

Learning LiNGAM based on data with more variables than observations

Shohei Shimizu*

Abstract

A very important topic in systems biology is developing statistical methods that automatically find causal relations in gene regulatory networks with no prior knowledge of causal connectivity. Many methods have been developed for time series data. However, discovery methods based on steady-state data are often necessary and preferable since obtaining time series data can be more expensive and/or infeasible for many biological systems. A conventional approach is causal Bayesian networks. However, estimation of Bayesian networks is ill-posed. In many cases it cannot uniquely identify the underlying causal network and only gives a large class of equivalent causal networks that cannot be distinguished between based on the data distribution. We propose a new discovery algorithm for uniquely identifying the underlying causal network of genes. To the best of our knowledge, the proposed method is the first algorithm for learning gene networks based on a fully identifiable causal model called LiNGAM. We here compare our algorithm with competing algorithms using artificially-generated data, although it is definitely better to test it based on real microarray gene expression data.

1 Introduction

An important topic in bioinformatics is developing computational methods to discover gene regulatory causal networks from static expression data [1–3]. Based on the estimated networks, one can compute intervention effects, *i.e.*, causal effects, which enable predicting to what extent the expression level of a gene changes when that of another gene is externally changed [4]. One can rank causal relations between genes according to the existence and/or strengths of causal effects. Such a ranking can be used as a priority list to efficiently conduct future interventional experiments and obtain solid evidence [5].

Many estimation techniques have been proposed for time series data [6–9]. Those techniques use temporal information to estimate the underlying gene network structure. However, it is not very feasible to obtain time series data for

*The Institute of Scientific and Industrial Research (ISIR), Osaka University, Mihogaoka 8-1, Ibaraki, Osaka 567-0047, Japan. Email: sshimizu@ar.sanken.osaka-u.ac.jp

many biological systems. In fact, many data that have been analyzed as time series data are not really longitudinal due to destructive sampling [10]. Then, discovery methods based on steady-state data could be better suited for such non-longitudinal ‘time series’ data.

A conventional approach for estimating gene causal networks based on steady state data is Bayesian networks [11, 12]. However, Bayesian networks suffers from the identifiability problem. In the framework of Bayesian networks, many networks with different structures give the same conditional independences between variables or genes, and in many cases one cannot uniquely estimate the underlying causal network without any prior knowledge on the structure [4, 13]. Thus, causal effects often are not uniquely estimated as well.

In this paper, we propose an attractive alternative approach that enables uniquely estimating a causal gene network. We first model the causal network of genes using a non-Gaussian causal model called LiNGAM [14]. LiNGAM is a fully identifiable causal model [14] unlike conventional Bayesian networks and has recently attracted much attention in machine learning [15, 16]. Then, we present a new algorithm for estimating LiNGAM based on data with more variables than observations. Finally, we test our approach using artificially-generated data.

2 Methods

We first define our model in Section 2.1 and then propose a new algorithm to estimate the model based on data with more variables than observations in Section 2.2.

2.1 Model

Let us denote by x_i the expression level of gene i ($i = 1, \dots, p$). We model causal relations of gene expression levels (variables) x_i ($i = 1, \dots, p$) using a linear non-Gaussian causal model called LiNGAM [14]:

$$x_i = \sum_{k(j) < k(i)} b_{ij} x_j + e_i, \quad (1)$$

where $k(i)$ is such a causal ordering of genes that they graphically form a directed acyclic graph (DAG). This means that no later gene directly or indirectly regulates any earlier gene, that is, has a directed path on any earlier gene. The e_i are the external influences or noises and b_{ij} are connection strengths of gene j on gene i . The zero/non-zero pattern of b_{ij} corresponds to the absence/existence pattern of directed edges. External influences e_i follow non-Gaussian continuous distributions with non-zero variances and are mutually independent.

The assumption that external influences e_i are non-Gaussian enables unique identification of a causal ordering $k(i)$ and connection strengths b_{ij} without using any background knowledge on the structure [14]. This feature is a big advantage [14–16] over conventional Bayesian network approaches based on

conditional independences and/or Gaussianity [17, 18]. Though the Gaussian approximation has been a common approach [19], real-world data could be considered more or less non-Gaussian. In fact, non-Gaussian data appear in many applications including bioinformatics [20, 21].

We note that each b_{ij} ($i \neq j$) represents the direct causal effect of x_j on x_i and each a_{ij} ($i \neq j$), the (i, j) -th element of the matrix $\mathbf{A} = (\mathbf{I} - \mathbf{B})^{-1}$, the total causal effect of x_j on x_i [4, 22]. Based on total causal effects a_{ij} , one can predict to what extent the expression level of gene i changes when that of gene j is externally changed. Further, based on direct causal effects b_{ij} , one can predict to what extent the expression level of gene i changes when that of gene j is externally changed while those of all other genes than gene i and gene j are held fixed.

Rigorously speaking, the linearity assumption would be more or less violated in real-world gene regulatory causal networks. Nonlinear approaches [23–25] might be better to model causal relations of genes. However, in general, linear methods can often give better results when it is more important to find quantitative relations since nonlinear methods usually require very large sample sizes [3].

2.2 Estimation of the model

Now, the problem of causal discovery is to estimate a causal ordering $k(i)$ and connection strengths b_{ij} based on data \mathbf{x} only. Several estimation methods for LiNGAM [14, 26] have been proposed that do not require to specify the distributions of the non-Gaussian external influences e_i . However, they are not applicable for such cases with more variables than observations that are typical in gene expression data analysis. Thus, we extend an LiNGAM estimation algorithm [26] to cases with more variables than observations.

In [26], a direct method was proposed to estimate a causal ordering $k(i)$. This leads to more algorithmically reliable results since it is not necessary to resort to iterative search in the parameter space. First, it estimates an exogenous variable. An exogenous variable is a variable with no parents and can be at the top of a causal ordering. One can estimate such a variable that is exogenous by finding a variable that minimizes a non-parametric independence measure called KGV [27] between the variable and its regression residuals [26]. Once an exogenous variable is found, then one subtract the effect of the exogenous variable from the other variables using linear least squares regression. One can find all the causal orders by iterating this [26]. Once a causal ordering $k(i)$ is estimated, one can prune or set to zero redundant connection strengths among b_{ij} by repeatedly applying a sparse regularization method called adaptive lasso [28] on each variable and its potential parents. The adaptive lasso [28] is a weighted version of a regularization technique for variable selection lasso [29] and assumes the same data generating process as LiNGAM:

$$x_i = \sum_{k(j) < k(i)} b_{ij} x_j + e_i.$$

The adaptive lasso assumes that the set of such potential parent variables x_j that $k(j) < k(i)$ is known, while LiNGAM estimates the set of such variables. The adaptive Lasso penalizes connection strengths b_{ij} in L_1 penalty by minimizing the objective function defined as:

$$\left\| x_i - \sum_{k(j) < k(i)} b_{ij} x_j \right\|^2 + \lambda \sum_{k(j) < k(i)} w_{ij} |b_{ij}|,$$

where λ is a regularization parameter and w_{ij} is a weight for b_{ij} . In [28], it was suggested to use the inverse of the absolute value of the ordinary least squares regression estimate or the ridge regression estimate of b_{ij} as w_{ij} . The adaptive lasso asymptotically selects the right set of such variables x_j that b_{ij} is not zero, where $k(j) < k(i)$.

To apply the direct method [26] on data with more variables than observations, we make three modifications. First, in cases with more variables than observations, the sample covariance matrix of variables is singular. Thus, we use ridge regression instead of linear least squares regression to compute regression residuals. Second, non-parametric independence measures, e.g., KGV [27] and HSIC [30], require large sample sizes and much computational time. They could give a similar performance as simple nonlinear correlation measures for small sample sizes [31]. Therefore, we use a nonlinear correlation [31, 32] to evaluate independence. We first evaluate pairwise independence between a variable and each of the residuals and then take the sum of the pairwise independence measures over the residuals. Third, before applying adaptive lasso, we reduce the dimension of data to at most $n-1$ using a dimension reduction method called iterative sure independence screening (ISIS) [33] so that the dimension is smaller than the sample size. This would be reasonable in estimation of gene networks since they are commonly assumed to be sparse. ISIS selects explanatory variables that have the first $n/\log(n)$ largest correlation coefficients with the explained variable in absolute value and apply lasso on the selected variables. It repeats this procedure on the residual of the explained variable over the selected explanatory variables until a desired number of variables, here $n-1$ variables, are selected. We then apply lasso on the selected variables by ISIS and finally adaptive lasso on the variables still surviving. The combination of ISIS, lasso and adaptive lasso was suggested by [33] to accurately identify non-zero coefficients when the dimension is larger than the sample size. We select the regularization parameter for each method using BIC based on Gaussianity [34, 35]. Though our model assumes non-Gaussianity, we optimistically assume that the effect of misspecification of the model distribution might not be so serious since the lasso and adaptive lasso estimation only involves means and covariances of observed variables. Moreover, we can estimate total causal effects a_{ij} by applying ISIS, lasso and adaptive lasso of gene i on the set of gene j and the parents of gene j according to the back-door criterion [36].

We now present our algorithm to estimate the LiNGAM in Equation (1) from data with more variables (or genes) than observations:

Input: Data matrix \mathbf{X}

1. Given a p -dimensional random vector \mathbf{x} , a set of its variable subscripts U and a $p \times n$ data matrix of the random vector as \mathbf{X} , initialize an ordered list of variables $K := \emptyset$ and $m := 1$.
2. Repeat until $p-1$ subscripts are appended to K :
 - (a) Denote by \mathbf{x}_K a vector that collects all the variables in K . Perform ridge regression of x_j on \mathbf{x}_K with the ridge parameter τ and compute the residual \tilde{x}_j for all $j \in U \setminus K$. Compute the matrix $\tilde{\mathbf{X}}$ that collects the values of \tilde{x}_j ($j \in U \setminus K$) from the data matrix \mathbf{X} .
 - (b) Perform ridge regression of x_i on $[x_j, \mathbf{x}_K^T]^T$ with the ridge parameter τ for all $i \in U \setminus K$ ($i \neq j$) and compute the residual vectors $\mathbf{r}^{(j)} = [r_i^{(j)}]$ and the residual data matrix $\mathbf{R}^{(j)}$ from the matrix \mathbf{X} for all $j \in U \setminus K$.
 - (c) Find a variable \tilde{x}_m that is most independent of the residuals $r_i^{(j)}$:

$$\tilde{x}_m = \arg \min_{j \in U \setminus K} \sum_{i \in U \setminus K, i \neq j} |\text{corr}\{g(\tilde{x}_j), r_i^{(j)}\}| + |\text{corr}\{\tilde{x}_j, g(r_i^{(j)})\}|, \quad (2)$$

where $g(\cdot)$ is $\tanh(\cdot)$.

- (d) Append m to the end of K .
3. Append the remaining variable subscript to the end of K .
4. Estimate connection strengths or direct causal effects b_{ij} by doing adaptive lasso of x_i on all the variables with earlier causal orders than x_i based on the data matrix \mathbf{X} for all $i \in U$. In case of more parent candidates than observations, the dimension is reduced to at most $n-1$ by ISIS and lasso before applying adaptive lasso. The weights w_{ij} for adaptive lasso are estimated by the inverses of the absolute values of ridge regression coefficients with the ridge parameter τ .
5. Estimate total causal effects a_{ij} by doing adaptive lasso of x_i on the set of x_j and its parent variables based on the data matrix \mathbf{X} for all $i, j \in U$ ($i \neq j$). In case of more parent candidates than observations, the dimension is reduced to at most $n-1$ by ISIS and lasso before applying adaptive lasso. The weights w_{ij} for adaptive lasso are estimated by the inverses of the absolute values of ridge regression coefficients with the ridge parameter τ .

Output: A causal network given by the zero/non-zero pattern of \mathbf{B} , direct causal effects b_{ij} and total causal effects a_{ij} between observed variables x_i and x_j ($i, j = 1, \dots, p, i \neq j$).

3 Experiments on artificial data

As a sanity check of our method, we performed an experiment with synthetic data. The experiment consisted of 1001 trials. In each trial, we generated a dataset with dimension $p = 100$ and sample size $n = 30$ and applied our estimation method on the data. For comparison, we also tested three methods: 1) random guessing, 2) lasso [29] and 3) elastic net [37]. Regarding lasso and elastic net, we applied lasso or elastic net on every variable taking the other variables as its explanatory variables and considered that explanatory variables with non-zero regression coefficients have non-zero direct causal effects and non-zero total causal effects on the explained variable as in [5] and [38]. Note that the lasso- and elastic net-based approaches do not give a DAG network. We used the coordinate descent algorithm in [39] of Matlab statistics toolbox to perform lasso and elastic net. Regarding elastic net, the weights of the lasso penalty and ridge penalty were set to be equal. All the regularization parameters were selected using BIC based on Gaussianity. The ridge parameter τ for ridge regression was set to 0.01. Regarding random guessing, we first generated a random ordering of observed variables and then randomly created as many non-zero direct causal effects and non-zero total causal effects as LiNGAM found while keeping the network being acyclic.

Each dataset was created as follows:

1. We constructed the $p \times p$ connection strength matrix with all zeros and replaced every element in the lower-triangular part by independent realizations of Bernoulli random variables with success probability s similarly to [40]. The probability s determines the sparseness of the model. The expected number of adjacent variables of each observed variable is given by $s(p-1)$. We randomly set the sparseness s so that the expected number of adjacent variables was 2 or 5.
2. We replaced each non-zero entry in the connection strength matrix by a value randomly chosen from the interval $[-1.5, -0.5] \cup [0.5, 1.5]$ and selected variances of the external influences from the interval $[1, 3]$. The resulting matrix was used as the data-generating connection strength matrix \mathbf{B} .
3. We generated data with sample size n by independently drawing the external influence variables e_i . We randomly selected the distribution of each e_i from a multimodal asymmetric mixture of two Gaussians, a multimodal symmetric mixture of two Gaussians and a Laplace distributions used in [27]. Then, we generated constants following the Gaussian distribution $N(0, 4)$ and added them to the external influence variables as their means.
4. The values of the observed variables x_i were generated using the connection strength matrix \mathbf{B} and external influences e_i . Finally, we permuted the variables according to a random ordering.




Figure 1: Box plot of accuracies of non-zero direct causal effects b_{ij} .

For every dataset, we computed the percentage of actually non-zero direct causal effects (or existing directed edges) in estimated ones, that is, the accuracy (or the true discovery rate) and the percentage of correctly estimated non-zero direct causal effects in actually non-zero ones, that is, the coverage. Fig. 1 and Fig. 2 show box plots of the accuracies and coverages of non-zero direct causal effects in \mathbf{B} . The median accuracies of non-zero direct causal effects for the four methods, random guessing, lasso, elastic net, and LiNGAM, were 0.017, 0.063, 0.047, and 0.104, respectively and the median coverages were 0.086, 0.829, 0.912, and 0.489, respectively. Fig. 3 and Fig. 4 show box plots of the accuracies and coverages of non-zero total causal effects in $\mathbf{A}(= (\mathbf{I} - \mathbf{B})^{-1})$. The median accuracies of non-zero total causal effects for random guessing, lasso, elastic net, and LiNGAM were 0.404, 0.462, 0.490, and 0.54, respectively and the median coverages were 0.085, 0.303, 0.473, and 0.142, respectively. The median computational times for the lasso, elastic net and LiNGAM were 60.13, 48.36, and 728 seconds on a standard PC.

In summary, our method gave better accuracies than the other methods. This implied that our method is more suitable for prioritizing future experiments.

4 Conclusions

We presented a new algorithm for estimating LiNGAM based on data with more variables than observations. Future works would include i) empirical comparison of our method and related algorithms on microarray gene expression datasets;




Figure 2: Box plot of coverages of non-zero direct causal effects b_{ij} .




Figure 3: Box plot of accuracies of non-zero total causal effects a_{ij} .

ii) extensions of our method to cases with latent confounding variables [22] and cyclic relations [41].




Figure 4: Box plot of coverages of non-zero total causal effects a_{ij} .

Acknowledgement

We would like to thank Satoshi Hara for discussions.

References

- [1] I. Cantone, L. Marucci, F. Iorio, M.A. Ricci, V. Belcastro, M. Bansal, S. Santini, M. Di Bernardo, D. Di Bernardo, and M.P. Cosma. A yeast synthetic network for in vivo assessment of reverse-engineering and modeling approaches. *Cell*, 137(1):172–181, 2009.
- [2] D. di Bernardo, M.J. Thompson, T.S. Gardner, S.E. Chobot, E.L. Eastwood, A.P. Wojtovich, S.J. Elliott, S.E. Schaus, and J.J. Collins. Chemogenomic profiling on a genome-wide scale using reverse-engineered gene networks. *Nature Biotechnology*, 23(3):377–383, 2005.
- [3] D. Pe’er and N. Hacohen. Principles and strategies for developing network models in cancer. *Cell*, 144:864–873, March 2011.
- [4] J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2000. (2nd ed. 2009).
- [5] M.H. Maathuis, D. Colombo, M. Kalisch, and P. Bühlmann. Predicting causal effects in large-scale systems from observational data. *Nature Methods*, 7(4):247–248, 2010.

- [6] 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(7):815–822, 2006.
- [7] C. Rangel, J. Angus, Z. Ghahramani, M. Lioumi, E. Sotharan, A. Gaiba, D.L. Wild, and F. Falciani. Modeling T-cell activation using gene expression profiling and state-space models. *Bioinformatics*, 20(9):1361–1372, 2004.
- [8] R. Opgen-Rhein and K. Strimmer. Using regularized dynamic correlation to infer gene dependency networks from time-series microarray data. In *Proceedings of the 4th International Workshop on Computational Systems Biology (WCSB 2006), Tampere*, volume 4, pages 73–76, 2006.
- [9] B.E. Perrin, L. Ralaivola, A. Mazurie, S. Bottani, J. Mallet, and F. dAlche Buc. Gene networks inference using dynamic Bayesian networks. *Bioinformatics*, 19(suppl 2):ii138–ii148, 2003.
- [10] C.J. Oates and S. Mukherjee. Network inference and biological dynamics. *Annals of Applied Statistics*, 2012. In press.
- [11] N. Friedman, M. Linial, I. Nachman, and D. Pe’er. Using Bayesian networks to analyze expression data. *Journal of Computational Biology*, 7(3-4):601–620, 2000.
- [12] D. Pe’er, A. Regev, G. Elidan, and N. Friedman. Inferring subnetworks from perturbed expression profiles. *Bioinformatics*, 17(suppl 1):S215–S224, 2001.
- [13] P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. Springer Verlag, 1993. (2nd ed. MIT Press 2000).
- [14] S. Shimizu, P. O. Hoyer, A. Hyvärinen, and A. Kerminen. A linear non-gaussian acyclic model for causal discovery. *J. Mach. Learn. Res.*, 7:2003–2030, 2006.
- [15] P. Spirtes, C. Glymour, R. Scheines, and R. Tillman. Automated search for causal relations: Theory and practice. In R. Dechter, H. Geffner, and J. Halpern, editors, *Heuristics, Probability, and Causality: A Tribute to Judea Pearl*, pages 467–506. College Publications, 2010.
- [16] A. Moneta, N. Chlaß, D. Entner, and P. Hoyer. Causal search in structural vector autoregressive models. In *Journal of Machine Learning Research: Workshop and Conference Proceedings, Causality in Time Series (Proc. NIPS2009 Mini-Symposium on Causality in Time Series)*, volume 12, pages 95–114, 2011.
- [17] P. Spirtes and C. Glymour. An algorithm for fast recovery of sparse causal graphs. *Social Science Computer Review*, 9:67–72, 1991.

- [18] D. Chickering. Optimal structure identification with greedy search. *Journal of Machine Learning Research*, 3:507–554, 2002.
- [19] A. Hyvärinen, J. Karhunen, and E. Oja. *Independent component analysis*. Wiley, New York, 2001.
- [20] A. Frigyesi, S. Veerla, D. Lindgren, and M. Höglund. Independent component analysis reveals new and biologically significant structures in micro array data. *BMC Bioinformatics*, 7(1):290, 2006.
- [21] Y. Sogawa, S. Shimizu, T. Shimamura, A. Hyvärinen, T. Washio, and S. Imoto. Estimating exogenous variables in data with more variables than observations. *Neural Networks*, 24(8):875–880, 2011.
- [22] P. O. Hoyer, S. Shimizu, A. Kerminen, and M. Palviainen. Estimation of causal effects using linear non-gaussian causal models with hidden variables. *International Journal of Approximate Reasoning*, 49(2):362–378, 2008.
- [23] S. Imoto, S. Kim, T. Goto, S. Aburatani, K. Tashiro, S. Kuhara, and S. Miyano. Bayesian network and nonparametric heteroscedastic regression for nonlinear modeling of genetic network. *Journal of Bioinformatics and Computational Biology*, 1(2):231–252, 2003.
- [24] P. O. Hoyer, D. Janzing, J. Mooij, J. Peters, and B. Schölkopf. Nonlinear causal discovery with additive noise models. In *Advances in Neural Information Processing Systems 21*, pages 689–696. 2009.
- [25] K. Zhang and A. Hyvärinen. On the identifiability of the post-nonlinear causal model. In *Proc. 25th Conference in Uncertainty in Artificial Intelligence (UAI2009)*, pages 647–655, 2009.
- [26] S. Shimizu, T. Inazumi, Y. Sogawa, A. Hyvärinen, Y. Kawahara, T. Washio, P. O. Hoyer, and K. Bollen. DirectLiNGAM: A direct method for learning a linear non-Gaussian structural equation model. *J. Mach. Learn. Res.*, 12:1225–1248, 2011.
- [27] F. R. Bach and M. I. Jordan. Kernel independent component analysis. *Journal of Machine Learning Research*, 3:1–48, 2002.
- [28] H. Zou. The adaptive Lasso and its oracle properties. *Journal of the American Statistical Association*, 101:1418–1429, 2006.
- [29] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of Royal Statistical Society: Series B*, 58(1):267–288, 1996.
- [30] A. Gretton, O. Bousquet, A. J. Smola, and B. Schölkopf. Measuring statistical dependence with Hilbert-Schmidt norms. In *Algorithmic Learning Theory: 16th International Conference (ALT2005)*, pages 63–77. 2005.

- [31] D. Entner and P. O. Hoyer. Discovering unconfounded causal relationships using linear non-gaussian models. In *New Frontiers in Artificial Intelligence, Lecture Notes in Computer Science*, volume 6797, pages 181–195, 2011.
- [32] S. Shimizu, A. Hyvärinen, Y. Kawahara, and T. Washio. A direct method for estimating a causal ordering in a linear non-gaussian acyclic model. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence (UAI2009), Montreal, Canada*, pages 506–513. AUAI Press, 2009.
- [33] J. Fan and J. Lv. Sure independence screening for ultrahigh dimensional feature space. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(5):849–911, 2008.
- [34] Y. Zhang, R. Li, and C.L. Tsai. Regularization parameter selections via generalized information criterion. *Journal of the American Statistical Association*, 105(489):312–323, 2010.
- [35] H. Zou and H.H. Zhang. On the adaptive elastic-net with a diverging number of parameters. *Annals of Statistics*, 37(4):1733, 2009.
- [36] J. Pearl. Causal diagrams for empirical research. *Biometrika*, 82(4):669–688, 1995.
- [37] H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.
- [38] T. Shimamura, S. Imoto, Y. Shimada, Y. Hosono, A. Niida, M. Nagasaki, R. Yamaguchi, T. Takahashi, and S. Miyano. A novel network profiling analysis reveals system changes in epithelial-mesenchymal transition. *PLoS One*, 6(6):e20804, 2011.
- [39] J. Friedman, T. Hastie, and R. Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software*, 33(1):1, 2010.
- [40] M. Kalisch and P. Bühlmann. Estimating high-dimensional directed acyclic graphs with the PC-algorithm. *Journal of Machine Learning Research*, 8:613–636, 2007.
- [41] G. Lacerda, P. Spirtes, J. Ramsey, and P. O. Hoyer. Discovering cyclic causal models by independent components analysis. In *Proceedings of the 24th Conference on Uncertainty in Artificial Intelligence (UAI2008)*, pages 366–374, 2008.