

Inferring signalling dynamics by integrating interventional with observational data

Mathias Cardner^{1,2}, Nathalie Meyer-Schaller^{3,†}, Gerhard Christofori³ and Niko Beerenwinkel^{1,2,*}

¹Department of Biosystems Science and Engineering, ETH Zurich, ²SIB Swiss Institute of Bioinformatics and

³Department of Biomedicine, University of Basel, Basel 4058, Switzerland

*To whom correspondence should be addressed.

†Present address: Institute of Pathology, University Hospital Basel, Basel 4031, Switzerland

Abstract

Motivation: In order to infer a cell signalling network, we generally need interventional data from perturbation experiments. If the perturbation experiments are time-resolved, then signal progression through the network can be inferred. However, such designs are infeasible for large signalling networks, where it is more common to have steady-state perturbation data on the one hand, and a non-interventional time series on the other. Such was the design in a recent experiment investigating the coordination of epithelial–mesenchymal transition (EMT) in murine mammary gland cells. We aimed to infer the underlying signalling network of transcription factors and microRNAs coordinating EMT, as well as the signal progression during EMT.

Results: In the context of nested effects models, we developed a method for integrating perturbation data with a non-interventional time series. We applied the model to RNA sequencing data obtained from an EMT experiment. Part of the network inferred from RNA interference was validated experimentally using luciferase reporter assays. Our model extension is formulated as an integer linear programme, which can be solved efficiently using heuristic algorithms. This extension allowed us to infer the signal progression through the network during an EMT time course, and thereby assess when each regulator is necessary for EMT to advance.

Availability and implementation: R package at <https://github.com/cbg-ethz/timeseriesNEM>. The RNA sequencing data and microscopy images can be explored through a Shiny app at <https://emt.bsse.ethz.ch>.

Contact: niko.beerenwinkel@bsse.ethz.ch

Supplementary information: [Supplementary data](#) are available at *Bioinformatics* online.

1 Introduction

In order to learn biological networks, gene silencing techniques such as RNA interference provide a means to infer causal interactions. This is often coupled with high-throughput sequencing technologies such as RNA sequencing to measure the mRNA expression of genes in perturbed cells. In the literature, methods for learning regulatory networks from gene expression data tend to rely on the assumption that the expression level of a given gene is a reliable proxy for the activity of the protein which it encodes (Friedman, 2004; Pe'er *et al.*, 2001; Segal *et al.*, 2003). The same holds true for more recent approaches incorporating perturbation data (Peters *et al.*, 2016) and time series (Chen *et al.*, 2017). This is appropriate in transcriptionally regulated networks, where changes in gene expression account for all causal interactions. However, in networks consisting of transcription factors and microRNAs, post-transcriptional regulation

and post-translational modifications play a major role in signalling. In such a case of non-transcriptional regulation, we can model the regulators as hidden variables, and attempt to infer their network structure through their perturbation effects on the transcriptome. The nested effects model (NEM) (Markowitz *et al.*, 2005) was designed for this very purpose.

In the original NEM framework (Markowitz *et al.*, 2007), the signalling network is considered static, and samples are assumed to be measured at a time when the perturbed system has reached a steady state. The dynamic NEM (Anchang *et al.*, 2009) is an attempt to include time series from perturbation experiments in the NEM framework. However, this method still relies on the steady state network as a basis for inference, and the main benefit is the ability to prune transitively closed edges from the static network. A more comprehensive treatment of dynamic NEM (Fröhlich *et al.*, 2011)

unrolls the NEM over time—similar to the idea of a dynamic Bayesian network—which can resolve feedback loops in the network. The inferred network is still static, but the time required for propagation of signals is allowed to vary. In contrast, the hidden Markov NEM (Wang et al., 2014) allows for the network structure itself to change over time.

All the above approaches require time series from perturbation experiments. In other words, the perturbed system must be measured at several time points until it reaches a steady state. While this experimental setup is possible for smaller studies containing few regulators, it becomes infeasible for larger networks. In the latter case, it is more common to have steady state data from perturbation experiments on the one hand, and a time series from an unperturbed system on the other. In such a setting we cannot reliably infer a dynamic network, nor resolve feedback loops in a static network. Instead, we propose to use the perturbation data to infer a steady state network, and then map the time series onto that network in order to assess the progression to the steady state. Given the differential expression profile at time point t versus the steady state, we try to find the most likely signalling state of the static network to yield such an effect profile.

Our method is motivated by a recent experiment investigating epithelial–mesenchymal transition (EMT) in murine mammary gland cells (Meyer-Schaller et al., 2019). In cancer, EMT plays an important role in the development of metastasis, which is the major cause of death for patients with cancer (Chaffer et al., 2016). Cells in the epithelial state tend to be stationary, whereas cells in the mesenchymal state are migratory and invasive, which enables cancer cells to enter the blood stream and to seed distant metastases. On the other hand, the reverse transition from a mesenchymal to an epithelial state (MET) is believed to boost the outgrowth of seeded metastatic cancer cells in distant organs (Brabletz, 2012). However, recent findings indicate that full EMT and MET processes may not be required for metastasis, but rather that transitions to intermediate EMT stages could be sufficient (Pastushenko et al., 2018).

EMT entails major changes in the transcriptional and epigenetic landscape of a cell, and multiple factors regulating this transition have been identified. In particular, the interplay between Snail and Zeb transcription factors with microRNAs miR-34 and miR-200 is believed to regulate the acquisition of intermediate EMT states (Nieto et al., 2016). A recent screening of >1500 gene regulators in normal murine mammary gland (NMuMG) cells led to the identification of 46 transcription factors and 13 microRNAs which were necessary for EMT to occur (Meyer-Schaller et al., 2019). Here, we aimed at inferring a signalling network of those 59 regulators and estimating the signal progression through the network during the transition from the epithelial to the mesenchymal state.

Treating epithelial cells with transforming growth factor beta (TGF- β) triggers a signalling cascade which causes the cells to transition into a mesenchymal state (Fig. 1a). If any of the crucial 59 transcription factors and microRNAs is perturbed before and during the TGF- β treatment, then the cells do not reach the mesenchymal state, but are instead halted at an intermediate stage (Fig. 1b). We hypothesized that perturbing a relevant transcription factor or microRNA would leave an informative record on the transcriptome. That is, when comparing the transcriptome measured in perturbation experiments versus a control experiment (all sampled 96 h after TGF- β treatment), regulators which are crucial early in the signalling cascade should show more extensive differential expression than regulators which play a role late in the cascade. Furthermore, we would expect to find a subset structure in the perturbation effects, due to signal propagation in the underlying pathway

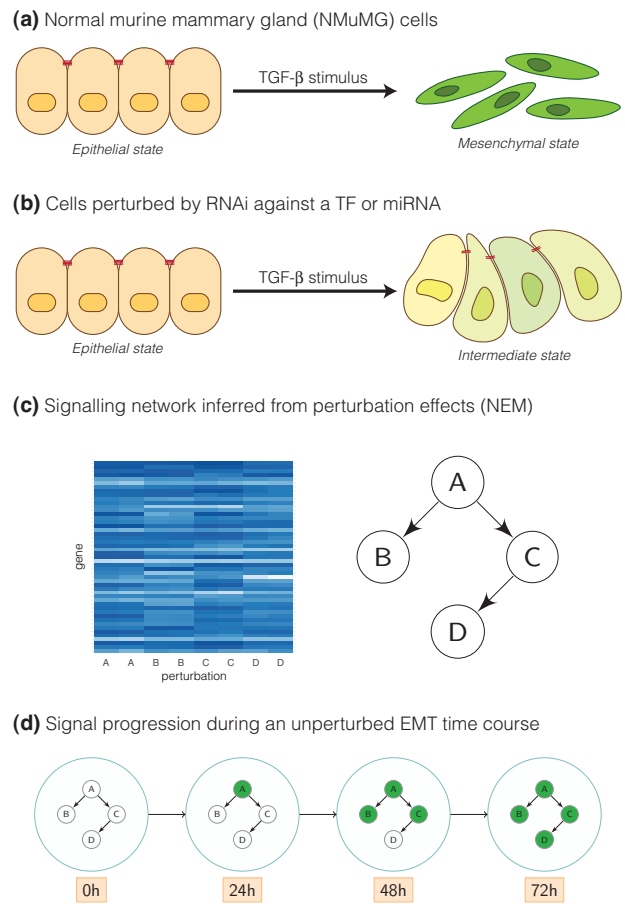


Fig. 1. Synopsis. (a) Normal cells undergo EMT induced by TGF- β . (b) EMT is halted in cells perturbed by RNAi against any of the 59 regulators under investigation. Despite TGF- β treatment to induce EMT, the perturbed cells do not reach the mesenchymal phenotype, but are halted at intermediate stages. (c) Nested effects model to infer a non-transcriptional signalling network from downstream perturbation effects on genes measured using RNA sequencing. (d) Inferring the signal progression during an unperturbed EMT time course

(Markowitz et al., 2005). We used a NEM (Fröhlich et al., 2008; Markowitz et al., 2007) to infer a signalling network from downstream perturbation effects based on RNA sequencing data from RNA interference experiments sampled after 96 h of TGF- β treatment (Fig. 1c). Following the hypothesis above, we expect that genes directly regulated by a transcription factor or microRNA crucial for instance 24 h, would show similar differential expression between perturbation and control (both sampled at 96 h), as they do in an unperturbed time course when comparing 24 h with 96 h. Based on this assumption, we developed a method for inferring signal progression through the pathway, in order to infer the time point at which each regulator is necessary for EMT to progress (Fig. 1d). Our model (Section 2.3) thus integrates a non-interventional time series with perturbation data, and is formulated as an integer linear programme which can be solved efficiently using heuristic algorithms.

2 Materials and methods

We analysed data gathered in the experiments described by Meyer-Schaller et al. (2019). In short, transforming growth factor beta (TGF- β) was used to induce EMT in epithelial NMuMG cells, but

RNA interference (RNAi) against certain regulators prevented the cells from reaching the mesenchymal state. A screening of approximately 1500 RNAi interventions identified 46 transcription factors (TFs) and 13 microRNAs (miRNAs) which were necessary for EMT to occur. Our aim was to infer the signalling network of those 59 regulators. Although direct regulation between TFs themselves and between TFs and miRNAs can occur, the signalling activity of the 59 regulators can also be regulated on post-transcriptional level, for instance by the miRNAs interfering with protein translation (Wang *et al.*, 2010). Moreover, TF activities can be affected indirectly through post-translational modification and nuclear localization. Therefore, modelling interactions between multiple TFs and miRNAs is complex and poorly understood (Le *et al.*, 2013).

We obtained RNA sequencing data from two separate experimental set-ups on epithelial normal murine mammary gland (NMuMG) cells; namely one EMT time course and one set of 59 RNAi perturbation experiments. For the time course, epithelial NMuMG cells were treated with TGF- β and subsequently sampled at 12 time points between 0 and 10 days. Already after 4 days, the cells exhibited a mesenchymal phenotype (Fig. 1a). The experiment was performed in biological duplicates.

For each of the 46 transcription factors, Ambion siRNAs were used to knock down the transcript coding for the given TF. For each of the 13 miRNAs, RNAi amounted to over-expression of the given miRNA. RNAi was performed against the regulators of interest, starting 2 days before TGF- β treatment, so that the RNAi would have time to take effect before EMT was induced. After 4 days of TGF- β treatment, the cells were sampled and sequenced. At this time point, the cells had been halted in an intermediate stage between epithelial and mesenchymal states (Fig. 1b). In addition, control samples with inert (non-targeting) siRNA or miRNA transfections were sampled after 0 and 4 days of TGF- β treatment. All experiments were performed in biological duplicates.

2.1 The standard NEM

We here outline the original nested effect model (Markowetz *et al.*, 2005, 2007), in order to introduce concepts and notation required to explain how we integrate perturbation data with a non-interventional time series (Section 2.3).

Figure 2 illustrates the structure of a NEM. The goal is to infer a signalling network Φ , represented as a directed acyclic graph involving the regulators, also referred to as signalling nodes (S-nodes), denoted S_j for $j \in \{1, \dots, n\}$. The classic NEM framework uses the term ‘S-genes’, which can be misleading, since we here refer to gene products, such as transcription factors. Furthermore, in our case 13 of the regulators are miRNAs, which is why we use the term




Fig. 2. Nested effects model. The S-nodes are modelled as hidden variables, and we aim to infer their causal graph Φ (solid arrows). In experiments separately perturbing each S-node, we observe the differential expression of all E-genes. Assuming that each E-gene is directly regulated by at most one S-node in Φ , we compute the maximum a posteriori attachment Θ (dashed arrows) of effect genes to S-nodes. We search for the signalling graph Φ which yields the most likely probabilistic nesting of effects

‘S-nodes’. The signalling nodes are individually perturbed in separate RNAi experiments, but their activity is not directly measured. Instead, we measure the expression level of effect reporter genes (E-genes) denoted E_i for $i \in \{1, \dots, m\}$, which are affected by the perturbation of S-nodes. The attachment of effect genes to regulators is encoded by Θ , and the network reconstruction is based on an assumption of parsimony. Namely, if an E-gene responds to the perturbation of two S-nodes, we assume that it is directly regulated by only one of them, which acts as a mediator for the other regulator. For instance, in Figure 2, S_1 and S_2 both regulate $\{E_1, E_2\}$. However, since the effects of perturbing S_1 are a subset of the effects of perturbing S_2 , we deduce that S_1 is downstream of S_2 in the signalling pathway, and we infer the edge $S_2 \rightarrow S_1$. Thus, the regulators are hierarchically ordered so that the effects of a child node are a subset of the effects of its parent. Due to this nesting of effects, we implicitly assume that the signalling network Φ is transitively closed; that is, if $S_1 \rightarrow S_2 \rightarrow S_3$, then $S_1 \rightarrow S_3$. In practice, we need to account for noise in the data, which requires a probabilistic model.

Let D be the data matrix such that D_{ij} records how E-gene E_i was affected by the perturbation of S-node S_j . (More details are given in Section 2.2) Suppose that we have a candidate network Φ , which is a directed acyclic graph (DAG) of S-nodes S_1, \dots, S_n . We are ultimately interested in the posterior probability of the model:

$$P(\Phi | D) = \frac{P(D | \Phi)P(\Phi)}{P(D)}$$

where the denominator does not depend on Φ , and can be disregarded for model comparison. Since almost nothing is known about the signalling network under our investigation, we use a uniform prior $P(\Phi)$. Thus we focus entirely on the likelihood $P(D | \Phi)$. It can be computed by marginalizing over E-gene attachments Θ , or by using the maximum a posteriori (MAP) estimate of Θ (Tresch and Markowetz, 2008). We will follow the latter approach, and thus solve the following:

$$\arg \max_{\Phi} \{ \max_{\Theta} P(D | \Phi, \Theta) \}.$$

The outer maximization is done by evaluating candidate DAGs in a heuristic way (see Section 2.2). The inner maximization over Θ is performed as follows.

Let $\Theta = (\theta_i)_{i=1}^m$ denote the attachment of E-genes E_1, \dots, E_m to the signalling nodes, such that $\theta_i = j$ if E_i is attached to S_j . We assume that the attachments of E-genes are conditionally independent given the signalling network Φ , so that $P(\Theta | \Phi) = \prod_{i=1}^m P(\theta_i | \Phi)$. Furthermore, we assume that E_1, \dots, E_m are conditionally independent given Φ and Θ . Thus, we do not assume genes to be marginally independent, but rather conditionally independent given their attachments to the regulators. Hence,

$$\max_{\Theta} P(D | \Phi, \Theta) = \prod_{i=1}^m \max_j P(D_i | \Phi, \theta_i = j). \quad (1)$$

The full derivations are provided in the [Supplementary Material](#).

2.2 Scoring and searching

In order to judge from noisy observations whether a certain E-gene is affected by the perturbation of a given S-node, we used the R (R Core Team, 2018) package DESeq2 (Love *et al.*, 2014) to perform differential expression analysis of RNA sequencing data. Each perturbed condition was compared with an unperturbed control, yielding a P -value for each gene, with the null hypothesis of no differential expression. The data matrix entry D_{ij} is the P -value

corresponding to gene E_i under perturbation of S_j . The alternative hypothesis of differential expression is central to our analysis. We therefore fit a beta-uniform mixture model to the distribution of P -values across genes in a given condition (Pounds and Morris, 2003). Since the null distribution of a P -value is uniform, the density function of the alternative hypothesis can be estimated from the fitted mixture components (Fröhlich et al., 2008); see details in [Supplementary Material](#).

For each perturbation experiment k , we estimated the probability density function f_k of P -values under the alternative hypothesis, such that:

$$P(D_{ik} | \Phi, \theta_i = j) := \begin{cases} f_k(D_{ik}) & \text{if } \Phi_{kj} = 1 \\ 1 & \text{if } \Phi_{kj} = 0 \end{cases} \quad (2)$$

where $(\Phi_{kj}) \in \{0, 1\}^{n \times n}$ denotes the adjacency matrix of the candidate network Φ . In words, if Φ predicts an effect on E_i when perturbing S_k (given $\theta_i = j$), we model the P -value D_{ik} using the alternative density f_k ; otherwise we model it with the uniform density 1. With this definition, we can evaluate [Equation \(1\)](#).

Since the number of DAGs grows super-exponentially with the number of nodes, an exhaustive enumeration is not feasible for signalling networks with more than four or five S -nodes. In order to infer larger networks, we must apply heuristic approaches. Specifically, we used a divide-and-conquer method called module network (Fröhlich et al., 2008), which is tailored to NEMs. It exploits the fact that the NEM objective is to find a subset structure among effect genes. Thus, the module network approach subdivides the problem by initially performing hierarchical clustering based on correlation distance of effect profiles, in order to group signalling nodes into modules containing at most four S -nodes. Then, an exhaustive search is performed within each module. Finally, modules are joined by the highest scoring pairwise connections between modules.

2.3 Mapping a time series onto a static network

Suppose that we have inferred a signalling network $\hat{\Phi}$ from a perturbed system sampled upon reaching a steady state at time t_s after applying a pathway stimulus. Now we aim to estimate the signal progression at a time point $t < t_s$ from an unperturbed time series sampled after application of the same stimulus. We can compute the maximum a posteriori attachment $\hat{\Theta}$ of E-genes to S -nodes in $\hat{\Phi}$ (see details in [Supplementary Material](#)). The *regulatory module* of a signalling node S_j is the set of E-genes whose MAP attachment is to S_j , that is the set $\{E_i : \hat{\theta}_i = j\}$. Using the unperturbed time series, we consider differential expression of E-genes in a given regulatory module as a proxy for the activity of its direct regulator.

For a given time point $t \in [0, t_s)$, we perform a differential gene expression analysis contrasted with the steady-state time point t_s . Let the P -value profile of the differential expression tests at time point t versus t_s be denoted by $T^{(t)} \in [0, 1]^m$, where m is the number of E-genes. Thus $T_i^{(t)}$ is the P -value of the i th gene from the differential expression test contrasting t with the steady state t_s . Based on this, we want to estimate the most likely state of the signalling network at time t . In other words, we are asking what signalling activities (perturbed or not) of the S -nodes in $\hat{\Phi}$ are most likely giving rise to the differential expression profile $T^{(t)}$ which we observe. This implicitly rests on the assumption that the n S -nodes in $\hat{\Phi}$ sufficiently explain the observed changes to the transcriptome, and that there is no substantive confounding by latent variables.

Let $K \in \{0, 1\}^n$ record the binary signalling states of the n regulators, where $K_j = 1$ if S_j is perturbed and $K_j = 0$ otherwise. Given

the above assumption, gene E_i is not differentially expressed if and only if $\text{pa}(E_i)$ is not perturbed under K . The former is the null hypothesis of the differential expression test, and the latter is a statement about the activity of S -nodes in $\hat{\Phi}$. In keeping with the transitive closure and propagation of signals in $\hat{\Phi}$, gene E_i is expected to be differentially expressed if and only if any of the ancestors of E_i in $\hat{\Phi}$ —denoted $\text{anc}(E_i)$ —is perturbed under K . Thus, based on the same reasoning as in [Equation \(2\)](#),

$$P(T_i^{(t)} | \hat{\Phi}, \hat{\Theta}, K) = \begin{cases} f_t(T_i^{(t)}) & \text{if } \exists S_j \in \text{anc}(E_i) : K_j = 1 \\ 1 & \text{otherwise} \end{cases} \quad (3)$$

$$= f_t(T_i^{(t)})^{\mathbb{1}(\exists S_j \in \text{anc}(E_i) : K_j = 1)} \quad (4)$$

where f_t is the alternative density of the differential expression test comparing time point t with the steady-state time point t_s . Again assuming conditional independence of E-genes given their attachment $\hat{\Theta}$ to the signalling network $\hat{\Phi}$, we have,

$$P(T^{(t)} | \hat{\Phi}, \hat{\Theta}, K) = \prod_{i=1}^m f_t(T_i^{(t)})^{\mathbb{1}(\exists S_j \in \text{anc}(E_i) : K_j = 1)}$$

which we seek to maximize with respect to K . This is equivalent to solving:

$$\arg \max_K \sum_{i=1}^m \mathbb{1}(\exists S_j \in \text{anc}(E_i) : K_j = 1) \log f_t(T_i^{(t)}). \quad (5)$$

Due to the transitively closed nature of $\hat{\Phi}$, we need only search for solutions K which are consistent with $\hat{\Phi}$, in the sense that if $K_j = 1$, then $K_k = 1$ for every node S_k downstream of S_j in $\hat{\Phi}$. Given this constraint, the sum in [Equation \(5\)](#) can be written as follows:

$$\sum_{j=1}^n K_j c_j \quad \text{where} \quad c_j = \sum_{i \in [m] : \hat{\theta}_i = j} \log f_t(T_i^{(t)}).$$

With $c = (c_1, \dots, c_n)^\top$ the problem in [Equation \(5\)](#) is therefore equivalent to maximizing $c^\top K$ given the constraint that if $K_j = 1$, then $K_k = 1$ for every descendant S_k of S_j in $\hat{\Phi}$. This constraint can be formulated as $AK \leq 0$ for a matrix A which we construct as follows. For each edge $S_j \rightarrow S_k$ in the transitive reduction of $\hat{\Phi}$, we create a row A_i such that $A_{ij} = 1$ and $A_{ik} = -1$, with $A_{i\ell} = 0$ for $\ell \neq j, k$. The optimization problem in [Equation \(5\)](#) can thus be formulated as a binary integer linear programme with canonical form:

$$\begin{aligned} & \text{minimize} && -c^\top K \\ & \text{subject to} && AK \leq 0 \end{aligned}$$

where K is a binary vector. We used the R package `lpSolve` (Berkelaar et al., 2015) to solve this problem.

2.4 Validation experiments

Luciferase reporter assays were used to assess the accuracy of the static NEM signalling network edge predictions in NMuMG epithelial cells which underwent EMT upon the addition of TGF- β . The assays were performed as previously reported by Meyer-Schaller et al. (2019). In short, RNAi was used to down-regulate a transcription (co)factor by siRNA, or to over-express a miRNAs. One day later, TGF- β was added to the cells to induce EMT, followed by retransfection (day 2) and transfection of a *Firefly* luciferase reporter 2 days after EMT induction (day 3) together with a *Renilla* luciferase reporter to normalize for cell counts. After 4 days of TGF- β

treatment (day 5), cells were assayed for luciferase activity. To measure Smad4 or Sox4 activities, *Firefly* luciferase reporters with Smad3/4 binding sequence (Denkler *et al.*, 1998) or seven specific Sox4 binding sites (Meyer-Schaller *et al.*, 2019) were used, respectively. The activity of the Notch co-activator Maml2 was assessed with a Notch reporter containing specific Rbpjk response elements (Promega, CS173601).

3 Results

We used the Bioconductor R package *nem* (Froehlich *et al.*, 2018) to infer a static signalling network from the RNAi perturbation data. In order to base the analysis on relevant genes, we included only genes which were differentially expressed (FDR < 5%) during the time course, provided that they were also differentially expressed in at least one perturbation experiment. This left us with 11 270 E-genes. Furthermore, the *nem* fitting procedure included a so-called null node, to which uninformative genes were attached. In order to assess the robustness of the inferred network, we performed bootstrap with 1000 resamples, each time resampling all E-genes with replacement, and inferring a network. Additionally, we ran a jackknife procedure, leaving out one signalling node at a time, and inferring the remaining network containing 58 regulators. We summarize the results in Figure 3, which aggregates the transitively closed adjacency matrices of the inferred networks.

For each edge, its bootstrap support is the proportion of inferred networks (out of 1000) in which the given edge is present. Analogously, the jackknife support of a given edge is the proportion of networks (out of 57) in which that particular edge is present. (The proportion is out of 57 rather than 59, since for each edge between a pair of nodes, there were precisely two jackknife runs where either of two nodes was left out.) Each cell in the adjacency heatmap corresponds to the average of the bootstrap and jackknife supports for that given edge. We used the average because we wanted both procedures to contribute equally to the consensus, and the bootstrap would otherwise dominate by virtue of having 1000 runs, whereas the jackknife is restricted by the number of S-nodes. Figure 4 shows the consensus network consisting of the edges whose average




Fig. 3. Adjacency matrix showing the average of bootstrap and jackknife supports for each edge in the graph. Parent nodes are topologically ordered (from root to leaf), and child nodes follow the same ordering. The entry in row i and column j records the average of the bootstrap and jackknife supports for the directed edge $S_i \rightarrow S_j$

support exceeds 50%. To aid visualization, the transitive reduction of the graph is displayed, meaning that we omit redundant edges implied by transitively closing the graph.

3.1 Validation of the static network

The transcription factors whose activity was measured by reporter assays were selected from the top (Sox4), middle (Smad4) and bottom (Maml2 via Rbpjk) of the inferred signalling hierarchy (Fig. 4). For each reporter assay, RNAi perturbation experiments were performed against regulators selected uniformly from the hierarchy. Each such experiment assessed the directed edge from the perturbed regulator (parent node) to the transcription factors measured by the reporter assay (child node). By comparing the validated effects (Fig. 5) with the predicted edges of the NEM (Fig. 3), we estimated precision and recall of the inferred signalling graph. In total, 53 NEM edge predictions were experimentally assessed. For the Smad4 and Sox4 reporter assays, the precision and recall is 75% and 62.5%, respectively. If including the Rbpjk reporter as a proxy for Maml2 activity, the overall precision and recall is 67.6% and 73.5%, respectively.

3.2 Signal progression during EMT

Principal component analysis of the time-course samples shows that the gradual progression of EMT during the time course was captured on a transcriptomic level (Fig. 6). The first principal component accounts for 68% of the variance in the data, and almost perfectly orders the samples chronologically. Exceptions are the samples taken at 60 h, which have scores on PC1 most similar to the scores of the samples taken at 36 h. The 60 h samples also have extreme scores with respect to PC3.

We used our method described in Section 2.3 to map the time-course data onto the inferred network. For a given time point $t \in \{0, 2, 6, 12, 24, 36, 48, 60, 72\}$, we computed the differential expression profile between t and 96 h, which represents the mesenchymal stage. Then, we estimated which states of the S-nodes are

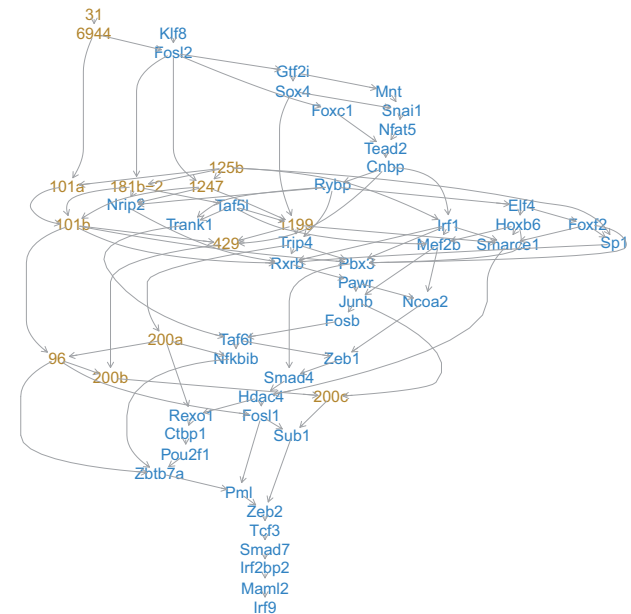


Fig. 4. Transitive reduction of the consensus network, which contains edges whose average of bootstrap and jackknife supports exceeds 50%. Nodes are topologically ordered so that all edges point downward. Transcription factors are coloured blue and miRNAs orange

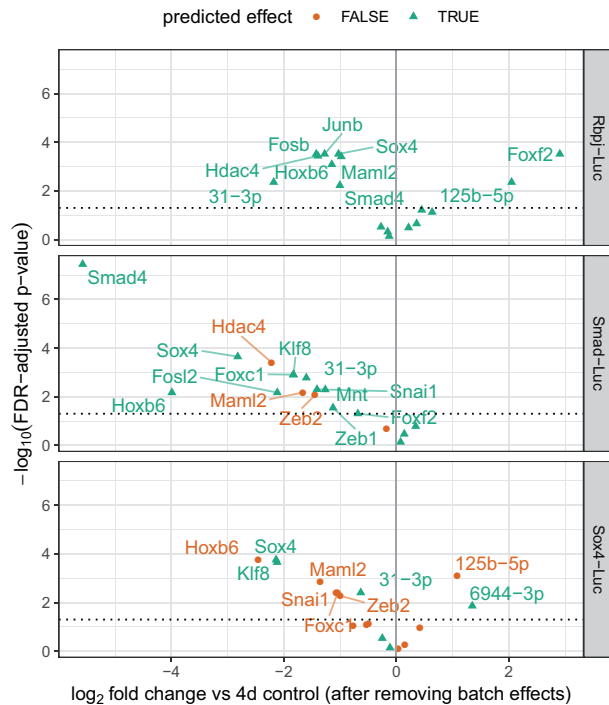


Fig. 5. Summary of luciferase reporter assays for Rbpj, Smad4 and Sox4. The volcano plots are based on Welch's t -tests of perturbation versus control, after removing batch effects using a linear model. Experiments above the dotted lines have FDR-adjusted P -values $< 5\%$, and are labelled. Predictions by the NEM consensus network (Fig. 4) are indicated by green triangles for a predicted effect, and orange circles for a prediction of no effect. Part of this validation dataset was previously published by Meyer-Schaller et al. (2019)

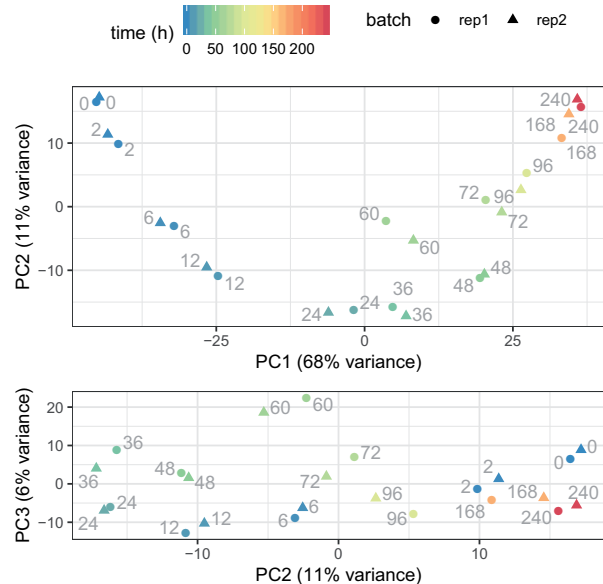


Fig. 6. Principal component analysis based on expression of 24 454 genes in biologically duplicated samples taken during an EMT time course. Top panel: PC1 versus PC2, bottom: PC2 versus PC3. The data was normalized using a regularized logarithmic transform (Love et al., 2014), and batch effects were removed using a linear model (Ritchie et al., 2015)

most likely to produce the differential expression profile which we observed. The procedure was performed on the inferred network $\hat{\Phi}^*$ and its MAP E-gene attachment $\hat{\Theta}^*$ for every bootstrap and jackknife resample. For each t , this yielded $1000 + 59$ estimates of K , which were aggregated by taking the component-wise (node-wise) average. The result is shown in Figure 7, where an S-node is uncoloured if it is inferred to be in a signalling state different to that at 96 h. Otherwise it is coloured green. Thus the colour of a given S-node indicates the time point at which its signalling activity is inferred to differ or agree with that at the mesenchymal state. The colour gradient was determined by running the integer linear programme on all networks inferred across bootstrap and jackknife resamples, followed by aggregating the results.

The progression to mesenchymal signalling activity starts with the S-nodes at the top of the hierarchy, as implied by enforcing the transitive closure in the integer linear programme. However, the same trend is observed even when not requiring the mapping to respect the transitive closure see Supplementary Material (Supplementary Fig. S1). As EMT progresses, the signal flows slowly from top to bottom of the network, except for the 60 h time point, and reaches the tail of the network by 72 h of TGF- β treatment. These results indicate that the NEM, although based on perturbation data from only one time point, can capture the temporal progression of signalling activities through the EMT system. Supplementary Figure S2 shows that signal progression would not be inferred by random chance, assessed by permuting gene labels before estimating signalling activity.

3.3 Computational complexity

The computational complexity of the module network algorithm for standard nested effect models has been addressed in previous articles (Fröhlich et al., 2008), so we will not study it here. In our case of inferring a NEM with 59 S-nodes and 11 270 E-genes, the average runtime was 12.4 CPU hours (SD 1.72) on an Intel® Xeon® Processor E5-2680 v3 (2.5–3.3 GHz). The average memory requirement was 36.3 GB (SD 3.37). In contrast, our method for mapping the EMT time-course data onto the inferred NEM takes less than a second. Thus the bottleneck is clearly the network inference.

In order to assess the computational complexity of our proposed method for inferring signalling activity in a NEM, we simulated synthetic NEMs with n S-nodes and m E-genes, where values of n and m were varied systematically. For each synthetic network, we randomly perturbed S-nodes, and generated synthetic E-gene data from a beta-uniform mixture model. Then, we ran our algorithm to map the effect profile onto the NEM, and recorded the elapsed run-time. To assess the complexity dependence on n , we fixed $m = 2000$ and varied $n = 10, 20, \dots, 100$. Each configuration was run 100 times each, and the resulting box plots are shown in the left panel of Figure 8. Since the run-time is plotted on a logarithmic scale, the linear trend (regression slope 0.042) suggests that the time complexity grows exponentially with the number of S-nodes. To estimate the complexity as a function of m , we fixed $n = 60$ and varied $m = 2000, 4000, \dots, 20000$. Each configuration was repeated 100 times, and the resulting box plots are shown in the right panel of Figure 8. Again the trend is roughly linear on the logarithmic scale (regression slope 5.8×10^{-5}), though not as steep as that for n . Thus we see a time complexity which grows exponentially with the number of E-genes.

4 Discussion

We extended the NEM framework to map an unperturbed EMT time series onto a steady state signalling network inferred from

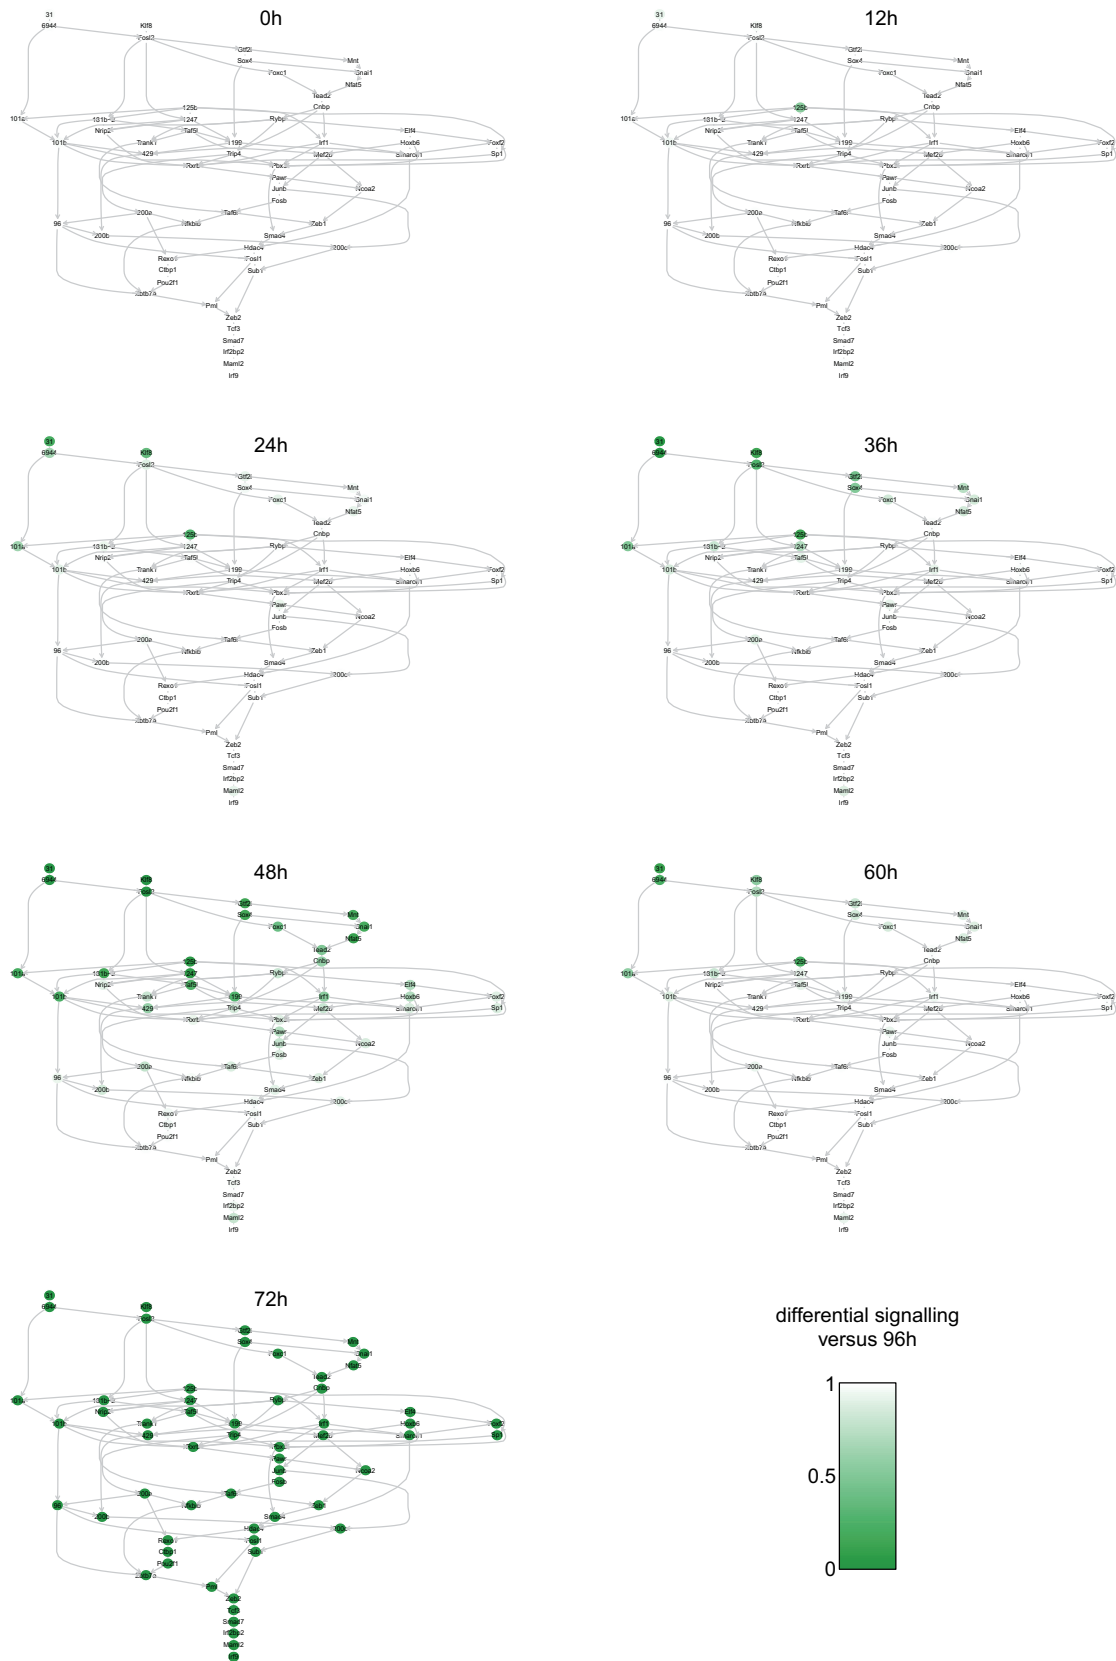


Fig. 7. Inferred signal progression at $t \in \{0, 12, 24, 36, 48, 60, 72\}$ hours after TGF- β treatment. An S-node is uncoloured if it is inferred to be in a different signalling state at t versus 96h; otherwise it is coloured green. The colour gradient comes from aggregating the inferred activity states over the networks resulting from 1000 bootstrap and 59 jackknife resamples. For illustration purposes, the signal progression is mapped onto the consensus network (Fig. 4)




Fig. 8. Computational complexity of our method for mapping an observational effect profile onto a static NEM. The CPU time in milliseconds is shown on a logarithmic scale on the y-axis. The number of S-nodes n varied from 10, 20, ..., 100 (left panel). The number of E-genes m varied from 2000, 4000, ..., 20 000 (right panel). Simulations were timed using the R package `microbenchmark` (Mersmann, 2018)

RNAi perturbation experiments. Based on the static network, our method estimates the progression of signalling activity of 59 transcription factors and miRNAs during EMT. The methodology integrates an unperturbed time series with perturbation data, which opens up new avenues in modelling signalling activities on large datasets where time-resolved perturbation experiments are infeasible.

EMT enables epithelial cells to become migratory and invasive, and requires a complex and coordinated flow of transcription (co)factor activities to reach a mesenchymal state. In the past, the acquisition of different EMT stages has been modelled using three- or four-factor systems, or based on activity predictions in the absence of perturbation data (Hong *et al.*, 2015; Steinway *et al.*, 2014; Zhang *et al.*, 2014). In contrast, in this study we made use of a large set of transcriptomic data from RNAi experiments of 59 transcription factors and miRNAs, which have been experimentally reported as essential EMT factors (Meyer-Schaller *et al.*, 2019). Using a NEM, we inferred a signalling network from the RNAi perturbation experiments. Extending the NEM framework, we mapped an unperturbed EMT time series onto the inferred network, in order to gauge the signal progression through the network during EMT. To validate the NEM network, we performed transcription factor reporter assays to measure TF signalling activities while perturbing upstream factors. These experiments yielded precision and recall estimates of 75% and 62.5%, respectively. Nevertheless, we were surprised that some important EMT factors, such as Smad4 or Zeb1 and Zeb2 (Lamouille *et al.*, 2014), appear rather low in the inferred signalling hierarchy, in the middle to the beginning of the ‘tail’ in Figure 4. The kinetic prediction of the signalling flow (Fig. 7) confirms an early activation of the top nodes in the network (up to 48 h TGF- β treatment) followed by a rapid activation of the bulk of the signalling network (72 h TGF- β treatment). Interestingly, the network flow reverts at the 60 h time point, coinciding with a similar divergence in the principal component analysis (Fig. 6), which otherwise captures the chronological EMT progression very accurately. TGF- β treatment of NMuMG cells results in TGF- β -induced cell-cycle arrest and induction of apoptosis before the cells convert to a mesenchymal state (Massagué, 2012). The partial reversion of TGF- β signalling after 60 h of treatment could reflect an important return point for the cells to overcome the apoptotic stimuli before continuing on the mesenchymal transition route.

Zeb transcription factors are engaged in a well characterized feedback loop with the miR-200 family (Wellner *et al.*, 2009). Since the NEM inference is constrained to acyclic graphs, we cannot infer feedback loops. However, the NEM infers directed edges into Zeb2

from the miR-200 family members miR-200a, miR-200b and miR-429, as well as an edge from Zeb1 to Zeb2, and the acyclic constraint forces one to be inferred upstream of the other. This may explain why Zeb1 is not inferred to be downstream of the miR-200 family, since this would otherwise create a feedback loop between Zeb1 and Zeb2. Due to the lack of time resolution in the perturbation data, feedback loops cannot be resolved, and we must consider the topology of the inferred network as a snapshot of the pathway at the time of sampling.

4.1 Limitations and possible extensions of the methodology

The NEM framework and our extension of it (Section 2.3) both make strong assumptions on the structure of the data generating mechanism. As discussed in the introduction, those assumptions are made because we have data only on the transcriptomic effects of perturbing signalling nodes, whereas we lack data on the activity of the signalling nodes themselves. However, assuming nested effects, we can infer signalling activities of the network nodes without information of whether this signalling is direct or mediated by unknown factors. The alternative—to model signalling nodes by the abundances of the mRNAs coding for them—would only work in systems that are transcriptionally regulated, that is, where regulators directly bind to each other’s promoter regions. Admittedly, the model assumption of nested effects may be too strict for our application since it does not accommodate feedback loops, which have been reported to exist between Zeb and the miR-200 families (Nieto *et al.*, 2016). However, feedback loops are notoriously difficult to model in structure learning, with only a few exceptions where the form of the feedback is described by linear structural equation models (Drton and Maathuis, 2017). This class of models is not applicable in our case, since the variables of interest are hidden.

An alternative approach to *de novo* network inference is to utilize prior knowledge of protein–DNA and protein–protein interactions (PPIs) (Yeang *et al.*, 2004). We did not pursue this approach for our EMT application in murine cells, because the current knowledge is too scarce. However, methods based on PPI networks have been developed with yeast as the model organism. For instance, Ourfali *et al.* (2007) utilize a directed PPI network to search for signalling pathways connecting a perturbed gene with an affected target. To utilize undirected PPI networks whose edges are weighted by confidence, Gitter *et al.* (2011) turn to the problem of orienting edges with the objective of maximizing the confidence in source-target paths of a bounded length. This method was also applied to incorporate time series (Gitter *et al.*, 2013). In a similar vein, Silverbush and Sharan (2014) aim to orient edges in an undirected PPI network while maximizing the number of directed source-target vertex pairs which are connected via shortest paths in the original graph. The above approaches all solve integer linear programmes (ILPs), though their objective is rather different from ours. Whereas those approaches solve ILPs to find paths or orient edges, our ILP takes the network inferred by the NEM as given, and instead aims to detect differential signalling based on a probabilistic measure of differential expression within regulatory modules.

In our analysis, we did not take into account the direction of differential expression; that is, whether the log-fold change was positive or negative. Similarly, we did not quantify whether relationships in the signalling network were activating or inhibiting. Rather, we considered the statistical significance of differential expression and the presence of interactions between regulators, regardless of their sign. This is in line with the standard NEM framework, but there exists an extension

called Factor Graph NEM (Vaske *et al.*, 2009) which natively models not only the presence of effects, but also their sign. This extension reformulates the NEM as a factor graph, and uses message passing to learn the optimal network. The authors reported that their extension performed better than the standard NEM in the presence of inhibitions in both the signalling graph and the effects. Since miRNAs tend to silence gene expression, and transcription factors can both activate and inhibit, it is possible that FG-NEM would yield different results for our data. However, in the absence of inhibition, the authors reported that FG-NEM, surprisingly, performs dramatically worse than the standard NEM. Furthermore, FG-NEM is designed only for knock-down experiments, and would therefore require some modification to appropriately model the over-expression of miRNAs, as were performed in our case.

Since the number of directed acyclic graphs grows super-exponentially with the number of nodes, the computational bottleneck is to infer the static NEM, which is a prerequisite for our method. A simulation study showed that the computational complexity of our method for mapping a time series onto a signalling network is negligible compared to inferring the static NEM.

Acknowledgements

We thank Jack Kuipers, David Seifert and Martin Pirkl for valuable suggestions. Furthermore, we thank Meera Saxena and Robert Ivanek for sharing the EMT time-course dataset with us.

Funding

This work was supported by SystemsX.ch [grant no. MRD 2014/267] to N.B.; and a Marie-Heim Vöglin grant from the Swiss National Science Foundation to N.M.-S.

Conflict of Interest: none declared.

References

- Anchang,B. *et al.* (2009) Modeling the temporal interplay of molecular signaling and gene expression by using dynamic nested effects models. *Proc. Natl. Acad. Sci. USA*, **106**, 6447–6452.
- Berkelaar,M. *et al.* (2015) *lpSolve: Interface to 'Lp_solve' v. 5.5 to Solve Linear/Integer Programs*. R Package Version 5.6.13.
- Brabletz,T. (2012) To differentiate or not—routes towards metastasis. *Nat. Rev. Cancer*, **12**, 425–436.
- Chaffer,C.L. *et al.* (2016) EMT, cell plasticity and metastasis. *Cancer Metastasis Rev.*, **35**, 645–654.
- Chen,S. *et al.* (2017) Network reconstruction from high-dimensional ordinary differential equations. *J. Am. Stat. Assoc.*, **112**, 1697–1707.
- Dennler,S. *et al.* (1998) Direct binding of smad3 and smad4 to critical TGF β -inducible elements in the promoter of human plasminogen activator inhibitor-type 1 gene. *EMBO J.*, **17**, 3091–3100.
- Drton,M. and Maathuis,M.H. (2017) Structure learning in graphical modeling. *Annu. Rev. Stat. Appl.*, **4**, 365–393.
- Friedman,N. (2004) Inferring cellular networks using probabilistic graphical models. *Science*, **303**, 799–805.
- Fröhlich,H. *et al.* (2018) *nem: (Dynamic) Nested Effects Models and Deterministic Effects Propagation Networks to Reconstruct Phenotypic Hierarchies*. R package version 2.54.0.
- Fröhlich,H. *et al.* (2008) Estimating large-scale signaling networks through nested effect models with intervention effects from microarray data. *Bioinformatics*, **24**, 2650–2656.
- Fröhlich,H. *et al.* (2011) Fast and efficient dynamic nested effects models. *Bioinformatics*, **27**, 238–244.
- Gitter,A. *et al.* (2011) Discovering pathways by orienting edges in protein interaction networks. *Nucleic Acids Res.*, **39**, e22–e22.
- Gitter,A. *et al.* (2013) Linking the signaling cascades and dynamic regulatory networks controlling stress responses. *Genome Res.*, **23**, 365–376.
- Hong,T. *et al.* (2015) An *ovol2-zeb1* mutual inhibitory circuit governs bidirectional and multi-step transition between epithelial and mesenchymal states. *PLoS Comput. Biol.*, **11**, 1–20.
- Lamouille,S. *et al.* (2014) Molecular mechanisms of epithelial–mesenchymal transition. *Nat. Rev. Mol. Cell Biol.*, **15**, 178–196.
- Le,T.D. *et al.* (2013) Inferring microRNA and transcription factor regulatory networks in heterogeneous data. *BMC Bioinformatics*, **14**, 1.
- Love,M.I. *et al.* (2014) Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2. *Genome Biol.*, **15**, 1.
- Markowetz,F. *et al.* (2005) Non-transcriptional pathway features reconstructed from secondary effects of RNA interference. *Bioinformatics*, **21**, 4026–4032.
- Markowetz,F. *et al.* (2007) Nested effects models for high-dimensional phenotyping screens. *Bioinformatics*, **23**, i305–i312.
- Massagué,J. (2012) TGF β signalling in context. *Nat. Rev. Mol. Cell Biol.*, **13**, 616–630.
- Mersmann,O. (2018) *microbenchmark: Accurate Timing Functions*. R package version 1.4-6.
- Meyer-Schaller,N. *et al.* (2019) A hierarchical regulatory landscape during the multiple stages of EMT. *Dev. Cell*, **48**, 539–553.e6.
- Nieto,M. *et al.* (2016) Emt: 2016. *Cell*, **166**, 21–45.
- Ourfali,O. *et al.* (2007) SPINE: a framework for signaling-regulatory pathway inference from cause-effect experiments. *Bioinformatics*, **23**, i359–i366.
- Pastushenko,I. *et al.* (2018) Identification of the tumour transition states occurring during EMT. *Nature*, **556**, 463–468.
- Pe'er,D. *et al.* (2001) Inferring subnetworks from perturbed expression profiles. *Bioinformatics*, **17**, S215–S224.
- Peters,J. *et al.* (2016) Causal inference by using invariant prediction: identification and confidence intervals. *J. R. Stat. Soc. B*, **78**, 947–1012.
- Pounds,S. and Morris,S.W. (2003) Estimating the occurrence of false positives and false negatives in microarray studies by approximating and partitioning the empirical distribution of p-values. *Bioinformatics*, **19**, 1236–1242.
- R Core Team (2018) *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Ritchie,M.E. *et al.* (2015) limma powers differential expression analyses for RNA-sequencing and microarray studies. *Nucleic Acids Res.*, **43**, e47.
- Segal,E. *et al.* (2003) Module networks: identifying regulatory modules and their condition-specific regulators from gene expression data. *Nat. Genet.*, **34**, 166–176.
- Silverbush,D. and Sharan,R. (2014) Network orientation via shortest paths. *Bioinformatics*, **30**, 1449–1455.
- Steinway,S.N. *et al.* (2014) Network modeling of TGF β signaling in hepatocellular carcinoma epithelial-to-mesenchymal transition reveals joint sonic hedgehog and wnt pathway activation. *Cancer Res*, **74**, 5963–5977.
- Tresch,A. and Markowetz,F. (2008) Structure learning in nested effects models. *Stat. Appl. Genet. Mol. Biol.*, **7**, 1–26.
- Vaske,C.J. *et al.* (2009) A factor graph nested effects model to identify networks from genetic perturbations. *PLoS Comput. Biol.*, **5**, 1–16.
- Wang,J. *et al.* (2010) Transmir: a transcription factor–microRNA regulation database. *Nucleic Acids Res.*, **38**, D119–D122.
- Wang,X. *et al.* (2014) Reconstructing evolving signalling networks by hidden Markov nested effects models. *Ann. Appl. Stat.*, **8**, 448–480.
- Wellner,U. *et al.* (2009) The EMT-activator ZEB1 promotes tumorigenicity by repressing stemness-inhibiting microRNAs. *Nat. Cell Biol.*, **11**, 1487–1495.
- Yeang,C.-H. *et al.* (2004) Physical network models. *J. Comput. Biol.*, **11**, 243–262.
- Zhang,J. *et al.* (2014) TGF- β -induced epithelial-to-mesenchymal transition proceeds through stepwise activation of multiple feedback loops. *Sci. Signal.*, **7**, ra91–ra91.