Note: This paper has been modified from its original form. Interest in the original should be directed to Peter Clote, clote@bc.edu or Jay Cahill, wjcahill@hotmail.com.

Part I: Motivation and Results

This paper is organized into three parts. The first part is intended to define the motivation and idea behind the project and present the results. The second part is intended as a user guide and documentation to enable the future maintenance and extension of the CAST software.
The third part offers conclusions and ideas for further study.

Section 1. Introduction

time series or under varying conditions can improve the understanding of a gene's function in an organism.

Advances in gene microarray technology allow scientists to produce experimental data at a much faster rate than data analysis can be performed by humans; this field is often referred to as "high-throughput" biology. Large data sets lend themselves to computer analysis and visualization.

One method of computer analysis of microarray data is clustering. Clustering is a machine learning technique that defines clusters of things with common attributes. Clustering as an approach to the analysis of microarray data groups together genes with "similar" expression levels across conditions or during cellular processes. This paper will focus on the analysis of gene microarray data using a clustering technique due to Ben-Dor *et al.*[3]

Clustering is not a new idea; clustering algorithms have been used in the field of statistics for decades[4]. Clustering problem solving approaches are used in many fields including economics, medicine, biology, and many others, but the application of clustering algorithms to microarray data for the purpose of functional genomics, cell classification and disease diagnostics is fairly recent. Several different clustering algorithms have been applied to gene expression data. Eisen *et al.* used a hierarchical clustering approach, similar to those constructing evolutionary trees[5]. Tamayo *et al.* used self-organizing maps to differentiate between leukemia tumor types for the purpose of choosing the most appropriate chemotherapy treatment, a result of particular medical significance.[6] Other popular clustering techniques include *single-linkage*, *complete-linkage*, *group-average*, *centroid*, and *k-means*.[7]

As can be observed from results like Tamayo's, the application of computer analysis to sources of large quantities of biological data could contribute to major improvements in the diagnosis and treatment of disease. In the long term, advances in this field could have a significant economic, medical and humanitarian impact. The possibility of this type of advance motivates the creation of clustering algorithms specifically designed to analyze biological data.

One of the main challenges of microarray data analysis is the level of noise that exists in the data. Keller *et al.* [8] classify potential sources of noise as follows: the data contain 'technical' noise that can be introduced at a number of different stages, such as production of the DNA array, preparation of

---

[3] Amir Ben-Dor and Zohar Yakhini. Clustering gene expression patterns. In Proceedings of the Third Annual International Conference on Computational Molecular Biology (RECOMB 99), 1999. http://citeseer.nj.nec.com/305183.html.
[4] Hartigan, John A., 1975, *Clustering Algorithms*, New York, John Wiley and Sons.
[5] M. B. Eisen, P. T. Spellman, P. O. Brown, and D. Botstein. Cluster analysis and display of genome-wide expression patterns. PNAS, 95:14863-14868, 1998.
[6] P Tamayo, D Slonim, J Mesirov, Q Zhu, S Kitareewan, E Dmitrovsky, ES Lander, TR Golub: Interpreting patterns of gene expression with self-organizing maps: methods and applications to hematopoietic differentiation. *Proc Natl Acad Sci USA* 1999, 96: 2907-2912
[7] B. Everitt. Cluster Analysis, Chapter 4. Edward Arnold, London, third edition, 1993.
[8] Keller, Andrew, et al. Bayesian Classification of DNA Array Expression Data. Technical Report UW-CSE-2000-08-01, August, 2000

Section 2. The CAST Software

Software was created to support this paper and will be referred to as CAST. CAST combines a simple microarray data file parser and an implementation of the Ben-Dor *et al.*'s CAST algorithm in C++.

The output of the CAST's clustering module consists of a visual representation of the clustering in both real-valued and boolean similarity matrices and a listing of cluster membership by gene name and optionally gene description.

Section 3: Results

Testing the efficacy of the CAST software was difficult because its performance on experimental data is highly dependant on the characteristics of that data. Some experiments may have a strong underlying cluster structure; others may not. Also, the experiments sensitivity and tendency for error affects the algorithms performance.

The software was tested on both artificial data and experimental data. *Saccharomyces cerevisiae* (yeast) data from Michael Eisen were used as an example of a multi-conditional microarray experiment[12] and are referred to here as the AffyYeast data set. Data from the Stanford Microarray Database[13] were used as an example of a time-series microarray experiment and are referred to here as the yeast_cycle dataset. Pearson's correlation coefficient was used as the similarity measure for AffyYeast and L1 distance was used as the similarity measure for yeast_cycle.

3.1 Measurements

Ben-Dor *et al.* use two quantitative criteria for measuring the distance or "closeness" between two boolean matrices. Given two binary matrixes A and B of the same dimensions, let Nij be the number of entries on that A and B have values *i* and *j*

ratio when "negative matches" (N_{00}) are ignored: $N_{11}/(N_{01} + N_{10} + N_{11})$. Jaccard coefficient is used to evaluate CAST's performance on artificial data.

To measure cluster stability, consider a *n x n* cluster membership matrix where *n* is the number of genes to be clustered and if two genes *i* and *j* are in the same cluster, then the *i*th/*j*th matrix entry is true. Otherwise, the entry is false. *Cluster stability* is defined as the Jaccard coefficient between two of these cluster membership matrices from different application of the CAST algorithm. When the algorithm is run twice on exactly the same data, you would expect cluster stability to be near 100% and this is verified in the results. For increasing levels of additional noise, cluster stability is reduced.

[13] Stanford Microarray Database, http://genome-www.stanford.edu/microarray/, April 16, 2002

3.2 Performance on Artificial Similarity Matrices

The algorithm's performance clustering artificial data is quite good both visually and quantitatively. Artificial similarity matrices meant to simulate a cluster structure that might be found in biological data were created with an arbitrary distribution of cluster membership where each gene has the following probability of being a member of the respective cluster: {1/4, 1/4, 1/4, 1/8, 1/16, 1/16}.

Increasing amounts of white noise were added to the similarity matrix to simulate false positives in the similarity measures on real data. Let alpha be the probability of error (a false positive or negative), then for each entry in the similarity matrix, that entry was flipped to be false with independent probability alpha.

The similarity matrices were then randomly permuted. These permuted matrices are the actual inputs to the algorithm. The permuted input to the human eye appears to be very close to white noise. The goal is to reconstruct the original cluster structure from the noisy, permuted data.

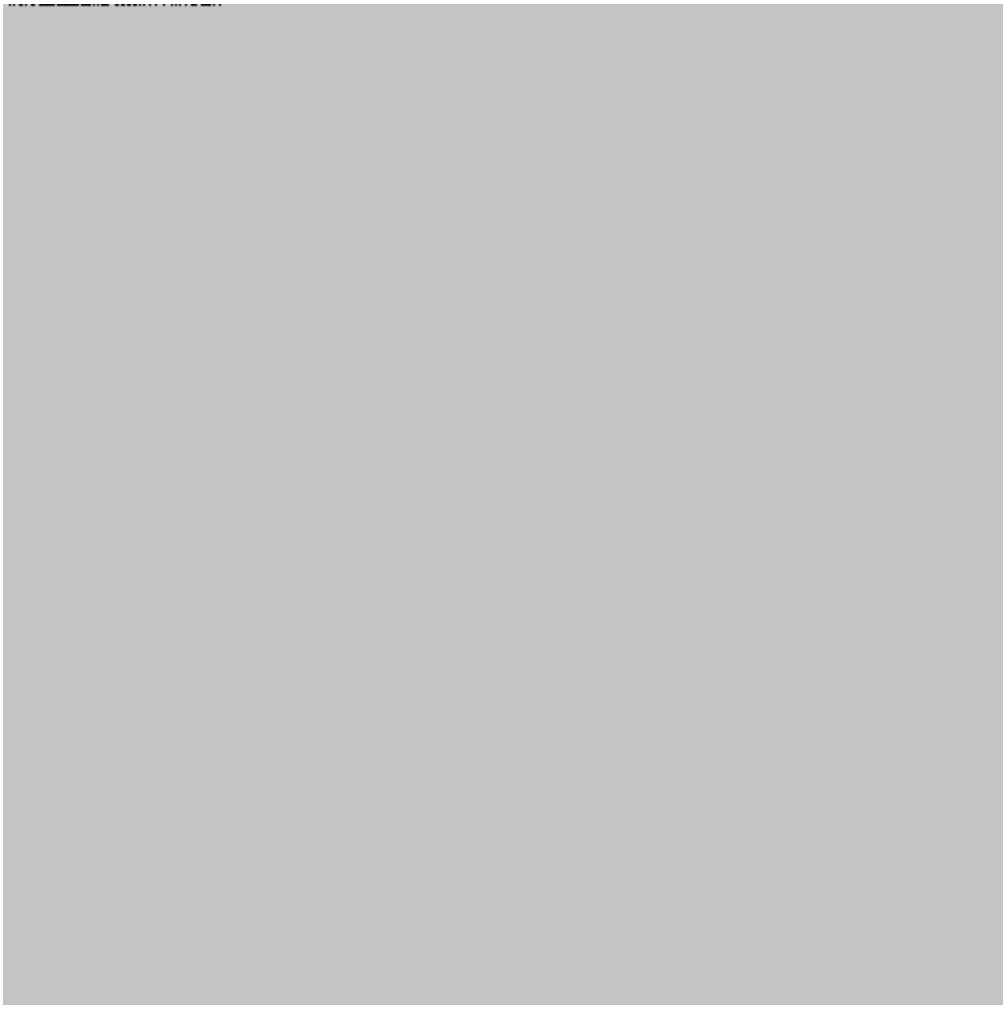

The clustering algorithm is successful in restoring the matrices to original form.

10% Error 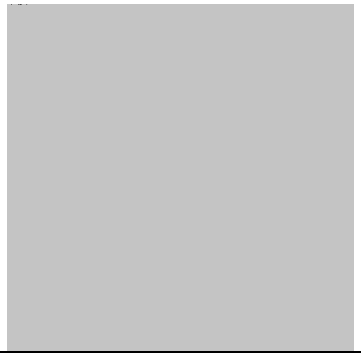	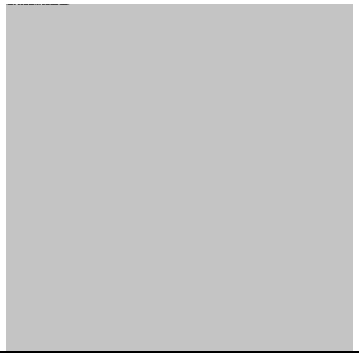	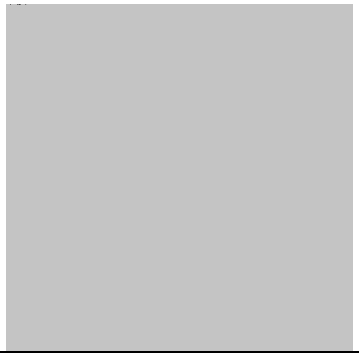
20% Error 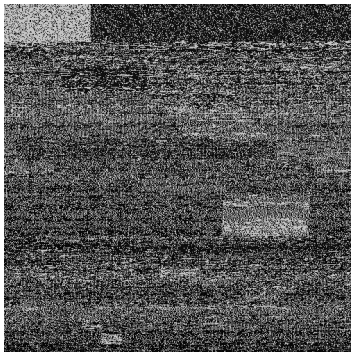	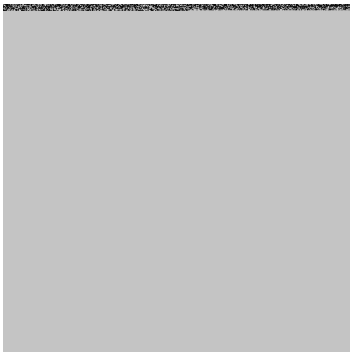	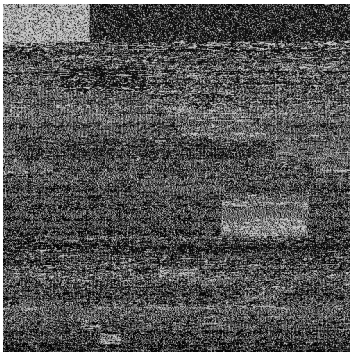

30% Error

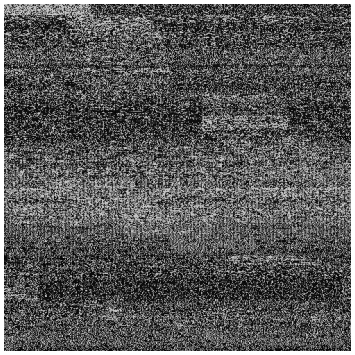

3.3 Cluster Stability

Stability was testing on the clustering produced with the AffyYeast.txt dataset. Both discretized and non-discretized clustering methods were tested at 80% and 90% affinity thresholds. The clusterings were more stable at the lower 80% cluster affinity threshold due to the creation of larger clusters. Non-discretized clusterings were more stable than clustering using discretized values.

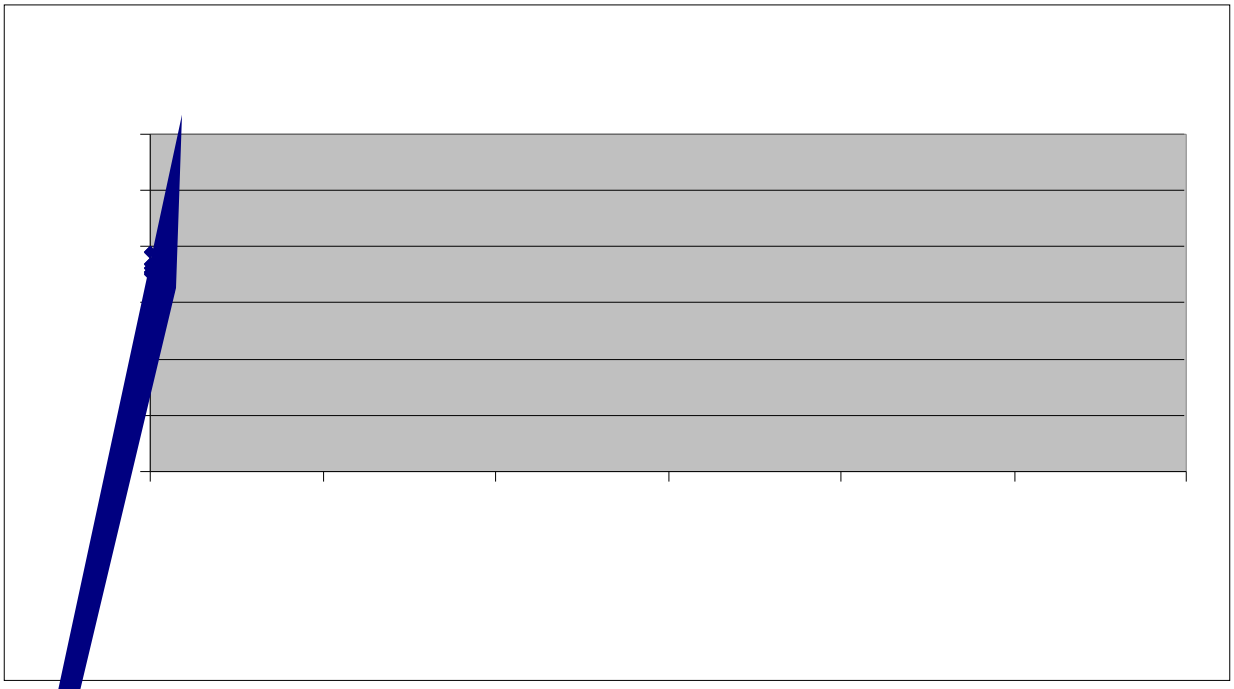

Part II: Technical Discussion and User Guide

Section 2: Clustering Implementation

Pseudocode of the CAST algorithm is as follows:

Let S(x,y) be the similarity measure between x any y.

Input:

An -by- similarity matrix , and an affinity threshold .

Initialization:

The collection of closed clusters: C = !

Elements not yet assigned to any cluster: U = {1, , }

Algorithm:

while (U ≠ !) do

Start a new cluster: C_{open} = !

Reset affinity: a(.) = 0

Repeat steps ADD and REMOVE as long as changes occur:

ADD: while max{a(u) | u # U} >= t|C_{open}| do

Pick an element u # U with maximum affinity.

$C_{open} = C_{open} \cup \{u\}$

U = U \ {u}

For all x # U ∪ C_{open} set a(x) = a(x) + S(x,u)

REMOVE:

Pick an element v # C_{open} with minimum affinity.

Remove v from C_{open}: $C_{open} = C_{open} \setminus \{v\}$

Insert v into U: U = U ∪ {v}

Update the affinity: For all x # U ∪ C_{open} set a(x) = a(x) – S(x,v)

Close the cluster: C = C ∪ {C_{open}}

The implementation of the CAST algorithm takes as input a list of *n* real valued vectors each corresponding to the microarray data for one gene. A *n x n* real valued similarity matrix is computed where the ith/jth value is the similarity measure between the ith and jth gene. Pearson's correlation coefficient, Euclidean distance, and L1 distance are available as similarity measures. The similarity matrix can either be left with its real valued measurements or discretized with some threshold of similarity significance transforming the similarity matrix to a boolean matrix.

The deciding factor for cluster membership is the affinity between the cluster and its constituent genes. The affinity of one gene to a cluster is the sum of the similarity measure between all of the genes in that cluster to that gene divided by the size of the cluster; to put it another way, affinity is the average of the similarity measures between genes in the cluster and a unclustered gene. The algorithm is given a

threshold value as the minimum average affinity to the cluster for any of its members. The algorithm adds genes to a cluster while affinity is greater than the threshold and removes genes from a cluster when affinity drops below the threshold. The algorithm terminates when all of the genes have been assigned to a cluster, even if that cluster is a singleton, meaning the gene is in a cluster by itself.

Having established cluster structure, the algorithm optionally performs a cleaning step that checks to make sure each gene is in the cluster to that it has maximum affinity; the algorithm guarantees that every gene will have affinity to it's cluster above the clustering threshold, but not that every gene is in the cluster to which it has maximum affinity. This cleaning step was implemented as a series of passes where, for each gene, the algorithm checks it's affinity to every cluster and moves the gene to the cluster of greatest affinity. As genes are swapped one at a time, cluster affinities change with each swap and the cleaning step can take many passes before a stable clustering is found.

Section 3: Cluster Visualization

To visualize the results of a clustering, the software displays the similarity matrix either as a monochromatic image for a discretized matrix or as a grayscale for a real-valued matrix. A white pixel represents similarity between two rows (genes) in the monochromatic and the white component of a grayscale pixel represents the magnitude of similarity between two rows (genes) in the grayscale image. The rows in the similarity matrix are reordered so that genes in the same cluster are next to each other and so that the clusters are sorted according to size, the largest first. This visualization produces an image with square clusters around the diagonal that are brighter than the surrounding image. This visualization is useful to get a rough idea of the quality of a clustering; a better clustering will have brighter squares around the diagonal and be darker in the regions outside the cluster.

Figure 3.1 is an example of the cluster visualization for the AffyYeast data set.

Figure 3.1:

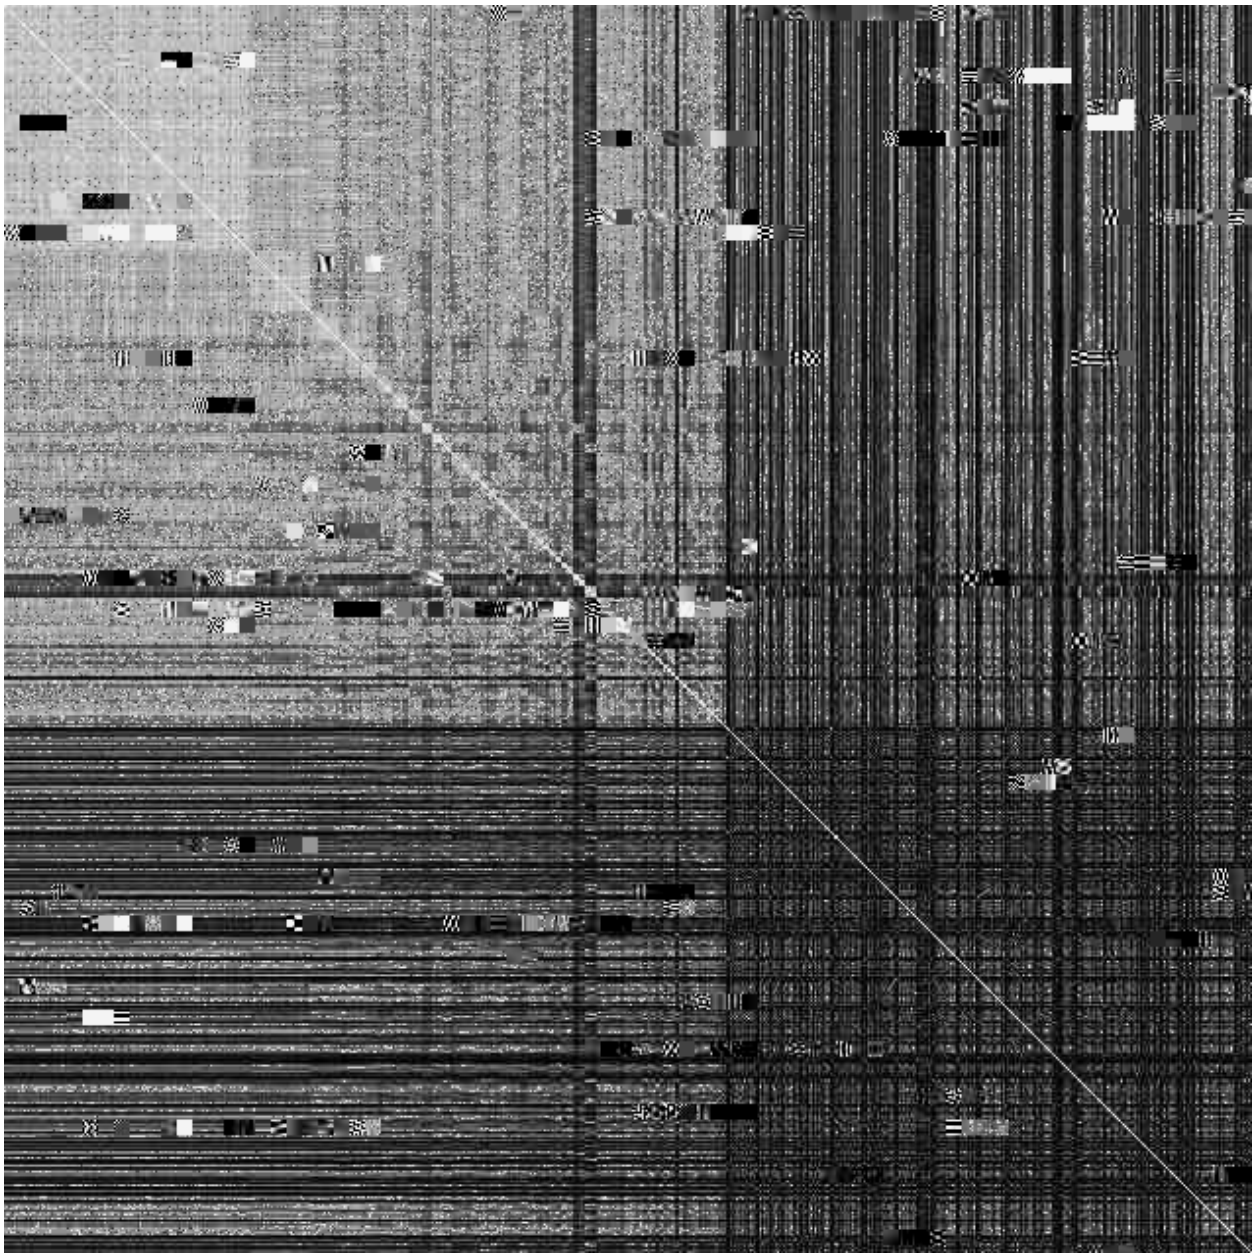

Part III: Conclusion and Further Study

Section 1. Conclusion

The performance of the CAST algorithm on artificial data is very good and visually quite impressive; however, the addition of uniform white noise to the similarity matrices may not accurately simulate error in biological data where the definition of the underlying cluster structure is fuzzier and alpha values may be higher than 50%.

Given the non-deterministic nature of the clustering algorithm, the stability of its output in the presence of error is satisfactory. Overall, the CAST software functions well as a clustering tool.

Section 2. Further Study

1) Back-end Database Interface. The most needed and obvious improvement necessary in CAST is a back-end database interface. Using flat files to represent microarray data is woefully inadequate. An interface to existing web-based database of microarray data such as the Stanford Microarray Database would be ideal. An intermediate step toward this goal would be the abstraction of the GeneArray class into a interface/abstract base case in a manner similar to the design of the Measure interface.

2) A standardized file format for microarray data. The author is not aware of a standardized format for microarray data, but the establishment of such a standard would greatly aid non-biologists conducting research into microarray analysis techniques. The variety of microarray file formats severely limits the applicability of the CAST software. At present, CAST requires a parsing script hack for each different file format. An intermediate step would be more generalized parsing script or perhaps using Python as a scripting language to address the algorithm as a C++ object.

3) Better handling of missing data values. Currently, CAST interpolates missing data values that introduce error for multi-conditional experiments. The Measure interface must be modified to allow for similarity measures to consider missing data points.