

Cluster analysis and display of genome-wide expression patterns

MICHAEL B. EISEN*, PAUL T. SPELLMAN*, PATRICK O. BROWN†, AND DAVID BOTSTEIN*‡

*Department of Genetics and †Department of Biochemistry and Howard Hughes Medical Institute, Stanford University School of Medicine, 300 Pasteur Avenue, Stanford, CA 94305

Contributed by David Botstein, October 13, 1998

ABSTRACT A system of cluster analysis for genome-wide expression data from DNA microarray hybridization is described that uses standard statistical algorithms to arrange genes according to similarity in pattern of gene expression. The output is displayed graphically, conveying the clustering and the underlying expression data simultaneously in a form intuitive for biologists. We have found in the budding yeast *Saccharomyces cerevisiae* that clustering gene expression data groups together efficiently genes of known similar function, and we find a similar tendency in human data. Thus patterns seen in genome-wide expression experiments can be interpreted as indications of the status of cellular processes. Also, coexpression of genes of known function with poorly characterized or novel genes may provide a simple means of gaining leads to the functions of many genes for which information is not available currently.

The rapid advance of genome-scale sequencing has driven the development of methods to exploit this information by characterizing biological processes in new ways. The knowledge of the coding sequences of virtually every gene in an organism, for instance, invites development of technology to study the expression of all of them at once, because the study of gene expression of genes one by one has already provided a wealth of biological insight. To this end, a variety of techniques has evolved to monitor, rapidly and efficiently, transcript abundance for all of an organism's genes (1–3). Within the mass of numbers produced by these techniques, which amount to hundreds of data points for thousands or tens of thousands of genes, is an immense amount of biological information. In this paper we address the problem of analyzing and presenting information on this genomic scale.

A natural first step in extracting this information is to examine the extremes, e.g., genes with significant differential expression in two individual samples or in a time series after a given treatment. This simple technique can be extremely efficient, for example, in screens for potential tumor markers or drug targets. However, such analyses do not address the full potential of genome-scale experiments to alter our understanding of cellular biology by providing, through an inclusive analysis of the entire repertoire of transcripts, a continuing comprehensive window into the state of a cell as it goes through a biological process. What is needed instead is a holistic approach to analysis of genomic data that focuses on illuminating order in the entire set of observations, allowing biologists to develop an integrated understanding of the process being studied.

A natural basis for organizing gene expression data is to group together genes with similar patterns of expression. The first step to this end is to adopt a mathematical description of similarity. For any series of measurements, a number of sensible measures of similarity in the behavior of two genes can

be used, such as the Euclidean distance, angle, or dot products of the two n -dimensional vectors representing a series of n measurements. We have found that the standard correlation coefficient (i.e., the dot product of two normalized vectors) conforms well to the intuitive biological notion of what it means for two genes to be “coexpressed;” this may be because this statistic captures similarity in “shape” but places no emphasis on the magnitude of the two series of measurements.

It is not the purpose of this paper to survey the various methods available to cluster genes on the basis of their expression patterns, but rather to illustrate how such methods can be useful to biologists in the analysis of gene expression data. We aim to use these methods to organize, but not to alter, tables containing primary data; we have thus used methods that can be reduced, in the end, to a reordering of lists of genes. Clustering methods can be divided into two general classes, designated supervised and unsupervised clustering (4). In supervised clustering, vectors are classified with respect to known reference vectors. In unsupervised clustering, no predefined reference vectors are used. As we have little *a priori* knowledge of the complete repertoire of expected gene expression patterns for any condition, we have favored unsupervised methods or hybrid (unsupervised followed by supervised) approaches.

Although various clustering methods can usefully organize tables of gene expression measurements, the resulting ordered but still massive collection of numbers remains difficult to assimilate. Therefore, we always combine clustering methods with a graphical representation of the primary data by representing each data point with a color that quantitatively and qualitatively reflects the original experimental observations. The end product is a representation of complex gene expression data that, through statistical organization and graphical display, allows biologists to assimilate and explore the data in a natural intuitive manner.


To illustrate this approach, we have applied pairwise average-linkage cluster analysis (5) to gene expression data collected in our laboratories. This method is a form of hierarchical clustering, familiar to most biologists through its application in sequence and phylogenetic analysis. Relationships among objects (genes) are represented by a tree whose branch lengths reflect the degree of similarity between the objects, as assessed by a pairwise similarity function such as that described above. In sequence comparison, these methods are used to infer the evolutionary history of sequences being compared. Whereas no such underlying tree exists for expression patterns of genes, such methods are useful in their ability to represent varying degrees of similarity and more distant relationships among groups of closely related genes, as well as in requiring few assumptions about the nature of the data. The computed trees can be used to order genes in the original data table, so that genes or groups of genes with similar expression patterns are adjacent. The ordered table can then be displayed graphically, as above, with a representation of the tree to indicate the relationships among genes.

The publication costs of this article were defrayed in part by page charge payment. This article must therefore be hereby marked “advertisement” in accordance with 18 U.S.C. §1734 solely to indicate this fact.

MATERIALS AND METHODS

Sources of Experimental Data. Data analyzed here were collected on spotted DNA microarrays (6, 7). Gene expression

pairs of genes. The matrix is scanned to identify the highest value (representing the most similar pair of genes). A node is created joining these two genes, and a gene expression profile



regulation. Finally, this result also indicates that noise present in single observations does not contribute significantly when genes are compared across even a relatively small number of nonidentical conditions. Therefore, when designing experiments, it may be more valuable to sample a wide variety of conditions than to make repeat observations on identical conditions.

Genes of Similar Function Cluster Together. A far more striking result is found when larger groups of clustered genes are examined, where we observe a strong tendency for these genes to share common roles in cellular processes. This relationship is clearest in data from experiments on the budding yeast *S. cerevisiae*, where arrays representing essentially all of the genes from this organism are available (8) and for which a large fraction of the identified genes (more than 35%) have been studied in some detail. Fig. 2A represents a clustering analysis of 2,467 genes, all the genes that currently have a functional annotation in the *Saccharomyces* Genome Database (12). As can be seen in Fig. 2B–K, numerous groups of coexpressed genes representing diverse expression patterns across the sampled conditions are involved in common cellular

Gene	Function	Gene	Function
SPU2	CYTOSKELETON	SPY2	SPY2
MEK1	MAP KINASE	SPY3	SPY3
MEK2	CYTOSKELETON	SPY4	SPY4
SPC42	CYTOSKELETON	SPY5	SPY5
CHS17	CYTOSKELETON	SPY6	SPY6
CEM1	CELL CYCLE	SPY7	SPY7
CDC10	CYTOSKELETON	SPY8	SPY8
CDC11	CYTOSKELETON	SPY9	SPY9
CDC12	CELL CYCLE	SPY10	SPY10
CDC13	CELL CYCLE	SPY11	SPY11
CDC14	CELL CYCLE	SPY12	SPY12
CDC15	CELL CYCLE	SPY13	SPY13
SPY14	PROTEIN DEGRADATION	SPY14	SPY14
SPY15	PROTEIN DEGRADATION	SPY15	SPY15
SPY16	PROTEIN DEGRADATION	SPY16	SPY16
SPY17	PROTEIN DEGRADATION	SPY17	SPY17
SPY18	PROTEIN DEGRADATION	SPY18	SPY18
SPY19	PROTEIN DEGRADATION	SPY19	SPY19
SPY20	PROTEIN DEGRADATION	SPY20	SPY20
SPY21	PROTEIN DEGRADATION	SPY21	SPY21
SPY22	PROTEIN DEGRADATION	SPY22	SPY22
SPY23	PROTEIN DEGRADATION	SPY23	SPY23
SPY24	PROTEIN DEGRADATION	SPY24	SPY24
SPY25	PROTEIN DEGRADATION	SPY25	SPY25
SPY26	PROTEIN DEGRADATION	SPY26	SPY26
SPY27	PROTEIN DEGRADATION	SPY27	SPY27
SPY28	PROTEIN DEGRADATION	SPY28	SPY28
SPY29	PROTEIN DEGRADATION	SPY29	SPY29
SPY30	PROTEIN DEGRADATION	SPY30	SPY30
SPY31	PROTEIN DEGRADATION	SPY31	SPY31
SPY32	PROTEIN DEGRADATION	SPY32	SPY32
SPY33	PROTEIN DEGRADATION	SPY33	SPY33
SPY34	PROTEIN DEGRADATION	SPY34	SPY34
SPY35	PROTEIN DEGRADATION	SPY35	SPY35
SPY36	PROTEIN DEGRADATION	SPY36	SPY36
SPY37	PROTEIN DEGRADATION	SPY37	SPY37
SPY38	PROTEIN DEGRADATION	SPY38	SPY38
SPY39	PROTEIN DEGRADATION	SPY39	SPY39
SPY40	PROTEIN DEGRADATION	SPY40	SPY40
SPY41	PROTEIN DEGRADATION	SPY41	SPY41
SPY42	PROTEIN DEGRADATION	SPY42	SPY42
SPY43	PROTEIN DEGRADATION	SPY43	SPY43
SPY44	PROTEIN DEGRADATION	SPY44	SPY44
SPY45	PROTEIN DEGRADATION	SPY45	SPY45
SPY46	PROTEIN DEGRADATION	SPY46	SPY46
SPY47	PROTEIN DEGRADATION	SPY47	SPY47
SPY48	PROTEIN DEGRADATION	SPY48	SPY48
SPY49	PROTEIN DEGRADATION	SPY49	SPY49
SPY50	PROTEIN DEGRADATION	SPY50	SPY50
SPY51	PROTEIN DEGRADATION	SPY51	SPY51
SPY52	PROTEIN DEGRADATION	SPY52	SPY52
SPY53	PROTEIN DEGRADATION	SPY53	SPY53
SPY54	PROTEIN DEGRADATION	SPY54	SPY54
SPY55	PROTEIN DEGRADATION	SPY55	SPY55
SPY56	PROTEIN DEGRADATION	SPY56	SPY56
SPY57	PROTEIN DEGRADATION	SPY57	SPY57
SPY58	PROTEIN DEGRADATION	SPY58	SPY58
SPY59	PROTEIN DEGRADATION	SPY59	SPY59
SPY60	PROTEIN DEGRADATION	SPY60	SPY60
SPY61	PROTEIN DEGRADATION	SPY61	SPY61
SPY62	PROTEIN DEGRADATION	SPY62	SPY62
SPY63	PROTEIN DEGRADATION	SPY63	SPY63
SPY64	PROTEIN DEGRADATION	SPY64	SPY64
SPY65	PROTEIN DEGRADATION	SPY65	SPY65
SPY66	PROTEIN DEGRADATION	SPY66	SPY66
SPY67	PROTEIN DEGRADATION	SPY67	SPY67
SPY68	PROTEIN DEGRADATION	SPY68	SPY68
SPY69	PROTEIN DEGRADATION	SPY69	SPY69
SPY70	PROTEIN DEGRADATION	SPY70	SPY70
SPY71	PROTEIN DEGRADATION	SPY71	SPY71
SPY72	PROTEIN DEGRADATION	SPY72	SPY72
SPY73	PROTEIN DEGRADATION	SPY73	SPY73
SPY74	PROTEIN DEGRADATION	SPY74	SPY74
SPY75	PROTEIN DEGRADATION	SPY75	SPY75
SPY76	PROTEIN DEGRADATION	SPY76	SPY76
SPY77	PROTEIN DEGRADATION	SPY77	SPY77
SPY78	PROTEIN DEGRADATION	SPY78	SPY78
SPY79	PROTEIN DEGRADATION	SPY79	SPY79
SPY80	PROTEIN DEGRADATION	SPY80	SPY80
SPY81	PROTEIN DEGRADATION	SPY81	SPY81
SPY82	PROTEIN DEGRADATION	SPY82	SPY82
SPY83	PROTEIN DEGRADATION	SPY83	SPY83
SPY84	PROTEIN DEGRADATION	SPY84	SPY84
SPY85	PROTEIN DEGRADATION	SPY85	SPY85
SPY86	PROTEIN DEGRADATION	SPY86	SPY86
SPY87	PROTEIN DEGRADATION	SPY87	SPY87
SPY88	PROTEIN DEGRADATION	SPY88	SPY88
SPY89	PROTEIN DEGRADATION	SPY89	SPY89
SPY90	PROTEIN DEGRADATION	SPY90	SPY90
SPY91	PROTEIN DEGRADATION	SPY91	SPY91
SPY92	PROTEIN DEGRADATION	SPY92	SPY92
SPY93	PROTEIN DEGRADATION	SPY93	SPY93
SPY94	PROTEIN DEGRADATION	SPY94	SPY94
SPY95	PROTEIN DEGRADATION	SPY95	SPY95
SPY96	PROTEIN DEGRADATION	SPY96	SPY96
SPY97	PROTEIN DEGRADATION	SPY97	SPY97
SPY98	PROTEIN DEGRADATION	SPY98	SPY98
SPY99	PROTEIN DEGRADATION	SPY99	SPY99
SPY100	PROTEIN DEGRADATION	SPY100	SPY100

applied to all of the approximately 6,200 genes of *S. cerevisiae*, the clusters of functionally related genes are maintained, but are usually expanded with the addition of uncharacterized genes (the results of this analysis will be the subject of a subsequent report)

2. Velculescu, V. E., Zhang, L., Vogelstein, B. & Kinzler, K. W. (1995) *Science* 270, 484–487.
3. Lockhart, D. J., Dong, H., Byrne, M. C., Follettie, M. T., Gallo, M. V., Chee, M. S., Mittmann, M., Wang, C., Kobayashi, M.