

Learning from Ambiguously Labeled Face Images

Ching-Hui Chen, Vishal M. Patel, *Senior Member, IEEE*, Rama Chellappa, *Fellow, IEEE*

Abstract—Learning a classifier from ambiguously labeled face images is challenging since training images are not always explicitly-labeled. For instance, face images of two persons in a news photo are not explicitly labeled by their names in the caption. We propose a Matrix Completion for Ambiguity Resolution (MCar) method for predicting the actual labels from ambiguously labeled images. This step is followed by learning a standard supervised classifier from the disambiguated labels to classify new images. To prevent the majority labels from dominating the result of MCar, we generalize MCar to a weighted MCar (WMCar) that handles label imbalance. Since WMCar outputs a soft labeling vector of reduced ambiguity for each instance, we can iteratively refine it by feeding it as the input to WMCar. Nevertheless, such an iterative implementation can be affected by the noisy soft labeling vectors, and thus the performance may degrade. Our proposed Iterative Candidate Elimination (ICE) procedure makes the iterative ambiguity resolution possible by gradually eliminating a portion of least likely candidates in ambiguously labeled face. We further extend MCar to incorporate the labeling constraints between instances when such prior knowledge is available. Compared to existing methods, our approach demonstrates improvement on several ambiguously labeled datasets.

Index Terms—Ambiguous learning, labeling imbalance, iterative candidate elimination, matrix completion, low-rank matrix recovery.

1 INTRODUCTION

LEARNING a classifier for naming a face requires a large amount of labeled face images and videos. However, labeling face images is expensive and time-consuming due to significant amount of human efforts involved. As a result, brief descriptions such as tags, captions and screenplays accompanying the images and videos become important for training the classifiers. Although such information is publicly available, it is not as explicitly labeled as human annotations. For instance, names in the caption of a news photo provide possible candidates for faces appearing in the image [1], [2] (see Figure 1). The names in the screenplays are only weakly associated with faces in the shots [3]. The problem in which instead of a single label per instance, one is given a candidate set of labels, of which only one is correct is known as ambiguously labeled learning¹ [4], [5], [6], [7], [8].

In recent years, the problem of completing a low-rank matrix with missing entries has gained a lot of attention. In particular, matrix completion methods have been shown to produce good results for multi-label image classification problems [9], [10]. In these methods, the underlying assumption is that the concatenation of feature vectors and their labels produce a low-rank matrix. Our work is motivated by these works. The proposed method, Matrix Completion for Ambiguity Resolution (MCar), takes the heterogeneous feature matrix, which is the concatenation of the labeling matrix and feature matrix, as input. We first show that the heterogeneous feature matrix is ideally low-rank in the absence of noise. This in turn, allows us to convert the labeling problem as a matrix completion problem by pursuing the underlying low-rank matrix of the heterogeneous feature matrix. In contrast to multi-



Fig. 1: The names in the captions are not explicitly associated with the face images appeared in the news photo.

label learning, ambiguous labeling provides the clue that one of the labels in the candidate label set is the true label. This knowledge is utilized to regularize the labeling matrix in the heterogeneous feature matrix. This is essentially the main difference between our work and some of the previously proposed matrix completion techniques [9], [10].

Although ambiguous learning techniques can take advantage of large-scale and diverse ambiguously labeled data, most of the methods cannot properly handle the labeling imbalance that is often present in publicly available training data. For instances, celebrities and leading actors usually dominate (appear more frequently) in the candidate label sets, and these majority labels can easily bias the results of ambiguity resolving methods. As the proposed method relies on low-rank approximation of the heterogeneous feature matrix, heterogeneous feature vectors associated with those majority labels can dominate the process of low-rank approximation and thus bias the recovery of the

- C.-H. Chen and R. Chellappa are with the Department of Electrical and Computer Engineering, University of Maryland, College Park, MD, 20742 USA (e-mail: {ching, rama}@umiacs.umd.edu).
- V. M. Patel is with the Department of Electrical and Computer Engineering, Rutgers University, Piscataway, NJ, 08901 USA (e-mail: vishal.m.patel@rutgers.edu).

1. also known as partially labeled learning and superset label learning

labeling matrix. We propose the weighted MCar (WMCAR) to overcome the labeling imbalance in ambiguously labeled data. Unlike conventional instance weighting techniques [11] that assign unequal instance weight to the cost function of instances, WMCAR performs unequal column-wise weighting on the heterogeneous feature vectors. Therefore, a heterogeneous feature vector associated with majority labels will contribute less to the process of low-rank approximation than that associated with minority labels.

The column-wise weighting in WMCAR can be computed by estimating the groundtruth label distribution from the recovered labeling matrix, but the recovered labeling matrix is not accessible without applying WMCAR to resolve the ambiguity in the original labeling matrix. Nevertheless, iteratively updating the column-wise weighting and recovering the labeling matrix with WMCAR is not reliable (see iterative WMCAR in Figure 11). An explanation is that there is some unresolved ambiguity in the soft labeling matrix recovered by WMCAR. The remaining ambiguity (noise) can be detrimental to the iterative process as we iteratively update WMCAR by substituting the labeling matrix with the recovered one from the previous iteration. Hence, we propose the Iterative Candidate Elimination (ICE) procedure to iteratively eliminate the least likely candidates from a portion of the ambiguously labeled data. This procedure iteratively suppresses the noise in the recovered labeling matrix and thus yields a better performance in the next iteration of WMCAR. Although WMCAR with ICE is an iterative approach, it is fundamentally different from previously suggested iterative methods [7], [12], [13]. Unlike previous works that iteratively construct class-specific models and update the labels, the iterative process of ICE is effective in sequential noise suppression. Besides, WMCAR concatenates the labels and features as a heterogeneous matrix to recover the labels in each iteration. This ensures that the information in the ambiguously labeled data is used as a whole in recovering the true labels.

Moreover, we generalize MCar to include the labeling constraints between the instances for practical applications. For instances, two persons in a news photo should not be identified as the same subject even though both of them are ambiguously labeled in the caption. As shown by the recent success in low-rank matrix recovery [14], several prior works have developed robust methods for classification [15], [16]. The proposed method inherits the benefit of low-rank matrix recovery and possesses the capability to resolve the label ambiguity via low-rank approximation of the heterogeneous matrix. As a result, our method is more robust compared to some of the existing discriminative ambiguous learning methods [5], [17]. The disambiguated labels from MCar are then used to learn a supervised learning classifier, which can be used to classify new data.

In this paper, we make the following contributions:

1. We propose a matrix completion method where instances and their associated ambiguous labels are jointly considered for disambiguating the class labels.
2. We provide a geometric interpretation of the matrix completion framework from the perspective of recovering the potentially-separable convex hulls of each class.
3. We propose WMCAR to resolve the label ambiguity in the presence of labeling imbalance.
4. We propose the ICE approach to improve the reliability of iterative WMCAR. The integration of WMCAR and ICE is effective in resolving the ambiguity and outperforms WMCAR in general.
5. Our method can handle the group constraints between instances for practical applications.

In this paper, we generalize our prior work in [18] to overcome labeling imbalance in ambiguously labeled data. The ICE procedure and the experimental analysis are extensions to [18].

The rest of this paper is organized as follows. In Section 2 we review some related work on ambiguously labeled learning methods. Section 3 describes the proposed MCar and WMCAR. The optimization procedure for WMCAR is described in Section 4. Section 5 describes the ICE procedure in detail. Section 6 presents the extension of MCar for incorporating the constraint between instances. In Section 7, we demonstrate the results on synthesized as well as real-world ambiguously labeled datasets. Finally, Section 8 concludes this work with a brief summary and discussion.

Notations: We use the following notations in this paper. The matrix element $a_{i,j}$ denotes the entity in the i^{th} row and j^{th} column of matrix \mathbf{A} . $\mathbf{1}_n$ represents a column vector of size $n \times 1$ consisting of 1's as its entries. \mathbf{v}_i is the canonical vector corresponding to the 1-of-K coding of i . $\|\cdot\|_1$ and $\|\cdot\|_0$ denote the ℓ_1 norm and ℓ_0 norm, respectively. The Frobenius norm and the nuclear norm of \mathbf{A} are defined as $\|\mathbf{A}\|_F = \left(\sum_{i,j}(a_{i,j})^2\right)^{\frac{1}{2}}$ and $\|\mathbf{A}\|_* = \sum_i \sigma_i(\mathbf{A})$, respectively where σ_i is the i^{th} singular value of \mathbf{A} . $(\cdot)^T$ denotes transposition operation. $|S|$ returns the cardinality in set S . $\mathcal{S}_a[b] = \text{sgn}(b) \max(|b| - a, 0)$ is the shrinkage operator. The concatenation of matrix \mathbf{A} and \mathbf{B} is defined as $\begin{bmatrix} \mathbf{A} \\ \mathbf{B} \end{bmatrix} = [\mathbf{A}; \mathbf{B}]$.

2 RELATED WORK

Various methods have been proposed in the literature for dealing with ambiguously labeled data. Some of these methods propose Expectation Maximization (EM)-like approaches to alternately disambiguate the labels and learn a discriminative classifier [19], [20]. **Berg *et al.* [1] proposed an EM-like approach to alternately disambiguate the labels by maximizing the likelihood of label assignment and estimate the parameters for the appearance model and language model.** Non-parametric methods have also been used to resolve the ambiguity by leveraging the inductive bias of learning methods [4]. For the ambiguously labeled training data the actual loss of mislabeling is not explicit. As a result, it is difficult to learn an effective discriminative model. Cour *et al.* [5], [21] proposed the partial 0/1 loss function for ambiguous labeling, which is a tighter upper bound for the actual loss as compared to the 0/1 loss [22]. Subsequently, a discriminative classifier can be learned from the ambiguous labels by minimizing the partial 0/1 loss. Several works have improved the learning of partial labels with the modeling of partial loss [23], error-correcting output codes [24], and iterative label propagation [25]. Liu *et al.* [6] proposed to learn a conditional multinomial mixture model for predicting the actual label from ambiguous labels.

Several dictionary-based methods have also been proposed for handling partially labeled datasets [7], [13], [26]. In particular, an EM-like dictionary learning approach was proposed in [7], where a confidence matrix and dictionary are updated in alternating iterations. **Although several methods have been utilizing the EM-like framework with robust appearance models [1], [7], [12], [13], [26], these methods can be very sensitive to the initialization of the model and may suffer from suboptimal performance. On the other hand, our proposed framework unifies the ambiguity resolution and appearance modeling into a single matrix completion framework, and thus it is more effective in ambiguity resolution.**

Luo *et al.* [17] generalize the ambiguously labeled learning problem addressed in [5] from single instances to a group of instances. The ambiguous loss considers the association between the group of identities and the candidate label vectors. The pairwise constraint between the instances (e.g. unique appearance of a subject) is accounted for when generating the candidate label vectors. Furthermore, Zeng *et al.* [12] use a Partial Permutation Matrix (PPM) to associate the identities in a group with ambiguous labels. The pairwise constraint is encoded by restricting the structure of PPM. Assuming that instances of the same subject inferred by PPM can ideally form a low-rank matrix, the actual identity of an instance can be predicted by alternatively updating the low-rank subspace and PPM. Xiao *et al.* [27] associate the identities in a group from ambiguous labels by minimizing the summation of the discriminative affinities in a group, where the affinities are learned from the low-rank reconstruction coefficient matrix and the weak supervision of ambiguous labels.

Recently, learning from weak annotations of labeling imbalance has received significant attention [28], [29]. Chen *et al.* [30] have employed the part-versus-part decomposition [31] to overcome the data imbalance in multi-label learning. Charte *et al.* [32] propose several methods to resample the multi-label training data to compensate the imbalance level. Wu *et al.* [33] incorporate the class cardinality bound constraints to deal with class imbalance. Although several prior works have addressed the issue of imbalanced data in the context of multi-label learning, the labeling imbalance in ambiguously labeled data remains to be investigated. We propose to estimate the groundtruth label distribution from ambiguous labels. With the estimated groundtruth label distribution, the instance weight of WMCAR can be computed to deal with labeling imbalance.

3 THE PROPOSED FRAMEWORK

The ambiguously labeled data is denoted as $\mathcal{L} = \{(\mathbf{x}_j, L_j), j = 1, 2, \dots, N\}$, where N is the number of instances. There are c classes, and the class labels are denoted as $\mathcal{Y} = \{1, 2, \dots, c\}$. Note that \mathbf{x}_j is the feature vector of the j^{th} instance, and its candidate labeling set $L_j \subseteq \mathcal{Y}$ consists of candidate labels associated with the j^{th} instance. The true label of the j^{th} instance is $l_j \in L_j$. In other words, one of the labels in L_j is the true label of \mathbf{x}_j . The objective is to resolve the ambiguity in \mathcal{L} such that each predicted label \hat{l}_j of \mathbf{x}_j matches its true label l_j . We associate the candidate labeling set L_j with a soft labeling vector \mathbf{p}_j , where $p_{i,j}$ indicates the probability that instance j belongs to class i . This allows us to quantitatively assign the likelihood of each class the instance belongs to if such information is provided. Given the ambiguous label of the j^{th} instance, we assign each entry of \mathbf{p}_j as

$$\begin{cases} p_{i,j} \in (0, 1] & \text{if } i \in L_j, \\ p_{i,j} = 0 & \text{if } i \notin L_j, \end{cases} \quad j = 1, 2, \dots, N, \quad (1)$$

where $\sum_{i=1}^c p_{i,j} = 1$. Without any prior knowledge, we assume equal probability for each candidate label. Let $\mathbf{P} \in \mathbb{R}^{c \times N}$ denote the ambiguous labeling matrix with \mathbf{p}_j in its j^{th} column. With this, one can model the ambiguous labeling as

$$\mathbf{P}^0 = \mathbf{P} - \mathbf{E}_P, \quad (2)$$

where \mathbf{P}^0 and \mathbf{E}_P denote the true labeling matrix and the labeling noise, respectively. The j^{th} column vector of \mathbf{P}^0 is $\mathbf{p}_j^0 = \mathbf{v}_{l_j}$,

where \mathbf{v}_{l_j} is the canonical vector corresponding to the 1-of-K coding of its true label l_j .

Similarly, assuming that the feature vectors are corrupted by some noise or occlusion, the feature matrix \mathbf{X} with \mathbf{x}_j in its j^{th} column can be modeled as

$$\mathbf{X}^0 = \mathbf{X} - \mathbf{E}_X, \quad (3)$$

where $\mathbf{X} \in \mathbb{R}^{m \times N}$ consists of N feature vectors of dimension m , \mathbf{X}^0 represents the feature matrix in the absence of noise and \mathbf{E}_X accounts for the noise. Concatenating (2) and (3), we obtain a unified model of ambiguous labels and feature vectors, which can be expressed as

$$\begin{bmatrix} \mathbf{P}^0 \\ \mathbf{X}^0 \end{bmatrix} = \begin{bmatrix} \mathbf{P} \\ \mathbf{X} \end{bmatrix} - \begin{bmatrix} \mathbf{E}_P \\ \mathbf{E}_X \end{bmatrix}. \quad (4)$$

Let

$$\mathbf{H}_{obs} = \begin{bmatrix} \mathbf{P} \\ \mathbf{X} \end{bmatrix} \text{ and } \mathbf{E} = \begin{bmatrix} \mathbf{E}_P \\ \mathbf{E}_X \end{bmatrix} \quad (5)$$

denote the heterogeneous feature matrix and its noise, respectively. If we can show that \mathbf{H}_{obs} is a low-rank matrix in the absence of noise, then we can use matrix completion methods for resolving the ambiguity in labeling. In the following section, we investigate the low-rank property of \mathbf{H}_{obs} .

3.1 Exploiting the Rank of \mathbf{H}_{obs}

The column vectors of \mathbf{X}_0 can be partitioned into sets S_1, S_2, \dots, S_c based on their true labels. We assume that the elements of S_k form a convex hull C_k of n_k vertices. It is clear that $n_k \leq |S_k|$. The representative matrix of the k^{th} class, $\mathbf{D}_k \in \mathbb{R}^{m \times n_k}$, consists of vertices of C_k as its column vectors, and each column vector is treated as a representative of the k^{th} class. Therefore, according to the definition of a convex hull, a noise-free instance \mathbf{x}_j^0 from class k ($\mathbf{x}_j^0 \in C_k$) can be represented as

$$\mathbf{x}_j^0 = \mathbf{D}_k \mathbf{a}_{k,j}, \text{ where } \mathbf{a}_{k,j}^T \mathbf{1}_{n_k} = 1, \mathbf{a}_{k,j} \in \mathbb{R}_+^{n_k \times 1}. \quad (6)$$

Note that $\mathbf{a}_{k,j} \in \mathbb{R}_+^{n_k \times 1}$ is the coefficient vector associated with the representative matrix of the k^{th} class. As the true label of an instance is not known in advance, we can represent \mathbf{x}_j^0 as

$$\begin{aligned} \mathbf{x}_j^0 &= \mathbf{D} \mathbf{q}_j, \quad \mathbf{D} = [\mathbf{D}_1 \ \mathbf{D}_2 \ \cdots \ \mathbf{D}_c], \\ \mathbf{q}_j &= [\mathbf{a}_{1,j}^T \ \mathbf{a}_{2,j}^T \ \cdots \ \mathbf{a}_{c,j}^T]^T, \quad \mathbf{q}_j^T \mathbf{1} = 1, \end{aligned} \quad (7)$$

where $\mathbf{D} \in \mathbb{R}^{m \times (\sum_{i=1}^c n_i)}$ is the collective representative matrix, and $\mathbf{q}_j \in \mathbb{R}_+^{(\sum_{i=1}^c n_i) \times 1}$ is the associated coefficient vector.

According to (7), we can decompose \mathbf{X}^0 as

$$\mathbf{X}^0 = \mathbf{D} \mathbf{Q}. \quad (8)$$

The coefficient matrix \mathbf{Q} in (8) is not unique as column vectors of \mathbf{D} are not necessarily linearly independent. However, we assume that an ideal decomposition $\mathbf{X}^0 = \mathbf{D} \mathbf{Q}^*$ satisfies the following condition

$$\begin{aligned} \mathbf{x}_j^0 &= \mathbf{D} \mathbf{q}_j^*, \text{ where } \mathbf{a}_{k,j}^{*T} \mathbf{1}_{n_k} = 1, \mathbf{x}_j^0 \in S_k, \\ &\mathbf{a}_{l,j}^{*T} \mathbf{1}_{n_l} = 0, \quad l \neq k, \end{aligned} \quad (9)$$

which implies that \mathbf{x}_j^0 is exclusively represented by \mathbf{D}_k even though it is possible that it can be written as a linear combination of any other vertices from different classes.

With this, we can recover the true labels from

$$\mathbf{P}^0 = \mathbf{T}\mathbf{Q}^*, \quad (10)$$

where $\mathbf{T} = [\mathbf{v}_1 \mathbf{1}_{n_1}^T \ \mathbf{v}_2 \mathbf{1}_{n_2}^T \ \dots \ \mathbf{v}_c \mathbf{1}_{n_c}^T]$ accumulates the coefficients associated with each matrix representative. Hence, the coefficient vector of dimension $\sum_{i=1}^c n_i$ is converted into labeling vector of dimension c . Concatenating $\mathbf{P}^0 = \mathbf{T}\mathbf{Q}^*$ and $\mathbf{X}^0 = \mathbf{D}\mathbf{Q}^*$, we further represent (4) as

$$\begin{bmatrix} \mathbf{P}^0 \\ \mathbf{X}^0 \end{bmatrix} = \begin{bmatrix} \mathbf{T} \\ \mathbf{D} \end{bmatrix} \mathbf{Q}^*. \quad (11)$$

It is clear that

$$\begin{aligned} \text{rank}([\mathbf{P}^0; \mathbf{X}^0]) &\leq \min(\text{rank}([\mathbf{T}; \mathbf{D}], \text{rank}(\mathbf{Q}^*)) \\ &\leq \min\left(c + m, \sum_{k=1}^c n_k, N\right). \end{aligned} \quad (12)$$

Since the representatives in \mathbf{D} only account for a subset of data samples, it is clear that $\sum_{k=1}^c n_k \leq N$. Therefore,

$$\text{rank}([\mathbf{P}^0; \mathbf{X}^0]) \leq \min\left(c + m, \sum_{k=1}^c n_k\right) \leq \sum_{k=1}^c n_k. \quad (13)$$

In the case of $N \gg \sum_{k=1}^c n_k$, the rank of $[\mathbf{P}^0; \mathbf{X}^0]$ is relatively smaller than N . From the above rank analysis and (4), we arrive at the following proposition:

Proposition 1. The heterogeneous feature matrix \mathbf{H}_{obs} is low-rank in the absence of noise.

Note that a similar result is also reported in [34] without making the convex hull assumption.

3.2 Matrix Completion for Ambiguity Resolution

According to (10), the true labeling matrix \mathbf{P}^0 can be recovered if \mathbf{D} and \mathbf{Q}^* are available. Nevertheless, obtaining \mathbf{D} and \mathbf{Q}^* based on the observed \mathbf{P} and \mathbf{X} is intractable by solving a matrix decomposition problem

$$\min_{\mathbf{T}, \mathbf{D}, \mathbf{Q}} \left\| \begin{bmatrix} \mathbf{P} \\ \mathbf{X} \end{bmatrix} - \begin{bmatrix} \mathbf{T} \\ \mathbf{D} \end{bmatrix} \mathbf{Q} \right\|_F^2, \quad (14)$$

subject to the conditions specified in (9)-(11). Following [9], we propose to resolve the ambiguity by recovering the underlying low-rank structure of the heterogeneous feature matrix. Hence, we transform the matrix decomposition problem to a matrix completion problem. For the ease of presentation, we start with solving a label assignment problem assuming that \mathbf{X} is noise-free, i.e. $\mathbf{X} = \mathbf{X}^0$. The predicted labeling matrix \mathbf{Y} can be estimated by solving the following rank minimization problem

$$\begin{aligned} \min_{\mathbf{Y}, \mathbf{E}_P} \text{rank} \left(\begin{bmatrix} \mathbf{Y} \\ \mathbf{X}^0 \end{bmatrix} \right) \\ \text{s.t.} \quad \begin{bmatrix} \mathbf{Y} \\ \mathbf{X}^0 \end{bmatrix} &= \begin{bmatrix} \mathbf{P} \\ \mathbf{X}^0 \end{bmatrix} - \begin{bmatrix} \mathbf{E}_P \\ \mathbf{0} \end{bmatrix}, \\ \mathbf{y}_j &\in \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_c\}, j = 1, 2, \dots, N, \\ y_{i,j} &= 0 \text{ if } i \notin L_j \ \forall j. \end{aligned} \quad (15)$$

The problem is to complete the labeling matrix \mathbf{Y} via pursuing a low-rank matrix $[\mathbf{Y}; \mathbf{X}^0]$ subject to the constraints given by the ambiguous labels. The first constraint defines the feasible region of label assignment and the second constraint implies that an instance can only be labeled among its candidate labels. We cannot guarantee that the optimal solution of (15) always yields

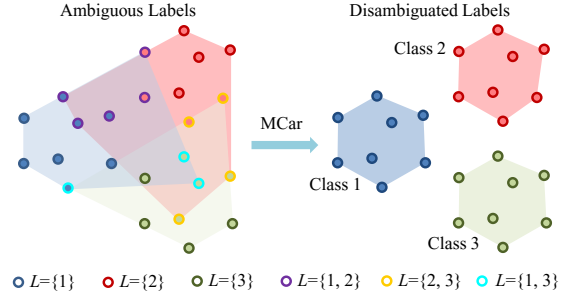


Fig. 2: MCar reassigns the labels for those ambiguously labeled instances such that instances of the same subjects cohesively form potentially-separable convex hulls. The vertices of each convex hull are the representatives of each class, forming \mathbf{D}_k . The interior and outline of the circles are color-coded to represent three different classes and various ambiguous labels, respectively.

a perfect recovery of ambiguous labeling such that $\mathbf{Y}^* = \mathbf{P}^0$. Several factors contribute to our inability to resolve the ambiguity. For instance, if label 1 is consistently present in the candidate labeling set of each instance, assigning \mathbf{v}_1 for each column vector of \mathbf{Y} yields a trivial solution. This issue is also addressed in [21], as learning from instances associated with two consistently co-occurring labels is impossible.

Note that $\mathbf{Y}^* = \mathbf{P}^0$ is one of the possible optimal solutions to (15). The solution may not be unique if any one of the instances belongs to more than one convex hull, i.e. the convex hulls from different classes overlap with each other. Hence, an instance can be ideally decomposed from either one of the convex hulls without further changing the rank of $[\mathbf{Y}; \mathbf{X}^0]$. This issue is analogous to the non-separable case of linear support vector machine (SVM). Nevertheless, it is our intention to seek $\mathbf{Y} = \mathbf{P}^0$ by solving (15) with the understanding that 1) the ambiguous labeling carries rational information, and 2) **the feature is sufficiently discriminative such that data lies in the feature subspace where convex hulls of each class are separable** [35].

Figure 2 illustrates the geometric interpretation of MCar using the convex hull representation. When each element in the candidate labeling set is trivially treated as the true label, the convex hulls of each class are erroneously expanded and the low-rank assumption of $[\mathbf{Y}; \mathbf{X}^0]$ does not hold. MCar exploits the underlying low-rank structure of $[\mathbf{Y}; \mathbf{X}^0]$, which is equivalent to reassigning the labels for those ambiguously labeled instances such that instances of the same class cohesively form a convex hull. Hence, each over-expanded convex hull shrinks to its actual contour, and the convex hulls become potentially separable. This is essentially different from discriminative ambiguous learning methods that construct the hyperplane between ambiguously labeled instances by minimizing the ambiguous loss.

When data is contaminated by sparse errors, the optimization problem in (15) can be reformulated as

$$\begin{aligned} \min_{\mathbf{H}, \mathbf{E}_X, \mathbf{E}_P} \text{rank}(\mathbf{H}) + \lambda \|\mathbf{E}_X\|_0 \\ \text{s.t.} \quad \mathbf{H} = \begin{bmatrix} \mathbf{Y} \\ \mathbf{Z} \end{bmatrix} = \begin{bmatrix} \mathbf{P} \\ \mathbf{X} \end{bmatrix} - \begin{bmatrix} \mathbf{E}_P \\ \mathbf{E}_X \end{bmatrix}, \\ \mathbf{y}_j \in \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_c\}, j = 1, 2, \dots, N, \\ y_{i,j} = 0 \text{ if } i \notin L_j \ \forall j, \end{aligned} \quad (16)$$

where \mathbf{H} is the heterogeneous feature matrix in the absence of noise, and \mathbf{Z} is the recovered feature matrix. The parameter

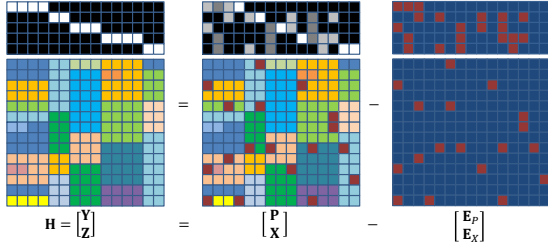


Fig. 3: Ideal decomposition of the heterogeneous feature matrix using MCar. The underlying low-rank structure and the ambiguous labeling are recovered simultaneously.

$\lambda \in \mathbb{R}_+$ controls the rank of \mathbf{H} and the sparsity of noise. The objective is to assign the predicted label \mathbf{Y} and extract the sparse noise of \mathbf{X} in pursuit of a low-rank \mathbf{H} . Figure 3 illustrates the ideal decomposition of the heterogeneous feature matrix, where the underlying low-rank structure and the ambiguous labels are recovered simultaneously.

As (16) is a combinatorial optimization problem, we relax each column vector of \mathbf{Y} in probability simplex in \mathbb{R}^c . The original formulation can be rewritten as

$$\begin{aligned} \min_{\mathbf{H}, \mathbf{E}_X, \mathbf{E}_P} \quad & \text{rank}(\mathbf{H}) + \lambda \|\mathbf{E}_X\|_0 + \gamma \|\mathbf{Y}\|_0 \\ \text{s.t.} \quad & \mathbf{H} = \begin{bmatrix} \mathbf{Y} \\ \mathbf{Z} \end{bmatrix} = \begin{bmatrix} \mathbf{P} \\ \mathbf{X} \end{bmatrix} - \begin{bmatrix} \mathbf{E}_P \\ \mathbf{E}_X \end{bmatrix}, \\ & \mathbf{1}_c^T \mathbf{Y} = \mathbf{1}_N^T, \quad \mathbf{Y} \in \mathbb{R}_+^{c \times N}, \\ & y_{i,j} = 0 \text{ if } i \notin L_j \quad \forall j, \end{aligned} \quad (17)$$

where $\gamma \in \mathbb{R}_+$ encourages the sparsity of \mathbf{Y} such that the original discrete feasible region can be well approximated. From the perspective of convex hull representation, such relaxation allows each instance to be represented from more than one set of representative matrix \mathbf{D}_k , while it will be penalized by the non-sparsity of \mathbf{Y} . Consequently, the predicted label of instance j can be obtained as

$$\hat{l}_j = \arg \max_{i \in L_j} y_{i,j}. \quad (18)$$

3.3 Ambiguously Labeled Data with Labeling Imbalance

The class imbalance may lead to performance degradation in SVM as a majority class with abundant training samples can bias the decision boundary toward a minority class with scarce training samples. Analogously, MCar may suffer from labeling imbalance when a majority label is frequently present among the candidate labels in the ambiguously labeled data. When we resolve the ambiguity using (17), the heterogeneous feature vectors associated with a majority label are more likely to dominate the low-rank approximation of the heterogeneous matrix than those associated with a minor label. Hence, the recovered soft labeling matrix will bias toward those soft labeling vectors associated with majority labels.

Class-weighted SVM applies unequal weighting to the cost function of different classes to mitigate the class imbalance [36]. Hence, instances from the minority label will be better emphasized than those from the dominant label to establish an objective decision boundary. However, the concept of class-weighted SVM cannot be directly applied to MCar to deal with label imbalances since each instance is not labeled as a particular class in the

ambiguously labeled data. Without the knowledge of the true labels, we formulate the instance-weighted objective function of (14) as

$$\min_{\mathbf{T}, \mathbf{D}, \mathbf{Q}} \sum_{j=1}^N \eta_j \left\| \begin{bmatrix} \mathbf{P}_j \\ \mathbf{x}_j \end{bmatrix} - \begin{bmatrix} \mathbf{T} \\ \mathbf{D} \end{bmatrix} \mathbf{q}_j \right\|_F^2, \quad (19)$$

where η_j is the instance weight of the j^{th} instance. In order to balance the square errors contributed by each class in (19), we aim to set instance weight η_j as $1/N_{l_j}$, where N_{l_j} is the number of the instances from the l_j class. Nevertheless, assigning a class weight for each instance is not feasible in the ambiguously labeled data since the true label l_j is not explicitly known. Moreover, N_i is intractable since the data is not explicitly labeled. Hence, we propose to set the instance weight as

$$\eta_j = \frac{1}{\sum_{i=1}^c p_{i,j} \hat{N}_i}, \quad (20)$$

where

$$\hat{N}_i = \sum_{j=1}^N p_{i,j} \quad (21)$$

is the estimated number of instances of the i^{th} class. The estimated number of instances of the i^{th} class accumulates the soft labeling scores corresponding to the i^{th} class across all the instances. With the soft labeling vector \mathbf{p}_j , we can compute the effective number of instances of the class that the j^{th} instance belongs to by $\sum_{i=1}^c p_{i,j} \hat{N}_i$. Hence, our proposed weighting scheme is eligible to compute the effective class weight of each ambiguously labeled instance even though the knowledge of true label is not available. The design of the instance weight is not unique, and readers may refer to [11], [37] for modeling the instance weight with respect to various objectives.

For the ease of presentation, we reformulate (19) as

$$\min_{\mathbf{T}, \mathbf{D}, \mathbf{Q}} \left\| \begin{bmatrix} \mathbf{P} \\ \mathbf{X} \end{bmatrix} \mathbf{W} - \begin{bmatrix} \mathbf{T} \\ \mathbf{D} \end{bmatrix} \mathbf{Q} \mathbf{W} \right\|_F^2, \quad (22)$$

where

$$\mathbf{W} = \sqrt{\text{diag}(\mathbf{1}_N^T \mathbf{P}^T \mathbf{P})}^{-1} \quad (23)$$

is a diagonal weighting matrix with $w_{j,j} = \sqrt{\eta_j}$. As post-multiplying \mathbf{W} does not increase the rank of a matrix, we claim that *Proposition 1* also applies to the weighted heterogeneous feature matrix $\mathbf{H}_{obs} \mathbf{W} = [\mathbf{P}; \mathbf{X}] \mathbf{W}$. We propose the weighted MCar (WMCAR) by generalizing (17) as

$$\begin{aligned} \min_{\mathbf{H}, \mathbf{E}_X, \mathbf{E}_P} \quad & \text{rank}(\mathbf{H} \mathbf{W}) + \lambda \|\mathbf{E}_X \mathbf{W}\|_0 + \gamma \|\mathbf{Y} \mathbf{W}\|_0 \\ \text{s.t.} \quad & \mathbf{H} \mathbf{W} = \begin{bmatrix} \mathbf{Y} \mathbf{W} \\ \mathbf{Z} \mathbf{W} \end{bmatrix} = \begin{bmatrix} \mathbf{P} \mathbf{W} \\ \mathbf{X} \mathbf{W} \end{bmatrix} - \begin{bmatrix} \mathbf{E}_P \mathbf{W} \\ \mathbf{E}_X \mathbf{W} \end{bmatrix}, \\ & \mathbf{1}_c^T \mathbf{Y} \mathbf{W} = \mathbf{1}_N^T \mathbf{W}, \quad \mathbf{Y} \mathbf{W} \in \mathbb{R}_+^{c \times N}, \\ & y_{i,j} = 0 \text{ if } i \notin L_j \quad \forall j. \end{aligned} \quad (24)$$

Let $\bar{\mathbf{H}}_{obs} = \mathbf{H}_{obs} \mathbf{W}$, $\bar{\mathbf{H}} = \mathbf{H} \mathbf{W}$, and $\bar{\mathbf{E}} = \mathbf{E} \mathbf{W}$, we reformulate (24) as

$$\begin{aligned} \min_{\bar{\mathbf{H}}, \bar{\mathbf{E}}_P, \bar{\mathbf{E}}_X} \quad & \text{rank}(\bar{\mathbf{H}}) + \lambda \|\bar{\mathbf{E}}_X\|_0 + \gamma \|\bar{\mathbf{Y}}\|_0 \\ \text{s.t.} \quad & \bar{\mathbf{H}} = \begin{bmatrix} \bar{\mathbf{Y}} \\ \bar{\mathbf{Z}} \end{bmatrix} = \begin{bmatrix} \bar{\mathbf{P}} \\ \bar{\mathbf{X}} \end{bmatrix} - \begin{bmatrix} \bar{\mathbf{E}}_P \\ \bar{\mathbf{E}}_X \end{bmatrix}, \\ & \mathbf{1}_c^T \bar{\mathbf{Y}} = \mathbf{1}_N^T \mathbf{W}, \quad \bar{\mathbf{Y}} \in \mathbb{R}_+^{c \times N}, \\ & \bar{y}_{i,j} = 0 \text{ if } i \notin L_j \quad \forall j. \end{aligned} \quad (25)$$

The predicted label can be retrieved from $\mathbf{Y} = \bar{\mathbf{Y}}\mathbf{W}^{-1}$ using (18). Interestingly, the instance-weighted MCar is equivalent to executing MCar with the weighted heterogeneous feature matrix. A larger weight on the heterogeneous feature vectors associated with minority labels provides those instances a stronger impact in the low-rank approximation of the heterogeneous matrix, and thus the labeling imbalance can be compensated. As (17) is generalized by (25) in consideration of labeling imbalance, WMCAR is identical to MCar in the special case of $\mathbf{W} = \mathbf{I}$.

Algorithm 1 The optimization algorithm for WMCAR (29)

Input: $\mathbf{P} \in \mathbb{R}^{c \times N}$, $\mathbf{X} \in \mathbb{R}^{m \times N}$, $\mathbf{W} \in \mathbb{R}^{N \times N}$, $L_j \forall j$, λ , and γ .

- 1: **Initialization:**
- 2: $\bar{\mathbf{P}} = \mathbf{P}\mathbf{W}$, $\bar{\mathbf{X}} = \mathbf{X}\mathbf{W}$, $\bar{\mathbf{H}}_{obs} = [\bar{\mathbf{P}}; \bar{\mathbf{X}}]$;
- 3: $\bar{\mathbf{Y}} = \mathbf{0}$, $\bar{\mathbf{Z}} = \mathbf{0}$, $\mu > 0$, $\mu_{max} > 0$, $\rho > 1$, $\Lambda = [\Lambda_P; \Lambda_X] = \bar{\mathbf{H}}_{obs} / \|\bar{\mathbf{H}}_{obs}\|_2$;
- 4: **while** not converged **do**
- 5: $\bar{\mathbf{E}}_P = \bar{\mathbf{P}} - \mathcal{S}_{\gamma\mu^{-1}}[\bar{\mathbf{Y}} - \mu^{-1}\Lambda_P]$;
- 6: $\bar{\mathbf{E}}_X = \mathcal{S}_{\lambda\mu^{-1}}[\bar{\mathbf{X}} - \bar{\mathbf{Z}} + \mu^{-1}\Lambda_X]$;
- 7: $(\mathbf{U}, \Sigma, \mathbf{V}) = \text{svd}(\bar{\mathbf{H}}_{obs} - \bar{\mathbf{E}} + \mu^{-1}\Lambda)$;
- 8: $\bar{\mathbf{H}} = \mathbf{U}\Sigma_{\mu^{-1}}[\Sigma]\mathbf{V}^T$;
- 9: $\Lambda = \Lambda + \mu(\bar{\mathbf{H}}_{obs} - \bar{\mathbf{H}} - \bar{\mathbf{E}})$;
- 10: $\mu = \min(\rho\mu, \mu_{max})$;
- 11: **Project** $\bar{\mathbf{Y}}$:
- 12: ▷ Line: 13: Projection for (31)
- 13: $\bar{y}_{i,j} = 0$ if $i \notin L_j \forall j$;
- 14: ▷ Line: 15-16: Projection for (30)
- 15: $\bar{\mathbf{Y}} = \max(\bar{\mathbf{Y}}, 0)$;
- 16: $\bar{\mathbf{y}}_j = w_{j,j} \bar{\mathbf{y}}_j / \|\bar{\mathbf{y}}_j\|_1, \forall j$;
- 17: **end while**
- 18: $\bar{\mathbf{H}} = \bar{\mathbf{H}}\mathbf{W}^{-1}$, $\bar{\mathbf{E}} = \bar{\mathbf{E}}\mathbf{W}^{-1}$

Output: $(\bar{\mathbf{H}}, \bar{\mathbf{E}})$

4 OPTIMIZATION

The augmented Lagrangian method (ALM) has been extensively used for solving low-rank problems [14], [38]. In this section, we propose to incorporate the ALM with the projection step [9], [10] to solve the optimization problem of WMCAR.

In order to decouple $\bar{\mathbf{Y}}$ in the first and third terms of the objective function in (25), we replace $\|\bar{\mathbf{Y}}\|_0$ with $\|\bar{\mathbf{P}} - \bar{\mathbf{E}}_P\|_0$ and rewrite (25) as

$$\begin{aligned} \min_{\bar{\mathbf{H}}, \bar{\mathbf{E}}_X, \bar{\mathbf{E}}_P} \quad & \text{rank}(\bar{\mathbf{H}}) + \lambda \|\bar{\mathbf{E}}_X\|_0 + \gamma \|\bar{\mathbf{P}} - \bar{\mathbf{E}}_P\|_0 \\ \text{s.t.} \quad & \bar{\mathbf{H}} = \begin{bmatrix} \bar{\mathbf{Y}} \\ \bar{\mathbf{Z}} \end{bmatrix} = \begin{bmatrix} \bar{\mathbf{P}} \\ \bar{\mathbf{X}} \end{bmatrix} - \begin{bmatrix} \bar{\mathbf{E}}_P \\ \bar{\mathbf{E}}_X \end{bmatrix}, \\ & \mathbf{1}_c^T \bar{\mathbf{Y}} = \mathbf{1}_N^T \mathbf{W}, \quad \bar{\mathbf{Y}} \in \mathbb{R}_+^{c \times N}, \\ & \bar{y}_{i,j} = 0 \text{ if } i \notin L_j \forall j. \end{aligned} \quad (26)$$

Following the procedure of ALM, we relax the first constraint in (26) and reformulate it as

$$\begin{aligned} \min_{\bar{\mathbf{H}}, \bar{\mathbf{E}}, \Lambda, \mu} \quad & \ell(\bar{\mathbf{H}}, \bar{\mathbf{E}}, \Lambda, \mu) \\ \text{s.t.} \quad & \mathbf{1}_c^T \bar{\mathbf{Y}} = \mathbf{1}_N^T \mathbf{W}, \quad \bar{\mathbf{Y}} \in \mathbb{R}_+^{c \times N}, \\ & \bar{y}_{i,j} = 0 \text{ if } i \notin L_j \forall j, \end{aligned} \quad (27)$$

where $\mu \in \mathbb{R}_+$ and $\Lambda \in \mathbb{R}^{(c+m) \times N}$. The Lagrangian is expressed as

$$\begin{aligned} \ell(\bar{\mathbf{H}}, \bar{\mathbf{E}}, \Lambda, \mu) = & \text{rank}(\bar{\mathbf{H}}) + \lambda \|\bar{\mathbf{E}}_X\|_0 + \gamma \|\bar{\mathbf{P}} - \bar{\mathbf{E}}_P\|_0 \\ & + \langle \Lambda, \bar{\mathbf{H}}_{obs} - \bar{\mathbf{H}} - \bar{\mathbf{E}} \rangle + \frac{\mu}{2} \|\bar{\mathbf{H}}_{obs} - \bar{\mathbf{H}} - \bar{\mathbf{E}}\|_F^2. \end{aligned} \quad (28)$$

In order to make the optimization problem feasible, we approximate the rank with the nuclear norm and the ℓ_0 norm with the ℓ_1 norm [39]. Thus, we solve the following formulation as the convex surrogate of (27)

$$\min_{\bar{\mathbf{H}}, \bar{\mathbf{E}}, \Lambda, \mu} \quad \ell_R(\bar{\mathbf{H}}, \bar{\mathbf{E}}, \Lambda, \mu) \quad (29)$$

$$\text{s.t.} \quad \mathbf{1}_c^T \bar{\mathbf{Y}} = \mathbf{1}_N^T \mathbf{W}, \quad \bar{\mathbf{Y}} \in \mathbb{R}_+^{c \times N}, \quad (30)$$

$$\bar{y}_{i,j} = 0 \text{ if } i \notin L_j \forall j, \quad (31)$$

where the Lagrangian is represented as

$$\begin{aligned} \ell_R(\bar{\mathbf{H}}, \bar{\mathbf{E}}, \Lambda, \mu) = & \|\bar{\mathbf{H}}\|_* + \lambda \|\bar{\mathbf{E}}_X\|_1 + \gamma \|\bar{\mathbf{P}} - \bar{\mathbf{E}}_P\|_1 \\ & + \langle \Lambda, \bar{\mathbf{H}}_{obs} - \bar{\mathbf{H}} - \bar{\mathbf{E}} \rangle + \frac{\mu}{2} \|\bar{\mathbf{H}}_{obs} - \bar{\mathbf{H}} - \bar{\mathbf{E}}\|_F^2. \end{aligned} \quad (32)$$

The ALM operates in the sense that $\bar{\mathbf{H}}$, $\bar{\mathbf{E}}_P$, and $\bar{\mathbf{E}}_X$ can be solved alternately by fixing other variables. In each iteration, we employ a similar projection technique used in [9], [10] to enforce $\bar{\mathbf{Y}}$ to be feasible. The entire procedure for solving (29) is summarized in Algorithm 1, and the details of the optimization algorithm are presented in the following paragraphs.

4.1 Solving for $\bar{\mathbf{E}}_P$

To update $\bar{\mathbf{E}}_P$, we fix $\bar{\mathbf{H}}$, $\bar{\mathbf{E}}_X$, Λ and μ obtained in the previous iteration. Hence, the problem for updating $\bar{\mathbf{E}}_P$ can be solved by first computing

$$\begin{aligned} \bar{\mathbf{E}}_P^* = \underset{\bar{\mathbf{E}}_P}{\text{argmin}} \quad & \gamma \|\bar{\mathbf{P}} - \bar{\mathbf{E}}_P\|_1 + \langle \Lambda_P, \bar{\mathbf{P}} - \bar{\mathbf{Y}} - \bar{\mathbf{E}}_P \rangle \\ & + \frac{\mu}{2} \|\bar{\mathbf{P}} - \bar{\mathbf{Y}} - \bar{\mathbf{E}}_P\|_F^2. \end{aligned} \quad (33)$$

For the ease of derivation, we let $\bar{\mathbf{B}} = \bar{\mathbf{P}} - \bar{\mathbf{E}}_P$ and update $\bar{\mathbf{B}}$ as surrogate. We can reformulate (33) as

$$\begin{aligned} \bar{\mathbf{B}}^* = \underset{\bar{\mathbf{B}}}{\text{argmin}} \quad & \gamma \|\bar{\mathbf{B}}\|_1 + \langle \Lambda_P, \bar{\mathbf{B}} - \bar{\mathbf{Y}} \rangle + \frac{\mu}{2} \|\bar{\mathbf{B}} - \bar{\mathbf{Y}}\|_F^2, \\ = \underset{\bar{\mathbf{B}}}{\text{argmin}} \quad & \gamma \|\bar{\mathbf{B}}\|_1 + \frac{\mu}{2} \|\bar{\mathbf{B}} - \bar{\mathbf{Y}} + \mu^{-1}\Lambda_P\|_F^2, \\ = \underset{\bar{\mathbf{B}}}{\text{argmin}} \quad & \gamma \|\bar{\mathbf{B}}\|_1 + \frac{\mu}{2} \|\bar{\mathbf{Y}} - \mu^{-1}\Lambda_P - \bar{\mathbf{B}}\|_F^2. \end{aligned} \quad (34)$$

Using the subgradient of (34), we can obtain the closed-form solution for updating $\bar{\mathbf{B}}$

$$\bar{\mathbf{B}}^* = \mathcal{S}_{\gamma\mu^{-1}}[\bar{\mathbf{Y}} - \mu^{-1}\Lambda_P]. \quad (35)$$

Consequently, we can update $\bar{\mathbf{E}}_P$ as

$$\bar{\mathbf{E}}_P^* = \bar{\mathbf{P}} - \bar{\mathbf{B}}^* = \bar{\mathbf{P}} - \mathcal{S}_{\gamma\mu^{-1}}[\bar{\mathbf{Y}} - \mu^{-1}\Lambda_P]. \quad (36)$$

4.2 Solve $\bar{\mathbf{E}}_X$

To update $\bar{\mathbf{E}}_X$, we fix $\bar{\mathbf{H}}$, $\bar{\mathbf{E}}_P$, Λ and μ obtained in the previous iteration. Thus, the problem for updating $\bar{\mathbf{E}}_X$ can be solved by

$$\begin{aligned} \bar{\mathbf{E}}_X^* = \underset{\bar{\mathbf{E}}_X}{\text{argmin}} \quad & \lambda \|\bar{\mathbf{E}}_X\|_1 + \langle \Lambda_X, \bar{\mathbf{X}} - \bar{\mathbf{Z}} - \bar{\mathbf{E}}_X \rangle \\ & + \frac{\mu}{2} \|\bar{\mathbf{X}} - \bar{\mathbf{Z}} - \bar{\mathbf{E}}_X\|_F^2, \\ = \underset{\bar{\mathbf{E}}_X}{\text{argmin}} \quad & \lambda \|\bar{\mathbf{E}}_X\|_1 + \frac{\mu}{2} \|\bar{\mathbf{X}} - \bar{\mathbf{Z}} + \mu^{-1}\Lambda_X - \bar{\mathbf{E}}_X\|_F^2. \end{aligned} \quad (37)$$

Using the subgradient of (37), we can obtain the closed-form solution for updating $\bar{\mathbf{E}}_X$

$$\bar{\mathbf{E}}_X^* = \mathcal{S}_{\lambda\mu^{-1}}[\bar{\mathbf{X}} - \bar{\mathbf{Z}} + \mu^{-1}\mathbf{\Lambda}_X]. \quad (38)$$

4.3 Solve $\bar{\mathbf{H}}$

To update $\bar{\mathbf{H}}$, we fix $\bar{\mathbf{E}}_P$, $\bar{\mathbf{E}}_X$, $\mathbf{\Lambda}$ and μ obtained in the previous iteration. The feasible region of $\bar{\mathbf{Y}}$ in $\bar{\mathbf{H}}$ is currently not considered but will be handled in the projection step of $\bar{\mathbf{Y}}$ (Section 4.4). Therefore, the problem for updating $\bar{\mathbf{H}}$ can be solved by

$$\bar{\mathbf{H}}^* = \operatorname{argmin}_{\bar{\mathbf{H}}} \|\bar{\mathbf{H}}\|_* + \langle \mathbf{\Lambda}, \bar{\mathbf{H}}_{obs} - \bar{\mathbf{H}} - \bar{\mathbf{E}} \rangle \quad (39)$$

$$+ \frac{\mu}{2} \|\bar{\mathbf{H}}_{obs} - \bar{\mathbf{H}} - \bar{\mathbf{E}}\|_F^2, \quad (40)$$

$$= \operatorname{argmin}_{\bar{\mathbf{H}}} \|\bar{\mathbf{H}}\|_* + \frac{\mu}{2} \|\mathbf{A}_H - \bar{\mathbf{H}}\|_F^2, \quad (41)$$

where $\mathbf{A}_H = \bar{\mathbf{H}}_{obs} - \bar{\mathbf{E}} + \mu^{-1}\mathbf{\Lambda}$. According to [40], the above problem can be solved by

$$\bar{\mathbf{H}}^* = \mathbf{U}\mathcal{S}_{\mu^{-1}}[\mathbf{\Sigma}]\mathbf{V}^T, \quad (42)$$

where $\mathbf{\Sigma}$ can be obtained from the singular value decomposition (SVD) of \mathbf{A}_H denoted as $(\mathbf{U}, \mathbf{\Sigma}, \mathbf{V}) = \operatorname{svd}(\mathbf{A}_H)$. Following the procedure in the augmented Lagrangian method (ALM), we can update $\mathbf{\Lambda}$ and μ as

$$\mathbf{\Lambda} = \mathbf{\Lambda} + \mu (\bar{\mathbf{H}}_{obs} - \bar{\mathbf{H}} - \bar{\mathbf{E}}), \quad (43)$$

where $\mu = \min(\rho\mu, \mu_{\max})$, in each iteration based on the updated $\bar{\mathbf{E}}_P$, $\bar{\mathbf{E}}_X$, and $\bar{\mathbf{H}}$.

4.4 Project $\bar{\mathbf{Y}}$

Since the SVD operation for solving $\bar{\mathbf{H}}$ does not always return a feasible $\bar{\mathbf{Y}}$, we use a projection technique similar to the one in [9], [10] to enforce $\bar{\mathbf{Y}}$ to be feasible in each iteration. The projection involves two steps. First, we enforce those entries of $\bar{\mathbf{Y}}$ that do not correspond to the candidate labels to be zeros since the actual label only comes from the candidate labeling set provided by the ambiguous labels. Second, each column vector of $\mathbf{Y} = \bar{\mathbf{Y}}\mathbf{W}^{-1}$ is constrained to be in the probability simplex. As a result, we replace those negative entries in $\bar{\mathbf{Y}}$ with zeros and then normalize each column $\bar{\mathbf{y}}_j$ so that the summation of the entries in $\bar{\mathbf{y}}_j$ is equal to $w_{j,j}$.

Algorithm 2 The algorithm for WMCAR-ICE

Input: $\mathbf{P} \in \mathbb{R}^{c \times N}$, $\mathbf{X} \in \mathbb{R}^{m \times N}$, $L_j \forall j$.

1: **while** $\mathcal{A} \neq \emptyset$ and within the maximum number of iterations **do**

2: $\mathbf{W} = \sqrt{\operatorname{diag}(\mathbf{1}_N^T \mathbf{P}^T \mathbf{P})}^{-1}$;

3: Obtain $\bar{\mathbf{Y}}$ using WMCAR (Algorithm 1);

4: Eliminate the least likely candidate in L_j , $j \in \mathcal{E}$ using (44)-(47);

5: \triangleright Line: 6-7: Project \mathbf{Y} to comply with $L_j, \forall j$

6: $y_{i,j} = 0$, if $i \notin L_j \forall j$;

7: $\mathbf{y}_j = \bar{\mathbf{y}}_j / \|\bar{\mathbf{y}}_j\|_1 \forall j$;

8: $\mathbf{P} \leftarrow \mathbf{Y}$;

9: **end while**

Output: (\mathbf{H}, \mathbf{E})

5 ITERATIVE CANDIDATE ELIMINATION FOR AMBIGUITY RESOLUTION

According to (23), the weighting matrix \mathbf{W} of WMCAR is a function of \mathbf{P} . As WMCAR resolves the label ambiguity in \mathbf{P} , the recovered soft labeling matrix \mathbf{Y} can provide a better estimate of \mathbf{W} than the original \mathbf{P} . This motivates us to iteratively resolve the ambiguity by alternating between recovering \mathbf{Y} and updating \mathbf{W} . Nevertheless, the performance of iterative WMCAR is not steady as shown in Figure 11. We propose WMCAR with ICE (WMCAR-ICE) to resolve the ambiguity by WMCAR and then remove the least likely candidate labels in each iteration. The least likely candidate label of the j^{th} instance is denoted as

$$m(j) = \operatorname{argmin}_{i \in L_j} y_{i,j}, \quad (44)$$

and its corresponding soft labeling score is denoted as $y_{m(j),j}$. As removing a candidate label, which is actually a true label, in the candidate set generates an irreversible error, we propose to iteratively remove a portion of the least likely candidate labels that have relatively low soft labeling scores than others.

Let \mathcal{A} denote the set consisting of the indices of those instances that have more than one candidate label, which is represented as

$$\mathcal{A} = \{j \mid |L_j| > 1, \forall j\}. \quad (45)$$

We define the elimination factor as f_e ($0 \leq f_e \leq 1$), which accounts for the proportion of instances in \mathcal{A} participating in the candidate elimination. We construct a subset \mathcal{E} of \mathcal{A} , which consists of the entries that correspond to the smallest f_e portion of $\{y_{m(j),j} \mid j \in \mathcal{A}\}$. We represent it as

$$\mathcal{E} = \{j \mid y_{m(j),j} \leq t, j \in \mathcal{A}\}. \quad (46)$$

Note that t is automatically determined such that $|\mathcal{E}| = \lceil f_e |\mathcal{A}| \rceil$. Hence, we can update the candidate labeling sets by

$$L_j \leftarrow L_j - \{m(j)\}, j \in \mathcal{E}. \quad (47)$$

We enforce the soft labeling matrix \mathbf{Y} to comply with the updated candidate labeling sets. We set $y_{i,j} = 0$, if $i \notin L_j \forall j$ and project each column vector of \mathbf{Y} in the probability simplex. The original \mathbf{P} will be replaced by \mathbf{Y} , which will serve as the input of WMCAR in the next iteration. The procedure of WMCAR-ICE is summarized in Algorithm 2. Note that updating the weighting matrix \mathbf{W} is an important step in WMCAR-ICE since it adaptively adjusts the importance among instances based on the updated \mathbf{Y} in the previous iteration. This ICE procedure can be utilized by other ambiguous learning techniques that adopt the soft labeling input/output similar to that of WMCAR.

6 LABELING CONSTRAINTS BETWEEN INSTANCES

In practical applications, several ambiguously labeled instances can appear in the same venue. As a result, pairwise relations between instances can be utilized to assist ambiguity resolution. For example, two persons in a news photo should not be identified as the same subject even though both of them are ambiguously labeled in the caption. Such prior knowledge can be easily incorporated by restricting the feasible region of the labeling matrix. Moreover, it is essential to handle the open set problem, where there are some instances whose identities never appear in the labels. These unrecognized instances can be treated as the null class.

In this section, we show how MCar’s formulation can be extended to associate the identities in news photos when the names are provided in captions. We assume all the instances (face images) are collected from the K groups (photos), and G_k is the set of indices of the instances (face images) appearing in the k^{th} group (photo). Note that instances (face images) from the same group (photo) share the same ambiguous labels provided by their associated caption. Without loss of generality, we assume that the c^{th} class corresponds to the null class. Considering the prior knowledge, the original formulation given in (17) can be reformulated as

$$\min_{\mathbf{H}, \mathbf{E}_X, \mathbf{E}_P} \text{rank}(\mathbf{H}) + \lambda \|\mathbf{E}_X\|_0 + \gamma \|\mathbf{Y}\|_0 \quad (48)$$

$$\text{s.t. } \mathbf{H} = \begin{bmatrix} \mathbf{Y} \\ \mathbf{Z} \end{bmatrix} = \begin{bmatrix} \mathbf{P} \\ \mathbf{X} \end{bmatrix} - \begin{bmatrix} \mathbf{E}_P \\ \mathbf{E}_X \end{bmatrix}, \quad (49)$$

$$\mathbf{1}_c^T \mathbf{Y} = \mathbf{1}_N^T, \quad \mathbf{Y} \in \mathbb{R}_+^{c \times N}, \quad (49)$$

$$y_{i,j} = 0 \text{ if } i \notin L_j, \quad i = 1, 2, \dots, c-1, \quad \forall j, \quad (50)$$

$$\sum_{j \in G_k} \sum_{i=1}^{c-1} y_{i,j} \geq 1 \text{ if } \bigcup_{j \in G_k} L_j \neq \{c\}, \quad \forall k, \quad (51)$$

$$\sum_{j \in G_k} y_{i,j} \leq 1, \quad i = 1, 2, \dots, c-1, \quad \forall k. \quad (52)$$

Constraints (49) and (50) are inherited from the original formulation. The constraint in (51), assumes that there is at least one non-null identity in a photo unless all the instances in a photo are explicitly labeled as null. This constraint is enforced to avoid the trivial solution that all the instances are treated as belonging to the null class. A similar constraint has been considered by [17] and [12] via restricting the candidate labeling set and confining the feasible space of PPM, respectively. The constraint in (52) enforces the uniqueness of non-null identities. Note that this framework can be easily tailored to handle other prior knowledge (e.g. must/cannot-link constraints, prior statistics) by regularizing the labeling matrix. This problem can be solved by following the similar relaxation procedures for solving (17). The optimization procedure is summarized in Algorithm 3.

Following the relaxation procedure in Section 4, we can reformulate (48) as

$$\min_{\mathbf{H}, \mathbf{E}, \mathbf{A}, \mu} \|\mathbf{H}\|_* + \lambda \|\mathbf{E}_X\|_1 + \gamma \|\mathbf{P} - \mathbf{E}_P\|_1 + \langle \mathbf{A}, \mathbf{H}_{obs} - \mathbf{H} - \mathbf{E} \rangle + \frac{\mu}{2} \|\mathbf{H}_{obs} - \mathbf{H} - \mathbf{E}\|_F^2, \quad (53)$$

$$\text{s.t. } \mathbf{1}_c^T \mathbf{Y} = \mathbf{1}_N^T, \quad \mathbf{Y} \in \mathbb{R}_+^{c \times N}, \quad (54)$$

$$y_{i,j} = 0 \text{ if } i \notin L_j, \quad i = 1, 2, \dots, c-1, \quad \forall j, \quad (55)$$

$$\sum_{j \in G_k} \sum_{i=1}^{c-1} y_{i,j} \geq 1 \text{ if } \bigcup_{j \in G_k} L_j \neq \{c\}, \quad \forall k, \quad (56)$$

$$\sum_{j \in G_k} y_{i,j} \leq 1, \quad i = 1, 2, \dots, c-1, \quad \forall k. \quad (57)$$

We use a similar procedure of Algorithm 1 presented in Section 4 to solve (53). We again use the projection method to guide the process of the matrix completion such that the constraints on \mathbf{Y} are satisfied. Additionally, the projection of \mathbf{Y} handles the group constraints such that the labeling constraints between instances are satisfied. Hence, we project \mathbf{Y} to the feasible regions indicated by (55), (56), and (57) one at a time, and each one is followed by the projection onto the feasible region indicated by (54) to ensure that

each column of \mathbf{Y} lies in the probability simplex. The detailed procedure is summarized in Algorithm 3. This algorithm can be easily extended to handle ambiguously labeled data with labeling imbalance by taking \mathbf{H} as input with proper manipulation on the projection steps of \mathbf{Y} .

Algorithm 3 The optimization algorithm for (53)

Input: $\mathbf{P} \in \mathbb{R}^{c \times N}$, $\mathbf{X} \in \mathbb{R}^{m \times N}$, $L_j \forall j$, $G_k \forall k$, λ , and γ .

- 1: **Initialization:** $\mathbf{Y} = \mathbf{0}$, $\mathbf{Z} = \mathbf{0}$, $\mu > 0$, $\mu_{\max} > 0$, $\rho > 1$,
 $\mathbf{A} = [\mathbf{A}_P; \mathbf{A}_X] = \mathbf{H}_{obs} / \|\mathbf{H}_{obs}\|_2$;
- 2: **while** not converged **do**
- 3: $\mathbf{E}_P = \mathbf{P} - \mathcal{S}_{\gamma\mu^{-1}}[\mathbf{Y} - \mu^{-1}\mathbf{A}_P]$;
- 4: $\mathbf{E}_X = \mathcal{S}_{\lambda\mu^{-1}}[\mathbf{X} - \mathbf{Z} + \mu^{-1}\mathbf{A}_X]$;
- 5: $(\mathbf{U}, \mathbf{\Sigma}, \mathbf{V}) = \text{svd}(\mathbf{H}_{obs} - \mathbf{E} + \mu^{-1}\mathbf{A})$;
- 6: $\mathbf{H} = \mathbf{U}\mathbf{\Sigma}\mu^{-1}[\mathbf{\Sigma}]\mathbf{V}^T$;
- 7: $\mathbf{A} = \mathbf{A} + \mu(\mathbf{H}_{obs} - \mathbf{H} - \mathbf{E})$;
- 8: $\mu = \min(\rho\mu, \mu_{\max})$;
- 9: **Project** \mathbf{Y} :
- 10: ▷ Line: 11-13: Projection for (55) and (54)
- 11: $\mathbf{Y} = \max(\mathbf{Y}, 0)$;
- 12: $y_{i,j} = 0$ if $i \notin L_j$, $i = 1, 2, \dots, c-1, \forall j$;
- 13: $\mathbf{y}_j = \mathbf{y}_j / \|\mathbf{y}_j\|_1, \forall j$;
- 14: ▷ Line: 15-22: Projection for (56) and (54)
- 15: **for** $k = 1 : K$ **do**
- 16: **if** $\bigcup_{j \in G_k} L_j \neq \{c\}$ **then**
- 17: **for** $i = 1 : c-1, j \in G_k$ **do**
- 18: $y_{i,j} = y_{i,j} / \min(\sum_{g \in G_k} \sum_{i=1}^{c-1} y_{i,g}, 1)$;
- 19: **end for**
- 20: **end if**
- 21: **end for**
- 22: $\mathbf{y}_j = \mathbf{y}_j / \|\mathbf{y}_j\|_1, \forall j$;
- 23: ▷ Line: 24-29: Projection for (57) and (54)
- 24: **for** $k = 1 : K$ **do**
- 25: **for** $i = 1 : c-1, j \in G_k$ **do**
- 26: $y_{i,j} = y_{i,j} / \max(\sum_{g \in G_k} y_{i,g}, 1)$;
- 27: **end for**
- 28: **end for**
- 29: $\mathbf{y}_j = \mathbf{y}_j / \|\mathbf{y}_j\|_1, \forall j$;
- 30: **end while**

Output: (\mathbf{H}, \mathbf{E})

7 EXPERIMENTAL RESULTS

We use the Labeled Faces in the Wild (LFW) dataset [41] and the CMU PIE dataset with synthesized ambiguous labels to evaluate the performance of our method under various controlled parameter settings. Furthermore, we use the *Lost* dataset [5] and the Labeled Yahoo! News dataset [1], [42] to demonstrate the effectiveness of our method in real-world applications. **For the LFW, CMU PIE, and Lost datasets**, we use face images in gray scale of range $[0, 1.0]$. Each instance is preprocessed with histogram equalization and converted into a column feature vector.

7.1 Parameters

It is interesting to observe that (16) becomes asymptotically similar to the formulation of Robust Principle Component Analysis (RPCA) [14] as the dimension of the data feature is far greater than the number of classes. Motivated by this fact, we fix λ as

$$\lambda_o = \frac{1}{\sqrt{\max(c+m, N)}}, \quad (58)$$

which is the tradeoff parameter suggested in RPCA. γ is a tuning parameter that controls the sparsity of the soft labeling vectors. For MCar-based methods, we use $\gamma = 2\lambda_0$ to encourage stronger sparsity of the labeling vector than that of feature noise. For ICE, we set the elimination factor f_e as 0.5, and set the maximum number of iterations as 5. These parameters yield good results in general, and we will investigate the sensitivity of parameters in Section 7.4.

7.2 Experiments with the Synthesized Dataset

We conduct two types of controlled experiments suggested in [21]. For the *inductive* experiment, the dataset is evenly split into ambiguously labeled training set and unlabeled testing set. The proposed methods, MCar/WMCar-SVM and WMCar-ICE-SVM, learn a multi-class linear SVM [43] with the disambiguated labels provided by MCar/WMCar and WMCar-ICE, respectively. The testing data is then classified using the learned classifier. For the *transductive* experiment, all the data is used as the ambiguously labeled training set.

We follow the ambiguity model defined in [21] to generate ambiguous labels in the controlled experiment. Note that α denotes the number of extra labels for each instance, and β represents the portion of the ambiguously labeled data among all the instances. The degree of ambiguity ϵ indicates the maximum probability that an extra label co-occurs with a true label, over all labels and instances. Each controlled experiment is repeated 20 times. We report the average testing (labeling) error rate for inductive (transductive) experiment, where the testing (labeling) error rate is the ratio of the number of erroneously labeled instances to the total number of instances in the testing (training) set. The standard deviations are plotted as error bars in the figures.

We compare the proposed MCar-based methods with several state-of-the-art ambiguous learning approaches for single instances with ambiguous labeling: CLPL [21], DLHD/DLSD [7], KDLS [13], and IPAL [25]. We report the performance of these methods when the experimental results are available in their papers. Otherwise, we use the configuration suggested in their papers to conduct the experiments. We use ‘naive’ [21] as the baseline method, which learns a classifier from minimizing the trivial 0/1 loss.

7.2.1 The LFW Dataset

The FIW(10b) dataset [21] consists of the top 10 most frequent subjects selected from the LFW dataset [41], and the first 50 face images of each subject are used for evaluation. We use the cropped and resized face images readily provided by the authors of [21], where the face images are of 45×55 pixels.

Figures 4a and 4b show the results of the inductive experiments. Figure 4a shows that the MCar-based methods significantly outperform all the other methods when the portion of ambiguously labeled data is larger than 0.2. The performance of WMCar is comparable to that of MCar since the ambiguously labeled data generated by this ambiguity model does not substantially result in labeling imbalance. WMCar-ICE demonstrates better performance than MCar and WMCar when more than 0.7 portions of the instances are ambiguously labeled. An explanation is that ICE eliminates the candidates based on the ordering of the least soft labeling score of each instance. This prioritization step can effectively benefit from a large portion of ambiguously labeled samples (i.e., large α) that usually carries a diverse aspect of soft

labeling scores. When the portion of ambiguously labeled samples is small, the improvement due to ICE becomes insignificant.

Figure 4b shows that MCar outperforms prior methods over various degrees of ambiguity except when $\epsilon > 0.7$. Thus, MCar yields improved performance at low and intermediate levels of ambiguity, but it becomes susceptible at high levels of ambiguity. One explanation is that both the true label and the extra labels of a subject will result in low-rank component of the labeling matrix when they are likely to co-occur in high degree of ambiguity. Consequently, separating the true label from the extra labels in MCar becomes challenging. Another explanation is that a high degree of ambiguity results in labeling imbalance, which causes the performance degradation of MCar. To verify this, we obtain the label distribution by counting the number of label occurrences in the candidate labeling sets for each class. We define the imbalance factor as the ratio of the maximum to the minimum value in the label distribution. The average imbalance factor varies from 1.33 to 3.58 as the degree of ambiguity increases. This confirms that WMCar outperforms MCar in high degree of ambiguity since WMCar is effective in mitigating the impact of labeling imbalance. Furthermore, WMCar-ICE outperforms WMCar by iteratively removing the least likely candidate labels from the candidate labeling sets. This experiment demonstrates that the labeling imbalance can cause performance degradation even though there is no class imbalance among the number of groundtruth faces per class.

In Figure 4c, MCar-based methods outperform the other approaches only when the number of extra labels is less than 5 in the transductive experiment. This shows that MCar-based methods cannot be effective when the labeling is severely cluttered such that the low-rank approximation of heterogeneous feature fails. Similar to the controlled parameter setting in Figure 4a, the ambiguously labeled data generated by this ambiguity model does not substantially result in labeling imbalance. Hence, the performance of WMCar is comparable to that of MCar, and WMCar-ICE slightly outperforms WMCar.

Figure 6a shows the intermediate results of low-rank decomposition of the feature matrix using MCar. Note that variations due to illumination, occlusions (e.g. eyeglasses, hand), and expressions are suppressed such that the low-rank component of a subject is preserved. In contrast to MCar-based methods, the discriminative methods (e.g. naive, CLPL) and IPAL are susceptible to such variations. Furthermore, it also demonstrates the robustness of our methods even though the face images are not perfectly aligned. The proposed method outperforms the dictionary-based methods [7], [13] for all cases except when there is severe ambiguity. Note the low-rank approximation of MCar operates on the feature matrix and ambiguous labeling matrix as a whole by concatenating them such that the actual labels and the low-rank component of feature matrix are recovered simultaneously. This essentially demonstrates the advantage of the proposed method over the DLHD/DLSD and KDLS methods that iteratively alternate between confidence and dictionary update.

7.2.2 The CMU PIE Dataset

The CMU PIE dataset contains face images from 68 subjects of different poses, illumination conditions, and expressions. Following the protocol presented in [7], we select the 18 subjects for evaluation. Each subject has 21 images under different illumination conditions, and the face images are resized to 40×48 pixels.

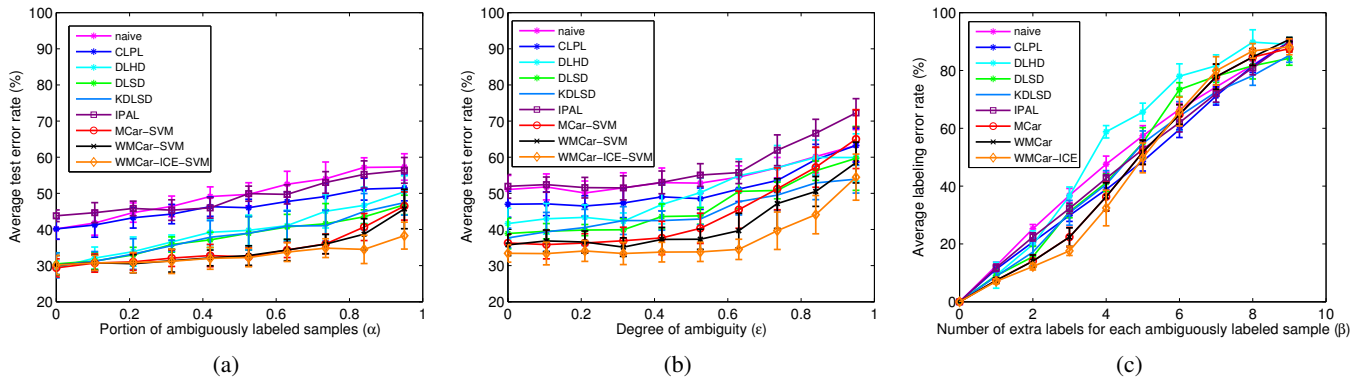


Fig. 4: Performance comparisons on the FIW(10b) dataset. (a) $\alpha \in [0, 0.95]$, $\beta = 2$, *inductive* experiment. (b) $\alpha = 1.0$, $\beta = 1$, $\epsilon \in [1/(c-1), 1]$, *inductive* experiment. (c) $\alpha = 1.0$, $\beta \in [0, 1, \dots, 9]$, *transductive* experiment.

We synthesize the ambiguous labels based on the controlled parameters. The results of two transductive experiments for CMU PIE dataset are shown in Figures 5a and 5b. In Figure 5a, MCar-based methods and IPAL recover all the label ambiguity for various portions of ambiguously labeled samples. In Figure 5b, our proposed methods consistently outperform most of the state-of-the-art methods except IPAL as we increase the number of extra labels for each ambiguously labeled sample. Since the CMU PIE dataset is collected in a constrained environment, the collective face images of a subject are well-modeled by low-rank approximation. Hence, MCar demonstrates marginally improvements over most of the methods in this dataset. This can be seen by visualizing the intermediate results of low-rank decomposition of the feature matrix using MCar as shown in Figure 6b. Besides, the IPAL method outperforms our methods when $\beta > 6$. Since the IPAL method utilizes the locally linear embedding for label propagation, which is effective in learning the underlying structure of data that has plenty of samples collected in the constrained environment. Hence, IPAL is able to recover the severely cluttered labels that MCar-based methods fail to approximate it as a low-rank matrix.

7.3 Experiments with Real-world Dataset

We conduct experiments on the *Lost* dataset and Labeled Yahoo! News dataset where the ambiguous labeling are collected in the real world. In the Labeled Yahoo! News dataset, we consider the labeling constraints between instances.

7.3.1 The Lost Dataset

The *Lost* dataset consists of face images and ambiguous labels automatically extracted using the screenplays provided in the TV series *Lost*. We use the *Lost* (16, 8) dataset released by the authors of [5] for evaluation. The *Lost* (16, 8) dataset consists of 1122 registered face images from 8 episodes, and the size of each is 60×90 pixels. The labels cover 16 subjects, but only 14 of them appear in the dataset. **Figure 7 illustrates the confusion matrix of the ambiguous labeling, which exhibits labeling imbalance.**

We compare MCar-based methods with the performance of ‘naive’, CLPL, MMS [17], and IPAL [25]. No labeling constraint between instances is considered in this experiment. Results are shown in Table 1. It can be seen from this table that MCar-based methods significantly outperform CLPL and MMS. This shows that MCar-based methods resolve the ambiguity and handles variations of instances in the TV series much better when compared to discriminative methods. Note that the performance of MMS is

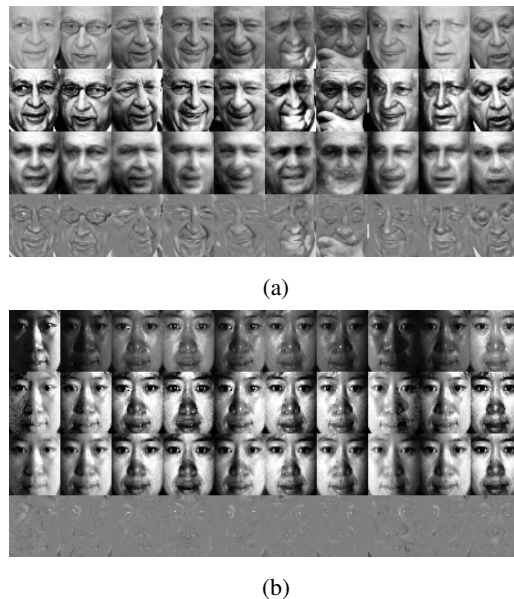


Fig. 6: Subsets of images from (a) FIW(10b) and (b) CMU PIE dataset demonstrate the low-rank decomposition of feature matrix in MCar: the original face images, histogram-equalized images X , low-rank component Z , and noisy component E_X , from the first row to the forth row, respectively.

close to that of CLPL since the ambiguous loss functions of both methods become similar when no labeling constraint between the instances is considered.

Figure 8 demonstrates that the groundtruth label distribution estimated by (21) is close to the groundtruth. Hence, WMCAR can effectively utilize this information to compensate the labeling imbalance. Although WMCAR slightly outperforms MCar, the collaboration of WMCAR and ICE (WMCAR-ICE) significantly outperforms MCar. On the other hand, MCar-ICE is inferior to WMCAR-ICE since the ICE procedure can inadvertently remove the candidates corresponding to minor labels without considering the labeling imbalance. It is challenging for IPAL to exploit the underlying structure of scarcity labeled data and deal with labeling imbalance. Hence, IPAL cannot successfully resolve the label ambiguity in this dataset. We tailor the RPCA [38] and MC-Pos [34] to solve the ambiguous learning problem by trivially taking the heterogeneous matrix as input and predicting the labels from

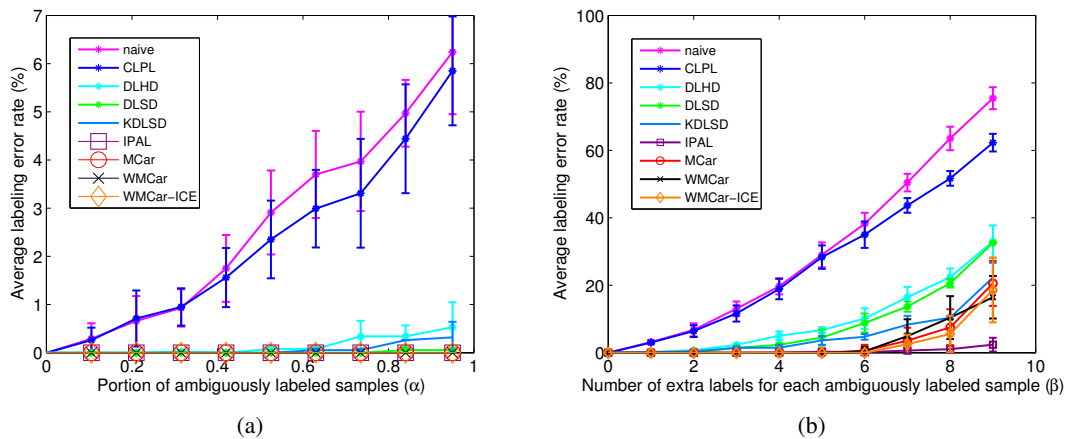


Fig. 5: Performance comparisons on the CMU PIE dataset. (a) $\alpha \in [0, 0.95]$, $\beta = 2$, *transductive* experiment. (b) $\alpha = 1.0$, $\beta \in [0, 1, \dots, 9]$, *transductive* experiment.

	204	75	20	29	71	24	7	2	1	15	5	4	1	0	0	1
1	99.17	30.83	7.00	11.17	33.00	9.50	3.00	0.67	0.33	5.17	2.17	1.33	0.33	0.00	0.00	0.33
2	83	198	19	89	12	31	6	2	3	0	5	0	0	0	0	0
3	32	5	142	0	54	1	40	2	0	26	0	0	0	0	12	1
4	28	72	2	103	0	16	2	0	1	4	0	1	3	0	1	1
5	11.33	2.00	67.17	0.00	25.67	0.33	19.50	0.83	0.00	9.33	0.00	0.00	0.00	0.00	5.50	0.33
6	34	6	39	0	88	4	1	0	0	11	0	0	2	0	0	0
7	15.67	2.00	19.00	0.00	43.50	1.33	0.33	0.00	0.00	0.00	5.50	0.00	0.00	0.67	0.00	0.00
8	34	35	1	19	9	103	2	7	0	6	5	4	19	5	0	0
9	13.33	13.00	0.50	6.33	3.67	47.67	0.67	2.33	0.00	2.00	1.67	1.33	8.83	1.67	0.00	0.00
10	5	11	59	1	0	2	76	0	0	0	0	0	0	0	0	0
11	2.33	5.17	28.33	0.33	0.00	0.67	39.17	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
12	7	1	1	0	0	6	0	33	14	7	2	1	0	6	0	0
13	2.33	0.50	0.33	0.00	0.00	2.00	0.00	14.83	7.00	2.50	0.67	0.33	0.00	2.50	0.00	0.00
14	3	0	0	0	0	0	0	29	61	0	0	23	0	0	0	0
15	11.83	72	2	103	0	16	2	0	0	0	0	0	0	0	0	0
16	4.00	0.00	5.33	0.00	1.83	0.00	0.00	13.83	35.50	0.00	0.00	10.50	0.00	1.00	0.00	0.33
1	5	11	59	1	0	2	76	0	0	0	0	0	0	0	0	0
2	2.50	1.33	0.00	0.00	7.00	1.50	0.00	0.00	0.00	0.00	12.33	0.00	1.33	0.00	0.00	0.00
3	0.67	0.00	0.00	0.00	0.00	0.67	0.00	10	7	2	0	18	0.00	0.00	0.00	0.00
4	0.00	1	1	4	0	12	0	1	0	3	11	0	25	0	0	0
5	0.00	0.33	0.33	1.67	0.00	5.67	0.00	0.33	0.00	1.00	4.50	0.00	11.17	0.00	0.00	0.00
6	0.00	0.00	0.00	5	1	2	0	12	5.00	0.00	1.67	1	3	0	20	0
7	0.00	0.00	0.00	2.33	0.33	0.67	0.00	5.00	0.00	4	0.33	1.00	0.00	8.67	0.00	0.00
8	0.00	0.00	0.00	0	0	0	0	0	0	0	0	0	0	0	0	0
9	0.00	0.00	0.00	0	0	0	0	0	0	0	0	0	0	0	0	0
10	0.00	0.00	0.00	0	0	0	0	0	0	0	0	0	0	0	0	0
11	0.00	0.00	0.00	0	0	0	0	0	0	0	0	0	0	0	0	0
12	0.00	0.00	0.00	0	0	0	0	0	0	0	0	0	0	0	0	0
13	0.00	0.00	0.00	0	0	0	0	0	0	0	0	0	0	0	0	0
14	0.00	0.00	0.00	0	0	0	0	0	0	0	0	0	0	0	0	0
15	0.00	0.00	0.00	0	0	0	0	0	0	0	0	0	0	0	0	0
16	0.00	0.00	0.00	0	0	0	0	0	0	0	0	0	0	0	0	0
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16

Fig. 7: The confusion matrix of the ambiguous labeling in *Lost* (16, 8) dataset. The upper number of each square accounts for the number of occurrences of a candidate label, whereas the lower number of each square is computed by accumulating the soft labeling score of each occurrence of a candidate label.

the soft labeling matrix of output with (18). The experimental result shows that existing low-rank approximation methods cannot substantially resolve the label ambiguity.

7.3.2 The Labeled Yahoo! News Dataset

The Labeled Yahoo! News dataset contains fully annotated faces in the images with names in the captions. It consists of 31147 detected faces from 20071 images. We use the precomputed SIFT feature of dimension 4992 extracted from that face images provided by Guillaumin *et al.* [42]. Following the protocol suggested in [17], we retain the 214 subjects with at least 20 occurrences in the captions. The remaining face images and names are treated as belonging to the additional null class. The ambiguous labeling is imbalanced in this dataset, where the number of labels present

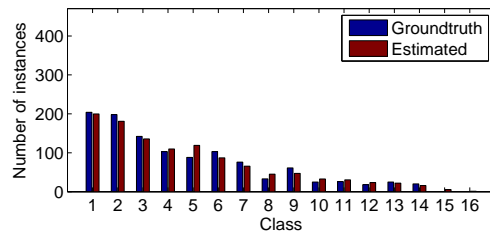


Fig. 8: The groundtruth label distribution of the *Lost* (16, 8) dataset. ‘Groundtruth’ denotes the number of instances per class counted from the groundtruth labels, and ‘Estimated’ denotes the estimate of the groundtruth label distribution from the ambiguous labels.

Method	Error Rate
naive	18.6 %
CLPL [21]	12.6 %
MMS [17]	11.4 %
IPAL [25]	22.9 %
RPCA [38]	29.9 %
MC-Pos [34]	23.6 %
MCar	8.5 %
WMC	8.2 %
MCar-ICE	8.0 %
WMC-ICE	5.2 %

TABLE 1: Labeling error rates for the *Lost* (16, 8) dataset (available at http://www.timotheecour.com/tv_data/tv_data.html).

in the captions ranges from 20 to 1917 with mean and standard deviations equal to 64.6 and 147.3, respectively. The top two subjects that are present most frequently in the captions are ‘george_w_bush’ and ‘saddam_hussein’. We conduct experiments on five training/testing splits by randomly selecting 80% of images and their associated captions as training set, and the rest are used as testing set. In each split, we also maintain the ratio between the number of training and testing instances from each subject.

The baseline approaches are CL-SVM, MIMLSVM [44], and MM-SVM [1]. The implementation details of CL-SVM and MIMLSVM are provided in [17]. We implemented the maximal

Method	Error Rate
CL-SVM	23.1 % \pm 0.6 %
MIMLSVM [44]	25.3 % \pm 0.3 %
MM-SVM [1]	21.9 % \pm 1.0 %
MMS [17]	14.3 % \pm 0.5 %
LR-SVM [12]	19.2 % \pm 0.4 %
MCar-SVM	14.5 % \pm 0.4 %
WMCar-SVM	13.6 % \pm 0.8 %
MCar-ICE-SVM	15.0 % \pm 1.0 %
WMCar-ICE-SVM	12.9 % \pm 0.8 %

TABLE 2: Average testing error rates for the Labeled Yahoo! News dataset (available at <http://lear.inrialpes.fr/data>).

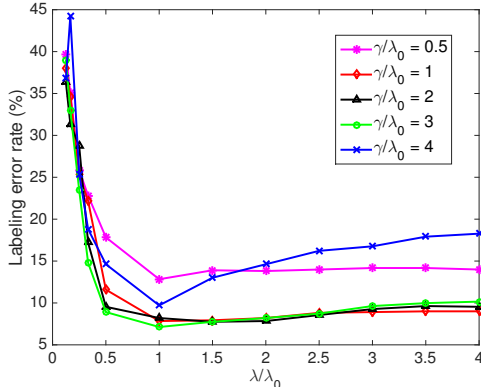


Fig. 9: Labeling error rates of WMCar evaluated with a set of parameters (λ, γ) in the *Lost* (16, 8) dataset.

likelihood assignment (MM) [1] to resolve the label ambiguity in the training data and train a multi-class linear SVM [43] to classify the testing data. For fair comparison, the language modeling of news captions in [1] is not included in the implementation. We compare with two state-of-the-art ambiguous labeling methods that consider labeling constraints between instances: MMS [17] and LR-SVM [12], which are based on discriminative model and low-rank framework, respectively. We resolve the ambiguity for the labels in the training set using (48) and train a multi-class linear SVM [43] to classify the testing data. Our MCar-SVM algorithm exhibits a slightly 0.2% higher error rate as compared to MMS. An explanation is that MCar relying on the low-rank approximation for ambiguity resolution is particularly sensitive to labeling imbalance. This results in performance degradation in the learned classifier since the output labels of MCar are potentially biased toward the majority labels.

Compared to the LR-SVM method, the MCar-SVM algorithm demonstrates 4.7% improvement on the testing accuracy. Since MCar assigns the labels across all instances via low-rank approximation of heterogeneous feature matrix, it is more effective than the LR-SVM method, which updates the PPM and the low-rank subspace of each class alternately. When we consider the labeling imbalance and utilize the ICE procedure, our proposed WMCar-ICE-SVM outperforms MCar-SVM by 1.6%

7.4 Sensitivity of Parameters

We use the *Lost* (16, 8) dataset to conduct the sensitivity analysis of MCar-based methods. In Figure 9, we evaluate the performance of WMCar over a set of parameters (λ, γ) . We observe that

the labeling error rate is relatively low when λ approaches λ_0 with respect to various γ . Hence, we conclude that the tradeoff parameter suggested in RPCA is applicable or at least a good reference for selecting λ .

In Figure 10, we evaluate the performance of MCar-based methods with various γ and a fixed $\lambda = \lambda_0$. For ICE, we set the elimination factor f_e as 0.5, and we set the maximum number of iterations as 5. Note that γ controls the sparsity of the soft labeling matrix \mathbf{Y} , and a larger γ will encourage a stronger sparsity on \mathbf{Y} . For $\gamma \in [1, 4]\lambda_0$, a moderate sparsity of \mathbf{Y} is helpful in predicting the actual label from the ambiguous labels. As γ (sparsity of \mathbf{Y}) significantly increases, the performance degrades. One explanation is that a very strong sparsity on \mathbf{Y} imposes an immediately hard decision on the labels, which may lead to adverse effect on low-rank approximation. On the other hand, a mild sparsity on \mathbf{Y} allows an unsure face image to be represented by the feature subspaces of those individuals that resemble the unsure face image. The selected parameter $\gamma = 2\lambda_0$ yields good performance for the MCar-based methods as illustrated in Figure 10. We observe that WMCar-ICE is less sensitive to γ since the ICE procedure intrinsically encourages the sparsity when removing the least likely candidate from a candidate labeling set.

We conduct the sensitivity analysis of WMCar-ICE with $\lambda = \lambda_0$ and $\gamma = 2\lambda_0$ and evaluate the performance with various f_e . In Figure 11, the performance of WMCar-ICE ($f_e = 0$) fluctuates since the ICE procedure becomes ineffective as $f_e = 0$. A small elimination factor ($f_e = 0.25$) yields better performance than large elimination factors, but it takes more iterations to converge. Since the candidate elimination step in WMCar-ICE can incur an irreversible error, a small elimination factor can conservatively eliminate the least likely candidates in the candidate labeling sets. Hence, abrupt decision resulting from large elimination factors can be avoided, and the soft labeling matrix can be gently updated to guide the low-rank approximation of heterogeneous matrix. Figure 11 confirms that the selected parameters ($f_e = 0.5$ and the maximum number of iterations equal to 5) yield good performance in terms of the rate of convergence and labeling error rate.

Owing to the greedy nature of the ICE procedure, we cannot guarantee that labeling error rates will monotonically decrease as the number of iterations of ICE increases. An incorrect removal of the actual label in the candidate label set leads to irreversible labeling error of this instance. Moreover, it can cause detrimental effect on resolving the ambiguity of other instances since the error of an incorrectly labeled instance can be propagated to other instances of similar appearance through low-rank approximation of the heterogeneous matrix.

7.5 Convergence

Since the projection method in Section 4.4 is not non-expansive, we cannot simply follow the rationale that the composition of gradient, shrinkage, and projection steps is non-expansive to prove convergence [34]. We attempt to replace the projection method of MCar with the Euclidean projection onto the simplex [45], which is a non-expansive projection, but the performance of the modified MCar degrades significantly. An explanation is that the Euclidean projection onto the simplex can inadvertently generate non-sparse entries, which conflicts with the original objective to encourage the sparsity of the soft labeling matrix in (25). On the other hand, our simple projection step normalizes the ℓ_1 norm of a soft labeling vector, which effectively restricts the soft labeling vector to lying

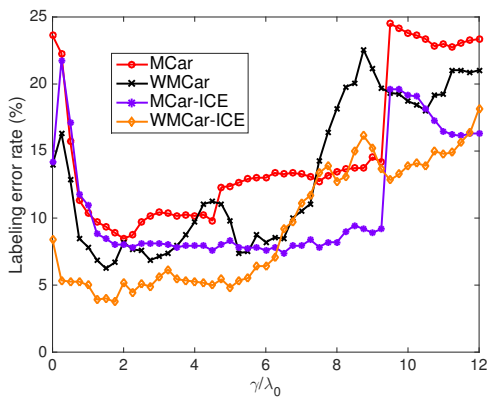


Fig. 10: Labeling error rates of MCar-based methods versus γ in the *Lost* (16, 8) dataset with $\lambda = \lambda_0$.

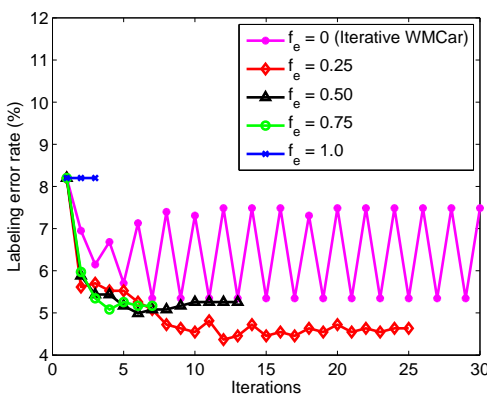


Fig. 11: Labeling error rate versus the number of iterations in WMCAR-ICE. The performance is evaluated in the *Lost* (16, 8) dataset with various elimination factors. The performance of WMCAR-ICE ($f_e = 0$) fluctuates since the ICE procedure becomes ineffective as $f_e = 0$.

on the ℓ_1 ball and maintains an identical sparsity. Although the convergence of MCar has been observed empirically in [18], a theoretical justification of convergence needs further investigation. Since the number of ambiguous labels is finite, the convergence of ICE is straightforward with $f_e > 0$

8 CONCLUSIONS

We introduced a novel matrix completion framework for resolving the ambiguity of labels. In contrast to existing iterative alternating approaches, the proposed MCar method ensures all the instances and their associated ambiguous labels are utilized as a whole for resolving the ambiguity. Since MCar is capable of discovering the underlying low-rank structure of subjects, it is robust to within-subject variations. Hence, MCar can serve as the counterpart of discriminative ambiguous learning methods. Besides, WMCAR generalizes MCar to compensate the labeling imbalance, and thus an instance associated with minority labels has a stronger impact than that associated with majority labels. The ICE procedure improves the performance of iterative WMCAR by eliminating a portion of the least likely candidates in each iteration. As demonstrated by the experiments on the synthesized ambiguous labels and two datasets collected from real world, our

proposed methods consistently resolve the ambiguity when single face images or group of face images are ambiguously labeled.

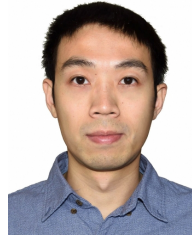
ACKNOWLEDGMENTS

This research is based upon work supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via IARPA R&D Contract No. 2014-14071600012. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon.

REFERENCES

- [1] T. L. Berg, A. C. Berg, J. Edwards, and D. A. Forsyth, “Who’s in the picture?” in *Neural Information Processing Systems (NIPS)*, 2004.
- [2] T. L. Berg, A. C. Berg, J. Edwards, M. Maire, R. White, Y.-W. Teh, E. Learned-Miller, and D. A. Forsyth, “Names and faces in the news,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2004.
- [3] M. Everingham, J. Sivic, and A. Zisserman, “Hello! My name is... Buffy - Automatic naming of characters in TV video,” in *Proceedings of the British Machine Vision Conference (BMVC)*, 2006.
- [4] E. Hüllermeier and J. Beringer, “Learning from ambiguously labeled examples,” in *Intelligent Data Analysis*, 2006.
- [5] T. Cour, B. Sapp, C. Jordan, and B. Taskar, “Learning from ambiguously labeled images,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [6] L.-P. Liu and T. G. Dietterich, “A conditional multinomial mixture model for superset label learning,” in *Neural Information Processing Systems (NIPS)*, 2012.
- [7] Y.-C. Chen, V. M. Patel, J. K. Pillai, R. Chellappa, and P. J. Phillips, “Dictionary learning from ambiguously labeled data,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [8] L.-P. Liu and T. G. Dietterich, “Learnability of the superset label learning problem,” in *International Conference on Machine Learning (ICML)*, 2014.
- [9] A. B. Goldberg, X. Zhu, B. Recht, J.-M. Xu, and R. D. Nowak, “Transduction with matrix completion: Three birds with one stone,” in *Neural Information Processing Systems (NIPS)*, 2010.
- [10] R. S. Cabral, F. D. la Torre, J. P. Costeira, and A. Bernardino, “Matrix completion for multi-label image classification,” in *Neural Information Processing Systems (NIPS)*, 2011.
- [11] H. He and E. A. Garcia, “Learning from imbalanced data,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 21, no. 9, pp. 1263–1284, Sep. 2009.
- [12] Z. Zeng, S. Xiao, K. Jia, T.-H. Chan, S. Gao, D. Xu, and Y. Ma, “Learning by associating ambiguously labeled images,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [13] Y.-C. Chen, V. M. Patel, R. Chellappa, and P. J. Phillips, “Ambiguously labeled learning using dictionaries,” *IEEE Transactions on Information Forensics and Security*, vol. 9, no. 12, pp. 2076–2088, Dec. 2014.
- [14] E. J. Candès, X. Li, Y. Ma, and J. Wright, “Robust principal component analysis?” *Journal of the ACM*, vol. 58, no. 3, pp. 11:1–11:37, Jun. 2011.
- [15] C.-F. Chen, C.-P. Wei, and Y.-C. F. Wang, “Low-rank matrix recovery with structural incoherence for robust face recognition,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [16] D. Huang, R. S. Cabral, and F. D. la Torre, “Robust regression,” in *European Conference on Computer Vision (ECCV)*, 2012.
- [17] J. Luo and F. Orabona, “Learning from candidate labeling sets,” in *Neural Information Processing Systems (NIPS)*, 2010.
- [18] C.-H. Chen, V. M. Patel, and R. Chellappa, “Matrix completion for resolving label ambiguity,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [19] C. Ambroise, T. Denoeux, G. Govaert, and P. Smets, “Learning from an imprecise teacher: Probabilistic and evidential approaches,” *Applied Stochastic Models and Data Analysis*, vol. 1, pp. 100–105, 2001.
- [20] R. Jin and Z. Ghahramani, “Learning with multiple labels,” in *Neural Information Processing Systems (NIPS)*, 2002.

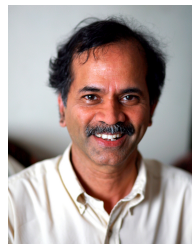
- [21] T. Cour, B. Sapp, and B. Taskar, "Learning from partial labels," *Journal of Machine Learning Research*, vol. 12, pp. 1501–1536, 2011.
- [22] T. Zhang, "Statistical analysis of some multi-category large margin classification methods," *Journal of Machine Learning Research*, vol. 5, pp. 1225–1251, 2004.
- [23] J. Cid-Sueiro, "Proper losses for learning from partial labels," in *Neural Information Processing Systems (NIPS)*, 2012.
- [24] M.-L. Zhang, "Disambiguation-free partial label learning," in *SIAM International Conference on Data Mining*, 2014.
- [25] M.-L. Zhang and F. Yu, "Solving the partial label learning problem: An instance-based approach," in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2015.
- [26] A. Shrivastava, V. M. Patel, and R. Chellappa, "Non-linear dictionary learning with partially labeled data," *Pattern Recognition*, vol. 48, no. 11, pp. 3283–3292, Nov. 2015.
- [27] S. Xiao, D. Xu, and J. Wu, "Automatic face naming by learning discriminative affinity matrices from weakly labeled images," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 26, no. 10, pp. 2440–2452, Oct. 2015.
- [28] M. Sahare and H. Gupta, "A review of multi-class classification for imbalanced data," *International Journal of Advanced Computer Research*, vol. 2, no. 3, pp. 160–164, 2012.
- [29] M.-L. Zhang, Y.-K. Li, and X.-Y. Liu, "Towards class-imbalance aware multi-label learning," in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2015.
- [30] K. Chen, B.-L. Lu, and J. T. Kwok, "Efficient classification of multi-label and imbalanced data using min-max modular classifiers," in *International Joint Conference on Neural Networks (IJCNN)*, 2006.
- [31] B.-L. Lu, K.-A. Wang, M. Utiyama, and H. Isahara, "A part-versus-part method for massively parallel training of support vector machines," in *International Joint Conference on Neural Networks (IJCNN)*, 2004.
- [32] F. Charte, A. J. Rivera, M. J. del Jesus, and F. Herrera, "Addressing imbalance in multilabel classification: Measures and random resampling algorithms," *Neurocomputing*, vol. 163, pp. 3–16, Sep. 2015.
- [33] B. Wu, S. Lyu, and B. Ghanem, "Constrained submodular minimization for missing labels and class imbalance in multi-label learning," in *AAAI Conference on Artificial Intelligence*, 2016.
- [34] R. Cabral, F. D. la Torre, J. P. Costeira, and A. Bernardino, "Matrix completion for weakly-supervised multi-label image classification," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 37, no. 1, pp. 121–135, Jan. 2015.
- [35] T. M. Cover, "Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition," *IEEE Transactions on Electronic Computers*, vol. EC-14, no. 3, pp. 326–334, Jun. 1965.
- [36] K. Veropoulos, C. Campbell, and N. Cristianini, "Controlling the sensitivity of support vector machines," in *International Joint Conference on AI*, 1999.
- [37] Y. Lin, Y. Lee, and G. Wahba, "Support vector machines for classification in nonstandard situations," *Machine Learning*, vol. 46, no. 1.
- [38] Z. Lin, M. Chen, L. Wu, and Y. Ma, "The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices," *UIUC Technical Report UILU-ENG-09-2215*, November 2009.
- [39] E. J. Candès and B. Recht, "Exact low-rank matrix completion via convex optimization," in *Allerton Conference on Communication, Control, and Computing*, 2008.
- [40] J.-F. Cai, E. J. Candès, and Z. Shen, "A singular value thresholding algorithm for matrix completion," *SIAM Journal on Optimization*, vol. 20, no. 4, pp. 1956–1982, Mar. 2004.
- [41] G. B. Huang, V. Jain, and E. Learned-Miller, "Unsupervised joint alignment of complex images," in *IEEE International Conference on Computer Vision (ICCV)*, 2007.
- [42] M. Guillaumin, J. Verbeek, and C. Schmid, "Multiple instance metric learning from automatically labeled bags of faces," in *European Conference on Computer Vision (ECCV)*, 2010.
- [43] C.-C. Chang and C.-J. Lin, "LIBSVM: A library for support vector machines," *ACM Transactions on Intelligent Systems and Technology*, vol. 2, pp. 27:1–27:27, 2011, software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- [44] Z. Zhou and M. Zhang, "Multi-instance multilabel learning with application to scene classification," in *Neural Information Processing Systems (NIPS)*, 2006.
- [45] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra, "Efficient projections onto the L1-ball for learning in high dimensions," in *International Conference on Machine Learning*, 2008.



Ching-Hui Chen received the B.S. degree in Electrical Engineering and Computer Science Undergraduate Honors Program and M.S. degree in Electronics Engineering both from National Chiao Tung University, Hsinchu, Taiwan. He received the Ph.D. degree in Electrical and Computer Engineering from the University of Maryland, College Park, MD, USA. His research interests are in computer vision and pattern recognition, with primary focuses on face recognition, stereo vision, and camera networks.



Vishal M. Patel received the B.S. degrees in electrical engineering and applied mathematics (Hons.) and the M.S. degree in applied mathematics from North Carolina State University, Raleigh, NC, USA, in 2004 and 2005, respectively, and the Ph.D. degree in electrical engineering from the University of Maryland College Park, MD, USA, in 2010. He is currently an A. Walter Tyson Assistant Professor in the Department of Electrical and Computer Engineering (ECE) at Rutgers University. Prior to joining Rutgers University, he was a member of the research faculty at the University of Maryland Institute for Advanced Computer Studies (UMIACS). His current research interests include signal processing, computer vision, and pattern recognition with applications in biometrics and imaging. He has received a number of awards including the 2016 ONR Young Investigator Award, the 2016 Jimmy Lin Award for Invention, A. Walter Tyson Assistant Professorship Award, the Best Paper Award at IEEE BTAS 2015, and Best Poster Awards at BTAS 2015 and 2016. He is an Associate Editor of the IEEE Signal Processing Magazine, IEEE Biometrics Compendium, and serves on the Information Forensics and Security Technical Committee of the IEEE Signal Processing Society. He is a member of Eta Kappa Nu, Pi Mu Epsilon, and Phi Beta Kappa.



Rama Chellappa is a Minta Martin Professor of Engineering and Chair of the ECE department at the University of Maryland. Prof. Chellappa received the K.S. Fu Prize from the International Association of Pattern Recognition (IAPR). He is a recipient of the Society, Technical Achievement and Meritorious Service Awards from the IEEE Signal Processing Society and four IBM faculty Development Awards. He also received the Technical Achievement and Meritorious Service Awards from the IEEE Computer Society. At UMD, he received college and university level recognitions for research, teaching, innovation and mentoring of undergraduate students. In 2010, he was recognized as an Outstanding ECE by Purdue University. Prof. Chellappa served as the Editor-in-Chief of PAMI. He is a Golden Core Member of the IEEE Computer Society, served as a Distinguished Lecturer of the IEEE Signal Processing Society and as the President of IEEE Biometrics Council. He is a Fellow of IEEE, IAPR, OSA, AAAS, ACM and AAAI and holds four patents.