Inference of Gene Coexpression Networks by Integrative Analysis across Microarray Experiments

Laura L. Elo [1,2,*] , Riitta Lahesmaa [2] , Tero Aittokallio [1,2,3]

[1]Department of Mathematics, University of Turku, FI-20014 Turku, Finland

[2]Turku Centre for Biotechnology, PO Box 123, FI-20521 Turku, Finland

[3]Systems Biology Unit, Institut Pasteur, FR-75724 Paris, France

Summary

We improve the reliability of detecting coexpressed gene pairs from microarray data by introducing a novel probe-level quality-weighted similarity measure for combining data across different Affymetrix experiments. In construction of gene coexpression networks, the proposed procedure is less sensitive to noise than the corresponding single-experiment approaches or the conventional integrative approaches, even when a relatively small number of samples and conditions is available. The present results indicate how the accumulated microarray data can be effectively exploited to increase the quality of the inferred networks. In particular, we demonstrate its biological relevance in identifying coexpressions in mouse T helper cell differentiation.

1 Introduction

The explosion of gene expression microarray data has opened great opportunities to explore the complex gene regulatory networks in higher eukaryotes, such as human and mouse. To effectively exploit the wealth of data accumulated in laboratories or in public repositories, there is an urgent call for systematic methods to integrate data across multiple experiments.

Most existing studies that combine data across multiple microarray experiments have focused on detecting differentially expressed genes between sample groups [6, 8]. Another type of integrative studies have been conducted to discover gene coexpressions across a multitude of diverse experiments. These global-scale meta-coexpression studies have shown their potential in predicting functional relatedness between gene pairs [18, 26], in inferring cooperativities between transcription factors [28], or in investigating the evolutionary conservation of coexpression [24]. To date, only few meta-analytic studies have been reported that investigate gene coexpressions involved under specific conditions or disease states. Aggarwal et al. [1] employed a rank-based statistic across four independent experiments to identify conserved gene coexpression interactions in gastric cancer. Choi et al. [7] identified coexpressed gene pairs in various cancer and normal tissues across several experiments and compared the induced cancer- and normal-networks to detect coexpression changes that might contribute to the malignant transformation in cancer cells. Such more specific meta-coexpression analyses would, however, be of practical importance to many researchers when improving the reliability of their

*Corresponding author. Tel: +358 2 333 8002; Fax: +358 2 333 8000; Email: laliel@utu.fi

inferences and discovering potential novel biological relationships that may be too weak to be detected in the individual experiments.

Gene regulatory networks provide a natural way to represent complex cellular interactions. Several approaches have been proposed to infer the relations between genes in such networks from microarray data or other data sources [22, 25, 26]. A rather simple and widely used approach is the relevance network, where pairwise expression correlations above a particular threshold value are considered as significant [4]. More sophisticated approaches are based on graphical models, such as the Bayesian networks, which can take conditional dependence relations into account [9]. However, due to their large number of parameters, the estimated models may contain many false interactions even in idealized situations. Further challenges are faced when applying the inference algorithms to short profiles of noisy data from real biological systems. Recently, Magwene and Kim introduced a novel network construction framework, called the first-order conditional independence model, which aims at exploiting the conditional interactions without excessive computational complexity [22].

In conventional meta-coexpression studies, similarity across experiments is assessed by calculating the correlation over all samples after merging the datasets [24], by averaging the single-experiment correlations [13], or by determining the reoccurrence of significant correlations in multiple experiments [18]. A major deficiency with such simplified approaches is that they do not consider the different characteristics of the experiments. Recently, Choi et al. [7] advanced the meta-correlation estimation by incorporating the sample sizes into the similarity calculations in the context of a random effects meta-analysis model [14]. We develop their approach further by considering also the qualities of the individual experiments performed with different generations of Affymetrix arrays. The quality weighing is motivated by the observation that a considerable number of probes on the arrays do not uniquely match their intended targets, which reduces the quality of the corresponding probesets [8, 23]. Other sources of discrepancies in the expression measurements originate from probes representing different splice variants of the same gene [17], or unreliable detections of the low-level signals [15].

In the present study, we introduce a systematic procedure to identify coexpressed genes by combining gene expression data across multiple Affymetrix experiments. We investigate the network inferences in a practical research setting where the datasets comprise several relatively small microarray experiments performed with different generations of Affymetrix arrays by different researches at different times. To improve the reliability of the inference, we present a novel probe-level quality-weighted similarity measure and compare its performance to that of the conventional integrative approaches as well as the corresponding single-experiment analyses. We concentrate on the relevance networks, although similar weighing scheme could be generalized to other network construction approaches as well. We show that the proposed procedure is robust against increasing the level of noise and applicable even with short profile-lengths. In particular, we apply the novel coexpression approach to identify gene coexpression networks involved in mouse T helper cell differentiation.

2 Materials and Methods

2.1 Microarray Data

In eight separate experiments denoted by E1-E8, CD4+ T cells were isolated from spleens of BALB/cJ mice (E1-E6) or C57BL/6 mice (E7-E8). Cells were harvested for RNA isolation either immediately (Thp) or after 2, 6, 24 or 48 h treatment. The treatments were Th0 (activation), Th1 (activation and polarization with IL-12) and Th2 (activation and polarization with IL-4). Full experimental details of the cell cultures, treatments and microarray processing in experiments E1-E3 are described in [21] and in experiments E4-E6 in [5]. Experiments E7-E8 were conducted similarly (unpublished). Experiments E1-E6 were performed with Affymetrix MG-U74Av2 arrays (12488 probesets), and experiments E7-E8 with Affymetrix MOE430A arrays (22690 probesets). Summary of the experiments used in the present study is shown in Table 1.

The Affymetrix microarrays utilize a set of different 25mer probes (the so-called probeset) to measure the expression of a transcript target [19]. Conventionally, these probe-level expression values are normalized and summarized into simple numerical estimates of the probeset-level gene activities. We used for the summarization the widely applied Robust Multiarray Average (RMA) procedure [16], implemented in the Bioconductor 'affy' package [10]. For the probe-level quality analysis, the perfect match (PM) probe intensity data were quantile normalized similarly as in RMA but no summarization was performed.

The expression measurements across the different Affymetrix generations were matched using the 'bestmatch' tables provided by the manufacturer (`www.affymetrix.com`). These best match pairs are based on the similarity between the target sequences of the probesets and are commonly used to compare different generations of Affymetrix arrays. To reduce the number of genes in the coexpression analysis, we carried out a non-specific filtering [11], where we required that a gene must be called Present (Affymetrix MAS5.0) and have an expression level greater than $\log_2 300$ in at least 25% of the samples. This left us with a total of 1221 genes.

Experiment	Number of conditions	Time points involved					Affymetrix array
E1	3					48h	MG-U74Av2
E2	10	0h	2h	6h		48h	MG-U74Av2
E3	8	0h	2h	6h		48h	MG-U74Av2
E4	13	0h	2h	6h	24h	48h	MG-U74Av2
E5	4	0h				48h	MG-U74Av2
E6	3					48h	MG-U74Av2
E7	4	0h				48h	MOE430A
E8	4	0h				48h	MOE430A

Table 1: Summary of the experiments used in the present study. In experiment E3, only Th1 samples were measured at 48 h. In the other experiments and time points (except 0h), all three treatments Th0, Th1 and Th2 were studied.

2.2 Conventional Expression Similarities

Perhaps the most frequently used similarity measure between two gene expression profiles $\mathbf{x}= (x_1, \ldots, x_N)$ and $\mathbf{y}= (y_1, \ldots, y_N)$ over N conditions (arrays) is the Pearson correlation, defined as

$$r(\mathbf{x}, \mathbf{y}) = \frac{\sum_{n=1}^{N}(x_n - \bar{x})(y_n - \bar{y})}{\sqrt{\sum_{n=1}^{N}(x_n - \bar{x})^2 \sum_{n=1}^{N}(y_n - \bar{y})^2}}, \tag{1}$$

where $\bar{x}$ and $\bar{y}$ are the profile averages. If we denote by $\mathbf{x}_k$ and $\mathbf{y}_k$ the expression profiles of genes g_x and g_y in an experiment k, $k = 1, \ldots, K$, then their single-experiment similarity is defined as

$$r_k(g_x, g_y) = r(\mathbf{x}_k, \mathbf{y}_k). \tag{2}$$

The Pearson correlation has extensively been used also in several meta-analytic similarity measures [7, 13, 18, 24]. Conventionally, the overall similarity between two genes across multiple experiments is determined by first merging the expression datasets and then calculating the correlation across all arrays [24]. We denote this by

$$r_M(g_x, g_y) = r(\mathbf{x}^*, \mathbf{y}^*), \tag{3}$$

where $\mathbf{x}^* = (\mathbf{x}_1, \ldots, \mathbf{x}_K)$ and $\mathbf{y}^* = (\mathbf{y}_1, \ldots, \mathbf{y}_K)$ are the merged expression profiles. Before merging the profiles across different experiments, the expression measurements were scaled to have the same average expression across the experiments. Another standard approach to assess the overall similarity between two genes is to determine the reoccurrence of their significant coexpressions over multiple datasets [18]. This gives the statistic

$$r_C(g_x, g_y) = \#\{k | r_k(g_x, g_y) \geq c_k, k = 1, \ldots, K\}, \tag{4}$$

where c_k is the significance threshold in the kth experiment and $\#S$ denotes the cardinality of set S. Alternatively, the average correlation over multiple datasets have been used [13], defined as

$$r_A(g_x, g_y) = \frac{1}{K}\sum_{k=1}^{K} r_k(g_x, g_y). \tag{5}$$

Recently, the meta-analytic random effects model of correlation [14] was applied in the context of gene coexpression analysis [7]. Under the model, the similarity between two particular genes is defined hierarchically as

$$z_k = \mu_k + \epsilon_k, \qquad \epsilon_k \sim N(0, s_k^2) \tag{6}$$

$$\mu_k = \mu + \delta_k, \qquad \delta_k \sim N(0, \tau^2),$$

where μ is the unknown similarity of interest, τ^2 represents the between-experiment variability in the coexpression, and z_k and s_k^2 are the observed similarity and the sampling error in experiment k, respectively. Before calculations, the Pearson correlations were converted to z-values using the Fisher's r-to-z transformation

$$z_k(g_x, g_y) = \frac{1}{2}\ln\left(\frac{1 + r_k(g_x, g_y)}{1 - r_k(g_x, g_y)}\right). \tag{7}$$

This statistic is approximately normally distributed with variance $s_k^2 = (N_k - 3)^{-1}$, where N_k is the profile length in experiment k [3]. The similarity was then estimated using the weighted average

$$z_R(g_x, g_y) = \frac{\sum_{k=1}^{K} w_k(g_x, g_y) \cdot z_k(g_x, g_y)}{\sum_{k=1}^{K} w_k(g_x, g_y)}. \quad (8)$$

The weights were defined as $w_k(g_x, g_y) = (s_k^2 + \tau^2(g_x, g_y))^{-1}$, where the coefficient $\tau^2(g_x, g_y)$ was estimated using the Cochran's Q-statistic as described in [6]. Finally, the Fisher's z-to-r transformation was applied to convert the results back into correlations

$$r_R(g_x, g_y) = \frac{e^{2z_R(g_x,g_y)} - 1}{e^{2z_R(g_x,g_y)} + 1}. \quad (9)$$

2.3 Quality-Weighted Expression Similarity

We determined the quality of a particular probeset g_x in experiment k based on its L_k probe-level expression profiles $\mathbf{x}_k^{(l)} = (x_{k1}^{(l)}, \ldots, x_{kN_k}^{(l)})$, $l = 1, \ldots, L_k$. The Pearson correlation was first calculated separately between each possible probe pair within a probeset, and the quality of the particular probeset was then estimated by averaging over the pairs

$$q_k(g_x) = \frac{2}{L_k(L_k - 1)} \sum_{i<j} r(\mathbf{x}_k^{(i)}, \mathbf{x}_k^{(j)}). \quad (10)$$

To consider the effect of the profile lengths, we transformed the obtained averages into p-values using the result that the statistic

$$t_k(g_x) = q_k(g_x) \cdot \sqrt{\frac{N_k - 2}{1 - q_k(g_x)^2}} \quad (11)$$

has a Student's t-distribution with $df = N_k - 2$ under the null hypothesis of no correlation [3]. The estimated value $1 - p_k(g_x)$ was then used as the quality weight of the probeset in experiment k. To incorporate the probeset qualities into the similarity estimation, we modified the meta-analytic measure (9) by using $u_k^2(g_x, g_y) = [(N_k - 3)(1 - p_k(g_x))(1 - p_k(g_y))]^{-1}$ instead of s_k^2 when calculating the weights in (8). The obtained quality-weighted meta-correlation between two genes was denoted by $r_W(g_x, g_y)$. The novel correlation r_W was compared with the previously applied meta-correlations r_M, r_A and r_R, as well as with the single-experiment correlations r_k when identifying coexpressed genes. With the meta-correlation r_C, only the largest datasets E2 and E4 produced any gene pairs with false discovery rate FDR$<$0.01 and hence it was not truly applicable in our study setting.

3 Results

3.1 Effect of Noise on Constructing Coexpression Networks

We tested whether the constructed gene coexpression networks were stable with respect to noise. Increasing levels of multiplicative Gaussian noise was added to the log-scaled signal intensities of a selected dataset. If we denote by x the original expression intensity of a particular

gene, then the corresponding noise-added expression intensity was defined as $\tilde{x} = x \cdot N(1, \alpha^2)$. Here, the coefficient $\alpha = 0.01, 0.05, 0.1, 0.2, 0.5$ determines the noise level. The effect of noise was evaluated by considering the sensitivity and specificity of the identifications in the noised data as compared to the identifications in the original data [24]. The sensitivity was defined as TP/(TP+FN) and the specificity as TN/(TN+FP), where the true positives TP were the coexpressed pairs selected in both the noised and the original data, the true negatives TN the pairs selected in neither of the data, the false negatives FN the pairs detected only in the original data, and the false positives FP the pairs detected only in the noised data. The results were summarized using the receiver operating characteristic (ROC) curves, which captured the performance of the methods across a range of correlation thresholds in the noise-added data. In the original data, a gene pair was considered as coexpressed if its absolute correlation was higher than 0.6 and it was among the top 1% of the absolute correlations. The cutoff 0.6 has been widely used in gene coexpression studies [27, 29] and corresponded to FDR<0.01 in the meta-correlation calculations as assessed by permuting the sample labels.

Figure 1 illustrates the performance of the different correlation measures in the presence of noise in experiment E5. The results obtained when adding noise to the other single experiments were similar (data not shown). The performance of the meta-analytic similarities r_R and r_W outperformed the other similarity measures. The novel quality-weighted correlation r_W showed consistently the best performance. The performance of the merged correlation r_M and the average correlation r_A decreased rapidly with the increasing noise level. As expected, the single-experiment identifications r_k were most affected by noise.

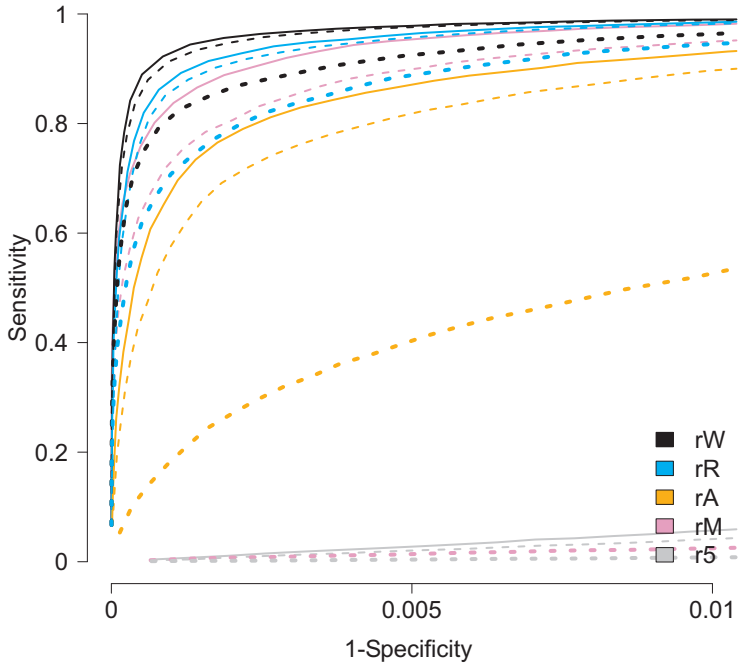

Figure 1: Performance of the different correlation measures when increasing levels of multiplicative Gaussian noise was added to the single experiment E5 (4 conditions). Average ROC curves over 10 datasets are shown at noise levels $\alpha = 0.01$ (solid lines), $\alpha = 0.2$ (dashed lines), and $\alpha = 0.5$ (dotted lines). The standard deviations were negligible and were omitted for the clarity of illustration. When comparing the traces, the one closest to the upper left corner shows the best performance. The results are shown for the individual correlation in experiment E5 (r5), the merged correlation r_M (rM), the average correlation r_A (rA), the random effects meta-correlation r_R (rR), and the novel quality-weighted correlation r_W (rW).

3.2 Coexpression Estimation Based on Short Profiles

To investigate the performance of the meta-correlations when only short expression profiles are available, we reduced the number of experiments and compared the identified coexpression pairs to those found in the meta-analysis of the two largest experiments (E2 and E4). The performance of the methods was evaluated by considering the sensitivity and precision of the identifications. Using the notation in the previous section, the precision was defined as TP/(TP+FP).

Figure 2 shows the performance of the different meta-correlation measures when 2-5 experiments with the shortest profile lengths were included into the analysis at the absolute correlation cutoff of 0.6. Cutoffs 0.7, 0.8, or 0.9 gave similar relative results (data not shown). The meta-analytic similarities r_R and r_W outperformed the other similarity measures in sensitivity. The highest sensitivity in each case was obtained with the quality-weighted meta-correlation r_W. The worst performing meta-correlation was the average correlation r_A. The relatively low level of precision of all the methods, however, demonstrates the limitations of using such short expression profiles alone.

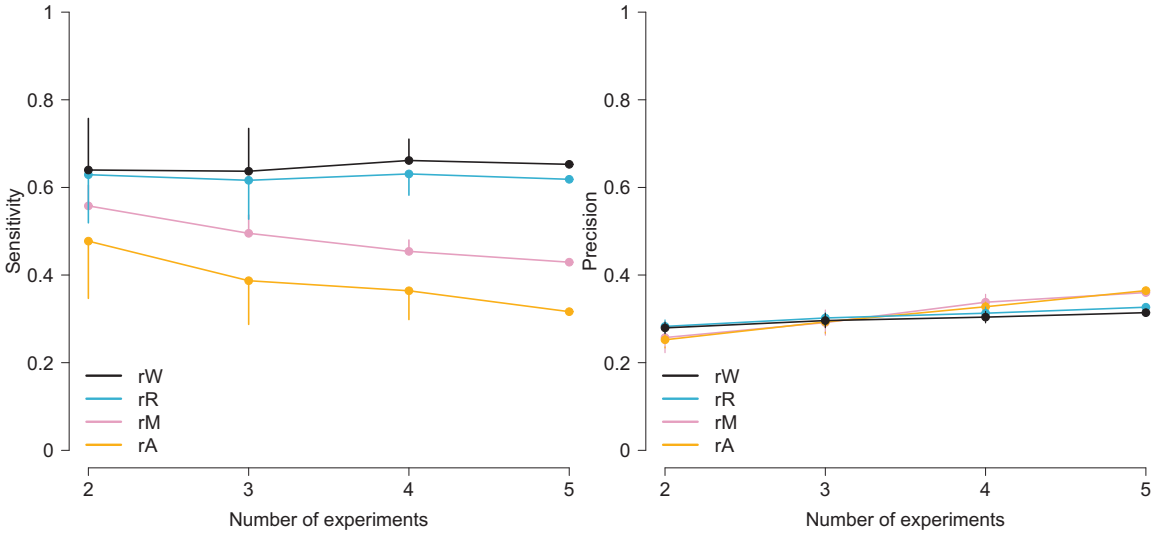

Figure 2: Sensitivity and precision as a function of the number of experiments. The meta-correlations over short-profile experiments were compared to the meta-correlations identified in the meta-analysis of the two largest experiments E2 and E4. The graphs show the average sensitivity (left graph) and the average precision (right graph) over several meta-analyses of 2-5 short-profile experiments at the absolute correlation cutoff of 0.6. The error bars represent the standard deviations.

3.3 The Reconstructed Network

To gain further insights into the coexpression networks predicted with the novel quality-weighted meta-correlation measure, we applied hierarchical clustering to the matrix representation of the network [18]. We restricted the clustering to those 596 genes showing absolute correlations over 0.9. The left graph of Figure 3 visualizes the identified clusters. We also investigated the enrichment of Gene Ontology terms [2] in the core clusters along the diagonal by using a hypergeometric test implemented in the Bioconductor 'GOstats' package [12]. Interestingly, the

largest cluster in the lower left corner of the graph showed significant enrichment of immune response genes ($p < 0.001$), which are known to be involved in T helper cell differentiation. The large cluster in the upper right corner contained genes involved in various levels of metabolism.

To investigate the network inferences of the quality-weighted meta-correlation at a gene level, we examined the identified neighbors of the transcription factor Gata3, which is known to play a key role in the normal T cell development and Th2 differentiation [20]. The right graph of Figure 3 illustrates the Gata3 subnetwork and its 14 first neighbors with absolute correlation over 0.6. The inferred subnetwork showed high connectivity, indicating the reliability of detection [4]. Moreover, the identified coexpressions included genes that are known to be involved in Th differentiation, such as Il4ra and Jak2 [20], suggesting the biological relevance of the network. As comparison, the merged correlation identified only one coexpressed gene with the same correlation threshold. The biological findings of this study will be further investigated and presented elsewhere.

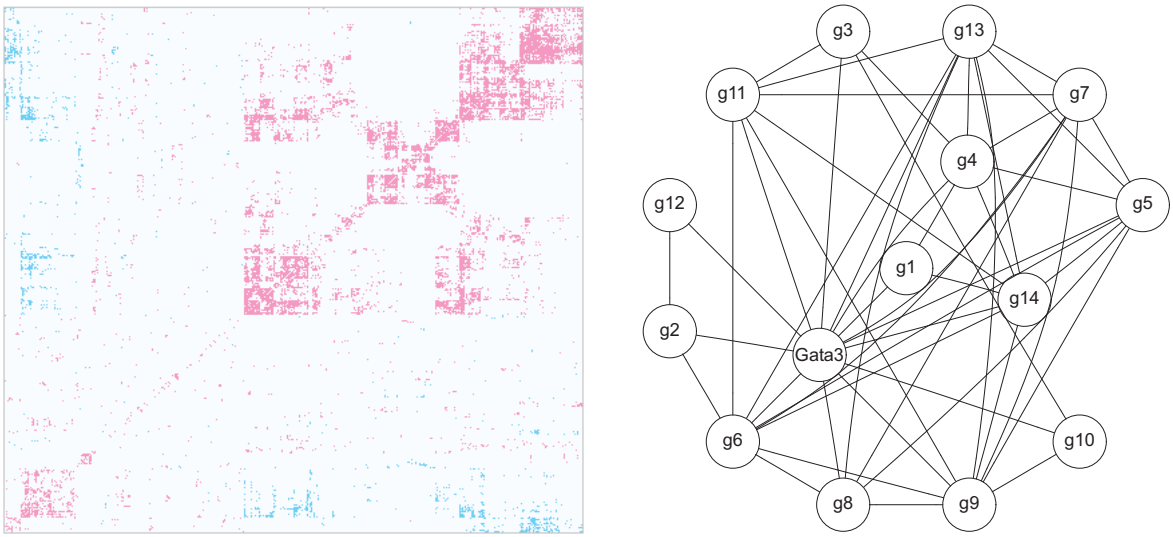

Figure 3: Visualization of the reconstructed coexpression network in the mouse T helper cell differentiation data. The left graph shows the hierarchical clustering of the coexpression network determined at the absolute correlation cutoff of 0.9. The rows and columns correspond to genes. Positive and negative correlations are indicated by red and blue, respectively. The right graph illustrates the subnetwork containing the transcription factor Gata3 and its first neighbors. The genes were associated an edge if their absolute correlation was over 0.6.

4 Conclusion

We have introduced a novel probe-level quality-weighted correlation to measure the similarity between gene expression profiles. The present results suggest that the proposed measure is robust against noise in individual experiments and applicable even when only short expression profiles are available. The novel approach outperformed the single-experiment correlations and the conventional meta-correlation measures which ignore the characteristics of the individual experiments. More specifically, the merged correlation r_M was highly sensitive to single

outliers; the average correlation r_A suffered from unreliable short-profile inferences; and the re-occurrence count r_C could not utilize the information from the short profiles at all, decreasing its applicability in many practical microarray studies. In general, the quality-weighted meta-correlation r_W performed at least as well as the previously introduced meta-correlation r_R, being slightly better in most cases. Moreover, the biological relevance of the inferred coexpression networks in the mouse T helper cell differentiation data imply the good performance of the novel approach. The biological findings of the networks may expand our understanding of the T helper cell differentiation.

Acknowledgements

We thank Riikka Lund, Zhi Chen, Sanna Filen and Pekka Ojala for providing the microarray data used in this study. The work was supported by the Academy of Finland (grant 203632) and the Graduate School in Computational Biology, Bioinformatics, and Biometry (ComBi).

References

[1] A. Aggarwal, D.L. Guo, Y. Hoshida, S.T. Yuen, K.M. Chu, S. So et al. Topological and Functional Discovery in a Gene Coexpression Meta-Network of Gastric Cancer. *Cancer Research*, 66(1):232-241, 2006.

[2] M. Ashburner, C.A. Ball, J.A. Blake, D. Botstein, H. Butler, J.M. Cherry et al. Gene Ontology: Tool for the Unification of Biology. *Nature Genetics*, 25(1):25-29, 2000.

[3] S. Bernstein and R. Bernstein. *Schaum's Outline of Theory and Problems of Elements of Statistics II*. McGraw-Hill, USA, 1999.

[4] A.J. Butte and I.S. Kohane. Mutual Information Relevance Networks: Functional Genomic Clustering Using Pairwise Entropy Measurements. *Pacific Symposium on Biocomputing*, 5:418-429, 2000.

[5] Z. Chen, R. Lund, T. Aittokallio, M. Kosonen, O. Nevalainen and R. Lahesmaa. Identification of Novel IL-4/Stat6-Regulated Genes in T Lymphocytes. *The Journal of Immunology*, 171(7):3627-3635, 2003.

[6] J.K. Choi, U. Yu, S. Kim and O.J. Yoo. Combining Multiple Microarray Studies and Modeling Interstudy Variation. *Bioinformatics*, 19:i84-i90, 2003.

[7] J.K. Choi, U. Yu, O.J. Yoo and S. Kim. Differential Coexpression Analysis Using Microarray Data and Its Application to Human Cancer. *Bioinformatics*, 21(24):4348-4355, 2005.

[8] L.L. Elo, L. Lahti, H. Skottman, M. Kyläniemi, R. Lahesmaa and T. Aittokallio. Integrating Probe-Level Expression Changes across Generations of Affymetrix Arrays. *Nucleic Acids Research*, 33(22):e193, 2005.

[9] 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.

[10] R.C. Gentleman, V.J. Carey, D.M. Bates, B. Bolstad, M. Dettling, S. Dudoit et al. Bioconductor: Open Software Development for Computational Biology and Bioinformatics. *Genome Biology*, 5(10):R80, 2004.

[11] R. Gentleman, B. Ding, S. Dudoit and J. Ibrahim. Distance Measures in DNA Microarray Data Analysis. In: R. Gentleman, V.J. Carey, W. Huber, R.A. Irizarry and S. Dudoit (eds.): *Bioinformatics and Computational Biology Solutions Using R and Bioconductor. Statistics for Biology and Health.* Springer, New York, 189-208, 2005.

[12] R. Gentleman, D. Scholtens, B. Ding, V.J. Carey and W. Huber. Case Studies Using Graphs on Biological Data. In: R. Gentleman, V.J. Carey, W. Huber, R.A. Irizarry and S. Dudoit (eds.): *Bioinformatics and Computational Biology Solutions Using R and Bioconductor. Statistics for Biology and Health.* Springer, New York, 369-394, 2005.

[13] O.L. Griffith, E.D. Pleasance, D.L. Fulton, M. Oveisi, M. Ester, A.S. Siddiqui and S.J. Jones. Assessment and Integration of Publicly Available SAGE, cDNA Microarray, and Oligonucleotide Microarray Expression Data for Global Coexpression Analyses. *Genomics*, 86(4):476-488, 2005.

[14] L.V. Hedges and I. Olkin. *Statistical Methods for Meta-analysis*. Academic Press, Orlando, 1985.

[15] P. Hu, C.M. Greenwood and J. Beyene. Integrative Analysis of Multiple Gene Expression Profiles with Quality-Adjusted Effect Size Models. *BMC Bioinformatics*, 6:128, 2005.

[16] R.A. Irizarry, B. Hobbs, F. Collin, Y.D. Beazer-Barclay, K.J. Antonellis, U. Scherf and T.P. Speed. Exploration, Normalization, and Summaries of High Density Oligonucleotide Array Probe Level Data. *Biostatistics*, 4(2):249-264, 2003.

[17] J.E. Larkin, B.C. Frank, H. Gavras, R. Sultana and J. Quackenbush. Independence and Reproducibility across Microarray Platforms. *Nature Methods*, 2(5):337-344, 2005.

[18] H.K. Lee, A.K. Hsu, J. Sajdak, J. Qin and P. Pavlidis. Coexpression Analysis of Human Genes Across Many Microarray Data Sets. *Genome Research*, 14(6):1085-1094, 2004.

[19] D.J. Lockhart, H. Dong, M.C. Byrne, M.T. Follettie, M.V. Gallo, M.S. Chee et al. Expression Monitoring by Hybridization to High-Density Oligonucleotide Arrays. *Nature Biotechnology*, 14(13):1675-1680, 1996.

[20] R. Lund. *Identification of Novel Genes Involved in the Early Differentiation of Th1 and Th2 Cells*. PhD Thesis, University of Turku, 2004.

[21] R. Lund, Z. Chen, J. Scheinin and R. Lahesmaa. Early Target Genes of IL-12 and STAT4 Signaling in Th Cells. *The Journal of Immunology*, 172(11):6775-6782, 2004.

[22] P.M. Magwene and J. Kim. Estimating Genomic Coexpression Networks Using First-Order Conditional Independence. *Genome Biology*, 5(12):R100, 2004.

[23] B.H. Mecham, G.T. Klus, J. Strovel, M. Augustus, D. Byrne, P. Bozso et al. Sequence-Matched Probes Produce Increased Cross-Platform Consistency and More Reproducible Biological Results in Microarray-Based Gene Expression Measurements. *Nucleic Acids Reseach*, 32(9):e74, 2004.

[24] J.M. Stuart, E. Segal, D. Koller and S.K. Kim. A Gene-Coexpression Network for Global Discovery of Conserved Genetic Modules. *Science*, 302(5643):249-255, 2003.

[25] L. Tari, C. Baral and P. Dasgupta. Understanding the Global Properties of Functionally-Related Gene Networks Using the Gene Ontology. *Pacific Symposium on Biocomputing*, 10:209-220, 2005.

[26] C.J. Wolfe, I.S. Kohane and A.J. Butte. Systematic Survey Reveals General Applicability of "Guilt-by-Association" within Gene Coexpression Networks. *BMC Bioinformatics*, 6:227, 2005.

[27] X. Zhou, M.C. Kao and W.H. Wong. Transitive Functional Annotation by Shortest Path Analysis of Gene Expression Data. *Proceedings of the National Academy of Sciences of the United States of America*, 99(20):12783-12788, 2002.

[28] X.J. Zhou, M.C. Kao, H. Huang, A. Wong, J. Nunez-Iglesias, M. Primig et al. Functional Annotation and Network Reconstruction through Cross-Platform Integration of Microarray Data. *Nature Biotechnology*, 23(2):238-243, 2005.

[29] D. Zhu, A.O. Hero, Z.S. Qin and A. Swaroop. High Throughput Screening of Co-Expressed Gene Pairs with Controlled False Discovery Rate (FDR) and Minimum Acceptable Strength (MAS). *Journal of Computational Biology*, 12(7):1029-1045, 2005.