(n)}$ by ℓ_2 norm. In this case, the objective function becomes

$$\log \mathcal{L}(\mathbf{J}, \mathbf{h}^{(n)}) = \sum_n^N \sum_{\mathbf{s}_{\text{Data}}^{(n)}} \log P(\mathbf{s}_{\text{Data}}^{(n)} | \mathbf{J}, \mathbf{h}^{(n)}, \mathbf{h}_0) - \lambda \sum_n^N (\mathbf{h}^{(n)} - \mathbf{h}_0 - \mathbf{h}_r^{(n)})^2, \quad (2.14)$$

where λ is the strength of the regularization, representing the uncertainty of the prior knowledge of the relative response estimation. For this objective function, the

gradient of \mathbf{J} stays the same and

$$\frac{\partial \log \mathcal{L}}{\partial \mathbf{h}_0} = \frac{1}{N} \sum_n^N \left(\langle s_i \rangle_{\text{Data}}^{(n)} - \langle s_i \rangle_{\text{Model}}^{(n)} \right) \quad (2.15)$$

$$\frac{\partial \log \mathcal{L}}{\partial \mathbf{h}^{(n)}} = \langle s_i \rangle_{\text{Data}}^{(n)} - \langle s_i \rangle_{\text{Model}}^{(n)} - 2\lambda(\mathbf{h}^{(n)} - \mathbf{h}_0 - \mathbf{h}_r^{(n)}) \quad (2.16)$$

2.5 Supplementary information

Key assumptions of D-SPIN

The D-SPIN framework is constructed based on the maximum entropy principle, which selects the model that has the highest entropy while satisfying all the constraints of the system (Jaynes, 1957). Several key assumptions underpin the current form of the model, and we discuss these assumptions, along with their biological rationale and limitations.

A unified regulatory network under perturbations The core concept of D-SPIN is that all different perturbation responses, or more generally, the single-cell data of the same biological system collected under various conditions – such as including signaling environment, genetic perturbations, drug treatments, and healthy/disease conditions – share the same underlying regulatory network architecture. The perturbations are presumed to only alter the activities of individual genes or gene programs within the network, without changing the regulatory relationship between them. In typical biological contexts of single-cell perturbation profiling such as genetic perturbations and drug treatments, this assumption is valid because these perturbations target individual elements of the regulatory network by increasing/decreasing their abundance or altering the molecular configuration for increased/decreased activities.

However, this assumption should be treated with care in scenarios where the changes in cell state distributions are not induced by altered activities, but by changed regulations. One specific example is the developmental processes, where epigenetic changes are constantly reshaping the interactions in the system. This dynamic nature of the regulatory network would require further detailed models of the network change such as sparse modification on a core network, or a dynamic model that evolves with time.

Discretizing gene and program expression levels The major reason for discretizing gene or gene program expression level is that the maximum entropy models for continuous variables are Gaussian distributions, which only allow the existence of

a single cell state. Mathematically, the probability density function of a Gaussian distribution is a convex function, that for any two cell states s_1 and s_2 , their average $(s_1 + s_2)/2$ is at least as probable as the two single cell states. However, in an immune cell population, for example, the expression average of a T-cell and a myeloid cell would not be a valid cell state.

There are a few potential choices of discretizing the expression level l : the number of discretized states m and the choice between $\{-\lfloor m/2 \rfloor, -\lfloor m/2 \rfloor + 1, \dots, \lfloor m/2 \rfloor - 1, \lfloor m/2 \rfloor\}$ or $\{0, 1, \dots, m\}$. We choose to discretize l into three states $\{-1, 0, 1\}$ due to the following reasons:

1. Larger m increases computational complexity, and makes the model closer to a Gaussian distribution that has a single probability density maximum as we discussed.
2. We find that $m = 2$ is insufficient to characterize phenomena in our drug profiling experiments. We observed different levels of program activation instead of on-and-off switching. For example, both glucocorticoid drugs and immune inhibitors activate the program P22 M2 macrophage, but with different expression levels.
3. The choice $\{-1, 0, 1\}$ is preferable because the self-interactions term J_{ii} has more clear biological interpretations. Considering a single program only, self-activation $J_{ii} > 0$ is similar to a bi-stable switch produced by nonlinear activation, where -1 and 1 states are preferable. Similarly, self-inhibition $J_{ii} < 0$ is similar to negative feedback where the 0 state is preferable.
4. On the network inference benchmarking on synthetic networks, there is no significant difference between the choice of $\{-1, 0, 1\}$ and $\{0, 1, 2\}$.

Partitioning genes into programs To reduce dimensionality, we group co-expressed genes into programs to reduce computational complexity and enhance interoperability. In the gene program discovery step, D-SPIN requires that gene programs be non-overlapping subsets of genes. This constraint prevents confounding interactions in the network, which could arise if genes were shared between programs. For example, if gene A is shared between program X and Y, a positive interaction identified between X and Y might merely reflect the influence of gene A in both programs.

In regulatory networks, key regulators such as transcription factors often play multiple roles in different pathways so the non-overlapping assumption might be unsuitable. However, the expression levels of these key regulators are typically lower than their downstream effect genes. As the gene program discovery aims to reconstruct the gene expression matrix with the reduced dimensionality, major gene components in each program are generally composed of these downstream genes with highly specific physiological functions, therefore the non-overlapping partition remains valid. To further explore the interactions between the key regulators with lower expression levels, it is recommended to construct single-gene level network models with the scalable inference algorithm we provided. This approach discovered detailed interactions between these regulators, complementing the global insights from the program-level analysis.

General data processing pipelines

The D-SPIN framework utilizes single-cell transcriptional profiling data from multiple conditions to construct a regulatory network model based on the gene programs identified, and the impacts of each perturbation to the network on the program level. The framework consists of three stages: gene program discovery, regulatory network model inference, and network analysis.

D-SPIN is compatible with different methods for gene program extraction, such as unsupervised matrix decomposition, differential expression, gene set enrichment, and manual curation. Users can also pre-assign a set of interesting single genes or gene programs and let D-SPIN partition the remaining genes into gene programs. For each gene program (set of genes $\{g_i\}$), the expression level l is determined by the weighted average of genes in the set that best approximates the expression of all genes in the set (Section 2.5). In our analysis, we use orthogonal non-negative matrix factorization (oNMF) to partition genes into gene programs (Section 2.5).

For the regulatory network inference step, D-SPIN applies a discretized spin-network modeling framework originating from statistical physics. The discretization can account for multiple different gene expression states while being resistant to counting noise in single-cell data. Further, discretization naturally generates multiple stable states and accommodates multiple cell types within a population and with a minimal interaction model, which can be interpreted as the steady state of a reaction system with nonlinear activations (Section 2.5). The number of possible states can be increased in a controlled way by changing the number of discretiza-

tions. The original Ising model from statistical physics uses two discretized states $\{-1, 1\}$, while we elected to use three discretized states $\{-1, 0, 1\}$ for both better interpretation and improved quantitative agreement between model and empirical cell state distribution by including two different levels of gene program expression (Section 2.3). In D-SPIN, each cell is represented as a state vector of gene programs $\mathbf{s} = [s_1, s_2, \dots, s_M], \mathbf{s} \in \{-1, 0, 1\}^M$. Thus, D-SPIN serves as a probabilistic graphical model of the observed distribution of cell states that can be interpreted as a regulatory network.

D-SPIN’s improved network inference accuracy is achieved by integrating information across multiple conditions, so as to capture different modes of the underlying regulatory network. If a large proportion of experimental conditions have very similar cell state distributions, the network inference would be biased towards such dominating conditions. In such cases, we recommend only including a subset of conditions that have different cell state distributions from each other to balance the conditions used for inference.

The regulatory network and response vectors inferred by D-SPIN can help define underlying principles of gene regulatory networks and identify classes of perturbations based on response signatures. From the network perspective, biological networks commonly exhibit a modular organization where the network contains a set of strongly interacting subnetworks, each of which is associated with a specific biological function. We analyze D-SPIN regulatory networks with Leiden community detection to identify modular subnetworks. Further, D-SPIN defines a probabilistic distribution of cell states only through the pairwise interaction between genes or gene programs, providing an interpretable framework for reasoning how cell states are constructed by these interactions, for example, the contribution of each network interaction to observed perturbation impacts.

From the perturbation perspective, the response vectors provide compact representations of the impact of perturbations on the network, enabling the clustering and visualization of perturbation responses across multiple experimental batches with a simple batch effect correction scheme. Due to the linear form of \mathbf{h} in the model, subtracting the response vector of the control sample calculates the up or downregulation of each gene program relative to the control. As the response vectors quantify the interactions between perturbations and the network that define the perturbation responses, joint graph clustering of network and perturbations can be applied to understand global perturbation action patterns and response strategies in the context

of the regulatory network.

Proof of convexity of the inference problem

We prove the multi-condition/multi- \mathbf{h} inference problem by D-SPIN is a convex optimization problem. Specifically, the log-likelihood objective function is concave. Therefore the only local optimum is the global optimum, and the solution to the optimization problem is unique.

For a single condition, D-SPIN formulation is the same as the inverse Ising problem, whose log-likelihood function is concave, and strictly concave if all coupling and fields are finite. We follow the proof from a review (Nguyen, Zecchina, and Berg, 2017). We denote $\boldsymbol{\lambda} = \{\mathbf{J}, \mathbf{h}\}$, and $Q_k(\mathbf{s}) = \{s_i s_j, s_i\}$ for the model parameters and the log-likelihood can be written as

$$\log \mathcal{L} = \sum_k \lambda_k Q_k - \log Z(\boldsymbol{\lambda}). \quad (2.17)$$

The second derivative of the log-likelihood function follows

$$\begin{aligned} -\frac{\partial^2 \log \mathcal{L}}{\partial \lambda_i \lambda_j} &= \frac{\partial^2 \log \sum_{\mathbf{s}} \exp \sum_k \lambda_k Q_k}{\partial \lambda_i \lambda_j} = \frac{\partial}{\partial \lambda_j} \frac{\sum_{\mathbf{s}} Q_i \exp \sum_k \lambda_k Q_k}{\sum_{\mathbf{s}} \exp \sum_k \lambda_k Q_k} \\ &= \frac{\sum_{\mathbf{s}} Q_i Q_j \exp \sum_k \lambda_k Q_k}{\sum_{\mathbf{s}} \exp \sum_k \lambda_k Q_k} - \frac{(\sum_{\mathbf{s}} Q_i \exp \sum_k \lambda_k Q_k)(\sum_{\mathbf{s}} Q_j \exp \sum_k \lambda_k Q_k)}{(\sum_{\mathbf{s}} \exp \sum_k \lambda_k Q_k)^2} \\ &= \langle Q_i Q_j \rangle - \langle Q_i \rangle \langle Q_j \rangle. \end{aligned} \quad (2.18)$$

The negative matrix of second derivatives $H = -\frac{\partial^2}{\partial \boldsymbol{\lambda}^2} \log \mathcal{L}$ is positive semidefinite because for any vector \mathbf{x} ,

$$\mathbf{x}^\top H \mathbf{x} = \sum_{ij} (\langle Q_i Q_j \rangle - \langle Q_i \rangle \langle Q_j \rangle) x_i x_j = \left\langle \left[\sum_k (x_k Q_k - \langle x_k Q_k \rangle) \right]^2 \right\rangle \geq 0. \quad (2.19)$$

Further, the matrix is positive definite if no observable Q_k has vanishing fluctuations, which holds when all the couplings and fields are finite. Therefore the log-likelihood function is strictly concave with finite couplings and fields, and the maximum likelihood estimation is a convex optimization problem.

For the formulation of D-SPIN of multiple conditions, we prove the case with two conditions and the proof of more conditions naturally follows. We start with proving a lemma:

Lemma: Let $A1, A2, B1, B2, C1, C2$ be matrices such that the block matrices $\begin{bmatrix} A1 & B1 \\ B1^\top & C1 \end{bmatrix}$ and $\begin{bmatrix} A2 & B2 \\ B2^\top & C2 \end{bmatrix}$ are both positive (semi)definite. Then the block

matrix

$$\begin{bmatrix} A1 + A2 & B1 & B2 \\ B1^\top & C1 & 0 \\ B2^\top & 0 & C2 \end{bmatrix} \quad (2.20)$$

is also positive (semi)definite.

Proof: For any test vector \mathbf{x} , we denote $\mathbf{x} = \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \end{bmatrix}$, where each component is compatible with the shape of $A_1 + A_2$, C_1 and C_2 , then we have

$$\begin{aligned} & \mathbf{x}^\top \begin{bmatrix} A1 + A2 & B1 & B2 \\ B1^\top & C1 & 0 \\ B2^\top & 0 & C2 \end{bmatrix} \mathbf{x} \\ &= \mathbf{x}_1^\top (A_1 + A_2) \mathbf{x}_1 + 2\mathbf{x}_1^\top B_1 \mathbf{x}_2 + 2\mathbf{x}_1^\top B_2 \mathbf{x}_3 + \mathbf{x}_2^\top C_1 \mathbf{x}_2 + \mathbf{x}_3^\top C_2 \mathbf{x}_3 \\ &= \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{bmatrix}^\top \begin{bmatrix} A1 & B1 \\ B1^\top & C1 \end{bmatrix} \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{bmatrix} + \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_3 \end{bmatrix}^\top \begin{bmatrix} A2 & B2 \\ B2^\top & C2 \end{bmatrix} \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_3 \end{bmatrix} (\geq) > 0 \end{aligned} \quad (2.21)$$

as both matrices are positive (semi)definite.

Similarly, we denote the parameter set $\boldsymbol{\lambda} = \{\mathbf{J}, \mathbf{h}^{(1)}, \mathbf{h}^{(2)}\}$, and the log-likelihood function of the two conditions are $\log \mathcal{L}^{(1)}$, $\log \mathcal{L}^{(2)}$, and the second order derivative of the objective function follows

$$\frac{\partial(\log \mathcal{L}^{(1)} + \log \mathcal{L}^{(2)})}{\partial \lambda_i \lambda_j} = \begin{bmatrix} \frac{\partial \log \mathcal{L}^{(1)}}{\partial J_i J_j} + \frac{\partial \log \mathcal{L}^{(2)}}{\partial J_i J_j} & \frac{\partial \log \mathcal{L}^{(1)}}{\partial J_i h_j^{(1)}} & \frac{\partial \log \mathcal{L}^{(2)}}{\partial J_i h_j^{(2)}} \\ \frac{\partial \log \mathcal{L}^{(1)}}{\partial J_i h_j^{(1)}}^\top & \frac{\partial \log \mathcal{L}^{(1)}}{\partial h_i h_j} & 0 \\ \frac{\partial \log \mathcal{L}^{(2)}}{\partial J_i h_j^{(2)}}^\top & 0 & \frac{\partial \log \mathcal{L}^{(2)}}{\partial h_i h_j} \end{bmatrix} \quad (2.22)$$

which is the rearrangement of the block matrices of the second-order derivative of the inverse Ising problem.

According to the lemma, the negative of the second-order derivative matrix is positive semidefinite, and positive definite if all couplings and fields are finite. This concludes our proof that the formulation of D-SPIN is a convex optimization problem.

Quantification of the inference uncertainty

Even though the optimization problem of D-SPIN inference has a unique solution, there is a possibility that there are other sub-optimal solutions that are distinct from

the true solution but have similar values of the objective function. These alternative solutions are phenomenologically referred as “flat valleys” of the optimization landscape, or “sloppy directions” of models. This flatness is quantified by Fisher information, which is a metric describing the curvature in the space of a parametric family of probability distributions. Large Fisher information indicates that small changes in the model parameter cause great changes in the distribution, while small Fisher information indicates that the distribution is insensitive to the parameter changes.

The Fisher information and uncertainty of the inference are connected by the Cramér-Rao bound, that for any unbiased estimator, the variance is lower bounded by the inverse of the Fisher information. Specifically, the bound is achieved by efficient estimators, and the maximum likelihood estimator is asymptotically efficient. Therefore, estimating the Fisher information matrix of the inference problem provides an estimation of the variance of the inferred network.

Specifically, using the same notation $\boldsymbol{\lambda} = \{\mathbf{J}, \mathbf{h}\}$, and $Q_k(\mathbf{s}) = \{s_i s_j, s_i\}$ as the previous section the Fisher information matrix of a single condition can be computed by definition

$$\begin{aligned} \mathcal{I}(\boldsymbol{\lambda}) &= -\mathbb{E} \left[\frac{\partial^2}{\partial \boldsymbol{\lambda}^2} \log P(\mathbf{s}|\boldsymbol{\lambda}) \right] = -\mathbb{E} \left[\frac{\partial^2}{\partial \boldsymbol{\lambda}^2} \left(\sum_k \lambda_k Q_k - \log Z(\boldsymbol{\lambda}) \right) \right] \\ &= \langle Q_i Q_k \rangle - \langle Q_i \rangle \langle Q_k \rangle, \end{aligned} \quad (2.23)$$

which has exactly the same form as Equation 2.18 as the first linear term $\sum_k \lambda_k Q_k$ vanishes under second order derivative. The bracket $\langle \cdot \rangle$ is average over $P(\mathbf{s}|\mathbf{J}, \mathbf{h})$, which is not directly available but can be estimated by the empirical distribution defined by samples. In practice, the Fisher information of a specific condition is estimated by

$$\mathcal{I}(J_{ij}, J_{kl}) = \langle s_i s_j s_k s_l \rangle_{\text{Data}} - \langle s_i s_j \rangle_{\text{Data}} \langle s_k s_l \rangle_{\text{Data}} \quad (2.24)$$

$$\mathcal{I}(J_{ij}, h_k) = \langle s_i s_j s_k \rangle_{\text{Data}} - \langle s_i s_j \rangle_{\text{Data}} \langle s_k \rangle_{\text{Data}} \quad (2.25)$$

$$\mathcal{I}(h_i, h_k) = \langle s_i s_k \rangle_{\text{Data}} - \langle s_i \rangle_{\text{Data}} \langle s_k \rangle_{\text{Data}}, \quad (2.26)$$

where the bracket $\langle \cdot \rangle_{\text{Data}}$ is average over all states in the data.

In the case of multiple samples, the overall Fisher information of the Fisher information has a similar form as Equation 2.22. We also only write for two conditions

for simplicity of notation.

$$\mathcal{I} = \begin{bmatrix} \mathcal{I}_{(\mathbf{J}, \mathbf{J})}^{(1)} + \mathcal{I}_{(\mathbf{J}, \mathbf{J})}^{(2)} & \mathcal{I}_{(\mathbf{J}, \mathbf{h})}^{(1)} & \mathcal{I}_{(\mathbf{J}, \mathbf{h})}^{(2)} \\ \mathcal{I}_{(\mathbf{J}, \mathbf{h})}^{(1)\top} & \mathcal{I}_{(\mathbf{h}, \mathbf{h})}^{(1)} & 0 \\ \mathcal{I}_{(\mathbf{J}, \mathbf{h})}^{(2)\top} & 0 & \mathcal{I}_{(\mathbf{h}, \mathbf{h})}^{(2)} \end{bmatrix} \quad (2.27)$$

As we have proved \mathcal{I} is positive definite when all couplings and fields are finite, $\mathcal{I}_{(\mathbf{J}, \mathbf{J})}$ and $\mathcal{I}_{(\mathbf{h}, \mathbf{h})}$ are all invertible. According to the Schur complement of the block matrix $M = \begin{bmatrix} A & B \\ B^\top & C \end{bmatrix}$, we have

$$\begin{bmatrix} A & B \\ B^\top & C \end{bmatrix}^{-1} = \begin{bmatrix} (M \setminus C)^{-1} & -(M \setminus C)^{-1} B C^{-1} \\ -C^{-1} B^\top (M \setminus C)^{-1} & C^{-1} + C^{-1} B^\top (M \setminus C)^{-1} B C^{-1} \end{bmatrix}, \quad (2.28)$$

where $M \setminus C := A - B C^{-1} B^\top$ is the Schur complement. Therefore the Cramér-Rao bound of the network estimation for N conditions is

$$\mathcal{I}_{(\mathbf{J}, \mathbf{J})}^{-1} = \left[\sum_n^N C_n \left(\mathcal{I}_{(\mathbf{J}, \mathbf{J})}^{(n)} - \mathcal{I}_{(\mathbf{J}, \mathbf{h})}^{(n)} \mathcal{I}_{(\mathbf{h}, \mathbf{h})}^{(n)} \mathcal{I}_{(\mathbf{J}, \mathbf{h})}^{(n)\top} \right) \right]^{-1}, \quad (2.29)$$

where C_n is the number of cells in the n -th condition. Specifically for each network edge J_{ij} , the diagonal term of the Cramér-Rao bound $\mathcal{I}_{(J_{ij}, J_{ij})}^{-1}$ quantifies the variance of the inference uncertainty of the specific edge.

The inverse of Fisher information $\mathcal{I}_{(\mathbf{J}, \mathbf{J})}^{-1}$ is the covariance matrix of the inferred network. Further, according to the Bernstein-von Mises theorem, the posterior distribution of the inferred network will converge to a normal distribution

$$P(\mathbf{J} | \{\mathbf{s}\}^{(1)}, \dots, \{\mathbf{s}\}^{(N)}) = \mathcal{N}(\mathbf{J}_0, \mathcal{I}_{(\mathbf{J}, \mathbf{J})}^{-1}), \quad (2.30)$$

where \mathbf{J}_0 is the result of inference. Therefore, we can sample from the posterior for alternative networks to validate the robustness of the inferred regulatory network.

Correspondence between D-SPIN and biochemical reaction systems

D-SPIN is closely connected to dynamical models of biochemical reactions as an approximation of the steady state of a chemical reaction system with saturating non-linear activation functions. For a biochemical reaction system with state $x \in \mathbb{R}^M$, the general dynamics can be written as

$$d\mathbf{X}_t = \boldsymbol{\mu}(\mathbf{X}_t, t) dt + \boldsymbol{\sigma}(\mathbf{X}_t, t) d\mathbf{W}_t \quad (2.31)$$

where $\mu(\mathbf{X}_t, t)$ are the reactions and $\sigma(\mathbf{X}_t, t) d\mathbf{W}_t$ is the stochastic noise. If the stochastic noise can be assumed a constant σ , then the state distribution $P(\mathbf{x})$ follows the Fokker-Planck equation

$$\frac{\partial P(\mathbf{x}, t)}{\partial t} = - \sum_{i=1}^N \frac{\partial}{\partial x_i} [\mu_i(\mathbf{x}, t) P(\mathbf{x}, t)] + \sum_{i=1}^N \sum_{j=1}^N \frac{\partial^2}{\partial x_i \partial x_j} \left[\frac{\sigma^2}{2} P(\mathbf{x}, t) \right]. \quad (2.32)$$

Specifically, if the reaction term can be written as the derivative of a potential function $\mu_i(\mathbf{x}) = -\partial V(\mathbf{x})/\partial x_i$, the equation has a Boltzmann distribution as the steady state solution for $\partial P(\mathbf{x}, t)/\partial t = 0$

$$P(\mathbf{x}) = \frac{1}{Z} \exp \frac{-2V(\mathbf{x})}{\sigma^2} \quad (2.33)$$

where Z is the normalization constant. A typical example is a linear expansion near the stable point of a dynamical system, where

$$\mu_i(\mathbf{x}) = \sum_j a_{ij} x_j + b_i = -\frac{\partial}{\partial x_i} \left(-\sum_{ij} a_{ij} x_i x_j - \sum_i b_i x_i \right). \quad (2.34)$$

This solution is a Gaussian distribution near the stable point (Sokolik et al., 2015; S. Chen et al., 2020). Such models only have a single minimum in the convex potential function, and are therefore not suitable for characterizing cell populations with various stable cell types and cell states.

The potential function, i.e., energy function Equation 2.2 in D-SPIN, is in a similar form to the Gaussian distribution as a second-order polynomial of variables. However, the energy in D-SPIN is a function of discretized state variables, which can be viewed as the discrete limit of composing the continuous state x_i with a saturating activation function $s_i = \phi(x_i)$, for example, sigmoid function for a two-state model. Given the potential function, the reaction term in the corresponding dynamical model is

$$\mu_i(\mathbf{x}) = -\frac{\partial}{\partial x_i} \left(-\sum_{ij} J_{ij} s_i s_j - \sum_i h_i s_i \right) \quad (2.35)$$

$$= \phi(x_i)' \left(\sum_j J_{ij} \phi(x_j) + h_i \right) \quad (2.36)$$

The first term $\phi(x_i)'$ is the derivative of the activation function, which can be approximated by a constant in the transition regime of the activation function, and close to 0 at the saturating regime of the activation function. The corresponding interpretation is that during the transition between states, the regulation received

by component i is a linear function of the activation $\phi(x_j)$ of all other components j . The nonlinearity is conceptually similar to the Hill function used in models of genetic circuits and allows the existence of multiple stable states controlled by the regulatory networks in the cell.

Gene program discovery by oNMF

We use orthogonal nonnegative matrix factorization (oNMF) for gene program discovery because oNMF generates a set of programs that are mathematically constrained to provide a high-accuracy representation of transcriptional states in the data but with no overlap between gene programs. Compared to typical matrix factorization methods like PCA, oNMF applies two constraints to the gene programs: non-negative weights and orthogonality. We used oNMF based on the following considerations:

1. Linearity: To ensure the interpretability of the model, each gene program is a linear combination of single genes
2. Non-negativity: The non-negative constraint avoids the ambiguity of interpreting negative weights, ensuring the programs are a set of co-expressed genes. Specifically in the context of regulatory network models, negative components in methods like PCA would complicate the interpretation of activation and inhibition interaction between gene programs.
3. Orthogonality: The orthogonality constraint makes it easier to interpret the data and aids model construction by forcing representations to be non-overlapping. Each gene program can be interpreted as a set of biological functions and each cell is represented by a collection of biological function activities. Without orthogonality, each cell is likely to be represented by a single program containing all genes expressed by the same cell type. Further, in the context of D-SPIN, shared genes between gene programs would cause confounding interactions between the programs. Therefore each gene program should be completely independent: in other words, orthogonal.

The formulation of oNMF is, given a non-negative gene matrix $X \in \mathbb{R}^{n_cell \times n_gene}$ and a program number K , oNMF solves the following optimization problem.

$$\begin{aligned} & \underset{W,H}{\text{minimize}} && \|X - WH\|_F^2 \\ & \text{subject to} && HH^\top = I, \quad W_{ij} \geq 0, \quad H_{ij} \geq 0 \end{aligned} \tag{2.37}$$

where $H \in \mathbb{R}^{K \times n_{\text{gene}}}$ is the gene program representation, $W \in \mathbb{R}^{n_{\text{cell}} \times K}$ is the cell state represented on the gene programs and $\|\cdot\|_F^2$ is the matrix Frobenius norm. We implemented the iterative algorithm proposed in this study (Choi, 2008), with W and H randomly initialized by uniform distribution on $[0, 1]$. The random initialization of H is orthogonalized using singular value decomposition (SVD) and takes absolute value before the start of iteration; we use 500 iterations in the computation by default.

According to the mathematical formulation of oNMF, the influence of each gene and cell state on the objective function is determined by the norm of the reconstruction error. To improve the discovery of gene programs, we utilized the following strategies to balance the contributions from each gene and cell state:

1. CV-based gene filtering.

We noted that genes with high expression levels tended to have greater variance than typically expected under Poisson distribution coefficient-of-variation (CV) filtering. Consequently, these genes often encode proteins involved in multiple housekeeping programs in the oNMF. Using a zero-inflated Poisson distribution for CV filtering, we accounted for the higher variance induced by dropouts, enabling better gene selection. We also removed genes that are expressed in a tiny fraction of cells in gene filtering, as genes with close to 0 expressions are not informative and risk introducing extra noise.

2. Gene normalization by standard deviation.

Genes with high variance contribute more to the reconstruction error and can thus dominate the oNMF gene program discovery. To mitigate this, we divided the expression of each gene by its standard deviation across cells, ensuring that each gene contributed equally. The expression of genes with only a few non-zero entries may be disproportionately large after the scaling. Ideally, such genes should be filtered out during gene filtering, but they can also be removed by setting a cap of the scaling factor, for example, no more than 5 times increase.

3. Cell-state balancing by cell clusters.

The gene program discovery will be dominated by reconstructing a specific cell state if that state constitutes a large proportion of the cell population. To address this, we used class-balance strategies frequently applied in machine learning practices. Specifically, we balanced different cell states by sub-

sampling cells based on cell clusters obtained by clustering or gene markers. We employed two different schemes: (1) Equal-sample balancing: to identify gene programs that are only expressed in a small fraction of the cell population, we sampled an equal number of cells from each cell cluster. (2) Square-root balancing: in machine learning, we found that taking samples from each data category proportional to the square root of the category size was effective (Mikolov et al., 2013). Similarly, we took samples from each cell cluster proportional to the square root of the cell number in that cluster.

The exact solution of non-negative matrix factorization is an NP-hard problem (Vavasis, 2010). But there are heuristic approximations of oNMF by iterative matrix update, and the solution varies with random initialization. Therefore, we run oNMF with different random seeds and compute a consensus gene program decomposition. For example, we take each row of H across different random seed repeats, and perform K-means clustering to compute the consensus composition of each gene program. The gene programs can be annotated with a combination of bioinformatics databases including DAVID, Enrichr, and String-db (Kuleshov et al., 2016; Szklarczyk et al., 2021; Sherman et al., 2022), and also manual lookup.

Evaluating gene program expression and discretization

A gene program is defined as a set of genes that co-express across the conditions in single-cell transcriptional profiling. Gene programs can either be identified through unsupervised learning techniques, such as oNMF (Section 2.5), or prior biological knowledge. In the D-SPIN framework, the expression level of gene program k of cell i is described by the variable $S_{ik} \in \{-1, 0, 1\}$. To transform the expression matrix of genes in the program into the discretized value for each cell, two steps are performed consecutively: computing a continuous weighted average of expression level L_{ik} , and discretizing into three levels.

Weighted average To reduce noise and to synthesize the expression of all genes in the program, the continuous expression level across all cells $\mathbf{l}_k = L_{\cdot, k}$ should be a weighted average of genes in the program. Denote the gene matrix as G , the index of genes in program k as $\{g^{(k)}\}$ (which is the index of nonzero elements in k -th row of H in the context of oNMF), the gene expression of cell i on the program is $G_{i, \{g^{(k)}\}}$. The optimal weight $w_{g^{(k)}}$ and the expression level L_{ik} that best characterize the gene

expression as the solution to the optimization problem

$$\min_{w_{g^{(k)}}, L_{ik}} \sum_i \sum_{\{g^{(k)}\}} (G_{i,\{g^{(k)}\}} - w_{g^{(k)}} L_{ik})^2 \quad (2.38)$$

which is summing over the error of reconstructing the gene expression with L_k and weight $w_{g^{(k)}}$ of every gene in the program. This is essentially a 1-component non-negative matrix factorization (NMF) problem and has mature implementations in Python scikit-learn packages.

Discretization Various methods have been empirically used to partition continuous expression into discrete levels, such as using percentiles or standard deviations. Discretization with K-means was found to perform well (Pe'er et al., 2001) and have a clear interpretation of minimizing the variance inside each expression level category. The 3-state K-means minimize the following objective function for each gene program k across all cells indexed by i

$$\begin{aligned} \min_{S_{ik}} & \sum_{r \in \{-1, 0, 1\}} \sum_{S_{ik}=r} \|L_{ik} - \text{Mean}\{L_{ik} | S_{ik} = r\}\|^2 \\ &= \sum_{r \in \{-1, 0, 1\}} \#\{S_{ik} = r\} \text{Var}\{L_{ik} | S_{ik} = r\}, \end{aligned} \quad (2.39)$$

which is assigning $-1, 0, 1$ to each of S_{ik} so to minimize the variance of L_{ik} inside each group of cells classified by S_{ik} . The K-means clustering also has mature implementations in Python scikit-learn packages.

Selection of gene program number in oNMF

The selection of the number of gene programs is a trade-off between model expressive power, computational complexity, and model interpretability. The optimal choice of program number typically hinges on the specific requirements of the application scenario. In the context of single-cell transcriptional profiling, the number of gene programs is typically set to $10 \sim 40$, depending on the desired resolution of the gene matrix decomposition. To aid in this decision, various model selection criteria have been proposed to facilitate this choice, among which the Bayesian information criterion (BIC) and elbow method are widely used examples.

BIC is derived from maximizing the model evidence $P(\text{Data}|\text{Model})$, where a penalty factor of the model dimension arises when integrating the parameter prior distribution over the whole parameter space (Hastie et al., 2009). Specifically, BIC is defined by maximized model log-likelihood $\hat{\mathcal{L}}$ with a linear penalty term of model

dimension.

$$\text{BIC} = k \log n - 2 \log \hat{\mathcal{L}}, \quad (2.40)$$

where k is the model dimension and n is the number of samples. Models with lower BIC are more preferred.

As oNMF is not a probabilistic model by definition, BIC is not directly applicable. However, we can model the residue of the gene expression matrix after the oNMF fitting. In general, if the model error can be described as independent and identically distributed normal distribution variables,

$$\log \hat{\mathcal{L}} = \log \prod_i^n \frac{1}{\sqrt{2\pi}\sigma} \exp \frac{(x_i - \hat{x}_i)^2}{2\sigma^2} = -n \log \sigma - n \log \sqrt{2\pi} + \sum_i^n \frac{(x_i - \hat{x}_i)^2}{2\sigma^2}, \quad (2.41)$$

where x_i and \hat{x}_i are data and model fitting, σ is the standard deviation of the error. As σ for a model is generally unknown and estimated from data, we have $\sigma^2 = \sum_i^n (x_i - \hat{x}_i)^2 / n$. Therefore the BIC is

$$\text{BIC} = k \log n - n \log \sigma^2 + C = k \log n - n \log (RSS/n) + C, \quad (2.42)$$

where C is a constant independent of model choice and $RSS = \sum_i^n (x - \hat{x})^2$ is the residual sum of squares. Motivated by this form, in oNMF we can use the matrix Frobenius norm of the fitting residual as RSS , number of cells as n , and number of programs as k to compute BIC to help to decide the number of gene programs.

Aside from statistical criteria like the BIC, the elbow method provides a heuristic approach to estimate the number of gene programs required in the oNMF. This method involves plotting a relevant cost function or objective function against the number of gene programs and looking for a point in the plot where the rate of changes drastically alters, resembling an “elbow”.

Gene matrices of single-cell profiling are known to be especially noisy; therefore traditional metrics such as model-explained variance may not be effective in revealing the elbow point. Instead, we use the number of genes that significantly correlate with corresponding gene programs as a more robust objective function. To consider the effect of noise, we can evaluate the pairwise gene-gene correlation as a reference distribution, which has a standard deviation in the order of 0.05 in our datasets. The correlation threshold between genes and programs can be set as 3σ or 4σ of the gene-gene correlations. Additionally, various methods (including MAGIC) have been proposed to denoise the gene expression matrix (Van Dijk et al., 2018). After denoising, both gene-gene correlation and gene-program correlation

drastically increase, allowing us to use a high correlation threshold, such as 0.5. These refined approaches would facilitate the decision of gene program numbers in D-SPIN.

BENCHMARKING AND ANALYZING THE INFERRED REGULATORY NETWORKS

3.1 Introduction

The regulatory networks are vital in theoretical biology, as they dictate plausible cell states and provide system-level explanations of information processing in cellular functions. The advancement of single-cell technologies provides unique opportunities to decode the regulatory network by depicting the cell states and their transitions with unprecedented resolution and data richness including differentiation, signal response, and disease progression. Multiplexed screening enables massive perturbation profiling at the single-cell level, including the knockdown of every single gene and activation of every transcriptional factor (Dixit et al., 2016; McGinnis et al., 2019; Replogle et al., 2022; Joung et al., 2023). As single-cell technologies advance, there is an increasing need for modeling frameworks to convert atlas-type databases into insights into the regulatory mechanisms of cellular functions. Nonetheless, current analyses are often limited to associating each perturbation with specific phenotypes and classifying perturbation. The critical challenge is to decode the regulatory network controlling these responses from thousands of perturbations and millions of cells.

From the early development of genomics with microarray and bulk sequencing to the current technologies of single-cell profiling, various computation methods have been proposed to identify regulatory relationships between genes. One major category of these methods infers regulations based on gene coexpression patterns, with various measures and mathematical models used. One early popular method was ARACNE, which selects regulatory interactions based on mutual information (MI) between genes and eliminates indirect interactions by data processing inequality (Margolin et al., 2006). Similarly, PIDC improved the pairwise MI measure with partial information decomposition to consider multiple variables together (Chan, Stumpf, and Babtie, 2017). Regression-based methods such as GENIE3 and GRN-Boost2 predict the expression of one gene with other genes, and higher prediction power indicates stronger interactions between the genes (Huynh-Thu et al., 2010; Moerman et al., 2019). There are also methods that use dynamical models where

time information is present such as SCODE (Matsumoto et al., 2017). Another category of methods enhances the inference with context-dependent biological information such as transcriptional factor (TF) binding, specifically searching for TF binding motifs in the promoter region of the candidate target genes, such as SCENIC (Aibar et al., 2017). The technology of single-cell ATAC-seq also facilitates the identification of activating TF binding sites by profiling the accessibility of chromatin regions on the genome. Methods like SCENIC+ and CellOracle were developed to leverage these multiomics measurements for improved network inference accuracy (Bravo González-Blas et al., 2023; Kamimoto et al., 2023). Although TF-DNA binding is one of the major gene regulatory mechanisms, many other regulatory mechanisms play important roles in cells including protein-protein interaction, kinase phosphorylation in signal transduction, and small molecule protein binding. The dependence on TF-motif-predicted interactions introduces limitations compared to coexpression-based network inference methods.

However, two significant challenges persist with the regulatory network inference on large-scale single-cell profiling data. Perturbations are widely used in genetics to uncover causal interactions from correlations, it is yet unclear whether perturbations are beneficial for large-scale network inference, and how the perturbation information can be formally incorporated into the network inference. Moreover, in existing analysis methods, the inferred regulatory networks are rarely connected to the observed cell states, namely how cell states are defined through the regulation between genes. Bridging the gap will lead to insight and control strategies for the organization of cell states, and open broad therapeutic applications from modulating the inflammatory response to advancing immunotherapy of cancer.

In this chapter, we show that the unique architecture of D-SPIN enables integrating information from multiple perturbation conditions to achieve high network inference accuracy and build a complete probabilistic model of cell states. With a benchmarking framework that can simulate single-cell data from Boolean regulatory relationships (Pratapa et al., 2020), we show D-SPIN achieves improved accuracy on a synthetic hematopoietic stem cell (HSC) network model over PIDC, GRNBoost2, and GENIE3, which are top-performers in the benchmarking study. Further, the probabilistic model built by D-SPIN connects the network with the distribution of altered cell states under perturbation, enabling us to develop a formal reasoning framework to evaluate how each individual edge contributes to the observed perturbation responses. Furthermore, the superior accuracy of D-SPIN

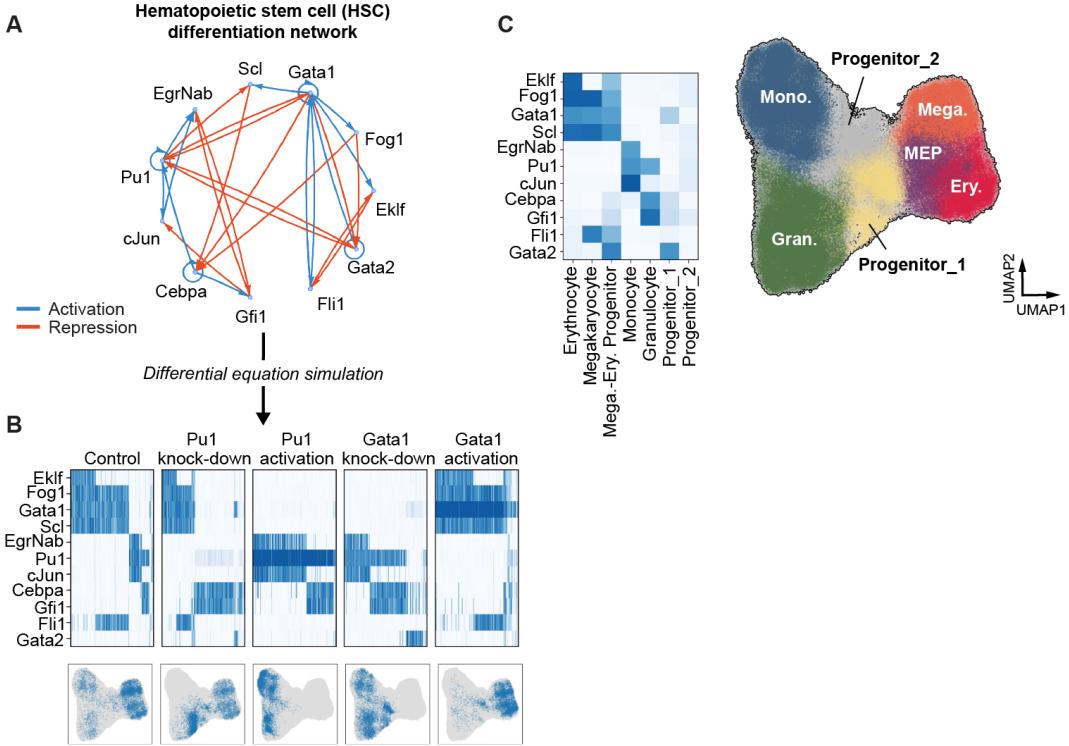


Figure 3.1: HSC network model generates various cell state distributions under perturbations. (A) HSC regulatory network model (Krumbsiek et al., 2011) contains 11 transcription factors (TFs) that interact through activation (blue) and repression (red) to modulate the differentiation of HSCs into different cell types. (B) Example simulated gene-expression profile heatmaps generated using a published computational framework BEELINE from the network across a series of simulated single-gene knockdown and activation conditions. (C) Heatmap shows average TF expression of seven clusters of cell states generated across all simulated conditions. Each cluster is annotated with its TF expression pattern based on the biological context of the HSC network model (Krumbsiek et al., 2011). All simulated cell states and clusters are visualized in the UMAP embedding.

extends to large networks with hundreds to a thousand nodes. D-SPIN outperforms existing methods by a large margin on two types of synthetic networks, modular networks, and Erdős-Rényi (ER) random networks (0.67 vs. 0.45 and 0.70 vs. 0.39 edge prediction accuracy). The improved accuracy of D-SPIN is enabled by integrating perturbation information into a unified network model, and the accuracy continues to rise with the increased number of perturbations.

3.2 State-of-the-art performance in inferring an HSC regulatory network model

To evaluate the accuracy of D-SPIN in performing network inference, we first applied D-SPIN to reconstruct and analyze a model of the hematopoietic stem cell (HSC) differentiation network using the network simulation and benchmarking framework BEELINE (Pratapa et al., 2020). The HSC network has two major groups of regulatory genes. The first group includes the transcription factor Pu1, controlling granulocyte/monocyte differentiation, and the second group includes Gata1, controlling megakaryocyte/erythrocyte differentiation (Krumsieck et al., 2011)(Figure 3.1 A). BEELINE simulates single-cell transcription data from biologically identified regulatory networks by generating and simulating ordinary differential equation (ODE) models with stochastic noise. BEELINE also includes existing network inference methods for benchmark accuracy evaluation. We used the BEELINE framework to generate synthetic gene-expression profiles from the HSC network for 22 perturbation conditions, encompassing knockdown/activation of each network node individually by turning down/up the transcription rate (Figures 3.1 BC, Section 3.5). For interpretation of the data we used the known biological functions of each transcription factor to classify the single-cell expression profiles into seven different cell state designations, including monocytes (Mono), granulocytes (Gran), erythrocytes (Ery) and megakaryocytes (Mega) (Krumsieck et al., 2011).

We applied D-SPIN to reconstruct a gene regulatory network model from the simulated data and compared the accuracy of the D-SPIN model to three different network reconstruction methods (PIDC, GRNBoost2, and GENIE3) that were the top performers in the BEELINE benchmarking study (Pratapa et al., 2020). We applied D-SPIN to generate a single network model encoded in the gene interaction matrix J and generate perturbation vectors $h^{(n)}$ for each perturbation condition (Figure 3.2 A, Figure 3.3 A). We applied the algorithms to individual genes, rather than gene programs, given the relatively compact size (11 nodes) of the HSC network used in the BEELINE benchmarking evaluation. Each method generates a set of inferred edges and provides confidence/importance estimates for each edge. We compared the methods by considering the accuracy of the top 10 and top 20 edges identified by each method.

On the simulated network reconstruction task, D-SPIN consistently achieved higher performance than comparable methods, primarily due to the natural way in which D-SPIN can accommodate perturbation data. For D-SPIN, on average 0.96 of the top

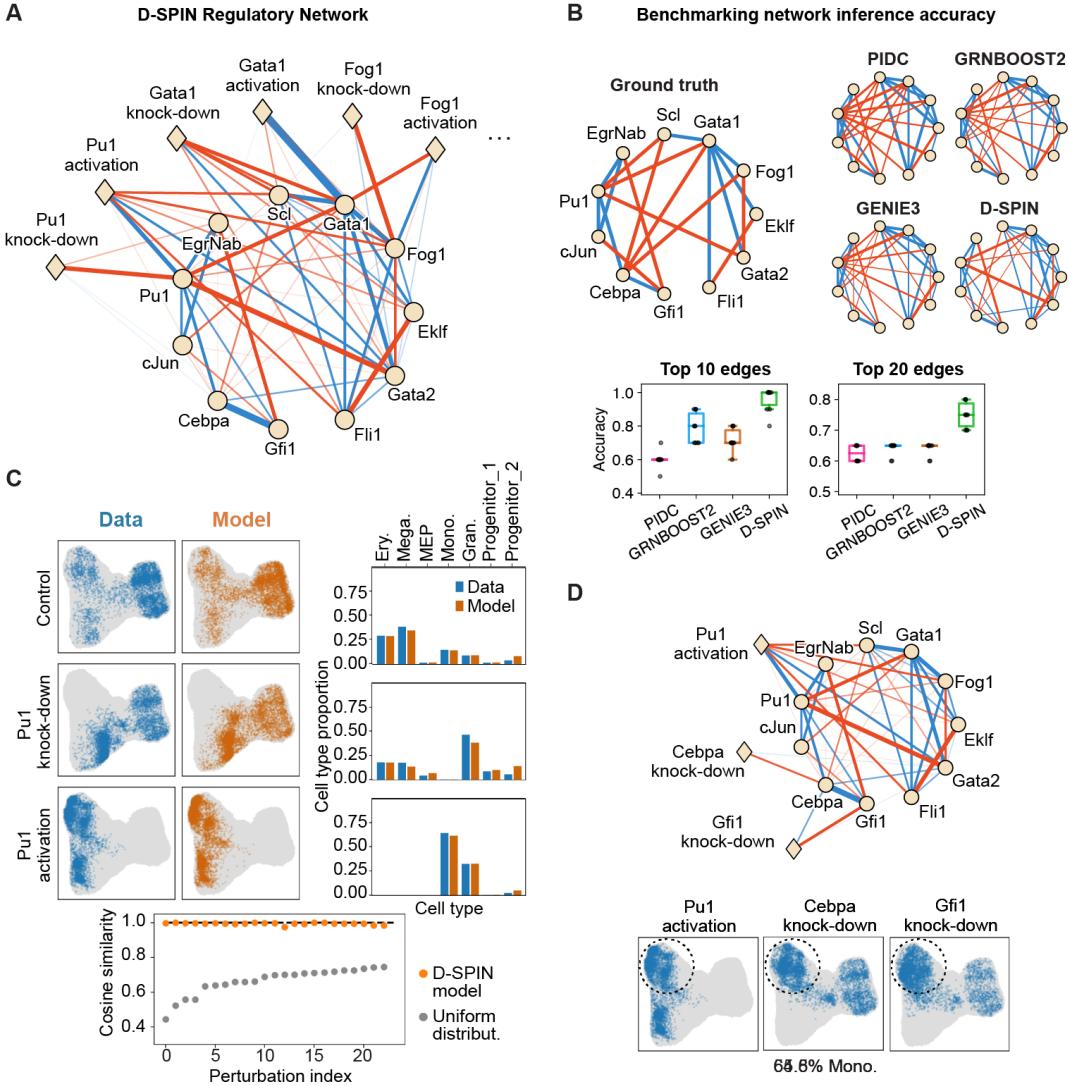


Figure 3.2: D-SPIN achieves state-of-the-art network inference accuracy and reveals network-level mechanisms of cell fate modulation by transcription factor perturbations. (A) Network diagram shows the inferred unified regulatory network model by D-SPIN that encodes inferred interactions between TFs as an interaction graph. Edges with diamond markers in the network show D-SPIN inferred perturbation vectors that estimate how knockdown or activation of TFs (e.g., Gata1 knockdown) impacts the regulatory network through up- or downregulation of TFs. (B) Diagram of the true network and inferred regulatory networks by D-SPIN and other state-of-the-art methods including PIDC, GRNBoost2, and GENIE3 (Pratapa et al., 2020). The inference accuracy is quantified by the box plot of the accuracy of the top 10 and 20 inferred edges across 10 random repeats of parameters in differential equation simulation using BEELINE. The ground truth network has a total of 20 edges.

Figure 3.2: (Continued) (C) (left) UMAP embedding comparisons of simulated single-cell data by BEELINE and state distributions generated by D-SPIN models for control and Pu1 activation/knockdown. (right) Bar plots quantify the proportion of HSC-derived cell types in each condition from simulated data and D-SPIN models. (Bottom) Cosine similarity between cell state distributions of data and model. D-SPIN models generate cell-state distribution highly consistent with simulated data by applying perturbations to the underlying network. (D) (top) Network diagram and (bottom) UMAP embedding of three different perturbations that generate increased monocyte population. The D-SPIN model reveals network edges that mediate the response of the network to different perturbations through the application of a formal reasoning framework to the regulatory network model. Both Pu1 activation and knockdown of Cebpa or Gfi1 generate increased monocyte states by directly interacting with monocyte genes or indirectly through the interaction between Gfi1 and EgrNab.

10 edges found by the model were correct across inference runs, as compared to 0.6, 0.77, and 0.7 for PIDC, GRNBoost2, and GENIE3, respectively. For D-SPIN, 0.75 of the top 20 edges were correct, compared with 0.625, 0.645, and 0.665 for PIDC, GRNBoost2, and GENIE3, respectively (Figure 3.2 B). The inference algorithm of D-SPIN is crucial to its accuracy improvement, as the Hebbian learning rule in Hopfield networks (Hopfield, 1982; Lang et al., 2014) only has 0.55 accuracy on the top 20 edges. In one example of D-SPIN outperforming existing methods, the networks inferred by PIDC, GRNBoost2, and GENIE3 all identified strong spurious activating interactions inside the Mega-Ery gene module. For example, Scl-Fog1 is predicted as one of the strongest edges in PIDC, GRNBoost2 and GENIE3 with weights 1.17, 1.34, and 1.70, while D-SPIN assigns an average weight of 0.35 to the edge and ranks it 24th. Similarly, Scl-Ekif is ranked within the top-12 edges in other methods with weights 0.99, 0.76, and 0.71, while D-SPIN assigns weight 0.4 to the edge and ranks it 22nd. D-SPIN reveals the central regulatory role of Gata1 by significantly reducing the weights of these false positive edges. D-SPIN also removed spurious inhibitory interactions between the two major groups of genes (Mono-Gran module and Mega-Ery module) from the edge predictions, e.g., Pu1-Ekif or Pu1-Fli1, revealing the actual regulatory target of Pu1. When spurious edges occurred, they were similar across the different inference frameworks: all four methods identified spurious edges in strongly connected regions of the HSC network, struggling to separate direct versus indirect connections. However, for D-SPIN uniquely, spurious edges can be further reduced through the sampling of higher-order perturbations. By including additional 96 double-gene perturbations,

D-SPIN achieves 100% accuracy on the top 10 edges and improves the accuracy of the top 20 edges from 0.75 to 0.81 (Figure 3.3 B).

Additionally, D-SPIN also achieves superior accuracy in inferring a directed network model using pseudolikelihood (Figure 3.3 C). The directed HSC network has 26 edges; the top 26 predicted edges of D-SPIN have an accuracy of 0.73, outperforming the other three methods (PIDC, GRNBoost2, and GENIE3) with accuracies of 0.41, 0.47, and 0.53, respectively. Despite the strength of D-SPIN in directed network inference, we focus our analysis on undirected networks due to the inherent challenge in defining consistent directed network models using steady-state cell state distributions (Section 2.4) (Pearl, 2022).

3.3 D-SPIN constructs an interpretable, generative model of HSC perturbation response

D-SPIN is a probabilistic, generative model that can generate the distribution of transcriptional states in a cell population subjected to a perturbation. In contrast, most existing network inference methods are not generative and use either regression or information-theoretic measures to identify and list candidate gene interactions (Pratapa et al., 2020). For the HSC network, the distribution generated by the D-SPIN model is highly concordant with the data distribution across all perturbation conditions, as visualized by the UMAP embedding and cell state distribution comparison (Figure 3.2 C, Figure 3.3 D). Quantitatively, the cell state distributions by D-SPIN are all above 96% cosine similarity with the simulated data. The quantitative agreement is especially remarkable as D-SPIN is a minimal and interpretable model that encodes each network interaction with a single signed strength parameter. Thus, for the HSC network, D-SPIN constructs a quantitative model of how the cell population is generated by the underlying gene regulatory network. We can apply the model to compute quantitative statements of gene expression change induced by each perturbation and to analyze the role of potential gene-gene interactions in mediating the perturbation response of the regulatory network.

As generative probabilistic models, D-SPIN models can be analyzed to gain insight into how edges within a D-SPIN network mediate the response to applied perturbations. We developed a formal framework to determine how the knockdown or activation of specific nodes in the network might impact the distribution of cell states generated by the regulatory network and to define how individual edges contribute to the perturbation response of the network (Section 3.5). Specifically, for

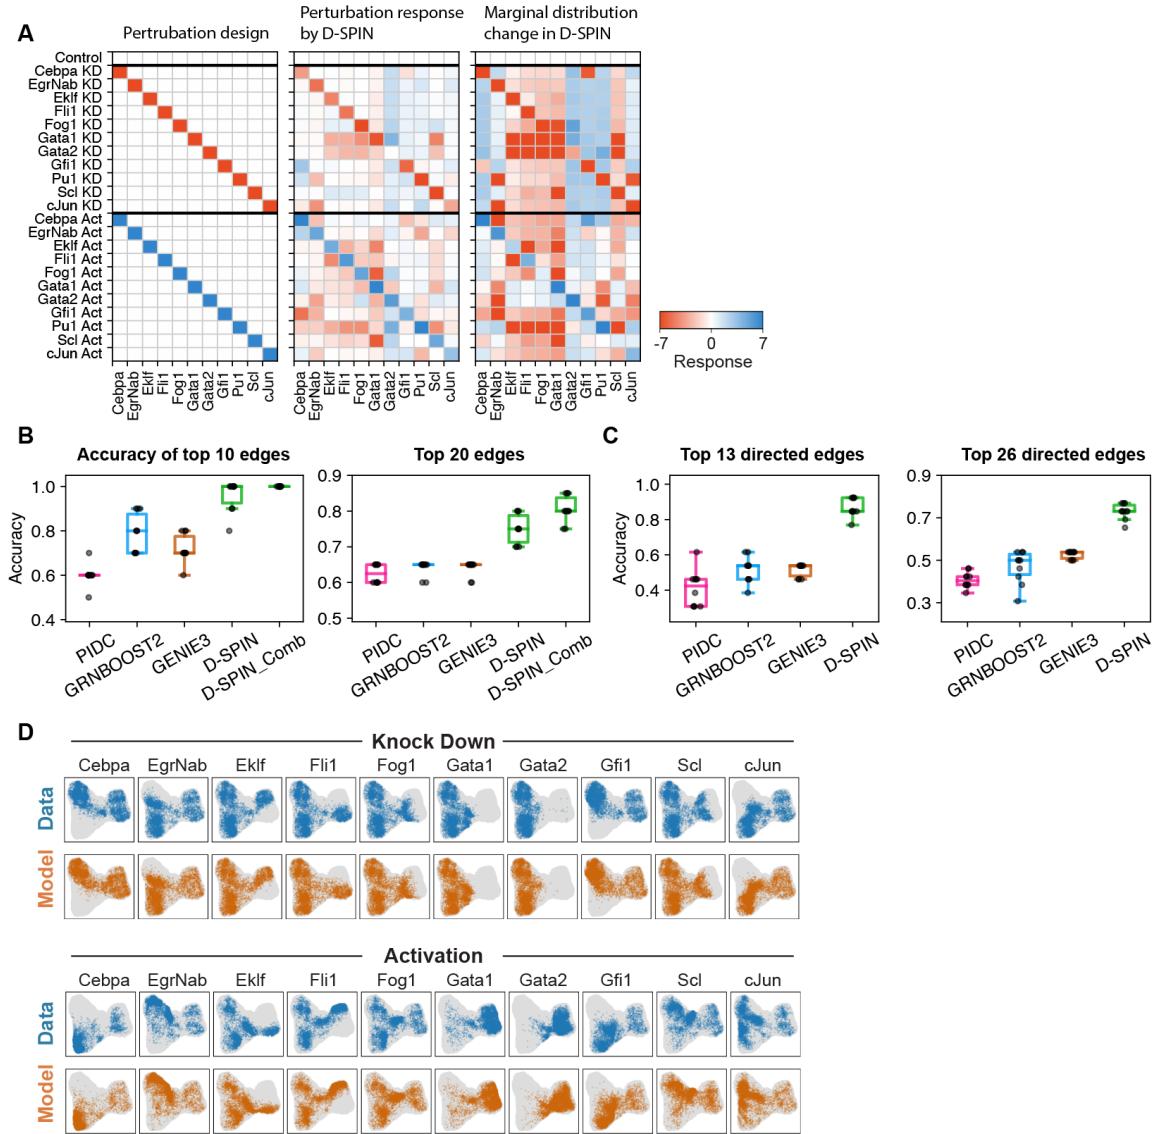


Figure 3.3: D-SPIN identifies perturbation responses and reproduces perturbed cell state distributions (A) Heatmaps indicating (left) applied perturbations to the network; (middle) perturbation response inferred by D-SPIN (right) marginal distribution changes of single genes. The heatmaps demonstrate that the D-SPIN response vectors can identify the applied perturbations and quantify their impact on single genes. The only discrepancy of inferred perturbation occurs on Gata2 knockdown, potentially because Gata2 is a transient regulator in the network, having low expression in differentiated cell states. (B) Bar plots quantify network inference accuracy of D-SPIN with combinatorial perturbations on top-predicted 10 and 20 edges. D-SPIN achieves further accuracy improvements with additional 96 double-gene perturbations, with exact top 10 predictions and 0.805 accuracy in the top 20 predictions.

Figure 3.3: (Continued) (C) Bar plots quantify the directed network inference accuracy of D-SPIN on top-predicted 13 and 26 edges where the ground truth network has 26 directed edges. D-SPIN is capable of inferring a directed network, as detailed in Section 2.4, with superior accuracy (0.73 on top 26 edges) compared with existing methods (0.41, 0.47, 0.53 on top 26 edges). (D) UMAP embedding of cell state distributions of (blue) simulated data by BEELINE and (orange) D-SPIN model distributions across all single-gene knockdown/activation perturbations. The D-SPIN model accurately generates cell-state distributions with the regulatory network and each perturbation across all these conditions.

a perturbation \mathbf{h} , we define the marginal activity \hat{h}_k of each gene k to be the ratio between the marginal probability of the node k being expressed $P(s_k = 1 \text{ given } \mathbf{h})$ or silenced $P(s_k = -1 \text{ given } \mathbf{h})$. Then we evaluate the influence of every edge J_{ij} in the network on \hat{h}_k to generate an edge sensitivity score $\epsilon(J_{ij}, k)$, which is the change of marginal activity \hat{h}_k before and after hypothetical deletion of edge J_{ij} from the network model. The edge sensitivity score allows us to rank the importance of network edges in modulating the response of the network to the knockdown or activation of a specific gene.

For example, in both the D-SPIN model and training data, the knockdown of the transcription factor Pu1 results in a complete loss of monocyte cell states and an increase in granulocyte cell states (Figure 3.2 C). The formal reasoning framework can provide insight into network edges that mediate the response of the network to Pu1 knockdown and the resulting change in population structure. According to the node marginal distribution and edge sensitivity analysis, under Pu1 knockdown the marginal activity of EgrNab and cJun decrease by -12.9 and -6.4. This is primarily induced by the positive interactions of Pu1-EgrNab and Pu1-cJun with the edge sensitivity score $\epsilon(\text{Pu1-EgrNab}, \text{EgrNab}) = -4.4$ and $\epsilon(\text{Pu1-cJun}, \text{cJun}) = -3.2$ (Figure 3.4 (i)). Pu1, EgrNab, and cJun are expressed in the monocyte state (Krumrieck et al., 2011), and the repression of these genes corresponds to the loss of the monocyte cell state. Conversely, under Pu1 knockdown, Gfi1 and Cebpa, genes expressed in granulocyte cell state (Krumrieck et al., 2011), have an increased marginal activity by 3.0 and 2.9, respectively, corresponding to the increased granulocyte cell states. Edge sensitivity scores show that the expression of Gfi1 and Cebpa is impacted both by the positive interactions of Pu1-Gfi1 and Pu1-Cebpa and by the negative interaction of Gfi1-EgrNab. For Gfi1, Gfi1-EgrNab has a stronger impact of $\epsilon(\text{Gfi1-EgrNab}, \text{Gfi1}) = 4.7$ compared with $\epsilon(\text{Pu1-Gfi1}, \text{Gfi1}) = -2.6$. The response of Cebpa is also primarily induced by the edge Gfi1-EgrNab with

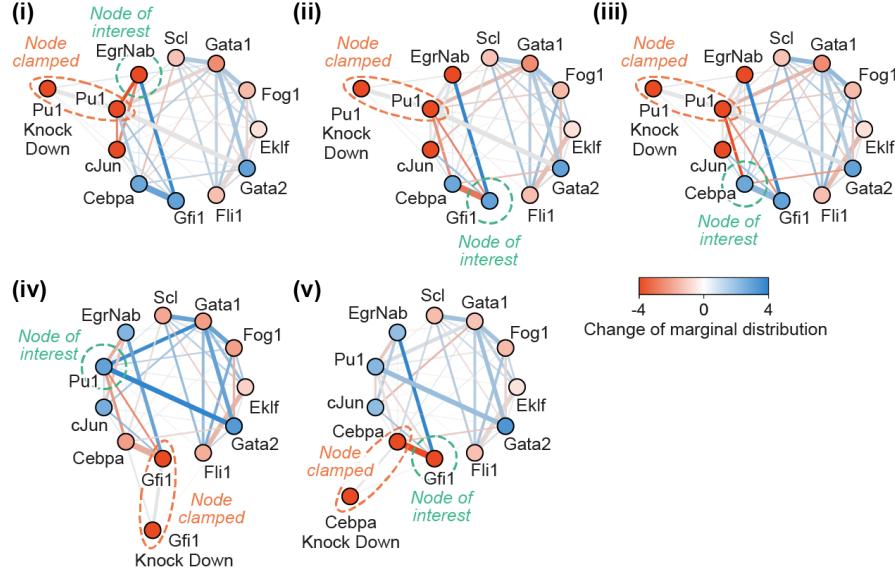


Figure 3.4: Edge sensitivity analysis reveals network-level mechanisms of cell state modulations. Network diagrams visualize edge sensitivity analysis in D-SPIN showing how each network edge positively (blue) or negatively (red) contributes to the marginal distribution change of the gene of interest under a clamped perturbation, as detailed in Section 3.5. Network nodes are colored by the upregulation or downregulation of their activities defined by marginal distributions, and edges are colored by their contribution to the activity of the node of interest. For example, panel (B) shows under Pu1 knockdown, the upregulation of Gfi1 is most significantly positively contributed by the edge EgrNab-Gfi1.

$\epsilon(Gfi1-EgrNab, EgrNab) = 4.3$ compared with $\epsilon(Pu1-EgrNab, EgrNab) = -3.9$ (Figures 3.4 (ii)(iii)). The Mega-Ery module (Eklf, Fog1, Gata1, Scl, and Fli1) is affected by both the negative interaction with Pu1 (which is downregulated) and Cebpa (which is unregulated), and the overall marginal activities exhibit weak inhibition (-0.7, -1.5, -2.6, -1.3, -1.4).

As another example of formal reasoning, in both simulated data and the D-SPIN model, both Pu1 activation and the knockdown of either Cebpa or Gfi1 individually lead to an increased prevalence of monocyte cell states, but with different inferred network mechanisms (Figure 3.2 D). Activating Pu1 causes the increase in marginal activity for other monocyte genes, EgrNab and cJun, by 2.4 and 2.8 which is mediated by the network edges Pu1-EgrNab and Pu1-cJun, with edge sensitivity scores of 4.4 and 3.1. Alternately, Gfi1 has both positive and negative interactions with monocyte genes (Pu1, EgrNab, and cJun), and edge sensitivity analysis on Pu1 under the Gfi1 knockdown condition shows that the negative edge Gfi1-EgrNab has a stronger effect

on Pu1 with score $\epsilon(\text{Gfi1-EgrNab}, \text{Pu1}) = 2.8$ compared with the positive edge Gfi1-Pu1 with a score $\epsilon(\text{Gfi1-Pu1}, \text{Pu1}) = -2.6$. Moreover, the marginal activities of monocyte genes Pu1, EgrNab, and cJun under Gfi1 knockdown increase by 3.0, 2.4, and 2.7, leading to an increased monocyte population (Figure 3.4 (iv)). Similarly, under Cebpa knockdown, the marginal activity of Gfi1 is strongly decreased by -6.4, which is primarily due to the interaction of Cebpa-Gfi1 with an edge sensitivity score of -6.3, therefore leading to a similar cell state change as in Gfi1 knockdown (Figure 3.4 (v)). Pu1 activation induces increased monocyte states through the positive interactions Pu1-EgrNab and Pu1-cJun, while increased monocyte states in Cebpa or Gfi1 knockdown are mediated by the negative edge Gfi1-EgrNab. Therefore similar cell-state distribution change of increased monocyte states can be achieved by two different network-level mechanisms in the D-SPIN model.

Thus, D-SPIN, using simulated data, can infer an accurate and generative regulatory network model of a single-cell population by integrating information across perturbation conditions. The D-SPIN model can be analyzed through formal reasoning—as in this example of tracking monocyte and granulocyte cell state distributions—to hypothesize network edges and network-level mechanisms that might be important for determining how the structure of a cell population shifts when individual network nodes are perturbed. The D-SPIN model demonstrates how activation of Pu1 and knockdown of Cebpa can engage two different sets of network edges to increase the prevalence of monocyte cell states. The inferred regulatory network reveals functional interactions between regulatory genes and provides a map of how the distribution of cell states in the underlying HSC cell population is generated through interactions between internal regulators.

3.4 D-SPIN achieves superior accuracy in large-scale regulatory network inference

To demonstrate the versatility of D-SPIN, we expanded the validation to infer the structure of large-scale simulated networks with hundreds to a thousand nodes, where we performed network inference using pseudolikelihood-based methods (Section 2.4). We explored two distinct network architectures: highly modular networks and Erdős-Rényi (ER) random networks (Figure 3.5 A, Figure 3.6 A). The modular networks are characterized by activating interactions between genes inside the same module, and inhibitory interactions between neighboring modules. Modular networks contain key structural properties of biological gene regulatory networks, where modules of coexpressed genes often perform coherent functions; further,

modular networks present challenges for network inference procedures because gene coexpression makes it difficult to define the precise regulators that control module activation and inhibition (Hartwell et al., 1999; Segal, Shapira, et al., 2003; G. P. Wagner, Pavlicev, and Cheverud, 2007; Jiang, Sivak, and Thomson, 2019). In ER networks, each gene has the same probability of interacting with any other gene; ER networks are considered the most general statistical model for generating unstructured random networks (Watts and Strogatz, 1998; Newman, 2018) and are often used for benchmarking inference methods.

Theoretical results (Gutenkunst et al., 2007; Machta et al., 2013; Jiang, Sivak, and Thomson, 2019) demonstrate that the inference of biological regulatory network models can be complicated by parameter degeneracy, where a set of network models can be consistent with a single set of measurements. Perturbations can be applied to break such degeneracy by driving a regulatory network into configurations that expose differences between network architectures that might generate very similar cell state distributions in an unperturbed context (Jiang, Sivak, and Thomson, 2019). Unlike many existing gene regulatory network inference methods, D-SPIN can naturally integrate and synthesize information collected across many perturbation experiments into a single network model, so that D-SPIN can exploit the information provided by perturbation experiments to increase the accuracy of model inference. As perturbation design for large networks remains an open challenge, we applied random perturbations to each network, where the perturbations are independent, normally distributed random variables on each network node (Section 3.5).

In numerical tests, we found that D-SPIN’s ability to integrate information from perturbations allows the algorithm to surpass the accuracy of other inference methods and to achieve increasing accuracy with increasing numbers of network perturbation experiments (Figure 3.5 B, Figure 3.6 B). The accuracy of network inference is measured by the accuracy of top K edge predictions, where K is the number of edges in the ground truth network. Without perturbation, the top K edge accuracy of D-SPIN is around 0.5, similar to existing methods such as PIDC, GRNBoost2, and GENIE3. However, when perturbation conditions are included, D-SPIN’s accuracy increases while the accuracy of other methods remains constant. By including 800 normally distributed, random perturbations, the accuracy of top K predictions rises from 0.51 to 0.67 on modular networks (Figure 3.5 C), and 0.40 to 0.70 on random networks (Figure 3.6 C). More significantly, the accuracy of top $K/4$ predictions rises from 0.57 to 0.93 on modular networks and from 0.59 to 0.93 on

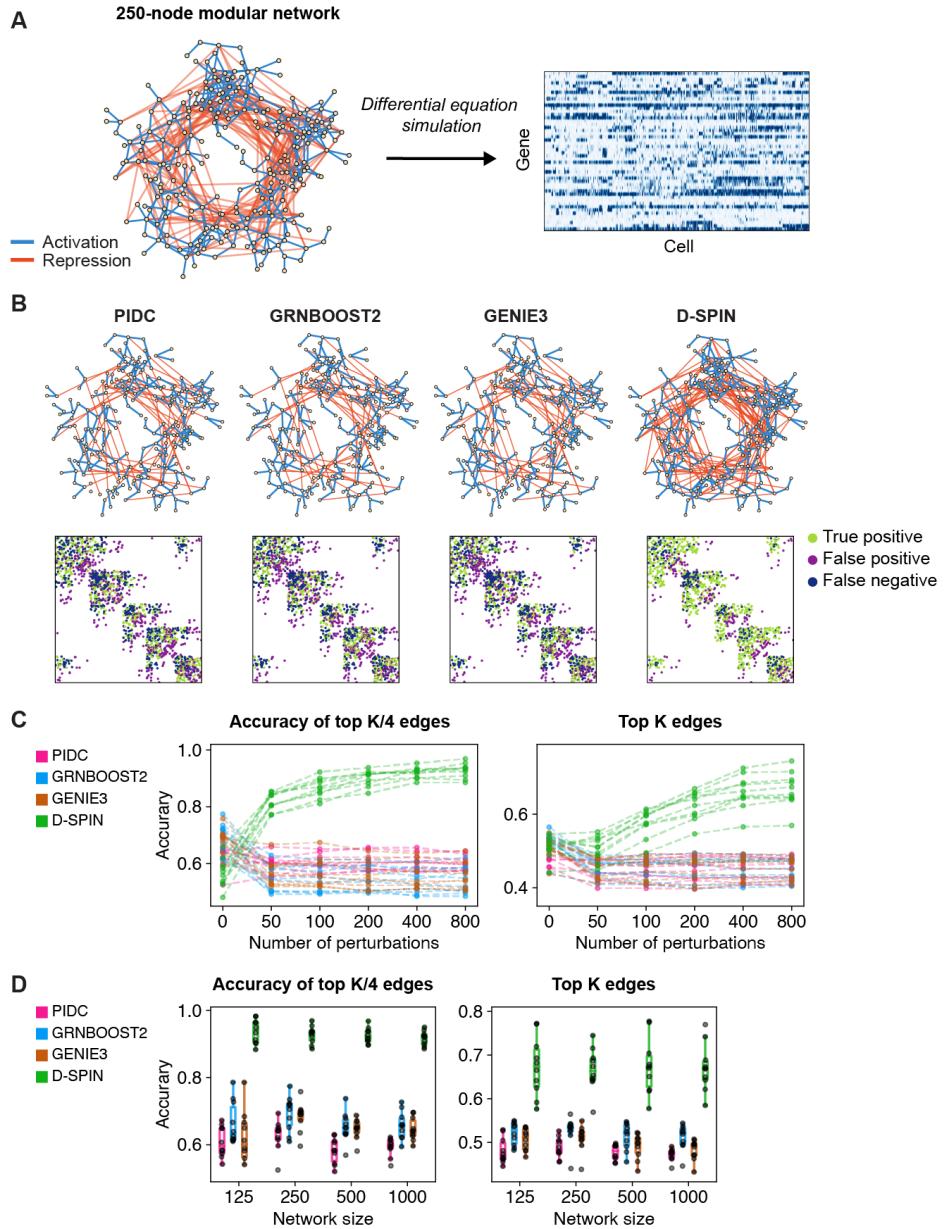


Figure 3.5: D-SPIN achieves superior accuracy on large modular networks with up to a thousand nodes. (A) (top) Network diagram of an example 250-node modular network with five mutually inhibiting modules and (bottom) simulated single-cell data from the network for evaluating the accuracy and scalability of network inference. (B) Diagrams of the inferred network of different methods as in (top) subnetwork of correct-inferred edges and (bottom) adjacency matrices with true positives, false positives, and false negatives. D-SPIN achieves superior accuracy compared to existing methods (C) The network inference accuracies are plotted with the number of perturbations in 10 randomly generated modular networks. The accuracy of D-SPIN continuously increases with perturbation number while other methods do not improve or slightly decrease. (D) Bar plots quantify the network inference accuracy of modular networks with different sizes of 125, 250, 500, and 1000 under 800 perturbations in 10 random repeats. D-SPIN achieves significantly increased accuracy compared with other methods across network sizes.

random networks. In contrast to the increased inference accuracy with perturbation number for D-SPIN, applying the alternative network inference methods (PIDC, GRNBoost2, GENIE3) to each perturbation condition and averaging the results does not increase the accuracy of inference and even slightly decreases the accuracy (Figure 3.5 C, Figure 3.6 C).

We also noticed that in both types of networks, existing methods have much lower accuracy on negative edges. For example, in the modular networks, the average accuracy of PIDC, GRNBoost2, and GENIE3 are all 0.53 on positive interactions, but are 0.37, 0.35, and 0.36, respectively on negative interactions. In D-SPIN with perturbations, the average accuracy is 0.7 on positive edges and 0.8 on negative edges. These results suggest that inhibitory interactions are more difficult to identify and require combining information from different perturbations to discover (Jiang, Sivak, and Thomson, 2019).

Furthermore, the superior accuracy of D-SPIN scales to networks with up to a thousand genes, the largest-sized network that we could test within BEELINE. We performed a similar edge accuracy analysis on modular networks and random networks with sizes of 125, 250, 500, and 1,000 genes, and observed significant accuracy superiority of D-SPIN compared with other methods. On 1,000-node modular networks, the average top $K/4$ and top K accuracy of D-SPIN is 0.92 and 0.67, while for PIDC, GRNBoost, and GENIE3, respectively, the top $K/4$ accuracy are 0.59, 0.66, 0.65, and the top K accuracy are 0.47, 0.51, 0.48 (Figure 3.5 D). On 1,000-node random networks, the average top $K/4$ and top K accuracy of D-SPIN is 0.93 and 0.70, while for the other three methods, the top $K/4$ accuracy are 0.58, 0.62, 0.60, and the top K accuracy are 0.38, 0.40, 0.38, respectively (Figure 3.6 D).

D-SPIN exhibits linear scaling of computation time with cell number

The unique model architecture and training algorithm of D-SPIN enables it to efficiently process large datasets containing hundreds of thousands of cells. The inference in D-SPIN is formulated as an optimization problem and can be solved in batches, thus circumventing the need to handle large matrix operations; as a result, the computation time of D-SPIN only increases linearly with the number of cells. We accessed the scalability of different network inference methods with the number of cells (Figure 3.6 E, Section 3.5). We evaluated the running time of D-SPIN, PIDC, GRNBoost2, and GENIE3 with 250-node Erdős-Rényi (ER) random networks, with cell numbers ranging from 250 to 256,000. While PIDC

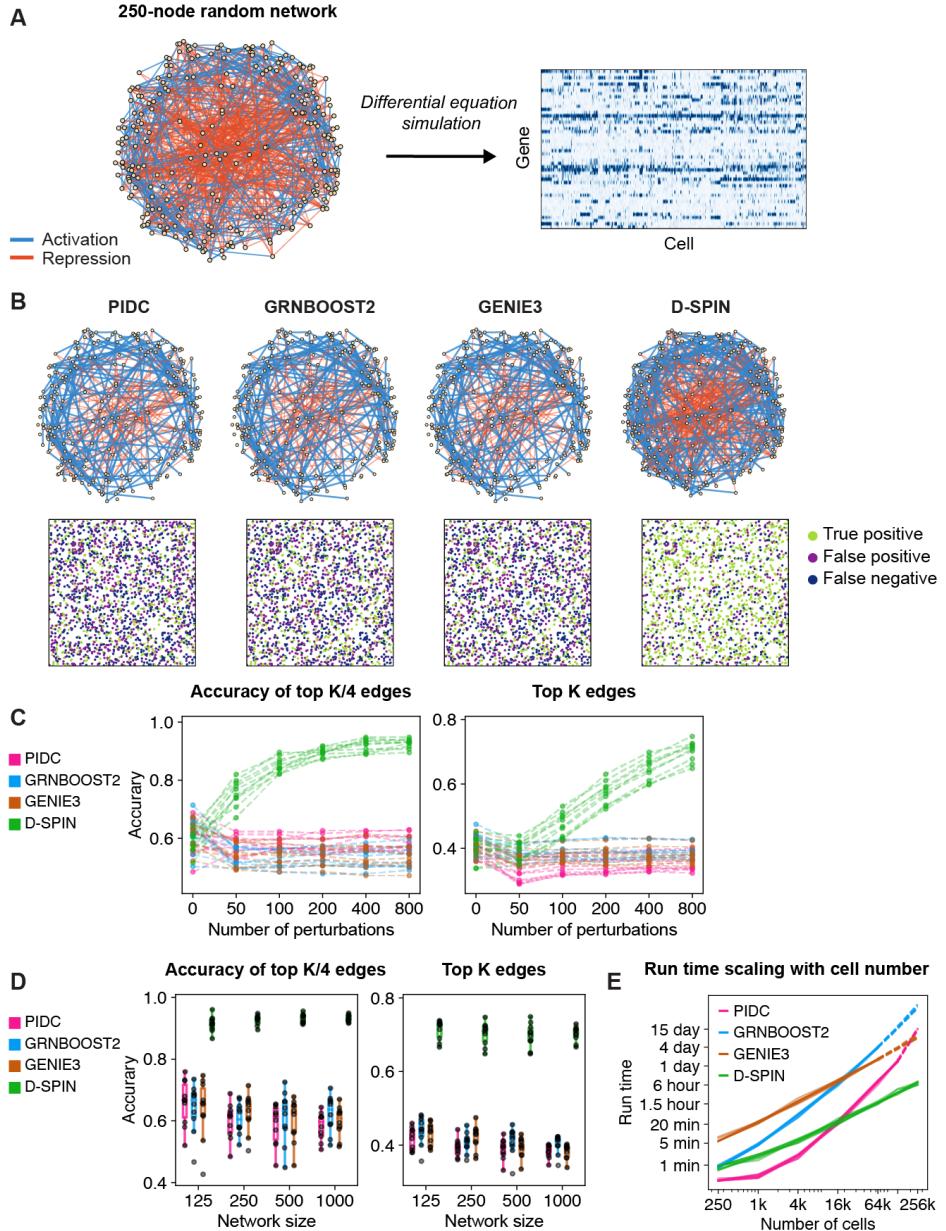


Figure 3.6: D-SPIN achieves superior accuracy on large random Erdős-Rényi networks with up to a thousand nodes. (A) (top) Network diagram of an example 250-node Erdős-Rényi (ER) network with equal probability of activating (blue) and inhibiting (red) edges and (bottom) simulated single-cell data from the network for evaluating the accuracy and scalability of network inference. (B) Diagrams of the inferred network of different methods as in (top) subnetwork of correct-inferred edges and (bottom) adjacency matrices with true positives, false positives, and false negatives. D-SPIN achieves superior accuracy compared to existing methods (C) The network inference accuracies are plotted with the number of perturbations in 10 randomly generated modular networks. The accuracy of D-SPIN continuously increases with perturbation number while other methods do not improve or slightly decrease. (D) Bar plots quantify the network inference accuracy of modular networks with different sizes of 125, 250, 500, and 1000 under 800 perturbations in 10 random repeats. D-SPIN achieves significantly increased accuracy compared with other methods across network sizes.

initially outpaced other methods at lower cell numbers, D-SPIN became the fastest method for datasets exceeding 16,000 cells, a number most current single-cell studies surpass. Remarkably, at 256,000 cells, D-SPIN finished in 6 hours on 2 CPU cores, while all other methods could not finish within a week.

3.5 Supplementary information

Simulations of the HSC network

The BEELINE framework provides a systematic tool to simulate single-gene expression data from Boolean regulatory networks and to benchmark the accuracy of network inference methods (Pratapa et al., 2020). Our simulation and evaluation mostly followed the formulation and code of BEELINE, but we complemented the framework with the simulation of gene knockdown and activation perturbations. BEELINE simulates the following stochastic differential equation (SDE) models of gene expression.

$$\frac{d[x_i]}{dt} = m_i f(R_i) - l_x[x_i] + s\sqrt{[x_i]}\Delta W_t \quad (3.1)$$

$$\frac{d[p_i]}{dt} = r[x_i] - l_p[p_i] + s\sqrt{[p_i]}\Delta W_t \quad (3.2)$$

$$\Delta W_t = \mathcal{N}(0, h) \quad (3.3)$$

where $[x_i]$ and $[p_i]$ are mRNA and protein concentration of the gene i , m_i and r is the mRNA and protein production rate, $f(R_i)$ is the regulatory term determined by the network, l_x and l_p are mRNA and protein degradation rate, s is thermal noise strength, W_t is the Wiener process representing thermal noises, and h is simulation time step. The default parameter values are $m_i = 20$, $l_x = 10$, $r = 10$, $l_p = 1$, $s = 10$. At $t = 0$ all mRNA concentrations $[x_i]$ are initialized as 1, and all protein concentrations $[p_i]$ are initialized as 0. Specifically, the mRNA default initial concentration 1 in the current BEELINE framework (v1.0) is different from 0.01 used in the version (v0.1) presented by the BEELINE paper (Pratapa et al., 2020), possibly for improved numerical stability.

To model the impact of gene knockdown and activation, we included two extra terms into the dynamics of mRNA: e_i , representing a constant external mRNA production not regulated by the network, and β_i , representing the proportion that the mRNA production rate is suppressed in addition to the network regulation. After the modification the dynamics of the mRNA is

$$\frac{d[x_i]}{dt} = e_i + \beta_i m_i f(R_i) - l_x[x_i] + s\sqrt{[x_i]}\Delta W_t \quad (3.4)$$

Without perturbation, we have $e_i = 0, \beta_i = 1$.

In single-gene and double-gene knockdown simulation, we used $\beta_i = 0.2$ for the reduced mRNA production rate. In single gene activation simulation, we set $e_i = 15$, corresponding to 0.75 of the expression of fully activated genes. For double perturbations, we found that higher external activation was required to drive the network, and we used $e_i = 22.5$ for gene activation in double perturbations. We simulated 2,000 cells for control conditions and 500 cells for each perturbation condition.

We simulated the system under 10 different sets of randomly sampled parameters. Each time the parameters m_i, l_x, l_p, r were assigned by the formula αD , where D represents the default value of the parameter and α is a random factor sampled from a truncated Gaussian distribution (mean = 1, std = 0.1, truncation to [0.9, 1.1]). The same type of parameters of all genes were set to the same value for numerical stability. We clustered the cell states and computed the UMAP embedding of the discretized simulated data with D-SPIN-generated samples for visualization, and annotated each cluster based on the research that proposed the HSC network model (Krumsieck et al., 2011).

To compare the network reconstruction accuracy with other methods, including PIDC, GRNBoost2, and GENIE3, we ran the Docker images of the three methods provided by the BEELINE framework. As the three methods do not explicitly model the impact of perturbations, we ran these methods independently on control and each perturbation condition and computed the average score of each network edge as the overall inferred network.

Clamping analysis of node marginal distribution and edge sensitivity

As a generative model of transcriptional state distribution, D-SPIN allows us to develop a procedure to compute perturbation impact on the marginal distribution of a node in the network, as well as the contribution of each interaction edge to the marginal distribution. The procedure allows us to identify critical interactions in the network that determine a specific perturbation response so that we can make quantitative statements about how cell state distribution and perturbation response are determined by the regulatory network.

To quantify the marginal distribution change of each node i , we define an effective marginal activity $\hat{h}_i(\mathbf{J}, \mathbf{h})$ by the ratio between the node being on ($s_i = 1$) and the

node being off ($s_i = -1$) in the marginal distribution $P(s_i|\mathbf{J}, \mathbf{h})$

$$\hat{h}_i(\mathbf{J}, \mathbf{h}) = \log \frac{P(s_i = 1|\mathbf{J}, \mathbf{h})}{P(s_i = -1|\mathbf{J}, \mathbf{h})}, \quad (3.5)$$

where positive \hat{h}_i indicates the node being activated and negative \hat{h}_i indicates the node being repressed. The definition is rationalized by parametrizing the complete three-state marginal distribution of node i by an effective 1-node D-SPIN model

$$P(s_i) = \frac{1}{Z} \exp(\hat{J}s_i^2 + \hat{h}_i s_i), \quad (3.6)$$

whose analytical solution is $\hat{h}_i = \frac{1}{2} \log P(s_i = 1)/P(s_i = -1)$, and we drop the constant coefficient for simplicity. As \hat{h}_i is defined as a linear term inside the energy function, the difference of \hat{h}_i between conditions represents the relative activation or inhibition of a node.

To quantify the contribution of each edge to the marginal distribution change, we compare the marginal distribution with and without the edge of interest. The edge sensitivity of J_{kl} to node i is defined as the marginal activity \hat{h}_i change after removing the edge J_{kl} from the network. Specifically, we define a clamped network $\mathbf{J}_{(\text{set } J_{kl}=0)}$ where J_{kl} is set to 0 while all other edges remain the same, and compute the marginal activity $\hat{h}_i(\mathbf{J}_{(\text{set } J_{kl}=0)}, \mathbf{h})$ of node i . The edge sensitivity $\epsilon(kl, i)$ is thus defined by

$$\begin{aligned} \epsilon(kl, i|\mathbf{J}, \mathbf{h}) &= \hat{h}_i(\mathbf{J}, \mathbf{h}) - \hat{h}_i(\mathbf{J}_{(\text{set } J_{kl}=0)}, \mathbf{h}) \\ &= \log \frac{P(s_i = 1|\mathbf{J}, \mathbf{h})}{P(s_i = -1|\mathbf{J}, \mathbf{h})} - \log \frac{P(s_i = 1|\mathbf{J}_{(\text{set } J_{kl}=0)}, \mathbf{h})}{P(s_i = -1|\mathbf{J}_{(\text{set } J_{kl}=0)}, \mathbf{h})} \end{aligned} \quad (3.7)$$

With the HSC network example, we demonstrate how this procedure identifies the impact of Pu1 knock-down on the node Gfi1, and which network edge contributes to the change of Gfi1. D-SPIN infers the network \mathbf{J} , response vector \mathbf{h}_0 without perturbation and \mathbf{h} under Pu1 knock-down. Without perturbation, by marginalizing the full distribution $P(\mathbf{s}|\mathbf{J}, \mathbf{h}_0)$, we have

$$P(\text{Gfi1} = -1) = 0.884, P(\text{Gfi1} = 1) = 0.024, \text{ thus } \hat{h}_0 = -3.60.$$

Similarly, under Pu1 knock-down, marginalizing the complete distribution $P(\mathbf{s}|\mathbf{J}, \mathbf{h}_0)$ gives

$$P(\text{Gfi1} = -1) = 0.381, P(\text{Gfi1} = 1) = 0.216, \text{ thus } \hat{h} = -0.566.$$

Therefore Pu1 knock-down has an overall impact of 3.04 activation on Gfi1. For edge sensitivity, take the edge Gfi1-EgrNab, for example. After clamping the edge to

0, the complete distribution $P(\mathbf{s}|\mathbf{J}_{(\text{set Gfi1-EgrNab} = 0)}, \mathbf{h})$ has an marginal distribution of

$$P(\text{Gfi1} = -1) = 0.904, P(\text{Gfi1} = 1) = 0.004, \text{ thus } \hat{h}_{(\text{set Gfi1-EgrNab} = 0)} = -5.26.$$

Therefore the edge sensitivity $\epsilon(\text{Gfi1-EgrNab}, \text{Gfi1}) = \hat{h} - \hat{h}_{(\text{set Gfi1-EgrNab} = 0)} = 4.69$. The edge Gfi1-EgrNab has a high contribution in activating the node Gfi1. The contribution of other edges can be computed similarly, and the top contributing edges to Gfi1 under Pu1 knock-down are Gfi1-EgrNab 4.69, Cebpa-Gfi1 -2.85, Gfi1-Pu1 -2.60 and Gfi1-cJun 2.31.

Simulation of large synthetic networks

To assess the accuracy and scalability of D-SPIN, we constructed two distinct sets of large regulatory network models, each comprising networks with 125, 250, 500, and 1000 nodes. These models represent two major categories of networks:

1. Modular Networks: These networks are organized into modules, characterized by activating interactions inside each module and inhibitory interactions between modules. Consequently, gene expression of these networks also exhibits a modular pattern, where genes within the same module are more likely to be activated at the same time.
2. Random Networks: One of the most representative random networks are Erdős-Rényi models (ER models). These networks are completely random in nature, where each node has the same probability of interacting with any other node.

Specifically, in our construction of modular networks, the networks contain Q modules that form a loop by connecting neighboring modules on a circle. The average module size is set to around 50 nodes, and the total number of modules should be an odd number to allow the pattern of alternatively activated modules on the loop. Therefore the module numbers Q for 125, 250, 500, and 1000-node networks are, respectively, 3, 5, 11, and 21. The exact number of nodes in each module is determined by sampling Poisson random numbers, with the distribution expectation being the expected module size. Within each module, each node has a 0.5 probability of being assigned as a core node. Each core node has a 0.0375 probability of activating another node in the same module, and a 0.05 probability of inhibiting another core node in neighboring modules. Each non-core node has

a 0.0125 probability of activating another node in the same module. Besides, following the structure of the classical toggle switch architecture, all inhibitory interactions are symmetric, and core nodes that have inhibitory interactions also have self-activating interactions.

The construction of a random network is straightforward. To ensure the same average degree of network nodes across different network sizes, each node has a $1.25/M$ probability to activate any other nodes and $1.25/M$ probability to inhibit any other nodes, where M is the number of nodes in the network. We ensure the network is fully connected by regenerating the network with a different random seed when the network has a disconnected structure.

Using these synthetic networks, we simulated the dynamics of each model with our modified BEELINE framework. Given the complexity and scope limitations of designing combinatorial perturbations on large networks, we opted for random perturbations, where each node was subject to a random perturbation following a truncated Gaussian distribution. For each node i , the perturbation h_i^* was taken from a Gaussian distribution with a mean of 0, a standard deviation of 0.4 and truncated to $[-1, 1]$. These perturbations dictated the biochemical parameters in the simulation by

$$e_i = \begin{cases} 20h_i^* & h_i^* > 0, \\ 0 & h_i^* \leq 0. \end{cases} \quad \beta_i = \begin{cases} 1 + h_i^* & h_i^* < 0, \\ 0 & h_i^* \geq 0. \end{cases} \quad (3.8)$$

For a perturbation $h_i^* = 1$, the strength of external activation is the same as the strongest internal expression in the network. For a perturbation $h_i^* = -1$, the gene i is completely deactivated. Therefore this perturbation assignment covered a spectrum of perturbation actions for accessing the capability of D-SPIN in constructing regulatory networks from perturbations.

Benchmarking time consumption with cell number

To assess the scalability of each method relative to cell number, we performed time benchmarks for D-SPIN together with PIDC, GENIE3, and GRNBoost2, using the dataset of 10 different random repeats of 250-node Erdős-Rényi networks from Section 3.5. Specifically for D-SPIN, the gradient evaluations were performed by batches for efficient computation, with each batch containing 2,000 cells. For PIDC, GENIE3, and GRNBoost2, we used the Docker images of the three methods provided by the BEELINE framework. The analyses were performed on the Caltech Resnick cluster, with each task taking 2 cores and 48G memory. The longest running

time of each task was 7 days. The time of unfinished tasks was interpolated using the running time of the same methods with smaller cell numbers using the function $y = Ax^B + C$ in the log-log space.

DISSECTING HOMEOSTASIS MAINTENANCE STRATEGIES FROM GENOME-WIDE PERTURB-SEQ

4.1 Introduction

The remarkable ability of cells to maintain homeostasis in response to external stimuli and perturbations is fundamental to their survival. Theories and models of such robustness have been established from various biological processes including chemotaxis, cell cycle, and development control, with design principles discovered (Stelling et al., 2004). However, our knowledge of the robustness is largely focused on specific functions controlled by local biomolecular circuits, while how the global organization of the regulatory network maintains the homeostasis of the cell remains unclear due to the lack of systematic experimental profiling and formal modeling frameworks.

Classic research in genetics has demonstrated that gene perturbation experiments can be used to reconstruct the architecture of gene regulatory networks by observing how the knockdown/activation of individual genes impacts the expression or activity of other genes (Ferguson, Sternberg, and Horvitz, 1987; L. S. Huang and Sternberg, 1995; Yuh, Bolouri, and Davidson, 1998; Z. Hu, Killion, and Iyer, 2007). Moreover, the cells being viable after perturbation provides operating examples of homeostasis maintenance of cell states. One common mechanism is genetic redundancy, where the loss of a gene function can be compensated by another gene with partial overlap of functions, typically homologs (Nowak et al., 1997; Kafri, Springer, and Pilpel, 2009). Nonetheless, the inhibition or activation of certain genes can induce global state transitions of cells, even completely reprogram the cell fate such as induced pluripotent stem cells by the Yamanaka factors Myc, Oct3/4, Sox2, and Klf4 (Takahashi and Yamanaka, 2006). It is yet unclear how the strategies of local redundancy and global state transition are deployed by the cellular regulatory networks to maintain homeostasis under perturbations.

The development of Perturb-seq enables highly parallelized profiling of gene perturbation effect at the transcriptome level. Perturb-seq introduces a large library of guide RNAs into a cell population, where each guide RNA targets a specific gene for editing via the CRISPR/Cas9 system. The identities of received perturbation

are simultaneously read out through single-cell mRNA sequencing (Dixit et al., 2016; Schraivogel et al., 2020; Replogle et al., 2022). Perturb-seq experiments can interrogate the impact of thousands of distinct gene knockdowns on cellular gene expression in parallel. D-SPIN provides a computational framework for the integration of data from thousands of perturbation experiments into a unified regulatory network model to generate insights into the architecture of cellular information flow and interactions between pathways.

In this chapter, we apply D-SPIN to dissect the response to gene knockdowns of a cancerous cell line, K562 cells. With a public dataset of genome-wide Perturb-seq, we build a global regulatory network model of the cell and partition the identified perturbation response into groups. The perturbation response groups reveal the organization of cellular pathways and the function partitioning of molecular complex subunits. Moreover, the regulatory network model suggests global strategies of homeostasis maintenance, where the loss of a gene function is compensated by the upregulation of another gene function that is not directly associated. Such long-range regulation suggests the existence of cellular information-processing mechanisms that sense the encountered stress, and coordinate other physiological functions to respond.

4.2 Constructing a regulatory network model from genome-wide Perturb-seq data

We used D-SPIN to construct a gene regulatory network model using data from a genome-wide Perturb-seq experiment that used CRISPRi to knock down genes in the human chronic myelogenous leukemia (CML) K562 cell line (Replogle et al., 2022). K562 is an erythrocyte precursor cell line with a BCR-ABL mutation, derived from a CML patient (C. B. Lozzio and B. B. Lozzio, 1975; Heisterkamp et al., 1983). In the genome-wide Perturb-seq experiment, 9,867 genes were knocked down individually across 2 million single cells using CRISPRi. Prior to D-SPIN analysis, we pre-processed the data to identify guide RNAs that have an impact on the K562-cell transcriptional state (Section 4.6). Guide RNA perturbation impact varies (Figure 4.1 A): 70.5% of perturbations had fewer than 10 differentially expressed genes (DEGs), 14.7% had 10–100 DEGs, and 14.8% had more than 100 DEGs. There is also great variation in the number of collected cells: 20.3% of perturbations had fewer than 100 cells collected.

To assess the network reconstruction quality of D-SPIN, we compared the network

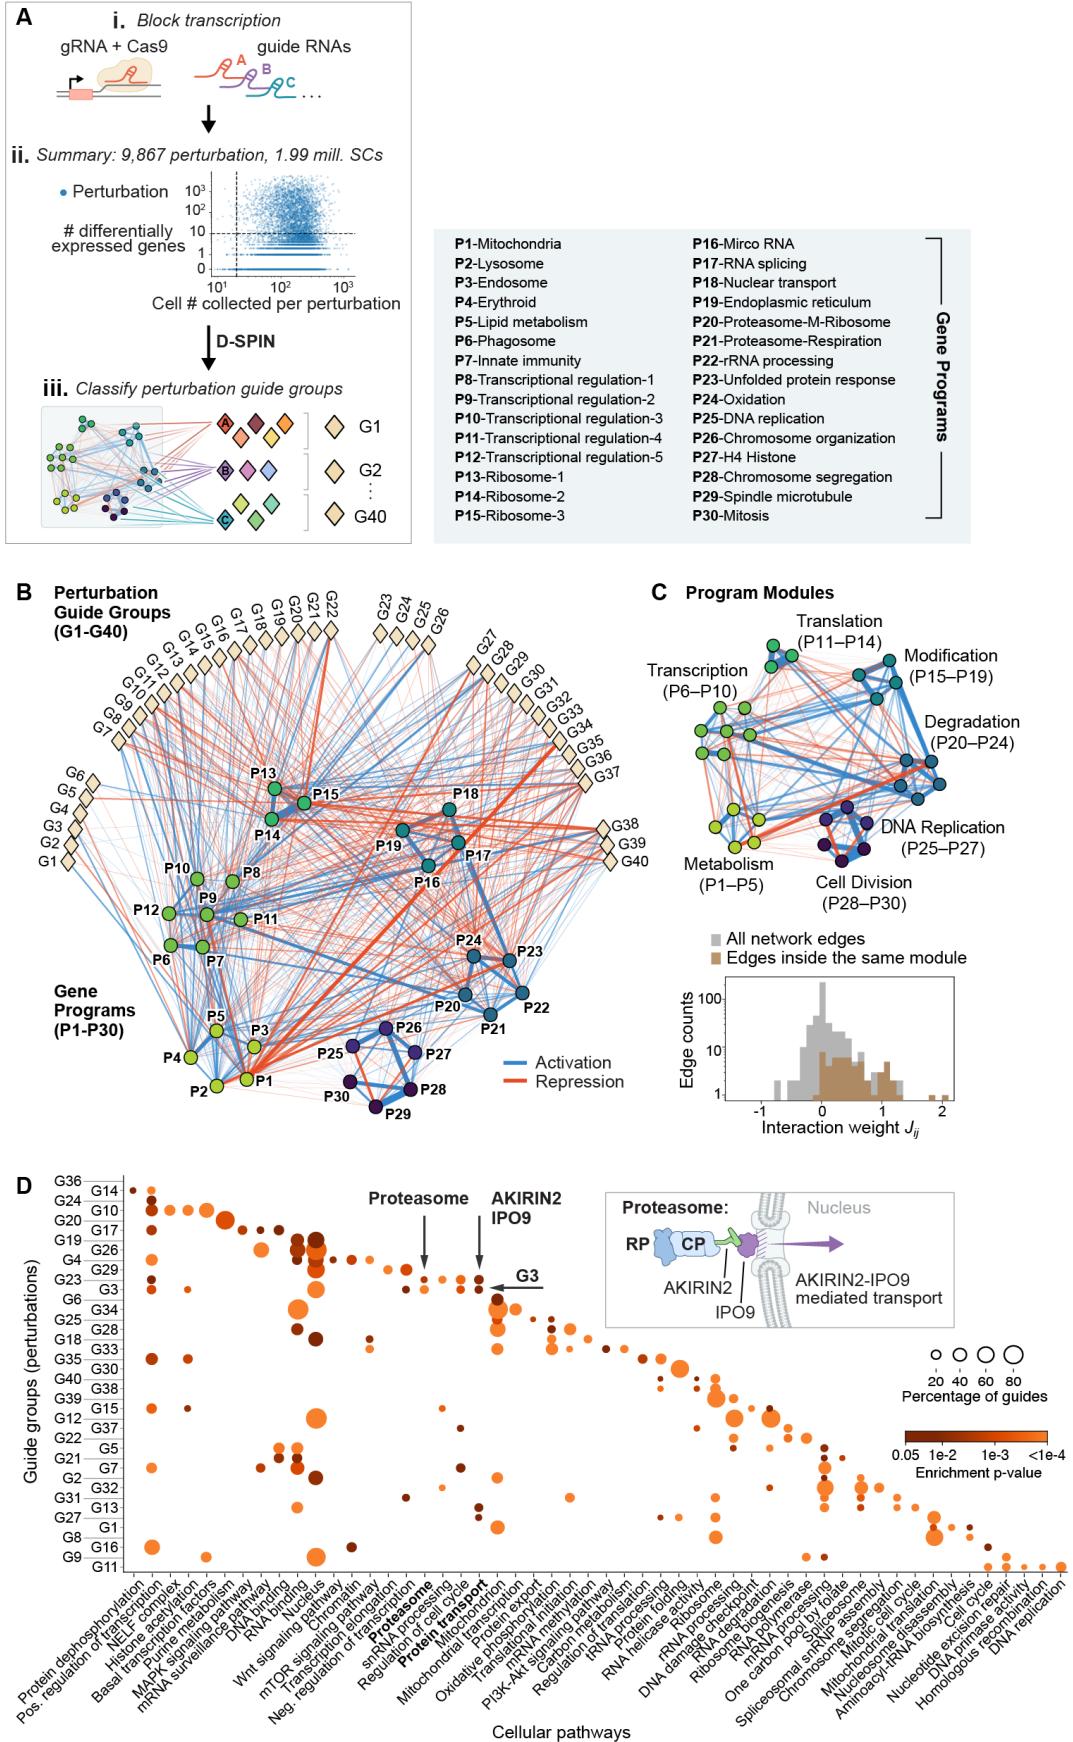


Figure 4.1: Caption on the following page.

Figure 4.1: D-SPIN generates a regulatory network model from genome-wide Perturb-seq data that reveals the organization of cellular pathways. (A) Schematics of the experiments of the genome-wide Perturb-seq dataset (Replogle et al., 2022). (i) 9,867 gene knockdown perturbations are individually delivered to $\sim 2M$ human K562 cells using guide RNAs. (ii) Scatter plots of the number of differentially expressed genes (DEGs) and the number of cells that passed quality filtering for each perturbation. D-SPIN model construction was applied to 3,136 perturbations that induced more than 10 DEGs with more than 20 cells. (iii) The D-SPIN model constructs a unified regulatory network model and inferred interactions between each perturbation and the network, which classify gene-knockdown perturbations into groups. (B) Network diagram shows the inferred regulatory network model (J matrix) between 30 gene-expression programs (P1–P30, circles), as well as the interactions between programs and 40 groups of gene-knockdown perturbations (G1–G40, diamond) encoded in response vectors (h vectors). Interactions are rendered as positive (blue) or negative (red) edges with thickness scales with the strength of interactions. (left box) Gene programs are functionally annotated through gene ontology annotation tools (Kuleshov et al., 2016; Szklarczyk et al., 2021; Sherman et al., 2022) and manual lookup. (C) (top) Network diagram of the regulatory network model by D-SPIN exhibits modular structure with tightly-connected subnetworks. The seven network modules can be assigned physiological functions based on the gene programs they encompass. (bottom) The histogram quantifies the distribution of all network edges (J matrix entries) and edges inside the same module. Edges inside modules are mostly positive interactions and contain the majority of strong positive interactions in the network. (D) Bubble plot shows enriched cellular pathways in gene ontology and KEGG pathway from databases (Sherman et al., 2022) for each guide RNA group where bubble size represents the percentage of enriched guides in the group. Gene targets in the same guide group identified by D-SPIN are involved in similar pathways or potentially have interactions, revealing associations between pathways, such as (arrow pointer)protein transport genes AKIRIN2 and IPO9 that were recently found to (inset schematics) mediate the import of proteasome into the nucleus (Almeida et al., 2021).

inference with independent biological measurements from ChIP-seq databases (Z.-P. Liu et al., 2015; Han et al., 2018; Garcia-Alonso et al., 2019). ChIP-seq data measure the binding of transcription factors (TFs) to target gene promoters and have been used in the literature to assess the quality of inferred gene regulatory network models. We applied a gene selection procedure from the BEELINE benchmarking framework (Pratapa et al., 2020) where we selected two gene subsets TFs+500 and TFs+1000 that contain filtered TFs and the top 500 or 1000 highly variable genes, respectively (Section 4.6). We reconstructed the single-gene regulatory networks and compared the networks with TF-target regulation data obtained from ChIP-seq databases (Figure 4.2 A, Section 4.6). The correspondence between inferred networks and ChIP-seq data is quantified by the early precision rate (EPR), which measures the proportion of inferred edges that can be associated with a TF binding event detected by ChIP-seq among the top K predictions, relative to the expectation of a random predictor. K is set to the total number of TF-target interactions in the ChIP-seq database. D-SPIN achieves leading correspondence with ChIP-seq data compared with other methods, with top K EPRs of 4.48 and 4.36 for the TFs+500 and TFs+1000 datasets, respectively, compared to 2.88 and 2.68 in PIDC, 3.58 and 3.92 in GRNBoost2, and 3.39 and 3.13 in GENIE3.

Interpreting a single-gene network with hundreds or thousands of nodes can be complex. To increase the interpretability of the D-SPIN model, we construct a regulatory network model at the gene-program level to reveal the global architecture and logic of the regulatory network implied by transcriptional changes induced across a large number of gene knockdown experiments. Transcriptome-scale gene expression changes in response to gene knockdown perturbation reflect how cells coordinate the activities of major cellular functions to maintain homeostasis under the stress induced by gene knockdowns. To reveal the global architecture of perturbation responses, we focused on 3,138 gene-knockdown perturbations that are each associated with more than 10 differentially expressed genes (DEGs) (Section 4.6). We coarse-grained the transcriptional profile into 30 gene programs, a number informed by both the Bayesian information criterion (BIC) and the elbow method, to balance the model’s representative power and complexity (Hastie et al., 2009) (Section 2.5). With 30 gene programs, 93.2% of the selected highly variable genes have a correlation larger than 0.5, with their corresponding gene programs in the gene matrix denoised by data diffusion (Figures 4.3 AB) (Van Dijk et al., 2018).

The extracted gene programs reflect both general cell biology and lineage-specific

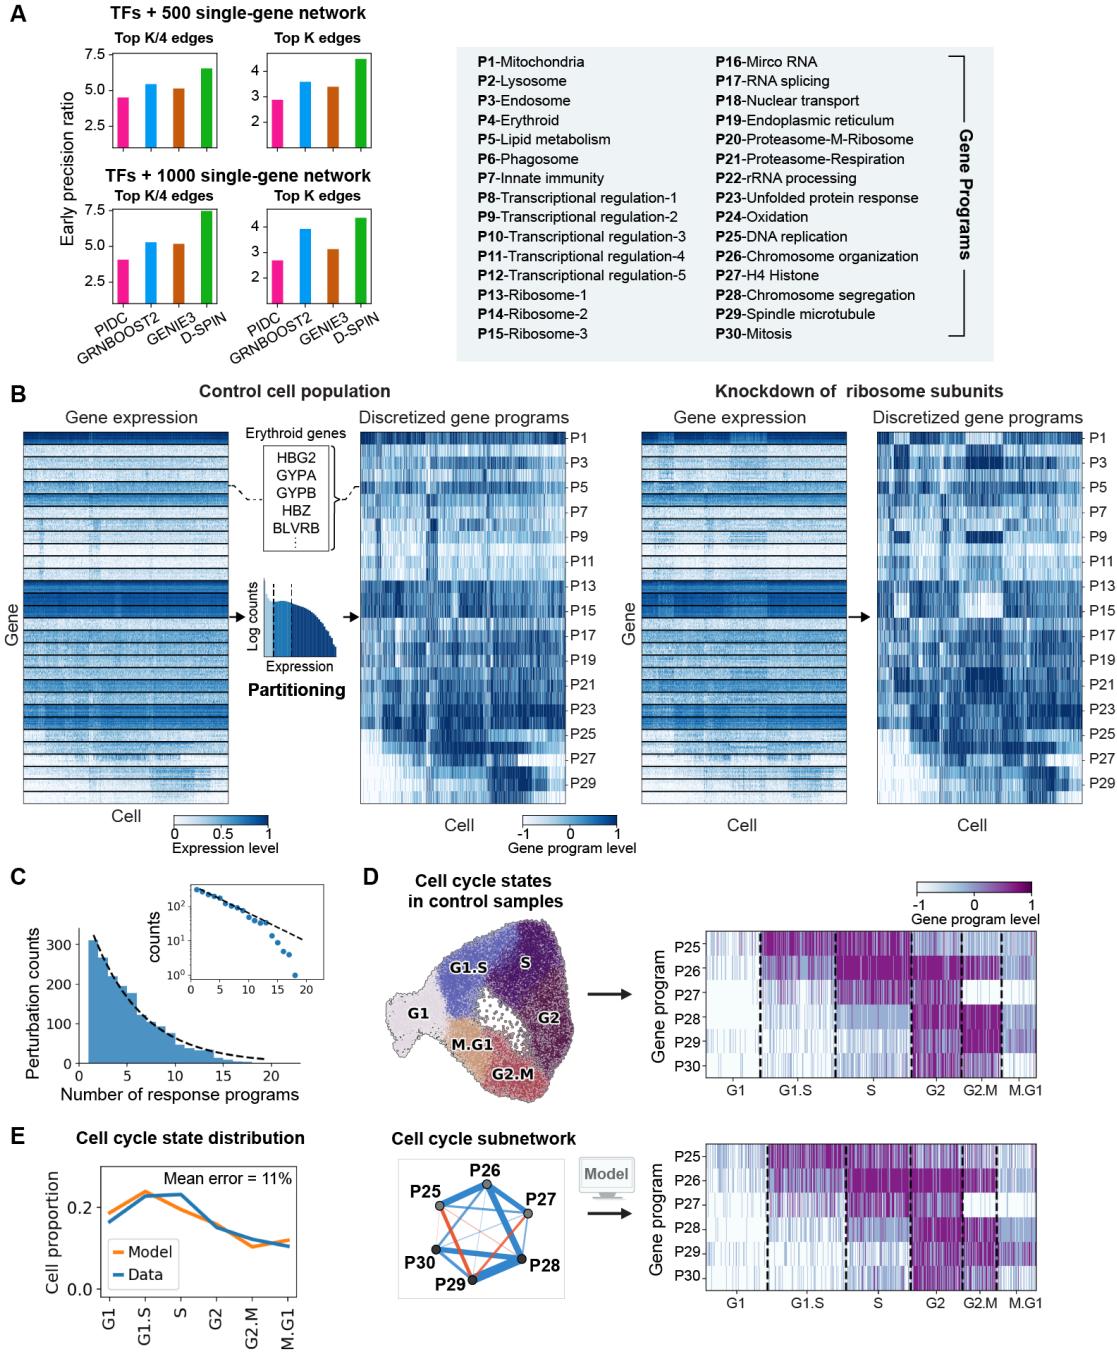


Figure 4.2: Evaluating single-gene network; D-SPIN reconstructs cell-cycle states. (A) Correspondence of D-SPIN and ChIP-seq data for single-gene network inference on two networks, TFs+500 and TFs+1000, derived from the Perturb-seq dataset. The correspondence is quantified early precision ratio, the fraction of top-inferred edges that can be associated with a known transcriptional factor (TF)-target binding data from ChIP-seq database compared with a random predictor, as detailed in Section 4.6. D-SPIN achieves better correspondence with ChIP-seq data on single-gene regulatory network inference compared with existing methods.

Figure 4.2: (Continued) (B) Heatmaps of gene expression and discretized gene program level for (left) control samples of non-targeting guide RNAs and (right) knockdown of the large or small ribosome subunits. The expression of each gene is normalized by its maximum expression and capped at 0.8 for visualization. The heatmaps demonstrate that discretized gene programs characterize and denoise major expression patterns in the gene matrix. (right box) Gene programs are functionally annotated through gene ontology annotation tools (Kuleshov et al., 2016; Sherman et al., 2022) and manual lookup. (C) Histogram showing the number of responding gene programs for each gene knockdown perturbation on a linear scale and (inset) log scale. The responding program number is exponentially distributed, suggesting that perturbations influencing a large number of perturbations are relatively rare. (D) (left-top) UMAP embedding of cell-cycle-related gene programs in control samples exhibits a circular structure. The clusters of cell-cycle states are annotated and ordered based on the gene program expression and sequential positions on the UMAP. (left-bottom) Network diagram of the D-SPIN cell-cycle network, which is taken as the subnetwork of 6 cell-cycle programs from the full 30-program network. (right-top) Heatmaps of cell-cycle programs for cell states from the Perturb-seq control samples and (right-bottom) the D-SPIN model under the control condition. The heatmaps demonstrate that D-SPIN reconstructs the distribution of cell states associated with cell cycle progression from the 6-node subnetwork. (E) The distribution of cell-cycle states in generated data from the D-SPIN model and experimental Perturb-seq control samples. The distribution generated by D-SPIN quantitatively agrees with the data distribution with an 11% mean error.

gene programs that are specific to the biology of K562 cells (Figure 4.1 B, Figure 4.2 B). We annotate the gene programs with a combination of bioinformatic databases including DAVID, Enrichr, and manual lookup (Kuleshov et al., 2016; Szklarczyk et al., 2021; Sherman et al., 2022). Through oNMF, we extracted a series of programs that are enriched for genes involved in specific core cell processes including transcription, translation, DNA replication, mitosis, and RNA processing. D-SPIN also identified lineage-specific programs including an erythroid-specific program with hemoglobin (HBG1, HBG2, HGZ) and glycophorin (GYPA, GYPB, GYPE) genes, as well as two myeloid-associated programs with phagosome/actin-organization (ACTB, ACTG1, ARPC3) and immune-response (LAPTM5, VASP, RAC2) genes, respectively, which agrees with the multi-lineage potential of K562 cells (Figure 4.1 B).

Given the set of 30 gene programs, D-SPIN generated a gene regulatory network model that provides a wiring diagram of K562 internal regulation and the response of the K562 regulatory network to gene knockdown perturbations (Figure 4.1 B,

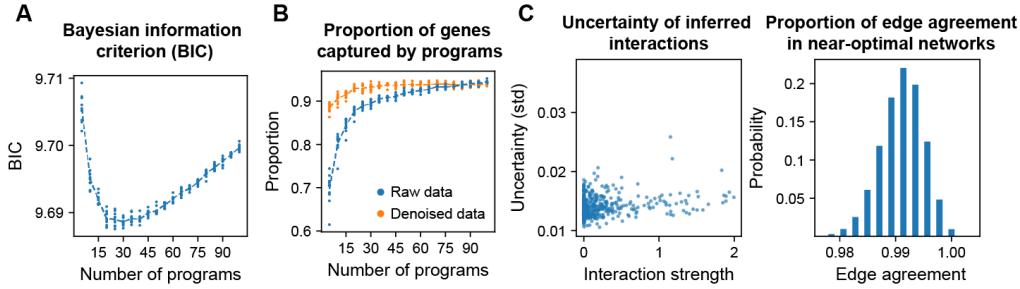


Figure 4.3: The choice of gene-program number; Uncertainty of the inferred network. (A) Bayesian information criterion (BIC) of the gene program decomposition as a function of gene program number. 30 is the optimal program number that minimizes BIC when the effective sample size is 20k cells. (B) The proportion of genes that have a high correlation with their corresponding gene programs as a function of gene program number. For raw data (without application of MAGIC), a high correlation is defined as exceeding 4 standard deviations of gene-gene correlation distribution. In the denoised data by MAGIC (Van Dijk et al., 2018), a high correlation is defined as larger than 0.5. The plot shows that 30 programs is an elbow point of information gain, as detailed in Section 2.5. (C) (left) Uncertainty of inferred network interactions as a function of interaction strength quantified by the standard deviation among alternative network solutions, as detailed in Section 2.5. The uncertainty is generally below 0.02 and much lower compared with interaction strengths. (right) Histogram showing the proportion of edge agreement between alternative solutions of the D-SPIN network inference problem. Alternative solutions are obtained by sampling from the Bayesian posterior distribution given the experimental data, as detailed in Section 2.5. Alternative solutions of D-SPIN are highly consistent with the inferred network, with typically more than 98% of the edges being in the same category of activation, inhibition, and non-interacting.

Figure 4.2 C). The network construction took fewer than 8 hours with parallelization on 500 CPU cores. The D-SPIN network contains gene programs as nodes, and interactions between gene programs as edges that represent inferred interactions. The inferred network has low edge uncertainty quantified by Fisher information, as edge uncertainties are generally below 0.02 while strong interaction strengths are above 1. More than 98% of edges remain structurally stable when resampled from the posterior distribution of networks that represent near-optimal solutions to the network inference problem (Figure 4.3 C, Section 2.5). D-SPIN identified 84 positive and 21 negative interactions between gene programs, corresponding respectively to 19.3% and 4.8% of all gene-program pairs. The D-SPIN regulatory network reconstruction contains a set of sub-networks or modules of gene programs that interact with one another through dense positive (blue) interactions (Figures 4.1 BC).

We automatically decomposed the D-SPIN network into seven modules through the Leiden community detection algorithm, and the module partition is robust to the hyperparameter choice (Traag, Waltman, and Van Eck, 2019) (Section 4.6). Each identified module reflects a core cellular function, including transcription, translation, protein-degradation, and cell-cycle-related functions.

The network contains a set of negative interactions between gene programs that are expressed in distinct cellular states. The strongest negative interaction is between gene programs expressed at different stages of mitosis. For example, the P29 Spindle microtubule has negative interactions with both P25 DNA replication and P27 Histone. The strong interactions inside the cell-cycle sub-network of P25-P30 are able to reconstruct the transcription state distribution during cell cycle progression (Figures 4.2 DE). We also find strong negative interactions between distinct stress-response programs, P23 Unfolded protein response, and P2 Lysosome. The Unfolded protein response program consists of heat-shock protein (HSP90AB1, HSPD1, HSP90AA1) and other molecular chaperones (CCT8, CCT3, CCT5 genes), while the Lysosome program contains genes involved in lysosome genesis (PSAP, ASAHI, PLD3). We also observe negative interactions between P4 Erythroid and P6 Phagosome, where the phagosome is an innate immune response mounted by macrophages and other myeloid cell types (H.-J. Lee et al., 2020). The negative interaction between the Erythroid and Phagosome programs is consistent with the presence of two mutually exclusive differentiation paths that lead to erythroid and myeloid cell fates.

4.3 D-SPIN reconstructs the architecture of core cellular pathways

In addition to inferring the structure of the core K562 gene regulatory network, D-SPIN also inferred interactions between each gene program and each gene knock-down perturbation. The interactions between perturbations and gene programs are represented by perturbation response vectors $\mathbf{h}^{(n)}$. The entries of $\mathbf{h}^{(n)}$ represent positive or negative interactions between perturbations and individual gene programs. The response vectors can be represented as directed edges that connect perturbations to the core regulatory network (Figure 4.1 B).

Similarly to the grouping of the K562 gene regulatory network into interacting gene modules, we grouped gene knockdown perturbations into a set of 40 perturbation classes that we refer to as G1-G40 (Guide RNA group 1 through 40) through unsupervised Leiden clustering of the $\mathbf{h}^{(n)}$ vectors (Section 4.6). Gene knockdowns

within the same guide cluster have similar perturbation responses, suggesting that these genes are involved in the same pathway or have potential interactions. Identified clusters reflect the pathway-level organization in the K562 cell, revealing both well-known cell biology and more novel or cryptic organization of pathways (Figure 4.1 D). Individual guide clusters are enriched for specific cell-biological functions, including DNA replication, the MAPK signaling pathway, and RNA degradation (Figure 4.1 D).

The guide clusters also provide detailed information regarding the more granular structure of cellular pathways. For example, the components of the proteasome core particle (20S proteasome) fall into a single perturbation cluster (G3). In addition to these proteasome components, group G3 also contains genes involved in protein transport, including AKIRIN2 and IPO9 (Figure 4.1 D), two protein components recently associated with the import of the proteasome into the cell nucleus through biochemical and genetic assays. In fact, in recent work, AKIRIN2 was found to directly bind to fully assembled 20S proteasomes to mediate their nuclear import with the nuclear import factor IPO9 (Almeida et al., 2021). This finding is consistent with the grouping of AKIRIN2 and IPO9 with proteasome components in the D-SPIN perturbation response clusters. Similarly, the D-SPIN guide clusters group the gene C7orf26 with the integrator subunits INTS10, INTS13 and INTS14, a key result of the original Perturb-seq study (Replogle et al., 2022).

Thus, analysis of the perturbation response vectors inferred by D-SPIN provides a strategy for generating hypotheses regarding the architecture of core cellular pathways directly from single-cell mRNA-seq data. D-SPIN-generated hypotheses can, then, motivate biochemical experimental validation. We highlight the example of the proteasome and AKIRIN2/IPO9, as well as C7orf26 and integrator, to demonstrate that D-SPIN response vectors are consistent with two regulatory associations that were only recently validated through biochemical experiments in the literature (Almeida et al., 2021; Replogle et al., 2022).

4.4 Generating hypotheses for macromolecular complex sub-functions

Macromolecular complexes like the ribosome, the proteasome, the spliceosome, and the integrator complex play a critical role in cell physiology and gene regulation (Beltrao, Cagney, and Krogan, 2010). Protein complexes can often exist in different configurations with different combinations of subunits. Defining the functional role of individual proteins within a complex and how protein components work together

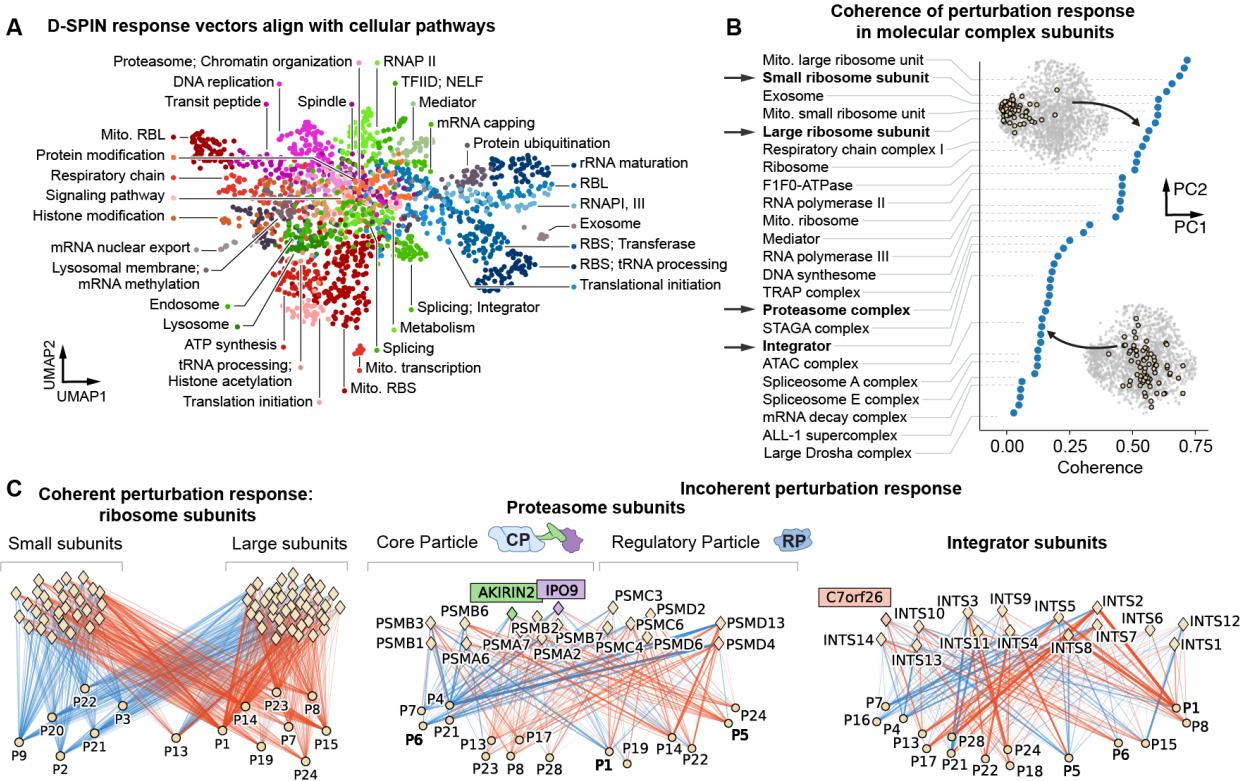


Figure 4.4: D-SPIN reveals potential function partitions for molecular complex subunits. (A) UMAP embedding of response vectors of each guide RNA where guides are colored by the major biological function that they are associated with. The UMAP is partitioned into regions corresponding to major cellular pathways including transcription, translation, respiration, and cell cycle. (B) The panel plots the regulatory coherence of macromolecular complex subunits, quantified by the average cosine similarity between the perturbation response vectors over all pairs of subunits. In high-coherence complexes, knockdowns of different subunits have similar interactions with the regulatory network encoded by the response vectors, as (top inset) visualized for an example complex in the PCA space of all response vectors. In low-coherence complexes, knockdowns of different subunits have different interactions with the regulatory network, as (bottom inset) shown in PCA for an example complex. The incoherence suggests distinct functions for complex subunits, (C) Diagrams of response vectors for subunits of the ribosome, proteasome, and integrator complexes, where diamond markers are targets of gene knockdowns and circle markers are gene programs of perturbation response. (left) The ribosome has a coherent signature where all subunits have similar inferred interactions with gene programs as encoded in response vectors, while (right) the proteasome and integrator each have signatures of sub-functionalization. In the example of proteasome, low coherence is observed as only a subset of subunits have an interaction on a gene program (P5, P6), or even display opposite interactions on a gene program (P1). The subunit partitions based on response vectors agree with the known structural partitions of the complexes, such as core and regulatory particles of the proteasome. Recently discovered genes AKIRIN2, IPO9, and C7orf26 are correctly grouped with the complex components that they interact with.

to achieve function is an important question in cell biology. Further, detailed insight into the architecture of specific proteins within protein complexes can identify therapeutic targets and has been a major focus in studying protein complexes in diseases including HIV, SARS-CoV-2, and cancer (Jaeger et al., 2012; D. E. Gordon et al., 2020; M. Kim et al., 2021; Belinda Wang et al., 2024). Single-cell mRNA-seq experiments can provide insight into the functional roles of proteins within macromolecular complexes to provide hypotheses that can be studied in detailed but technically challenging biochemical experiments, including affinity purification and mass spectrometry (Hiatt et al., 2022).

Therefore, we used D-SPIN to gain insight into sub-functions of complex components by comparing cosine similarity between the response vectors $h^{(n)}$ for components of known complexes in the CORUM database (Giurgiu et al., 2019). When subunits of a complex interact to achieve a common function, disruption of individual subunits will disrupt that common function, leading to high cosine similarity between subunit response vectors that approaches 1. When complexes can achieve distinct functions through combinations of distinct protein components, the knock-down of these components will generate different responses and the average cosine similarity approaches 0, which is the random expectation. We compiled a list of 46 core cellular complexes from the CORUM database by filtering out complexes that are not present in the dataset and merging variants of the same complex (Section 4.6). For each complex in the list, we defined a coherence score by computing the average cosine similarity between response vectors, $h^{(n)}$, for all pairs of complex subunits (Section 4.6).

Across the 46 complexes in the CORUM database, our analysis revealed a wide range of coherence scores across different complexes, ranging from a low coherence score of 0.01 for the large Drossha complex to a high coherence of 0.73 for the mitochondrial large ribosome (Figure 4.4 B, Figure 4.5 A). Notably, complexes that play critical roles in specific biological functions including translation and metabolism, such as the mitochondrial ribosome, ribosome, exosome, DNA repair, and respiratory chain, exhibited the highest coherence scores. In contrast, complexes with known regulatory roles, including the large Drossha complex, DGCR8 complex, ALL-1 supercomplex (involved in chromatin regulation), the ATAC complex (histone acetyltransferase), and spliceosome all had lower coherence scores below 0.25. The low-coherence complexes are involved in versatile regulatory roles in epigenetic regulation and microRNA and mRNA processing. D-SPIN provides a list of the

candidate cellular sub-processes (gene programs) that are controlled by individual sub-units of each macromolecular regulatory complex, providing a specific framework for generating experimental hypotheses to determine how modulation of the subunit composition of these macromolecular regulatory complexes might enable the complex to achieve a range of regulatory functions.

As an example, the ribosome behaves as a coherent complex; the knockdown of individual genes in the large or small ribosome subunit has a similar impact on D-SPIN gene programs, as illustrated by the monochromatic edge connections between subunit perturbations and gene programs in the interaction diagram (Figure 4.4 C). Perturbations to ribosome subunits activate protein degradation-related gene programs including P2 Lysosome, P3 Endosome, and P20 P21 Proteasome. The perturbations inhibit ribosome programs, P1 Mitochondria, P19 Endoplasmic reticulum, P23 Unfolded protein response, and P24 Oxidation. Thus, perturbing any ribosome subunits induces the cell into a stress response state with upregulated protein degradation and inhibition of protein production and secretion processes.

Unlike the ribosome, the proteasome and integrator behave as low-coherence complexes, and have response vectors that split subunits into smaller sub-groups that have distinct inferred interactions with downstream gene expression programs. The integrator complex subunits are partitioned into four sub-groups that agree with the original Perturb-seq study (Figure 4.4 C): an INTS10-13-14-C7orf26 module, a cleavage module (INTS3, INTS4, INTS8, INTS11), and two shoulder and backbone modules (INTS2, INTS5, INTS7, INTS8; INTS1, INTS6, INTS12). Knockdown of the cleavage module activates P21 Proteasome and inhibits P22 rRNA processing, P18 Nuclear transport, P24 Oxidation, and P5 Lipid metabolism, while knockdown of the shoulder and backbone modules activates P4 Erythroid and P7 Innate immunity and inhibits P21 Proteasome, P28 Chromosome segregation, P17 RNA splicing, and P13 Ribosome programs.

Similarly, the proteasome complex subunits are split into two groups of core particles (20S proteasome) (PSMA6, PSMB1, PSMB3, PSMB6; PSMA2, PSMA7, PSMB2, PSMB7) and two regulatory particles (PSMC3, PSMC3, PSMC3, PSMD2, PSMD6; PSMD4, PSMD13). D-SPIN finds that the knockdown of proteasome sub-unit PSMC6 (also known as TBP7 and Rpt4) has a negative impact on respiration; this result is consistent with the finding that PSMC6 interacts with mitochondrial proteins including TRAP1 in the endoplasmic reticulum to assist the folding of mitochondrial proteins (Amoroso et al., 2012). We also identified signatures of protein

complex sub-modules in the large Drosha complex, ALL-1 histone methyltransferase complex, and spliceosome E complex. This indicates that specific subgroups of these complexes have different regulatory roles in controlling downstream gene expression programs. For example, the response vectors of spliceosome E complex subunits split into four sub-modules. Example genes in the four sub-modules (M1–M4) include heterogeneous nuclear ribonucleoproteins (hnRNPs)(M1), small nuclear ribonucleoprotein particles (snRNPs)(M2), core splicing factors (M3), and serine/arginine-rich (SR) proteins (M4) (Figure 4.5 B). Different sub-modules can induce opposite perturbation responses on the same gene program: M1 activates P1 Respiration while M4 inhibits P1; M1 and M4 activate P4 Erythroid while M2 inhibits P4. These results demonstrate that the response vectors reveal the impacts of genetic perturbations, and provide a compact signature for identifying functional partitioning of molecular complexes.

4.5 Coarse-grained D-SPIN models provide insight into global perturbation response strategies

All cells must regulate a large number of internal processes (e.g., transcription, translation, cell cycle, metabolism, and differentiation) to maintain homeostasis in response to damage and perturbation. Understanding how cells control their physiological state is essential for an integrated understanding of cell regulation, but it remains poorly understood due to the lack of global integrated models. The D-SPIN gene-program model (Figure 4.1 B) allows us to analyze the distributed regulation of core cellular processes in response to gene knockdown perturbations, a form of cellular damage.

While the D-SPIN model is compact, we find that the model can be simplified further through an automated coarse-graining strategy based on graph clustering (Figure 4.6 A) that reduces the D-SPIN network by grouping gene programs into sets that we call *modules* and single gene perturbations into perturbation clusters. We call this strategy *network coarse-graining* because of its similarity to coarse-graining strategies in physical models, which group strongly correlated degrees of freedom into a single variable to aid computation and interpretation. The coarse-graining strategy produces a minimal network that we analyze to reveal four classes of stress-response strategies, patterns of gene activation, and inhibition generated by the cell when specific cellular processes are disrupted. The four strategies are consistent with the hyperparameter choice of the graph clustering (Section 4.6).

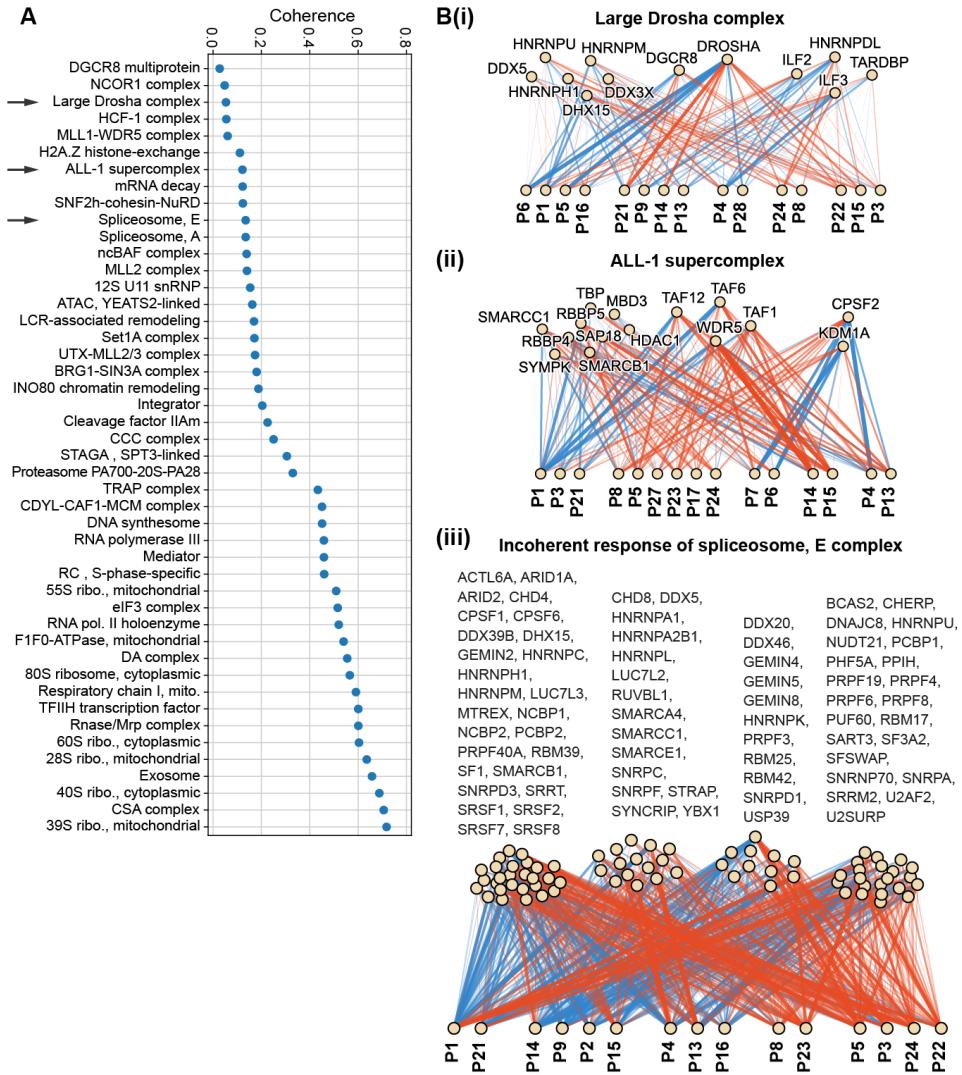


Figure 4.5: Extended examples of molecular complex partitions identified by D-SPIN. (A) Full plot of regulatory coherence of macromolecular complex subunits of the 46 filtered molecular complexes in the CORUM database (Giurgiu et al., 2019). In low-coherence complexes, knockdowns of different subunits have different interactions with the regulatory network, suggesting distinct functions for complex subunits. (B) Diagrams of response vectors, for example, incoherent molecular complexes that have signatures of sub-functionalization. The subunits of (i) large Drosha complex and (ii) ALL-1 supercomplex are partitioned into three groups. The subunits of (iii) spliceosome E complex are partitioned into four groups.

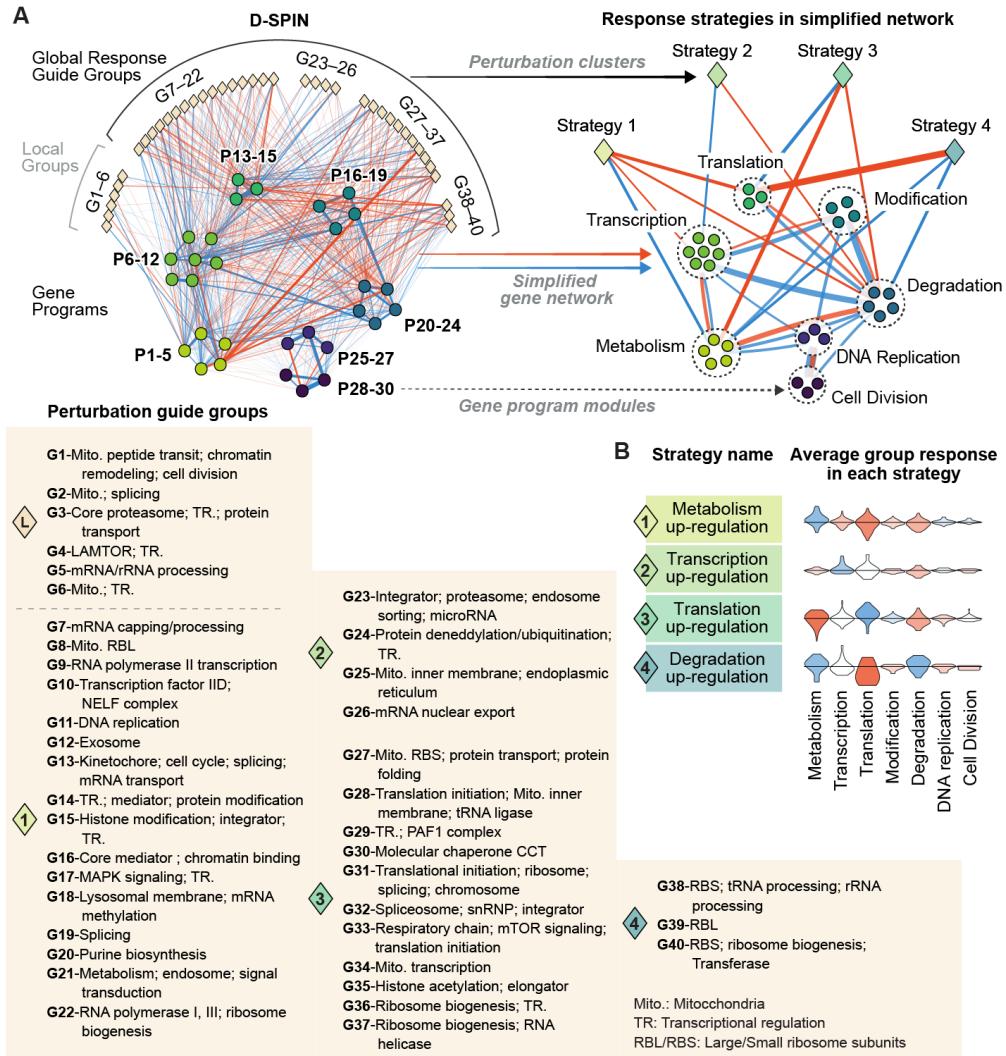


Figure 4.6: Coarse-graining of D-SPIN network model identifies global perturbation response strategies in K562 cells for distinct classes of gene knockdowns.

(A) Diagram of network coarse-graining by grouping gene programs into seven identified gene program modules and grouping guide RNA classes into four identified strategies. The resulting coarse-grained model enables global analysis of cellular regulatory responses. (B) Violin plots of averaged perturbation response vectors on programs in guide groups in each response strategy. K562 cells contain four distinct classes of global response strategies, and we name each strategy by the upregulated characteristic biological function.

The D-SPIN-identified regulatory strategies suggest non-trivial forms of homeostatic regulation. We find that each strategy features the upregulation of a major biological function and downregulation of other functions; the only exception is that in the fourth strategy, both metabolism and degradation are upregulated (Figure 4.6 B). The upregulation indicates that the cell compensates for the loss of a gene by upregulating a potentially compensating cellular function. The upregulation suggests the existence of active regulatory feedback within the regulatory network under the applied perturbations to maintain homeostasis. Therefore, the four global response strategies are named by the characteristic program upregulation as metabolism upregulation, transcription upregulation, translation upregulation, and degradation upregulation. In general, the K562 cells may be capable of mounting a wide range of homeostatic response strategies, but D-SPIN identified the four highlighted strategies given the specific experimental conditions in which the Perturb-seq experiment was performed.

For example, in the “translation upregulation” response strategy, perturbation to components impacting translation initiation (G28) leads to upregulation of translation and downregulation of protein degradation, suggesting that the cell compensates for decreased translational flux through a regulatory response that both upregulates translation and downregulates processes that degrade proteins. Interestingly, perturbations to ribosome subunits and ribosomal rRNA (G38, G39, G40) production lead to a distinct response where the cell downregulates translation while upregulating protein degradation, suggesting a cellular attempt to contain the impact of unassembled ribosome subunits and damaged proteins (Figure 4.7 A). The program-level perturbation responses are confirmed at the single-gene level. Under the knockdown of translation initiation factors (EIF2B1, EIF3H, EIF4G2), all ribosome subunits are coherently upregulated and proteasome subunits are coherently downregulated, as quantified by the Z-score relative to the control distribution of non-targeting guide RNAs (Figure 4.7 B). Similarly, the knockdown of single large/small ribosome subunits or ribosome genesis genes (FCF1, RIOK1, POP1) induces coherent downregulation of all ribosome subunits and upregulation of proteasome subunits.

The regulatory strategies also provide insight into the flow of information in the cell. The compensatory functions are typically not directly associated with the gene being knocked down, suggesting long-range regulation or coupling between seemingly distinct processes. For example, a large proportion of metabolism upregulation responses are induced by perturbations to transcription processes, including RNA

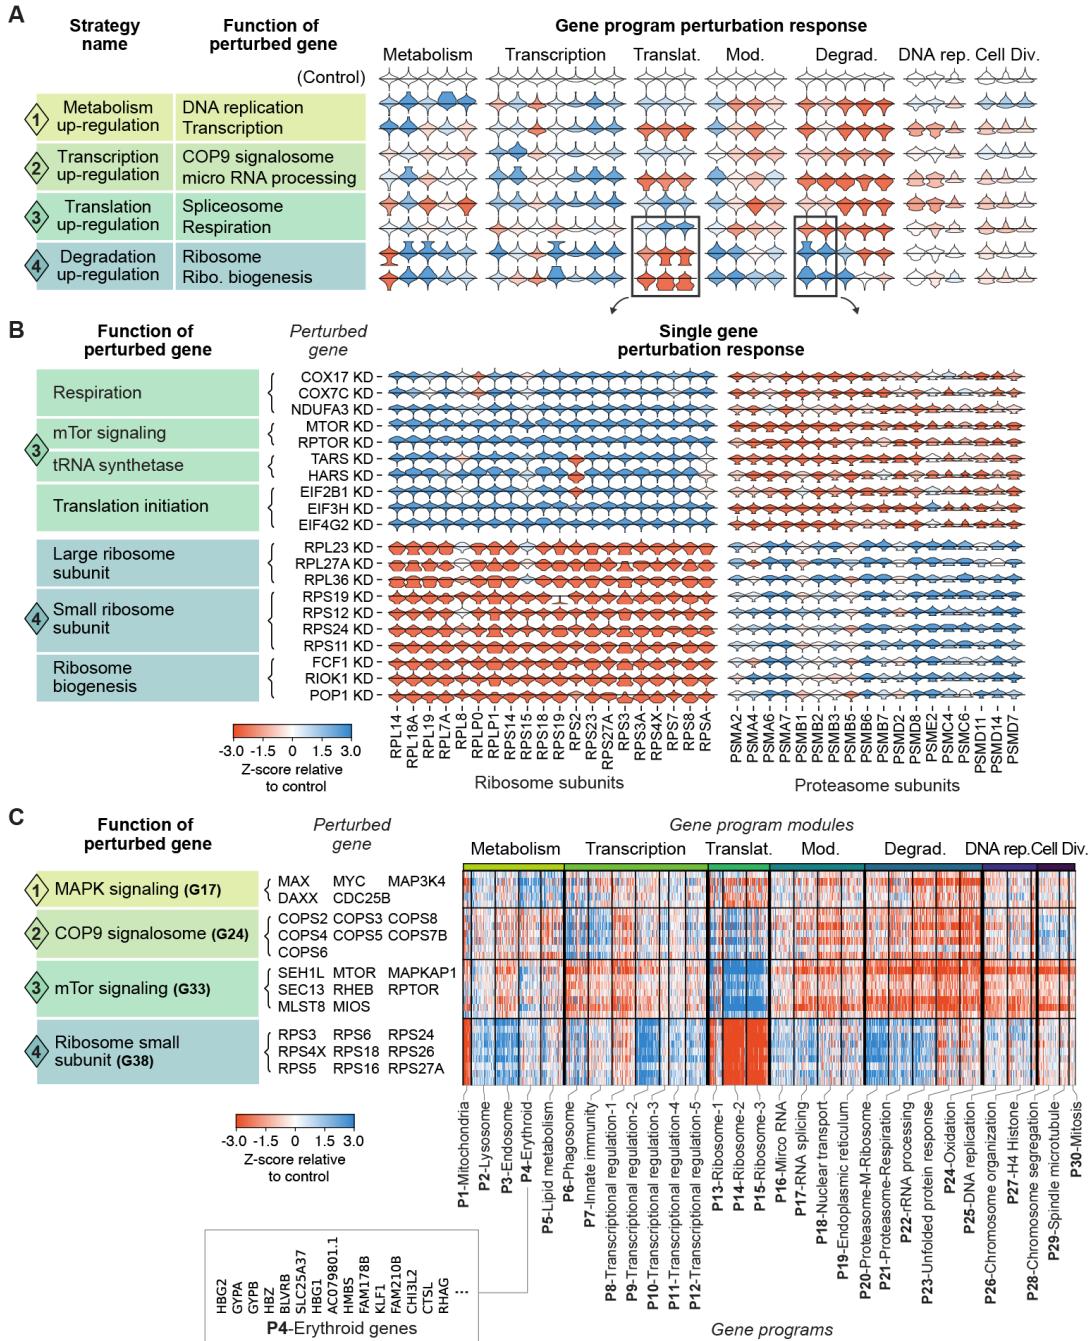


Figure 4.7: Response strategies exhibit as coherent up and down-regulation of genes of relevant functions. (A) Violin plots show gene program expression under knockdown of example pathway components relative to the control samples where a non-targeting guide RNA is introduced to cells. The upregulated biological function typically is distinct from the genes being perturbed. For example, when RNA polymerase is knocked down, cells generate the metabolism upregulation response, where the cell upregulates metabolism and downregulates transcription, translation, and degradation. As a second example, when ribosome subunits are knocked down, cells upregulate protein degradation and metabolism while downregulating translation.

Figure 4.7: (Continued) (B) Violin plots show the single-cell mRNA expression distribution for ribosome and proteasome subunits under knockdown of example genes in translation upregulation and degradation upregulation strategies relative to control samples. The gene-program level response strategies are reflected at the single-gene level. For example, when respiration-related gene COX17 or COX7C are knocked down, all ribosome subunit genes are upregulated and proteasome subunit genes are downregulated. (C) Gene-expression heatmap showing the upregulation (blue) or downregulation (red) of representative genes in all gene programs to the knockdown of individual genes in a series of cellular pathways relative to control samples. The program-level response strategies are shown as coherent upregulation/downregulation patterns of single genes in gene programs.

polymerase (G9, G22), Transcription Factor IID, NELF complex (G10), and mediator (G16). The knockdown of these genes leads to upregulation of metabolism, and downregulation of transcription, translation, and degradation. It has been found that transcription stress would induce an elevated ATP pool and rewired metabolism states (Milanese et al., 2019). As another example, the COP9 signalosome (CSN) is a protein complex involved in the deubiquitination and deneddylation of proteins. Knockdown of components including COPS2, COPS4, and COPS6 leads to upregulation of the phagosome and downregulation of protein-degradation machinery (Figure 4.7 C). As another example, the knockdown of mTOR signaling components (which control cell growth) leads to the upregulation of ribosomal programs in a translation upregulation response.

These results demonstrate that D-SPIN-generated regulatory networks provide insights into global control strategies used by cells to regulate a distributed set of processes to achieve homeostasis in response to various forms of internal damage. Knockdown of translation initiation factors and ribosome subunits both affect the translation process, while the cells employ two distinct strategies (translation upregulation vs. degradation upregulation) to maintain homeostasis, indicating the presence of a cellular sensing and information processing mechanism to distinguish and respond to different causes of stress. The correspondence between gene function damage and compensatory functions may lead to new therapeutic strategies, such as combinatorial inhibition of both a gene function and its associated compensatory pathways, for example, for improved killing of cancer cells. Such combinatorial strategies can be difficult to design due to the exponentially large space of potential combinatorial genetic interventions. Conceptually, the global perspective of D-SPIN illuminates classes of regulatory strategies, suggesting that compensatory

mechanisms and information flow between cellular processes that might otherwise not be connected, thereby providing insights into principles of information processing and homeostatic control.

4.6 Supplementary information

Filtering, normalization, and clustering of the Perturb-seq dataset

The Perturb-seq dataset is a genome-wide screening study where every single gene is knocked down individually (Replogle et al., 2022). In Perturb-seq, a library of guide RNAs is introduced by lentiviral transduction into engineered cells that express the dCAS9 CRISPRi effector, after which single RNA sequencing (scRNA-seq) is conducted to simultaneously read out the transcription profile and guide RNA for each cell. Because guide RNA transfection is a stochastic process, the number of cells receiving each perturbation varies vastly. Here, the dataset contains perturbations with a sample size ranging from 2 to 2,000 cells. To reduce noise introduced by the small sample size, we excluded perturbations with fewer than 20 cells (metadata entry “number of cells (filtered)”), which represented 92 perturbations, 0.86% of all perturbations.

More than half of the perturbations do not yield significant single gene-level expression changes (apart from the perturbed gene), possibly due to genome redundancies that ensure the robustness of cellular activities, and that some genes may not function in the profiled experimental condition. To avoid the model overfitting the measurement noise on these perturbations with no significant effect, we excluded perturbations with fewer than 10 differentially expressed genes (metadata entry “Number of DEGs (anderson-darling)”), which represented 7,523 perturbations, 70.5% of all perturbations. Similarly, we only selected a subset of the 585 control samples with non-targeting guide RNAs to avoid overfitting on measurement noise. We selected control samples that were labeled as core control (metadata entry “core_control”) and had no fewer than 200 collected cells (metadata entry “number of cells (filtered)”). In total, we selected 3,136 perturbation conditions and 105 control conditions.

Normalization of single-cell gene expression data is required because (1) the total number of transcripts captured for each cell can vary from 1,000 to 100,000 due to technical variability in reagents and library preparation steps, and (2) transcript expression levels across genes can span five orders of magnitude. We normalized the raw counts for each cell by the total counts detected in each cell, times a scaling

factor that is roughly the median of the total counts of transcripts in a cell. We used the scaling factor 10^4 as the target total count, that is, computing $g'_i = 10^4 * g_i / \sum_i g_i$ for every cell, where g_i is the count of gene i . We also used a log transformation to mitigate the great variety of expression levels across genes $g''_i = \log(g'_i + 1)$. Except for gene filtering, all the following analysis was performed on the log-transformed data.

Due to biological noise, measurement noise, and gene drop-out, genes that have constant expression across cells do not contribute to biological impact, but only increase noise in the data. Therefore, single-cell data analyses are performed on a subset of genes that have different expression levels across cells, i.e., highly variable genes. Here, we developed a highly variable gene filtering model adapted from the coefficient-of-variation (CV) filtering using the zero-inflated Poisson distribution or negative binomial distribution as the statistics of gene counts. In both distributions, the variance has a $AX^2 + BX$ form of dependence with the mean so the resulting filtering procedure is identical. The model was motivated by our observation that genes with high expression generally have higher CV than the prediction of the Poisson model, and can be explained by the heterogeneous RNA counts in the cell population or drop-out during sequencing. Drop-out means every gene count has a random probability of being mistakenly sampled as 0, which is also called zero inflation in statistics. For notation simplicity, we derive the filtering using the zero-inflated Poisson distribution. For a zero-inflated Poisson distribution,

$$P(X = x) = \begin{cases} \pi + (1 - \pi)e^{-\lambda}, & \text{if } x = 0, \\ (1 - \pi)e^{-\lambda}\lambda^x/x!, & \text{if } x > 0, \end{cases} \quad (4.1)$$

where λ is the parameter of Poisson distribution and π is the probability of drop-out or zero inflation. The mean and CV of the distribution is

$$\mu = (1 - \pi)\lambda \quad (4.2)$$

$$\frac{\sigma}{\mu} = \frac{\sqrt{1 + \pi\mu/(1 - \pi)}}{\sqrt{\mu}} \quad (4.3)$$

The zero-inflated Poisson distribution suggests another type of log CV (y) - log mean (x) relationship where the expected log CV is higher than the prediction of a Poisson model for high-expression genes.

$$y = -\frac{1}{2} \left[x - \log\left(1 + \frac{\pi}{1 - \pi} \exp x\right) \right] \quad (4.4)$$

Therefore, we followed the form of the equation and obtained the empirical log CV - log mean relationship by curve fitting (`scipy.optimize.curve_fit`) of the following function

$$\hat{y} = -k(x - \log(1 + \pi \exp x)) + b \quad (4.5)$$

where we denoted $\pi/(1 - \pi)$ into a single parameter π for numerical stability. The genes are ranked by $\log \text{CV} - \hat{y}$, the difference between $\log \text{CV}$ and curve fitting. A higher value indicates that the gene is more variable.

We selected highly variable genes as the union of the following three sets: (1) we performed the Wilcoxon test of log-transformed data for each guide RNA and control (`sc.tl.rank_genes_groups`) and selected the top 10 differentially expressed genes for each guide that passed $p < 0.05$ under Bonferroni multi-test correction (we selected fewer than 10 if no enough genes pass the test), (2) the top 1,000 highly variable genes selected by our zero-inflated Poisson CV filtering, and (3) for comparison with the original Perturb-seq study, we included the genes that appeared in the gene programs in the original Perturb-seq study. In total, we selected 3,273 genes for the analysis using the D-SPIN framework.

To visualize and cluster single cells in the Perturb-seq dataset, we corrected the batch effect by constructing latent representations with scVI on raw gene counts (`dispersion='gene-batch', n_latent=10, n_layers=2, max_epoch=20, batch_size=2048`) (Lopez et al., 2018). We constructed a 15-nearest-neighbor graph (`scanpy.pp.neighbors`) on the latent representation and computed Leiden clustering for cell clusters (`scanpy.tl.leiden`). As K562 is a purified cell line, it contains primarily a single cell type. Thus, the UMAP embedding of the cell population is a circular structure, representing the progression of the cell cycle, which is less relevant to the study of perturbation responses. Therefore, we only show the UMAP embedding on control samples along with D-SPIN-generated samples using cell-cycle-associated gene programs (Section 5.6).

oNMF decomposition of single-cell dataset

In the Perturb-seq study, we used a subset of 10^5 single cells, with the Square-root balancing scheme in Section 2.5 for the Leiden cluster of cell states. In the Perturb-seq study, we ran oNMF 100 times with $K = 30$ with different random seeds and computed a consensus gene program set by K-means.

Evaluating single-gene network reconstruction with ChIP-seq data

To further assess the performance of D-SPIN on regulatory network reconstruction with biological measurement datasets, we applied different network inference methods to the genome-wide Perturb-seq dataset, and compared the results with biological validations obtained from ChIP-seq data. ChIP-seq data measure the binding of transcription factors (TFs) to target gene promoters and have been used in the literature to assess the quality of inferred gene regulatory network models. We evaluate the correspondence between inferred network and ChIP-seq data by the number of inferred edges that can be associated with a measured TF-target gene binding event.

Specifically, we generated two datasets from the Perturb-seq data, similar to the BEELINE framework (Pratapa et al., 2020), labeled as TFs+500 and TFs+1000. We selected TFs that expressed in over 5% of cells and the top 500 or 1000 differentially expressed genes that were identified with the zero-inflated Poisson distribution model we developed (Section 4.6). We only included genes that had corresponding knockdown perturbation conditions that satisfied the following criteria.

1. More than 20 cells collected (metadata entry “number of cells (filtered)”).
2. More than two differentially expressed genes identified (metadata entry “Number of DEGs (anderson-darling)”).
3. A knockdown efficiency exceeding 5% (metadata entry “percent knockdown”).

Additionally, we incorporated 46 control conditions with non-targeting guide RNA, with a minimum of 250 cells each. After filtering, the TFs+500 dataset comprised 103k cells, 421 genes, and 467 conditions, while the TFs+1000 dataset included 150k cells, 624 genes, and 670 conditions.

For existing network inference methods, considering the scalability challenges faced by PIDC, GENIE3 and GRNBoost2 with datasets exceeding 100k cells, as previously benchmarked in Section 3.5, we divided the cells into approximately 2k-cell chunks. Each method was then applied to these chunks independently, with the resulting networks averaged to produce an overall inference.

For D-SPIN, we partitioned the data based on each perturbation condition. We did not partition by condition for other methods, because many perturbation conditions only had a few tens of cells, which would be insufficient for their network inference.

D-SPIN can handle perturbations with small cell numbers as information is shared through inferring a unified regulatory network. As we knew the target of each knockdown perturbation, we used the prior $h_i^* = -3$ for the knockdown target genes. We used pseudolikelihood inference because of the large network size. We applied ℓ_1 regularization on the network with $\lambda_1 = 5e - 3$, and perturbation prior regularization with $\lambda = 0.5$. Compared with the inference of the large synthetic network (Section 2.4), we used a smaller perturbation prior regularization coefficient, as the effect of experimental single-gene knockdown is more noisy and uncertain.

The inferred networks were compared with processed ChIP-seq data of TF-target interactions from human ChIP-seq datasets. We used the interactions from DoRothEA (Garcia-Alonso et al., 2019), RegNetwork (Z.-P. Liu et al., 2015), and TRRUST (Han et al., 2018), which were integrated together by the BEELINE framework. As the ChIP-seq data only contained TF-target regulations, we removed all inferred interactions between non-TF genes, which was consistent with the BEELINE framework. The correspondence with ChIP-seq data is quantified by the early precision rate (EPR), which measures the proportion of inferred edges that can be associated with a TF binding event detected by ChIP-seq among the top K predictions, relative to the expectation of a random predictor. K is set to the total number of TF-target interactions in the ChIP-seq database.

Discovering regulatory networks modules and clustering perturbation responses
 To gain insight into the structure and modular organization of the regulatory networks inferred by D-SPIN, we applied the Leiden algorithm for graph community detection to the inferred networks. The Leiden algorithm identifies densely connected subnetworks, or communities, by optimizing a modularity score that reflects the density of positive interactions within communities compared to a random network. For networks with both positive and negative interactions, the following score function is employed (Traag, Waltman, and Van Eck, 2019).

$$Q = Q(J_+) - Q(J_-) \quad (4.6)$$

$$Q := \frac{1}{2m} \sum_{i,j} \left[\left(J_{ij} - \gamma \frac{k_i k_j}{2m} \right) \delta(c_i, c_j) \right] \quad (4.7)$$

where the regulatory network is split into J_+ for positive and J_- for negative interactions. In the score function Q , m is the sum of all edge weights, k_i denotes the sum of weights for edges connected to node i , c_i is the community to which i

belongs, and γ is the resolution parameter. We set $\gamma = 2$ to achieve a finer resolution of module discovery, while the presented result is robust with a range of γ from 1.05 to 2.5. The Leiden algorithm partitions the network into distinct communities, whose biological functions were inferred from the gene programs they contained.

In addition to analyzing the network architecture analysis, we also clustered the perturbation response vectors to uncover interaction patterns between perturbations and the network. In general, we computed relative response $\tilde{\mathbf{h}}$ for each \mathbf{h} by subtracting the average \mathbf{h} of control samples. Effective perturbations were then selected by comparing with the distribution of relative responses in control samples.

In the Perturb-seq dataset, we selected 1,755 effective perturbations with $p < 0.05$ under Bonferroni multi-test correction using non-targeting guide RNA as null distribution. We clustered perturbations into groups, because the perturbation impacts are noisy due to the low cell number in each condition and varying knockdown efficiencies. We used the Leiden algorithm for better clustering of the noisy perturbation responses. We built an 8-nearest-neighborhood graph and ran Leiden algorithms with resolution $\gamma = 3$ to partition the perturbation response vectors into 40 groups. The relative responses of guide groups were defined as the average relative responses in the group.

Automated identification of regulatory strategies and network coarse-graining

The response vectors can further reveal the information processing of the regulatory network by elucidating how perturbation responses are organized across network modules. The response vectors not only identify perturbation impacts on individual programs, but also provide a global view of how cells respond to the stress induced by the perturbation. To further consolidate perturbation groups, typical measures such as Euclidean distance are not sufficient, as each gene program is not an independent dimension but interacts closely with other gene programs through the regulatory network. Therefore, we used a graph clustering strategy to integrate information of program interactions into the clustering.

Specifically, we built an extended network by also including perturbations as nodes of the network. The extended network encompassed both interactions between gene programs and interactions between perturbation and programs. Formally, the adjacency matrix of the extended network is

$$\mathbf{J}_{\text{ext}} = \begin{bmatrix} \mathbf{J} & [\tilde{\mathbf{h}}_1, \tilde{\mathbf{h}}_2, \dots] \\ [\tilde{\mathbf{h}}_1, \tilde{\mathbf{h}}_2, \dots]^T & \mathbf{0} \end{bmatrix}, \quad (4.8)$$

where $\tilde{\mathbf{h}}$ is the average relative response for each perturbation group defined in Section 4.6. We then ran the Leiden algorithm on the extended network to detect perturbations that had similar impacts at the network level. In Leiden algorithms, modules of gene programs were fixed, and we only partitioned the perturbations. Perturbations were partitioned into the same community if they influenced gene programs that play analogous roles within the regulatory network. We set $\gamma = 1$ in the Leiden algorithm, while the algorithm always generates four perturbation clusters for γ from 0.5 to 2.5, and the presented cluster assignment result is robust with a range of γ from 0.9 to 1.45.

Through graph clustering, perturbation groups were assigned to modules in the regulatory network. Based on the score function Equation 4.7 of the Leiden algorithm, this suggested that the perturbations predominantly activated gene programs within the module and were more likely to inhibit the gene programs outside the module. Conceptually, this correspondence between perturbations and modules established a connection of homeostasis regulation logic, that the cell upregulated the gene programs in the module in response to the loss of function caused by the perturbations. Therefore, we named the global strategy of perturbation responses according to the gene module that the perturbations were grouped with. For example, perturbations that were assigned to the Translation module were named “Translation up-regulation strategy”.

We further visualized the global organization of biological functions and perturbation response strategies with a minimal network derived from the D-SPIN model by coarse-graining. In this simplified representation, both gene program modules and response strategies were condensed into single network nodes, as programs in the same module mostly have positive interactions inside and collectively contribute to a specific biological function. In the minimal network, the interactions between modules were determined by the most significant positive or negative interactions across their constituent gene programs. Similarly, the sign of the interaction between a gene module and a perturbation strategy was based on whether the activating or inhibiting interactions were stronger on average across the constituent gene programs and perturbation groups, and the strength is the average strength. As a visualization tool, the minimal network spotlighted the major response strategies employed by the network. For example, the impairment of translation machinery induced downregulation of the Translation module and upregulation of the Degradation module.

Coherence of molecular complexes in CORUM

To obtain a comprehensive list of molecular complexes whose subunit knockdown effects were captured by the Perturb-seq dataset, we filtered the list of protein complexes in the CORUM database (Giurgiu et al., 2019) by two criteria: (1) each complex should have no fewer than 10 subunits, and (2) at least 2/3 of subunits are contained in our filtered list of guide RNA, as in Section 4.6. Due to the presence of closely related complexes with many shared subunits in the database, we constructed a similarity graph using the overlap of subunits as edges. Through clustering, we distilled these into 46 distinct categories, and selected the most representative complex from each class that had the greatest number of subunits.

To assess the consistency of perturbation effects on subunits of each complex, we computed a coherence score as the average cosine similarity across all pairs of subunit perturbation response vectors.

$$\text{coherence} = \frac{2}{N(N-1)} \sum_{i < j} \frac{\tilde{\mathbf{h}}_i^\top \cdot \tilde{\mathbf{h}}_j}{\|\tilde{\mathbf{h}}_i\| \|\tilde{\mathbf{h}}_j\|}, \quad (4.9)$$

where N denotes the number of subunits in the complex, and $\tilde{\mathbf{h}}_i$ represents the relative response vector for knockdown of subunit i .

The clustering of perturbation response vectors further revealed functional partitions of subunits within molecular complexes. Selecting an optimal number of clusters does not have a universal rule and is inherently subjective based on the context of the analysis. Among various metrics available to guide this choice, we employed the Silhouette score, which evaluates clustering quality based on the average distances within clusters and between neighboring clusters. A higher Silhouette score indicates a better quality of clustering.

For each complex, we performed K-means clustering, with the number of clusters K ranging from 2 to 15, and calculated the Silhouette score for each. We considered a complex to exhibit distinct functional subunit groups if the peak Silhouette score was reached for a K larger than 3. Examples of molecular complexes that exhibited modular organization included Proteasome, Integrator, Large Drosha complex, ALL-1 complex, and Spliceosome E complex, as presented in Figure 4.4 C and Figure 4.5 B.

ANALYZING IMMUNOMODULATORY DRUG AND COMBINATION ACTION ON HUMAN IMMUNE CELL POPULATION

5.1 Introduction

The rational programming of the human immune system is a major goal for the treatment of diseases including cancer, auto-immunity, and neurodegeneration (Marrack, Kappler, and Kotzin, 2001; De Visser, Eichten, and Coussens, 2006; Labzin, Heneka, and Latz, 2018). The human immune system is composed of a diversity of cell types (T-cells, macrophages, and B-cells), and each cell type adopts specialized sub-states of signaling and gene expression in response to environmental cues. In diseases including cancer and auto-immunity individual cell types including macrophages and T-cells can become locked into aberrant states of immune suppression or hyper-activation that actually drive disease progression. Signaling interactions between cell types reinforce aberrant immune states and generate a major obstacle for the treatment of conditions like hyper-inflammation and immunotherapy resistance cancer. An ability to reprogram immune cell populations between states of activation and suppression using combinations of drugs and other therapeutics could allow the rational design of therapeutic strategies for conditions that currently defy treatment. However, immune state programming remains challenging because we lack a detailed understanding of how drugs and other therapeutics act on the large diversity of different cell states that execute immune function.

Single-cell genomics methods now enable full transcriptome profiling of heterogeneous cell populations across large numbers of therapeutic interventions (Adamson et al., 2016; Dixit et al., 2016; Subramanian et al., 2017; S. Chen et al., 2020; McFarland et al., 2020; Srivatsan et al., 2020; Burkhardt et al., 2021; Replogle et al., 2022). Single-cell mRNA-seq measurements can quantify how a drug alters the transcriptional state of individual cell types, and also reveal how drug action across many cell types rebalances the abundance of different cell states in an immune population. However, conventional single-cell data analysis methods typically focus on cell-state identification and differential gene expression and do not allow the construction of quantitative and predictive models that will be required for ra-

tional therapeutic design. Immune state programming will require data modeling frameworks that can dissect and predict therapeutic responses in a heterogeneous cell population across thousands of expressed genes cell types.

Conceptually, quantitative models might reveal principles that simplify the biological interpretation of combinatorial drug responses while also aiding rational treatment design. For example, previous work has demonstrated that combinations of drugs can induce additive responses on protein expression and dynamics (Geva-Zatorsky et al., 2010) where additivity means that the response of a specific gene to a drug combination can be predicted as a simple linear combination of single drug effects. The concept of drug additivity has also enabled the identification of non-additive drug responses (Yeh, Tschumi, and Kishony, 2006; Nichols et al., 2011). Non-additive drug responses are associated with interactions between pathways targeted by individual drugs, and therefore, reveal points of information integration in cellular regulatory networks. While principles like drug additivity have been explored for small numbers of genes and relatively simple phenotypes like growth rate, drug response principles at the transcriptome scale in diverse cell populations remain poorly understood.

In order to develop an experimental, conceptual, and mathematical framework for rational programming of the immune cell population, we combined large-scale single-cell genomics with a physics-inspired mathematical modeling framework to dissect combinatorial drug responses at the transcriptome scale in the human immune system. Using multiplexed single-cell mRNA-seq, we profile the state of 1.5 million human immune cells responding to more than 500 different immunomodulatory drugs alone and in combination. We constructed a mathematical model that predicts the transcriptome scale response of individual cell types including macrophages, T-cells, and B-cells while also predicting how drugs re-balance the relative proportion of cell types and, thus, alter population structure. The model reveals that combinatorial drug responses can be predicted as emerging as linear and non-linear combinations of single drug responses on individual gene expression programs. Our model quantitatively predicts the progression of a heterogeneous immune cell population between states of activation, inhibition, and rest across drug combinations allowing us to predict optimal drug doses for immune state control. Broadly, our work provides an experimental and mathematical foundation for the design of therapeutic strategies that can rationally reprogram the state of the human immune system.

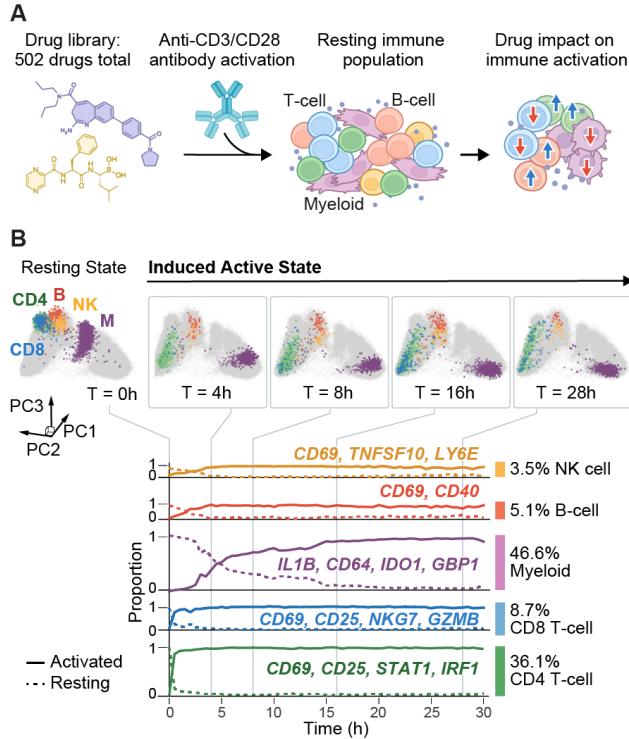


Figure 5.1: Experimental system of drug profiling in T-cell-induced immune activation. (A) Schematic of the experiment design for profiling drug responses on T-cell-mediated immune activation. Peripheral blood mononuclear cells (PBMCs) were harvested from a human donor. The cell population was treated by anti-CD3/CD28 antibodies that specifically activate T-cells and immunomodulatory drugs drawn from a library. The cell population was profiled after 24 hours of drug and antibody treatment. (B) (top) PCA projections derived from a 30-hour time-course experiment of T-cell-mediated immune activation, where samples were taken every 30 minutes for single-cell mRNA-seq. (bottom) The plot shows the time courses of the proportion of activated and resting cells in each cell type with example activation gene markers. T-cell reaches activated states first in 2 hrs and myeloid cell activation lasts 16 hrs.

5.2 Modeling immunomodulatory drug responses in primary human immune cells

D-SPIN models can contribute to the design of therapeutic strategies by providing global insights into how different underlying cell types respond to a given small molecule or other therapeutic intervention. Given the role of the human immune system in many diseases, we applied D-SPIN to dissect the response of human immune cells to a library of small-molecule immunomodulatory drugs. We developed an experimental dataset in which we applied single-cell mRNA-seq to profile

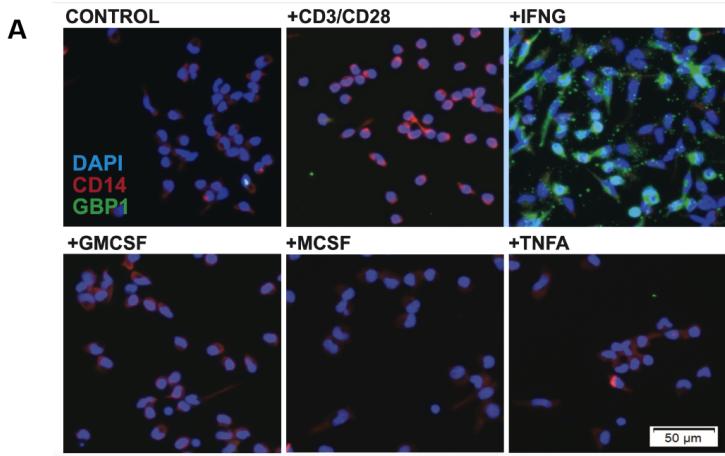


Figure 5.2: Immunostaining validates CD3/CD28 antibody does not activate monocyte alone. (A) Immunostaining of primary human monocytes cultured in CD3/CD28 antibody (25 μ L/mL cells), IFNG (5 ng/mL), GMCSF (5 ng/mL), MCSF (5 ng/mL), or TNFA (5 ng/mL). The monocyte activation marker GBP1 is only expressed under IFNG activation. The results show that anti-CD3/CD28 antibody does not activate monocytes alone, and the observed monocyte activation in the PBMC is due to signaling between cell types.

the response of human immune cells to 502 different immunomodulatory drugs. Our experimental system was designed to probe the impact of drug modulation on a cell population undergoing T-cell driven hyper-activation, as is observed in auto-immune and hyperinflammatory states (Jamiloux et al., 2019; Fajgenbaum and June, 2020). We cultured a heterogeneous population of primary donor-derived human immune cells (peripheral blood mononucleated cells, PBMC) that contain T-cells, B-cells, myeloid cells, and NK-cells. T-cells were activated using T-cell receptor antibodies anti-CD3/CD28 (Figure 5.1 A, Figure 5.2 A, Section 5.6), leading to the activation of macrophage, B-cell, and NK-cell cell types (Figure 5.1 B, Figures 5.3 A-C, Section 5.6). We applied the system to analyze the action of different immunomodulatory drugs on the activation of the immune cell population.

Single-cell mRNA-seq was used to measure the effect of the 502 immunomodulatory small molecules on T-cell-driven-immune activation. The drug library contains a diverse set of small molecules targeting pathways including mTOR, MAPK, glucocorticoids, JAK/STAT, and histone deacetylases (HDAC) (Figure 5.4 A). The library contains drugs used for treating auto-immune disease (e.g., Tacrolimus, Budesonide, Tofacitinib) as well as FDA-approved anti-cancer drugs (e.g., Bosutinib, Crizotinib). We added each compound at 1 μ M at the same time as the anti-CD3/CD28 antibody,

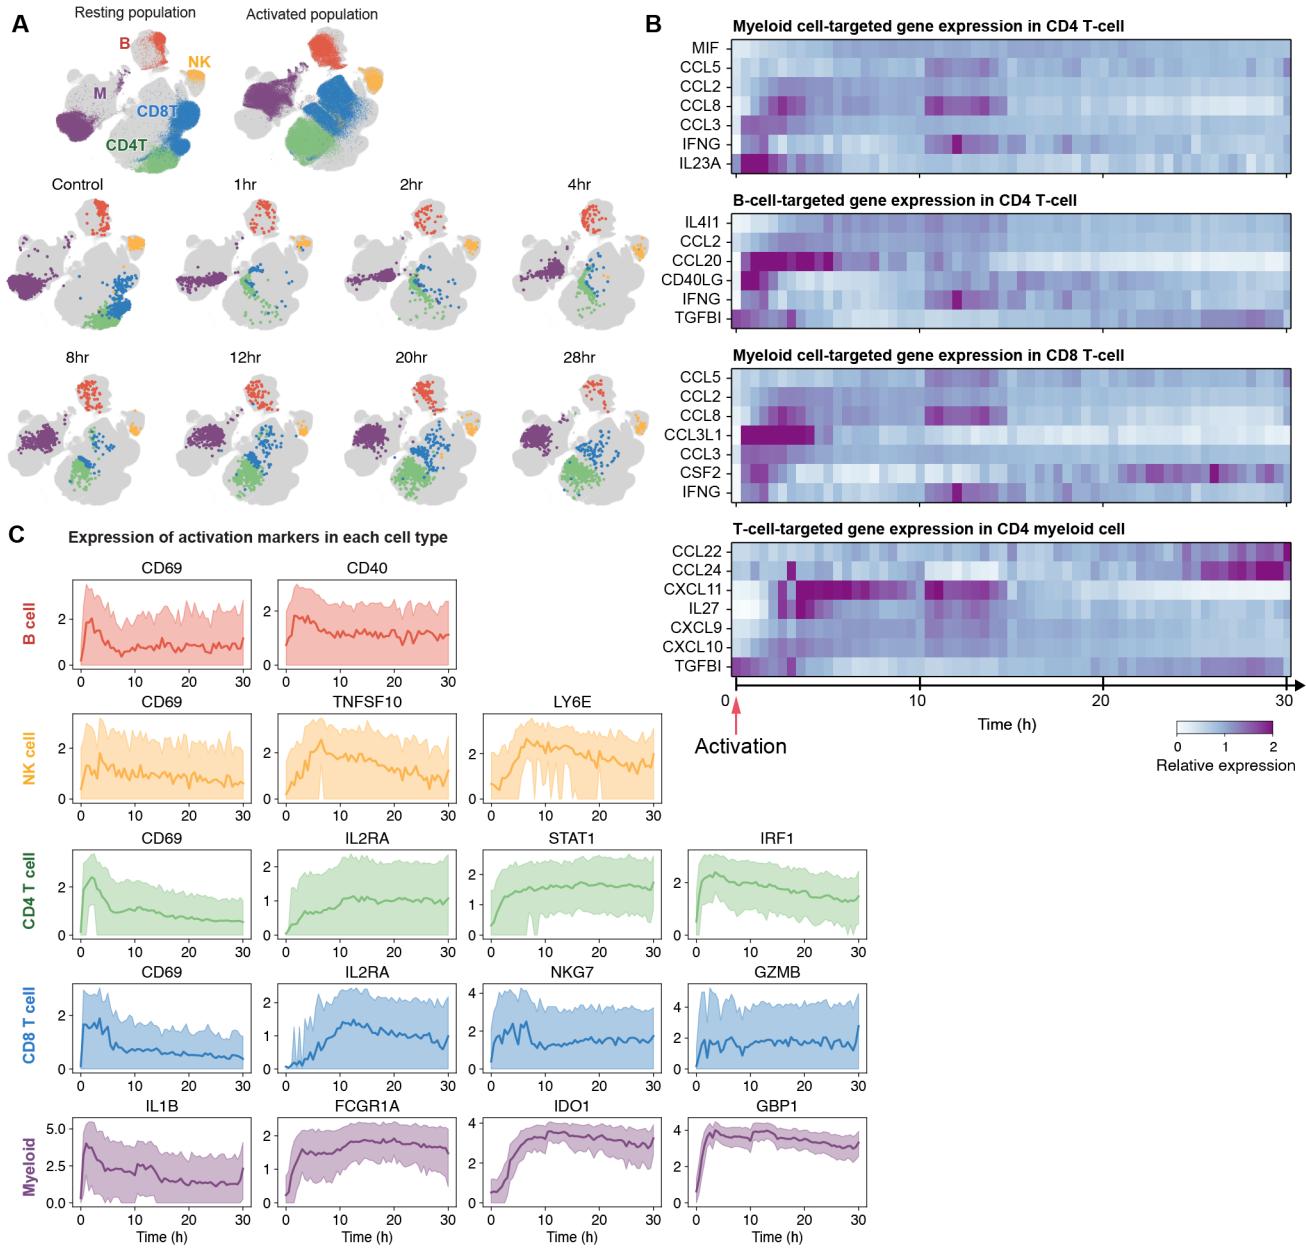


Figure 5.3: Timecourse of T-cell mediated immune activation shows various dynamics of inflammatory and signaling genes. (A) UMAP embeddings of example time points in the 30-hour time-course experiment of T-cell-mediated immune activation, where samples were taken every 30 minutes for single-cell mRNA-seq. The immune population gradually moves from the resting state to the activated state. (B) Signaling gene expression over time for select signaling genes in four cell types. Expression is normalized by the mean expression across the time course. Many signaling genes are immediately upregulated in T-cells after antibody activation, while in myeloid cells signaling genes such as IL27, and CXCL11 are upregulated after a few hours, showing rich dynamics of communications between different cell types.

Figure 5.3: (Continued) (C) The dynamics of gene markers for immune activation in each major cell type. The color-shaded range is the 10th and 90th percentile of gene expression. After T-cell activation by the anti-CD3/CD28 antibody, major cell types in PBMC are activated with different time dynamics.

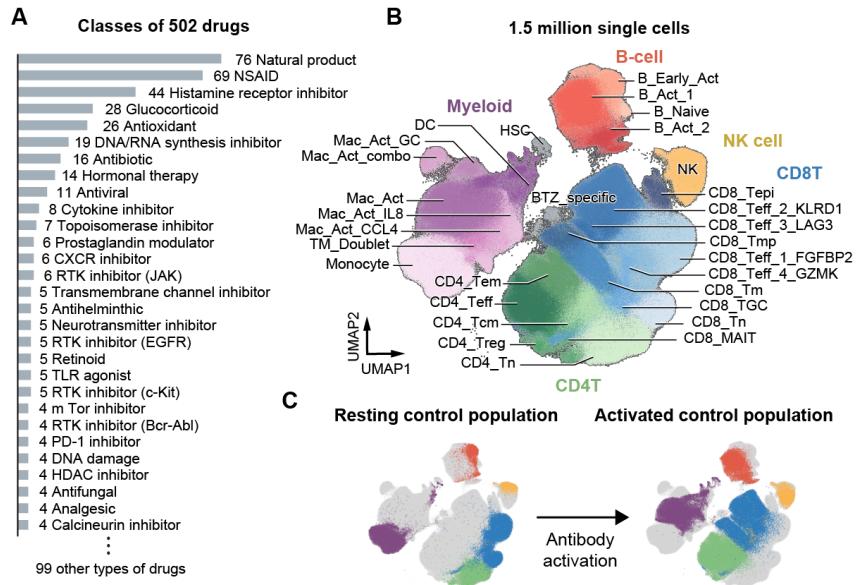


Figure 5.4: Large-scale drug profiling revealed a diversity of immune cell states under the modulation of 502 small molecule drugs. (A) The histogram of classes of profiled drugs shows the selected drugs have a variety of biochemical properties and target pathways. (B) UMAP embedding of 1.5 million filtered single cells obtained from the drug profiling experiments (Lopez et al., 2018). In the profiled cell population we identified 32 cell states in the major cell types of T-, B-, NK, and myeloid cells, and each cell state is curated by marker genes and gene differential expression. (C) UMAP embedding of the resting control cell population (without antibody activation) and activated cell population. The resting and activated cell states compose the major partition on the UMAP.

and profiled the cell population using single-cell mRNA-seq at 24 hours (Holthaus et al., 2018; McGinnis et al., 2019). For additional context, we also profiled drug responses in resting cell populations without CD3/CD28 activation for a subset of drugs. In total, we profiled 1.5 million single cells in resting and activated conditions, with over 1,200 total conditions and 31 different immune cell states, including 4 CD4 T-cell states, 10 CD8 T-cell states, an NK-cell state, 4 B-cell states, and 8 myeloid cell states (Figures 5.4 BC, Figure 5.5 A, Section 5.6).

Using the single-cell mRNA-seq data, we applied D-SPIN to construct a regulatory network model of the cell population and the interactions of individual small

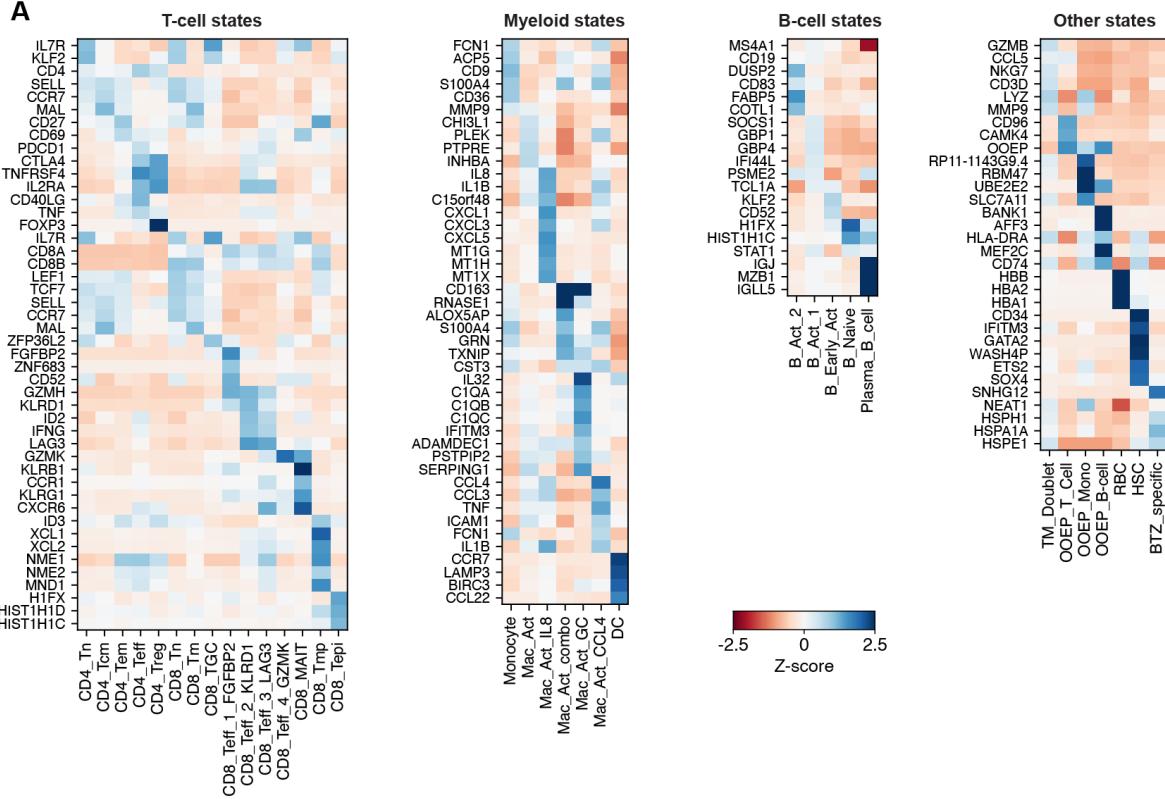


Figure 5.5: Expression of cell-typing markers in each major cell type. (A) The heatmap shows the Z-score of gene expression for cell type markers. Marker genes are identified by differential expression analysis of each Leiden cluster. The Z-score is computed from the mean marker gene expression in each Leiden cluster compared to the mean and standard deviation in all the cells of the same major cell type (T-cell, Myeloid, B-cell, and others).

molecules with the network to drive the observed drug impacts. The heterogeneous population contains various immune cell types as well as resting, activated, and drug-specific states. We applied D-SPIN in an unsupervised manner, using the algorithm to construct a single, integrated regulatory network model that can capture all cell types and cell states by pairwise interactions between gene-expression programs and interactions between gene programs and applied drugs.

For network construction, we chose 30 gene programs as informed by both the Bayesian information criterion (BIC) and the elbow method on gene information capturing (Hastie et al., 2009) (Figure 5.6 A, Section 2.5). With 30 gene programs, 89.0% of the selected highly variable genes have correlations larger than 0.5 with corresponding gene programs in the gene matrix denoised by data diffusion (Van

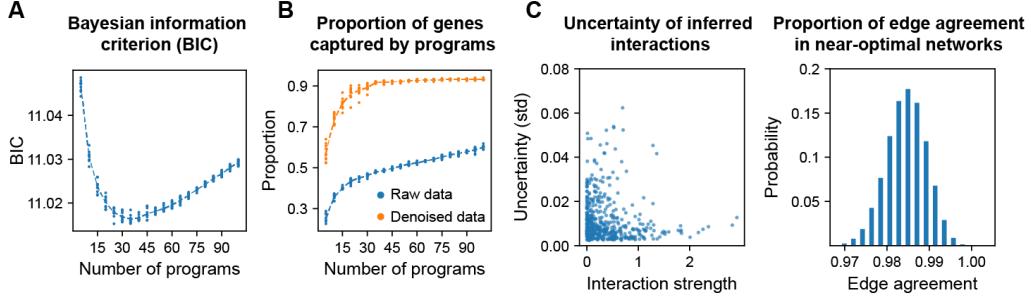


Figure 5.6: The choice of gene-program number; Uncertainty of the inferred network. (A) Bayesian information criterion (BIC) of the gene program decomposition as a function of gene program number. 30 is the optimal program number that minimizes BIC when the effective sample size is 15k cells. (B) The proportion of genes that have a high correlation with their corresponding gene programs as a function of gene program number. In raw data, a high correlation is defined as exceeding 3 standard deviations of gene-gene correlation distribution. In the denoised data by MAGIC (Van Dijk et al., 2018), a high correlation is defined as larger than 0.5. The plot shows that 30 programs is an elbow point of information gain, as detailed in Section 2.5. (C) (left) Uncertainty of inferred network interactions as a function of interaction strength quantified by the standard deviation among alternative network solutions, as detailed in Section 2.5. The uncertainty is generally below 0.06 and much lower compared with interaction strengths. (right) Histogram showing the proportion of edge agreement between alternative solutions of the D-SPIN network inference problem. Alternative solutions are obtained by sampling from the Bayesian posterior distribution given the experimental data, as detailed in Section 2.5. Alternative solutions of D-SPIN are highly consistent with the inferred network, with typically more than 97% of the edges being in the same category of activation, inhibition, and non-interacting.

Dijk et al., 2018). We performed oNMF with seeding to include drug-specific programs, and the programs were sufficient to describe the transcriptional state of all cells in the population (Figures 5.8 AB, Section 5.6). The decomposition yielded known T-cell and macrophage gene-expression programs, including T-cell resting (P5) and activation programs (P7), macrophage anti-inflammatory (P22 M2 macrophage), pathogen response (P27), antigen presentation (P19), and global cell-state programs for T-cells (P6), B-cells (P15), NK cells (P14), and myeloid cells (P20). The function of each program was annotated by a combination of informatics databases (Kuleshov et al., 2016; Szklarczyk et al., 2021; Sherman et al., 2022) and manual annotation, and subsequently validated on a single-cell atlas of human immune cells (Domínguez Conde et al., 2022).

Using the experimental data and the gene expression programs, D-SPIN inferred a

single interaction network \mathbf{J} and a set of drug-specific response vectors $\mathbf{h}^{(n)}$ across all treatment conditions (Figure 5.7 A). The inferred network has low uncertainty, as edge uncertainties are generally below 0.06, while strong interactions have strength above 1. Further, more than 97% of edges remain structurally stable when resampled from the posterior distribution of near-optimal networks (Figure 5.6 C, Section 2.5). Like the K562 cell network, the immunomodulatory drug-response network has a distinct modular structure, with a set of core network modules composed of tightly interacting gene programs (Figure 5.7 A, Figure 5.8 C). The inside-module interactions are mostly positive and account for 76.2% of strong activating edges in the network with strength larger than 1. The modules in the drug data contain groups of gene programs that are expressed together in individual cell types. For example, a module corresponding to anti-inflammatory myeloid cells contains P21 Monocyte, P22 M2 macrophage, and P23 Myeloid growth and a network hub P20 Myeloid cell. Similarly, in T-cells, another module containing P11 Granzyme K T-cell, P12 T-cell cytotoxicity, P13 Effector CD8 T-cell, and P14 NK cell corresponds to the population of CD8 T cells and NK cells. In the D-SPIN model, negative interactions primarily occur between gene programs expressed in different cell types, e.g., P2 T-cell maintenance and P19 Antigen presentation, as antigen presentation happens in B-cells and myeloid cells but not T-cells (Figure 5.7 A).

As a maximum entropy model with a minimal number of parameters, the D-SPIN model only contains a total of 465 network interaction parameters and 30 parameters for each drug condition. However, it can generate samples that are consistent with all profiling experiments of 502 drugs (Section 5.6). Qualitatively, we found visual similarity between the distribution generated by the D-SPIN model on the UMAP embedding and cell state distributions (Figure 5.7 B). Quantitatively, we computed the cosine similarity between the model and data distribution on cell states, and in 92.4 % of the samples the model is $> 90\%$ similar to data (Figure 5.7 C). The cell state distributions produced by D-SPIN also exhibit significantly higher similarity with data distributions compared to the reference distribution obtained by pooling all experimental conditions together to reflect the relative abundance of cell types ("all samples" in Figure 5.7 C). The drugs with lower cosine similarity tended to be drugs that drove the cell population into a few highly specialized cell states, like the proteasome-inhibitor Bortezomib, which produced a strong signature of cell stress in all cells. For these drugs, the D-SPIN model produced a cell-population distribution that had a qualitative agreement with the data but differed quantitatively, likely due to the extremely focused cell type distribution for these drugs. To further access the

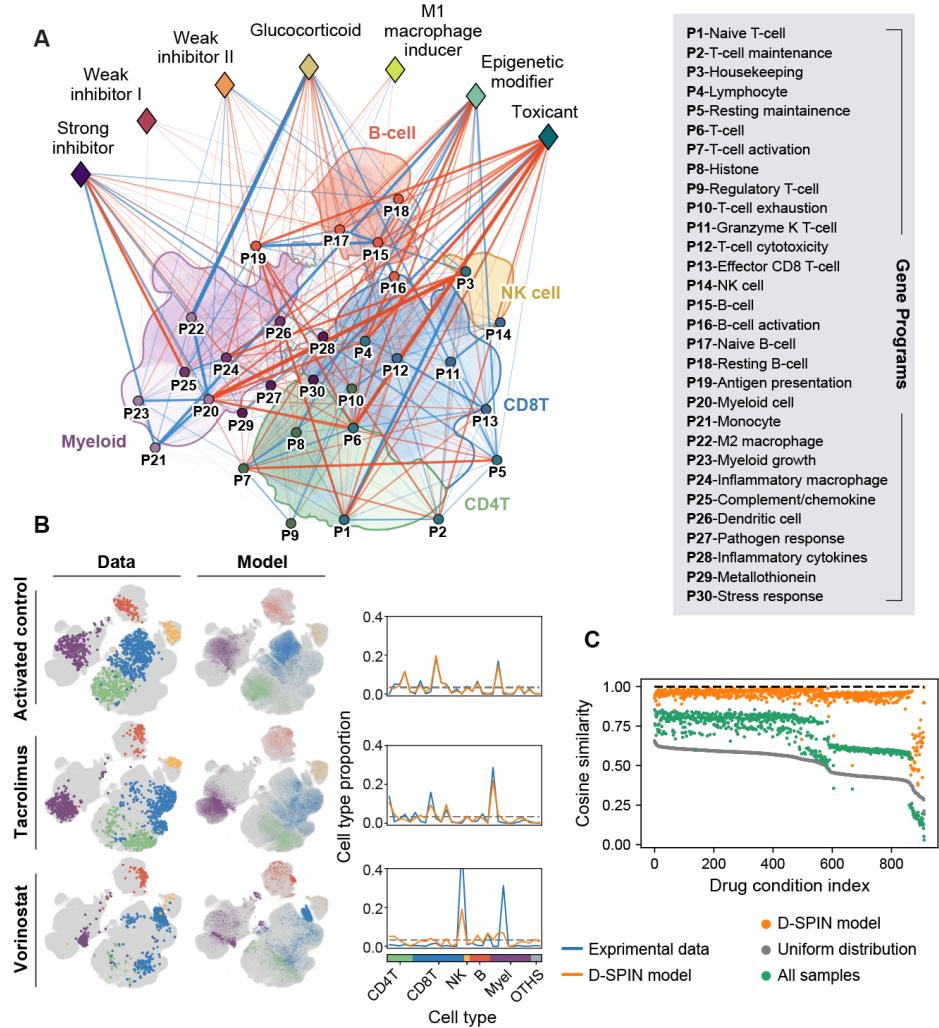


Figure 5.7: D-SPIN derives a drug-response network model from human immunomodulatory drug-response single-cell mRNA-seq profiling experiments.

(A) Network diagram shows the inferred regulatory network model between 30 gene programs (P1-P30, circles), as well as interactions between programs and seven drug classes (diamonds) identified by D-SPIN through clustering the response vectors. In the network rendering, each program is positioned on the UMAP schematic upon the cell state where the program is highly expressed. (right box) Gene programs are functionally annotated through gene ontology annotation tools (Kuleshov et al., 2016; Szklarczyk et al., 2021; Sherman et al., 2022) and manual lookup. (B) (left) UMAP embeddings of experimental cell state distribution compared with the state distribution generated by the D-SPIN model in example conditions including activated control, Tacrolimus treatment, and Vorinostat treatment. (right) Line plots quantify the distribution of cell states of experimental data and D-SPIN models in each example condition. The dashed line in cell state distribution is uniform distribution as reference. D-SPIN models agree well with experimental data in the control and Tacrolimus-treated samples. The model fits less well in the Vorinostat-treated sample but still captures qualitative features of the distribution.

Figure 5.7: (Continued) (C) Scatter plots show cosine similarity between experimental data and cell-state distributions generated by D-SPIN, compared with two reference null models. One null model is a uniform distribution over cell states, and the other null model is pooling all cells from samples of all experimental conditions to reflect the relative abundance of cell states. The cosine similarities between cell state distributions of D-SPIN models and experimental data are higher than 0.9 for 92.4% of conditions.

distributions generated by the D-SPIN model, we employed several other metrics for comparing high-dimensional distributions. Specifically, the errors between data and model for the sample mean and cross-correlation by D-SPIN are, respectively, 3.9% and 7.2% compared to the uniform distribution over all possible states, and the optimal transport distance by D-SPIN distributions is 12.4% of the uniform distribution as a distribution-wise metric (Figure 5.9 A). At the single-state level, the probabilities of states in D-SPIN models correlate with the empirical probabilities of states in experimental data with a correlation of 0.54, 0.58, and 0.48 in example conditions of activated control, Tacrolimus treatment and Vorinostat treatment; the probabilities are not in exact numerical agreement due to the entire high-dimensional space of $3^{30} \approx 10^{14}$ states being far under-sampled by the thousands of cells in each condition in the experimental data (Figure 5.9 B, Section 5.6).

5.3 Classification of small molecules using D-SPIN

In addition to the regulatory network model, D-SPIN also inferred a perturbation response vector $\mathbf{h}^{(n)}$ that quantifies interactions between each drug and the regulatory network. By comparing the inferred response vector with control experiments, we identified 70 out of 502 drugs that had statistically significant interactions with the core gene network (Section 4.6). Moreover, as in the case of the Perturb-seq data, we applied clustering to group drugs into a set of seven classes, including strong inhibitor, weak inhibitor I, weak inhibitor II, glucocorticoid, M1 macrophage inducer, epigenetic regulator, and toxicant, based on their interactions with the underlying regulatory network (Figure 5.7 A, Figure 5.10 A). We named the classes by analyzing the structure of the cell population induced by drugs within the class. In addition to inducing similar transcriptional responses in the underlying cell population, the drug classes identified by clustering of D-SPIN perturbation response vectors, in many cases, belong to specific chemical families with common mechanisms of action.

Three of the seven drug classes are drugs that inhibit activation of the immune-cell population (Figure 5.10 A, Figure 5.12 A), leaving T-, B-, and myeloid cells in the

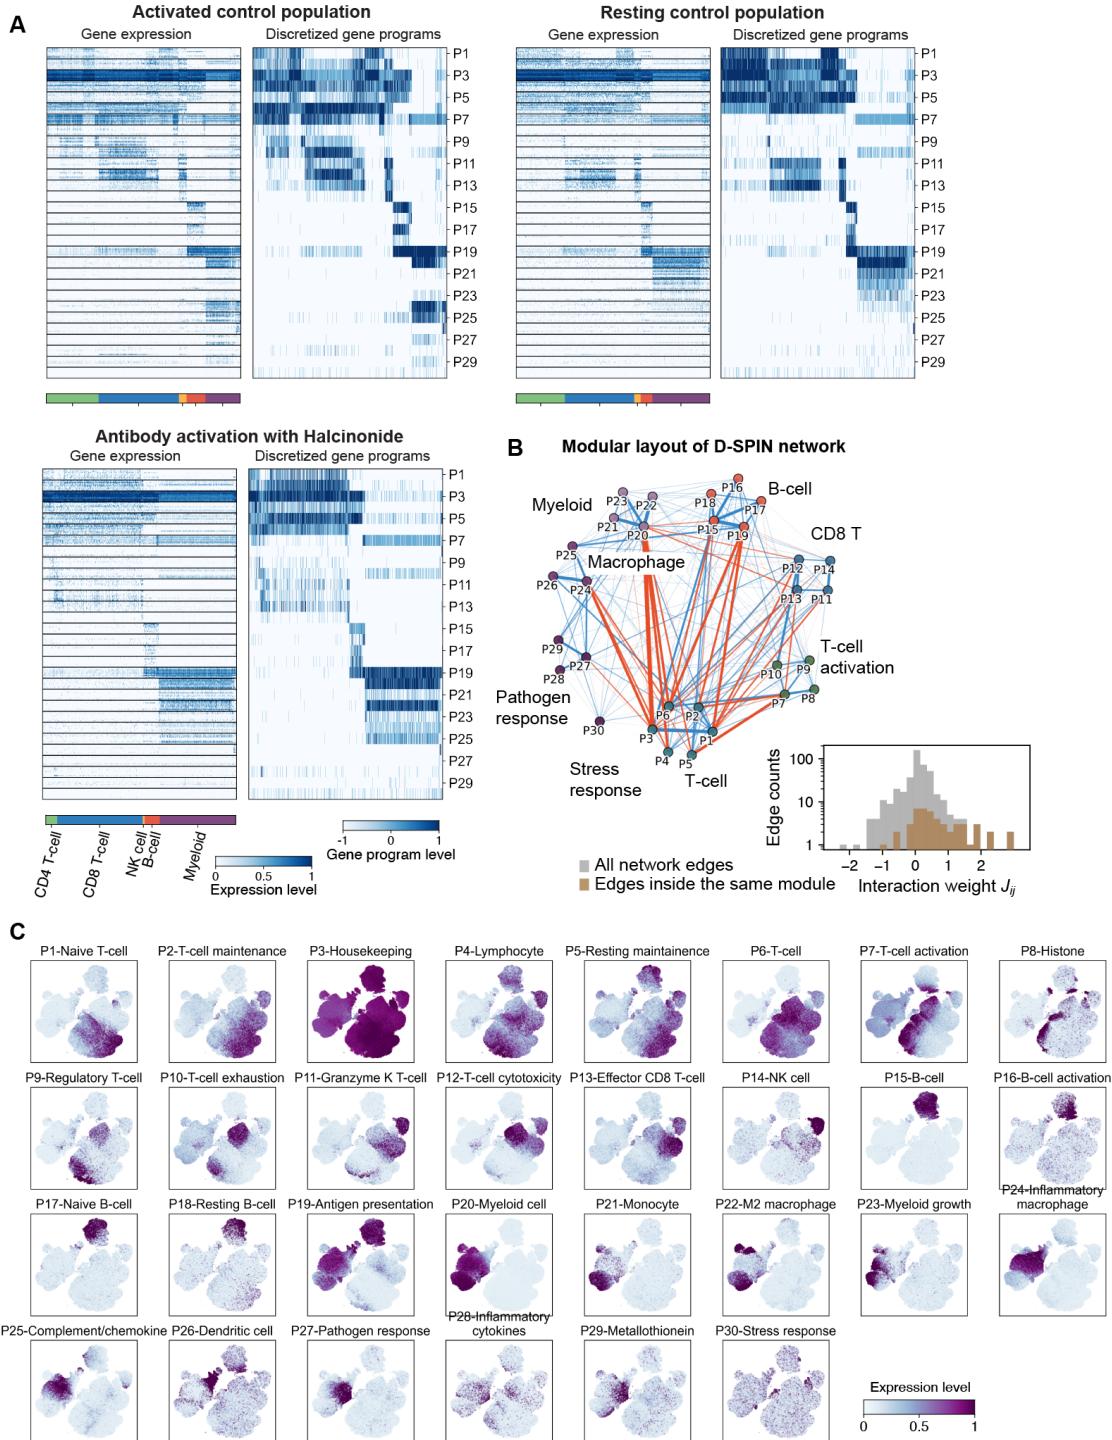


Figure 5.8: Examples of gene program discretization and gene program visualization on immune cell population UMAP. (A) Heatmaps of gene expression and discretized gene program level on (left) activated control population (middle) resting control population, and (right) T-cell receptor activation with Halcinonide treatment. The expression of each gene is normalized by its maximum expression and capped at 0.8 for visualization. The heatmaps demonstrate that discretized gene programs characterize and denoise major expression patterns in the gene matrix.

Figure 5.8: (Continued) (B) UMAP rendering of the expression level of each gene program. Each cell is colored by its expression level on the gene program. The expression level is normalized by the 95th percentile for visualization only. The plots demonstrate that gene programs are specifically expressed by certain cell states, which are localized in separate regions on the UMAP. (C) (top) Network diagram of the regulatory network model by D-SPIN exhibits modular structure with tightly connected subnetworks. The eight network modules correspond to major cell types or cell type functions in the PBMC population, such as T-cells, B-cells, and myeloid cells. (bottom) The histogram quantifies the distribution of all network edges (J matrix entries) and edges inside the same module. Edges inside modules are mostly positive interactions and contain the majority of strong positive interactions in the network.

resting state, including drugs we classified as strong inhibitors, weak inhibitor I, and weak inhibitor II. Across these drug classes, analysis of cell state distributions and UMAP visualizations indicate that these drugs block the transition of T-, B-, and myeloid cell types to the activated state, shifting the balance of the cell population to resting cell states, but acting with a spectrum of inhibition strengths (Figures 5.10 B-D). Very strong inhibitors, including the immunosuppressive drug Tacrolimus and the cancer drug Dasatinib, completely block the activation of T-cells and myeloid cells, inducing a cell population with a similar cell-state profile to unstimulated PBMC population (Figure 5.10 D). Weak inhibitors, such as Temsirolimus, only slightly bias the cell population towards resting states, compared with the activated control population.

The individual drugs that fall into the inhibitor drug classes act on a variety of molecular targets and signaling pathways (Wishart et al., 2006; Skuta et al., 2017) (Figure 5.11 A). The strong-inhibitor class targets include JAK (Tofacitinib, Cerdulatinib, Ruxolitinib), BCR-ABL and Src family kinases (Bosutinib, Dasatinib), and calcineurin (Cyclosporine, Tacrolimus), which are all pathways downstream of T-cell receptor signaling. The drugs in the weak inhibitor I class target another group of receptors including ALK (Crizotinib), FLT3 (Sunitinib), and various growth-factor receptors (Crizotinib, Sunitinib, Nilotinib, Tivozanib). The drugs in the weak inhibitor II class target protein kinases including mTOR (Everolimus, Temsirolimus, Sirolimus) and ATR (Elisentib). Broadly, the analysis suggests that inhibition can be achieved via a range of distinct biochemical pathways and mechanisms. However, the strongest inhibitors in our data have a high affinity for signaling molecules immediately downstream of T-cell receptor activation, including the Src family kinase Lck.

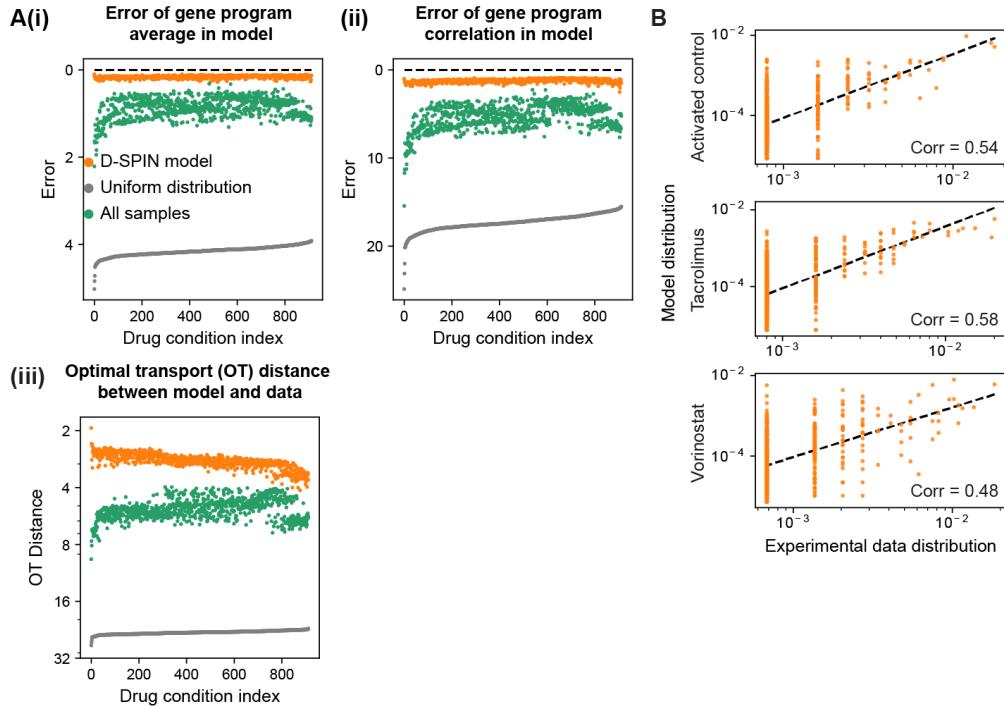


Figure 5.9: D-SPIN generates cell state distributions that are highly similar to experimental data under various metrics. (A) Scatter plots of evaluating how the experimental data distribution is different from the D-SPIN model distribution and two reference null distributions by (i) Euclidean distance between the average of gene program expression, (ii) Frobenius norm of the difference between the cross-correlation of gene program expression, and (iii) Optimal transport distance of probability distributions. The two null distributions are (gray) uniform distribution over all states and (green) the distribution of pooling samples from all experimental conditions to reflect relative cell type abundance, as detailed in Section 5.6. Across all experimental conditions, the distributions defined by D-SPIN quantitatively agree better with the data distribution compared with the two reference distributions. (B) Scatter plot comparisons between state distributions of experimental data and the D-SPIN model. Even though the experimental observations are sparse samples of a few thousand cells in the total $3^{30} \approx 2 \times 10^{14}$ possible states, the probabilities of empirical distribution and D-SPIN model probabilities are moderately correlated with correlation 0.54, 0.58, and 0.48. The plot shows that high-probability states in the D-SPIN model are also high-probability states in experimental data.

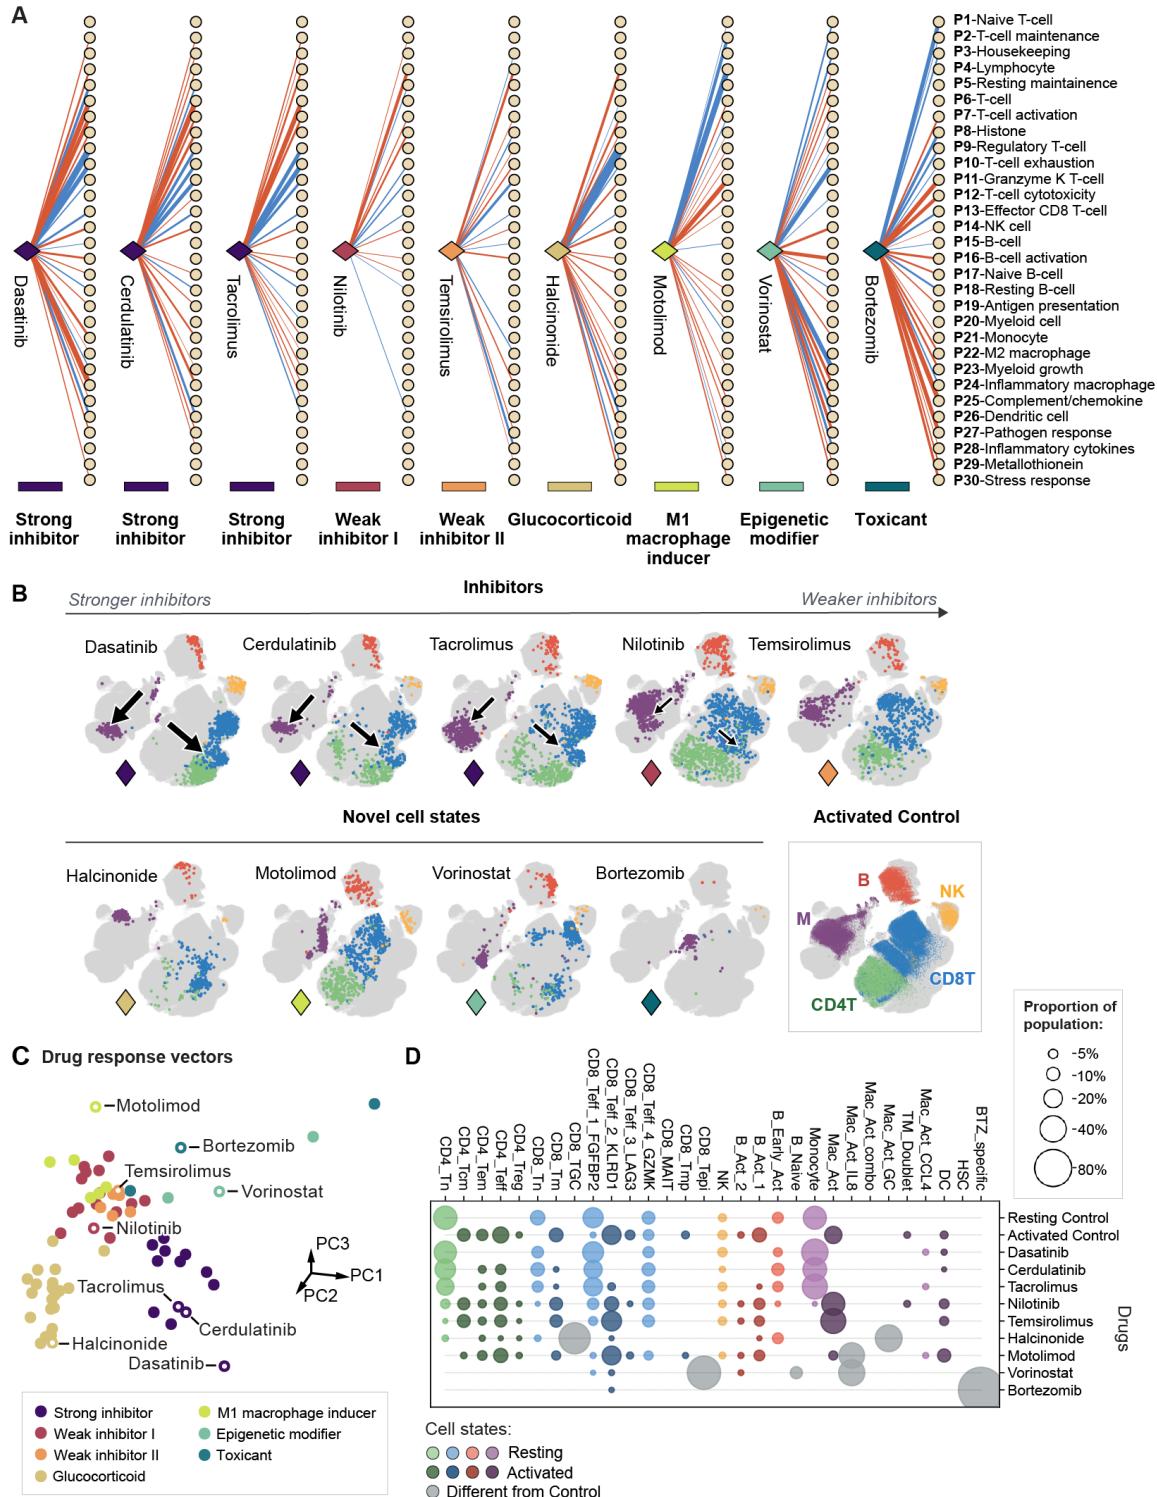


Figure 5.10: **Drug response vectors identified by D-SPIN group drug into categories of phenotypic class.** (A) Diagrams of response vectors inferred by D-SPIN, for example, drugs show positive/negative interactions between the drug and gene-expression programs as blue/red edges. Edge thickness scales with interaction strength. Gene program names for all diagrams are indicated at the bottom of the plot. D-SPIN identifies 7 phenotypical classes of drugs based on their response vectors.

Figure 5.10: (Continued) (B) UMAP embeddings of immune cell population after treatment of antibody and example drugs from different classes. The top row shows 5 immune inhibitors with decreasing strength from left to right, from complete suppression of immune activation (Dasatinib) to slightly increased resting cell population (Temirolimus). The bottom row shows 4 example drugs that each induce novel cell states distinct from both the resting and activated control cell population. Cell state distribution of the activated control population (bottom-right) is shown for comparison. Using response vectors, D-SPIN identifies immune inhibitors with a spectrum of strengths and drugs that produce novel cell states. (C) Scatter plot visualizes the response vectors of identified effective drugs that are statistically different from control in the space of leading principal components (PCs). Each drug is colored by its drug class. (D) Bubble plots show cell-subtype distributions induced by selected drugs as well as the control cell population. Bubbles have sizes scaling with the proportion in the population, and are colored by major cell types and resting/activated classification. With decreasing inhibitor strength, the proportion of the activated immune-cell population (deep colors) gradually increases. Some drugs induce cell states that are different from both resting- and activated-control populations.

While many of the drugs in our dataset have been previously reported as immunosuppressive, including well-known drugs such as Tacrolimus, our single-cell profiling measurements and D-SPIN analysis organized the drugs into a quantitative series and provided a single integrated framework for comparing drug action at the transcriptome scale across a large number of inhibitors with different chemical mechanisms of action. Specifically, immunomodulatory drugs vary in the breadth and specificity with which they interact with biochemical targets, and a therapeutic question is to understand how differences in biochemical preferences translate into transcriptional changes and thus differences in therapeutic response (Schwartz et al., 2017; Zhou et al., 2019; Van de Sande et al., 2023). Our dataset provides the opportunity to perform a detailed comparison between the effects of drugs with the same primary target but different breadth and specificity of secondary targets. Thus, the single-cell assay and D-SPIN model together can identify phenotypic patterns that capture potential differences in the outcome of drug action (Figure 5.11 A).

To quantify and rank the strength of immune inhibition by drugs, we computed an inhibitor strength for each immune inhibitor drug by projecting the D-SPIN drug response vector onto the leading singular vector of all inhibitors in our dataset (Figure 5.12 C, Section 5.6). The resulting quantification of inhibitor strength, when combined with information on biochemical targets (Figure 5.11 A), revealed potential principles of immunosuppression. For example, a specific set of BCR-ABL

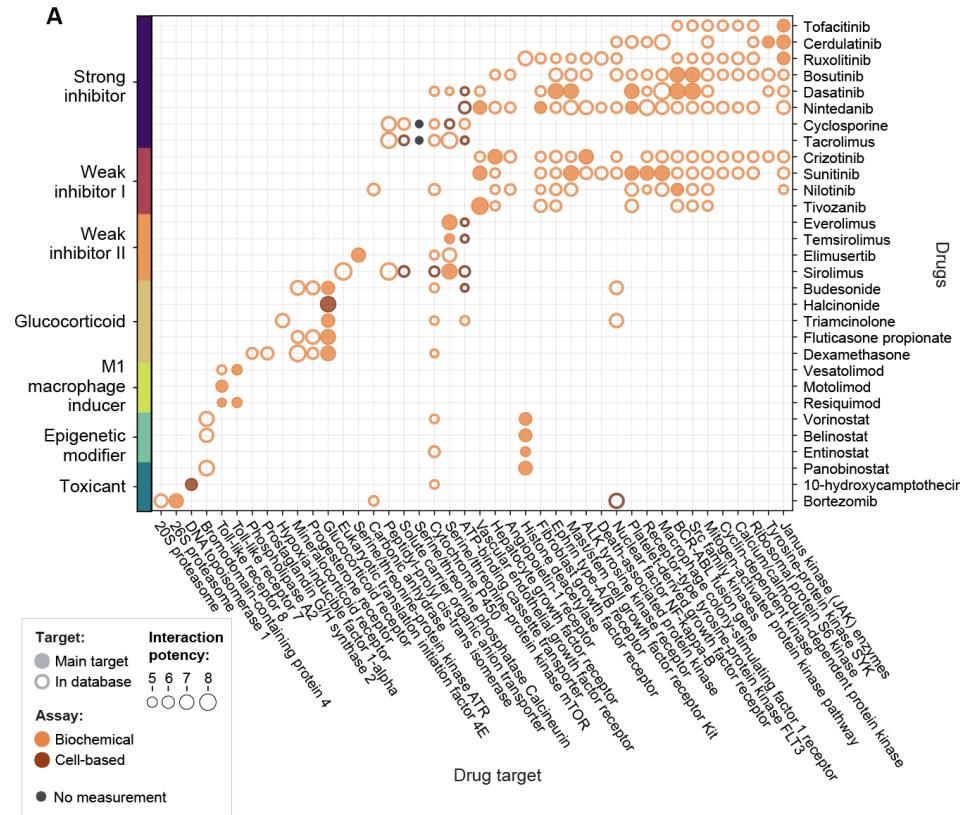


Figure 5.11: Drug classification derived from D-SPIN aligns with known drug targets and mechanisms of action. (A) Bubble plots display the biochemical targets, primary targets, and affinities of example drugs obtained from database (Wishart et al., 2006; Skuta et al., 2017). Bubbles have sizes scaling with interaction potency quantified by biochemical (orange) or cell-based (red) assays. Primary targets are represented in solid dots and other targets in open dots. Drugs in the same phenotypical class identified by D-SPIN can act on a variety of molecular targets and pathways.

kinase inhibitors (Dasatinib, Bosutinib) developed for the treatment of leukemia generated strong inhibition with strengths of 6.57 and 4.61, while another set of BCR-ABL inhibitors, including Nilotinib, fall into the weak-inhibitor class with a strength of 1.67. Biochemically, the BCR-ABL strong inhibitors can be distinguished from the BCR-ABL weak inhibitors by their increased Lck/Src kinase specificity (Figure 5.11 A). The Lck kinase is a central signaling node directly downstream of T-cell receptor activation. Similarly, both as strong inhibitors, the Janus Kinase (JAK) inhibitor Cerdulatinib has a stronger inhibitory strength of 5.31 than the FDA-approved anti-inflammatory JAK inhibitor Tofacitinib with a strength of 3.75. While both drugs have JAK as their primary target, Cerdulatinib

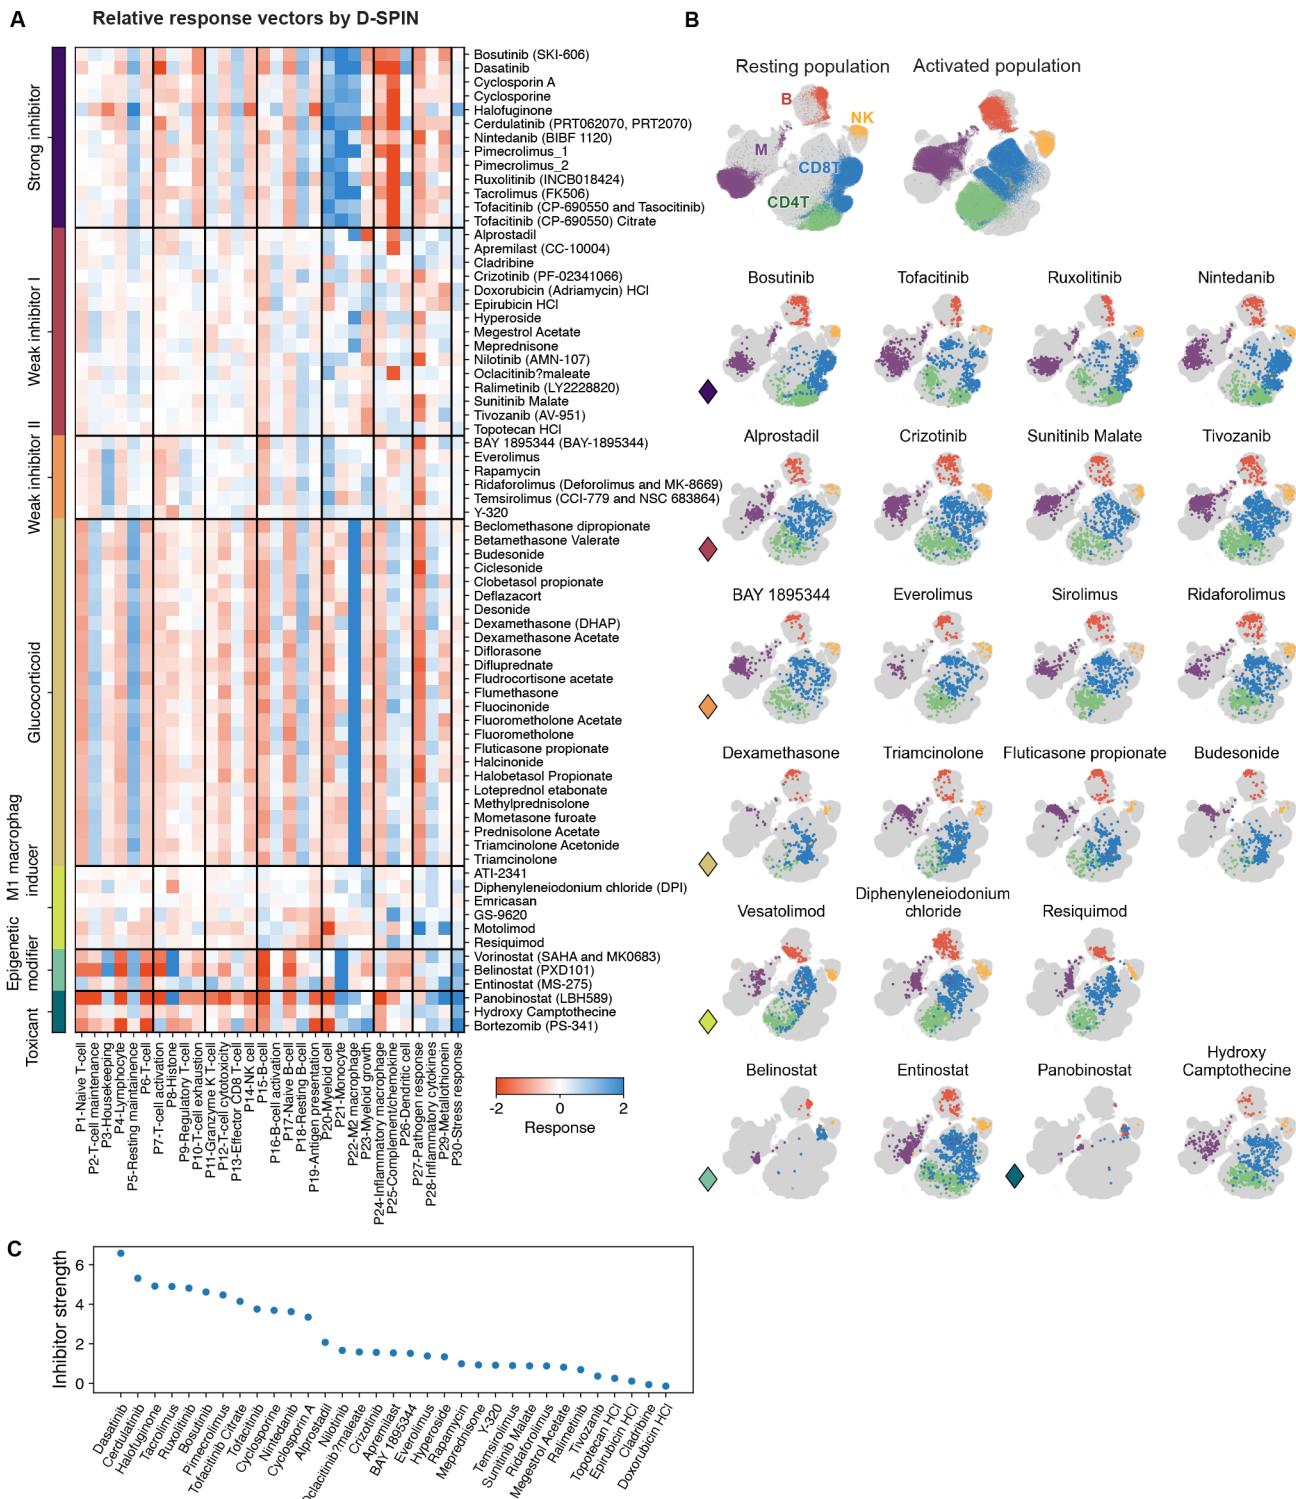


Figure 5.12: Caption on the following page.

Figure 5.12: D-SPIN identifies phenotypical classes of drug categories and strength of immune inhibitors. (A) Heatmap of drug relative response vectors inferred by D-SPIN. The relative response vectors are defined as the change of drug response vector over control samples. D-SPIN identifies 70 effective drugs on the immune population that have statistically significantly different relative response vectors from control samples, as detailed in Section 5.6. We group these drugs into seven phenotypical classes: strong inhibitor, weak inhibitor I, weak inhibitor II, glucocorticoid, M1 macrophage inducer, epigenetic modifier, and toxicant. (B) UMAP embeddings of example drugs from each class. Drugs in the same phenotypical class identified by D-SPIN exhibit similar cell population shifts on the UMAP embedding, while the cell population states vary greatly between different classes. (C) Strength of immune inhibition of all D-SPIN-identified immune inhibitors. Inhibitor strengths are computed by projecting the relative response vector towards the leading singular vector of all inhibitors, as detailed in Section 5.6.

also targets the kinase SYK. Therefore, the response vectors generated by D-SPIN, when combined with data on drug-target binding (Figure 5.11 A), allowed us to identify potential signaling nodes like Lck and SYK that differentiate the transcriptional responses of otherwise related drugs like the JAK inhibitors Cerdulatinib and Tofacitinib (Schwartz et al., 2017).

Beyond the strong and weak inhibitors, D-SPIN identifies a class of drugs that contains glucocorticoids (GCs), steroid-derived small molecules that activate the glucocorticoid receptor (Figures 5.10 A-D). The GC class of drugs includes well-known immunosuppressive drugs including Halcinonide, Budesonide, Triamcinolone, and Dexamethasone. The GCs generate a population response that is similar to the strong inhibitors like Dasatinib. However, the two classes of drugs also have important differences in their interaction with the D-SPIN regulatory network. Specifically, D-SPIN infers that the GCs not only inhibit activation programs in T-cells and myeloid cells, but also activate a specific immunosuppressive gene program in macrophages. The program is named P22 M2 macrophages, as it contains genes associated with the immunosuppressive M2 macrophage including CD163 (Y.-C. Liu et al., 2014; Desgeorges et al., 2019). Therefore, D-SPIN infers important differences in the immunosuppressive effect of GCs and other inhibitors. GCs have weaker suppression of programs associated with chemokine secretion (P25), but induce the M2 macrophage program that is more weakly activated by drugs in the strong-inhibitor classes.

In addition to the suppressive drug classes, D-SPIN also identifies a group of drugs that induce the activation of inflammatory, pathogen-responsive M1 macrophages.

The class includes small-molecule activators of Toll-like receptors, receptors in macrophages that sense the presence of pathogens and whose activation induces an innate immune response. The M1 macrophage inducer class contains TLR7 agonists (Vesatolimod, Resiquimod) and TLR8 agonists (Motolimod, Resiquimod), and produces macrophage states related to host defense that highly express P27 Pathogen response and P29 Metallothionein.

Finally, D-SPIN identifies two classes of drugs, which we call epigenetic modifiers and toxicants, that the model associates with activated cell-stress programs. Toxicants include the proteasome inhibitor Bortezomib, potent non-selective histone deacetylase inhibitor Panobinostat, and DNA topoisomerase inhibitor 10-Hydroxycamptothecin. Toxicants strongly activate the P30 Stress response and have mostly inhibitory interactions with other gene programs, especially generic cell type programs such as P6 T cell, P20 Myeloid cell, and P4 Lymphocyte. The epigenetic-modifiers class consists of histone deacetylases (HDAC) inhibitors which generate an epigenetically disrupted T-cell state (CD8_Tepi) that has elevated expression of histone-related genes such as histone component genes and DNA topoisomerase TOP2A.

5.4 Drug combinations that generate hyper-suppression

Immunomodulatory drugs are often used in combinations, and drug combinations could in principle be used to tune the transcriptional cell-state distribution of the immune cell population to meet therapeutic goals. However, the design of drug combinations at the transcriptome scale is challenging due to a large number of potential drug-gene interactions. D-SPIN models provide a framework to compare the action of individual drugs on gene expression programs within the regulatory network in order to identify drug combinations with potentially useful therapeutic applications.

Therefore, we selected 10 drugs from different drug classes identified by D-SPIN, profiled all pairwise combinations experimentally, and applied D-SPIN to interpret the mechanisms of combination action (Figure 5.13 A). By comparing the D-SPIN inferred response vectors of drug combination with single drugs, we found that most of the drug interactions are additive or subadditive on the gene program level where the effect of drug combination is equal to, or slightly weaker than, the linear sum of single drug effect. Other types of drug interaction include dominant (single drugs have different effects on the gene program and the combination follows one

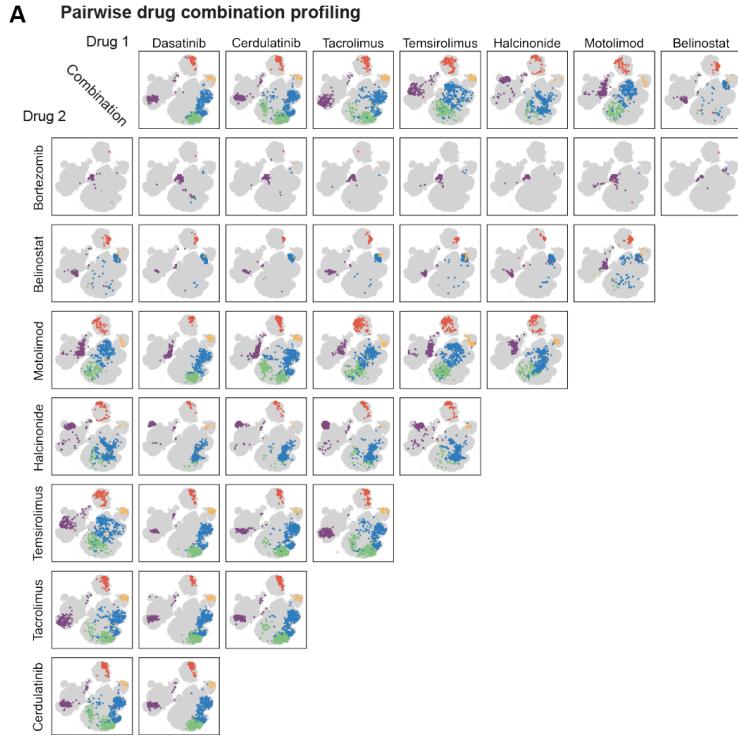


Figure 5.13: Drug combination profiling discovers novel cell states. (A) UMAP embeddings of selected single drugs and pairwise drug combinations. We performed pair-wise drug combination profiling for 12 selected drugs (8 shown here) from different categories. (a). Strong inhibitor: broad-spectrum kinase inhibitor Dasatinib, BCR-ABL, and Src tyrosine kinase inhibitor Bosutinib, SYK/JAK kinase inhibitor Cerdulatinib, calcineurin phosphatase inhibitor Tacrolimus; (b). Weak inhibitor: mTOR kinase inhibitor Temsirolimus; (c). Glucocorticoid: Halcinonide; (d). M1-macrophage inducer: Motolimod; (e). Epigenetic modifier: Vorinostat, Belinostat; (f). Toxicant: Bortezomib; (g). Others: Forsythin, Naproxen sodium. Forsythin and Naproxen cause the loss of macrophage population in single-drug profiling, which is not observed in repeats in drug combination profiling.

of them), synergistic (the combination is stronger than the linear sum of single drug effects), and antagonistic (the combination is weaker than each single drug effect). Quantitatively, 83.0% of drug interactions are additive or subadditive on gene programs, 9.4% are dominant, 6.2% are synergistic, and 1.4% are antagonistic (Section 5.6). The additive interactions between drugs can recruit a combination of transcriptional programs from single drugs, creating novel cell states or population states, especially between drugs that have distinct impacts.

To determine how the additivity between single drugs on different programs generates novel biological outcomes, we first investigated the network-level similarities

and differences between glucocorticoids (GCs) and strong inhibitors. In our experiments, GC action on the D-SPIN regulatory network was similar to the action of strong inhibitors, with both classes of drugs suppressing the activation of T-cells, macrophages, B-cells, and NK-cells. However, GCs induce weaker suppression, allowing cells to activate inflammatory cytokines, including CXCL10 and CXCL11 in the P25 complement/chemokine program, whereas strong inhibitors, including the Src family inhibitor Dasatinib, strongly suppress these cytokine programs. GCs uniquely induce an M2 macrophage gene-expression program, which could be useful in some therapeutic settings, due to the expression of anti-inflammatory signaling molecules. The difference in impact between GCs and strong inhibitors is seen in the comparison of the inferred interactions with the network of the GC Halcinonide with the inhibitor Dasatinib (Figure 5.15 B). Gene programs P22 M2 Macrophage, P25 Complement/chemokine, and P21 Monocyte are induced or repressed more strongly by one drug than the other.

Therefore, we investigated whether combinations of glucocorticoids and strong inhibitors could induce both the M2 macrophage activation and cytokine-repression effects, by performing single-cell mRNA-seq on cell populations treated with the glucocorticoid Halcinonide in combination with the strong inhibitor Dasatinib across a range of doses (Figure 5.16 A). We used D-SPIN to analyze the impact of the drug combination on the regulatory network and underlying cell population. We found that the drug combination induces both the strong immune suppression of Dasatinib as well as the induction of M2 macrophage gene-expression programs for a specific range of concentrations (Figures 5.15 BC). The response vectors for the drug combinations show that the impact of Dasatinib and Halcinonide are additive at the gene-program level across combinations of different concentrations. The response vectors for the drug combination indicate significantly higher P22 M2 macrophage induction than the single drug effect of either Dasatinib or Halcinonide (Figure 5.15 D), while also exhibiting increased suppression of P24 Inflammatory macrophage over the single drugs. The response vectors also maintain strong activation of the resting program P21 Monocyte. However, the combinatorial responses do not exceed the sum of individual drug effects, so we classify the effect as additive or subadditive rather than synergistic. Similarly, on T-cell programs, the drug combination causes a subadditive activation on P5 Resting maintenance, and additive inhibition on P7 T-cell activation and P12 T-cell cytotoxicity (Figure 5.15 A).

The combination of Dasatinib and Halcinonide, therefore, generated a hyper-repressed

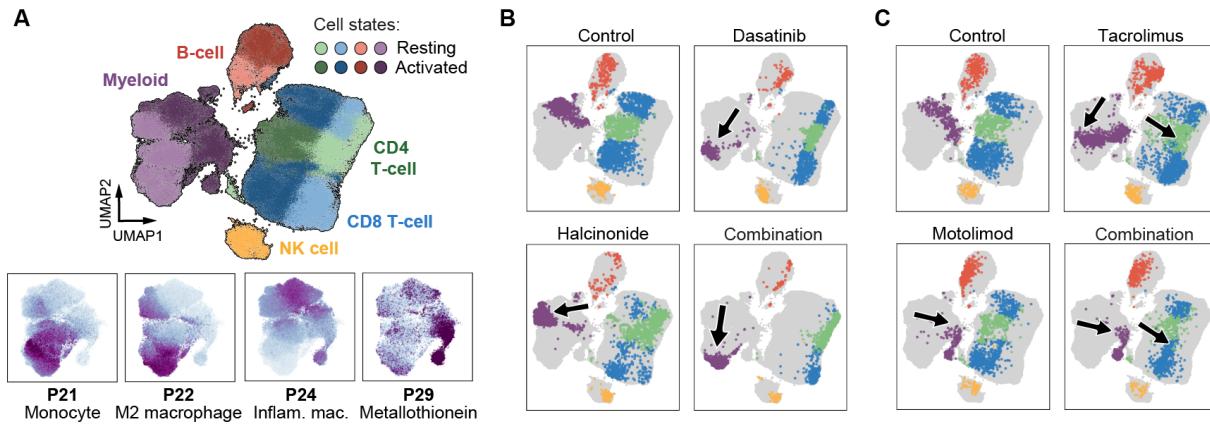


Figure 5.14: Two drug combinations generate novel cell states distinct from single drugs. (A) UMAP embeddings of activated control, single-drug-treated, and drug-combination-treated cell population, with arrows indicating different myeloid cell state changes. The drug combination induces a novel macrophage state that occupies a distinct region on the UMAP. (B) UMAP embedding of cell states of the two drug dosage titration experiments, Dasatinib plus Halcinonide and Tacrolimus plus Motolimod. Cell states are colored by major cell types and resting/activated classification. (bottom) UMAP rendering of example myeloid programs, where cells are colored by the expression level of each gene program. (C) UMAP embeddings of activated control, single-drug-treated, and drug-combination-treated cell population, with arrows indicating different cell state changes. The drug combination induces a novel cell population state with inhibited T-cell states and Motolimod-induced-like myeloid states.

cell population with a repressed T-cell population and a macrophage cell population with activation of both resting myeloid program P21 and M2 macrophage program P22 (Figure 5.14 B). The modulation is monotonic with doses as reflected on the response vectors, and manipulation of the drug doses allowed conversion of the cell population from activated to intermediate suppression (resting-like) to hyper-repressed M2 macrophage-like states (Figures 5.14 CD, Figure 5.16 C, Section 5.6). Conceptually, by combining the two drugs, we are effectively able to fine-tune the transcriptional state of the macrophage population by exploiting a superposition interaction. Here, the drug combination recruits a combination of gene-expression programs in a dose-dependent fashion, which is also observed in a previous study performed at the single-protein level (Geva-Zatorsky et al., 2010).

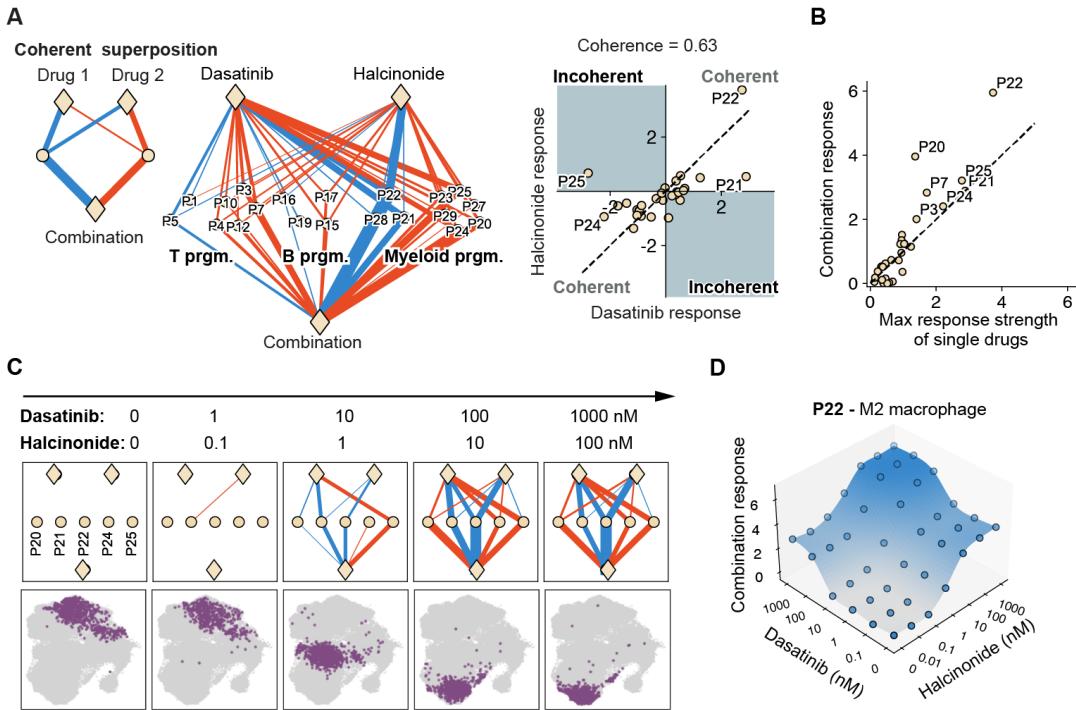


Figure 5.15: The combination of Dasatinib and Halcinonide produces novel cell states with coherent superposition. (A) (left) Schematic diagram of coherent superposition where the single drugs both have activation/inhibition effects on downstream gene programs and the combination drug response is the superposition of single drug effects. (middle) Diagram of response vectors for Dasatinib and Halcinonide alone and their combination. The two single drugs and their combination have coherent effects on gene programs by activating/inhibiting the same set of programs. (right) Scatter plots compare the response vectors on each gene program for Dasatinib and Halcinonide. The two drugs have a high coherence score (0.63) quantified by the cosine similarity, but different strengths of interactions with gene programs. (B) Scatter plots of the response vectors on each gene program for the drug combination and maximum response strength of two single drugs. The combination responses on gene programs are a superposition of single-drug responses, as they are generally higher than the maximum strength of the single-drug responses. (C) (top) Diagram of response vectors for two single drugs and their combination on most impacted myeloid programs under different dosages. The combination response is a superposition of coherent single-drug interactions across dosages. (bottom) UMAP embeddings of myeloid states for each drug combination dosage. The novel macrophage state is induced by a combination of gene program recruitment of single drugs. (D) Surface plot shows response vectors on program P22 under different dosage combinations, where dots are from experimental data and the surface is spline interpolation. On the program, each single drug response exhibits a sigmoid-shaped curve with the logarithm concentration but with different saturating strengths, and the combination response is approximately the sum of the two sigmoid-shaped single drug responses, showing the signature of superposition/additivity.

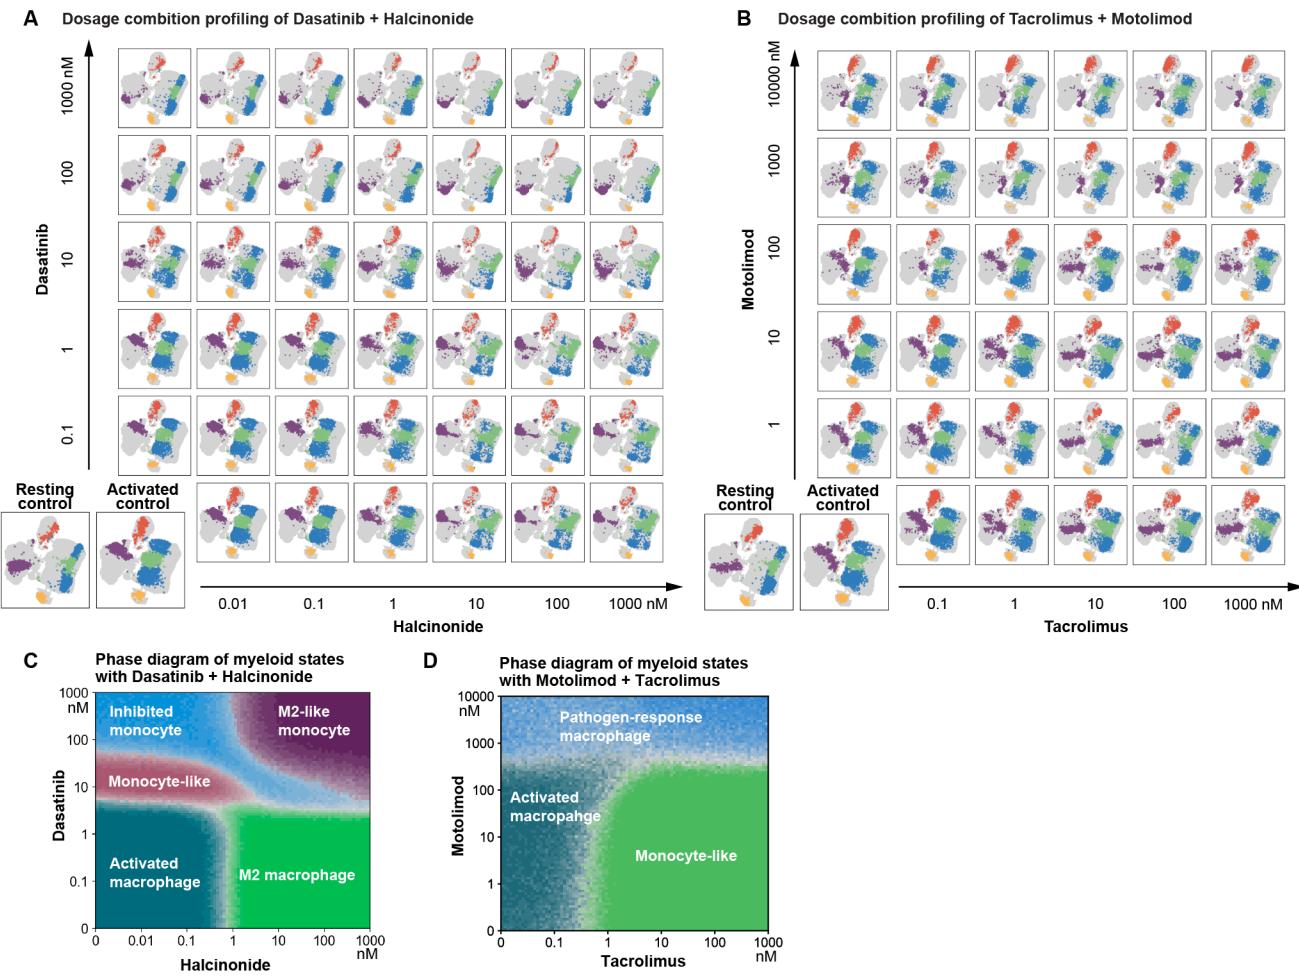


Figure 5.16: Drug dosage combination profiling reveals phase diagrams of drug combination effect. (A) UMAP embeddings of dosage combination profiling of Dasatinib plus Halcinonide. The two drugs create a Monocyte-like resting myeloid population at a medium level of combination (10 nM Dasatinib, 1 nM Halcinonide) and a novel combinatorial myeloid state expressing both Monocyte and M2 macrophage gene programs. On T-cells, the drug combination also produces inhibited T-cell states. (B) UMAP embeddings of dosage combination profiling of Tacrolimus plus Motolimod. The drug combination produces myeloid states that are similar to Motolimod-induced pathogen-responding myeloid states, and intermediate T-cell states that have gene program expression signatures of both drugs. (C) The phase diagram of myeloid states under different drug dosage combinations of Dasatinib and Halcinonide depicted by D-SPIN. The phase diagram is computed by evaluating the likelihood of each reference cell state (activated macrophage, M2 macrophage, etc.) at different dosage combinations using spline interpolation of response vectors of profiled dosages as detailed in Section 5.6. At a low dosage of Dasatinib, the myeloid state transitions from the activated macrophage state to the M2 macrophage state as the Halcinonide dosage increases. At high dosages of Dasatinib, increasing Halcinonide dosage produces monocyte-like, inhibited monocyte, and M2-like monocyte states.

Figure 5.16: (Continued) (D) The phase diagram of myeloid states under different drug dosage combinations of Tacrolimus and Motolimod depicted by D-SPIN. The phase diagram shows the dominance of Motolimod-induced pathogen-response macrophage states over Tacrolimus-induced monocyte-like states.

5.5 Drug combinations that activate the innate immune system while repressing T-cell activation

In a variety of therapeutic settings, it can be advantageous to modulate the balance between adaptive (T-cell-driven) and innate (macrophage-driven) immune responses. For example, T-cell responses may react adversely with transplanted organs and are therefore suppressed in transplant patients. Unfortunately, infections become a significant cause of mortality as a result of immunosuppression. Thus, the induction of an innate immune response in the background of immunosuppression provides a strategy for combating infection (Jay A Fishman and Rubin, 1998; J. Fishman, 2017; Pereira et al., 2020). To understand whether D-SPIN could be applied to tune the balance between different immune cell states at the population level, we use D-SPIN model-identified M1 macrophage-inducing small molecule Motolimod, as well as standard immunosuppressive small molecules, including Tacrolimus, which are revealed to be mostly incoherent with Motolimod in their impact on the underlying gene regulatory network (Figure 5.17 A). Specifically, the global cosine similarity between Motolimod and Tacrolimus responses was only 0.01 and the two drugs are biased to impact T-cell (Tacrolimus) or macrophage (Motolimod) gene-expression programs, with Tacrolimus activating resting T-cell programs and Motolimod activating a pathogen-responsive macrophage program, P27 (Figure 5.17 B).

To exploit the incoherence between Tacrolimus and Motolimod, we added the drugs in combination across a series of doses, applied single-cell mRNA-seq to profile the resulting cell population, and analyzed the combinatorial drug response with D-SPIN (Figure 5.16 B). D-SPIN identifies a concentration regime (Figure 5.14 C, Figure 5.16 D) in which the cell population contains activated macrophages with activated pathogen response program P27 while maintaining T-cell repression with inhibited T-cell activation program P5, with higher than 1000 nM Motolimod and higher than 100 nM Tacrolimus (Figure 5.17 C). Specifically, the two single drugs have different impacts on myeloid programs: Tacrolimus activates the resting program P21 Monocyte, and Motolimod activates the pathogen response programs P27. The drug combination exhibits non-additive, dominant interactions on the programs P21, where the combination follows the no-effect of Motolimod instead of the acti-

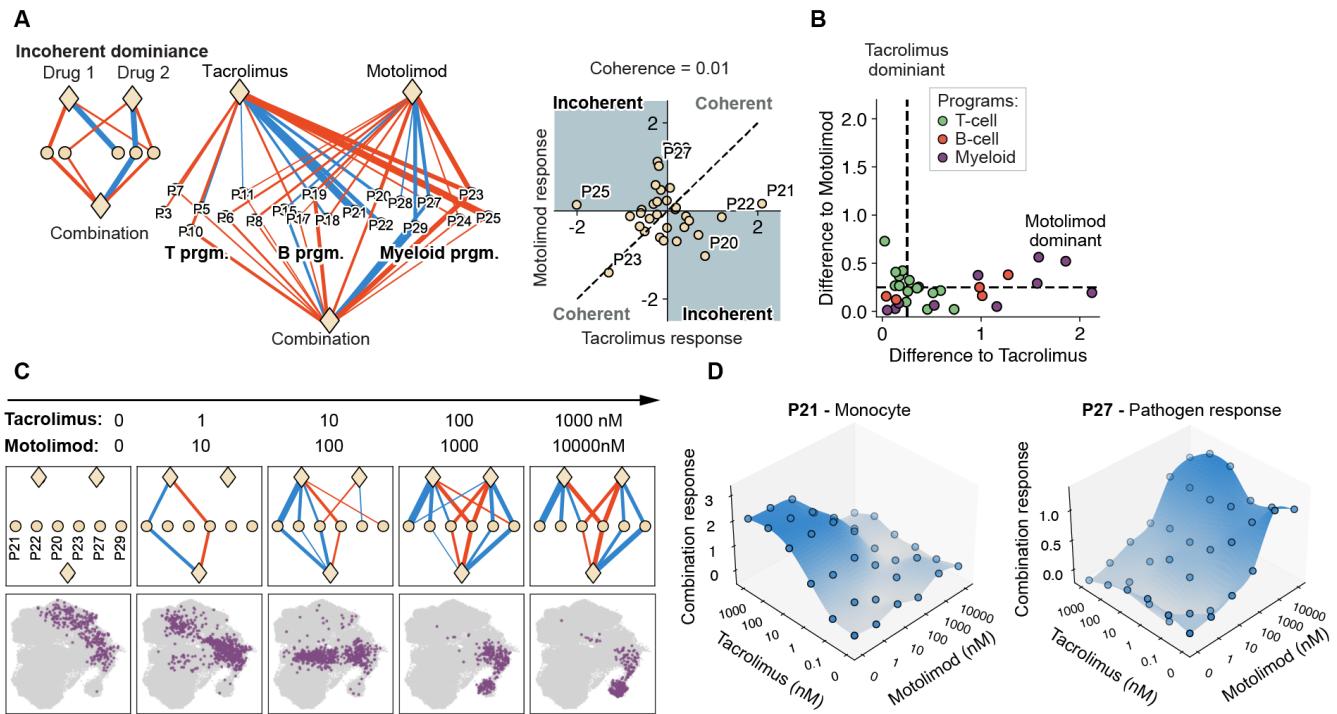


Figure 5.17: The combination of Tacrolimus and Motolimod produces novel cell states with incoherent dominance. (A) (left) Schematic diagram of incoherent dominance where the single drugs have different activation/inhibition effects on gene programs and the combination response is dominated by one of the single drugs. (middle) Diagram of response vectors for Tacrolimus and Motolimod alone and their combination. The two drugs have incoherent effects as they activate/inhibit different sets of gene programs, and the combination effects follow one of the single drugs. (right) Scatter plots compare the incoherent response vectors of the two drugs with a low coherence score (0.01). (B) Scatter plots of the difference between combination response and two single drug responses. On myeloid and B-cells, the response is more similar to, therefore dominated by Motolimod. On T-cells, some programs are dominated by Tacrolimus and some are dominated by Motolimod. (C) (top) Diagram of response vectors for two single drugs and their combination on most impacted myeloid programs under different dosages. At high concentrations of Motolimod, the activations of Tacrolimus on P21 and P22 are dominated by the no-effect of Motolimod. (bottom) UMAP embeddings of myeloid states for each drug combination dosage. The dominance of Motolimod over Tacrolimus on myeloid programs induces myeloid states similar to the effect of Motolimod alone. (D) Surface plot shows response vectors on (left) program P21 and (right) program P27 under different dosage combinations. On the program P21, the activation effect of Tacrolimus is abolished by increasing the dosage of Motolimod. On the program P27, the activation of Motolimod is not influenced by the dosage of Tacrolimus.

vation of Tacrolimus (Figure 5.17 D). The non-additive interactions are important because Motolimod's dominance over Tacrolimus on P21 leads to the pathogen-responding M1 macrophage state, similar to the state induced by Motolimod alone. Conceptually, the Motolimod and Tacrolimus drug combination demonstrates that drugs impacting separate cell types can be combined to tune the structure of an immune cell population.

5.6 Supplementary information

High-throughput drug perturbation experiments on PBMCs

All high-throughput drug perturbation experiments were performed in cell-repellent treated 96-well plates (Greiner Bio-One 07-000-640). We selected the drug library from the SelleckChem Immunology and Inflammation-related library (L4100), or the FDA-approved drug library (L2000), which were supplied at a stock concentration of 10mM. The layout for drugs and/or drug combinations for each 96-well plate was pre-defined before starting the experiment. The final working concentrations for drugs varied between 0.01 nM to 10 uM. To stimulate PBMCs into an activated state, some wells included the use of ImmunoCult Human CD3/CD28 T Cell Activator (Stemcell Technologies 10971), which was used at the commercially recommended final concentration of 5 uL/200 uL. After completing drug dilutions, we dispensed drugs into a 96-well plate layout, loaded 200,000 PBMCs into each well, and cultured for 24 hours in a tissue culture incubator at 37 °C and 5% CO₂. All working drug dilutions and cell suspensions were prepared in RPMI-1640 + 1% Pen/strep + 10% FBS. The final concentration of DMSO within these experiments was always below 0.1%, which has been shown to produce minimal effects on immune cell activation (Holthaus et al., 2018).

High-throughput single-cell profiling using sample multiplexing techniques

After 24 hours of drug exposure, we harvested cells for single-cell profiling. Cells were enzymatically dissociated into a single-cell suspension and multiplexed using MULTI-seq lipid-modified oligos (McGinnis et al., 2019). One experiment (MULT-19) was performed using 10X Chromium Cellplex Kit (1000261), and manufacturer protocols were followed. Although plates were treated to be cell repellent, activation of certain cell types within the PBMC mixture can produce strong adhesive tendencies to the plate and each other. Thus, we enzymatically digested both the cell suspension and the plate surface with 100 uL of TrypLE (Thermo Fisher 12604021) to ensure the complete recovery of all cells. Briefly, the digestion was performed

as follows: Using a multichannel pipette, all cells from each plate were transferred to a conical V-bottom 96-well plate (Thermo Fisher 249935), and then 100 μ L of TrypLE was added to each well of the original plate and kept at room temperature. The conical V-bottom well plate, which contains the cells originally in suspension, was centrifuged at 400 rcf for 4 minutes. The supernatant was removed, and 100 μ L of TrypLE was added to each well of pelleted cells and triturated gently. Both the original plate and the v-bottom plate were then incubated at 37 °C for 2 minutes simultaneously. After checking that any adherent cells were removed from the original plate, the cells in each plate were triturated and pooled back together. Cells were then labeled with MULTI-seq lipid-modified oligos according to the original protocol (McGinnis et al., 2019) and profiled using the 10x Chromium 3' chemistry (v3.1, PN-1000121) or using the HT version (HT Kit v3.1, 1000370). Cells were super-loaded onto each lane of a Chromium chip at 25,000–40,000 cells per lane, or onto a Chromium HT chip at 60,000 cells/lane. Sample tags were purified and amplified according to the MULTI-seq procedure or according to manufacturer instructions for Cellplex. Libraries were sequenced using Illumina Novaseq S4 chips at the UCSF CAT Sequencing facility to a target depth of 40,000 transcriptome reads/cell barcode and 5000 sample tag reads/cell barcode. Cells that were multiply labeled by sample tags were computationally removed from the analysis.

Time-course profiling of T-cell mediated immune activation

A reverse time course of CD3/CD28 T Cell Activator was performed using a Tecan Evo robot (Caltech Protein Expression Center). Briefly, cryopreserved PBMCs from healthy donors were thawed and rested for 16 hours as described above. Cells were then pelleted and dissociated with gentle trituration and seeded into 96-well plates at 200,000 cells/well in 150 μ L of RPMI-1640 + 10% FBS + 1% P/S. The cells were then brought to a Tecan Freedom Evo2 liquid handling robot, which was programmed to dispense CD3/CD28 activator into each well at 30-minute intervals for 30 hours. The CD3/CD28 activator was diluted 1:10 in RPMI-1640 + 1% Pen/strep + 10% FBS, and 50 μ L of the dilution was added at each time point. One well which did not receive any activator served as timepoint 0. Between dispenses, the plate was held in an onboard incubator at 37 °C and 5% CO₂. After the conclusion of the time course, cells were immediately harvested for single-cell profiling using multiplexing tags (MULTI-seq) as described above.

Alignment and de-multiplexing

We used the Cell Ranger software (10x Genomics) to align single-cell RNA sequencing reads to the transcriptome. For experiment batches MULT-6, 7, 8, 9, 10, 12, 13, 14, 15, 16, TC, we used Cell Ranger 3.1.0. For experiment batches MULT-17, and 18, we used Cell Ranger 5.0.1. For experiment batches MULT-19 and MULT-20, we used Cell Ranger 6.1.2. The 10x Genomics hg19-1.2.0 genome build was used as a reference transcriptome.

Except for MULT-19, where we used the 10X Cellplex pipeline to demultiplex the conditions of each cell, we used a custom Python pipeline to assign cells to each condition based on the sequenced tags following the pipeline in MULTI-seq (McGinnis et al., 2019).

Filtering, normalization, and clustering of the drug profiling dataset

The single-cell gene expression normalization in the drug profiling dataset is generally identical to the Perturb-seq dataset. We normalized the raw counts by each cell with the same scaling factor 10^4 , and gene expressions were log-transformed for the following analyses except for the highly variable gene selection.

The highly variable gene selection was based on the zero-inflated Poisson expression model we developed (Section 4.6). The drug profiling dataset contains (1) a large dataset of single drug and drug combination profiling (Drug Profiling) and (2) a dataset of follow-up drug dosage combinations (Dosage Combination). The gene filtering was performed on the Drug Profiling dataset, and Dosage Combination follows the same filtering. To avoid noise in low-expression genes, we excluded genes that only expressed in fewer than 0.1% of the cells in the Drug Profiling dataset. (In the Perturb-seq dataset, the low-expression genes were already filtered out by the authors in the published data). We selected 3,287 highly-variable genes in the remaining genes, which is $\log 10(1.1)$ above the curve fitting Equation 4.5 for the zero-inflated Poisson distribution model.

Collected single-cell profiling data have batch effects across experiments, likely due to variations in cell culture conditions and technical differences across sequencing machine runs. To visualize, cluster, and cell type single-cells in the Drug Profiling dataset, we corrected batch effects by constructing latent representations on raw gene counts of 5,000 selected genes (scipy.pp.highly_variable_genes) with scVI (dispersion='gene-batch', n_latent=10, n_layers=2, max_epoch=200, batch_size=5120) (Lopez et al., 2018). We constructed a 15-nearest-neighbor graph (scipy.pp.neighbors)

on the latent representation, computed Leiden clustering for cell clusters (scipy.tl.leiden), and conducted UMAP embedding for visualization (scipy.tl.umap). The cell subtype identities of each Leiden cluster were annotated manually by comparing the most significant up-regulated genes of the cluster with the remaining cells. We removed three clusters that mostly came from the experimental batch “MULT-6” and were putatively cell debris, as they expressed the maternal effect gene OOEP and disrupted transcriptional states were present. We also excluded red blood cells (RBC) and plasma B-cells for D-SPIN analysis, as they are small and isolated cell populations with distinct transcriptional profiles.

The Dosage Combination dataset followed the same normalization and logarithm transformation pre-processing. We visualized the data by batch effect correction with scVI (dispersion='gene-batch', n_latent=10, n_layers=2, max_epoch=100, batch_size=5120) followed by the same neighborhood graph, Leiden clustering, and UMAP embedding procedures.

oNMF decomposition of single-cell dataset

In the drug profiling study, we used a subset of 3.3×10^5 single cells using the equal-balancing scheme, sampling 1,000 single cells from each Leiden cluster of cell states. In drug profiling data, the sampling scheme was to better capture gene programs in drug-specific cell states that composed a small proportion of all the cells. In the drug profiling study, as specific drug-induced cell states were a small fraction of the whole dataset, we increased the resolution of oNMF to $K = 50$. During preliminary runs of oNMF, we identified drug-specific response programs such as Pathogen response and Stress response. We then manually curated random seeding to select the oNMF decomposition that included these target gene programs. We then combined the 50 programs into 30 by merging programs with similar biological functions and distribution of cell-subtype expression.

Embedding and comparison of D-SPIN and data distributions

To access the alignment between samples generated from D-SPIN and experimental data (or synthetic data as in Section 3.5 and Section 3.5), we employed several measures, including UMAP, cell state distributions, distribution moments, optimal transport and exact probability distribution.

For small networks, the number of possible states s was relatively small and the distribution defined by D-SPIN could be explicitly computed, such as the HSC network and the cell-cycle subnetwork of the Perturb-seq data. We directly took

samples from the D-SPIN distributions, computed joint embedding of D-SPIN samples and experimental data with UMAP, and computed Leiden clustering on the nearest-neighborhood graph obtained during the computation of UMAP. As states in the model are discretized states $[-1, 0, 1]$, to avoid the neighborhood graph being disconnected due to frequent states forming cliques between themselves, we added a Gaussian noise on sample states with zero mean and 0.1 standard deviation.

For larger networks, the total 3^M possible states were impossible to enumerate, and the network encompassed a more complex landscape of cell states, such as the network of the drug profiling dataset. For a comprehensive evaluation of state distribution by D-SPIN and experimental data, we sampled 10^6 samples by MCMC (Section 2.4) for each experimental condition as an empirical distribution of the D-SPIN model, and performed the following analysis.

1. UMAP embeddings of samples from the D-SPIN model. In the drug profiling dataset, there were more than 1,000 conditions and the total number of D-SPIN-generated states was more than 10^9 (10^6 for each condition) and intractable for joint UMAP embedding. Therefore, we projected the states generated by D-SPIN to the batch-corrected UMAP of the experimental data with over 1 million cells (Section 5.6). Specifically, for each D-SPIN sample, we mapped it to the nearest neighbor of the discretized states in the experimental data. In the case of multiple nearest neighbors, the probability was evenly distributed between these cells. In this way, the 10^6 samples from D-SPIN for each condition defined a distribution over the entire observed dataset, and could be visualized on the UMAP. We sampled 10^5 cells from the distribution to render the UMAP for D-SPIN samples.
2. Similarity of distribution on cell subtypes. The projected distribution over the whole dataset also defined a distribution over the cell subtypes as the clusters in the UMAP embedding. We computed the cosine similarity between cell subtype distributions of experimental data and the D-SPIN model for each condition. For reference comparison, we also defined two null distributions of cell state: the first null distribution is the uniform distribution over all possible states in the discrete space of $[-1, 0, 1]^M$, the second null distribution is the uniform distribution over the drug profiling dataset by pooling all conditions together to reflect the relative cell type abundance.

3. Gene program expression mean and cross-correlation. The first- and second-order moments of distributions are their fundamental characteristics. We computed the Euclidean (Frobenius) norm of the difference between the mean (cross-correlation) of the experimental data and D-SPIN model and presented in Figures 5.9 A(i)(ii)
4. Optimal transport (OT) distance between experimental data and D-SPIN model. OT offers a robust method to quantify the distance between two probability distributions, taking into account the underlying geometry of the data space. For computational tractability, we computed the entropy-regularized Sinkhorn OT distance with Python package POT (pot.sinkhorn(reg=0.5)) between experimental data and D-SPIN model distributions and presented in Figure 5.9 A(iii). Specifically, as the samples of the D-SPIN model were obtained through MCMC, they were highly noisy at the single-state level. States that appeared only once or twice were likely by chance and led to skewed representations of the true distributions. This noise could be mitigated with neighborhood projection or distribution average in previous metrics, but had large impacts on state-level distribution comparisons such as OT. Therefore we filtered out states in the D-SPIN model that appeared fewer than 5 times, which were roughly half of the 10^6 samples.
5. Direct state probability comparison. We further directly compared the probability of each single state, presented as scatter plots as in Figure 5.9 B. The probability of the D-SPIN model also filtered out states that appeared fewer than 5 times, as in optimal transport distance computation.

Discovering regulatory networks modules and clustering perturbation responses

In the drug profiling dataset, due to some batches only having a single control sample, a unified statistical test for the significance of perturbation response vectors is not viable. We scored each drug by the negative log probability density function at \tilde{h} of a multivariate Gaussian distribution with zero mean and diagonal covariance. The covariance was determined by the variance of control samples of the same experimental batch if there were more than one control sample, otherwise by the variance of all control samples in the dataset. We selected 158 candidate drugs with a score above 20, and performed K-means clustering with 20 clusters. We found that some drug clusters were putative experimental batch effects, as these contained biochemically different drugs from the same experimental batch. We also removed

the drug cluster where a complete loss of macrophage population was observed, as the loss of macrophage population was not reproducible in other experiments, and likely resulted from cell adhesion to the culture plate. The remaining drug clusters were then manually reviewed and consolidated based on their biological impacts, resulting in 70 drugs in seven major categories.

Identification of drug target and potency from database

The drug target information was obtained by querying the probes & drugs database (Skuta et al., 2017). Repetitive entries in the database were consolidated with manual annotations. The primary target of drugs is based on Drug Bank (Wishart et al., 2006) as well as manual lookups.

Computing strength of immune inhibitors

The strength of immune inhibitors in the D-SPIN model is quantified by projecting each inhibitor's relative response vector \tilde{h} onto a principal direction, denoted by e with the expression Strength = $\tilde{h} \cdot e$. The principal direction, e , is derived from the leading left singular vector of the matrix composed of stacked relative response vectors $[\tilde{h}_1, \tilde{h}_2, \dots]$ from all immune inhibitors, including strong inhibitors, weak inhibitors I, and weak inhibitors II. The leading singular vector, determined through singular value decomposition (SVD), captures the primary axis of variance of the inhibitor response vectors. Essentially, it represents the consensus direction of gene program combinations that all immune inhibitors act on.

Categories of drug interactions on gene program level

We define five major types of drug interactions on the gene-program level: additive, subadditive, dominant, synergistic, and antagonistic. Denoting the response of two single drugs and drug combinations on a specific gene program as h_1, h_2, h_c , we classify different drug interactions based on the sign and magnitude of the single drug and combination responses. We define a tolerance $\delta = 0.5$, and the sign of a response h is defined as the sign after soft thresholding

$$\text{sgn}(h) = \begin{cases} -1 & \text{if } h < -\delta, \\ 0 & \text{if } -\delta \leq h \leq \delta, \\ 1 & \text{if } h > \delta. \end{cases} \quad (5.1)$$

Additive interactions are defined as

$$|h_c - h_1 - h_2| < \delta \quad (5.2)$$

Otherwise, if the two drug responses have the same sign $\text{sgn}(h_1)\text{sgn}(h_2) > 0$, we assume $h_1 > 0, h_2 > 0$ for notation simplicity and the situation of $h_1 < 0, h_2 < 0$ naturally follows. The interactions are defined by

$$\text{interaction} = \begin{cases} \text{synergistic} & \text{if } h_c \geq \max(h_1, h_2) + \delta \\ \text{antagonistic} & \text{if } h_c < \min(h_1, h_2) - \delta \\ \text{subadditive} & \min(h_1, h_2) - \delta \leq h_c < \max(h_1, h_2) + \delta \end{cases} \quad (5.3)$$

If the two drug responses have opposite signs $\text{sgn}(h_1)\text{sgn}(h_2) < 0$, we assume $h_1 > 0, h_2 < 0, h_c \geq 0$ for notation simplicity. The interactions are defined by

$$\text{interaction} = \begin{cases} \text{synergistic} & \text{if } h_c \geq h_1 + \delta \\ \text{dominant} & \text{if } h_1 - \delta \leq h_c < h_1 + \delta \\ \text{subadditive} & 0 \leq h_c < h_1 - \delta \end{cases} \quad (5.4)$$

If the two drug responses both have sign 0, the interactions are defined by

$$\text{interaction} = \begin{cases} \text{subadditive} & \text{if } \text{sgn}(h_c) = 0 \\ \text{synergistic} & \text{if } \text{sgn}(h_c) \neq 0 \end{cases} \quad (5.5)$$

Lastly, if only one of the drug responses has sign 0, we assume $h_1 > 0, \text{sgn}(h_2) = 0$ for notation simplicity. The interactions are defined by

$$\text{interaction} = \begin{cases} \text{dominance} & \text{if } \text{sgn}(h_c) = 0 \\ \text{synergistic} & \text{if } h_c \geq h_1 + \delta \\ \text{antagonistic} & \text{if } h_c < h_1 - \delta \\ \text{subadditive} & \text{if } h_1 - \delta \leq h_c < h_1 + \delta \end{cases} \quad (5.6)$$

Coherence between drug responses

To evaluate the similarity in responses induced by the two drugs, we analyzed both global coherence and program-level coherence. The global coherence was defined by the cosine similarity between the relative response vector of the two drugs, that coherence = $\tilde{h}_1^\top \cdot \tilde{h}_2 / \|\tilde{h}_1\| \|\tilde{h}_2\|$. The program-level coherence was determined by comparing the sign of perturbation effect on each gene program. If the relative response vectors of the two single drugs share the same sign on a gene program, the two drugs are coherent on the program. Conversely, opposite signs indicate they are incoherent on the program.

Phase diagrams of dosages of drug combination

The regulatory network of the D-SPIN model is able to convert response vectors into distributions of cell states, therefore enabling the interpolation of cell state transitions between experimental conditions. In the dosage combination experiments with Dasatinib and Halcinonide, we observed the transition between activated macrophage states, monocyte-like resting states, and combinatorial states with both M2 and monocyte transcriptional programs. We used D-SPIN to depict the transitions between the different cell states induced by the change in drug combination dosages.

To precisely tailor the regulatory networks controlling cell state changes in drug dosage combinations, we exclusively used samples from two drug dosage combination experiments to fit a regulatory network \mathbf{J} and response vectors \mathbf{h} . These experiments, including Dasatinib and Halcinonide in one and Tacrolimus and Motolimod in the other, spanned a range of dosages and provided a rich dataset to explore the effect of drug dosage combinations. The first experiment included varying concentrations of Dasatinib and Halcinonide in the following dosages: Dasatinib at 0.1, 1, 10, 100, and 1000 nM, and Halcinonide at 0.01, 0.1, 1, 10, 100, and 1000 nM. The second experiment included a similar dosage design, using Tacrolimus (dosages of 0.1, 1, 10, 100, and 1000 nM), and Motolimod (dosages of 1, 10, 100, 1000, and 10000 nM). In both experiments, we profiled all combinations of these two drug dosages.

To elucidate the modulation of cell states with drug combinations, we constructed a phase diagram of cell states by estimating the average log-likelihood of major cell subtypes under every drug dosage combination. For any given drug combination concentration (c_1, c_2) , we interpolated the response vector $\hat{\mathbf{h}}(c_1, c_2)$ (MATLAB fit ‘thinplateinterp’) and evaluated the score u of the target cell subtype as the average log-likelihood of individual states $\{\mathbf{s}\}_{\text{Target}}$ as

$$u_{\text{Target}} = \frac{1}{\#\{\mathbf{s}\}_{\text{Target}}} \sum_{\{\mathbf{s}\}_{\text{Target}}} \log P(\mathbf{s}; \mathbf{J}, \hat{\mathbf{h}}(c_1, c_2)) \quad (5.7)$$

Then, we used softmax across all target cell subtypes to compute the estimated likelihood of each state

$$\hat{p}_{\text{Target i}} = \frac{\exp u_{\text{Target i}}}{\sum_i \exp u_{\text{Target i}}} \quad (5.8)$$

and rendered the phase diagram by the target state with the largest probability as presented in Figures 5.16 CD.

In the analysis of Dasatinib and Halcinonide combinations, we distinguished five myeloid states of interest, including (a) Activated macrophage: activated control, (b) Monocyte-like: resting control, (c) M2 macrophage: 1000 nM Halcinonide, (d) Inhibited monocyte: 1000 nM Dasatinib, and (e) M2-like monocyte: 1000 nM Dasatinib plus 1000 nM Halcinonide. Similarly, in the analysis of Tacrolimus and Motolimod combinations, we investigated three target myeloid states: (a) Activated macrophage: activated control, (b) Monocyte-like: 1000 nM Tacrolimus, and (c) Pathogen-response macrophage: 10000 nM Motolimod.

CONCLUSIONS AND DISCUSSIONS

Here, we introduce D-SPIN, a probabilistic modeling framework, that constructs quantitative, generative models for gene regulatory networks from large-scale single-cell mRNA-seq datasets collected from cell populations across a series of perturbation conditions. The D-SPIN framework is general and can be applied to study the response of cells to genetic perturbations, small molecules, and signaling conditions. D-SPIN models can be treated both as network models that are analyzed using the tools of graph theory as well as complete probabilistic models that can predict the structure of a cell population across a series of different conditions. The mathematical structure of D-SPIN allowed us to develop a computationally efficient, parallel inference procedure that can be run on hundreds of CPU cores to perform network inference on datasets with thousands of perturbations and millions of cells. Our work establishes an interpretable model framework for transcriptome-scale perturbation datasets. The D-SPIN framework is a maximum-entropy modeling framework adapted from the analysis of systems in condensed-matter physics (Jaynes, 1957). D-SPIN provides a low-dimensional, compact yet predictive framework for analyzing population-level effects in single-cell data through a regulatory network model. Biologically, we apply D-SPIN to generate global insights into network-level mechanisms of cellular homeostasis, to define hypotheses regarding the sub-functions of protein complexes, to demonstrate the additive logic of combinatorial drug responses at the transcriptome scale, and thus to develop a framework for applying the additivity of immunomodulatory drugs to program immune cell-states and immune population structure.

Single-cell mRNA-seq methods enable large-scale perturbation response studies across cell types and organisms. D-SPIN provides a computational framework for integrating information from such studies into regulatory network models that can be analyzed and compared to reveal principles of cellular regulation. We have demonstrated that D-SPIN provides insight into damage response strategies in K562 cells. Stress response strategies play a role in many diseases including cancer, and diseases of ageing and degeneration. Further, cell response strategies change over the lifetime of organisms and also over evolutionary time scales. Our understanding of gene regulatory network architecture, logic, and evolution has

been limited by the limited number of network reconstructions that exist. D-SPIN provides a framework to construct, analyze, and compare regulatory networks to understand how they change over time.

Cells are distributed control systems that modulate many internal processes to maintain homeostasis and execute cell-fate transitions. However, the principles of distributed control at the transcriptome scale are poorly understood. D-SPIN models help reveal the principles of distributed control. We show that K562 cells mount a set of regulatory responses when different internal processes are perturbed through gene knockdown. The response strategies point to a broadly distributed control where long-range information flow between distinct cellular processes, such as coherent up or downregulation of proteasome/ribosome subunits under ribosome subunit knockdown, points to the presence of sensing and controlling mechanisms in the cell. Identifying these feedback control points with combinatorial knockdown screening may lead to the discovery of new regulators or regulating mechanisms.

D-SPIN can also assist the design of interventions that modulate networks and cell states. We showed that D-SPIN can be applied to a model of the HSC network to construct strategies for controlling the differentiation of a stem cell population. D-SPIN also reveals how small molecules and small-molecule combinations can generate novel cell states and population structures, suggesting that D-SPIN could be used to study a broader set of interventions in both stem cells and immune cell populations. In both cases, a major goal is to control the distribution of cell states in a cell population, such as generating a specific set of progeny from stem cells or modulating the state of the immune system in cancer or autoimmunity. D-SPIN can uncover non-intuitive connections between different cellular processes, revealing how drugs might recruit a set of gene programs to generate a particular cell state. While this principle has been studied on a small scale for a small number of drugs (Geva-Zatorsky et al., 2010), our results suggest that such principles of super-position might hold at the transcriptome scale.

In our analysis of drug responses, we demonstrate that drug combinations can act through coherent and incoherent superpositions to generate novel cell states. In both cases, drugs act in combination to recruit a combination of gene-expression programs within a single-cell type (Dasatinib and Halcinonide) or across a cell population (Motolimod and Tacrolimus). Superposition refers to the ability of drug combinations to function additively, resulting in the union of the gene-expression programs, each activated by the individual drugs. D-SPIN enables us to dissect

these combinatorial mechanisms and interpret them through the inferred D-SPIN regulatory network and its interaction with drug response vectors. The results provide a conceptual framework for interpreting and predicting the effect of drug combinations. Although we have shown that additivity alone is insufficient to accurately predict combinatorial response due to the presence of non-additive interactions, it could still serve as a guide to narrow down potential drug-combination targets for therapeutic objectives. Additivity could arise from the modularity of gene-regulatory circuits, such that different pathways impact gene expression levels independently. Further work is needed to reveal the specific conditions where additivity holds or breaks down (Geva-Zatorsky et al., 2010).

The program-level or module-level description has been widely used in literature to provide interpretable insights into cellular regulation, but mainly from the view of the differential expression of single modules (Segal, Shapira, et al., 2003; Segal, Friedman, et al., 2004; Amit et al., 2009). D-SPIN extends the program-level description by inferring interactions between programs. By representing the cell as a set of interacting programs, D-SPIN allows us to gain insight into the information flow between core cellular processes like transcription, translation, and protein degradation in response to single-gene knockdowns, as well as how therapeutic interventions combine additively to recruit modules of immune function (monocyte and M2 macrophage programs) to create novel cell states. How program-level regulation is implemented through biomolecular interactions is a future topic of research, which D-SPIN can facilitate as we show that D-SPIN can identify regulators of inhibition between mutually repressing modules by incorporating perturbations.

The predictive power of D-SPIN and the connection between inferred interactions and biochemical interactions within the cell could be further enhanced by integrating data from other single-cell measurement modalities like ATAC-seq, known regulatory interactions, as well as predicted biochemical interactions. For example, single-cell ATAC-seq measures chromatin accessibility along the genome in single cells, providing information about the accessibility of transcription factor (TF) binding sites. Single-cell ATAC-seq data and transcription factor binding predictions can be combined with transcriptional profiling to identify TF-target interactions such as in methods like SCENIC+ and CellOracle (Bravo González-Blas et al., 2023; Kamimoto et al., 2023). CellOracle in particular predicts the response of transcription factor regulatory networks to TF perturbations by integrating TF motif binding predictions with ATAC-seq data and linear regression. As D-SPIN

is a statistical model, measured or predicted TF-target interactions can be included as a prior distribution during network inference, enhancing the model accuracy while still remaining flexible for biological contexts not dominantly controlled by TF-target interactions, such as signal transduction in T-cells. Further, techniques from deep neural networks and statistical learning have been used to learn representations and predict the result of perturbation experiments, and future versions of D-SPIN could apply neural networks to construct energy functions that include higher-order interactions between more than two components and more realistic concentration-dependent regulatory functions (Lotfollahi, F. A. Wolf, and Theis, 2019; Dong et al., 2023; Lotfollahi, Klimovskaia Susmelj, et al., 2023; Roohani, K. Huang, and Leskovec, 2023). Incorporating more comprehensive terms into the energy function in Equation 2.2 would enhance model fitting and predictive power, while the trade-off with model interpretability would be an important topic of further research.

D-SPIN’s connection with spin network models provides insights into the fundamental nature of cellular regulation. The spin network model is an equilibrium model used to study physical systems at or near thermal equilibrium. Theoretically, D-SPIN depicts the cell population as a collection of points residing in an energy landscape that can be tilted by the perturbation vector to shift the distribution of cell states in a cell population. Energy-landscape models represent a highly simplified class of dynamical systems, as their behavior can be captured within a single energy (potential) function. Equilibrium spin network models have been used to study a much broader range of systems that are far from thermal equilibrium, including neural networks and bird flocks (Hopfield, 1982; Schneidman et al., 2006; Bialek et al., 2012). However, it remains unclear why equilibrium models can yield such significant predictive power for strongly non-equilibrium systems. The ability of equilibrium models to produce low-error reconstructions of cell population gene-expression states suggests that cells, in certain situations, may be effectively modeled as equilibrium systems driven through various configurations by a biasing drive. Such models have been demonstrated for cell-fate regulation and might represent a simplifying principle (Sokolik et al., 2015). The ability to model a cell as an equilibrium system driven through different states presents a powerful simplification, as also explored in other chemical systems, offering a potential path toward more global theories of gene regulation (Bintu, Buchler, Garcia, Gerland, Hwa, Kondev, Kuhlman, et al., 2005; Gunawardena, 2012).

The D-SPIN framework has important current limitations that represent targets for extension in future work, although these possibilities fall beyond the scope of the present study. First, for simplicity and interpretability, D-SPIN only considers pairwise interactions between genes or gene programs in the network. These interactions correspond to the second-order terms in the energy function. For enhanced accuracy and predictive capabilities, the inclusion of higher-order multi-body interactions is a viable future modification. Second, D-SPIN is an equilibrium model and does not account for the dynamics of the system. Spin network models have natural extensions that incorporate dynamics, and studies of spin glass dynamics have been important in condensed matter physics (Glauber, 1963; Sompolinsky and Zippelius, 1982; Fisher and Huse, 1988). Incorporating temporal dynamics would also allow the inclusion of directed edges, as directed network models in general cannot define consistent stationary distributions when the network contains feedback loops (Pearl, 1987; Arnold and Press, 1989; Pearl, 2022). Third, D-SPIN assumes that the interactions within the core regulatory network J are not altered by the perturbations, which is a reasonable approximation in perturbation response scenarios such as gene knockdown or small molecule action. However, in scenarios such as cellular differentiation or disease progression, the regulatory network may undergo changes under different conditions due to shifts in epigenetic regulation. Epigenetic reorganization can be included in future versions of D-SPIN by allowing interactions encoded in J to also be condition-dependent.

BIBLIOGRAPHY

Adamson, Britt et al. (Dec. 2016). “A Multiplexed Single-Cell CRISPR Screening Platform Enables Systematic Dissection of the Unfolded Protein Response”. In: *Cell* 167.7, 1867–1882.e21.

Aibar, Sara et al. (2017). “SCENIC: single-cell regulatory network inference and clustering”. In: *Nature methods* 14.11, pp. 1083–1086.

Almeida, Melanie de et al. (2021). “AKIRIN2 controls the nuclear import of proteasomes in vertebrates”. In: *Nature* 599.7885, pp. 491–496.

Amit, Ido et al. (2009). “Unbiased reconstruction of a mammalian transcriptional network mediating pathogen responses”. In: *Science* 326.5950, pp. 257–263.

Amoroso, Maria Rosaria et al. (2012). “TRAP1 and the proteasome regulatory particle TBP7/Rpt3 interact in the endoplasmic reticulum and control cellular ubiquitination of specific mitochondrial proteins”. In: *Cell Death & Differentiation* 19.4, pp. 592–604.

Arnold, Barry C and S James Press (1989). “Compatible conditional distributions”. In: *Journal of the American Statistical Association* 84.405, pp. 152–156.

Aurell, Erik and Magnus Ekeberg (2012). “Inverse Ising inference using all the data”. In: *Physical review letters* 108.9, p. 090201.

Beltrao, Pedro, Gerard Cagney, and Nevan J Krogan (2010). “Quantitative genetic interactions reveal biological modularity”. In: *Cell* 141.5, pp. 739–745.

Besag, Julian (1974). “Spatial interaction and the statistical analysis of lattice systems”. In: *Journal of the Royal Statistical Society: Series B (Methodological)* 36.2, pp. 192–225.

Bialek, William et al. (2012). “Statistical mechanics for natural flocks of birds”. In: *Proceedings of the National Academy of Sciences* 109.13, pp. 4786–4791.

Bintu, Lacramioara, Nicolas E Buchler, Hernan G Garcia, Ulrich Gerland, Terence Hwa, Jané Kondev, Thomas Kuhlman, et al. (2005). “Transcriptional regulation by the numbers: applications”. In: *Current opinion in genetics & development* 15.2, pp. 125–135.

Bintu, Lacramioara, Nicolas E Buchler, Hernan G Garcia, Ulrich Gerland, Terence Hwa, Jané Kondev, and Rob Phillips (2005). “Transcriptional regulation by the numbers: models”. In: *Current opinion in genetics & development* 15.2, pp. 116–124.

Bravo González-Blas, Carmen et al. (2023). “SCENIC+: single-cell multiomic inference of enhancers and gene regulatory networks”. In: *Nature methods* 20.9, pp. 1355–1367.

Bray, Dennis (1995). "Protein molecules as computational elements in living cells". In: *Nature* 376.6538, pp. 307–312.

Burkhardt, Daniel B et al. (2021). "Quantifying the effect of experimental perturbations at single-cell resolution". In: *Nature biotechnology*.

Castellana, Michele and William Bialek (2014). "Inverse spin glass and related maximum entropy problems". In: *Physical review letters* 113.11, p. 117204.

Chan, Thalia E, Michael PH Stumpf, and Ann C Babtie (2017). "Gene regulatory network inference from single-cell data using multivariate information measures". In: *Cell systems* 5.3, pp. 251–267.

Chari, Tara and Lior Pachter (2023). "The specious art of single-cell genomics". In: *PLoS Computational Biology* 19.8, e1011288.

Chen, Sisi et al. (2020). "Dissecting heterogeneous cell populations across drug and disease conditions with PopAlign". In: *Proceedings of the National Academy of Sciences* 117.46, pp. 28784–28794.

Chen, Xiaoqiao, Sisi Chen, and Matt Thomson (2022). "Minimal gene set discovery in single-cell mRNA-seq datasets with ActiveSVM". In: *Nature computational science* 2.6, pp. 387–398.

Choi, Seungjin (2008). "Algorithms for orthogonal nonnegative matrix factorization". In: *2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence)*. IEEE, pp. 1828–1832.

Consortium*, Tabula Sapiens et al. (2022). "The Tabula Sapiens: A multiple-organ, single-cell transcriptomic atlas of humans". In: *Science* 376.6594, eabl4896.

Davidson, Eric H (2010a). "Emerging properties of animal gene regulatory networks". In: *Nature* 468.7326, pp. 911–920.

Davidson, Eric H (2010b). *The regulatory genome: gene regulatory networks in development and evolution*. Elsevier.

Davidson, Eric H and Douglas H Erwin (2006). "Gene regulatory networks and the evolution of animal body plans". In: *Science* 311.5762, pp. 796–800.

De Visser, Karin E, Alexandra Eichten, and Lisa M Coussens (2006). "Paradoxical roles of the immune system during cancer development". In: *Nature reviews cancer* 6.1, pp. 24–37.

Desgeorges, Thibaut et al. (2019). "Glucocorticoids shape macrophage phenotype for tissue repair". In: *Frontiers in immunology* 10, p. 1591.

Dixit, Atray et al. (2016). "Perturb-Seq: dissecting molecular circuits with scalable single-cell RNA profiling of pooled genetic screens". In: *Cell* 167.7, pp. 1853–1866.

Domínguez Conde, C et al. (2022). "Cross-tissue immune cell analysis reveals tissue-specific features in humans". In: *Science* 376.6594, eabl5197.

Dong, Mingze et al. (2023). “Causal identification of single-cell experimental perturbation effects with CINEMA-OT”. In: *Nature methods* 20.11, pp. 1769–1779.

Fajgenbaum, David C and Carl H June (2020). “Cytokine storm”. In: *New England Journal of Medicine* 383.23, pp. 2255–2273.

Ferguson, Edwin L, Paul W Sternberg, and H Robert Horvitz (1987). “A genetic pathway for the specification of the vulval cell lineages of *Caenorhabditis elegans*”. In: *Nature* 326.6110, pp. 259–267.

Fisher, Daniel S and David A Huse (1988). “Nonequilibrium dynamics of spin glasses”. In: *Physical review B* 38.1, p. 373.

Fishman, JA (2017). “Infection in organ transplantation”. In: *American Journal of Transplantation* 17.4, pp. 856–879.

Fishman, Jay A and Robert H Rubin (1998). “Infection in organ-transplant recipients”. In: *New England Journal of Medicine* 338.24, pp. 1741–1751.

Friedman, Nir et al. (2000). “Using Bayesian networks to analyze expression data”. In: *Proceedings of the fourth annual international conference on Computational molecular biology*, pp. 127–135.

Garcia-Alonso, Luz et al. (2019). “Benchmark and integration of resources for the estimation of human transcription factor activities”. In: *Genome research* 29.8, pp. 1363–1375.

Gehring, Jase et al. (2020). “Highly multiplexed single-cell RNA-seq by DNA oligonucleotide tagging of cellular proteins”. In: *Nature Biotechnology* 38.1, pp. 35–38.

Geva-Zatorsky, Naama et al. (2010). “Protein dynamics in drug combinations: a linear superposition of individual-drug responses”. In: *Cell* 140.5, pp. 643–651.

Giurgiu, Madalina et al. (2019). “CORUM: the comprehensive resource of mammalian protein complexes—2019”. In: *Nucleic acids research* 47.D1, pp. D559–D563.

Glauber, Roy J (1963). “Time-dependent statistics of the Ising model”. In: *Journal of mathematical physics* 4.2, pp. 294–307.

Gordon, David E et al. (2020). “A SARS-CoV-2 protein interaction map reveals targets for drug repurposing”. In: *Nature* 583.7816, pp. 459–468.

Gunawardena, Jeremy (2012). “A linear framework for time-scale separation in nonlinear biochemical systems”. In: *PLoS one* 7.5, e36321.

Gutenkunst, Ryan N et al. (2007). “Universally sloppy parameter sensitivities in systems biology models”. In: *PLoS computational biology* 3.10, e189.

Han, Heonjong et al. (2018). “TRRUST v2: an expanded reference database of human and mouse transcriptional regulatory interactions”. In: *Nucleic acids research* 46.D1, pp. D380–D386.

Harbison, Christopher T et al. (2004). “Transcriptional regulatory code of a eukaryotic genome”. In: *Nature* 431.7004, pp. 99–104.

Hartwell, Leland H et al. (1999). “From molecular to modular cell biology”. In: *Nature* 402.Suppl 6761, pp. C47–C52.

Hastie, Trevor et al. (2009). *The elements of statistical learning: data mining, inference, and prediction*. Vol. 2. Springer.

Heimberg, Graham et al. (2016). “Low dimensionality in gene expression data enables the accurate extraction of transcriptional programs from shallow sequencing”. In: *Cell systems* 2.4, pp. 239–250.

Heisterkamp, Nora et al. (1983). “Localization of the c-abl oncogene adjacent to a translocation break point in chronic myelocytic leukaemia”. In: *Nature* 306.5940, pp. 239–242.

Hiatt, Joseph et al. (2022). “A functional map of HIV-host interactions in primary human T cells”. In: *Nature communications* 13.1, p. 1752.

Hnisz, Denes et al. (2013). “Super-enhancers in the control of cell identity and disease”. In: *Cell* 155.4, pp. 934–947.

Holthaus, Lisa et al. (2018). “CD4+ T cell activation, function, and metabolism are inhibited by low concentrations of DMSO”. In: *Journal of immunological methods* 463, pp. 54–60.

Hopfield, John J (1982). “Neural networks and physical systems with emergent collective computational abilities”. In: *Proceedings of the National Academy of Sciences* 79.8, pp. 2554–2558.

Hu, Zhanzhi, Patrick J Killion, and Vishwanath R Iyer (2007). “Genetic reconstruction of a functional transcriptional regulatory network”. In: *Nature genetics* 39.5, pp. 683–687.

Huang, Linda S and Paul W Sternberg (1995). “Genetic dissection of developmental pathways”. In: *Methods in cell biology* 48, pp. 97–122.

Huynh-Thu, Vân Anh et al. (2010). “Inferring regulatory networks from expression data using tree-based methods”. In: *PLoS one* 5.9, e12776.

Ireland, William T et al. (2020). “Deciphering the regulatory genome of Escherichia coli, one hundred promoters at a time”. In: *Elife* 9, e55308.

Jaeger, Stefanie et al. (2012). “Global landscape of HIV–human protein complexes”. In: *Nature* 481.7381, pp. 365–370.

Jamilloux, Yvan et al. (2019). “JAK inhibitors for the treatment of autoimmune and inflammatory diseases”. In: *Autoimmunity reviews* 18.11, p. 102390.

Jaynes, Edwin T (1957). “Information theory and statistical mechanics”. In: *Physical review* 106.4, p. 620.

Jiang, Jialong, David A Sivak, and Matt Thomson (2019). “Active Learning of Spin Network Models”. In: *arXiv preprint arXiv:1903.10474*.

Joung, Julia et al. (2023). “A transcription factor atlas of directed differentiation”. In: *Cell* 186.1, pp. 209–229.

Kafri, Ran, Michael Springer, and Yitzhak Pilpel (2009). “Genetic redundancy: new tricks for old genes”. In: *Cell* 136.3, pp. 389–392.

Kamimoto, Kenji et al. (2023). “Dissecting cell identity via network inference and in silico gene perturbation”. In: *Nature* 614.7949, pp. 742–751.

Kim, Minkyu et al. (2021). “A protein interaction landscape of breast cancer”. In: *Science* 374.6563, eabf3066.

Kingma, Diederik P and Jimmy Ba (2014). “Adam: A method for stochastic optimization”. In: *arXiv preprint arXiv:1412.6980*.

Krumsiek, Jan et al. (2011). “Hierarchical differentiation of myeloid progenitors is encoded in the transcription factor network”. In: *PLoS one* 6.8, e22649.

Kueh, Hao Yuan et al. (2013). “Positive feedback between PU. 1 and the cell cycle controls myeloid differentiation”. In: *Science* 341.6146, pp. 670–673.

Kuleshov, Maxim V et al. (2016). “Enrichr: a comprehensive gene set enrichment analysis web server 2016 update”. In: *Nucleic acids research* 44.W1, W90–W97.

Labzin, Larisa I, Michael T Heneka, and Eicke Latz (2018). “Innate immunity and neurodegeneration”. In: *Annual review of medicine* 69, pp. 437–449.

Lang, Alex H et al. (2014). “Epigenetic landscapes explain partially reprogrammed cells and identify key reprogramming genes”. In: *PLoS computational biology* 10.8, e1003734.

Lee, Hyo-Ji et al. (2020). “Formation and maturation of the phagosome: a key mechanism in innate immunity against intracellular bacterial infection”. In: *Microorganisms* 8.9, p. 1298.

Lee, Tong Ihn et al. (2002). “Transcriptional regulatory networks in *Saccharomyces cerevisiae*”. In: *science* 298.5594, pp. 799–804.

Liu, Yan-Cun et al. (2014). “Macrophage polarization in inflammatory diseases”. In: *International journal of biological sciences* 10.5, p. 520.

Liu, Zhi-Ping et al. (2015). “RegNetwork: an integrated database of transcriptional and post-transcriptional regulatory networks in human and mouse”. In: *Database* 2015, bav095.

Lopez, Romain et al. (2018). “Deep generative modeling for single-cell transcriptomics”. In: *Nature methods* 15.12, pp. 1053–1058.

Lotfollahi, Mohammad, Anna Klimovskaia Susmelj, et al. (2023). “Predicting cellular responses to complex perturbations in high-throughput screens”. In: *Molecular Systems Biology*, e11517.

Lotfollahi, Mohammad, F Alexander Wolf, and Fabian J Theis (2019). “scGen predicts single-cell perturbation responses”. In: *Nature methods* 16.8, pp. 715–721.

Lozzio, Carmen B and Bismarck B Lozzio (1975). “Human chronic myelogenous leukemia cell-line with positive Philadelphia chromosome”. In: *Blood* 45, pp. 321–334.

Machta, Benjamin B et al. (2013). “Parameter space compression underlies emergent theories and predictive models”. In: *Science* 342.6158, pp. 604–607.

Maerkl, Sebastian J and Stephen R Quake (2007). “A systems approach to measuring the binding energy landscapes of transcription factors”. In: *Science* 315.5809, pp. 233–237.

Margolin, Adam A et al. (2006). “ARACNE: an algorithm for the reconstruction of gene regulatory networks in a mammalian cellular context”. In: *BMC bioinformatics*. Vol. 7. Springer, pp. 1–15.

Marrack, Philippa, John Kappler, and Brian L Kotzin (2001). “Autoimmune disease: why and where it occurs”. In: *Nature medicine* 7.8, pp. 899–905.

Matsumoto, Hirotaka et al. (2017). “SCODE: an efficient regulatory network inference algorithm from single-cell RNA-Seq during differentiation”. In: *Bioinformatics* 33.15, pp. 2314–2321.

McFarland, James M et al. (Aug. 2020). “Multiplexed single-cell transcriptional response profiling to define cancer vulnerabilities and therapeutic mechanism of action”. In: *Nature communication* 11.1, p. 4296.

McGinnis, Christopher S et al. (2019). “MULTI-seq: sample multiplexing for single-cell RNA sequencing using lipid-tagged indices”. In: *Nature methods* 16.7, pp. 619–626.

Mikolov, Tomas et al. (2013). “Distributed representations of words and phrases and their compositionality”. In: *Advances in neural information processing systems* 26.

Milanese, Chiara et al. (2019). “DNA damage and transcription stress cause ATP-mediated redesign of metabolism and potentiation of anti-oxidant buffering”. In: *Nature communications* 10.1, p. 4887.

Milo, Ron et al. (2002). “Network motifs: simple building blocks of complex networks”. In: *Science* 298.5594, pp. 824–827.

Moerman, Thomas et al. (2019). “GRNBoost2 and Arboreto: efficient and scalable inference of gene regulatory networks”. In: *Bioinformatics* 35.12, pp. 2159–2161.

Monod, Jacques (1949). “The growth of bacterial cultures”. In: *Annual review of microbiology* 3.1, pp. 371–394.

Newman, Mark (2018). *Networks*. Oxford university press.

Nguyen, H Chau, Riccardo Zecchina, and Johannes Berg (2017). “Inverse statistical problems: from the inverse Ising problem to data science”. In: *Advances in Physics* 66.3, pp. 197–261.

Nichols, Robert J et al. (2011). “Phenotypic landscape of a bacterial cell”. In: *Cell* 144.1, pp. 143–156.

Nowak, Martin A et al. (1997). “Evolution of genetic redundancy”. In: *Nature* 388.6638, pp. 167–171.

Pardee, Arthur B, François Jacob, and Jacques Monod (1959). “The genetic control and cytoplasmic expression of “inducibility” in the synthesis of β -galactosidase by *E. coli*”. In: *Journal of Molecular Biology* 1.2, pp. 165–178.

Pe'er, Dana et al. (2001). “Inferring subnetworks from perturbed expression profiles”. In: *BIOINFORMATICS-OXFORD-* 17, S215–S224.

Pearl, Judea (1987). “Evidential reasoning using stochastic simulation of causal models”. In: *Artificial intelligence* 32.2, pp. 245–257.

Pearl, Judea (2022). “Reverend Bayes on inference engines: A distributed hierarchical approach”. In: *Probabilistic and Causal Inference: The Works of Judea Pearl*, pp. 129–138.

Pereira, Marcus R et al. (2020). “COVID-19 in solid organ transplant recipients: initial report from the US epicenter”. In: *American journal of transplantation* 20.7, pp. 1800–1808.

Pratapa, Aditya et al. (2020). “Benchmarking algorithms for gene regulatory network inference from single-cell transcriptomic data”. In: *Nature methods* 17.2, pp. 147–154.

Ravikumar, Pradeep, Martin J Wainwright, and John D Lafferty (2010). “High-dimensional Ising model selection using ℓ_1 -regularized logistic regression”. In: *The Annals of Statistics*, pp. 1287–1319.

Regev, Aviv and Ehud Shapiro (2002). “Cellular abstractions: Cells as computation”. In: *Nature* 419.6905, pp. 343–343.

Regev, Aviv and Ehud Shapiro (2004). “The π -calculus as an abstraction for biomolecular systems”. In: *Modelling in Molecular Biology*. Springer, pp. 219–266.

Replogle, Joseph M et al. (2022). “Mapping information-rich genotype-phenotype landscapes with genome-scale Perturb-seq”. In: *Cell*.

Roguev, Assen et al. (2008). “Conservation and rewiring of functional modules revealed by an epistasis map in fission yeast”. In: *Science* 322.5900, pp. 405–410.

Roohani, Yusuf, Kexin Huang, and Jure Leskovec (2023). “Predicting transcriptional outcomes of novel multigene perturbations with gears”. In: *Nature biotechnology*, pp. 1–9.

Santhanam, Narayana P and Martin J Wainwright (2012). “Information-theoretic limits of selecting binary graphical models in high dimensions”. In: *IEEE Transactions on Information Theory* 58.7, pp. 4117–4134.

Schneidman, Elad et al. (2006). “Weak pairwise correlations imply strongly correlated network states in a neural population”. In: *Nature* 440.7087, pp. 1007–1012.

Schraivogel, Daniel et al. (2020). “Targeted Perturb-seq enables genome-scale genetic screens in single cells”. In: *Nature methods* 17.6, pp. 629–635.

Schwartz, Daniella M et al. (2017). “JAK inhibition as a therapeutic strategy for immune and inflammatory diseases”. In: *Nature reviews drug discovery* 16.12, pp. 843–862.

Segal, Eran, Nir Friedman, et al. (2004). “A module map showing conditional activity of expression modules in cancer”. In: *Nature genetics* 36.10, pp. 1090–1098.

Segal, Eran, Michael Shapira, et al. (2003). “Module networks: identifying regulatory modules and their condition-specific regulators from gene expression data”. In: *Nature genetics* 34.2, pp. 166–176.

Shen-Orr, Shai S et al. (2002). “Network motifs in the transcriptional regulation network of *Escherichia coli*”. In: *Nature genetics* 31.1, pp. 64–68.

Sherman, Brad T et al. (2022). “DAVID: a web server for functional enrichment analysis and functional annotation of gene lists (2021 update)”. In: *Nucleic acids research* 10.

Sivak, David A and Matt Thomson (2014). “Environmental statistics and optimal regulation”. In: *PLoS computational biology* 10.9, e1003826.

Skuta, Ctibor et al. (2017). “Probes & drugs portal: an interactive, open data resource for chemical biology”. In: *Nature methods* 14.8, pp. 759–760.

Sokolik, Cameron et al. (2015). “Transcription factor competition allows embryonic stem cells to distinguish authentic signals from noise”. In: *Cell systems* 1.2, pp. 117–129.

Sompolinsky, Haim and Annette Zippelius (1982). “Relaxational dynamics of the Edwards-Anderson model and the mean-field theory of spin-glasses”. In: *Physical review B* 25.11, p. 6860.

Srivatsan, Sanjay R et al. (Jan. 2020). “Massively multiplex chemical transcriptomics at single-cell resolution”. In: *Science* 367.6473, pp. 45–51.

Stelling, Jörg et al. (2004). “Robustness of cellular functions”. In: *Cell* 118.6, pp. 675–685.

Subramanian, Aravind et al. (Nov. 2017). “A Next Generation Connectivity Map: L1000 Platform and the First 1,000,000 Profiles”. In: *Cell* 171.6, 1437–1452.e17.

Szklarczyk, Damian et al. (2021). “The STRING database in 2021: customizable protein–protein networks, and functional characterization of user-uploaded gene/measurement sets”. In: *Nucleic acids research* 49.D1, pp. D605–D612.

Takahashi, Kazutoshi and Shinya Yamanaka (2006). “Induction of pluripotent stem cells from mouse embryonic and adult fibroblast cultures by defined factors”. In: *Cell* 126.4, pp. 663–676.

Teschendorff, Andrew E and Andrew P Feinberg (2021). “Statistical mechanics meets single-cell biology”. In: *Nature reviews genetics* 22.7, pp. 459–476.

Traag, Vincent A, Ludo Waltman, and Nees Jan Van Eck (2019). “From Louvain to Leiden: guaranteeing well-connected communities”. In: *Scientific reports* 9.1, p. 5233.

Van de Sande, Bram et al. (2023). “Applications of single-cell RNA sequencing in drug discovery and development”. In: *Nature reviews drug discovery*, pp. 1–25.

Van Dijk, David et al. (2018). “Recovering gene interactions from single-cell data using data diffusion”. In: *Cell* 174.3, pp. 716–729.

Vavasis, Stephen A (2010). “On the complexity of nonnegative matrix factorization”. In: *SIAM Journal on Optimization* 20.3, pp. 1364–1377.

Wagner, Günter P, Mihaela Pavlicev, and James M Cheverud (2007). “The road to modularity”. In: *Nature reviews genetics* 8.12, pp. 921–931.

Wang, Belinda et al. (2024). “A foundational atlas of autism protein interactions reveals molecular convergence”. In: *bioRxiv*. doi: 10.1101/2023.12.03.569805. eprint: [https://www.biorxiv.org/content/early/2024/02/08/2023.12.03.569805](https://www.biorxiv.org/content/early/2024/02/08/2023.12.03.569805.full.pdf).

Watts, Duncan J and Steven H Strogatz (1998). “Collective dynamics of ‘small-world’ networks”. In: *Nature* 393.6684, pp. 440–442.

Wishart, David S et al. (2006). “DrugBank: a comprehensive resource for in silico drug discovery and exploration”. In: *Nucleic acids research* 34.suppl_1, pp. D668–D672.

Yeh, Pamela, Ariane I Tschumi, and Roy Kishony (2006). “Functional classification of drugs by properties of their pairwise interactions”. In: *Nature genetics* 38.4, pp. 489–494.

Yosef, Nir et al. (2013). “Dynamic regulatory network controlling TH17 cell differentiation”. In: *Nature* 496.7446, pp. 461–468.

Yuh, Chiou-Hwa, Hamid Bolouri, and Eric H Davidson (1998). “Genomic cis-regulatory logic: experimental and computational analysis of a sea urchin gene”. In: *Science* 279.5358, pp. 1896–1902.

Zhou, Wen et al. (2019). "Single-cell analysis reveals regulatory gene expression dynamics leading to lineage commitment in early T cell development". In: *Cell systems* 9.4, pp. 321–337.