

Reconstruction of Metabolic Networks from High-Throughput Metabolite Profiling Data

In Silico Analysis of Red Blood Cell Metabolism

ILYA NEMENMAN,^{a,b} G. SEAN ESCOLA,^{a,c} WILLIAM S. HLAVACEK,^{b,d}
PAT J. UNKEFER,^e CLIFFORD J. UNKEFER,^e
AND MICHAEL E. WALL^{a,b,e}

^a*Computer, Computational and Statistical Sciences Division, Los Alamos*

^b*Center for Nonlinear Studies, National Laboratory*

^c*Columbia University Medical Center, New York, New York 10032, USA*

^d*Theoretical Division, Los Alamos, New Mexico 87545, USA*

^e*Bioscience Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*

ABSTRACT: We investigate the ability of algorithms developed for reverse engineering of transcriptional regulatory networks to reconstruct metabolic networks from high-throughput metabolite profiling data. For benchmarking purposes, we generate synthetic metabolic profiles based on a well-established model for red blood cell metabolism. A variety of data sets are generated, accounting for different properties of real metabolic networks, such as experimental noise, metabolite correlations, and temporal dynamics. These data sets are made available online. We use ARACNE, a mainstream algorithm for reverse engineering of transcriptional regulatory networks from gene expression data, to predict metabolic interactions from these data sets. We find that the performance of ARACNE on metabolic data is comparable to that on gene expression data.

KEYWORDS: metabolism; network reverse engineering; red blood cells; synthetic data

Addresses for correspondence: Ilya Nemenman, CCS-3, MS-B256, Los Alamos National Laboratory, Los Alamos, NM 87545. Voice: 505-665-8250; fax: 505-667-1126. Michael E. Wall, CCS-3, MS-B256, Los Alamos National Laboratory, Los Alamos, NM 87545. Voice: 505-665-4209; fax: 505-667-1126.

nemenman@lanl.gov; mewall@lanl.gov

Ann. N.Y. Acad. Sci. 1115: 102–115 (2007). © 2007 New York Academy of Sciences.
doi: 10.1196/annals.1407.013

THE NEED FOR BENCHMARK DATA

In recent years, high-throughput (HTP) microarray profiling has yielded large data sets that characterize the simultaneous activities of essentially all genes in a cell. These data sets have been used successfully to reverse-engineer (RE) cellular transcriptional regulatory networks.^{1–3} Similar developments are expected in the emerging field of metabolomics, in which sensitive HTP measurements of (relative or absolute) concentrations of many metabolites in a sample of cells are becoming possible.^{4–6} Anticipating the resulting data sets, there is a strong interest in development of computational tools that, unlike more traditional approaches based on sequence information⁷ or chemical reactivity and conservation laws,^{8,9} would use the relevant HTP data to expand our knowledge of metabolic networks, which is extensive but incomplete. Because metabolic networks share features with transcriptional regulatory networks, it is tempting to transfer successful methods developed in the context of transcriptional networks, such as those in Basso *et al.*,¹ Faith *et al.*,² Bussemaker *et al.*,³ and Margolin *et al.*,¹⁰ to inference of metabolic networks.

An advantage of transferring these methods is that only minimal modifications are required to the very extensive RE code base. However, it is not obvious that the existing methods will perform well on metabolic networks. Despite the superficial similarity, metabolic and transcriptional networks are quite different. In the transcriptional case, a transcription factor (parent) causes a change in the expression of its target gene (child), without any direct effects on its own activity. This leads to correlations among expressions of transcription factors (TFs) and their targets, and these can be readily discovered by various statistical techniques. Conversely, in metabolism, a substrate (parent) is *transformed* into a product (child). Thus, an increase in the child's abundance comes at the cost of decreasing the abundance of the parent. We therefore expect that the statistical associations in metabolic data will differ from those in gene expression data sets in as yet unknown ways. Furthermore, the experimental noise has a tendency to mask interactions of low-mean or low-variance species. This has been a problem even in transcriptional analysis (e.g., spurious interactions in the ribosomal complex¹), where the expression levels and the involved characteristic time scales of reaching steady states are largely uniform across all genes. On the other hand, kinetic rates in a metabolic network can vary over many orders of magnitude for different species. Thus the time required for an organism to achieve a metabolic steady state can vary from milliseconds to hundreds of hours.¹¹ Furthermore, many metabolites are short-lived and low-abundance, and a “fully expressed” metabolite can mean anywhere from a few molecules to a few million molecules per cell, making consideration of the measurement noise very important.

Because of these differences between transcription and metabolism, the fidelity of standard transcriptional RE algorithms for metabolic networks cannot be assumed. It is therefore useful to test these methods on benchmark data that

resemble real metabolic measurements, and for which the ground truth structure of the network is known. We are unaware of the existence of experimental data sets of this kind, and therefore we turn to numerical simulations. Existing synthetic data sets have focused on realistic modeling of transcriptional regulation,¹² not metabolism. Therefore, in this work, we undertake the task of generating synthetic benchmark metabolic data by using a published kinetic model of red blood cell (RBC) metabolism,¹¹ which involves 39 metabolites connected by 44 individual reactions. These data have been made publicly available at http://www.menem.com/~ilya/wiki/index.php/RBC_Metabolic_Network. We then use ARACNE, a modern transcriptional network RE algorithm, which was developed and validated for gene expression analysis,¹⁰ to infer metabolic interactions from these synthetic metabolic data, and we find that its performance is comparable to that on gene expression data. The outcome of this exercise suggests the possibility that, with minimal changes, other RE algorithms designed for use on HTP gene expression data might be fruitfully applied to metabolic data as well.

THE RBC METABOLIC BENCHMARK DATA

In generating synthetic benchmark data, our goal is not to accurately simulate a real system. Rather, our goal is to exercise transcriptional RE algorithms by generating data that are complex enough to incorporate different features of metabolism (dynamic ranges, temporal properties, correlations among chemical species, noises, etc.), but are still simple enough to analyze in detail. Specifically, we generated four data sets to account for increasingly complex scenarios of realistic profiling of RBC metabolism (see below). As the majority of transcriptional RE methods take steady-state abundance data as inputs, we focused on steady-state metabolic profiling in the first three of these synthetic data sets, and the fourth data set simulates dynamic metabolic profiling.

To generate the simulated data, we modified a publicly available Mathematica workbook implementation of the RBC model.¹¹ The model has five parameters that can be controlled externally: the Donnan ratio, R (which determines the difference in the pH inside and outside of the cell); glucose concentration, G ; total intracellular magnesium concentration, both free and bound, Mg ; intracellular inorganic phosphorus concentration, Pi ; and extracellular (plasma) sodium concentration, Na . For the first three data sets, these external control parameters were sampled at random 1000 times from specified probability distributions, representing different experimental setups, and the steady-state values of the metabolic network were found by using the methods in the RBC workbook. In a significant number of situations (up to 30% or more, depending on the data set), the randomly selected parameters did not lead to steady-state solutions. These samples were removed from the data sets.

Data Set 1 (Chemostat)

This data set simulates RBC steady-state measurements from chemostat culture experiments. All the parameters are uncorrelated, uniformly distributed variables, with the ranges indicated below (the numbers in parentheses are the values of the parameters assumed in the RBC workbook model¹¹). The ranges were established by a literature search for conditions of various culture experiments that did not lead to an immediate cell death. We emphasize again that our aim is not to accurately model specific biochemical experiments—instead, our aim is to provide test data with realistic features. Hence the crude specification of the parameter ranges below.

1. $R = 0.2 \dots 1.6$ (0.6). The natural value of R seems to be hard to pinpoint^{13,14} (see also discussion of data set 2 below), but experiments on prepared/perturbed cultures achieve R as high as 1.6.^{13,15} The lowest value of 0.69 comes from Kirk *et al.*¹³ and is higher than the value of 0.6, used by Jamshidi *et al.*¹¹ However, Kemp *et al.*¹⁶ suggest that the internal-to-external P_i concentration ratio (which is closely related to R) can go down to 0.2 for pH near 8.0. We chose this value as the lower limit on R , even though it is probably too low in the context of the RBC (pH = 7.4). For $R > 0.8$, the RBC dynamic system often does not have accessible steady-state solutions (depending on the other control parameters).
2. $G = 2.0 \dots 30.0$ mM (5.0).¹⁷
3. $Mg = 0.1 \dots 20$ mM (2.7).¹⁵ For larger values of Mg , steady states are hard to find, and we do not include such parameter combinations in the data set.
4. $P_i = 0.6 \dots 1.8$ mM (1.2).¹⁶
5. $Na = 100 \dots 200$ mM (140). Identifying this parameter from culture experiments is difficult since most data are about internal, rather than plasma sodium. However, Gill *et al.*,¹⁸ and McDonnell *et al.*¹⁹ note clinical cases with Na down to ~ 110 and up to ~ 180 , in which the patient still survived. In view of this, the range of 100–200 mM for culture experiments seems reasonable.

Additionally, we observed that an external pH of 7.55 is normal for culture conditions, and values down to pH = 7.0¹⁵ and up to 8.0¹⁶ have been recorded.

Data Set 2 (Natural)

This data set represents the variability of RBC metabolite concentrations in blood samples from healthy humans. The control parameters are taken as uncorrelated normal variables with means μ_i and standard deviations σ_i (indicated as $\mu_i \pm \sigma_i$ below), where $i = 1, \dots, 5$ denotes the identity of the

parameter. We take physiologically allowable intervals found in the literature as $\pm 2\sigma$ around the mean.

1. $R = 0.75 \pm 0.1$. Walser¹⁴ gives 0.825 for the normal human. The RBC model¹¹ uses 0.6, citing,¹³ which suggests $R = 0.69$. At the normal pH = 7.4, Raftos *et al.*¹⁵ suggest that the internal-to-external P_i ratio (and hence R) is between 0.4 and 1.0, with the median about 0.8. Given cell preparations for all of these analyses, neither of the values may be anatomically relevant, and the real value is likely unknown for *in vivo* conditions. Hence we've chosen the distribution for which the mean is roughly the average of the reported human data, and the range of the normal data is about $\pm 2\sigma$ around the mean.
2. $G = 5 \pm 0.6$ mM.^{20,21}
3. $Mg = 3.3 \pm 0.2$ mM.^{15,22}
4. $P_i = 0.9 \pm 0.15$ mM. This estimate is based on P_i values between 0.6 and 1.2,¹⁶ obtained for an external pH of 7.4 (the default value of the RBC model). Values reported in alternative sources (1.0 in Jacobasch *et al.*,²³ 0.8 in Katz *et al.*,²⁴ and 0.98 in Garay and Garrahan²⁵) differ from these by less than 2σ . It is important to realize that the RBC model uses intracellular concentration as the control parameter, while most investigators^{20,22} focus on plasma concentrations, leading to large discrepancies.
5. $Na = 140 \pm 2.5$ mM.^{20,22} Values reported in alternative sources differ from these by less than 2σ .

Additionally, the following information was collected: normal external pH of 7.24¹⁵ or 7.4¹⁶ for an unperturbed cell.

Data Set 3 (Correlated)

This data set attempts to model the *in vivo* metabolite concentrations more faithfully by incorporating physiological correlations among the controls. Using this data set, one may study effects of the correlations (and thus reduction of the dynamic ranges of the response variables) on the performance of RE algorithms. For most of the parameters, we were unable to find quantitative measurements of the correlation coefficients in the literature, and instead only trends were reported. We summarized the trends into correlation coefficients ρ_{ij} of 0 (no trend or no data available), ± 0.3 (weak correlation), and ± 0.5 (strong correlation). Then the data set was generated by sampling the control parameters from multivariate normal distributions with means and variances as in data set 2, and with the correlation coefficients summarized in TABLE 1.

1. P_i - R and Mg - R : The P_i -pH and R -pH correlation coefficients are -0.85 and $-0.6 \dots -0.76$ (for different species), respectively.¹⁶ Thus, it is reasonable to assume that the P_i - R correlation is positive and large ($+0.5$

TABLE 1. Correlation coefficients of the five control parameters of the RBC model^a

	R	G	Mg	Pi	Na
R		-0.3 (see text)	-0.5 ¹⁵	+0.5 ^{15,16}	-0.3 ²⁶
G			+0.3 ²⁹	-0.5 ²⁷	+0.3 ²⁶
Mg				0	-0.3 ²⁸
Pi					-0.3 ²⁵
Na					

^aReferences used to set the values are listed as well.

in our notation). Also notice that the Donnan ratio should have correlations of similar magnitudes with all internal ionic concentrations, with positive/negative correlations for anions/cations, respectively. An agreement with the Mg - R value¹⁵ is encouraging.

2. Na - R : Recall that Na is an extracellular concentration and the correlation with R is not obvious. Nonetheless, equation 2 in Nguyen and Kurtz²⁶ suggests a negative correlation. This, however, may be affected by fluctuations of the total sodium level and of cell volumes. Therefore, we choose a value of -0.3 for this correlation.
3. Na - G : A small positive correlation is reported in Nguyen and Kurtz.²⁶
4. G - R : While we found no explicit data relevant for estimating this correlation, G is positively correlated with Mg and Na ions, which are, in turn, negatively correlated with R (see above). Thus, a small negative value for the G - R correlation is assumed.
5. Na - Pi : Garay and Garrahan²⁵ suggest strong positive pair-wise correlations between the internal Na and Na efflux, between the internal Pi and the inverse of the Na efflux, and between the external and internal Na . Overall, we deduce a weak negative correlation between external Na and internal Pi ; this correlation is further supported by the opposite sign charges of these particles.
6. G - Pi : Significant negative correlation is reported by Rose and Warms.²⁷
7. Na - Mg : Weak competitive behavior between these species is reported by Frenkel *et al.*²⁸

Data Set 4 (Evolving Parameters)

The RBC model takes up to 100 hours or more to reach a steady state.¹¹ However, in a natural environment, the control parameters fluctuate on time scales less than an hour. For example, it takes only tens of seconds for blood to circulate. Further, the same drop of blood visits the liver every 20 minutes or so, and this may completely change the concentration of various ions in the cells and around them. Since we found no explicit data about the temporal variability of the five control parameters in humans, we decided to model

each of them as correlated Ornstein–Uhlenbeck processes with the means, the standard deviations, and the species-species correlations as in data set 3, and the correlation time $\tau = 20$ min for each process. This data set required the most extensive changes to the RBC model Mathematica script, enabling dynamic variation of the control parameters during the temporal evolution of the system. The resulting time-series data represent 20 hours of evolution of the RBC model, sampled every 10 sec (for a total of 7201 samples); researchers may subsample the series and/or shorten it to better match the sampling scheme of an actual experiment. This data set is designed to test the application of network reconstruction to non-equilibrated time-series data.

NOISE

We simulate experimental errors for each metabolite concentration by adding Gaussian zero-mean random noise to the output of RBC model. The noise variance is given by $A^2 + B^2x^2$, where x is the abundance, and A and B describe the contribution of the absolute and the relative noise components. Each of the four simulated data sets is available from our web site with a multitude of A and B values. This noise model is suitable for transcriptional data,³⁰ and here we assume it is relevant for metabolic data as well. However, specialized noise studies will have to be performed on real experimental data, once available, to verify this model. Notice, in particular, that this model does not take into the account errors that may emerge due to the overlap of peaks in mass spectrometry–based metabolite profiling.

REVERSE ENGINEERING OF METABOLIC NETWORKS WITH ARACNE ALGORITHM

Using our simulated metabolic data sets, we tested whether the ARACNE algorithm¹⁰ could be used for analysis of metabolic networks.

Like many other network reconstruction methods, ARACNE models dependencies among activity variables (e.g., gene expressions or metabolite concentrations), $\{g_i\}$, as probabilistic, defined by an unknown probability distribution $P(\{g_i\})$. A probabilistic description reflects the effects of unobserved molecular species and of experimental noise. A bona fide biological interaction corresponds to a nonzero statistical dependency between activity profiles, measured by the mutual information $I(g_i, g_j) = \langle \log P(g_i, g_j) / P(g_i)P(g_j) \rangle$. Evaluating the mutual information and identifying its value above a certain threshold with an interaction is the basis of the relevance networks (RN) method.³¹ However, as signals propagate through the networks, many non-interacting species also become correlated and result in a positive I (e.g., two non-interacting targets of the same TF may be highly statistically dependent). This is a major problem

for most RE algorithms. To isolate statistical interactions that have the highest chance to correspond to real biological interactions, ARACNE then uses the data processing inequality (DPI)¹⁰ after statistically significant values of the mutual information have been identified. The DPI states that if stochastic variables g_1 and g_3 interact only through a third one, g_2 , then $I(g_1, g_3) \leq \min [I(g_1, g_2); I(g_2, g_3)]$. Thus ARACNE analyzes each gene triplet and designates the link with the lowest MI value as indirect. To minimize the effect of incorrect estimations of the MI, this designation is made only when a link is at least $\tau\%$ below the second weakest one. For $\tau \sim 5 \dots 15\%$, the method has been validated in synthetic¹⁰ and *in vitro*¹ transcriptional networks.

To establish a metric for the fidelity of a reconstruction, we note that the RBC interaction network is specified by a system of first-order differential equations $dx_i/dt = f(\{x_{j \in Ne(i)}\})$, where $Ne(i)$ is a set of neighbors of the node i , including all immediate parents, children, and modulators (effectors) of the reactions. Thus, the steady-state probability distribution is $P(i | \text{rest of network}) = P(i | Ne(i))$, which corresponds to linking metabolites in the interaction graph to all of their neighbors.¹⁰ This results in a “gold standard” adjacency matrix with 107 pair-wise interactions among metabolites, to which the ARACNE reconstruction is to be compared.

As a first check, we reconstructed the RBC metabolic network using just 19 conditions from data set 1 ($\tau = 0$). ARACNE was applied using default parameter values.³² A total of 14 interactions are predicted, 11 of which are substantiated by the model; this is a recall of 10% (14 out of 107) at a precision of 78% (11 out of 14). We then performed a systematic study using data set 1 with added measurement noises of different levels, modeling real experiments. For each run, we fine-tuned the *kernel width*,^{10,32} an internal ARACNE parameter essential for estimation of the mutual information. In addition, we removed metabolites for which the variance across different steady states was comparable to that due to experimental noise (establishing interactions of such metabolites is impossible by any statistical method). All indirect interactions mediated by the deleted metabolites were then considered to be direct in the validation. For a small and dense network, like the RBC one, where 14% of all metabolite pairs are connected by interactions, the node removal sets a limit on realistic values of the noise: at high noise variance, the network becomes a small and almost fully connected cluster, making precision of the algorithm artificially high.

Precision vs. recall curves (PRCs) for noisy and noiseless data and for two different DPI tolerance values corresponding to the RN and the ARACNE algorithms are shown in FIGURE 1. These curves are generated by adjusting the significance threshold for mutual information estimation so that metabolite pairs with a mutual information below the threshold are not allowed to participate in an interaction. Higher thresholds decrease the number of putative interactions, which eliminates most of the false positives, thus increasing the precision. Lower thresholds admit more pairs for consideration, and lead to




FIGURE 1. Precision vs. recall curves for different noise levels (*thick lines*, no noise; *thin lines*, $A = 0.001$, $B = 0$), indexed by the number of nodes with the variance above the experimental noise (39 and 29, respectively). *Dashed lines* were constructed with the DPI tolerance of 100% (no DPI applied, equivalent to the relevance networks algorithm), and solid lines have 0% DPI tolerance (pure ARACNE algorithm). Curves corresponding to other values of the tolerance generally fall between these two extremes. At recall of 1.0, the precision is 14% and 22% for thick/thin lines, respectively, and it corresponds to 107/88 true positives out of 741/406 possible metabolite pairs, all indexed as putative interactions if no DPI is applied, and if all mutual information values are treated as significant.

higher recalls. Because the RBC network is small, dense, and loopy, we expect the precision to drop quickly as the recall increases, as seen in FIGURE 1. Further, because random assignment of interaction edges leads to precisions of 15–20%, only the top left corner of FIGURE 1 shows significant improvement over random guessing.

The most significant aspect of FIGURE 1 is that, for all curves, there is always a range with a non-negligible recall and with precision ~ 1 . Thus the algorithms can be tuned to produce a (small) number of predictions that are highly likely to be true. Furthermore, in the relevant region of high precision, the low tolerance (ARACNE) lines are substantially higher than the high tolerance (RN) ones. Just as in the case of transcription,¹⁰ this indicates an improvement from using ARACNE over RNs on metabolic data. Thus, a minimally modified ARACNE algorithm can be used to accurately predict metabolic interactions.

In FIGURE 2, we compare how smaller variability and correlations in the external control parameters affect the validity of the reconstruction. These effects can be observed only for large noise levels, which cause the noise to be larger than the signal for some metabolites. Here we used a noise level




FIGURE 2. Precision vs. recall curves for data sets with different distributions of the control parameters and the same noise ($A = 0.01$, $B = 0$). For the three data sets, the number of true interactions (chance precision) is 32%, 35%, and 34% of the total number of possible metabolite pairs. Note unconventional scaling of the axes.

such that the effective number of nodes in data set 1 (*chemostat*) is 19. At this noise level, both data sets, 2 and 3 (*natural* and *correlated*) have a smaller number of effective nodes (16 and 17, respectively), implying that they have a smaller variability of responses. The correlated data set has more effective nodes than the natural one, implying that parameter correlations are synergistic in their effect on variability of responses. At first glance, the *natural* data set PRC seems to indicate the best performance. However, the effective network is smaller for this data set and the random-chance precision (35%) is higher. By comparison, the *chemostat* network has a smaller random-chance precision (32%) than the *correlated* one (34%), even though its PRC is higher. Overall, the results imply that the decreased response variance and/or spurious correlations among metabolites introduced by correlations among control parameters hinder reverse engineering of metabolic interactions.

Finally, in FIGURE 3 we examine applicability of ARACNE to time-dependent metabolic data. Specifically, we would like to understand how temporal correlations among subsequent samples affect reconstruction. In order not to confound dynamics with the data set size, here we always reconstruct the network from 400 samples. However, in different runs, these samples are spaced every 10, 40, and 160 sec apart, with the temporal correlations among subsequent samples decreasing with the dilution. For comparison, we also plot




FIGURE 3. Precision vs. recall curves for dynamic data using 400 samples sampled at different intervals. Steady-state reconstruction with the same number of samples is shown for comparison. For these plots, the effective number of nodes is 37, 38, 38, and 38, respectively. Note unconventional scaling of the axes.

reconstruction based on 400 independently sampled steady states from data set 3. We clearly see that, while reconstruction with temporally correlated data is possible, the corresponding PRCs are generally lower than for the steady-state case; they also drop more sharply. However, we note that ARACNE is designed for steady-state data, and better performance would be expected using methods that explicitly consider temporal structure in the data,^{3,33} As an exception, at very large precision values, the steady-state PRC is lower than its counterparts for some of the temporal data. While this better performance might just be an artifact, it also hints at the existence of transient metabolite interactions that are unobservable in steady state, and that provide useful information for reverse engineering. Finally, the relatively low initial precision of the 10-s sampling curve might be a result of small abundance changes between subsequent samples at fine temporal discretization—the noise might easily mask such small changes.

CONCLUSIONS

We have generated benchmark synthetic metabolic data sets and analyzed them with ARACNE, a representative transcriptional networks reverse-engineering algorithm. The performance of the algorithm for metabolic networks is comparable to that for transcriptional ones. This finding may be

considered as a basis for an optimistic view that transcriptional networks RE algorithms may be transferred *en masse* with relatively few modifications to metabolic applications. However, this ease of transfer must be verified on a case-by-case basis. Importantly, we now have synthetic benchmarking data sets, which can aid in this verification. The most important limitation of the data sets is the relatively small size of the RBC metabolic network, which limits the ability to ascertain statistical significance of many findings. For example, because of this problem, we cannot verify whether the MINDY algorithm,³⁴ which is an extension of ARACNE capable of elucidating interactions that are statistically insignificant overall, but become apparent in subsets of data, provides an advantage over the original ARACNE when applied to this metabolic data.

The steady-state values in data sets 1–3 are ideal for application of ARACNE, RN, and other steady-state algorithms. However, because it may take many hours for a controlled culture to equilibrate, time-resolved assays might enable more rapid inference of metabolic interactions. Data set 4 provides a means for testing metabolic network inference using such experiments, and can potentially help designing efficient sampling schemes that minimize the cost of reconstructing metabolic networks.

Finally, we notice that performance of transcriptional RE algorithms on metabolic data can be substantially improved beyond that observed in FIGURES 1–3. Indeed, not every biochemical reaction is possible in nature. For example, metabolic reactions conserve mass, and, unlike the case in transcription, this places strict limits on which species can be metabolically coupled. Similarly, atomic species are also conserved by metabolism, which places even more constraints on allowable reactions, akin to that reported by Kauffman *et al.*⁸ Metabolic profiling almost always involves detailed determination of metabolite masses, and frequently provides information about the chemical structure of the compounds using isotopic labeling. It is, therefore, essential to incorporate these constraints into HTP profiling-based methods, such as ARACNE, in the future.

ACKNOWLEDGMENTS

This work was supported by DOE/NNSA under contract No. DE-AC52-06NA25396. I.N. and W.S.H. were further supported by NIH award 1 R21 GM080216 01.

REFERENCES

1. BASSO, K. *et al.* 2005. Reverse engineering of regulatory networks in human B cells. *Nat. Genet.* **37**: 382–390.

2. FAITH, J. *et al.* 2007. Large-scale mapping and validation of *Escherichia coli* transcriptional regulation from a compendium of expression profiles. *PLoS Biol.* **5**: 54–66.
3. BUSSEMAKER, H., B. FLOAT & L. WARD. 2007. Predictive modeling of genome-wide mRNA expression: from modules to molecules. *Annu. Rev. Biophys. Biomol. Struct.* **36**: 329–347.
4. DE VOS, R. *et al.* 2007. Untargeted large-scale plant metabolomics using liquid chromatography coupled to mass spectrometry. *Nat. Protoc.* **2**: 778–791.
5. STEINHAUSER, D. & J. KOPKA. 2007. Methods, applications and concepts of metabolite profiling: primary metabolism. *EXS* **97**: 171–194.
6. RABINOWITZ, J. 2007. Cellular metabolomics of *Escherichia coli*. *Expert Rev. Proteomics* **4**: 187–198.
7. PALEY, S. & P. KARP. 2002. Evaluation of computational metabolic-pathway predictions for *Helicobacter pylori*. *Bioinformatics* **18**: 715–724.
8. KAUFFMAN, K., P. PRAKASH & J. EDWARDS. 2003. Advances in flux balance analysis. *Curr. Opin. Biotech.* **14**: 491–496.
9. MU, F. *et al.* 2006. Prediction of oxidoreductase-catalyzed reactions based on atomic properties of metabolites. *Bioinformatics* **22**: 3082–3088.
10. MARGOLIN, A. *et al.* 2006. ARACNE: an algorithm for the reconstruction of gene regulatory networks in a mammalian cellular context. *BMC Bioinformatics* **7**(Suppl 1): S7.
11. JAMSHIDI, N. *et al.* 2001. Dynamic simulation of the human red blood cell metabolic network. *Bioinformatics* **17**: 286–287.
12. MENDES, P., W. SHA & K. YE. 2003. Artificial gene networks for objective comparison of analysis algorithms. *Bioinformatics* **19**(Suppl 2): II122–II129.
13. KIRK, K., P.W. KUCHEL & R.J. LABOTKA. 1988. Hypophosphite ion as a ^{31}P nuclear magnetic resonance probe of membrane potential in erythrocyte suspensions. *Biophys. J.* **54**: 241–247.
14. WALSER, M. 1960. Protein-binding of inorganic phosphate in plasma of normal subjects and patients with renal disease. *J. Clin. Invest.* **39**: 501–506.
15. RAFTOS, J., V. LEW & P. FLATMAN. 1999. Refinement and evaluation of a model of Mg^{2+} buffering in human red cells. *Eur. J. Biochem.* **263**: 635–645.
16. KEMP, G., A. BEVINGTON & R.G.G. RUSSELL. 1988. Erythrocyte phosphate metabolism and pH in vitro: a model for clinical phosphate disorders in acidosis and alkalosis. *Miner. Electrolyte Metab.* **14**: 266–270.
17. Wikipedia. 2006. The Free Encyclopedia: Glucose meter. http://en.wikipedia.org/wiki/Glucose_meter.
18. GILL, G. *et al.* 2006. Characteristics and mortality of severe hyponatraemia: a hospital-based study. *Clin. Endocrinol.* **65**: 246.
19. MCDONNELL, C. *et al.* 2005. Diabetic ketoacidosis, hyperosmolarity and hypernatremia: are high-carbohydrate drinks worsening initial presentation? *Pediatr. Diabetes* **6**: 90–94.
20. Medline Plus Medical Encyclopedia. <http://www.nlm.nih.gov/medlineplus/encyclopedia.html>. 2006. U.S. National Library of Medicine.
21. Lab Tests Online: Comprehensive Metabolic Panel. <http://www.labtestsonline.org/understanding/analytes/cmp/cmp.html>. 2004. American Association for Clinical Chemistry.
22. Blood test results—normal ranges. 2006. bloodbook.com.

23. JACOBASCH, G., S. MINAKAMI & S.M. RAPOPORT. 1974. Glycolysis of the erythrocyte. *In Cellular and Molecular Biology of Erythrocytes*. H. Yoshikawa & S.M. Rapoport, Eds.: 55–92. University Park Press. Baltimore, MD.
24. KATZ, L. *et al.* 1988. Intracellular pH and inorganic phosphate content of heart in vivo: a ³¹P-NMR study. *Am. J. Physiol.* **255**: H189–H196.
25. GARAY, R. & P. GARRAHAN. 1975. The interaction of adenosinetriphosphate and inorganic phosphate with the sodium pump in red cells. *J. Physiol.* **249**: 51–67.
26. NGUYEN, M. & I. KURTZ. 2004. New insights into the pathophysiology of the dysnatremias: a quantitative analysis. *Am. J. Physiol. Renal Physiol.* **287**: F172–F180.
27. ROSE, I. & J. WARMS. 1966. Control of glycolysis in the human red blood cell. *J. Biol. Chem.* **241**: 4848.
28. FRENKEL, E., M. GRAZIANI & H. SCHATZMANN. 1989. ATP requirement of the sodium-dependent magnesium extrusion from human red blood cells. *J. Physiol.* **414**: 385–397.
29. BARBAGALLO, M. *et al.* 1999. Effects of glutathione on red blood cell intracellular magnesium: relation to glucose metabolism. *Hypertension* **34**: 76–82.
30. TU, Y., G. STOLOVITZKY & U. KLEIN. 2002. Quantitative noise analysis for gene expression microarray experiments. *Proc. Natl. Acad. Sci. USA* **99**: 14031–14036.
31. BUTTE, A.J. & I.S. KOHANE. 2000. Mutual information relevance networks: functional genomic clustering using pairwise entropy measurements. *Pac. Symp. Biocomput.* **5**: 415–426.
32. MARGOLIN, A. *et al.* 2006. Reverse engineering cellular networks. *Nat. Protocols* **1**: 663–672.
33. WIGGINS, C. & I. NEMENMAN. 2003. Process pathway inference via time series analysis. *Experimental Mech.* **43**: 361–370.
34. WANG, K. *et al.* 2006. Genome-wide discovery of modulators of transcriptional interactions in human B lymphocytes. *Lect. Notes Comput. Sci.* **3909**: 348–362.