Co-ClusterD: Data Co-Clustering using Distributed Framework for Sequential Restores

Prasitha Unnikrishnan[1] E. K. Girisan[2]
[1]Research scholar [2]Assistant Professor
[1,2]Department of Computer Science and Engineering
[1,2]Sree Narayana Guru College Coimbatore

Abstract— Co-clustering is an effective data mining tool for existence and diploid data. As data sets become increasingly large, the scalability of existence becomes increasingly crucial. In this paper, we propose two approaches to parallelize co-clustering with subsequent amend in a distributed environment. Based on these two approaches, we present a new distributed framework, Co-ClusterD, that backing effective utilization of co-clustering algorithms with sequential updates. We developed and applied Co-ClusterD, and give its ability through two co-clustering algorithms: information theoretic co-clustering (ITCC), fast nonnegative matrix tri-factorization (FNMTF). We evaluate our framework on both a local cluster of the systems and the Amazon EC2 cloud. Our assessment gives that co-clustering algorithms implemented in Co-ClusterD can attain excellent outputs and run faster than their traditional concurrent counterparts.

Key words: Co-Clustering, Sequential Updates, Cloud Computing, Distributed Framework, FNMTF, ITCC

I. INTRODUCTION

Co-clustering is an effective data mining tool for two-dimensional existence and diploid data. It has practical importance in a vast extend of utilization such as text mining [1], recommendation systems [2], and the analysis of gene expression data [3]. Usually, clustering techniques leverage an iterative refinement method to group input points into clusters. The cluster tasks are accomplished established on the current cluster information (e.g., the centroids of clusters in k-means clustering). The resulted cluster assignments can be utilized to further update the cluster information. Such a refinement method is repeated till the cluster task become stable. Depending on how frequently the cluster information is restored, clustering techniques can be largely classified into two classes. The first class updates the cluster information after all input points have restored their cluster tasks. We refer to this class of algorithms as clustering algorithms with simultaneous restores. In comparison, the second class restores the cluster information whenever a point changes its cluster task. We invoke to this class of techniques as clustering algorithms with sequential updates.

Clustering techniques with sequential restores apparently outperform their concurrent counterparts, since they always advantage the better current cluster data to associate input points. A number of existing studies (e.g., [4], [5]) have backed this allegation. In spite of the possible benefit of sequential updates, parallelizing co-clustering algorithms with sequential restores are claim and not straight-backed by the existing distributed frameworks (e.g., Spark [7] or Hadoop [6]). Particularly, if we let every worker machine update the cluster information sequentially, it ability output in irregular cluster data across worker machines and thus the convergence properties of co-clustering techniques cannot be assured; if we coordinate the cluster information whenever a cluster assignment is changed, it will incur vast coordination overhead and thus output in poor performance in a distributed environment. Consequently, co-clustering techniques with sequential restores cannot be easily performed in a distributed manner.

Toward this end, we recommended two ways to complement sequential updates for co-clustering algorithms. The first method is introduced to as dividing clusters. It separates the problem of clustering rows (or columns) into independent tasks and each of which is assigned to a worker. In order to make tasks independent, we randomly divide row (or column) clusters into various coinciding subsets at the starting of each iteration, and let each worker perform row (or column) clustering with sequential updates on one of these subsets.

The second method is mention to as contrive points. Relaxing the stringent requirement of sequential updates, it complements sequential restores by operating batch restores. Instead of updating the cluster information after each change in cluster task, batch restore accomplished a batch of row (or column) cluster assignments, and then update the cluster data. We regularly separate rows and columns of the input data matrix into several batches and let all workers accomplish row (or column) clustering with simultaneous restores on each batch.

Based on these two approaches, we design and implement a distributed framework, Co-ClusterD, to support efficient utilization of co-clustering techniques with sequential restores. Co-ClusterD provides an abstraction for co-clustering techniques with sequential restores and grant developer to specify the sequential update operations via simple APIs. We decide Co-ClusterD through two co-clustering techniques: fast nonnegative matrix trifactorization (FNMTF) [8] and information theoretic co clustering (ITCC) [9]. Investigate on a local cluster of machines and the Amazon EC2 cloud, we give that co clustering technique realize in Co-ClusterD can obtain better results while running faster than their traditional concurrent counterparts.

II. RELATED WORKS

As large information sets become accepted, developing the scalability of clustering techniques has drawn more and more attention. Lot of scalable clustering techniques is recommended recently. Dave et al. [15] propose a scheme of implementing k-means with simultaneous restores on Microsoft's Windows Azure cloud. Ene et al. [16] design a method of implementing k-center and k-median on Mapreduce. Yin et al. [17] design a distributed framework called FreEM for parallelizing EM algorithms. Whenever, this research is vary from ours as they are devoted to scaling up one-sided clustering algorithms. Folino et al. [18] recommended a complement effective result to the high-order

co-clustering problem (i.e., the problem of together clustering composite types of domain) [19]. George et al. [20] design a parallel version of the weighted Bregman co-clustering algorithm [11] and utilized it to build an efficient real-time collaborative filtering framework. Deodhar et al. [21] design a complement task of the simultaneous co-clustering and learning algorithm [22] established on Mapreduce. Papadimitriou et al. [13] recommend the distributed co-clustering (DisCo) framework, under which multiple co-clustering techniques can be evaluated. However, while these studies focus on parallelizing co-clustering with concurrent updates, our work is devoted to parallelizing co-clustering with sequential updates.

III. CO-CLUSTERING AND UPDATE STRATEGIES

A. Definitions and Overview:

Co-clustering is also known as bi-clustering, block clustering or direct clustering [10]. Formally, given a *m x n* matrix *Z*, a *co-clustering* can be defined by two maps ρ and γ , which groups rows and columns of *Z* into *k* and *l* disjoint or hard clusters respectively. Specifically, $\rho : \{u_0, u_1, \dots . u_m\} \rightarrow \{p_1, p_2, \dots . . p_h\}$ and $\gamma : \{v_0, v_1, \dots . v_n\} \rightarrow \{q_0, q_1, \dots . q_l\}$ where ρ(u) = p means that row u is in row cluster p, and γ(v) = q indicates that column v is in column cluster q. If we reorder rows and columns of Z and let rows and columns of the same cluster be close to each other, we obtain k x l correlated sub-matrices. Each sub-matrix is referred to as a co-cluster.

Typically, the goal of data co-clustering is to find (ρ, γ) such that the following objective function is minimized.

$C(Z, \tilde{Z}) = \sum_{u=1}^{m} \sum_{v=1}^{n} w_{uv}\, d_{\varphi}(Z_{uv}, \tilde{Z}_{uv}) = \sum_{p=1}^{k} \sum_{\{u|\rho(u)=p\}} \sum_{q=1}^{1} \sum_{\{v|\gamma(v)=q} w_{uv}\, (Z_{uv}, S_{pq})$ (1)

Where $C(Z,\tilde{Z})$ is the approximation error between the original matrix Z and the approximation matrix $\tilde{Z}$ uniquely determined by (ρ, γ) w_{uv} denotes the pre-specified weight of pair (u, v), dϕ is a given distance measure (e.g., Euclidean distance), Z_{uv} and $\tilde{Z}_{uv}$ are the elements of Z and $\tilde{Z}$ respectively, S_{pq} is the cluster information that gives the statistic on co-cluster (p, q).

To find the optimal (ρ,γ) a broadly applicable approach is to leverage an iterative process, which monotonically reduce the main function high by associate both row and column clustering iterations. Such kind of coclustering techniques can be invoke to as alternate minimization based co-clustering algorithms [11], which are considered as our main focus in this paper.

Mostly, substitute minimization based co-clustering algorithms repeat the following four steps till convergence.

Step 1. Keep γ fixed, for every row u, find its new row cluster assignment by the following equation.

$$\rho(u) = argmin_p \sum_{\{u|\rho(u)=p\}} \sum_{q=1}^{1} w_{uv}\, (Z_{uv}, S_{pq}) \quad (2)$$

Step 2. With respect to (ρ,γ), update the cluster information (i.e., the statistic of each co-cluster) by the following equation.

$$S_{pq} = argmin_{s_{pq}} \sum_{\{u|p(u)=p\}} \sum_{\{v|\gamma(v)=q\}} w_{uv}\, d_{\varphi}(Z_{uv}, S_{pq}) \quad (3)$$

Step 3. Keep ρ fixed, for every column v, find its

$$\gamma(v) = argmin_q \sum_{p=1}^{k} \sum_{\{u|\rho(u)=p\}} w_{uv}\, d_{\varphi}(Z_{uv}, S_{pq}) \quad (4)$$

Step 4. The same as Step 2. In the above general algorithm, some implementations might combine Step 2 and Step 4 into one step.

B. Co-Clustering with Sequential Updates

Stimulated by the case that sequential restores can attain faster convergence and better results than concurrent updates for clustering techniques, we proposed sequential restore for alternate minimization based co-clustering algorithms. Unlike concurrent updates that perform the cluster information update after all rows (or columns) have updated their cluster tasks, sequential restore execute the cluster data update after each change in cluster assignments.

Specifically, alternate minimization based co-clustering algorithms with sequential updates repeat the following six steps till convergence.

Step 1. Keep γ fixed, pick a row u in some order, find its new row cluster assignment by Eq. (2).

Step 2. With respect to (ρ,γ), update the involved statistics of co-clusters by Eq. (3) once u changes its row cluster assignment.

Step 3. Repeat Step 1 and Step 2 until all rows have been processed.

Step 4. Keep ρ fixed, pick a column v in some order, find its new column cluster assignment by Eq. (4).

Step 5. With respect to (ρ,γ), update the involved statistics of co-clusters by Eq. (3) once v changes its column cluster assignment.

Step 6. Repeat Step 4 and Step 5 until all columns have been processed.

When accomplishing sequential restores, a row (or column) re-task only gives high to the restore of the enumeration of co-clusters connected to the rescheduled row (or column). In include, if the statistic can be updated incrementally (e.g., the statistic are the mean value of the co-cluster or summation of the co-cluster), we can update the statistics by subtracting or adding the belongings of the moved row (or column). Therefore, restoring the cluster information simultaneously does not certainly propose more computational overhead.

IV. PARALLELIZING CO-CLUSTERING WITH SEQUENTIAL UPDATES

A. Dividing Clusters Approach:

Assume in a distributed situation which exists of a number of worker machines, each worker independently performs sequential restores during the repetitive process. The statistics of co-clusters (S_cc) should be updated whenever a row (or column) alteration its cluster task. Whenever, since the workers run concurrently, it may result in inconsistent S_cc across workers. Thus the union farms of co-clustering techniques cannot be managed. Therefore, we recommend dividing clusters method to determine this problem.

The design of the separate clusters method is labeled as follows. Suppose we want to group the input data matrix into l column clusters and k row clusters, the amount of workers is p (p < k and p < l), and each worker w_i holds a subset of rows R_i and a subset of columns C_i. When performing row clustering, we randomly divide row clusters S^r into p non-overlapping row subsets $S_1^r, S_2^r, \dots . S_p^r$. These subsets are distributed to each worker in a one-to-one manner.

When worker w_i receives S_i^r, it can perform row clustering with sequential updates for its rows R_i among the subset of row clusters S_i^r . For example, assume that S_i^r is {1,3,6}, w_i will perform row clustering for its rows whose current cluster assignments are in S_i^r, and allow these rows to change their cluster assignments among row clusters 1, 3, and 6. Since w_i updates only a non-overlapping subset of S_{cc} , the sequential updates on worker w_i will never affect the updates on other workers. The subsets of S_{cc} and cluster indicators updated by each worker will be combined and synchronized over redundancy. Here we have pictorial how to execute row clustering. Column clustering can be done in a similar way.

B. Batching Points Approach:

The separating clusters method removes the reliance on the cluster information for each worker and enables coclustering with sequential updates in a parallel and distributed manner. However, it assumes that the number of workers is less than the amount of clusters. Specific hypothesis might confine the scalability of this approach. For example, when the number of workers is larger than the number of clusters, this approach cannot utilize the extra workers to perform data clustering. Therefore, by relaxing the stringent constraint of sequential updates, we introduce batch updates for alternate minimization based co-clustering algorithms. The difference between batch and sequential updates is that batch updates perform the cluster information update after a batch of rows (or columns) have updated their cluster assignments, rather than after each change in cluster assignments.

The details of the batching points approach are described as follows. Suppose each worker w_i holds a subset of rows R_i and a subset of columns C_i. When performing row clustering, we randomly divide R_i into p non-overlapping subsets R_i^1, R_i^2,….., R_i^p. We refer to R_i^j as a batch. Each worker processes only one of its batches with concurrent updates in each iteration. A synchronization process for the cluster information update is initiated at the end of each iteration. Such iteration continues until all batches have been processed, then it switches to column clustering which is performed in a similar way. During the synchronization process, each worker computes the statistics of co-clusters on R_i (or C_i). We refer to such statistics obtained by each worker as one slice of S_{cc}. All of the slices will be combined to obtain a new S_{cc} used for the next iteration.

V. CO-CLUSTERD FRAMEWORK

A. Design and Implementation:

Since existing distributed frameworks cannot directly support co-clustering algorithms with sequential updates, it calls for a distributed framework that inherently supports sequential updates for co-clustering algorithms. Based on the proposed approaches for parallelizing sequential updates in the previous section, we design and implement Co-ClusterD, a distributed framework for co-clustering algorithms with sequential updates. Co-ClusterD consists of a number of basic workers and a leading worker. Each basic worker performs row and column clustering independently. The leading worker plays a coordination role during the data co-clustering process.

For a given co-clustering job, Co-ClusterD proceeds in two stages: cluster information initialization and data coclustering. In the cluster information initialization stage, assuming there are w workers in the distributed environment, Co-ClusterD first partitions the input data matrix into w row and w column subsets. Next, each worker loads one row subset, one column subset, and the initial cluster assignments. Then, each worker calculates its slice of S_{cc} and sends it to the leading worker. Finally, the leading worker combines all slices of S_{cc} and thus the initial S_{cc} is obtained. In the data co-clustering stage, Co-ClusterD works on the coclustering algorithm implemented by users. The algorithm can be easily implemented by overriding a number of APIs provided by Co-ClusterD. It alternatively performs row and column clusterings until the number of iterations exceeds a user-defined threshold. In particular, for the dividing clusters approach, users can specify the number of iterations repeated for row (or column) clustering before switching to the other side of clustering. For the batching points approach, users can specify the number of row (or column) batches each work holds. To reduce the communication overhead, the input data matrix will not be repartitioned during the data co-clustering stage. In other words, row and column subsets held by each worker will not be shuffled and only the restore cluster tasks and the cluster data will be synchronized among workers over iterations through the network.

Co-ClusterD is realized established on iMapreduce [12], which is a distributed framework based on Hadoop and has built-in backing for repetitive techniques. In fact, Co-ClusterD is independent of the underlying frameworks. We choose iMapreduce since it can better support the iterative processes of co-clustering algorithms.

B. API:

Co-ClusterD allows users without much knowledge on distributed computing to write distributed co-clustering algorithms. Users need to implement only a set of well defined APIs provided by Co-ClusterD. In fact, these APIs are callback functions, which will be automatically invoked by the framework during the data co-clustering process. The descriptions of these APIs are as follows.

1) cProto genClusterProto(bRow, pointS, Ind): Users specify how to generate the cluster prototype, which plays the role of "centroid" in row (or column) clustering, and it can be constructed by the cluster indicators and the statistics of co-clusters S_{cc}. The parameter bRow indicates whether row clustering is performed right now. If bRow is true, pointS is one row of S_{cc}, and Ind is the column-cluster indicators of the input data matrix. Otherwise, pointS is one column of S_{cc} , and Ind is the row-cluster indicators of the input data matrix.
2) double disMeasure(point, cProto): Given a point and a cluster prototype cProto, users specify a measure to quantify the distance between them. point denotes a row or a column.
3) Scc updateSInc(bRow, point preCID, curCID, Scc, rInd, cInd): Users specify how to incrementally update *Scc* when a point changes its cluster assignment from previous cluster preCID to current cluster curCID. rInd and cInd are row-cluster and column-cluster indicators of the input data matrix respectively.

4) slice updateSliceInc(bRow, point, preCID, curCID, slice, subInd, Ind): Users specify how to incrementally update a slice of *Scc* when a point changes its cluster assignment from previous cluster preCID to current cluster curCID. If bRow is true, subInd is the row-cluster indicators of the subset of rows the worker holds, and Ind is the column-cluster indicators of the input data matrix. Otherwise, subInd is the column-cluster indicators of the subset of columns the worker holds, and Ind is the row-cluster indicators of the input data matrix.
5) slice buildOneSlice(bRow, subInd, Ind): Users specify how to build one slice of Scc. The parameters in this function have the same meanings as the parameters in function (4).
6) Scc combineSlices(slices, rInd, cInd): Users specify how to combine the slices of *Scc* sent by workers. Slices is the slices of *Scc* given by workers. rInd and cInd are row-cluster and columncluster indicators of the input data matrix respectively.

The API sets used by the dividing clusters approach and the batching points approach are summarized in table 1.1.

Approach	Initialization	Data Co-clustering
Dividing Clusters	(5), (6)	(1), (2), (3)
Batching Points	(5), (6)	(1), (2), (4), (6)

Table 1: API sets for different approaches

C. *Fault Tolerance and Load Balance:*

In Co-ClusterD, fault tolerance is implemented by a global checkpoint/restore mechanism, which is performed at a user-defined time interval. The cluster information and cluster assignments in the leading worker are dumped to a reliable file system every period of time. If any worker fails (including the leading worker), the computation will roll back to the most recent iteration checkpoint and resume from that iteration. In addition, since Co-ClusterD is based on iMapreduce [12], it also inherits all the salient features of Mapreduce's style fault tolerance.

The synchronous computation model used by our Co- ClusterD framework also makes load balance become an important issue. This is because the running time of each iteration is dependent on the slowest worker. Therefore, if the capacity of each computer in the distributed environment is homogeneous, the workload should be evenly distributed. Otherwise, the workload should be distributed according to the capacity of each worker.

VI. EXPERIMENTAL RESULTS

In this portion, we classify the capability and ability of Co-ClusterD in the context of two clustering algorithms ITCC and FNMTF on various real-time data sets. We analyse Co-ClusterD to a state-of-the-art Hadoop based co-clustering framework DisCo [13]. Because DisCo does not backing co-clustering techniques with sequential updates, we implement only concurrent restore in DisCo. Whenever, in Co-ClusterD, both simultaneous and sequential restores are implemented. The experiments are performed on both smallscale and large-scale clusters.

A. *Experiment Setup*

We use real world data sets downloaded from UCI Machine Learning Repository [14] to evaluate the co-clustering algorithms. These data sets are summarized in Table 1.2.

Data sets	Samples	Features	non-zeros
KOS	3430	6906	467714
NIPS	1500	12419	1900000
ENRON	39861	28102	6400000

Table 2: Description of data sets

We build a small-scale cluster of local machines and a large-scale cluster on the Amazon EC2 cloud to run experiments. The small-scale cluster consists of 4 machines. Each machine has Intel Core 2 Duo E8200 2.66GHz processor, 3GB of RAM, 1TB hard disk, and runs 32-bit Linux Debian 6.0 OS. These 4 machines are connected to a switch with communication bandwidth of 1Gbps. The large-scale cluster consists of 100 High-CPU medium instances on the Amazon

EC2 Cloud. Each instance has 1.7GB memory and 5 EC2 compute units.

B. *Small-Scale Experiments:*

For small-scale experiments, data sets KOS and NIPS are used for evaluation. The number of row (or column) clusters is set to 40. For the dividing clusters approach, each worker is randomly assigned a subset of row (column) clusters with size 10 in each iteration. In addition, the number of iterations repeated for row (or column) clustering is set to 4. For the batching points approach, each worker divides its subset of rows (or columns) into 16 batches.

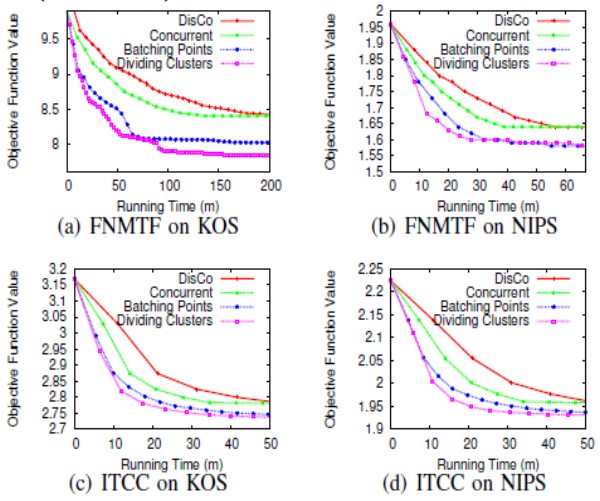

Fig. 1: Cost function vs. running time comparisons for co-clustering algorithms with different update strategies

As shown in Figure 1, we can observe that co-clustering algorithms with concurrent updates implemented in Co- ClusterD (denoted by "Concurrent") converge faster than those implemented in DisCo (denoted by "DisCo") although they converge to the same values. This is because Co- ClusterD leverages a persistent job for the iterative process of the co-clustering algorithm rather than using one job for one row (or column) clustering iteration which is adopted by DisCo. Hence, Co-ClusterD reduces the repeated job initialization overhead in each iteration and achieves fast convergence. In addition, we can also observe that algorithms parallelized by the dividing clusters approach and the batching points approach converge faster and obtain better results than their concurrent counterparts implemented in DisCo and Co-ClusterD.

C. *Large-Scale Experiments:*

To validate the scalability of our framework, we also run experiments on the Amazon EC2 cloud. The data set ENRON is used for evaluation. The number of row (or column) clusters is set to 20. Since the scalability of the dividing clusters approach is dependent on the relationship between the number of workers and the number of clusters, only the batching points approach is evaluated. The number of row (or column) batches each worker holds is set to 4.

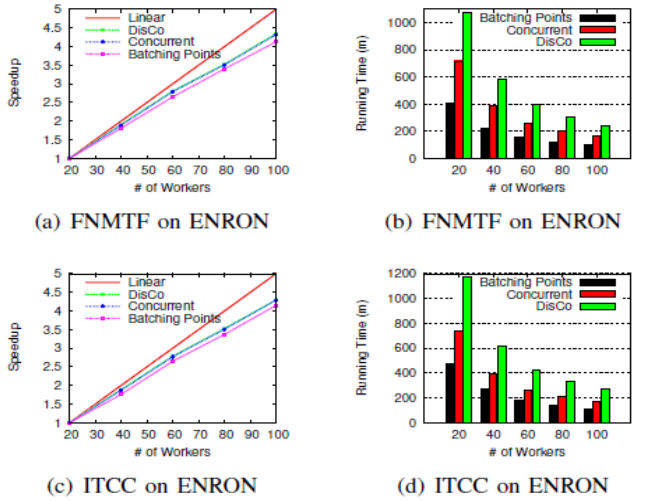

Fig. 2: Speedup and performance comparisons

As shown in Figure 2(a) and Figure 2(c), algorithms implemented in DisCo (denoted by "DisCo") and Co-ClusterD with concurrent updates (denoted by "Concurrent") obtain the same speedup. This is because both of them perform concurrent updates. We can also observe that their speedups are better than Co-ClusterD using the batching points approach. The reason lies in that concurrent updates performs less cluster information updates than the batching points approach and thus results in less synchronization overhead. However, since the bases of computing speedups are different, a better speedup does not necessarily lead to a shorter running time. As shown in Figure 2(b) and Figure 2(d), co-clustering algorithms parallelized by the batching points approach still converge much faster than their concurrent counterparts implemented in DisCo and Co-ClusterD.

VII. Conclusion

In this paper, we propose dividing clusters and batching point's approaches to parallelize co-clustering with sequential updates. Based on these two approaches, we design and implement a distributed framework referred to as Co- ClusterD, which supports effective utilization of coclustering techniques with sequential restores. Experimental results show that co-clustering techniques utilized in Co-ClusterD can obtain better results and run faster than their traditional concurrent counterparts.

Reference

[1] I. Dhillon, "Co-clustering documents and words using bipartite spectral graph partitioning," in KDD. ACM, 2001, pp. 269–274.

[2] S. Daruru, N. M. Marin, M. Walker, and J. Ghosh, "Pervasive parallelism in data mining: dataflow solution to co-clustering large and sparse netflix data," in KDD. ACM, 2009, pp. 1115–1124.

[3] H. Cho, I. Dhillon, Y. Guan, and S. Sra, "Minimum sumsquared residue co-clustering of gene expression data," in SIAM, vol. 114, 2004.

[4] N. Slonim, N. Friedman, and N. Tishby, "Unsupervised document classification using sequential information maximization," in SIGIR. ACM, 2002, pp. 129–136.

[5] I. Dhillon, Y. Guan, and J. Kogan, "Iterative clustering of high dimensional text data augmented by local search," in ICDM, 2003, pp. 131 – 138.

[6] "Hadoop," http://http://hadoop.apache.org/.

[7] M. Zaharia, M. Chowdhury, M. Franklin, S. Shenker, and I. Stoica, "Spark: cluster computing with working sets," in HotCloud. USENIX Association, 2010, pp. 10–10.

[8] H. Wang, F. Nie, H. Huang, and F. Makedon, "Fast nonnegative matrix tri-factorization for large-scale data co-clustering," in IJCAI, 2011.

[9] I. Dhillon, S. Mallela, and D. Modha, "Information-theoretic co-clustering," in KDD. ACM, 2003, pp. 89–98.

[10] J. Hartigan, "Direct clustering of a data matrix," Journal of the American Statistical Association, pp. 123–129, 1972.

[11] A. Banerjee, I. Dhillon, J. Ghosh, S. Merugu, and D. Modha, "A generalized maximum entropy approach to bregman coclustering and matrix approximation," in KDD. ACM, 2004, pp. 509–514.

[12] Y. Zhang, Q. Gao, L. Gao, and C. Wang, "iMapreduce: A distributed computing framework for iterative computation," in IPDPSW. IEEE, 2011, pp. 1112–1121.

[13] S. Papadimitriou and J. Sun, "Disco: Distributed co-clustering with map-reduce: A case study towards petabyte-scale endto- end mining," in ICDM. IEEE, 2008, pp. 512–521.

[14] "UCI Machine Learning Repository," http://archive.ics.uci.edu/ml/datasets.html.

[15] A. Dave, W. Lu, J. Jackson, and R. Barga, "Cloudclustering: Toward an iterative data processing pattern on the cloud," in IPDPSW. IEEE, 2011, pp. 1132–1137.

[16] A. Ene, S. Im, and B. Moseley, "Fast clustering using MapReduce," in KDD. ACM, 2011, pp. 681–689.

[17] J. Yin, Y. Zhang, and L. Gao, "Accelerating expectationmaximization algorithms with frequent updates," in CLUSTER. IEEE, 2012, pp. 275–283.

[18] F. Folino, G. Greco, A. Guzzo, and L. Pontieri, "Scalable parallel co-clustering over multiple heterogeneous data types," in HPCS. IEEE, 2010, pp. 529–535.

[19] G. Greco, A. Guzzo, and L. Pontieri, "Coclustering multiple heterogeneous domains: Linear combinations and agreements," TKDE, vol. 22, no. 12, pp. 1649–1663, 2010.

[20] T. George and S. Merugu, "A scalable collaborative filtering framework based on co-clustering," in ICDM. IEEE, 2005 pp. 4–pp.

[21] M. Deodhar, C. Jones, and J. Ghosh, "Parallel simultaneous co-clustering and learning with map-reduce," in GrC. IEEE, 2010, pp. 149–154.

[22] M. Deodhar and J. Ghosh, "A framework for simultaneous coclustering and learning from complex data," in KDD. ACM, 2007, pp. 250–259.

