
Identifying biological perturbation targets through causal differential networks

Menghua Wu¹ Umesh Padia¹ Sean H. Murphy¹ Regina Barzilay¹ Tommi Jaakkola¹

Abstract

Identifying variables responsible for changes to a biological system enables applications in drug target discovery and cell engineering. Given a pair of observational and interventional datasets, the goal is to isolate the subset of observed variables that were the targets of the intervention. Directly applying causal discovery algorithms is challenging: the data may contain thousands of variables with as few as tens of samples per intervention, and biological systems do not adhere to classical causality assumptions. We propose a causality-inspired approach to address this practical setting. First, we infer noisy causal graphs from the observational and interventional data. Then, we learn to map the differences between these graphs, along with additional statistical features, to sets of variables that were intervened upon. Both modules are jointly trained in a supervised framework, on simulated and real data that reflect the nature of biological interventions. This approach consistently outperforms baselines for perturbation modeling on seven single-cell transcriptomics datasets. We also demonstrate significant improvements over current causal discovery methods for predicting soft and hard intervention targets across a variety of synthetic data.¹

1. Introduction

Cells form the basis of biological systems, and they take on a multitude of dynamical states throughout their lifetime. In addition to natural factors like cell cycle, external perturbations (e.g. drugs, gene knockdown) can alter a cell’s state. While perturbations can affect numerous downstream

variables, identifying the root causes, or *targets*, that induce these transitions has vast therapeutic implications, from cellular reprogramming (Cherry & Daley, 2012) to mechanism of action elucidation (Schenone et al., 2013).

Machine learning methods have primarily been designed to infer the effects of perturbations on cells, with the goal of generalizing to unseen perturbations (Roohani et al., 2023), or unseen cell distributions (Lotfollahi et al., 2023; Bunne et al., 2023). In principle, these models can be used within an active learning framework to discover intervention targets (Huang et al., 2024; Zhang et al., 2023a). However, the number of candidate sets scales exponentially with the number of variables under consideration, and there are limited training data for combinatorial perturbations (< 150 pairs in Norman et al. (2019)). Thus, this option may be impractical for larger search spaces.

Alternatively, some methods predict targets by comparing pairs of “before” and “after” cell states, similar to this paper. They do so by attributing observed differences to sparse, mechanistic changes. Specifically, relationships between variables are represented through a graph, e.g. a gene regulatory network (Babu et al., 2004). Existing methods often use data-mined relations as priors for this graph, and search for modifications to the set of edges that best explain the interventional data (Cosgrove et al., 2008; Gonzalez et al., 2024). However, biological knowledge graphs are incomplete and noisy priors. They contain heterogeneous, potentially conflicting information, since the reported findings were observed under diverse contexts (Ashburner et al., 2000). Moreover, different cell populations may exhibit distinct gene regulatory behavior (Huttlin et al., 2021), leading to incomplete knowledge, especially of less well-characterized systems.

In this paper, we propose causal differential networks (CDN): a causality-inspired approach to identify variables that drive desired shifts in cell state, while estimating their mechanistic structure directly from data. Given a pair of observational and interventional datasets, we train a causal structure learner to predict causal graphs that could have generated each dataset. The pair of graphs is input to an attention-based classifier, which predicts whether each variable was subject to intervention. To address the challenges of data

¹Department of Computer Science, Massachusetts Institute of Technology, Cambridge, MA, USA. Correspondence to: Menghua Wu <rmwu@mit.edu>.

Proceedings of the 42nd International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).

¹Code is available at <https://github.com/rmwu/cdn>

noise and sparsity, and to simulate the structure of biological interventions, these models are trained jointly in a supervised (amortized) framework (Ke et al., 2023; Petersen et al., 2023), over thousands of synthetic or real datasets.

We evaluate CDN on real transcriptomic data and synthetic settings. CDN outperforms the state-of-the-art in perturbation modeling (deep learning and statistical approaches), evaluated on the five largest Perturb-seq datasets at the time of publication (Replogle et al., 2022; Nadig et al., 2024) without using any external knowledge. Furthermore, CDN generalizes with minimal performance drop to unseen cell lines, which have different supports (genes), causal mechanisms (gene regulatory networks), and data distributions. On synthetic settings, CDN outperforms causal discovery approaches for estimating unknown intervention targets.

2. Background and related work

2.1. Problem statement

We consider a system of N random variables $X_i \in X$, whose relationships are represented by directed graph $G = (V, E)$, where nodes $i \in V$ map to variables X_i , and edges $(i, j) \in E$ to relationships from X_i to X_j . Samples from the data distribution P_X are known as *observational*. We can perform *interventions* by assigning new conditionals

$$P(X_i | X_{\pi_i}) \leftarrow \tilde{P}(X_i | X_{\pi_i}), \quad (1)$$

where π_i denotes the parents of node i in G . *Hard* interventions remove all dependence between X_i and π_i , while *soft* interventions maintain the relationship with a different conditional. We denote the joint *interventional* distribution as \tilde{P}_X . Given an observational dataset $D_{\text{obs}} \sim P_X$ and an interventional dataset $D_{\text{int}} \sim \tilde{P}_X$, our goal is to predict the set of nodes I for which

$$P(X_i | X_{\pi_i}) \neq \tilde{P}(X_i | X_{\pi_i}), \forall i \in I. \quad (2)$$

In the context of molecular biology, a *perturbation* is a hard or soft intervention, and I is the set of perturbation *targets*, where $\|I\| \geq 1$.

2.2. Modeling perturbations

Perturbation experiments result in datasets of the form $D_{\text{obs}}, \{(I^k, D_{\text{int}}^k)\}$, where $D_{\text{obs}} \sim P_X, D_{\text{int}}^k \sim \tilde{P}_X^k$. In this paper, we focus on the *transcriptomic* data modality, where variables X_i represent the levels of gene i ; the number of variables N ranges from hundreds to thousands; D_{obs} contains thousands of samples; and each D_{int} contains tens to hundreds of samples (Replogle et al., 2022).

Each experiment is conducted in a single *cell line*, which represents a distinct set of variables X , distribution P_X , and graph G . These three objects may share similarities across

cell lines, as certain genes or pathways are essential for survival. However, the empirical discrepancy can be quite large between experiments. For example, Nadig et al. (2024) reports that the median correlation in log-fold change of genes under the same perturbation, across two experiments in the same cell line, by the same lab, was 0.16.

Early efforts to map the space of cells measured the distribution of gene levels as a surrogate for cell state (Lamb et al., 2006). Diseases, drugs, or genetic perturbations that induce similar changes in this distribution were thought to act through similar mechanisms. However, it is intractable to map all combinations of starting cell states and perturbations through experiment alone. This bottleneck has led to two directions in the *in-silico* modeling of perturbations.

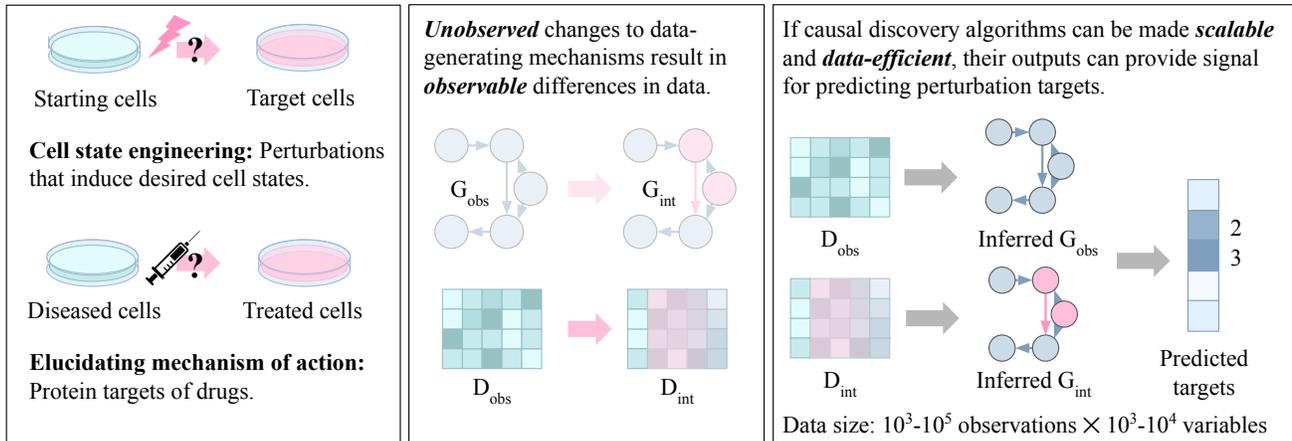
To reduce experimental costs, which scale with the number of cell lines and perturbations, one option is to infer the post-intervention distribution of cells. That is, given $D_{\text{obs}} \sim P_X$ and I , the goal is to infer \tilde{P}_X . Works such as scGen (Lotfollahi et al., 2019), CPA (Lotfollahi et al., 2023), ChemCPA (Hetzel et al., 2022), and CellOT (Bunne et al., 2023) aim to generalize to new settings P_X , e.g. unseen cell lines or patients. Alternatively, models like GEARS (Roohani et al., 2023), graphVCI (Wu et al., 2023), AttentionPert (Bai et al., 2024), and GIM (Schneider et al., 2024) infer the effects of unseen perturbations. This setting is also a common application for single cell foundation models (Theodoris et al., 2023; Cui et al., 2024; Hao et al., 2024). These latter methods generalize by modeling the relationships between seen and unseen perturbations. For example, GEARS leverages literature-derived relationships in the Gene Ontology knowledge graph (Ashburner et al., 2000), and graphVCI learns to refine the gene regulatory relationships implied by ATAC-seq data (Grandi et al., 2022).

Approaches for inferring the effects of interventions do *not* explicitly predict targets, but if the space of perturbations \mathcal{I} can be efficiently enumerated, these models are useful for comparing different I (Zhang et al., 2023a; Huang et al., 2024). For example, given an oracle $\phi : D_{\text{obs}}, I \rightarrow \tilde{P}$, the I^* associated with \tilde{P}^* is

$$I^* = \arg \min_{I \in \mathcal{I}} d(\tilde{P}^*, \phi(D_{\text{obs}}, I)) \quad (3)$$

for some dissimilarity d . However, recent works have reported that the predictive ability of current models does not exceed that of naive baselines (Kernfeld et al., 2023; Ahlmann-Eltze et al., 2024; Märtens et al., 2024), reflecting that phenotype prediction is currently unsolved.

Perturbation datasets can also be used to train models for predicting targets of interest, as is the focus of this work. A common strategy is to relate observed changes (e.g. disease, gene levels) to mechanistic explanations (e.g. pathways), in which the targets are involved. To draw this connection, current approaches use external information such as



(A) Perturbation target prediction (B) Causal interpretation (C) Causality-inspired approach

Figure 1: (A) Biological applications. (B) Differences between observation dataset D_{obs} and interventional dataset D_{int} can be attributed to changes in the underlying causal mechanisms. (C) We use an amortized causal discovery model to predict G_{obs} and G_{int} , whose hidden representations are input to a perturbation target classifier. Both modules are jointly trained.

knowledge graphs (Cosgrove et al., 2008; Gonzalez et al., 2024) and natural language (Roohani et al., 2025). However, these information are collected from highly inhomogenous sources (Ashburner et al., 2000), and it is difficult to quantify their correctness and completeness in each new setting. An alternative is to infer the mechanistic structure directly from data, either explicitly (di Bernardo et al., 2005) or as a model parameter (Noh & Gunawan, 2016). In this work, we follow the latter strategy, while allowing for a richer (non-linear) class of relationships.

2.3. Causal discovery

The task of inferring causal mechanisms from data is known as causal discovery, or causal structure learning (Spirtes et al., 2001). In the simplest setting, the goal is to infer G from D_{obs} . The earliest works were discrete search algorithms, which operated over the combinatorial space of potential graphs (Spirtes et al., 1995; Tian & Pearl, 2001; Chickering, 2002; Shimizu et al., 2006). More recently, a number of works have reframed the task as a constrained optimization over continuous adjacency matrices (Zheng et al., 2018; Lachapelle et al., 2020; Sanchez et al., 2023).

If data from multiple environments are available, they can provide additional identifiability regarding the system (Jaber et al., 2020; Zhang et al., 2023b). Differences between these environments, e.g. intervention targets I , are either given as input or part of the inference task. As in the observational case, discovery algorithms can be categorized into discrete search (Hauser & Bühlmann, 2012; Ghassami et al., 2018; Ke et al., 2019; Huang et al., 2020; Mooij et al., 2020) or continuous optimization (Brouillard et al., 2020; Hägele et al., 2023). Additionally, Varici et al. (2022) and Yang et al. (2024) frame intervention target prediction as its own task, similar to this paper. However, the former strictly as-

sumes linearity, while the latter requires at least as many environments as $\|I\|$ (we only consider before and after perturbation). Finally, several works aim to infer differences between causal mechanisms (edges), rather than identifying the full causal structure (Wang et al., 2018; Belyaeva et al., 2021; Chen et al., 2023; Malik et al., 2024). Our work is also motivated by the idea that differences between summary statistics (e.g. regression coefficients in DCI) are informative of changes between systems. Thus, CDN can be viewed as implementing similar operations within an end-to-end, supervised framework.

2.4. Amortized causal inference

While the majority of causal discovery algorithms must be fit or run from scratch for each data distribution, amortized causal discovery algorithms have addressed this task as a supervised machine learning problem. Using large numbers of synthetic datasets, generated by synthetic graphs, a neural network is trained to map D directly to G (Li et al., 2020; Ke et al., 2023; Lorch et al., 2022; Petersen et al., 2023; Wu et al., 2025). Compared to traditional approaches, these algorithms trade explicit assumptions for implicit characteristics of the data simulation process (Montagna et al., 2024). These models are also fast and robust to noise. In this work, we simulate interventions that reflect biological perturbations, and our model architecture builds upon that of Wu et al. (2025), as it scales easily to hundreds of variables.

3. Methods

We introduce causal differential networks (CDN), a causality inspired method for predicting intervention targets. On a high level, CDN is motivated by the idea that differences in observations D_{obs} and D_{int} can be explained by the changes

to their underlying causal mechanisms (Figure 1B).

Our model is composed of two modules: a *causal structure learner* (Section 3.1), and a *differential network classifier* (Section 3.2). Given a single dataset D , the causal structure learner predicts a putative G . The pair of datasets D_{obs} and D_{int} yields graphs G_{obs} and G_{int} , whose features are input to the differential network, to predict targets I (Figure 1C). Both modules are trained jointly, supervised by classification losses for the graph and predicted targets. Finally, individual samples are known to be noisy in single-cell biology and are often analyzed in “pseudo-bulk” (Blake et al., 2003). Thus, instead of operating over raw data, we also represent datasets in terms of their summary statistics.

3.1. Causal structure learner

The causal structure learner takes as input dataset D and outputs the predicted G . In principle, this module can be implemented with any data-efficient causal discovery model. Here, we based the architecture on SEA (Wu et al., 2025), as its runtime complexity does not explicitly depend on the number of samples (which can be large for D_{obs}).

Input featurization We represent dataset D through two types of statistics. First, we compute global statistics $\rho \in \mathbb{R}^{N \times N}$ (pairwise correlation) over all variables. This statistic is an inexpensive way for the model to filter out edges that are not likely to occur. Second, we obtain estimates of local causal structure to orient edges, by running classical causal discovery algorithms over small subsets of variables. Specifically, we run the FCI algorithm T times over subsets of k variables (Spirtes et al., 1995), which are sampled proportional to their pairwise correlation. We denote the set of estimates as $E' \in \mathcal{E}^{k \times k \times T}$, where \mathcal{E} is the space of possible edge types inferred by FCI (e.g. \leftrightarrow , $-$, \rightarrow). The size and number of local estimates, k and T , are hyperparameters that trade off accuracy vs. runtime (Section 3.3).

To compare ρ and E' in the same space, we linearly project elements of ρ from \mathbb{R} to \mathbb{R}^d , and map categorical edge types into \mathbb{R}^d using a learned embedding, i.e. like token embeddings in language models (Devlin et al., 2019). Thus, the dataset D is summarized into a matrix of size $\mathbb{R}^{N \times N \times T \times d}$.

Model architecture Both the causal structure learner and differential network are implemented with a series of attention-based blocks. Since scaled dot-product attention scales quadratically as sequence length (Vaswani et al., 2017), we attend over one dimension at a time. For example, given a N^2 entries in an adjacency matrix, naive self-attention would cost $O(N^4)$ operations. Instead, we apply self-attention along all nodes in the “outgoing edge” direction, with the “incoming edges” as a batch dimension, followed by the opposite – resulting in two operations, each of cost $O(N \times N^2)$. Following Rao et al. (2021), we use

pre-layer normalization on each self-attention, followed by dropout and residual connections:

$$h \leftarrow h + \text{Dropout}(\text{Self-Attn}(\text{LayerNorm}(h))) \quad (4)$$

where $h \in \mathbb{R}^{\dots \times d}$ denotes an intermediate hidden representation. This design takes after Transformer models for higher dimensional data in other domains, e.g. Axial Transformers for videos (Ho et al., 2019) and MSA Transformers for aligned protein sequences (Rao et al., 2021). A key difference is the graph modality: when making edge or node-level predictions, the model should be equivariant to node labeling. This is accomplished by adding randomly-permuted “positional” embeddings to each node and edge representation. For more details, please see Appendix A.

Edge-level outputs Let $h_{\text{obs}}, h_{\text{int}} \in \mathbb{R}^{N \times N \times d}$ denote the final layer representations of each dataset, after pooling over the T dimension. These representations are input to the differential network. In addition, we introduce a graph prediction head,

$$z_{i,j} = \text{FFN}([h_{i,j}, h_{j,i}]) \quad (\text{logits}) \quad (5)$$

$$P(E_{i,j} = e) = \sigma(z_{i,j})_e \quad (\text{softmax}) \quad (6)$$

where σ denotes softmax over the set of output edge types (forwards, backwards, or no edge). The graph prediction loss \mathcal{L}_G is simply the cross entropy loss between the predicted and true edge types, summed over the entire graph.

3.2. Differential network

Given a pair of graph representations $h_{\text{obs}}, h_{\text{int}}$, the differential network predicts I , the set of intervention targets.

Input featurization In addition to pairwise representations h , we compute node-level statistics (mean, variance), which are mapped to \mathbb{R}^d via a feed-forward network, and concatenated along “incoming edge” dimension. This results in graph representation $h' \in \mathbb{R}^{N \times (N+1) \times d}$.

We combine the graph representations either through concatenation or subtraction:

$$h_{\text{cat}} = [h'_{\text{obs}}, h'_{\text{int}}] \in \mathbb{R}^{N \times (N+1) \times 2d} \quad (7)$$

$$h_{\text{diff}} = h'_{\text{int}} - h'_{\text{obs}} \in \mathbb{R}^{N \times (N+1) \times d}. \quad (8)$$

The former can express richer relationships between the two graphs, while the latter is more efficient on large graphs.

Model architecture The differential network closely follows the architecture of the causal structure learner. However, since paired graph representation lacks the T dimension (introduced by marginal estimates E'), the model only needs to attend over two “length” dimensions (rows and columns of adjacency matrix). After the attention layers, we mean over all “incoming edges” and linearly project to binary node-level predictions, where a normalized 1 indicates an intervention target, and 0 indicates otherwise.

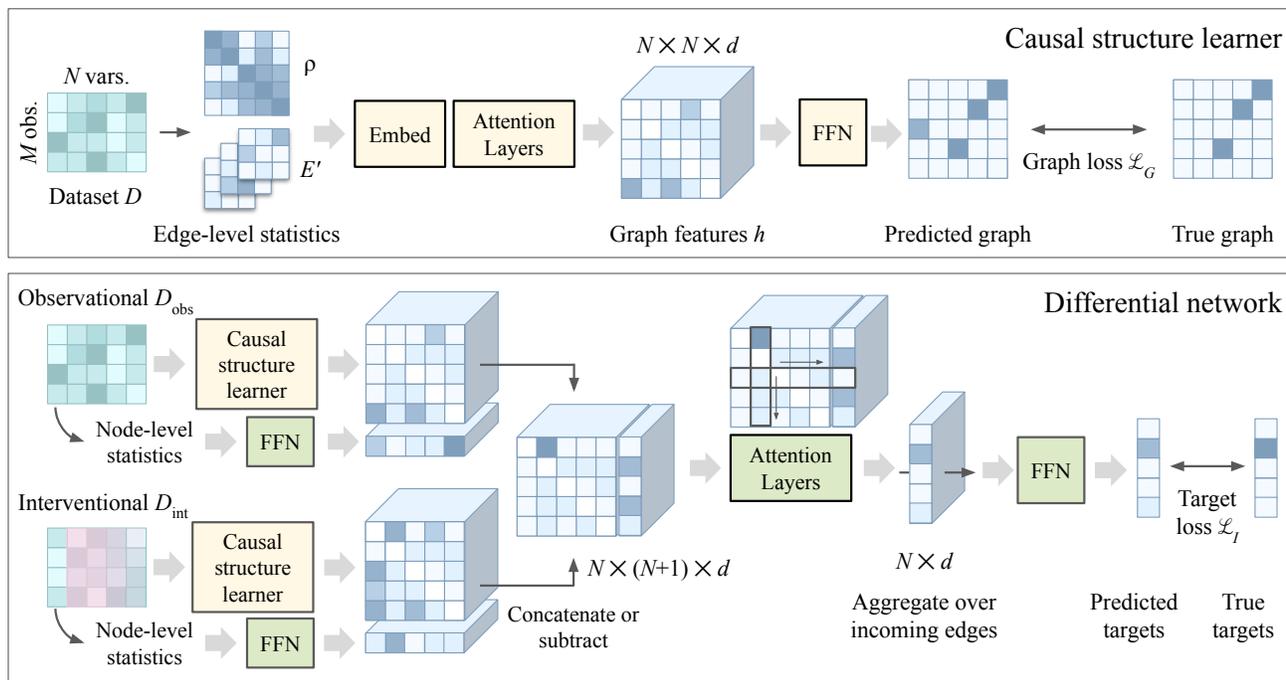


Figure 2: Our model is composed of a causal structure learner and a differential network. Learned components in yellow (causal structure learner) and green (differential network). Both modules are trained jointly to optimize $\mathcal{L} = \mathcal{L}_G + \mathcal{L}_I$.

The target prediction loss \mathcal{L}_I is the binary cross entropy loss between the predicted and true targets, computed independently for each node. This is to accommodate interventions with multiple targets, e.g. drugs with off-target effects.

Remark We take an amortized causal discovery approach because it is hard to specify what assumptions a biological system ought to follow. This makes it difficult to provide identifiability guarantees without simplifying aspects of either the data or modeling framework. Instead, in Appendix C, we show that the differential network is a well-specified class of models. That is, there exist parametrizations of its attention architecture that map predicted graphs and global statistics (for hard and soft interventions) to the set of targets.

3.3. Implementation details

Training procedure The overall training loss is $\mathcal{L} = \mathcal{L}_G + \mathcal{L}_I$ with L2 weight regularization. Both \mathcal{L}_G and \mathcal{L}_I supervise the causal structure learner, while only the latter supervises the differential network. We use the synthetic and real datasets as follows.

1. Initialize the causal structure learner with pretrained weights from Wu et al. (2025) (except ablations).
2. Train both modules jointly on synthetic data.
3. Finetune both modules jointly with only \mathcal{L}_I (no ground truth graphs), on each training split of the real datasets.

For the synthetic training data, we generated 8640 datasets

with hard and soft interventions, which vary in causal mechanism, intervention type, and graph topology. To emulate transcriptomics data, in which perturbation effect sizes are quantified in (log) fold change, we introduced soft interventions: “shift” $x \leftarrow f(\pi_x) + c$, where $c \in \mathbb{R}$, and “scale” $x \leftarrow cf(\pi_x)$, where c is a positive scaling factor, with equal probability $c \leq 1$. We also sampled interventions with multiple targets, to simulate off-target effects. Details are available in Section 4.2 and Appendix D.2.

Model parameters Marginal estimates E' are computed using the classic FCI algorithm (Spirtes et al., 1995) with $k = 5$ and $T = 100$, and global statistic ρ is pairwise correlation, as it is fast to compute and numerically stable. We swept over the number of differential network layers (Figure 4) on synthetic data, and we used 3 layers for h_{cat} and 2 layers for h_{diff} . Following SEA, we adopted hidden dimension $d = 64$, the AdamW optimizer (Loshchilov & Hutter, 2019), learning rate $1e-4$, batch size 16, and weight decay $1e-5$. On the real data, where $N = 1000$, we changed to a batch size of 1, decreased the learning rate to $5e-6$, and finetuned the models with half precision (FP16). We trained both h_{cat} and h_{diff} models on synthetic data, but only finetuned the h_{diff} architecture on the real data (memory constraint).

4. Experiments

While our method is intended to be used in cases where true targets are unknown, we can only evaluate on datasets

where the target is known. Thus, we demonstrate that CDN can recover perturbation targets on seven transcriptomics datasets, with comparisons to state-of-the-art models for these applications (Section 4.1). For completeness, we also benchmarked CDN against multiple causal discovery algorithms for unknown interventions in variety of controlled settings (Section 4.2).

4.1. Biological experiments

Datasets We validate CDN on five Perturb-seq (Dixit et al., 2016) datasets (genetic perturbations) from Replogle et al. (2022) and Nadig et al. (2024); as well as two Sci-Plex (Srivatsan et al., 2020) datasets (chemical perturbations) from McFaline-Figueroa et al. (2024). Each dataset is a real-valued matrix of gene expression levels: the number of examples M is the number of cells, the number of variables N is the number of genes, and each entry is a log-normalized count of how many copies of gene j was measured from cell i . In Perturb-seq datasets, we aim to recover the gene whose promoter was targeted by the CRISPR guide, and in Sci-Plex datasets, we aim to identify the gene that corresponds to the drug’s intended target (discussion in Appendix D.1).

Biologically, many perturbations do not affect the cell’s production of mRNA (they may affect other factors like protein activity, which transcriptomics does not measure). To ensure that P_{obs} and P_{int} are distinct, we filtered perturbations to those that induced over 10 differentially-expressed genes (statistically significant change, compared to control), of which the true target should be present. This excludes perturbations with insufficient cells (low statistical power), with minimal to no effect (uninteresting), and those that did not achieve the desired effect (low CRISPR efficiency). In the Sci-Plex data, several drugs have known off-targets, but we did not observe differential expression of any of them (adjusted p-value ≈ 1).

Evaluation We consider two splits: seen and unseen cell lines. In the former, models may be trained on approximately half of the perturbations from each cell line, and are evaluated on the unseen perturbations. In the latter, we hold out one cell line at a time, and models may be trained on data from the remaining cell lines. To ensure that our train and test splits are sufficiently distinct, we cluster perturbations based on their log-fold change and assign each cluster to the same split (Figure 5). Since the Sci-Plex data are limited, we use them only as test sets without finetuning.

We limited the set of candidate targets to the top 1000 differentially expressed genes by log-fold change, per perturbation. If too few genes passed the significance threshold ($p < 0.05$), we added additional candidates with the largest log-fold change until we reached a minimum of 100 genes. Genetic perturbations have highly specific effects compared to chemical perturbations, so we focused on the subset for

which the intended target was *not* trivially identifiable as the gene with the largest log-fold change.²

While perturbation target prediction appears to be a simple classification task, standard metrics are not well-suited for these real data. Not all genes are present in the baselines’ domain knowledge (e.g. isolated node in graph), so their effects as perturbations cannot be predicted. In addition, due to genetic redundancy, it is common for multiple perturbations to elicit similar responses (Kernfeld et al., 2023). An “incorrect” top 1 prediction may not necessarily reflect poor performance. Thus, we propose the following metrics.

- We order genes based on the similarity of their predicted effect to the ground truth target effect, or based on their predicted probability of being a target. The **rank** is the index of the true target, normalized by the number of candidate genes. Rank ranges from 0 (worst, bottom of list) to 1 (best, top of list).
- We measure the similarity between the perturbation effect (log-fold change) of the top 1 predicted target and the true target with the **Spearman rank correlation** ρ and the **Pearson correlation** r . These range from -1 to 1, where 1 is best.
- To emulate “virtual screening,” we plot **recall at p** : the fraction of targets recovered within the $p \in (0, 1)$ candidates. Recall ranges from 0 to 1, where 1 is best. Rank is equivalent to the expected recall at p .

Baselines We compare to both algorithms that infer the effects of perturbations on cells (used as scoring functions, Equation 3), and algorithms that directly predict targets. All baselines were run with their official implementations and/or latest releases. For details, please see Appendix B.2.

GEARS (Roohani et al., 2023) and **GENEPT** (Chen & Zou, 2023) predict the effects of unseen genetic perturbations. GEARS is a graph neural network that regresses log-fold change upon perturbation, using the Gene Ontology knowledge graph as an undirected backbone. GENEPT is a set of large language model-derived embeddings that have been shown to achieve state-of-the-art performance (Märtens et al., 2024) with simple downstream models like logistic regression (used here). Both are trained per cell line.

PDGRAPHER (Gonzalez et al., 2024) is a graph neural network that predicts the perturbation targets of genetic or chemical perturbations, on seen *or* unseen cell lines, with the Human Reference Interactome (Luck et al., 2020) as the knowledge graph. To the best of our knowledge, PDGRAPHER is the only published deep learning baseline for the target prediction task on transcriptomic data. Therefore, we compare against additional statistical and naive baselines.

²CDN predicts trivial perturbations nearly perfectly (Table 9).

DGE stands for differential gene expression (Love et al., 2014), i.e. the process of identifying statistically significant changes in genes between settings. Specifically, we run the Wilcoxon ranked-sum test with Benjamini-Hochberg correction (Wilcoxon, 1945; Benjamini & Hochberg, 2000) between control and perturbed cells, as implemented by Wolf et al. (2018). Genes are ranked by adjusted p-value. **LINEAR** and **MLP** take as input the mean expression of all perturbation targets, plus the top 2000 highly-variable genes (Wolf et al., 2018). They are trained to predict a binary label for each gene on each cell line independently.

Results CDN consistently outperforms all baselines on the five Perturb-seq datasets, in both the seen and unseen cell line settings (Table 1). There is minimal drop in between seen and unseen cell lines, suggesting that the functional characteristics of genetic perturbations are highly transferrable, even to different data distributions. Notably, no other deep learning method exceeds simple statistical tests (DGE). This is more evident in Figure 3, in which CDN achieves higher recall at p at nearly all points. Finally, if we ablate the finetuning step (“synthetic” vs. “Perturb-seq”), we find that finetuning leads to significant improvements.

Chemical perturbations are much less specific than genetic perturbations (Figures 6 and 7), and it is less clear whether drug effects *should* be reflected in mRNA levels (via feedback mechanisms), as drugs act upon their protein products. Still, on the two chemical perturbation datasets, CDN finds the true target within the top 100 candidates for 3/6 cases (Table 2). Due to the low performance of DGE on several cases, we hypothesized that gene-level changes may not be as predictive. Indeed, an ablation of CDN, trained without node-level statistics, improves where DGE fails (“no μ ”).

4.2. Synthetic experiments

While it would be ideal to evaluate all algorithms on real data, current causal discovery algorithms that support unknown interventions are not tractable on datasets with more than tens of variables. In fact, many baselines require hours on even $N = 10$ datasets, and they do not scale favorably (Table 6). Our transcriptomics datasets contain hundreds of genes, even after filtering to those that are differentially expressed (Figure 8). At the same time, existing models for biological perturbations all rely on some form of domain knowledge, and their performance is inseparable from the choice and quality of these external data. Synthetic data allow us to assess the model’s capacity to predict intervention targets in isolation.

Datasets Each evaluation setting consists of 5 unseen, i.i.d. sampled graphs. To assess the generalization capacity of our framework, we tested the model on seen (linear) and unseen causal mechanisms (polynomial, sigmoid), and unseen soft interventions. To generate observational data, we

sampled Erdős-Rényi graphs with $N = 10, 20$ nodes and $E = N, 2N$ expected edges; causal mechanism parameters; and observations of each variable, in topological order. We sampled $3N$ distinct subsets of 1-3 nodes (N each) as intervention targets. Hard interventions set $x \leftarrow z$, where z is uniform. For soft interventions (Section 3.3), we trained on shift and tested on scale for synthetic experiments (see Table 10 for explanation). Real experiments trained on all interventions.

Baselines We compare against discrete and continuous causal discovery algorithms for unknown interventions. **UT-IGSP** (Squires et al., 2020) infers causal graphs and unknown targets by greedily selecting the permutation of variables that minimizes their proposed score function. **DCDI** (Brouillard et al., 2020) and **BACADI** (Hägele et al., 2023) are continuous causal discovery algorithms that fit generative models to the data, where the causal graph and intervention targets are model parameters. DCDI’s -G and -DSF suffixes correspond to Gaussian and deep sigmoidal flow parametrizations of the likelihood. BACADI’s -E and -M suffixes indicate empirical (standard) and mixture (bootstrap) variants.

MB+CI is inspired by Huang et al. (2020). We introduce a domain index (indicator of environment), compute its Markov boundary using graph lasso, and use conditional mutual information to identify true children of the domain index (intervention targets). **DCI** with stability selection (Belyaeva et al., 2021) predicts edge-level differences between two causal graphs. We take each node’s proportion of changed edges as its likelihood of being an intervention target. While DCI was also motivated by biological applications, it only scales to around a hundred variables at most, so we evaluate it alongside other causal discovery methods here.

Evaluation In the synthetic case, we are not constrained by biological redundancy or incomplete predictions, so we report standard classification metrics: mean average precision (**mAP**) and area under the ROC curve (**AUC**). Both metrics are computed independently for each variable and averaged over all regimes of the same number of targets. Their values range from 0 to 1 (perfect). The mAP random baseline depends on the positive rate, while the AUC random baseline is 0.5 (per edge).

Results On synthetic data, CDN achieves high performance across intervention types and data-generating mechanisms (Table 3), while running in seconds (Table 6). As an ablation study, we investigated removing the graph loss and the pretrained weights, with no change to the graph-centric architecture. We found that supervising based on the graph is almost always helpful (“no \mathcal{L}_G ” row), and pretraining leads to more stable training dynamics (Figure 10).

DCI and MB+CI are the best baselines, which is encourag-

Table 1: Results on 5 Perturb-seq datasets. Top: Train on all cell lines jointly. Bottom: Leave one cell line out of training. Test sets (unseen perturbations) are identical in both settings. CDN (synthetic) and DGE do *not* require training on real Perturb-seq data. Metrics, from left to right: normalized rank of ground truth, top 1 Spearman and Pearson correlations. Uncertainty quantification in Table 7. Runtimes in Table 11.

Setting	Model	K562 (gw)			K562 (es)			RPE1			HepG2			Jurkat		
		rank	ρ	r												
Seen cell line	LINEAR	0.51	0.18	0.23	0.48	0.22	0.25	0.49	0.42	0.52	0.51	0.35	0.43	0.50	0.20	0.23
	MLP	0.45	0.18	0.21	0.49	0.08	0.11	0.53	0.32	0.41	0.48	0.33	0.40	0.45	0.24	0.29
	GENEPT	0.52	0.29	0.32	0.42	0.16	0.19	0.41	0.37	0.46	0.31	0.32	0.40	0.51	0.31	0.36
	GEARS	0.56	0.19	0.22	0.45	0.18	0.20	0.49	0.35	0.44	0.54	0.32	0.40	0.50	0.22	0.26
	PDG	0.54	0.25	0.29	0.49	0.20	0.23	0.48	0.41	0.52	0.54	0.37	0.43	0.52	0.31	0.35
	CDN (Perturb-seq)	0.92	0.72	0.74	0.95	0.81	0.82	0.92	0.67	0.76	0.94	0.69	0.75	0.98	0.83	0.84
Unseen cell line	DGE	0.82	0.50	0.54	0.89	0.60	0.62	0.79	0.56	0.64	0.86	0.51	0.51	0.84	0.45	0.50
	PDG	0.39	0.10	0.13	0.44	0.24	0.28	0.54	0.38	0.48	0.48	0.34	0.41	0.57	0.31	0.35
	CDN (synthetic)	0.79	0.46	0.49	0.80	0.59	0.62	0.82	0.61	0.71	0.84	0.56	0.63	0.87	0.64	0.67
	CDN (Perturb-seq)	0.91	0.69	0.71	0.95	0.81	0.82	0.92	0.65	0.74	0.93	0.67	0.72	0.97	0.84	0.85

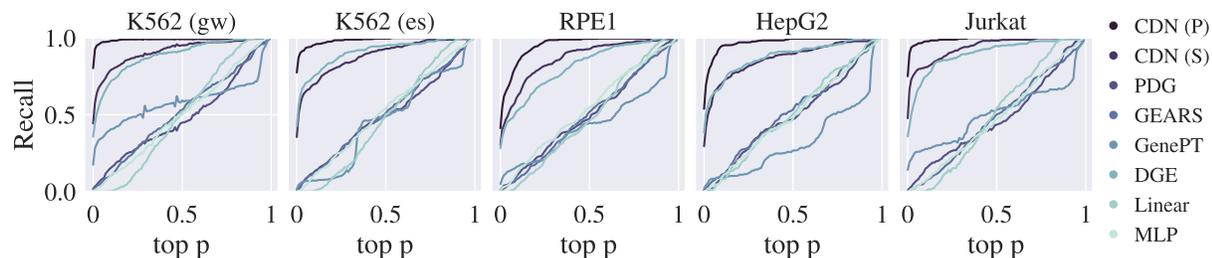


Figure 3: Recall at p (normalized over number of candidate genes) on 5 Perturb-seq datasets. For CDN, (P) and (S) denote ‘‘Perturb-seq’’ and ‘‘synthetic’’ training data.

Table 2: Rank of drug targets on chemical perturbation datasets (McFaline-Figueroa et al., 2024). A172 and T98G are *unseen* cell lines. Highlighted: Within top 100. Drug names and details in Appendix D.1.

Model	A172			T98G		
	infig.	nint.	palb.	doxo.	palb.	vola.
DGE	0.12	0.64	0.28	0.94	0.24	0.83
PDG	0.72	0.35	0.39	0.06	0.34	0.42
CDN	0.14	0.81	0.01	0.89	0.93	0.53
CDN (no μ)	0.88	0.44	0.58	0.06	0.54	0.65

ing, as they are designed with the intuition that detecting differences is easier than reproducing the entire causal graph. Surprisingly, most other baselines perform poorly at recovering intervention targets. In the case of DCDI and BACADI, this may be because it is hard to select a single sparsity threshold for varying sizes of intervention sets, and to balance the sparsity regularizer with the generative modeling objective.

5. Discussion

This work introduces causal differential networks (CDN), which focuses on the application of predicting perturbation targets for efficient experimental design. On a high level,

our goal is to lower the cost and scale of experiments, by reducing the search space. Our key technical contribution is the idea of using supervised causal discovery as a pretraining objective, to obtain dataset representations useful for downstream tasks like inferring targets. Specifically, CDN trains a *causal graph learner* that maps a pair of observational and interventional datasets to their data-generating mechanisms, and a *differential network classifier* to identify variables whose conditional independencies have changed. We demonstrated that our approach outperforms state-of-the-art methods in perturbation modeling and causal discovery, across seven transcriptomics datasets and a variety of synthetic settings. Overall, these results highlight the potential of our method to design more efficient experiments, by reducing the search space of perturbations.

There are several limitations of this work and directions for future research. First, our implementation does not address cyclic, feedback relationships or time-resolved, dynamic systems – both of which are common in biology. Modeling these cases may require data simulation schemes that can sample from the steady-state of a cyclic system, or architectures that directly predict the parameters of a system of stochastic differential equations (Lorch et al., 2024). In addition, current supervised causal discovery algorithms, which we use for the causal graph learner, tend to assume causal

Table 3: Intervention target prediction results on synthetic datasets with $N = 20$, $E = 40$. Top: hard interventions (uniform); bottom: soft interventions (scale). Neither polynomial nor “scale” soft interventions were seen during training. Number in parentheses indicates number of intervention targets. Uncertainty is standard deviation over 5 i.i.d. datasets. Runtimes in Table 6. Extended results in Table 12.

Type	Model	Linear (1)		Linear (3)		Polynomial (1)		Polynomial (3)	
		mAP \uparrow	AUC \uparrow						
Hard	UT-IGSP	0.10 \pm .01	0.63 \pm .03	0.18 \pm .01	0.52 \pm .03	0.19 \pm .02	0.70 \pm .02	0.23 \pm .03	0.56 \pm .03
	DCDI-G	0.16 \pm .02	0.50 \pm .02	0.27 \pm .05	0.50 \pm .04	0.15 \pm .03	0.45 \pm .04	0.26 \pm .02	0.51 \pm .04
	DCDI-DSF	0.15 \pm .02	0.46 \pm .03	0.26 \pm .03	0.47 \pm .04	0.18 \pm .05	0.48 \pm .08	0.28 \pm .04	0.52 \pm .05
	BACADI-E	0.18 \pm .09	0.71 \pm .06	0.28 \pm .02	0.69 \pm .03	0.20 \pm .06	0.84 \pm .04	0.37 \pm .07	0.82 \pm .05
	BACADI-M	0.10 \pm .02	0.66 \pm .05	0.21 \pm .01	0.63 \pm .02	0.09 \pm .02	0.71 \pm .05	0.24 \pm .03	0.70 \pm .05
	DCI	0.58 \pm .07	0.85 \pm .04	0.63 \pm .06	0.84 \pm .03	0.55 \pm .08	0.82 \pm .04	0.56 \pm .08	0.79 \pm .05
	MB+CI	0.47 \pm .12	0.68 \pm .08	0.58 \pm .09	0.71 \pm .06	0.77 \pm .05	0.88 \pm .03	0.79 \pm .04	0.88 \pm .02
	CDN	0.82 \pm .05	0.96 \pm .02	0.88 \pm .04	0.95 \pm .02	0.85 \pm .05	0.97 \pm .01	0.90 \pm .04	0.96 \pm .02
	CDN (no \mathcal{L}_G)	0.65 \pm .07	0.86 \pm .06	0.72 \pm .09	0.84 \pm .07	0.81 \pm .05	0.95 \pm .02	0.86 \pm .03	0.94 \pm .01
Soft	UT-IGSP	0.10 \pm .01	0.69 \pm .03	0.19 \pm .01	0.56 \pm .03	0.18 \pm .02	0.77 \pm .02	0.22 \pm .01	0.60 \pm .01
	DCDI-G	0.17 \pm .03	0.50 \pm .05	0.23 \pm .01	0.44 \pm .03	0.14 \pm .03	0.48 \pm .03	0.30 \pm .04	0.54 \pm .04
	DCDI-DSF	0.17 \pm .04	0.46 \pm .04	0.24 \pm .02	0.47 \pm .03	0.16 \pm .02	0.50 \pm .05	0.26 \pm .02	0.48 \pm .02
	BACADI-E	0.25 \pm .07	0.65 \pm .03	0.40 \pm .13	0.68 \pm .08	0.43 \pm .22	0.81 \pm .09	0.49 \pm .15	0.77 \pm .05
	BACADI-M	0.11 \pm .01	0.61 \pm .02	0.26 \pm .07	0.64 \pm .06	0.25 \pm .10	0.80 \pm .09	0.38 \pm .10	0.76 \pm .06
	DCI	0.48 \pm .02	0.78 \pm .02	0.53 \pm .06	0.77 \pm .04	0.62 \pm .09	0.89 \pm .04	0.55 \pm .04	0.80 \pm .04
	MB+CI	0.46 \pm .10	0.65 \pm .06	0.49 \pm .04	0.64 \pm .03	0.97 \pm .02	0.98 \pm .01	0.92 \pm .05	0.94 \pm .03
	CDN	0.83 \pm .10	0.95 \pm .05	0.76 \pm .07	0.89 \pm .04	0.97 \pm .03	1.00 \pm .00	0.96 \pm .02	0.99 \pm .01
	CDN (no \mathcal{L}_G)	0.77 \pm .06	0.95 \pm .03	0.81 \pm .08	0.93 \pm .04	0.94 \pm .03	0.99 \pm .00	0.96 \pm .02	0.99 \pm .01
exclude scale	0.32 \pm .07	0.64 \pm .04	0.34 \pm .09	0.65 \pm .08	0.06 \pm .02	0.63 \pm .02	0.08 \pm .02	0.61 \pm .06	
exclude scale	0.44 \pm .04	0.92 \pm .01	0.43 \pm .09	0.90 \pm .03	0.06 \pm .01	0.72 \pm .02	0.10 \pm .02	0.73 \pm .02	

sufficiency, i.e. that there are no unseen confounders. This assumption is unrealistic in biology, as it is not physically possible to measure all relevant variables. Thus, there is an opportunity to extend this work for the latent confounding setting. Finally, since it is difficult to simulate biology, there will be an inevitable domain shift between synthetic and real data. However, there does exist plentiful unlabeled single-cell data (Consortium* et al., 2022), which could be used with self-supervised objectives to bridge the gap (Sun et al., 2020). In conclusion, we hope that this work will enable efficient experimental design and interpretation, as well as future machine learning efforts towards these tasks.

Impact statement

Our work proposes a machine learning method to improve understanding of large-scale perturbation experiments, in the context of molecular biology. Towards broader impact, this work may facilitate better understanding of drug mechanisms; inform protocol design for cellular reprogramming; or identify causal targets of disease. While this work has numerous applications to the life sciences, it does not directly concern the design of molecules or other potentially harmful agents.

Acknowledgements

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1745302. We would like to acknowledge support from the NSF Expeditions grant (award 1918839: Collaborative Research: Understanding the World Through Code), Machine Learning for Pharmaceutical Discovery and Synthesis (MLPDS) consortium, and the Abdul Latif Jameel Clinic for Machine Learning in Health.

References

- Ahlmann-Eltze, C., Huber, W., and Anders, S. Deep learning-based predictions of gene perturbation effects do not yet outperform simple linear methods. *bioRxiv*, 2024. doi: 10.1101/2024.09.16.613342.
- Ashburner, M., Ball, C. A., Blake, J. A., Botstein, D., Butler, H. L., Cherry, J. M., Davis, A. P., Dolinski, K., Dwight, S. S., Eppig, J. T., Harris, M. A., Hill, D. P., Issel-Tarver, L., Kasarskis, A., Lewis, S. E., Matese, J. C., Richardson, J. E., Ringwald, M., and Rubin, G. M. Gene ontology: tool for the unification of biology. *Nature Genetics*, 25: 25–29, 2000.
- Babu, M. M., Luscombe, N. M., Aravind, L., Gerstein, M., and Teichmann, S. A. Structure and evolution of transcriptional regulatory networks. *Current opinion in structural biology*, 14(3):283–291, 2004.
- Bai, D., Ellington, C. N., Mo, S., Song, L., and Xing, E. P. AttentionPert: accurately modeling multiplexed genetic perturbations with multi-scale effects. *Bioinformatics*, 40:i453–i461, 06 2024. ISSN 1367-4811. doi: 10.1093/bioinformatics/btae244.
- Belyaeva, A., Squires, C., and Uhler, C. DCI: learning causal differences between gene regulatory networks. *Bioinformatics*, 37(18):3067–3069, 03 2021. ISSN 1367-4803. doi: 10.1093/bioinformatics/btab167.
- Benjamini, Y. and Hochberg, Y. On the adaptive control of the false discovery rate in multiple testing with independent statistics. *Journal of Educational and Behavioral Statistics*, 25(1):60–83, 2000. ISSN 10769986, 19351054.
- Blake, W. J., Kærn, M., Cantor, C. R., and Collins, J. J. Noise in eukaryotic gene expression. *Nature*, 422(6932): 633–637, 2003.
- Brouillard, P., Lachapelle, S., Lacoste, A., Lacoste-Julien, S., and Drouin, A. Differentiable causal discovery from interventional data. *Advances in Neural Information Processing Systems*, 33:21865–21877, 2020.
- Bunne, C., Stark, S. G., Gut, G., Del Castillo, J. S., Levesque, M., Lehmann, K.-V., Pelkmans, L., Krause, A., and Rättsch, G. Learning single-cell perturbation responses using neural optimal transport. *Nature Methods*, 20(11):1759–1768, 2023.
- Chen, S., Chandra Tjin, C., Gao, X., Xue, Y., Jiao, H., Zhang, R., Wu, M., He, Z., Ellman, J., and Ha, Y. Pharmacological inhibition of pi5p4ka/β disrupts cell energy metabolism and selectively kills p53-null tumor cells. *Proceedings of the National Academy of Sciences*, 118(21):e2002486118, 2021.
- Chen, T., Bello, K., Aragam, B., and Ravikumar, P. iSCAN: identifying causal mechanism shifts among nonlinear additive noise models. *Advances in Neural Information Processing Systems*, 36:44671–44706, 2023.
- Chen, Y. T. and Zou, J. GenePT: A simple but hard-to-beat foundation model for genes and cells built from ChatGPT. *bioRxiv*, 2023. doi: 10.1101/2023.10.16.562533.
- Cherry, A. and Daley, G. Reprogramming cellular identity for regenerative medicine. *Cell*, 148(6):1110–1122, 2012. ISSN 0092-8674. doi: <https://doi.org/10.1016/j.cell.2012.02.031>.
- Chickering, D. M. Optimal structure identification with greedy search. *Journal of Machine Learning Research*, 3: 507–554, November 2002.
- Consortium*, T. T. S., Jones, R. C., Karkanas, J., Krasnow, M. A., Pisco, A. O., Quake, S. R., Salzman, J., Yosef, N., Bulthaupt, B., Brown, P., et al. The Tabula Sapiens: A multiple-organ, single-cell transcriptomic atlas of humans. *Science*, 376(6594):eabl4896, 2022.
- Cosgrove, E. J., Zhou, Y., Gardner, T. S., and Kolaczyk, E. D. Predicting gene targets of perturbations via network-based filtering of mRNA expression compendia. *Bioinformatics*, 24(21):2482–2490, 09 2008. ISSN 1367-4803. doi: 10.1093/bioinformatics/btn476.
- Cui, H., Wang, C., Maan, H., Pang, K., Luo, F., Duan, N., and Wang, B. scGPT: toward building a foundation model for single-cell multi-omics using generative AI. *Nature Methods*, pp. 1–11, 2024.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*, pp. 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423.
- di Bernardo, D., Thompson, M. J., Gardner, T. S., Chobot, S. E., Eastwood, E. L., Wojtovich, A. P., Elliott, S. J., Schaus, S. E., and Collins, J. J. Chemogenomic profiling on a genome-wide scale using reverse-engineered gene networks. *Nature Biotechnology*, 23(3):377–383, 2005.
- Dixit, A., Parnas, O., Li, B., Chen, J., Fulco, C. P., Jerby-Arnon, L., Marjanovic, N. D., Dionne, D., Burks, T., Raychowdhury, R., Adamson, B., Norman, T. M., Lander, E. S., Weissman, J. S., Friedman, N., and Regev, A. Perturb-seq: Dissecting molecular circuits with scalable single-cell RNA profiling of pooled genetic screens. *Cell*, 167(7):1853–1866.e17, 2016. ISSN 0092-8674. doi: <https://doi.org/10.1016/j.cell.2016.11.038>.

- Ghassami, A., Kiyavash, N., Huang, B., and Zhang, K. Multi-domain causal structure learning in linear systems. *Advances in Neural Information Processing Systems*, 31, 2018.
- Gonzalez, G., Herath, I., Veselkov, K., Bronstein, M., and Zitnik, M. Combinatorial prediction of therapeutic perturbations using causally-inspired neural networks. *bioRxiv*, 2024. doi: 10.1101/2024.01.03.573985.
- Grandi, F. C., Modi, H., Kampman, L., and Corces, M. R. Chromatin accessibility profiling by ATAC-seq. *Nature Protocols*, 17(6):1518–1552, 2022.
- Hägele, A., Rothfuss, J., Lorch, L., Somnath, V. R., Schölkopf, B., and Krause, A. BaCaDI: Bayesian causal discovery with unknown interventions. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206, pp. 1411–1436. PMLR, 25–27 Apr 2023.
- Hao, M., Gong, J., Zeng, X., Liu, C., Guo, Y., Cheng, X., Wang, T., Ma, J., Song, L., and Zhang, X. Large-scale foundation model on single-cell transcriptomics. *Nature Methods*, 2024. doi: 10.1038/s41592-024-02305-7.
- Hauser, A. and Bühlmann, P. Characterization and greedy learning of interventional markov equivalence classes of directed acyclic graphs. *Journal of Machine Learning Research*, 13(79):2409–2464, 2012.
- Hetzel, L., Boehm, S., Kilbertus, N., Günnemann, S., Theis, F., et al. Predicting cellular responses to novel drug perturbations at a single-cell resolution. *Advances in Neural Information Processing Systems*, 35:26711–26722, 2022.
- Hilberg, F., Roth, G. J., Krssak, M., Kautschitsch, S., Sommergruber, W., Tontsch-Grunt, U., Garin-Chesa, P., Bader, G., Zoephel, A., Quant, J., et al. BIBF 1120: triple angiokinase inhibitor with sustained receptor blockade and good antitumor efficacy. *Cancer Research*, 68(12):4774–4782, 2008.
- Ho, J., Kalchbrenner, N., Weissenborn, D., and Salimans, T. Axial attention in multidimensional transformers. *arXiv preprint arXiv:1912.12180*, 2019.
- Huang, B., Zhang, K., Zhang, J., Ramsey, J., Sanchez-Romero, R., Glymour, C., and Schölkopf, B. Causal discovery from heterogeneous/nonstationary data. *Journal of Machine Learning Research*, 21(89):1–53, 2020.
- Huang, K., Lopez, R., Hütter, J.-C., Kudo, T., Rios, A., and Regev, A. Sequential optimal experimental design of perturbation screens guided by multi-modal priors. In *International Conference on Research in Computational Molecular Biology*, pp. 17–37. Springer, 2024.
- Huttlin, E. L., Bruckner, R. J., Navarrete-Perea, J., Cannon, J. R., Baltier, K., Gebreab, F., Gygi, M. P., Thornock, A., Zarraga, G., Tam, S., et al. Dual proteome-scale networks reveal cell-specific remodeling of the human interactome. *Cell*, 184(11):3022–3040, 2021.
- Isik, Z., Baldow, C., Cannistraci, C. V., and Schroeder, M. Drug target prioritization by perturbed gene expression and network information. *Scientific Reports*, 5(1):17417, 2015.
- Jaber, A., Kocaoglu, M., Shanmugam, K., and Bareinboim, E. Causal discovery from soft interventions with unknown targets: Characterization and learning. *Advances in Neural Information Processing Systems*, 33: 9551–9561, 2020.
- Ke, N. R., Bilaniuk, O., Goyal, A., Bauer, S., Larochelle, H., Pal, C., and Bengio, Y. Learning neural causal models from unknown interventions. *arXiv preprint arXiv:1910.01075*, 2019.
- Ke, N. R., Chiappa, S., Wang, J. X., Bornschein, J., Goyal, A., Rey, M., Weber, T., Botvinick, M., Mozer, M. C., and Rezende, D. J. Learning to induce causal structure. In *International Conference on Learning Representations*, 2023.
- Kernfeld, E., Yang, Y., Weinstock, J. S., Battle, A., and Cahan, P. A systematic comparison of computational methods for expression forecasting. *bioRxiv*, 2023. doi: 10.1101/2023.07.28.551039.
- Knox, C., Wilson, M., Klinger, C. M., Franklin, M., Oler, E., Wilson, A., Pon, A., Cox, J., Chin, N. E., Strawbridge, S. A., et al. DrugBank 6.0: the DrugBank knowledgebase for 2024. *Nucleic Acids Research*, 52(D1):D1265–D1275, 2024.
- Lachapelle, S., Brouillard, P., Deleu, T., and Lacoste-Julien, S. Gradient-based neural DAG learning. In *International Conference on Learning Representations*, 2020.
- Lamb, J., Crawford, E. D., Peck, D., Modell, J. W., Blat, I. C., Wrobel, M. J., Lerner, J., Brunet, J.-P., Subramanian, A., Ross, K. N., Reich, M., Hieronymus, H., Wei, G., Armstrong, S. A., Haggarty, S. J., Clemons, P. A., Wei, R., Carr, S. A., Lander, E. S., and Golub, T. R. The connectivity map: Using gene-expression signatures to connect small molecules, genes, and disease. *Science*, 313:1929 – 1935, 2006.
- Li, H., Xiao, Q., and Tian, J. Supervised whole DAG causal discovery. *arXiv preprint arXiv:2006.04697*, 2020.
- Lorch, L., Sussex, S., Rothfuss, J., Krause, A., and Schölkopf, B. Amortized inference for causal structure learning. *Advances in Neural Information Processing Systems*, 35:13104–13118, 2022.

- Lorch, L., Krause, A., and Schölkopf, B. Causal modeling with stationary diffusions. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238, pp. 1927–1935. PMLR, 02–04 May 2024.
- Loshchilov, I. and Hutter, F. Decoupled weight decay regularization, 2019.
- Lotfollahi, M., Wolf, F. A., and Theis, F. J. scGen predicts single-cell perturbation responses. *Nature Methods*, 16: 715 – 721, 2019.
- Lotfollahi, M., Klimovskaia Susmelj, A., De Donno, C., Hetzel, L., Ji, Y., Ibarra, I. L., Srivatsan, S. R., Naghipourfar, M., Daza, R. M., Martin, B., et al. Predicting cellular responses to complex perturbations in high-throughput screens. *Molecular Systems Biology*, pp. e11517, 2023.
- Love, M. I., Huber, W., and Anders, S. Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2. *Genome Biology*, 15(12):550–550, 2014. doi: 10.1186/s13059-014-0550-8.
- Luck, K., Kim, D.-K., Lambourne, L., Spirohn, K., Begg, B. E., Bian, W., Brignall, R., Cafarelli, T., Campos-Laborie, F. J., Charlotteaux, B., et al. A reference map of the human binary protein interactome. *Nature*, 580 (7803):402–408, 2020.
- Luecken, M. D. and Theis, F. J. Current best practices in single-cell RNA-seq analysis: a tutorial. *Molecular systems biology*, 15(6):e8746, 2019.
- Malik, V., Bello, K., Ghoshal, A., and Honorio, J. Identifying causal changes between linear structural equation models. In *The 40th Conference on Uncertainty in Artificial Intelligence*, 2024.
- Märtens, K., Donovan-Maiye, R., and Ferkinghoff-Borg, J. Enhancing generative perturbation models with LLM-informed gene embeddings. In *ICLR 2024 Workshop on Machine Learning for Genomics Explorations*, 2024.
- McFaline-Figueroa, J. L., Srivatsan, S., Hill, A. J., Gasperini, M., Jackson, D. L., Saunders, L., Domcke, S., Regalado, S. G., Lazarchuck, P., Alvarez, S., Monnat, R. J., Shendure, J., and Trapnell, C. Multiplex single-cell chemical genomics reveals the kinase dependence of the response to targeted therapy. *Cell Genomics*, 4(2):100487, 2024. ISSN 2666-979X. doi: <https://doi.org/10.1016/j.xgen.2023.100487>.
- Montagna, F., Cairney-Leeming, M., Sridhar, D., and Locatello, F. Demystifying amortized causal discovery with transformers. *arXiv preprint arXiv:2405.16924*, 2024.
- Mooij, J. M., Magliacane, S., and Claassen, T. Joint causal inference from multiple contexts. *J. Mach. Learn. Res.*, 21(1), Jan 2020. ISSN 1532-4435.
- Nadig, A., Replogle, J. M., Pogson, A. N., McCarroll, S. A., Weissman, J. S., Robinson, E. B., and O’Connor, L. J. Transcriptome-wide characterization of genetic perturbations. *bioRxiv*, 2024. doi: 10.1101/2024.07.03.601903.
- Newman, J. R., Ghaemmaghami, S., Ihmels, J., Breslow, D. K., Noble, M., DeRisi, J. L., and Weissman, J. S. Single-cell proteomic analysis of *S. cerevisiae* reveals the architecture of biological noise. *Nature*, 441(7095): 840–846, 2006.
- Noh, H. and Gunawan, R. Inferring gene targets of drugs and chemical compounds from gene expression profiles. *Bioinformatics*, 32(14):2120–2127, 2016.
- Norman, T. M., Horlbeck, M. A., Replogle, J. M., Ge, A. Y., Xu, A., Jost, M., Gilbert, L. A., and Weissman, J. S. Exploring genetic interaction manifolds constructed from rich single-cell phenotypes. *Science*, 365(6455):786–793, 2019. doi: 10.1126/science.aax4438.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Petersen, A. H., Ramsey, J., Ekstrøm, C. T., and Spirtes, P. Causal discovery for observational sciences using supervised machine learning. *Journal of Data Science*, 21(2):255–280, 2023. ISSN 1680-743X. doi: 10.6339/23-JDS1088.
- Rao, R. M., Liu, J., Verkuil, R., Meier, J., Canny, J., Abbeel, P., Sercu, T., and Rives, A. MSA transformer. In *International Conference on Machine Learning*, pp. 8844–8856. PMLR, 2021.
- Replogle, J. M., Saunders, R. A., Pogson, A. N., Hussmann, J. A., Lenail, A., Guna, A., Mascibroda, L., Wagner, E. J., Adelman, K., Lithwick-Yanai, G., Iremadze, N., Oberstrass, F., Lipson, D., Bonnar, J. L., Jost, M., Norman, T. M., and Weissman, J. S. Mapping information-rich genotype-phenotype landscapes with genome-scale Perturb-seq. *Cell*, 185(14):2559–2575, Jul 2022.
- Roohani, Y., Huang, K., and Leskovec, J. Predicting transcriptional outcomes of novel multigene perturbations with gears. *Nature Biotechnology*, 2023.
- Roohani, Y. H., Vora, J., Huang, Q., Liang, P., and Leskovec, J. BioDiscoveryAgent: An AI agent for designing genetic perturbation experiments. In *International Conference on Learning Representations*, 2025.

- Sanchez, P., Liu, X., O’Neil, A. Q., and Tsafaris, S. A. Diffusion models for causal discovery via topological ordering. In *The Eleventh International Conference on Learning Representations*, 2023.
- Schenone, M., Dancík, V., Wagner, B. K., and Clemons, P. A. Target identification and mechanism of action in chemical biology and drug discovery. *Nature Chemical Biology*, 9 4:232–40, 2013.
- Schneider, N., Lorch, L., Kilbertus, N., Schölkopf, B., and Krause, A. Generative intervention models for causal perturbation modeling. *arXiv preprint arXiv:2411.14003*, 2024.
- Shimizu, S., Hoyer, P. O., Hyvarinen, A., and Kerminen, A. A linear non-gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(72):2003–2030, 2006.
- Spirites, P., Meek, C., and Richardson, T. Causal inference in the presence of latent variables and selection bias. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, UAI’95, pp. 499–506, San Francisco, CA, USA, 1995. Morgan Kaufmann Publishers Inc. ISBN 1558603859.
- Spirites, P., Glymour, C., and Scheines, R. *Causation, Prediction, and Search*. MIT Press, 2001. doi: <https://doi.org/10.7551/mitpress/1754.001.0001>.
- Squires, C., Wang, Y., and Uhler, C. Permutation-based causal structure learning with unknown intervention targets. In Peters, J. and Sontag, D. (eds.), *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI)*, volume 124, pp. 1039–1048. PMLR, 03–06 Aug 2020.
- Srivatsan, S. R., McFaline-Figueroa, J. L., Ramani, V., Saunders, L., Cao, J., Packer, J., Pliner, H. A., Jackson, D. L., Daza, R. M., Christiansen, L., Zhang, F., Steemers, F., Shendure, J., and Trapnell, C. Massively multiplex chemical transcriptomics at single-cell resolution. *Science*, 367 (6473):45–51, 2020. doi: 10.1126/science.aax6234.
- Su, J., Ahmed, M., Lu, Y., Pan, S., Bo, W., and Liu, Y. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.
- Sun, Y., Wang, X., Liu, Z., Miller, J., Efros, A., and Hardt, M. Test-time training with self-supervision for generalization under distribution shifts. In *International Conference on Machine Learning*, pp. 9229–9248. PMLR, 2020.
- Theodoris, C. V., Xiao, L., Chopra, A., Chaffin, M. D., Al Sayed, Z. R., Hill, M. C., Mantineo, H., Brydon, E. M., Zeng, Z., Liu, X. S., et al. Transfer learning enables predictions in network biology. *Nature*, 618(7965):616–624, 2023.
- Tian, J. and Pearl, J. Causal discovery from changes. In *Proceedings of the Seventeenth conference on Uncertainty in artificial intelligence*, pp. 512–521, 2001.
- Varici, B., Shanmugam, K., Sattigeri, P., and Tajer, A. Intervention target estimation in the presence of latent variables. In *Uncertainty in Artificial Intelligence*, pp. 2013–2023. PMLR, 2022.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. Attention is all you need. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- Wang, Y., Squires, C., Belyaeva, A., and Uhler, C. Direct estimation of differences in causal graphs. *Advances in Neural Information Processing Systems*, 31, 2018.
- Wilcoxon, F. Individual comparisons by ranking methods. *Biometrics Bulletin*, 1(6):80–83, 1945. ISSN 00994987.
- Wolf, F. A., Angerer, P., and Theis, F. J. Scanpy: large-scale single-cell gene expression data analysis. *Genome Biology*, 19:1–5, 2018.
- Wu, M., Bao, Y., Barzilay, R., and Jaakkola, T. Sample, estimate, aggregate: A recipe for causal discovery foundation models. *Transactions on Machine Learning Research*, 2025.
- Wu, Y., Barton, R., Wang, Z., Ioannidis, V. N., Donno, C. D., Price, L. C., Voloch, L. F., and Karypis, G. Predicting cellular responses with variational causal inference and refined relational information. In *The Eleventh International Conference on Learning Representations*, 2023.
- Yang, Y., Salehkaleybar, S., and Kiyavash, N. Learning unknown intervention targets in structural causal models from heterogeneous data. In Dasgupta, S., Mandt, S., and Li, Y. (eds.), *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238, pp. 3187–3195. PMLR, 02–04 May 2024.
- Yun, C., Bhojanapalli, S., Rawat, A. S., Reddi, S. J., and Kumar, S. Are Transformers universal approximators of sequence-to-sequence functions? *CoRR*, abs/1912.10077, 2019.
- Zhang, J., Cammarata, L., Squires, C., Sapsis, T. P., and Uhler, C. Active learning for optimal intervention design in causal models. *Nature Machine Intelligence*, 5:1066–1075, 2023a.
- Zhang, J., Greenewald, K., Squires, C., Srivastava, A., Shanmugam, K., and Uhler, C. Identifiability guarantees for

causal disentanglement from soft interventions. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023b.

Zheng, X., Aragam, B., Ravikumar, P. K., and Xing, E. P. DAGs with no tears: Continuous optimization for structure learning. *Advances in Neural Information Processing Systems*, 31, 2018.

A. Model architecture

A.1. Attention layers

To balance expressivity and efficiency, our model is designed around series of “axial” attention layers, similar to Ho et al. (2019); Rao et al. (2021); Wu et al. (2025). The causal structure learner closely resembles the architecture of SEA (Wu et al., 2025), while the differential network is slightly different (since the input space is different). Briefly, the input to the causal structure learner is a matrix of shape $N \times N \times T \times d$. However, the (dense) global statistic does not span the T dimension, and the local graph estimates are generally sparse. Thus, the causal graph learner separately attends over the $N \times N \times d$ and $K \times T \times d$ components (K is the number of unique edges in sparse format), with “messages” to align the representations between each attention layer. In contrast, the differential network only operates over the dense $N \times (N + 1) \times d$ (or $2d$) representation. Thus, it omits the message passing layers and directly stacks dense attention layers.

A.2. Edge embeddings

Since attention itself is invariant to input ordering, Transformer-style architectures use absolute (Devlin et al., 2019) or relative (Su et al., 2024) positional embeddings to distinguish input elements. Images and sequences have a natural ordering among their elements. This structure inherently induces distances between positions, which provides signal for learning models. In contrast, models over graphs should be equivariant to the labeling of nodes. That is, if $1 \rightarrow 2$ is an edge, while 2 and 3 are not connected, then re-labeling the nodes $(1, 2, 3) \mapsto (2, 3, 1)$ should result in edge $2 \rightarrow 3$ and no edge between 3 and 1.

If we assign an arbitrary ordering to each graph, the Transformer will be able to distinguish between nodes. However, there are several practical issues. Since synthetic data are generated in the topological ordering of the nodes, always assigning the root nodes to low positions may lead to unintended information leakage. Furthermore, if the model is trained on graphs of up to $N = 100$, but tested on larger graphs, the positional embeddings associated with $N > 100$ will be entirely random and out of distribution. Following Wu et al. (2025), we randomly sample a permutation of the maximum graph size to each graph. Each edge embedding is the concatenation of its constituent nodes’ embeddings, combined via a feed-forward network:

$$\text{Embed}(i, j) = \text{FFN}([\text{Embed}(i), \text{Embed}(j)]) \quad (9)$$

where i, j are arbitrarily assigned, but consistent within the graph. This embedding is added to the hidden representation $h_{i,j}$ before the first layer of attention.

In the differential network, the additional “column” associated with node-level statistics is not subject to any embedding (which implicitly differentiates it from the remaining inputs), while the “row” identifies the node these statistics were computed for.

B. Implementation details

B.1. CDN training

We trained until convergence (no improvement for 50 epochs on synthetic; 10 epochs on Perturb-seq) and selected the best model based on validation mAP (Figure 4). The validation distribution was identical to training and did *not* contain test causal mechanisms. During training, we used 15 CPU workers (primarily for local graph estimates) and 1 A6000 GPU. CDN took around 6 hours to train on synthetic data, and 1 hour to finetune on real data. Unless otherwise noted, the causal structure learner was initialized from the SEA inverse covariance, FCI pretrained weights (Wu et al., 2025).

B.2. Baselines

We used the latest releases of all baselines.

GEARS is a graph neural network that predicts the effects of unseen genetic perturbations, based on the Gene Ontology knowledge graph (Ashburner et al., 2000). We trained GEARS on each cell line separately, and predicted the effects of perturbing every expressed gene. Then we ranked candidates based on cosine similarity to the interventional data. Figure 9 depicts the low, but non-trivial proportion of differentially expressed genes for which GEARS was unable to make a prediction, due to lack of node coverage in their processed gene ontology graph (Ashburner et al., 2000). These genes were not considered in the rankings or evaluations for GEARS.

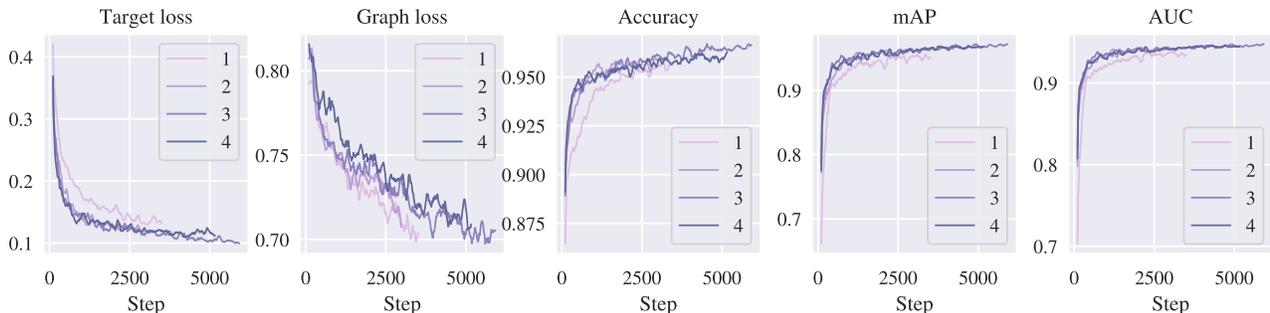


Figure 4: Hyperparameter search: number of differential network layers. We selected 3 layers based on validation metrics.

GENEPT utilized the v2 March 2024 update, which used newer models and additional protein data, compared to their initial paper (GenePT_gene_protein_embedding_model_3_text). We concatenated each perturbation’s gene embedding to the log-fold change of the top 5000 highly-variable genes (Wolf et al., 2018) and trained a logistic regression model on each cell line to predict true vs. decoy perturbations. Candidates were ranked based on predicted probability. Of the 2,842 unique genetic perturbations, only 2,819 (99.1%) mapped to GENEPT embeddings. The remainder used the mean gene embedding (within the dataset) as the language-based embedding.

PDGRAPHER was published on the union of three distinct knowledge graphs, but their harmonized graphs were not available publicly, and the authors did not respond to requests for data sharing. As a result, we relied on solely the human reference interactome (Luck et al., 2020) graph, as it was the only one of the three that could be easily processed. All test perturbation targets could be inferred through PDGRAPHER.

UT-IGSP used the official implementation from commit 7349396, since the repository is not back-compatible.

BACADI is evaluated using the fully-connected implementation, since it performed better than the linear version. However, since the linear version is significantly faster, we include its runtime in Table 6.

MB+CI was implemented using scikit-learn (Pedregosa et al., 2011).

C. Theoretical motivation

While this paper focuses on biological applications, we show that the differential network is well-specified as a class of models, i.e. a correct solution is contained within the space of possible parametrizations.

Preliminary note In this analysis, we assume that the “causal” representations h contains enough information both to retain global statistics and recover the true graph E . The former is reasonable due to high model capacity, and the latter is based on high empirical performance in graph reconstruction (Wu et al., 2025). We emphasize that the latter is an empirical judgment, which may not hold on all datasets in practice.

Attention-based architecture Our differential network is implemented using an attention layer, which is composed of two self-attention layers (one along each axis of the adjacency matrix) and a feed-forward network. For simplicity, we follow prior work (Yun et al., 2019) and ignore layer normalization and dropout.

Our inputs $h \in \mathbb{R}^{2d \times N \times N}$ are $2d$ -dimension features, which represent a pair of $N \times N$ causal graphs. We use $h_{\cdot,j}$ to denote a length N row for a fixed column j , and $h_{i,\cdot}$ to denote a length N column for a fixed row i . The attention layer implements:

$$\begin{aligned} \text{Attn}_{\text{row}}(h_{\cdot,j}) &= h_{\cdot,j} + W_O W_V h_{\cdot,j} \cdot \sigma \left[(W_K h_{\cdot,j})^T W_Q h_{\cdot,j} \right], \\ \text{Attn}_{\text{col}}(h_{i,\cdot}) &= h_{i,\cdot} + W_O W_V h_{i,\cdot} \cdot \sigma \left[(W_K h_{i,\cdot})^T W_Q h_{i,\cdot} \right], \\ \text{FFN}(h) &= h + W_2 \cdot \text{ReLU}(W_1 \cdot h + b_1) + b_2, \end{aligned}$$

where $W_O \in \mathbb{R}^{2d \times 2d}$, $W_V, W_K, W_Q \in \mathbb{R}^{2d \times 2d}$, $W_2 \in \mathbb{R}^{2d \times m}$, $W_1 \in \mathbb{R}^{m \times 2d}$, $b_2 \in \mathbb{R}^{2d}$, $b_1 \in \mathbb{R}^m$, and m is the FFN hidden dimension. We have omitted the i and j subscripts on the W s, but they use separate parameters. Any self attention can take on the identity mapping by setting W_O, W_V, W_K, W_Q to $2d \times 2d$ matrices of zeros.

Hard interventions Let $G = (V, E)$ be a causal graphical model associated with data distribution P_X . Let $G' = (V, E')$ and \tilde{P}_X denote the causal graph and data distribution after an unknown intervention, with ground truth targets $I \subsetneq V$. For convenience, we use E, E' both to denote sets of edges, as well as the equivalent adjacency matrices.

In the case of perfect interventions,

$$f(x_i) \leftarrow z_i, \forall i \in I \quad (10)$$

where $z_i \perp\!\!\!\perp X$ are independent random variables. \tilde{P}_X is associated with mutilated graph E' , where

$$E' = E \setminus \bigcup_{i \in I} \{(j, i)\}_{(j, i) \in E}. \quad (11)$$

In terms of the associated adjacency matrices, $E - E'$ has 1s in each column $i \in I$ and 0s elsewhere.

Here, the attention layer should implement $h - h'$, so that when we collapse over the incoming edges, the output is non-zero only at the edge differences. Suppose the first dimension of the $2d$ feature stores E , and the second dimension stores E' . The row self-attention implements the identity (in the first two dimensions). Then we can set $W_{K, Q}$ to zero, W_V to the identity, and W_O to

$$W_O = \begin{bmatrix} 1 & -1 \\ 0 & 0 \end{bmatrix} - \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad (12)$$

to account for the residual. The FFN implements the identity, so that when we take the mean over along the rows, we recover non-zero elements at all nodes whose incoming edges were removed.

Soft interventions We also study soft interventions the context of causal models with non-multiplicative noise, in which intervention targets are scaled by constant factors,

$$f(x_i) \leftarrow c_i f(x), c_i > 0 \quad (13)$$

where c_i are sampled at random per synthetic dataset. This choice is inspired by the fact that biological perturbation effects are measured in fold-change. Here, the adjacency matrices are the same, but global statistics differ. In particular, we focus on two statistics: the correlation matrix R and the covariance matrix Σ . Note that while we do not explicitly provide the covariance matrix as input, we do provide marginal variances, from which the same information can be derived.

Suppose x is an intervention target.

- $R - R'$ is non-zero in all entries i, j and j, i where i is a descendent of x , and j is any node for which $R_{i, j} \neq 0$ (e.g. ancestors, descendants, and x , if P_X is faithful to G).
- $\Sigma - \Sigma'$ is non-zero in all entries i, j and j, i where i is a descendent of x or $i = x$, and j is any node for which $\Sigma_{i, j} \neq 0$ (e.g. ancestors, descendants, and x , if P_X is faithful to G).

All descendants are always affected, due to the non-multiplicative noise term. These two differ in the row and column that correspond to x since

$$\text{Corr}(c \cdot x, y) = \text{Corr}(x, y) \quad (14)$$

$$\text{Cov}(c \cdot x, y) = c \cdot \text{Cov}(x, y). \quad (15)$$

Therefore, to identify x , we should find the index in which Σ differs but not R . Suppose that dimensions 3-6 of h encode R, R', Σ, Σ' . Following the same strategy as the hard interventions, we can use the row attention to compute $R - R', \Sigma - \Sigma'$ and store them in dimensions 3, 4. Then we use the column attention to filter out variables that are independent from x by storing the sum of each column in dimensions 5, 6. While not strictly impossible, it is unlikely that a variable dependent on x would result in a column that sums to exactly 0. Thus, all columns with non-zero sums are either ancestors, descendants, or x . The feedforward network implements

$$\text{FFN}(h_{\cdot, 3-6}) = \begin{cases} 1 & h_{\cdot, 3} = 0, h_{\cdot, 4} \neq 0, h_{\cdot, 5} \neq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (16)$$

This results in 1s in the rows and columns where ΔR and $\Delta \Sigma$ differ. After collapsing over incoming edges and normalizing to probabilities, the maximum probabilities can be found at the intervention targets.

Supporting both intervention types Recall that the final output layer is a linear projection from $2d$ to 1. If this layer implements a simple summation over all $2d$, the predicted intervention targets are consistent with both hard and soft interventions. For soft interventions, $E = E'$, so the hard intervention dimensions will be 0. Likewise, for hard interventions, both R and Σ will differ at the same locations, as the underlying variable has changed, so the soft intervention dimensions will be 0. Since the two techniques produce mutually exclusive predictions, this means that both hard and soft interventions can co-exist and be detected on different nodes.

D. Datasets

D.1. Biological datasets

Data processing We converted all single cell datasets to log-normalized, transcripts per 10,000 UMIs. Perturb-seq genes were mapped to standard identifiers and filtered by the authors. Sci-Plex dataset variables represented genes that appeared in at least 5,000 cells (threshold chosen to achieve a similar number of genes). We performed differential expression analysis via the `scanpy` package (Wolf et al., 2018), using the Wilcoxon signed-rank test (Wilcoxon, 1945) with Benjamini-Hochberg p-value correction and a threshold of adjusted p-value < 0.05 . Table 4 reports statistics of the raw, unprocessed datasets.

For Perturb-seq datasets, we kept perturbations with > 10 differentially-expressed genes (DEGs), and clustered them using k-means with $k = 200$, chosen heuristically based on log-fold change heatmaps (Figure 5). These clusters were used to inform data splits for seen cell lines, where the largest cluster was allocated to the training set, and all remaining clusters were split equally among train and test. The largest cluster(s) appear to contain perturbations with smaller effects. For Sci-Plex datasets, only 6 drug perturbations across 2 cell lines resulted in differential expression of their known protein targets (Supplementary Table 8 from McFaline-Figueroa et al. (2024)). Therefore, we used these exclusively as test sets. Table 5 reports statistics of the final, processed datasets. Figure 8 plots the full distribution of number of cells and DEGs per perturbation.

Top differentially expressed genes Here, we limited our analysis to the top 1000 DEGs. In the literature, it is quite common to restrict analysis to subsets of genes, though the threshold and criteria may vary (Luecken & Theis, 2019). Selecting the top 1000 genes (based on differential expression p-value) is reasonably permissive, since these genes are selected per perturbation. For a sense of scale, Nadig et al. (2024) writes:

“In a genome-scale Perturb-seq screen, we find that a typical gene perturbation affects an estimated 45 genes, whereas a typical essential gene perturbation affects over 500 genes.”

A classic approach for identifying drug targets argues that targets are not highly differentially expressed, but network-based approaches are useful for predicting them from gene expression (Isik et al., 2015). This work also samples 1000 candidate targets as decoys.

Gene co-expression Figures 6 and 7 depict the differences in gene co-expression matrices between control and perturbed cells. Genetic perturbations tend to have much clearer phenotypes compared to chemical perturbations, whose effects are more diffuse.

Sci-Plex drug targets The full names of the Sci-Plex drugs and their targets are as follows. These targets were collated by the authors (McFaline-Figueroa et al., 2024), and are each well-documented in the literature.

- Infigratinib (“infig”) binds target FGFR1 with nano-molar affinities (Knox et al., 2024).
- Nintedanib (“nint”) and target FGFR1 have been probed through structural studies (Hilberg et al., 2008).
- Palbociclib (“palb”) has been co-crystallized (PDB 5L2I) with its target CDK6.
- Doxorubicin (“doxo”) and target TOP2A are associated in a number of works (Knox et al., 2024).
- Volasertib (“vola”) has been co-crystallized (PDB 3FC2) with PLK1 and a number of works confirm specificity for PLK2 (Chen et al., 2021).

Identifying biological perturbation targets through causal differential networks

Table 4: Extended biological dataset statistics (raw).

Type	Source	Accession	Cell line	# Perts	# Genes	# NTCs	# Cells
Genetic	Replogle et al. (2022)	Figshare 20029387	K562 gw	9,866	8,248	75,328	1,989,578
			K562 es	2,057	8,563	10,691	310,385
			RPE1	2,393	8,749	11,485	247,914
	Nadig et al. (2024)	GSE220095	HepG2	2,393	9,624	4,976	145,473
			Jurkat	2,393	8,882	12,013	262,956
Chemical	McFaline-Figueroa et al. (2024)	GSM7056151	A172	23	8,393	8,660	58,347
			T98G	23	8,393	6,921	58,347

Table 5: Extended biological dataset statistics (processed).

Dataset	Perturbations				Genes		Cells
	Train	Test	Trivial	Non-trivial	Unique # DE	Median # DE	
K562 gw	1089	678	587	91	7,378	81	492,096
K562 es	640	420	348	72	8,492	226	213,552
RPE1	564	397	233	164	8,641	399	179,696
HepG2	364	263	162	101	9,282	271	79,309
Jurkat	679	333	262	71	8,432	162	174,698
A172	—	3	—	—	445	324	18,196
T98G	—	3	—	—	1,644	508	13,126

It is important to note that these drugs are also associated with known off-targets. However, we did not observe differential expression of any off-target, adjusted $p \approx 1$, perhaps because the impact on their mRNA expression is negligible. Protein activity would be more suitable for quantifying the impact of a drug on a cell. However, this cannot be measured at scale. In fact, even protein abundance is difficult to measure at single-cell resolution, as it relies on mass spectrometry to differentiate proteins (Newman et al., 2006), while mRNA can be exhaustively enumerated and sequenced (based on the known human genome). Therefore, this chemical perturbation study is primarily a proof of concept, to probe whether targets can be inferred from mRNA expression, and if so, to what extent.

D.2. Synthetic data

Synthetic data were generated using code modified from DCDI (Brouillard et al., 2020) for soft interventions. We used the following implementations for causal mechanisms $f(x)$, where x is the variable in question, M is a binary mask for the parents of x , X contains measurements of all variables, E is independent Gaussian noise, and W is a random weight matrix.

- Linear: $f(x) = MXW + E$.
- Neural network, non-additive $f(x) = \text{Tanh}((X, E)W_{\text{in}})W_{\text{out}}$
- Neural network, additive: $f(x) = \text{Tanh}(XW_{\text{in}})W_{\text{out}} + E$
- Polynomial: $f(x) = W_0 + MXW_1 + MX^2W_2 + E$
- Sigmoid: $f(x) = \sum_{i=1}^d W_i \cdot \text{sigmoid}(X_i) + E$

Linear and the neural network variants were used for training. Polynomial and sigmoid were only used for testing.

Root causal mechanisms are uniform. For hard interventions, we set $f(x) \leftarrow z$, where $z \sim \text{Uniform}(-1, 1)$. For soft “scale” interventions, we set

$$\begin{aligned}
 f(x) &\leftarrow z_1^{\text{Sign}(z)} \cdot f(x) \\
 z_1 &\sim \text{Uniform}(2, 4) \\
 z &\sim \text{Uniform}(-1, 1).
 \end{aligned}$$

That is, we multiply $f(x)$ by a scaling factor that is equal probability ≤ 1 (constant across all observations).

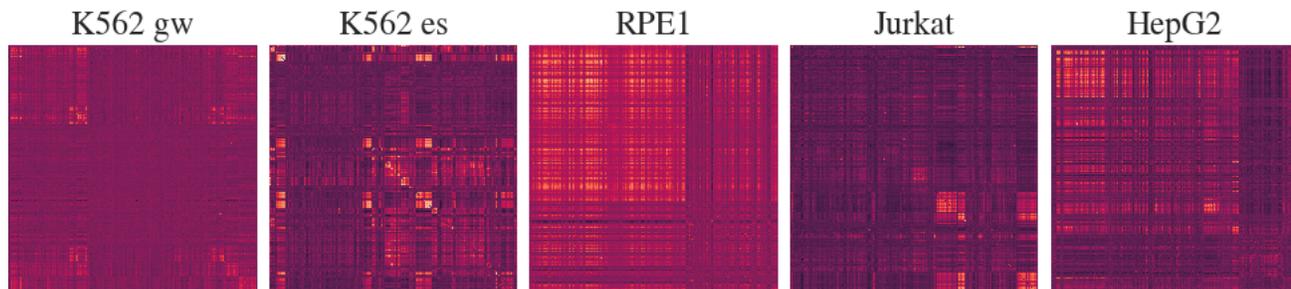


Figure 5: Heatmap of correlation between log-fold change, sorted by cluster, used for data splits.

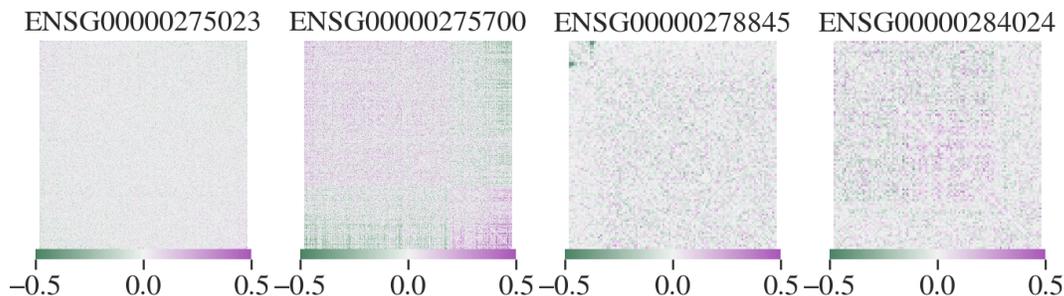


Figure 6: Difference between observational and interventional correlation matrices on K562 genome-wide. In some cases, the changes are minimal and specific; in others, they are more diffuse.

For soft “shift” interventions, we set

$$\begin{aligned} f(x) &\leftarrow f(x) + \text{Sign}(z) \cdot z_1 \\ z_1 &\sim \text{Uniform}(2, 4) \\ z &\sim \text{Uniform}(-1, 1). \end{aligned}$$

That is, we translate $f(x)$ by a scaling factor that is equal probability ≤ 0 (constant across all observations).

E. Additional analyses

Runtime We compared the runtimes of various algorithms on the Perturb-seq and synthetic datasets. On the Perturb-seq data, all models finished running within minutes with the exception of GEARS (Figure 11). Runtimes of the various causal algorithms varied significantly (Table 6). The slowest method was DCDI, with an average runtime of around 10 hours on $N = 20$ datasets, while the fastest were UT-IGSP and the MLP variant of CDN. All models were benchmarked on equivalent hardware (A6000 GPU, 1 CPU core).

Perturb-seq uncertainty quantification Due to space limitations, we report uncertainty estimates in Tables 7 and 8. Multiple baselines produce deterministic results (LINEAR, DGE, GENEPT), so instead of model randomness, we report uncertainties over the sampling of single cells (Table 7). Specifically, for each perturbation with M cells, we sample

$$M' = \min(M, \max(50, 0.8M)) \quad (17)$$

cells uniformly at random, repeated 5 times. We also consider the robustness of our model to samplings of the test set (Table 8). Typically, the largest cluster (assigned to train, e.g. nearly 1/3 of K562) contains weak perturbations with minimal phenotype, so different clusterings don’t introduce significant heterogeneity. Instead, we can partition all perturbations uniformly into 5 sets for the unseen cell line setting, across which the variation in results is minimal.

Perturb-seq trivial perturbations For the main results in Table 1, we focus on genetic perturbations for which the intended gene target was *not* the gene with the greatest log-fold change. This is because genetic perturbations are (unrealistically) specific, compared to drugs or transcription factors. In Table 9, we show that our model performs nearly perfectly in this easy setting.

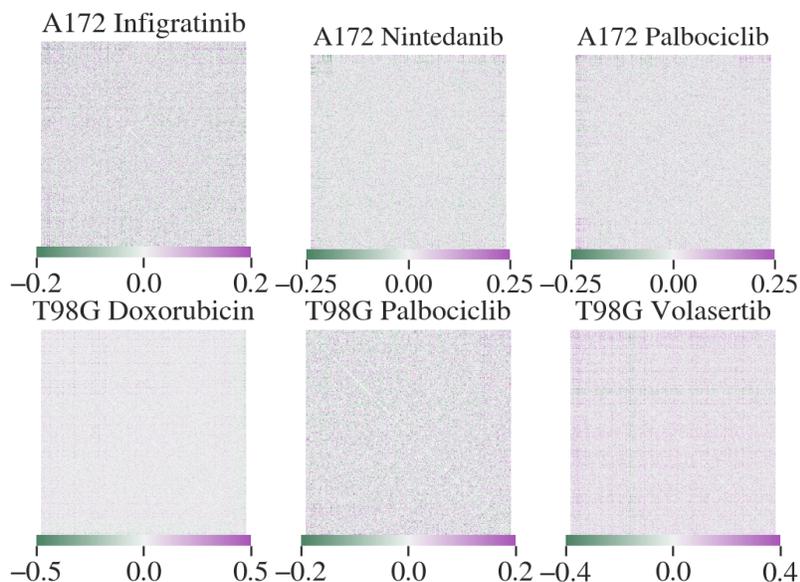


Figure 7: Difference between observational and interventional correlation matrices on Sci-Plex datasets. Effects are much more diffuse than in genetic perturbations (Figure 6).

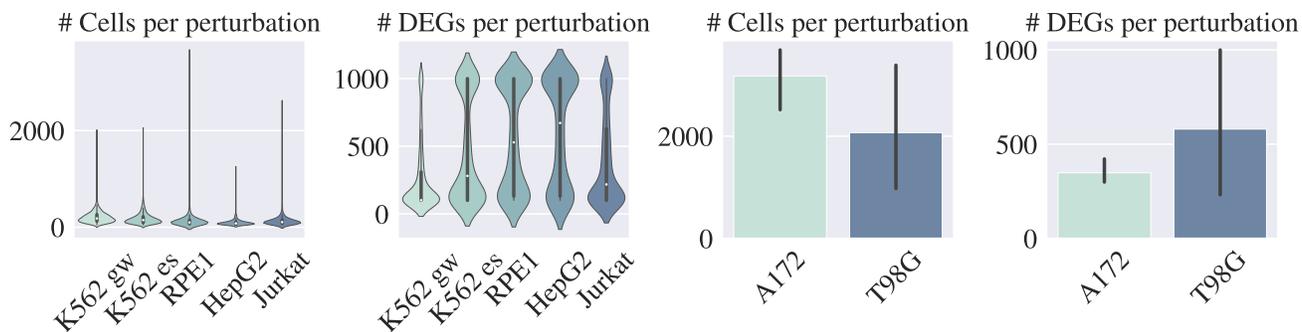


Figure 8: Transcriptomics dataset statistics, after processing. See Table 5 for details.

Graph loss We assess the value of predicting the causal graph, as opposed to only training a supervised classifier (Figure 10). Without modifying the architecture (which has inductive biases towards the graph modality), we train a version of CDN without any graph labels (causal graph learner no longer pretrained; remove graph loss \mathcal{L}_G). Across the validation loss and various metrics, we find that the full CDN is more stable, and leads to better results. This model corresponds to the “no \mathcal{L}_G ” test results in Table 3.

Additional synthetic results During model development, we realized that if the model is allowed to train on all types of perturbations, it can easily overfit and learn these distributional shifts very well. For example, in Table 10, the models trained on both types of soft interventions performs perfectly (unrealistic) in many cases, even though the graphs and causal mechanism parameters are completely novel. Therefore, to assess whether our model can generalize, we chose to hold out a class of soft intervention for our main results in Table 3.

Table 11 reports results on larger graphs. There is some drop in performance, especially Polynomial. This could be attributed to error propagation from the causal graph predictor, which performs less well (at predicting graphs) on larger graphs, and could potentially be ameliorated by finetuning on larger synthetic data (similar to how Perturb-seq models were finetuned to be larger). Linear + Soft remains quite good, perhaps because the theory (Appendix C) suggests that the summary statistics are sufficient for predicting targets, without the graph prediction.

Table 12 reports results on all synthetic test datasets, averaging over all interventions for each dataset. We also include the

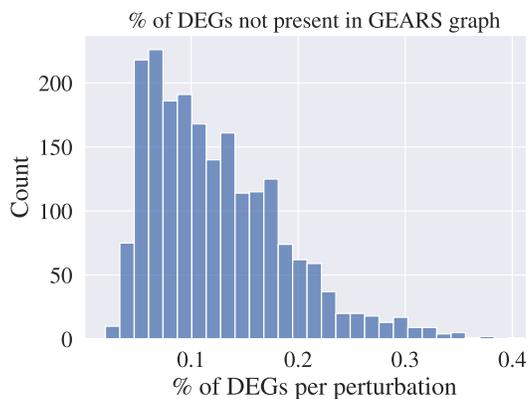


Figure 9: Percentage of differentially expression genes (per perturbation) that were *not* in the GEARS knowledge graph, and whose effects could not be predicted.

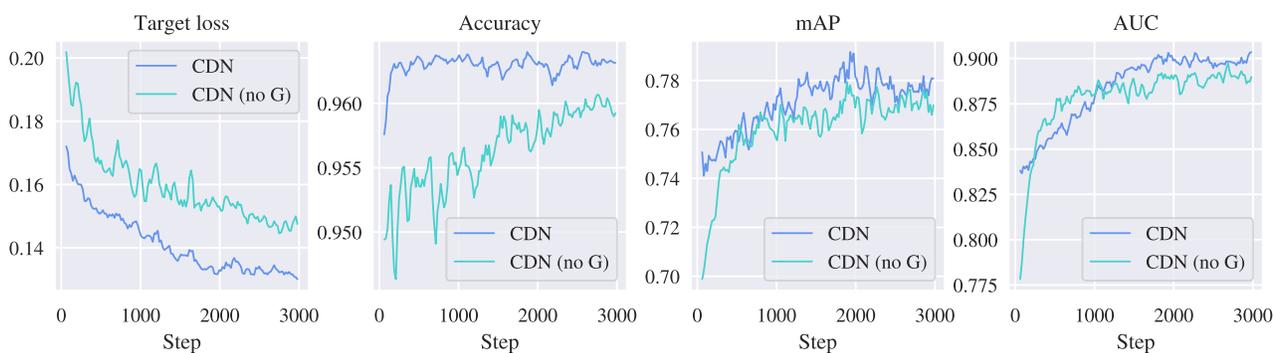


Figure 10: Ablation study: We compare CDN trained on synthetic data, to the same model without any ground truth graph information (no pretraining or \mathcal{L}_G). We find that graph information leads to more stable and better results. Validation metrics and loss. One step is one batch (16 samples).

“full” version of CDN, i.e. trained on all mechanisms, in this analysis.

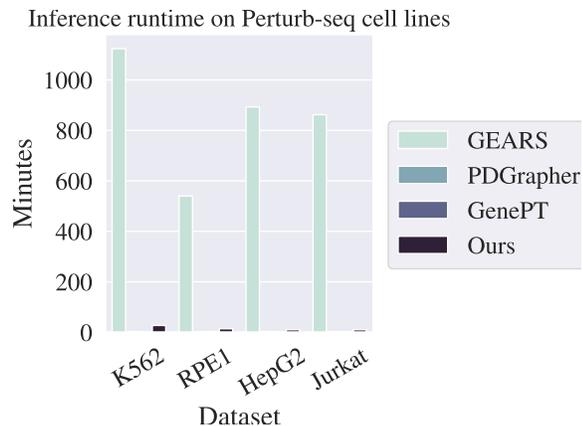


Figure 11: Inference runtimes on Perturb-seq datasets. K562 datasets are reported together. All models run on a single A6000 GPU, no constraint on memory (up to 500G). Only GEARs required over ~ 20 G of memory. To the best of our ability, we normalized batch size to 1.

Table 6: Runtimes on synthetic datasets (sec).

Nodes	Model	Min	Max	Mean	Std
10	UT-IGSP	1	98	26	31
	DCDI-G	31	27908	18731	9857
	DCDI-DSF	7118	44440	23798	5688
	DCI	41	1438	404	365
	BACADI	4005	9483	6284	1584
	BACADI-L	1076	1403	1272	76
	CDN (MLP)	1	5	1	1
	CDN	17	147	39	27
20	UT-IGSP	1	78	19	26
	DCDI-G	45	46281	35032	16068
	DCDI-DSF	23181	55406	30076	7200
	DCI	325	21414	4415	4707
	BACADI	15082	50737	26020	9044
	BACADI-L	2862	3819	3461	252
	CDN (MLP)	1	3	1	0
	CDN	36	86	54	13

Table 7: Uncertainty quantification on Perturb-seq datasets, seen cell lines, by sub-sampling to 80% of cells per perturbation (or 50, whichever is higher). Note that results may be slightly different from Table 1 due to sub-sampling. Standard deviation reported over 5 seeds. GENEPT performance is highly variable. DGE is quite sensitive to sub-sampling on K562 genome-wide.

Model	K562 gw			K562 es			RPE1		
	rank	ρ	r	rank	ρ	r	rank	ρ	r
LINEAR	0.50 \pm .00	0.17 \pm .01	0.20 \pm .01	0.50 \pm .01	0.25 \pm .01	0.28 \pm .01	0.48 \pm .00	0.40 \pm .01	0.49 \pm .01
MLP	0.48 \pm .00	0.17 \pm .00	0.20 \pm .00	0.48 \pm .00	0.16 \pm .00	0.19 \pm .00	0.53 \pm .00	0.30 \pm .00	0.39 \pm .00
DGE	0.72 \pm .01	0.48 \pm .02	0.52 \pm .02	0.86 \pm .00	0.59 \pm .02	0.61 \pm .02	0.74 \pm .00	0.56 \pm .00	0.64 \pm .00
GENEPT	0.54 \pm .09	0.26 \pm .08	0.29 \pm .08	0.50 \pm .08	0.36 \pm .10	0.39 \pm .10	0.51 \pm .10	0.43 \pm .08	0.51 \pm .07
GEARs	0.50 \pm .01	0.19 \pm .01	0.22 \pm .01	0.51 \pm .00	0.20 \pm .01	0.23 \pm .00	0.46 \pm .00	0.33 \pm .01	0.41 \pm .02
PDG	0.49 \pm .01	0.22 \pm .00	0.25 \pm .00	0.50 \pm .01	0.28 \pm .01	0.32 \pm .01	0.49 \pm .01	0.38 \pm .01	0.47 \pm .01
CDN	0.91 \pm .01	0.71 \pm .02	0.72 \pm .02	0.95 \pm .00	0.80 \pm .03	0.81 \pm .03	0.92 \pm .00	0.65 \pm .02	0.74 \pm .01
Model	HepG2			Jurkat					
	rank	ρ	r	rank	ρ	r			
LINEAR	0.49 \pm .00	0.34 \pm .01	0.42 \pm .01	0.50 \pm .00	0.23 \pm .01	0.27 \pm .01			
MLP	0.50 \pm .00	0.34 \pm .00	0.41 \pm .00	0.47 \pm .00	0.25 \pm .00	0.31 \pm .00			
DGE	0.83 \pm .00	0.49 \pm .01	0.50 \pm .01	0.78 \pm .01	0.44 \pm .02	0.49 \pm .02			
GENEPT	0.50 \pm .13	0.39 \pm .07	0.45 \pm .06	0.51 \pm .07	0.36 \pm .11	0.40 \pm .11			
GEARs	0.51 \pm .00	0.38 \pm .00	0.45 \pm .00	0.52 \pm .01	0.24 \pm .00	0.30 \pm .00			
PDG	0.49 \pm .00	0.34 \pm .01	0.41 \pm .01	0.49 \pm .01	0.29 \pm .00	0.34 \pm .00			
CDN	0.93 \pm .00	0.67 \pm .01	0.73 \pm .01	0.96 \pm .01	0.83 \pm .04	0.85 \pm .03			

Table 8: To quantify the uncertainty across draws of the test set (perturbations, instead of cells), we partition all perturbations uniformly into 5 sets for the unseen cell line setting on K562 genome-wide. The variation across test sets is minimal.

Test Split	rank	ρ	r	Recall@1	Recall@5	Recall@20
Original ($n = 678$)	0.974	0.915	0.921	0.699	0.900	0.957
New splits ($n = 353, \times 5$)	0.973 \pm 0.009	0.917 \pm 0.017	0.923 \pm 0.016	0.696 \pm 0.031	0.908 \pm 0.032	0.961 \pm 0.021

Table 9: Results on 5 Perturb-seq datasets, *trivial* perturbations (intended target has largest log-fold change, though not necessarily smallest adjusted p-value). Top: Train on all cell lines jointly. Bottom: Leave one cell line out of training. Test sets (unseen perturbations) are identical in both settings. CDN (synthetic) and DGE do *not* require training on real Perturb-seq data. Metrics, from left to right: normalized rank of ground truth, top 1 Spearman and Pearson correlations.

Setting	Model	K562 (gw)			K562 (es)			RPE1			HepG2			Jurkat		
		rank	ρ	r												
Seen cell line	LINEAR	0.49	0.18	0.21	0.50	0.28	0.31	0.48	0.37	0.45	0.48	0.34	0.43	0.48	0.24	0.29
	MLP	0.48	0.16	0.20	0.47	0.17	0.21	0.54	0.29	0.37	0.51	0.35	0.42	0.48	0.26	0.32
	GENEPT	0.60	0.36	0.39	0.47	0.23	0.26	0.57	0.49	0.55	0.44	0.35	0.41	0.55	0.48	0.53
	GEARS	0.49	0.15	0.18	0.51	0.25	0.27	0.49	0.32	0.40	0.48	0.32	0.40	0.47	0.20	0.24
	PDG	0.49	0.23	0.26	0.52	0.31	0.34	0.51	0.34	0.42	0.46	0.34	0.40	0.49	0.30	0.35
	CDN (Perturb-seq)	0.99	0.97	0.97	0.99	0.98	0.98	0.99	0.93	0.94	0.99	0.95	0.96	0.99	0.96	0.96
Unseen cell line	DGE	0.87	0.68	0.70	0.92	0.77	0.79	0.85	0.73	0.77	0.89	0.59	0.60	0.91	0.73	0.77
	PDG	0.40	0.12	0.15	0.41	0.22	0.24	0.48	0.28	0.37	0.47	0.33	0.41	0.53	0.28	0.33
	CDN (synthetic)	0.91	0.72	0.74	0.92	0.83	0.84	0.94	0.83	0.86	0.90	0.73	0.77	0.95	0.82	0.84
	CDN (Perturb-seq)	0.98	0.95	0.95	0.99	0.95	0.96	0.99	0.92	0.94	0.99	0.95	0.96	0.99	0.96	0.96

Table 10: Intervention target prediction results on synthetic datasets with $N = 20, E = 40$. Top: hard interventions (uniform); bottom: soft interventions (scale). We found that if the model is allowed to see the class of soft intervention (“all mechanisms”) during training, it can easily learn (overfit) the distribution shift perfectly, even though test datasets and graphs are novel. Thus, our main results in Table 3 treat “scale” as an unseen soft intervention.

Type	Setting	Linear (1)		Linear (3)		Polynomial (1)		Polynomial (3)	
		mAP \uparrow	AUC \uparrow						
Hard	all mechanisms (h_{cat})	0.86 \pm .09	0.97 \pm .02	0.91 \pm .06	0.97 \pm .03	0.86 \pm .04	0.97 \pm .01	0.91 \pm .03	0.97 \pm .01
	all mechanisms (h_{diff})	0.78 \pm .05	0.95 \pm .03	0.89 \pm .04	0.96 \pm .02	0.83 \pm .05	0.96 \pm .01	0.89 \pm .03	0.96 \pm .02
	exclude scale	0.82 \pm .05	0.96 \pm .02	0.88 \pm .04	0.95 \pm .02	0.85 \pm .05	0.97 \pm .01	0.90 \pm .04	0.96 \pm .02
Soft	all mechanisms (h_{cat})	1.0 \pm .00	1.0 \pm .00	0.99 \pm .01	1.0 \pm .00				
	all mechanisms (h_{diff})	0.99 \pm .01	1.0 \pm .00	0.98 \pm .01	0.99 \pm .00	1.0 \pm .00	1.0 \pm .00	1.0 \pm .00	1.0 \pm .00
	exclude scale	0.83 \pm .10	0.95 \pm .05	0.76 \pm .07	0.89 \pm .04	0.97 \pm .03	1.0 \pm .00	0.96 \pm .02	0.99 \pm .01

Table 11: While the majority of classic causal discovery models struggle to scale to larger graphs, we show that CDN still performs reasonably on $N = 100, E = 200$ graphs. AUC reported over 10 i.i.d. graphs.

Intervention	Linear (1)	Linear (3)	Sigmoid (1)	Sigmoid (3)	Polynomial (1)	Polynomial (3)
Hard	.64 \pm .04	.65 \pm .08	.63 \pm .02	.61 \pm .06	.50 \pm .01	.49 \pm .03
Soft	.92 \pm .01	.90 \pm .03	.72 \pm .02	.73 \pm .02	.53 \pm .03	.52 \pm .05

Table 12: Intervention target prediction results on synthetic datasets, extended results. Uncertainty is standard deviation over 5 i.i.d. datasets. Metrics are averaged over all $3N$ perturbations for a given dataset (1-3 targets). CDN (full) denotes that we trained on all intervention mechanisms jointly, i.e. as in Table 10.

N	E	Model	Linear (Hard)		Linear (Soft)		Poly. (Hard)		Poly. (Soft)		Sigmoid. (Hard)		Sigmoid. (Soft)	
			mAP \uparrow	AUC \uparrow	mAP \uparrow	AUC \uparrow	mAP \uparrow	AUC \uparrow						
10	10	UT-IGSP	.29 \pm .02	.57 \pm .03	.30 \pm .01	.63 \pm .02	.33 \pm .04	.59 \pm .03	.32 \pm .01	.64 \pm .02	.28 \pm .03	.58 \pm .04	.26 \pm .02	.59 \pm .03
		DCDI-G	.39 \pm .04	.52 \pm .04	.40 \pm .04	.51 \pm .04	.38 \pm .03	.49 \pm .02	.36 \pm .03	.49 \pm .03	.37 \pm .03	.49 \pm .03	.40 \pm .04	.52 \pm .06
		DCDI-DSF	.36 \pm .02	.51 \pm .04	.40 \pm .05	.53 \pm .05	.37 \pm .02	.50 \pm .01	.38 \pm .02	.49 \pm .03	.37 \pm .02	.50 \pm .02	.37 \pm .02	.48 \pm .03
		BACADI-E	.26 \pm .03	.61 \pm .05	.42 \pm .13	.63 \pm .11	.25 \pm .03	.61 \pm .05	.30 \pm .05	.61 \pm .04	.28 \pm .04	.64 \pm .06	.43 \pm .12	.69 \pm .09
		BACADI-M	.23 \pm .01	.56 \pm .03	.37 \pm .11	.62 \pm .11	.22 \pm .02	.56 \pm .04	.30 \pm .05	.61 \pm .04	.25 \pm .03	.61 \pm .05	.42 \pm .11	.69 \pm .09
		DCI	.55 \pm .08	.79 \pm .04	.50 \pm .08	.70 \pm .05	.47 \pm .09	.71 \pm .06	.41 \pm .02	.67 \pm .01	.59 \pm .07	.80 \pm .05	.51 \pm .05	.76 \pm .04
		MB+CI	.53 \pm .05	.64 \pm .03	.57 \pm .10	.66 \pm .06	.70 \pm .09	.79 \pm .07	.85 \pm .07	.89 \pm .05	.84 \pm .08	.88 \pm .06	.96 \pm .05	.97 \pm .04
		CDN	.86 \pm .02	.93 \pm .01	.74 \pm .06	.85 \pm .06	.83 \pm .09	.91 \pm .05	.89 \pm .05	.94 \pm .03	.88 \pm .06	.94 \pm .04	.85 \pm .07	.91 \pm .05
		CDN (no \mathcal{L}_G)	.67 \pm .09	.78 \pm .07	.82 \pm .05	.89 \pm .05	.82 \pm .08	.91 \pm .05	.91 \pm .04	.96 \pm .02	.89 \pm .05	.95 \pm .03	.94 \pm .03	.97 \pm .02
		CDN (full)	.87 \pm .01	.94 \pm .01	.99 \pm .01	1.0 \pm .00	.82 \pm .08	.91 \pm .05	1.0 \pm .00	1.0 \pm .00	.95 \pm .05	.98 \pm .02	1.0 \pm .00	1.0 \pm .00
10	20	UT-IGSP	.28 \pm .02	.56 \pm .03	.26 \pm .01	.59 \pm .03	.28 \pm .03	.58 \pm .04	.27 \pm .01	.60 \pm .01	.25 \pm .00	.55 \pm .03	.25 \pm .02	.56 \pm .05
		DCDI-G	.42 \pm .02	.54 \pm .02	.40 \pm .03	.51 \pm .04	.40 \pm .06	.53 \pm .05	.39 \pm .02	.51 \pm .03	.37 \pm .02	.50 \pm .01	.35 \pm .03	.49 \pm .04
		DCDI-DSF	.41 \pm .03	.52 \pm .03	.39 \pm .05	.52 \pm .05	.39 \pm .02	.51 \pm .02	.39 \pm .03	.50 \pm .05	.36 \pm .03	.50 \pm .05	.36 \pm .03	.49 \pm .03
		BACADI-E	.34 \pm .04	.68 \pm .03	.53 \pm .08	.71 \pm .04	.33 \pm .03	.72 \pm .04	.64 \pm .07	.78 \pm .05	.34 \pm .05	.71 \pm .04	.51 \pm .07	.73 \pm .04
		BACADI-M	.27 \pm .02	.63 \pm .03	.48 \pm .09	.71 \pm .04	.27 \pm .04	.64 \pm .08	.59 \pm .09	.77 \pm .06	.28 \pm .03	.65 \pm .05	.45 \pm .08	.71 \pm .05
		DCI	.59 \pm .03	.78 \pm .03	.57 \pm .04	.77 \pm .03	.68 \pm .07	.82 \pm .04	.65 \pm .08	.84 \pm .05	.70 \pm .03	.84 \pm .02	.66 \pm .06	.83 \pm .05
		MB+CI	.54 \pm .11	.66 \pm .09	.49 \pm .06	.62 \pm .04	.83 \pm .04	.87 \pm .03	.96 \pm .02	.96 \pm .02	.80 \pm .09	.84 \pm .08	.96 \pm .02	.96 \pm .02
		CDN	.88 \pm .04	.95 \pm .02	.56 \pm .07	.71 \pm .07	.96 \pm .02	.98 \pm .01	.95 \pm .04	.97 \pm .02	.94 \pm .04	.98 \pm .02	.68 \pm .02	.79 \pm .02
		CDN (no \mathcal{L}_G)	.79 \pm .06	.90 \pm .04	.77 \pm .04	.86 \pm .02	.94 \pm .01	.97 \pm .01	.95 \pm .05	.96 \pm .04	.87 \pm .02	.93 \pm .02	.94 \pm .02	.96 \pm .01
		CDN (full)	.89 \pm .03	.94 \pm .02	.98 \pm .02	.99 \pm .01	.97 \pm .01	.99 \pm .00	1.0 \pm .00	1.0 \pm .00	.96 \pm .03	.98 \pm .01	.98 \pm .02	.99 \pm .01
20	20	UT-IGSP	.15 \pm .01	.54 \pm .01	.18 \pm .01	.65 \pm .02	.20 \pm .02	.60 \pm .02	.21 \pm .02	.67 \pm .02	.15 \pm .01	.57 \pm .02	.17 \pm .01	.65 \pm .01
		DCDI-G	.24 \pm .02	.49 \pm .02	.22 \pm .02	.49 \pm .02	.22 \pm .03	.50 \pm .03	.22 \pm .02	.49 \pm .03	.22 \pm .02	.50 \pm .04	.22 \pm .02	.50 \pm .03
		DCDI-DSF	.24 \pm .02	.50 \pm .02	.23 \pm .02	.50 \pm .03	.22 \pm .01	.50 \pm .02	.20 \pm .01	.48 \pm .02	.22 \pm .02	.50 \pm .04	.21 \pm .02	.47 \pm .03
		BACADI-E	.14 \pm .01	.61 \pm .03	.25 \pm .10	.60 \pm .06	.22 \pm .06	.76 \pm .06	.30 \pm .06	.74 \pm .05	.15 \pm .02	.65 \pm .03	.17 \pm .05	.57 \pm .05
		BACADI-M	.12 \pm .00	.58 \pm .02	.17 \pm .03	.60 \pm .05	.14 \pm .02	.66 \pm .05	.22 \pm .05	.71 \pm .06	.12 \pm .01	.59 \pm .04	.14 \pm .02	.56 \pm .05
		DCI	.43 \pm .03	.77 \pm .01	.45 \pm .08	.75 \pm .04	.48 \pm .05	.76 \pm .03	.50 \pm .02	.79 \pm .02	.46 \pm .05	.75 \pm .02	.38 \pm .12	.73 \pm .06
		MB+CI	.48 \pm .09	.66 \pm .05	.57 \pm .08	.70 \pm .05	.72 \pm .10	.84 \pm .06	.92 \pm .04	.94 \pm .02	.68 \pm .06	.80 \pm .04	.89 \pm .03	.92 \pm .02
		CDN	.74 \pm .02	.90 \pm .01	.85 \pm .06	.96 \pm .02	.78 \pm .09	.92 \pm .04	.92 \pm .02	.98 \pm .00	.75 \pm .06	.91 \pm .03	.78 \pm .06	.92 \pm .04
		CDN (no \mathcal{L}_G)	.50 \pm .08	.73 \pm .07	.87 \pm .03	.96 \pm .01	.76 \pm .09	.90 \pm .05	.92 \pm .03	.98 \pm .01	.71 \pm .04	.89 \pm .02	.93 \pm .02	.99 \pm .00
		CDN (full)	.75 \pm .03	.91 \pm .02	.99 \pm .01	1.0 \pm .00	.79 \pm .08	.93 \pm .04	1.0 \pm .00	1.0 \pm .00	.76 \pm .05	.92 \pm .02	1.0 \pm .00	1.0 \pm .00
20	40	UT-IGSP	.14 \pm .01	.57 \pm .01	.14 \pm .01	.61 \pm .02	.20 \pm .01	.62 \pm .02	.19 \pm .01	.65 \pm .01	.13 \pm .01	.56 \pm .02	.14 \pm .00	.60 \pm .01
		DCDI-G	.22 \pm .03	.49 \pm .03	.21 \pm .02	.48 \pm .03	.22 \pm .02	.49 \pm .02	.22 \pm .02	.50 \pm .02	.21 \pm .01	.48 \pm .02	.23 \pm .01	.51 \pm .02
		DCDI-DSF	.21 \pm .02	.46 \pm .03	.21 \pm .02	.47 \pm .02	.23 \pm .02	.50 \pm .03	.21 \pm .01	.49 \pm .03	.21 \pm .02	.49 \pm .02	.21 \pm .03	.49 \pm .02
		BACADI-E	.23 \pm .05	.70 \pm .04	.33 \pm .10	.67 \pm .04	.29 \pm .08	.83 \pm .04	.45 \pm .18	.78 \pm .06	.20 \pm .05	.73 \pm .05	.51 \pm .13	.79 \pm .04
		BACADI-M	.15 \pm .01	.64 \pm .02	.19 \pm .04	.63 \pm .02	.17 \pm .02	.70 \pm .05	.31 \pm .10	.77 \pm .07	.14 \pm .01	.65 \pm .03	.32 \pm .06	.77 \pm .04
		DCI	.60 \pm .06	.85 \pm .03	.50 \pm .02	.77 \pm .03	.56 \pm .08	.81 \pm .05	.59 \pm .06	.84 \pm .04	.64 \pm .06	.87 \pm .02	.66 \pm .05	.88 \pm .03
		MB+CI	.53 \pm .09	.70 \pm .06	.47 \pm .06	.64 \pm .03	.78 \pm .04	.88 \pm .02	.95 \pm .03	.96 \pm .02	.83 \pm .03	.89 \pm .02	.96 \pm .04	.97 \pm .03
		CDN	.86 \pm .04	.96 \pm .02	.79 \pm .09	.93 \pm .05	.88 \pm .03	.97 \pm .01	.96 \pm .02	.99 \pm .00	.88 \pm .03	.97 \pm .01	.71 \pm .05	.87 \pm .04
		CDN (no \mathcal{L}_G)	.69 \pm .09	.85 \pm .07	.80 \pm .06	.94 \pm .03	.84 \pm .03	.94 \pm .01	.95 \pm .02	.99 \pm .00	.84 \pm .05	.93 \pm .03	.91 \pm .02	.98 \pm .00
		CDN (full)	.89 \pm .06	.97 \pm .02	.99 \pm .01	1.0 \pm .00	.89 \pm .03	.97 \pm .01	1.0 \pm .00	1.0 \pm .00	.93 \pm .03	.98 \pm .01	1.0 \pm .00	1.0 \pm .00