

Discrimination of Metabolic Flux Profiles Using a Hybrid Evolutionary Algorithm

Stefan Bleuler

Computer Engineering and Networks Laboratory
ETH Zurich
CH-8092 Zurich, Switzerland
bleuler@tik.ee.ethz.ch

Eckart Zitzler

Computer Engineering and Networks Laboratory
ETH Zurich
CH-8092 Zurich, Switzerland
zitzler@tik.ee.ethz.ch

ABSTRACT

Studying metabolic fluxes is a crucial aspect of understanding biological phenotypes. However, it is often not possible to measure these fluxes directly. As an alternative, fluxome profiling provides indirect information about fluxes in a high-throughput setting. In this paper, we consider a scenario where fluxome profiling is used to investigate characteristic differences between a number of bacterial mutant strains. The goal is to identify groups of mutants that show maximally different fluxome profiles. We propose an evolutionary algorithm for this optimization problem and demonstrate that it outperforms alternative methods based on principle component analysis and independent component analysis on both real and synthetic data sets.

Categories and Subject Descriptors

I.2 [Artificial Intelligence]: Miscellaneous

General Terms

Algorithms

Keywords

Evolutionary Algorithm, Biological Application, Fluxome Analysis

1. MOTIVATION

For a given organism the metabolic network describes the set of biochemical reactions and their connections. Often, information about the structure of the metabolic network is available but in order to understand biological phenotypes knowledge about the activities in the metabolic network is important [9]. These activities are typically characterized by the molecular fluxes through the network. No suitable method exists for measuring the fluxes directly but fluxome

profiling provides an indirect measurement. This is achieved by measuring the accumulation of labeled inputs at different nodes in the network. Abstracting from the underlying biology, the resulting data can be described as a real valued vector representing the metabolic fingerprint of the respective organism, see Section 2.1 for more details.

Recent advances in measurement technology enable metabolism-wide determination of fluxome profiles in relatively short time [4]. Often, such experiments are used to study the effects of mutations in bacteria. In such a scenario, a large number (several ten to several hundred) of different mutants are subjected to fluxome profiling. The corresponding goal for the analysis of the measurements is to identify distinct groups of mutants and determine the characteristic differences in their flux profiles. Or, from an opposite point of view, to identify characteristic features in the fingerprints that exhibit a good separation of groups of mutants.

The approach followed in [12] and [10] is to search for features in the flux profiles which can be directly related to the biological differences between the mutants. To this end, principal component analysis (PCA) and independent component analysis (ICA) have been used and it was found that in some cases these components could be related to underlying biological properties such as flux ratios or biological replicas. However, such strong relations cannot be reliably inferred in general. In this study, we take a different approach which consists of identifying groups of mutants that have distinct fingerprints. More specifically, we are searching for various pairs of mutant groups that are well separated with respect to their flux profiles. In contrast to the methods proposed in [12, 10], here the goal is not to identify specific features in the flux profiles but we try to discriminate the mutants based on their flux profiles directly. The resulting mutant groups can then direct the further investigations of the biologist experimenter towards potentially interesting biological differences.

Several EA-based methods have been proposed for grouping data points, e.g., [3, 5]. All these methods are targeted to a cluster analysis, i.e., they generate a partition by assigning each data point to one group. However, in the context of the fluxome analysis we are interested in identifying pairs of groups that are well separated. To this end, we propose i) a general problem formulation for this separation which is independent of the distance measure used and ii) an optimization framework specifically adapted to this problem. The method is a combination of an evolutionary algorithm used for global search and a greedy heuristic used for locally optimizing solutions. We verify the ability of the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO '07, July 7–11, 2007, London, England, United Kingdom.
Copyright 2007 ACM 978-1-59593-697-4/07/0007 ...\$5.00.

method to find diverse sets of well separated mutant groups and show that these groups are better separated than those found by heuristics based on PCA or ICA. Additionally, we demonstrate the validity of the approach by showing that well separated mutants also have distinct flux ratios on a set of simulated fluxome measurements.

The next section provides a short overview of fluxome profiling and of methods that have been previously used for the analysis of fluxome data. Following on that, we propose a problem formulation for the discrimination of fluxome profiles in Section 3 and a corresponding optimization method in Section 4. The aforementioned analysis on real and synthetic data is presented in Section 5.

2. BACKGROUND

2.1 Fluxome Profiling

A metabolic flux describes the number of molecules that are going through a biochemical reaction per time unit. Fluxome profiling is a method used to characterize the entirety of molecular fluxes through the metabolic network [9]. For an improved understanding of many biological processes, knowledge of these fluxes is crucial because, unlike gene or protein expression, they directly determine the cellular phenotype. No suitable method exists for measuring the fluxes directly but fluxome profiling provides an indirect measurement.

Several methods exist for determining flux profiles. The only one which is suitable for high-throughput experiments uses isotope markers [9]. The basic idea is to provide an organism, e. g., *Bacillus subtilis* bacteria with a labeled substrate such as glucose built of ^{13}C isotopes or alternative heavy isotopes and then determine in which metabolites the labeled isotopes end up [4]. For these measurements, a number of amino acids are extracted from the metabolites and gas chromatography mass spectrometry (GC-MS) is used to determine the proportion of atoms in these amino acids that has been replaced by the marker isotopes. The resulting data specify for each measured amino acid what proportion of the molecules contain zero, one, two, etc. marked atoms. Note that the same amino acid is often contained in multiple metabolites included in the analysis and can thus not be uniquely linked to one position in the metabolic network. For a detailed description, the reader is referred to [4].

2.2 Fluxome Profile Analysis

Basically, two different approaches exist for the analysis of such data sets. If detailed knowledge of the metabolic reaction network, the substrate uptake, etc. is available flux ratios can be calculated directly [4]. In the general case, techniques from multivariate statistics can be applied to discriminate different mutant strains. From this category mainly two methods have been applied to the analysis of fluxome profiles, namely PCA and ICA, which we will now discuss in some more detail.

PCA is a method mainly used for dimensionality reduction. The first principle component represents the direction in which the variance of the data is highest. The following principle components are chosen such that each captures the largest amount of remaining variance in the data. All principle components are orthogonal. Two studies [12, 10] which investigate the effectiveness of different methods find that the principle components do not correspond to the under-

lying biological differences. For example in [10] PCA is not able to identify groups of biological replicas. This is probably due to the fact that high variance does in general not coincide with good group discrimination [7] and PCA was designed to maintain the variance of the data in a reduced dimensionality.

ICA, in contrast, was developed for the discrimination of independent signals, a task also known as blind source separation [7]. It considers the inputs as different additive mixtures of the unknown signals. The data model requires that the source signals be mutually independent and non-Gaussian. Given that, any mixture of the source signals are more Gaussian than the source signals themselves. The ICA algorithm thus tries to identify non-correlated components that are maximally non-gaussian. By itself, ICA does not change the number of dimensions but an additional method like PCA can be used to achieve a dimensionality reduction. Two studies [12, 10] have used ICA for the analysis of fluxome profiles. In [10] independent components clearly separated groups of biological replicas which PCA was not able to do. In [12] the projections of the isotope profiles on the independent components were correlated to the flux ratios at critical points in the network. A few combinations of independent components and flux ratios exhibited high correlations. However, in general it is not possible to extract the flux ratios by means of ICA as these flux ratios at different points in the network are connected and thus not independent.

3. PROBLEM DEFINITION

In contrast to ICA, which performs a coordinate transform, in this study we search for well discriminated groups of mutants in the original space of isotope profiles. The underlying idea is that two groups of mutants showing distinct isotope profiles probably exhibit two different characteristic flux distributions. The association of mutants with the two groups can serve as a starting point for further biological investigation of the characteristic differences and similarities of the mutant strains. Correspondingly, we are not only interested in finding two well separated groups but multiple diverse pairs of mutant groups. In the following, we describe the criteria for these two levels of optimization more formally by specifying when two groups are well separated and how to define diversity of multiple pairs of groups.

Given m mutants and c amino acids which are included in the measurements, the fluxome profiles are given by the $m \times n$ matrix \mathbf{P} where each row represents the measurements for one mutant. Such a row vector contains 3–9 elements l_j^i for each amino acid j according to the number of carbon atoms it contains, i. e., $i \in \{0, 1, 2, \dots, d_j\}$ where d_j is the number of carbon atoms in amino acid j . These values specify what proportion of the total amino acid contained i heavy isotopes, c. f. Figure 1. It follows that

$$\sum_{i=0 \dots d} l_i^j = 1 \quad \forall j. \quad (1)$$

For the current study we consider each mutant as a point in euclidean space given by the corresponding row of \mathbf{P} .

As a first step, we look at the problem of selecting two groups from the set of m mutants such that the isotope profiles are similar within the two groups but clearly distinct between them. This selection corresponds to a partitioning

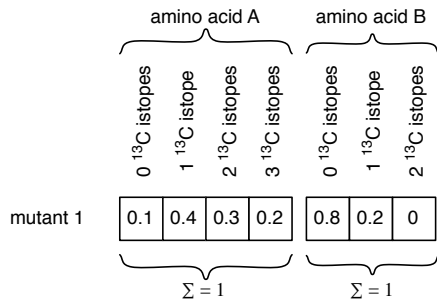


Figure 1: A hypothetical isotope profile for one mutant. Amino acid A contains 3 carbon atoms and amino acid B contains 2 carbon atoms. The values in the array specify the proportion of amino acid molecules containing the given number of heavy isotopes

of mutants into three parts: the mutants in each group and the mutants not included in either group. It can be defined as a pair of sets (G_1, G_2) where $G_1, G_2 \subset \{1, \dots, m\}$ and $G_1 \cap G_2 = \emptyset$. The size of the induced search space is 3^m as there are three possible assignments for each mutant. A solution (G_1, G_2) is evaluated based on two criteria, namely the separation with respect to the isotope profiles and the total size of the two groups.

The separation of two sets of points can be defined as the silhouette width [8]. This is a measure often used in the validation of clustering methods, as it combines information about the intra-cluster distances and the inter-cluster distances into a scalar [6]. The silhouette width w is defined as the mean silhouette value over all points. The silhouette value for each point measures the confidence in the point's assignment to the group and is calculated as

$$s(i) = \frac{b_i - a_i}{\max(b_i, a_i)}, \quad (2)$$

where a_i denotes the mean distance between i and all points in the same group and b_i denotes the mean distance between i and all points in the opposite group. Due to the scaling factor in (2) the silhouette width w is in $[-1, 1]$. Note that higher values of w correspond to better separation. The silhouette width calculation can be applied to arbitrary distance measures and in the present study we have focussed on euclidean distance since large euclidean distances in the isotope space corresponds to a reasonable degree to large distances in the space of flux ratios as will be shown in Section 5.3.

In general, choosing small groups of mutants leads to better separation values than choosing larger groups. To counterbalance this tendency when evaluating the quality of a pair of groups we need to take into account the total size of the two groups. In summary, one pair of mutant groups $H = (G_1, G_2)$ is evaluated by two criteria: it's silhouette width $w(H)$ and its size $t(H) = |G_1| + |G_2|$.

Based on these two criteria, size and separation, different optimization problems can be formulated. They could be aggregated into one objective function or a Pareto-based multiobjective approach could be used. But as we are only interested in the solutions where the groups are highly separated, we impose a user defined constraint σ on the separation w and maximize the group size t .

So far we have considered the problem of identifying two distinct groups of mutants. However, for the biological analysis we are interested in several pairs of well separable groups in order to capture the major biological differences in the set of mutants under investigation. To this end, the different pairs should not be highly similar to each other.

The extended problem is to identify a set of k group pairs $\{H_1, H_2, \dots, H_k\}$ for which the average group size should be high and the overlap between the group pairs should be low. For the present study, we have defined the overlap between H_i and H_j as follows.

$$o(H_i, H_j) = \max(o_{\text{normal}}(H_i, H_j), o_{\text{flip}}(H_i, H_j)) \quad (3)$$

$$o_{\text{normal}}(H_i, H_j) = \frac{|G_{1,i} \cap G_{1,j}|}{|G_{1,i}|} \cdot \frac{|G_{2,i} \cap G_{2,j}|}{|G_{2,i}|} \quad (4)$$

$$o_{\text{flip}}(H_i, H_j) = \frac{|G_{2,i} \cap G_{1,j}|}{|G_{2,i}|} \cdot \frac{|G_{1,i} \cap G_{2,j}|}{|G_{1,i}|} \quad (5)$$

The multiplication of the overlaps of the single groups reduces the total overlap to zero if either of the groups have no overlap. Three non-overlapping groups for example can form two pairs where one group is identical in both pairs and nevertheless have an overlap of zero.

As we are only interested in solutions with low overlap we set a constraint δ on the maximal pairwise overlap. This leads to the following optimization problem:

$$\begin{aligned} \max \quad & \frac{1}{k} \sum_{i=1}^k t(H_i) \\ \text{s.t.} \quad & w(H_i) \geq \sigma \quad \forall i \in \{1, \dots, k\} \\ & \max_{i,j} o(H_i, H_j) \leq \delta \quad \forall i \neq j, i, j \in \{1, \dots, k\} \end{aligned} \quad (6)$$

Note that the described problem formulation is general in the following sense: arbitrary distance measures can be used and any feature selection method or coordinate transform deemed appropriate can be applied to the isotope data as a preprocessing step.

4. AN EA FOR THE DISCRIMINATION OF FLUXOME PROFILES

Given the scoring scheme described in the previous section the primary optimization goal is to find the largest two groups which meet the user-specified constraint σ on their separation w . Additionally, we search for several pairs that are diverse. This section presents an optimization method for this problem which is based on a combination of an evolutionary algorithm with a local search heuristic and is based on the method used in [1].

4.1 Basic Architecture

The optimization method consists of two main components: a global search method and a local search heuristic. The global search is used to explore the space of possible groupings systematically while the local search method improves the solutions identified by the global search. The global search is performed by an evolutionary algorithm in which each individual represents one pair of mutant groups, i. e., two non-overlapping subsets of mutants.

In order to improve the performance of the global search method regarding the huge search space, a greedy strategy

Algorithm 1 Environmental Selection for Diversity Maintenance

N : number of individuals to select
 P : old population to select from
 M : size of P
 $Q = \emptyset$: new population
sort P in decreasing order of $t(H)$
 $s = 0$
for $k = 0$ to M **do**
 if $\max(o(P[k], Q[i]) \forall i < \delta)$ **then**
 add $P[k]$ to Q
 $s = s + 1$
 end if
end for
 $k = 0$
while $s < N$ **do**
 if $P[k] \notin Q$ **then**
 add $P[k]$ to Q
 $s = s + 1$
 end if
end while

is incorporated which is able to rapidly improve the solutions locally. The basic idea of the local search heuristic is to iteratively add or remove mutants until the separation constraint σ is met. While this method can produce valid solutions to the optimization problem at hand by itself it quickly gets stuck in local optima and the global search ensures that the local optima can be overcome.

According to the extended optimization problem, we are not only interested in two well separated mutant groups but we would like to identify several diverse pairs. The basic strategy to achieve this is to distribute the population of the EA and ensure it does not converge to a single point. We use a special kind of environmental selection to enforce diversity. This method, along with further details of the algorithm is described in the next section.

4.2 Algorithmic Details

4.2.1 Representation

Each individual represents two groups of mutants. For reasons of simplicity we have chosen to use a ternary representation with a string of length m . An element is set to 1 or 2 if the corresponding mutant is assigned to the first or the second group, respectively. Unassigned mutants are represented by a 0.

4.2.2 Initialization

The initial population should be generated such that a high diversity of groupings is attained. A simple strategy, for example, which randomly assigns each mutant to either group or leaves it unassigned with equal probability produces groups containing different mutants but all groups will be of similar sizes. To avoid this, we do not use equal probabilities but for each individual randomly choose the probabilities of assigning mutants to the different groups. This is done, by first randomly sampling the probability for a mutant to be in any of the two groups from a uniform distribution and in a second step sampling the probability for a selected mutant to be in the first of the two groups from a uniform distribution.

4.2.3 Variation

Each element of the string undergoes mutation with a certain probability p_{mut} . Mutation changes a 0 into a 1 or 2 with equal probability and vice versa. As recombination operator we apply uniform crossover which for each element picks the value of either of the parents with equal probability.

4.2.4 Selection

For mating selection, tournament selection is used, i. e., τ individuals are chosen from the population with replacement and the fittest one is copied to the pool of parents. In choosing the value of τ the selection pressure can be influenced: A higher τ results in more pressure towards fit solutions.

As described, we propose a specific environmental selection to maintain diversity in the population. In this process the constraint δ on the maximal overlap is used as a soft constraint, i. e., solutions cannot be guaranteed to fulfill the constraint. The algorithm proceeds as follows. First the individuals are sorted by the total size of their groups $t(H)$. Starting with the largest one all individuals are selected which do not overlap more than a user-defined threshold δ with any of the previously selected individual. The overlap $o(H_i, H_j)$ is calculated as defined in Equation (3). If not enough non-overlapping individuals are found the new population is filled using the largest of the previously omitted individuals. The entire procedure is detailed in Algorithm 1.

4.2.5 Local Search and Fitness

Before the evaluation of an individual a local search is performed to improve the solution. If the separation w is above the user specified threshold σ mutants are added to the groups in a greedy strategy which in each step adds the mutant with the maximal silhouette value $s(i)$. The algorithm terminates when no more mutant can be added without violating the constraint. Conversely, if $w < \sigma$ mutants are removed applying the opposite greedy strategy, i. e., in each step the mutant with the lowest silhouette value $s(i)$ is removed. Note that the local search guarantees constraint satisfaction with respect to the separation w as it is always possible to reduce the groups to one mutant each and thus reach the maximal w .

An individual is evaluated based on the result of the local search. Since the objective is to find large groups the fitness $f(H)$ of an individual H is calculated as the inverse of the total number of mutants included in the two groups $f(H) = \frac{1}{t(H)}$. This fitness is to be minimized. The result of the local search can either replace the original individual or just be used to determine the fitness of the original individual while the latter one remains unchanged. In this study the second strategy, called Baldwinian evolution, is used since it is able to generate a more diverse set of solutions.

5. RESULTS

In the simulation runs mainly two questions were investigated: (i) Is the proposed EA effective compared to random search and two grouping methods based on PCA and ICA, respectively, and (ii) do the mutant groups which have well separated isotope profiles identify any biological differences?

5.1 Data Preparation and Experimental Setup

The proposed approach was evaluated on a large-scale measurement of fluxome profiles for the central carbon me-

Table 1: Default parameter settings for this study.

σ	0.8 (0.5 for real data)
δ	0.2
mutation rate	0.01
crossover rate	0.1
tournament size	2
population size	100
number of generations	100

tabolism of the microorganism *Bacillus subtilis* [4] and on two artificial data sets attained by simulating the same metabolism [2]¹. For this simulation random flux maps were generated by random sampling within the polytope constrained by the carbon stoichiometry with glucose as unique substrate and CO₂ and acetoin as unique allowed products. Carbon labeling experiments were simulated for each flux map using a Matlab-based implementation of the algorithm described in [11].

The EA parameter settings used in the following simulations are described in Table 1. The crossover rate refers to the percentage of parents involved in crossover. The mutation rate is the probability for an element in the ternary string to undergo mutation. 11 replicates with different random number generator seeds were performed for each parameter setting.

5.2 Evaluation of the Evolutionary Algorithm

As a first step in the evaluation, we compared the EA results to randomly chosen groups. For each of the individuals in the final population 100 randomly chosen pairs of groups of the same size were generated. Figure 2 shows a histogram of the resulting separation values w for the randomly chosen groups in comparison to the threshold set for the EA ($\sigma = 0.8$). As expected the optimization leads to far better separation for the same group size than randomly choosing groups.

As an alternative to optimization with the EA, one can focus on a direction in the isotope space in which the mutants are known or thought to be well separated and form two groups by picking mutants on the two extremes of this direction. We have used this strategy to compare the EA results to ICA, PCA and random components. More specifically, we have projected the isotope data on each component and then greedily picked mutants from the extremes of this projection as long as the constraint on separation ($w > \sigma$) on the original data was satisfied.

Basically, ICA and PCA both produce the same number of components as the dimension of the input data. But in most applications the number of components is restricted to a few, e.g., ten in the present study. In order to compare these results to the EA an additional postprocessing is necessary to pick the same number of individuals from the final EA population. This is done by the same algorithm as used for the environmental selection, i.e., by picking the largest ones as long as they do not overlap more than a specified amount with any of the previously selected individuals. In this procedure one can either choose the same overlap constraint as during the optimization or be more or

¹The data sets are available on www.tik.ee.ethz.ch/sop/fluxome.

less restrictive. Thus, it is possible to adjust the selection of the final set within the trade-off between larger groups and larger overlap on the one hand and smaller groups but smaller overlap on the other hand.

The resulting pairs of mutant groups all satisfy the separation constraint σ and an equal number of pairs has been chosen for each method. Thus, the different methods can be directly compared by comparing the respective group sizes, the overlaps or the number of mutants that are included in any of the groupings, in the following referred to as coverage. Figure 4 shows the histogram of group sizes, overlap and coverage for random components and the corresponding results for PCA, ICA and the EA. For both the random components and the ICA, which is also a stochastic method, 1000 sets of ten components were sampled to build the histograms. It can be clearly seen that the EA achieves higher group size, higher coverage and lower overlap than the other methods. While the EA results are not necessarily significantly better for each run on each criterion separately, they are highly significant with respect to the combination of all criteria. In the whole analysis no random set of components was found to be superior or equal to any EA result in all three criteria simultaneously. It is interesting to note that the relative performance of the EA is much better on the real data set than on the synthetic data sets. Whether this difference is due to general characteristics of synthetic and real data sets remains to be seen.

These results demonstrate that the proposed combination of an evolutionary algorithm and a local search heuristic is successful in solving the optimization problem of finding a diverse set of well separated group pairs.

5.3 Validation using Flux Ratios

The last section provided results which demonstrated that the proposed method is successful in identifying mutant groups which are well separated with respect to the isotope data. While this is the main focus of this study the use of synthetic data enables us to test whether the similarities of the fluxome profiles are representative for the underlying fluxes. For the synthetic data sets flux ratios can be calculated exactly. For the analysis, we calculated the silhouette width for the EA solutions based on the flux ratios and compared them to the corresponding silhouette widths for random groups of the same size. As Figure 3 shows, the groups that were optimized for separation on the isotope values are far better separated on the flux ratios than random groups of the same size. It is thus possible, using the proposed problem formulation, to extract information about biological differences and similarities without calculating the exact flux ratios.

6. CONCLUSIONS

Fluxome profiles provide indirect information about the molecular fluxes through a metabolic network. Extracting the real fluxes or flux ratios from these data is often not possible. Thus, we propose the identification of distinct groups of profiles as a possible first step in the fluxome analysis. To this end, we have

- formalized the problem in terms of the separation of two groups as measured by the silhouette width,
- proposed a flexible optimization method for this prob-

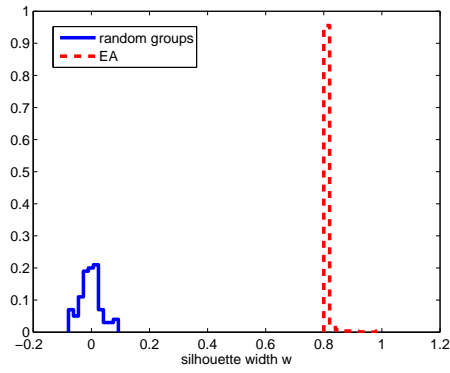


Figure 2: Comparison of the silhouette widths for the EA results and those for random groups of the same size.

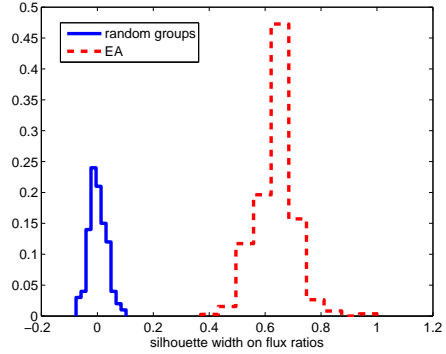


Figure 3: Comparison of EA groups to random groups with respect to separation on the true flux ratios.

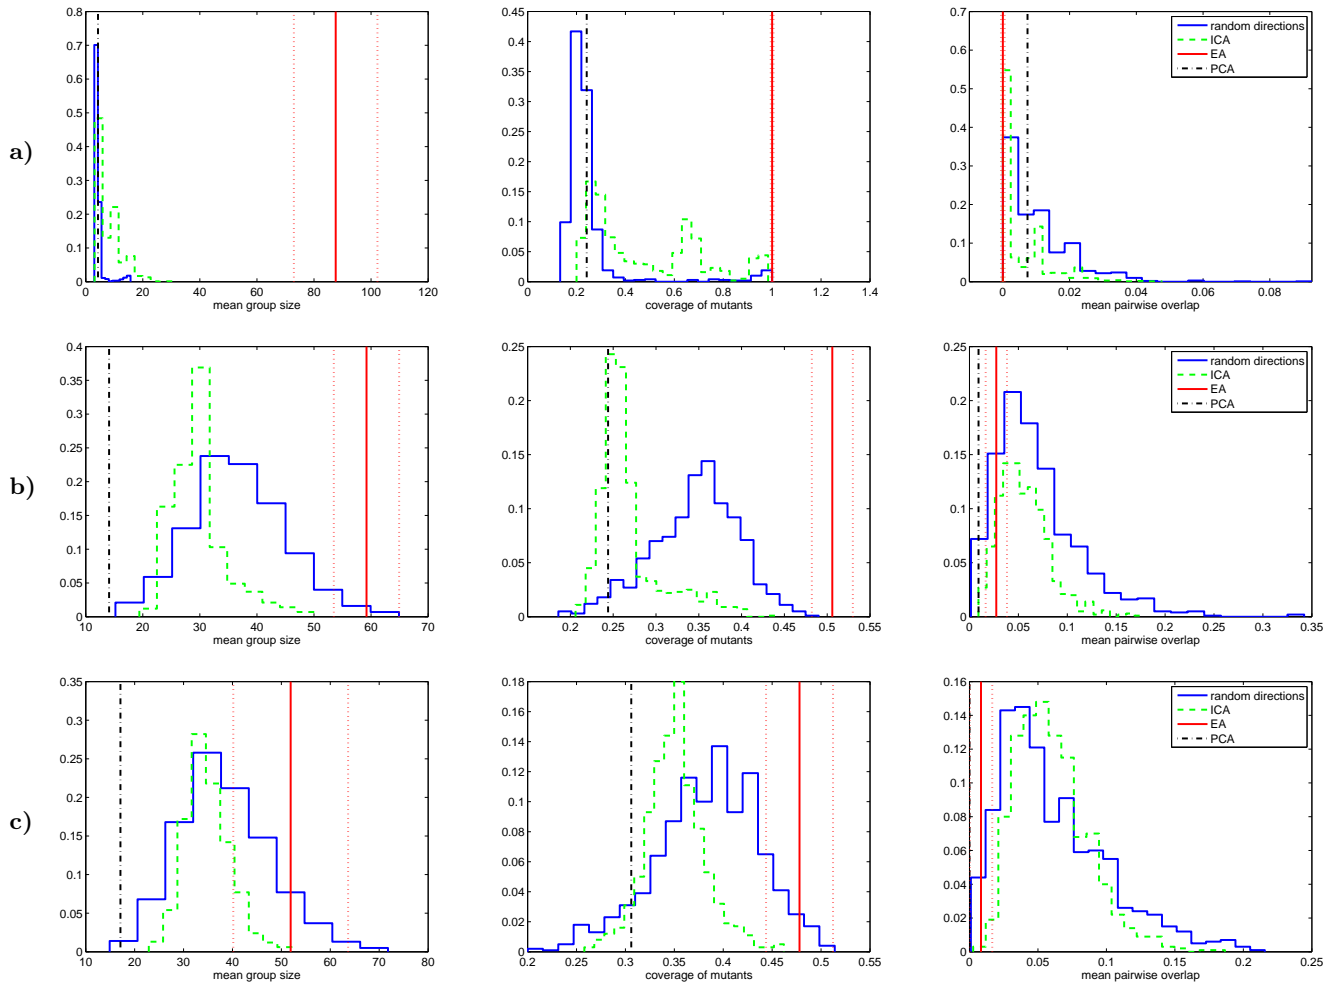


Figure 4: Comparison of the EA results to mutant groups extracted from predefined directions given by random directions, ICA and PCA. The results are compared with respect to the three criteria: average size of mutant groups, number of mutants included in any of the groups (coverage) and average overlap of all pairs of mutant groups. The analysis was performed on a real world data set (a), and two synthetic data sets (b and c). EA results are given by the median (solid line) of 11 runs and the width of 2 standard deviations (dotted lines).

lem based on a combination of an evolutionary algorithm and a local search strategy, and

- verified both the usefulness of the problem formulation with respect to the actual flux ratios and the effectiveness of the evolutionary algorithm for this optimization problem.

The problem formulation and the optimization method are flexible with respect to the actual distance measure used. While euclidian distance was used in this study any other distance measure could be used instead and even feature selection methods can be easily included in the flow of the analysis. However, our analysis of an artificial data set simulating fluxes in *Bacillus subtilis* mutants showed that even euclidean distance leads to a good correspondance between highly separated groups on the profiles and distinct flux ratios. Further simulations using both artificial and real-world data sets showed that the proposed approach compares favorably to alternative heuristics based on PCA or ICA. The EA is able to identify multiple pairs of mutant groups where the individual groups are larger (for a fixed separation) and the different pairs are more diverse than those found by the alternative methods.

Potential directions for future research in this young field include the investigation of alternative distance measures and the potential inclusion of a feature selection step. Additionally, more biological validation is indubitably necessary.

7. ACKNOWLEDGMENTS

The authors would like to thank Nicola Zamboni from the Institute of Molecular Systems Biology at ETH Zurich for providing the fluxome simulation framework and the real data set and for many helpful discussions. This work has been supported by the SEP program at ETH Zurich under the poly project TH-8/02-2.

8. REFERENCES

- [1] S. Bleuler, A. Prelić, and E. Zitzler. An EA framework for biclustering of gene expression data. In *Congress on Evolutionary Computation (CEC-2004)*, pages 166–173, Piscataway, NJ, 2004. IEEE.
- [2] M. Dauner, J. E. Bailey, and U. Sauer. Metabolic flux analysis with a comprehensive isotopomer model in *Bacillus subtilis*. *Biotechnol Bioeng*, 76:144–156, 2001.
- [3] E. Falkenauer. *Genetic Algorithms and Grouping Problems*. John Wiley & Sons, 1998.
- [4] E. Fischer and U. Sauer. Metabolic flux profinig of *Escherichia coli* mutants in central carbon metabolism using gc-ms. *Eur. J. Biochem.*, 270:880–891, 2003.
- [5] J. Handl and J. Knowles. Evolutionary multiobjective clustering. In *Parallel Problem Solving from Nature (PPSN VIII)*, volume 3242 of *LNCS*, pages 1081–1091. Springer, 2004.
- [6] J. Handl, J. Knowles, and D. B. Kell. Computational Cluster Validation in Post-Genomic Data Analysis. *Bioinformatics*, 21(15):3201–3212, 2005.
- [7] A. Hyvärinen. Survey on independent component analysis. *Neural Computing Surveys*, 2:94–128, 1999.
- [8] P. J. Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53–65, 1987.
- [9] U. Sauer. High-throughput phenomics: experimental methods for mapping fluxomes. *Curr. Opin. Biotechnol.*, 15(1):58–63, 2004.
- [10] M. Scholz, S. Gatznek, A. Sterling, O. Fiehn, and J. Selbig. Metabolite fingerprinting: detecting biological features by independent component analysis. *Bioinformatics*, 20(15):2447–2454, 2004.
- [11] W. Wiechert, M. Möllney, N. Isermann, M. Wurzel, and A. A. de Graaf. Bidirectional reaction steps in metabolic networks: III. explicit solution and analysis of isotopomer labeling systems. *Biotechnol Bioeng*, 66:69–85, 1999.
- [12] N. Zamboni and U. Sauer. Model-independent fluxome profiling from 2H and 13C experiments for metabolic variant discrimination. *Genome Biology*, 5(R99), 2004.