

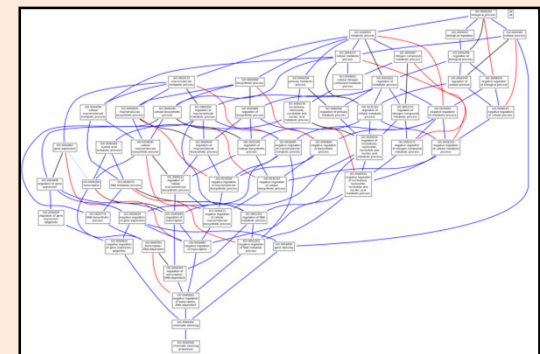

Gene Set Enrichment Analysis

Genome 559: Introduction to Statistical and
Computational Genomics


Elhanan Borenstein

A quick review

- Gene expression profiling
 - Which molecular processes/functions are involved in a certain phenotype (e.g., disease, stress response, etc.)
- The Gene Ontology (GO) Project
 - Provides shared vocabulary/annotation
 - GO terms are linked in a complex structure
- Enrichment analysis:
 - Find the “most” differentially expressed genes
 - Identify functional annotations that are **over-represented**
 - Modified Fisher's exact test




A quick review: Modified Fisher's exact test



Do I have a surprisingly high number of blue genes?

Null model: the 8 genes/balls are selected randomly





So, if you have 50 balls, 10 of them are blue, and you pick 8 balls randomly, what is the probability that k of them are blue?

A quick review: Modified Fisher's exact test

Hypergeometric distribution

$$\mathbb{P}(\sigma_t = k) = \frac{\binom{m_t}{k} \binom{m-m_t}{n-k}}{\binom{m}{n}}$$

$m=50, m_t=10, n=8$





So ... do I have a surprisingly high number of blue genes?

Can such high numbers (4 or above) occur by chance?

What is the probability of getting at least 4 blue genes in the null model?


$P(\sigma_t \geq 4)$

Enrichment Analysis



Enrichment Analysis

Genes ranked by expression correlation to Class A



Biological function?

Cutoff

Function 1
(e.g., metabolism)


2 / 10

Function 2
(e.g., signaling)

5 / 11

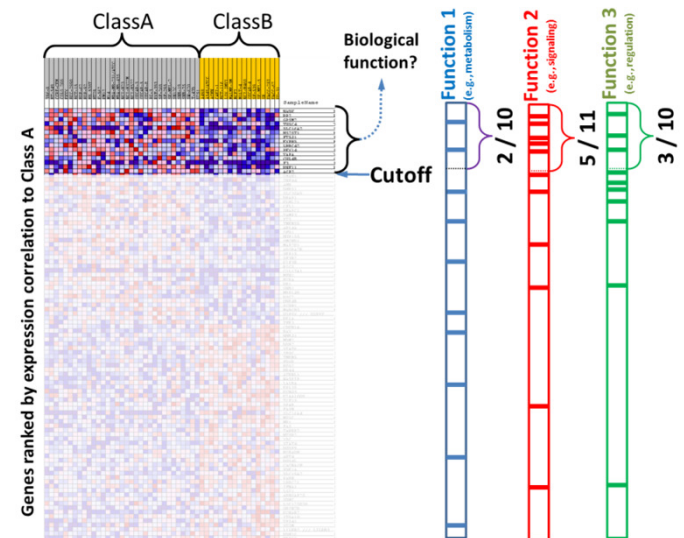
Function 3
(e.g., regulation)

3 / 10



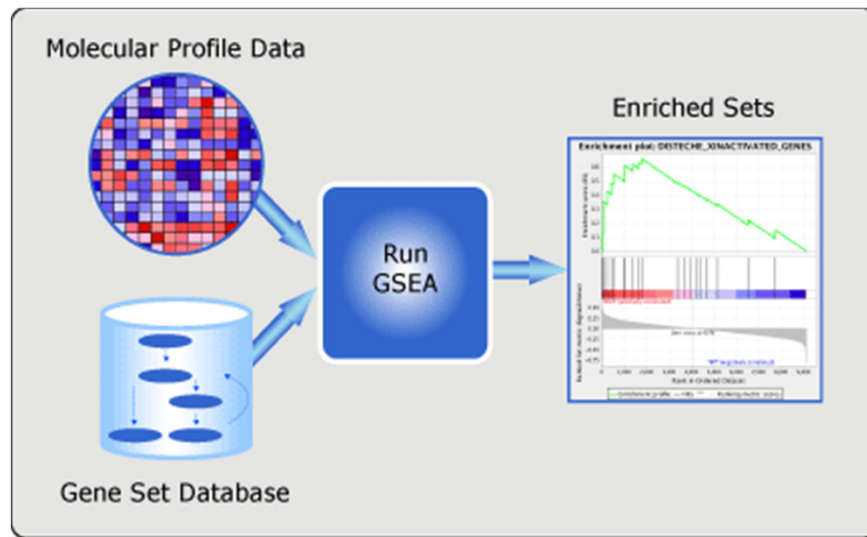
Problems with cutoff-based analysis

- After correcting for multiple hypotheses testing, no individual gene may meet the threshold due to noise.
- Alternatively, one may be left with a long list of significant genes without any unifying biological theme.
- The cutoff value is often arbitrary!
- **We are really examining only a handful of genes, totally ignoring much of the data**



Gene Set Enrichment Analysis

- MIT, Broad Institute
- V 2.0 available since Jan 2007



Gene set enrichment analysis: A knowledge-based approach for interpreting genome-wide expression profiles

Aravind Subramanian^{1,2}, Pablo Tamayo^{3,4}, Vamsi K. Mootha^{5,6}, Sayan Mukherjee⁷, Benjamin L. Ebert^{8,9}, Michael A. Gillette¹⁰, Amanda Paulovich¹¹, Scott L. Pomeroy¹², Todd R. Golub^{13,14}, Eric S. Lander^{15,16}, and Jill P. Mesirov^{17,18}

¹Broad Institute of Massachusetts Institute of Technology and Harvard, 320 Charles Street, Cambridge, MA 02141; ²Department of Systems Biology, Albert Einstein College of Medicine, 1400 Morris Avenue, Roseton, MA 02840; ³Institute for Genome Sciences and Policy, Center for Interdisciplinary Engineering, Medicine, and Applied Sciences, Duke University, 101 Science Drive, Durham, NC 27708; ⁴Department of Medical Oncology, Dana-Farber Cancer Institute, 44 Binney Street, Boston, MA 02115; ⁵Division of Pulmonary and Critical Care Medicine, Massachusetts General Hospital, 75 Fruit Street, Boston, MA 02114; ⁶Harvard Medical School, 77 Avenue Louis Pasteur, Boston, MA 02115; ⁷Department of Biology, Massachusetts Institute of Technology, Cambridge, MA 02139; ⁸Department of Pathology, Brigham Young University, Provo, UT 84602; ⁹Department of Pathology, Brigham Young University, Provo, UT 84602; ¹⁰Department of Pathology, Brigham Young University, Provo, UT 84602; ¹¹Department of Pathology, Brigham Young University, Provo, UT 84602; ¹²Department of Pathology, Brigham Young University, Provo, UT 84602; ¹³Department of Pathology, Brigham Young University, Provo, UT 84602; ¹⁴Department of Pathology, Brigham Young University, Provo, UT 84602; ¹⁵Department of Pathology, Brigham Young University, Provo, UT 84602; ¹⁶Department of Pathology, Brigham Young University, Provo, UT 84602; ¹⁷Department of Pathology, Brigham Young University, Provo, UT 84602; ¹⁸Department of Pathology, Brigham Young University, Provo, UT 84602

Contributed by Eric S. Lander, August 2, 2005

Although genome-wide RNA expression analysis has become a routine tool in biomedical research, extracting biological insight from such information remains a major challenge. Here, we describe a powerful analytical method called Gene Set Enrichment Analysis (GSEA) for interpreting gene expression data. The method derives its power by focusing on gene sets, that is, groups of genes that share common biological function, chromosomal location, or regulation. We demonstrate how GSEA yields insights into several cancer-related data sets, including leukemia and lung cancer. Notably, where single-gene analysis finds little similarity between two independent studies of patient survival in lung cancer, GSEA reveals many biological pathways in common. The GSEA method is embodied in a freely available software package, together with an initial database of 1,325 biologically defined gene sets.

microarray

Genome-wide expression analysis with DNA microarrays has become a mainstay of genomics research (1,2). The challenge no longer lies in obtaining gene expression profiles, but rather in interpreting the results to gain insights into biological mechanisms.

In a typical experiment, mRNA expression profiles are generated for thousands of genes from a collection of samples belonging to one of two classes, for example, tumors that are sensitive vs. resistant to a drug. The genes can be ordered in a ranked list L , according to their differential expression between the classes. The challenge is to extract meaning from this list.

A common approach involves focusing on a handful of genes at the top and bottom of L (i.e., those showing the largest difference) to discern salient biological clues. This approach has a few major limitations.

(i) After correcting for multiple hypotheses testing, no individual gene may meet the threshold for statistical significance, because the relevant biological differences are modest relative to the noise inherent to the microarray technology.

(ii) Alternatively, one may be left with a long list of statistically significant genes without any unifying biological theme. Interpretation can be daunting and ad hoc, being dependent on a biologist's area of expertise.

(iii) Single-gene analysis may miss important effects on pathways. Cellular processes often affect sets of genes acting in concert. An increase of 20% in all genes encoding members of a metabolic pathway may dramatically alter the flux through the pathway and may be more important than a 20-fold increase in a single gene.

(iv) When different groups study the same biological system, the list of statistically significant genes from the two studies may show distressingly little overlap (3).

To overcome these analytical challenges, we recently developed a method called Gene Set Enrichment Analysis (GSEA) that

evaluates microarray data at the level of gene sets. The gene sets are defined based on prior biological knowledge, e.g., published information about biochemical pathways or coexpression in previous experiments. The goal of GSEA is to determine whether members of a gene set S tend to occur toward the top (or bottom) of the list L , in which case the gene set is correlated with the phenotypic class distinction.

We used a preliminary version of GSEA to analyze data from muscle biopsies from diabetic vs. healthy controls (4). The method revealed that genes involved in oxidative phosphorylation show reduced expression in diabetes, although the average decrease per gene is only 20%. The results from this study have been independently validated by other microarray studies (5) and by *in vivo* functional studies (6).

Given this success, we have developed GSEA into a robust technique for analyzing molecular profiling data. We studied its characteristics and performance and substantially revised and generalized the original method for broader applicability.

In this paper, we provide a full mathematical description of the GSEA methodology and illustrate its utility by applying it to several diverse biological problems. We have also created a software package, called GSEA-P, and an initial inventory of gene sets (Molecular Signature Database, MSigDB), both of which are freely available.

Methods

Overview of GSEA. GSEA considers experiments with genome-wide expression profiles from samples belonging to two classes, labeled 1 or 2. Genes are ranked based on the correlation between their expression and the class distinction by using any suitable metric (Fig. 1A).

Given an *a priori* defined set of genes S (e.g., genes encoding products in a metabolic pathway, located in the same cytogenetic band, or sharing the same GO category), the goal of GSEA is to determine whether the members of S are randomly distributed throughout L or primarily found at the top or bottom. We expect

Freely available online through the PNAS open access option.

Abbreviations: AML, acute myeloid leukemia; ALL, acute lymphoid leukemia; AM, acute myeloid leukemia; ES, enrichment score; GSEA, Gene Set Enrichment Analysis; GSEA-P, Gene Set Enrichment Analysis; MSigDB, Molecular Signature Database; NES, normalized enrichment score.

See Commentary on page 15278.

*A.S. and P.T. contributed equally to this work.

*To whom correspondence may be addressed. E-mail: lander@broad.mit.edu or mesirov@broad.mit.edu.

© 2005 by The National Academy of Sciences of the USA.


(Subramanian et al. PNAS. 2005.)

GSEA key features

- Calculates a score for the enrichment of a **entire set of genes** rather than single genes!
- Does not require setting a cutoff!
- Identifies the set of relevant genes as part of the analysis!
- Provides a more robust statistical framework!

Gene Set Enrichment Analysis

Genes ranked by expression correlation to Class A



Biological function?

Cutoff

Function 1
(e.g., metabolism)


2 / 10

Function 2
(e.g., signaling)

5 / 11


Function 3
(e.g., regulation)

3 / 10



Gene Set Enrichment Analysis


Genes ranked by expression correlation to Class A



Function 1
(e.g., metabolism)

Function 2
(e.g., signaling)

Function 3
(e.g., regulation)

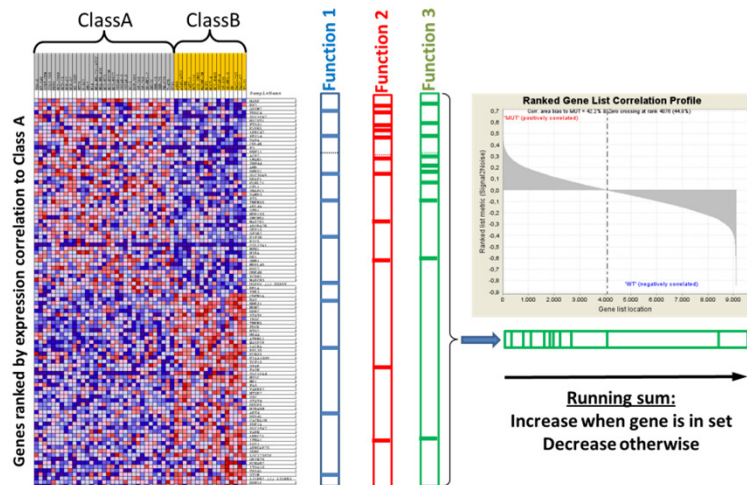


Running sum:
Increase when gene is in set
Decrease otherwise

Gene Set Enrichment Analysis

What would you expect if the hits were randomly distributed?

What would you expect if most of the hits cluster at the top of the list?


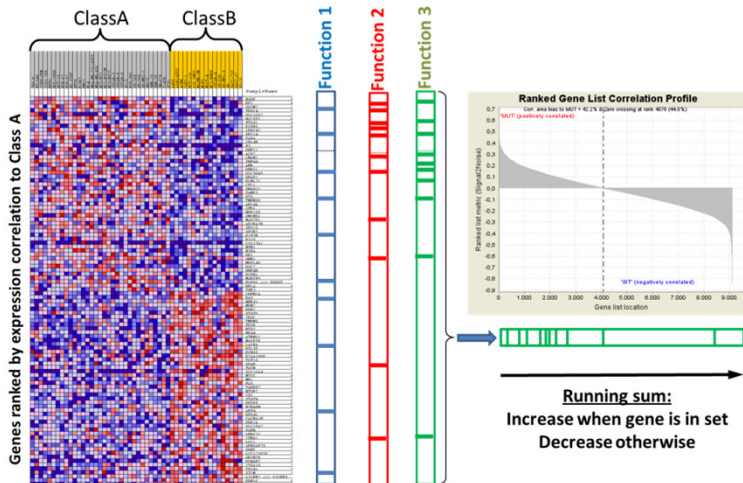


Gene Set Enrichment Analysis

Enrichment score (ES) =
max deviation from 0

Leading
Edge genes

Genes within
functional set
(hits)

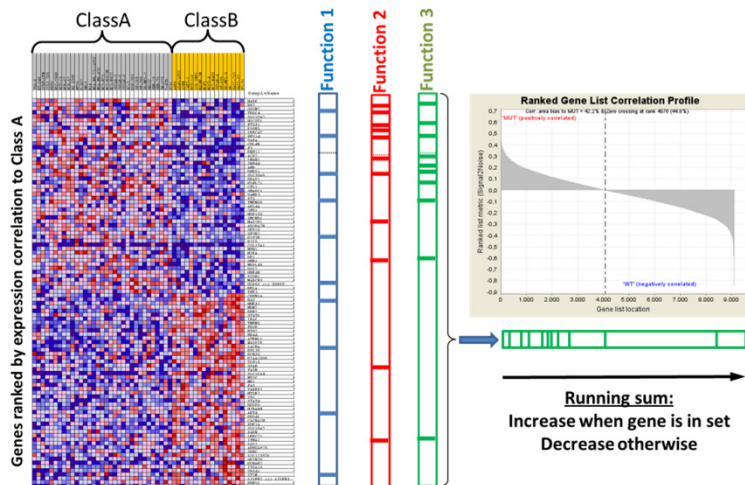





Gene Set Enrichment Analysis

ES = 0.43

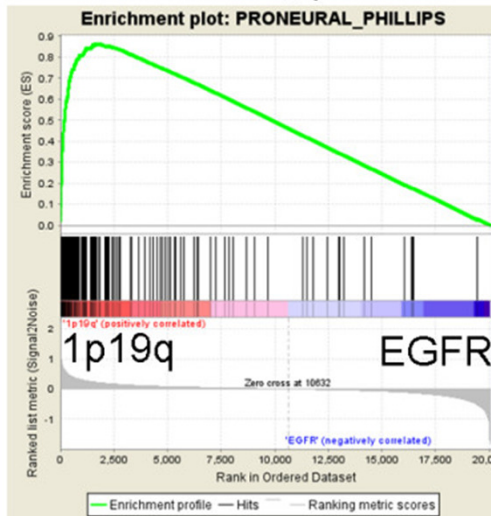
ES = -0.45

Low ES (evenly distributed)




Gene Set Enrichment Analysis

A ES=0.86, $p<0.001$



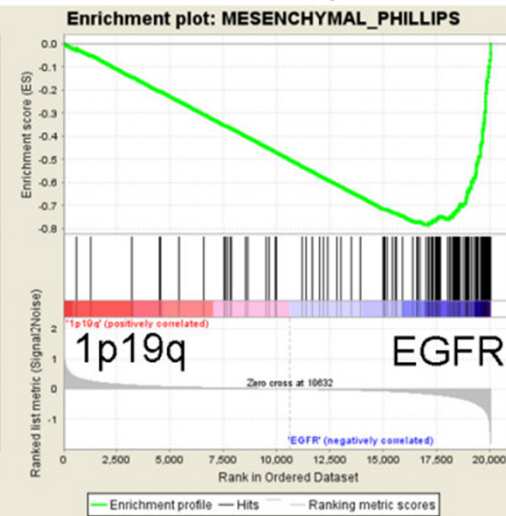
SNR

B ES= -0.79, $p<0.001$

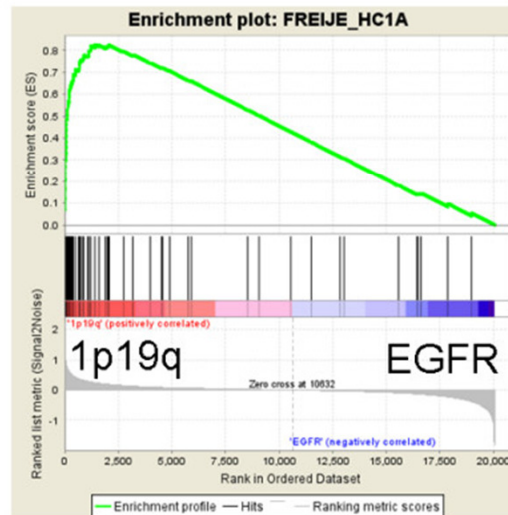


SNR

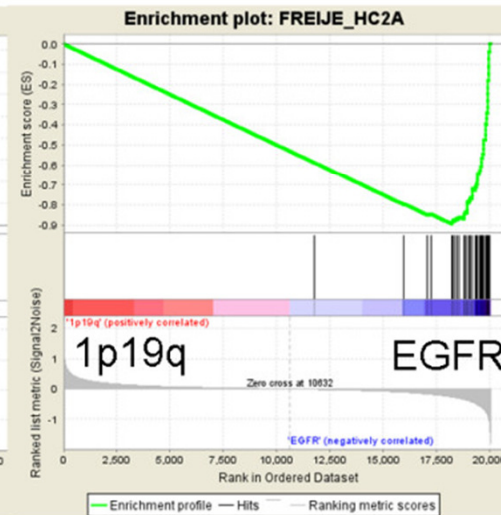
C ES= -0.78, $p<0.001$



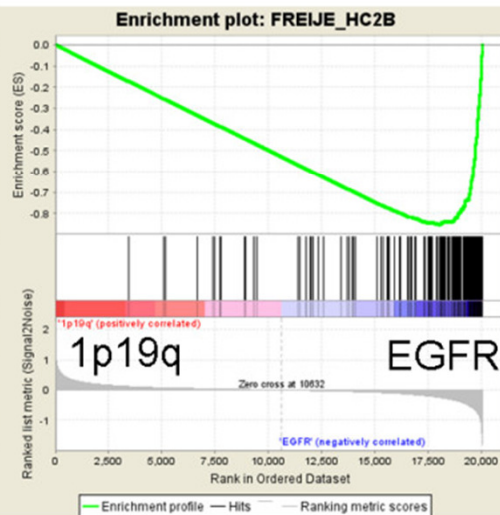
SNR



D ES=0.82, $p<0.001$




E ES= -0.89, $p<0.001$



F ES= -0.85, $p<0.001$


GSEA Steps

1. Calculation of an enrichment score (ES) for each functional category



2. Estimation of significance level of the ES

- An empirical permutation test
- Phenotype labels are shuffled and the ES for this functional set is recomputed. Repeat 1000 times.
- Generating a null distribution



3. Adjustment for multiple hypotheses testing

- Necessary if comparing multiple gene sets (i.e., functions)
- Computes FDR (false discovery rate)

