NIR imaging has been used for many years to look at functional parameters in the human brain (for example, saturations of hemoglobin)[12,13]. Although it is unlikely that NIR imaging will permit imaging of the entire brain, it is possible to image several centimeters below the human skull, potentially enabling, for example, detection of amyloid in the cortex. If amyloid imaging were used to diagnose Alzheimer disease, it is likely that simply detecting plaques in areas of the brain known to be affected by Alzheimer disease would be good enough.

The development of amyloid probes that can be imaged *in vivo* is almost certain to expedite the preclinical and clinical evaluation of novel Alzheimer-disease therapeutics that target amyloid β. Such ligands may also be useful in the diagnosis of atypical Alzheimer-disease cases. But current clinical diagnosis of Alzheimer disease is reasonably accurate, so it is unlikely that amyloid imaging will become a routine diagnostic modality unless it were relatively quick, safe and inexpensive. With further advances in the technology and ligands, NIR imaging of amyloid may fulfill these criteria[14].

There is evidence to suggest that amyloid deposition predates the clinical signs of Alzheimer disease by years or even decades; however, the exact temporal relationship between amyloid deposition and cognitive dysfunction remains to be established. The utility of existing amyloid probes for detecting very early stages of amyloid deposition in the brain of humans has not yet been determined, although most believe that significant improvements in sensitivity will be needed. As it is almost certain that Alzheimer disease will be easier to prevent than treat, a refined version of current amyloid imaging methods may ultimately be the diagnostic tool used to determine both who needs prophylactic treatment and when that treatment should be initiated.

1. Hintersteiner, M. *et al. Nat. Biotechnol.* **23**, 577–583 (2005).
2. Golde, T.E. *J. Clin. Invest.* **111**, 11–18 (2003).
3. Jack, C.R. Jr. *et al. Magn. Reson. Med.* **52**, 1263–1271 (2004).
4. Lee, S.P., Falangola, M.F., Nixon, R.A., Duff, K. & Helpern, J.A. *Magn. Reson. Med.* **52**, 538–544 (2004).
5. Klunk, W.E. *et al. Ann. Neurol.* **55**, 306–319 (2004).
6. Poduslo, J.F. *et al. Biochemistry* **43**, 6064–6075 (2004).
7. Skovronsky, D.M. *et al. Proc. Natl. Acad. Sci. USA* **97**, 7609–7614 (2000).
8. Wadghiri, Y.Z. *et al. Magn. Reson. Med.* **50**, 293–302 (2003).
9. Higuchi, M. *et al. Nat. Neurosci.* **8**, 527–533 (2005).
10. Frangioni, J.V. *Curr. Opin. Chem. Biol.* **7**, 626–634 (2003).
11. Graves, E.E., Ripoll, J., Weissleder, R. & Ntziachristos, V. *Med. Phys.* **30**, 901–911 (2003).
12. Strangman, G., Boas, D.A. & Sutton, J.P. *Biol. Psychiatry* **52**, 679–693 (2002).
13. Pouratian, N. *et al. Magn. Reson. Med.* **47**, 766–776 (2002).
14. Skoch, J., Dunn, A., Hyman, B.T. & Bacskai, B.J. *J. Biomed. Opt.* **10**, 011007 (2005).

Reverse engineering gene regulatory networks

Alexander J Hartemink

An information theoretic algorithm that prunes away potentially indirect interactions allows for improved reconstruction of biological networks.

Biological systems are wondrously and notoriously complex. Over the last fifty years, molecular biology has helped to reveal the vast and stunning array of components in biological systems. Now, we face the even more daunting challenge of systems biology: determining how all these puzzle pieces come together to create living systems. A recent paper by Basso *et al.*[1] published in *Nature Genetics* describes a statistical algorithm for more compactly and more accurately reverse engineering networks describing pair-wise interactions among genes and thin protein products. The network they recover from gene expression profiles of a variety of human B-cell populations suggests that the B-cell regulatory network has both a scale-free and hierarchical architecture, implying the presence of a few 'hubs' that are highly connected and preferentially connected to one another.

Reverse engineering is the process of elucidating the structure of a system by reasoning backwards from observations of its behavior. In reverse engineering biological networks, one of the first hurdles to overcome is semantic. The term 'network' has come to mean different things throughout biology, and the semantic overload is magnified when computational and statistical interpretations are added. Even in networks whose nodes are ostensibly the same objects (for examples, genes or their protein products), the network edges can mean vastly different things and should be interpreted with care. As just one example, edges can either be undirected (without an orientation) to capture relations that are symmetric or directed (with an orientation) to capture relations that are asymmetric.

An undirected edge between two genes may indicate that the genes are coexpressed or coregulated, participate in a common pathway or regulatory 'module' or share a common biological function, location or process; or that their protein products coprecipitate, directly bind one another, or assemble into the same complex (a problematic term in its own right). On the other hand, a directed edge between two genes may be used to represent a step in a metabolic pathway, signal transduction cascade, or stage of development; or it may indicate a causal control or a regulatory relationship.

This semantic caveat is important in trying to understand the myriad methods that have been proposed in the last decade for reverse engineering biological networks from system-wide data, especially gene expression data. Within this broader context, the ARACNe algorithm of Basso *et al.* is most closely related to an earlier method for producing 'relevance networks'[2,3]. Both sets of authors use a pair-wise mutual information criterion across gene expression profiles to recover edges that are undirected, but ARACNe improves on this somewhat by using the data processing inequality to prune out interactions suspected to be indirect.

After using synthetic data to assess the accuracy of their ARACNe algorithm, Basso *et al.* apply it to a rather sizable set of gene expression array data, collected from human B-cell populations with a variety of phenotypes, including both normal and malignantly transformed cells at different stages in the germinal center reaction process, from naive cells in the mantle zone to differentiated memory or plasma cells. This results in a network with about 129,000 undirected interactions between pairs of genes. Owing to the obvious complexity of such a network, the authors choose to focus on two simpler aspects: a statistical summary of the (global) connectivity distribution among all the nodes in the network, an approach that is quite in vogue; and a more detailed look at a specific (local) portion of the network centered around the proto-oncogene *MYC*, chosen both because of its clinical importance and because of the wealth of information available for corroboratory purposes.

Alexander J. Hartemink is in the Department of Computer Science, Duke University, Box 90129, Durham, North Carolina 27708-0129, USA.
e-mail: amink@cs.duke.edu

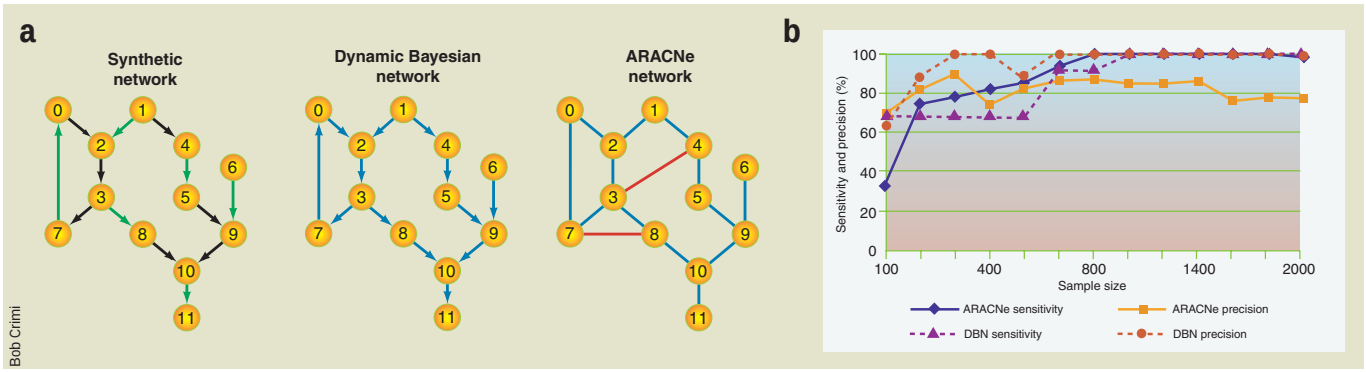

Figure 1 Comparison of the performance of ARACNe and a Bayesian network inference algorithm in reverse engineering a synthetic gene regulatory network. (**a**) The synthetic gene regulatory network used for assessing the reconstruction accuracy of network inference algorithms[7]. The network has 8 disconnected nodes serving as a negative control (not shown) and 12 interconnected nodes, including a cyclic loop formed by nodes 0, 2, 3 and 7 (regulatory interactions between two genes (nodes) are shown as arrows (edges); black and green arrows represent up- and downregulation, respectively). Reverse engineering using a Bayesian network inference algorithm to recover a dynamic Bayesian network (DBN) on the full data set results in a reconstructed network with 100% accuracy, as reported earlier[7] (blue arrows indicate correct edges with correct orientation; no incorrect edges were recovered). Reverse engineering using ARACNe on the full data set results in a reconstructed network with the same 13 correct edges as the DBN reconstruction, but without orientations; it includes two incorrect edges, between nodes 3 and 4 and nodes 7 and 8 (correct and incorrect edges are represented by blue and red, respectively; the ARACNe network is reproduced from Basso *et al.*). (**b**) Performance of ARACNe and a Bayesian network inference algorithm on subsets of the full data set. Sensitivity and precision are plotted as a function of the number of samples used for the analysis. At roughly the same sensitivity, the Bayesian network inference algorithm appears to exhibit better precision over a wide range of sample sizes (the ARACNe plots are reproduced from Basso *et al.*).

Analyzing the connectivity distribution of a network is currently popular for two reasons. First, it is a sensible first step in reasoning about networks so large that they are difficult to understand otherwise; for all intents and purposes, the interactions recovered by a tool like ARACNe are impossible to visualize directly in a way that facilitates insight. Second, articles and books[4] suggesting that many kinds of networks—biological, social, and engineered—are scale-free have recently been published in a flurry. Indeed, the network recovered by ARACNe from B-cell expression profiles has a connectivity distribution that suggests that it, too, is scale-free. Basso *et al.* appropriately caution that the reported connectivity distribution is not conclusive because an explainable saturation occurs in the 'low interaction count' portion of the curve, resulting in a distribution that is scale-free over only one order of magnitude. Nevertheless, the results are consistent with a hypothesis that this network is scale-free.

As for the subnetwork centered around *MYC*, it contains 56 genes adjacent to *MYC*, termed 'first neighbors,' along with 2,007 genes adjacent to these first neighbors, termed 'second neighbors.' Even a comparatively small subnetwork of this size is still a challenge to visualize insightfully, so the authors assess its quality in two ways. First, they determine whether the genes of the subnetwork are enriched for specific cellular process categories in the Gene Ontology database[5], which they are. Second—and this is a wonderful strength of the paper—the authors experimentally validate some of the first neighbors of *MYC*.

The list of *MYC* first neighbors was pruned to exclude those with lowest mutual information scores, those that do not contain *MYC* binding sites near the transcription initiation site, and those already known to be bound directly by *MYC*. The remaining 12 genes were tested for direct *MYC* binding using a standard chromatin immunoprecipitation assay, and 11 predictions were positive. Although the authors' resultant claim of over 90% specificity for ARACNe is perhaps optimistic as they excluded predictions with lowest mutual information scores and, more important, predictions not known to contain a *MYC* binding site, the results are still extremely encouraging. The success of this kind of experimental validation lends credence both to ARACNe and also to computational approaches more generally.

In closing, two further points should be made. First, this paper provides evidence confirming a simple intuition that many in this field have had, namely that gene expression data need not necessarily be collected from perturbation experiments for reverse engineering to be successful. Although perturbation experiments are certainly useful for network inference, they are also costly, and in some cases infeasible for either technical or ethical reasons. Basso *et al.* demonstrate that as long as the available data explore a wide range in the 'expression space' of the system, biologically meaningful interactions can be recovered by computational algorithms.

Second, the authors of this paper should be commended for evaluating the performance of ARACNe on synthetic data[6], and indeed, it performs nobly. However, they seem to misrepresent the performance of Bayesian networks on the same synthetic data. They report that ARACNe offers "substantially higher precision" in comparison with Bayesian networks, whereas we have observed exactly the opposite (**Fig. 1**). The discrepancy is most likely due to the fact that Basso *et al.* used a static Bayesian network in place of a more appropriate dynamic Bayesian network. This is only a minor quibble because the B-cell expression profiles examined in the remainder of the study are of quite a different character from the synthetic expression data in many regards, and it is not clear which method would be best suited to network inference in the B-cell context. Indeed, given the earlier caveat that the networks recovered by these and other methods typically have different semantics, it is likely that multiple methods will be needed to completely understand the regulation and dysregulation of B-cell differentiation, as well as other similar problems in systems biology.

1 Basso, K. *et al. Nat. Gen.* **37**, 382–390 (2005).
2. Butte, A.J. & Kohane, I.S. *Pac. Symp. Biocomput.* 2000, 418–429 (2000).
3. Butte, A.J., Tamayo, P., Slonim, D., Golub, T.R. & Kohane, I.S. *Proc. Natl. Acad. Sci. USA* **97**, 12182–12186 (2000).
4. Barabasi, A.L. *Linked: How Everything Is Connected to Everything Else and What It Means* (Plume, New York, NY, 2003).
5. The Gene Ontology Consortium. Gene Ontology: tool for the unification of biology. *Nat. Gen.* **25**, 25–29 (2000).
6. Smith, V.A., Jarvis, E.D. & Hartemink, A.J. *Bioinformatics* **18**, S216–S224 (2002).
7. Yu, J., Smith, V.A., Wang, P.P., Hartemink A.J. & Jarvis, E.D. in *3rd International Conference on Systems Biology* (Karolinska Institute, Stockholm, Sweden, 2002).