

A POSTERIOR PROBABILITY APPROACH FOR GENE REGULATORY NETWORK INFERENCE IN GENETIC PERTURBATION DATA

WILLIAM CHAD YOUNG AND ADRIAN E. RAFTERY*

University of Washington
Department of Statistics
Box 354322
Seattle, WA 98195-4322, USA

KA YEE YEUNG

University of Washington
Institute of Technology
Box 358426
1900 Commerce Street
Tacoma, WA 98402-3100, USA

ABSTRACT. Inferring gene regulatory networks is an important problem in systems biology. However, these networks can be hard to infer from experimental data because of the inherent variability in biological data as well as the large number of genes involved. We propose a fast, simple method for inferring regulatory relationships between genes from knockdown experiments in the NIH LINC dataset by calculating posterior probabilities, incorporating prior information. We show that the method is able to find previously identified edges from TRANSFAC and JASPAR and discuss the merits and limitations of this approach.

1. Introduction. Gene regulatory networks are very important in understanding the biological functioning of cells. Identifying the interactions between genes can aid biologists in their attempts to understand how the cell functions both in steady state and in reaction to external stimuli. Unfortunately, due to the complexity of the cells and the large number of genes involved, discovering the true networks is very difficult. In most cases the number of genes measured far exceeds the number of observations, as is typical in microarray or sequencing experiments. Any method for analyzing such data must take this fact into account. Often this is done by enforcing a sparsity constraint, either via an added penalty on non-sparseness or via priors placed on the model. Even with these constraints, the ability to make valid inference is limited in these high-dimensional regimes as the number of genes grows compared to the number of observations [55, 54].

Many methods have been developed for inferring relationships between genes from gene expression data. One approach is to model the network holistically using

2010 *Mathematics Subject Classification.* Primary: 62P10, 92D10; Secondary: 92C42.

Key words and phrases. Bayesian analysis, gene regulatory network, statistics, statistical computation.

Research supported by NIH grant U54-HL127624-01.

* Corresponding author.

Bayesian networks [40, 21, 22, 62, 24, 48]. This yields good interpretable models, but often does not scale well and is hard to apply at the whole-genome level. Regression-based methods, where the expression level of a target gene is modeled as a function of the expression level of that gene's regulators, can be applied to much larger sets of genes but do not provide an overarching model of the entire network. Inference for these models generally becomes a statistical variable selection or model selection problem. Common methods for this include Significance Analysis of Microarrays [53], Least Absolute Shrinkage and Selection Operator (LASSO) [52, 17, 36] and Bayesian Model Averaging (BMA) [42, 19, 57, 29, 59]. Another class of methods looks at mutual information among the measured genes [2, 34, 13, 37].

When looking for regulatory relationships, it has been found that the information available from knockout experiments, where a single gene is fully suppressed, can be highly informative since they give a way to identify a causal pathway, direct or indirect. In the DREAM4 *in silico* network challenge [32, 33, 10], the winning method used only the data from the knockout experiments to infer the true networks, ignoring the time-series data entirely [41]. In real biological experiments, full knockout experiments are not possible for many essential genes, but knockdown experiments, where the target gene is partially suppressed, are often available. Methods for analyzing knockdown or knockout data include correlation-based approaches [45], implicit latent variable scoring [58], and Bayesian network scoring [14, 44, 15]. In addition, there has been some work in combining steady-state data with knockdown data to improve results [51, 5].

In this paper, we propose a simple, fast method for inferring gene regulatory relationships from knockdown data alone. Our method uses a simple linear regression model focusing on single regulator-target gene pairs based on knockdown data. This method allows the incorporation of prior knowledge about the relationships and generates posterior probabilities which can be used to form a ranked edgelist or as a part of a more expansive analysis.

2. Data. Our data come from the National Institute of Health (NIH) Library of Integrated Network-based Cellular Signatures (LINCS) program [28, 11]. The aim of this program is to generate genetic and molecular signatures of cells in response to various perturbations. One thrust of this program is the large-scale generation of gene expression profiles using L1000 technology. This technology has resulted in measurements from over one million experiments to date on over fifty human cell lines. These cell lines are populations of cells descended from an original source cell and having the same genetic makeup, kept alive by growing them in a culture separate from their original source.

Each of the L1000 experiments measures the expression levels of 1000 landmark genes in the human genome. These genes were selected specifically to cover as much of the gene expression variation in cellular expression as possible, since all 20,000+ genes cannot be measured. These experiments have measured cellular responses to more than 20,000 different chemical perturbagens. In addition, knockdown and over-expression experiments, where a single gene is targeted to control its expression level, have been performed on thousands of individual genes, both within and outside of the 1000 landmark genes.

The L1000 experiments were performed using Luminex Bead technology [12], in which color-coded microspheres were coded to attach to specific RNA sequences corresponding to a landmark gene and to fluoresce according to the level of that

gene's expression. Sets of beads for measuring the 1000 landmark genes were added to the solution for a single experiment along with the perturbing agent. The experiment was left for a specified period of time and then the gene expression levels were measured.

Experiments were done in sets on a single plate having individual wells for 384 experiments. This minimizes some external sources of error, such as environmental conditions, across these experiments. A small set of these experiments were used as controls with no perturbation. This gives a baseline distribution of expression level for each gene from which to measure deviations in other experiments. A common approach in this setup is to look at deviations in the perturbation experiments from the controls on the same plate, again recognizing that experiments on the same plate are likely to be more similar than those on different plates. Multiple plates, typically three or four, are prepared and analyzed together as a batch. These plates are prepared as technical replicates, with a given perturbation being prepared and then put into the same well of each plate. This gives additional power in removing systematic biases. Any given perturbation also is performed in multiple different batches, resulting in biological replicates since the sets of experiments were not prepared together.

3. Methods. We want to use the L1000 data to infer gene regulatory networks. This means that we need a method for inferring causality. One way to do this would be to use a causal time-series model, but the L1000 data include a very small number of time points (drug perturbation experiments include only one to two time points). Instead, we use knockdown experiments to identify a single gene as a putative causal agent. Although this limits the amount of information we can gain from a single experiment, it allows us to use a straightforward model with a clearly defined regulatory gene.

When looking at the knockdown experiments, it is important to understand that not all experiments are equally useful. The efficacy of the perturbation varies between target genes and even between experiments for the same target. The experimental setup of the LINCS data is helpful in identifying these differences. We use the control experiments on a plate to get an estimate of the normal variability of a gene. This eliminates some of the variability due to effects we cannot control or even measure, including differences in experimental conditions such as ambient temperature and the scientist performing the experiment, since these are captured in plate-level effects.

To take advantage of this aspect of the LINCS data, we calculate plate-level z -values for each gene in a knockdown experiment. To do this, we first calculate the baseline mean, \bar{x}_{hp} , and standard deviation, s_{hp} , for each gene h across all control experiments on plate p . Then the z -value for gene h in knockdown experiment i on plate p is

$$x_{hi}^* = \frac{x_{hi} - \bar{x}_{hp}}{s_{hp}}.$$

Once we have transformed the data in this way, we use a simple linear regression model to model the change in a target gene t as dependent on the change in the knockdown gene h :

$$x_{ti}^* = \beta_0 + \beta_1 x_{hi}^* + \epsilon_i, \quad i = 1, \dots, n_h \quad (1)$$

$$\epsilon_i \stackrel{\text{iid}}{\sim} N(0, \sigma^2). \quad (2)$$

Here, n_h is the number of available knockdown experiments for gene h . This model specifies a linear relationship between the z -score of the knockdown gene h and that of the target gene t . This is a simplification of the true process underlying the relationship between genes h and t , but it can still be effective for discovering relationships.

We estimate this model with a Bayesian approach using Zellner's g -prior [60] for the model parameters $(\beta_0, \beta_1, \sigma^2)$, namely

$$\begin{aligned}\beta_1 | \sigma^2, g &\sim N\left(0, \frac{g\sigma^2}{\sum_i x_{hi}^2}\right), \\ p(\beta_0, \sigma^2) &\propto 1/\sigma^2.\end{aligned}$$

The parameter g specifies the expected size of the regression parameter β_1 relative to the standard error of the OLS estimate of β_1 . The choice $g = 1$ indicates that the regression parameter is expected to be nearly indistinguishable from the noise, and thus gives a practical lower bound for g . Also, n_h/g is the effective number of data points in the prior, with $g = n_h$ corresponding to a unit information prior and giving similar results to BIC. We do not want a prior that has more spread than a unit information prior [43], and thus we choose g in the range $1 \leq g \leq n_h$. In this case, we used $g = \sqrt{n_h}$. We have found this to be a good compromise between the extremes. We found that when we estimated g using an Expectation-Maximization algorithm [7, 59], the estimated value was often close to $\sqrt{n_h}$.

The regression model with the g -prior allows us to quickly calculate the posterior probability that gene h regulates gene t [6]. We first calculate the ratio of the posterior probability that h regulates t given the data \mathbf{x} , $\Pr(h \rightarrow t | \mathbf{x})$, to the posterior probability that there is no regulatory relationship, \Pr_0 . Further, we can incorporate a prior probability of regulation, π_{ht} , which reflects prior information about a regulatory relationship between genes h and t . This gives us

$$\begin{aligned}T_{ht} &\equiv \frac{\Pr(h \rightarrow t | \mathbf{x})}{\Pr_0} \\ &= \frac{\pi_{ht}}{1 - \pi_{ht}} \exp\left[(n_h - 2) \log(1 + g)/2 - (n_h - 1) \log(1 + g(1 - R^2))/2\right],\end{aligned}$$

where R^2 is the coefficient of determination for the simple linear regression model (1). From this we can get the posterior probability that h regulates t , or posterior edge probability:

$$p_{ht} = \frac{\Pr(h \rightarrow t | \mathbf{x})}{\Pr(h \rightarrow t | \mathbf{x}) + \Pr_0} = \frac{T_{ht}}{1 + T_{ht}}.$$

We use this posterior edge probability to rank potential edges and find likely edges for further investigation.

Two advantages of this method are its speed and simplicity. To compute the z -scores, we first get plate-level means and standard deviations, which can be done in a single read through the baseline data by keeping track of sums and sums of squares. From there, standardization of the knockdown data is quick and we need only to calculate correlations between the knockdown gene and each other gene to get the posterior edge probabilities. Additionally, including external knowledge through the prior edge probability can provide a significant boost in accuracy and precision [59]. Finally, these posterior probabilities have a straightforward interpretation, namely the probability that a given regulator-target pair is a true relationship given the data.

4. Results. We computed posterior probabilities for edges on the LINCS data for cell line A375. Cell line A375 is a human skin melanoma cell line with over 100,000 experiments in the LINCS data. That includes approximately 15,000 knockdown experiments on landmark genes. This gives a good set of data to evaluate our method. We set the prior probability of any edge being present to 0.0005, reflecting the average expected number of regulators (parents) for each node determined by Guelzim et al. [16] for yeast and the assumption that transcription factors will regulate approximately the same proportion of target genes regardless of the total number of genes available.

To assess our results, we need a reference standard. In our case, we looked at the Enrichr website [4], which has collected numerous gene-set libraries, including some that list gene regulatory relationships. We used the TRANSFAC and JASPAR lists of edges; these list transcription factors as well as putative binding sites on other genes using a position weight matrix [56, 47]. This is not a comprehensive gold standard for assessment since these regulatory relationships are limited to well-studied transcription factors. However, assessment of gene networks in the mammalian systems is non-trivial due to incomplete knowledge.

The TRANSFAC and JASPAR (T&J) dataset includes 37 transcription factors that overlap with the LINCS landmark genes. Thus we limit our assessment to only those genes as potential regulators. For these 37 transcription factors, the T&J dataset has 4,303 regulation-target pairs where both the transcription factor and the target are among the landmark genes. For each transcription factor in T&J, there is a variable number of knockdown experiments in the L1000 data targeting that particular gene, generally between 10 and 20 per gene. These are used to calculate posterior probabilities for about 42,000 pairs of genes to be compared with the T&J reference dataset.

To further evaluate our method, we compared our results with results from applying Significance Analysis of Microarrays (SAM) [53] and mutual information methods to the data. SAM is an adaptable method for identifying significant changes in gene expression level while estimating the false discovery rate. It is widely used to evaluate microarray data and is implemented in the R package `samr`. Mutual information methods, based in information theory, have also been used extensively to identify relationships among genes. We used the `minet` package in R [38] to analyze our data with three different mutual information methods: Context Likelihood of Relatedness (CLR) [13], Algorithm for the Reconstruction of Accurate Cellular Networks (ARACNE) [34], and minimum redundancy - maximum relevance (MRMR) [9, 37]. These three methods offer differing approaches for identifying relationships between genes.

Each method produces a list of gene pairs along with some measure of the strength of their relationship. SAM returns p-values for each relationship, the mutual information methods produce weights indicating the strength of the relationship, and our method gives posterior probabilities. We can sort these to produce a ranked list and evaluate these lists against the reference dataset.

First, we looked at two-by-two tables from each method. For the posterior probability method, we used probability cutoffs of 0.5 and 0.95 to define found edges. SAM provides a list of relationships found to be significant. The mutual information methods do not define a particular cutoff for significance, and so all relationships returned with non-zero weight were included. These two-by-two tables are produced in Table 1.

To assess whether the lists and the T&J dataset are related, we also report approximate (non-Bayesian) p -values by using the probability of getting at least the number of true positives found using a binomial(n, p) distribution, where n is the number of pairs in the inferred list and p is the probability of selecting a true edge from the total number of possible edges. Note that these p -values are not the same as the posterior probability thresholds, and inference is based on the posterior probability thresholds. The p -values turn out to be equal to 0.02 for both thresholds, indicating that the posterior edge probabilities are related to the T&J results at conventional levels of significance. The competing methods are not able to accurately identify a small number of edges as true, returning many more than the posterior probability method.

TABLE 1. Assessment results showing 2x2 tables for cell line A375 using knockdown experiments for finding edges via posterior probability calculation and compared to approximately 4,200 edges from TRANSFAC and JASPAR across 37 transcription factors. PP indicates the posterior probability method at the given threshold. When looking at edges with posterior probability of 0.5 or greater (top left table), 41 of the 292 candidate edges are found in TRANSFAC and JASPAR, and 14 of the 76 candidate edges at a cutoff of 0.95 (top center table) are true edges. Binomial approximate p -values are 0.02 (left) and 0.02 (right). Lower p -values indicate better results. The competing methods return many more edges as true but are not as precise, resulting in higher p -values.

		PP 0.5		PP 0.95		SAM	
		Yes	No	Yes	No	Yes	No
T&J	Yes	41	4262	14	4289	1193	3110
	No	251	37566	62	37755	11151	26666
		p-value: 0.02		p-value: 0.02		p-value: 0.98	
		CLR		ARACNE		MRMR	
		Yes	No	Yes	No	Yes	No
T&J	Yes	1651	2652	34	4269	1530	2773
	No	14671	23146	910	36907	13533	24284
		p-value: 0.67		p-value: 1.00		p-value: 0.60	

We expect errors from the assessment results in the form of both false positives and false negatives due not only to limitations of our method, but also due to the nature of the data and the TRANSFAC and JASPAR reference standard by which we evaluate our edges. This is in part because the T&J reference standard is not specific for the given cell line A375. Also, false positives might arise because the expression levels of target genes change due to indirect effects. The path from the transcription factor to the target gene may go through intermediate genes. In fact, since only about 5% of the human genes are measured by the LINCS experiments, there are likely to be many genes in relevant pathways that are not measured. If we had measurements for all 22,000 genes, using a link removal procedure could be very useful [23, 41]. We also expect false negatives since the T&J dataset is not a set of confirmed regulatory relationships. Rather, it is informed from attributes of the transcription factor as well as the target gene. This means that many of the

true relationships as designated by the T&J dataset may not in fact reflect true interaction at the cellular level. In general, we do not expect a transcription factor to affect 10% of the possible targets, which is what the T&J dataset reports, so it is likely that the T&J data is overestimating the number of regulatory relationships.

Another way of looking at the results is via the precision-recall curve. Precision and recall are both calculated by truncating our ranked list of edges and looking only at those proposed edges. Precision is the proportion of the proposed edges which are true edges. Recall is the proportion of true edges which are in the proposed set. The precision-recall curve takes a ranked list of edges from a procedure and shows how the precision varies as more and more edges are included from that list. High precision at low recall indicates that the procedure is good at identifying true edges at the highest probability. This is important in many cases, particularly genetic studies, because it gives researchers good information on where to focus their efforts in subsequent studies.

Figure 1 shows the precision-recall curves generated by the different methods. We do see that the edges most highly ranked by posterior probability yield better results than expected from random guessing by a factor of 1.5 to 2. The precision declines as we add more edges, returning to hover near random guessing. The MRMR and ARACNE results fare worse than random guessing, and although CLR ranks a few true edges highly, it returns to random much faster than the posterior probability edges. The ranked list from SAM performs comparably to the posterior probability method, but it is unable to differentiate between the edges at the very top of its list, with 168 edges yielding the same lowest p -value. From a scientific point of view, it is important to have high precision among the edges ranked most highly, since there are limited resources for designing and executing experiments investigating particular edges more closely. Of course it would be preferable to see even better precision, but our previous discussion has shown why that may not be achievable with this dataset and standard.

5. Discussion. We have demonstrated a straightforward approach to inferring gene regulatory network edges from knockdown data. This approach is simple to apply to large datasets and includes the ability to incorporate prior information when available. This approach is able to find confirmable regulatory relationships between genes from the L1000 data. We showed that our method performs comparably to or better than popular approaches for identifying important regulatory relationships as found in the TRANSFAC & JASPAR evaluation dataset.

One key benefit of this approach is that it can be applied to extremely large datasets, requiring only one read through the data to compile all sufficient statistics for computing the posterior probabilities. There is no need to retain all the data after reading it and no iterative methods, such as Expectation-Maximization or Markov Chain Monte Carlo, are used. Methods which model the entire network [24, 48, 31, 46] may yield more comprehensive network results but are also generally restricted to a smaller set of genes due to computational constraints. Our approach is complementary to these other approaches in potentially narrowing down the set of edges for further investigation.

We have applied the method to knockdown data in order to identify causal regulatory relationships. This method can also be applied to over-expression data or even steady-state data, although for steady-state data the resulting edges would

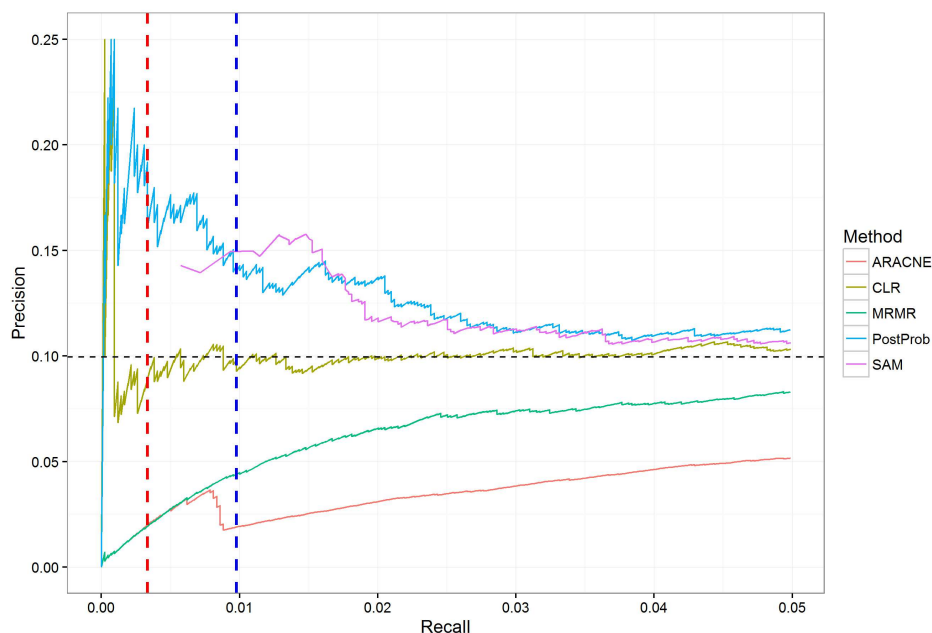


FIGURE 1. Precision-recall curves for cell line A375 using knock-down experiments for finding edges via posterior probability calculation evaluated against approximately 4,200 edges from TRANSFAC and JASPAR across 37 transcription factors. To the left of the red line are those edges with a posterior probability of at least 0.95, while the blue line shows the cutoff for edges with a posterior probability of 0.5. The horizontal dashed line shows the expected precision of 0.1 that would result from randomly ranking edges. The color version of this plot is available in the online version of the paper as well as on arXiv (<http://arxiv.org/abs/1603.04835>).

lack directionality [39]. This method could also be used to infer differential expression for a perturbation such as a drug treatment. This could be done using a 2-class model where the predictor variable indicates whether the expression measurements come from a perturbation experiment or a control experiment. An implementation of our method is available as an R package, `BayesKnockdown`, on Bioconductor, including functions for both knockdown and perturbation data.

We considered using an edge reduction technique, such as that used by Pinna et al. [41], but ultimately decided against it. Since we used only 37 genes as regulators due to our assessment data, the resulting networks did not tend to have multiple pathways from one gene to another. In cases where the resulting networks have more multi-gene pathways, using an edge reduction method may be appropriate.

Another possible use of this method is to use the resulting edge probabilities as an informative prior for another method using a different type of data. This could allow the integration of multiple data sources and increase the usefulness of knockdown data that provide only a small amount of evidence within a larger experimental context.

Acknowledgments. This research was supported by National Institutes of Health grants [R01 HD054511 and R01 HD070936 to A.E.R., U54 HL127624 to A.E.R. and K.Y.Y.]; Microsoft Azure for Research Award to K.Y.Y.; and Science Foundation Ireland ETS Walton visitor award 11/W.1/I207 to A.E.R.

REFERENCES

- [1] M. Bansal, G. Della Gatta and D. Di Bernardo, [Inference of gene regulatory networks and compound mode of action from time course gene expression profiles](#), *Bioinformatics*, **22** (2006), 815–822.
- [2] K. Basso, A. A. Margolin, G. Stolovitzky, U. Klein, R. D. Favera and A. Califano, [Reverse engineering of regulatory networks in human B cells](#), *Nature Genetics*, **37** (2005), 382–390.
- [3] P. Bühlmann, M. Kalisch and L. Meier, High-dimensional statistics with a view towards applications in biology, *Annual Review of Statistics and Its Application*, **1** (2014), 255–278.
- [4] E. Y. Chen, C. M. Tan, Y. Kou, Q. Duan, Z. Wang, G. V. Meirelles, N. R. Clark and A. Ma’ayan, [Enrichr: Interactive and collaborative HTML5 gene list enrichment analysis tool](#), *BMC Bioinformatics*, **14** (2013), 128–141.
- [5] S. Christley, Q. Nie and X. Xie, [Incorporating existing network information into gene network inference](#), *PLoS One*, **4** (2009), e6799.
- [6] M. Clyde and E. I. George, [Model uncertainty](#), *Statistical Science*, **19** (2004), 81–94.
- [7] A. P. Dempster, N. M. Laird and D. B. Rubin, Maximum likelihood from incomplete data via the EM algorithm, *Journal of the Royal Statistical Society. Series B (Methodological)*, **39** (1977), 1–38.
- [8] P. D’haeseleer, X. Wen, S. Fuhrman and R. Somogyi, [Linear modeling of mRNA expression levels during CNS development and injury](#), *Pacific Symposium on Biocomputing*, **4** (1999), 41–52.
- [9] C. Ding and H. Peng, [Minimum redundancy feature selection from microarray gene expression data](#), *Bioinformatics Conference, 2003. CSB 2003. Proceedings of the 2003 IEEE*, (2003), 523–528.
- [10] DREAM4 In Silico Network Challenge, website, <http://www.synapse.org/#!/Synapse:syn3049712/wiki/74628>
- [11] Q. Duan, C. Flynn, M. Niepel, M. Hafner, J. M. Muhlich, N. F. Fernandez, A. D. Rouillard, C. M. Tan, E. Y. Chen, T. R. Golub, P. K. Sorger, A. Subramanian and A. Ma’ayan, LINC Canvas Browser: Interactive web app to query, browse and interrogate LINC L1000 gene expression signatures, *Nucleic Acids Research*, **42** (2014), w449–w460.
- [12] S. A. Dunbar, Applications of Luminex[®] xMAP[™] technology for rapid, high-throughput multiplexed nucleic acid detection, *Clinica Chimica Acta*, **363** (2006), 71–82.
- [13] J. J. Faith, B. Hayete, J. T. Thaden, I. Mogno, J. Wierzbowski, G. Cottarel, S. Kasif, J. J. Collins and T. S. Gardner, [Large-scale mapping and validation of Escherichia coli transcriptional regulation from a compendium of expression profiles](#), *PLoS Biol*, **5** (2007), e8.
- [14] N. Friedman, M. Linial, I. Nahman and D. Pe’er, [Using bayesian networks to analyze expression data](#), *RECOMB ’00 Proceedings of the fourth annual international conference on Computational molecular biology*, (2000), 127–135.
- [15] H. Fröhlich, M. Fellmann, H. Suelmann, A. Poustka and T. Beissbarth, Large scale statistical inference of signaling pathways from RNAi and microarray data, *BMC Bioinformatics*, **8** (2007), 386.
- [16] N. Guelzim, S. Bottani, P. Bourguin and F. Képés, [Topological and causal structure of the yeast transcriptional regulatory network](#), *Nature Genetics*, **31** (2002), 60–63.
- [17] M. Gustafsson, M. Hörnquist, J. Lundström, J. Björkegren and J. Tegnér, [Reverse engineering of gene networks with LASSO and nonlinear basis functions](#), *Annals of the New York Academy of Sciences*, **1158** (2009), 265–275.
- [18] M. Hecker, S. Lambeck, S. Toepfer, E. Someren and R. Guthke, [Gene regulatory network inference: Data integration in dynamic models, A review](#), *Biosystems*, **96** (2009), 86–103.
- [19] J. A. Hoeting, D. Madigan, A. E. Raftery and C. T. Volinsky, [Bayesian model averaging: A tutorial](#), *Statistical Science*, **14** (1999), 382–417.
- [20] R. E. Kass and A. E. Raftery, [Bayes factors](#), *Journal of the American Statistical Association*, **90** (1995), 773–795.
- [21] S. Y. Kim, S. Imoto and S. Miyano, [Inferring gene networks from time series microarray data using dynamic Bayesian networks](#), *Briefings in Bioinformatics*, **4** (2003), 228–235.

- [22] S. Y. Kim, S. Imoto and S. Miyano, [Dynamic Bayesian network and nonparametric regression for nonlinear modeling of gene networks from time series gene expression data](#), *Computational Methods in Systems Biology*, **2602** (2003), 104–113.
- [23] S. Klamt, R. J. Flassig and K. Sundmacher, [TRANSWESD: inferring cellular networks with transitive reduction](#), *Bioinformatics*, **26** (2010), 2160–2168.
- [24] S. Lèbre, J. Becq, F. Devaus, M. Stumpf and G. Lelandais, Statistical inference of the time-varying structure of gene-regulation networks, *BMC Systems Biology*, **4** (2010), article 130.
- [25] W. Lee and W. Tzou, [Computational methods for discovering gene networks from expression data](#), *Briefings in Bioinformatics*, **10** (2009), 408–423.
- [26] J. Li and R. Tibshirani, [Finding consistent patterns: A nonparametric approach for identifying differential expression in RNA-Seq data](#), *Statistical Methods in Medical Research*, **22** (2013), 519–536.
- [27] Z. Li, P. Li, A. Krishnan and J. Liu, [Large-scale dynamic gene regulatory network inference combining differential equation models with local dynamic Bayesian network analysis](#), *Bioinformatics*, **27** (2011), 2686–2691.
- [28] Library of Integrated Network-based Cellular Signatures (LINCS), website, <http://lincsproject.org/>
- [29] K. Lo, A. E. Raftery, K. M. Dombeck, J. Zhu, E. E. Schadt, R. E. Bumgarner and K. Y. Yeung, [Integrating external biological knowledge in the construction of regulatory networks from time-series expression data](#), *BMC Systems Biology*, **6** (2012), p101.
- [30] F. M. Lopes, E. A. de Oliveira and R. M. Cesar, [Inference of gene regulatory networks from time series by Tsallis entropy](#), *BMC Systems Biology*, **5** (2011), p61.
- [31] M. J. McGeachie, H. Chang and S. T. Weiss, [CGBayesNets: Conditional Gaussian Bayesian network learning and inference with mixed discrete and continuous data](#), *PLoS Computational Biology*, **10** (2014), e1003676.
- [32] D. Marbach, T. Schaffter, C. Mattiussi and D. Floreano, Generating realistic in silico gene networks for performance assessment of reverse engineering methods, *Journal of Computational Biology*, **16** (2009), 229–239.
- [33] D. Marbach, R. J. Prill, T. Schaffter, C. Mattiussi, D. Floreano and G. Stolovitzky, [Revealing strengths and weaknesses of methods for gene network inference](#), *Proceedings of the National Academy of Sciences*, **107** (2010), 6286–6291.
- [34] A. A. Margolin, I. Nemenman, K. Basso, C. Wiggins, G. Stolovitzky, R. D. Favera and A. Califano, [ARACNE: an algorithm for the reconstruction of gene regulatory networks in a mammalian cellular context](#), *BMC Bioinformatics*, **7** (2006), S7.
- [35] F. Markowetz and R. Spang, [Inferring cellular networks: A review](#), *BMC Bioinformatics*, **8** (2007), S5.
- [36] P. Menéndez, Y. Kourmpetis, C. J. ter Braak and F. A. van Eeuwijk, Gene regulatory networks from multifactorial perturbations using Graphical Lasso: Application to the DREAM4 challenge, *PLoS One*, **5** (2010), e14147.
- [37] P. E. Meyer, K. Kontos, F. Lafitte and G. Bontempi, [Information-theoretic inference of large transcriptional regulatory networks](#), *EURASIP Journal on Bioinformatics and Systems Biology*, **2007** (2007), 79879.
- [38] P. E. Meyer, F. Lafitte and G. Bontempi, [minet: A R/Bioconductor package for inferring large transcriptional networks using mutual information](#), *BMC Bioinformatics*, **9** (2008), article 461.
- [39] G. Michailidis and F. d’Alché-Buc, [Autoregressive models for gene regulatory network inference: Sparsity, stability and causality issues](#), *Mathematical Biosciences*, **246** (2013), 326–334.
- [40] K. Murphy and S. Mian, *Modelling Gene Expression Data Using Dynamic Bayesian Networks*, Vol. 104. Technical report, Computer Science Division, University of California, Berkeley, CA, 1999.
- [41] A. Pinna, N. Soranzo and A. De La Fuente, [From knockouts to networks: Establishing direct cause-effect relationships through graph analysis](#), *PLoS One*, **5** (2010), e12912.
- [42] A. E. Raftery, D. Madigan and J. A. Hoeting, [Bayesian model averaging for linear regression models](#), *Journal of the American Statistical Association*, **92** (1997), 179–191.
- [43] A. E. Raftery, Bayes factors and BIC, *Sociological Methods & Research*, **27** (1999), 411–417.
- [44] S. Rogers and M. Girolami, [A Bayesian regression approach to the inference of regulatory networks from gene expression data](#), *Bioinformatics*, **21** (2005), 3131–3137.

- [45] F. H. M. Salleh, M. A. Arif, S. Zainudin and M. Firdaus-Raih, Reconstructing gene regulatory networks from knockout data using Gaussian Noise Model and Pearson Correlation Coefficient, *Computational Biology and Chemistry*, **59** (2015), 3–14.
- [46] M. Sanchez-Castillo, I. Tienda-Luna, D. Blanco, M. C. Carrion-Perez and Y. Huang, Bayesian sparse factor model for transcriptional regulatory networks inference, *Signal Processing Conference (EUSIPCO), 2013 Proceedings of the 21st European*, (2013), 1–4.
- [47] A. Sandelin, W. Alkema, P. Engström, W. W. Wasserman and B. Lenhard, JASPAR: an open-access database for eukaryotic transcription factor binding profiles, *Nucleic Acids Research*, **32** (2004), D91–D94.
- [48] M. Scutari, Learning Bayesian Networks with the bnlearn R Package, *Journal of Statistical Software*, **35** (2010), 1–22.
- [49] A. Shojaie and G. Michailidis, [Analysis of gene sets based on the underlying regulatory network](#), *Journal of Computational Biology*, **16** (2009), 407–426.
- [50] A. Shojaie and G. Michailidis, [Discovering graphical Granger causality using the truncating lasso penalty](#), *Bioinformatics*, **26** (2010), i517–i523.
- [51] A. Shojaie, A. Jauhiainen, M. Kallitsis and G. Michailidis, [Inferring regulatory networks by combining perturbation screens and steady state gene expression profiles](#), *PLoS One*, **9** (2014), e82393.
- [52] R. Tibshirani, Regression shrinkage and selection via the lasso, *Journal of the Royal Statistical Society. Series B (Methodological)*, **58** (1996), 267–288.
- [53] V. G. Tusher, R. Tibshirani and G. Chu, Significance analysis of microarrays applied to the ionizing radiation response, *Proceedings of the National Academy of Sciences*, **98** (2001), 5116–5121.
- [54] N. Verzelen, [Minimax risks for sparse regressions: Ultra-high dimensional phenomena](#), *Electronic Journal of Statistics*, **6** (2012), 38–90.
- [55] M. J. Wainwright, [Sharp thresholds for high-dimensional and noisy sparsity recovery using-constrained quadratic programming \(Lasso\)](#), *IEEE Transactions on Information Theory*, **55** (2009), 2183–2202.
- [56] E. Wingender, X. Chen, R. Hehl, H. Karas, I. Liebich, V. Matys, T. Meinhardt, M. Prüß, I. Reuter and F. Schacherer, [TRANSFAC: an integrated system for gene expression regulation](#), *Nucleic Acids Research*, **28** (2000), 316–319.
- [57] K. Y. Yeung, K. M. Dombeck, K. Lo, J. E. Mittler, J. Zhu, E. E. Schadt, R. E. Bumgarner and A. E. Raftery, [Construction of regulatory networks using expression time-series data of a genotyped population](#), *Proceedings of the National Academy of Sciences*, **108** (2011), 19436–19441.
- [58] C. Yoo, V. Thorsson and G. F. Cooper, [Discovery of causal relationships in a gene regulation pathway from a mixture of experimental and observational DNA microarray data](#), *Pacific Symposium on Biocomputing*, **7** (2002), 498–509.
- [59] W. C. Young, A. E. Raftery and K. Y. Yeung, [Fast Bayesian inference for gene regulatory networks using ScanBMA](#), *BMC Systems Biology*, **8** (2014), 47–57.
- [60] A. Zellner, On assessing prior distributions and Bayesian regression analysis with g -prior distributions, *Bayesian Inference and Decision Techniques: Essays in Honor of Bruno De Finetti*, **6** (1986), 233–243.
- [61] P. Zoppoli, S. Morganella and M. Ceccarelli, [TimeDelay-ARACNE: Reverse engineering of gene networks from time-course data by an information theoretic approach](#), *BMC Bioinformatics*, **11** (2010), p154.
- [62] M. Zou and S. D. Conzen, [A new dynamic Bayesian network \(DBN\) approach for identifying gene regulatory networks from time course microarray data](#), *Bioinformatics*, **21** (2005), 71–79.

Received September 29, 2015; Accepted May 16, 2016.

E-mail address: wmchad@uw.edu

E-mail address: raftery@uw.edu

E-mail address: kayee@uw.edu