
Pruning neural network models for gene regulatory dynamics using data and domain knowledge

Intekhab Hossain^{*,§}

Department of Biostatistics
Harvard T.H. Chan School of Public Health
Boston, MA 02115
ihossain@g.harvard.edu

Jonas Fischer^{*}

Dep. for Computer Vision and Machine Learning
Max Planck Institute for Informatics
Saarbrücken, Germany
jonas.fischer@mpi-inf.mpg.de

Rebekka Burkholz[†]

Helmholtz Center CISA
for Information Security
Saarbrücken, Germany
burkholz@cispa.de

John Quackenbush[†]

Department of Biostatistics
Harvard T.H. Chan School of Public Health
Boston, MA 02115
johnq@hsph.harvard.edu

Abstract

The practical utility of machine learning models in the sciences often hinges on their interpretability. It is common to assess a model’s merit for scientific discovery, and thus novel insights, by how well it aligns with already available domain knowledge—a dimension that is currently largely disregarded in the comparison of neural network models. While pruning can simplify deep neural network architectures and excels in identifying sparse models, as we show in the context of gene regulatory network inference, state-of-the-art techniques struggle with biologically meaningful structure learning. To address this issue, we propose DASH[‡], a generalizable framework that guides network pruning by using domain-specific structural information in model fitting and leads to sparser, better interpretable models that are more robust to noise. Using both synthetic data with ground truth information, as well as real-world gene expression data, we show that DASH, using knowledge about gene interaction partners within the putative regulatory network, outperforms general pruning methods by a large margin and yields deeper insights into the biological systems being studied.

1 Introduction

With ever-growing neural network architectures encouraged by the success of overparametrization, with over a trillion parameters in a single model such as GPT4, there is a similarly growing demand for sparser, more parameter-efficient neural networks that are more resource-friendly and interpretable. The Lottery Ticket Hypothesis (LTH) provides an empirical existence proof of sparse, trainable network architectures [18], that eventually achieve a similar performance as their dense counterpart. Subsequent work introduced *structured* pruning approaches, facilitating group-wise neuron- [32, 63], or channel-sparsity [34, 25, 26], which are, however, focused on the structure of the *architecture design*, aiming for better alignment with hardware implementations to eliminate operations, rather than structure that reflects *relevant domain information*. Especially for scientific discovery, an

[‡]Code is available at <https://github.com/QuackenbushLab/DASH>.

^{*,†} These authors contributed equally to this work.

[§] IH is currently employed by Analysis Group, Inc.

alignment of learned structure with such domain knowledge is, however, essential for interpretability, as only then the model represents meaningful domain-relevant relations. Such problem settings often occur for example in physics or biology where a learned model should give an explanation to be able to form a hypothesis. This poses the question: How can we select among multiple predictive models and promote the search for *meaningful* neural network structures?

To guide the learning process, we argue that we need additional problem-specific structural information, and should leverage any available - and reliable - domain knowledge. One of the fundamental tasks of molecular biology is to understand the *gene regulatory dynamics* in health and disease. Gene regulatory dynamics describe the changes of the expression of a gene—the generation of small copies of a DNA segment that can serve among others as blueprint for proteins—dependent on other regulatory factors such as transcription factors, which are proteins that bind next to the gene (DNA segment) to modulate its expression. Yet, the exact dynamics are far from understood and changes in these dynamics can be drivers for diseases such as cancer. As such, improving an understanding of the mechanics behind these dynamics increases the understanding of the disease and can ultimately inform therapy design. This problem setting of estimating gene regulatory dynamics requires high interpretability, as an understanding of the true biological mechanics—the relationship between regulatory factors and a gene’s expression—is needed, as well as sample-efficiency, as generating time-course data even for a few patients is extremely expensive. While models to estimate gene regulatory dynamics have been suggested [14, 1, 61, 27], none of these is particularly sparse or interpretable. We propose a new approach of network sparsification that *guides pruning by domain knowledge* implemented for a neural model for estimating gene regulatory dynamics, which yields networks that are very sparse, align with underlying biology, while accurately predicting dynamics.

This idea of *prior-informed* or domain-aware pruning is at the heart of this paper. In particular, we propose DASH (**D**omain-**A**ware **S**parsity **H**euristic, Fig. 1), an iterative pruning approach that scores parameters taking into account structural domain-knowledge. With DASH, it is possible to control the level of prior information taken into account for pruning and it *automatically finds an optimal sparsity level* aligned with both the prior and the data. Considering the task of estimating gene regulatory dynamics, we first show in synthetic experiments that DASH generally outperforms standard (task- and architecture agnostic) pruning as well as task-specific pruning approaches. On real data with a reference gene regulatory network (GRN) derived from gold standard biological experiments, we show that DASH better recovers the reference GRN and reflects more biologically plausible information. On recent single cell data on blood differentiation, we show that DASH, in contrast to existing work, identifies biologically relevant pathways that can be used to generate new insights and inform domain experts. We anticipate that our work serves both for future benchmarking on how well pruning approaches are in structure learning, as well as a blueprint for guided network pruning in fields where domain knowledge is readily available, such as in other hard sciences including physics or material science, where knowledge about variables, e.g., associations between atoms or molecules or equations relating quantities in a system, is available.

2 Related work

The Lottery Ticket Hypothesis (LTH) [18] provides an empirical existence proof of sparse, trainable neural network architectures. It conjectures that dense, randomly initialized neural networks contain subnetworks that can be trained in isolation with the same training algorithm that is successful for the dense networks. However, a strong version of this hypothesis [52, 66], which has also been proven theoretically [43, 49, 47, 17, 5, 12, 15], suggests that the identified initial parameters are not only specific to the sparse structure but also the learning task and benefit from information about the larger dense network that has been pruned [48]. Acknowledging this strong relation, other works have proposed to combine mask and parameter learning directly in continuous sparsification approaches [54] that employ regularization strategies that approximate L0 penalties [54, 41, 31, 56]. In the following, we recap the key ideas behind the methods most relevant to ours.

Explicit pruning-based approaches

Magnitude pruning (MP) In magnitude pruning (MP) a neural network is trained and then (post-hoc) pruned to a desired sparsity level by masking the corresponding proportion of *smallest magnitude* weights. This smaller, masked network is then further trained, reminiscent of fine-tuning [22].

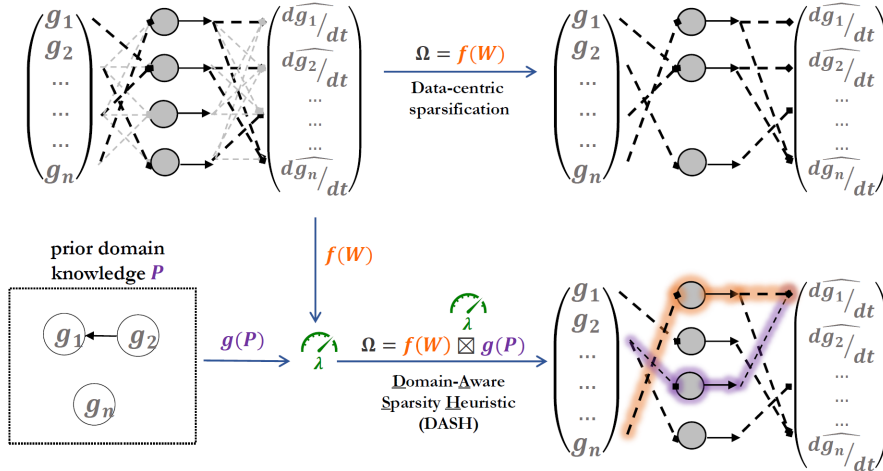


Figure 1: *DASH*. A NN, here a neural ODE for gene regulatory dynamics, is traditionally sparsified in a data-centric way (top). Pruning is done based on data alone, the pruning score Ω is a function of the learned weights W . Such sparsified models often do not learn plausible relationships in the data domain. We propose *DASH* (bottom), which additionally incorporates domain knowledge P into the pruning score Ω , yielding sparse networks giving meaningful and useful insights into the domain.

More formally, we start with a fully connected neural network NN_{Θ} with L layers parametrized by $\Theta := \{(W^l, b^l)\}_{l=1}^L$. MP is performed after training is complete (i.e. post hoc). For a suitable threshold of $p\%$, MP “prunes” the trained Θ by setting the smallest (absolute value) $p\%$ of weights in Θ to 0. The choice of parameters to prune can either be made in an unstructured way, by choosing the the lowest $p\%$ across all $\{W^l\}_{l=1}^L$, semi-structured, by pruning the lowest weights per layer, or in a structured way such that, e.g., neurons with lowest $p\%$ average outgoing weight are pruned. Once pruning is complete the $p\%$ -sparse $\text{NN}_{\Theta'}$ is fine-tuned on the data so that Θ' is learned appropriately. **Iterative magnitude pruning (IMP)** [18] suggest to alternate between training and magnitude pruning, iteratively sparsifying the network, to date still of the most successful pruning strategies. In practice, a pruning schedule is used to *iteratively* sparsify Θ , until a Θ' with desired target sparsity or predictive performance plateau is reached. Importantly, weights are reset to initial pre-training values, either after each round of pruning or once after the target sparsity has been reached. **Pruning with rewinding** [53] have demonstrated that rewinding weights to an earlier training point—a compromise between fine-tuning of MP and reset to initialization of IMP— provides good performance, which also has been suggested in the context of Neural ODEs as SparseFlow [37].

Implicit penalty-based approaches

C-NODE [2] seek to reduce the overall number of input-output dependencies (i.e. paths of contribution from input neuron i to output neuron j) through $\{W^l\}_{l=1}^L$. The approach can result in both feature and weight sparsity in NeuralODEs.

L_0 [40] incorporate a differentiable L_0 norm regularizer term in the objective. It implicitly prunes the network by encouraging weights to get exactly to zero. The L_0 regularizer is operationalized using non-negative stochastic gates which act as masks on the weights.

PathReg [1] innovatively combines the strengths of both C-NODE and the L_0 approach to promote both weight and feature sparsity in NeuralODEs. It uses stochastic gates similar to [40] and add a C-NODE-inspired penalty terms that constrain the overall number of input-output paths by regularizing the *probability* of any path from input i to output j being non-zero.

Modeling gene regulatory dynamics As application, we consider estimation of gene regulatory dynamics. Early work, such as COPASI [45] use a fully parametric modeling approach that are limited in their prediction capabilities. With recent advances in machine learning, tools such as Dynamo [51], PROB [58], and RNA-ODE [39] aim to learn regulatory ODEs using sparse kernel regression, Bayesian Lasso, and random forests, respectively. Leveraging high flexibility and performance of neural models, PRESCIENT [65] uses a simple NN to learn regulatory ODEs, whereas tools such




Figure 2: *Results on simulated data.* We visualize performance of pruning strategies in comparison to original PHOENIX (baseline) in terms of achieved sparsity (x-axis) and balanced accuracy (y-axis) of the recovered gene regulatory network against the ground truth on the SIM350 data with 5% noise. Error bars are omitted when error is smaller than depicted symbol. ✓ indicate methods that leverage prior information. Top left is best: recovering true, inherently sparse biological relationships.

as DeepVelo [9] and sctour [35] have a variational autoencoder as backbone. The latest line of research [14, 1, 61, 27] uses neural ordinary differential equations [8]. However, a key limitation of these methods is the lack of interpretability arising from non-sparse dynamics that do not align with ground truth biology. Consequently, the induction of sparsity in gene regulatory ODEs has been an active area of research with C-NODE and PathReg as most recent advances[2, 1], the achieved sparsity levels are, however, not yet sufficient to capture the relevant biology.

3 Domain-aware pruning with DASH

While the above sparsification strategies have shown to perform well in various settings, the resulting models are often either not particularly sparse or do not reflect meaningful domain knowledge. We hypothesize that this is due to two reasons: (1) the difficulty of identifying a good sparse network and (2) the current focus on *hardware-centric* rather than *task-centric* pruning, valuing structural pruning of a model in terms of groups of neurons (layers, channels) over structural pruning reflecting task-specific knowledge. To overcome these problems, we suggest to ground the model search (here, the network training) with existing domain-specific knowledge, which eases identifiability due to the introduced constraints and enables task-aware pruning to identify meaningful domain knowledge.

In the following, we propose DASH (**D**omain-**A**ware **S**parsity **H**euristic), an iterative pruning-based approach that accounts for prior knowledge by scoring parameters in terms of their alignment with this prior, and show its usefulness for a neural model for the inference of gene regulatory dynamics. We assume the domain knowledge to be given as input-output relations (e.g., known protein—gene interactions in molecular biology), for which we want the network flow between any input and output to align with. Suppose our domain knowledge for a task from this domain is given as a real-valued relationship graph $G = (V, E)$, where nodes are relevant entities from the domain, e.g. genes or proteins, and edges are a strength of association between these entities, e.g. evidence of association derived from literature or experiments. Examples for such relational information in case of protein—gene associations can be derived from protein binding profiles [57]. For the rest of the paper, we will assume that this domain knowledge is given as a matrix $\mathbf{P} \in \mathbb{R}^{k \times r}$, which encodes the strength of association between the k inputs and r outputs for our task of interest, such as known proxies of protein—gene interactions. Intuitively, for a one-layer neural network, we encourage pruning scores for a (neural) network edge to be proportional to the corresponding edge in the prior knowledge graph G while still taking into account the data-specific knowledge, thus enabling learning

of new knowledge and robustness to wrong or missing information in the prior. We begin with this simple base case of task-aware pruning of a fully connected neural network with $L = 1$ layer and extend to more layers below.

DASH for $L = 1$. For a single layer NN, with k input and r output neurons and corresponding weight matrix $\mathbf{W} \in \mathbb{R}^{r \times k}$, we compute non-negative **pruning scores** $\Omega \in \mathbb{R}^{r \times k}$ by leveraging the **domain knowledge** $\mathbf{P} \in \mathbb{R}^{r \times k}$. In practice, we allow balancing between *data-driven* and *prior-knowledge-driven* pruning, implemented through a convex combination of the learned weights \mathbf{W} and prior domain knowledge \mathbf{P} controlled by the parameter $\lambda \in [0, 1]$. Alternating between training and pruning akin to Iterative Magnitude Pruning [18], we set the following pruning score during a pruning phase:

$$\Omega^{(t)} := (1 - \lambda) |\widetilde{\mathbf{W}^{(t)}}| + \lambda |\mathbf{P}|,$$

where $|\widetilde{\mathbf{W}^{(t)}}|$ represents the appropriately normalized matrix (details in Appendix B.2) of absolute weights as learned up to epoch t . We then prune the parameters in $\mathbf{W}^{(t)}$ corresponding to the lowest absolute $p_t\%$ of entries in $\Omega^{(t)}$, where p_t is the desired sparsity level at time t given by a schedule.

DASH for $L = 2$. For two-layered NNs with weights $\mathbf{W}_1 \in \mathbb{R}^{m \times k}$, $\mathbf{W}_2 \in \mathbb{R}^{r \times m}$, i.e k inputs, r outputs, and m hidden neurons, we consider knowledge about input-output relationships $\mathbf{P} \in \mathbb{R}^{r \times k}$ as before. We can additionally use further knowledge about input-input relationships $\mathbf{C} \in \mathbb{R}^{k \times k}$. In molecular biology this could be information about binding or interaction partners available in databases such as STRINGDB [59], which has also been employed to guide static gene regulatory network inference [62], or co-regulators, derived from co-occurrence of proteins. Intuitively, we now project pruning scores for the first layer to the prior knowledge about input-input relations, encouraging closeness to this prior, while projecting the product of pruning scores of first and second layer to known input-output relations, reflecting the flow of information from input to output through these two layers. Given that $\mathbf{W}_1^{(t)}$ represents how the k inputs are encoded by m neurons, and $\Omega_1^{(t)}$ are the corresponding pruning scores, we surmise that the matrix product $\Omega_1^{(t)\top} \Omega_1^{(t)} \in \mathbb{R}^{k \times k}$ should approximately align with the prior knowledge \mathbf{C} . Since solving $\Omega_1^{(t)\top} \Omega_1^{(t)} = \mathbf{C}$ is not directly feasible we initialize $\Omega_1^{(0)}$ randomly, and resort to solving a recurrence relation version of problem, that is $\Omega_1^{(t-1)\top} \Omega_1^{(t)} = \mathbf{C}$. Using the left and right pseudo-inverse to obtain a solution to the above, defined as $\text{PInv}_L(X) = (X^\top X)^{-1} X^\top$ and $\text{PInv}_R(X) = X^\top (X X^\top)^{-1}$ respectively:

$$\Omega_1^{(t)} := (1 - \lambda_1) |\widetilde{\mathbf{W}_1^{(t)}}| + \lambda_1 \left| \text{PInv}_L \left(\Omega_1^{(t-1)\top} \right) \cdot \mathbf{C} \right|.$$

$\text{PInv}_L \left(\Omega_1^{(t-1)\top} \right) \cdot \mathbf{C}$ encourages $\Omega_1^{(t)\top} \Omega_1^{(t)}$ to iteratively align with \mathbf{C} as t increases (i.e. as training progresses). With $\Omega_1^{(t)}$ fixed, we can update scores $\Omega_2^{(t)}$ of the second layer parameters $\mathbf{W}_2^{(t)}$. Since the product $\mathbf{W}_2^{(t)} \cdot \mathbf{W}_1^{(t)} \in \mathbb{R}^{r \times k}$ represents the overall flow of information from inputs to outputs at epoch t , we surmise that $(\Omega_2^{(t)} \Omega_1^{(t)}) \in \mathbb{R}^{r \times k}$ should reflect \mathbf{P} . We thus get

$$\Omega_2^{(t)} := (1 - \lambda_2) |\widetilde{\mathbf{W}_2^{(t)}}| + \lambda_2 \left| \mathbf{P} \cdot \text{PInv}_R \left(\Omega_1^{(t)} \right) \right|.$$

Similar to the case for DASH for $L = 1$ layer, we can now prune the parameters of $\mathbf{W}_1^{(t)}$ and $\mathbf{W}_2^{(t)}$ based on the magnitude of pruning scores $\Omega_1^{(t)}$ and $\Omega_2^{(t)}$, respectively.

DASH for $L > 2$ For many interpretability-centric tasks, including our application to gene regulatory networks, small architectures of $L = 2$ are common, as domain experts are interested in understanding the exact flow of information through the network. Furthermore, we know that two-layer neural networks exhibit universal approximation [11]. We however hypothesize that the technique of computing pruning scores by fixing those of preceding layers can be extended to a larger number L of fully connected layers and elaborate in App. B.6.

Flexibility While the λ_l can be tuned using cross-validation (see App. B.4), we note that it allows for flexibly encoding different pruning philosophies. Specifically, when $\lambda_l = 0 \forall l$, DASH corresponds to SparseFlow, and when $\lambda_l = 1 \forall l$, DASH represents fully prior-based sparsification (which we term "BioPrune" and consider as experimental baseline).

Table 1: *Synthetic data results.* We give model sparsity, balanced accuracy with respect to edges in the ground truth gene regulatory network, mean squared error of predicted gene regulatory dynamics on the test set, and number of epochs (till validation performance plateaus) as proxy of runtime. \checkmark is used to indicate methods that leverage prior information. Results are on SIM350 data with 5% noise.

Strategy (\checkmark = prior-informed)		Sparsity(%)	Bal. Acc.(%)	MSE (10^{-3})
None/Baseline [27]		7.5 ± 0.1	51.8 ± 0.03	3.0 ± 0.4
Penalty-based (implicit)	L_0 [40]	33.8 ± 4.7	55.0 ± 0.5	8.5 ± 1.0
	C-NODE [2]	6.2 ± 0.5	55.9 ± 0.1	2.8 ± 0.6
	PathReg [1]	56.5 ± 1.5	61.9 ± 1.0	8.0 ± 1.8
	PINN [27] \checkmark	9.9 ± 0.4	58.6 ± 0.7	2.5 ± 0.2
	DST[38]	92.8 ± 0.3	71.9 ± 0.5	4.0 ± 0.5
Pruning-based (explicit)	IMP [18]	81.9 ± 6.6	61.7 ± 0.7	4.7 ± 1.1
	Iter. SynFlow [60]	79.3 ± 1.2	58.4 ± 0.6	7.0 ± 2.1
	SparseFlow [37]	96.0 ± 0.01	70.9 ± 1.5	3.6 ± 0.6
	BioPrune (Ours, see 3) \checkmark	83.5 ± 1.9	87.3 ± 0.8	2.6 ± 0.9
	DASH (Ours) \checkmark	94.6 ± 1.2	90.7 ± 0.4	2.4 ± 1.2
Hybrid	PINN + MP (Ours) \checkmark	87.0 ± 0.01	82.4 ± 0.2	2.3 ± 0.3

4 Task-aware pruning for sparse gene regulatory dynamics

Perhaps one of the most interesting applications of Machine Learning is in the field of Molecular Biology with the goal of understanding human health and disease. A central mechanism in humans is the process of gene expression in each cell. There, copies of short segments of our genome are produced. These copies are among other things the blueprint for the production of different proteins, which are needed virtually everywhere in our bodies. If this tightly regulated process of gene expression goes wrong, for example because of a mutation in our genome, this can have profoundly bad effects, such as in the case of cancer. As such, studying this process is of great interest to understand and improve human health and discover new therapeutic targets.

Here, we consider the task of predicting the regulatory dynamics of gene expression. To be able to understand the model and transfer it to clinical practice, interpretability is key. The most recent developments in modeling gene regulatory systems allow to model actual (temporal) regulatory dynamics, but require complex models, such as NeuralODEs, that hinder interpretability. While state-of-the-art results are now achieved with shallow architectures [27] that are more tractable than deep, heavily over-parameterized networks, these models still encode information across many thousands of weights and we show experimentally that such information does not reflect true biology well. In fact, true gene regulatory networks and hence their underlying dynamics are inherently sparse [6]. This sparsity should be properly reflected by neural dynamics models. The PHOENIX NeuralODE model will serve as our base model for applying sparsification strategies and we show that pruning aligned with prior domain knowledge improves interpretability as well as quality of inferred (new) knowledge.

In a nutshell, given a time series gene expression sample for k genes, PHOENIX uses NeuralODEs to construct the predicted trajectory between gene expression $\mathbf{g}(t) \in \mathbb{R}^k$ (inputs) at time $t = t_i$ to any future expression $\hat{\mathbf{g}}(t_{i+1})$ (outputs), by implicitly modeling the RNA velocity ($d\mathbf{g}/dt$). PHOENIX uses biokinetics-inspired activation functions to separately model additive and multiplicative co-regulatory effects. The trained model encodes the ODEs governing the dynamics of gene expression, which can be directly extracted for biological insights. We apply DASH to PHOENIX and give a brief review of the PHOENIX architecture in App. B.10 and a detailed account on how to apply DASH to this architecture in App. B.11. Next, we provide experiments on synthetic and real data showing the advantages of prior-informed pruning on the task of predicting gene-regulatory dynamics.




Figure 3: *Reconstruction of ground truth relationships.* Estimated effect of gene g_j (x-axis) on the dynamics of gene g_i (y-axis) in SIM350 for different levels of noise (rows). Ground truth is given on the left, our suggested approach and baselines (DASH, BioPrune, and PINN+MP) on the right with mean squared error between inferred regulatory relationships and ground truth in purple.

5 Experiments

For evaluation we consider synthetic data from an established simulator tool [4], as well as real world data of gene expression from breast cancer tissue [13], from yeast with synchronized cell cycle [50], and from human bone marrow [1]. In case of synthetic data, we use the ground truth regulatory system from the generating model for validation. For breast cancer and yeast cell cycle data we use additional experimental data (ChIP-seq) from the corresponding studies, which are independent gold-standard biological experiment measuring sample-specific TF–gene interactions, to evaluate inferred regulatory relationships. We measure the correctness of a GRN learned by a model (see App. B.10.4) in terms of balanced accuracy, which measures whether an edge is correctly reconstructed weighted by the sparsity of the aforementioned ground truth graph. To evaluate predictive performance for real data, we use a 6% hold-out test set for breast cancer and one of the biological replicates hold out from training for the yeast data. As prior knowledge we leverage general information of transcription factor binding to gene promoter regions as prior information, which can be computed from binding motif matches with the corresponding genome (human respectively yeast). The result is a matching score that can be thresholded to get a 0, 1-based matrix encoding which (TF-encoding) gene has a relationship with which other gene. We follow the approach of Guebila et al. [3] to obtain matrix P . As prior C , we use the STRING database [59], which gives a general (i.e., not tissue-specific) graph of protein-protein interaction. Here, we use the interactions based on experimental evidence only and employ a cutoff of .6 to get a binary adjacency matrix. (for more details, see App. B.3).

To compare pruning strategies, we consider the PHOENIX model as a basis, which is the state-of-the-art NeuralODE for estimating gene regulatory dynamics [27] and provide an ablation on a standard MLP architecture (see App. Tab. 7, App. Tab. 10, and App B.12). We compare DASH against the PHOENIX model without additional pruning as performance reference, and suggest two simple yet powerful baselines, which is post-hoc magnitude pruning of weights followed by finetuning (PINN+MP), and BioPrune, a fully prior-based pruning (cf. Sec. 3). From the literature, we consider L_0 -regularized pruning [40], C-NODE [2], and PathReg [1], which have been recently proposed for the inference of sparse gene-regulatory relationships, PHOENIX with biological regularization [27], and dynamic sparse training (DST) [38], all of which are implicit pruning approaches. We further consider explicit, iterative score-based pruning approaches including Iterative Magnitude Pruning (IMP) [18], the flow-based model-agnostic pruning method SynFlow [60], and the flow-based Neural ODE pruning SparseFlow [37]. We tune hyperparameters, including λ for DASH, on a validation set. Unlike other methods (e.g., L_0) DASH does not prolong the runtime of training much. A common pruning schedule where pruning scores are computed once every 10 epochs only increases runtime by <2% for the full model fitting process.

Simulated gene regulatory systems We simulate gene expression time-series data from a fixed dynamical system, the ground truth was thus known (see App. B.1). In short, we adapt SimulatorGRN [4] to generate noisy time-series expression data from two synthetic gene regulatory systems (SIM350 and SIM690, consisting of 350 and 690 genes, respectively). We split trajectories into training (88%), validation (6% for tuning λ), and testing (6%). We evaluate all methods in terms of achieved sparsity and MSE of predicted gene expression values on the test set (App. B.2, B.4) and investigate biological plausibility by calculating balanced accuracy of regulatory relationships extracted from the PHOENIX model (for details, see App. B.5, B.10.4). We here report the results for the data of 350 genes and 5% noise, noting that results are consistent across different noise levels and with more number of genes (see App. Sec. A.1).

A general trend across all experiments that aligns with our initial motivation is that dense models (sparsity $< 50\%$) have a significantly worse reconstruction of the underlying biology – the ground truth GRN – than sparse models (sparsity $> 80\%$) (see Fig. 2). Furthermore, we see that DASH retrieves not only among the sparsest networks, but also reflects the underlying GRN best across all methods, outperforming comparably sparse IMP by about 20 percentage points accuracy in different settings, even with decrease in quality of the prior (see sensitivity analysis in App. Tab. 6). Due to the prior-informed structured pruning, it is able to occupy the sweet spot of highly sparse at the same time biologically meaningful models.

Consistent with the literature, PathReg outperforms L_0 as well as C-NODE in terms of sparsity [1], we additionally find evidence that it also delivers more biologically meaningful results. Yet, IMP as well as prior-informed pruning approaches outperform PathReg by a large margin. The MSE of predicted gene expression of DASH is among the best, within one standard error of the best overall method. The only better approach is our suggested baseline, a combination of posthoc magnitude pruning of PHOENIX combined with additional finetuning (PINN+MP), which is, however, impractical as it requires to train and prune many PHOENIX models along a grid of sparsity levels (see App. B.9.3).

Visualizing the estimated against ground truth regulatory effects (i.e., functional relationships between variables), we observe that DASH captures the effects much better than competitors (see Fig. 3). Virtually all existing approaches discover spurious regulatory effects, whereas prior-informed pruning identify the main regulatory effects correctly. Moreover, with increasing levels of noise in the simulation, we observe that both BioPrune as well as PINN+MP start finding spurious dependencies, while DASH still recovers the overall structure well. While not perfect, as seemingly there are more dependencies than in the sparse ground truth, potentially introduced by correlations between features, DASH provides a sparse estimation of regulatory effects that most closely resembles the ground truth relationships among existing work.

Pseudotime-ordered breast cancer samples To investigate the performance of DASH on real data, we consider gene expression measurements from a cross-sectional breast cancer study [13]. This data of 198 breast cancer patients with measurements for 22000 genes has been preprocessed and ordered in pseudotime [58], which we use as basis for our experiments (cf App. B.7). Across methods, we observe that implicit sparsification methods generally perform poorly in terms of sparsity and accuracy of recovered relationships (see Tab. 2). While pruning-based sparsification approaches achieve greater sparsity and performance in predicted gene expression, with SparseFlow reaching the highest sparsity (95.7%) among all methods, the recovered biological relations are not better than random chance, which renders the underlying models useless for scientific discovery. DASH in contrast finds a comparably sparse network (92.7% sparsity) while having top of the line performance in terms of test MSE and high alignment with true biology (95.7% balanced accuracy). For this particular dataset, we observe that DASH primarily builds on the prior knowledge, not surprisingly performing similarly as BioPrune, which is our suggested baseline pruning approach taking only the prior into account. We will see for other real-world data that this weight of domain knowledge is highly task-specific and BioPrune yields sub-optimal results on different data.

To better understand whether the inferred gene regulatory dynamics align with meaningful biology, we additionally perform a pathway analysis (see App. B.8). Such pathway analysis are a standard approach for domain experts to distill information for example for therapeutic design. The genes that show the highest impact on the dynamics within the derived model are tested whether they enrich in a specific higher level biological pathways. For the top-20 most significantly enriched pathway per model (App. Fig. 5), we observe that in contrast to prior-informed methods, the existing pruning approaches show only very few significant pathways, consistent with our quantitative results on

Table 2: Results on breast cancer and yeast data. Balanced accuracy is based on reference gold standard experiments (transcription factor binding ChIP-seq) available for this data. DASH found optimal λ -values of (0.995, 0.95) respectively (0.75, 0.75) for breast cancer and yeast. * marks our suggested baselines and method, \checkmark marks methods that use prior information for sparsification.

Data Strategy	Breast cancer in pseudotime			Yeast cell cycle		
	Sparsity	Bal. Acc.	MSE (10^{-5})	Sparsity	Bal. Acc.	MSE (10^{-2})
None/Baseline	0.03%	49.99%	7.78	0.10%	49.87%	4.84
L_0	10.77%	50.15%	7.90	34.43%	48.43%	5.33
C-NODE	11.20%	50.01%	8.06	10.89%	50.04%	4.87
PathReg	14.09%	50.24%	7.92	12.09%	50.11%	5.35
PINN \checkmark	0.11%	49.99%	7.82	0.17%	49.93%	5.77
DST	67.02%	50.42%	7.78	77.80%	49.92%	5.18
IMP	36.02%	50.34%	7.77	83.22%	49.99%	5.46
Iter. SynFlow	91.93%	49.37%	7.78	85.65%	49.57%	5.41
SparseFlow	95.70%	49.70%	7.76	95.22%	49.89%	5.38
BioPrune *, \checkmark	93.44%	95.67%	7.80	94.69%	79.23%	5.94
DASH *, \checkmark	92.71%	95.69%	7.76	97.18%	88.43%	5.27
PINN + MP *, \checkmark	92.00%	54.02%	7.79	95.01%	55.39%	6.09

inferred regulatory relations. Moreover, disease-relevant pathways such as TP53 activity or FOXO mediated cell death, both of which are highly relevant in cancer [44, 28], are only visible in models pruned with prior information. This provides evidence that pruning informed by a biological prior recovers biological signals that are relevant in the disease and which can not be picked up otherwise. Furthermore, we find Heme-signaling as a pathway uniquely identified as relevant in our approaches (cf. App. Fig. 5). Heme as a signaling molecule has key roles in the gene regulatory system [46], and turns out to have an anti-tumor role in breast cancer specifically [19]. Subsequent approaches pharmaceutically targeting Heme signaling showed success [30], with one of the key regulators affected being Bach1. To suggest further targets for e.g. combination treatment, we hence examined the top-5 regulatory factors in terms of weights in our estimated gene regulatory dynamics. These factors include PBX1 and FOXM1, for which a drug repurposing of existing compounds, such as [55], could lead to a potential new treatment for this specific cancer.

Yeast cell-cycle data We next consider real data of synchronized yeast cell [50] (see App. B.2 for training setup). We observe an overall trend similar to the breast cancer study in terms of achieved sparsity and balanced accuracy (cf. Tab. 2), with implicit sparsification methods generally finding significantly less sparse models and all methods that do not incorporate prior knowledge having inferred relationships that are not better than random chance. For this data, however, DASH finds an optimal lambda value that incorporates more data-specific knowledge ($\lambda = 0.75$) compared to the breast cancer study above. This shows the advantage of DASH over our BioPrune baseline model (prior-only pruning), as here we gain about 10% points in balanced accuracy over BioPrune for ChIP-seq validation data [23] and 2% points over BioPrune for an independent TF perturbation network curated to derive a "true" causal GRN [21] (see App. Tab. 8). We also retrieve a 3% points sparser model. Comparing inferred biological knowledge between BioPrune and DASH through a pathway analysis, we see that DASH recovers more significantly enriched pathways related to cell cycle processes (cf. App. Fig. 6).

Cell differentiation in human bone marrow Lastly, we investigate the performance of DASH in an exploratory setting with single cell data of human bone marrow ordered in pseudotime [1]. Here, we are interested in better understanding the gene regulatory dynamics of blood cell differentiation, the process of hematopoietic stem cells specializing into cells taking over roles such as immune response (e.g., B- and T-cells). This process is called hematopoiesis. We follow the steps of [1] to first split samples (i.e., cells) into the three different lineages (paths of differentiation), and train separate models for each (see App. B.7). We will here focus on the analysis of the Erythroid lineage.

As before, DASH yields highly sparse (95%) networks, the most sparse among all competitors (App. Tab. 9). IMP shows similarly strong sparsification as DASH while PathReg achieves much less sparsity (14%). In terms of performance, all methods achieve similar MSE of predicted gene expression dynamics on the test set, meaning that even though much sparser, both DASH and IMP predict equally well as an order of magnitude more dense network. For this data, there are no gold-standard experiments for regulatory relations available, we hence focus on analysing the network topology. From the literature, we would expect sparser networks to be better align with biology [6]. DASH indeed shows the lowest out-degree in the inferred regulatory network, less than half of what IMP recovers. PathReg shows an order of magnitude larger average out-degree. To confirm the biological plausability, we again do a pathway analysis. DASH seems to find significant enrichment in biologically relevant pathways (App. Fig. 7) that can directly be linked to hematopoiesis, such as *heme signaling* or *RUNX1 regulates differentiation of hematopoietic stem cells*, which neither BioPrune nor SparseFlow—the only other method yielding a proper sparse model—could recover.

6 Discussion & Conclusion

We considered the problem of identifying sparse neural networks in the context of interpretability with a focus on the application to gene regulatory systems modeling. In domains such as biology and contexts when the true underlying systems are sparse, interpretability is key for experts, rendering the use of the common over-parametrized and complex neural network architectures difficult. Although NNs do not directly encode e.g. the regulatory relationship between genes its deep architecture is necessary to model complex functional relationships while ensuring stable learning. Yet, we can ensure that recent advances in neural network pruning, such as those around the Lottery Ticket Hypothesis [18], promise sparse and well-performing models, yet, hardness results prove finding optimally sparse models to be challenging [43], which is also reflected by recent benchmarking results [16]. Our experiments confirmed that general pruning strategies provide sub-optimal sparsity, moreover, the underlying biological relationships are not properly reflected in the model. We proposed *to guide pruning by domain knowledge*, leveraging existing prior information to improve the interpretability and meaningfulness of pruned models.

In case studies on gene regulatory dynamic inference, a key task in molecular biology with high relevance in cancer research, we showed based on simulated as well as real world data that our method, DASH, in contrast to a wide range of state-of-the-art methods, is able to recover neural networks that are both very sparse and at the same time biologically meaningful, allowing for direct extraction of a sparse gene regulatory network. On real data, DASH not only better aligns with gold-standard experimental evidence of regulatory interactions, but also uniquely reflects the data-specific biological pathways, which can be used by domain experts to generate new insights.. It thus serves as a proof of concept that in critical domains, where interpretability is essential and domain knowledge exists, pruning can be heavily improved by alignment with prior knowledge. While our guided pruning approach is in principle agnostic to the type of neural network and task, we here focused on a specific case study that we deemed important. In the future, it would be interesting to apply DASH to different cases and domains, including other biomedical tasks, but also to physics or material sciences, where interpretability is also key and domain knowledge exists in the form of physical constraints and models. For any application, an important consideration to apply DASH is on the one hand the availability of prior knowledge, but on the other hand its quality; while we show here that even with incomplete and noisy prior knowledge we receive good results, factually wrong priors could steer the solution towards a wrong model. We hence assume that DASH will be of primary use in classical hard sciences mentioned above, with priors that stood the test of time over several decades. Another line of future work includes different architectural designs, such as convolution or attention mechanisms, where input-output relationships are less straightforward to project to across several layers.

In summary, we make a case for pruning informed by domain knowledge and provide evidence that such approaches can massively improve sparsity along with domain specific interpretability.

Acknowledgements RB received funding from the European Research Council (ERC) under the Horizon Europe Framework Programme (HORIZON) for proposal number 101116395 SPARSE-ML. JQ was supported by a grant from the US National Cancer Institute (R35CA220523) and additional funding from the National Human Genome Research Institute (R01HG011393).

References

- [1] Hananeh Aliee, Till Richter, Mikhail Solonin, Ignacio Ibarra, Fabian Theis, and Niki Kilbertus. Sparsity in continuous-depth neural networks. *Advances in Neural Information Processing Systems*, 35:901–914, 2022.
- [2] Hananeh Aliee, Fabian J Theis, and Niki Kilbertus. Beyond predictions in neural odes: Identification and interventions. *arXiv preprint arXiv:2106.12430*, 2021.
- [3] Marouen Ben Guebila, Camila M Lopes-Ramos, Deborah Weighill, Abhijeet Rajendra Sonawane, Rebekka Burkholz, Behrouz Shamsaei, John Platig, Kimberly Glass, Marieke L Kuijjer, and John Quackenbush. GRAND: a database of gene regulatory network models across human conditions. *Nucleic Acids Research*, 50(D1):D610–D621, 09 2021.
- [4] Dharmesh D Bhuva, Joseph Cursons, Gordon K Smyth, and Melissa J Davis. Differential co-expression-based detection of conditional relationships in transcriptional data: comparative analysis and application to breast cancer. *Genome biology*, 20(1):1–21, 2019.
- [5] Rebekka Burkholz. Most activation functions can win the lottery without excessive depth. In *Advances in Neural Information Processing Systems*, 2022.
- [6] Daniel M Busiello, Samir Suweis, Jorge Hidalgo, and Amos Maritan. Explorability and the origin of network sparsity in living systems. *Scientific reports*, 7(1):12323, 2017.
- [7] Yang Cao, Shengtai Li, Linda Petzold, and Radu Serban. Adjoint sensitivity analysis for differential-algebraic equations: The adjoint dae system and its numerical solution. *SIAM journal on scientific computing*, 24(3):1076–1089, 2003.
- [8] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- [9] Zhanlin Chen, William C King, Aheyon Hwang, Mark Gerstein, and Jing Zhang. Deepvelo: Single-cell transcriptomic deep velocity field learning with neural ordinary differential equations. *Science Advances*, 8(48):eabq3745, 2022.
- [10] Jeanne Chèneby, Marius Gheorghe, Marie Artufel, Anthony Mathelier, and Benoit Ballester. Remap 2018: an updated atlas of regulatory regions from an integrative analysis of dna-binding chip-seq experiments. *Nucleic acids research*, 46(D1):D267–D275, 2018.
- [11] George Cybenko. Approximation by superpositions of a sigmoidal function. *Math. Control. Signals Syst.*, 2(4):303–314, 1989.
- [12] Arthur da Cunha, Emanuele Natale, and Laurent Viennot. Proving the lottery ticket hypothesis for convolutional neural networks. In *International Conference on Learning Representations*, 2022.
- [13] Christine Desmedt, Fanny Piette, Sherene Loi, Yixin Wang, Françoise Lallemand, Benjamin Haibe-Kains, Giuseppe Viale, Mauro Delorenzi, Yi Zhang, Mahasti Saghatchian d’Assignies, et al. Strong time dependence of the 76-gene prognostic signature for node-negative breast cancer patients in the transbig multicenter independent validation series. *Clinical cancer research*, 13(11):3207–3214, 2007.
- [14] Rossin Erbe, Genevieve Stein-O’Brien, and Elana J Fertig. Transcriptomic forecasting with neural ordinary differential equations. *Patterns*, 4(8), 2023.
- [15] Damien Ferbach, Christos Tsirigotis, Gauthier Gidel, and Joey Bose. A general framework for proving the equivariant strong lottery ticket hypothesis. In *International Conference on Learning Representations*, 2023.
- [16] Jonas Fischer and Rebekka Burkholz. Plant ’n’ seek: Can you find the winning ticket? In *International Conference on Learning Representations*, 2022.
- [17] Jonas Fischer, Advait Harshal Gadhikar, and Rebekka Burkholz. Towards strong pruning for lottery tickets with non-zero biases. *arXiv preprint arXiv:2110.11150*, 2021.

- [18] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *arXiv preprint arXiv:1803.03635*, 2018.
- [19] Norberto Ariel Gandini, Eliana Noelia Alonso, Mar´ıa Eugenia Fermento, Marilina Mascar´o, Mart´ın Carlos Abba, Georgina Pamela Col´o, Juli´an Ar´evalo, Mar´ıa Julia Ferronato, Josefina Alejandra Guevara, Myriam N´uñez, Alejandro Carlos Pichel, Pamela and Curino, and Mar´ıa Marta Facchinetti. Heme oxygenase-1 has an antitumor role in breast cancer. *Antioxidants & Redox Signaling*, 30(18):2030–2049, 2019. PMID: 30484334.
- [20] Kimberly Glass, Curtis Huttenhower, John Quackenbush, and Guo-Cheng Yuan. Passing messages between biological networks to refine predicted interactions. *PLoS one*, 8(5):e64832, 2013.
- [21] S. R. Hackett, E. A. Baltz, M. Coram, B. J. Wranik, G. Kim, A. Baker, M. Fan, D. G. Hendrickson, M. Berndl, and R. S. McIsaac. Learning causal networks using inducible transcription factors and transcriptome-wide time series. *Mol Syst Biol*, 16(3):e9174, Mar 2020.
- [22] Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. *Advances in neural information processing systems*, 28, 2015.
- [23] Christopher T Harbison, D Benjamin Gordon, Tong Ihn Lee, Nicola J Rinaldi, Kenzie D Macisac, Timothy W Danford, Nancy M Hannett, Jean-Bosco Tagne, David B Reynolds, Jane Yoo, et al. Transcriptional regulatory code of a eukaryotic genome. *Nature*, 431(7004):99–104, 2004.
- [24] Trevor Hastie, Robert Tibshirani, Jerome H Friedman, and Jerome H Friedman. *The elements of statistical learning: data mining, inference, and prediction*, volume 2. Springer, 2009.
- [25] Yang He, Guoliang Kang, Xuanyi Dong, Yanwei Fu, and Yi Yang. Soft filter pruning for accelerating deep convolutional neural networks. In *International Joint Conference on Artificial Intelligence*, 2018.
- [26] Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 1398–1406, 2017.
- [27] Intekhab Hossain, Viola Fanfani, Jonas Fischer, John Quackenbush, and Rebekka Burkholz. Biologically informed neuralodes for genome-wide regulatory dynamics. *Genome Biology*, 25(127), 2024.
- [28] Y. Jiramongkol and E. W. Lam. FOXO transcription factor family in cancer and metastasis. *Cancer Metastasis Rev*, 39(3):681–709, Sep 2020.
- [29] W Evan Johnson, Cheng Li, and Ariel Rabinovic. Adjusting batch effects in microarray expression data using empirical bayes methods. *Biostatistics*, 8(1):118–127, 2007.
- [30] Pritpal Kaur, Shreya Nagar, Madhura Bhagwat, Mohammad Uddin, Yan Zhu, Ivana Vancurova, and Ales Vancura. Activated heme synthesis regulates glycolysis and oxidative metabolism in breast and ovarian cancer cells. *PLOS ONE*, 16(11):1–19, 11 2021.
- [31] Aditya Kusupati, Vivek Ramanujan, Raghav Somani, Mitchell Wortsman, Prateek Jain, Sham Kakade, and Ali Farhadi. Soft threshold weight reparameterization for learnable sparsity. In *International Conference on Machine Learning*, pages 5544–5555. PMLR, 2020.
- [32] Vadim Lebedev and Victor Lempitsky. Fast convnets using group-wise brain damage. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [33] Jeffrey T Leek, W Evan Johnson, Hilary S Parker, Andrew E Jaffe, and John D Storey. The sva package for removing batch effects and other unwanted variation in high-throughput experiments. *Bioinformatics*, 28(6):882–883, 2012.
- [34] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. In *International Conference on Learning Representations*, 2017.

- [35] Qian Li. sctour: a deep learning architecture for robust inference and accurate prediction of cellular dynamics. *Genome Biology*, 24(1):1–33, 2023.
- [36] Arthur Liberzon, Chet Birger, Helga Thorvaldsdóttir, Mahmoud Ghandi, Jill P Mesirov, and Pablo Tamayo. The molecular signatures database hallmark gene set collection. *Cell systems*, 1(6):417–425, 2015.
- [37] Lucas Liebenwein, Ramin Hasani, Alexander Amini, and Daniela Rus. Sparse flows: Pruning continuous-depth models. *Advances in Neural Information Processing Systems*, 34:22628–22642, 2021.
- [38] Junjie Liu, Zhe Xu, Runbin Shi, Ray C. C. Cheung, and Hayden K.H. So. Dynamic sparse training: Find efficient sparse network from scratch with trainable masked layers. In *International Conference on Learning Representations*, 2020.
- [39] Ruishan Liu, Angela Oliveira Pisco, Emelie Braun, Sten Linnarsson, and James Zou. Dynamical systems model of rna velocity improves inference of single-cell trajectory, pseudo-time and gene regulation. *Journal of Molecular Biology*, 434(15):167606, 2022.
- [40] Christos Louizos, Max Welling, and Diederik P Kingma. Learning sparse neural networks through L_0 regularization. *arXiv preprint arXiv:1712.01312*, 2017.
- [41] Christos Louizos, Max Welling, and Diederik P Kingma. Learning sparse neural networks through L_0 regularization. In *International Conference on Learning Representations*, 2018.
- [42] Malte D Luecken, Daniel Bernard Burkhardt, Robrecht Cannoodt, Christopher Lance, Aditi Agrawal, Hananeh Aliee, Ann T Chen, Louise Deconinck, Angela M Detweiler, Alejandro A Granados, et al. A sandbox for prediction and integration of dna, rna, and proteins in single cells. In *35th Conference on Neural Information Processing Systems (NeurIPS 2021) Track on Datasets and Benchmarks*, 2021.
- [43] Eran Malach, Gilad Yehudai, Shai Shalev-Schwartz, and Ohad Shamir. Proving the lottery ticket hypothesis: Pruning is all you need. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 6682–6691. PMLR, 13–18 Jul 2020.
- [44] C. Marvalim, A. Datta, and S. C. Lee. Role of p53 in breast cancer progression: An insight into p53 targeted therapy. *Theranostics*, 13(4):1421–1442, 2023.
- [45] Pedro Mendes, Stefan Hoops, Sven Sahle, Ralph Gauges, Joseph Dada, and Ursula Kummer. Computational modeling of biochemical networks using copasi. *Systems Biology*, pages 17–59, 2009.
- [46] S. M. Mense and L. Zhang. Heme: a versatile signaling molecule controlling the activities of diverse regulators ranging from transcription factors to MAP kinases. *Cell Res*, 16(8):681–692, Aug 2006.
- [47] Laurent Orseau, Marcus Hutter, and Omar Rivasplata. Logarithmic pruning is all you need. In *Advances in Neural Information Processing Systems*, 2020.
- [48] Mansheej Paul, Feng Chen, Brett W. Larsen, Jonathan Frankle, Surya Ganguli, and Gintare Karolina Dziugaite. Unmasking the lottery ticket hypothesis: What’s encoded in a winning ticket’s mask? In *International Conference on Learning Representations*, 2023.
- [49] Ankit Pensia, Shashank Rajput, Alliot Nagle, Harit Vishwakarma, and Dimitris Papailiopoulos. Optimal lottery tickets via subset sum: Logarithmic over-parameterization is sufficient. In *Advances in Neural Information Processing Systems*, volume 33, pages 2599–2610, 2020.
- [50] T. Pramila, W. Wu, S. Miles, W. S. Noble, and L. L. Breeden. The Forkhead transcription factor Hcm1 regulates chromosome segregation genes and fills the S-phase gap in the transcriptional circuitry of the cell cycle. *Genes Dev*, 20(16):2266–2278, Aug 2006.
- [51] Xiaojie Qiu, Yan Zhang, Jorge D Martin-Rufino, Chen Weng, Shayan Hosseinzadeh, Dian Yang, Angela N Pogson, Marco Y Hein, Kyung Hoi Joseph Min, Li Wang, et al. Mapping transcriptomic vector fields of single cells. *Cell*, 185(4):690–711, 2022.

- [52] Vivek Ramanujan, Mitchell Wortsman, Aniruddha Kembhavi, Ali Farhadi, and Mohammad Rastegari. What’s hidden in a randomly weighted neural network? In *Conference on Computer Vision and Pattern Recognition*, 2020.
- [53] Alex Renda, Jonathan Frankle, and Michael Carbin. Comparing rewinding and fine-tuning in neural network pruning. *arXiv preprint arXiv:2003.02389*, 2020.
- [54] Pedro Savarese, Hugo Silva, and Michael Maire. Winning the lottery with continuous sparsification. In *Advances in Neural Information Processing Systems*, 2020.
- [55] Yao-An Shen, Jin Jung, Geoffrey D. Shimberg, Fang-Chi Hsu, Yohan Suryo Rahmanto, Stephanie L. Gaillard, Jiaxin Hong, Jürgen Bosch, Ie-Ming Shih, Chi-Mu Chuang, and Tian-Li Wang. Development of small molecule inhibitors targeting pbx1 transcription signaling as a novel cancer therapeutic strategy. *iScience*, 24(11):103297, 2021.
- [56] Kartik Sreenivasan, Jy yong Sohn, Liu Yang, Matthew Grinde, Alliot Nagle, Hongyi Wang, Eric Xing, Kangwook Lee, and Dimitris Papailiopoulos. Rare gems: Finding lottery tickets at initialization. In *Advances in Neural Information Processing Systems*, 2022.
- [57] G. D. Stormo. Modeling the specificity of protein-DNA interactions. *Quant Biol*, 1(2):115–130, Jun 2013.
- [58] X. Sun, J. Zhang, and Q. Nie. Inferring latent temporal progression and regulatory networks from cross-sectional transcriptomic data of cancer samples. *PLoS Comput Biol*, 17(3):e1008379, Mar 2021.
- [59] D. Szklarczyk, R. Kirsch, M. Koutrouli, K. Nastou, F. Mehryary, R. Hachilif, A. L. Gable, T. Fang, N. T. Doncheva, S. Pyysalo, P. Bork, L. J. Jensen, and C. von Mering. The STRING database in 2023: protein-protein association networks and functional enrichment analyses for any sequenced genome of interest. *Nucleic Acids Res*, 51(D1):D638–D646, Jan 2023.
- [60] Hidenori Tanaka, Daniel Kunin, Daniel L. Yamins, and Surya Ganguli. Pruning neural networks without any data by iteratively conserving synaptic flow. In *Advances in Neural Information Processing Systems*, 2020.
- [61] Alexander Tong, Manik Kuchroo, Shabarni Gupta, Aarthi Venkat, Beatriz Perez San Juan, Laura Rangel, Brandon Zhu, John G Lock, Christine Chaffer, and Smita Krishnaswamy. Learning transcriptional and regulatory dynamics driving cancer cell plasticity using neural ode-based optimal transport. *bioRxiv*, pages 2023–03, 2023.
- [62] Deborah Weighill, Marouen Ben Guebila, Camila Lopes-Ramos, Kimberly Glass, John Quackenbush, John Platig, and Rebekka Burkholz. Gene regulatory network inference as relaxed graph matching. *AAAI Conference on Artificial Intelligence*, 2021.
- [63] Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Learning structured sparsity in deep neural networks. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS’16*, page 2082–2090, Red Hook, NY, USA, 2016. Curran Associates Inc.
- [64] YH Yang and AC Paquet. Preprocessing two-color spotted arrays. In *Bioinformatics and Computational Biology Solutions Using R and Bioconductor*, pages 49–69. Springer, 2005.
- [65] Grace Hui Ting Yeo, Sachit D Saksena, and David K Gifford. Generative modeling of single-cell time series with prescient enables prediction of cell trajectories with interventions. *Nature communications*, 12(1):3222, 2021.
- [66] Hattie Zhou, Janice Lan, Rosanne Liu, and Jason Yosinski. Deconstructing lottery tickets: Zeros, signs, and the supermask. In *Advances in Neural Information Processing Systems*, pages 3597–3607, 2019.

A Supplementary results

A.1 Synthetic data

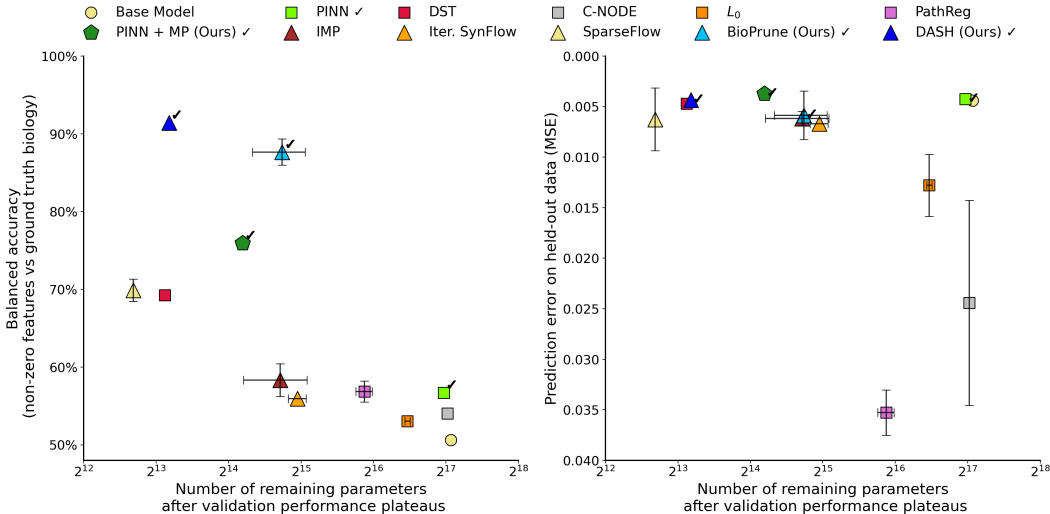


Figure 4: *SIM690* data with 5% noise. We visualize performance of pruning strategies in comparison to original PHOENIX (baseline) in terms of achieved sparsity (x-axis) and balanced accuracy (y-axis) of the recovered gene regulatory network against the ground truth. Error bars are omitted when error is smaller than depicted symbol. Checkmarks (✓) are used to indicate methods that leverage prior information. Ideal models are in the top left quadrant; they recover the true, inherently sparse biological relationships.

Table 3: *Simulation study results – 0% noise*. We provide achieved model sparsity, balanced accuracy of inferred gene regulatory network, and MSE of predicted gene expression dynamics on test data at 0% noise level for SIM350 (the synthetic system of 350 genes). * marks our suggested baselines and method, ✓ marks methods that use prior information for sparsification.

Strategy		Sparsity(%)	Bal. Acc.(%)	MSE (10^{-3})	Epochs
None/Baseline [27]		11.5 ± 0.3	54.8 ± 0.6	3.6 ± 1.7	69 ± 2
Penalty-based (implicit)	L_0 [40]	34.7 ± 2.4	61.3 ± 0.1	6.1 ± 1.8	119 ± 43
	C-NODE [2]	10.7 ± 0.2	60.5 ± 0.1	1.9 ± 0.5	213 ± 5
	PathReg [1]	59.7 ± 1.5	64.2 ± 0.8	6.1 ± 2.3	213 ± 7
	PINN [27] ✓	11.3 ± 0.3	60.3 ± 0.3	2.3 ± 0.4	206 ± 7
	DST[38]	94.3 ± 0.5	72.3 ± 1.2	4.2 ± 1.4	216 ± 86
Pruning-based (explicit)	IMP [18]	86.1 ± 5.1	63.2 ± 1.7	4.1 ± 0.6	251 ± 7
	Iter. SynFlow [60]	79.1 ± 2.1	60.0 ± 0.9	5.8 ± 1.6	323 ± 37
	SparseFlow [37]	95.8 ± 0.3	72.8 ± 0.7	2.9 ± 0.5	195 ± 16
	BioPrune *, ✓	83.5 ± 3.5	88.0 ± 0.5	3.6 ± 0.8	94 ± 22
	DASH *, ✓	92.6 ± 1.2	91.1 ± 1.2	1.9 ± 0.6	164 ± 22
Hybrid	PINN + MP *, ✓	90.0 ± 0.01	89.8 ± 0.3	2.6 ± 0.5	1788 ± 77

Table 4: *Simulation study results – 10% noise.* We provide achieved model sparsity, balanced accuracy of inferred gene regulatory network, and MSE of predicted gene expression dynamics on test data at 10% noise level for SIM350 (the synthetic system of 350 genes). * marks our suggested baselines and method, \checkmark marks methods that use prior information for sparsification.

Strategy		Sparsity(%)	Bal. Acc.(%)	MSE (10^{-3})	Epochs
None/Baseline [27]		1.3 ± 0.3	50.1 ± 0.3	3.5 ± 0.03	55 ± 5
Penalty-based (implicit)	L_0 [40]	31.9 ± 1.6	50.2 ± 0.3	13.3 ± 1.8	156 ± 11
	C-NODE [2]	7.2 ± 0.3	50.5 ± 0.3	36.6 ± 12.5	189 ± 11
	PathReg [1]	47.8 ± 1.8	50.2 ± 1.2	12.7 ± 1.4	224 ± 8
	PINN [27] \checkmark	2.9 ± 0.5	51.3 ± 0.3	5.5 ± 1.9	211 ± 15
	DST[38]	93.3 ± 2.0	67.2 ± 2.6	4.1 ± 1.1	286 ± 33
Pruning-based (explicit)	IMP [18]	79.8 ± 0.1	59.5 ± 2.1	6.7 ± 1.5	240 ± 26
	Iter. SynFlow [60]	79.7 ± 1.3	55.9 ± 1.2	7.8 ± 0.03	165 ± 5
	SparseFlow [37]	89.9 ± 5.0	63.8 ± 3.0	5.1 ± 0.8	142 ± 13
	BioPrune *, \checkmark	79.7 ± 1.3	85.2 ± 0.3	5.6 ± 0.3	67 ± 18
	DASH *, \checkmark	90.8 ± 4.7	85.4 ± 4.2	5.8 ± 1.3	154 ± 4
Hybrid	PINN + MP *, \checkmark	92.0 ± 0.01	83.8 ± 0.7	4.1 ± 1.8	1813 ± 94

Table 5: *Comparison of systems with different number of genes N_g .* We provide achieved model sparsity, balanced accuracy of inferred gene regulatory network, and MSE of predicted gene expression dynamics on test data at 5% noise level for SIM350 and SIM690. * marks our suggested baselines and method, \checkmark marks methods that use prior information for sparsification.

N_g	Strategy		Sparsity(%)	Bal. Acc.(%)	MSE (10^{-3})	Epochs	
350	None/Baseline [27]		7.5 ± 0.1	51.8 ± 0.03	3.0 ± 0.4	67 ± 6	
	Penalty-based (implicit)	L_0 [40]	33.8 ± 4.7	55.0 ± 0.5	8.5 ± 1.0	134 ± 37	
		C-NODE [2]	6.2 ± 0.5	55.9 ± 0.1	2.8 ± 0.6	214 ± 4	
		PathReg [1]	56.5 ± 1.5	61.9 ± 1.0	8.0 ± 1.8	200 ± 29	
		PINN [27] \checkmark	9.9 ± 0.4	58.6 ± 0.7	2.5 ± 0.2	160 ± 7	
		DST[38]	92.8 ± 0.3	71.9 ± 0.5	4.0 ± 0.5	244 ± 61	
	Pruning-based (explicit)	IMP [18]	81.9 ± 6.6	61.7 ± 0.7	4.7 ± 1.1	308 ± 13	
		Iter. SynFlow [60]	79.3 ± 1.2	58.4 ± 0.6	7.0 ± 2.1	271 ± 38	
		SparseFlow [37]	96.0 ± 0.01	70.9 ± 1.5	3.6 ± 0.6	220 ± 3	
		BioPrune *, \checkmark	83.5 ± 1.9	87.3 ± 0.8	2.6 ± 0.9	79 ± 2	
		DASH *, \checkmark	94.6 ± 1.2	90.7 ± 0.4	2.4 ± 1.2	192 ± 24	
	Hybrid	PINN + MP *, \checkmark	87.0 ± 0.01	82.4 ± 0.2	2.3 ± 0.3	1721 ± 50	
	690	None/Baseline [27]		1.1 ± 0.3	50.6 ± 0.1	4.4 ± 0.4	188 ± 17
		Penalty-based (implicit)	L_0 [40]	34.9 ± 1.1	53.0 ± 0.3	12.8 ± 3.1	100 ± 26
			C-NODE [2]	4.3 ± 0.4	54.0 ± 0.2	24.4 ± 10.2	210 ± 2
PathReg [1]			57.7 ± 0.5	57.4 ± 0.2	35.3 ± 2.2	244 ± 17	
PINN [27] \checkmark			7.7 ± 0.1	56.7 ± 0.03	4.3 ± 0.3	166 ± 29	
DST[38]			94.1 ± 0.2	69.2 ± 0.9	4.7 ± 0.4	188 ± 28	
Pruning-based (explicit)		IMP [18]	81.1 ± 5.2	58.3 ± 2.1	6.2 ± 0.7	319 ± 24	
		Iter. SynFlow [60]	77.6 ± 1.4	55.9 ± 0.4	6.7 ± 0.05	215 ± 18	
		SparseFlow [37]	95.8 ± 0.3	69.9 ± 1.5	6.3 ± 3.1	205 ± 10	
		BioPrune *, \checkmark	80.8 ± 4.3	87.7 ± 1.7	5.9 ± 2.4	71 ± 23	
		DASH *, \checkmark	93.9 ± 0.01	91.4 ± 0.2	4.3 ± 0.5	178 ± 2	
Hybrid		PINN + MP *, \checkmark	87.0 ± 0.01	75.9 ± 0.2	3.7 ± 0.2	1657 ± 106	

Table 6: *Sensitivity of DASH to noise in prior.* To understand the impact of the quality of the prior knowledge on the performance of DASH, we show results for different levels of prior corruption in the synthetic data (SIM 350). We keep expression noise constant at 0% to understand the impact of prior corruption alone.

Strategy	Prior corruption	Sparsity(%)	Bal. Acc.(%)	MSE (10^{-3})
None/Baseline	<i>Does not use prior</i>	11.5	54.8	3.6
L_0	<i>Does not use prior</i>	34.7	61.3	6.1
C-NODE	<i>Does not use prior</i>	10.7	60.5	1.9
PathReg	<i>Does not use prior</i>	59.7	64.2	6.1
DST	<i>Does not use prior</i>	94.3	72.3	4.2
IMP	<i>Does not use prior</i>	86.1	63.2	4.1
Iter. SynFlow	<i>Does not use prior</i>	79.1	60.0	2.3
SparseFlow	<i>Does not use prior</i>	95.8	72.8	2.9
PINN	0%	11.3	60.3	2.3
	20%	12.4	60.8	3.1
	40%	11.2	60.6	2.7
BioPrune	0%	83.5	88.0	3.6
	20%	80.9	81.5	7.6
	40%	86.8	80.1	11.1
DASH	0%	92.6	91.1	1.9
	20%	92.4	86.2	6.7
	40%	85.9	79.5	6.1

Table 7: *Prior-informed pruning on an MLP for simulated data.* We compare sparsification strategies on PHOENIX base model and a simple 2-layer MLP base model with ELU activations. We tested on SIM350 with 5% noise. Balanced Accuracy is included since ground truth regulatory structure is known.

Strategy	PHOENIX base model			MLP base model (with ELU)		
	Sparsity	Bal Acc	Test MSE (10^{-3})	Sparsity	Bal Acc	Test MSE (10^{-3})
None	7.5%	51.8%	3.0	0%	50.0%	5.2
L_0	33.8%	55.0%	8.5	29.6%	51.2%	8.1
C-NODE	6.2%	55.9%	2.8	48.8%	50.3%	5.3
PathReg	56.5%	61.9%	8.0	47.1%	55.3%	8.2
DASH	94.6%	90.7%	2.4	89.4%	88.4%	3.8

A.2 Breast cancer data




Figure 5: *BRCA* pathway analysis. We visualize the top-20 significant pathways for each method, showing the pathway z-score (x-axis) and indicate significant results after FWER correction (Bonferroni, p-value cutoff at .05) with *.

A.3 Yeast data



Figure 6: *Yeast pathway analysis*. We visualize the top-20 significant pathways for each method, showing the pathway z-score (x-axis) and indicate significant results after FWER correction (Bonferroni, p-value cutoff at .05) with *.

Strategy	Sparsity	Bal. Acc. (ChIPseq)	Bal. Acc. (TF Perturb)	MSE (10^{-2})
None/Baseline	0.10%	49.87%	49.92%	4.84
L_0	34.43%	48.43%	49.28%	5.33
C-NODE	10.89%	50.04%	50.17%	4.87
PathReg	12.09%	50.11%	49.92%	5.35
PINN ✓	0.17%	49.93%	50.01%	5.77
DST	77.80%	49.92%	50.33%	5.18
IMP	83.22%	49.99%	48.45%	5.46
Iter. SynFlow	85.65%	49.57%	49.77%	5.41
SparseFlow	95.22%	49.89%	51.58%	5.38
BioPrune *, ✓	94.69%	79.23%	64.50%	5.94
DASH *, ✓	97.18%	88.43%	66.79%	5.27
PINN + MP *, ✓	95.01%	55.39%	52.95%	6.09

Table 8: *Balanced accuracies for experiments on yeast data.* Balanced accuracy is based on **1)** transcription factor binding ChIP-seq available for this data [23] and **2)** A TF perturbation network created by Hackett *et al.* based on their TF perturbation experiments [21]. * marks our suggested baselines and method, ✓ marks methods that use prior information for sparsification.

A.4 Bone marrow data

Table 9: *Results on Hematopoiesis data for the Erythroid lineage.* We give sparsity of pruned model and test MSE on predicted gene expression dynamics. No reference gene regulatory network is available to compute the accuracy of the recovered network we hence resort to reporting the average out-degree of nodes in the recovered network. DASH found an optimal λ -value of (0.80, 0.80). * marks our suggested baselines and method, \checkmark marks methods that use prior information for sparsification.

Strategy	Sparsity	OutDeg	Test MSE (10^{-4})
None/Baseline	0.25%	529	2.12
L_0	12.03%	518	2.14
C-NODE	1.33%	529	2.17
PathReg	13.94%	522	2.12
PINN \checkmark	5.93%	529	2.14
DST	90.54%	328	2.23
IMP	73.00%	404	2.26
Iter. SynFlow	82.28%	350	2.18
SparseFlow	94.44%	123	2.12
BioPrune *, \checkmark	87.09%	319	2.14
DASH *, \checkmark	95.95%	54	2.12
PINN + MP *, \checkmark	92.00%	218	2.14

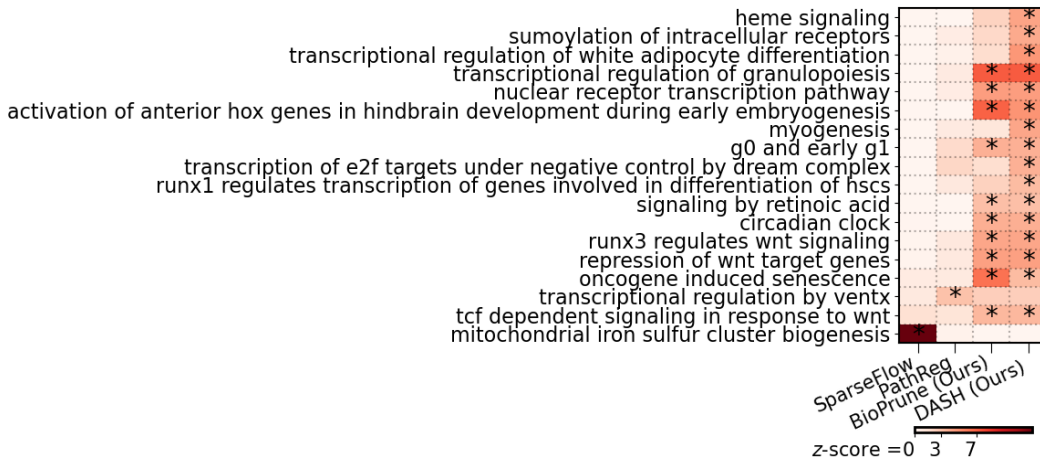


Table 10: *Prior-informed pruning on an MLP for bone marrow data.* We compare sparsification strategies on PHOENIX base model and a simple 2-layer MLP base model with ELU activations. We tested on Erythroid lineage of the bone marrow data.

Strategy	PHOENIX base model		MLP base model (with ELU)	
	Sparsity	Test MSE (10^{-4})	Sparsity	Test MSE (10^{-4})
None	0.25%	2.12	0.00%	2.11
L_0	12.03%	2.14	8.41%	2.29
C-NODE	1.33%	2.17	2.55%	2.22
PathReg	13.94%	2.12	21.48%	2.15
DASH	95.95%	2.12	84.06%	2.13

B Supplementary methods

B.1 Synthetic data generation

The purpose of simulation based data is so that the the underlying dynamical system that produced the this gene expression was known. Do this end, we closely follow the steps outlined by the simulation pipeline provided by [27] to generate reliable synthetic time-series gene expression data from two ground truth gene regulatory networks G_{350} and G_{690} consisting of 350 and 690 genes, respectively.

The pipeline adapts SimulatorGRN [4] to generate from two synthetic *S. cerevisiae* gene regulatory systems (SIM350 and SIM690, consisting of 350 and 690 genes respectively). For every noise setting $\in \{0\%, 5\%, 10\%\}$, the connectivity structure of each *in silico* system is used to synthesize 160 noisy expression trajectories for each gene across $t \in T = \{0, 2, 3, 7, 9\}$. We split up the trajectories into training (88%), validation (6% for tuning λ), and testing (6%). Since the average simulated expression value is ≈ 0.5 , adding Gaussian noise of $\mathcal{N}(0, \sigma^2)$ using $\sigma \in \{0, \frac{1}{40}, \frac{1}{20}\}$ corresponds roughly to average noise levels of $\{0\%, 5\%, 10\%\}$.

B.2 Setup for model training

B.2.1 Model complexity

Since the number of genes k in each problem is different, the number of neurons m in PHOENIX’s hidden layer is chosen to roughly scale with this k according to the original paper [27]:

- SIM350: $k = 350, m = 40$
- SIM690: $k = 690, m = 50$
- Bone marrow data: $k = 529, m = 50$
- Yeast data: $k = 3551, m = 120$
- Breast cancer data: $k = 11165, m = 300$

B.2.2 Initialization and optimizers

For initialization values for each of NN_{sums} , NN_{prods} , $\text{NN}_{\Sigma combine}$, and $\text{NN}_{\Pi combine}$, as well as that of v_i s we choose the default provided by the PHOENIX implementation [27]. The ODESolver (dopri5) and optimizer (Adam) are also chosen as the PHOENIX defaults across all experiments.

B.2.3 Pruning details

We use iterative **pruning schedules** that are initially very aggressive and then become much more gradual. We found this approach to achieve high sparsity without adversely affecting the training dynamics (and subsequently the validation performance).

- SIM350: prune 70% at epoch 3, and then 10% every 10 epochs
- SIM690: prune 70% at epoch 3, and then 10% every 10 epochs
- Bone marrow data: prune 70% at epoch 3, and then 10% every 10 epochs
- Yeast data: prune 90% at epoch 10, and then 10% every 20 epochs
- Breast cancer data: prune 90% at epoch 10, and then 10% every 20 epochs

Weight normalization for DASH pruning scores As described in Section 3, the weight matrices of PHOENIX need to be normalized to $|\widetilde{\mathbf{W}}_{\Sigma}^{(t)}|$, $|\widetilde{\mathbf{W}}_{\Pi}^{(t)}|$, $|\widetilde{\mathbf{U}}_{\Sigma}^{(t)}|$, and $|\widetilde{\mathbf{U}}_{\Pi}^{(t)}|$ in the formula for calculating DASH pruning scores Ω_{Σ} , Ω_{Π} , Ψ_{Σ} , Ψ_{Π} . We perform the following normalizations:

- For $|\widetilde{\mathbf{W}}_{\Sigma}^{(t)}|$ we simply normalize by taking elementwise absolute values of $\mathbf{W}_{\Sigma}^{(t)}$ and dividing all entries by the overall sum of absolute values.

- For $|\widetilde{\mathbf{W}}_{\Pi}^{(t)}|$ we approach similarly, with the only modification that the weights are element-wise exponentiated instead of elementwise absolute value, given that $\widetilde{\mathbf{W}}_{\Pi}^{(t)}$ operates on the log-space.
- For $|\widetilde{\mathbf{U}}_{\Sigma}^{(t)}|$ and $|\widetilde{\mathbf{U}}_{\Pi}^{(t)}|$ we approach again similarly, with the important modification that the gene-specific multipliers (from Section B.10.2) are row-wise multiplied into the weight matrices prior to normalization. This allows the effect of gene multipliers to appropriately be considered when performing pruning.

B.2.4 Learning rates

The learning rate is used as the PHOENIX default of 10^{-3} . We reduce the learning rate by 10% every 3 epochs, unless the validation set performance shows reasonable improvement. **Importantly**, we reset the learning rate back to 10^{-3} immediately after a pruning step is completed, thereby allowing the newly sparsified model to start learning with a higher learning rate.

B.2.5 Stopping criteria

We train for up to 500 epochs on an AWS c5.9xlarge instance, where each epoch consisted of the entire training set being fed to the model, preceded by any pruning step that is prescribed by the pruning schedule. Training is terminated if the validation set performance fails to improve in 40 consecutive epochs. Upon training termination, we have obtained a model that has been iteratively sparsified to an extent that fails to improve the validation set performance. Hence this training procedure **automatically finds an optimal sparsity level** using the validation set.

B.3 Prior knowledge to obtain DASH pruning scores

As mentioned in Section 3, DASH can leverage prior matrices \mathbf{P} and \mathbf{C} to inform its pruning score. We use the following in our experiments:

- SIM350 and SIM690:
 - for synthetic experiments we choose $\mathbf{P} = \mathbf{A}^{\sigma\%}$ to be noisy/corrupted versions (see B.3.1) of the adjacency matrices of ground truth networks G_{350} and G_{690} to reflect that transcription factor binding to target genes can itself be a noisy process in real life. A 1 $\mathbf{A}^{\sigma\%}$ represents prior knowledge of an interaction existing between two genes, and a 0 represented no interaction.
 - For \mathbf{C} we use the outer product $\mathbf{C} = \mathbf{A}^{\sigma\%}(\mathbf{A}^{\sigma\%})^{\top}$, to represent prior knowledge of coregulation. We again applied the corruption/mis-specification procedure from B.3.1 so that \mathbf{C} is also noisy.
- Breast cancer data:
 - For the prior domain knowledge, we set $\mathbf{P} = \mathbf{W}_0$, where \mathbf{W}_0 was a motif map derived from the human reference genome, for the breast tissue specifically, which we obtained through GRAND [3]. \mathbf{W}_0 is a binary matrix with $\mathbf{W}_{0,i,j} \in \{0, 1\}$ where 1 indicates a likely occurrence of a TF sequence motif in the promoter of the target gene, and hence indicating a putative interaction. More simply put, it indicates that whether there is (likely) a binding interface for the protein close to the target gene.
 - Based on information from the STRING database [59], we obtained a protein-protein interaction matrix (PPI) which could use to operationalize our \mathbf{C} matrix, since a PPI is again a binary matrix that is suggestive of which transcription factors have combined (or coregulatory) effects. In a nutshell, the STRING database is a graph with proteins as vertices and knowledge about interactions between two proteins in the graph specifying edges. We set an entry $P_{i,j}$ to 1 if the experimental evidence score on the edge between protein i and j is larger than 0.6, and set $P_{i,j}$ to 0 otherwise.
- Bone marrow data:
 - For the prior domain knowledge, we followed a similar strategy as the breast cancer analysis, and set \mathbf{P} and \mathbf{C} based on the motif map and PPI matrix used in [62]. We

appropriately subsetting \mathbf{P} and \mathbf{C} to only be limited to the $k = 529$ genes that were selected by the PathReg authors [1] in the analysis.

- Yeast data:
 - For prior domain knowledge model we set \mathbf{P} to reflect the regulatory network structure of a motif map. The map is based on predicted binding sites for 204 yeast transcription factors (TFs) [23]. These data include 4360 genes with tandem promoters. 3551 of these genes are also covered on the yeast cell cycle gene expression array. 105 total TFs in this data set target the promoter of one of these 3551 genes. The motif map between these 105 TFs and 3551 target genes provides the adjacency matrix \mathbf{A} of 0s and 1s, representing whether or not a prior interaction is likely between TF and gene.
 - We set \mathbf{C} to be the PPI matrix used for the same data in the PANDA paper [20]. We appropriately subsetting \mathbf{C} to only be limited to the $k = 3551$ genes that were in the data.

B.3.1 Creating corrupted/misspecified prior models for synthetic data

For each noise level $\sigma\% \in \{0\%, 5\%, 10\%\}$ in our *in silico* experiments, we created a shuffled version of G_{350} (and similarly G_{690}) where we shuffled $\sigma\%$ of the edges by relocating those edges to new randomly chosen origin and destination genes within the network. This yielded the shuffled network $G_{350}^{\sigma\%}$ (and similarly $G_{690}^{\sigma\%}$) with corresponding adjacency matrix $\mathbf{A}^{\sigma\%}$. We additionally performed sensitivity analyses using $\sigma\% \in \{20\%, 40\%\}$ to investigate the effect of **even** higher levels of prior corruption. We then used $\mathbf{A}^{\sigma\%}$ to obtain “corrupted“ \mathbf{P} and \mathbf{C} as described in B.3.

B.4 Validation and testing

The choice of $\boldsymbol{\lambda} = (\lambda_1, \lambda_2)$ is important for optimally combining prior information with model weights. Hence we implement a K -fold cross validation approach to choose $\boldsymbol{\lambda}$. Test set performance is measured as the mean squared error between predictions and held-out expression values in the test set.

B.5 Measuring biological alignment of sparsified models

To validate biological alignment of trained and sparsified models, we extracted GRNs from each models (as explained in B.10.4), and compared back to the **validation networks**. Specifically, once we extracted a GRN from the trained model, we looked at how well 0s vs non-zeros in that network aligned with 0s vs non-zeros in the validation network. Our comparison metric was **balanced accuracy**, which is the average of the true positive and true negative rates. The validation networks were as follows:

- SIM350, SIM690: We used the ground truth networks G_{350} and G_{690} .
- Breast cancer data: We used ChIP-seq data from the MCF7 cell line (breast cancer) in the ReMap2018 database [10] to create a validation network of TF-target interactions.
- Bone marrow data: As also noted by [1], validation network was not available, so we resorted to the pathway analysis.
- Yeast data: We used two kinds of validation networks
 1. ChIP-seq data [23] to create a network of TF-target interactions, and used this as a validation network to test explainability. The targets of transcription factors in this ChIP-chip data set were filtered using the criterion $p < 0.001$.
 2. A TF perturbation network created by Hackett *et al.*, who fit dynamical systems to their TF perturbation experiments [21].

B.6 Strategy to potentially extend DASH to arbitrary number of layers

Supposing we only have access to prior knowledge in the form of putative prior effect sizes between the n inputs and o outputs $\mathbf{P} \in \mathbb{R}^{o \times n}$. Then, for an NN with $L - 1$ layers $\mathbf{W}_1 \in \mathbb{R}^{m_1 \times n}$, $\mathbf{W}_2 \in \mathbb{R}^{m_2 \times m_1}, \dots, \mathbf{W}_L \in \mathbb{R}^{o \times m_L}$, we can adopt a strategy where we consider the pruning scores to be fixed for all but one layer.

Since the product $\mathbf{W}_L^{(t)} \dots \mathbf{W}_2^{(t)} \cdot \mathbf{W}_1^{(t)} \in \mathbb{R}^{o \times n}$ represents the overall flow of information from inputs to outputs at epoch t , we surmise that $\Omega_L^{(t)} \dots \Omega_2^{(t)} \cdot \Omega_1^{(t)} \in \mathbb{R}^{o \times n}$ should reflect \mathbf{P} . We can thus prune as follows:

1. Starting with the last layer, we **fix** the pruning scores of all other layers and compute as follows:

$$\Omega_L^{(t)} := (1 - \lambda_L) |\widetilde{\mathbf{W}}_L^{(t)}| + \lambda_L \left| \mathbf{P} \cdot \text{PInv}_R \left(\Omega_{L-1}^{(t)} \dots \Omega_2^{(t)} \Omega_1^{(t)} \right) \right|.$$

2. For the middle layers $l \in \{2, 3, \dots, L-1\}$, we do:

$$\Omega_l^{(t)} := (1 - \lambda_l) |\widetilde{\mathbf{W}}_l^{(t)}| + \lambda_l \left| \text{PInv}_L \left(\Omega_L^{(t)} \dots \Omega_{l+2}^{(t)} \Omega_{l+1}^{(t)} \right) \cdot \mathbf{P} \cdot \text{PInv}_R \left(\Omega_{l-1}^{(t)} \dots \Omega_2^{(t)} \Omega_1^{(t)} \right) \right|.$$

3. The first layer can be pruned using:

$$\Omega_1^{(t)} := (1 - \lambda_1) |\widetilde{\mathbf{W}}_1^{(t)}| + \lambda_1 \left| \text{PInv}_L \left(\Omega_L^{(t)} \dots \Omega_3^{(t)} \Omega_2^{(t)} \right) \cdot \mathbf{P} \right|.$$

B.7 Processing steps for real data

B.7.1 Breast cancer

The original data set comes from a cross-sectional breast cancer study (GEO accession **GSE7390** [13]) consisting of microarray expression values for 22000 genes from 198 breast cancer patients, that is sorted along a pseudotime axis. We note that the same data set was also ordered in pseudotime by [58] in the PROB paper. For consistency in pseudotime inference, we obtained the same version of this data that was already preprocessed and sorted by PROB. We normalized the expression values to be between 0 and 1. We limited our analysis to the genes that had measurable expression and appeared in the aforementioned motif map and PPI matrices. This resulted in a pseudotrajectory of expression values for 11165 genes across 186 patients. We removed a contiguous interval of expression across 8 time points for testing (5%), and split up the remaining 178 time points into training (170, 90%) and validation for tuning λ_{prior} (8, 5%).

B.7.2 Yeast

GPR files were downloaded from the Gene Expression Omnibus (accession **GSE4987** [50]), and consisted of two dye-swap technical replicates measured every five minutes for 120 minutes. Each of two replicates were separately ma-normalized using the `maNorm()` function in the `mararray` library in R/Bioconductor [64]. The data were batch-corrected [29] using the `ComBat()` function in the `sva` library [33] and probe-sets mapping to the same gene were averaged, resulting in expression values for 5088 genes across fifty conditions. Two samples (corresponding to the 105 minute time point) were excluded for data-quality reasons, as noted in the original publication, and genes without motif information were then removed, resulting in an expression data set containing 48 samples (24 time points in each replicate) and 3551 genes.

B.7.3 Bone marrow

The data is originally from [42] (GEO accession code = **GSE194122**). The cleaning, preprocessing, and pseudotime analysis and was appropriately performed by [1] in the PathReg paper, and made publicly available, allowing us to access the processed version. Importantly, [1] split up the data into 3 different lineages (Erythroid, Monocyte, and B-Cell), and we fit a separate PHOENIX model on each lineage. The set contains 5 separate batches of data for each lineage, we used 1 for training (batch S1D2), 1 for validation (batch S1D1) and 3 for testing (batches S1D1, S2D4, and S3D6).

B.8 Pathway analyses for breast cancer, yeast, and bone marrow datasets

We followed very closely the steps below from the Methods section of the PHOENIX paper [27] in order to compute pathway scores, with the only difference that we compute scores between different sparsification strategies Pru.

B.8.1 Gene influence scores

Given \mathcal{M}_{PrU} a PHOENIX model trained on a dataset consisting of k genes, and sparsified using the pruning strategy PrU (for $\text{PrU} \in \{\text{DASH}, \text{IMP}, \text{C-NODE}, \dots, \text{PathReg}\}$), we performed perturbation analyses to compute gene influence scores $\mathcal{IS}_{\text{PrU},j}$. We randomly generated 200 initial ($t = 0$) expression vectors via i.i.d standard uniform sampling $\{\mathbf{g}(\mathbf{0})_r \in \mathbb{R}^k\}_{r=1}^{200}$. Next, for each gene j in \mathcal{M}_{PrU} , we created a perturbed version of these initial value vectors $\{\mathbf{g}^j(\mathbf{0})_r\}_{r=1}^{200}$, where only gene j was perturbed in each unperturbed vector of $\{\mathbf{g}(\mathbf{0})_r\}_{r=1}^{200}$. We then fed both sets of initial values into \mathcal{M}_{PrU} to obtain two sets of predicted trajectories $\{\{\widehat{\mathbf{g}}(t)_r \in \mathbb{R}^k\}_{t \in T}\}_{r=1}^{200}$ and $\{\{\widehat{\mathbf{g}}^j(t)_r \in \mathbb{R}^k\}_{t \in T}\}_{r=1}^{200}$ across a set of time points T . We calculated influence as the average absolute difference between the two sets of predictions, that represented how changes in initial ($t = 0$) expression of gene j affected subsequent ($t > 0$) predicted expression of all other genes in the PrU-dimensional system

$$\mathcal{IS}_{\text{PrU},j} = \frac{1}{200} \sum_{r=1}^{200} \left[\frac{1}{|T|} \sum_{\substack{t \in T \\ t \neq 0}} \left(\frac{1}{k} \sum_{\substack{i=1 \\ i \neq j}}^k |\widehat{g}_i(t)_r - \widehat{g}_i^j(t)_r| \right) \right].$$

B.8.2 Pathway influence scores

Having computed gene influence scores $\mathcal{IS}_{\text{PrU},j}$ for each gene j in each dynamical system of dimension k genes sparsified with method PrU, we translated these gene influence scores into pathway influence scores. We used the Reactome pathway data set, GO biological process terms, and GO molecular function terms from MSigDB [36], that map each biological pathway/process, to the genes that are involved in it. For each system sparsified by PrU, we obtained the pathway (p) influence scores ($\mathcal{PS}_{\text{PrU},p}$) as the sum of the influence scores of all genes involved in pathway p

$$\mathcal{PS}_{\text{PrU},p} = \sum_{j \in p} \mathcal{IS}_{\text{PrU},j}.$$

We statistically tested whether each pathway influence score is higher than expected by chance using empirical null distributions. We randomly permuted the gene influence scores across the genes to recompute "null" values $\mathcal{PS}_{\text{PrU},p}^0$. For each pathway, we performed $Q = 1000$ permutations to obtain a null distribution $\{\mathcal{PS}_{\text{PrU},p,q}^0\}_{q=1}^Q$ that can be compared to $\mathcal{PS}_{\text{PrU},p}$. We could then compute an empirical p -value as $p = \frac{1}{Q} \sum_{q=1}^Q \mathbb{I}_{\mathcal{PS}_{\text{PrU},p,q}^0 > \mathcal{PS}_{\text{PrU},p}}$, where \mathbb{I} is the indicator function. Finally, we used the mean ($\mu_{0(\text{PrU},p)}$) and variance ($\sigma_{0(\text{PrU},p)}^2$) of the null distribution $\{\mathcal{PS}_{\text{PrU},p,q}^0\}_{q=1}^Q$ to obtain and visualize pathway z -scores that are now comparable across pathways (p) and sparsification strategies (PrU)

$$z_{(\text{PrU},p)} = \frac{\mathcal{PS}_{\text{PrU},p} - \mu_{0(\text{PrU},p)}}{\sqrt{\sigma_{0(\text{PrU},p)}^2}}.$$

B.9 Implementation details for other sparsification strategies on the PHOENIX architecture

B.9.1 Iterative magnitude pruning

As discussed in Section 3, IMP can be operationalized as a special case of DASH by setting $\lambda_1 = \lambda_2 = 0$.

B.9.2 PINN

This is simply the PHOENIX model equipped with the prior-informed loss term. This loss-term in the original PHOENIX paper [27] is inspired by Physics-informed neural networks (PINNs).

B.9.3 PINN + MP

Once a PHOENIX model (**trained including the prior-informed loss term, i.e. the PINN term**) is fully trained (without any pruning), we inspect the trained model and pruned the lowest $p\%$ of parameters in each of \mathbf{W}_Σ , \mathbf{W}_Π , \mathbf{U}_Σ , \mathbf{U}_Π based on on the **normalized** weights (see B.2) to 0. We then fine-tune (i.e. retrain **without** training the pruned parameters) this $p\%$ sparsified model and

calculate its performance on the validation set. We repeat this process for a grid of values for $p \in \{0.50, 0.75, 0.83, 0.87, 0.90, 0.92, 0.95, 0.97, 0.99\}$. The validation set can then inform the best value of p . We repeated this entire procedure 3 times, so that we could apply the 1 standard error rule [24] and choose the optimal p as the **sparsest** fine-tuned model whose validation MSE is within 1 standard error of lowest obtained average validation MSE.

B.9.4 Penalty based methods

C-NODE, PathReg, and L_0 implementations were obtained from the code associated with the PathReg paper [1]. We adapted the code so that the base NN architecture was exactly that of the PHOENIX model, including an implementation of the gene-specific multipliers (from Section B.10.2). Finally, we tuned the relevant parameters λ_0 and λ_1 in the objective function using the validation set.

The code for DST was obtained from: <https://github.com/junjieliu2910/DynamicSparseTraining>

B.10 A brief overview of PHOENIX NeuralODE model

The following are adapted from [27] and provided here for the reader’s convenience.

B.10.1 Neural ordinary differential equations

NeuralODEs [8] learn dynamical systems by parameterizing the underlying derivatives with neural networks: $\frac{d\mathbf{g}(t)}{dt} = f(\mathbf{g}(t), t) \approx \text{NN}_\Theta(\mathbf{g}(t), t)$. Given an initial condition $\mathbf{g}(t_0)$, the output $\mathbf{g}(t_i)$ at any given time-point t_i can now be approximated using a numerical ODE solver of adaptive step size:

$$\widehat{\mathbf{g}(t_1)} = \mathbf{g}(t_0) + \int_{t_0}^{t_1} \text{NN}_\Theta(\mathbf{g}(t), t) dt.$$

A loss function $L(\mathbf{g}(t_1); \mathbf{g}(t_0) + \int_{t_0}^{t_1} \text{NN}_\Theta(\mathbf{g}(t), t) dt)$ is then optimized for Θ via back propagation, using the adjoint sensitivity method [7] to carry the backpropagation through the integration steps of the ODESolver.

B.10.2 PHOENIX - overview

PHOENIX models gene expression dynamics using NeuralODEs. Notably, for an expression vector of r genes, PHOENIX models both additive and multiplicative regulatory effects using two parallel linear layers with m neurons each: NN_{sums} (with weights $\mathbf{W}_\Sigma \in \mathbb{R}^{m \times r}$, and biases $\mathbf{b}_\Sigma \in \mathbb{R}^m$) and NN_{prods} ($\mathbf{W}_\Pi \in \mathbb{R}^{m \times r}$, $\mathbf{b}_\Pi \in \mathbb{R}^m$). Here, NN_{sums} and NN_{prods} are equipped with activation functions that model the Hill equation

$$\phi_\Sigma(x) = \frac{x - 0.5}{1 + |x - 0.5|} \quad \text{and} \quad \phi_\Pi(x) = \log(\phi_\Sigma(x) + 1).$$

The Hill equation is a classical formula in biochemistry that models molecular binding in dependence of concentration. This results in outputs of the two parallel layers

$$\begin{aligned} \mathbf{c}_\Sigma(\mathbf{g}(t)) &= \mathbf{W}_\Sigma \phi_\Sigma(\mathbf{g}(t)) + \mathbf{b}_\Sigma \quad \text{and} \\ \mathbf{c}_\Pi(\mathbf{g}(t)) &= \exp \circ (\mathbf{W}_\Pi \phi_\Pi(\mathbf{g}(t)) + \mathbf{b}_\Pi). \end{aligned}$$

As shown above, $\phi_\Pi(x)$ yields the output of NN_{prods} in the log space and is subsequently exponentiated in \mathbf{c}_Π to represent multiplicative effects in the linear space. The outputs $\mathbf{c}_\Sigma(\mathbf{g}(t))$ and $\mathbf{c}_\Pi(\mathbf{g}(t))$ are then separately fed into two more parallel linear layers $\text{NN}_{\Sigma combine}$ (with weights $\mathbf{U}_\Sigma \in \mathbb{R}^{r \times m}$) and $\text{NN}_{\Pi combine}$ ($\mathbf{U}_\Pi \in \mathbb{R}^{r \times m}$), respectively. The outputs of $\text{NN}_{\Sigma combine}$ and $\text{NN}_{\Pi combine}$ are summed to obtain

$$\mathbf{c}_\cup(\mathbf{g}(t)) = \mathbf{U}_\Sigma \mathbf{c}_\Sigma(\mathbf{g}(t)) + \mathbf{U}_\Pi \mathbf{c}_\Pi(\mathbf{g}(t)).$$

Finally, PHOENIX includes gene-specific multipliers $\mathbf{v} \in \mathbb{R}^r$ for modeling steady states of genes that do not exhibit any temporal variation ($\frac{d\mathbf{g}_i(t)}{dt} = 0, \forall t$). Accordingly, the output for each gene i in $\mathbf{c}_\cup(\mathbf{g}(t))$ is multiplied with $\text{ReLU}(v_i)$ in the final estimate for the local derivative

$$\text{NN}_\Theta(\mathbf{g}(t), t) = \text{ReLU}(\mathbf{v}) \odot [\mathbf{c}_\cup(\mathbf{g}(t)) - \mathbf{g}(t)].$$

Although PHOENIX achieves some sparsity in its weight matrices ($\mathbf{W}_\Sigma, \mathbf{W}_\Pi, \mathbf{U}_\Sigma, \mathbf{U}_\Pi$) **without any external sparsification strategy**, the achieved sparsity level is, however, at most 12% (see Figure 2, and Tables 1, 3, 4).

B.10.3 Prior knowledge incorporation in base PHOENIX model itself

PHOENIX has the option to promote the NeuralODE to flexibly align with structural domain knowledge, while still explaining the observed gene expression data. This is operationalized via a modified loss function

$$\begin{aligned} \mathcal{L}_{\text{mod}}(\mathbf{g}(t_1), \widehat{\mathbf{g}}(t_1)) = & \overbrace{\tau L(\mathbf{g}(t_1); \mathbf{g}(t_0) + \int_{t_0}^{t_1} \text{NN}_{\Theta}(\mathbf{g}(t_1), t) dt)}^{\text{loss based on matching observed gene expression data}} \\ & + \overbrace{(1 - \tau) L(\mathcal{P}^*(\mathbf{g}(t_1)); \text{NN}_{\Theta}(\mathbf{g}(t_1), t))}^{\text{loss based on matching domain-knowledge}} \end{aligned}$$

that incorporates the effect of any user-provided prior model \mathcal{P}^* , using a tuning parameter τ , and the original loss function $L(\mathbf{x}, \widehat{\mathbf{x}})$. PHOENIX implements \mathcal{P}^* as a simple linear model $\mathcal{P}^*(\gamma) = \mathbf{A} \cdot \gamma - \gamma$, where \mathbf{A} is the adjacency matrix of likely connectivity structure based on prior domain knowledge (such as experimentally validated interactions) with $\mathbf{A}_{ij} \in \{+1, -1, 0\}$ representing an activating, repressive, or no prior interaction, respectively.

For synthetic experiments, we used the simple linear model: $\mathcal{P}^*(\gamma) = \mathbf{A}^{\sigma\%} \cdot \gamma - \gamma$, where we chose $\mathbf{A}^{\sigma\%}$ to be noisy/corrupted versions of the adjacency matrices of ground truth networks G_{350} and G_{690} (details in B.3.1). We set activating and repressive edges in $\mathbf{A}^{\sigma\%}$ to 1, while “no interaction” was represented using 0.

B.10.4 Algorithm for efficiently retrieving encoded GRN from trained PHOENIX model

We start with PHOENIX’s prediction for the local derivative given a gene expression vector $\mathbf{g}(t) \in \mathbb{R}^r$ in an r -gene system:

$$\begin{aligned} \frac{d\widehat{\mathbf{g}}(t)}{dt} &= \text{ReLU}(\mathbf{v}) \odot [\mathbf{W}_{\cup} \{ \mathbf{c}_{\Sigma}(\mathbf{g}(t)) \oplus \mathbf{c}_{\Pi}(\mathbf{g}(t)) \} - \mathbf{g}(t)], \text{ where} \\ \mathbf{c}_{\Sigma}(\mathbf{g}(t)) &= \mathbf{W}_{\Sigma} \phi_{\Sigma}(\mathbf{g}(t)) + \mathbf{b}_{\Sigma} \text{ and } \mathbf{c}_{\Pi}(\mathbf{g}(t)) = \exp \circ (\mathbf{W}_{\Pi} \phi_{\Pi}(\mathbf{g}(t)) + \mathbf{b}_{\Pi}) \end{aligned}$$

A trained PHOENIX model encodes interactions *between* genes primarily within the gene-specific multipliers $\mathbf{v} \in \mathbb{R}^r$, and the weight parameters from its neural network blocks $\mathbf{W}_{\Pi}, \mathbf{W}_{\Sigma} \in \mathbb{R}^{m \times r}$ and $\mathbf{W}_{\cup} \in \mathbb{R}^{r \times 2m}$. This inspired an efficient means of projecting the estimated dynamical system down to a gene regulatory network (GRN) \widehat{G}_n . In particular a matrix $\mathbf{D} \in \mathbb{R}^{r \times r}$ is calculated, where \mathbf{D}_{ij} approximated the *absolute contribution* of gene j to the derivative of gene i ’s expression

$$\mathbf{D} = \mathbf{W}_{\cup} \begin{bmatrix} \mathbf{W}_{\Sigma} \\ \mathbf{W}_{\Pi} \end{bmatrix}.$$

Gene-specific multipliers \mathbf{v} are applied, before adapting the marginal attribution approach described by Hackett *et al.* [21]. This resulted in the **dynamics matrix** $\widetilde{\mathbf{D}}$ where $\widetilde{\mathbf{D}}_{ij}$ was scaled according to the *relative contribution* of gene j to the rate of change in gene i ’s expression:

$$\widetilde{\mathbf{D}}_{ij} = \frac{\mathbf{v}_i \mathbf{D}_{ij}}{\sum_{j'=1}^n |\mathbf{v}_i \mathbf{D}_{ij'}|}.$$

B.11 Subjecting PHOENIX to DASH pruning

With its simple yet powerful architecture, PHOENIX provides an ideal base setting for applying and testing the merits of the discussed neural network sparsification strategies (Section 2), including DASH. We discuss the DASH implementation here and provide implementation details for other sparsification strategies in Appendix B.9.

For DASH, we follow the steps of the $L = 2$ case in Section 3. Note that in PHOENIX we have $k = r$, since we are modeling how a set of k genes (inputs) affects each others derivatives (i.e. k outputs). We apply the steps from Section 3 on each parallel member of the first $\{\mathbf{W}_{\Sigma}, \mathbf{W}_{\Pi}\}$ and second $\{\mathbf{U}_{\Sigma}, \mathbf{U}_{\Pi}\}$ layers. For the first layer, we leverage prior knowledge of transcription

factor coregulation (often available in the form of protein-protein interaction matrices) to formulate $C \in \mathbb{R}^{k \times k}$. This is then used to calculate pruning scores $\Omega_\Sigma, \Omega_\Pi$ for W_Σ, W_Π . Similarly, for U_Σ, U_Π in the second layer, we utilize $P \in \mathbb{R}^{r \times k}$ which are easily obtainable motif map matrices encoding prior knowledge of transcription factor binding sites around genes to calculate pruning scores Ψ_Σ, Ψ_Π .

Algorithm 1 Computing $\Omega_\Sigma^{(t)}, \Omega_\Pi^{(t)}, \Psi_\Sigma^{(t)}, \Psi_\Pi^{(t)}$

Inputs: weights $W_\Sigma^{(t)}, W_\Pi^{(t)}, U_\Sigma^{(t)}, U_\Pi^{(t)}$; priors C, P ; epoch t ;
previous scores $\Omega_\Sigma^{(t-1)}, \Omega_\Pi^{(t-1)}$; tuning λ_1, λ_2

- Normalize $W_\Sigma^{(t)}$ to get $|\widetilde{W_\Sigma^{(t)}}|$ (see Appendix B.2)
 - Similarly obtain $|\widetilde{W_\Pi^{(t)}}|, |\widetilde{U_\Sigma^{(t)}}|, \text{ and } |\widetilde{U_\Pi^{(t)}}|$
 - If $t = 0$
 - Initialize $\Omega_\Sigma^{(t)}$ and $\Omega_\Pi^{(t)}$ randomly with values from a standard Gaussian
 - Else
 - $\Omega_\Sigma^{(t)} := (1 - \lambda_1)|\widetilde{W_\Sigma^{(t)}}| + \lambda_1 \left| \text{PInv}_L \left(\Omega_\Sigma^{(t-1)\top} \right) \cdot C \right|$
 - $\Omega_\Pi^{(t)} := (1 - \lambda_1)|\widetilde{W_\Pi^{(t)}}| + \lambda_1 \left| \text{PInv}_L \left(\Omega_\Pi^{(t-1)\top} \right) \cdot C \right|$
 - $\Psi_\Sigma^{(t)} := (1 - \lambda_2)|\widetilde{U_\Sigma^{(t)}}| + \lambda_2 \left| P \cdot \text{PInv}_R \left(\Omega_\Sigma^{(t)} \right) \right|$
 - $\Psi_\Pi^{(t)} := (1 - \lambda_2)|\widetilde{U_\Pi^{(t)}}| + \lambda_2 \left| P \cdot \text{PInv}_R \left(\Omega_\Pi^{(t)} \right) \right|$
-

B.12 Ablation study using a plain MLP instead of PHOENIX as the base model

We believe that DASH should remain performant even when using a base model **that is different from PHOENIX**. To be more explicit, the base model would give a new expression for $\text{NN}_\Theta(g(t), t)$ in B.10.2.

$$\text{For PHOENIX as base model we have: } \text{NN}_\Theta(g(t), t) = \text{ReLU}(v) \odot \left[c_\cup(g(t)) - g(t) \right].$$

For alternative base model we have: $\text{NN}_\Theta(g(t), t) = \text{alternative base model}(g(t)) - g(t)$

In our ablation experiments, we used a simple **two-layer MLP with ELU activation functions** as the alternative base model. The choice of ELU activation is motivated by the PathReg paper [1]. We chose the number of hidden neurons such that the total number of parameters was comparable between the MLP and the PHOENIX base models. We tested a few sparsification strategies against DASH on this new base model, for both synthetic data and the bone marrow data.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We explain the new concept of prior-informed pruning (Fig. 1) in Section 3, introduce it to the state-of-the-art model of gene regulatory dynamics in Section 4 and show in Extensive Experiments in Sec. 5 and App. A that it finds more meaningful and domain relevant models.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We provide a short discussion in Sec. 6, pointing to potential new direction of future work applying DASH to different domains, and provide critical analysis in the Experiments (see e.g. closing remarks of synthetic data study).

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: There are no theorems nor proofs in this paper.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide discussion of synthetic data generation (App. B.1), setup of training (App. B.2), and information about how prior knowledge was defined (App. B.3). We further provide information on which reference gold standard was used for evaluation of the inferred gene regulatory network (App. B.5), any pre-processing steps for the real world data (App. B.7) and how analyses were carried out (App. B.8). We further discuss all necessary details on how different pruning methods were implemented (App. B.9) and how regulatory networks were extracted (App. B.10.4). We provide an implementation of DASH as supplementary material.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).

- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide code and data via GitHub: <https://github.com/QuackenbushLab/DASH> and reference where all data can be found in the text (note that all real data is publicly available through the respective original authors).

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: See questions above.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We provide error ranges for all synthetic data experiments and provide statistical significance (multiple test corrected) for pathway analyses.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [No]

Justification: All experiments were carried out on the same machine. As these are comparably small-scale experiments (i.e., network parameters are in the thousands) most standard hardware will easily reproduce these experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: We follow the NeurIPS Ethics Guidelines.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [No]

Justification: We do not see a direct harmful application of our suggested approach, it is a foundational line of research and our suggested application is in molecular biology. As such, it will rather bring a positive societal impact by improving understanding of health and disease and thereby improving therapy design.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Not applicable

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [No]

Justification: We properly cite the work for the compared methods and all data used within the scope of this work. We also checked that each license is applicable. **But we do not list each license and terms explicitly.**

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The primary contribution of this work is foundational, yet we do provide code and data-generating scripts as supplementary material.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: All data used within the experiments was from **public sources and databases**, which have undergone reviews to ensure no risks are involved prior to their respective publication.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.