D-SPIN constructs gene regulatory network models from multiplexed scRNA-seq data revealing organizing principles of cellular perturbation response

Jialong Jiang[1,*], Sisi Chen[1,2,11*], Tiffany Tsou[1,2], Christopher S. McGinnis[3], Tahmineh Khazaei[1], Qin Zhu[3], Jong H. Park[1,2], Inna-Marie Strazhnik[1], John Hanna[4], Eric D. Chow[5,6], David A. Sivak[7], Zev J. Gartner[3,8,9,10], and Matt Thomson[1,2†]

[1]Division of Biology and Biological Engineering, California Institute of Technology, Pasadena, California, 91125, USA

[2]Beckman Single-Cell Profiling and Engineering Center, California Institute of Technology, Pasadena, CA, 91125, USA

[3]Department of Pharmaceutical Chemistry, University of California San Francisco, San Francisco, CA, 94143, USA

[4]Department of Pathology, Harvard Medical School and Brigham and Women's Hospital, Boston, MA, 02115, USA

[5]Department of Biochemistry and Biophysics, University of California San Francisco, San Francisco, CA, 94143, USA

[6]Center for Advanced Technology, University of California San Francisco, San Francisco, CA, 94143, USA

[7]Department of Physics, Simon Fraser University, Burnaby, BC V5A 1S6, Canada

[8]Helen Diller Family Comprehensive Cancer Center, San Francisco, CA, 94115, USA

[9]Chan Zuckerberg BioHub, University of California San Francisco, San Francisco, CA, 94143, USA

[10]Center for Cellular Construction, University of California San Francisco, San Francisco, CA, 94143, USA

[11]Apertura Gene Therapy, 345 Park Ave South, New York, NY 10010

[*]These authors contributed equally to this work

[†]mthomson@caltech.edu

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. Here, we introduce a computational framework, D-SPIN, that generates quantitative models of gene-regulatory networks from single-cell mRNA-seq data sets 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.

Introduction

The human genome encodes more than $30,000$ genes, but a typical human cell expresses only 5,000 genes at any given time [1]. Gene regulatory networks within cells modulate gene expression based upon environmental cues and the cell's internal state [2, 3, 4, 5]. In gene regulatory networks, transcription factors and associated regulatory proteins interact with one another and gene promoters to activate or repress gene transcription [6, 7, 8]. Gene regulatory networks play a central role in cellular decision-making processes [9, 10, 11, 12]. 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 [13, 14, 15, 16, 17, 18, 19]. 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 [20]. 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 [10, 11, 12]. 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 cellular 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 perform bottom-up

gene network reconstructions through pairwise binding measurements requiring order-M^2 biochemical assays for a network with M components [21, 22]. 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 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 [23, 24, 25, 26]. For global analysis of gene regulatory networks, 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 perturbation has, like biochemistry, been limited by experimental scale when genes were knocked out in bulk assays.

Developments in single-cell genomics and perturbation bar-coding circumvent some of the conventional limitations of perturbation-driven network reconstruction. Perturbation bar-coding 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 [27, 28, 29, 30]. 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 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.

Here, we introduce 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 Markov random fields, to infer regulatory interactions between gene-expression programs and applied perturbations. To enable scaling of D-SPIN to large data sets with 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 two steps, construction of a unified gene regulatory network model and inference of how each perturbation interacts with the gene programs within the unified network. To maximize the interpretability of D-SPIN models and computational efficiency, D-SPIN operates on gene-expression programs, rather than single genes, enabling D-SPIN to derive a coarse-grained description of a cellular regulatory network in terms of physiological units—gene-expression programs—and their interactions with applied perturbations. In benchmarking tests, we demonstrate that D-SPIN exhibits state-of-the-art performance on simulated gene network inference tasks.

We apply D-SPIN to construct gene regulatory network models from experimental data sets with thousands of perturbations and millions of cells, including two of the largest single-cell mRNA-seq data sets in existence, a genome-wide Perturb-seq experiment [28] and a new human immune cell drug-response experiment we performed. 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 response to perturbation. Applied to profile the response of human immune cells to combinations of immunomodulatory drugs, D-SPIN models remonstrate that drug combinations can generate novel cell states through additive recruitment or superposition of drug-specific gene expression programs. 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.

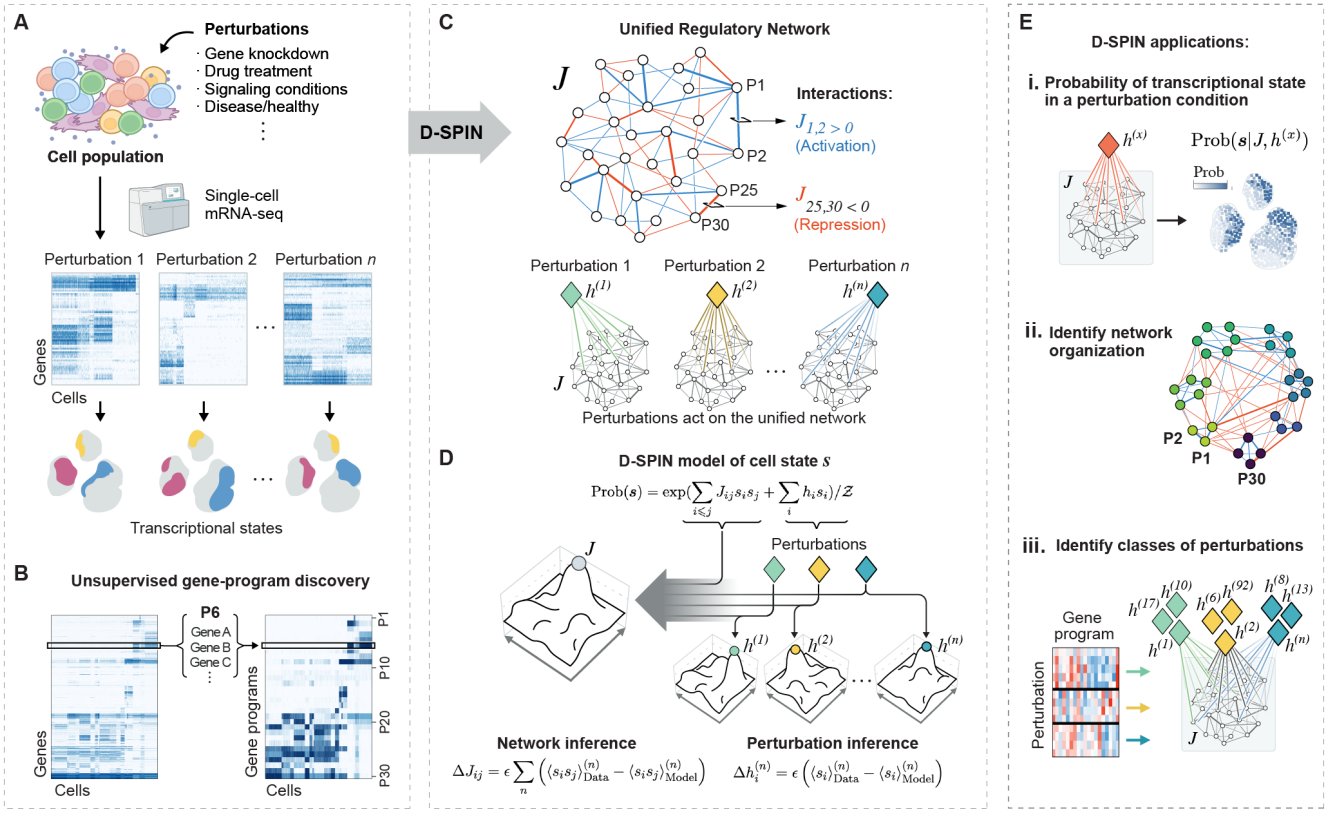

Figure 1: **D-SPIN constructs unified gene regulatory network models from single-cell mRNA-seq data collected across perturbation conditions.** (A) D-SPIN accepts as input single-cell mRNA-seq data from a cell population profiled across a series of different perturbed experimental conditions such as genetic perturbations, drug treatments, signaling conditions, or disease/healthy conditions. (B) To reduce the dimensionality of the network-inference task, D-SPIN extracts gene-expression programs from the data using orthogonal non-negative matrix factorization. (C) D-SPIN then constructs a unified regulatory network $\boldsymbol{J}$, 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 $\boldsymbol{J}$ and sample-specific perturbation response vectors $\boldsymbol{h}^{(n)}$. D-SPIN scales to large data sets with millions of cells through a paralleled inference 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 $\boldsymbol{J}$ and each perturbation response vector $\boldsymbol{h}^{(n)}$. (E) D-SPIN yields a probabilistic model that can be interpreted as a gene regulatory network. The D-SPIN model can be applied to (i) estimate the distribution of transcriptional state defined by the regulatory network model under a specific condition, (ii) reveal principles of gene regulatory network organization including modularity and sub-network architecture, and (iii) classify perturbations into groups that reflect biological function.

D-SPIN infers a unified regulatory network model from single-cell perturbation data

Here, we develop mathematical and computational strategies to integrate data across many different single-cell perturbation experiments into a network model that provides qualitative insights into 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 (Fig. 1, SI Text Sec. 1.1). Specifically, 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. The models have been generalized and applied to many different systems of interacting elements including neural networks, bird flocks, and proteins [31, 32, 33]. 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 $P(s_1, \ldots, s_M)$ of each potential configuration of the network of elements based upon inferred interactions between elements, as quantified in a coupling matrix $\boldsymbol{J}$.

In general, D-SPIN can be applied to model networks of individual genes; however, cells express thousands of genes, and gene-level models are both computationally expensive and also difficult to interpret. Therefore, we designed D-SPIN to construct reduced or low-dimensional spin-network models by modeling interactions between gene programs, i.e., co-regulated groups of genes [34, 35, 36]. Cells regulate their transcriptional state by modulating transcription factors that impact sets of genes related to differentiation state, cellular pathways, and core physiological processes. By focusing on gene 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 identifies co-regulated groups of genes, gene programs, through unsupervised orthogonal nonnegative matrix factorization (oNMF) and phenotypic gene-program annotation (Fig. 1A, B, SI Text Sec. 1.4) [37, 38].

Following the extraction of gene programs, D-SPIN applies a maximum likelihood statistical inference procedure to infer a network of regulatory interactions between gene-expression programs and between programs and applied perturbations (Fig. 1C). Regulatory interactions between gene programs are represented in a gene-program interaction matrix $\boldsymbol{J}$ 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 $\boldsymbol{h}^{(n)}$ where $h_i^{(n)}$ quantifies the interaction of the perturbation n with gene program i.

Mathematically, an energy function $E(\boldsymbol{s})$ computes the effective energy of a given transcriptional state $\boldsymbol{s}$ by balancing the influence of the regulatory network (encoded in $\boldsymbol{J}$) and the perturbation-specific inputs (encoded in each $\boldsymbol{h}^{(n)}$). The resulting probability of a given transcriptional state $\boldsymbol{s}$ is $P(\boldsymbol{s}) = \frac{1}{Z}\exp[-E(\boldsymbol{s}; \boldsymbol{J}, \boldsymbol{h})]$, where

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

and Z is a normalizing constant ensuring all probabilities sum to 1 (SI Text Sec. 1.5) [39, 40]. The regulatory network accounts for positive ($J_{ij} > 0$) and negative ($J_{ij} < 0$) interactions between gene programs. Perturbations exert a bias on each gene program to increase the likelihood of being on ($h_i > 0$) or off ($h_i < 0$).

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 $\boldsymbol{J}$ and $\boldsymbol{h}^{(n)}$ given the data (Fig. 1D). For spin-network models, the optimization procedure yields closed-form parameter-update rules

$$\Delta J_{ij} = \epsilon \sum_n \left(\langle s_i s_j \rangle_{\text{Data}}^{(n)} - \langle s_i s_j \rangle_{\text{Model}}^{(n)} \right) \tag{1}$$

$$\Delta h_i^{(n)} = \epsilon \left(\langle s_i \rangle_{\text{Data}}^{(n)} - \langle s_i \rangle_{\text{Model}}^{(n)} \right) \tag{2}$$

where $\langle s_i s_j \rangle_{\text{Data}}^{(n)}$ are gene-program cross-correlations computed from data for each of n perturbation conditions, $\langle s_i s_j \rangle_{\text{Model}}^{(n)}$ are gene-program cross-correlations computed from the model, and ϵ is a learning rate (SI Text Sec. 1.7). At each step of model inference, the procedure samples from the model for each perturbation condition and updates the regulatory network $\boldsymbol{J}$ and each perturbation vector $\boldsymbol{h}^{(n)}$. In general, spin-network model inference is a convex optimization problem, so the optimization has a single unique solution.

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 on 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 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 data sets. Practically, we estimate the network using a subset of experimental conditions with high diversity, and utilize the inferred network to identify response vectors of each condition independently (SI Text Sec. 1.7, 1.8).

Following inference of $\boldsymbol{J}$ and $\boldsymbol{h}$, D-SPIN models can be analyzed as network models that reveal regulatory connections between cellular processes and identify interactions between applied perturbations and cellular gene-expression programs. (Fig. 1E, ii-iii). D-SPIN models can also be used to predict the distribution of transcriptional states in a cell population under normal and perturbed conditions and 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 UMAP with mechanistic modeling approaches that view the cell as a regulatory network with links between cellular processes, pathways, and perturbations. D-SPIN 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 UMAP.

D-SPIN achieves state-of-the-art performance on gene regulatory network inference

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 [41]. The HSC network has two major groups of regulatory genes, the first group includes Pu1 controlling granulocyte/monocyte differentiation and the second group includes Gata1 controlling megakaryocyte/erythrocyte differen-

tiation [42]. 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 (Fig. 2B, C, SI Text Sec. 1.9, SI Table 1). For interpretation of the data we used the biological functions of each transcription factor to classify the single-cell expression profiles into 7 different cell state designations including monocytes (Mono), granulocytes (Gran), erythrocytes (Ery) and megakaryocytes (Mega) [42].

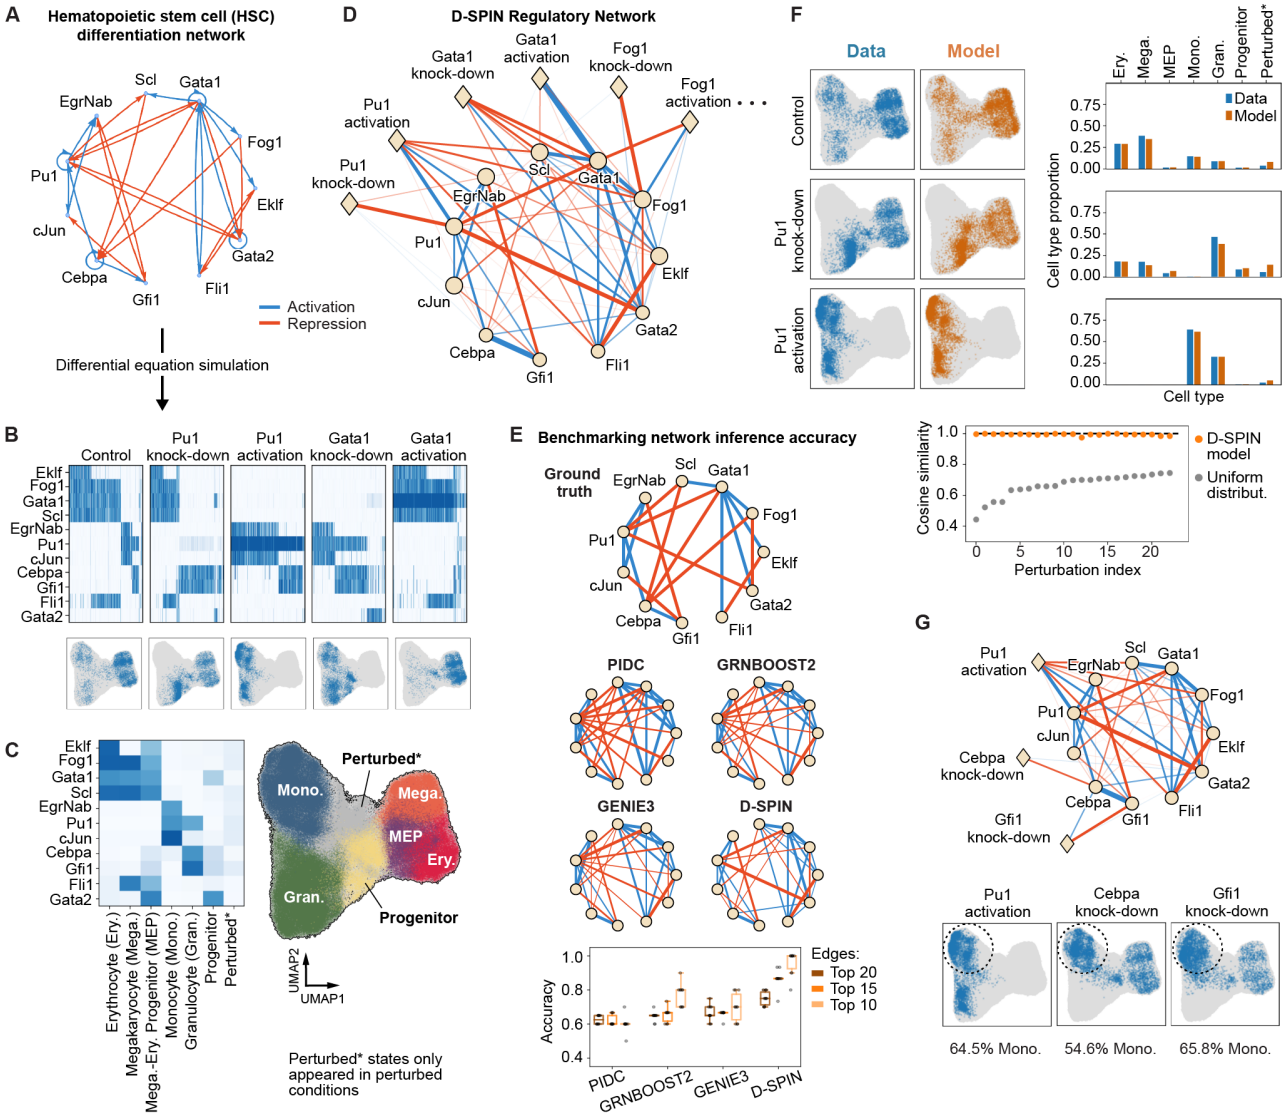

Figure 2: **D-SPIN achieves state-of-the-art performance on network inference benchmarking task, revealing cell-fate control strategies for hematopoietic stem cell differentiation network** (A) HSC regulatory network model [42] contains 11 transcription factors that interact through positive (blue) and negative interactions (red) to modulate the differentiation of HSCs into different cell types. (B) Example simulated gene-expression profile heatmaps generated using published computational model of the network for a series of simulated single-gene knockdown and activation conditions. (C) The simulated network generates a distribution of cell states that can be classified by their transcription factor expression into 7 cell states. The heatmap shows transcription factor expression patterns that define each cell state generated across all simulated conditions, visualized as a UMAP. (D) Using the simulated data, D-SPIN infers an integrated regulatory network model that encodes inferred interactions between transcription factors as an interaction graph. D-SPIN also infers perturbation vectors that estimate how knockdown or activation of transcription factors (e.g., Gata1 knockdown) impacts the regulatory network through up- or down-regulation of transcription factors. (E) The D-SPIN model has greater accuracy than state-of-the-art methods [41], as shown through the rendering of the network inferred from data by each method as well as a box plot of the number of correctly inferred edges across different repeats of simulated gene-expression data. The box plot shows the accuracy of top k edges for different values of k where the ground truth network has a total of 20 edges. (F) The D-SPIN model predicts the cell-fate distribution generated through perturbations to the underlying network. UMAP plots of cell state distribution generated for control and Pu1 activation/knockdown are shown for simulated data and the D-SPIN model. Bar graphs quantify the proportion of HSC-derived cell states in each condition from simulated data and the D-SPIN model. The data and model distribution are highly consistent across all 22 perturbations as quantified by the cosine similarity between cell state distributions of data and model. (G) The model can be used to understand how perturbations to distinct transcription factors can generate similar cell populations. Both Pu1 activation and knockdown of Cebpa or Gfi1 generate an increased Mono state by directly activating Mono genes or indirectly through the interaction between Gfi1 and EgrNab.

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 3 different network reconstruction methods (PIDC, GENIE3, and GRNBoost2) that were the top performers in the BEELINE benchmarking study [41]. We applied D-SPIN to generate a single network model encoded in the gene interaction matrix $\boldsymbol{J}$ and generate perturbation vectors $\boldsymbol{h}^{(n)}$ for each perturbation condition (Fig. 2D, SI Fig. 1A). We applied the algorithms to individual genes, rather than gene programs, given the relatively compact size (11 nodes) of the network used in the BEELINE benchmarking evaluation. Each method generates a set of inferred edges and provides confidence estimates for each edge. We compared the methods by considering the top 10, top 15, 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 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 (Fig. 2E). 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. As an example where D-SPIN outperforms existing methods, the networks inferred by PIDC, GENIE3, and GRNBoost2 all identified a spurious activating interaction inside the Mega.-Ery. gene module, e.g., Scl-Fog1 or Scl-Eklf, while D-SPIN pruned these false positive edges and revealed the central regulatory role of Gata1. D-SPIN also pruned the spurious inhibiting interactions between the two modules, e.g., Pu1-Eklf or Pu1-Fli1, revealing the actual regulatory target of Pu1. When spurious edges occur they were similar across the different inference frameworks: all 4 methods identified spurious edges in strongly connected regions of the HSC network, struggling to separate direct versus indirect connections [43]. However, for D-SPIN uniquely, spurious edges can be further reduced through the sampling of higher-order perturbations (SI Fig. 1B, SI Table 1).

D-SPIN predicts the impact of simulated perturbations on differentiation

D-SPIN is a probabilistic or generative model that can predict the distribution of transcriptional states in a cell population subjected to a perturbation, unlike the other models, which use either regression or information-theoretic measures to find gene linkages. D-SPIN infers a quantitative model of the cell population generated by the HSC network that can be sampled to generate quantitative predictions regarding the distribution of cell types induced by each gene knockdown or activation. Unlike existing methods, D-SPIN allowed us to ask how specific gene knockdown or over-expression strategies could be applied to modulate the differentiation trajectory of HSC cells.

We applied D-SPIN to compute the population structure of a simulated HSC cell population when different transcription factors are perturbed. The model distributions quantitatively agree with data, and capture important biological impacts of the perturbations (Fig. 2F, SI Fig. 1C). Analysis of the model allowed us to gain insight into how disruption of different underlying network components might impact the differentiation trajectory of the cells. For example, when Pu1 is knocked down, it results in a complete loss of Monocyte cell states and an increase in Granulocyte cell states. The change in population structure occurs because Pu1 directly activates EgrNab and cJun, and the three factors together induce the Monocyte cell state while EgrNab represses Gfi1, a key regulator of the Granulocyte cell state. Knocking down Pu1 turns off the Pu1-cJun-EgrNab group leading to loss of the Monocyte cell state as well as activation of Gfi1. Gfi1 activation leads to the activation of Cebpa and the Granulocyte differentiation program. The Mega-Ery module (Eklf, Fog1, Gata1, and Scl) is affected by both the reduced inhibition from Pu1 and increased inhibition from Cebpa. Consequently, the cell states Mega and Ery remain present, albeit in reduced proportions.

Both Pu1 activation and the knockdown of Cebpa and Gfi1 lead to an increased prevalence of Monocyte cell states (Fig. 2G). Pu1 directly activates the monocyte regulatory network that contains EgrNab and cJun. Gfi1 represses monocyte differentiation through repression of EgrNab, and knockdown of Gfi leads to activation of EgrNab which in turn activates Pu1 and cJun, resulting in increased Monocyte states. The knockdown of Cebpa turns off Gfi1 because of the positive interaction between the two, leading to the activation of Monocyte states in a similar mechanism as Gfi1 knockdown.

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. D-SPIN prunes the ambiguity of alternative models, and builds a coherent integrated network upon which perturbations act. 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.

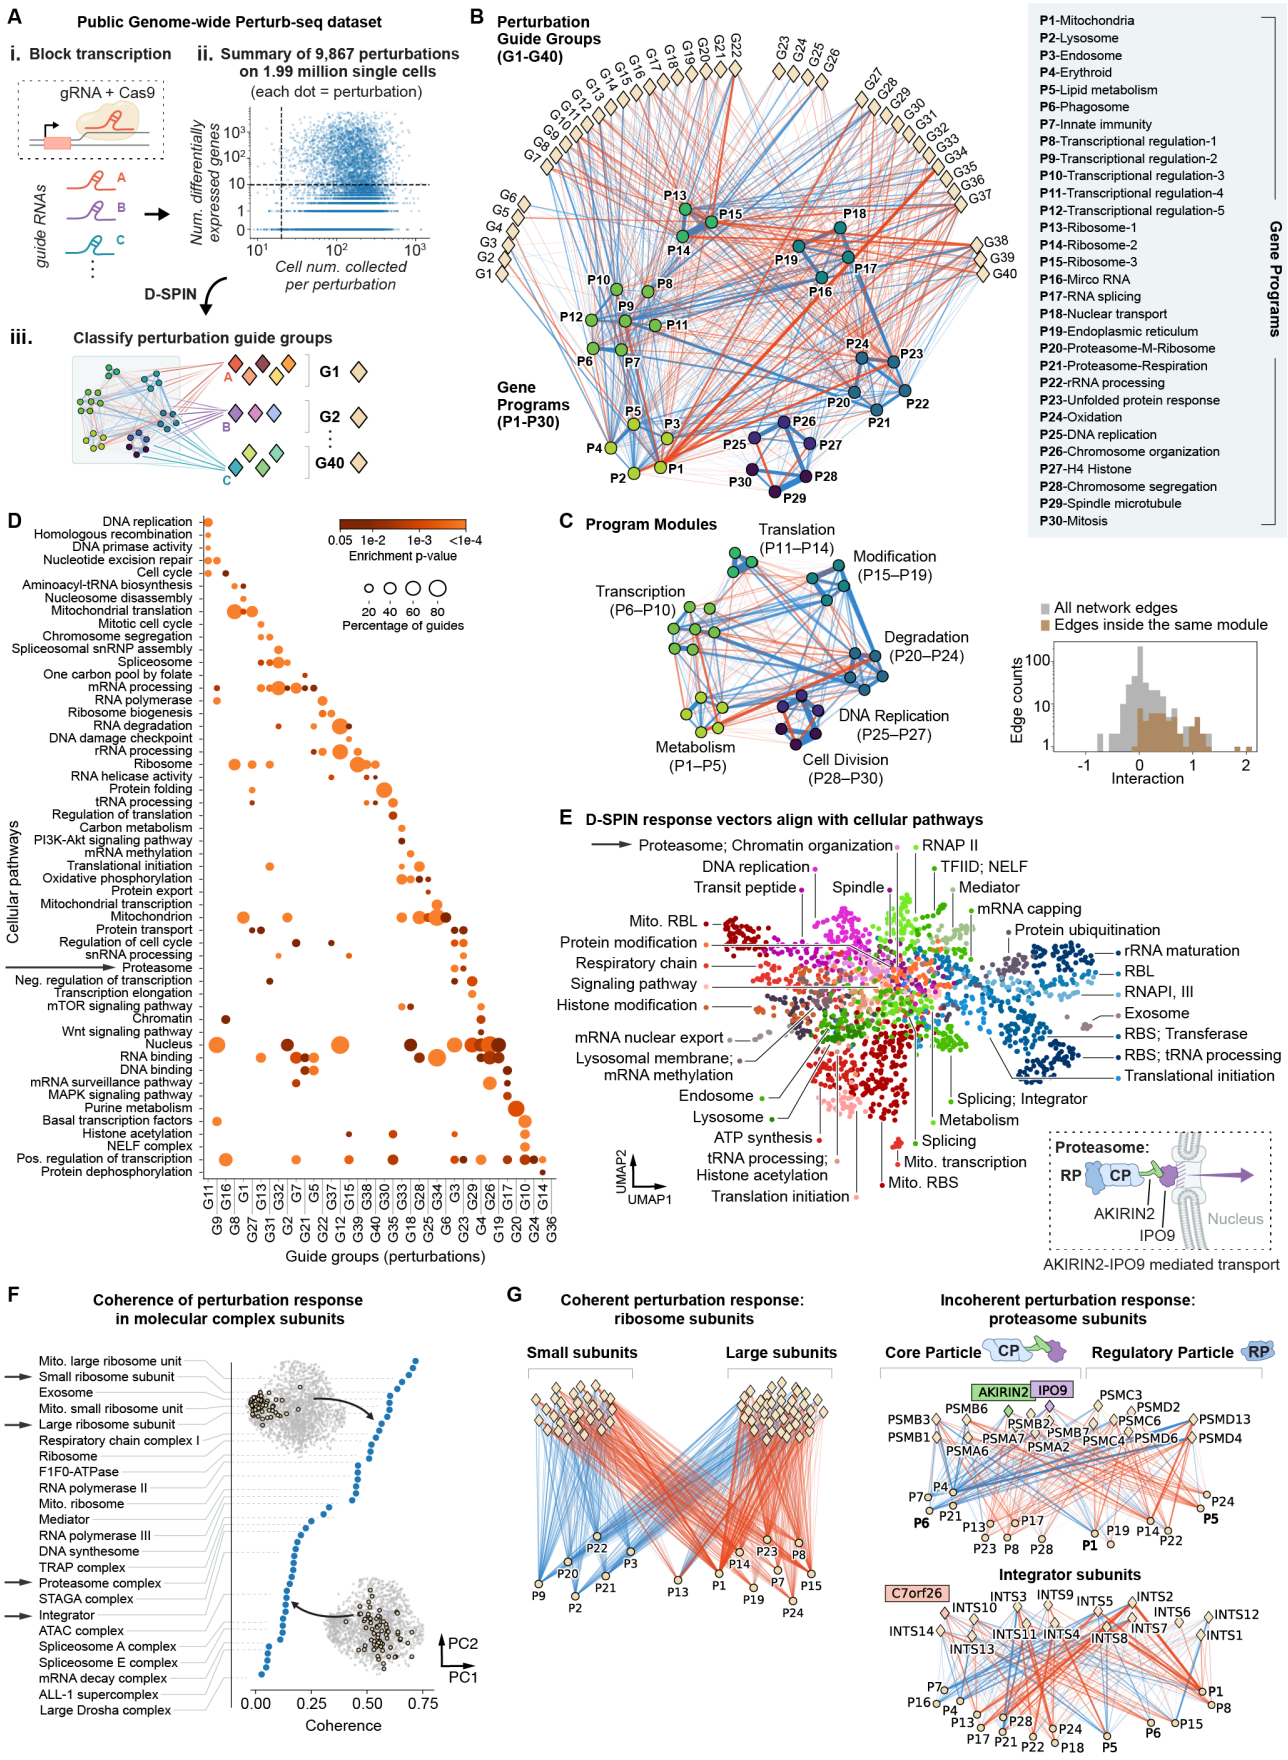

Figure 3: Caption on the following page

Figure 3 **D-SPIN generates a regulatory network model from genome-wide Perturb-seq data that reveals the architecture of cellular pathways and sub-functions for pathway components.** (A) Genome-wide Perturb-seq experiment [28] delivered 9,867 gene knockdown perturbations individually to ∼2M human K562 cells using guide RNAs. (ii) D-SPIN model construction was applied to knockdown perturbations that induced more than 10 differentially expressed genes and more than 20 cells passed quality filtering. (iii) D-SPIN model was applied to generate a regulatory network model, inferring interactions between gene-expression programs and impacts of perturbations that classify guides into groups. (B) Inferred D-SPIN model contains a core gene regulatory network composed of 30 gene-expression programs and the response of each gene knockdown. Gene programs are annotated by a combination of function annotation tools and manual lookup. D-SPIN classified gene knockdowns into 40 different guide classes with Leiden clustering. The interactions between gene-expression programs and perturbations are rendered as positive (blue) or negative (red) edges, and edge thickness scales with the absolute value of edge weight. (C) D-SPIN network model exhibits modular architecture with 7 core network modules that can be assigned physiological functions: metabolism, transcription, translation, modification, protein degradation, DNA replication, and cell division. Interactions between programs in the same modules are mostly activation and contain the majority of strong activation interactions. (D) Guide RNA classes are each enriched for physiological functions and cellular pathways in gene ontology and KEGG pathway databases. The bubble size represents the percentage of enriched guides in the class. (E) UMAP representation of the response vector of each guide RNA agrees with the global organization of major cellular pathways including transcription, translation, respiration, and cell cycle. (F) Regulatory coherence of macromolecular complex subunits, quantified by the average pairwise cosine similarity between the perturbation response vector for knockdown of complex sub-components. High-coherence complexes have subunits with similar impacts on the D-SPIN gene regulatory network. In low-coherence complexes, knocking down different subunits leads to different impacts, suggesting distinct functions for complex subunits. (G) Perturbation vectors for subunits of the ribosome, proteasome, and integrator complexes. The ribosome has a coherent signature where components of the complex have similar inferred response vectors, while 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 impact on a gene program (P5, P6), or even display opposite effects on a gene program (P1). D-SPIN also identified recently annotated pathway components IPO9 and AKIRIN2 (proteasome) and C7orf26 (integrator).

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

Peturb-seq (or its derivatives) involves the knockdown of individual genes, using Cas9 and a targeting guide RNA. Single-cell mRNA-seq is used to assess the transcriptional state of each cell and to identify the guide RNA delivered to the cell. The technique can be applied at scale to interrogate the impact of thousands of distinct gene knockdowns on cellular gene expression. We used D-SPIN to construct a gene regulatory network model using data from a genome-wide Perturb-seq experiment on the erythrocyte leukemia K562 cell line [28]. K562 is an erythrocyte precursor cell line with a BCR-ABL mutation, derived from a chronic myelogenous leukemia (CML) patient [44, 45]. 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 (SI Text Sec. 1.2). Guide RNA perturbation impact varies (Fig. 3A): 70.5% of perturbations had less 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 less than 100 cells collected. We focused our analysis on 3,138 perturbations with more than 10 DEGs and 20 collected cells, and D-SPIN was able to construct the network model in 8 hours with parallelization on 500 CPU cores.

We constructed a D-SPIN model using a set of 30 gene programs that D-SPIN extracted with its automated oNMF procedure (Fig. 3B, SI Fig. 2A, SI Text Sec. 1.4, SI Table 2). The extracted gene programs reflect both general cell biology as well as lineage-specific gene programs that reflect the specific biology of K562 cells. We annotate the gene programs with a combination of bioinformatic databases including DAVID, Enrichr, and manual lookup [46, 47]. Among the programs, D-SPIN 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 (Fig. 3B).

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 (Fig. 3B, SI Fig. 2B, SI Table 2). The D-SPIN network contains gene programs, as nodes, and interactions between gene programs as edges that represent inferred interactions. D-SPIN identified 84 activating and 21 inhibiting 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 (Figs. 3B, C). We automatically decomposed the D-SPIN network into seven modules through the Leiden community detection algorithm [48] (SI Text Sec. 1.10). Each identified module reflects a core cellular function, including transcription, translation, protein-degradation, and cell-cycle-related functions.

The network contains a set of inhibitory interactions between gene programs that are expressed in distinct cellular states. The strongest inhibitory interaction is between gene programs expressed at different stages of mitosis, for example, P29 Spindle microtubule inhibits both P25 DNA replication and P27 Histone. The strong positive interactions inside the cell cycle sub-network of P25-P30 reconstruct the transcription state distribution during cell cycle progression (SI Fig. C, D). We also find strong inhibiting 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, ASAH1, PLD3). We also observe inhibition between P4 Erythroid and P6 Phagosome, where the phagosome is an innate immune response mounted by macrophages and other myeloid cell types [49]. This indicates the presence of two, mutually exclusive differentiation paths that lead to erythroid and myeloid cell fates.

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 knockdown perturbation. The interactions between perturbations and gene programs are represented by perturbation response vectors $\boldsymbol{h}^{(n)}$. The entries of $\boldsymbol{h}^{(n)}$ 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 (Fig. 3B).

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 $\boldsymbol{h}^{(n)}$ vectors (SI Table 2). 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 (Figs. 3D, E). Guide clusters are enriched for specific cell-biological functions, including DNA replication, MAPK signaling pathway, and RNA degradation (Fig. 3D). The guide clusters also provide detailed information regarding the structure of cellular pathways. As a specific example, the core components of the proteasome fall into a single perturbation cluster. The cluster also contains AKIRIN2 and IPO9, two protein components recently associated with the proteasome. Similarly, D-SPIN identifies C7orf26 as associated with the integrator subunits INTS10, INTS13 and INTS14, a key result of the original primary Perturb-seq study [28].

Predicting macro-molecular complex component function

Macromolecular complexes like the ribosome, the proteasome, and the integrator complex play a critical role in cell physiology and gene regulation. We used D-SPIN perturbation vectors to analyze the functions of individual macromolecular-complex components. D-SPIN generates a perturbation vector $\boldsymbol{h}^{(n)}$ for each subunit of each macromolecular complex that was knocked down in the Perturb-seq experiment. We applied D-SPIN to gain insight into sub-functions of complex components by comparing cosine similarity between $\boldsymbol{h}^{(n)}$ for different components of known complexes in the CORUM database [50]. We selected complexes that have more than two-thirds of subunits in the dataset and picked representative ones from subunits with more than half overlapping subunits, resulting in a list of 46 core cellular complexes. For each complex in

the list, we defined a coherence score by computing the average cosine similarity between response vectors, $\boldsymbol{h}^{(n)}$, for all pairs of complex components (SI Text Sec. 1.13). 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 and an average cosine similarity across complex components approaching 1. When complexes can achieve distinct functions through combinations of distinct protein components, then such complexes have with response vectors with average cosine similarity closer to 0 which is the random expectation.

Across the 46 complexes in CORUM, our analysis revealed a wide range of coherence scores across different complexes, ranging from 0.01 (large Drosha complex) to 0.73 (mitochondrial large ribosome) (Fig. 3F, SI Fig. 2E). 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 Drosha complex, DGCR8 complex, ALL-1 supercomplex (involved in chromatin regulation), the ATAC complex (histone acetyltransferase), and spliceosome all had lower coherence scores $< .25$. The low-scoring 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 (SI Table 2) providing a specific framework for generating experimental hypothesis to determine how modulation of subunits composition of these macromolecular regulatory complexes might enable the complex to achieve a range of regulatory functions.

The ribosome behaves as a coherent complex where 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 (Fig. 3G). Perturbations to ribosome subunits activate protein digestion-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 digestion and inhibition of protein production and secretion processes.

The proteasome and integrator behave as incoherent complexes, and have response vectors that split subunits into smaller sub-groups. The integrator complex subunits are partitioned into 4 sub-groups that agree with the original Perturb-seq study (Fig. 3J): 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 inhibited 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 subunit PSMC6 (also known as TBP7 and Rpt4) has a negative impact on respiration, and 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 [51]. Further, AKIRIN2 and IPO9 are grouped together with the 20S proteasome subunits by the response vectors. These results demonstrate that the response vectors reveal the impacts of genetic perturbations, and provide a compact signature for identifying functional associations between genes [52].

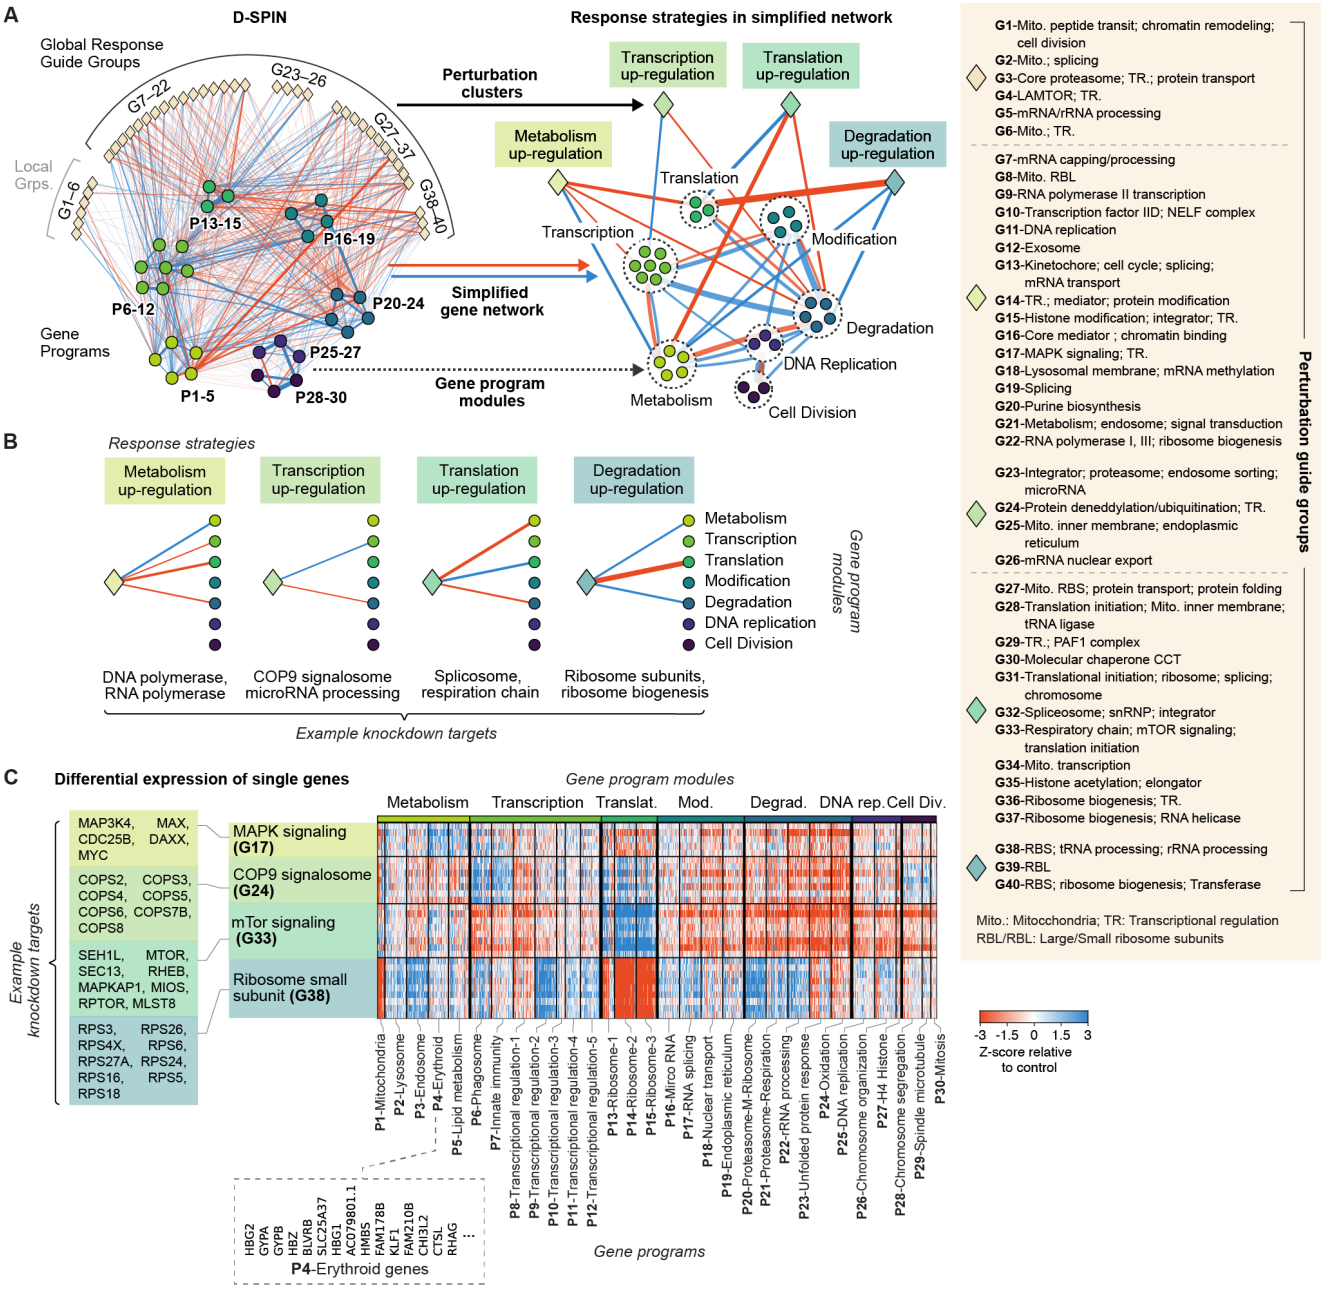

Figure 4: **D-SPIN network model identifies global regulatory response strategies mounted by K562 cells to distinct classes of gene perturbation.** (A) A coarse-grained network model is generated by grouping gene programs into 7 identified gene modules and grouping guide RNA gene knockdown classes based on network impact. The resulting coarse-grained model enables global analysis of cellular regulatory responses. (B) The minimal model reveals that K562 cells mount 4 distinct classes of global regulatory responses in response to gene knockdown perturbations. 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. (C) Gene-expression heatmap showing the upregulation (blue) or downregulation (red) of genes in gene programs to perturbation of individual genes in a series of cellular pathways. Heatmap indicates the change in gene expression for perturbed cell population compared to the control where a non-targeting guide RNA is introduced to cells.

D-SPIN coarse-graining reveals global regulatory strategies

While the D-SPIN model in Fig. 3 is compact, we found that the model could be simplified further through an automated coarse-graining strategy (diagrammed in Fig. 4A, SI Text Sec. 1.11) 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 with coarse-graining strategies in physical models that 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 4 classes of damage-response strategies, patterns of gene activation, and inhibition generated by the cell when specific cellular processes are disrupted.

The D-SPIN-identified regulatory strategies suggest non-trivial forms of homeostatic regulation. We name the strategies by the most prominent upregulation response to the perturbations, because in the context of knockdown perturbations, gene activation represents the active regulatory feedback of cells to encounter the perturbed or stressed condition. Therefore the four global response strategies are metabolism upregulation, transcription upregulation, translation upregulation, and degradation upregulation. For example, 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 while also downregulating processes that degrade proteins. Interestingly, perturbations to ribosome subunits and ribosomal rRNA (G38, G39, G40) production actually 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.

The regulatory strategies also provide insight into the flow of information in the cell. Knockdown of genes often impacted processes that suggest long-range regulation or coupling between seemingly distinct processes. For example, a large proportion of metabolism upregulation responses are related to transcription processes, including RNA polymerase (G9, G22), Transcription Factor IID, NELF complex (G10), and mediator (G16). The knockdown of these genes led 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 [53]. 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 (Fig. 4C). Knockdown of mTOR signaling components leads to the upregulation of ribosomal programs in a translation upregulation response. Finally, the knockdown of ribosomal subunits leads to the downregulation of translation and upregulation of transcription and protein degradation.

These results demonstrate that D-SPIN-generated regulatory networks provide insights into control strategies used by cells to regulate a distributed set of processes to achieve homeostasis in response to various forms of internal damage. D-SPIN models can provide global insights into classes of regulatory strategies, suggest mechanisms of compensation and feedback, and also identify information flow between cellular processes that might not be intuitively connected, thereby providing insights into principles of information processing and homeostatic control.

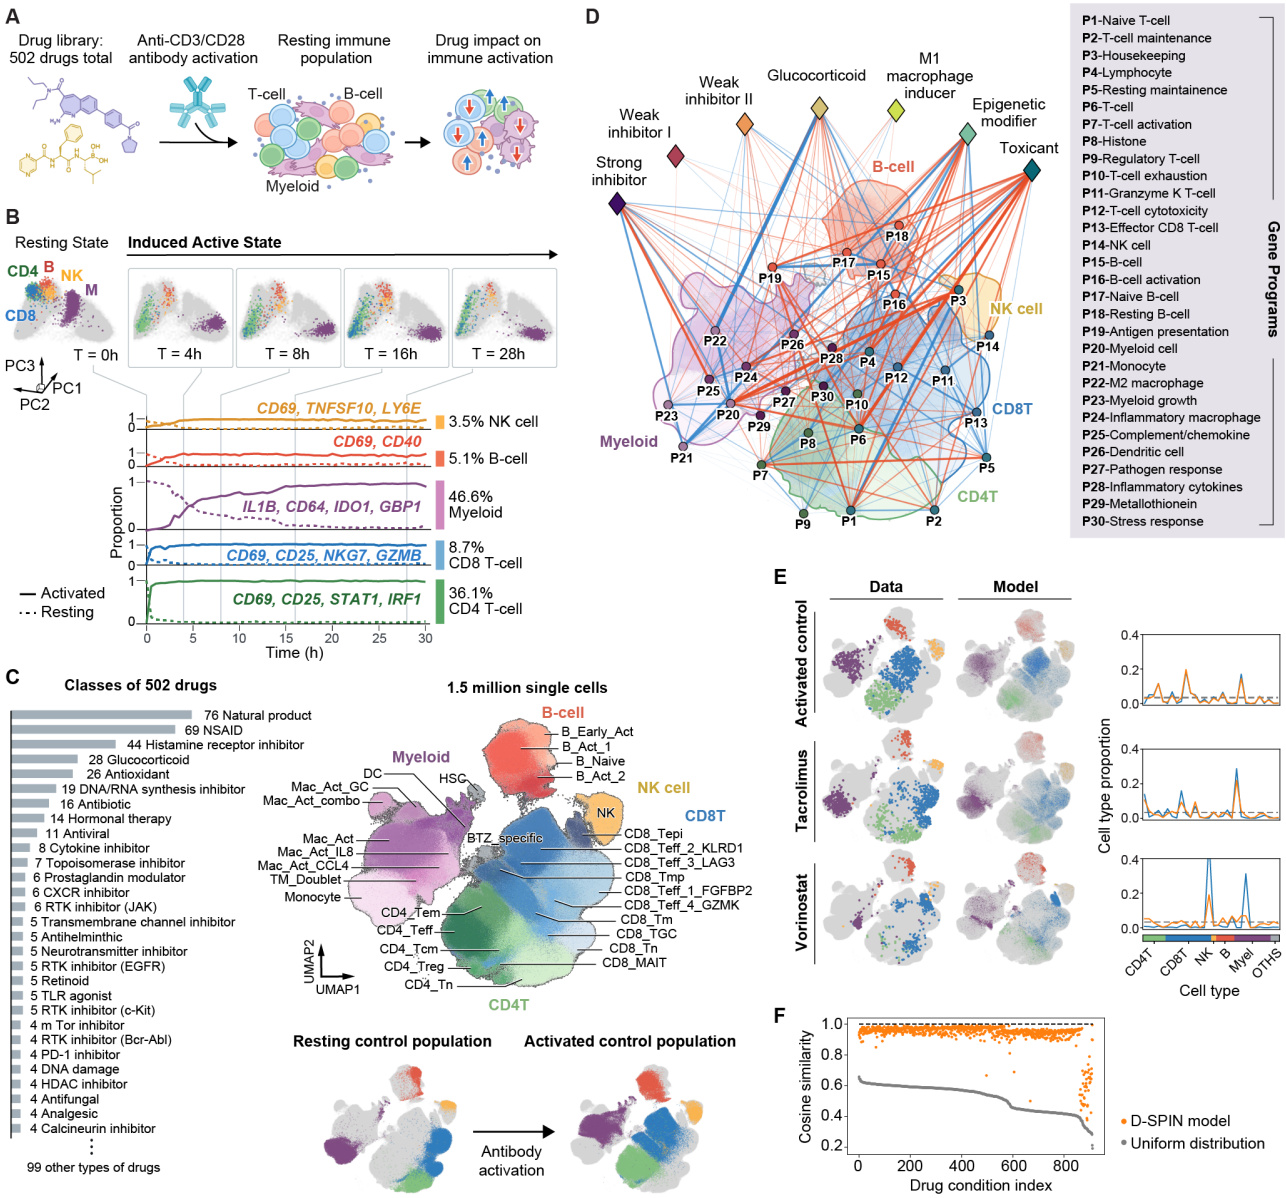

Figure 5: Caption on the following page

Figure 5 **D-SPIN derives drug-response network model from human immunomodulatory drug-response single-cell mRNA-seq profiling experiments.** (A) Illustration of drug-response experiment design. 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) PCA projections derived from single-cell mRNA-seq of T-cell mediated immune activation at selected time points within a 30-hr time-course experiment. Time-courses of the proportion of activated and resting cells in each cell type with activation gene markers. T-cells reach activated states first in 2 hrs and myeloid cell activation lasts until 16 hr. (C) Histogram of classes of profiled drugs. The selected drugs have a variety of biochemical properties and target pathways. 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. The UMAP embedding is primarily divided into two distinct regions of resting and activated cell population. (D) The regulatory network extracted by D-SPIN contains 30 gene programs that are identified from data with oNMF and annotated with cellular function by a combination of bioinformatic database and manual lookup. Drugs are classified into seven drug classes based on response vectors identified by D-SPIN. D-SPIN network is rendered on top of a UMAP schematic of the cell population: the UMAP shows the underlying cell states, and D-SPIN nodes are positioned on the UMAP based upon the cell state in which each program is highly expressed. (E) Examples of model-generated cell state distribution compared with data. The model and data agree well 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. The model UMAPs are rendered by directly sampling from the model and mapping each state to the nearest data cell state. The dashed line in cell state distribution is uniform distribution as reference. (F) Model-generated cell-state distributions are highly similar to experimental data in cosine similarity measures. The cosine similarity is higher than 0.9 for the majority of conditions. For comparison, the null model is a uniform distribution over cell subtypes. Drug conditions are ordered by cosine similarity with uniform distribution for visualization.

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 cell types respond to therapeutic interventions. Given the role of the human immune system in many diseases, we apply 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 the response of human immune cells to 502 different immunomodulatory drugs. Our experimental system was designed to probe the modulation of a cell population undergoing T-cell driven hyper-activation as is observed in auto-immune and hyperinflammatory states [54, 55]. We cultured a heterogeneous population of donor-derived human immune cells (peripheral blood mononucleated cells, PBMC) that contain T-cells, B-cells, macrophages/monocytes, and NK-cells. T-cells were activated using a pair of T-cell receptor antibodies (Fig. 5A, SI Fig. 3A), leading to the activation of macrophage, B-cell, and NK-cell cell types (Fig. 5B, SI Fig. 3B-D). We applied the system to analyze the action of different immunomodulatory drugs on the activating immune cell population.

Single-cell mRNA-seq was used to measure the effect of 502 immunomodulatory small molecules on T-cell driven-immune activation (Fig. 5C). The drug library contains a diverse set of small molecules targeting pathways including mTor, MAPK, glucocorticoids, JAK/STAT, and histone deacetylases (HDAC) (Fig. 5C, SI Table 3). 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 1μM at the same time as the anti-CD3/CD28 antibody, and profiled the cell population using single-cell mRNA-seq at 24 hours [56]. 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 (Fig. 5C, SI Fig. 3E, SI Table 3).

Classification of small molecules using D-SPIN

Using the single-cell mRNA-seq data, we applied D-SPIN to construct a regulatory network that models the interaction of individual small molecules in the drug library with the immune cell population, which contains various cell types (T-cells, B-cells, myeloid cells, and NK-cells) as well as states of activation and rest. 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 modeling interactions between gene-expression programs and interactions between gene-expression programs and applied drugs. D-SPIN extracted a set of 30 gene programs by oNMF of the transcription data with seeding to include drug-specific programs and the programs were sufficient to describe the transcriptional state of all cells in the population (SI Fig. 4A, B, SI Text Sec. 1.4, SI Table 3). 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), and pathogen-response programs (P27), an antigen-presentation program (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 database [46] and manual annotation, and validated on a single-cell atlas of human immune cells [57].

D-SPIN inferred a single interaction network $\boldsymbol{J}$ and a set of 502 drug-specific biasing vectors $\boldsymbol{h}^{(n)}$ (Fig. 5D, SI Table 3). The D-SPIN model contains a total of 465 network parameters and 30 parameters for each drug condition but was able to describe all profiling experiments of 502 drugs with low quantitative and qualitative error (Fig. 5E, F, SI Text Sec. 1.12). Quantitatively, we computed the cosine similarity between the model and data distribution, and in 92.4% of samples the model is $> 90\%$ similar to data (Fig. 5F). The drugs with lower cosine similarity tended to be drugs that drove the cell population into a small number of 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 to do the extremely focused cell-type distribution for these drugs.

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 (Fig. 5D). 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.

In addition to the regulatory network model, D-SPIN also inferred a response vector $\boldsymbol{h}^{(n)}$ 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. Moreover, as in the case of the Perturb-seq data, we applied clustering to group drugs into a set of 7 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 (Fig. 5C, SI Fig. 5A, B, SI Table 4). We named the classes by analyzing the structure of the cell population induced by drugs within the class. While we defined the drug classes based only on D-SPIN's inferences regarding how each drug interacts with the core network, the classes contained drugs from specific chemical families and mechanisms of action.

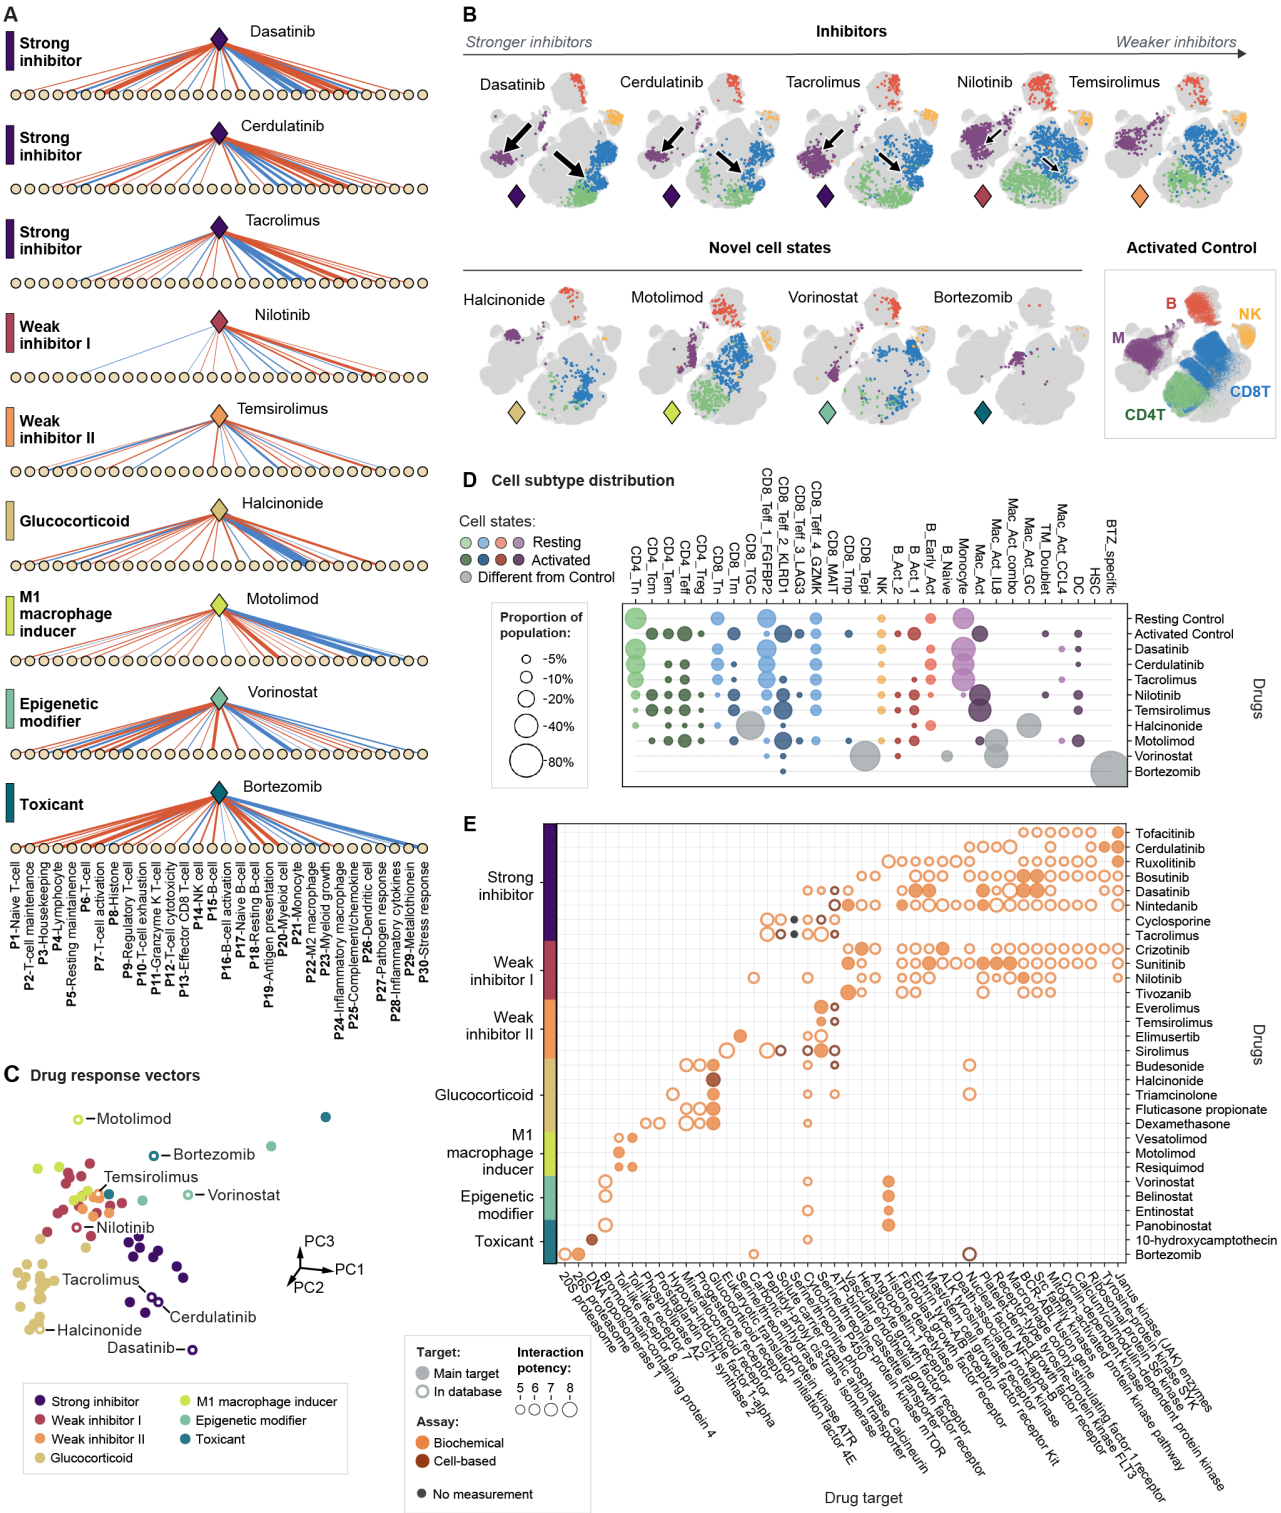

Figure 6: **Drug classification derived from D-SPIN aligns with known drug targets and mechanisms of action.** (A) Interactions between drugs and gene programs inferred from the D-SPIN model for 7 drug categories. (B) UMAPs of selected drugs from different classes. We find immune inhibitors with a spectrum of strengths, and drugs producing novel cell states. (C) Leading principal components of response vectors of identified effective drugs visualize the relationship between different drug groups. (D) Bubble plot of cell-subtype distributions of control and selected drugs. With decreasing inhibitor strength, the proportion of the activated immune-cell population (deep colors) gradually increases. Some drugs produce cell states different from both resting- and activated-control populations. (E) Biochemical targets for selected drugs in the database. Drugs in the same population response category can act on a variety of molecular targets and pathways.

Three of the seven classes include drugs that inhibit activation of the immune-cell population (Fig. 5D, SI Fig. 4A, B), including drugs we classified as strong inhibitors, weak inhibitor I, and weak inhibitor II. Across these drug classes, analysis of cell-population fractions 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 (Fig. 6A-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 cells (Fig. 6D). Weak inhibitors, such as Temsirolimus, only slightly bias the cell population towards resting, compared with the control population.

The individual drugs that fall into the inhibitor drug classes act on a variety of molecular targets and signaling pathways [58] (Fig. 6E, SI Table 4). 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 (Elimusertib). 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. Further, while many of the drugs in our data set have been previously reported as immunosuppressive including well-known drugs such as Tacrolimus, D-SPIN organized the drugs into a quantitative series and provides a single integrated framework for comparing drug action at the transcriptome scale across a large number of inhibitors with different chemical mechanisms of action.

A detailed comparison between the effects of drugs with the same primary target demonstrates that the single-cell assay and D-SPIN model can identify phenotypic patterns that capture potential underlying differences in the biochemical mechanism of drug action (Fig. 6E). A specific set of BCR-ABL kinase inhibitors (Bosutinib, Dasatinib) developed for the treatment of leukemia generated strong inhibition, while another set of BCR-ABL inhibitors, including Nilotinib, fall into the weak-inhibitor class (Fig. 6E). Biochemically, the BCR-ABL strong inhibitors can be distinguished from the BCR-ABL weak inhibitors by their increased Lck/Src kinase specificity (Fig. 6E). The Lck kinase is a central signaling node directly downstream of T-cell receptor activation. Similarly, the JAK-kinase inhibitor Cerdulatinib generated a stronger inhibitory response than the FDA-approved anti-inflammatory JAK-kinase inhibitor Tofacitinib. While both drugs have JAK-family kinases as their primary target, Cerdulatinib also targets the kinase SYK. Therefore, the response vectors generated by D-SPIN allowed us to identify potential signaling nodes that differentiate responses from related drugs.

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 (Fig. 5D, 6A-D). The glucocorticoid 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 between the classes. 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 P22 M2 macrophage named due to the presence of genes associated with immunosuppressive M2 macrophages including CD163 [59, 60]. 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 induces the activation of inflammatory, pathogen-responsive or 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 D-SPIN identifies to have activated cell-stress programs. Toxicants include the proteasome inhibitor Borte-

zomib and DNA damage agents 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 HDAC inhibitors which generate an epigenetically disrupted T-cell state (CD8_Tepi) that have elevated expression of histone-related genes such as histone component genes and DNA topoisomerase TOP2A.

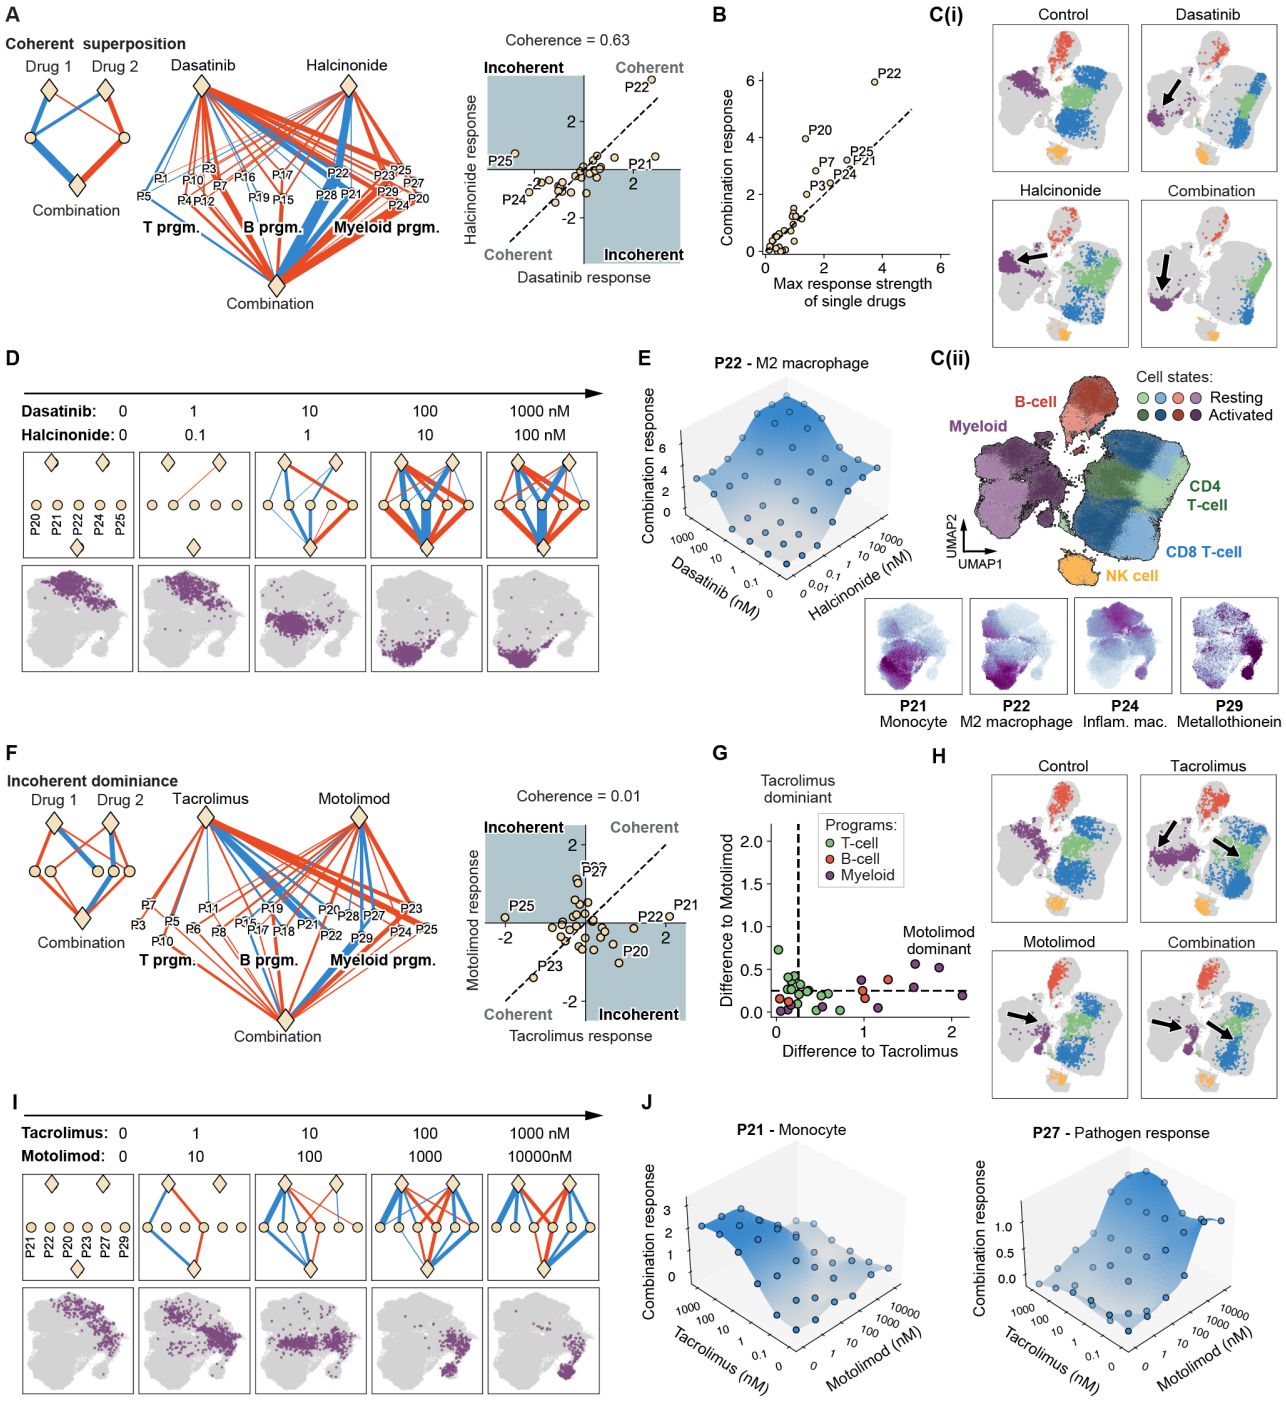

Figure 7: Caption on the following page

Figure 7 **D-SPIN reveals network-level mechanisms of combinatorial drug action.** (A) Schematic diagram of coherent superposition and the combination perturbation response of Dasatinib and Halcinonide. The two single drugs and their combination have a coherent effect on gene programs by activating/inhibiting the same set of programs. The two drugs have a high coherence score (0.63), but different-strength interactions with gene programs. (B) The combination responses on gene programs are a superposition of single-drug responses as they are generally higher than the max strength of the single-drug responses. (C)(i) Comparison of single drugs and combination effect with UMAP rendering. The combinatorial recruitment of gene programs created a novel macrophage state that occupies a distinct region on the UMAP. (ii) UMAP rendering of cell states and gene-program expression of the two dosage-titration experiments. Gene programs are rendered on UMAP based on cell expression intensity. (D) Combinatorial response of selected dosage series and most impacted myeloid programs. The mechanism of coherent superposition is consistent across drug dosage combinations. (E) On the P22 M2 macrophage program, coherent superposition produced a response function close to the sum of the sigmoid-shaped single-drug dosage responses. (F) Schematic diagram of incoherent dominance and the combination perturbation response of Tacrolimus and Motolimod. The two drugs have an incoherent effect on gene programs and have a low coherence score (0.01). The two drugs act on different gene programs and the combination response is dominated by different drugs (G) On myeloid and B-cells, the response is dominated by Motolimod. On T-cells, some programs are dominated by Tacrolimus and some are dominated by Motolimod. (H) The gene-program dominance of drug effects produced novel cell states, with inhibited T-cell and activated macrophage states induced by Motolimod. (I) Across dosage combinations strengths, Motolimod dominates over Tacrolimus on the key myeloid programs. (J) On the P22 Monocyte program, the activation effect of Tacrolimus is eliminated with increasing Motolimod dosage. On the P27 Pathogen response program, Tacrolimus dosage has no impact on the activation effect of Motolimod.

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 state of immune cell populations. However, the design and analysis of drug combinations are challenging at the transcriptome scale due to a large number of potential drug-gene interactions. D-SPIN models provide a framework to compare the action of individual drugs on the underlying gene regulatory network in order to identify drug combinations with potentially useful therapeutic applications. To analyze the drug response classes identified by D-SPIN to identify potentially useful drug combinations, we profiled these combinations experimentally and applied D-SPIN to interpret the mechanisms of combination action (SI Fig. 6A).

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 induced weaker suppression, allowing cells to activate inflammatory cytokines, including CXCL10 and CXCL11 in the P25 compliment/chemokine program whereas these cytokine programs are strongly suppressed by strong inhibitors, including the Src family inhibitor Dasatinib. GCs uniquely induced 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 network impact of the GC Halcinonide with the inhibitor Dasatinib (Fig. 7B). Gene-expression programs P22 M2 Macrophage, P25 Complement/chemokine, and P21 Monocyte are induced or repressed more strongly by one drug than the other.

Therefore, we investigate 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 glucocorticoids in combination with strong inhibitors (including Dasatinib) across a range of doses (SI Fig. 6A, B). We use D-SPIN to analyze the impact of the drug combination on the gene regulatory network and underlying cell population. We find that the drug combination induces both the strong suppression of Dasatinib as well as the induction of macrophage gene-expression programs for specific relative concentrations (Fig. 7B-D). 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 for Dasatinib or Halcinonide alone (Fig. 7E), while also exhibiting increased

suppression of macrophage activation over the single drugs. The interaction vector also maintained strong activation of the resting program P21 Monocyte.

The combination of Dasatinib and Halcinonide, therefore, generated a hyper-repressed cell population with a T-cell population resembling resting T-cells (Fig. 7C) and a macrophage cell population with activation of both resting myeloid program P21 and M2 macrophage program P22. The modulation is monotonic with dose, and manipulation of the drug doses allowed conversion of the cell population from activated to intermediate suppression to hyper-suppression M2 macrophage-like states (Fig. 7D, SI Fig. 6D, SI Text Sec. 1.16). 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 where the drug combination recruits a super-set of gene-expression programs in a dose-dependent fashion that is also observed in a previous study on single-protein level [61].

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 are suppressed in transplant patients because T-cells can react adversely with transplanted organs, however, infection then becomes a significant cause of mortality. Thus, induction of an innate immune response in the background of immuno-suppression provides a strategy for combating infection [62, 63, 64]. 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 molecules, including Motolimod, as well as standard immunosuppressive small molecules, including Tacrolimus, which are revealed to be mostly incoherent in their impact on the underlying gene regulatory network (Fig 7F). 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 (Fig 7G).

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 (SI Fig. 6C). We identify a concentration regime (Fig 7H, SI Fig. 6E) in which the cell population contains activated macrophages that exhibit strong induction of the pathogen response gene-expression program while maintaining T-cell repression, with higher than 1000nM Motolimod and higher than 100nM Tacrolimus (Fig 7I). D-SPIN infers that the specific combination of concentrations enables the activation of the pathogen response macrophage program while maintaining the activation of resting T-cell gene-expression programs. 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.

Discussion

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 data sets 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 proteins. D-SPIN models can be treated both as network models that are analyzed using the tools of graph theory as well as fully 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 data sets 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 matter physics [65]. The framework provides a low-dimensional, compact yet predictive framework for analyzing population-level effects in single-cell data through a regulatory network model.

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 aging 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 have been limited by the limited number of network reconstructions that exist. D-SPIN provides a framework in which 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 can 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 network where the cell, for example, upregulates protein degradation programs and downregulates translation when the ribosome is perturbed. D-SPIN models provide a global wiring diagram of cellular regulation that can provide insight into the higher-order logic of sub-process regulation. D-SPIN models can identify potential points of feedback control and long-range information flow in cells facilitating the study of control strategies used to maintain homeostasis.

D-SPIN can also be used to design interventions that modulate networks and cell states. We show 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 (Fig. 2). D-SPIN also reveals how small molecules and small-molecule combinations can generate novel cell states and population structures (Figs 5-7), suggesting D-SPIN could be used to study a broader set of interventions in stem cells and immune cell populations. In both cases, a major goal is to control the distribution of cell states in a cell population, for example, generating a specific set of progeny from stem cells or modulating the state of the immune system in cancer or auto-immunity. D-SPIN can reveal non-intuitive connections between different cellular processes, revealing how drugs can recruit a set of gene programs to generate a cell state. This principle has been studied on a small scale for a small number of drugs [61], and 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 super-position to generate novel cell states. In both cases, drugs act to recruit super-sets of gene-expression programs. Super-position refers to the ability of drug combinations to function additively, resulting in the union of the gene-expression programs activated by single drugs. D-SPIN allows us to dissect these combinatorial mechanisms and interpret them through drug response vectors. Drug-combination responses at the gene program level provide a means for interpreting and predicting the effect of drug combinations. If additivity extends beyond our studied system, our model could facilitate the programming of target states in cell populations. Even though we have shown additivity alone is not able to accurately predict cell-population structure, 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 conditions where additivity holds or breaks down [61].

D-SPIN has a number of limits. First, the model D-SPIN generates is an equilibrium model and so does not account for network dynamics. In future iterations of D-SPIN, we aim to convert the network into a dynamical systems model that will allow a cell to move between different steady-states. Second, D-SPIN models produce undirected rather than directed graphs, because this allows natural mapping to cells that contain feedback regulation and other forms of cyclic regulation, which are not typically allowed in directed acyclic graph modeling frameworks. Further, undirected graphical models have a powerful inference procedure which we have also extended to consider situations like active learning and perturbation design [43]. Future versions of D-SPIN might relax the undirected framework to consider directed connections between gene programs.

D-SPIN's connection with spin-network models provides some 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 models the cell population as a collection of points moving on an energy landscape that is being titled through perturbation vectors. Theoretically, energy-landscape models represent a highly simplified class of dynamical systems, because the behavior of the system can be captured within a single energy (potential) function. Maximum entropy methods have applied equilibrium spin network

models to study a much broader range of systems that are far from thermal equilibrium, including neural networks and bird flocks. However, it remains unclear why equilibrium models can provide such significant predictive power for strongly non-equilibrium systems. The strategy of using an equilibrium model to produce low-error probabilistic reconstructions of the gene expression states of a cell population suggests that cells, in certain situations, might be effectively modeled as equilibrium systems that are being driven through a series of different configurations by a biasing drive. Such models have been demonstrated for cell-fate regulation and might represent a simplifying principle [10]. The ability to model a cell as an equilibrium system that is being driven through states represents a powerful simplification that has also been studied in other chemical systems and might provide a path to more global theories of gene regulation [66, 67].

Data and code availability

Gene counts and metadata of drug profiling experiments are available at Caltech Research Data Repository https://doi.org/10.22002/2cjss-wgh69

MATLAB, Python implementations and Jupyter Notebook demonstrations of D-SPIN are available on GitHub https://github.com/JialongJiang/DSPIN

Acknowledgements

We would like to acknowledge the NIH (TR01 GM150125, R01HD100039), the Heritage Medical Research Institute, the Shurl and Kay Curci Foundation, the Merkin Institute for Translational Research, and Amgen. The single-cell profiling experiments were performed at the Beckman Institute Single-cell Profiling and Engineering Center (SPEC). Sequencing was performed at the UCSF CAT, supported by UCSF PBBR, RRP IMIA, and NIH 1S10OD028511-01 grants. We acknowledge Jost Vielmetter for the use of the liquid-handling robot. We acknowledge Dr. Guy Riddihough for editorial assistance with the manuscript. DAS acknowledges the support of an NSERC Discovery Grant and a Tier-II Canada Research Chair. The authors acknowledge Rob Phillips and Venkat Chandrasekaran for insightful scientific discussion. JH acknowledges the support of NIH 5R01GM135337 grants.

Declaration of Interests

The work was supported in part by funds from a Caltech Amgen Research Collaboration Award and reagent gifts from 10x Genomics. MWT received funding from Adaptive Biotechnologies for unrelated work. MWT is a member of the advisory board of Cell Systems. MWT is a co-founder of CognitiveAI and YurtsAI. ZJG is a co-founder of Scribe Biosciences. The Regents of the University of California with ZJG and CSM as inventors have filed patent applications related to MULTI-seq.

References

[1] Tabula Sapiens Consortium*, Robert C Jones, Jim Karkanias, Mark A Krasnow, Angela Oliveira Pisco, Stephen R Quake, Julia Salzman, Nir Yosef, Bryan Bulthaup, Phillip Brown, et al. The tabula sapiens: A multiple-organ, single-cell transcriptomic atlas of humans. *Science*, 376(6594):eabl4896, 2022.

[2] Eric H Davidson. *The regulatory genome: gene regulatory networks in development and evolution*. Elsevier, 2010.

[3] Dennis Bray. Protein molecules as computational elements in living cells. *Nature*, 376(6538):307–312, 1995.

[4] Aviv Regev and Ehud Shapiro. The π-calculus as an abstraction for biomolecular systems. In *Modelling in Molecular Biology*, pages 219–266. Springer, 2004.

[5] Aviv Regev and Ehud Shapiro. Cellular abstractions: Cells as computation. *Nature*, 419(6905):343–343, 2002.

[6] Eric H Davidson. Emerging properties of animal gene regulatory networks. *Nature*, 468(7326):911–920, 2010.

[7] Eric H Davidson and Douglas H Erwin. Gene regulatory networks and the evolution of animal body plans. *Science*, 311(5762):796–800, 2006.

[8] Lacramioara Bintu, Nicolas E Buchler, Hernan G Garcia, Ulrich Gerland, Terence Hwa, Jané Kondev, and Rob Phillips. Transcriptional regulation by the numbers: models. *Current opinion in genetics & development*, 15(2):116–124, 2005.

[9] David A Sivak and Matt Thomson. Environmental statistics and optimal regulation. *PLoS computational biology*, 10(9):e1003826, 2014.

[10] Cameron Sokolik, Yanxia Liu, David Bauer, Jade McPherson, Michael Broeker, Graham Heimberg, Lei S Qi, David A Sivak, and Matt Thomson. Transcription factor competition allows embryonic stem cells to distinguish authentic signals from noise. *Cell systems*, 1(2):117–129, 2015.

[11] Hao Yuan Kueh, Ameya Champhekar, Stephen L Nutt, Michael B Elowitz, and Ellen V Rothenberg. Positive feedback between pu. 1 and the cell cycle controls myeloid differentiation. *Science*, 341(6146):670–673, 2013.

[12] Nir Yosef, Alex K Shalek, Jellert T Gaublomme, Hulin Jin, Youjin Lee, Amit Awasthi, Chuan Wu, Katarzyna Karwacz, Sheng Xiao, Marsela Jorgolli, et al. Dynamic regulatory network controlling th17 cell differentiation. *Nature*, 496(7446):461–468, 2013.

[13] Jacques Monod. The growth of bacterial cultures. *Annual review of microbiology*, 3(1):371–394, 1949.

[14] Arthur B Pardee, François Jacob, and Jacques Monod. The genetic control and cytoplasmic expression of "inducibility" in the synthesis of β-galactosidase by e. coli. *Journal of Molecular Biology*, 1(2):165–178, 1959.

[15] Shai S Shen-Orr, Ron Milo, Shmoolik Mangan, and Uri Alon. Network motifs in the transcriptional regulation network of escherichia coli. *Nature genetics*, 31(1):64–68, 2002.

[16] Ron Milo, Shai Shen-Orr, Shalev Itzkovitz, Nadav Kashtan, Dmitri Chklovskii, and Uri Alon. Network motifs: simple building blocks of complex networks. *Science*, 298(5594):824–827, 2002.

[17] Günter P Wagner, Mihaela Pavlicev, and James M Cheverud. The road to modularity. *Nature Reviews Genetics*, 8(12):921–931, 2007.

[18] Tong Ihn Lee, Nicola J Rinaldi, François Robert, Duncan T Odom, Ziv Bar-Joseph, Georg K Gerber, Nancy M Hannett, Christopher T Harbison, Craig M Thompson, Itamar Simon, et al. Transcriptional regulatory networks in saccharomyces cerevisiae. *science*, 298(5594):799–804, 2002.

[19] Denes Hnisz, Brian J Abraham, Tong Ihn Lee, Ashley Lau, Violaine Saint-André, Alla A Sigova, Heather A Hoke, and Richard A Young. Super-enhancers in the control of cell identity and disease. *Cell*, 155(4):934–947, 2013.

[20] William T Ireland, Suzannah M Beeler, Emanuel Flores-Bautista, Nicholas S McCarty, Tom Röschinger, Nathan M Belliveau, Michael J Sweredoski, Annie Moradian, Justin B Kinney, and Rob Phillips. Deciphering the regulatory genome of escherichia coli, one hundred promoters at a time. *Elife*, 9:e55308, 2020.

[21] Christopher T Harbison, D Benjamin Gordon, Tong Ihn Lee, Nicola J Rinaldi, Kenzie D Macisaac, Timothy W Danford, Nancy M Hannett, Jean-Bosco Tagne, David B Reynolds, Jane Yoo, et al. Transcriptional regulatory code of a eukaryotic genome. *Nature*, 431(7004):99–104, 2004.

[22] Sebastian J Maerkl and Stephen R Quake. A systems approach to measuring the binding energy landscapes of transcription factors. *Science*, 315(5809):233–237, 2007.

[23] Zhanzhi Hu, Patrick J Killion, and Vishwanath R Iyer. Genetic reconstruction of a functional transcriptional regulatory network. *Nature genetics*, 39(5):683–687, 2007.

[24] Linda S Huang and Paul W Sternberg. Genetic dissection of developmental pathways. *Methods in cell biology*, 48:97–122, 1995.

[25] Edwin L Ferguson, Paul W Sternberg, and H Robert Horvitz. A genetic pathway for the specification of the vulval cell lineages of caenorhabditis elegans. *Nature*, 326(6110):259–267, 1987.

[26] Chiou-Hwa Yuh, Hamid Bolouri, and Eric H Davidson. Genomic cis-regulatory logic: experimental and computational analysis of a sea urchin gene. *Science*, 279(5358):1896–1902, 1998.

[27] Daniel Schraivogel, Andreas R Gschwind, Jennifer H Milbank, Daniel R Leonce, Petra Jakob, Lukas Mathur, Jan O Korbel, Christoph A Merten, Lars Velten, and Lars M Steinmetz. Targeted perturb-seq enables genome-scale genetic screens in single cells. *Nature methods*, 17(6):629–635, 2020.

[28] Joseph M Replogle, Reuben A Saunders, Angela N Pogson, Jeffrey A Hussmann, Alexander Lenail, Alina Guna, Lauren Mascibroda, Eric J Wagner, Karen Adelman, Gila Lithwick-Yanai, et al. Mapping information-rich genotype-phenotype landscapes with genome-scale perturb-seq. *Cell*, 2022.

[29] Christopher S McGinnis, David M Patterson, Juliane Winkler, Daniel N Conrad, Marco Y Hein, Vasudha Srivastava, Jennifer L Hu, Lyndsay M Murrow, Jonathan S Weissman, Zena Werb, et al. Multi-seq: sample multiplexing for single-cell rna sequencing using lipid-tagged indices. *Nature methods*, 16(7):619–626, 2019.

[30] Jase Gehring, Jong Hwee Park, Sisi Chen, Matthew Thomson, and Lior Pachter. Highly multiplexed single-cell rna-seq by dna oligonucleotide tagging of cellular proteins. *Nature Biotechnology*, 38(1):35–38, 2020.

[31] Elad Schneidman, Michael J Berry, Ronen Segev, and William Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007–1012, 2006.

[32] William Bialek, Andrea Cavagna, Irene Giardina, Thierry Mora, Edmondo Silvestri, Massimiliano Viale, and Aleksandra M Walczak. Statistical mechanics for natural flocks of birds. *Proceedings of the National Academy of Sciences*, 109(13):4786–4791, 2012.

[33] John J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558, 1982.

[34] Assen Roguev, Sourav Bandyopadhyay, Martin Zofall, Ke Zhang, Tamas Fischer, Sean R Collins, Hongjing Qu, Michael Shales, Han-Oh Park, Jacqueline Hayles, et al. Conservation and rewiring of functional modules revealed by an epistasis map in fission yeast. *science*, 322(5900):405–410, 2008.

[35] Eran Segal, Michael Shapira, Aviv Regev, Dana Pe'er, David Botstein, Daphne Koller, and Nir Friedman. Module networks: identifying regulatory modules and their condition-specific regulators from gene expression data. *Nature genetics*, 34(2):166–176, 2003.

[36] Xiaoqiao Chen, Sisi Chen, and Matt Thomson. Minimal gene set discovery in single-cell mrna-seq datasets with activesvm. *Nature Computational Science*, 2(6):387–398, 2022.

[37] Graham Heimberg, Rajat Bhatnagar, Hana El-Samad, and Matt Thomson. Low dimensionality in gene expression data enables the accurate extraction of transcriptional programs from shallow sequencing. *Cell systems*, 2(4):239–250, 2016.

[38] Sisi Chen, Paul Rivaud, Jong H Park, Tiffany Tsou, Emeric Charles, John R Haliburton, Flavia Pichiorri, and Matt Thomson. Dissecting heterogeneous cell populations across drug and disease conditions with popalign. *Proceedings of the National Academy of Sciences*, 117(46):28784–28794, 2020.

[39] Grace Hui Ting Yeo, Sachit D Saksena, and David K Gifford. Generative modeling of single-cell time series with prescient enables prediction of cell trajectories with interventions. *Nature communications*, 12(1):1–12, 2021.

[40] Caleb Weinreb, Samuel Wolock, Betsabeh K Tusi, Merav Socolovsky, and Allon M Klein. Fundamental limits on dynamic inference from single-cell snapshots. *Proceedings of the National Academy of Sciences*, 115(10):E2467–E2476, 2018.

[41] Aditya Pratapa, Amogh P Jalihal, Jeffrey N Law, Aditya Bharadwaj, and TM Murali. Benchmarking algorithms for gene regulatory network inference from single-cell transcriptomic data. *Nature methods*, 17(2):147–154, 2020.

[42] Jan Krumsiek, Carsten Marr, Timm Schroeder, and Fabian J Theis. Hierarchical differentiation of myeloid progenitors is encoded in the transcription factor network. *PloS one*, 6(8):e22649, 2011.

[43] Jialong Jiang, David A Sivak, and Matt Thomson. Active learning of spin network models. *arXiv preprint arXiv:1903.10474*, 2019.

[44] Nora Heisterkamp, John R Stephenson, John Groffen, Pamela F Hansen, Annelies de Klein, Claus R Bartram, and Gerard Grosveld. Localization of the c-abl oncogene adjacent to a translocation break point in chronic myelocytic leukaemia. *Nature*, 306(5940):239–242, 1983.

[45] CB LOZZIO. Human chronic myelogenous leukemia cell-line with positive philadelphia chromosome. *Blood*, 45:321–334, 1975.

[46] Brad T Sherman, Ming Hao, Ju Qiu, Xiaoli Jiao, Michael W Baseler, H Clifford Lane, Tomozumi Imamichi, and Weizhong Chang. David: a web server for functional enrichment analysis and functional annotation of gene lists (2021 update). *Nucleic Acids Res*, 10, 2022.

[47] Maxim V Kuleshov, Matthew R Jones, Andrew D Rouillard, Nicolas F Fernandez, Qiaonan Duan, Zichen Wang, Simon Koplev, Sherry L Jenkins, Kathleen M Jagodnik, Alexander Lachmann, et al. Enrichr: a comprehensive gene set enrichment analysis web server 2016 update. *Nucleic acids research*, 44(W1):W90–W97, 2016.

[48] Vincent A Traag, Ludo Waltman, and Nees Jan Van Eck. From louvain to leiden: guaranteeing well-connected communities. *Scientific reports*, 9(1):5233, 2019.

[49] Hyo-Ji Lee, Yunseo Woo, Tae-Wook Hahn, Young Mee Jung, and Yu-Jin Jung. Formation and maturation of the phagosome: a key mechanism in innate immunity against intracellular bacterial infection. *Microorganisms*, 8(9):1298, 2020.

[50] Madalina Giurgiu, Julian Reinhard, Barbara Brauner, Irmtraud Dunger-Kaltenbach, Gisela Fobo, Goar Frishman, Corinna Montrone, and Andreas Ruepp. Corum: the comprehensive resource of mammalian protein complexes—2019. *Nucleic acids research*, 47(D1):D559–D563, 2019.

[51] Maria Rosaria Amoroso, Danilo Swann Matassa, Gabriella Laudiero, Anastasia V Egorova, Roman S Polishchuk, Francesca Maddalena, Annamaria Piscazzi, Simona Paladino, Daniela Sarnataro, Corrado Garbi, et al. Trap1 and the proteasome regulatory particle tbp7/rpt3 interact in the endoplasmic reticulum and control cellular ubiquitination of specific mitochondrial proteins. *Cell Death & Differentiation*, 19(4):592–604, 2012.

[52] Melanie de Almeida, Matthias Hinterndorfer, Hanna Brunner, Irina Grishkovskaya, Kashish Singh, Alexander Schleiffer, Julian Jude, Sumit Deswal, Robert Kalis, Milica Vunjak, et al. Akirin2 controls the nuclear import of proteasomes in vertebrates. *Nature*, 599(7885):491–496, 2021.

[53] Chiara Milanese, Cíntia R Bombardieri, Sara Sepe, Sander Barnhoorn, César Payán-Goméz, Donatella Caruso, Matteo Audano, Silvia Pedretti, Wilbert P Vermeij, Renata MC Brandt, et al. Dna damage and transcription stress cause atp-mediated redesign of metabolism and potentiation of anti-oxidant buffering. *Nature communications*, 10(1):4887, 2019.

[54] Yvan Jamilloux, Thomas El Jammal, Lucine Vuitton, Mathieu Gerfaud-Valentin, Sébastien Kerever, and Pascal Sève. Jak inhibitors for the treatment of autoimmune and inflammatory diseases. *Autoimmunity reviews*, 18(11):102390, 2019.

[55] David C Fajgenbaum and Carl H June. Cytokine storm. *New England Journal of Medicine*, 383(23):2255–2273, 2020.

[56] Christopher S McGinnis, David M Patterson, Juliane Winkler, Daniel N Conrad, Marco Y Hein, Vasudha Srivastava, Jennifer L Hu, Lyndsay M Murrow, Jonathan S Weissman, Zena Werb, Eric D Chow, and Zev J Gartner. MULTI-seq: sample multiplexing for single-cell RNA sequencing using lipid-tagged indices. *Nat. Methods*, 16(7):619–626, July 2019.

[57] C Domínguez Conde, C Xu, LB Jarvis, DB Rainbow, SB Wells, T Gomes, SK Howlett, O Suchanek, K Polanski, HW King, et al. Cross-tissue immune cell analysis reveals tissue-specific features in humans. *Science*, 376(6594):eabl5197, 2022.

[58] Ctibor Skuta, Martin Popr, Tomas Muller, Jindrich Jindrich, Michal Kahle, David Sedlak, Daniel Svozil, and Petr Bartunek. Probes & drugs portal: an interactive, open data resource for chemical biology. *Nature methods*, 14(8):759–760, 2017.

[59] Thibaut Desgeorges, Giorgio Caratti, Rémi Mounier, Jan Tuckermann, and Bénédicte Chazaud. Glucocorticoids shape macrophage phenotype for tissue repair. *Frontiers in immunology*, 10:1591, 2019.

[60] Yan-Cun Liu, Xian-Biao Zou, Yan-Fen Chai, and Yong-Ming Yao. Macrophage polarization in inflammatory diseases. *International journal of biological sciences*, 10(5):520, 2014.

[61] Naama Geva-Zatorsky, Erez Dekel, Ariel A Cohen, Tamar Danon, Lydia Cohen, and Uri Alon. Protein dynamics in drug combinations: a linear superposition of individual-drug responses. *Cell*, 140(5):643–651, 2010.

[62] JA Fishman. Infection in organ transplantation. *American Journal of Transplantation*, 17(4):856–879, 2017.

[63] Jay A Fishman and Robert H Rubin. Infection in organ-transplant recipients. *New England Journal of Medicine*, 338(24):1741–1751, 1998.

[64] Marcus R Pereira, Sumit Mohan, David J Cohen, Syed A Husain, Geoffrey K Dube, Lloyd E Ratner, Selim Arcasoy, Meghan M Aversa, Luke J Benvenuto, Darshana M Dadhania, et al. Covid-19 in solid organ transplant recipients: initial report from the us epicenter. *American journal of transplantation*, 20(7):1800–1808, 2020.

[65] Edwin T Jaynes. Information theory and statistical mechanics. *Physical review*, 106(4):620, 1957.

[66] Lacramioara Bintu, Nicolas E Buchler, Hernan G Garcia, Ulrich Gerland, Terence Hwa, Jane Kondev, Thomas Kuhlman, and Rob Phillips. Transcriptional regulation by the numbers: applications. *Current opinion in genetics & development*, 15(2):125–135, 2005.

[67] Jeremy Gunawardena. A linear framework for time-scale separation in nonlinear biochemical systems. *PloS one*, 7(5):e36321, 2012.

Supplementary material

Supplementary Figures

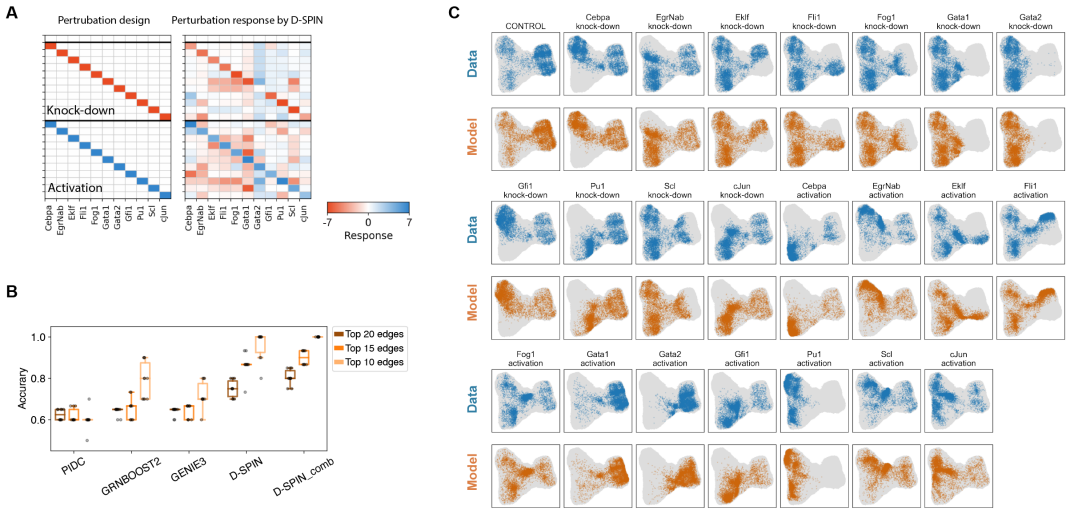

Figure S1: **Perturbation responses and predicted cell state distributions by D-SPIN; network inference of D-SPIN is further improved with double perturbations; related to Fig. 2** (A) On the single-gene knockdown and activation perturbations, the response vectors by D-SPIN agree with the actual perturbation except for the Gata2 knockdown. This is potentially due to Gata2 being a transient regulator in the network and does not express in the differentiated states. (B) D-SPIN achieves further improved accuracy with 96 more double gene perturbations. With extra double perturbations, the top 10, 15, and 20 accuracies are 1, 0.9, and 0.805. (C) The D-SPIN model predicts the cell-fate distribution generated through perturbations to the underlying network.

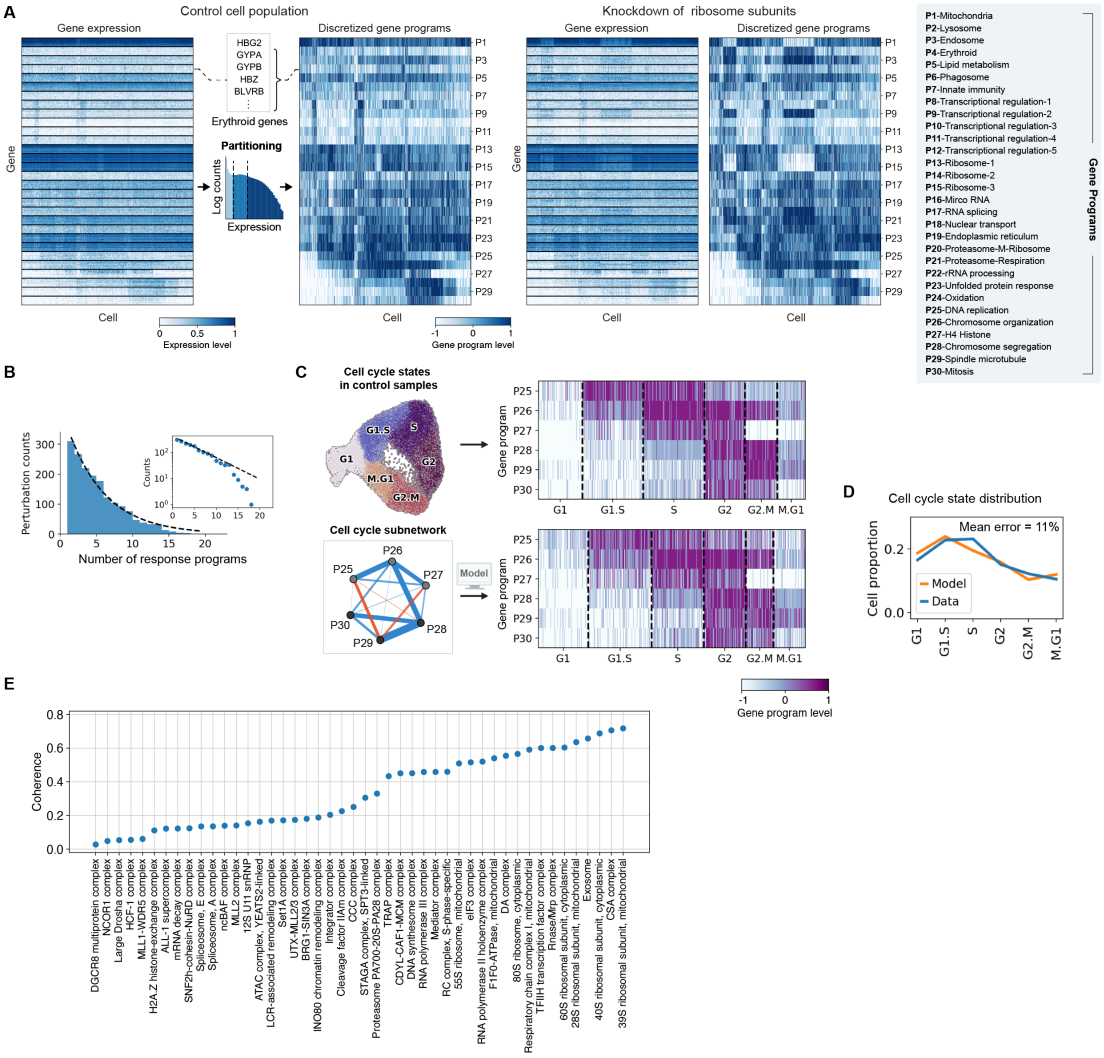

Figure S2: **Exams of gene program discretization; distribution of responding program number of perturbations; D-SPIN reconstruction of cell cycle states; Coherence of macromolecular complex; related to Fig. 3** (A) Heatmap of gene expression and discretized gene program level on two examples: control samples of non-targeting guide RNAs and knockdown of ribosome large or small subunits. The expression of each gene is normalized by its maximum expression and capped at 0.8 for visualization. (B) The number of responding gene programs of each perturbation shows an exponential distribution. Perturbations that influence a large number of perturbations are relatively rare. (C) D-SPIN reconstructs the distribution of cell cycle progression with a 6-node subnetwork taken from the inferred network. The response vector of the subnetwork is inferred by the mean gene program expression on control samples conditioned on the network. Samples from the control population and D-SPIN model samples are embedded together on UMAP, and we annotate and order the cell cycle state of each Leiden cluster by the gene program expression. (D) The cell cycle state distribution of D-SPIN is similar to data distribution with an 11% mean error. (E) Regulatory coherence of macromolecular complex subunits of the 46 filtered molecular complexes in the CORUM database.

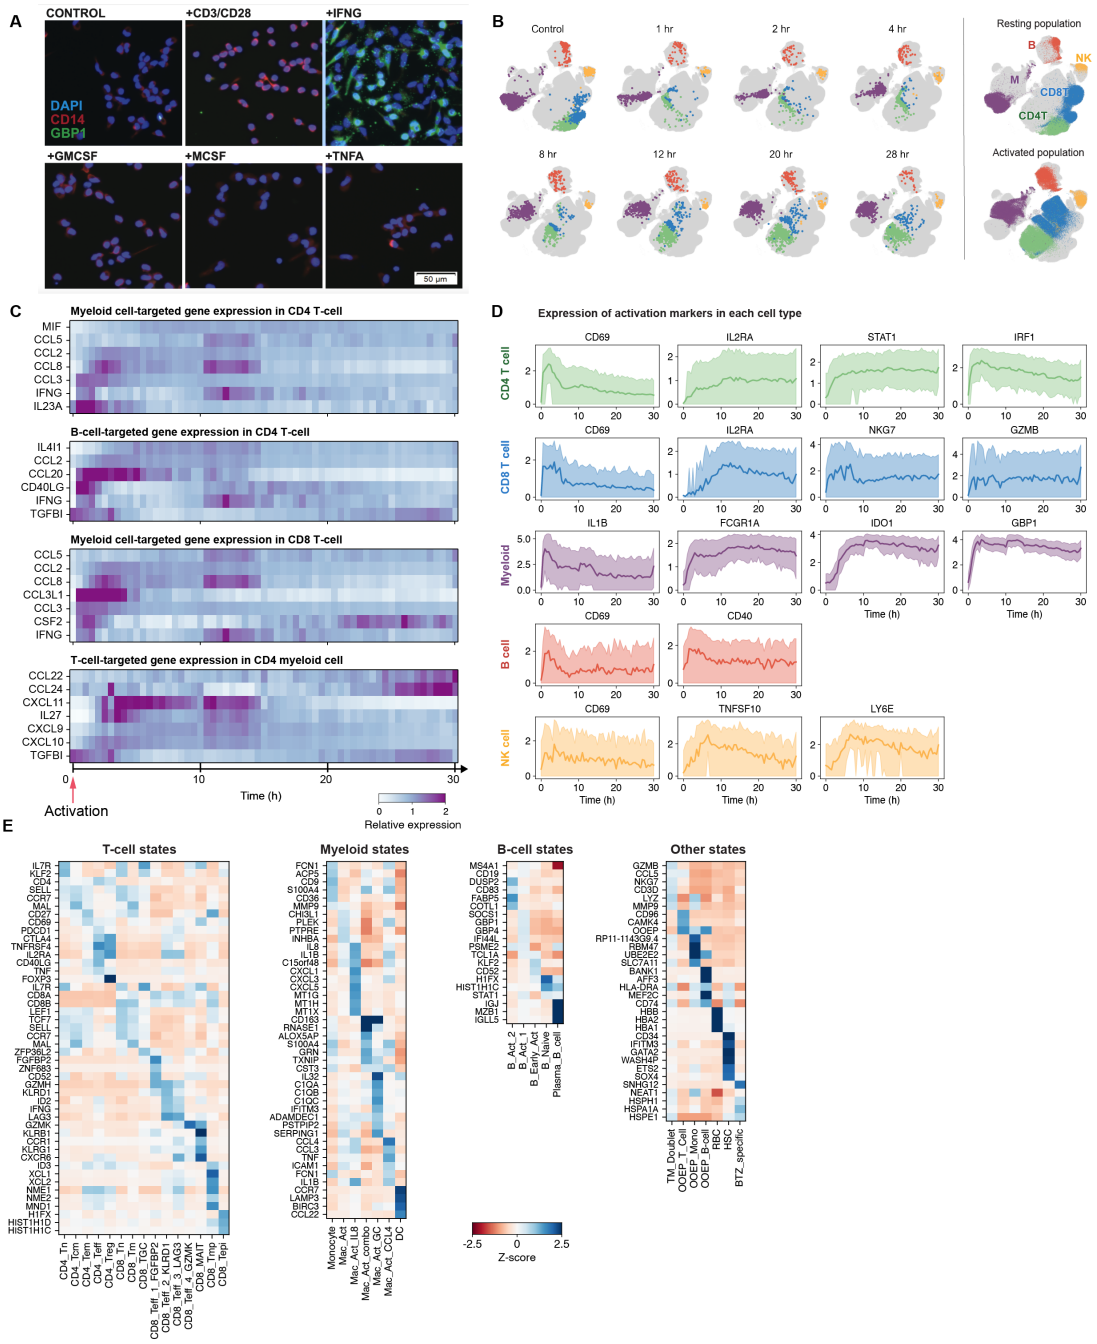

Figure S3: **Timecourse of T-cell mediated immune activation; cell typing of the drug profiling experiments; related to Fig. 5** (A) CD3/CD28 antibody does not activate monocytes. Immunostaining of primary human monocytes cultured in CD3/CD28 antibody (25uL/mL cells), IFNG (5ng/mL), GMCSF (5ng/mL), MCSF (5ng/mL), or TNFA (5ng/mL) reveals that GBP1 is only expressed in IFNG-activated macrophages. (B) UMAP rendering of the timecourse of T-cell mediated immune activation for a series of example time points. The immune population gradually moves from the resting state to the activated state. (C) The dynamics of signaling gene expression in each cell type show active communication between different cell types. Expression is normalized by the mean across the timecourse. (D) The dynamics of gene markers for immune activation in each major cell type. The color-shaded range is 10 and 90 percentile of gene expression. (E) The Z-score of cell typing gene marker expression. Gene markers are identified by differential expression analysis of each Leiden cluster. The Z-score is computed in reference to all the cells in the same major cell type.

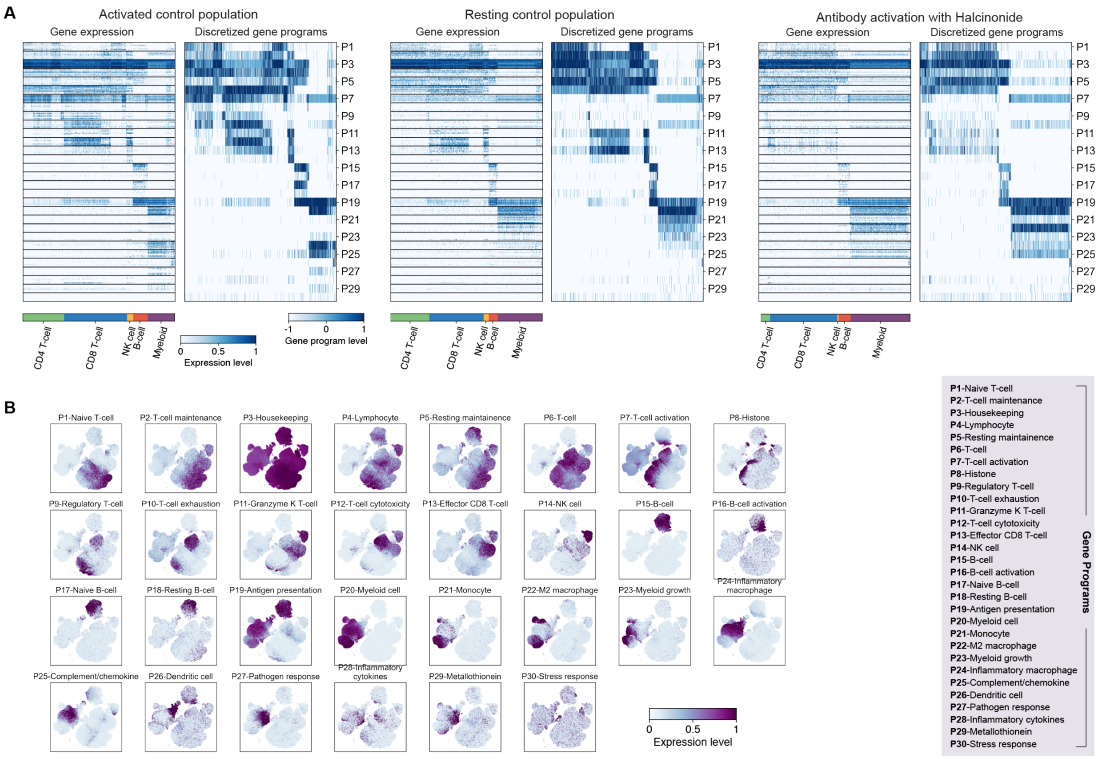

Figure S4: **Examples of gene program discretization and gene program expression on the UMAP; related to Fig. 5** (A) Heatmap of gene expression and discretized gene program level on three examples: activated control population, resting control population, and antibody activation with Halcinonide treatment. The expression of each gene is normalized by its maximum expression and capped at 0.8 for visualization. (B) The expression level of each gene program is rendered on the UMAP embedding. Each cell is colored by its expression level on the gene program. The expression level is normalized by the 95 percentile for visualization only.

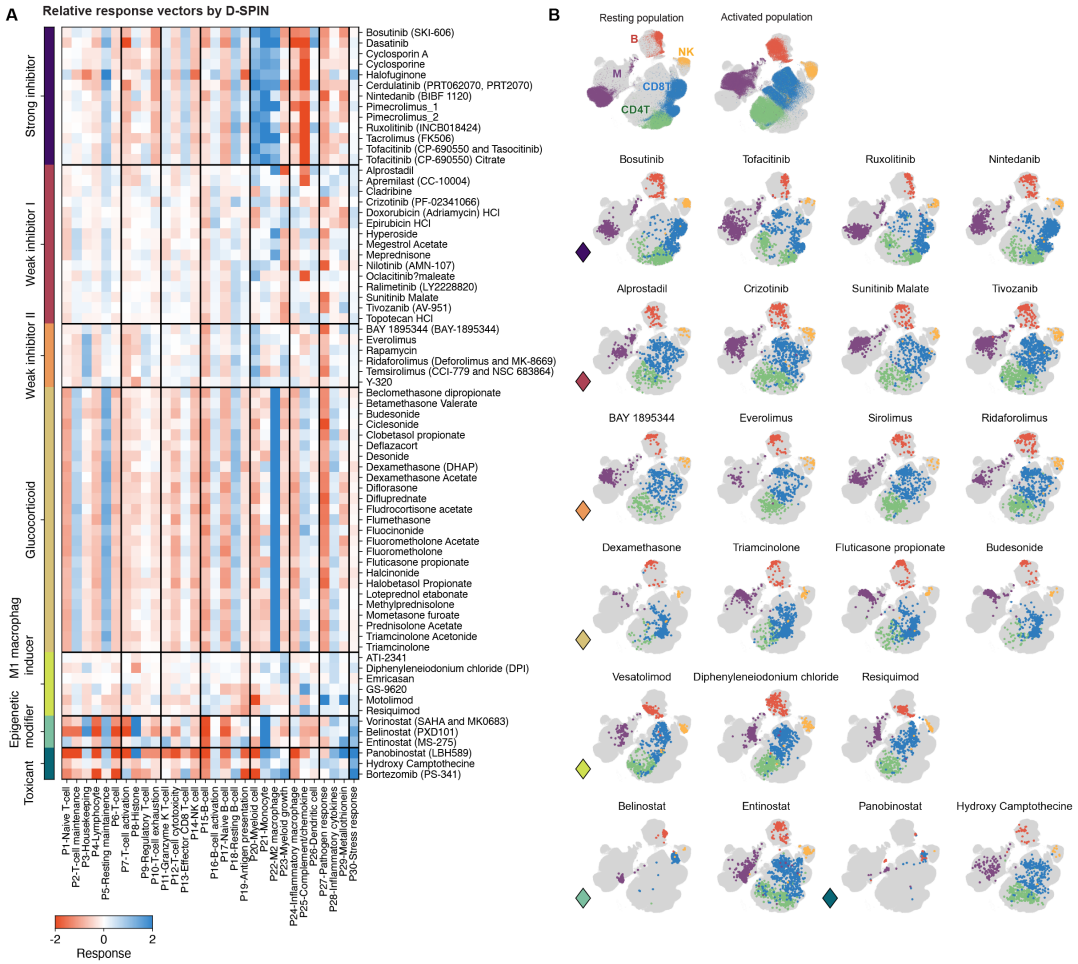

Figure S5: **D-SPIN identifies phenotypical classes of drug categories; related to Fig. 6** (A) Heatmap of relative response vector inferred by D-SPIN. D-SPIN identifies 71 effective drugs on the immune population and groups the drugs into 7 phenotypical classes. (B) UMAP embedding of example drugs from each class. Drug categories exhibit different signatures of immune cell population change.

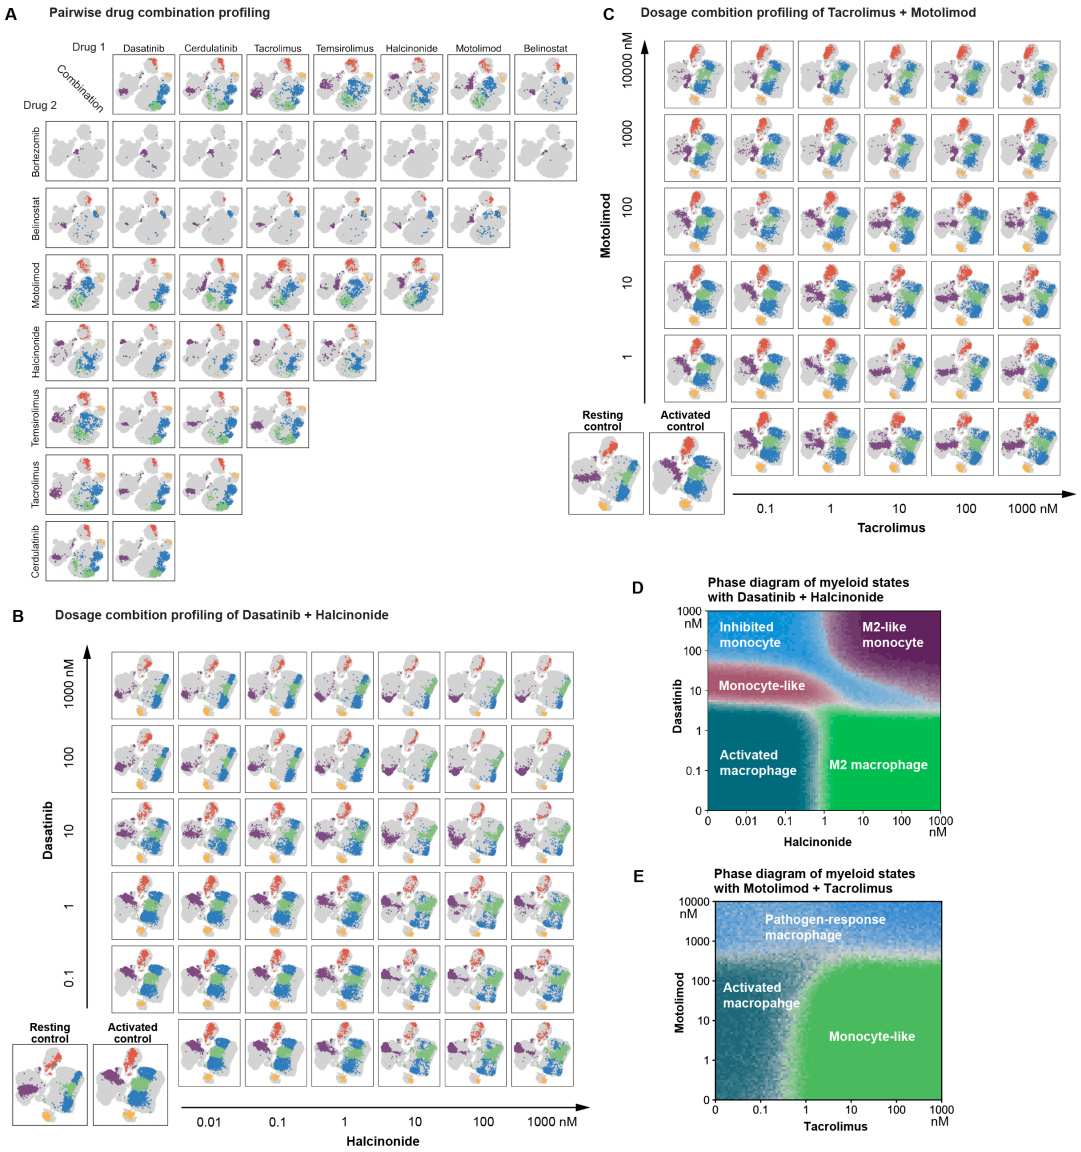

Figure S6: **Drug combination profiling discovers novel cell states; dosage combination reveals phase diagrams of drug combination effect; related to Fig. 7** (A) We perform pair-wise drug combination profiling for 12 selected drugs (8 shown here) from different categories. a. Strong inhibitor: Dasatinib, Bosutinib, Cerdulatinib, Tacrolimus; b. Weak 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. (B) Dosage combination profiling of Dasatinib plus Halcinonide. The two drugs create a Monocyte-like resting myeloid population at a medium level of combination (10nM Dasatinib, 1nM Halcinonide) and a novel combinatorial myeloid state expressing both Monocyte and M2 macrophage gene programs. On T-cells, the drug combination produces stronger inhibition than single drugs. (C) Dosage combination profiling of Tacrolimus plus Motolimod. The drug combination produces Motolimod-induced pathogen response myeloid states, and intermediate T-cell states similar to both drugs. (D) D-SPIN depicts the phase diagram of drug combination effects by interpolating response vectors of profiled dosages. At a low dosage of Dasatinib, the myeloid transit from the activated macrophage state to the M2 macrophage state as the increase of Halcinonide dosage. High dosages of Dasatinib produces Monocyte-like, inhibited monocyte, and M2-like monocyte states as the increase of Halcinonide dosage. (E) The phase diagram of myeloid states with Tacrolimus plus Motolimod shows the dominance of Motolimod-induced pathogen-response state over Tacrolimus-induced monocyte-like states.

1 Methods details

1.1 Overview of D-SPIN framework

The D-SPIN framework utilizes single-cell transcriptional profiling data from multiple conditions as input. It constructs a set of gene programs based on this data, generates a regulatory network model for these gene programs, and identifies the impact of each perturbation on the gene programs. The framework consists of three stages: gene program discovery, regulatory network model inference, and network analysis. In this section, we present general principles and delve into the implementation details in the subsequent sections.

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 of D-SPIN can include gene programs of interest and D-SPIN can infer the interactions between the gene programs and how these programs are regulated by different conditions in the context of the regulatory network. For a gene program containing the gene set $\{g_i\}$, we compute the weighted average expression level $\boldsymbol{l}$ of the program by solving the optimal ℓ_2 approximation of the gene set expression (Sec. 1.4). In our analysis, we use orthogonal non-negative matrix factorization (oNMF) for gene program identification (Sec. 1.4).

D-SPIN applies a discretized spin-network modeling framework originated 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 in a population with a minimal interaction model which can be interpreted as the steady state of a reaction system with nonlinear activations (Sec. 1.6). The number of possible states can be increased in a controlled way by changing the number of discretizations. The original model, the Ising model from statistical physics uses two discretized states $\{-1, 1\}$, while we use three discretized states $\{-1, 0, 1\}$ as it improved quantitative agreement between model and empirical cell state distribution by including two different levels of gene program expression (Sec. 1.5). In D-SPIN, each cell is represented as a state vector of gene programs $\boldsymbol{s} = [s_1, s_2, \cdots, s_M], \boldsymbol{s} \in \{-1, 0, 1\}^M$. D-SPIN identifies a probabilistic graphical model of the observed distribution of cell states that can be interpreted as a regulatory network.

As further explored in another work [1], the improved network inference of D-SPIN is empowered by the ability to integrate information across multiple conditions that capture different aspects of the 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 balancing the conditions used for inference by taking a subset of conditions that have different cell state distributions from each other (Sec. 1.7).

The regulatory network and response vectors inferred by D-SPIN can be analyzed to define the underlying principles of the gene regulatory network and to identify classes of perturbations with similar responses. Biological networks commonly exhibit a modular organization where the network contains a set of strongly interacting sub-networks that each is associated with a specific biological function. We analyze the regulatory network with Leiden community detection to identify these sub-networks (Sec. 1.10). The response vectors enable clustering of perturbation response across multiple experimental batches with a simple batch effect correction scheme. Due to the linear form of response vector $\boldsymbol{h}$ in the model, subtracting the response vector of the control sample produces the up or down-regulation of each gene program relative to the control. Response vectors can be clustered and embedded as feature vectors of perturbations, or grouped by graph clustering with the network to reveal their impact on the regulatory network (Sec. 1.11).

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

Perturb-seq introduces a library of guide RNAs by lentiviral transduction to engineered cells that express the dCAS9 CRISPRi effector, and simultaneously readout the transcription profile and guide RNA for each cell using single-cell RNA sequencing [2]. Due to the stochastic nature of guide RNA transfection, the number of cells receiving each perturbation vastly varies. In the Perturb-seq, the perturbation with the smallest sample size only had 2 cells collected, while the perturbation with the largest sample size had close to 2,000 cells collected. To reduce noise introduced by the small sample size, we exclude 17 perturbations with less than 20 cells collected (metadata entry "number of cells

(filtered)"), which is 0.16% of all conditions.

The Perturb-seq dataset is a genome-wide screening study where every single gene is knockdown individually. More than half of the perturbations do not yield significant single-gene-level expression changes except the perturbed gene, possibly because the genome contains redundancies to ensure the robustness of cellular activities, and 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 7,523 perturbations with less than 10 differentially expressed genes detected (metadata entry "Number of DEGs (anderson-darling)"), which is 70.5% of all perturbations. Similarly, we only select a subset of the 585 control samples with nontargeting guide RNAs to avoid overfitting on measurement noise. We select control samples that are identified as core control (metadata entry "core_control") and have no less than 200 collected cells (metadata entry "number of cells (filtered)"). In total, we select 3,138 perturbation conditions and 105 control conditions.

Single-cell gene expression data need normalization as (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 (2) across genes the transcript expression level can span 5 orders of magnitude. We normalize the raw counts by each cell by the total counts detected in each cell, and times a scaling factor that is roughly the median of the total counts of transcripts in a cell. We use 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 use a log transformation to mitigate the great variety of expression levels across genes $g_i'' = \log(g_i' + 1)$. Except for gene filtering, all following analysis is performed on the log-transformed data.

Due to measurement noise and gene drop-out, genes that have constant expression across cells significantly increase the noise in the data. Therefore single-cell single-cell data analyses are performed on a subset of genes that have different expression levels across cells, i.e., highly-variable genes. We develop a highly-variable gene filtering model adapted from the coefficient-of-variation (CV) filtering. We use the zero-inflated Poisson distribution as the statistics of gene counts because we observed that genes with high expression generally have higher CV than a Poisson model and can be explained by a zero inflation term. For a zero-inflated Poisson distribution,

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

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

$$\mu = (1-\pi)\lambda \tag{2}$$

$$\frac{\sigma}{\mu} = \frac{\sqrt{1 + \pi\mu/(1-\pi)}}{\sqrt{\mu}} \tag{3}$$

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

$$y = -\frac{1}{2}\left[x - \log(1 + \frac{\pi}{1-\pi}\exp x)\right] \tag{4}$$

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

$$\hat{y} = -k\left(x - \log(1 + \pi \exp x)\right) + b \tag{5}$$

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

We select highly-variable genes as the union of the following three sets (1) We perform the Wilcoxon test of log-transformed data for each guide RNA and control (sc.tl.rank_genes_groups) and select Top 10 differentially expressed genes for each guide that pass $p < 0.05$ under Bonferroni multi-test correction (select less than 10 if not enough genes pass the test). (2) Top 1,000 highly-variable genes selected by our zero-inflated Poisson CV filtering. (3) For comparison with the original Perturb-seq study, we

include the genes that appeared in the gene programs in the original Perturb-seq study. In total, we select 3,273 genes for the D-SPIN framework.

To visualize and cluster single cells in the Perturb-seq dataset, we correct 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) [3]. We construct a 15-nearest-neighbor graph (scanpy.pp.neighbors) on the latent representation and compute Leiden clustering for cell clusters (scanpy.tl.leiden). As K562 is a purified cell line, it contains primarily a single cell type. The UMAP embedding of the cell population is a circular structure representing the progress of the cell cycle, which is less relevant to the study of perturbation response. Therefore we only show the UMAP embedding on control samples along with D-SPIN-generated samples using cell-cycle-associated gene programs (Sec. 1.12).

1.3 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 normalize the raw counts by each cell with the same scaling factor 10^4, and gene expressions are log-transformed for the following analyses except for the highly-variable gene selection.

The highly-variable gene selection is based on the zero-inflated Poisson expression model we develop (Sec. 1.2). 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 is performed on the Drug Profiling dataset and Dosage Combination follows the same filtering. To avoid noise in low-expression genes, we exclude genes that only have a non-zero expression in less than 0.1% of cells in the Drug Profiling dataset. (In the Perturb-seq dataset the published data low-expression genes are already filtered out by the authors). We select 3,287 highly-variable genes in the remaining genes, which is $\log 10(1.1)$ above the curve fitting Eq. 5 for the zero-inflated Poisson distribution model.

Collected single-cell profiling data have batch effects across experimental batches potentially due to the state of cell culture and different sequencing machine runs. To visualize, cluster, and cell-typing single-cells in the Drug Profiling dataset, we correct batch effect by constructing latent representations on raw gene counts of selected 5,000 genes (scanpy.pp.highly_variable_genes) with scVI (dispersion='gene-batch', n_latent=10, n_layers=2, max_epoch=200, batch_size=5120) [3]. We construct a 15-nearest-neighbor graph (scanpy.pp.neighbors) on the latent representation, compute Leiden clustering for cell clusters (scanpy.tl.leiden), and UMAP embedding for visualization (scanpy.tl.umap). The cell subtype identities of each Leiden cluster are annotated manually by comparing up-regulated differentially expressed genes of the cluster with the remaining cells. We remove 3 clusters that mostly come from the experimental batch "MULT-6" and are putative as cell debris, as they express a maternal effect gene OOEP and show disrupted transcriptional states. We also exclude Red blood cells (RBC) and plasma B-cell for D-SPIN analysis, as they are small and isolated cell populations with distinct transcriptional profiles.

The Dosage Combination dataset follows the same normalization and log-transformation prepossessing. We visualize 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 processing.

1.4 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 with no overlapping between gene programs. Compared to typical matrix factorization methods like principal components analysis (PCA), oNMF applies two constraints to the gene programs: non-negative weights and orthogonality. The non-negative constraint avoids the ambiguity of interpreting negative weights, ensuring the programs are a set of co-expressed genes. The orthogonality constraint further enhances interpretability and aids model construction by forcing representations to be non-overlapping. Thus 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 that cell type. Further, in the context of D-SPIN, overlapping gene programs can result in confounding activating interactions between gene programs by the shared genes.

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

$$\begin{aligned} \underset{W,H}{\text{minimize}} \quad & \|X - WH\|_F^2 \\ \text{subject to} \quad & W^T W = I, \quad W_{ij} \geq 0, \quad H_{ij} \geq 0 \end{aligned} \tag{6}$$

where $W \in \mathbb{R}^{n\text{-gene} \times k}$ is the gene program representation, $H \in \mathbb{R}^{k \times n\text{-cell}}$ 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 [4], with W and H randomly initialized by uniform distribution on $[0, 1]$. W is orthogonalized using singular value decomposition (SVD) and takes absolute value before the starting of iteration and we use 500 iterations.

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 utilize the following strategies to balance the contributions from each gene and cell state:

1. Zero-inflated Poisson distribution gene CV filtering (Sec. 1.2).
   We noted that genes with high expression levels tend to have greater variance than typically expected under Poisson distribution CV filtering. Consequently, these genes often compose multiple housekeeping programs in the oNMF. Using a zero-inflated Poisson distribution for CV filtering, we account for the higher variance induced by dropouts, enabling better gene selection. We also remove genes that are expressed in a tiny fraction of cells in gene filtering. Genes with mostly 0 expressions are not informative and can introduce 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 divide the expression of each gene by its standard deviation across cells, ensuring that each gene contributes equally. 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 can also be mitigated 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 it constitutes a large proportion of the cell population. To address this, we use class-balance strategies frequently applied in machine learning practices. Specifically, we balance different cell states by subsampling cells based on cell clusters obtained by clustering or gene markers. We employ two different schemes. (1) Equal-sample balancing: to emphasize gene programs that only express in a small fraction of the cell population, we sample an equal number of cells from each cell cluster. (2) Square-root balancing: in machine learning, taking samples from each data category proportional to the square root of the category size was shown to be effective [5]. Similarly, we take samples from each cell cluster proportional to the square root of the cell number in that cluster.

In the Perturb-seq study, we use a subset of 10^5 single cells, with the second scheme, Square-root balancing for the Leiden cluster of cell states. In the drug profiling study, to better capture programs in novel cell states induced by specific drugs, we use sampling 1,000 single cells from each Leiden cluster.

The exact solution of nonnegative matrix factorization is an NP-hard problem [6]. But there are heuristic approximations of oNMF by iterative matrix update, and the solution varies with random initialization. So we run oNMF with different random seeds and use clustering to generate a consensus decomposition. In the Perturb-seq study, we run oNMF 100 times with $k = 30$ and cluster the results by k-means from the 100 runs. In the drug profiling study, as specific drug-induced cell states are a small fraction of the whole population, we increase the resolution of oNMF with $k = 50$. During preliminary runs of oNMF, we identify drug-specific response programs such as Pathogen response, Stress response. We perform manual curation of random seeding to select oNMF decomposition that

includes these target gene programs. We then combine the 50 programs into 30 by merging programs with similar biological functions and expression distribution on cell subtypes. The gene programs are annotated with a combination of bioinformatics databases including DAVID, Enrichr, and String-db [7, 8, 9], and also manual lookup.

Given a set of genes $\{g_i\}$, either from user input or the result of oNMF, we compute the weighted average expression level $\boldsymbol{l}$ of the gene set by solving the optimization problem

$$\min_{w_{g_i}, \boldsymbol{l}} \quad \sum_{g_i} \|\boldsymbol{x}_{g_i} - w_{g_i} \boldsymbol{l}\|^2 \tag{7}$$

where $\boldsymbol{x}_{g_i}$ is the expression profile of gene g_i, and w_{g_i} is the averaging weight to be determined. This is a 1-component non-negative matrix factorization (NMF) problem (sklearn.decomposition.NMF).

1.5 Formulation of D-SPIN framework

The probabilistic graphical model we used originates from the Ising model in physics, also named spin glass, spin network, and Markov random field in different disciplines. The model defines a probability distribution on a discretized state vector $\boldsymbol{s} = [s_1, s_2, \cdots, s_M], s_i \in \{-1, 0, 1\}$, where M is the number of gene or gene programs. The probability distribution defined by the model is

$$P(\boldsymbol{s}) = \frac{1}{Z} \exp[-E(\boldsymbol{s}; \boldsymbol{J}, \boldsymbol{h}^{(n)})] \tag{8}$$

$$E(\boldsymbol{s}; \boldsymbol{J}, \boldsymbol{h}^{(n)}) = -\sum_{ij} J_{ij} s_i s_j - \sum_i h_i^{(n)} s_i \tag{9}$$

$$Z = \sum_{\boldsymbol{s}} \exp[-E(\boldsymbol{s}; \boldsymbol{J}, \boldsymbol{h}^{(n)})]\,, \tag{10}$$

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 across perturbation conditions, the interaction network $\boldsymbol{J}$ stays the same and the bias vector $\boldsymbol{h}^{(n)}$ is modulated by each perturbation condition n to activate different modes in the network and produced different cellular transcription state distributions. We learn the model from data by gradient ascent of the log maximum likelihood function of the model.

$$\log \mathcal{L}(\boldsymbol{J}, \boldsymbol{h}^{(n)}) = \sum_n \sum_{\boldsymbol{s}^{(n)}_{\text{Data}}} \log P(\boldsymbol{s}^{(n)}_{\text{Data}} | \boldsymbol{J}, \boldsymbol{h}^{(n)}) \tag{11}$$

We compute the gradient of the objective function

$$\frac{\partial \log \mathcal{L}}{\partial J_i j} = \sum_n \sum_{\boldsymbol{s}^{(n)}_{\text{Data}}} \frac{\partial}{\partial J_i j} \left[-E(\boldsymbol{s}^{(n)}_{\text{Data}}) - \log Z \right] \tag{12}$$

$$= \sum_n \sum_{\boldsymbol{s}^{(n)}_{\text{Data}}} \left\{ s^{(n)}_{i\text{Data}} s^{(n)}_{j\text{Data}} - \frac{1}{Z} \sum_{\boldsymbol{s}} s_i s_j \exp[-E(\boldsymbol{s})] \right\} \tag{13}$$

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 parameter $\boldsymbol{J}, \boldsymbol{h}$. Normalize the gradient by sample number so it is independent of the sample number, and with similar derivation for $\partial \log \mathcal{L} / \partial h_i$, the gradient of the objective function is in the following form.

$$\frac{\partial \log \mathcal{L}}{\partial \boldsymbol{J}} = \sum_n \left(\langle s_i s_j \rangle^{(n)}_{\text{Data}} - \langle s_i s_j \rangle^{(n)}_{\text{Model}} \right) \tag{14}$$

$$\frac{\partial \log \mathcal{L}}{\partial \boldsymbol{h}^{(n)}} = \langle s_i \rangle^{(n)}_{\text{Data}} - \langle s_i \rangle^{(n)}_{\text{Model}} \tag{15}$$

where $\langle\cdot\rangle^{(n)}_{\text{Data}}$ is the average over the data of n-th conditions, and $\langle\cdot\rangle^{(n)}_{\text{Model}}$ is the expectation on the distribution defined by current model parameters $\boldsymbol{J}$, $\boldsymbol{h}^{(n)}$. 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 a few unique advantages

1. From a statistical perspective, the spin network is the maximum entropy model given the mean and pairwise cross-correlation of data [10]. Entropy describes how a distribution "spread out" over all possible states, and the principle of maximum entropy states that the best model describing a system is the model with maximum entropy while satisfying all required statistics. Because such a model has no extra assumption of the system apart from the known 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. So optimization techniques to avoid traps of local minimums are not necessary. Though it is worth noting that the optimization is still a difficult problem, primarily due to the computational complexity of accurate gradient estimation and large condition numbers from model identifiability issues [11]. Inferring the model across perturbation conditions can mitigate these issues, and is further explored in another work [1].

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 (Sec. 1.6)

There are a few potential choices of discretizing the expression level $\boldsymbol{l}$: the number of discretized states m and the choice between $\{-\lfloor m/2\rfloor, \ldots, 0, \lfloor m/2\rfloor\}$ or $\{0, 1, \ldots, m\}$. We choose to discretize $\boldsymbol{l}$ into 3 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.

2. we find $m = 2$ 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 bistable 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 $\{0, 1, 2\}$.

1.6 Correspondence between D-SPIN and biochemical reaction systems

D-SPIN is closely connected to the dynamical models of biochemical reactions as an approximation to 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

$$d\mathbf{X}_t = \boldsymbol{\mu}(\mathbf{X}_t, t)\, dt + \boldsymbol{\sigma}(\mathbf{X}_t, t)\, d\mathbf{W}_t \tag{16}$$

where $\boldsymbol{\mu}(\mathbf{X}_t, t)$ are the reactions and $\boldsymbol{\sigma}(\mathbf{X}_t, t)\, d\mathbf{W}_t$ is the stochastic noise. Assume the stochastic noise is a constant, and the distribution of state $P(\boldsymbol{x})$ follows the Fokker-Planck equation

$$\frac{\partial P(\mathbf{x}, t)}{\partial t} = -\sum_{i=1}^{N} \frac{\partial}{\partial x_i} \left[\mu_i(\mathbf{x}, t) P(\mathbf{x}, t)\right] + \sum_{i=1}^{N}\sum_{j=1}^{N} \frac{\partial^2}{\partial x_i\, \partial x_j} \left[\frac{\sigma^2}{2} P(\mathbf{x}, t)\right]. \tag{17}$$

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

$$P(\boldsymbol{x}) = \frac{1}{Z} \exp \frac{-2V(\boldsymbol{x})}{\sigma^2} \tag{18}$$

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

$$\mu_i(\boldsymbol{x}) = \sum_j a_{ij}x_j + b_i = -\frac{\partial}{\partial x_i}(-\sum_{ij} a_{ij}x_ix_j - \sum_i b_ix_i). \tag{19}$$

This solution is a Gaussian distribution near the stable point [12, 13]. Such models only have a single minimum in the convex potential function, therefore not able to characterize cell populations with various stable cell types and cell states.

The potential function, i.e. energy function Eq. 9 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

$$\mu_i(\boldsymbol{x}) = -\frac{\partial}{\partial x_i}(-\sum_{ij} J_{ij}s_is_j - \sum_i h_is_i) \tag{20}$$

$$= \phi(x_i)'(\sum_j J_{ij}\phi(x_j) + h_i) \tag{21}$$

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.

1.7 Computational details of D-SPIN

The choice of gene program number is a trade-off between computational complexity, interpretability, and resolution. As a graphical model, the number of all possible states grows exponentially with the number of nodes, so less than 50 total nodes are preferred for the accuracy of model estimation. The number of gene programs is typically set to $10 \sim 40$ depending on the desired resolution of gene matrix decomposition. In both Perturb-seq and drug profiling studies, we choose 30 programs for presentation and interpretability.

D-SPIN discretized the continuous expression level $\boldsymbol{l}$ on a gene program into three discrete states $\{-1, 0, 1\}$. We find the partition that minimizes the sum of variance inside each group

$$\min_{s(\boldsymbol{l})} \quad \mathrm{Var}[\boldsymbol{l}|(s(\boldsymbol{l}) = -1)] + \mathrm{Var}[\boldsymbol{l}|(s(\boldsymbol{l}) = 0)] + \mathrm{Var}[\boldsymbol{l}|(s(\boldsymbol{l}) = 1)] \tag{22}$$

which is equivalent to k-means clustering with 3 clusters. In the Perturb-seq study, as a large proportion of perturbations induce insignificant effects on the gene program level, we balance the conditions by selecting 128 samples to infer the regulatory network (Sec. 1.1). The 128 samples are selected by taking one sample from 128 k-means clusters of all samples with vectorized cross-correlations. We then learn the response vectors of all other samples independently conditioned on the inferred network. In the drug profiling study, we infer the regulatory network with all experimental conditions on the Drug Profiling dataset. At each optimization step, D-SPIN computes the gradient of $\boldsymbol{J}$ and $\boldsymbol{h}^{(n)}$ by estimating Eq. 14 and Eq. 15 with Markov Monte Carlo (MCMC) sampling (node number > 16) or explicit computing the distribution (node number ≤ 16). To make step size independent of the number of conditions used, the gradient of $\boldsymbol{J}$ is computed as the average of the gradient in each sample instead of the sum in Eq. 14. In the case of 30 gene programs, we use 5×10^7 MCMC samples in each iteration for each condition with 10^3 samples warm-up for network inference and 10^7 samples for perturbation response vector inference.

We use Adam optimizer with a learning rate of 0.02 for MCMC gradient (node number > 16), a learning rate of 0.2 for explicit gradient (node number ≤ 16), and with a decay rate $\beta 1 = 0.9, \beta 2 = 0.999, \epsilon = 10^{-6}$ [14]. Adam optimizer can significantly speed up the convergence rate of optimization

but suffer from instability during the final stage of optimization. We deploy a backtracking scheme where we roll back 20 iterations and half the step size when the current gradient is larger than twice of gradient 20-iteration before. We use $300 \sim 1000$ learning interactions depending on the convergence rate and select the network with the smallest gradient in the learning trajectories.

1.8 Parallelization of D-SPIN inference

The gradient estimation of each summation term in Eq. 14 and each condition in Eq. 15 are independent of each other. Leveraging the property, D-SPIN performs parallelized model inference by parceling computation for each condition to different CPUs. We use MATLAB "parfor" function on the Caltech Resnick cluster with each task taking 2 cores and 32G memory. The inference of response vector $\boldsymbol{h}$ given on a fixed network $\boldsymbol{J}$ is also independent for each condition and is parallelized with Slurm on Caltech Resnick cluster with 2 cores and 32G memory.

1.9 Simulation of perturbations on HSC network

We follow the formulation and code in the BEELINE benchmarking study [15], with the modification of adding an external mRNA production input term as activation perturbations and a factor of mRNA production blocking as knockdown perturbations. Specifically, we simulate the following dynamical systems

$$\frac{d[x_i]}{dt} = e_i + \beta_i m_i f(R_i) - l_x[x_i] + s\sqrt{[x_i]}\Delta W_t \tag{23}$$

$$\frac{d[p_i]}{dt} = r[x_i] - l_p[p_i] + s\sqrt{[p_i]}\Delta W_t \tag{24}$$

$$\Delta W_t = \mathcal{N}(0, h) \tag{25}$$

where $[x_i]$ and $[p_i]$ are mRNA and protein concentration of gene i, e_i is external mRNA production by perturbation, β_i is a factor of reduced production rate by perturbation, m_i and r is the mRNA and protein production rate, $f(R_i)$ is the regulation received from the network, l_x and l_p are mRNA and protein degradation rate, s is thermal noise strength, W_t is the Winner process representing thermal noises, and h is simulation time step. The default parameter values are $e_i = 0, \beta_i = 1, m_i = 20, l_x = 10, r = 10, l_p = 1, s = 10$.

In single gene and double gene knockdown simulation, we use $\beta_i = 0.2$ for reduced mRNA production rate. In single gene knockdown simulation, we set $e_i = 15$, corresponding to 0.75 of the expression of fully activated genes. We find for double perturbations higher external activation is required to drive the network, so we use $e_i = 22.5$ for gene activation in double perturbations.

We simulate the system under 10 different sets of randomly sampled parameters. Each time the parameter m_i, l_x, l_p, r are set to a value αD, where D is the default value and $alpha$ is a random number sampled from a truncated Gaussian distribution (mean = 1, std = 0.1, truncation to $[0.9, 1.1]$). m_i of all genes are set to the same value for numerical stability. We compute the UMAP embedding of discretized simulated data with D-SPIN generated samples and annotate the cell type of each cluster in the UMAP following the study originally proposed the HSC network [16]. The other three methods PIDC, GRNBOOST2, and GENIE3 do not model perturbation explicitly, so we run the Docker image of the three methods provided by BEELINE on control and each perturbation condition independently and take the average score for each network edge as the overall inference.

1.10 Discovering regulatory networks modules and clustering perturbation responses

To gain insight into the structure and modularity of the D-SPIN inferred regulatory network, we apply the Leiden graph community detection algorithm to both Perturb-seq and drug profiling datasets. The Leiden algorithm identifies sub-networks that have more positive connections within a set of nodes than expected in a random network. Formally, the Leiden algorithm for a network with positive

and negative edges maximizes the following score function [17].

$$Q = Q(J_+) - Q(J_-) \tag{26}$$

$$Q := \frac{1}{2m} \sum_{i,j} \left[\left(J_{ij} - \gamma \frac{k_i k_j}{2m} \right) \delta(c_i, c_j) \right] \tag{27}$$

where the regulatory network is split into two networks of positive interactions J_+ and negative interactions J_-. In the quality function Q, m is the sum of all edge weights, k_i is the sum of edge weights of node i, c_i is the community index of node i, and γ is the resolution parameter. We set $\gamma = 2$ for a finer resolution of module discovery. The Leiden algorithm partitions the network into a set of non-overlapping node sets, and we annotate the biological function of these modules based on the gene programs they contain.

In addition to analyzing the architecture of the D-SPIN network, we also cluster the perturbation response vectors to identify patterns of interaction between perturbations and the network. In general, we compute relative effect $\tilde{\boldsymbol{h}}$ for each $\boldsymbol{h}$ by subtracting the average $\boldsymbol{h}$ of control samples, and perform statistical testing of whether any entry of $\tilde{\boldsymbol{h}}$ significantly different from control samples to identify effective perturbations. We then cluster $\tilde{\boldsymbol{h}}$ for effective perturbations to define a set of perturbation classes.

In the Perturb-seq study, we select 1,755 effective perturbations with $p < 0.05$ under Bonferroni multi-test correction using non-targeting guide RNA as control distribution. Because (1) some perturbations in the dataset are noisy as the number of collected cells is low (2) the knockdown efficiency of each guide RNA varies, we use the Leiden algorithm for better clustering using the information of similar perturbations. Using an 8-nearest-neighborhood graph and Leiden resolution $\gamma = 3$, we group the perturbation response vectors into 40 clusters. The guide class response vectors are defined as the average of relative responses for perturbations in the class.

In the drug profiling study, we compute a score of each drug as the negative log probability density function at $\tilde{\boldsymbol{h}}$ of multivariate Gaussian distribution with zero mean and diagonal covariance of (1) the control samples of the same experimental batch if there are no less than 2 control samples (2) the control samples of all experimental batch if there is only a single control sample. We select 158 candidate drugs with a score larger than 20, and performed k-means clustering with 20 clusters. We find some drug clusters are putative experimental batch effects, as it contains biochemically different drugs in the same experimental batch. We also remove the drug cluster where the complete loss of macrophage population is observed, as the loss of macrophage is not reproducible in other experiments, potentially due to cell adhesion to the plate. We then manually annotate the drug clusters and combine clusters of drugs with similar impacts and biology. In the end, we present 71 drugs with 7 major categories.

1.11 Automated identification of regulatory strategies and network coarse-graining

The response vectors not only provide single-program-level perturbation impacts but also enable a global perspective of how cells coordinate a variety of cellular activities to adapt to the stressed condition induced by the perturbation. To group perturbation classes into global strategies, typical measures such as Euclidean distance are not sufficient because each gene program is not an independent dimension, but closely interacts with each other through the regulatory network. Therefore we include detailed information about relationships between gene programs by applying a graph clustering strategy where we construct an extended network that incorporates both edges between perturbations and gene programs as well as interactions between gene programs themselves. By community detection on the extended network, perturbations that have similar impacts on closely-related gene programs are grouped together. As the regulatory network is partitioned into subnetworks of gene modules, the graph clustering strategy can unravel the information processing of the regulatory network by dissecting how perturbation responses are organized across modules.

Specifically, we perform the Leiden community detection algorithm on the extended network given the modules of gene programs. The extended network includes both gene programs and perturbation classes as nodes, and the edges are interactions between gene programs plus edges connecting each perturbation class to gene programs with weights being the average relative response $\tilde{\boldsymbol{h}}$ in the class.

Formally, the adjacency matrix of the extended network is

$$\boldsymbol{J}_{\text{ext}} = \begin{bmatrix} \boldsymbol{J} & [\tilde{\boldsymbol{h}}_1, \tilde{\boldsymbol{h}}_2, \dots] \\ [\tilde{\boldsymbol{h}}_1, \tilde{\boldsymbol{h}}_2, \dots]^\intercal & \boldsymbol{0} \end{bmatrix} \tag{28}$$

The graph clustering assigns each perturbation class to a network module. Based on the score function Eq. 27 of the Leiden algorithm, a perturbation class assigned to a gene program module indicates that the perturbation has more activating interactions with programs in the module and more inhibiting interactions with programs outside the module compared to random network. Conceptually, the perturbation being grouped with a module means that the response to the perturbation is generally activating programs in that module. Therefore, we name the global types of perturbation responses by the gene module that they are grouped with. For example, perturbations that are grouped with the Translation module are named "Translation up-regulation response".

We further visualize the global organization of biological functions and perturbation response strategies with a minimal network coarse-grained from the D-SPIN model. As the response strategies are defined on the module level, and modules of gene programs typically correspond to a biological function with only positive interactions inside, each module and perturbation response strategy can be simplified to a single node in the minimal network. Specifically, in the minimal network interactions between modules are the strongest positive and negative edges between all programs in the two modules. The edges between a module and a response strategy are set to the average strength of the positive or negative edge between the module and the strategy, taking the stronger one. As a visualization tool, the minimal network clearly shows the major regulatory connections between modules of biological functions such as Translation and Degradation, as well as how each response strategies activate a major responding gene module and represses other modules.

1.12 Embedding and comparison of D-SPIN and data distributions

To compare samples from D-SPIN models and empirical data, we embed model-generated data and discretized single-cell data in a common space, or project model-generated data onto the space of single-cell data.

As there are relatively few cell states and moderate cell numbers in HSC simulation and Perturb-seq control samples, we embed D-SPIN samples with empirical data to the same UMAP. D-SPIN samples are directly generated with Markov Chain Monte Carlo. Empirical data samples are discretized by the D-SPIN framework so model-generated data and empirical data are in the same space. For UMAP embedding only, a small Gaussian noise is added to data to avoid the neighborhood graphs only having connections between the same state.

In the drug profiling study, as the UMAP is computed over 1 million cells and is batch corrected (Sec. 1.3), we project D-SPIN samples to the data UMAP. For each D-SPIN sample, we map it to the corresponding single-cell sample whose discretized state is the nearest neighbor of the D-SPIN sample. Such that a set of D-SPIN samples defines a distribution over all observed single-cell states. We sample 10^5 cells from the distribution to render the UMAP of D-SPIN samples.

1.13 Coherence of molecular complexes in CORUM

We filter the protein complex in the CORUM database [18] by (1) having no less than 10 subunits and (2) having no less than 2/3 of subunits in the selected guide RNA list. Many complexes in the database are variation forms of the same complex and contain a large number of shared components. We construct a similarity graph between these complexes using shared subunits as edges, cluster the complexes into 46 classes, and pick the complex with the highest number of subunits from each class. The coherence score is computed as the average of all pair-wise cosine similarity of response vectors of subunit perturbations.

$$\text{coherence} = \frac{2}{N(N-1)} \sum_{i<j} \frac{\tilde{\boldsymbol{h}}_i^\intercal \cdot \tilde{\boldsymbol{h}}_j}{\|\tilde{\boldsymbol{h}}_i\| \|\tilde{\boldsymbol{h}}_j\|} \tag{29}$$

where N is the number of subunits in the complex, and $\tilde{\boldsymbol{h}}_i$ is the relative response vector for knock-down of subunit i.

The D-SPIN can reveal the function partition of molecular complex subunits by clustering the perturbation response vectors. There is no universal rule for choosing the number of clusters, and the choice can be somewhat subjective depending on the application scenario. Various metrics have been proposed to facilitate the choice of the cluster number, and we use the Silhouette score for demonstration purposes. Silhouette score examines the quality of clustering by computing the average distance inside a cluster and between a neighboring cluster, and clustering with a higher score is considered better. For each of the molecular complexes, we run K-means clustering for k ranging from 2 to 15 and compute the Silhouette score for each k. A molecular complex is considered to contain multiple functional groups of subunits if the maximum Silhouette score is achieved by a k larger than 3. We find the following molecular complexes contain multiple functional groups: Proteasome, Integrator, Spliceosome A complex, Spliceosome E complex, mRNA decay complex, and NCOR1 complex. We note that other complexes may also be partitioned into functional groups, and these complexes are the ones selected by the Silhouette score.

1.14 Identify drug target and potency from database

The drug target information is obtained by querying the probes & drugs database [19]. The primary target of drugs is based on Drug Bank [20] as well as manual lookup.

1.15 Coherence between drug responses

We quantify the coherence between two drugs with both global coherence and program-level coherence. The global coherence is defined by the cosine similarity between the relative response vector of the two drugs. coherence $= \tilde{\boldsymbol{h}}_1^\intercal \cdot \tilde{\boldsymbol{h}}_2 / \|\tilde{\boldsymbol{h}}_1\| \|\tilde{\boldsymbol{h}}_2\|$. The program-level coherence is defined by the sign function of single drug effects. When the relative response vectors of the two single drugs have the same sign on a gene program, the two drugs are coherent on the program. Different signs mean incoherent on the program.

1.16 Phase diagrams of dosages of drug combination

The regulatory network inferred by D-SPIN is able to convert a response vector into a distribution of cell states, therefore providing a natural way of interpolating cell state distributions between experimental conditions. In the dosage combination experiments of Dasatinib and Halcinonide, we observe the transition between activated macrophage states, monocyte-like resting states, and combinatorial states with both M2 and monocyte transcriptional programs. We use D-SPIN to depict the transition 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 combination dosage titration, we exclusively used samples from two drug dosage combination experiments to fit a regulatory network $\boldsymbol{J}$ and response vectors $\boldsymbol{h}$. The two drug combinations compose a variety of cell states so we can leverage the capacity of D-SPIN to integrate information across different perturbations. The first experiment includes varying concentrations of two drugs, Dasatinib and Halcinonide, in the following dosages: Dasatinib at 0.1, 1, 10, 100, and 1000nM, and Halcinonide at 0.01, 0.1, 1, 10, 100, and 1000nM. We profile all possible combinations of these two drug dosages. The second experiment includes a similar dosage design, using Tacrolimus and Motolimod: Tacrolimus is administered at dosages of 0.1, 1, 10, 100, and 1000nM, and Motolimod at dosages of 1, 10, 100, 1000, and 10000nM. Again, all combinations of these two drug dosages are profiled.

To understand how drug combination dosage modulates cell states, we obtain the phase diagram of cell states by estimating the average log-likelihood of a set of experimentally observed states. Specifically, we are interested in the myeloid state in the cell population. At any given drug combination concentration (c_1, c_2), we interpolate the response vector $\hat{\boldsymbol{h}}(c_1, c_2)$ (MATLAB fit 'thinplateinterp') and compute the score u of the state as average log-likelihood of target state $\{\boldsymbol{s}\}_{\text{Target}}$ as

$$u_{\text{Target}} = \frac{1}{\#\{\boldsymbol{s}\}_{\text{Target}}} \sum_{\{\boldsymbol{s}\}_{\text{Target}}} \log P(\boldsymbol{s}; \boldsymbol{J}, \hat{\boldsymbol{h}}(c_1, c_2)) \tag{30}$$

Then we use softmax across all target states to compute the estimated likelihood of each state

$$\hat{p}_{\text{Target i}} = \frac{\exp u_{\text{Target i}}}{\sum_i \exp u_{\text{Target i}}} \tag{31}$$

and render the phase diagram by the target state with the largest probability.

In the Dasatinib plus Halcinonide experiment, we use 5 target myeloid states of (a) Activated macrophage: activated control (b) Monocyte-like: resting control (c) M2 macrophage: 1000nM Halcinonide (d) Inhibited monocyte: 1000nM Dasatinib (e) M2-like monocyte: 1000nM dasatinib plus 1000nM Halcinonide. In the Tacrolimus plus Motolimod experiment, we use 3 taget myeloid states of (a) Activated macrophage: activated control (b) Monocyte-like: 1000nM Tacrolimus (c) Pathogen-response macrophage: 10000nM Motolimod.

2 Experimental methods details

2.1 High-throughput drug perturbation experiments on PBMCs

Cryopreserved human peripheral blood mononuclear cells (PBMCs) from healthy donors (Hemacare) were thawed in a 37C water bath and transferred to a conical vial. Pre-warmed RPMI-1640 was then added at a 10:1 ratio relative to the original cell suspension to dilute the cryoprotectant DMSO. Cells were pelleted for 2 min at 300 rcf and then resuspended in pre-warmed RPMI-1640 (Gibco 22400089) + 1% Pen/strep (Gibco 15070063) + 10% Fetal Bovine Serum (Thermo Fisher A3840001) to a final cell density of 10^6 cells/mL. Cells were plated into CELLSTAR CellRepellent Surface Cell Culture Dishes (Greiner 7000687/664970) and then rested in a tissue culture incubator at 37C, 5% CO2 for 16 hours before being exposed to small molecule drugs.

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 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/200uL. After drug dilutions we prepared and dispensed into a 96-well plate layout, 200,000 PBMCs were loaded into each well of the 96-well plate and cultured for 24 hours in a tissue culture incubator at 37C 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[21].

2.2 High-throughput single-cell profiling using sample multiplexing techniques

After 24 hours of exposure to drugs, cells were harvested for single-cell profiling. Cells were enzymatically dissociated into a single-cell suspension and multiplexed using MULTI-seq lipid-modified oligos [22]. 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 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 (249935), and then 100 uL 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 uL 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 37C 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 [22] 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.

2.3 Timecourse profiling of T-cell mediated immune activation

A reverse timecourse 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 uL of RPMI-1640 + 10% FBS + 1% P/S. The cells were then brought to a Tecan Evo 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 + 10% FBS + 1% P/S, and 50 uL 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 37C and 5% CO2. After the conclusion of the timecourse, cells were immediately harvested for single-cell profiling using multiplexing tags (MULTI-seq) as described above.

2.4 Alignment and de-multiplexing

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

Excpet MULT-19 where 10X Cellplex pipeline was used to demultiplex the conditions of each cell, we use a custom Python pipeline to assign cells to each condition based on the sequenced tags following the pipeline in MULTI-seq [22].

References

[1] Jialong Jiang, David A Sivak, and Matt Thomson. Active learning of spin network models. *arXiv preprint arXiv:1903.10474*, 2019.

[2] Joseph M Replogle, Reuben A Saunders, Angela N Pogson, Jeffrey A Hussmann, Alexander Lenail, Alina Guna, Lauren Mascibroda, Eric J Wagner, Karen Adelman, Gila Lithwick-Yanai, et al. Mapping information-rich genotype-phenotype landscapes with genome-scale perturb-seq. *Cell*, 185(14):2559–2575, 2022.

[3] Romain Lopez, Jeffrey Regier, Michael B Cole, Michael I Jordan, and Nir Yosef. Deep generative modeling for single-cell transcriptomics. *Nature methods*, 15(12):1053–1058, 2018.

[4] Seungjin Choi. Algorithms for orthogonal nonnegative matrix factorization. In *2008 ieee international joint conference on neural networks (ieee world congress on computational intelligence)*, pages 1828–1832. IEEE, 2008.

[5] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems*, 26, 2013.

[6] Stephen A Vavasis. On the complexity of nonnegative matrix factorization. *SIAM Journal on Optimization*, 20(3):1364–1377, 2010.

[7] Brad T Sherman, Ming Hao, Ju Qiu, Xiaoli Jiao, Michael W Baseler, H Clifford Lane, Tomozumi Imamichi, and Weizhong Chang. David: a web server for functional enrichment analysis and functional annotation of gene lists (2021 update). *Nucleic acids research*, 50(W1):W216–W221, 2022.

[8] Maxim V Kuleshov, Matthew R Jones, Andrew D Rouillard, Nicolas F Fernandez, Qiaonan Duan, Zichen Wang, Simon Koplev, Sherry L Jenkins, Kathleen M Jagodnik, Alexander Lachmann, et al. Enrichr: a comprehensive gene set enrichment analysis web server 2016 update. *Nucleic acids research*, 44(W1):W90–W97, 2016.

[9] Damian Szklarczyk, Annika L Gable, Katerina C Nastou, David Lyon, Rebecca Kirsch, Sampo Pyysalo, Nadezhda T Doncheva, Marc Legeay, Tao Fang, Peer Bork, et al. The string database in 2021: customizable protein–protein networks, and functional characterization of user-uploaded gene/measurement sets. *Nucleic acids research*, 49(D1):D605–D612, 2021.

[10] H Chau Nguyen, Riccardo Zecchina, and Johannes Berg. Inverse statistical problems: from the inverse ising problem to data science. *Advances in Physics*, 66(3):197–261, 2017.

[11] Narayana P Santhanam and Martin J Wainwright. Information-theoretic limits of selecting binary graphical models in high dimensions. *IEEE Transactions on Information Theory*, 58(7):4117–4134, 2012.

[12] Sisi Chen, Paul Rivaud, Jong H Park, Tiffany Tsou, Emeric Charles, John R Haliburton, Flavia Pichiorri, and Matt Thomson. Dissecting heterogeneous cell populations across drug and disease conditions with popalign. *Proceedings of the National Academy of Sciences*, 117(46):28784–28794, 2020.

[13] Cameron Sokolik, Yanxia Liu, David Bauer, Jade McPherson, Michael Broeker, Graham Heimberg, Lei S Qi, David A Sivak, and Matt Thomson. Transcription factor competition allows embryonic stem cells to distinguish authentic signals from noise. *Cell systems*, 1(2):117–129, 2015.

[14] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[15] Aditya Pratapa, Amogh P Jalihal, Jeffrey N Law, Aditya Bharadwaj, and TM Murali. Benchmarking algorithms for gene regulatory network inference from single-cell transcriptomic data. *Nature methods*, 17(2):147–154, 2020.

[16] Jan Krumsiek, Carsten Marr, Timm Schroeder, and Fabian J Theis. Hierarchical differentiation of myeloid progenitors is encoded in the transcription factor network. *PloS one*, 6(8):e22649, 2011.

[17] Vincent A Traag, Ludo Waltman, and Nees Jan Van Eck. From louvain to leiden: guaranteeing well-connected communities. *Scientific reports*, 9(1):5233, 2019.

[18] Andreas Ruepp, Brigitte Waegele, Martin Lechner, Barbara Brauner, Irmtraud Dunger-Kaltenbach, Gisela Fobo, Goar Frishman, Corinna Montrone, and H-Werner Mewes. Corum: the comprehensive resource of mammalian protein complexes—2009. *Nucleic acids research*, 38(suppl_1):D497–D501, 2010.

[19] Ctibor Skuta, Martin Popr, Tomas Muller, Jindrich Jindrich, Michal Kahle, David Sedlak, Daniel Svozil, and Petr Bartunek. Probes & drugs portal: an interactive, open data resource for chemical biology. *Nature methods*, 14(8):759–760, 2017.

[20] David S Wishart, Craig Knox, An Chi Guo, Savita Shrivastava, Murtaza Hassanali, Paul Stothard, Zhan Chang, and Jennifer Woolsey. Drugbank: a comprehensive resource for in silico drug discovery and exploration. *Nucleic acids research*, 34(suppl_1):D668–D672, 2006.

[21] Lisa Holthaus, Daniel Lamp, Anita Gavrisan, Virag Sharma, Anette-Gabriele Ziegler, Martin Jastroch, and Ezio Bonifacio. Cd4+ t cell activation, function, and metabolism are inhibited by low concentrations of dmso. *Journal of immunological methods*, 463:54–60, 2018.

[22] Christopher S McGinnis, David M Patterson, Juliane Winkler, Daniel N Conrad, Marco Y Hein, Vasudha Srivastava, Jennifer L Hu, Lyndsay M Murrow, Jonathan S Weissman, Zena Werb, et al. Multi-seq: sample multiplexing for single-cell rna sequencing using lipid-tagged indices. *Nature methods*, 16(7):619–626, 2019.