

Statistical File Matching of Flow Cytometry Data

Gyemin Lee, William Finn, and Clayton Scott

Abstract

Flow cytometry is a technology that rapidly measures antigen-based markers associated to cells in a cell population. Although analysis of flow cytometry data has traditionally considered one or two markers at a time, there has been increasing interest in multidimensional analysis. However, flow cytometers are limited in the number of markers they can jointly observe, which is typically a fraction of the number of markers of interest. For this reason, practitioners often perform multiple assays based on different, overlapping combinations of markers. In this paper, we address the challenge of imputing the high dimensional jointly distributed values of marker attributes based on overlapping marginal observations. We show that simple nearest neighbor based imputation can lead to spurious subpopulations in the imputed data, and introduce an alternative approach based on nearest neighbor imputation restricted to a cell's subpopulation. This requires us to perform clustering with missing data, which we address with a mixture model approach and novel EM algorithm. Since mixture model fitting may be ill-posed, we also develop techniques to initialize the EM algorithm using domain knowledge. We demonstrate our approach on real flow cytometry data.

Index Terms

statistical file matching, flow cytometry, mixture model, probabilistic PCA, EM algorithm, imputation, incomplete data, clustering

I. INTRODUCTION

Flow cytometry is a technique for quantitative cell analysis [1]. It provides simultaneous measurements of multiple characteristics of individual cells. Typically, a large number of cells are analyzed in a short

G. Lee and C. Scott are with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI, USA. E-mail: {gyemin, cscott}@eecs.umich.edu. W. Finn is with the Department of Pathology, University of Michigan, Ann Arbor, MI, USA. E-mail: wgfinn@umich.edu.

period of time – up to thousands of cells per second. Since its development in the late 1960s, flow cytometry has become an essential tool in various biological and medical laboratories. Major applications of flow cytometry include hematological immunophenotyping and diagnosis of diseases such as acute leukemias, chronic lymphoproliferative disorders, and malignant lymphomas [2].

Flow cytometry data has traditionally been analyzed by visual inspection of one-dimensional histograms or two-dimensional scatter plots. Clinicians will visually inspect a sequence of scatter plots based on different pairwise marker combinations, and perform gating, the manual selection of marker thresholds, to eliminate certain subpopulations of cells. They identify various pathologies based on the shape of cell subpopulations in these scatter plots. There has been recent work, reviewed below, on automatic cell gating or classification of pathologies based on multidimensional analysis of flow cytometry data.

Despite the promise of multidimensional analysis, this direction is limited by the number of markers that can be simultaneously measured, which is typically a fraction of the number of markers of interest. It is therefore common in practice to perform multiple assays based on different, overlapping combinations of markers. We may view these combinations as different marginals of the joint distribution of all observed markers. However, even when analysis is based on visual inspection of scatter plots, problems arise when the desired marker pair was not jointly measured. This situation arises frequently in the analysis of historical data.

To address these issues and to facilitate higher dimensional analysis, we present a statistical method for file matching, which imputes higher dimensional flow cytometry data from multiple lower dimensional data files. While [3] proposed a simple approach based on Nearest Neighbor (NN) imputation, this method is prone to induce spurious clusters, as we demonstrate below. Our method can improve the file matching of flow cytometry and is less likely to generate false clusters.

In the following, we explain the principles of flow cytometry and introduce the file matching problem in the context of flow cytometry data. We then present an approach to file matching which imputes a cell's missing marker values with the values of the nearest neighbor among cells of the same type. To implement this approach we develop a method for clustering with missing data. We model flow cytometry data with a latent variable Gaussian mixture model, where each Gaussian corresponds to a cell type, and develop an expectation-maximization (EM) algorithm to fit the model. Since a large majority of all values are unobserved, most covariances cannot be estimated from the data. However, domain experts possess considerable knowledge about the characteristics of different cell types, and we incorporate this knowledge into the initialization of the EM algorithm. We compare our method with simple nearest neighbor imputation on real flow cytometry data, and show that our method offers improved performance.




Fig. 1. A flow cytometer system. As a stream of cells passes through a laser beam, the photo-detectors detect forward angle light scatter, side angle light scatter, and light emissions from fluorochromes. Then the digitized signals are analyzed in a computer.

II. BACKGROUND AND MOTIVATION

In this section, we explain the principles of flow cytometry. We also define the statistical file matching problem in the context of flow cytometry data, and motivate the need for an improved solution.

A. Flow Cytometry

In flow cytometry analysis, a cell suspension is first prepared from peripheral blood, bone marrow, or lymph node. The suspension of cells is then mixed with a solution of fluorochrome-labeled antibodies. Typically, each antibody is labeled with a different fluorochrome. As the stream of suspended cells passes through a focused laser beam, they either scatter or absorb the light. If the labeled antibodies are attached to proteins of a cell, the associated fluorescent markers absorb the laser and emit light with the corresponding wavelength (color). Then a set of photo-detectors in the line of the light beam and perpendicular to the light capture the scattered and emitted light. The signals from the detectors are digitized and stored in a computer system. Forward scatter (FS) and side scatter (SS) signals as well as various fluorescence signals are collected for each cell (see Fig. 1).

In a flow cytometer that is capable of measuring d attributes, called *markers*, the measurements of each cell can be represented with a d -dimensional vector $\mathbf{x} = (x^{(1)}, x^{(2)}, \dots, x^{(d)})$ where $x^{(1)}$ is FS, $x^{(2)}$ is SS, and $x^{(3)}, \dots, x^{(d)}$ are the fluorescent markers. Thus, the accumulation of N cells forms a $N \times d$ matrix.

The detected signals provide information about the physical and chemical properties of each cell analyzed. FS is related to the relative size of the cell and SS is related to its internal granularity or complexity. The fluorescence signals reflect the abundance of expressed antigens on the cell surface. These various attributes are used for identification and quantification of cell populations. FS and SS are always measured, while the marker combination is a part of the experimental design.

Flow cytometry data is usually analyzed using a sequence of one dimensional histograms and two or three dimensional scatter plots by choosing a subset of one, two or three markers. The analysis typically involves manually selecting and excluding cell subpopulations, called gating, by thresholding and drawing boundaries on the scatter plots. Clinicians routinely diagnose by visualizing the scatter plots.

Recently, some attempts have been made to analyze directly in high dimensional spaces by mathematically modeling flow cytometry data. In [4], [5], a mixture of Gaussian distributions is used to model cell populations, while a mixture of t -distributions with a Box-Cox transformation is used in [6]. A mixture of skew t -distributions is studied in [7]. The knowledge of experts is sometimes incorporated as prior information [8]. Instead of using finite mixture models, some recent approaches proposed information preserving dimension reduction to analyze high dimensional flow cytometry data [9], [10]. However, standard techniques for multi-dimensional flow cytometry analysis are not yet established.

B. Statistical File Matching

The number of dimensions in flow cytometry is limited by the number of light sources and detectable fluorochrome markers, and available reagent combinations. Even though recent innovations have enabled measuring near 20 cellular attributes, there are typically dozens or hundreds of markers of interest in a given flow cytometry experiment. Furthermore, instruments deployed in clinical laboratories still only measure 5-7 markers simultaneously [11].

Being unable to simultaneously measure all markers of interest, it is common to divide a sample into several “tubes” and stain each tube separately with a different set of markers [12]. In practice, partially overlapping marker combinations are used to help identify cell populations (see Fig. 2). The marker combinations are designed based on which markers need to be observed together. However, it is not always possible to anticipate all marker combinations of potential interest.

In the sequel, we present a method that generates flow cytometry data in which all the markers of interest are available for the union of cells. Thus, we obtain a single higher dimensional dataset beyond the current limits of instrumentation. Then pairs of markers that are not measured together can still be visualized through scatter plots, and methods of multidimensional analysis may be applied to the full

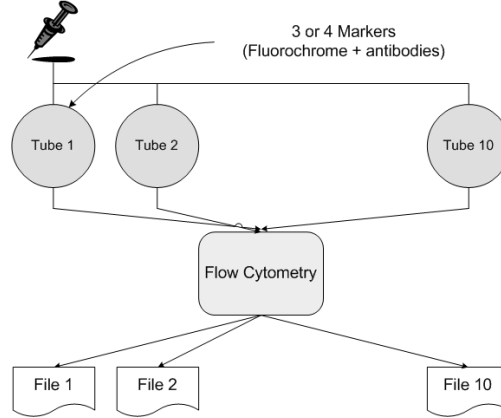


Fig. 2. Flow cytometry analysis on a large number of antibody reagents within a limited capacity of a flow cytometer. A sample from a patient is separated into multiple tubes with which different combinations of fluorochrome labeled antibodies are stained. Each output file contains at least two variables, FS and SS, in common as well as some variables that are specific to the file.




Fig. 3. Data structure of two incomplete data files. Two files have some overlapping variables c , and some variables s_1 and s_2 that are never jointly observed. File matching combines the two files by completing the missing blocks of variables.

dataset.

This technique, called file matching, merges two or more datasets that have some commonly observed variables as well as some variables unique to each dataset. An exemplary two file case is drawn in Fig. 3. Each unit (cell) \mathbf{x}_n is a vector in \mathbb{R}^d and belongs to one of the data files (tubes) \mathcal{X}_1 or \mathcal{X}_2 , where each file has N_1 and N_2 units, respectively. While variables c are observed in all the units, units in \mathcal{X}_1 have variables s_2 missing and units in \mathcal{X}_2 have variables s_1 missing, where s_1 , s_2 , and c represent specific and common variable sets. If the observed and missing components of a unit \mathbf{x}_n are denoted by o_n and m_n , then $o_n = c \cup s_1$ and $m_n = s_2$ for $\mathbf{x}_n \in \mathcal{X}_1$, and $o_n = c \cup s_2$ and $m_n = s_1$ for $\mathbf{x}_n \in \mathcal{X}_2$.

The file matching problem is a missing data problem where blocks of missing data need to be imputed. Among imputation methods, algorithms using conditional mean or regression are most common. As shown

in Fig. 4, however, these imputation algorithms tend to shrink the variance of data. Thus, these approaches are inappropriate in flow cytometry where the shapes of cell populations are important in analysis, and the preservation of variability after file matching is highly desired. More discussions on missing data analysis and file matching can be found in [13] and [14].




Fig. 4. Examples of imputation methods: NN, conditional mean, and regression. The NN method relatively well preserves the distribution of imputed data, while other imputation methods such as conditional mean and regression significantly reduce the variability of data.

[3] proposed to use Nearest Neighbor (NN) imputation to match flow cytometry data files. In their approach, missing variables of one unit, called the recipient, are imputed with observed variables from a unit in the other file, called the donor, that is most similar. If \mathbf{x}_i is a unit in \mathcal{X}_1 , the missing variables are set as follows

$$\mathbf{x}_i^{s_2} = \mathbf{x}_j^{*s_2} \text{ where } \mathbf{x}_j^* = \arg \min_{\mathbf{x}_j \in \mathcal{X}_2} \|\mathbf{x}_i^c - \mathbf{x}_j^c\|_2.$$

Note that the similarity is based on the distance in the projected space of jointly observed variables. This algorithm is advantageous over other imputation algorithms, based on conditional mean or regression, as displayed in Fig. 4. It generally preserves the distribution of cells, while the other methods cause the variance structure to shrink toward zero.

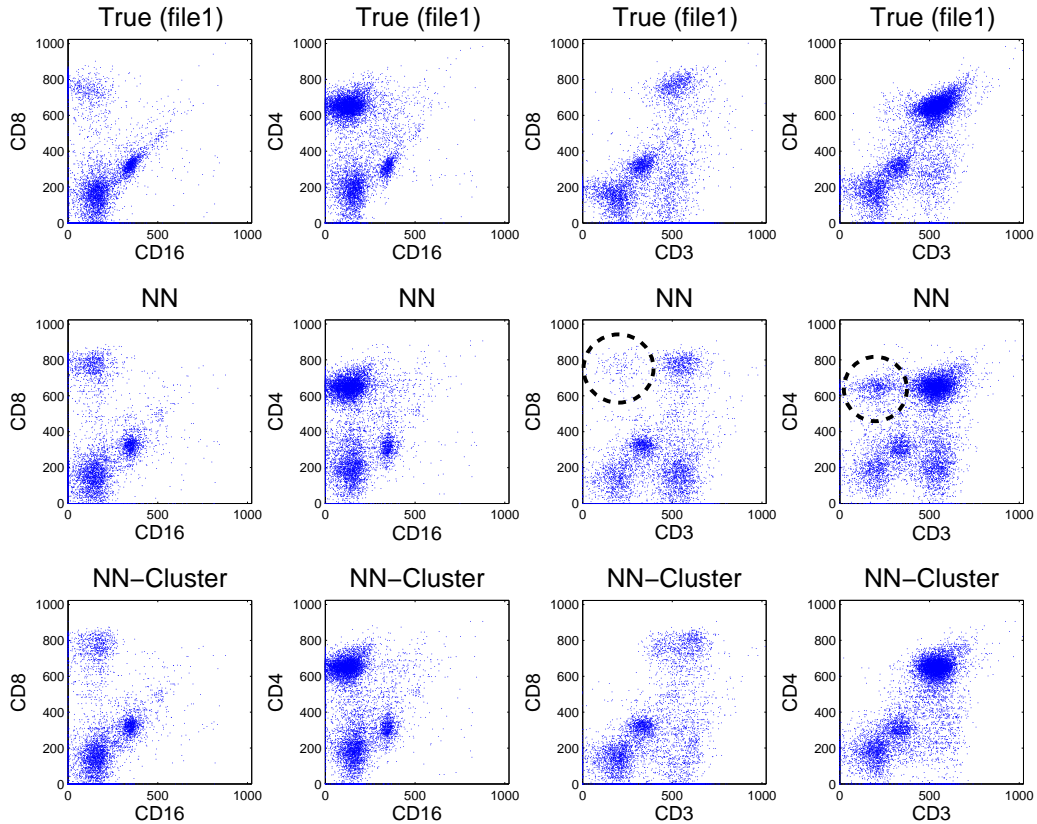


Fig. 5. Comparison of results for two imputations methods to the ground truth cell distribution. Figures show scatter plots on pairs of markers that are not jointly observed. The middle row and the bottom row shows the imputation results from the NN and Cluster-based NN, respectively. The results from the NN method show spurious clusters conspicuously in the right two panels. The false clusters are indicated by dotted circles in CD3 vs. CD8 and CD3 vs. CD4 scatter plots. On the other hand, the results from our proposed approach better resemble the true distribution on the top row.

However, the NN method sometimes introduces spurious clusters into the imputed results and fails to replicate the true distribution of cell populations. Fig. 5 shows an example of false clusters from the NN imputation algorithm (for detailed experimental setup, see Section V). We present a toy example to explain why the NN imputation can fail, and to motivate our approach.

C. Motivating Toy Example

Fig. 6 shows a toy example dataset in \mathbb{R}^3 . In the two data files, two of three features of these points are observed: c and s_1 in file 1, and c and s_2 in file 2. Each data point belongs to one of two clusters, but its label is unavailable.

When imputing feature s_1 of units in file 2, the NN algorithm produces four clusters whereas there should be two, as shown in Fig. 6 (d). This is because the NN method uses only one feature, and fails to leverage the information about the joint distribution of variables that are not observed together. On the other hand, if we can infer the cluster membership of data points, the NN imputation can be applied within the same cluster. Hence, we search a donor from the subgroup (1) for the data points in (3), and likewise we search a donor from (2) for the points in (4) in the example. Then the file matching result greatly improves and better replicates the true distribution as in Fig. 6 (e).

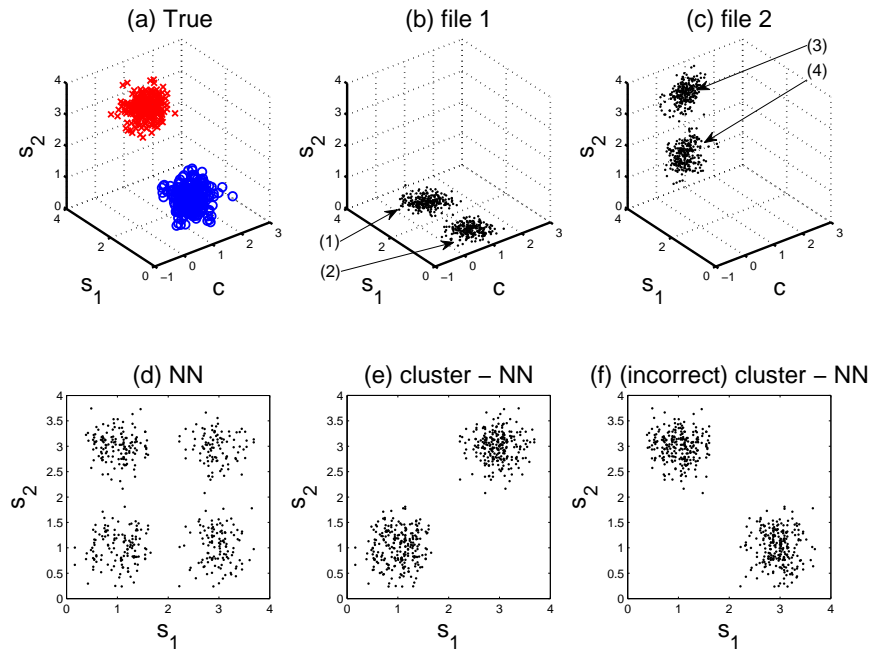


Fig. 6. Toy example of file matching. Two files (b) and (c) provide partial information of data points (a) in \mathbb{R}^3 . The variable c is observed in both files while s_1 and s_2 are specific to each file. The NN method creates false dot populations in the s_1 vs. s_2 scatter plot in (d). On the other hand, the NN applied within the same cluster successfully replicated the true distribution. If the cluster are incorrectly paired, however, the Cluster-NN approach fails, as in (f).

In this example, as in real flow cytometry data, there is no way to infer cluster membership from the

Input: two files \mathcal{X}_1 and \mathcal{X}_2 to be matched

1. Cluster the units in \mathcal{X}_1 and \mathcal{X}_2 .
2. Perform NN imputation within the same cluster.

Output: statistically matched complete files $\widehat{\mathcal{X}}_1$ and $\widehat{\mathcal{X}}_2$

Fig. 7. The description of the Cluster-based NN algorithm for two files. For two input flow cytometry data files \mathcal{X}_1 and \mathcal{X}_2 , specific variables are imputed using NN method after clustering each cell into one of K clusters.

data alone, and incorrect labeling can lead to poor results (Fig. 6 (f)). Fortunately, in flow cytometry we can incorporate prior knowledge to achieve an accurate clustering.

III. CLUSTER-BASED IMPUTATION OF MISSING VARIABLES

We first focus on the case of matching two files. The case of more than two files is discussed in Section VI. For the present section, we assume that \mathcal{X}_1 and \mathcal{X}_2 have both been partitioned into K clusters. Let \mathcal{X}_1^k and \mathcal{X}_2^k denote the cells in \mathcal{X}_1 and \mathcal{X}_2 from the k th cluster, respectively.

Suppose that the data is configured as in Fig. 3. In order to impute the missing variables of a unit in file 1, we locate a donor among the data points in file 2 that have the same cluster label as the recipient. When imputing incomplete units in file 2, the roles change. The similarity between two units is evaluated on the projected space of jointly observed variables, while constraining both units to belong to the same cluster. Then we impute the missing variables of the recipient by patching the corresponding variables from the donor. More specifically, for $\mathbf{x}_i \in \mathcal{X}_1^k$, we impute the missing variables by

$$\mathbf{x}_i^{m_i} = \mathbf{x}_j^{*m_i} \text{ where } \mathbf{x}_j^* = \arg \min_{\mathbf{x}_j \in \mathcal{X}_2^k} \|\mathbf{x}_i^c - \mathbf{x}_j^c\|_2$$

The proposed Cluster-based NN imputation algorithm is summarized in Fig. 7.

In social applications such as survey completion, file matching is often performed on the same class such as gender, age, or county of residence. However, this information that is used to label each unit is available in data, and the inference as in our algorithm is not necessary [14].

IV. CLUSTERING WITH MISSING DATA

To implement the above approach, we view \mathcal{X}_1 and \mathcal{X}_2 as a single data set and cluster its elements. We propose a method for clustering with missing data based on a finite Gaussian mixture model. Mixture models are common models for flow cytometry where each component corresponds to a cell type.

While non-Gaussian models might provide a better fit, there is a trade-off between estimation error and approximation error. More complicated models tend to be more challenging to fit. Furthermore, even with an imperfect data model, we may still achieve an improved file matching.

Thus, clustering amounts to fitting the parameters of the mixture model. In general, fitting such a model is ill-posed. For example, in the toy example, there is no way to know the correct cluster inference based solely on the data. However, we can leverage domain knowledge to select the number of components and initialize model parameters.

A. Mixture of PPCA

In a mixture model framework, the probability distribution of a d -dimensional data vector \mathbf{x} takes the form

$$p(\mathbf{x}) = \sum_{k=1}^K \pi_k p_k(\mathbf{x})$$

where K is the number of components in the mixture and π_k is a mixing weight.

In flow cytometry, mixture models are common models of cell subpopulations. Mixture models with Gaussian components are common [4], [5], [8], although distributions with more parameters, such as t -distributions or skew t -distributions, have been proposed [6], [7]. However, these models require estimating a large number of parameters, and it becomes difficult to obtain reliable estimates when the number of components or the dimensions of the data increase. In this application, the model needs not be perfect to get improved imputation. We adopt a probabilistic principal component analysis (PPCA) mixture model as a way to model cell populations with fewer parameters. Without PPCA, our experience has revealed that even a Gaussian mixture model may have too many parameters to be accurately fit.

PPCA was proposed by [15] as a probabilistic interpretation of PCA. While conventional PCA lacks a probabilistic formulation, PPCA specifies a generative model. It is a latent variable model, in which a data vector is linearly related to a latent variable. The latent variable space is generally lower dimensional than the ambient variable space, so the latent variable provides an economical representation of the data.

The PPCA model is built by specifying a conditional distribution of a data vector \mathbf{x} in \mathbb{R}^d , given a latent variable \mathbf{t} in \mathbb{R}^q :

$$p(\mathbf{x}|\mathbf{t}) = \mathcal{N}(\mathbf{W}\mathbf{t} + \boldsymbol{\mu}, \sigma^2\mathbf{I})$$

where $\boldsymbol{\mu}$ is a d -dimensional vector and \mathbf{W} is a $d \times q$ linear transform matrix. The latent variable is also

assumed to be Gaussian with $p(\mathbf{t}) = \mathcal{N}(\mathbf{0}, \mathbf{I})$. Then the marginal distribution of \mathbf{x} is also Gaussian:

$$p(\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}, \mathbf{C})$$

with a covariance matrix $\mathbf{C} = \mathbf{W}\mathbf{W}^T + \sigma^2\mathbf{I}$. The posterior distribution can be shown to be Gaussian as well:

$$p(\mathbf{t}|\mathbf{x}) = \mathcal{N}(\mathbf{M}^{-1}\mathbf{W}^T(\mathbf{x} - \boldsymbol{\mu}), \sigma^2\mathbf{M}^{-1})$$

where $\mathbf{M} = \mathbf{W}^T\mathbf{W} + \sigma^2\mathbf{I}$ is a $q \times q$ matrix.

The PPCA mixture model is a combination of multiple PPCAs. Each PPCA component explains local data structure or cell subpopulation. The model is defined by the collection of each component parameters $\theta_k = \{\pi_k, \boldsymbol{\mu}_k, \mathbf{W}_k, \sigma_k^2\}$. From a flow cytometry dataset $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$, an EM algorithm can learn the mixture model by iteratively computing these parameters. More details on the PPCA mixture and the EM algorithm are explained in [16]

The mixture of PPCA offers a way of controlling the number of parameters to be estimated without completely sacrificing the flexibility of model. In mixture model framework, a more common choice is the standard Gaussian mixture model. In the Gaussian mixture model, however, each Gaussian component requires $d(d+1)/2$ covariance parameters to be estimated if a full covariance matrix is used. Thus, as the data dimension increases, more data points are needed for reliable estimation of those parameters. The number of parameters can be reduced by constraining the covariance matrix to be isotropic or diagonal. These are too restrictive, however, since an isotropic or diagonal covariance makes the Gaussian component spherical or, respectively, elliptical aligned along the data axes; hence, the correlation structure between variables cannot be captured. On the other hand, the PPCA mixture model lies between those two extremes, and allows to control the number of parameters by specifying q , the dimension of the latent variable.

B. Mixture of PPCA with Missing Data

Even though our file matching problem has a particular pattern of missing variables, we develop a more general algorithm that allows for an arbitrary pattern of missing variables. Our development assumes values are “missing at random,” meaning that whether a variable is missing or not, is independent of its value [13]. Our algorithm may be viewed as an extension of the algorithm of [17] to PPCA, or the algorithm of [16] to data with missing values.

Denoting the observed and missing components by o_n and m_n , each data point can be divided $\mathbf{x}_n = (\mathbf{x}_n^{o_n}, \mathbf{x}_n^{m_n})$. In a missing data problem, a set of partial observations $\{\mathbf{x}_1^{o_1}, \dots, \mathbf{x}_N^{o_N}\}$ is given. Similar

to EM algorithms for Gaussian mixture models, we introduce indicator variables \mathbf{z}_n . One and only one entry of \mathbf{z}_n is nonzero, and $z_{nk} = 1$ indicates that the k th component is responsible for generating \mathbf{x}_n . We also include the missing components $\mathbf{x}_n^{m_n}$ and the set of latent variables \mathbf{t}_{nk} for each component to form the complete data $(\mathbf{x}_n^o, \mathbf{x}_n^{m_n}, \mathbf{t}_{nk}, \mathbf{z}_n)$ for $n = 1, \dots, N$ and $k = 1, \dots, K$.

We derive an iterative EM algorithm for the PPCA mixture model with missing data. The key difference from the EM algorithm for completely observed data is that the conditional expectation is taken with respect to \mathbf{x}^o as opposed to \mathbf{x} in the expectation steps.

To develop an EM algorithm, we employ and extend the two step procedure as described in [16]. In the first stage of the algorithm, the component weights π_k and the component center $\boldsymbol{\mu}_k$ are updated:

$$\hat{\pi}_k = \frac{1}{N} \sum_n \langle z_{nk} \rangle, \quad (1)$$

$$\hat{\boldsymbol{\mu}}_k = \frac{\sum_n \langle z_{nk} \rangle \begin{bmatrix} \mathbf{x}_n^{o_n} \\ \langle \mathbf{x}_n^{m_n} \rangle \end{bmatrix}}{\sum_n \langle z_{nk} \rangle} \quad (2)$$

where $\langle z_{nk} \rangle = P(z_{nk} = 1 | \mathbf{x}_n^{o_n})$ is the responsibility of mixture component k for generating the unit \mathbf{x}_n , and $\langle \mathbf{x}_n^{m_n} \rangle = \mathbb{E}[\mathbf{x}_n^{m_n} | z_{nk} = 1, \mathbf{x}_n^{o_n}]$ is the conditional expectation. Note that we are not assuming the vectors in the bracket are stackable. This notation can be replaced by the true component ordering without difficulty.

In the second stage, we update \mathbf{W}_k and σ_k^2 :

$$\widehat{\mathbf{W}}_k = \mathbf{S}_k \mathbf{W}_k (\sigma_k^2 \mathbf{I} + \mathbf{M}_k^{-1} \mathbf{W}_k^T \mathbf{S}_k \mathbf{W}_k)^{-1}, \quad (3)$$

$$\hat{\sigma}_k^2 = \frac{1}{d} \text{tr} \left(\mathbf{S}_k - \mathbf{S}_k \mathbf{W}_k \mathbf{M}_k^{-1} \widehat{\mathbf{W}}_k^T \right) \quad (4)$$

from local covariance matrix \mathbf{S}_k :

$$\mathbf{S}_k = \frac{1}{N \hat{\pi}_k} \sum_n \langle z_{nk} \rangle \langle (\mathbf{x}_n - \hat{\boldsymbol{\mu}}_k)(\mathbf{x}_n - \hat{\boldsymbol{\mu}}_k)^T \rangle.$$

These update rules boil down to the update rules for completely observed data when there are no missing variables. We derive the EM algorithm in detail in Appendix A.

After model parameters are estimated, the observations are divided into groups according to their posterior distribution:

$$\arg \max_{k=1, \dots, K} p(z_{nk} = 1 | \mathbf{x}_n^{o_n}),$$

so each unit (cell) is classified into one of K cell populations. Note that this posterior probability is computed in the E-step.

Cell Type		CD markers
granulocyte		CD45+, CD15+
monocyte		CD45+, CD14+
lymphocyte	helper T cell	CD45+, CD3+
	cytotoxic T cell	CD45+, CD3+, CD8+
	B cell	CD45+, CD19+ or CD45+, CD20+
	NK cell	CD16+, CD56+, CD3-

Fig. 8. Types of white blood cells. Each cell type is characterized by a set of expressed CD markers. The cluster of differentiation (CD) markers are commonly used to identify cell surface molecules on white blood cells. The ‘+/-’ signs indicate whether a certain cell type has corresponding antigens on the cell surface.

C. Domain Knowledge and Initialization of EM algorithm

In file matching of flow cytometry data, domain knowledge is critical. First, as explained above, the incompletely observed data is insufficient to determine the correct cluster labeling. Second, the initial conditions of the EM algorithm affect its performance and convergence rate. Domain knowledge allows us to choose the number of components, and to initialize the algorithm so that it converges to the best local maximum.

In flow cytometry, from the design of fluorochrome marker combinations and the knowledge about the blood sample composition, we can anticipate certain properties of cell subpopulations. For example, Fig. 8 summarizes white blood cell types and their characteristic cluster of differentiation (CD) marker expressions. That these are six cell types suggests choosing $K = 6$ when analyzing white blood cells.

The CD markers indicated are commonly used in flow cytometry to identify cell surface molecules on leukocytes [18]. However, this information is qualitative, and needs to be quantified.

To achieve this, we use one dimensional histograms. In a histogram, two large peaks are generally expected depending on the expression level of the corresponding CD marker. If a cell subpopulation expresses a CD marker, denoted by ‘+’, then it forms a peak on the right side of the histogram. On the other hand, if a cell population does not express the marker, denoted by ‘-’, then a peak can be found on the left side of the histogram. We use the locations of the peaks to quantify the expression levels.

These quantified values can be combined with the CD marker expression levels of each cell type to specify the initial cluster centers. Thus, each component of μ_k of a certain cell type is initialized by either the positive quantity or the negative quantity from the histogram. In our implementation, these are set manually based on visual inspection of histograms. Then we initialize the mixture model parameters $\{\pi_k, \mu_k, \mathbf{W}_k, \sigma_k^2\}$ as described in Fig. 9.

Input: $\mathcal{X}_1, \mathcal{X}_2$ data files ; K the number of components ; q the dimension of latent variable space ; $\boldsymbol{\mu}_k$ for initial component mean.

for $k = 1$ to K **do**

1. using distance $\|\mathbf{x}_n^{O_n} - \boldsymbol{\mu}_k^{O_n}\|$, find the set of data points \mathcal{X}^k whose nearest component mean is $\boldsymbol{\mu}_k$
2. initialize a covariance matrix \mathbf{C}_k with random entries
3. replace submatrices of \mathbf{C}_k with sample covariance of data points in \mathcal{X}^k
4. make \mathbf{C}_k positive definite by enforcing the eigenvalues to be positive
5. set $\pi_k = \frac{|\mathcal{X}^k|}{N_1 + N_2}$
6. set \mathbf{W}_k with the q principal eigenvectors of \mathbf{C}_k
7. set σ_k^2 with the average of remaining eigenvalues of \mathbf{C}_k

end for

Output: $\{\pi_k, \boldsymbol{\mu}_k, \mathbf{W}_k, \sigma_k^2\}$ for $k = 1, \dots, K$

Fig. 9. Parameter initialization of an EM algorithm for missing data. Cell populations are partitioned into K groups based on the distance to each component center. The component weight π_k is initialized according to the size of each partition. From the covariance matrix estimate $b\mathbf{C}_k$, parameters $b\mathbf{W}_k$ and σ_k^2 are initialized by taking eigen-decomposition.

	c	s_1	s_2
c			
s_1			
s_2			

Fig. 10. Structure of covariance matrix \mathbf{C} . The sub-matrices $\mathbf{C}_k^{s_1, s_2}$ and $\mathbf{C}_k^{s_2, s_1}$ cannot be estimated from a sample covariance matrix because these variables are never jointly observed.

An important issue in file matching arises from the covariance matrix. When data is completely observed, a common way of initialization of a covariance matrix is using a sample covariance matrix. In the case of file matching, however, it cannot be evaluated since some sets of variables are never jointly observed (see Fig. 10). We chose to build a covariance matrix \mathbf{C}_k from variable to variable with sample covariances. For example, we can set \mathbf{C}_k^{c, s_1} with the sample covariance for variables c and s_1 based on cases for which both variables c and s_1 are present. On the other hand, the submatrix $\mathbf{C}_k^{s_1, s_2}$ cannot be built based on the observation. In our implementation, we set the submatrix $\mathbf{C}_k^{s_1, s_2}$ with arbitrary

ID	N_1	N_2	N_e
Patient1	10000	10000	5223
Patient2	7000	7000	4408
Patient3	3000	3000	3190

Fig. 11. Three flow cytometry datasets from three different patients. Each dataset is divided into two data files and an evaluation set. N_1 and N_2 denote the size of two data files and N_e is the size of evaluation set.




Fig. 12. File structure used in the experiment. FS, SS, and CD56 are common in both files, and a pair of CD markers are observed in only one of the files. The blank blocks correspond to the unobserved variables. The blocks in file 1 are matrices with N_1 rows, and the blocks in file 2 are matrices with N_2 rows.

values. However, the resulting matrix may not be positive definite. Thus, \mathbf{C}_k is made positive definite by replacing negative eigenvalues with a small positive value. Once a covariance matrix \mathbf{C}_k is obtained, we can initialize \mathbf{W}_k and σ_k^2 by taking eigen-decomposition of \mathbf{C}_k .

V. EXPERIMENTS AND RESULTS

We apply the proposed file matching technique to real flow cytometry datasets, and present experimental results.

Three flow cytometry datasets are prepared from lymph node samples of three patients. These datasets were provided by the Department of Pathology at the University of Michigan. The measurements are of different sizes and have seven attributes: FS, SS, CD56, CD16, CD3, CD8, and CD4. Each dataset is randomly permuted ten times and divided into two data files and a separate evaluation set. In Fig. 11, the cell counts of the two files and the held-out set are denoted N_1 , N_2 , and N_e , respectively. Two attributes from each file are made hidden to construct hypothetical files with missing data. Thus, CD16 and CD3 are available only in file 1, and CD8 and CD4 are available only in file 2, while FS, SS, and CD56 are common. The pattern of the constructed data files is illustrated in Fig. 12 where the blocks of missing variables are left blank.

For each white blood cell type, its expected marker expressions (CD markers), relative size (FS), and relative granularity (SS) are presented in Fig. 13. Because it is from a lymph node sample, the majority of cell population is lymphocytes, while the most common white blood cells in a human body

Cell type	FS	SS	CD56	CD16	CD3	CD8	CD4
granulocyte	+	+	-	+	-	-	-
monocyte	+	-	-	+	-	-	-
helper T cell	-	-	-	-	+	-	+
cytotoxic T cell	-	-	-	-	+	+	-
B lymphocyte	-	-	-	-	-	-	-
Natural Killer cell	-	-	+	+	-	-	-

Fig. 13. Cell types in the dataset and their corresponding marker expressions. ‘+’ or ‘-’ indicates whether a certain cell type expresses the CD marker or not.

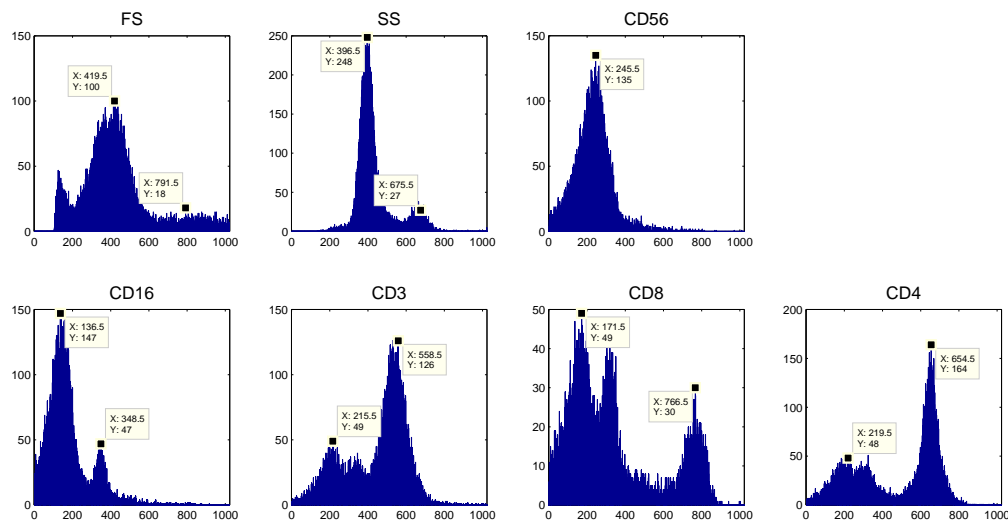


Fig. 14. Histogram of each marker in the dataset. The peaks are hand-picked and are indicated in each panels.

are granulocytes. The ‘+/-’ signs indicate whether a certain cell type expresses the markers or not. For example, helper T cells express both CD3 and CD4 but not others. This qualitative knowledge is quantified with the help of single dimensional histograms as explained in Section IV-C. Two dominant peaks are picked from each histogram and their corresponding measurement values are set to the positive and negative expression levels. Fig. 14 and Fig. 15 summarize this histogram analysis.

Two incomplete data files are completed following the procedure as described in Fig. 7. A mixture of PPCA is fitted with six components because six cell types are expected from this dataset. The latent variable dimension of each PPCA component is fixed to two.

	FS	SS	CD56	CD16	CD3	CD8	CD4
+	800	680	500	350	550	750	650
-	400	400	240	130	200	170	200

Fig. 15. The positive and negative expression levels are summarized.

The synthesized data after file matching is displayed in Fig. 5. The figure shows scatter plots of specific variables: CD16, CD3, CD4, and CD8. Note that these marker pairs are not jointly observed from the two incomplete data files. The imputation results from the NN and the Cluster-based NN methods are compared in the figure. For reference, scatter plots from the original complete dataset (ground truth) are also presented. As can be seen, the results from the Cluster-based NN are far more similar to the true distributions. On the other hand, the results from the NN method generates spurious clusters in the CD3-CD8 and CD3-CD4 scatter plots. In Fig. 5, these false clusters are indicated.

A. Evaluation method

To quantitatively evaluate the previous results, we use Kullback-Leibler (KL) divergence. The KL divergence between two distribution $f(\mathbf{x})$ and $g(\mathbf{x})$ is defined by

$$KL(g \parallel f) = \mathbb{E}_g [\log g - \log f].$$

Let f denote a true distribution responsible for the observations, and g denote its estimate.

The KL divergence is not symmetric, so $KL(f \parallel g)$ and $KL(g \parallel f)$ have different meanings. For a given distribution f , a distribution g minimizes $KL(f \parallel g)$ when g takes nonzero values in the region where f takes nonzero values; hence, it overestimates the support of f . On the other hand, $KL(g \parallel f)$ is minimized for g that is close to zero in the region where f is near zero. A distribution g that minimizes $KL(g \parallel f)$ tends to have smaller support. Therefore, $KL(g \parallel f)$ is a better evaluation method for detecting spurious clusters in an estimate.

Then the empirical estimate of the KL divergence is evaluated by

$$KL(g \parallel f) \approx KL(\hat{g} \parallel \hat{f}) \approx \frac{1}{N_e} \sum_{n=1}^{N_e} \left[\log \hat{g}(\hat{\mathbf{x}}_n) - \log \hat{f}(\hat{\mathbf{x}}_n) \right].$$

where the distributions f and g are replaced by their corresponding density estimates, and the expectation is approximated by a finite sum over imputed results $\hat{\mathbf{x}}_n$ on the held-out validation set of size N_e .

We used kernel density estimation on the ground truth data and the imputed data for \hat{f} and \hat{g} , respectively. The KL divergences are computed for ten random permutations, and their averages and

ID	NN (file 1)	Cluster-NN (file 1)	NN (file 2)	Cluster-NN (file 2)
Patient1	2.90 ± 0.05	1.55 ± 0.05	2.66 ± 0.03	1.12 ± 0.04
Patient2	4.54 ± 0.07	1.22 ± 0.03	4.12 ± 0.08	0.92 ± 0.03
Patient3	4.46 ± 0.10	2.40 ± 0.11	4.18 ± 0.11	2.30 ± 0.07

Fig. 16. The KL divergences are computed for ten permutations of each flow cytometry dataset. The averages and standard errors are reported in the table. For both the NN and Cluster-based NN algorithm, the file matching results are evaluated. The KL divergences of Cluster-based NN are closer to zero than those of NN. Thus the results from Cluster-based NN better replicated the true distribution.

standard errors are reported in Fig. 16. As can be seen, the KL divergences from Cluster-based NN are substantially smaller than those from NN. Therefore, the Cluster-based NN yields a better replication of true distribution.

VI. DISCUSSION

In this paper, we demonstrated the use of a cluster-based nearest neighbor imputation method for file matching in flow cytometry data. We applied the proposed algorithm on real flow cytometry data to generate a dataset of higher dimensions by merging two data files of lower dimensions. The resulting matched file can be used for visualization and high-dimensional analysis of cellular attributes.

While the presented imputation method focused on the case of two files, it can be generalized to more than two files. For each missing component of a recipient cell, we can find a donor cell among files that have the component of interest. We envision two extensions of the clustering-based imputation method. The first is training a PPCA mixture model on all the data files. This approach involves the entire data points for model fitting. The second method considers a pair of files at a time. In this approach, we first select a donor file in which the missing component of the recipient file is available. Then we apply method of this paper to the pair of files. This approach involves smaller number of data points in training, but mixture models of smaller dimensions need to be fitted multiple times. After training of a mixture model and clustering of each cell, the similarity between cells can be computed. The Euclidean distance on the projected space of commonly observed variables can be used to find the similarity under the constraint that both units should have the same cluster label. The missing components are then imputed from the donor.

Future research directions include finding ways of automatic prior information extraction. The construction of covariance matrices from incomplete dataset in the initialization of the EM algorithm is also an interesting problem. We expect that better covariance structure estimation will be helpful for better

replication of non-symmetric and non-elliptic cell populations in the imputed results.

A limitation of this work is that it has only been validated on lymphocyte data, where, for certain marker combinations, cell types tend to form relatively well-defined clusters. However, for other samples and marker combinations, clusters may be more elongated or less well-defined due to cells being at different stages of physiologic development. Future work may also consider more flexible models for clustering such data, and associated inference algorithms.

APPENDIX

Suppose that we are given an incomplete observation set. We can divide each unit \mathbf{x}_n as $\mathbf{x}_n = \begin{bmatrix} \mathbf{x}_n^{o_n} \\ \mathbf{x}_n^{m_n} \end{bmatrix}$ by separating the observed components and the missing components. Note that we do not assume that the observed variables are first, and the notation can be replaced by the actual ordering of components without difficulty.

In the PPCA mixture model, the probability distribution of \mathbf{x} is

$$p(\mathbf{x}) = \sum_{k=1}^K \pi_k p(\mathbf{x}|k).$$

where K is the number of components in the mixture and π_k is a mixing weight corresponding to the component density $p(\mathbf{x}|k)$. We estimate the set of unknown parameters $\theta = \{\pi_k, \boldsymbol{\mu}_k, \mathbf{W}_k, \sigma_k^2\}$ using an EM algorithm from the partial observations $\{\mathbf{x}_1^{o_1}, \dots, \mathbf{x}_N^{o_N}\}$.

To develop an EM algorithm, we introduce indicator variables $\mathbf{z}_n = (z_{n1}, \dots, z_{nK})$ for $n = 1, \dots, N$. One and only one entry of \mathbf{z}_n is nonzero, and $z_{nk} = 1$ indicates that the k th component is responsible for generating \mathbf{x}_n . We also include a set of the latent variables \mathbf{t}_{nk} for each component, and missing variables $\mathbf{x}_n^{m_n}$ to form the complete data $(\mathbf{x}_n^{o_n}, \mathbf{x}_n^{m_n}, \mathbf{t}_{nk}, \mathbf{z}_n)$ for $n = 1, \dots, N$ and $k = 1, \dots, K$. Then the corresponding complete data likelihood function has the form:

$$\begin{aligned} \mathcal{L}_C &= \sum_n \sum_k z_{nk} \ln [\pi_k p(\mathbf{x}_n, \mathbf{t}_{nk})] \\ &= \sum_n \sum_k z_{nk} \left[\ln \pi_k - \frac{d}{2} \ln \sigma_k^2 - \frac{1}{2\sigma_k^2} \text{tr} [(\mathbf{x}_n - \boldsymbol{\mu}_k)(\mathbf{x}_n - \boldsymbol{\mu}_k)^T] \right. \\ &\quad \left. + \frac{1}{\sigma_k^2} \text{tr} [(\mathbf{x}_n - \boldsymbol{\mu}_k) \mathbf{t}_{nk}^T \mathbf{W}_k^T] - \frac{1}{2\sigma_k^2} \text{tr} [\mathbf{W}_k^T \mathbf{W}_k \mathbf{t}_{nk} \mathbf{t}_{nk}^T] \right] \end{aligned}$$

where terms independent of the parameters are not included in the second equality. Instead of developing an EM algorithm directly on this likelihood function \mathcal{L}_C , we extend the strategy in [16] and build a two-stage EM algorithm, where each stage is a two-step process.

In the first stage of the two stage EM algorithm, we update the component weight π_k and the component mean $\boldsymbol{\mu}_k$. We form a complete data log-likelihood function with the component indicator variables \mathbf{z}_n and missing variables \mathbf{x}_n^m while ignoring the latent variables \mathbf{t}_{nk} . Then we have the following likelihood function:

$$\begin{aligned}\mathcal{L}_1 &= \sum_{n=1}^N \sum_{k=1}^K z_{nk} \ln[\pi_k p(\mathbf{x}_n^{o_n}, \mathbf{x}_n^{m_n} | k)] \\ &= \sum_n \sum_k z_{nk} \left[\ln \pi_k - \frac{1}{2} \ln |\mathbf{C}_k| - \frac{1}{2} \text{tr} [\mathbf{C}_k^{-1} (\mathbf{x}_n - \boldsymbol{\mu}_k)(\mathbf{x}_n - \boldsymbol{\mu}_k)^T] \right]\end{aligned}$$

where terms unrelated to the model parameters are omitted in the second line. We take the conditional expectation with respect to $p(\mathbf{z}_n, \mathbf{x}_n^{m_n} | \mathbf{x}_n^{o_n})$. Since the conditional probability factorizes as

$$p(\mathbf{z}_n, \mathbf{x}_n^{m_n} | \mathbf{x}_n^{o_n}) = p(\mathbf{z}_n | \mathbf{x}_n^{o_n}) p(\mathbf{x}_n^{m_n} | \mathbf{z}_n, \mathbf{x}_n^{o_n}),$$

the next conditional expectations follow

$$\begin{aligned}\langle z_{nk} \rangle &= p(k | \mathbf{x}_n^{o_n}) = \frac{\pi_k p(\mathbf{x}_n^{o_n} | k)}{\sum_{k'} \pi_{k'} p(\mathbf{x}_n^{o_n} | k')}, \\ \langle z_{nk} \mathbf{x}_n^{m_n} \rangle &= \langle z_{nk} \rangle \langle \mathbf{x}_n^{m_n} \rangle, \\ \langle \mathbf{x}_n^{m_n} \rangle &= \boldsymbol{\mu}_k^{m_n} + \mathbf{C}_k^{m_n o_n} \mathbf{C}_k^{o_n o_n}{}^{-1} (\mathbf{x}_n^{o_n} - \boldsymbol{\mu}_k^{o_n}), \\ \langle z_{nk} \mathbf{x}_n^{m_n} \mathbf{x}_n^{m_n T} \rangle &= \langle z_{nk} \rangle \langle \mathbf{x}_n^{m_n} \mathbf{x}_n^{m_n T} \rangle, \\ \langle \mathbf{x}_n^{m_n} \mathbf{x}_n^{m_n T} \rangle &= \mathbf{C}_k^{m_n m_n} - \mathbf{C}_k^{m_n o_n} \mathbf{C}_k^{o_n o_n}{}^{-1} \mathbf{C}_k^{o_n m_n} + \langle \mathbf{x}_n^{m_n} \rangle \langle \mathbf{x}_n^{m_n T} \rangle\end{aligned}$$

where $\langle \cdot \rangle$ denote the conditional expectation. Maximizing $\langle \mathcal{L}_1 \rangle$ with respect to π_k , using a Lagrange multiplier, and with respect to $\boldsymbol{\mu}_k$ give the parameter updates

$$\hat{\pi}_k = \frac{1}{N} \sum_n \langle z_{nk} \rangle, \quad (5)$$

$$\hat{\boldsymbol{\mu}}_k = \frac{\sum_n \langle z_{nk} \rangle \begin{bmatrix} \mathbf{x}_n^{o_n} \\ \langle \mathbf{x}_n^{m_n} \rangle \end{bmatrix}}{\sum_n \langle z_{nk} \rangle}. \quad (6)$$

In the second stage, we include the latent variable \mathbf{t}_{nk} as well to formulate the complete data log-likelihood function. The new values of $\hat{\pi}_k$ and $\hat{\boldsymbol{\mu}}_k$ are used in this step to compute sufficient statistics. Taking the conditional expectation on \mathcal{L}_C with respect to $p(\mathbf{z}_n, \mathbf{t}_{nk}, \mathbf{x}_n^{m_n} | \mathbf{x}_n^{o_n})$, we have

$$\begin{aligned}\langle \mathcal{L}_C \rangle &= \sum_n \sum_k \langle z_{nk} \rangle \left[\ln \hat{\pi}_k - \frac{d}{2} \ln \sigma_k^2 - \frac{1}{2\sigma_k^2} \text{tr} [\langle (\mathbf{x}_n - \hat{\boldsymbol{\mu}}_k)(\mathbf{x}_n - \hat{\boldsymbol{\mu}}_k)^T \rangle] \right] \\ &\quad + \frac{1}{\sigma_k^2} \text{tr} [\langle (\mathbf{x}_n - \hat{\boldsymbol{\mu}}_k) \mathbf{t}_{nk}^T \rangle \mathbf{W}_k^T] - \frac{1}{2\sigma_k^2} \text{tr} [\mathbf{W}_k^T \mathbf{W}_k \langle \mathbf{t}_{nk} \mathbf{t}_{nk}^T \rangle].\end{aligned}$$

Since the the conditional probability factorizes

$$p(\mathbf{z}_n, \mathbf{t}_{nk}, \mathbf{x}_n^{m_n} | \mathbf{x}_n^{o_n}) = p(\mathbf{z}_n | \mathbf{x}_n^{o_n}) p(\mathbf{x}_n^{m_n} | \mathbf{z}_n, \mathbf{x}_n^{o_n}) p(\mathbf{t}_{nk} | \mathbf{z}_n, \mathbf{x}_n^{o_n}, \mathbf{x}_n^{m_n}),$$

we can evaluate the conditional expectations as follows :

$$\begin{aligned} \langle (\mathbf{x}_n - \hat{\boldsymbol{\mu}}_k)(\mathbf{x}_n - \hat{\boldsymbol{\mu}}_k)^T \rangle &= \left(\begin{bmatrix} \mathbf{x}_n^{o_n} \\ \langle \mathbf{x}_n^{m_n} \rangle \end{bmatrix} - \hat{\boldsymbol{\mu}}_k \right) \left(\begin{bmatrix} \mathbf{x}_n^{o_n} \\ \langle \mathbf{x}_n^{m_n} \rangle \end{bmatrix} - \hat{\boldsymbol{\mu}}_k \right)^T + \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{Q}_{nk} \end{bmatrix}, \\ Q_{nk} &= \mathbf{C}_k^{m_n m_n} - \mathbf{C}_k^{m_n o_n} \mathbf{C}_k^{o_n o_n^{-1}} \mathbf{C}_k^{o_n m_n}, \\ \langle \mathbf{t}_{nk} \rangle &= \mathbf{M}_k^{-1} \mathbf{W}_k^T (\mathbf{x}_n - \hat{\boldsymbol{\mu}}_k), \\ \langle (\mathbf{x}_n - \hat{\boldsymbol{\mu}}_k) \mathbf{t}_{nk}^T \rangle &= \langle (\mathbf{x}_n - \hat{\boldsymbol{\mu}}_k)(\mathbf{x}_n - \hat{\boldsymbol{\mu}}_k)^T \rangle \mathbf{W}_k \mathbf{M}_k^{-1}, \\ \langle \mathbf{t}_{nk} \mathbf{t}_{nk}^T \rangle &= \sigma_k^2 \mathbf{M}_k^{-1} + \mathbf{M}_k^{-1} \mathbf{W}_k^T \langle (\mathbf{x}_n - \hat{\boldsymbol{\mu}}_k)(\mathbf{x}_n - \hat{\boldsymbol{\mu}}_k)^T \rangle \mathbf{W}_k \mathbf{M}_k^{-1}. \end{aligned}$$

Remember that the $q \times q$ matrix $\mathbf{M}_k = \mathbf{W}_k^T \mathbf{W}_k + \sigma_k^2 \mathbf{I}$. Then the maximization of $\langle \mathcal{L}_C \rangle$ with respect to \mathbf{W}_k and σ_k^2 leads to the parameter updates,

$$\widehat{\mathbf{W}}_k = \left[\sum_n \langle z_{nk} \rangle \langle (\mathbf{x}_n - \hat{\boldsymbol{\mu}}_k) \mathbf{t}_{nk}^T \rangle \right] \left[\sum_n \langle z_{nk} \rangle \langle \mathbf{t}_{nk} \mathbf{t}_{nk}^T \rangle \right]^{-1}, \quad (7)$$

$$\begin{aligned} \widehat{\sigma}_k^2 &= \frac{1}{d \sum_n \langle z_{nk} \rangle} \left[\sum_n \langle z_{nk} \rangle \text{tr} [\langle (\mathbf{x}_n - \hat{\boldsymbol{\mu}}_k)(\mathbf{x}_n - \hat{\boldsymbol{\mu}}_k)^T \rangle] \right. \\ &\quad - 2 \sum_n \langle z_{nk} \rangle \text{tr} [\langle (\mathbf{x}_n - \hat{\boldsymbol{\mu}}_k) \mathbf{t}_{nk}^T \rangle \mathbf{W}_k^T] \\ &\quad \left. + \sum_n \langle z_{nk} \rangle \text{tr} [\mathbf{W}_k^T \mathbf{W}_k \langle \mathbf{t}_{nk} \mathbf{t}_{nk}^T \rangle] \right]. \end{aligned} \quad (8)$$

Substituting the conditional expectations simplifies the M-step equations

$$\widehat{\mathbf{W}}_k = \mathbf{S}_k \mathbf{W}_k (\sigma_k^2 \mathbf{I} + \mathbf{M}_k^{-1} \mathbf{W}_k^T \mathbf{S}_k \mathbf{W}_k)^{-1}, \quad (9)$$

$$\widehat{\sigma}_k^2 = \frac{1}{d} \text{tr} \left(\mathbf{S}_k - \mathbf{S}_k \mathbf{W}_k \mathbf{M}_k^{-1} \widehat{\mathbf{W}}_k^T \right) \quad (10)$$

where

$$\mathbf{S}_k = \frac{1}{N \widehat{\pi}_k} \sum_n \langle z_{nk} \rangle \langle (\mathbf{x}_n - \hat{\boldsymbol{\mu}}_k)(\mathbf{x}_n - \hat{\boldsymbol{\mu}}_k)^T \rangle.$$

Each iteration of the EM algorithm updates the set of old parameters $\{\pi_k, \boldsymbol{\mu}_k, \mathbf{W}_k, \sigma_k^2\}$ with the set of new parameters $\{\widehat{\pi}_k, \widehat{\boldsymbol{\mu}}_k, \widehat{\mathbf{W}}_k, \widehat{\sigma}_k^2\}$ as in (5), (6), (9), and (10). The algorithm terminates when the value of the log-likelihood function no longer changes.

REFERENCES

- [1] H. Shapiro, *Practical Flow Cytometry*, 3rd ed. Wiley-Liss, 1994.
- [2] M. Brown and C. Wittwer, “Flow cytometry: Principles and clinical applications in hematology,” *Clinical Chemistry*, vol. 46, pp. 1221–1229, 2000.
- [3] C. E. Pedreira, E. S. Costa, S. Barrena, Q. Lecrevisse, J. Almeida, J. J. M. van Dongen, and A. Orfao, “Generation of flow cytometry data files with a potentially infinite number of dimensions,” *Cytometry A*, vol. 73A, pp. 834–846, 2008.
- [4] M. J. Boedigheimer and J. Ferbas, “Mixture modeling approach to flow cytometry data,” *Cytometry Part A*, vol. 73, pp. 421 – 429, 2008.
- [5] C. Chan, F. Feng, J. Ottinger, D. Foster, M. West, and T. Kepler, “Statistical mixture modeling for cell subtype identification in flow cytometry,” *Cytometry Part A*, vol. 73, pp. 693–701, 2008.
- [6] K. Lo, R. R. Brinkman, and R. Gottardo, “Automated gating of flow cytometry data via robust model-based clustering,” *Cytometry Part A*, vol. 73, pp. 321 – 332, 2008.
- [7] S. Pyne, X. Hu, K. Wang, E. Rossin, T.-I. Lin, L. M. Maier, C. Baecher-Allan, G. J. McLachlan, P. Tamayo, D. A. Hafler, P. L. D. Jager, and J. P. Mesirov, “Automated high-dimensional flow cytometric data analysis,” *PNAS*, vol. 106, pp. 8519–8524, 2009.
- [8] J. Lakoumentas, J. Drakos, M. Karakantza, G. C. Nikiforidis, and G. C. Sakellaropoulos, “Bayesian clustering of flow cytometry data for the diagnosis of b-chronic lymphocytic leukemia,” *Journal of Biomedical Informatics*, vol. 42, pp. 251–261, 2009.
- [9] K. Carter, R. Raich, W. Finn, and A. O. Hero, “Information preserving component analysis: data projections for flow cytometry analysis,” *Journ. of Selected Topics in Signal Processing*, vol. 3, pp. 148–158, 2009.
- [10] —, “Fine: Fisher information non-parametric embedding,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 2009, to appear.
- [11] S. Perfetto, P. K. Chattopadhyay, and M. Roederer, “Seventeen-colour flow cytometry: unravelling the immune system,” *Nature Reviews Immunology*, vol. 4, pp. 648–655, 2004.
- [12] M. Sánchez, J. Almeida, B. Vidriales, M. López-Berges, M. García-Marcos, M. Moro, A. Corrales, M. Calmuntia, J. S. Miguel, and A. Orfao, “Incidence of phenotypic aberrations in a series of 467 patients with b chronic lymphoproliferative disorders: basis for the design of specific four-color stainings to be used for minimal residual disease investigation,” *Leukemia*, vol. 16, pp. 1460–1469, 2002.
- [13] R. Little and D. Rubin, *Statistical analysis with missing data*, 2nd ed. Wiley, 2002.
- [14] S. Rässler, *Statistical matching: a frequentist theory, practical applications, and alternative Bayesian approaches*, ser. Lecture Notes in Statistics. Springer, 2002, no. 168.
- [15] M. Tipping and C. Bishop, “Probabilistic principal component analysis,” *Journal of the Royal Statistical Society, B*, vol. 6, pp. 611–622, 1999.
- [16] —, “Mixtures of probabilistic principal component analysis,” *Neural Computation*, vol. 11, pp. 443–482, 1999.
- [17] Z. Ghahramani and M. Jordan, “Supervised learning from incomplete data via an em approach,” *Advances in Neural Information Processing Systems*, vol. 6, 1994.
- [18] H. Zola, B. Swart, I. Nicholson, B. Aasted, A. Bensussan, L. Boumsell, C. Buckley, G. Clark, K. Drbal, P. Engel, D. Hart, V. Horejsi, C. Isacke, P. Macardle, F. Malavasi, D. Mason, D. Olive, A. Saalmueller, S. F. Schlossman, R. Schwartz-Albiez, P. Simmons, T. F. Tedder, M. Ugucioni, and H. Warren, “Cd molecules 2005: human cell differentiation molecules,” *Blood*, vol. 106, pp. 3123–3126, 2005.