

Finding biclusters by random projections[☆]

Stefano Lonardi^{a,*}, Wojciech Szpankowski^b, Qiaofeng Yang^a

^a*Department of Computer Science and Engineering, University of California, Riverside, CA, USA*

^b*Department of Computer Sciences, Purdue University, West Lafayette, IN, USA*

Abstract

Given a matrix X composed of symbols, a bicluster is a submatrix of X obtained by removing some of the rows and some of the columns of X in such a way that each row of what is left reads the same string. In this paper, we are concerned with the problem of finding the bicluster with the largest area in a large matrix X . The problem is first proved to be **NP**-complete. We present a fast and efficient randomized algorithm that discovers the largest bicluster by random projections. A detailed probabilistic analysis of the algorithm and an asymptotic study of the statistical significance of the solutions are given. We report results of extensive simulations on synthetic data.

© 2006 Elsevier B.V. All rights reserved.

Keywords: Biclustering; Two-way clustering; Co-clustering; Biclques; Textual matrices

1. Introduction

Clustering refers to the problem of finding a partition of a set of input vectors, such that the vectors in each subset (cluster) are “close” to one another (according to some predefined distance). A common limitation to the large majority of clustering algorithms is their inability to perform on high dimensional spaces (see, e.g., [1,2]).

Recent research has focused on the problem of finding hidden sub-structures in large matrices composed by thousands of high dimensional vectors (see, e.g., [4,30,29,3,32,7,27,21,25,20]). This problem is known as *biclustering*. In biclustering, one is interested in determining the similarity in a *subset* of the dimensions (subset that has to be determined as well). Although there exists several definitions of biclustering, it can be informally described as the problem of finding a partition of the vectors and a subset of the dimensions such that the projections along those directions of the vectors in each cluster are close to one another. The problem requires to cluster the vectors and the dimensions simultaneously, thus the name “biclustering”.

Biclustering has important applications in several areas, such as data mining, machine learning, computational biology, and pattern recognition. Data arising from text analysis, market-basket data analysis, web logs, etc., is usually arranged in a contingency table or co-occurrence table, such as, a word-document table, a product-user table, a cpu-job table or a webpage-user table. Discovering a large bicluster in a product-user matrix indicates, for example, which users share the same preferences. Finding biclusters has therefore applications in recommender systems and collaborative filtering, identifying web communities [17], load balancing, discovering association rules, among others.

[☆] A preliminary version of this work was presented at the *Symposium on Combinatorial Pattern Matching*, Istanbul, Turkey, 2004.

* Corresponding author. Tel.: +1 951 827 2203; fax: +1 951 827 4643.

E-mail address: stelo@cs.ucr.edu (S. Lonardi).

In computational biology, this problem is associated with the analysis of gene expression data obtained from microarray experiments. Gene expression data is typically arranged in a table with rows corresponding to genes, and columns corresponding to patients, tissues, time points, etc. The classical approach to analyze microarray data is clustering. The process of clustering partitions genes into mutually exclusive clusters under the assumption that genes that are involved in the same genetic pathway behave similarly across all the testing conditions. The assumption might be true when the testing conditions are associated with time points. However, when the testing conditions are heterogeneous, such as patients or tissues, the previous assumption is not appropriate anymore. One would expect that a group of genes would exhibit similar expression patterns only in a subset of conditions, such as the subset of patients suffering from the same type of disease. Under this circumstance, biclustering becomes the alternative to the traditional clustering paradigm. The results of biclustering may enable one to discover hidden structures in gene expression data in which many genetic pathways might be embedded. It might also allow one to uncover unknown genetic pathways, or to assign functions to unknown genes in already known genetic pathways.

Biclustering is indeed, not a new problem. In fact, it is also known under several other names, namely “co-clustering”, “two-way clustering” and “direct clustering” (see, e.g., [6,10,11]). Because of this and the wide range of application of this problem, the theoretical and practical results are spread across different communities in Computer Science such as computational biology, machine learning, databases and also in Statistics. In the following, we review some of the results, without pretending to be exhaustive.

The problem was first introduced in the seventies in a paper by Hartigan [12]. Almost 30 years later, Cheng and Church [4] raised the interest on this problem for applications in gene expression data analysis (see, e.g., [11,13,15,18,19,31]).

Several other researchers studied the problem recently. Wang et al. propose the *pCluster* model that is capable of discovering shifting or scaling patterns from raw data sets [30]. Tanay et al. [29] combine a graph-theoretic approach with a statistical modeling of the data to discover biclusters in large gene expression data sets. Ben-Dor et al. [3] introduce a new notion of a bicluster called *order preserving submatrix*, which is a group of genes whose expression level induces a linear ordering across a subset of the conditions. Murali and Kasif [22] (see also [25]) propose the concept of *xmotif*, which is defined as a subset of genes whose expression is simultaneously conserved for a subset of samples.

As we were writing this document, we became aware of two other contributions to the subject, by Sheng et al. [27], and Mishra et al. [21], that use a randomized approach similar with the work described here. Sheng et al. [27] propose a randomized algorithm based on Gibbs sampling to discover large biclusters in gene expression data. Their model of a bicluster is probabilistic, that is, each entry of the matrix is associated with a probability. Mishra et al. [21] are concerned with the problem of finding ε -biclques which maximizes the number of edges.¹ Given a bipartite graph (U, V, E) , a subgraph (U', V') is ε -biclque if each vertex in U' is a neighbor of at least $(1 - \varepsilon)$ fraction of vertices in V' . The authors give an efficient randomized algorithm that finds the largest ε -biclque, but no experimental results are reported. Finally, in Koyuturk et al. [16] dense biclusters were found using a combination of randomization and linear algebra algorithms.

As shown in papers [22,27], the problem of biclustering gene expression data can be formulated on a discrete domain, by first discretizing the gene expression matrix into a matrix over a finite alphabet. The simplifying assumption is that the set of states in which each gene operates is finite, such as up-regulated, down-regulated or unchanged. Once the data is discretized into strings where each symbol corresponds to a state, the biclustering problem reduces to the problem of finding a subset of the rows and a subset of the columns such that the submatrix induced has the property that each row reads the same string. Such a submatrix would therefore correspond to a group of genes that exhibit a coherent pattern of states over a subset of conditions. This is indeed the formulation of the problem that we define in Section 2, which is first proved to be **NP**-complete. In Section 3 we present a randomized algorithm which is efficient and easy to understand and implement. Section 4 presents an asymptotic analysis that allows one to determine the statistical significance of the solution. Finally, in Section 5 we report simulation results on synthetic data.

2. Notations and problem definition

We use standard concepts and notation about strings. The set Σ denotes a non-empty *alphabet of symbols* and a *string* over Σ is an ordered sequence of symbols from the alphabet. We use the variable a as a shorthand

¹ The connection between biclques and bicluster will be explained in detail in Section 2.

for the cardinality of the set Σ , that is, $a = |\Sigma|$. Given a string x , the number of symbols in x defines the *length* $|x|$ of x .

Similarly, we can define a two-dimensional $n \times m$ string (or matrix) $X \in \Sigma^{n \times m}$ over the alphabet Σ . The element (i, j) of X is denoted by $X_{[i,j]}$. A *row selection* of size k of X is defined as the subset of the rows $R = \{i_1, i_2, \dots, i_k\}$, where $1 \leq i_s \leq n$ for all $1 \leq s \leq k$. Similarly, a *column selection* of size l of X is defined as a subset of the columns $C = \{j_1, j_2, \dots, j_l\}$, where $1 \leq j_t \leq m$ for all $1 \leq t \leq l$.

The submatrix $X_{(R,C)}$ induced by the pair (R, C) is defined as the matrix

$$X_{(R,C)} = \begin{pmatrix} X_{[i_1,j_1]} & X_{[i_1,j_2]} & \cdots & X_{[i_1,j_l]} \\ X_{[i_2,j_1]} & X_{[i_2,j_2]} & \cdots & X_{[i_2,j_l]} \\ \cdots & \cdots & \cdots & \cdots \\ X_{[i_k,j_1]} & X_{[i_k,j_2]} & \cdots & X_{[i_k,j_l]} \end{pmatrix}.$$

Given a selection of rows R , we say that a column j , $1 \leq j \leq m$, is *clean* with respect to R if the symbols in the j th column of X restricted to the rows R , are identical.

Example 1. Given the 6×6 matrix

$$X = \begin{pmatrix} \mathbf{001020} \\ \mathbf{100100} \\ \mathbf{301120} \\ \mathbf{201020} \\ \mathbf{131111} \\ \mathbf{110120} \end{pmatrix}$$

over the alphabet $\Sigma = \{\mathbf{0}, \mathbf{1}, \mathbf{2}, \mathbf{3}\}$, a selection $(R, C) = (\{2, 5, 6\}, \{1, 4, 6\})$ results in the matrix

$$X_{(R,C)} = \begin{pmatrix} \mathbf{110} \\ \mathbf{111} \\ \mathbf{110} \end{pmatrix}.$$

Given the row selection $R = \{2, 5, 6\}$, columns 1 and 4 are clean.

The problem addressed in this paper is defined as follows.

LARGEST BICLUSTER(f) problem
Instance: A matrix $X \in \Sigma^{n \times m}$ over the alphabet Σ .
Question: Find a row selection R and a column selection C such that the rows of $X_{(R,C)}$ are identical strings and the objective function $f(X_{(R,C)})$ is maximized.

Some examples of objective functions are the following:

- $f_1(X_{(R,C)}) = |R| + |C|$;
- $f_2(X_{(R,C)}) = |R|$ provided that $|C| = |R|$; and
- $f_3(X_{(R,C)}) = |R||C|$.

Example 2. Assume X to be the matrix defined in Example 1. If we choose the objective functions f_2 there are three solutions, namely $(\{1, 3, 4\}, \{2, 3, 5\})$ (corresponding to the string $\mathbf{012}$), $(\{1, 3, 4\}, \{2, 3, 6\})$ (corresponding to $\mathbf{010}$) and $(\{1, 3, 4\}, \{3, 5, 6\})$ (corresponding to $\mathbf{120}$). If we choose the objective functions f_1 or f_3 , the solution is $(\{1, 3, 4\}, \{2, 3, 5, 6\})$ (corresponding to $\mathbf{0120}$).

The problem in general may have multiple solutions which optimize the objective function. The solutions may also “overlap”, that is, they may share some elements of the original matrix.

The computational complexity of this family of problems depends on the objective function f . In the literature, the problem has been studied mostly from a graph-theoretical viewpoint which corresponds to the special case $\Sigma = \{0, 1\}$.

In fact, observe that a matrix $X \in \{0, 1\}^{n \times m}$ is the adjacency matrix of a bipartite graph $G = (V_1, V_2, E)$ with $|V_1| = n$ and $|V_2| = m$. An edge $(i, j) \in E$ connects node $i \in V_1$ to node $j \in V_2$ if $X_{i,j} = 1$. Thus, a submatrix of 1s in X corresponds to a subgraph of G which is completely connected. Such a subgraph is called a *biclique*. Because of this relation, we use the terms “submatrix”, “biclique”, and “bicluster” interchangeably.

When the alphabet is binary and we are looking for the largest submatrix composed only by 1s,² as for example in [16], the LARGEST BICLUSTER reduces to well-known problems on bipartite graphs. More specifically, the LARGEST BICLUSTER problem associated with objective function f_1 is known as the MAXIMUM VERTEX BICLIQUE problem, and it can be solved in polynomial time because it is equivalent to the maximum independent set in bipartite graphs which, in turn, can be solved by a minimum cut algorithm (see, e.g., [14]). The same problem with objective function f_2 over a binary alphabet is called BALANCED COMPLETE BIPARTITE SUBGRAPH problem or BALANCED BICLIQUE problem and it is listed as GT24 among the **NP**-complete problems in Garey and Johnson’s book [8] (see also [9]).

The LARGEST BICLUSTER problem with objective function f_3 and $\Sigma = \{0, 1\}$ is called MAXIMUM EDGE BICLIQUE problem. The problem requires to find the biclique which has the maximum number of edges. The problem is proved to be **NP**-complete in [24] by reduction from 3SAT. The weighted version of this problem is shown **NP**-complete by Dawande et al. [5].

In [14] Hochbaum studies a problem related to MAXIMUM EDGE BICLIQUE, which is the problem of finding the number of edges that need to be deleted so that the resulting graph is a biclique. Hochbaum describes a two-approximation algorithm based on LP-relaxation. According to Pasechnik [23] this approximation ratio does not hold for the original MAXIMUM EDGE BICLIQUE problem. Pasechnik shows a semidefinite relaxation, and claims that his relaxation is in general better than [14].

The following theorem establishes the hardness of the problem of finding the largest area bicluster over a general alphabet. For lack of space the proof is omitted.

Theorem 1. *The decision problem associated with LARGEST BICLUSTER(f_3) is **NP**-complete.*

Proof. The problem is trivially in **NP**. We reduce MAXIMUM EDGE BICLIQUE to LARGEST BICLUSTER(f_3). Observe that LARGEST BICLUSTER is more general than the biclique problems on graphs, for two reasons: (1) in the LARGEST BICLUSTER problem the alphabet is not necessarily binary and (2) a biclique corresponds to a submatrix of 1s in the corresponding adjacency matrix, while LARGEST BICLUSTER can return submatrices with 0s as well.

Assume that we are given an instance of the biclique problem represented as an adjacency matrix $X \in \{0, 1\}^{n \times m}$. Process X as follows. Replace each 0 in X by a unique new symbol. Now LARGEST BICLUSTER will be forced to return a matrix of 1s, which corresponds to a biclique in the bipartite graph. \square

By the same approach, LARGEST BICLUSTER(f_2) can also be proved to be **NP**-complete. In the rest of this paper we will concentrate our attention on the problem of finding the largest-area bicluster. For practical reasons that will become apparent in Section 3, the objective function that we are maximizing is

$$\tilde{f}_3(X_{(R,C)}, \hat{r}, \hat{c}) = |R||C| \quad \text{provided that } |R| \geq \hat{r} \text{ and } |C| \geq \hat{c},$$

where \hat{r} and \hat{c} are two input parameters.

3. Randomized search

Given that LARGEST BICLUSTER(f_3) problem is **NP**-complete, it is unlikely that a polynomial time algorithm could be found. In this paper, we present a randomized algorithm which finds a maximal solution with probability $1 - \varepsilon$, where $0 < \varepsilon < 1$.

Assume that we are given a large matrix $X \in \Sigma^{n \times m}$ in which a submatrix $X_{(R^*, C^*)}$ is implanted. Assume also that the submatrix $X_{(R^*, C^*)}$ is maximal. To simplify the notation, let $r^* \equiv |R^*|$ and $c^* \equiv |C^*|$.

The idea behind the algorithm comes from the following simple observation. Observe that if we knew R^* , then C^* could be determined by selecting the clean columns with respect to R^* . If instead we knew C^* , then R^* could be obtained by taking the maximal set of rows which read the same string. Unfortunately, neither R^* nor C^* is known. Our

² In general, a solution of the largest bicluster can contain a column of zeros, as long as they appear in all rows of the submatrix.

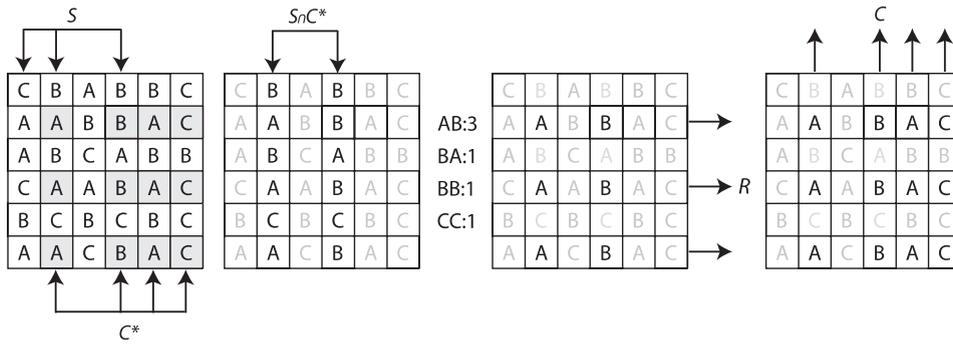


Fig. 1. An illustration of a recovery of the embedded matrix by random projections. C^* is the set of columns containing the embedded submatrix. S is a random selection of columns. By following the steps described in the text, the correct solution can be easily retrieved.

approach is to “sample” the matrix by random projections, with the expectation that at least some of the projections will overlap with the solution (R^*, C^*) . Clearly, one can project either rows or columns. In what follows we describe how to retrieve the solution by sampling columns.

The algorithm works as follows. Select a random subset S of size k uniformly from the set of columns $\{1, 2, \dots, m\}$. Assume for the time being that $S \cap C^* \neq \emptyset$. If we knew $S \cap C^*$, then (R^*, C^*) could be determined by the following three steps: (1) select the string(s) w that appear exactly r^* times in the rows of $X_{[1:n, S \cap C^]}$, (2) set R^* to be the set of rows in which w appears and (3) set C^* to be the set of clean columns corresponding to R^* .

Example 3. For example, consider the matrix in Fig. 1. The shaded cells contain the embedded submatrix, and $r^* = 3$, $c^* = 4$. Once $S \cap C^*$ is determined, the string AB is chosen since it is the only one that has frequency equal to r^* . Then, R is determined by the location of AB and C is determined by selecting the clean columns of R .

The algorithm would work, but there are a few problems that are still unresolved. First, the set $S \cap C^*$ could be empty. The solution is to try different random projections S , relying on the argument that the probability that $S \cap C^* \neq \emptyset$ at least once will approach one with more and more projections. The second problem is that we do not really know $S \cap C^*$. But, certainly $S \cap C^* \subseteq S$, so our approach is to check all possible subsets $U \subseteq S$ such that $|U| \geq k_{\min}$, where $1 \leq k_{\min} \leq k$ is a user-defined parameter. The final problem is that we assumed that we knew r^* , but we do not. The solution is to introduce a row threshold parameter, called \hat{r} , that replaces r^* .

As it turns out, we need another parameter to avoid producing solutions with too few columns. The column threshold \hat{c} is used to discard submatrices whose number of columns is smaller than \hat{c} . The algorithm considers all the submatrices which satisfy the user-defined row and column threshold as candidates. Among all candidate submatrices, only the ones that maximize the total area are kept. A sketch of the algorithm is shown in Fig. 2. As noted in the introduction, a very similar strategy was developed independently and concurrently by Mishra et al. [21].

The algorithm depends on five key parameters, namely the projection size k , the minimum subset size k_{\min} , the row threshold \hat{r} , the column threshold \hat{c} , and the number of iterations t . We discuss how to choose each of these in the rest of the section.

Parameter selection: The projection size k is determined by a probabilistic argument. To simplify our presentation in this section, we assume that all symbols are generated with the same probability $1/a$ (for a generalization see the next section). It is well-known that in a random string of size m over an alphabet of size a , the number of occurrences of substrings has two different probabilistic regimes: (1) Gaussian distributed for strings shorter than $\log_a m$ and (2) Poisson distributed for strings longer than $\log_a m$ (see, e.g., [26]).³ Based on this observation, when $k_{\min} = k$ we argue that $k = \log_a m$ is the optimal trade-off between generating too many trivial solutions (k too small) and potentially missing the solution (k too large). At the end of this section we will see that $k = \log_a m$ corresponds to the minimum in the graph associated with the probability of missing a solution (see Fig. 3.) This value of k has also been confirmed to be the optimal choice in our simulations. When $k_{\min} = 1$, then k can be chosen significantly larger, but this will

³ In general, the cut of depends whether $mP(\text{string}) = O(1)$ or $mP(\text{string}) \rightarrow \infty$.

```

LARGEST_BICLUSTER_C( $X, t, k, k_{\min}, \hat{r}, \hat{c}$ )
INPUT:  $X$  is a  $n \times m$  matrix over  $\Sigma$ 
       $t$  is the number of iterations
       $k$  is the projection size
       $k_{\min}$  is the size of the smallest subset of the projection
       $\hat{r}, \hat{c}$  are the “thresholds” on the number of rows and columns, resp.
1  repeat  $t$  times
2  select randomly a subset  $S$  of columns such that  $|S| = k$ 
3  for all subsets  $U \subseteq S$  such that  $|U| \geq k_{\min}$  do
4     $D \leftarrow$  all strings induced by  $X_{[1:n,U]}$  that appear at least  $\hat{r}$  times
5    for each string  $w$  in  $D$ 
6       $V \leftarrow$  rows corresponding to  $w$ 
7       $Z \leftarrow$  all “clean” columns corresponding to  $V$ 
8      if  $|Z| \geq \hat{c}$  then save  $(V, Z)$ 
9  return the  $(V, Z)$  that maximizes  $f$ 

```

Fig. 2. A sketch of the algorithm that discovers large biclusters (sampling columns).

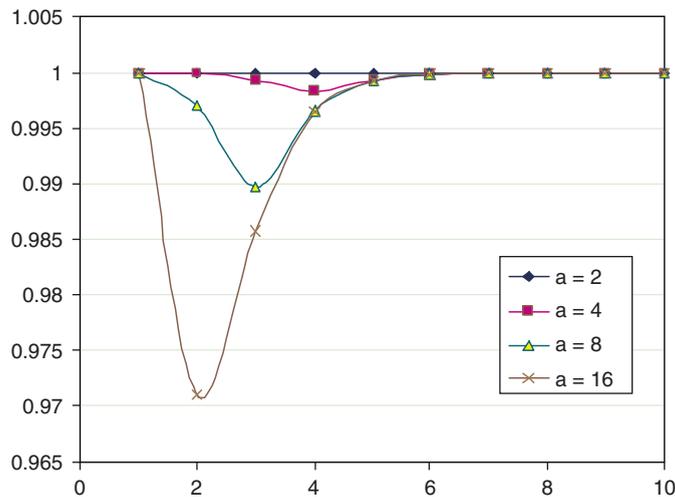


Fig. 3. The graph of the function $\tilde{\alpha}$ for different choices of the projection size k (x -axis) and alphabet size a . The fixed parameters are $n = 256$, $m = 256$, $r^* = 64$, $c^* = 64$. Note that the minimum of the function $\tilde{\alpha}$ is at $k \approx \log_a m$.

adversely affect the running time. An experimental comparison between $k_{\min} = k$ (i.e., no subsets), and $k_{\min} = 1$ (i.e., all subsets) is reported in Section 5.1.

The thresholds \hat{r} and \hat{c} are associated with the uncertainty on the size of the largest submatrix r^* , c^* for a particular input instance. There may be situations in which the user has already a good idea about r^* , c^* . If however r^* and c^* are completely unknown, then our target will be to find “statistically significant” biclusters. In Section 4 we will present a theorem (Theorem 2) which gives the expected number of columns of the largest submatrix in a random matrix, when the number of rows is fixed. Based on this, we propose the following trial-and-error strategy. Set \hat{r} to some value between 1 and n , and use Theorem 2 to set the value \hat{c} . Run the algorithm. If the algorithm returns too many solutions, try to increase \hat{r} and update \hat{c} correspondingly. If there are no solutions, lower the value of \hat{r} and repeat. Observe that the number of choices for \hat{r} is finite since $\hat{r} \in [1, n]$. By using Theorem 2 to set the threshold \hat{c} , we are trying to filter out submatrices whose size is small enough that they could appear in totally random matrices.

The worst case time complexity of the algorithm LARGEST_BICLUSTER_C is $O(tk \sum_{j=k_{\min}}^k \binom{k}{j} (kn + nm))$, where k_{\min} is size of the smallest subset of S that the algorithm generates. If $k_{\min} = 1$, the complexity becomes $O(tk2^k(kn + nm))$. Although the complexity is exponential in k , the optimal value of k is $O(\log_a m)$, which makes the algorithm run in $O(tkm(kn + nm))$ time.

Because of the randomized nature of the approach, there is no guarantee that the algorithm will find the solution after a given number of iterations. We therefore need to choose t so that the probability that the algorithm will recover the solution in at least one of the t trials is $1 - \varepsilon$, where $0 < \varepsilon < 1$ is a user-defined parameter.

Let $\alpha(n, m, k, r^*, c^*, a)$ be the probability of missing the solution in one of the trials assuming that r^* and c^* are known and that $k_{\min} = 1$. There are two disjoint cases in which the algorithm can miss (R^*, C^*) . The first is when the random projection S misses completely C^* , i.e., $S \cap C^* = \emptyset$. The second is when $S \cap C^* = U \neq \emptyset$ but the string w chosen by the algorithm among the rows $X_{[1:n,U]}$ also appears in another row that does not belong to the set R^* of the real solution. In this case, the algorithm will select a set of rows larger than R^* and thus miss the solution. Hence, we have

$$\begin{aligned} \alpha(n, m, k, r^*, c^*, a) &= \Pr\{S \cap C^* = \emptyset\} + \sum_{i=1}^k \Pr\{|S \cap C^*| = i \text{ and } |R| > r^*\} \\ &= \Pr\{S \cap C^* = \emptyset\} + \sum_{i=1}^k \Pr\{|R| > r^* \text{ given } |S \cap C^*| = i\} \Pr\{|S \cap C^*| = i\}. \end{aligned}$$

Let Y be the random variable associated with the size of the set $S \cap C^*$, that is, $Y = |S \cap C^*|$. Since we are sampling without replacement, Y follows the hyper-geometric distribution.

$$\Pr\{Y = 0\} = \binom{m - c^*}{k} / \binom{m}{k} \quad \text{and} \quad \Pr\{Y = i\} = \binom{c^*}{i} \binom{m - c^*}{k - i} / \binom{m}{k}.$$

In order to compute the probability of missing the solution given that $|S \cap C^*| = i$, we have to estimate how likely a string w belonging to some of the rows of $X_{[1:n,U]}$ is more frequent than r^* . In fact, if the string w appears more than r^* times, the algorithm will select a larger set of rows, and it will miss the solution.

Hereafter, we assume that the pattern is known to be w and its probability is $p(w)$. For example, if the symbols in the matrix X are generated by a symmetric Bernoulli i.i.d. model, then simply $p(w) = a^{-i}$. The probability that w will never appear in the other $n - r^*$ rows is $(1 - p(w))^{n-r^*}$, and therefore

$$\Pr\{|R| > r^* \text{ given } |S \cap C^*| = i\} = 1 - (1 - p(w))^{n-r^*}.$$

Combining all together, the probability of missing the solution in one iteration is given by

$$\alpha(n, m, k, r^*, c^*, a) = \frac{\binom{m - c^*}{k} + \sum_{i=1}^k (1 - (1 - p(w))^{n-r^*}) \binom{c^*}{i} \binom{m - c^*}{k - i}}{\binom{m}{k}}.$$

Now suppose we want the probability of missing the solution to be smaller than a given ε , $0 < \varepsilon < 1$. We can obtain the number of iterations t by solving the inequality $(\alpha(n, m, k, r^*, c^*, a))^t \leq \varepsilon$, which gives

$$t \geq \frac{\log \varepsilon}{\log \alpha(n, m, k, r^*, c^*, a)}. \tag{1}$$

This bound on the number of iterations has been verified by our experimental results (compare Table 1 with our experimental results shown in Fig. 5). For example, by setting $a = 4$, $k = 4$, $\varepsilon = 0.7$, Eq. (1) gives $t = 90$ iterations whereas the experimental results show that with 90 iterations we obtain a performance of $\varepsilon = 0.689$.

The worst-case time complexity of LARGEST_BICLUSTER_C is bounded by $O(t \sum_{j=k_{\min}}^k \binom{k}{j} (kn + nm))$. If $k_{\min} = 1$, then the time complexity becomes $O(t 2^k (kn + nm))$. Although the complexity is exponential in k , choosing k to be $O(\log_a m)$ makes the algorithm run in $O(tm^{1/\log_2 a} (kn + nm))$ time.

The probability of missing the solution changes significantly when we set $k_{\min} = k$. In this case, we are not checking any of the subsets of S , but we simply rely on the fact that eventually one of the random projections S will end up completely contained in C^* , in which case we have a chance to find the solution.

Since we avoid checking the $O(2^k)$ subsets of S , the number of iterations t to achieve the same level of performance of the case $k_{\min} = 1$ must be significantly larger. Indeed, by a similar argument as we did for $k_{\min} = 1$, the probability

Table 1

The estimated number of iterations for a matrix 256×256 with a submatrix 64×64 , for different choices of ε , alphabet size a , and projection size k (sampling columns)

ε	$a = 2, k = 8$	$a = 4, k = 4$	$a = 8, k = 3$	$a = 16, k = 2$	$a = 32, k = 2$
0.005	18,794	1342	306	179	99
0.05	10,626	759	173	101	56
0.1	8168	583	133	78	43
0.2	5709	408	93	54	30
0.3	4271	305	70	41	23
0.4	3250	232	53	31	17
0.5	2459	176	40	23	13
0.6	1812	129	29	17	10
0.7	1265	90	21	12	7
0.8	792	57	13	8	4
0.9	374	27	6	4	2

```

LARGEST_BICLUSTER_R( $X, t, k, k_{\min}, \hat{r}, \hat{c}$ )
INPUT:  $X$  is a  $n \times m$  matrix over  $\Sigma$ 
       $t$  is the number of iterations
       $k$  is the projection size
       $k_{\min}$  is the size of the smallest subset of the projection
       $\hat{c}, \hat{r}$  are the “thresholds” on the number of columns and rows, resp.
1  repeat  $t$  times
2  select randomly a subset  $S$  of rows such that  $|S| = k$ 
3  for all subsets  $U \subseteq S$  such that  $|U| \geq k_{\min}$  do
4     $Z \leftarrow$  all columns induced by  $X_{[U,1:m]}$  that are clean
5     $w \leftarrow$  string induced by any of the row of  $X_{[U,Z]}$ 
6    if  $|Z| \geq \hat{c}$  then
7       $V \leftarrow$  rows of  $X_{[1:n,Z]}$  which contain  $w$ 
8      if  $|V| \geq \hat{r}$  then save  $(V, Z)$ 
9  return the  $(V, Z)$  that maximizes  $f$ 

```

Fig. 4. A sketch of the algorithm that discovers large biclusters (sampling rows).

of missing the solution when $k_{\min} = k$ can be estimated by the following formula

$$\begin{aligned}
\tilde{\alpha}(n, m, k, r^*, c^*, a) &= \Pr\{|S \cap C^*| < k\} + \Pr\{|S \cap C^*| = k \text{ and } |R| > r^*\} \\
&= 1 - \Pr\{|S \cap C^*| = k\} + \Pr\{|S \cap C^*| = k \text{ and } |R| > r^*\} \\
&= 1 - \left(\binom{c^*}{k} / \binom{m}{k} \right) + \left((1 - (1 - p(w))^{n-r^*}) \binom{c^*}{k} / \binom{m}{k} \right) \\
&= 1 - \left((1 - p(w))^{n-r^*} \binom{c^*}{k} / \binom{m}{k} \right),
\end{aligned}$$

where w is the string of size k selected by the algorithm. In the symmetric model, $p(w) = a^{-k}$.

In Fig. 3 we plotted the function $\tilde{\alpha}$ for different choices of k and a ($n = 256, m = 256, r^* = 64, c^* = 64$). As discussed in the beginning of this section, the value of k which minimizes the value of $\tilde{\alpha}$ is $k \approx \log_a m$.

As mentioned above, we also have the option to project the rows instead of the columns, which would result in a slightly different algorithm that we called LARGEST_BICLUSTER_R. We illustrate the algorithm in Fig. 4.

The worst case time complexity of the algorithm LARGEST_BICLUSTER_R is $O(t \sum_{j=k_{\min}}^k \binom{k}{j} (km + nm))$, which becomes $O(t2^k (km + nm))$ when $k_{\min} = 1$. The probability of missing a solution can be obtained using the same

argument as above, and it is given by

$$\beta(n, m, k, r^*, c^*, a) = \frac{\binom{n-r^*}{k} + \sum_{i=1}^k (1 - (1 - p(w))^{n-c^*}) \binom{r^*}{i} \binom{n-r^*}{k-i}}{\binom{n}{k}}.$$

Both strategies were implemented and tested extensively. Results are reported in Section 5.

4. Statistical analysis

We now analyze the statistical significance of finding a large submatrix of size $r \times c$ hidden into a random $n \times m$ matrix over an alphabet of cardinality a . More specifically, we randomly generate a matrix $X \in \Sigma^{n \times m}$ using a memoryless source with parameters $\{p_1, \dots, p_a\}$ where p_i is the probability of the i th symbol in Σ . Given X , the goal is to characterize asymptotically the size of the largest submatrix in X .

For convenience of notation, let us call $P_r = p_1^r + p_2^r + \dots + p_a^r$ the probability of observing a clean column over r rows, and let us define $H(x) = -x \ln x - (1-x) \ln(1-x)$. We observe that $P_r = p_{\max}(1 + O(\rho^r))$ as $r \rightarrow \infty$ where $\rho < 1$ and $p_{\max} = \max\{p_1, \dots, p_a\}$.

The first result characterizes the number of columns of the largest bicluster, when the number of rows is set to be r . More precisely, let $\mathbf{Z}_{j_1, \dots, j_r}$ be the number of clean columns in rows $1 \leq j_1 < \dots < j_r \leq n$. Our goal is to find

$$C_{n,m,r,a} = \max_{1 \leq j_1 < \dots < j_r \leq n} \{\mathbf{Z}_{j_1, \dots, j_r}\}. \tag{2}$$

The following characterize the typical behavior of $C_{n,m,r,a}$.

Theorem 2. Let $C_{n,m,r,a}$ be the random variable associated with the number of columns of the submatrix with the largest area in a matrix $X \in \Sigma^{n \times m}$ generated from a memoryless source, once the number of rows r is fixed as defined in (2). Then

$$C_{n,m,r,a} \leq m P_r + \sqrt{2 P_r (1 - P_r) m F(n, r)} \equiv C_{\max}$$

with high probability and as $n \rightarrow \infty$, where

$$F(n, r) = \begin{cases} r \log n & \text{if } r = o(n), \\ n H(\alpha) & \text{if } r = \alpha n \text{ and } 0 < \alpha < 1, \end{cases}$$

where $P_r \sim p_{\max}$ when r is large.

Proof. Observe that

$$\Pr\{\mathbf{Z}_{j_1, \dots, j_r} = k\} = \binom{m}{k} P_r^k (1 - P_r)^{m-k},$$

and in particular, $\mathbf{E}[\mathbf{Z}_{j_1, \dots, j_r}] = m P_r$ and $\text{Var}[\mathbf{Z}_{j_1, \dots, j_r}] = m P_r (1 - P_r)$. We now proceed as follows:

$$\begin{aligned} C_{n,m,r,a} &= \max_{1 \leq j_1 < \dots < j_r \leq n} \{\mathbf{Z}_{j_1, \dots, j_r}\} \\ &= \sqrt{\text{Var}[\mathbf{Z}_{j_1, \dots, j_r}]} \max_{1 \leq j_1 < \dots < j_r \leq n} \left\{ \frac{\mathbf{Z}_{j_1, \dots, j_r} - \mathbf{E}[\mathbf{Z}_{j_1, \dots, j_r}]}{\sqrt{\text{Var}[\mathbf{Z}_{j_1, \dots, j_r}]}} \right\} + \mathbf{E}[\mathbf{Z}_{j_1, \dots, j_r}] \\ &= \sqrt{\text{Var}[\mathbf{Z}_{j_1, \dots, j_r}]} \max_{1 \leq j_1 < \dots < j_r \leq n} \{\mathbf{Y}_{j_1, \dots, j_r}\} + \mathbf{E}[\mathbf{Z}_{j_1, \dots, j_r}], \end{aligned} \tag{3}$$

where $\mathbf{Y}_{j_1, \dots, j_r}$ are asymptotically normally distributed $N(0, 1)$. Observe that $\mathbf{Y}_{j_1, \dots, j_r}$ are not independent.

We now estimate $C'_{n,m,r,a} = \max_{1 \leq j_1 < \dots < j_r \leq n} \{\mathbf{Y}_{j_1, \dots, j_r}\}$. Since $\mathbf{Y}_{j_1, \dots, j_r}$ are normally distributed we have, when $r = o(n)$

$$\Pr\{C'_{n,m,r,a} > x\} \leq \binom{n}{r} \Pr\{\mathbf{Y}_{j_1, \dots, j_r} > x\} \sim \frac{n^r e^{-x^2/2}}{r! x}, \tag{4}$$

where we approximate $\binom{n}{r} = n^r / r!$ for $r = o(n)$. Hence for $x = \sqrt{2 \log n^r}$ we find that

$$\Pr\{\mathbf{C}'_{n,m,r,a} > \sqrt{2 \log n^r}\} \leq O(1 / \log n) \rightarrow 0$$

as $n \rightarrow \infty$. We have

$$\mathbf{C}_{n,m,r,a} \leq m P_r + \sqrt{2 P_r (1 - P_r) m r \log n}$$

with high probability.

If instead $r = \alpha n$, then we must use (cf. [28])

$$\binom{n}{\alpha n} \sim \frac{e^{nH(\alpha)}}{\sqrt{2\pi n\alpha(1-\alpha)}},$$

where $H(\alpha) = -\alpha \log \alpha - (1 - \alpha) \log(1 - \alpha)$. In this case, Eq. (4) should be replaced by

$$\Pr\{\mathbf{C}'_{n,m,r,a} > x\} \leq \binom{n}{r} \Pr\{\mathbf{Y}_{j_1, \dots, j_r} > x\} \leq \frac{1}{\sqrt{2\pi n\alpha(1-\alpha)}} \frac{e^{nH(\alpha) - x^2/2}}{x}.$$

Thus $x = \sqrt{2nH(\alpha)}$ makes the above probability small and $\mathbf{C}'_{n,m,r,a} \sim \sqrt{2nH(\alpha)}$ with high probability leading finally to

$$\mathbf{C}_{n,m,r,a} \leq m P_r + \sqrt{2 P_r (1 - P_r) m n H(\alpha)}$$

with high probability. \square

The practical implications of Theorem 2 are twofold. First, the expected number of columns can be used to set the column threshold parameter $\hat{c} \gg \max\{C_{\max}, 1\}$. That allows the algorithm to avoid considering statistically non-significant submatrices. Second, observe that when $\log n = o(m)$, then the dominant term of C_{\max} is the average, say $\mathbf{E}[C]$, of the number of clean columns, that is, $\mathbf{E}[C] = m P_r$. This implies $C_{\max} / \mathbf{E}[C] \leq 1 + o(1)$ for $\log n = o(m)$, and therefore with high probability any algorithm is asymptotically optimal. Clearly, this is not true for $r = \alpha n$.

Now we deal with the optimization problem associated with f_3 (in a probabilistic setting). More precisely, we define

$$\mathbf{A}_{n,m,a} = \max_{\mathbf{i}, \mathbf{j}} \{A_{\mathbf{i}, \mathbf{j}}\}, \tag{5}$$

where $\mathbf{i} = (i_1, \dots, i_r)$ and $\mathbf{j} = (j_1, \dots, j_c)$ such that $1 \leq i_1 < \dots < i_r \leq n$ and $1 \leq j_1 < \dots < j_c \leq m$. Our second result characterizes the typical behavior of $\mathbf{A}_{n,m,a}$. It is easy to realize that $\mathbf{A}_{n,m,a} = O(\log nm)$ when $r = o(n)$ and $c = o(m)$, thus for the largest area either $r = O(n)$ or $c = O(m)$, as shown below.

Theorem 3. *Let $\mathbf{A}_{n,m,a}$ be the random variable associated with the area of the largest submatrix in a matrix $X \in \Sigma^{n \times m}$, $m \leq n$, generated from a memoryless source. Then, with high probability for any $\varepsilon > 0$ and as $n \rightarrow \infty$*

$$\mathbf{A}_{n,m,a} \leq (1 + \varepsilon) r c,$$

where $r = n/2$ and $c = 2 \ln 2 / \ln p_{\max}^{-1}$.

Proof. First, observe that

$$\Pr\{\mathbf{A}_{n,m,a} > r \cdot c\} \leq \binom{n}{r} \binom{m}{c} P_r^c. \tag{6}$$

Now one needs to select r and c such that the right-hand side of Eq. (6) is small. Since we know that the largest bicluster is of order $O(n)$, we assume $r = \alpha n$ and $c = o(n)$. As above, we have that

$$\binom{n}{\alpha n} \sim \frac{e^{nH(\alpha)}}{\sqrt{2\pi n\alpha(1-\alpha)}},$$

Table 2
The statistics of large submatrices in a random $\{0, 1\}$ -matrix of size 256×256

Rows	Columns observed	Columns predicted
1	256	256
2	160	165.6771209
3	100	103.9626215
4	67	67.24371945
5	45	44.84053788
6	31	30.70906224
7	23	21.48364693
8	16	15.26873716

The second column reports the number of columns of the submatrices observed in a random matrix, whereas the third reports the prediction based on Theorem 2.

where $H(\alpha)$ is a natural entropy. Since $P_{\alpha n} \sim p_{\max}^{\alpha n}$, we arrive at

$$\Pr\{\mathbf{A}_{n,m,a} > \alpha cn\} \leq \frac{e^{nH(\alpha)+c \ln m - rc \ln p_{\max}^{-1}}}{\sqrt{2\pi n\alpha(1-\alpha)}}.$$

Now we choose α^* such that

$$H(\alpha^*) = c\alpha^* \ln p_{\max}^{-1}. \tag{7}$$

If we select $r = (1 + \varepsilon)\alpha^*n$ we can make Eq. (6) small. In fact, since $m \leq n$, we have

$$\Pr\{\mathbf{A}_{n,m,a} > \alpha^*cn\} \leq e^{-\delta n} \rightarrow 0$$

for some $\delta > 0$. Thus the largest submatrix has $\mathbf{A}_{n,m,a} \leq (1 + \varepsilon)n\alpha^*c$ provided $c = O(1)$ (since otherwise (7) does not have a solution).

But we may still want to find the largest α^*c subject to constraint (7), that is, $\max \alpha^*c$ under the condition that $H(\alpha^*) = c\alpha^* \ln p_{\max}^{-1}$. This can be solved by setting $c = H(\alpha^*)/(\alpha^* \ln p_{\max}^{-1})$ and the maximizing

$$\alpha^*c = \alpha^* \frac{H(\alpha^*)}{\alpha^* \ln p_{\max}^{-1}} = \frac{H(\alpha^*)}{\ln p_{\max}^{-1}}. \tag{8}$$

Eq. (8) reaches the maximum when $\alpha^* = 0.5$. Then the optimal number of columns is $c = 2 \ln 2 / \ln p_{\max}^{-1}$. \square

The intuition behind Theorem 3 is that on random matrices one should expect the largest submatrix to be “skinny”, that is, a few columns and lots of rows, or vice versa. For example, we expect the largest submatrix in a random $\{0, 1\}$ -matrix of size 256×256 to be size 2×160 (see Table 2).

5. Implementation and experiments

We implemented column- and row-sampling algorithms in C++ and tested the programs on a desktop PC with a 1.2 GHz Athlon CPU and 1 GB of RAM, under Linux. Although the algorithms do not require sophisticated data structures, in order to carry out step 4 in the algorithm of Fig. 2, one needs a data structure to store the strings and their frequencies. Since k and a are usually not very large, our experience shows that a simple hash table (of size a^k) is a good choice. If a^k becomes too large, a trie would be a better data structure. If one uses the hash table, it is important to keep track of the non-zero entries in another balanced data structure. That would avoid the algorithm to spend $O(a^k)$ to search for the frequently occurring strings. Observe also that row-sampling algorithm (Fig. 4) does not require any hash table, or any other data structure. However, our experiments show that in order to get the same level of performance of the column sampling, the row sampling strategy needs a significantly larger projection k which adversely affects the running time.

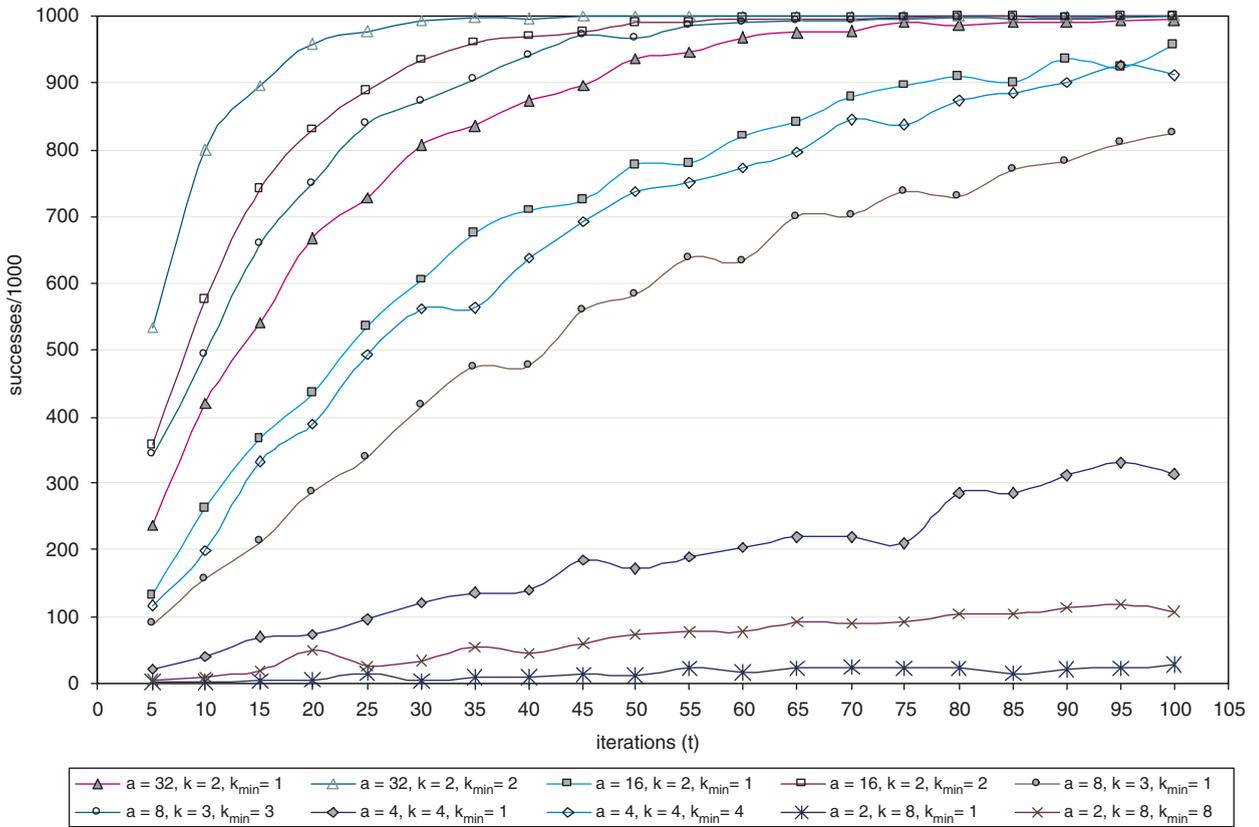


Fig. 5. Comparing the performance of the randomized algorithm LARGEST_BICLUSTER_C when $k_{\min} = k$ versus $k_{\min} = 1$, for different choices of the alphabet size a . The projection size is $k = \log_a m$.

Another issue is whether one should keep track of the projections generated so far to avoid generating duplicates. We studied this matter experimentally, and found that it is worthwhile to keep track of the projections in some balanced data structure only when k is small. If k is large, the overhead required to keep the data structure updated is much higher than the time wasted in processing the same projection multiple times.

5.1. Simulations

In order to evaluate the performance of the algorithms, we designed several simulation experiments. In these experiments we randomly generated one thousand 256×256 matrices of symbols drawn from an symmetric i.i.d. distribution over an alphabet of cardinality $a = 2, 4, 8, 16, 32$. Then, in each matrix we embedded a random 64×64 submatrix at random columns and random rows. We ran the algorithms for a few tens of iterations ($t = 5, \dots, 100$), and for each choice of t we measured the number of successes out of the 1000 distinct instances. Fig. 5 summarizes the performance of LARGEST_BICLUSTER_C, for several choices of alphabet size a and projection size k , and minimum subset size k_{\min} . Fig. 6 summarizes the performance of LARGEST_BICLUSTER_R under the same conditions.

In order to make a fair comparison between $k_{\min} = k$ and $k_{\min} = 1$, the number of iterations for the case $k_{\min} = k$ was multiplied by $2^k - 1$. Note that by doing so, we are assuming that one projection for $k_{\min} = 1$ takes about the same time as one projection for $k_{\min} = k$, which is not necessarily very accurate. Under this assumption, however, $k_{\min} = k$ outperforms $k_{\min} = 1$ (see Fig. 5). This not necessarily true in the row sampling strategy (see Fig. 6).

By comparing the performance of row sampling against column sampling, one can observe that if one uses the same set of parameters, column sampling always outperforms row sampling.

Unfortunately, we were unable to compare the performance of our randomized approach to other biclustering algorithms (e.g. [4,30,3,29,7,22,27,20]), because their notion of bicluster is generally different from ours.

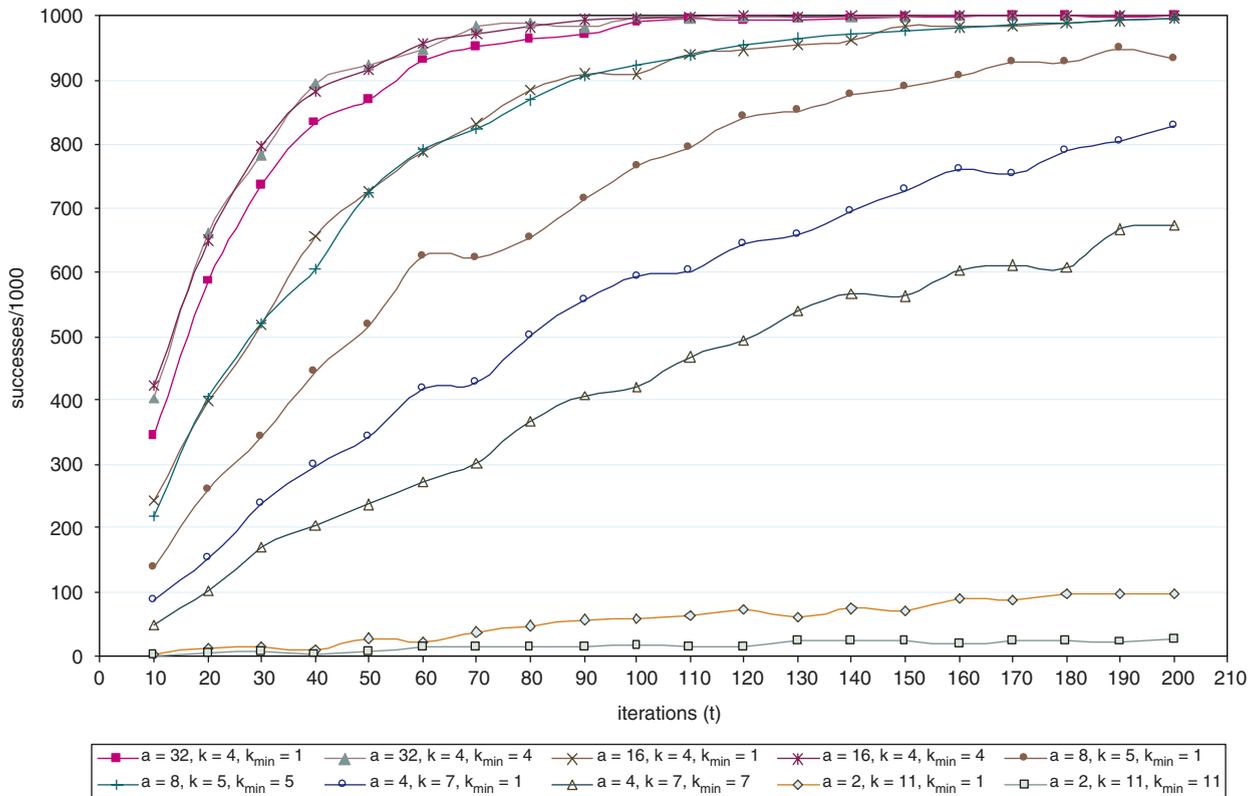


Fig. 6. Comparing the performance of the randomized algorithm LARGEST_BICLUSTER_R for different choices of the alphabet size a , the projection size k , and the smallest subset k_{\min} .

6. Conclusions

In this paper we have introduced the LARGEST BICLUSTER problem. This problem has a variety of applications ranging from computational biology to data mining. As far as we know, the pattern matching community has not looked yet at this problem from a combinatorial perspective. Unfortunately, the problem is generally NP complete.

Here we presented a rather simple algorithm based on random projections. Its performance with respect to the number of projection was carefully analyzed. We have also presented a probabilistic analysis of the LARGEST BICLUSTER problem, which allows one to determine the statistical significance of a solution.

Our approach performs remarkably well on synthetic data. On large alphabets, thirty or so iterations are enough to give a performance close to 100%. With respect to other biclustering algorithms (see e.g., [4,22,27]), our algorithm simultaneously discovers multiple solutions which satisfy the user-defined parameters without masking or changing the original data. In addition to this, the algorithm will never report solutions which are completely contained in other solutions.

Acknowledgments

This research was supported in part by NSF Grants CCR-0208709, DBI-0321756, NSF CAREER IIS-0447773 and the NIH Grant R01 GM068959-01.

References

- [1] C.C. Aggarwal, C. Procopiuc, J.L. Wolf, P.S. Yu, J.S. Park, Fast algorithms for projected clustering, in: Proc. ACM SIGMOD Internat. Conf. on Management of Data (SIGMOD-99), SIGMOD Record, Vol. 28(2), ACM Press, New York, 1999, pp. 61–72.

- [2] R. Agrawal, J. Gehrke, D. Gunopulos, P. Raghavan, Automatic subspace clustering of high dimensional data for data mining applications, in: Proc. ACM SIGMOD Internat. Conf. on Management of Data (SIGMOD-98), ACM SIGMOD Record, Vol. 27(2), ACM Press, New York, 1998, pp. 94–105.
- [3] A. Ben-Dor, B. Chor, R. Karp, Z. Yakhini, Discovering local structure in gene expression data: the order-preserving submatrix problem, in: Proc. Sixth Internat. Conf. on Computational Molecular Biology (RECOMB 2002), ACM Press, New York, 2002, pp. 45–55.
- [4] Y. Cheng, G.M. Church, Biclustering of expression data, in: Proc. Eighth Internat. Conf. on Intelligent Systems for Molecular (ISMB-00), AAAI Press, Menlo Park, CA, 2000, pp. 93–103.
- [5] M. Dawande, P. Keskinocak, J.M. Swaminathan, S. Tayur, On bipartite and multipartite clique problems, *J. Algorithms* 41 (2001) 388–403.
- [6] I. Dhillon, Co-clustering documents and words using bipartite spectral graph partitioning, in: Proc. Seventh ACM SIGKDD Internat. Conf. on Knowledge Discovery and Data Mining (KDD-01), ACM Press, New York, 2001, pp. 269–274.
- [7] I. Dhillon, S. Mallela, D. Modha, Information-theoretic co-clustering, in: Proc. Seventh ACM SIGKDD Internat. Conf. on Knowledge Discovery and Data Mining (KDD-03), ACM Press, New York, 2001, pp. 89–98.
- [8] M.R. Garey, D.S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, Freeman, New York, NY, 1979.
- [9] M. Grigni, F. Manne, On the complexity of the generalized block distribution, in: Proc. of Irregular'96, The Third Internat. Workshop on Parallel Algorithms for Irregularly Structured Problems, Lecture Notes in Computer Science, Vol. 1117, Springer, Berlin, 1996, pp. 319–326.
- [10] M. Gu, H. Zha, C. Ding, X. He, H. Simon, Spectral relaxation models and structure analysis for k -way graph clustering and bi-clustering, Technical Report CSE-01-007, Department of Computer Science and Engineering, Pennsylvania State University, 2001.
- [11] D. Hanisch, A. Zien, R. Zimmer, T. Lengauer, Co-clustering of biological networks and gene expression data, in: Proc. Eighth Internat. Conf. on Intelligent Systems for Molecular (ISMB-02), AAAI Press, Menlo Park, CA, 2002, pp. 145–154.
- [12] J.A. Hartigan, Direct clustering of a data matrix, *J. Amer. Statist. Assoc.* 67 (337) (1972) 123–129.
- [13] T. Hastie, R. Tibshirani, M. Eisen, A. Alizadeh, R. Levy, L. Staudt, W. Chan, D. Botstein, P. Brown, 'Gene shaving' as a method for identifying distinct sets of genes with similar expression patterns, *Genome Biol.* 1 (2) (2000) 1–21.
- [14] D.S. Hochbaum, Approximating clique and biclique problems, *J. Algorithms* 29 (1) (1998) 174–200.
- [15] Y. Kluger, R. Basri, J. Chang, M. Gerstein, Spectral biclustering of microarray data: coclustering genes and conditions, *Genome Res.* 13 (4) (2003) 703–716.
- [16] M. Koyuturk, W. Szpankowski, A. Grama, Biclustering gene-feature matrices for statistically significant dense patterns, in: IEEE Computer Society Bioinformatics Conf., Stanford, CA, 2004, pp. 480–483.
- [17] R. Kumar, P. Raghavan, S. Rajagopalan, A. Tomkins, Trawling the web for emerging cyber-communities, *Comput. Networks* 31 (11–16) (1999) 1481–1493.
- [18] L. Lazzeroni, A. Owen, Plaid models for gene expression data, *Statist. Sinica* 12 (1) (2002) 61–86.
- [19] J. Liu, J. Yang, W. Wang, Biclustering in gene expression data by tendency, in: IEEE Computational Systems Bioinformatics Conf. (CSB'04), 2004, pp. 224–235.
- [20] A.A. Melkman, E. Shaham, Sleeved coclustering, in: Proc. 2004 ACM SIGKDD Internat. Conf. on Knowledge Discovery and Data Mining, Seattle, WA, 2004, pp. 635–640.
- [21] N. Mishra, D. Ron, R. Swaminathan, On finding large conjunctive clusters, in: Proc. ACM Conf. on Computational Learning Theory (COLT'03), 2003, pp. 448–462.
- [22] T.M. Murali, S. Kasif, Extracting conserved gene expression motifs from gene expression data, in: Proc. Pacific Symp. on Biocomputing (PSB'03), 2003, pp. 77–88.
- [23] D.V. Pasechnik, Bipartite sandwiches, Technical Report, available at (<http://arXiv.org/abs/math.CO/9907109>), 1999.
- [24] R. Peeters, The maximum-edge biclique problem is NP-complete, Technical Report 789, Faculty of Economics and Business Administration, Tilberg University, 2000.
- [25] M. Procopiuc, M. Jones, P. Agarwal, T.M. Murali, A Monte Carlo algorithm for fast projective clustering, in: Proc. 2002 Internat. Conf. on Management of Data (SIGMOD'02), 2002, pp. 418–427.
- [26] G. Reinert, S. Schbath, M.S. Waterman, Probabilistic and statistical properties of words: an overview, *J. Comput. Biol.* 7 (2000) 1–46.
- [27] Q. Sheng, Y. Moreau, B.D. Moor, Biclustering microarray data by Gibbs sampling, in: Proc. of European Conf. on Computational Biology (ECCB'03), 2003, pp. 196–205.
- [28] W. Szpankowski, *Average Case Analysis of Algorithms on Sequences*, Wiley Interscience, New York, 2001.
- [29] A. Tanay, R. Sharan, R. Shamir, Discovering statistically significant biclusters in gene expression data, in: Proc. 10th Internat. Conf. on Intelligent Systems for Molecular Biology (ISMB'02), Bioinformatics, Vol. 18, 2002, pp. S136–S144.
- [30] H. Wang, W. Wang, J. Yang, P.S. Yu, Clustering by pattern similarity in large data sets, in: Proc. 2002 ACM SIGMOD Internat. Conf. on Management of Data (SIGMOD-02), ACM Press, New York, 2002, pp. 394–405.
- [31] J. Yang, H. Wang, W. Wang, P.S. Yu, Enhanced biclustering on gene expression data, in: IEEE Symp. on Bioinformatics and Bioengineering (BIBE'03), 2003, pp. 321–327.
- [32] L. Zhang, S. Zhu, A new clustering method for microarray data analysis, in: Proc. First IEEE Computer Society Bioinformatics Conf. (CSB'02), IEEE Press, New York, 2002, pp. 268–275.