

Query-based Biclustering using Formal Concept Analysis

Faris Alqadah*

Joel S. Bader[†]

Rajul Anand[‡]

Chandan K. Reddy[§]

Abstract

Biclustering methods have proven to be critical tools in the exploratory analysis of high-dimensional data including information networks, microarray experiments, and bag of words data. However, most biclustering methods fail to answer specific questions of interest and do not incorporate prior knowledge and expertise from the user. To this end, query-based biclustering algorithms that are recently developed in the context of microarray data utilize a set of seed genes provided by the user which are assumed to be tightly co-expressed or functionally related to prune the search space and guide the biclustering algorithm. In this paper, a novel **Query-Based Bi-Clustering** algorithm, *QBBC*, is proposed by a new formulation that combines the advantages of low-variance biclustering techniques and Formal Concept Analysis. We prove that statistical dispersion measures that are order-preserving induce an ordering on the set of biclusters in the data. In turn, this ordering is exploited to form query-based biclusters in an efficient manner. Our novel approach provides a mechanism to generalize query-based biclustering to sparse high-dimensional data such as information networks and bag of words. Moreover, the proposed framework performs a local approach to query-based biclustering as opposed to the global approaches that previous algorithms have employed. Experimental results indicate that this local approach often produces higher quality and precise biclusters compared to the state-of-the-art query-based methods. In addition, our results on the performance evaluation illustrate the efficiency and scalability of QBBC compared to full biclustering approaches and other existing query-based approaches.

Keywords: Biclustering; Formal Concept Analysis.

1 Introduction

The abundance of high-dimensional data in applications ranging from text mining to bioinformatics prompted the development of biclustering, co-clustering and subspace clustering algorithms [5, 13]. All of these approaches attempt to identify clusters of objects in conjunction with the subset of features in high-dimensional datasets to avoid the curse of dimensionality. Although these methods have proven to be useful tools in exploratory analysis, most of these methods do not answer specific questions of interest and fail to incorporate prior knowledge and expertise from the user. For example, biologists often know that specific sets of genes are related to shared biological functions or pathways. Based on this prior knowledge, experts may want to enlist additional genes involved in that function in a microarray data [7]. In the domain of information networks, consider a network

linking key terms with research papers; a database researcher may wish to uncover which authors have been performing research in the field of ‘collaborative filtering’. Seeding an ideal query based biclustering algorithm with the seed terms *collaborative* and *filtering* would not only unveil the authors who have addressed this topic but also additional key terms that are related to collaborative filtering.

Query-based biclustering algorithms have been originally developed in the bioinformatics community [11, 22, 16, 4, 24, 7] specifically targeting microarray data. These algorithms utilize a set of seed genes provided by the user with the assumption that these seeds are tightly co-expressed or functionally related. In turn, the seed set is employed to prune the search space and guide the biclustering algorithm. Query-based biclustering algorithms characteristically attempt to keep biclusters centered around the seed set. However, they must also be robust and recognize an incoherent or partially incoherent seed set. Bio-inspired algorithms tend to adhere to these requirements, but are highly expensive and do not generalize well to large-scale datasets such as information networks, social networks and bag of words. In general, matrix representations of these data sources tend to be orders of magnitude larger than microarray data and much sparser. Furthermore, the state-of-the-art methods utilize probabilistic relation models [24] or Bayesian methods [7] which naturally allow for expert input into the algorithm through the use of prior distributions. However, generally users do not have this level of expertise and can only provide an intuitive seed set.

A novel formulation of query-based biclustering is proposed, in this paper, to generalize previous approaches to sparse and very high-dimensional data. Combining low-variance biclustering techniques and Formal Concept Analysis (FCA), the QBBC algorithm is developed. We prove that statistical dispersion measures that are order-preserving induce an ordering on the set of biclusters; consequently, this ordering is exploited to mine query-based biclusters in an efficient manner. Additionally, we capitalize on this ordering to identify neighboring biclusters that admit minimal noise when joining the clusters. In this manner, biclusters may be combined to enhance query results while still centering on the seed set.

1.1 Contributions The main contributions of our work are summarized as follows:

*Johns Hopkins University, Baltimore, MD. Email: faris.alqadah@jhu.edu

[†]Johns Hopkins University, Baltimore, MD. Email: joel.bader@jhu.edu

[‡]Wayne State University, Detroit, MI. Email: rajulanand@wayne.edu

[§]Wayne State University, Detroit, MI. Email: reddy@cs.wayne.edu

- Novel formulation of query-based biclustering (and biclustering in general) through a combination of low-variance biclustering and Formal Concept Analysis. In this formulation, we prove that order-preserving statistical measures of dispersion can induce an ordering that permits efficient mining.
- Development of the QBBC algorithm to efficiently mine query-based biclusters and approximate their ordering relation. The QBBC algorithm extends the seminal CHARM [23] algorithm (traditionally utilized to mine closed itemsets) by making use of an original operator termed “range intersection”.
- Formulation of two data-driven evaluation measures that capture the notion of coherence in biclusters.
- Experimental study with six real-world datasets from a wide range of real-world applications and performance comparison of the results with existing state-of-the-art approaches.

Following a review of existing works in Section 2, the theoretical clustering formulation is described in Section 3. Section 4 describes the QBBC algorithm, while Section 5 presents the results of our performance experiments. Finally, Section 6 offers concluding remarks, shortcomings, and avenues for future work.

2 Related Work

Existing works that are most relevant to the proposed approach fall into three categories: (i) bio-inspired query-based biclustering, (ii) semi-supervised clustering, and (iii) pattern-based biclustering. In this section, prior works in each of these categories are succinctly summarized.

As mentioned in the introduction, query-based biclustering methods have been developed in the bioinformatics community. The two works that are most relevant to the proposed one are PROBIC [24] and QDB [7] algorithms. QDB encompasses a Bayesian framework in which a conditional maximization is utilized for model estimation. Intuitively, biclusters are defined as sub-matrices of the original matrix whose expression values are modeled by a bicluster distribution as opposed to a “background” distribution (the rest of the data). Domain knowledge is encoded in the form of prior probability distributions. Finally, a resolution sweep method determines the ideal resolution that biclusters should be displayed at. In our method, we believe the issue of resolution sweep is naturally addressed by making use of a bicluster ordering. The PROBIC method, which is a follow up to QDB, is conceptually similar but adapts a probabilistic relation model as an extension to the Bayesian framework. Hard assignment of biclusters is assigned with the Expectation Maximization (EM) algorithm used to learn the

model. Another method in this category is an earlier approach, namely, the Iterative Signature Algorithm (ISA) [4] which utilizes the mean expression profile of the seed set to initialize the biclustering. Biclusters are then defined as fixed points with significant over or under expression. ISA does not deal with missing values, making it highly unlikely to be effective with sparse data. In addition, the algorithm is not purely query-based; there is no guarantee that a bicluster does not completely drift away from the original seed set. GeneRecommender [16] primarily focuses on prioritizing genes, and hence requires additional post-processing steps to convert to a biclustering approach. QDB was shown to outperform both GeneRecommender and ISA on synthetic data, while producing biologically relevant results in real data. More recently, PROBIC was shown to be more effective compared to ISA and QDB.

Semi-supervised clustering has mainly been characterized by constraint-based mono-dimensional clustering [3] which is a well-investigated research topic. In these works, pairwise constraints on objects such as “cannot-link” and “must-link” are utilized to integrate domain knowledge and thus improve the quality of the clustering solutions. Few methods have extended such constraint-based formulations to even biclustering settings [18, 20, 15]. However, this problem is substantially different from the problem that we are addressing in this paper. *The query seed set provided as input does not impose an explicit constraint, rather, it represents a user preference which may in fact be ignored by the algorithm if such a set is determined to be incoherent.*

Recently, some preliminary efforts have been made to extend closed pattern and association rule analysis to biclustering and multi-way clustering of real-valued data [17, 12]. The advantage of these methods is the ability to exhaustively search the set of biclusters and locate smaller and finer grain clusters often missed or masked by other methods. The primary disadvantage of pattern based biclustering algorithms is their potentially prohibitive computational cost. To address this issue, both [17] and [12] apply order preserving dispersion measures allowing efficient pruning of the search space. In this work, we build upon and extend the theoretical foundation of pattern based biclustering. The notion of an ordering preserving statistic is explicitly introduced and is proven that in addition to permitting effective pruning, such statistics impose an ordering upon the set of biclusters. As real-world data is typically dominated by a small number of very strong biclusters, this ordering is critical in facilitating the identification of neighboring biclusters to a seed set that enhance the final result. Finally, [17] and [12] measure coherence of biclusters through comparison of the range to a constant user-selected threshold. In this work, we introduce two original data-dependent measures for the evaluation of coherence.

3 Clustering Formulation

A context $\mathbb{K} = (\mathbf{G}, \mathbf{M}, \mathbf{K})$ is a triple where \mathbf{G} and \mathbf{M} are sets of objects and \mathbf{K} is a $|\mathbf{G}|$ by $|\mathbf{M}|$ matrix relating the objects of \mathbf{G} and \mathbf{M} . We assume that if an object $g \in \mathbf{G}$ is not related to object $m \in \mathbf{M}$ then $\mathbf{K}[g, m] = -\infty$. \mathbf{K} may also be thought of as the adjacency matrix of a bipartite graph with vertex sets \mathbf{G} and \mathbf{M} and edge set $\{(g, m) | \mathbf{K}[g, m] \neq -\infty\}$ and edge weighting function $w(g, m) = \mathbf{K}[g, m]$. $\Gamma(g)$ denotes the set of adjacent vertices to g (dually $\Gamma(m)$). An **object-set** G or M is a subset of objects from \mathbf{G} or \mathbf{M} . A **subspace** is any pair of object-sets (G, M) which also maybe thought of as a submatrix $\mathbf{K}[G, M]$. A **query set** is any object-set M^q (dually G^q) that is input by a user querying \mathbb{K} . Given M^q , our goal is to identify a subspace (G, M) , where $M \supseteq M^q$, that exhibit consistent values across the rows (columns) of $\mathbf{K}[G, M]$. In the terminology of biclustering [13], the desired result is to produce constant value biclusters in terms of rows or columns with the given constraint of a user query set. In order to quantify consistency of values in a subspace, statistical measures of dispersion such as standard deviation, inter-quantile range, range, and mean difference are utilized. For a given query set M^q , the **supporting set** of M^q are those objects in \mathbf{G} that are jointly adjacent to M^q and exhibit consistent values. Formally, define d as a dispersion measure that maps a subspace $(G, M) \mapsto \mathbb{R}$. Moreover, define a **consistency function**, f , that serves as a standard on what constitutes a consistent subspace; $f : (\mathbb{K}, g, M, \alpha) \mapsto \mathbb{R}$, where α is a user-selected parameter. Figure 2 displays several dispersion and consistency functions.

DEFINITION 1. Given $\mathbb{K} = (\mathbf{G}, \mathbf{M}, \mathbf{K})$, query set M^q , and user selected parameter α , the **supporting set** of M^q , denoted as $\psi_f^\alpha(M^q)$, is defined as

$$\{g \in \mathbf{G} \mid \Gamma(g) \subseteq M^q \wedge d(\mathbf{K}[g, M^q]) \leq f(\mathbb{K}, g, M, \alpha)\}$$

where f is a **consistency function** and d is a statistical measure of dispersion.

A general definition for a constant valued bicluster follows naturally from the supporting set formulation.

DEFINITION 2. Given $\mathbb{K} = (\mathbf{G}, \mathbf{M}, \mathbf{K})$, consistency function f , dispersion measure d , and parameter α , a **bicluster** or **α -concept** of \mathbb{K} is a subspace (G, M) such that

1. $\psi_f^\alpha(M) = G$
2. There does not exist $m \in \mathbf{M} \setminus M$ such that $\psi_f^\alpha(M \cup m) = G$.

The second condition of Definition 2 is referred to as the **closure** condition, which ensures that the maximum number of rows (columns) have been included in the bicluster without violating the consistency conditions.

DEFINITION 3. Given $\mathbb{K} = (\mathbf{G}, \mathbf{M}, \mathbf{K})$, consistency function f , dispersion measure d , and query set $Q \subset \mathbf{M}$ (dually $Q \subset \mathbf{G}$), then an **ideal query-based bicluster** of \mathbb{K} is an α -concept (G, Q') , where $Q' \supseteq Q$.

In real datasets, we expect the structure and nature of query-based biclusters to be highly sensitive to the choices of α and f . Non-stringent measures cause biclustering algorithms to mask or miss small but relevant biclusters [17]; on the other hand, too stringent parameter settings may cause the algorithm to conclude that a query set is incoherent and thus will not return any bicluster. As a result, given a query we do not expect to locate an ideal query-based bicluster due to the resolution problem. We propose a computationally efficient scheme to account for the parameter selection problem that does include varying parameter settings. We advocate setting stringent parameter settings and utilizing several small, localized and similar α -concepts to construct larger query-based biclusters that still center around the query. We show in the sequel that utilizing order preserving dispersion measures induce an ordering on the set of α -concepts in the data. In turn, this ordering is exploited to identify α -concept neighborhoods that consist of similar α -concepts centered around the query set.

3.1 Formal Concept Analysis As applied to binary relations, Formal Concept Analysis (FCA) stipulates that biclusters in binary valued data (maximal bi-cliques) are ordered by the **hierarchical** order [10]. In this section, we prove that this ordering also applies to α -concepts given that the dispersion measures and consistency functions adhere to certain ordering properties.

DEFINITION 4. Given α -concepts (G_1, M_1) and (G_2, M_2) , then $(G_1, M_1) \leq (G_2, M_2)$ if and only if $G_1 \subseteq G_2$ and $M_1 \supseteq M_2$. This ordering relation is referred to as the **hierarchical** ordering.

Undoubtedly, the selected dispersion measure and consistency function determine if the set of α -concepts are in fact ordered by the hierarchical ordering.

DEFINITION 5. d is order-preserving if $d(\mathbf{K}[g, M]) \leq d(\mathbf{K}[g, M \cup m])$, where $m \in \mathbf{M} \setminus M$

DEFINITION 6. f is anti-monotone if $f(\mathbb{K}, g, M, \alpha) \geq f(\mathbb{K}, g, M \cup m, \alpha)$, where $m \in \mathbf{M} \setminus M$

Clearly, if d and f are order-preserving and anti-monotone then $\psi_f^\alpha(M^q)$ is also anti-monotone; this in turn implies that α -concepts defined in terms of an order-preserving dispersion statistics and anti-monotone consistency functions are ordered by the hierarchical order.

THEOREM 3.1. Given a context \mathbb{K} , α , d , and f , then if d and f are order-preserving and anti-monotone respectively the α -concepts of \mathbb{K} are ordered by the hierarchical order.

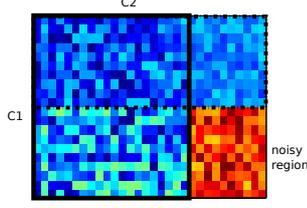


Figure 1: Noisy region induced by combining two neighboring concepts

Proof. See Appendix.

FCA stipulates that concepts ordered by the hierarchical order form a complete lattice; therefore we conclude that the set of α -concepts also form a complete lattice. The α -concept lattice forms the basis for defining α -concept neighborhoods.

3.1.1 Neighboring Concepts Consider concepts (G_1, M_1) and (G_2, M_2) such that $(G_1, M_1) \geq (G_2, M_2)$. If there is no concept (G_3, M_3) fulfilling $(G_1, M_1) \geq (G_3, M_3) \geq (G_2, M_2)$ then (G_1, M_1) is an **upper neighbor** of (G_2, M_2) ; dually (G_2, M_2) is a **lower neighbor** of (G_1, M_1) . For any concept C_1 , its set of upper neighbors is denoted as $\Upsilon(C_1)$. Intuitively, neighboring concepts are assumed to be similar. This intuition has been formalized and exploited to extract knowledge in binary contexts [1, 14]. In terms of query-based biclustering, neighboring α -concepts provide a mechanism to explore and combine closely related α -concepts in order to enhance or broaden a seed bicluster. We formally show that combining α -concept $C_1 = (G_1, M_1)$ and any upper (lower) neighbor $C_2 = (G_2, M_2)$ will result in the minimal degree of inconsistency among all possible pairings of larger α -concepts. Figure 1 depicts the result of combining α -concept C_1 and an upper neighbor C_2 . By definition, the values encompassed by C_1 and C_2 are consistent, however, the values induced by the difference between C_1 and C_2 may not be (noisy region). One possible measure of dissimilarity between two α -concepts is the degree of inconsistency introduced by the noisy region when combining the two α -concepts. Given below, the *dist* score assesses the dissimilarity by computing the ratio of the consistency function as measured in the original concept to the noisy region.

$$(3.1) \quad \text{dist}((G_1, M_1), (G_2, M_2)) = \frac{1}{|G_1 \setminus G_2|} \sum_{g \in G_1 \setminus G_2} \frac{d(\mathbf{K}[g, M_2 \setminus M_1])}{d(\mathbf{K}[g, M_2])} \times s_g$$

where

$$s_g = 1 + |(\Gamma(g) \cap M_2) \setminus M_1|$$

Name	Computation	Order Preserving?
Range	$r(\mathbf{K}[m, G]) = \max \mathbf{K}[g, M] - \min \mathbf{K}[g, M]$	yes
Standard Deviation	$\sigma(\mathbf{K}[m, G])$	no
Inter-quantile Range	$Q_3(\mathbf{K}[m, G]) - Q_1(\mathbf{K}[m, G])$	no
Mean difference	$\frac{1}{ M (M -1)} \sum_{i=1}^{ M } \sum_{j=1}^{ M } \mathbf{K}[g, m_i] - \mathbf{K}[g, m_j] $	no
Coefficient of variation	$\frac{\sigma(\mathbf{K}[g, M])}{\bar{\mathbf{K}}(\mathbf{K}[g, M])}$	no
Quartile coefficient	$\frac{Q_3(\mathbf{K}[g, M]) - Q_1(\mathbf{K}[g, M])}{Q_3(\mathbf{K}[g, M]) + Q_1(\mathbf{K}[g, M])}$	no

(a) Dispersion Functions

Name	Computation	Anti-monotone ?
Constant threshold	$f(\mathbb{K}, g, M, \alpha) = c$	yes
Min range	$f(\mathbb{K}, g, M, \alpha) = \min \mathbf{K}[g, M]$	yes

(b) Consistency Functions

Figure 2: Dispersion and consistency functions

In Equation (3.1), the fact that d is order preserving bounds the left hand side ratio inside the summation to 1. The right hand side term, s_g , is introduced to account for sparse data. In the case that \mathbf{K} is full (\mathbb{K} is a complete bipartite graph) then s_g evaluates to 1 and the ratios are simply summed up. On the other hand, if the row $\mathbf{K}[g, M_2 \setminus M_1]$ contains missing values, then a penalty term is proportionally imposed on the ratio. Finally, the consistency ratios are averaged over the total number of rows in the join of the concepts.

THEOREM 3.2. For any α -concept (G_1, M_1) , then

$$\operatorname{argmax}_{(G_2, M_2) \geq (G_1, M_1)} \text{dist}((G_1, M_1), (G_2, M_2)) \in \Upsilon(G_1, M_1)$$

Proof. See appendix.

Theorem 3.2 provides a theoretical basis for combining neighboring α -concepts to form a query-centered bicluster.

3.2 Dispersion and Consistency Functions Figure 2 shows several standard statistical measures of dispersion along with some select consistency functions that have been utilized recently in biclustering algorithms [17, 12]. As can be seen, the only dispersion measure that is order-preserving is range, which justifies its use previously. Unfortunately, range is not a robust statistic; nevertheless, its use as a dispersion measure is justified because it is order-preserving and it bounds the standard deviation as follows:

$$(3.2) \quad 2 \times \sigma(\mathbf{K}[g, M]) \leq r(\mathbf{K}[m, G])$$

Furthermore, relatively few consistency functions exist; previous methodologies have simply imposed hard thresholds. In the sequel, two novel data-dependent, anti-monotone consistency functions are developed.

3.2.1 α -sigma Consistency Function Given $\mathbb{K} = (\mathbf{G}, \mathbf{M}, \mathbf{K})$ assume that the rows and columns of \mathbf{K} are i.i.d and normally distributed. That is, for all $g \in \mathbf{G}$, $\mathbf{K}[g, \mathbf{M}]$ is normally distributed with $\sigma(\mathbf{K}[g, \mathbf{M}])$ and

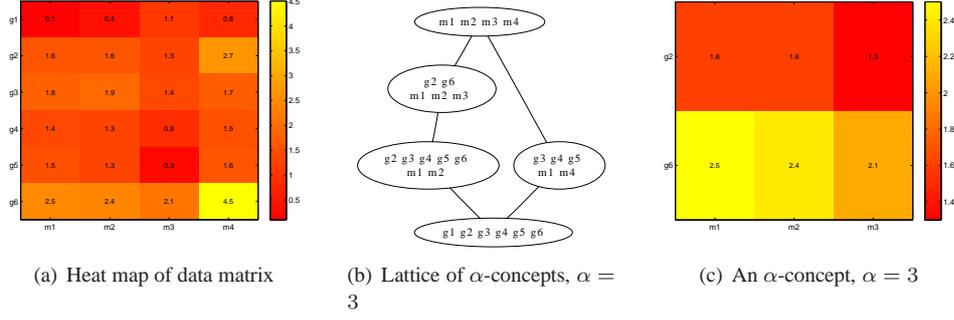


Figure 3: Concepts of a sample data matrix

$\mu(\mathbf{K}[g, \mathbf{M}])$ respectively. The 3-sigma rule states that about 68.26 % of the values for each $\mathbf{K}[g, \mathbf{M}]$ lie within a single standard deviation of the mean while 95.44 % of the values lie within two standard deviations from the mean. Hence, randomly selecting query object-set M , the majority of values in the subspace $\mathbf{K}[g, M]$ are also expected to lie within the $\mu(\mathbf{K}[g, \mathbf{M}]) \pm 2\sigma(\mathbf{K}[g, \mathbf{M}])$ with mean $\mu(\mathbf{K}[g, \mathbf{M}])$ and standard deviation $\sigma(\mathbf{K}[g, \mathbf{M}])$. A subspace $\mathbf{K}[g, M]$ where $\sigma(\mathbf{K}[g, M]) < \sigma(\mathbf{K}[g, \mathbf{M}])$ indicates that the values in the subspace are more consistent than expected. This can be generalized to subspaces where $\alpha \times \sigma(\mathbf{K}[g, M]) < \sigma(\mathbf{K}[g, \mathbf{M}])$, where $\alpha \geq 1$; then the α -sigma consistency function is formulated as

$$(3.3) \quad \alpha \times \sigma(\mathbf{K}[g, M]) < \sigma(\mathbf{K}[g, \mathbf{M}])$$

$$\alpha \frac{r_g}{2} < \sigma(\mathbf{K}[g, \mathbf{M}])$$

$$(3.4) \quad r_g < \frac{2\sigma(\mathbf{K}[g, \mathbf{M}])}{\alpha}$$

where r_g denotes $r(\mathbf{K}[g, M])$ and equation (3.3) follows from the stringent assumption imposed by inequality (3.2).

DEFINITION 7. For context $\mathbb{K} = (\mathbf{G}, \mathbf{M}, \mathbf{K})$ and given subspace $\mathbf{K}[g, M]$, the α -sigma consistency function is defined as

$$(3.5) \quad f(\mathbb{K}, g, M, \alpha) = \frac{2\sigma(\mathbf{K}[g, \mathbf{M}])}{\alpha}$$

The α -sigma consistency function is clearly anti-monotone.

3.2.2 Maximum Spacing Uniform Estimator A less stringent consistency function may seek subspaces whose range is α times smaller than the range of a uniformly distributed random variable. Assume that the rows and columns of \mathbf{K} are i.i.d and uniformly distributed with end points $\max(\mathbf{K}[g, \mathbf{M}])$ and $\min(\mathbf{K}[g, \mathbf{M}])$. Let $\{x_1, \dots, x_n\}$ be an ordered sample from uniform distribution $U(a, b)$ with unknown endpoints a and b . A known uniformly minimum

variance unbiased estimator of the end points is the maximum spacing uniform estimator given as follows:

$$(3.6) \quad \hat{a} = \frac{nx_1 - x_n}{n-1}$$

$$(3.7) \quad \hat{b} = \frac{nx_n - x_1}{n-1}$$

Clearly, the end points of the distribution can be estimated utilizing only the end points of the sample; hence, the range of the distribution can also be estimated as $\hat{b} - \hat{a}$. The actual range can then be compared to an estimated range as follows:

$$\hat{b} - \hat{a} < \alpha(b - a)$$

$$\frac{nx_n - x_1}{n-1} - \frac{nx_1 - x_n}{n-1} < \alpha(b - a)$$

$$x_n - x_1 + n(x_n - x_1) < \alpha(n-1)(b - a)$$

$$x_n - x_1 < \alpha(n-1)(b - a) - n(x_n - x_1)$$

By the above reasoning, a subspace $\mathbf{K}[g, M]$ is considered consistent if its maximum space estimated range is α times smaller than the range of the distribution.

DEFINITION 8. For context $\mathbb{K} = (\mathbf{G}, \mathbf{M}, \mathbf{K})$ and given subspace $\mathbf{K}[g, M]$, the **maximum spacing uniform estimator consistency function** is

$$(3.8) \quad f(\mathbb{K}, g, M, \alpha) = \alpha(|\mathbf{M}| - 1)r(\mathbf{K}[g, \mathbf{M}]) - |\mathbf{M}|r(\mathbf{K}[g, M])$$

Clearly the maximum spacing uniform estimator consistency function is anti-monotone.

EXAMPLE 1. Consider the data matrix in Figure 3(a). The α -concept lattice of the data is depicted in Figure 3(b) utilizing the α -sigma consistency function with $\alpha = 3$. Figure 3(c) displays a heat map of the concept $(\{g_2, g_6\}, \{m_1, m_2, m_3\})$.

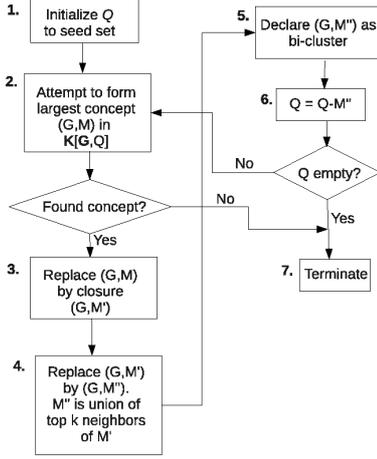


Figure 4: Overview of QBBC algorithm

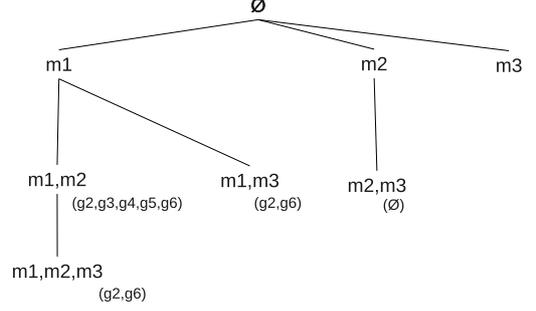


Figure 5: Prefix based search for concepts

4 QBBC Algorithm

4.1 Overview Figure 4 describes the main steps of the QBBC algorithm. Initially, an ideal query-based bicluster is sought for the given query set $Q \subset M$ (steps 2-3) by enumerating all α -concepts in the sub-matrix $K[G, Q]$. Exact details of the search and enumeration procedure are presented in Section 4.2. If an ideal bicluster is located, then the closure of that bicluster is computed in the full matrix and the algorithm is terminated. In the more likely scenario that an ideal bicluster is not found, then QBBC retains the α -concept, (G, M) , with most query objects. Next, QBBC assumes that Q is a coherent query set, but an ideal bicluster was not located due to the resolution of the consistency function. Theorem 3.2 stipulates that if the query set Q was indeed coherent, then the objects in the query set not found in (G, M) would appear in α -concepts in the neighborhood of (G, M) . Hence, guided by Theorem 3.2, the top k upper neighbors, ranked in ascending order by $dist$, are augmented to (G, M) to form the bicluster (G, M'') . At this point, any objects in Q but not in M'' are considered incoherent with the current bicluster, hence, a new query set is formed as $Q \setminus M''$ (step 6). The entire procedure is repeated with the new query set until Q is empty or no α -concepts are enumerated in step 2. In the sequel, we describe the algorithmic and implementation details for completing steps 2, 3 and 4 by taking advantage of the properties of order-preserving dispersion measures and building upon the seminal CHARM algorithm for closed itemset mining.

4.2 Identifying α -concepts Step 2 of QBBC calls for enumerating α -concepts in the sub-matrix $K[G, Q]$. The search space for this task is formulated as a prefix tree as shown in Figure 5. Recalling the dataset presented in example 1, Figure 5 depicts the search tree for query set $Q = \{m1, m2, m3\}$. The tree utilizes the idea of prefix-based

equivalence classes in order to break up the search tree into independent sub-problems. Viewing any objectset $P \subseteq Q$ as a string, two object-sets are in the same prefix-class if they share a common-length prefix. Each node of the tree, P , represents an object-set $P \subseteq Q$ and associated supporting set of P . Any node of the search tree with a non-empty supporting set is either an α -concept or a non-closed α -concept. For a given node P , the next level of the search is generated by computing $\psi_f^\alpha(P \cup P_r)$ for nodes P_r located to the right of P under the same branch. For example, in Figure 5 consider node $\{m1, m2\}$. Only node $\{m1, m3\}$ is located to the right of $\{m1, m2\}$ under the same branch; hence, the next level of the search tree consists of node $\{m1, m2, m3\}$ and supporting set $\psi_f^\alpha(\{m1, m2, m3\})$. Moreover, the order-preserving and anti-monotone properties of d and f ensure that the search tree may be pruned when an empty supporting set is encountered. Computing $\psi_f^\alpha(P \cup P_r)$ is a matter of identifying the supporting objects common to both P and P_r with the added restriction that the common supporting objects also contain consistent values; hence, we introduce a new operator: *range intersection*. Given two object sets P and P_r , coupled with supporting sets S and S_r define range intersection as:

$$\begin{aligned} \psi_f^\alpha(P \cup P_r) &= S \cap S_r \\ &= \{s \in S \cap S_r \mid d(K[s, P \cup P_r]) \leq f(K[s, P \cup P_r], \alpha)\} \end{aligned}$$

where d is the range dispersion function. A naive strategy to identify α -concepts is to simply enumerate the entire prefix-based search tree. On the other hand, several closed itemset enumeration algorithms follow this strategy with additional pruning steps that vastly shrink the search space. We exploit the main theorem guiding the seminal closed item set enumeration algorithm, CHARM, which offers further opportunity to prune the search tree by investigating the result of $\psi_f^\alpha(P \cup P_r)$ when generating the next level of the search tree. Given that f is defined to be a constant, the α -

sigma function, or the maximum spacing uniform estimator and the range dispersion function is utilized, we prove that the CHARM theorem is in fact generalizable to α -concepts.

THEOREM 4.1. *Given the range dispersion function d and consistency function f that is either constant, the α -sigma function, or the maximum spacing uniform estimator then the following holds: Let PP_l and PP_r be two nodes under the same branch in the search tree with PP_r located to the right of PP_l . Moreover let S_l and S_r be the supporting sets of PP_l and PP_r respectively. Given that*

1. For any prefix node PP_x under the same branch
 $\forall s \in S_x \quad d(\mathbf{K}[s, PP_x]) \leq \frac{f(\mathbf{K}, s, PP_x, \alpha)}{2}$.
2. $SS = S_l \cap S_r \neq \emptyset$

then the following properties hold:

1. If $|SS| = |S_l| \wedge |SS| = |S_r|$ then every occurrence of PP_l maybe replaced by $PP_l \cup PP_r$ and the node PP_r and all its children may be pruned from the search tree.
2. If $|SS| = |S_r|$ then every occurrence of PP_l may be replaced with $PP_l \cup PP_r$.
3. If $|SS| = |S_l|$ then every occurrence of PP_r and its children maybe pruned from search tree, but a new child node $PP_l \cup PP_r$ must be formed with the supporting set $\psi_f^\alpha(PP_l \cup PP_r)$.
4. If $|SS| \neq |S_l| \wedge |SS| \neq |S_r|$ then no condensation of the search tree is possible and new child node $PP_l \cup PP_r$ with supporting set $\psi_f^\alpha(PP_l \cup PP_r)$ must be formed.

Proof. See Appendix.

The actualization of the above theorem yields three important outcomes: 1) computing the closure of PP_l , 2) several possible condensations of the search tree, and 3) enumerating the children of PP_l .

EXAMPLE 2. *Consider the branch under node $\{m1\}$ with nodes $P_l = \{m1, m2\}$, $P_r = \{m1, m3\}$ along with supporting sets $S_l = \{g2, g3, g4, g5, g6\}$ and $S_r = \{g2, g6\}$ as depicted in Figure 5. Condition 1 of Theorem 1 is satisfied due to the fact that no other prefix nodes exist under this branch. Let f be set to the α -sigma consistency function with $\alpha = 3$ then $SS = S_l \cap S_r = \{g2, g6\}$, hence condition 2 is also satisfied. In this case $|SS| = |S_r|$, hence case 2 of the theorem is applied and P_r is replaced by $\{m1, m2, m3\}$. In effect, the closure of the subspace (S_r, P_r) is computed and the need to form a new child in the search tree is eliminated resulting in computational savings.*

Input: Prefix node PP_l , supporting set S_l
Result: Update PP_l by its closure
Result: Prune search space
Result: Compute children of PP_l at this branch

```

1 begin
2   C ← ∅ ;
   // children
3   B ← {(PP_r^1, S_r^1), ..., (PP_r^n, S_r^n)} ;
   // nodes to the right
4   flg ← ∨(PP_r^i, S_r^i) ∀s ∈ S_r^i d(K[s, PP_r^i]) ≤
      f(K, s, PP_r^i, α) ;
5   for (PP_r, S_r) ∈ B do
6     SS_r ← S_r ∩ S ;
7     if |SS_r| = |S_r| ∧ |SS| = |S| then
8       P ← P ∪ P_r ;
9       if flg then
10        Remove (PP_r, S_r) from tree ;
11    else if |SS_r| = |S_r| then
12      P ← PP_l ∪ PP_r ;
13    else if |SS_r| = |S| then
14      C ← C ∪ (PP_l ∪ PP_r, SS_r) ;
15      if flg then
16        Remove (PP_r, SS_r) from tree ;
17    else
18      C ← C ∪ (PP_l ∪ PP_r, SS_r) ;

```

Algorithm 1: A single search step corresponding to a single branch

As pointed out in [23] checking for each of the four cases during is a constant computational cost due to the fact that set intersection must be computed to generate the next level of the search in any case. Computing range intersection of supporting sets S_l and S_r is performed in $O(|S|)$ time as it entails set intersection and a constant time operation to compute range if the indices of maximum and minimum elements are maintained throughout the search. Augmenting Theorem 4.1 with the range intersection operation, a procedure to generate children prefix nodes and prune the tree at any branch of the search tree is given in Algorithm 1. The procedure is an exact implementation of Theorem 4.1.

4.2.1 Closure and Upper Neighbors Computing the closure of a subspace (G, M) (step 3 of QBBC) is accomplished by applying a single step of CHARM where M is the prefix node and objects $m_r \in \mathbf{M} \setminus M$ constitute nodes to the right of M . A similar strategy is utilized to determine the upper neighbors of (G, M') (step 4); in this case two rounds of CHARM are applied. Unfortunately, applying CHARM in this manner does not guarantee the exact set of upper neighbors

Name	Domain	Size	Density	Num.classes
<i>mer</i>	Bag of words	1,990 x 21,258	0.003	2
<i>allpc</i>	Bag of words	4,966 x 26,323	0.001	5
<i>allsci</i>	Bag of words	3,975 x 30,440	0.001	4
<i>papers</i>	Information network	28,564 x 16,891	0.0003	4
<i>emap</i>	Microarray	3,300 x 3,300	0.113	na
<i>hughes</i>	Microarray	4,684 x 300	1.0	na

Figure 6: Six real-world datasets used in our experiments.

in the α -concept lattice; rather, a superset of the upper neighbors may in fact be enumerated [23]. On the other hand, Theorem 3.2 guarantees that the top ranked concepts in the generated set are upper neighbors of (G, M') .

5 Experimental Results

5.1 Datasets and evaluation criteria Six real-world datasets were used in our experimental study and are listed in Figure 6. The first three datasets came from the large 20Newsgroups dataset [2], *papers* [21] is a subset of the DBLP database linking authors with paper titles, and *emap* [19] and *hughes* [9] are both microarray datasets. *mer* represents the news feeds from the Middle East politics and Religion forum, *allpc* is a combination of all news feeds pertaining to computers, and *allsci* is an aggregation of documents associated with science. In all text-based datasets, stop words were removed and TF-IDF weights were computed. For each dataset, two categories of seed sets were constructed. The first category of seed sets was manually assembled and generally contained 2-5 objects per query; these objects were determined to be coherent. For example, in *mer* query terms such as $\{israel, palestine\}$ formed a query while in *papers* the terms $\{query, optimization\}$ were utilized. For *emap* and *hughes*, query genes were assembled by consulting the Biological Process hierarchy of the Gene Ontology [8] and only those sets that were annotated by the same functional class were retained. The second category of seed sets consisted of randomly selecting objects from each dataset; typically 10-50 objects were selected per query. A total of 10 manually created and 50 random queries were generated for each dataset. Evaluation of our experimental results was based on three criterion:

1. Evaluation using mean square error to measure the cluster quality.
2. Average purity of biclusters using class labels available in most of the datasets.
3. Visual assessment of cluster quality.

For comparison, we used the R implementation of QDB, available at <http://homes.esat.kuleuven.be/~kmarchal>, with the default parameter settings. The QBBC algorithm was implemented in C++ with both the α -sigma (QBBC-alpha) and maximum spacing uniform estimator consistency functions (QBBC-max) and is available

at <http://faris-alqadah.herokuapp.com>. The α parameter was set to 3 on both QBBC-sigma and QBBC-max and both methods were set to augment initial query clusters with the top 20 neighbors in the α -concept lattice. At the time of this writing, no implementation was available for PROBIC. Although QDB is designed to handle datasets with missing values, when attempting to find query based biclusters in the sparse datasets of this study (all but *hughes*), no biclusters were ever produced. This may be due to the extreme sparsity levels affecting the probabilistic model that QDB is based upon. As a result, when executing QDB, the sparse datasets were filled in with randomly generated values or zeros. In effect, this created a background distribution from which QDB should discern actual biclusters.

5.2 Cluster quality The first evaluation criterion utilized was mean square error (MSE) as given by Cheng and Church [6]:

$$MSE(G, M) = \frac{1}{|G||M|} \sum_{g \in G, m \in M} (\mathbf{K}[g, m] - \mu_{g, M} - \mu_{G, m} + \mu_{G, M})^2$$

where $\mu_{g, M}$ is the mean of the g^{th} row under M columns, $\mu_{G, m}$ is the mean of the m^{th} column respectively under the G rows and $\mu_{G, M}$ is the overall mean of the subspace. Under this formulation the minimum value of MSE is 0 when all values in the subspace are equal. To accommodate sparse data, missing or non-edge values were either ignored or computed as usual with the missing values filled by the generated background distribution; the best results are reported here. Average purity of a set of biclusters $\mathcal{C} = \{(G_1, M_1), \dots, (G_n, M_n)\}$ and set of class labels $\mathcal{L} = \{L_1, \dots, L_m\}$ is defined as

$$\frac{1}{|\mathcal{C}|} \sum_k \max_j \frac{|M_k \cap L_j|}{|M_k|}$$

In other words, each cluster is assigned to the class which is most frequent in the cluster and the purity measure is computed as the precision of the cluster with respect to this class; the average of the purity scores is then computed.

Figures 7(a)-7(c) displays the distribution of MSE scores for biclusters mined with both manually created and randomly generated query sets. In sparse data (Figures 7(a)-7(c)) the QDB cluster MSE scores were strikingly larger than those of both QBBC methods. This trend is explained by the fact that QDB tended to produce very large clusters even when a small seed set was input. For example, a seed set in *mer* contained only two terms, yet QDB returned a cluster containing 10,256 words and 1,600 documents. Clearly, the extreme sparsity levels of these datasets confound the QDB algorithm and it is unable to discriminate between the

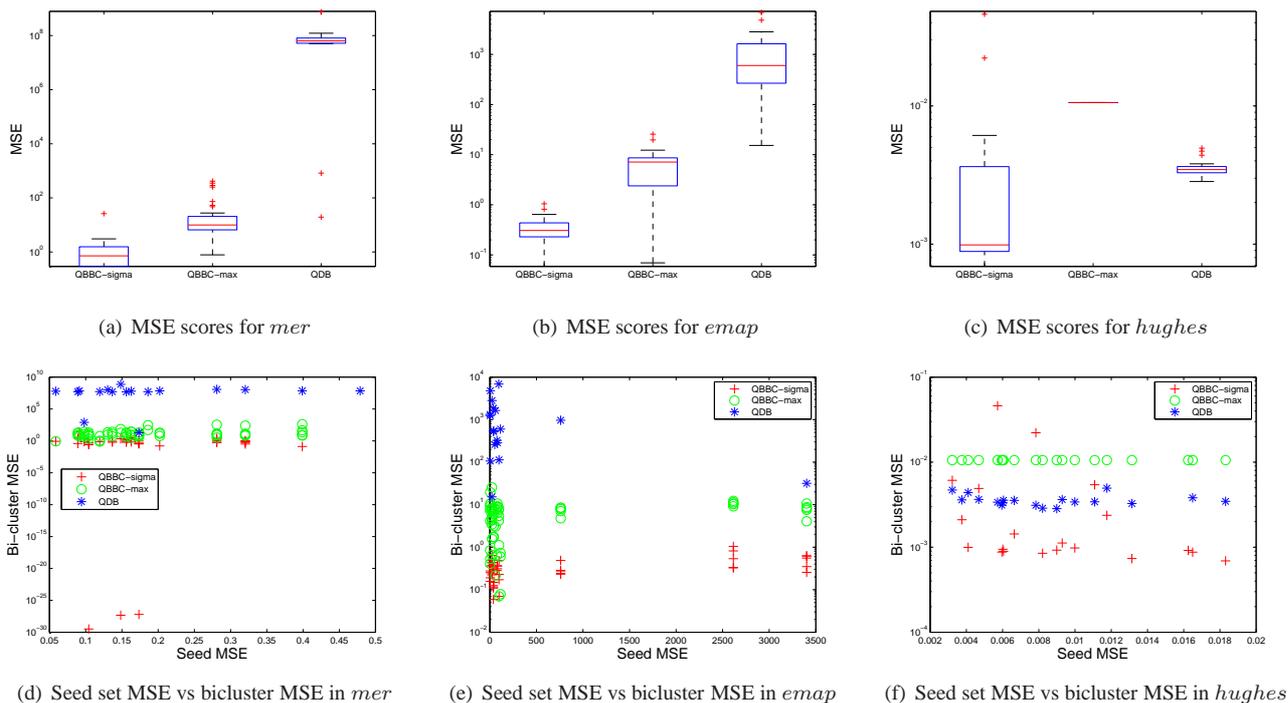


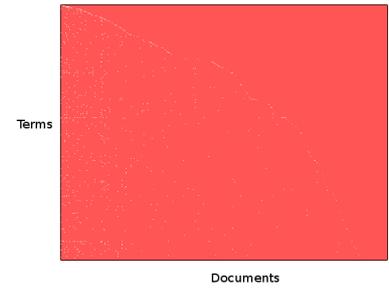
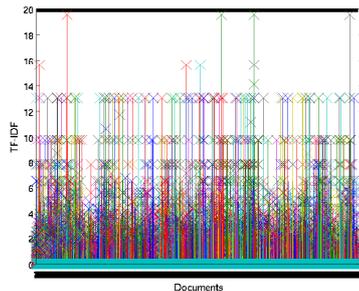
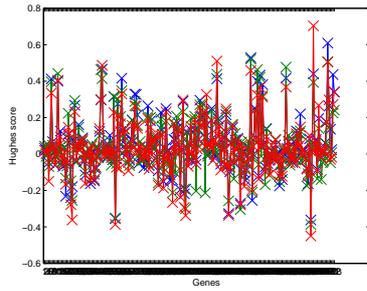
Figure 7: MSE distributions of query-based biclusters with QBBC and QDB

background distribution and high quality query centered biclusters. These facts illustrate the clear advantage of QBBC in sparse high-dimensional data; nonetheless, both methods have similar performance on standard non-sparse microarray data (Figure 7(c)). The lack of variance in MSE scores of clusters produced by QBBC-max in *hughes* was attributed to the uniform distribution assumption; all queries resulted in similar biclusters indicating that this consistency function is more suited to sparse data.

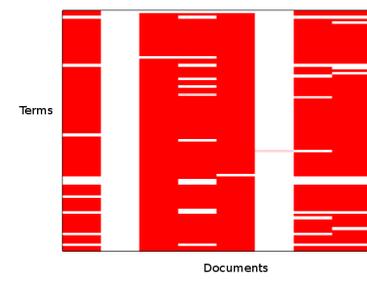
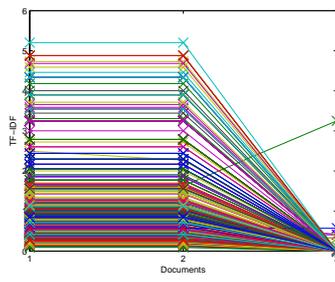
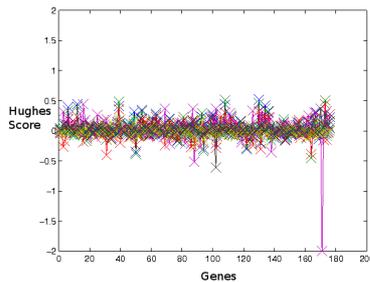
Figures 7(d)-7(f) display the relationship between MSE scores of the initial query sets and the final biclusters. The score of the query set was computed as the MSE of the subspace encompassing the query set in conjunction with all the rows (or columns) of the dataset. Random seeds that lead to no cluster being enumerated are not displayed. As can be seen, both QBBC algorithms' bicluster scores tend to be lower than the query set while QDB scores are higher. As expected, the difference in seed scores and bicluster scores are much closer for manually created query sets. Interestingly, for random query-sets, the QBBC scores tended to be orders of magnitude lower than those of the initial query set. In most cases, QBBC broke down the random query and produced several biclusters that were consistent in a small subspace of the initial query set, while QDB failed to do this. Once again, in all sparse datasets the performance of the QBBC algorithms is clearly superior to QDB while the results are comparable in *hughes* dataset.

Sample clusters mined by algorithms are depicted in Figure 8. Visually inspecting these Figures, biclusters mined by QDB and QBBC (Figures 8(a) and 8(d)) appear to be of similar quality. On the other hand, once again, biclusters produced by QDB in sparse data (Figures 8(b) and 8(c)) are clearly less informative compared to those mined by QBBC (Figures 8(e) and 8(f)). These images manifest the ability of QBBC to filter out the background distribution and zoom in on query-centered biclusters in both sparse and dense datasets. Figure 8(g) illustrates a sampling of manually created query sets and the resulting word clusters. In general, the resulting biclusters contained most of the original query set and greatly expanded the cluster to include mostly pertinent terms with a few noisy terms. The most coherent bicluster came from *papers*; an argument can be made that all the terms in this bicluster are relevant, while the entire seed set was preserved. This was expected as *papers* is extremely sparse and only contains paper titles. In the case of *allsci*, a much noisier dataset than *papers*, the terms *hacker* and *checksum* were dropped while several noisy terms were introduced.

Finally, the average purity of all biclusters resulting from manual query sets are presented in figure 8(h). Due to the inability of QDB to scale to the three larger datasets, *papers*, *allpc*, and *allsci*, the average purity scores were computed by repeating experiments on small subsets of these datasets with appropriate query terms and computing the av-



(a) QDB cluster in *hughes* with manually selected query (b) QDB cluster in *mer* with manually selected query (c) QDB cluster in *allpc* with manually selected query



(d) QBBC cluster in *hughes* with manually selected query (e) QBBC cluster in *mer* with manually selected query (f) QBBC cluster in *allpc* with manually selected query

Dataset	Seed set	bicluster
<i>papers</i>	query optimization	databases semantic models based analysis dynamic evaluation systems information algorithms query design xml database distributed processing multiple queries efficient relational optimization temporal
<i>allsci</i>	cryptography hacker checksum algorithm cipher	algorithm cipher cryptography states rsa signed plaintext freedom scheme text number exists men recover selected application create united versions comments archive included documentation approved attempt licensed internet broad claimed recommended newsgroups
<i>mer</i>	israel palestine	arab palestine israelis israel israeli opposition zionist ground zionism international peace problem state wrote meant states human feel necessity statements creation guess occupation statement forms disregard refers minister racism fully intervention

(g) Sample seed sets and corresponding bicluster

Dataset	Algorithm	Avg. Purity	Avg. num documents
<i>mer</i>	QBBC-sigma	0.89	28.125
	QBBC-max	0.89	38.75
	QDB	0.47	602.4
<i>allpc</i>	QBBC-sigma	0.60	51.4
	QBBC-max	0.54	66.2
	QDB	0.24	838
<i>allsci</i>	QBBC-sigma	0.75	19.4
	QBBC-max	0.75	26
	QDB	0.25	804.25
<i>papers</i>	QBBC-sigma	0.72	102.8
	QBBC-max	0.72	103.2
	QDB	0.31	1021.2

(h) Cluster purity

Figure 8: Sample clusters and precision-recall of clusters

erage. Each experiment was repeated ten times on the sampled dataset with all algorithms and the average purity scores being reported. QBBC-sigma and QBBC-max produced clusters with approximately the same purity levels. Due to the large size of QDB clusters, the purity scores are near random. On the other hand, the QBBC scores were not induced by extreme cluster sizes as is demonstrated by the average

number of documents in each cluster.

5.3 Performance Tests Performance tests were conducted to evaluate the practical running time and scalability of the QBBC. All experiments were conducted on a 3.33 GHz Intel I7 quad core CPU with 16 GB of RAM. Due to its implementation in R, QDB scaled extremely poorly in the

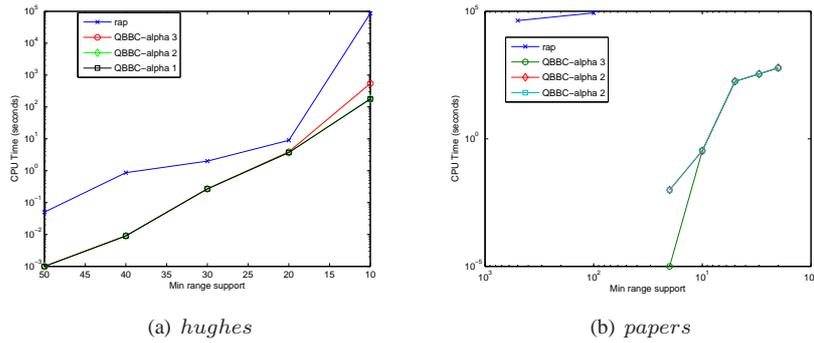


Figure 9: Performance comparisons between QBBC and RAP algorithms.

case of the three largest datasets and ran out of memory on a 16GB main memory machine in every instance. On the other hand, the memory requirements of QBBC never exceeded 10 MB. Moreover, running times on these datasets exceeded an hour even with query sets containing fewer than 5 objects. In order to conduct a more fair comparison, and evaluate the true effectiveness of the pruning rules introduced, we re-implemented QBBC specifying it to enumerate all α -concepts in an entire dataset; we compared the running times to the state-of-the-art pattern-based biclustering algorithm RAP on all datasets. An implementation of RAP in C++ was downloaded from <http://vk.cs.umn.edu/gaurav/rap/>. Due to space limitations, only the results of two of the performance tests are displayed in Figure 9. The min range support is a user defined parameter instructing the algorithms to only retain clusters with a minimum degree of supporting sets. Due to min range support also being anti-monotone it can also be utilized to prune the search space. As shown by these results, the additional pruning steps introduced by Theorem 4.1 clearly result in much improved performance. In non-sparse data (Figure 9(a)), we observed well over three orders of magnitude speed up at the lowest minimum support levels. On the largest sparse data set, *papers*, rap was unable to complete within 48 hours at the minimum support levels specified for QBBC (Figure 9(b)) while QBBC completed in under 700 seconds at the lowest support levels.

6 Conclusion

In this paper, a novel query-based biclustering algorithm, QBBC, was developed. It was shown that statistical dispersion measures that are order-preserving induce an ordering on the set of biclusters in the data; in turn this ordering is exploited to mine similar biclusters centered around a query seed set. Making use of an original operator, range intersection, it was further shown that the seminal CHARM algorithm for mining closed itemsets is generalizable to the computational framework for mining query-based biclustering. Experimental results unveiled that in high dimensional sparse

data, QBBC has a clear advantage over the current state-of-the-art query-based biclustering methods while similar performance was observed on standard microarray datasets. Moreover, we illustrated that the pruning measures introduced in the QBBC algorithm resulted in significant computational savings compared to both QDB and RAP. Shortcomings of QBBC include the inability to compute exact neighbors of a bicluster in the concept lattice. We believe that this fact leads to significantly higher computational cost especially in dense datasets. Moreover, no data-driven criterion was derived for determining if and when neighboring biclusters should be merged to enhance the seed bicluster. Finally, additional order-preserving dispersion methods and anti-monotone consistency functions should be developed in future in order to accommodate some specific application demands.

References

- [1] F. Alqadah and R. Bhatnagar. Discovering substantial distinctions among incremental bi-clusters. In *SDM'09*, 2009.
- [2] A. Asuncion and D. Newman. UCI machine learning repository, 2007.
- [3] S. Basu, I. Davison, and K. Wagstaff. *Constrained Clustering: Advances in Algorithms and Theory*. Chapman & Hall / CRC, 2008.
- [4] S. Bergmann, J. Ihmels, and N. Barkai. Iterative signature algorithm for the analysis of large-scale gene expression data. *Phys. Rev. E*, 67(3):031902, Mar 2003.
- [5] S. Busygin, O. Prokopyev, and P. M. Pardalos. Biclustering in data mining. *Comput. Oper. Res.*, 35(9):2964–2987, 2008.
- [6] Y. Cheng and G. Church. Biclustering of expression data. In *8th International Conference on Intelligent Systems for Molecular Biology*, 2000.
- [7] T. Dhollander, Q. Sheng, K. Lemmens, B. D. Moor, K. Marchal, and Y. Moreau. Query-driven module discovery in microarray data. *Bioinformatics*, pages 2573–2580, 2007.
- [8] M. A. et al. Gene ontology: tool for the unification of biology. *Nature Genetics*, 25(1):25–29, 2000.
- [9] T. H. et al. Functional discovery via a compendium of expression profile. *Cell*, 102(1):109–126, 2000.

- [10] B. Gamter and R. Wille. *Formal Concept Analysis: Mathematical Foundations*. Springer-Verlag, Berlin, 1999.
- [11] M. A. Hibbs, D. C. Hess, C. L. Myers, C. Huttenhower, K. Li, and O. G. Troyanskaya. Exploring the functional landscape of gene expression: directed search of large microarray compendia. *Bioinformatics*, 23(20):2692–2699, 2007.
- [12] Z. Hu and R. Bhatnagar. Algorithm for discovering low-variance 3-clusters from real-valued datasets. In *ICDM '10*, pages 236–245, 2010.
- [13] S. C. Madeira and A. L. Oliveira. Biclustering algorithms for biological data analysis: A survey. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 1 (1):24–45, 2004.
- [14] M. O. Mohammed J. Zaki. Theoretical foundations of association rules. *3rd SIGMOD'98 Workshop on Research Issues in Data Mining and Knowledge Discovery (DMKD)*, 1998.
- [15] O. Odibat, C. K. Reddy, and C. N. Giroux. Differential biclustering for gene expression analysis. In *BCB '10*, 2010.
- [16] A. B. Owen, J. Stuart, K. Mach, A. M. Villeneuve, and S. Kim. A gene recommender algorithm to identify coexpressed genes in *c. elegans*. *Genome Research*, 13(8):1828–1837, 2003.
- [17] G. Pandey, G. Atluri, M. Steinbach, C. L. Myers, and V. Kumar. An association analysis approach to biclustering. In *KDD '09*, pages 677–686, 2009.
- [18] R. G. Pensa and B. Jean-Francois. Constrain co-clustering of gene expression data. In *SDM '08*, 2008.
- [19] M. Schuldiner, S. Collins, J. Weissman, and N. Krogan. Quantitative genetic analysis in *saccharomyces cerevisiae* using epistatic miniarray profiles (e-maps) and its application to chromatin functions. *Methods*, 40(4):344–352, 2006.
- [20] Y. Song, S. Pan, S. Liu, F. Wei, M. X. Zhou, and W. Qian. Constrained coclustering for textual documents. In *AAAI '10*, 2010.
- [21] Y. Sun, Y. Yu, and J. Han. Ranking-based clustering of heterogeneous information networks with star network schema. In *KDD'09*, 2009.
- [22] C.-J. Wu and S. Kasif. Gems: a web server for biclustering analysis of expression data. *Nucleic Acids Research*, 33(suppl 2):W596–W599, 2005.
- [23] M. J. Zaki and C.-J. Hsiao. Efficient algorithms for mining closed itemsets and their lattice structure. *IEEE Transactions on Knowledge and Data Engineering*, 17 (4):462–478, 2005.
- [24] H. Zhao, L. Cloots, T. Van den Bulcke, Y. Wu, R. De Smet, V. Storms, P. Meysman, K. Engelen, and K. Marchal. Query-based biclustering of gene expression data using probabilistic relational models. *BMC Bioinformatics*, 12(Suppl 1):S37, 2011.

Appendix

Proof of Theorem 3.1

Proof. Let (G_1, M_1) and (G_2, M_2) be α -concepts of \mathbb{K} , such that $M_2 \supseteq M_1$. Assume that $G_2 \not\subseteq G_1$. Two cases arise:

1. $G_2 \supset G_1$. In this case, $\exists g \in G_2 \setminus G_1$ s.t.

$$\begin{aligned} d(K[g, M_2]) &\leq f(\mathbb{K}, g, M_2, \alpha) \\ d(K[g, M_1]) &\leq f(\mathbb{K}, g, M_1, \alpha) \end{aligned}$$

however, this implies that (G_1, M_1) is not an α -concept which contradicts the original assumption.

2. $G_2 \cap G_1 = \emptyset \wedge |G_2| > 0$. Same argument as above by the fact that $M_2 \supseteq M_1$.

Hence, we conclude that $G_2 \subseteq G_1$ implying that $(G_2, M_2) \geq (G_1, M_1)$.

Proof of Theorem 3.2

Proof. By definition, for any $(G_3, M_3) \notin \Upsilon((G_1, M_1)) \wedge (G_3, M_3) \geq (G_1, M_1)$ there exists $(G_2, M_2) \in \Upsilon((G_1, M_1))$ such that $(G_3, M_3) \geq (G_2, M_2)$. This implies that :

$$\begin{aligned} &dist((G_1, M_1), (G_3, M_3)) = \\ (1) \quad &\frac{1}{|G_1 \setminus G_3|} \left(\sum_{g \in G_1 \setminus G_2} \frac{d(\mathbb{K}[g, M_3 \setminus M_1]) \times s_g}{d(\mathbb{K}[g, M_3])} + \sum_{g \in G_2 \setminus G_3} \frac{d(\mathbb{K}[g, M_3 \setminus M_1]) \times s_g}{d(\mathbb{K}[g, M_3])} \right) \\ (2) \quad &\sum_{g \in G_1 \setminus G_2} \frac{d(\mathbb{K}[g, M_3 \setminus M_1])}{d(\mathbb{K}[g, M_3])} > \sum_{g \in G_1 \setminus G_2} \frac{d(\mathbb{K}[g, M_2 \setminus M_1])}{d(\mathbb{K}[g, M_3])} \\ (3) \quad &\sum_{g \in G_2 \setminus G_3} \frac{d(\mathbb{K}[g, M_3 \setminus M_1])}{d(\mathbb{K}[g, M_3])} > \sum_{g \in G_2 \setminus G_2} \frac{d(\mathbb{K}[g, M_2 \setminus M_1])}{d(\mathbb{K}[g, M_3])} \\ (4) \quad &(1 + |\Gamma(g) \cap M_3| \setminus M_1|) > (1 + |\Gamma(g) \cap M_2| \setminus M_1|) \end{aligned}$$

where inequality (2) follows from the order-preserving property of d , inequality (3) follows from the order-preserving property of d , the anti-monotone property of f and the definition of a concept. Inequality (4) follows by the properties of set difference. Hence, by the fact that $|G_1 \setminus G_3| = |G_1 \setminus G_2| + |G_2 \setminus G_3|$ and inequalities (2), (3) and (4).

$$dist((G_1, M_1), (G_3, M_3)) > dist((G_1, M_1), (G_2, M_2))$$

Proof of Theorem 4.1

Proof. By the triangle inequality and definition of consistency we have

$$\begin{aligned} &d(\mathbb{K}[s, PP_l \cup PP_r \cup PP_x]) \leq \\ (5) \quad &d(\mathbb{K}[s, PP_l \cup PP_r]) + d(\mathbb{K}[s, PP_r \cup PP_x]) \\ (6) \quad &d(\mathbb{K}[s, PP_l \cup PP_r \cup PP_x]) \leq f(\mathbb{K}, s, PP_l \cup PP_r \cup PP_x, \alpha) \end{aligned}$$

for any prefix node PP_x and any $s \in S_l \cap S_r \cap S_x$, and d is the range dispersion statistic. Considering each case:

1. $|SS| = |S_l| \wedge |SS| = |S_r|$. In this case $SS = S_l$ and $SS = S_r$ and the subspace $(SS, PP_l \cup PP_r)$ is consistent by definition of range intersection. Hence to ensure closure, PP_l should be replaced by $PP_l \cup PP_r$. Moreover, by equation (6) and the properties of set intersection

$$\begin{aligned} S_l \cap S_x &= S_r \cap S_x \\ &= S_l \cup S_r \cap S_x \\ &\text{implying} \\ \psi_f^\alpha(PP_l \cup PP_x) &= \psi_f^\alpha(PP_r \cup PP_x) \\ &= \psi_f^\alpha(PP_l \cup PP_r \cup PP_x) \end{aligned}$$

for any prefix node PP_x under the same branch.

2. $|SS| = |S_r|$. In this case $S_r \supset S_l$ and the subspace $(SS, PP_l \cup PP_r)$ is consistent by definition of range intersection. Hence, as shown above, it is possible to combine PP_l and PP_r into a single node. On the other hand, the formation of a subspace involving PP_r but not PP_l is possible hence both nodes $PP_l \cup PP_r$ and PP_r must be maintained.
3. $|SS| = |S_l|$. In this case $S_l \supset S_r$ and the subspace $(SS, PP_l \cup PP_r)$ is consistent by definition of range intersection. By the definition of closure and as shown above every consistent subspace involving PP_r will also involve PP_l , therefore PP_r and its children may be pruned. On the other hand, the formation of a subspace involving PP_l but not PP_r is possible hence both nodes $PP_l \cup PP_r$ and PP_l must be maintained.
4. No pruning possible, hence new nodes must be formed.