Published in final edited form as:

J Am Stat Assoc. 2016 ; 111(513): 73–92. doi:10.1080/01621459.2015.1110523.

Perturbation Detection Through Modeling of Gene Expression on a Latent Biological Pathway Network: A Bayesian hierarchical approach

Lisa M. Pham, Luis Carvalho, Scott Schaus, and Eric D. Kolaczyk

Boston University, Boston, MA 02215

Lisa M. Pham: lisa@capellabio.com

Abstract

Cellular response to a perturbation is the result of a dynamic system of biological variables linked in a complex network. A major challenge in drug and disease studies is identifying the key factors of a biological network that are essential in determining the cell's fate.

Here our goal is the identification of perturbed pathways from high-throughput gene expression data. We develop a three-level hierarchical model, where (i) the first level captures the relationship between gene expression and biological pathways using confirmatory factor analysis, (ii) the second level models the behavior within an underlying network of pathways induced by an unknown perturbation using a conditional autoregressive model, and (iii) the third level is a spike-and-slab prior on the perturbations. We then identify perturbations through posterior-based variable selection.

We illustrate our approach using gene transcription drug perturbation profiles from the DREAM7 drug sensitivity predication challenge data set. Our proposed method identified regulatory pathways that are known to play a causative role and that were not readily resolved using gene set enrichment analysis or exploratory factor models. Simulation results are presented assessing the performance of this model relative to a network-free variant and its robustness to inaccuracies in biological databases.

Keywords

Bayesian Factor Models; Confirmatory Factor Analysis; Conditional Autoregressive Models; MCMC; Network Biology; Drug Target Prediction; Microarray

1. INTRODUCTION

With the influx of high-throughput genomic data, understanding biological mechanisms of action that cause changes in cellular homeostasis has become a reachable challenge in the fields of bioinformatics, computational biology, and statistics. High-throughput measurement techniques such as transcriptional profiling allow us to measure gene transcript levels across thousands of genes simultaneously. However, analyzing individual gene

transcriptional profiles, cannot, by themselves, elucidate biological mechanisms that are responsible for the changes observed in gene expression. Rather, gene transcription provides only one facet of a multifaceted system of biological variables that culminate to form a cellular response.

Mechanisms of action (MoA) that drive cellular dysregulation can arise in several biological contexts. Chemotherapeutic compounds for instance, alter very distinct mechanisms, such as changing the topology of DNA structure induced by specific topoisomerase inhibitors (e.g. Malik et al., 2006; Nakada et al., 2006), or inhibiting cellular motility mechanisms caused by myosin II inhibitor compounds (e.g. (Allingham et al., 2005)). In another example, cancer metastasis is also caused by aberrant pathways that disrupt normal cellular regulation, resulting in cancer proliferation. The identification of such key mechanisms is important as they can provide unique signatures not readily apparent by directly analyzing gene expression without added structure or biological context. Methods that leverage biological information by incorporating added structure or biology can be used as a diagnostic tool in a clinical context.

Our primary goal in this paper is to identify drug targets and mechanisms of action in drug perturbation experiments. That is, we aim to find pathways significantly related to each drug's MoA, which can enhance both therapeutic benefit and assessment of efficacy. Moreover, understanding a drug's MoA and biological pathway information can provide a better assessment of cellular response to drugs, possibly providing a therapeutic profile for each drug selection.

1.1 The drug target problem: an inverse problem

A cellular response is the result of a culmination of interactions between genes/proteins. A single perturbed pathway, for instance, can cause a rippling effect across a global network of interactions, leaving behind a cascade of transcriptional dysregulation across the genome. For instance, Fig. 1 is a schematic illustration of how a drug perturbs a system of interacting pathways. In Fig. 1, the inhibition/activation of a protein (s) in a single pathway consequently alters several downstream pathways, leading to changes in gene expression. To better understand the primary targets of a perturbation source, we propose a method that formally filters out these regulatory dependencies between biological factors (pathways) to uncover the primary target underlying a perturbation response.

From a mathematical perspective, the nature of the problem we address in this paper is not unlike that of 'deconvolution' in image processing, a similarity that has been noted by others in this area (e.g., 'drug target deconvolution' (Terstappen et al., 2007)). In the image processing version of this problem, an image, say f , is of interest but one has available only blurred and noisy measurements, say $y = Kf + e$. While denoising y can be relatively straightforward, it only leaves one with an estimate of the blurred image, Kf . In order to recover f itself, the effect of the blurring operator K must be inverted. However, even in the ideal case where K is known this inversion can be ill-posed and the recovery of f can be severely degraded by the corresponding inflation of the noise e . When K is unknown or only partially known, as is analogous to what we face in the drug target prediction problem, the degradation can be arbitrarily worse.

1.2 Identifying Pathway Targets in the DREAM 7 Drug Sensitivity Prediction Challenge Data Set

For our purposes of target pathway identification in drug perturbation experiments, we explore the NCI DREAM7 drug sensitivity prediction challenge dataset (Bansal et al., 2014) which is a part of the Dialogue for Reverse Engineering Assessments and Methods (DREAM) challenge series (Marbach et al., 2012; Prill et al., 2011). To assess the performance of our method, we focus our attention on the DREAM7 drug sub-challenge 2 dataset (Bansal et al., 2014) which consists of microarray gene expression profiles from the LY3 cancer cell line. Exactly 14 drugs were tested at different concentrations and durations, and were compared to their mock control counterparts. These high quality, methodical, and carefully designed experiments serve well in testing methods that are designed to predict drug mode of actions because their cellular effects have been well studied, spanning a variety of mechanisms from DNA-damaging agents (e.g. etoposide (Nakada et al., 2006)) or cellular motility inhibitors (e.g. blebbistatin (Allingham et al., 2005)) to compounds that disrupt regulatory signaling mechanisms (e.g geldanamycin (Neckers et al., 1999; Grenert et al., 1997)).

Differential gene expression analyses and other gene enrichment methods may provide insight into dysregulated genes or gene sets (e.g. biological pathways) resulting from a drug perturbation propagating through a system of interacting genes or proteins. However, identifying the primary source of perturbation that can explain the global variation in gene expression is often difficult to discern from differential gene expression alone. For instance, DNA damaging agents that induce cell cycle arrest initiate a series of biological processes such as cell death pathways (apoptosis), protein degradation pathways (e.g. RNA degradation, ubiquitin mediated proteolysis), and possibly DNA-repair pathways. As a result, genes associated with these downstream pathways may be upregulated and consequently, identified by differential gene analyses such as gene set enrichment analysis (GSEA, Subramanian et al., 2005). Rather than detecting the residual effects of such a perturbation, we aim to identify upstream pathways positioned to cause changes in gene expression. In fact, in the case of DNA damaging agents such as the drug camptothecin, we identified P53 signaling in the DREAM7 dataset while GSEA has not (see Section 6.4 for details); P53 signaling may be causally linked to cell cycle arrest induced by DNA damage (Jaks et al., 2001; Gupta et al., 1997; Wang et al., 2004).

Moreover, comparing drug profiles from two different exposure times, we show that certain drugs are more sensitive on the LY3 cancer cell line than others. We also identified drug-induced pathways that were consistently identified across varying conditions. Lastly, we found that drugs having similar mechanisms (e.g. DNA damaging agents) clustered together using profiles generated by our method.

1.3 Organization of this paper

In Section 2, we discuss related work. In Section 3, we describe the hierarchical model in detail including priors and model identification constraints. In Section 4, we outline the aims of posterior inference and the steps to our sampler. We assess the performance of our model compared to an exploratory factor analysis (EFA) model using simulated data sets in Section

5. We discuss our results after applying our method to a drug perturbation dataset (DREAM7, Bansal et al., 2014) and compare our method against an EFA model and gene set enrichment analysis (GSEA, Subramanian et al., 2005) in Section 6. Lastly, in Section 7, we summarize our method as well as our results.

2. PRIOR AND RELATED WORK

Currently, there are roughly three ways the task of identifying disease or drug targets have been approached. In the simplest approach, statistical methods are used to find differential expression between a binary phenotype (Dudoit et al., 2003; Storey and Tibshirani, 2003). The goal with this approach is to identify individual genes whose expression levels are associated with a trait relative to some control conditions. These methods are rooted in multiple testing and usually control either the family-wise error rate (Dudoit et al., 2003) or the false discovery rate (Dudoit et al., 2003; Storey and Tibshirani, 2003). Although relatively straightforward, these methods have two main drawbacks: they can only identify downstream transcriptional effects since the goal is to detect differential expression and not possible sources of perturbations; and they operate on a gene-by-gene basis and so the phenotypes for each gene are compared in isolation instead of in coordination across gene sets.

The second approach, however, uses biological information to group genes with the goal of capturing coordinated expression changes within a gene set that might not have been detected if the analysis were conducted in individual genes. The aim, therefore, is to look for statistical significance in differential expression across entire gene sets. These gene sets represent a system's level view of the cell such as cellular pathways or transcription factor targets. The most simple and straightforward method to test if a gene set contains a significant amount of transcriptionally altered genes is a hypergeometric test. This method and other similarly straightforward methods are reviewed in (Khatri and Draghici, 2005; Rivals et al., 2007). Subramanian et al. (2005) developed a method called gene set enrichment analysis (GSEA) with the same goal, but this method does not require a user to choose a threshold to separate differentially expressed genes from non-differentially expressed genes. Instead, GSEA measures how well a gene set clusters at the top or bottom of a *ranked* list of genes (e.g. ranked using fold ratios or p-values from two sample tests when comparing binary phenotypes). Extensions of GSEA and other threshold-free methods are described in (Jiang and Gentleman, 2007; Nam and Kim, 2008). GSEA is the most widely used gene set analysis tool to date, however it does not use any gene network information.

The third and final approach is network-based and uses information from cellular regulation or gene-protein interactions to provide a better understanding of the molecular mechanisms underlying a response. On an individual protein-gene level, many supervised methods have been proposed to predict interactions between drugs and proteins. For instance, using existing information of known protein-drug interactions, machine learning methods (Faulon et al., 2008; Yamanishi et al., 2008) have been used to predict interactions of individual proteins and drugs using a combination of protein sequence and chemical data. Network based approaches that identify gene sets over individual genes or proteins have also been

proposed. Gu et al. (2010), for instance, use a two step method where in the first step they identify a set of differentially expressed (DE) genes that are then mapped to a biological network; in the second step, a clustering method is used to identify genes that form a connected subgraph of DE genes. Other techniques, including pathway-express (Draghici et al., 2007) and gene network enrichment analysis (Liu et al., 2010) have also used network topology. Pathway-express (Draghici et al., 2007) scores genes in a network using differential expression such that differentially expressed genes that are connected are scored higher than those that are not connected. Pre-defined gene sets are then represented using the scores of its genes in the network. Gene network enrichment analysis (Liu et al., 2010) first identifies a connected subnetwork of differentially expressed genes of a global protein-protein interaction network, and then identifies pre-defined gene sets, such as biological pathways, that are enriched with respect to the genes in a differentially expressed subnetwork.

Although these increasingly complex approaches serve well in identifying genes or gene sets that are differentially expressed, they are not designed to identify which genes or gene sets, of the collection of genes or gene sets identified by their method, are primary targets of a specific response, and thus they do not prescribe a mechanism of action. To this end, it is essential to explicitly model the cascading effects of a perturbation by considering biological relationships in a gene network. Our proposed method is such a model-based approach: a factor model, called confirmatory factor analysis (CFA) with a conditional autoregressive (CAR) component designed to identify primary targets underlying gene expression. Spatial conditional autoregressive models have recently been explored in biology, specifically with functional networks (e.g. (Wei and Pan, 2008)). The motivation behind such models stems from the principle that genes that work together (i.e. genes that are functionally related) are usually co-expressed/inhibited, co-(de)regulated, or co-(de)activated.

Factor analysis models have also been used extensively in biology (Lucas et al., 2006; Carvalho et al., 2008; Lucas et al., 2010; Ma and Zhao, 2012). Ma and Zhao (2012) use a Bayesian factor analysis model to identify drug-pathway interactions by analyzing paired gene expression and drug sensitivity data and their relationship to common pathways, represented as latent factors. In their model all factors (pathways), however, are independent. These network-free models are referred to as exploratory factor analysis (EFA) models. West et al. (Lucas et al., 2006; Carvalho et al., 2008) also use exploratory factor analysis models but with an aim to identify biomarkers using gene expression data. However, in (Lucas et al., 2006; Carvalho et al., 2008) the factors are not structurally informed by known biological gene sets and instead are completely data-driven. Lastly, EFA or traditional CFA models assume a zero prior mean for the latent factors. Conversely, our proposed method is a novel approach to perturbation target identification on a latent scale which give rise to non-trivial prior means on the latent factors. Together, our proposed method is a synthesis of several key principles from both confirmatory factor models and conditional autoregressive models that allow us to explain the variation observed in transcriptional data as a function of perturbations to a latent network of biological pathways.

Our proposed model was first motivated by the work of di Bernardo et al. (2005), where a mathematical model adapted from Hill-type transcription kinetics (Liao et al., 2003) was

used to describe the relationships among gene transcripts in a cell. The model captures both internal regulatory influences among p gene transcripts and effects due to external perturbation effects:

$$\log_{10} \left(\frac{\nu_i}{\nu_{ib}} \right) = \log_{10} \left(\frac{\mu_i}{\mu_{ib}} \right) + \sum_{j \neq i}^p \eta_{ij} \log_{10} \left(\frac{\nu_j}{\nu_{jb}} \right), \quad (1)$$

where ν_i for $i = 1, \dots, p$, represents the expression level of transcript i , μ_i represents the direct influence of the perturbation on transcript i , ν_{ib} and μ_{ib} represent a set of baseline values of ν_i and μ_i , respectively, and η_{ij} represents the influence of transcript j on transcript i . This can be re-expressed as

$$\psi = B\psi + \phi \quad (2)$$

where $\psi = (\psi_1, \dots, \psi_p)'$ with $\psi_i = \log_{10}(\nu_i/\nu_{ib})$, and $\phi = (\phi_1, \dots, \phi_p)'$ with $\phi_i = [1/(1 - \eta_{ii})] \log_{10}(\mu_i/\mu_{ib})$, a scaled version of the log-relative direct influence of the perturbation. Here, B is a $p \times p$ matrix representing the network interaction effects among gene transcripts with $B_{ii} = 0$ and, for $i \neq j$, $B_{ij} = \eta_{ij}/(1 - \eta_{ii})$.

The work in (Cosgrove et al., 2008) is a statistical extension of the mathematical model in Eq. (2) in the form of

$$y = B\psi + \phi + e \quad (3)$$

where y is now the $p \times 1$ vector of measures of ψ . Here, each measured expression level y_i is potentially influenced by other transcripts, where this influence is captured through B . The perturbations are additive components represented by ϕ .

Here, we further extend this model to factor analysis models which are frequently used to analyze high-dimensional gene expression data (Ma and Zhao, 2012; Lucas et al., 2010; Carvalho et al., 2008; Lucas et al., 2006). Specifically, we assume that gene expression is linked to an overall factor defined by biological pathways in the cell. Our proposed model can be described as a hierarchical model where the first (gene) level is a regression of gene expression on biological pathways, and the second (pathway) level is an auto-regressive model that models the behavior of *a system of interacting biological pathways* underlying external perturbation effects. We describe the proposed model in Section 3.

3. MODEL FRAMEWORK

We developed a statistical hierarchical model to uncover perturbations to biological pathways using high-throughput measurements of gene expression on individual genes. The data $Y = \{Y_e, Y_c\}$ is a $p \times n$ matrix of transcriptional profiles of p genes across n varying

conditions, and consists of a set of cases, Y_c , and controls, Y_c . To avoid using an intercept in our model, we subtract the mean of the control group Y_c from the matrix Y . As a result of this normalization, the mean value of expression is zero if there is no difference between cases and controls. Our goal is to identify latent factors (e.g. biological pathways) that can explain the variation observed in the experimental data, Y_c , relative to the control conditions, Y_c .

Our model is built on three hierarchical components:

$$y_{ki}|\Lambda, \omega_i, \psi_k \stackrel{\text{ind}}{\sim} N(\Lambda_k \omega_i, \psi_k) \quad (4)$$

$$\omega_{ji}|\omega_{[-j]i}, \rho_{ji}, \sigma^2 \stackrel{\text{ind}}{\sim} N\left(\sum_{j'=1, j' \neq j}^q B_{jj'} \omega_{j'i} + \rho_{ji}, \sigma^2\right) \quad (5)$$

$$\rho_{ji}|\theta_{ji}, \tau^2 \stackrel{\text{ind}}{\sim} N(0, \theta_{ji} \tau^2 + (1 - \theta_{ji}) v_0 \tau^2) \quad (6)$$

for $i = 1, \dots, n$, $k = 1, \dots, p$, $j = 1, \dots, q$, where n , p , and q represent the number of samples, genes, and pathways, respectively, and

- y_{ki} is the observed gene expression level of gene k in sample i .
- ω_i is a $q \times 1$ vector of latent common factors representing inherent pathway effects for sample i .
- Λ is a $p \times q$ factor loadings matrix in the CFA model, relating gene expression to pathway effects. Λ_k is the k -th row of Λ .
- ρ_{ji} is a random “external” perturbation effect for pathway j in sample i .
- θ_{ji} is an indicator for pathway j in sample i being perturbed.

We describe and motivate the various aspects of this model in more detail in the following subsections.

3.1 Modelling Biological Replicates

It is common in biological studies to have replicates. In this case we allow perturbations and latent factors to vary for each replicate, yielding the model

$$y_{kir}|\Lambda, \omega_{ir}, \psi_k \stackrel{\text{ind}}{\sim} N(\Lambda_k \omega_{ir}, \psi_k) \quad (4')$$

$$\omega_{jir} | \omega_{[-j]ir}, \rho_{jir}, \sigma^2 \stackrel{\text{iid}}{\sim} N \left(\sum_{j'=1, j' \neq j}^q B_{jj'} \omega_{j'ir} + \rho_{jir}, \sigma^2 \right) \quad (5')$$

$$\rho_{jir} | \theta_{jir}, \tau^2 \stackrel{\text{iid}}{\sim} N(0, \theta_{ji} \tau^2 + (1 - \theta_{ji}) v_0 \tau^2) \quad (6')$$

for replicates $r = 1, \dots, m_i$ in each study. Note that while ρ and ω vary, we assume the same perturbation targets and gene-pathway loadings across replicates and so θ and Λ do not depend on index r . In the next sections we present each of these levels in detail, but we drop the replicate indices to simplify the notation.

3.2 Confirmatory Factor Level

In the first component, we focus on the likelihood of individual gene expression measurements. It has been shown that the transcriptional activity of a gene is, in part, regulated by pathways positioned upstream of gene regulation (e.g. (Vogelstein and Kinzler, 2004)). We describe this relationship between a gene and its pathways using a *confirmatory factor analysis* model, as shown in Eq. (4). Here, the gene expression measurements, Y , are regressed on a set of biologically structured latent factors ω (biological pathways).

Both the structure of the latent factors ω and the dependency between individual genes y and these latent factors are encoded, a priori, in the $p \times q$ loading factor matrix Λ . We integrate biological information into the model through Λ , by limiting some of the elements of Λ to zero, where the non-zero elements in each *column* of Λ correspond to the genes of a known canonical pathway. That is, a gene k loads onto a pathway j , (that is, $\lambda_{kj} \neq 0$) if gene k (or its gene product) is a member of pathway j . If a gene (or its gene product) is not in a pathway, the corresponding loading factor is constrained to 0. Thus genes are regressed only on the pathways to which they belong.

We use the following priors for the loading factors Λ :

$$\lambda_{kj} \stackrel{\text{iid}}{\sim} \begin{cases} N(0, 0.1), & \text{if gene } k \text{ is in pathway } j \\ \delta_0(\cdot), & \text{otherwise} \end{cases}$$

where δ_0 is the Dirac delta function and we set the variance of Λ to be small to prevent the latent factors from shrinking towards zero in practice, and keep the pathway noise variance weakly informative.

We give the gene variances Ψ inverse gamma priors,

$$\psi_k \stackrel{\text{iid}}{\sim} \text{IG}(\zeta, \zeta - 1), \quad k = 1, \dots, p,$$

where the shape is chosen so that the prior mean is 1. We now control the prior strength of the gene noise variance parameters Ψ through the hyper-parameter ζ to avoid overfitting the model. This prior regularization is important in cases where the data size is not large enough to fit complex hierarchical models and can lead to instability, poor convergence, multiple local modes, and empirically non-identified models when the prior on each ψ_k is too flexible. The extent to which an inverse gamma prior is informative depends on the strength of the data and is translated as a variance tradeoff between the prior and the data in the conditional posterior distribution of ψ_k :

$$\psi_k | \Omega, \Lambda, Y \sim \text{IG} \left(\zeta + \frac{n}{2}, \zeta - 1 + \frac{1}{2} \sum_{i=1}^n (Y_{ki} - \Lambda_k, \omega_i)^2 \right). \quad (7)$$

Since an inverse gamma distribution with shape α corresponds to a χ^2 distribution with 2α degrees of freedom (Gelman, 2006), the prior provides 2ζ pseudo-observations and the conditional in (7) has $2\zeta + n$ observations. We thus constrain the prior variances of Ψ by setting κ times data observations as prior pseudo-observations, $\zeta = \kappa n/2$, and so the conditional in (7) has $n(\kappa + 1)$ degrees of freedom. In this manner, we fix $\kappa < 1$ so that the prior does not overpower the likelihood. In our implementation, we fixed κ to 0.5 to have a reasonably informed prior.

3.3 Conditional Autoregressive Level

The second component of the model is designed to characterize a perturbed system of interacting pathways. The dependency between pathways is captured by a *conditional autoregressive model* (Cressie, 1993; de Oliveira, 2012) in Eq. (5). The reasoning is that, under normal stationary conditions, each pathway can be explained, on average, by other pathways in a pathway-pathway network through the coefficients B . However, when a pathway is targeted by some external perturbation source (e.g. from disease, or drug perturbation), then the expected mean of this pathway is shifted by an additional term ρ due to the perturbation. In the context of drug perturbations, one can effectively imagine that after a drug hits a specific pathway, this drug-induced effect propagates across the network, affecting many pathways by mere association, and altering cellular phenotype.

In the CAR level of the model, Eq. (5), the $q \times q$ design matrix B represents a system of pathway-pathway interactions. More specifically, let us initially set $B := \gamma W$, where γ is referred in the CAR literature as the *spatial scaling* parameter, and W is a symmetric zero-diagonal matrix. We briefly describe the pathway network W in Section 3.4 and provide a technical description in Appendix A.

If we define $\Phi := (I - \gamma W)^{-1}$, where I represents the identity matrix of order q , the joint distribution of ω can be written as the following (Cressie, 1993):

$$\omega | \rho, \Phi, \sigma^2 \sim N(\Phi \rho, \sigma^2 \Phi). \quad (8)$$

Thus, if $\gamma = 0$ the CAR model is reduced to an exploratory factor analysis model. However, to attain an identifiable model (please refer to Section 3.6 for details), we require that the factors have equal variance s^2 and so $\Phi = s^2 R(\gamma)$, where $R(\gamma)$ is the correlation induced by $G(\gamma) := (I - \gamma W)^{-1}$,

$$R(\gamma) = \text{Diag}(G(\gamma))^{-\frac{1}{2}} G(\gamma) \text{Diag}(G(\gamma))^{-\frac{1}{2}}. \quad (9)$$

In addition, to ensure that Φ is positive definite, we constrain γ to the interval $(1/\iota_1, 1/\iota_q)$, where ι_1 and ι_q are the minimum and maximum eigenvalues of W , respectively. We note that if W is not positive definite then $\iota_1 < 0$ and so γ can take negative values which, in turn, induce negative entries in $R(\gamma)$ (Wall, 2004). These negative correlations can represent a net inhibitory effect between genes in two pathways, an expected pattern if the pathways contain regulatory genes. The prior on γ is then uniform,

$$\gamma \sim \text{Unif}\left(\frac{1}{\iota_1} + \delta, \frac{1}{\iota_q} - \delta\right),$$

where $\delta > 0$ ensures that the maximum correlation in $R(\gamma)$ is not near 1, which can lead to a non-identified model. We set $\delta = 0.005$, which, in our experience, bounds the maximum correlation at roughly 0.90. Finally, we set a weakly informative prior on σ^2 , $\sigma^2 \sim \text{IG}(0.001, 0.001)$.

3.4 Interaction Network in the CAR Model

The pathway interaction network lies at the core of our model. We specifically use a network, where nodes are pathways, constructed in a manner that implicitly links pathways by their common function in the cell. To date there are several ways of building a pathway interaction network. The most ad-hoc approach is to simply define a link between two pathways where there is a non-empty intersection of gene members. However, this rule overlooks interactions that may occur *between non-overlapping genes* of two pathways. An alternative approach is to use a protein-protein interaction database to identify physical interactions *between* members of two pathways, and using an aggregate score to define the overall interaction link. However, using PPI interactions is often very noisy given the high variability in PPI data (von Mering et al., 2002; Reguly et al., 2006; Gandhi et al., 2006) resulting from methods that produce high-coverage of the proteome inherently having a high false positive rate.

To better capture interactions between pathways, we used Gene Ontology's (GO) biological processes (The Gene Ontology Consortium, 2000) to define functional links between pathways. The advantage of GO over a PPI database is that a GO gene set is manually curated, thus containing much fewer false positives. Functional networks have been used extensively in biology (Ideker et al., 2002; Chuang et al., 2007; Roguev et al., 2008), which

are generally motivated by the principle that genes that work together to accomplish a task in the cell are usually co-activated/inhibited or co-(de)regulated.

To construct the pathway-pathway interaction network W , we regard the set of pathways as a weighted network where the nodes represent canonical pathways and the edge weights in W reflect the degree of functional similarity between two pathways. As an example, in Fig. 1 the nodes would represent Pathways A, B, and C with a link connecting Pathways A to B and A to C. We then create W in a manner similar to an algorithm by Pham et al. (2011) using KEGG regulatory and signaling pathways (Kanehisa et al., 2006; Kanehisa and Goto, 2000), and GO biological processes (The Gene Ontology Consortium, 2000) obtained from the MSigDB collection (Subramanian et al., 2005). For a technical description of the network construction see Section A.

3.5 Spike-and-Slab Prior

Our goal of elucidating primary targets of a perturbation reduces to identifying perturbation effects after accounting for all pathway-pathway interactions (i.e. via network filtering) as described by a known biological network W . Importantly, we assume that relatively few pathways are primary targets of a perturbation. We use a spike-and-slab prior on ρ (George and McCulloch, 1997) and identify perturbations through posterior-based variable selection. Thus our parameters of interest are variables θ_{ji} that indicate whether ρ_{ji} represents a non-zero perturbation for pathway j in sample i , as in Eq. (6).

Because the scale of the perturbations is unknown, we choose a mixture of two normals (George and McCulloch, 1997), one approximating the spike and the other the slab, in such a way that is invariant to the scale of the perturbations. That is, we define the spike variance to be some fraction v_0 of the slab variance (Braunstein et al., 2011) and let the slab variance be random. We fixed $v_0 = 0.01$, meaning that a non-perturbed pathway has only one-hundredth of the variance of a perturbed pathway, while keeping the prior on τ^2 weakly informative, similarly to the prior on σ^2 , $\tau^2 \sim \text{IG}(0.001, 0.001)$.

Lastly, the prior probability of a pathway being perturbed is captured by a hyper-parameter α , $\theta_{ji} \stackrel{\text{iid}}{\sim} \text{Bern}(\alpha)$, for $j = 1, \dots, q$ and $i = 1, \dots, n$. We expect α to be small to reflect the targeted effect of drug perturbations to a few pathways because, for instance, the FDA requires perturbation selectivity for drug approval. Thus, to induce sparsity we set $\alpha = 0.1$ for the expectation of perturbed pathways. While we expect *a priori* that on average 10% of the pathways are selectively perturbed, we found in a prior sensitivity analysis that small changes in α do not significantly change our inference on perturbation detection.

3.6 Model Identification

Model identification has always been a non-trivial task with factor models (see Chapters 4, 7, 8 in (Bollen, 1989)). In CFA models, constraints are usually applied to elements of both Φ and Λ . This can be done, for example, by fixing the factor variances and the first q rows of Λ to a particular structure (Lucas et al., 2006). These methods prevent non-trivial rotational or scale transformations of Λ and Φ by forcing the transformation matrix U to be the

identity matrix. However, since in our model $\Phi := s^2 R$, where R is a correlation matrix (Eq. (9)), Φ is already rotationally unique.

This model, however, is still not identified in a few additional ways. Firstly, for the i -th sample we have $Y_i | \Lambda, \omega_i, \Psi \sim N(\Lambda \omega_i, \Psi)$ and so, if we marginalize ω_i we have the following conditional distribution on Y_i :

$$Y_i | \Lambda, \rho_i, \sigma^2, \gamma, \Psi \sim N(s^2 \Lambda R(\gamma) \rho_i, \Psi + \sigma^2 s^2 \Lambda R(\gamma) \Lambda^\top). \quad (10)$$

Thus, if we re-scale $\tilde{\sigma}^2 = k_\sigma \sigma^2$, $\tilde{s}^2 = k_s s^2$, $\tilde{\rho}_i = k_\rho \rho_i$ and $\tilde{\Lambda} = k_\Lambda \Lambda$ such that

$$k_s k_\Lambda k_\rho = 1 \quad \text{and} \quad k_\sigma k_s k_\Lambda^2 = 1$$

we obtain the same conditional distribution on Y_i since its mean and variance remain unaltered, respectively. For example, re-scaling Λ , σ^2 and ρ_i with any scalar $a \neq 0$ such that $\tilde{\Lambda} = a\Lambda$, $\tilde{\sigma}^2 = \sigma^2/a^2$, and $\tilde{\rho}_i = \rho_i/a$ does not affect the distribution of Y_i . To solve these forms of rescaling, we fix $s^2 = 1$ as well as the prior variance of Λ to 0.1, a small value that avoids shrinking the scale on ω towards 0 in practice. Moreover, we note that since the prior scale of Ψ is centered at 1 and the prior on σ^2 is weakly informative, σ^2 controls, in effect, how the variance of Y_i in (10) is partitioned into pathway effects via Λ and residual effects via Ψ .

Finally, there is an additional case that can lead to a non-identified model. This last case arises from the latent factor correlation matrix R . In the simplest case, suppose $\gamma = 0$ and W is a connected network (i.e. R can only be arranged as a single block matrix). Under these conditions, there exist exactly two solutions where

$$\tilde{\Lambda} = -\Lambda, \quad \tilde{\omega} = -\omega, \quad \tilde{\rho} = -\rho.$$

This is simply a sign flip across all pathways, which would not affect the underlying factor covariance structure Φ . In fact, if R is rearranged into a block diagonal matrix, then each set of pathways corresponding to a block in R can be flipped in this manner without affecting Θ and yielding the same joint posterior density. In other words, if R is rearranged into a block diagonal matrix with b blocks then there exist exactly 2^b solutions. Therefore, perturbations are identified up to their signs; that is, we can infer if two pathways are perturbed in the same direction, but if the perturbation is positive or negative is arbitrary. Moreover, since our goal is to infer perturbations, we are more interested in the magnitude of a perturbation—whether it is significantly close to zero or not—rather than its sign.

4. POSTERIOR INFERENCE

Our primary goal is the posterior inference on $\Theta = \{\theta_{ji}\}$, which is used to identify perturbation targets. We obtained posterior estimates to our model parameters via a partially collapsed hybrid Gibbs sampler (van Dyk and Park, 2008) with an adaptive Metropolis step

(Gelman et al., 2004). We call the sampler “partially collapsed” because we sample some of the parameters (namely, ρ and θ) from a *marginalized* conditional posterior. We found that using this partially collapsed sampler facilitates the sampler to move into regions of high posterior mass concentration faster.

First, the steps for an ordinary, non-collapsed Gibbs sampler (dropping iteration indices and irrelevant conditional parameters) are:

$$\begin{aligned} & [\Lambda|\omega, \psi, Y], & [\rho|\theta, \Phi, \omega, \sigma^2, \tau^2]^*, & [\theta|\rho, \tau^2]^*, \\ & [\gamma|\omega, \rho, \Phi, \sigma^2], & [\omega|\Lambda, \rho, \psi, \Phi, \sigma^2, Y], & [\psi|\Lambda, \omega, Y], \\ & [\sigma^2|\omega, \rho, \Phi], & [\tau^2|\rho, \theta]. \end{aligned}$$

Note that, even though $\Phi = s^2 R(\gamma)$ is a matrix that depends on γ , we use them interchangeably to make the conditional distributions clearer. Now, to improve the rate of convergence, we marginalize the two steps starred above: (i) instead of sampling from $[\rho | \theta, \Phi, \omega, \sigma^2, \tau^2]$, we integrate out ω from $[\rho, \omega | \Lambda, \theta, \Phi, \psi, \sigma^2, \tau^2, Y]$; and (ii) instead of sampling from $[\theta | \rho, \tau^2]$, we marginalize ρ from $[\theta, \rho | \omega, \Phi, \tau^2, \sigma^2]$. The collapsed sampler has then updated steps

$$[\rho|\Lambda, \theta, \Phi, \psi, \sigma^2, \tau^2, Y] \quad \text{and} \quad [\theta|\omega, \Phi, \tau^2, \sigma^2].$$

We note that, as van Dyk and Park (2008) pointed out, marginalizing does not alter the stationary distribution of the full posterior nor the compatibility of the conditional distributions. In the next sections we provide details on each of these steps.

4.1 Initializing MCMC chains

We implemented a tempering algorithm at the beginning of our MCMC chains to find reasonable starting points such that the chains were less likely to get stuck at a local mode. These local modes are caused, in part, by each likelihood in the CFA level, Eq. (4), being close to non-identifiable up to the mean $\Lambda\omega$. To do this, we run our sampler as usual, but for Ψ , Λ , and Ω , we temper the likelihood distributions $Y|\Psi$, Λ , Ω when obtaining their conditional posteriors. Similarly, we do the same for the likelihood of $Y|\rho$ when sampling ρ . At hot temperatures, the likelihood of Y would be flatter, allowing the chains to move more freely around the space. We slowly cool the temperature to $T=1$ to obtain the targeted posterior and thus beginning the true sampler.

4.2 Sampling Λ

Recall that we have constrained some elements in Λ to 0 such that genes are regressed only on pathways to which they belong. That is, if gene k is not in pathway j , then we constrain the element λ_{kj} to 0. Consider the k -th row of Λ , which we denote as Λ_k , and suppose some of the elements in Λ_k are constrained to zero. Let c_k be the corresponding $1 \times q$ row vector such that

$$c_{kj} := I(\text{gene } k \notin \text{pathway } j),$$

that is, $c_{kj} = 0$ if λ_{kj} is a structural zero and so $\lambda_{kj} \stackrel{\text{iid}}{\sim} N(0, (1 - c_{kj})0.1)$. If $r_k = \sum_j c_{kj}$ is the number of pathways containing gene k and Λ_k^* is the $1 \times r_k$ row vector that contains the unknown parameters in Λ_k then the prior on Λ_k^* is $N(0, H_k)$ where $H_k = 0.1I$, I the identity matrix of order r_k .

Similarly, let Ω_k^* be the $r_k \times n$ sub-matrix of Ω such that for $j = 1, \dots, r_k$, all the rows corresponding to $c_{kj} = 0$ are deleted and also let Y_k be the $1 \times n$ vector of observations for gene k . Then, from a well known result of Lindley and Smith (1972), the posterior conditional distribution for Λ_k^* , $k = 1, \dots, p$, is

$$\Lambda_k^* | Y_k, \Omega_k^*, \psi_k \sim N(A_k^{-1} a_k, A_k^{-1})$$

where $A_k = \psi_k^{-1} \Omega_k^* \Omega_k^{*\top} + H_k^{-1}$ and $a_k = \psi_k^{-1} \Omega_k^* Y_k^\top$.

4.3 Sampling ρ

We define $\Sigma(\theta_j)$ as a diagonal matrix whose j -th diagonal element is $\Sigma(\theta_j)_{jj} = \tau^2[\theta_j + v_0(1 - \theta_j)]$. Then, again exploiting the result of Lindley and Smith (1972) on the marginalized conditional distribution of Y_j after integrating out ω_j in (10), we sample ρ_j from the following conditional posterior distribution:

$$\rho_i | \theta_i, \Lambda, \Psi, Y_i \sim N(A_i^{-1} a_i, A_i^{-1}),$$

where

$$\begin{aligned} A_i &= \Phi^\top \Lambda^\top (\sigma^2 \Lambda \Phi \Lambda^\top + \Psi)^{-1} \Lambda \Phi + \Sigma(\theta_i)^{-1}, \\ a_i &= \Phi^\top \Lambda^\top (\sigma^2 \Lambda \Phi \Lambda^\top + \Psi)^{-1} Y_i. \end{aligned}$$

4.4 Sampling θ

As in the previous section, we compute the marginalized conditional posterior of θ after integrating out ρ . For computational efficiency, we store the marginal variance $V(\theta_i^{(t)})$ of ω_i conditional on θ_j at each iteration t of our sampler, where index i runs over samples. That is,

$$V(\theta_i) = \sigma^2 \Phi + \Phi \sum (\theta_i) \Phi^\top.$$

To simplify the notation, we drop the sample index i on all parameters including θ .

At this step we sample iteratively from $\theta_j | \theta_{[-j]}, \omega, \Phi, \tau^2$ for $j = 1, \dots, q$. There are then two cases: when $\theta_j^{(t)} = 0$ and when $\theta_j^{(t)} = 1$, for which we define $V_0 = V(\theta_j^{(t)} = 0, \theta_{[-j]}^{(t)})$ and $V_1 = V(\theta_j^{(t)} = 1, \theta_{[-j]}^{(t)})$. Let ϕ_j denote the j^{th} column of Φ . If $\theta_j^{(t)} = 0$, then we define $\delta_{j0} = 1 + (\tau^2 - v_0 \tau^2) \phi_j^\top V_0^{-1} \phi_j$ and $\Delta_{j0} = (\tau^2 - v_0 \tau^2) / \delta_{j0} (V_0^{-1} \phi_j) (V_0^{-1} \phi_j)^\top$ to obtain:

$$\text{logit } \mathbb{P}(\theta_j^{(t+1)} = 1 | \theta_{[-j]}, \omega, \Phi) = -\frac{1}{2} \log \delta_{j0} + \frac{1}{2} \omega^\top \Delta_{j0} \omega + \text{logit}(\alpha).$$

If $\theta_j^{(t)} = 1$ we define $\delta_{j1} = 1 - (\tau^2 - v_0 \tau^2) \phi_j^\top V_1^{-1} \phi_j$ and $\Delta_{j1} = (\tau^2 - v_0 \tau^2) / \delta_{j1} (V_1^{-1} \phi_j) (V_1^{-1} \phi_j)^\top$ in a similar way to obtain:

$$\text{logit } \mathbb{P}(\theta_j^{(t+1)} = 1 | \theta_{[-j]}, \omega, \Phi) = \frac{1}{2} \log \delta_{j1} + \frac{1}{2} \omega^\top \Delta_{j1} \omega + \text{logit}(\alpha).$$

For the full derivation of these posterior probabilities, please refer to Appendix B.

4.5 Sampling γ

As with most spatial models, the scaling parameter γ lacks conjugacy. We used a random walk adaptive Metropolis algorithm such that the proposal density is normal and centered at the current sample, $\gamma^{(\ell)}$, with variance ξ^2 . This variance is tuned to adjust for the acceptance rate and fixed post burn-in.

We propose $\gamma^* \sim \mathcal{N}(\gamma^{(\ell)}, \xi^2)$. Let $\Phi^* = s^2 R(\gamma^*)$ with the correlation computed as in Eq. (9). Then, if

$$l(\Phi) := -\frac{n}{2} \log |\Phi| - \frac{1}{2} \sum_{i=1}^n (\omega_i - \Phi \rho_i)^\top (\sigma^2 \Phi)^{-1} (\omega_i - \Phi \rho_i),$$

the acceptance ratio of the Metropolis step is just $r(\gamma^*, \gamma^{(\ell)}) = \exp\{l(\Phi^*) - l(\Phi^{(\ell)})\}$ since we have a flat prior on γ .

4.6 Sampling ω, ψ, σ^2 , and τ^2

We sampled ω according to the linear conditional posterior:

$$\omega_i | Y_i, \Lambda, \Psi, \rho_i, \gamma \stackrel{\text{ind}}{\sim} N(A_i^{-1} a_i, A_i^{-1})$$

where

$$\begin{aligned} A_i &= \Lambda^\top \Psi^{-1} \Lambda + \sigma^{-2} \Phi^{-1} \\ a_i &= \Lambda^\top \Psi^{-1} Y_i + \sigma^{-2} \rho_i. \end{aligned}$$

The conditional distributions of variance parameters ψ and σ^2 follow from conjugacy:

$$\begin{aligned}\psi_k|\omega, \Lambda, Y &\sim \text{IG}\left(\zeta + \frac{n}{2}, \zeta - 1 + \frac{1}{2} \sum_{i=1}^n (Y_{ki} - \Lambda_k, \omega_i)^2\right) \\ \sigma^2|\omega, \Phi, \rho, &\sim \text{IG}\left(0.001 + \frac{qn}{2}, 0.001 + \frac{1}{2} \sum_{i=1}^n (\omega_i - \Phi \rho_i)^\top \Phi^{-1} (\omega_i - \Phi \rho_i)\right).\end{aligned}$$

where $\zeta = n/4$ according to the discussion in Section 3.2.

The conditional distribution of τ^2 is also conjugate; if $v_{ij} = 1$ when $\theta_{ij} = 1$ and $v_{ij} = v_0$ when $\theta_{ij} = 0$, then

$$\tau^2|\rho, \theta \sim \text{IG}\left(0.001 + \frac{qn}{2}, 0.001 + \sum_{i=1}^q \sum_{j=1}^n \frac{\rho_{ij}^2}{2v_{ij}}\right).$$

4.7 Posterior Inference

We infer perturbations based on the centroid estimator of θ for some threshold t (Carvalho and Lawrence, 2008):

$$\hat{\theta}_{ji}(t) := I(\mathbb{P}(\theta_{ji}|Y) > t). \quad (11)$$

We choose the threshold t by controlling a Bayesian false discovery rate. We define the Bayesian false discovery rate (BFDR) of an estimator as

$$\text{BFDR}(\hat{\theta}(t)) := \mathbb{E}_{\theta|Y} \left[\frac{\sum_{i=1}^n \sum_{j=1}^q \hat{\theta}_{ji}(1 - \theta_{ji})}{\sum_{i=1}^n \sum_{j=1}^q \hat{\theta}_{ji}} \right] = \frac{\sum_{i=1}^n \sum_{j=1}^q \hat{\theta}_{ji}(1 - \mathbb{P}(\theta_{ji} = 1|Y))}{\sum_{i=1}^n \sum_{j=1}^q \hat{\theta}_{ji}}. \quad (12)$$

5. SIMULATION STUDY

We conducted a simulation study to (i) assess the impact of including biological network information into a factor model, and (ii) test the robustness of our model to inaccuracies in the biological databases used to construct pathways and the pathway-pathway interaction network.

5.1 Assessing the impact of incorporating biological network information

We tested the CFA-CAR model under various signal-to-noise (SNR) ratios and compared its performance with an exploratory factor analysis (EFA) model. This comparison model can be obtained from our original CFA-CAR model by letting the pathway-pathway network be empty. That is, an EFA model is a factor model such that the factor covariance matrix, Φ , is diagonal, which is equivalent to setting $\gamma = 0$ in the CFA-CAR model. In effect, this

constraint removes the network filtering effect of the CAR model. The EFA model can thus be written as (Eq. 13):

$$y_{ki}|\Lambda, \omega_i, \psi_k \stackrel{\text{ind}}{\sim} N(\Lambda_k \omega_i, \psi_k) \quad (13)$$

$$\omega_{ji}|\omega_{[-j]i}, \rho_{ji}, \sigma^2 \stackrel{\text{ind}}{\sim} N(\rho_{ji}, \sigma^2) \quad (14)$$

$$\rho_{ji}|\theta_{ji}, \tau^2 \stackrel{\text{ind}}{\sim} N(0, \theta_{ji}\tau^2 + (1-\theta_{ji})v_0\tau^2) \quad (15)$$

We simulated toy-scale networks with graph densities similar to the original KEGG pathway network. For each network, we randomly selected $q = 10$ pathways from our set of KEGG pathways. In this paper, we describe the results obtained from one of these networks. We obtained the pathway-pathway weighted networks by taking the subnetwork from our original pathway-pathway network induced by the ten randomly selected pathways. The number of distinct genes in the union of the selected pathways was 878. We used the real gene to pathway membership to create the mask of Λ . We sampled the non-zero loading factors in Λ from a Gaussian with zero mean and variance 0.1. The spatial scaling parameter γ was relatively high to emphasize non-trivial links in the network W . Each analysis was run using R on a Linux 8-core 3.6GHz CPU machine with 8Gb of memory and took, on average, 55 minutes for 8000 iterations.

We define the SNR in each simulated data set using the SNR of a perturbed pathway (which is the same for any perturbed case). If pathway k was perturbed with effect ρ , we define $\text{SNR} = \rho/\sigma$, and if we further fix the pathway noise variances (σ) to 1, then the $\text{SNR} = \rho$. In each case sample, we simulated an experiment where only a single pathway was perturbed, with 5 “biological” replicates per experiment (perturbation). Therefore, if we perturbed pathway k in sample $i \in \{1, 2, 3, 4, 5\}$ in experiment j , we set $\rho_{ki}^j = m$, where m is the signal to noise ratio in the data set, and $\rho_{k'i}^j = 0$ for $k' \neq k$. Across all data sets, we fixed the gene variances (Ψ) to 1.

For this network, we created six data sets of different signal to noise ratio in ten simulation replicates. In each data set we fixed SNR to the following values: 0.50, 1.50, 2.50, 3.50, 4.50, and 5.50. These values were chosen based on the distribution of the SNRs in the NCI-DREAM drug sensitivity data sets (see Section 6 and Fig. 7). For each of these data sets, we sampled 50 cases and 50 controls.

When running either models, we conditioned the perturbations of all replicates, ρ_i^j , where j indexes the experiment and i indexes the experimental replicate, on the same indicator vector θ_j . For the controls, we fixed the corresponding indicators to zero. We ran both our CFA model and the EFA model on each data set. Fig. 2 is a series of ROC curves across data sets

of different signal to noise ratios. Both models performed similarly under low SNRs. However, the curves begin to separate as the signal in the data sets increased. This is further emphasized by the AUC boxplots in Fig. 3. The CFA-CAR model consistently exhibits a low false positive rate across varying signal to noise ratios. Unlike the CFA model, the false positive rate of the EFA model increases as the signal to noise ratios increase. When $\text{SNR} = 4.5$ the AUC of the EFA model begins to decrease dramatically. This was expected because given a strong enough perturbation, the EFA model confuses a perturbation with the residual downstream effects of a perturbation. That is, pathways associated with the perturbed pathway are wrongly identified by the EFA model.

5.2 Assessing robustness to inaccuracies in the pathway database

We performed sensitivity analysis on the CFA-CAR model to test how possible inaccuracies in KEGG can affect the performance of CFA-CAR relative to an EFA model. To this end, we define the mask of Λ , $\Lambda^{(\text{mask})}$ as the $p \times q$ binary matrix (for p genes and q pathways), where $\lambda_{jk}^{(\text{mask})} = 1$ when gene j is in pathway k and $\lambda_{jk}^{(\text{mask})} = 0$ otherwise. With misannotated genes, the inaccuracies encoded in $\Lambda^{(\text{mask})}$ would effect both the CFA-CAR and EFA models. However, inaccuracies in the pathway network, W , would affect only the CFA-CAR model.

In this simulation, we tested six levels of gene set perturbations in the set of pathways. For a fixed SNR of 3.5, we chose 10 real pathways. Next, we simulated 10 data sets using the real network induced by these 10 pathways with a fixed SNR of 3.5 (same approach to the data simulations in Section 5.1).

To assess how LFIA performs under different degrees of database inaccuracy, we ran the LFIA model using a different weighted pathway-pathway network W and a different mask for Λ , where we randomly perturbed the assignment of genes to pathways for these false networks. We reassigned x percent of randomly selected genes to randomly selected pathways. Then, we re-created the pathway-pathway network W from the new Pathway-GO bipartite graph. Similarly, the mask of Λ is generated from the new gene to pathway assignment. For each percentage of genes tested, we generated a new W network and $\Lambda^{(\text{mask})}$ for each simulated data set.

We perturbed $x = 1\%, 2\%, 4\%, 8\%, 16\%, 32\%$ and 64% of all genes in our set of ten pathways. For each of these perturbations, we ran ten replicates where for each replicate we simulated a new “pathway” database and ran CFA-CAR on the original data set, using the newly constructed $\Lambda^{(\text{mask})}$ and W .

To assess the changes in W and $\Lambda^{(\text{mask})}$ relative to the true W and true $\Lambda^{(\text{mask})}$, we computed the Frobenius norm of the difference between the real and false networks scaled by the Frobenius norm of the real network (see Figure 4.) The changes increase almost linearly with the percentage of perturbed genes in a logarithmic scale, doubling at each iteration.

Fig. 5 is a series of ROC curves and Fig. 6 are AUC boxplots corresponding to the curves in Fig. 5 for both the CFA-CAR and EFA model across various levels of gene to pathway reassignments. We begin to see a small departure from the original ROC curve where the network and loading factor matrices were not randomized when 8% of the genes were

randomly reassigned to different pathways. At this level, the specificity of the CFA-CAR model begins to decrease. However, even at 32% randomness, the specificity and sensitivity of the model is still significantly higher than what you expect to see by random chance alone. We conclude that the CFA-CAR model is rather robust to possible inaccuracies in the biological databases, performing as well or better than an EFA version of this model.

6. CASE STUDY: DRUG TARGET PREDICTION IN DREAM 7 DRUG SENSITIVITY CHALLENGE DATA

We applied the CFA-CAR model on the NCI-DREAM drug sensitivity challenge data set. As a comparison, we also implemented an EFA model (see Section 5 for a description of the EFA model) and gene set enrichment analysis (GSEA) (Subramanian et al., 2005). The core algorithm of GSEA is a threshold-free variant of a hypergeometric test; however, unlike the CFA-CAR model, network interactions are not used.

In the NCI-DREAM drug sensitivity data set, 14 compounds (see Table 1) were tested at various concentrations and exposure times on the LY3 (Lymphoma) cancer cell line. We focused our attention on two data sets of different drug exposure times. The first data set consisted of drug perturbations at a concentration of (IC20) at 12 hours. The second data set consisted of drug perturbations at the same concentration but exposed for a duration of 24 hours. The IC20 concentration value is a measure of how much a drug is needed to inhibit a given biological process by 20%. These data sets were compared to mock control data sets where no drugs were exposed to the cells for 12 or 24 hour durations. There were eight mock replicates and each distinct experiment included three replicates, yielding a total of 50 samples per data set (Data Set 1: IC20 at 24 hours vs Mock at 24 hours and Data Set 2: IC20 at 12 hours vs Mock at 12 hours). Each analysis was run using R on a Linux 8-core 3.6GHz CPU machine with 8Gb of memory and took, on average, 8 hours for 8000 iterations.

When running CFA-CAR and EFA models, we conditioned the perturbations of all replicates, $\rho_1^j, \rho_2^j, \rho_3^j$, where j indexes the experiment, on the same indicator vector θ_j . In the control cases, we fixed the corresponding indicators θ to 0. For each dataset, we ran two chains of 4000 iterations each, with a burn-in of 2000 samples. We assessed convergence using the Gelman-Rubin statistic on the absolute values of all the continuous parameters. Using the same hardware and software configuration as the simulation studies in Section 5.1, these iterations took roughly 230 minutes to run.

Fig. 7 summarizes the posterior distributions of the model parameters γ , σ^2 , τ^2 , and the signal to noise ratios (SNRs) across the different drugs for each data set (IC20 at 12 hours and IC20 at 24 hours). We define the SNR of a pathway j in a sample i as the mean ratio, across replicates, of the absolute perturbation ρ to either σ , if the pathway is perturbed, or $\sigma \sqrt{v_0}$ otherwise; that is $\text{SNR}_{ji} = \sum_{r=1}^3 (\theta_{ji} |\rho_{jir}| / \sigma + (1 - \theta_{ji}) |\rho_{jir}| / (\sigma \sqrt{v_0})) / 3$, where r indexes the replicate of a drug experiment (with 3 replicates per drug). To assess SNR in a drug case, we pool the signal to noise ratio across all pathways.

For GSEA, we called a pathway “dysregulated” if the false discovery rate (FDR) adjusted p-value is less than 5%. For the CFA-CAR and EFA models we identify perturbations based on the centroid estimator of θ in Eq. (11) where the threshold t was chosen to control BFDR at 5% (see Eq. (12)). Restricting the BFDR at 5%, we used a perturbation threshold of 0.72 for IC20 at 12 hours and 0.50 for IC20 at 24 hours for the CFA-CAR model. For the EFA model, we used thresholds 0.80 and 0.70 for IC20 at 12 hours and IC20 at 24 hours, respectively.

6.1 Assessing model fit

We assessed model fit by first standardizing the data with respect to its marginalized likelihood given Θ . If, for sample i ,

$$Y_i | \Lambda, \Phi, \theta_i, \sigma^2, \Psi \sim N(0, \Lambda \Phi \sum (\theta_i) \Phi^\top \Lambda^\top + \sigma^2 \Lambda \Phi \Lambda^\top + \Psi),$$

then, if C_i denotes the Cholesky factor of the above marginalized covariance of Y_i , we have that $Y_i \sim N(0, C_i C_i^\top)$ and so $\tilde{Y}_i := C_i^{-1} Y_i \sim N(0, I_p)$, where I_p is the p -th order identity matrix. Thus, to assess model fit we take posterior mean estimates $\hat{C}_i$ of C_i , consider standardized gene expressions $\tilde{Y}_i = \hat{C}_i^{-1} Y_i$ for each sample $i = 1, \dots, n$, and then check two assumptions: zero-mean gene expressions and normality.

Figure 8 summarizes our assessment, with samples from IC20 at 12 and 24 hours in top and bottom panels, respectively. In the first sub-figure in the left, we plot pooled standardized gene expressions over both samples and genes to assess the zero-mean assumption. As can be seen from the smoother fit, this assumption is well met, but samples might have slightly different variance scales. For this reason, since we are focused on gene expressions, we assess normality via average standardized gene expressions over samples in the two last sub-figures. We found that the IC20 at 24 hours data has larger departures from normality in the tails when compared to the IC20 at 12 hours dataset. This departure at IC20 at 24 hours is specifically due to two of the drugs, H-7 Dihydrochloride (DHCL) and Mitomycin C (MMC), which showed a drastic change in variability between the IC20 at 12 hours and IC20 at 24 hours. This effect is explained in further detail in Section 6.5.

6.2 Assessing control of false positives via cross-validation

To verify the accuracy of our method, we tested our model against the control (mock) data by leaving three mock samples (microarray), out of the control group and treating them as replicates of an experiment, where the perturbations are left unknown. Because each drug experiment consisted of three replicates, we chose to treat three control samples as replicates of a single condition to maintain the same power. We applied this leave-three-out cross-validation 56 times to each data set to account for every possible combination of the 8 control samples. The aim here is to verify that the perturbed pathways identified in the case samples are not similarly identified against control samples.

We assessed pathway perturbations in the “case” group using the estimated posterior mean probabilities of a pathway being perturbed. We conducted this leave-three-out cross-validation procedure to both CFA and EFA models. We found that in each leave-three-out cross-validation test, in either model, we did not identify any pathways as being perturbed in the mock sample (Fig. 9).

6.3 Comparing EFA and CFA-CAR model results

Overall, we found that for both data sets the EFA model was less sensitive than the CFA model, finding a subset of mechanistically relevant pathways that were identified by the CFA model. Because the EFA model identified fewer pathways linked to drug mechanisms, we choose to focus our comparison with GSEA in the following subsections. We note, however, that the smaller number of relevant pathways found by the EFA model relative to the CFA model is due to the BFDR control forcing more stringent thresholds on the posterior means of θ for the EFA model. Results obtained from the EFA model can be found in Appendix C.

6.4 IC20 concentration at 12 hours

For the CFA-CAR model results, we ran a hierarchical bi-clustering analysis on the posterior means $\hat{\Theta}$ of θ . More specifically, we independently cluster samples (drugs) and pathways using complete linkage and Euclidean distances as dissimilarity metrics. Each dimension has their observations swapped to match the clustering hierarchy, as can be seen in Fig 10A. Fig. 10B shows the results from a similar analysis on adjusted p-values obtained from GSEA for comparison.

At IC20 at 12 hours, the most commonly identified pathway across drugs by the CFA-CAR model was P53-signaling: it was identified for etoposide (ETP), mitomycin-C (MMC), camptothecin (CPT), and doxorubicin (DOX). Interestingly, all four of these compounds are DNA damaging agents and were clustered into two groups (ETP/MMC and CPT/DOX) (Fig. 10), whereas some of the compounds that affect protein function unrelated to DNA damage, such as geldanamycin (GA) and rapamycin (RPM), also formed their own cluster. P53 is an important tumor suppressor that can regulate various cellular processes including apoptosis, cell cycle, and DNA repair; furthermore, it is mutated in over 50% of human cancers (reviewed in Ling and Wei-Guo, 2006). This loss of P53-signaling can lead to cancer growth and propagation. In many cases, P53-signaling is induced as a response to specific DNA damaging agents (Gupta et al., 1997), playing key roles in cell cycle arrest and apoptosis. All four drugs have been shown to play significant roles in activating P53-dependent mechanisms [ETP:(Karpnich et al., 2002; Grandela et al., 2007), MMC:(Abbas et al., 2002; Fritsche et al., 1993; Verweij and Pinedo, 1990), CPT:(Jaks et al., 2001; Gupta et al., 1997; Wang et al., 2004), DOX:(Kurz et al., 2004; Ling et al., 1996; Zhou et al., 2002)]. Furthermore, for some of these compounds including CPT and DOX, cell cycle arrest at specific check points occurs through the induction of a P53 signaling mechanism. For instance, P53 signaling was shown to play a critical role in the G1 checkpoint of the cell cycle under CPT-induced DNA damage (Jaks et al., 2001; Gupta et al., 1997). DOX is another chemotherapeutic DNA-damaging drug that causes cell cycle arrest in the G2/M

phase (Ling et al., 1996). Furthermore, DOX was shown to cause an accumulation of P53 that lead to a depletion of cells in the G2/M phase and apoptosis (Zhou et al., 2002).

GSEA did not identify P53-signaling or cell cycle for any of the 14 compounds. Moreover, the connection between the set of pathways identified by GSEA and some of the drug's primary MoAs were not readily apparent. Both GSEA and the CFA-CAR model identified DNA repair pathways for Rapamycin (RPM), with GSEA identifying mismatch repair and CFA-CAR identifying base excision repair, although the connection between these repair pathways and RPM's MoA is unclear.

6.5 IC20 concentration at 24 hours

At longer exposure times, both methods picked up more pathways (see Fig. 11 for heatmaps of the CFA-CAR and GSEA results, respectively.) This is expected from a biological point of view because as time of exposure increases the cell needs to respond to potentially more complex effects caused by the loss of cell homeostasis. Therefore, cells activate more pathways that, for example, help in cell survival, increase energy output, or signal cell demise or death. With the CFA-CAR model, we found that at IC20 concentrations at 24 hours of exposure, H-7 Dihydrochloride (DHCL) and Mitomycin C (MMC) were overly perturbed (i.e. many pathways had a significantly high posterior mean probability of having a non-zero perturbation). Base excision repair was again identified for RPM but was also picked by several other drugs such as monastrol (MNS), cycloheximide (CHX), trichostatin-A (TSA), DOX, and ETP. As mentioned earlier, the connection between DNA repair pathways and these compounds is unclear. Moreover, GSEA identified various DNA damaging repair or DNA related pathways at an FDR level of 5%. Such pathways include base excision repair, homologous recombination, mismatch repair, and DNA replication.

P53-signaling remains one of the most perturbed pathways among several of DNA-damaging or inhibiting compounds at the longer exposure time, including ETP, CPT, and DOX, all of which clustered together. P53-signaling was also picked up for methotrexate (MTX) which was something that was not picked up at the shorter exposure time but is very much linked to P53 induction (Krause et al., 2002; Huang et al., 2011). Krause et al. (2002) showed that treating HepG2 cells with methotrexate increases levels of p53. Huang et al. (2011) also demonstrated that methotrexate can induce apoptosis by the induction of the P53 targeted genes, DR5, P21, Puma and Noxa. Similarly, GSEA also identified P53 signaling for CHX, an inhibitor of protein biosynthesis at the longer exposure time, but was missed by CFA-CAR. Cell cycle was picked up for DOX and CPT (as in the shorter exposure time), but under these stronger experimental conditions, cell cycle was further identified for ETP and RPM as well.

7. DISCUSSION

Identifying mechanisms of action from transcriptional data alone is challenging. In disease phenotypes, differences in gene expression cannot, by themselves, elucidate aberrant causal pathways. Drug perturbations are even more difficult because often times perturbations occur directly on a proteomic level rather than a gene level. In many cases, transcriptional

responses are the residual aftermath of a perturbation rather than being representative of the direct target.

GSEA and other gene set methods attempt to boost the signal on transcriptional data by focusing on entire gene sets with the assumption that related genes may contribute a greater signal than individual genes alone. Furthermore, network-based methods use cellular regulation or gene/protein interaction data to better understand the MoAs underlying a transcriptional response.

Here, we link transcriptional data to unobserved mechanisms of action using confirmatory factor analysis in conjunction with a conditional autoregressive model. In the CFA level of the model we link gene expression to the biological pathways that contain their gene products. In the CAR model, we seek to explain the variation observed in these latent pathways by a pathway-pathway network induced by “external” perturbations.

The core goal of this model is to uncover these perturbation targets. We have shown through simulation and real data that providing a network filtering method on expression data is a significant improvement over non-network approaches. In simulation, we show that the CFA-CAR model exhibits a high specificity across various signal to noise ratios, unlike an EFA model (network-free) which has a higher false positive rate as the signal in the data set increases. Furthermore, we have shown that the CFA-CAR model is very robust to inaccuracies in the databases used to create the pathways and the pathway-pathway interaction network.

Moreover, our method was able to identify pathways significantly related to some of the drugs’ MoAs in the NCI DREAM 7 data sets. Specifically, we were able to identify signaling/regulatory pathways that play a causative role for many of the DNA-damaging compounds. These pathways were generally not identified using gene set enrichment analysis, or less so with EFA. Some of the known targets were missed by both the CFA-CAR model and GSEA across both data sets. For instance, rapamycin targets the mTOR protein (Alqurashi et al., 2013) but mTOR signaling was not identified. In another example, blebbistatin inhibits the myosin II protein (Allingham et al., 2005), and we expected to see some of the cellular motility pathways such as regulation of cytoskeleton pathway. Moreover, some of the results we found with both methods were also unclear (e.g. identifying DNA repair pathways across many of these compounds).

While our model arguably yields better results than other approaches in Section 6 and the simulation study in Section 5 shows that these results are fairly robust to gene-pathway network annotation inaccuracies, we recognize that other, possibly more severe, types of misspecification can occur and impact reproducibility. Unfortunately, it is out of the scope of this work to perform a thorough investigation of possible sources of model misspecification, but we facilitate such future endeavor by providing our software as supplemental material.

There are several natural directions for extension of this model. For instance, other forms of biological variables such as transcription factors or metabolic pathways could be used as latent factors. In our more recent work, we have been focusing on transcription factors (TFs) as latent factors, and we have created a network connecting two TFs based on the set of

shared targets. In this application, two TFs are linked if they jointly regulate a non-empty intersection of genes, that is, the biological function used here is co-regulation.

Another extension of this model would be to integrate other forms of high-throughput data in conjunction with gene expression. A temporal component could also be incorporated to accommodate temporal data and define a dynamic model. All these extensions, however, are at the cost of increased model complexity, which can greatly affect the fit and convergence of the model.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This research was supported, in part, by NIH grant GM078987. EK is partially supported by AFOSR grant 12RSL042. LC was supported by NSF grant DMS-1107067. LP is supported by NIH/NIDDK grant R01 DK078616.

References

- Abbas T, Olivier M, Lopez J, Houser S, Xiao G, Suresh-Kumar G, Tomasz M, Bargonetti J. Differential activation of p53 by the various adducts of mitomycin C. *Biol Chem.* 2002; 277:40513–40519.
- Allingham JS, Smith R, Rayment I. The structural basis of blebbistatin inhibition and specificity for myosin II. *Nature Structural & Molecular Biology.* 2005; 12:378–379.
- Alqurashi N, Hashimi SM, Wei MQ. Chemical Inhibitors and microRNAs (miRNA) Targeting the Mammalian Target of Rapamycin (mTOR) Pathway: Potential for Novel Anticancer Therapeutics. *Int J Mol Sci.* 2013; 14:3874–3900. [PubMed: 23434669]
- Bansal M, et al. 2014 “(To be determined),” (*To be determined*), to be submitted.
- Bollen, KA. Structural Equations with Latent Variables. New Jersey: Wiley-Interscience; 1989.
- Braunstein A, McShane B, Piette J, Jensen S. A Hierarchical Bayesian Variable Selection Approach to Major League Baseball Hitting Metrics. *JQAS.* 2011; 7
- Carvalho C, Chang J, Lucas J, Wang Q, Nevins J, West M. High-dimensional sparse factor modelling: Applications in gene expression genomics. *J Am Stat Assoc.* 2008; 103:1438–1456. [PubMed: 21218139]
- Carvalho LE, Lawrence CE. Centroid estimation in discrete high-dimensional spaces with applications in biology. *Proceedings of the National Academy of Sciences.* 2008; 105:3209–3214.
- Chuang HY, Lee E, Liu YT, Lee D, Ideker T. Network-based classification of breast cancer metastasis. *Molecular Systems Biology.* 2007; 3:140. [PubMed: 17940530]
- Cosgrove EJ, Yingchun Z, Gardner TS, Kolaczyk ED. Predicting gene targets of perturbations via network-based filtering of mRNA expression compendia. *BMC Bioinformatics.* 2008; 24:2482–2490.
- Cressie, NAC. Statistics for Spatial Data. New York: Wiley; 1993.
- de Oliveira V. Bayesian Analysis of Conditional Autoregressive Models. *Ann Inst Stat Math.* 2012; 64:107–133.
- di Bernardo D, Thompson M, Gardner T, et al. Chemogenomic profiling on a genome-wide scale using reverse-engineered gene networks. *Nature Biotechnology.* 2005; 23:377–383.
- Draghici S, Khatri P, Tarca AL, Amin K, Done A, Voichita C, Georgescu C, Romero R. A systems biology approach for pathway level analysis. *Genome Res.* 2007; 17:1537–1545. [PubMed: 17785539]

- Dudoit S, Shaffer JP, Boldrick JC. Multiple hypothesis testing in microarray experiments. *Stat Science*. 2003; 18:71–103.
- Faulon JL, Misra M, Martin S, Sale K, Sapra R. Genome scale enzyme-metabolite and drug-target interaction predictions using the signature molecular descriptor. *Bioinformatics*. 2008; 24:225–233. [PubMed: 18037612]
- Figueiredo-Pereira ME, Chen WE, Li J, Johdo O. The antitumor drug aclacinomycin A, which inhibits the degradation of ubiquitinated proteins, shows selectivity for the chymotrypsin-like activity of the bovine pituitary 20 S proteasome. *Journal of Biological Chemistry*. 1996; 271:16455–9. [PubMed: 8663210]
- Fritsche M, Haessler C, Brandner G. Induction of nuclear accumulation of the tumor-suppressor protein p53 by DNA damaging agents. *Oncogene*. 1993; 8:307–318. [PubMed: 8426740]
- Gandhi T, Zhong J, Mathivanan S, et al. Analysis of the human protein interactome and comparison with yeast, worm and fly interaction datasets. *Nature genetics*. 2006; 38:285–293. [PubMed: 16501559]
- Gelman A. Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis*. 2006; 1:515–533.
- Gelman, A.; Carlin, JB.; Stern, HS.; Rubin, DB. *Bayesian Data Analysis*. New York: Chapman & Hall/CRC; 2004.
- George EI, McCulloch RE. Approaches for Bayesian variable selection. *Statistica Sinica*. 1997; 7:339–373.
- Goodsell DS. The Molecular Perspective: Methotrexate. *The Oncologist*. 1999; 4:340–341. [PubMed: 10476546]
- Grandela C, Pera MF, Grimmond SM, Kolle G, Wolvetang EJ. p53 is required for etoposide-induced apoptosis of human embryonic stem cells. *Stem Cell Res*. 2007; 1:116–128. [PubMed: 19383392]
- Grenert JP, et al. The amino-terminal domain of heat shock protein 90 (hsp90) that binds geldanamycin is an ATP/ADP switch domain that regulates hsp90 conformation. *J Biol Chem*. 1997; 272:23843–23850. [PubMed: 9295332]
- Gu J, Chen Y, Li S, Li Y. Identification of responsive gene modules by network-based gene clustering and extending: application to inflammation and angiogenesis. *BMC Sys Bio*. 2010; 4:47.
- Gupta M, Fan S, Zhan Q, Kohn KW, O'Connor PM, Pommier Y. Inactivation of p53 increases the cytotoxicity of camptothecin in human colon HCT116 and breast MCF-7 cancer cells. *Clin Cancer Res*. 1997; 3:1653–1660. [PubMed: 9815856]
- Hidaka, et al. Isoquinolinesulfonamides, novel and potent inhibitors of cyclic nucleotide dependent protein kinase and protein kinase C. *Biochemistry*. 1984; 23:5036. [PubMed: 6238627]
- Huang WY, Yang PM, Chang YF, Marquez VE, Chen CC. Methotrexate induces apoptosis through p53/p21-dependent pathway and increases E-cadherin expression through downregulation of HDAC/EZH2. *Biochem Pharmacol*. 2011; 81:510–517. [PubMed: 21114963]
- Ideker T, Ozier O, Schwikowski B, Siegel AF. Discovering regulatory and signalling circuits in molecular interaction networks. *Bioinformatics*. 2002; 18(Suppl 1):S233–40. [PubMed: 12169552]
- Jaks V, Jers A, Kristjuhan A, Maimets T. p53 protein accumulation in addition to the transactivation activity is required for p53-dependent cell cycle arrest after treatment of cells with camptothecin. *Oncogene*. 2001; 20:1212–1219. [PubMed: 11313865]
- Jiang Z, Gentleman R. Extensions to gene set enrichment. *Bioinformatics*. 2007; 23:306–313. [PubMed: 17127676]
- Kanehisa M, Goto S. KEGG: Kyoto Encyclopedia of Genes and Genomes. *Nucleic Acids Res*. 2000; 28:27–30. [PubMed: 10592173]
- Kanehisa M, et al. From genomics to chemical genomics: new developments in KEGG. *Nucleic Acids Res*. 2006; 34:354–357.
- Karpinich NO, Tafani M, Rothman RJ, Russo MA, Farber JL. The course of etoposide-induced apoptosis from damage to DNA and p53 activation to mitochondrial release of cytochrome c. *J Biol Chem*. 2002; 277:16547–16552. [PubMed: 11864976]
- Khatri P, Draghici S. Ontological analysis of gene expression data: current tools, limitations, and open problems. *Bioinformatics*. 2005; 21:3587–3595. [PubMed: 15994189]

- Krause K, Wasner M, Reinhard W, et al. The tumour suppressor protein p53 can repress transcription of cyclin B. *Nucl Acids Res.* 2002; 22:4410–4441.
- Kurz EU, Douglas P, Lees-Miller SP. Doxorubicin Activates ATM-dependent Phosphorylation of Multiple Downstream Targets in Part through the Generation of Reactive Oxygen Species. *J Biol Chem.* 2004; 279:53272–53281. [PubMed: 15489221]
- Liao J, Boscolo R, Yang Y, et al. Network component analysis: reconstruction of regulatory signals in biological systems. *Proceedings of the National Academy of Sciences.* 2003; 100:15522–15527.
- Lindley DV, Smith AF. Bayes estimates for the linear model. *Journal of the Royal Statistical Society. Series B (Methodological).* 1972:1–41.
- Ling B, Wei-Guo Z. p53: Structure, Function and Therapeutic Applications. *J Cancer Mol.* 2006; 2:141–153.
- Ling YH, el Naggat AK, Priebe W, Perez-Soler R. Cell cycle-dependent cytotoxicity, G2/M phase arrest, and disruption of p34cdc2/cyclin B1 activity induced by doxorubicin in synchronized P388 cells. *Mol Pharmacol.* 1996; 49:832–841. [PubMed: 8622633]
- Liu M, et al. Network-Based Analysis of Type 2 Diabetes. *PLoS Genet.* 2010; 3:e96. [PubMed: 17571924]
- Lucas J, Carvalho CM, Wang Q, Bild A, Nevins JR, West M. Sparse statistical modelling in gene expression genomics. *Bayesian Inference for Gene Expression and Proteomics.* 2006; 14:155–176.
- Lucas JE, Kung HN, Chi JT. Latent Factor Analysis to Discover Pathway-Associated Putative Segmental Aneuploidies in Human Cancers. *PLoS Comput Biol.* 2010; 6:e1000920. [PubMed: 20824128]
- MT. Mitomycin C: small, fast and deadly (but very selective). *Chemical Biology.* 1995; 2:575–579.
- Ma H, Zhao H. iFad: an integrative factor analysis model for drug-pathway association inference. *Bioinformatics.* 2012; 28:1911–1918. [PubMed: 22581178]
- Maliga Z, Kapoor TM, JMT. Evidence that monastrol is an allosteric inhibitor of the mitotic kinesin Eg5. *Chemical Biology.* 2002; 9:989–996.
- Malik M, Nitiss KC, Enriques-Rios V, Nitiss JL. Roles of nonhomologous endjoining pathways in surviving topoisomerase II-mediated DNA damage. *Mol Cancer Ther.* 2006; 5:1405. [PubMed: 16818498]
- Marbach D, CCJ, Kffner R, NMV, Prill R, Camacho D, Allison K, Kellis M, Collins JJ, Stolovitzky G. Consortium TD. Wisdom of crowds for robust gene network inference. *Nature Methods.* 2012; 9:796–804. [PubMed: 22796662]
- Nakada S, Katsuki Y, Imoto I, Yokoyama T, Nagasawa M, Inazawa J, Mizutani S. Early G2/M checkpoint failure as a molecular mechanism underlying etoposide-induced chromosomal aberrations. *J Clin Invest.* 2006; 116:80–89. [PubMed: 16357944]
- Nam D, Kim SY. Gene-set approach for expression pattern analysis. *Brief Bioinform.* 2008; 9:189–197. [PubMed: 18202032]
- Neckers L, Schulte TW, Mimnaugh E. Geldanamycin as a potential anti-cancer agent: its molecular target and biochemical activity. *Invest New Drugs.* 1999; 17:361–373. [PubMed: 10759403]
- Obrig TG, Culp WJ, McKeen WL, Hardesty B. The mechanism by which cycloheximide and related glutarimide antibiotics inhibit peptide synthesis on reticulocyte ribosomes. *Journal of Biological Chemistry.* 1971; 246:174–181. [PubMed: 5541758]
- Pham L, Christadore L, Schaus S, Kolaczyk ED. Network-based prediction for sources of transcriptional dysregulation via latent pathway identification analysis. *Proc Nat Acad Sci.* 2011; 108:13347–13352. [PubMed: 21788508]
- Pommier Y, Leo E, Zhang H, Marchand C. DNA topoisomerases and their poisoning by anticancer and antibacterial drugs. *Chemical Biology.* 2012; 17:421–433.
- Prill RJ, Saez-Rodriguez J, Alexopoulos LG, Sorger PK, Stolovitzky G. Crowdsourcing Network Inference: The DREAM4 Predictive Signaling Network Challenge. *Science.* 2011; 4:mr7.
- Rao CV, Kurkjian CD, Yamada HY. Mitosis-targeting natural products for cancer prevention and therapy. *Curr Drug Targets.* 2012; 13:1820–1830. [PubMed: 23140292]
- Reguly T, Breitkreutz A, Boucher L, et al. Comprehensive curation and analysis of global interaction networks in *Saccharomyces cerevisiae*. *Journal of Biology.* 2006; 5:11. [PubMed: 16762047]

- Rivals I, Personnaz L, Taing L, et al. Enrichment or depletion of a GO category within a class of genes: which test? *Bioinformatics*. 2007; 23:401–407. [PubMed: 17182697]
- Roguev A, Bandyopadhyay S, et al. Conservation and rewiring of functional modules revealed by an epistasis map in fission yeast. *Science*. 2008; 322:405–410. [PubMed: 18818364]
- Storey JD, Tibshirani R. Statistical significance for genome-wide studies. *Proc Natl Acad Sci*. 2003; 100:9440. [PubMed: 12883005]
- Subramanian A, et al. Gene set enrichment analysis: A knowledge-based approach for interpreting genome-wide expression profiles. *Proc Natl Acad Sci USA*. 2005; 102:15545–15550. [PubMed: 16199517]
- Terstappen GC, Schlüpen C, Raggiaschi R, Gaviraghi G. Target deconvolution strategies in drug discovery. *Nature Reviews*. 2007; 6:891–903.
- The Gene Ontology Consortium. Gene Ontology: tool for the unification of biology. *Nature Genet*. 2000; 25:25–29. [PubMed: 10802651]
- van Dyk DA, Park T. Partially collapsed Gibbs samplers: Theory and methods. *J Amer Stat Soc*. 2008; 482:790–796.
- Vanhaecke T, Papeleu P, Elaut G, Rogiers V. Trichostatin A-like hydroxamate histone deacetylase inhibitors as therapeutic agents: toxicological point of view. *Curr Med Chem*. 2004; 11:1629–1643. [PubMed: 15180568]
- Verweij J, Pinedo HM. Mitomycin C: mechanism of action, usefulness and limitations. *Anticancer Drugs*. 1990; 1:5–13. [PubMed: 2131038]
- Vogelstein B, Kinzler KW. Cancer genes and the pathways they control. *Nature Medicine*. 2004; 10:789–799.
- von Mering C, Krause R, Snel B, Cornell M, et al. Comparative assessment of large-scale data sets of protein-protein interactions. *Nature*. 2002; 417:6887.
- Wall MM. A close look at the spatial structure implied by the CAR and SAR models. *Journal of Statistical Planning and Inference*. 2004; 121:311–324.
- Wang S, Konorev E, Kotamraju S, Joseph J, Kalivendi S, Kalyanaraman B. Doxorubicin Induces Apoptosis in Normal and Tumor Cells via Distinctly Different Mechanisms. *J Biol Chem*. 2004; 279:25535–25543. [PubMed: 15054096]
- Wei P, Pan W. Incorporating gene networks into statistical tests for genomic data via a spatially correlated mixture model. *Bioinformatics*. 2008; 24:404–411. [PubMed: 18083717]
- Yamanishi Y, Araki M, Gutteridge A, Honda W, Kanehisa M. Prediction of drug-target interaction networks from the integration of chemical and genomic spaces. *Bioinformatics*. 2008; 24:232–240.
- Zhou M, Gu L, Li F, Zhu Y, Woods WG, Findley HW. DNA damage induces a novel p53-survivin signaling pathway regulating cell cycle and apoptosis in acute lymphoblastic leukemia cells. *Mol Pharmacol*. 2002; 303:124–31.

APPENDIX A. NETWORK CONSTRUCTION

Our network W was created in a similar manner to Pham et al. (2011). We constructed a weighted network of pathways W where the nodes represent canonical pathways and the edges represent functional links. We describe the details of the construction of W below.

First, we classified genes using two sources of biological information: KEGG regulatory/signaling pathways (Kanehisa et al., 2006; Kanehisa and Goto, 2000), and GO biological processes (The Gene Ontology Consortium, 2000) obtained from the MSigDB collection (Subramanian et al., 2005) (version 3.0). We removed disease specific pathways and any metabolic pathways, focusing our attention to a regulatory/signaling network. We also removed 4 pathways that did not represent a single pathway but a collection of signaling molecules that are themselves part of other pathways (e.g. Cell Adhesion Molecules and

Cytokine-Cytokine Receptor Interactions). In the end, we had 76 distinct regulatory/signaling KEGG pathways containing a total of 3636 genes.

To define functional links between two pathways, we constructed a bipartite network, where the two sets of nodes represent KEGG pathways and GO functions, respectively. A pathway P and a GO term G were linked together in this bipartite network if the intersection of P and G (as gene sets) is non-empty. Furthermore, the edge was weighted by the Jaccard index between G and P , measuring the relative overlap between function and pathway. Edges with Jaccard index smaller than 3% were removed to reduce the possibility of false positives in the database. We can represent this bipartite graph as an incidence matrix M where the rows represent KEGG pathways and the columns represent GO terms.

Our final network of pathways was obtained by translating the information from this two-mode network onto a one mode network formed by the KEGG node set alone. The projection of M onto a single mode network A was obtained by $A = M^T M$. As a result, two pathways are linked in this network if and only if they share at least one biological process. Furthermore, edges between pathways that heavily contribute to the same GO terms will be heavier than those pathways that contribute less to the same GO terms. Finally, we standardized A in the following manner:

$$W_{ij} = \frac{A_{ij}}{\sqrt{A_{ii}A_{jj}}}(1 - \delta_{ij}),$$

where δ is the Kronecker delta.

APPENDIX B. MARGINALIZED CONDITIONAL POSTERIOR DISTRIBUTION OF θ

In this section, we derive the marginalized posterior of θ_i for sample i . After integrating out ρ_i from Eq. (8) we have that

$$\omega_i | \theta_i, \Phi \stackrel{\text{ind}}{\sim} N(0, V(\theta_i)),$$

where $V(\theta_i) = \sigma^2 \Phi + \Phi \Sigma(\theta_i) \Phi^T$ and the index runs over samples. To simplify the notation, we set $\tau_0^2 := v_0 \tau^2$ and drop the indices, that is, $\omega | \theta, \Phi \sim N(0, V(\theta))$.

Since $\theta_j \stackrel{\text{iid}}{\sim} \text{Bern}(\alpha)$ from the prior and with $\Sigma(\theta)$ similarly defined as in Section 4.3, we have that

$$\begin{aligned} V(\theta) &= \sigma^2 \Phi + \tau_0^2 \Phi \Phi^T + (\tau^2 - \tau_0^2) \Phi \text{Diag}(\theta) \Phi^T \\ &= \sigma^2 \Phi + \tau_0^2 \Phi \Phi^T + (\tau^2 - \tau_0^2) \sum_{j=1}^q \theta_j \phi_j \phi_j^T, \end{aligned}$$

where ϕ_j is the j -th column of Φ .

We want

$$\mathbb{P}(\theta_j=1|\theta_{[-j]}, \omega, \Phi) = \frac{\mathbb{P}(\theta_j=1, \theta_{[-j]}, \omega, \Phi)}{\sum_{b \in \{0,1\}} \mathbb{P}(\theta_j=b, \theta_{[-j]}, \omega, \Phi)}.$$

By the distributions above,

$$\mathbb{P}(\theta_j, \theta_{[-j]}, \omega, \Phi) \propto \prod_j \alpha^{\theta_j} (1-\alpha)^{1-\theta_j} (2\pi)^{-q/2} |V(\theta)|^{-1} \exp \left\{ -\frac{1}{2} \omega^\top V(\theta)^{-1} \omega \right\}.$$

Now, we define $V_0 = V(\theta_j=0, \theta_{[-j]})$ and $V_1 = V(\theta_j=1, \theta_{[-j]})$. Then,

$$\begin{aligned} V_1 &= \sigma^2 \Phi + \tau_0^2 \Phi \Phi^\top + (\tau^2 - \tau_0^2) \phi_j \phi_j^\top + (\tau^2 - \tau_0^2) \sum_{k \neq j} \theta_k \phi_k \phi_k^\top \\ &= V_0 + (\tau^2 - \tau_0^2) \phi_j \phi_j^\top, \end{aligned}$$

and so, by the Sherman-Morrison formula,

$$\begin{aligned} V_1^{-1} &= (V_0 + (\tau^2 - \tau_0^2) \phi_j \phi_j^\top)^{-1} \\ &= V_0^{-1} - \frac{(\tau^2 - \tau_0^2) V_0^{-1} \phi_j \phi_j^\top V_0^{-1}}{1 + (\tau^2 - \tau_0^2) \phi_j^\top V_0^{-1} \phi_j}. \end{aligned}$$

Moreover, by the matrix determinant lemma,

$$|V_1| = (1 + (\tau^2 - \tau_0^2) \phi_j^\top V_0^{-1} \phi_j) |V_0|.$$

Using these relations we can proceed to sample $\theta_j^{(t+1)}$ based on $\theta^{(t)}$ from the current iteration t ; we keep the inverse of $V(\theta)$ and only update it when θ_j is flipped. We have

$$\mathbb{P}(\theta_j^{(t+1)}=1|\theta_{[-j]}, \omega, \Phi) = \frac{|V_1|^{-1/2} \exp\{-\frac{\omega^\top V_1^{-1} \omega}{2}\} \alpha}{|V_1|^{-1/2} \exp\{-\frac{\omega^\top V_1^{-1} \omega}{2}\} \alpha + |V_0|^{-1/2} \exp\{-\frac{\omega^\top V_0^{-1} \omega}{2}\} (1-\alpha)},$$

or, taking logits,

$$\text{logit} \mathbb{P} \left(\theta_j^{(t+1)} = 1 | \theta_{[-j]}, \omega, \Phi \right) = -\frac{1}{2} \log \frac{|V_1|}{|V_0|} - \frac{1}{2} \omega^\top (V_1^{-1} - V_0^{-1}) \omega + \text{logit}(\alpha). \quad (\text{A.1})$$

There are two cases: when $\theta_j^{(t)} = 0$ and when $\theta_j^{(t)} = 1$. In the first case we already have V_0^{-1} , and so, by defining $\delta_{j0} = 1 + (\tau^2 - \tau_0^2) \phi_j^\top V_0^{-1} \phi_j$ and $\Delta_{j0} = (\tau^2 - \tau_0^2) / \delta_{j0} (V_0^{-1} \phi_j) (V_0^{-1} \phi_j)^\top$ we can apply Eq. (A.1) to obtain

$$\text{logit} \mathbb{P} \left(\theta_j^{(t+1)} = 1 | \theta_{[-j]}, \omega, \Phi \right) = -\frac{1}{2} \log \delta_{j0} + \frac{1}{2} \omega^\top \Delta_{j0} \omega + \text{logit}(\alpha).$$

In the sampler, we update $(V(\theta)^{-1})^{(t+1)} = (V(\theta)^{-1})^{(t)} - \Delta_{j0}$ if $\theta_j^{(t+1)} = 1$, that is, when θ_j is flipped.

If $\theta_j^{(t)} = 1$ we define $\delta_{j1} = 1 - (\tau^2 - \tau_0^2) \phi_j^\top V_1^{-1} \phi_j$ and $\Delta_{j1} = (\tau^2 - \tau_0^2) / \delta_{j1} (V_1^{-1} \phi_j) (V_1^{-1} \phi_j)^\top$ in a similar way to obtain:

$$\text{logit} \mathbb{P} \left(\theta_j^{(t+1)} = 1 | \theta_{[-j]}, \omega, \Phi \right) = \frac{1}{2} \log \delta_{j1} + \frac{1}{2} \omega^\top \Delta_{j1} \omega + \text{logit}(\alpha),$$

and update $(V(\theta)^{-1})^{(t+1)} = (V(\theta)^{-1})^{(t)} + \Delta_{j1}$ if θ_j gets flipped to $\theta_j^{(t+1)} = 0$.

APPENDIX C. EFA RESULTS ON THE DREAM 7 DATA SETS

Fig 12 includes heatmaps of the results obtained by the EFA model for the IC20 at 12 hours and IC20 at 24 hours exposure times, respectively.

At IC20 12 hours, the EFA model identified fewer pathways than the CFA-CAR model, identifying P53-signaling for CPT (also identified by CFA-CAR in addition to other DNA-damaging agents ETP, MMC, and DOX). At IC20 24 hours, the EFA and CFA-CAR models both identified P53-signaling for the same pathways (MMC, DOX, MTX, and CPT), as well as cell cycle for ETP, a pathway mechanistically linked to p53-signaling. However, CFA-CAR further identified cell cycle for CPT, DOX, and RPM. Moreover, both EFA and CFA-CAR both found MMC and DHCL overly perturbed at IC20 24 hours.

Figure 1.
A representation of a drug perturbation to a pathway A and its downstream effects on associated pathways and gene expression.

Figure 2.

Mean ROC curves with standard error bars for the CFA-CAR and EFA models under different signal to noise ratios. The k -th set of error bars represents the mean and standard deviations in the k -th FPR and TPR across ten data replicates for each method. Each sub-figure corresponds to a different signal to noise ratio in the data. The solid line represents an expected ROC curve under random guessing. Note that we truncated the x-axis at 0.42 because there were no additional plot points until $(x = 1, y = 1)$.

Figure 3.
Boxplot of the AUCs of the ROC plots in Fig. 2.

Figure 4. Boxplots of the relative error across different percentages of database inaccuracy for W (left), the mask of Λ (middle), and the posterior means of Λ (right.)

Figure 5.

Mean ROC curves with standard error bars for the CFA-CAR and EFA models under different percentages of randomly reassigned genes to pathways. The k -th set of error bars represents the mean and standard deviations in the k -th FPR and TPR across ten data replicates of W and Λ . The standard error bars on the EFA model represent standard errors across ten replicates of Λ . Each sub-figure corresponds to a different percentage of randomly reassigned genes. The solid line represents an expected ROC curve under random guessing.

Figure 6.
Boxplot of the AUCs of the ROC plots in Fig. 5.

Figure 7.

Beanplots with bean averages of the marginal posterior distributions of γ , σ^2 , τ^2 , and the SNR across each drug for the IC20 at 12 hours and IC20 at 24 hours data sets.

Figure 8.

Top and bottom panels show model fit assessment for the IC20 at 12 hours and IC20 at 24 hours datasets, respectively. Left: scatterplot of standardized data with zero expectation and smoother (lowess) fit shown by a red line and green curve, respectively. Middle: histogram, density, and rug plot of the standardized dataset averaged over samples. Right: quantile-quantile normal plot of standardized data averaged over samples.

Figure 9.

Boxplots representing the posterior means of θ for each of the 76 pathways in the control samples in the IC20 12 hours data set (Top) and the IC20 24 hours data set (Bottom) in the leave-three-out cross-validation using the CFA-CAR model (red) and the EFA model (blue).

Figure 10.

A. Heatmap of the CFA-CAR posterior means of Θ for IC20 at 12 hours. Green cells indicate high posterior probabilities (above the $x = 0.72$ threshold or equivalently, less than 5% BFDR). B. Heatmap of (1-FDR adjusted P-values from) GSEA for IC20 at 12 hours. Green cells indicate small FDR values (above a 5% threshold). In both heatmaps, x-axes indicate drug experiments (samples) and y-axes indicate pathways. Each axis is clustered hierarchically using complete linkage and Euclidean distances as dissimilarities.

Figure 11.

A. Heatmap of the CFA-CAR posterior means of Θ for IC20 at 24 hours. Green cells indicate high posterior probabilities (above the $\alpha = 0.50$ threshold or equivalently, less than 5% BFDR). B. Heatmap of (1-FDR adjusted P-values from) GSEA for IC20 at 24 hours. Green cells indicate small FDR values (above a 5% threshold). In both heatmaps, x-axes indicate drug experiments (samples) and y-axes indicate pathways. Each axis is clustered hierarchically using complete linkage and Euclidean distances as dissimilarities.

Figure 12.

A. Heatmap of the EFA posterior means of Θ for IC20 at 12 hours. Green cells indicate high posterior probabilities (above the $x = 0.80$ threshold or equivalently, less than 5% BFDR). B. Heatmap of the EFA posterior means of Θ for IC20 at 24 hours. Green cells indicate high posterior probabilities (above the $x = 0.70$ threshold or equivalently, less than 5% BFDR). In both heatmaps, x-axes indicate drug experiments (samples) and y-axes indicate pathways. Each axis is clustered hierarchically using complete linkage and Euclidean distances as dissimilarities.

Table 1

Compounds tested on the LY3 cancer cell line in the NCI-DREAM data challenge, the abbreviations used in the figures and description of drug and known mechanisms of action.

DRUG	Abbrev.	Description
Aclacinomycin A	ACLA-A	Inhibits the degradation of ubiquitinated proteins, affecting the proteasome (Figueiredo-Pereira M et al., 1996)
Blebbistatin	BLEBB	Inhibits the myosin II protein affecting cellular motility pathways (Allingham et al., 2005)
Camptothecin	CPT	DNA-damaging agent targeting DNA topoisomerase I, affecting cell-cycle and p53-signaling (Gupta et al., 1997; Jaks et al., 2001; Wang et al., 2004)
Cycloheximide	CHX	Inhibits protein biosynthesis pathways (Obrig et al., 1971)
Doxorubicin Hydrochloride	DOX	DNA-damaging agent targeting DNA topoisomerase II (Pommier et al., 2012), and causes cell cycle arrest (Ling et al., 1996). Implicated in a P53-dependent mechanism of action (Ling et al., 1996; Zhou et al., 2002; Kurz et al., 2004).
Etoposide	ETP	DNA-damaging agent causing DNA strands to break causing apoptosis and errors in DNA synthesis (Pommier et al., 2012). Involved in P53-dependent mechanisms of action (Karpnich et al., 2002; Grandela et al., 2007)
Geldanamycin	GA	Disrupts regulatory signaling mechanisms via HSP90 inhibition (Grenert et al., 1997; Neckers et al., 1999)
H-7, Dihydrochloride	DHCL	Protein kinase C inhibitor (Hidaka et al., 1984)
Methotrexate	MTX	Inhibits DNA-synthesis via inhibition of purine and thymine synthesis (Goodsell, 1999)
Mitomycin C	MTC.	DNA-damaging agent that is a potent DNA crosslinker (M, 1995)
Monastrol	MNS	Inhibits ATPase activity of the kinesin (Maliga et al., 2002)
Rapamycin	RPM	Targets the mTOR protein (Alqurashi et al., 2013) affecting mTOR signaling pathway
Trichostatin A	TSA.	Inhibits the class of histone deacetylase (HDAC) families of enzymes, inhibiting DNA transcription pathways (Vanhaecke et al., 2004)
Vincristine	VCR	Causes cell-cycle arrest via inhibiting the assembly of microtubule structures during mitosis (Rao et al., 2012)