

# Consensus Clusterings

Nam Nguyen, Rich Caruana  
Department of Computer Science, Cornell University  
Ithaca, New York 14853  
{nhnguyen,caruana}@cs.cornell.edu

## Abstract

*In this paper we address the problem of combining multiple clusterings without access to the underlying features of the data. This process is known in the literature as clustering ensembles, clustering aggregation, or consensus clustering. Consensus clustering yields a stable and robust final clustering that is in agreement with multiple clusterings. We find that an iterative EM-like method is remarkably effective for this problem. We present three iterative algorithms for finding clustering consensus. An extensive empirical study compares our proposed algorithms with eleven other consensus clustering methods on four data sets using six different clustering performance metrics. The experimental results show that the new ensemble clustering methods produce clusterings that are as good as, and often better than, these other methods.*

## 1. Introduction

Clustering often is a first step in data analysis. Many different clustering methods have been developed [9, 21] such as hierarchical agglomerative clustering, mixture densities, graph partitioning, and spectral clustering. Most clustering methods focus on finding a single optimal or near-optimal clustering according to some specific clustering criterion. Consensus clustering can provide benefits beyond what a single clustering algorithm can achieve. Consensus clustering algorithms often: generate better clusterings; find a combined clustering unattainable by any single clustering algorithm; are less sensitive to noise, outliers or sample variations; and are able to integrate solutions from multiple distributed sources of data or attributes.

In addition to the benefits outlined above, consensus clustering can be useful in a variety of domains. For example, clustering categorical data (where it is more difficult to define useful distance metrics) can be considered as a consensus clustering problem where each discrete feature is viewed as a simple clustering of the data. Then the consensus clustering algorithm can be applied to the ensemble

of all clusterings produced by discrete features of the data set. An example of categorical data is a *Movie* database where some of the discrete attributes are *Director*, *Actor*, *Actress*, *Genre*, *Year*, etc. As another example, consensus clustering can be employed in “privacy-preserving” scenarios where it is not possible to centrally collect all of the underlying features for all data points, but only how the data points are grouped together. Such a situation might arise when different companies or governmental agencies have information about individuals that they can not share, but still need to find high-quality clusters of individuals. For such cases, the consensus clustering algorithm offers a natural model for clustering the data maintained in separate sites in a privacy-preserving manner, that is, without the need for different sites to fully-reveal their data to each other, and without the need for relying on a trusted authority.

In this paper, we propose three EM-like consensus clustering algorithms which utilize a feature map constructed from the set of base level clusterings. Our experiments show that the proposed algorithms produce results as good as or better than other existing consensus clustering algorithms. The paper is structured as follows: in section 2 we introduce the consensus clustering framework; in section 3 we give a brief description of related works in this area; in section 5 we describe in detail our proposed algorithms; and the experimental results and conclusion are given in section 8 and 11.

## 2. Consensus Clustering Framework

We are given a set of  $N$  data points  $\mathbf{X} = \{x_1, x_2, \dots, x_N\}$  and a set of  $C$  clusterings  $\mathbf{\Pi} = \{\pi_1, \pi_2, \dots, \pi_C\}$  of the data points in  $\mathbf{X}$ . Each clustering  $\pi_i$  is a mapping from  $\mathbf{X}$  to  $\{1, \dots, n_{\pi_i}\}$  where  $n_{\pi_i}$  is the number of clusters in  $\pi_i$ . The problem of clustering consensus is to find a new clustering  $\pi^*$  of the data  $\mathbf{X}$  that best summarizes the clustering ensemble  $\mathbf{\Pi}$ .

Similar to the approach of [19], our proposed algorithms utilize the space of constructed features induced by the ensemble of clusterings  $\mathbf{\Pi}$ . For each data point  $x$ , we construct a corresponding  $C$ -dimensional feature vector  $y$ , where the

$i^{\text{th}}$  feature is simply the cluster label from the clustering  $\pi_i$ . The set of vectors  $\mathbf{Y} = \{y_1, y_2, \dots, y_N\}$  where

$$y_i = \langle \pi_1(x_i), \pi_2(x_i), \dots, \pi_C(x_i) \rangle,$$

is exploited in our proposed algorithms to find the consensus clustering.

### 3. Related Work

Consensus clustering has recently attracted the interest of a number of researchers in the machine learning community. This section provides a brief summary of the work in this area. (This is by no means a complete survey of consensus clusterings.)

#### 3.1. Pairwise Similarity Approach

In this approach, a measure of similarity between a pair of data points can be estimated as the ratio of a number of clusterings in which the two data points are shared the same clusters to the total number of clusterings in the ensemble. More precisely, the similarity between two data points  $x_i$  and  $x_j$  is defined as,

$$S_{ij} = S(x_i, x_j) = \frac{1}{C} \sum_{c=1}^C \mathcal{I}(\pi_c(x_i) = \pi_c(x_j)), \quad (1)$$

where  $\mathcal{I}$  is the indicator function. Thus, any similarity-based clustering algorithms can be applied to the similarity matrix  $S$  to find a consensus clustering of the ensemble.

We experiment with two similarity-based clustering algorithms: Furthest Consensus (FC) [7] and Hierarchical Agglomerative Clustering Consensus (HAC) [5, 6, 12]. In both of these algorithms, the matrix  $S$  is used as the similarity measure.

**Furthest Consensus (FC):** The goal of the algorithm is to find  $K$  cluster centers that are furthest from each other. The algorithm starts with finding a pair of data points that are furthest apart, and assigning them as cluster centers. Then it repeatedly finds a next cluster center that is furthest apart from previous found centers. At the end, all data points are assigned to its closest center.

**Hierarchical Agglomerative Clustering Consensus (HAC):** HAC algorithm is a standard bottom-up algorithm for the correlation clustering problem. It starts by placing all data points into singleton clusters. Then it repeatedly merges two clusters that have the largest averaged similarity measure which is given in the similarity matrix  $S$ . The algorithm stops when there are  $K$  remaining clusters.

#### 3.2. Graph-based Approach

In [17] the authors propose three consensus clustering algorithms: Cluster-based Similarity Partitioning Algorithm

(CSPA), HyperGraph Partitioning Algorithm (HGPA), and Meta-Clustering Algorithm (MCLA). All algorithms first transform the ensemble of clusterings into a graph representation.

**CSPA:** First, the similarity matrix is computed as in equation 1. Then, an induced similarity graph, where vertices correspond to data points and edges' weights to similarity measures, is partitioned into  $K$  clusters using METIS [11].

**HGPA:** The hypergraph is constructed using the ensemble of  $C$  clusterings,  $\mathbf{\Pi} = \{\pi_1, \pi_2, \dots, \pi_C\}$ , where vertices correspond to data points and each hyperedge is represented by a cluster from one of the clusterings in the ensemble. All hyperedges have the same weight. This algorithm looks for a hyperedge separator that partitions the hypergraph into  $K$  unconnected components of approximately the same size. The hypergraph partitioning package HMETIS [10] is used to perform the partitioning.

**MCLA:** This algorithm is based on clustering clusters where each cluster is also represented as a hyperedge. The algorithm groups and collapses related hyperedges into  $K$  clusters; and then assigns each data point to the collapsed hyperedge in which it participates most strongly.

#### 3.3. Mutual Information Approach

In [18], an objective function is formulated as the mutual information between the target consensus clustering and the clustering ensemble. Then the quadratic mutual information is maximized via an EM algorithm in the space of constructed features as defined in Section 2 to find the target consensus function. This algorithm usually requires multiple restarts in order to avoid convergence to low quality local minima.

#### 3.4. Mixture Model Approach

In [19], the authors propose a probabilistic model of consensus using a finite mixture of multinomial distributions in a space of constructed features as defined in Section 2. A combined clustering is found as a solution to the corresponding maximum likelihood problem using the EM algorithm.

#### 3.5. Cluster Correspondence Approaches

In [2], the authors first solve the problem of finding the correspondence between clusters of different clusterings in the ensemble either by constrained and unconstrained search via optimizing a linear programming formulation (CCC and CCU), or with Singular Value Decomposition (CCSVD). Then a simple voting procedure is applied to assign data points into clusters.

### 3.6. Objective Dependence Approach

The representative example of this approach is the **BEST** algorithm [7]. Given an objective function, BEST chooses among all clusterings in its ensemble one that maximizes the objective function.

In our experiments, we compare the proposed algorithms with these eleven algorithms: Furthest Consensus (FC), Hierarchical Agglomerative Clustering Consensus (HAC), Cluster-based Similarity Partitioning Algorithm (CSPA), HyperGraph Partitioning Algorithm (HGPA), Meta-Clustering Algorithm (MCLA), Quadratic Mutual Information Algorithm (QMI), EM Mixture Model Algorithm (EM), BEST, Cluster Correspondence Constrained Search (CCC), Cluster Correspondence Unconstrained Search (CCU), and Cluster Correspondence SVD (CCSVD).

### 4. Motivations

In this section, we point out some drawbacks of consensus clusterings algorithms and advantages of EM-like approach to consensus clustering problem. In the consensus clustering setting, pairwise similarity often does not reflect a good measure of similarity between data points, especially when the number of base-clusterings are limited. Consider a 5-items data set,  $\{x_i\}_{i=1}^5$ , where the ground-truth partition of the data is  $\{(x_1, x_2, x_3), (x_4, x_5, x_6)\}$ . The four clusterings of the data set and the similarity matrix are shown in Figure 1. We notice that similarity-based clustering algorithms will not be able recover the ground-truth partition of the data set since the similarities of the pairs  $\{(x_1, x_6), (x_3, x_6)\}$  have the highest values. However, an EM-like approach (described in the next section) can recover the ground-truth partition since it is a local minimum in the EM process.

Clustering	$x_1$	$x_2$	$x_3$	$x_4$	$x_5$	$x_6$
I	1	1	1	2	2	1
II	2	1	1	2	2	2
III	2	1	2	1	2	2
IV	1	1	2	1	2	2

	$x_1$	$x_2$	$x_3$	$x_4$	$x_5$	$x_6$
$x_1$			1/2	1/2	1/2	1/2
$x_2$			1/2	1/2	0	1/4
$x_3$				0	1/2	3/4
$x_4$					1/2	1/4
$x_5$						3/4
$x_6$						

Figure 1. Example of Consensus Clustering

Most of the current consensus clustering algorithms return a single consensus clustering as the final result. EM-like approaches have an advantage over these because they can generate multiple consensus clusterings with different restarts, and the best consensus clustering with respect the evaluation criteria can be selected.

### 5. Consensus Clustering Algorithms

In this section we present in detail our EM-like consensus clustering algorithms. Our proposed algorithms can be applied to any group of clusterings, including clusterings in which each individual clustering may have different numbers of clusters. In addition, our new algorithms conform to two constraints which most other consensus clustering algorithms follow. The first constraint is that the number of clusters  $K$  in the target clustering  $\pi^*$  is given. We are not trying to solve the difficult problem of determining the correct number of clusters. The second constraint is to only use the information provided by the ensemble of clusterings, i.e.  $\Pi = \{\pi_1, \pi_2, \dots, \pi_C\}$ , without any access to the underlying features of data points in  $\mathbf{X}$ .

In the following sections, we describe in details the three iterative consensus clustering algorithms.

#### 5.1. Iterative Voting Consensus (IVC)

The algorithm is an iterative process which utilizes the constructed feature-vectors  $\mathbf{Y} = \{y_1, y_2, \dots, y_N\}$  where

$$y_i = \langle \pi_1(x_i), \pi_2(x_i), \dots, \pi_C(x_i) \rangle$$

induced by the ensemble of clusterings  $\Pi$ . Each cluster in the target consensus clustering has a cluster center which is also a  $C$ -dimensional vector.

Each iteration of the algorithm involves two steps: the first step computes the cluster center of each cluster in the target consensus clustering, and second step reassigns each data point to its closest cluster center. In the first step, the  $i^{th}$  feature of the cluster center vector is just the majority values of all  $i^{th}$  features of data points belonging to the considered cluster. This is possible since all  $i^{th}$  features of data points are coming from the same clustering  $\pi_i$  in the clustering ensemble. In the second step, we assign each data point to its closest center by finding the center of the minimum distance to the considered data point. The distance of a point to a cluster center is just the Hamming distance between the two vector representations. The pseudo-code of the IVC algorithm is described in Algorithm 1.

#### 5.2. Iterative Probabilistic Voting Consensus (IPVC)

Similar to the previous algorithm, this iterative algorithm also utilizes the constructed feature-vectors  $\mathbf{Y} = \{y_1, y_2, \dots, y_N\}$ , induced by the ensemble of clusterings  $\Pi$ . In each iteration of the algorithm, each data point is reassigned to different clusters based on a defined distance function between the considered data point and the previous established clusters in the target consensus clustering. The main operation in each iteration is to find a closest cluster

---

**Algorithm 1** Iterative Voting Consensus

---

**Input:** a set of  $N$  data points  $\mathbf{X} = \{x_1, x_2, \dots, x_N\}$   
a set of  $C$  clusterings  $\mathbf{\Pi} = \{\pi_1, \pi_2, \dots, \pi_C\}$   
 $K$  is a desired number of clusters

**Output:** a consensus clustering  $\pi^*$  with  $K$  clusters

Initialize  $\pi^*$

**repeat**

Let  $P_i = \{y \mid \pi^*(y) = i\}$  be the  $i^{\text{th}}$  cluster

Compute the representation of each cluster:  
 $y_{P_i} = \langle \text{majority}\{(P_i)_1\}, \dots, \text{majority}\{(P_i)_C\} \rangle$ ,  
where  $(P_i)_j$  is the set of the  $j^{\text{th}}$  features of all data points in  $P_i$

**for**  $y$  in  $\mathbf{Y}$  **do**

Re-assign  $\pi^*(y) \leftarrow \text{argmin}_i D(y, y_{P_i})$ , where  
 $D(y, y_{P_i}) = \sum_{j=1}^C \mathcal{I}((y)_j \neq (y_{P_i})_j)$

**end for**

**until**  $\pi^*$  does not change

---

---

**Algorithm 2** Iterative Probabilistic Voting Consensus

---

**Input:** a set of  $N$  data points  $\mathbf{X} = \{x_1, x_2, \dots, x_N\}$   
a set of  $C$  clusterings  $\mathbf{\Pi} = \{\pi_1, \pi_2, \dots, \pi_C\}$   
 $K$  is a desired number of clusters

**Output:** a consensus clustering  $\pi^*$  with  $K$  clusters

Initialize  $\pi^*$

**repeat**

Let  $P_i = \{y \mid \pi^*(y) = i\}$  be the  $i^{\text{th}}$  cluster,  $n_i = |P_i|$

**for**  $y$  in  $\mathbf{Y}$  **do**

Re-assign  $\pi^*(y) \leftarrow \text{argmin}_i D(y, P_i)$ , where

$$D(y, P_i) = \sum_{j=1}^C \frac{\sum_{y' \in P_i} \mathcal{I}((y)_j \neq (y')_j)}{n_i}$$

**end for**

**until**  $\pi^*$  does not change

---

for each data point via a defined distance measure between them. Formally, the distance between a data point  $y$  and a cluster of  $c$  data points  $\{y_1, y_2, \dots, y_c\}$ , is defined as

$$D(y, \{y_1, y_2, \dots, y_c\}) = \sum_{j=1}^C \frac{\sum_{i=1}^c \mathcal{I}((y)_j \neq (y_i)_j)}{c}. \quad (2)$$

The pseudo-code of the IPVC algorithm is described in algorithm 2. This algorithm can be viewed as a refinement of the IVC algorithm. Particularly, the distance function takes

into account the proportion that each feature of a point differs from those of points in a cluster of the target clustering.

### 5.3. Iterative Pairwise Consensus (IPC)

This iterative algorithm utilizes the similarity matrix  $S$  as defined in equation 1. In each iteration, each data point is reassigned to different clusters based on the similarity measure between the data point and the previous established clusters in the target consensus clustering. Formally, the similarity measure between a data point  $x$  and a cluster of  $c$  data points  $\{x_1, x_2, \dots, x_c\}$ , is defined as

$$S(x, \{x_1, x_2, \dots, x_c\}) = \frac{\sum_{i=1}^c S(x, x_i)}{c}. \quad (3)$$

---

**Algorithm 3** Iterative Pairwise Consensus

---

**Input:** a set of  $N$  data points  $\mathbf{X} = \{x_1, x_2, \dots, x_N\}$   
a set of  $C$  clusterings  $\mathbf{\Pi} = \{\pi_1, \pi_2, \dots, \pi_C\}$   
 $K$  is a desired number of clusters

**Output:** a consensus clustering  $\pi^*$  with  $K$  clusters

Compute the  $n \times n$  similarity matrix  $S$ , where  
 $S(x_i, x_j) = \frac{1}{C} \sum_{c=1}^C \mathcal{I}(\pi_c(x_i) = \pi_c(x_j))$

Initialize  $\pi^*$

**repeat**

Let  $P_i = \{x \mid \pi^*(x) = i\}$  be the  $i^{\text{th}}$  cluster of the data points, and  $n_i = |P_i|$

**for**  $x$  in  $\mathbf{X}$  **do**

Re-assign  $\pi^*(x) \leftarrow \text{argmax}_i S(x, P_i)$ , where  
 $S(x, P_i) = \frac{\sum_{x' \in P_i} S(x, x')}{n_i}$

**end for**

**until**  $\pi^*$  does not change

---

Different from the previous two, this algorithm is a pairwise similarity approach. We can view this algorithm as applying a variation of k-means method to the similarity matrix  $S$ . The pseudo-code of the IPC algorithm is described in algorithm 3.

## 6. Evaluation Criteria

Evaluating the quality of a clustering is a nontrivial and ill-posed task [16]. In supervised learning, model performance is assessed by comparing model predictions to targets. In clustering we do not have targets and usually do not know *a priori* what groupings of the data are best. This hinders discerning when one clustering is better than another, or when one clustering algorithm outperforms another. In

general, there are two main approaches to evaluate consensus clusterings: consensus criteria measure how the target consensus clustering is in agreement with all clusterings in the ensemble, and clustering criteria measure how well the target consensus clustering performs on the underlying features of the data points  $\mathbf{X}$ , possibly with respect to hidden true labels on these points.

## 6.1. Consensus Criteria

Given an ensemble of clusterings  $\mathbf{\Pi} = \{\pi_1, \pi_2, \dots, \pi_C\}$  and a target consensus clustering  $\pi^*$ , the consensus performance is computed as

$$Perf(\pi^*, \mathbf{\Pi}) = \frac{1}{C} \sum_{i=1}^C f(\pi^*, \pi_i), \quad (4)$$

where  $f(\pi^*, \pi_i)$  is a measure of how similar or dissimilar the two clusterings are. In our experiment, three different distance measures and the normalized mutual information between the two clusterings  $\pi$  and  $\pi^*$  are used in place of the function  $f(\pi^*, \pi_i)$ .

To talk about the different measures of distance and the normalized mutual information between two clusterings, we first introduce some notation following convention in [13]. A clustering  $\pi$  is a partition of the data points  $\mathbf{X} = \{x_1, x_2, \dots, x_N\}$  into sets  $C_1, C_2, \dots, C_K$  called clusters such that  $C_k \cap C_l = \emptyset$  and  $\bigcup_{k=1}^K C_k = \mathbf{X}$ . Let the number of data points in cluster  $C_k$  be  $n_k$ . Hence, we have that  $N = \sum_{k=1}^K n_k$ . Let the second clustering of the same data points  $\mathbf{X}$  be  $\pi' = \{C'_1, C'_2, \dots, C'_{K'}\}$ , with cluster size  $n'_{k'}$ . Next, we compute the so-called confusion matrix, or contingency table  $M$  of the clustering pair  $(\pi, \pi')$ , whose the  $m_{kk'}$  element is the number of points in the intersection of clusters  $C_k$  of  $\pi$  and  $C'_{k'}$  of  $\pi'$ ,  $m_{kk'} = |C_k \cap C'_{k'}|$ . Next, we define the entropy associated with clustering  $\pi$  as

$$H(\pi) = - \sum_{k=1}^K P(k) \log P(k), \quad (5)$$

where  $P(k) = \frac{n_k}{N}$ ; and the mutual information between the pair of clusterings  $(\pi, \pi')$  as

$$MI(\pi, \pi') = \sum_{k=1}^K \sum_{k'=1}^{K'} P(k, k') \log \frac{P(k, k')}{P(k)P(k')}, \quad (6)$$

where  $P(k, k') = \frac{|C_k \cap C'_{k'}|}{N}$ .

The first clustering distance is based on set matching. This symmetric distance metric was introduced by [20]:

$$f(\pi, \pi') = \frac{1}{2N} (2N - \sum_{k=1}^K \max_{k'} m_{kk'} - \sum_{k'=1}^{K'} \max_k m_{kk'}) \quad (7)$$

The second clustering distance is based on counting pairs of points on which two clusterings agree or disagree. The (adjusted) Rand distance was introduced by [8] of Rand's [15] criterion,

$$f(\pi, \pi') = \frac{N_{10} + N_{01}}{N_{11} + N_{10} + N_{01} + N_{00}}, \quad (8)$$

where  $N_{11}$  and  $N_{00}$  are the number of point pairs that are in the same cluster, in the different clusters respectively under both  $\pi$  and  $\pi'$ ;  $N_{10}$  and  $N_{01}$  are the number of point pairs that are in the same cluster under  $\pi$  but not under  $\pi'$ , and the same cluster under  $\pi'$  but not under  $\pi$  respectively.

The third clustering distance is based on the variation of information. The variation of information (VI) distance [13] measures the amount of information that is lost or gained in changing from clustering  $\pi$  to clustering  $\pi'$ ,

$$f(\pi, \pi') = H(\pi) + H(\pi') - 2MI(\pi, \pi'). \quad (9)$$

Finally, the normalized mutual information between the pair of clusterings  $(\pi, \pi')$  was used in [17],

$$f(\pi, \pi') = \frac{MI(\pi, \pi')}{\sqrt{H(\pi)H(\pi')}}. \quad (10)$$

## 6.2. Clustering Criteria

In addition to evaluate a clustering based on how well it combines multiple clusterings in the ensemble, there are also two additional clustering criteria to measure how well a clustering partitions the data into its natural groupings. In our experiment, we examine two clustering criteria: compactness and accuracy. Compactness measures the average pairwise distances between points in the same cluster:

$$Compactness(\pi) = \frac{1}{N} \sum_{k=1}^K n_k \left( \frac{\sum_{x_i, x_j \in C_k} d(x_i, x_j)}{n_k(n_k - 1)/2} \right), \quad (11)$$

where  $d(x_i, x_j)$  is the distance between  $x_i$  and  $x_j$ .

All of our test data sets have external true labels which are not used by the clustering process. The second clustering criterion is accuracy, which measures how well the target clustering perform in comparison to the external true labels of the data points:

$$Accuracy(\pi) = \frac{\sum_{k=1}^K majority(C_k | L_k)}{N}, \quad (12)$$

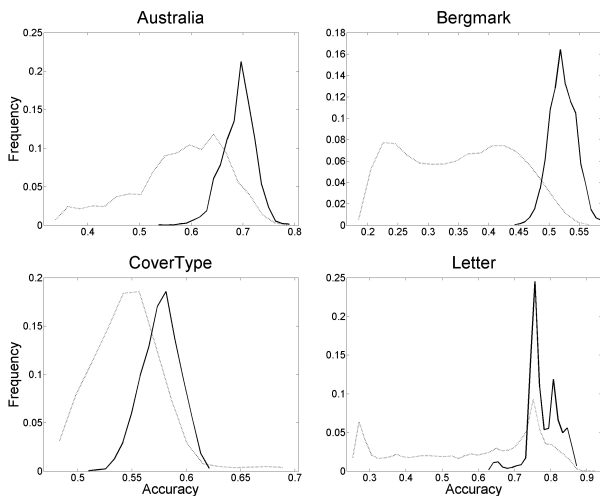
where  $majority(C_k | L_k)$  is the number of points with the plurality label in the  $C_k$  cluster (if label  $l$  appeared in cluster  $k$  more often than any other label, then  $majority(C_k | L_k)$  is the number of points in  $C_k$  with label  $l$ ).

Data Set	features	cases	classes	clusters
Australia	17	245	10	10
Bergmark	254	1000	25	25
Coverttype	49	1000	7	15
Letters	617	514	7	10

**Table 1. Description of Data Sets**

## 7. Data Sets

We evaluate our consensus clustering algorithms on four data sets: Australia, Bergmark, Coverttype, and Letter. The Australia Coastal data is a subset of the data available from the Biogeoinformatics of Hexacorals environmental database [1]. The Bergmark data was collected using 25 focused web crawls, each with a different keyword. The variables are counts in a bags-of-words model describing the web pages. The Coverttype data is from the UCI Machine Learning Repository [14]. It contains cartographic variables sampled at  $30 \times 30$  meter grid cells in four wilderness areas in Roosevelt National Forest in northern Colorado. The letters data is a subset of the isolet spoken letter data set from the UCI Machine Learning Repository [14].



**Figure 2. Two different accuracy distributions (dotted line: Feature Weighting K-means, solid line: K-means) of four data sets.**

In our experiments, different clusterings of each data set have the same number of clusters, and the same number of clusters is chosen for the target consensus clusterings as well. The clusterings that will be combined are generated for the data sets using two of the methods described in [3]: feature weighting k-means and k-means with different random restarts. The accuracy distributions of these two meth-

ods are quite different as shown in Figure 2. The clusterings of the feature weighting k-means category are generated by applying a Zipf feature weighting method and tend to spread out from low accuracy to high accuracy. The clusterings of the k-means category are generated by applying the k-means algorithm with different random initializations and usually concentrate on the high range of accuracy.

## 8. Experiments

Similar to k-means our proposed algorithms are sensitive to the initial partition of the data. Hence, for each ensemble of clusterings our algorithms are repeated with 100 different initializations. The initial partition is randomly selected as one of the clusterings in the ensemble. (We tried initializing the new algorithms both with random partitions of the data and with clusterings from the base level clusterings. We observe that better initializations usually lead to better results, so the random partition initialization does not work as well as initializing with a base-level clustering.)

In our first experiment, the number of clusterings in the ensemble is 200. In this experiment, all consensus clustering algorithms are evaluated with respect to six different performance criteria as described in Section 6. For the VI distance, Dongen Distance, Rand distance, and compactness, lower values indicate better performance. For normalized mutual information and accuracy, higher values indicate better performance. Since different evaluation criteria of the consensus clustering results under two ensemble distributions (i.e. k-means and feature weighting k-means) are very similar, we simply interpretation of results by showing the averaged scores.

Figure 3 shows the performance of 14 consensus clustering methods on the six metrics averaged across the four problems and two ensemble distributions. Table 2 shows the number of times that each algorithm comes in first for a given metric across four different data sets and two clustering distributions.

The three new algorithms perform comparably to each other. Across the six different performance criteria, the new algorithms produce results that usually are equal to, and often better than, other algorithms. Of the existing methods, the graph-based approach tends to have the worst performance compared to others, especially the HyperGraph Partitioning Algorithm (HGPA). In addition, we notice that the HAC and CCSVD algorithms sometimes are competitive in comparison to the new algorithms on some performance metrics with some data sets. However, none of the existing algorithms consistently produce results as good as the new methods across different performance metrics and different data sets.

In the second experiment, we evaluate how different performance criteria changes as the algorithm iterates. In Fig-

	VI Distance	Dongen Distance	Rand Distance	Normalize MI	Accuracy	Compactness	Total
IVC	1/8	3/8	2/8	1/8	3/8	0/8	10/48
IPVC	4/8	2/8	2/8	3/8	0/8	0/8	11/48
IPC	2/8	3/8	4/8	2/8	5/8	0/8	16/48
BEST	0/8	0/8	0/8	0/8	0/8	0/8	0/48
FC	0/8	0/8	0/8	0/8	0/8	0/8	0/48
HAC	1/8	0/8	0/8	2/8	0/8	2/8	5/48
CSPA	0/8	0/8	0/8	0/8	0/8	0/8	0/48
HGPA	0/8	0/8	0/8	0/8	0/8	0/8	0/48
MCLA	0/8	0/8	0/8	0/8	0/8	0/8	0/48
EM	0/8	0/8	0/8	0/8	0/8	0/8	0/48
QMI	0/8	0/8	0/8	0/8	0/8	2/8	2/48
CCC	0/8	0/8	0/8	0/8	0/8	0/8	0/48
CCU	0/8	0/8	0/8	0/8	0/8	0/8	0/48
CCSVD	0/8	0/8	0/8	0/8	0/8	4/8	4/48

**Table 2. The win statistic for a given method and a performance metric across four data sets and two different clustering distributions.**

ure 4, we show the evaluation scores of the IVC algorithm on the CoverType data set. The scores of two other algorithms, and on the other data sets, have similar behavior. We notice that as iterations increase, performance improves. Even though the proposed algorithms do not aim directly at optimizing with respect to these different performance criteria, the algorithms do achieve good scores, suggesting that the heuristics are effective.

In the third experiment, there are 1000 clusterings in the ensembles. In this experiment, we want to visualize where the consensus clusterings fall in the space of base-level clusterings. One way to achieve the visualization is the scatter plot of accuracy versus compactness of the ensemble of clusterings as well as the consensus clusterings found by all algorithms. The results for the Bergmark and CoverType data sets are shown in Figure 5. Consistent with the results above, the new algorithms tend to produce better consensus clusterings in terms of accuracy. Interestingly, the consensus clusterings found with the new algorithms sometimes fall outside of the space of clusterings in the original ensembles.

## 9 Complexity Analysis

Similar to k-means, the space complexity of IVC and IPVC is  $O(NC)$  where  $N$  is the number of data points and  $C$  is the number of base-level clusterings. The space complexity of IPC is  $O(N^2)$  since it needs to keep track of all pair-wise distances.

Since our proposed algorithm are EM algorithms, the number of iterations that the algorithms update the target clustering before getting to a local minima is varied and depends on the data sets. In our experiments, we observe that

the algorithms always terminate with less than 40 iterations. There are many speed-up techniques (i.e. using inequality triangle to accelerate k-means [4]) developed for k-means which may also be applied to our algorithms to improve running time.

## 10 Comparison of the three algorithms

In term of the evaluation criteria, the three algorithms perform similarly to each other. There is no statistical significant difference among the three algorithms. In comparison to eleven other consensus clustering methods, the three algorithms consistently perform as well as or even better in all evaluation criteria. However, among the three algorithms IVC is the simplest one and easy to implement. These properties makes IVC a desirable algorithm for practical usage.

## 11. Conclusion

We present three EM-like algorithms for the consensus clustering problem. An extensive empirical study shows that the proposed algorithms yield results as good as, and often better than, eleven other clustering consensus methods. Our proposed algorithms are in essence variations of k-means algorithm using different distance measures applied to the vector of base-level clusterings. The results indicate that this approach works well for consensus clustering. It is interesting that this simple, and somewhat *obvious*, approach performs better than other more complicated methods in the literature.

In future work, we intend to enhance the new algorithms by relaxing the two constraints of the new algorithms mentioned in Section 5. First, we want to use consensus to de-

termine the best number of clusters in the consensus clustering. Second, we will develop methods to leverage the underlying features of the original data.

## Acknowledgments

This work was supported by NSF CAREER Grant # 0347318.

## References

- [1] Biogeoinformatics of Hexacorals. Environmental database, <http://www.kgs.ku.edu/hexacoral/>, 1998.
- [2] C. Boulis and M. Ostendorf. Combining multiple clustering systems. In *The 8th European conference on Principles and Practice of Knowledge Discovery in Databases(PKDD), LNAI 3202*, pages 63–74, 2004.
- [3] R. Caruana, M. Elhawary, N. Nguyen, and C. Smith. Meta clustering. In *Proceedings IEEE International Conference on Data Mining*, 2006.
- [4] C. Elkan. Using the triangle inequality to accelerate k-means. In *Proceedings of the 20th International Conference on Machine Learning*, 2003.
- [5] X. Z. Fern and C. E. Brodley. Random projection for high dimensional data clustering: A cluster ensemble approach. In *Proceedings of the 20th ICML*, 2003.
- [6] A. L. Fred and A. K. Jain. Data clustering using evidence accumulation. In *Proceedings of the 16th International Conference on Pattern Recognition*, 2002.
- [7] A. Gionis, H. Mannila, and P. Tsaparas. Clustering aggregation. In *Proceedings of the 21st ICDM*, 2005.
- [8] L. Hubert and P. Arabie. Comparing partitions. *Journal of Classification*, 2:193–218, 1985.
- [9] A. K. Jain, M. N. Murty, and F. P. J. Data clustering: a review. *ACM Computing Surveys*, 31:264–323, 1999.
- [10] G. Kharypis, R. Aggarwal, V. Kumar, and S. Shekhar. Multilevel hypergraph partitioning: Applications in vlsi domain. In *Proceedings of the Design and Automation Conference*, 1997.
- [11] G. Kharypis and V. Kumar. Multiple k-way partitioning scheme for irregular graphs. *SIAM Journal on Scientific Computing*, 20(1):359–392, 1998.
- [12] G. N. Lance and W. T. Williams. A general theory of classificatory sorting strategies. i. hierarchical systems. *Computer Journal*, 9:373–380, 1967.
- [13] M. Meila. Comparing clusterings by the variation of information. In *Proceedings of the Sixteenth Annual Conference of Computational Learning Theory (COLT)*, 2003.
- [14] D. J. Newman, S. Hettich, C. L. Blake, and C. J. Merz. UCI repository of machine learning databases, <http://www.ics.uci.edu/~mllearn/mlrepository.html>, 1998.
- [15] W. M. Rand. Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical Association*, 66:846–850, 1971.
- [16] N. Slonim and N. Tishby. Agglomerative information bottleneck. In *Proceedings of Neural Information Processing Systems*, 1999.
- [17] A. Strehl and J. Ghosh. Cluster ensembles - a knowledge reuse framework for combining multiple partitions. *Journal of Machine Learning Research*, 3:583–617, 2002.
- [18] A. Topchy, A. K. Jain, and W. Punch. Combining multiple weak clusterings. In *Proceedings IEEE International Conference on Data Mining*, 2003.
- [19] A. Topchy, A. K. Jain, and W. Punch. A mixture model for clustering ensembles. In *Proceedings SIAM International Conference on Data Mining*, 2004.
- [20] van Dongen. Performance criteria for graph clustering and markov cluster experiments. Technical Report INS-R0012, Centrum voor Wiskunde en Informatica, 2000.
- [21] D. Verma and M. Meila. A comparison of spectral clustering algorithms. Technical Report submitted, Department of Computer Science, University of Washington, 2001.



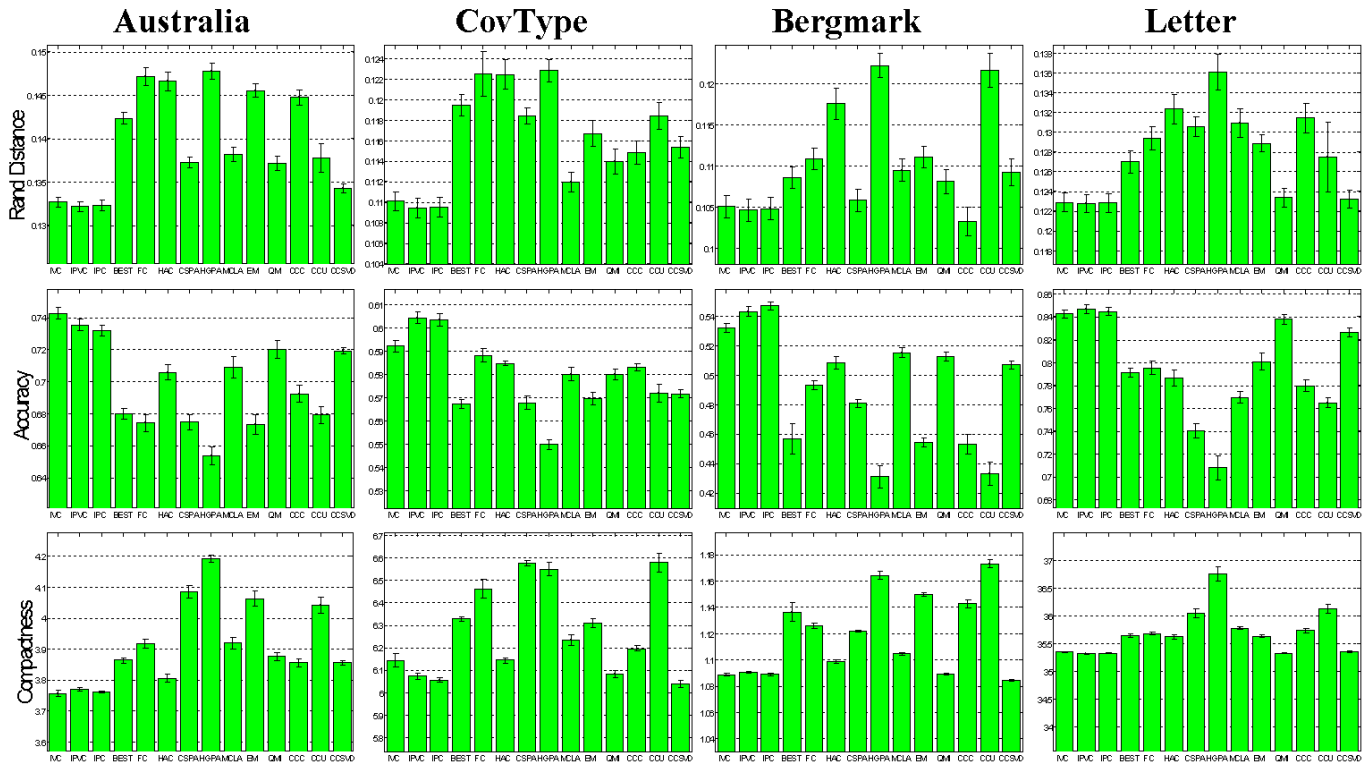


Figure 3. Consensus clustering results for different evaluation criteria.

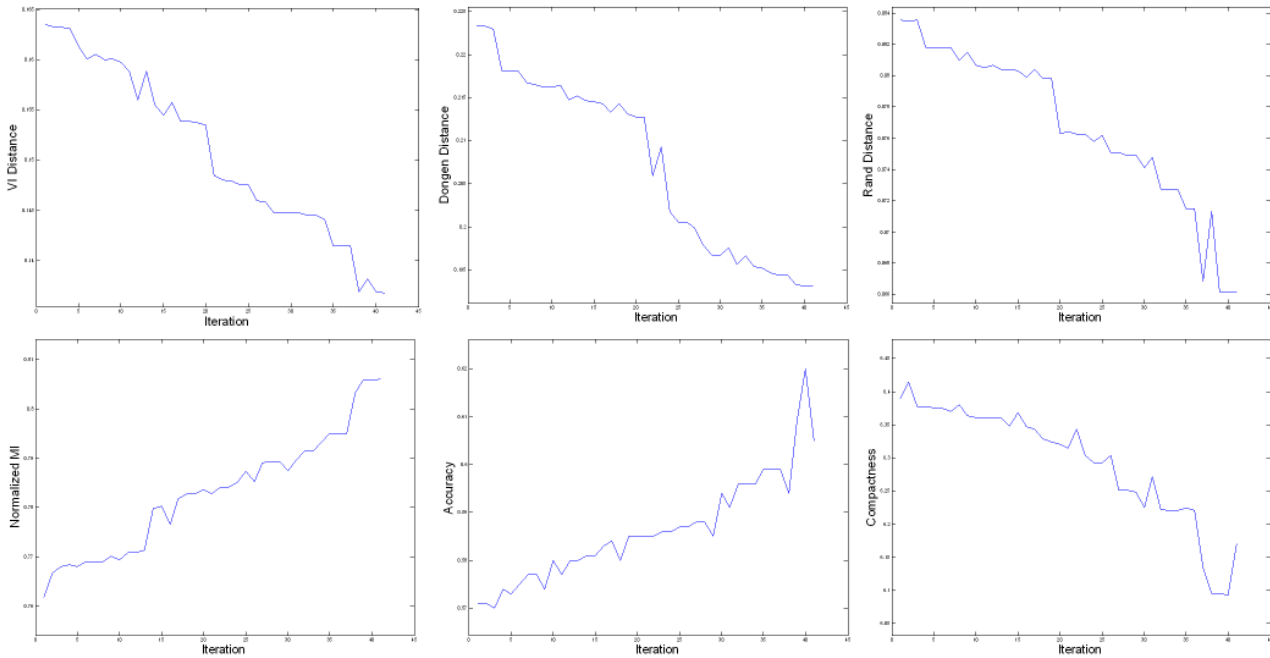


Figure 4. IVC Convergence of difference evaluation scores of the CoverType data set.

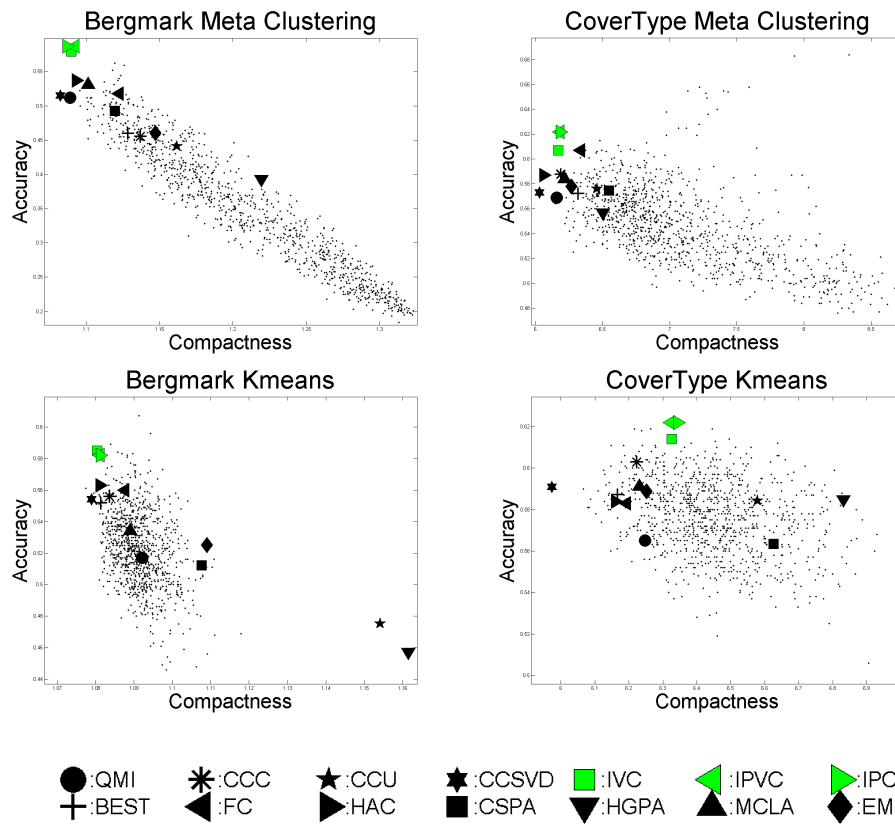


Figure 5. Accuracy and compactness scatter plots of consensus clusterings for four data sets