Perturbation-based gene regulatory network inference to unravel oncogenic mechanisms

Daniel Morgan[1], Matthew Studham[1], Andreas Tjärnberg[1,2], Holger Weishaupt[3], Fredrik J. Swartling[3], Torbjörn E. M. Nordling[4], Erik L.L. Sonnhammer[1]*

[1]Department of Biochemistry and Biophysics, Stockholm University, Science for Life Laboratory, Box 1031, 17121 Solna, Sweden. [2]Lionel Lab, Center for Developmental Genetics, New York University. [3]Rudbeck Laboratory, Department of Immunology, Genetics and Pathology, Science for Life Laboratory, Uppsala University, Uppsala, Sweden. [4]Department of Mechanical Engineering, National Cheng Kung University, Taiwan. * Corresponding author Erik.Sonnhammer@SciLifeLab.se

Abstract

The gene regulatory network (GRN) of human cells encodes mechanisms to ensure proper functioning. However, if this GRN is dysregulated, the cell may enter into a disease state such as cancer. Understanding the GRN as a system can therefore help identify novel mechanisms underlying disease, which can lead to new therapies. Reliable inference of GRNs is however still a major challenge in systems biology.

To deduce regulatory interactions relevant to cancer, we applied a recent computational inference framework to data from perturbation experiments in squamous carcinoma. GRNs were inferred using several methods, and the false discovery rate was controlled by the NestBoot framework. We developed a novel approach to assess the predictiveness of inferred GRNs against validation data, despite the lack of a gold standard. The best GRN was significantly more predictive than the null model, both in crossvalidated benchmarks and for an independent dataset of the same genes under a different perturbation design. It agrees with many known links, in addition to predicting a large number of novel interactions, a subset of which were experimentally validated. The inferred GRN captures regulatory interactions central to cancer-relevant processes and thus provides mechanistic insights that are useful for future cancer research.

Data available at GSE125958.
Inferred GRNs and inference statistics available at https://dcolin.shinyapps.io/CancerGRN/
Software available at https://bitbucket.org/sonnhammergrni/genespider/src/BFECV/
Keywords: cellular regulation, gene regulatory network inference, oncogenic mechanism

Author Summary

Cancer is the second most common cause of death globally, and although cancer treatments have improved in recent years, we need to understand how regulatory mechanisms are altered in cancer to combat the disease efficiently. By applying gene perturbations and inference of gene regulatory networks to 40 genes known or suspected to have a role in cancer, we deduce their underlying regulatory interactions. Using a recent computational framework for inference together with a novel method for cross validation, we infer a reliable regulatory model of this system in a completely data driven manner, not reliant on literature or priors. The novel interactions add to the understanding of the progressive oncogenic regulatory process and may provide new targets for therapy.

Introduction

Cancer can be seen as an altered state of the regulatory systems that control cell proliferation and cell death. Such systems are generally not sensitive to individual gene malfunctions, but an aggregation of aberrations can lead to sufficient dysregulation to cause cancer. Reliable models of these regulatory interactions would offer insight into key mechanistic alterations for therapeutic targeting. Cancer subtype-specific gene regulatory networks (GRN) encode intracellular dynamics [1], and offer understanding into the functional changes driving disease development. Inference of such models generally exploits certain aspects of the experimental setup, such as pooling among replicates to amplify signal, or makes use of prior knowledge [2,3]. The experimental techniques, setup, and data collection quality determine the quality of an inferred GRN model. However, practical limitations of experimentation, such as high noise levels and few experiments relative to the vast combinatorial landscape of possible regulatory interactions, often prevent any GRN inference methods from inferring a correct GRN[4]. Methods using data from known and systematic perturbations have shown greater accuracy among inference techniques since more information is available to determine regulatory causal mechanisms in the system[5].

GRN inference has proven its value to unravel novel regulatory links of biological significance. For instance, ARACNe was applied to gene expression profiles to predict a glioma-specific GRN, revealing that C/EBPbeta and STAT3 are master regulators of mesenchymal transformation, which was validated experimentally [6]. In another study, eight key genes were knocked down by siRNA, and the gene expression together with prior knowledge were used to infer a GRN network in the RAS pathway with good validation performance [7].

A large number of GRN inference algorithms exist. In a survey by the DREAM5 project [8] of 35 inference methods, it was shown on simulated and *E. coli* data that some methods performed better than random predictions. However, many methods did not outperform random, and on yeast data no method performed much better than random. Since this study, the community has convergently evolved methods of integrating various priors: literature/database, TF, ChIP-Seq, or protein data to increase information about the system, increasing accuracies by some 30%[9]. Similarly, in another benchmark ARACNe2 [10] performed only marginally better than random with an average AUROC of 0.58 across different sized networks, which is consistent with another benchmark study [11]. These benchmarks testify to the fact that network inference remains a very challenging task, especially when constrained to a single data modality. While GRN inference has been shown to be extremely valuable, it has also been shown that existing algorithms leave much room for improvement, especially in terms of limiting inclusion of false positives [12].

In this study, we deployed perturbations through siRNA gene knockdown of each gene in our literature-curated set of cancer-related regulator genes, each followed by transcriptomics measurements of all genes, in order to measure the global influence of each individual gene. Knockdown experiments are more informative about the system than irreversible and complete knockout which may cause drastic rewiring of the network into another system.

Assuming a linear dynamical system (LDS) [13] model, once the system has reached a steady-state equilibrium, a GRN can be inferred by solving a set of first order ordinary differential equations (ODEs)[14] in the form of our linear model (eq.1). Importantly, however, our linear model is reliant upon a known perturbation design, which additionally informs our model. In this way a selected set of 40 genes relevant to cancer were perturbed, and the transcriptomic response data were used to construct a model of underlying regulatory interactions. We inferred these interactions by relating the effect of the gene perturbations to the expression of the readout genes, using three GRN inference algorithms specifically tuned to our linear model and incorporating perturbation design: LASSO, LSCO, and TLSCO.

A drawback of all GRN inference algorithms is that they generally produce erroneous GRNs if the noise level is high [11,15]. To ensure inference of reliable GRNs, we employed NestBoot [12], a recent framework implementing nested bootstrapping, wrapped around any individual GRN inference method to better account for sample variation and noise. Contained within the GeneSPIDER package [11], NestBoot generates bootstrap support distributions for links inferred from measured as well as shuffled data[16], and minimizes false links by comparing them. This way Nestboot is able to discard links even if they have high bootstrap support, if they also have this in the null distribution. NestBoot has been shown to give substantially increased inference accuracy across both synthetic and experimental datasets when compared to the methods in their native implementation.

Lacking a gold standard, we first measured the inference accuracy by the GRN's ability to predict the data compared to a distribution of GRNs with the same topology as the inferred one but whose links have been shuffled. We complemented this performance evaluation by measuring the GRN's ability to predict the data compared to a distribution of shuffled data. Finally, we present the best performing GRN as the most reliable one, although the other inferred GRNs are largely subsets of each other and mostly perform well too. Two of the novel links of the best GRN were experimentally validated. The presented GRN captures regulatory interactions central to cancer-relevant processes and we foresee that it can provide mechanistic insights that can help to guide future cancer research. For instance, many cancers are caused by dysregulation of the MYC oncogene, hence our finding of a new regulator of MYC may potentially lead to new therapies.

Methods

Knockdown Data Collection

A set of genes was assembled from different pathways and complexes, each interacting to some degree with the oncogene MYC [17] (Table S2 and S3). Each readout gene is perturbed in the human squamous carcinoma cell line A431 via transfection with short interfering RNAs (siRNAs). We then harvest, purify and prepare libraries using the Ambion Library Construction Kit[18]. A precise record of perturbations is key to modeling (next section). In order to minimize siRNA off target effects, two to three siRNA are used per target

(Table S3), which are then averaged to purify the effects of the targeted siRNA perturbation. Cells were collected at 72 hours after siRNA knockdown and washed of Phosphate-Buffer Solution (PBS), and lysed using CelluLyser[19]. Cell counts were calculated using the resazurin fluorescence assay. Since no endogenous gene can be assumed to be free of MYC regulation, which is thought to be a universal transcriptional amplifier [20], a spike-in RNA transcript was added to each sample to act as a reference gene for the quantitative polymerase chain reaction (qPCR) analysis, added in proportion to the cell count before RNA isolation. It consisted of a 1000-base sequence with a 5' cap and a polyA tail. This was only used for normalization of mRNA level across samples[21]. Negative controls are included, which siRNA does not map to human genes, as well as an untreated control absent of any siRNA. The cDNA was prepared from the RNA and preamplified in preparation for the high-throughput qPCR screening. Finally, the transcript profiles with respect to the 40 genes were determined with TaqMan qPCR assays (Table S4) using Fluidigm Biomark 96x96 Dynamic Array integrated fluidic circuits. Raw qPCR output was then processed as a function of the samples' initial quantity returning a measure of fold change[22] via the experimental controls. Three experimental replicates were made per targeted perturbation and five outlying replicates were discarded due to clear machine read error, thus the dataset is composed of 40 genes (N) and 115 samples. Including all controls, a total of 18432 qPCRs were performed on 192 samples. Two technical replicates were performed to ensure minimal machine error. They generated very similar values up to 25 qPCR cycles (Fig. S1).

Experimental validation of individual interactions was performed on GTML2 brain tumor cells, which were cultured in serum-free stem cell medium as previously described [23] and treated for 2 hours with DMSO or JQ1 (500 nM). RNA was purified using the RNeasy Kit (Qiagen). RNA sequencing was performed using the Ion Proton™ System for Next-Generation Sequencing at NGI, SciLifeLab, Uppsala Biomedical Center (BMC), Sweden. All treatment conditions were performed in triplicates. All RNA sequence reads were processed and the differentially expressed genes were analyzed as previously described [23].

Model Generation

The fold change is calculated in comparison to the spike-in for all knockdown experiments. It is used in combination with the collective experimental design matrix (describing the location of perturbed and readout genes) to determine the GRN, i.e. the interaction matrix A, of instantaneous regulatory effects. We use a linear ODE model, similar to [13], which simplifies to

$$Y = -A^{\dagger}(P - F) + E \qquad (1)$$

where Y is an expression matrix of calculated fold changes, with N genes and M experiments. In Eq. 1, *P* is the design matrix if we solve for $A^{\dagger}$, where the Moore-Penrose generalized inverse, denoted *†*, is used throughout in place of the inverse due to computational intractability wherein sparse GRNs might be rank deficient. However as we want to solve for A and not $A^{\dagger}$, we reformulate Eq. 1 to a traditional regression problem on

errors-in-variables form, $-(P-F) = A(Y-E)$. The error in *Y* and *P* are represented as measurement error *E* and process error *F*, respectively, as defined in Table 1.

Table 1. Description of matrix contents used in the linear model for GRN inference.

matrix	**dimension**	**contents**	**description/ value range**
Y	115 x 40	Gene expression	SNR=0.014, max=3.3, min=-8.7, condition number=55.3
A	40 x 40	Interaction matrix (GRN)	Range dependent on BFE per init; initial inference returns variously sparse, weighted A, which is reweighted (see BFE section)
P	115 x 50	Perturbation/ design	Targeted perturbations (-1=KD, 0=no KD)
F	115 x 40	Input error	Range dependent on BFE per init
E	115 x 40	Output error	Range dependent on BFE per init

Three methods are employed based on their ability to mix model selection and parameter estimation. LSCO (least squares with a cut off to produce variably sparse networks)[25] for its value as a standard method, LASSO (least absolute shrinkage and selection operator)[26] was chosen for its proven ability to find the solution of minimum error, and TLSCO (total least squares with the same sparsity inducing cut off as LSCO,[27]) for its ability to model error in both the dependent and independent variables [15]. Each method is encapsulated within the nested bootstrapping framework to estimate the linear model in an accurate and reproducible manner by limiting false discovery rate (FDR), in their native configuration.

Model Selection & Validation

Only the measurement error can be estimated from the expression matrix Y, while the error in P cannot be estimated directly from data. This optimization occurs during the leave one out procedure (Algorithm 1), using the CVX convex optimization package (v1.22)[28] for MATLAB , where the left out gene (g) is expressed as a linear combination of the other experiments (crossvalidation) as follows:

Algorithm 1 : BalanceFitError (BFE) with Crossvalidation

1. For each perturbed gene *g*:
   a. Remove all experiments perturbing g from $Y, P \rightarrow Y_{!g}, P_{!g}$
   b. While $(F_{rel} - E_{rel} > t)$
      i. Minimize $\left\|E_{!g}\right\|_{Fro}^2 + \left\|F_{!g}\right\|_{Fro}^2$ subject to
      $$A\,(Y_{!g} - E_{!g}) = -(P_{!g} - F_{!g})$$
      $$\left\|E_{!g}\right\|_{Fro}^2 <= \frac{1}{\chi^2(\alpha,N)}$$
      $$\left\|F_{!g}\right\|_{Fro}^2 <= w\,\|A\,Y_{!g}\|_{Fro}^2$$
      ii. $E_{rel} = \frac{\|E_{!g}\|_{Fro}^2}{\|Y_{!g}\|_{Fro}^2}$ and $F_{rel} = \frac{\|F_{!g}\|_{Fro}^2}{\|P_{!g}\|_{Fro}^2}$
      iii. Use bisectioning to find the value of w that minimises $F_{rel} - E_{rel}$.
   c. $F_g = F_{!g}\,(P_{!g}^{\dagger}\,P_g)$
   d. $wRSSP_g = \frac{\|\widehat{P}_g - P_g\|_{Fro}^2}{cov(P_g)}$, where $\widehat{P}_g = P_g - F_g$
   e. $wRSSY_g = \frac{\|\widehat{Y}_g - Y_g\|_{Fro}^2}{cov(Y_g)}$, where $\widehat{Y}_g = -A^{\dagger}\,\widehat{P}_g$

2. $wRSS = \Sigma wRSSY_g + \Sigma wRSSP_g$

Here *A* contains the GRN structure (topology), with each non-zero value representing an interaction and each zero a lack of interaction. For each experiment in the dataset one creates new *Y* and *P* matrices by leaving out said gene's perturbations. The noise of Y, represented as *E,* is used as an estimate of the variance inherent to the cells' expression as well as error in plate reading[29]. In addition, since there is no true GRN to measure inference accuracy, part of this read error is distributed as a proxy for methodological error, balanced from E_g to process error, F_g. Since error is ultimately a function of the degrees of freedom of the given matrix, relative error ($E_{rel,}$ F_{rel}) is used to more equally balance these input and output errors. All norm calculations are of the frobenius variant. Each inference run returns a GRN for each parameter setting in the hyper parameter space. From step i. to iii. all matrices have the perturbation experiments of the left out gene *g* removed and are thus denoted *!g*. However, *A* maintains all genes, remaining square throughout, and later the left out experiments can be inferred from the remaining data. All inference methods used here have a regularization parameter that determines the number of nonzero parameters in the models, which is varied to span the complete range from empty to full network.

Because our leave out procedure assesses individual gene prediction errors, we assembled null GRN performance distributions by shuffling GRN links and fitting these new networks to the data to create both a standardized and fairly conservative link weight using a constrained least squares (CLS) algorithm [25,30]. To this end we implement a Monte Carlo sampling method, sampling links to maintain the node indegree and preserving hubs thereby approximating an estimated link null distribution based on the inferred GRN judged to generate conservative and fair null GRNs. For a fair comparison, both the inferred and shuffled GRNs are fit to the original data. However topology and sign are preserved in the GRNs. To obtain a measure of the goodness of fit of both inferred and shuffled GRNs, crossvalidation was used to calculate the weighted residual sum of squares (wRSS) of the original training data. We are able to predict a left out gene in step *c* and *d* by expressing it as a linear combination of the other genes. This goodness of fit measure was also made of the inferred GRN's ability to under crossvalidation predict the original data compared to the distribution of prediction errors using shuffled data. The relative error metric comparing measured and shuffled wRSS (Fig. 2, 3, S3, S4, S7, S8) is complemented by R^2 values (Fig. S6). Before calculating these, the condition number of each GRN was reduced to ensure that the predicted response remain similarly bounded as the observed gene expression. This was done by performing singular value decomposition of the GRN, setting singular values below a cutoff to zero, and then reconstructing the GRN without the smallest singular values. The GRN with reduced condition number was then fit to training data under crossvalidation to generate predicted expression responses in the same way as described above. The cutoff on the minimum singular value was set independently for each GRN to ensure that the predicted expression values were within the range of the measured values. The small singular values generally represent noise if the data is ill-conditioned and removing them reduces the effect of noise.

To further verify both accuracy and generalizability, these GRNs are also applied to a second independent, validation dataset based on the same genes knocked down in pairs, in single replicates. While this data is not used to infer GRNs, we apply the same cross-validation strategy as for the original data to validate the GRNs. This is necessary for parameter fitting, since the process error is different from the single knockdown data. Furthermore, by running the same pipeline we obtain a comparable measure of how well the independent data fits our inferred GRNs, and build null distributions of expected error from shuffled GRNs to examine an inferred GRNs' ability to predict the data.

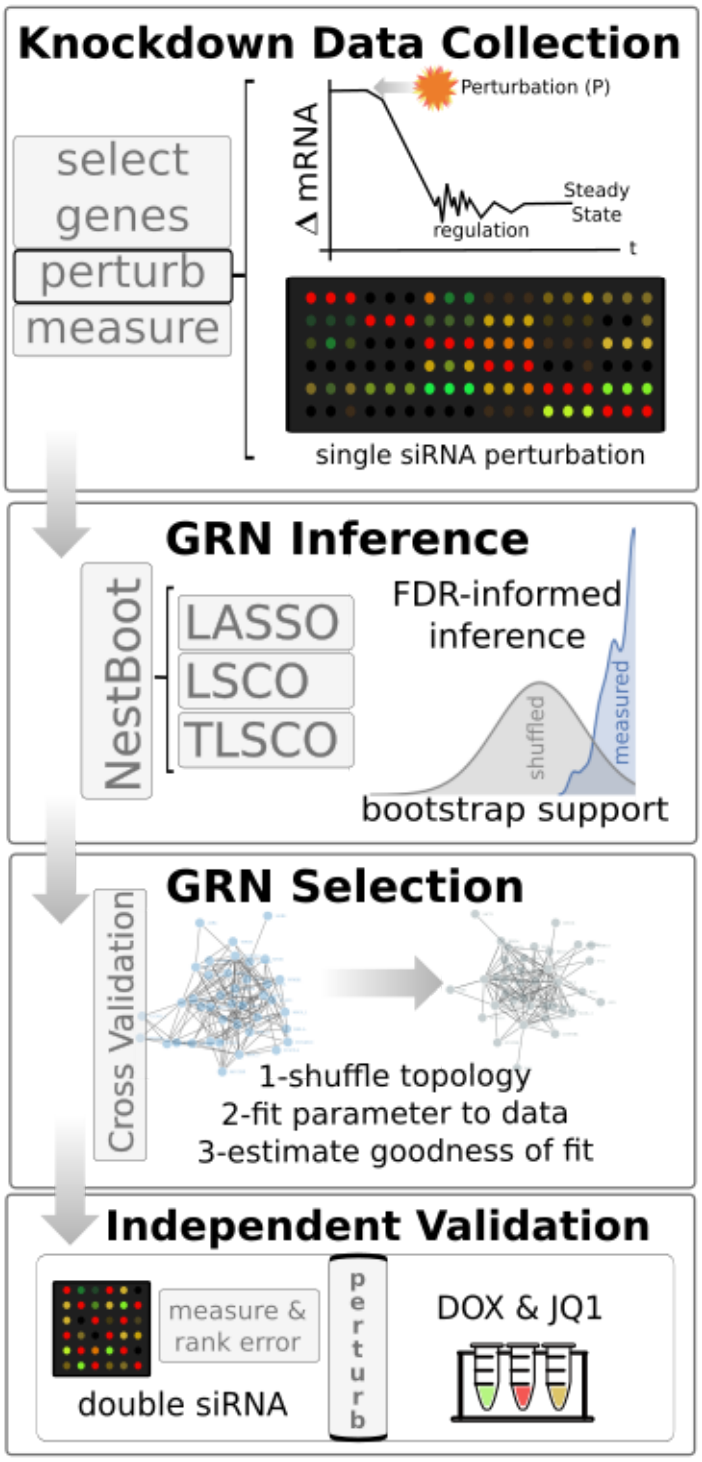

Fig. 1. Workflow of project. siRNA perturbation experiments were carried out systematically per gene, resulting in a change in mRNA level which elicits a regulatory response over time before reaching steady state when gene expression was measured.

GRN inference: Nested bootstrapping was applied to three inference methods, producing GRNs with an FDR set to 5%.

GRN Selection: Estimate goodness of fit under crossvalidation and error balancing, and compare to shuffled topologies (Algorithm 1).

Independent Validation: Each inferred GRNs' ability to predict an independent dataset was evaluated in comparison to a distribution of shuffled topologies. Finally, the overall most accurate GRN was selected. Two novel links were experimentally validated.

Results

A four step procedure was implemented to generate a cancer-centric GRN oriented towards the MYC oncogene (Fig. 1). First, a list of 340 cancer associated genes was ranked heuristically by what was known of each, giving preference to genes with known

associations to both cancer and MYC. The criteria were as follows, in decreasing order of rank: i. members of a complex with MYC affecting or not affecting transcripts, ii. genes directly affecting MYC or MYC transcripts (activating/repressing), iii. genes affecting MYC targeted transcripts, and iv. genes indirectly affecting MYC transcripts. Only genes expressed in the used cell line were further considered. The 40 top ranked genes were perturbed by siRNA in the well characterized human A431 squamous carcinoma cell line (Tables S2 and S3). Of the selected genes, 31 are transcriptional regulators, 10 are oncogenes, and 7 are tumor suppressors.

RNA silencing experiments were carried out to knock down the expression of each individual gene, whereafter the expression of all genes in response to the perturbation was measured. The experiments were carried out with three biological replicates per gene. At steady-state, gene expression was measured using high-throughput RT-qPCR, totalling 18432 quantifications. Most targeted genes were seen to have dramatic reduction in expression, generally a stronger effect than for the other genes (Fig. S2). In fact 31 targets were significantly downregulated ($p < 0.1$ and $\log_2$ fold change < -2).

The perturbation response gene expression data were used for network inference with three methods, LASSO, LSCO, and TLSCO, each run in conjunction with nested bootstrapping. Each GRN inference method was run with varying parameters to produce GRNs in a range of different sparsities. Nested bootstrapping was then used to select the significantly supported links in each GRN, resulting in final sparsities that tend to never exceed 4 links/gene even if the natively inferred GRN had 40 links/gene, for example. In order to select the GRN that has the links most likely to exist in reality, we compared how well each inferred GRN's topology fits the data compared to shuffled topologies of the same GRN. The reason for validating the topology rather than the complete GRN including the actual parameter values is that those parameters are optimized to fit the data and therefore no suitable null model exists. Note that by topology we mean the structure of the GRN, i.e. the inferred links and their sign. For each topology, the parameters were fit to the observed data using crossvalidation, and the error was measured as the difference between predicted and observed values of the hold-out samples, after assembling the individual gene predictions into the full predicted matrix via the iterative BFE. This was done with the novel BalanceFitError algorithm (Algorithm 1), which ensures that the error is balanced between input and output, i.e. that the error is not merely pushed from the output to the input or vice versa.

Each inferred GRN was shuffled a hundred times, and the data was fit the same way to these topologies to estimate a null distribution of expected inference error. Note that since the parameters of the shuffled topologies are fit to minimize the error, this is a very stringent test of the inferred topology and for a suboptimal GRN one expects some of its shuffled topologies to by chance have lower error. Yet, several of the inferred GRNs greatly outperformed their null model, both when using the original training data (Fig. 2) and the independent validation data with double knockdown design in the same cell line (Fig. S3). This is shown to be a fair proxy for accuracy, which link-calling based measure is not attainable in gold-standardless investigations such as this, in Fig S7 and S8. These figures demonstrate that when an inferred GRN outperforms its shuffled topologies, it tends to also

have high prediction accuracy. We also calculated R^2 values to show the proportion of the variation that our GRNs explain (Fig S6). All GRNs are available at https://dcolin.shinyapps.io/CancerGRN/. Five of the inferred topologies had 1000 times lower error than the median of the shuffled null model. The most accurate GRNs were inferred by Lasso, in terms of outperforming their null distributions. All but one GRN outperformed the median of their null distributions, and eight of the nineteen GRNs across all three methods outperformed all shuffled GRNs in their respective null distribution.

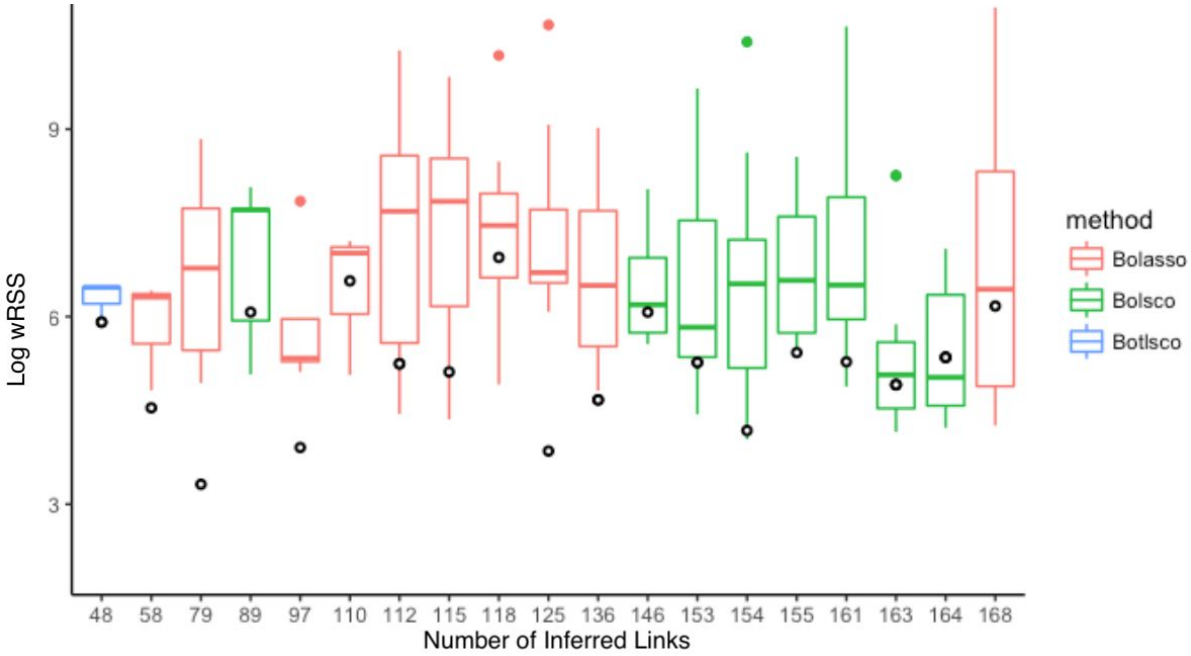

Fig. 2. Validation of inferred GRNs' topologies. Each x-axis tick mark shows the prediction performance in terms of the wRSS error of each inferred GRN topology (circles) fit to training data under crossvalidation, compared to its shuffled topologies. The box displays the median and interquartile range, and whiskers bound points maximally extending 1.5 times this range. Beyond this, outlier points are shown.

We also applied another null model to test how well the data fits the inferred GRN. Here we shuffled the original data one hundred times and fit these datasets to the inferred GRN in order to generate a null distribution. For many inferred GRNs the error was significantly lower than the median of the null distribution, both for the original training data (Fig. 3) and for independent validation data(Figs. 3 and S4).

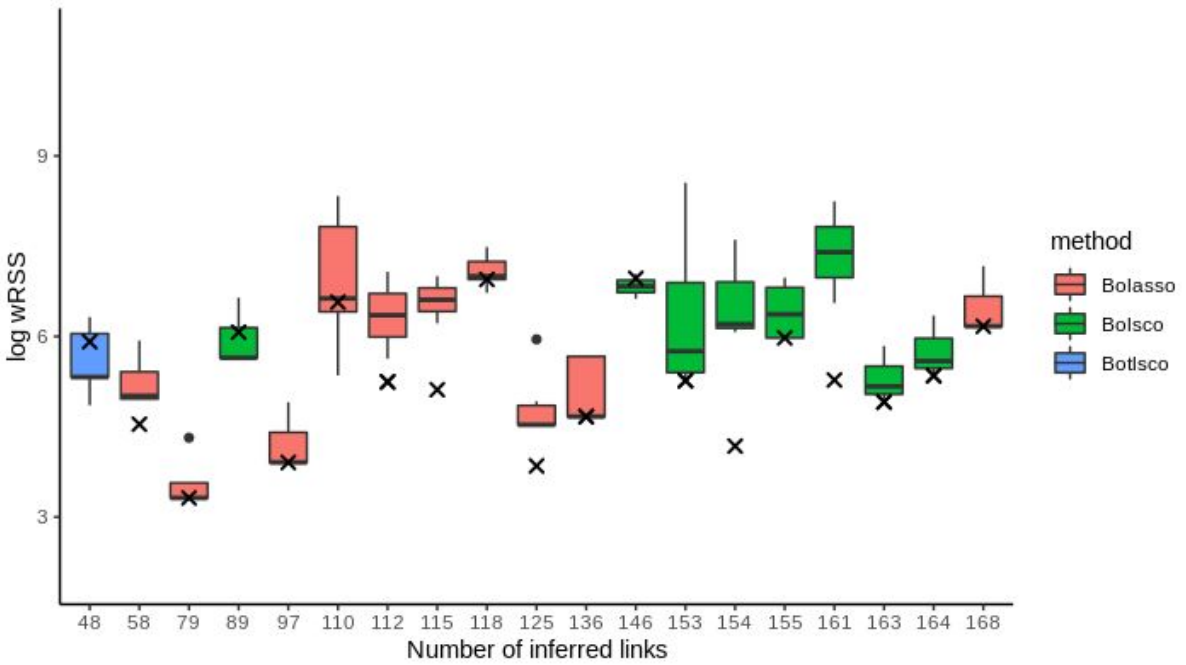

Fig. 3. Validation of the inferred GRNs' fit to the measured data. Each x-axis tick mark shows the prediction performance in terms of the wRSS of an inferred GRN topology fit to training data under crossvalidation, compared to its ability to fit shuffled data. X marks represent the inferred GRNs. The filled color box displays the median and interquartile range, and whiskers bound points maximally extending 1.5 times this range. Beyond this, outlier points are shown.

The GRN that outperformed its null distribution by the largest margin was Bolasso_network_L1145_M115_support97.5_1.52e-03, which we will refer to as the best GRN, with 125 links, including 39 self-links, between 39 genes (Fig. 4) and a sparsity of 3.2 links/gene. The full name indicates certain properties, namely that 1145 links were natively inferred before NestBoot, 115 experiments were used, 97.5% bootstrap support was attained at FDR=0.05, and 1.52e-03 is the sparsity penalty parameter used. In this GRN's overlap plot (Fig. 5) one can see the distributions of bootstrap values for both measured and shuffled data. The frequency of bootstrap support for measured data increases sharply at the right end above 98%, suggesting that this part of the distribution represents real and therefore highly reproducible links. In contrast, the shuffled data decreases towards support=1. The fact that some links inferred from shuffled data can attain such high bootstrap values can be attributed to the fact that the inference is done at a sparsity that yields very dense GRNs which may result in spurious links with high bootstrap support. However, the nested bootstrap framework monitors the distribution of spurious links and takes them into account when calculating FDR. The plot shows how FDR varies for different bootstrap support cutoffs.

One can also see the level of overlap between one hundred nested bootstrap runs in Fig. 5. Each run yields a bootstrap support for every link, which can be converted to a GRN for a given cutoff. For the measured data, the overlap (Jaccard) between runs stays relatively high (0.6) all the way to links with 100% bootstrap support, indicating that the reproducibility is high. In contrast, for the shuffled data not a single link with bootstrap support above 70%

overlaps with another nested run, indicating poor reproducibility despite relatively high bootstrap support.

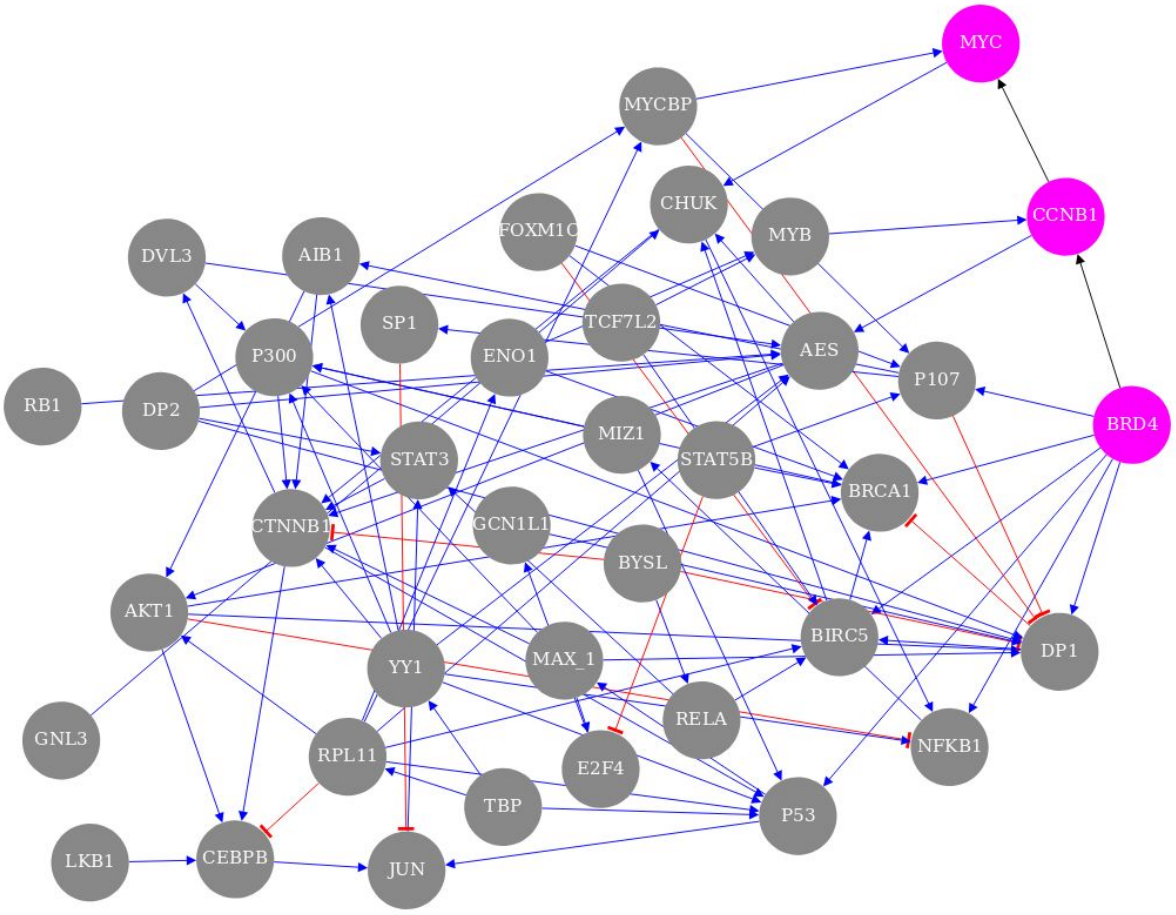

Fig. 4. The overall best performing GRN. Each of the 125 links has at least 97.5% bootstrap support, and the sparsity is 3.2 links/gene among its 39 genes. The 39 self links are not shown. The genes involved in the external validation experiment, BDR4, CCNB1, and MYC, are highlighted pink. Blue links reflect up regulation while red reflect negative down regulation. The visualization was made by the provided shiny app.

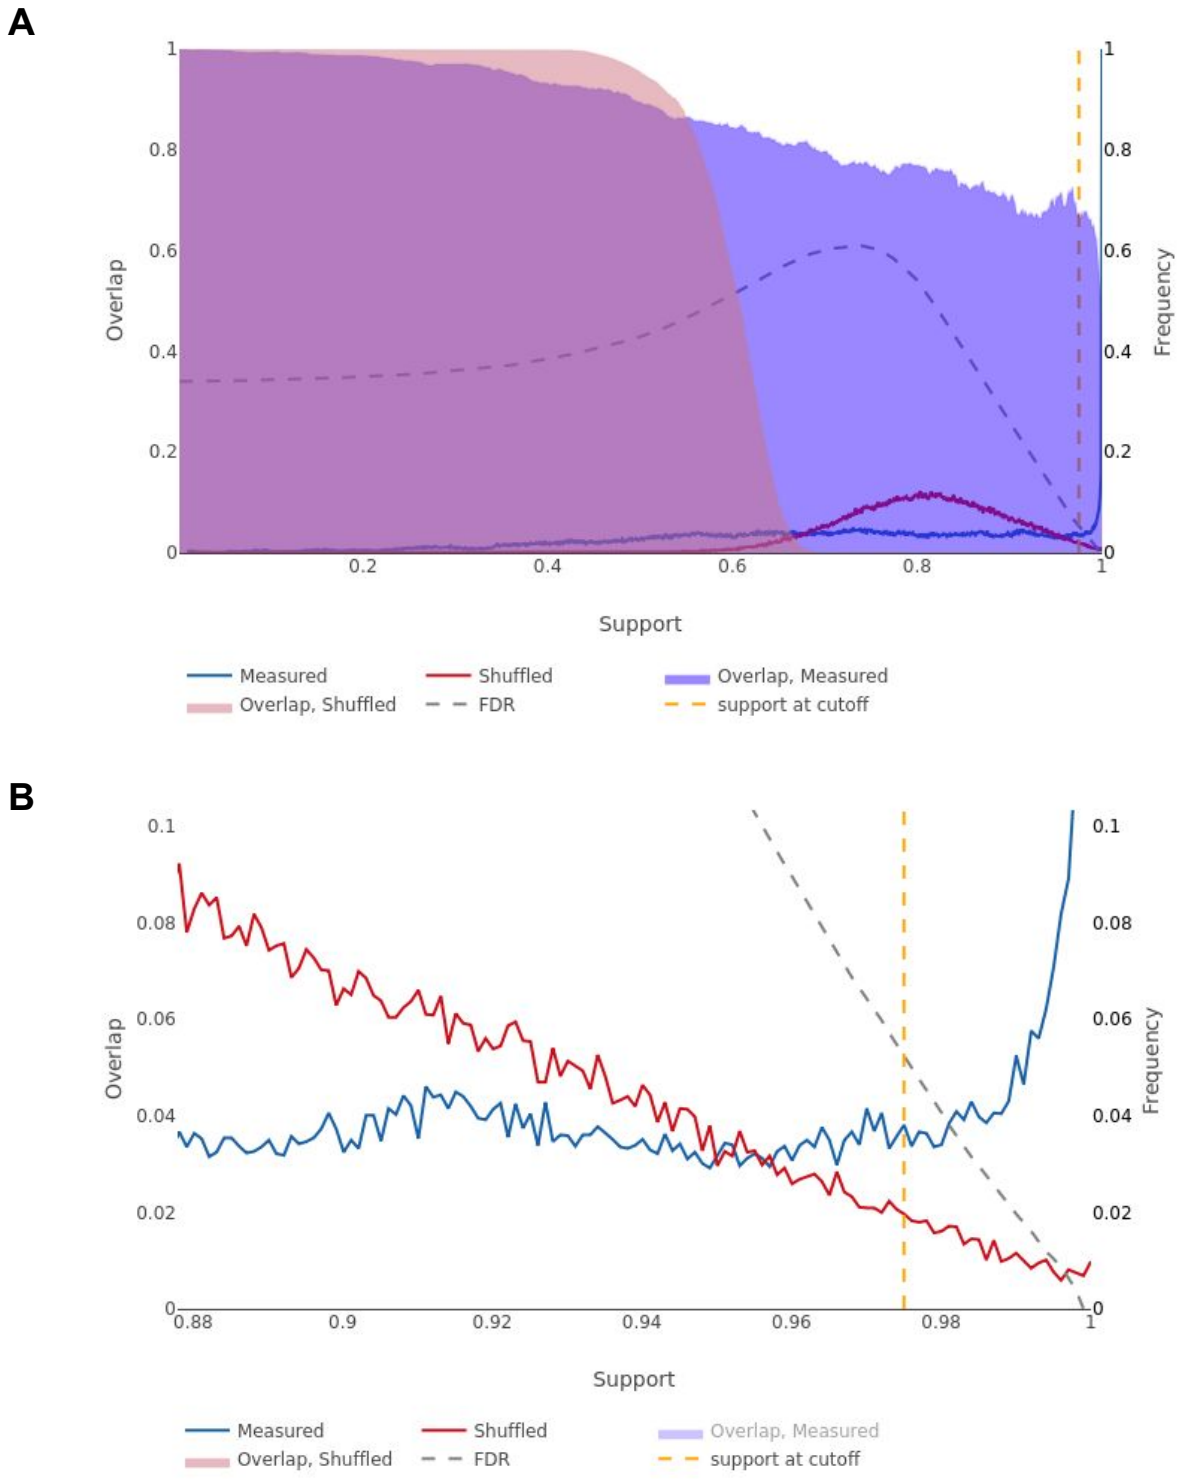

Fig. 5. NestBoot output for the overall best performing GRN. **A** shows the entire bootstrap support range from 0 to 1, as well as overlap between all bootstrap GRNs for measured (blue) and shuffled (red) data. The FDR is estimated via a null background model based on networks inferred from shuffled data. This is done to restrict inclusion of false links by setting FDR e.g. to 5%. The dashed orange line represents the cutoff where this is reached, here at 97.5% bootstrap support. The dashed grey line shows how the FDR behaves as a function of the bootstrap support. **B** shows the fine detail of the curves in the support levels between 0.9 and 1. These visualizations were made by the provided shiny app.

Validation of the best GRN

Of the 125 links inferred in the top performing GRN, two novel MYC-related links were experimentally validated. The novel regulatory relationships BRD4→CCNB1 and CCNB1→MYC (Fig. 3) were examined in an independent study [23] in which JQ1 was used to inhibit BRD4 in the GTML2 cell line, a mouse brain tumor cell line that overexpresses human MYCN. The inferred activation of CCNB1 by BRD4 was supported by a significant reduction of CCNB1 expression when BRD4 was inhibited, from 7.12 to 7.04 average log(CPM) after 6h (Fig. S5). However, in order to study immediate effects of BRD4 inhibition we here performed a new analysis in which GTML2 cells were treated with JQ1 for just 2 hours. Again, CCNB1 expression was significantly reduced, from 7.48 to 7.04 average log(CPM). Support for the inferred activation of MYC by CCNB1 was found by the fact that the CCNB1 expression changed from normal newborn mouse brain to adult mouse brain (from FPKM 20.3 to 0.2) which agrees with the change of MYC (from FPKM 9.2 to 0.8) [23].

These validation experiments support a novel mechanism for MYC regulation inferred in the best GRN. While it is well known that BRD4 can activate MYC in some cancer types ([31]), the best GRN presents a regulatory route that goes via CCNB1 (Cyclin B1). Bound with cyclin-dependent protein kinases, CCNB1 is involved in controlling the cell cycle at mitosis. The findings here suggest that CCNB1's role in regulating biological processes such as proliferation and oncogenesis can proceed via the activation of MYC.

Another type of validation is comparison to known links in public network resources. The links in the best GRN were searched for in the databases TRRUST [32], FunCoup [33], HumanNet [34] and STRING [35] as well as in our prior network from data mining. Where these reference networks contained undirected links, we compared them to an undirected version of our GRN. Many known interactions were witnessed in the best GRN, speaking to its ability to accurately infer what is known. The overlap with the other GRNs was significant ($p < 0.1$) in a hypergeometric test in all cases but one (Table S1).

Discussion

This study carries out a complete workflow for inferring reliable GRNs, from the selection of genes, experimental perturbation, data collection, and GRN inference, to validation of the inferred GRNs. It was applied to 40 cancer-related genes whose GRN was studied in a human squamous carcinoma cell line. The collected dataset was used to infer variously sparse GRN using three inference techniques within the NestBoot framework. The predictive accuracy of the inferred GRNs was estimated using the novel BalanceFitError algorithm under crossvalidation. Almost all inferred GRNs were more predictive than expected by chance, and some were vastly more predictive. These top performing GRN were also able to predict an independent pairwise-gene perturbation validation dataset significantly better than expected by chance. The best GRN contains many known links as well as proposes many novel links, two of which were verified experimentally.

The performed gene perturbations caused a range of fold changes in both targeted and readout genes (Fig. S2). Knockdown is advantageous for GRN inference compared to complete inhibition through knockout, as that could alter the gene functioning within the cells to such an extent as to potentially drive the cell to any number of non-native states by activating alternative pathways to cope with the loss of the knocked out one. This would result in measurement of an altogether different cellular GRN which lacks the knocked out gene. With knockdown, a gene's effect is lowered in the hope of measuring an otherwise wild-type GRN from the perspective of the single gene perturbation, across the gene repertoire.

The knockdown efficiency of each siRNA is unknown, and varies between genes. It may seem desirable to know the siRNA efficiency since this is a parameter in the perturbation design matrix that is used in the mathematical modelling. However, its value does not affect the inferred GRN's topology, and since the topology is the main outcome of the inference, and what we compare to null, this lack of knowledge is inconsequential. Prior information, whether literature-curated, ChIP-seq, ATAC-seq has been shown to be of value in modern GRN investigations, and may also be helpful. Such integration could be built into the model as a method of constraining spurious link additions much the same as NestBoot restricts links based on shuffled link distributions. We find many accuracy improvements in this sampling and anticipate similar accuracy improvements from the addition of priors. However, such experimental information is not available for this study and we therefore pursued a strictly data driven approach.
Despite our efforts to measure absolute mRNA levels using spiked-in RNA as qPCR reference, MYC was not found to be a universal amplifier as previously claimed [20,36,37]. Our observation agrees with the results of [38]. In both their and our study, measurements were done after 72 hours. It is possible that MYC knockdown activates a response leading to rapid restoration of MYC expression, so that cells return to their original state within that time span, instead of reaching a new steady state. In our study, the targeted MYC transcript was not significantly repressed by the MYC siRNA. This may be caused by its unusually high turnover rate [39], which can make it difficult to knock down with siRNA [40]. The same lack of observed knockdown for the target was noted for three other transcription factors: SP1, LMYC, and JUN. This seems to suggest a need for optimizing the experimental protocol to obtain perturbed steady state conditions when knocking down certain transcription factors.

During the NestBoot procedure the sparsity of the native GRNs is varied from almost a full to almost an empty network. However, as NestBoot selects the strongest supported links only, the sparsity of the GRNs output by NestBoot does not vary much for the denser GRNs, and even less in gene makeup. We observe consistency across different sparsities, i.e. the smaller GRNs are mostly a subset of the larger ones. This consistency among sparsities adds further confidence beyond the GRNs' predictiveness relative to a null distribution of shuffled topologies. Selecting the GRN with the optimal sparsity can be done in several ways. Here we followed the strategy of selecting the GRN with the best combination of coverage and predictiveness. Another criterion to select GRNs is the biological rationale that natural systems usually contain 3-5 links per gene[8].

In this study we face the problem of how to measure accuracy in the absence of a true network. Lacking such a gold standard it is impossible to determine if an inferred link is true or false. Instead, we compared each inferred network to a null distribution of GRNs with the same sparsity and indegree distribution. Since the prediction error depends on the weights of the links, it is crucial to fit each shuffled-link GRN to the data to give it reasonable weight estimates. To make the comparison fair, both the inferred GRN and the shuffled-link GRNs are refit to the data. By showing that the inferred GRN outperform their shuffled counterparts in terms of ability to explain the data, measured both by the wRSS and R^2, we know that they have a topology closer to the unknown real GRN. The exact same procedure can then be applied to other data, such as the independent validation dataset. With enough repeated shuffled-link GRNs to produce a sufficient null distribution, this results in an unbiased estimate of how predictive a given GRN is compared to what is expected, despite lacking a known gold standard network. Benchmarking on data with a known gold standard shows that increased predictiveness measured this way generally agrees with higher accuracy.

References

1. Shulman LP. Analysis of microarray experiments of gene expression profiling. Yearbook of Obstetrics, Gynecology and Women's Health. 2007;2007: 58–59.

2. Haury A-C, Mordelet F, Vera-Licona P, Vert J-P. TIGRESS: Trustful Inference of Gene REgulation using Stability Selection. BMC Syst Biol. 2012;6: 145.

3. Mordelet F, Vert J-P. SIRENE: supervised inference of regulatory networks. Bioinformatics. 2008;24: i76–82.

4. Guo S, Jiang Q, Chen L, Guo D. Gene regulatory network inference using PLS-based methods. BMC Bioinformatics. 2016;17: 545.

5. Tegnér J, Björkegren J. Perturbations to uncover gene networks. Trends in Genetics. 2007;23: 34–41.

6. Carro MS, Lim WK, Alvarez MJ, Bollo RJ, Zhao X, Snyder EY, et al. The transcriptional network for mesenchymal transformation of brain tumours. Nature. 2010;463: 318–325.

7. Olsen C, Fleming K, Prendergast N, Rubio R, Emmert-Streib F, Bontempi G, et al. Inference and validation of predictive gene networks from biomedical literature and gene expression data. Genomics. 2014;103: 329–336.

8. Marbach D, Costello JC, Küffner R, Vega NM, Prill RJ, Camacho DM, et al. Wisdom of crowds for robust gene network inference. Nat Methods. 2012;9: 796–804.

9. Castro DM, de Veaux NR, Miraldi ER, Bonneau R. Multi-study inference of regulatory networks for more accurate models of gene regulation. PLoS Comput Biol. 2019;15: e1006591.

10. Schaffter T, Marbach D, Floreano D. GeneNetWeaver: in silico benchmark generation

and performance profiling of network inference methods. Bioinformatics. 2011;27: 2263–2270.

11. Tjärnberg A, Morgan DC, Studham M, Nordling TEM, Sonnhammer ELL. GeneSPIDER - gene regulatory network inference benchmarking with controlled network and data properties. Mol Biosyst. 2017;13: 1304–1312.

12. Morgan D, Tjärnberg A, Nordling TEM, Sonnhammer ELL. A Generalized Framework for Controlling FDR in Gene Regulatory Network Inference. Bioinformatics. 2018; doi:10.1093/bioinformatics/bty764

13. Gardner TS, di Bernardo D, Lorenz D, Collins JJ. Inferring genetic networks and identifying compound mode of action via expression profiling. Science. 2003;301: 102–105.

14. Bonneau R, Reiss DJ, Shannon P, Facciotti M, Hood L, Baliga NS, et al. The Inferelator: an algorithm for learning parsimonious regulatory networks from systems-biology data sets de novo. Genome Biol. 2006;7: R36.

15. Tjärnberg A, Nordling TEM, Studham M, Nelander S, Sonnhammer ELL. Avoiding pitfalls in L1-regularised inference of gene networks. Mol Biosyst. 2015;11: 287–296.

16. Gobbi A, Iorio F, Dawson KJ, Wedge DC, Tamborero D, Alexandrov LB, et al. Fast randomization of large genomic datasets while preserving alteration counts. Bioinformatics. 2014;30: i617–23.

17. Hsieh AL, Walton ZE, Altman BJ, Stine ZE, Dang CV. MYC and metabolism on the path to cancer. Semin Cell Dev Biol. 2015;43: 11–21.

18. Ambion RNA-Seq Library Construction Kit [Internet]. ThermoFisher; 01/2017. Available: http://tools.thermofisher.com/content/sfs/manuals/4452440C.pdf

19. CelluLyser ™ Lysis and cDNA Synthesis Kit [Internet]. TATTA; 10/2012. Available: http://www.tataa.com/wp-content/uploads/2012/10/prodblad_v03_tataa-CelluLyser.pdf

20. Lovén J, Orlando DA, Sigova AA, Lin CY, Rahl PB, Burge CB, et al. Revisiting global gene expression analysis. Cell. 2012;151: 476–482.

21. Biocenter T. TATAA Universal RNA Spike I [Internet]. 2017 Oct. Available: https://webshop.tataa.com/dokument/Manual_TATAA%20Universal%20RNA%20Spike%20I%20SYBR-%20Probe_v1.3.pdf

22. Jitao David Zhang, Rudolf Biczok, and Markus Ruschhaupt. The ddCt Algorithm for the Analysis of Quantitative Real-Time PCR (qRT-PCR) [Internet]. 2017. Available: https://www.bioconductor.org/packages/release/bioc/html/ddCt.html

23. Bolin S, Borgenvik A, Persson CU, Sundström A, Qi J, Bradner JE, et al. Combined BET bromodomain and CDK2 inhibition in MYC-driven medulloblastoma. Oncogene. 2018;37: 2850–2862.

24. Nordling T. Robust inference of gene regulatory networks. Jacobsen E, editor. PhD, KTH Royal Institute of Technology. 2013.

25. Tjärnberg A, Nordling TEM, Studham M, Sonnhammer ELL. Optimal sparsity criteria for

network inference. J Comput Biol. 2013;20: 398–408.

26. Tibshirani R. Regression Shrinkage and Selection Via the Lasso [Internet]. Journal of the Royal Statistical Society: Series B (Methodological). 1996. pp. 267–288. doi:10.1111/j.2517-6161.1996.tb02080.x

27. de Groen PPN. An introduction to total least squares. 1998.

28. Michael Grant and Stephen Boyd. CVX: Matlab Software for Disciplined Convex Programming, version 2.1 [Internet]. mar, 2014. Available: \url{http://cvxr.com/cvx}

29. Cantone I, Marucci L, Iorio F, Ricci MA, Belcastro V, Bansal M, et al. A yeast synthetic network for in vivo assessment of reverse-engineering and modeling approaches. Cell. 2009;137: 172–181.

30. Chang LY, Pollard NS. Constrained least-squares optimization for robust estimation of center of rotation. J Biomech. 2007;40: 1392–1400.

31. Shi J, Wang Y, Zeng L, Wu Y, Deng J, Zhang Q, et al. Disrupting the interaction of BRD4 with diacetylated Twist suppresses tumorigenesis in basal-like breast cancer. Cancer Cell. 2014;25: 210–225.

32. Han H, Cho J-W, Lee S, Yun A, Kim H, Bae D, et al. TRRUST v2: an expanded reference database of human and mouse transcriptional regulatory interactions. Nucleic Acids Res. 2018;46: D380–D386.

33. Ogris C, Guala D, Sonnhammer ELL. FunCoup 4: new species, data, and visualization. Nucleic Acids Res. 2018;46: D601–D607.

34. Zeng X, Lin J, Lin C, Liu X, Rodriguez-Paton A. Structural Hole Spanner in HumanNet Identifies Disease Gene and Drug targets. IEEE Access. 2018; 1–1.

35. Szklarczyk D, Morris JH, Cook H, Kuhn M, Wyder S, Simonovic M, et al. The STRING database in 2017: quality-controlled protein-protein association networks, made broadly accessible. Nucleic Acids Res. 2017;45: D362–D368.

36. Lin CY, Lovén J, Rahl PB, Paranal RM, Burge CB, Bradner JE, et al. Transcriptional amplification in tumor cells with elevated c-Myc. Cell. 2012;151: 56–67.

37. Nie Z, Hu G, Wei G, Cui K, Yamane A, Resch W, et al. c-Myc is a universal amplifier of expressed genes in lymphocytes and embryonic stem cells. Cell. 2012;151: 68–79.

38. Nishiyama A, Sharov AA, Piao Y, Amano M, Amano T, Hoang HG, et al. Systematic repression of transcription factors reveals limited patterns of gene expression changes in ES cells. Sci Rep. 2013;3: 1390.

39. Jones TR, Cole MD. Rapid cytoplasmic turnover of c-myc mRNA: requirement of the 3' untranslated sequences. Mol Cell Biol. 1987;7: 4513–4521.

40. Larsson E, Sander C, Marks D. mRNA turnover rate limits siRNA and microRNA efficacy. Mol Syst Biol. 2010;6: 433.

41. Delmore JE, Issa GC, Lemieux ME, Rahl PB, Shi J, Jacobs HM, et al. BET bromodomain inhibition as a therapeutic strategy to target c-Myc. Cell. 2011;146:

904–917.

42. Basso K, Margolin AA, Stolovitzky G, Klein U, Dalla-Favera R, Califano A. Reverse engineering of regulatory networks in human B cells. Nat Genet. 2005;37: 382–390.

43. Peck B, Ferber EC, Schulze A. Antagonism between FOXO and MYC Regulates Cellular Powerhouse. Front Oncol. 2013;3: 96.

44. Zwolinska AK, Heagle Whiting A, Beekman C, Sedivy JM, Marine J-C. Suppression of Myc oncogenic activity by nucleostemin haploinsufficiency. Oncogene. 2012;31: 3311–3321.