

Automatic Face Naming by Learning Discriminative Affinity Matrices From Weakly Labeled Images

Shijie Xiao, Dong Xu, *Senior Member, IEEE*, and Jianxin Wu, *Member, IEEE*

Abstract—Given a collection of images, where each image contains several faces and is associated with a few names in the corresponding caption, the goal of face naming is to infer the correct name for each face. In this paper, we propose two new methods to effectively solve this problem by learning two discriminative affinity matrices from these weakly labeled images. We first propose a new method called regularized low-rank representation by effectively utilizing weakly supervised information to learn a low-rank reconstruction coefficient matrix while exploring multiple subspace structures of the data. Specifically, by introducing a specially designed regularizer to the low-rank representation method, we penalize the corresponding reconstruction coefficients related to the situations where a face is reconstructed by using face images from other subjects or by using itself. With the inferred reconstruction coefficient matrix, a discriminative affinity matrix can be obtained. Moreover, we also develop a new distance metric learning method called ambiguously supervised structural metric learning by using weakly supervised information to seek a discriminative distance metric. Hence, another discriminative affinity matrix can be obtained using the similarity matrix (i.e., the kernel matrix) based on the Mahalanobis distances of the data. Observing that these two affinity matrices contain complementary information, we further combine them to obtain a fused affinity matrix, based on which we develop a new iterative scheme to infer the name of each face. Comprehensive experiments demonstrate the effectiveness of our approach.

Index Terms—Affinity matrix, caption-based face naming, distance metric learning, low-rank representation (LRR).

I. INTRODUCTION

IN SOCIAL networking websites (e.g., Facebook), photo sharing websites (e.g., Flickr) and news websites (e.g., BBC), an image that contains multiple faces can be associated with a caption specifying who is in the picture. For instance, multiple faces may appear in a news photo with a caption that briefly describes the news. Moreover, in TV serials, movies, and news videos, the faces may also appear in a video clip with scripts. In the literature, a few methods were developed for the face naming problem (see Section II for more details).

Manuscript received May 28, 2014; revised September 4, 2014 and November 26, 2014; accepted December 11, 2014.

S. Xiao is with the School of Computer Engineering, Nanyang Technological University, Singapore 639798 (e-mail: xiao0050@e.ntu.edu.sg).

D. Xu was with the School of Computer Engineering, Nanyang Technological University, Singapore 639798. He is now with the School of Electrical and Information Engineering, The University of Sydney, Sydney, NSW 2006, Australia. (e-mail: dongxudongxu@gmail.com).

J. Wu is with the National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210023, China (e-mail: wujx2001@nju.edu.cn).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNNLS.2014.2386307

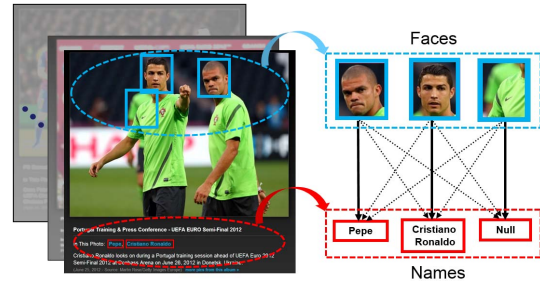


Fig. 1. Illustration of the face-naming task, in which we aim to infer which name matches which face, based on the images and the corresponding captions. The solid arrows between faces and names indicate the ground-truth face-name pairs and the dashed ones represent the incorrect face-name pairs, where null means the ground-truth name of a face does not appear in the candidate name set.

In this paper, we focus on automatically annotating faces in images based on the ambiguous supervision from the associated captions. Fig. 1 gives an illustration of the face-naming problem. Some preprocessing steps need to be conducted before performing face naming. Specifically, faces in the images are automatically detected using face detectors [1], and names in the captions are automatically extracted using a name entity detector. Here, the list of names appearing in a caption is denoted as the *candidate name set*. Even after successfully performing these preprocessing steps, automatic face naming is still a challenging task. The faces from the same subject may have different appearances because of the variations in poses, illuminations, and expressions. Moreover, the candidate name set may be noisy and incomplete, so a name may be mentioned in the caption, but the corresponding face may not appear in the image, and the correct name for a face in the image may not appear in the corresponding caption. Each detected face (including falsely detected ones) in an image can only be annotated using one of the names in the candidate name set or as null, which indicates that the ground-truth name does not appear in the caption.

In this paper, we propose a new scheme for automatic face naming with caption-based supervision. Specifically, we develop two methods to respectively obtain two discriminative affinity matrices by learning from weakly labeled images. The two affinity matrices are further fused to generate one fused affinity matrix, based on which an iterative scheme is developed for automatic face naming.

To obtain the first affinity matrix, we propose a new method called regularized low-rank representation (rLRR) by incorporating weakly supervised information into the low-rank representation (LRR) method, so that the affinity matrix can be




Fig. 2. Coefficient matrix \mathbf{W}^* according to the groundtruth and the ones obtained from LRR and rLRR. (a) \mathbf{W}^* according to the groundtruth. (b) \mathbf{W}^* from LRR. (c) \mathbf{W}^* from our rLRR.

obtained from the resultant reconstruction coefficient matrix. To effectively infer the correspondences between the faces based on visual features and the names in the candidate name sets, we exploit the *subspace structures among faces* based on the following assumption: the faces from the same subject/name lie in the same subspace and the subspaces are linearly independent. Liu *et al.* [2] showed that such subspace structures can be effectively recovered using LRR, when the subspaces are independent and the data sampling rate is sufficient. They also showed that the mined subspace information is encoded in the reconstruction coefficient matrix that is block-diagonal in the ideal case. As an intuitive motivation, we implement LRR on a synthetic dataset and the resultant reconstruction coefficient matrix is shown in Fig. 2(b) (More details can be found in Sections V-A and V-C). This near block-diagonal matrix validates our assumption on the *subspace structures among faces*. Specifically, the reconstruction coefficients between one face and faces from the same subject are generally larger than others, indicating that the faces from the same subject tend to lie in the same subspace [2]. However, due to the significant variances of in-the-wild faces in poses, illuminations, and expressions, the appearances of faces from different subjects may be even more similar when compared with those from the same subject. Consequently, as shown in Fig. 2(b), the faces may also be reconstructed using faces from other subjects. In this paper, we show that the candidate names from the captions can provide important supervision information to better discover the subspace structures.

In Section III-C2, we first propose a method called rLRR by introducing a new regularizer that incorporates caption-based weak supervision into the objective of LRR, in which we penalize the reconstruction coefficients when reconstructing the faces using those from different subjects. Based on the inferred reconstruction coefficient matrix, we can compute an affinity matrix that measures the similarity values between every pair of faces. Compared with the one in Fig. 2(b), the reconstruction coefficient matrix from our rLRR exhibits more obvious block-diagonal structure in Fig. 2(c), which indicates that a better reconstruction matrix can be obtained using the proposed regularizer.

Moreover, we use the similarity matrix (i.e., the kernel matrix) based on the Mahalanobis distances between the faces as another affinity matrix. Specifically, in Section III-D, we develop a new distance metric learning method called ambiguously supervised structural metric learning (ASML) to learn a discriminative Mahalanobis distance metric based on weak supervision information. In ASML, we consider the

constraints for the label matrix of the faces in each image by using the feasible label set, and we further define the image to assignment (I2A) distance that measures the incompatibility between a label matrix and the faces from each image based on the distance metric. Hence, ASML learns a Mahalanobis distance metric that encourages the I2A distance based on a selected feasible label matrix, which approximates the ground-truth one, to be smaller than the I2A distances based on infeasible label matrices to some extent.

Since rLRR and ASML explore the weak supervision in different ways and they are both effective, as shown in our experimental results in Section V, the two corresponding affinity matrices are expected to contain complementary and discriminative information for face naming. Therefore, to further improve the performance, we combine the two affinity matrices to obtain a fused affinity matrix that is used for face naming. Accordingly, we refer to this method as regularized low rank representation with metric learning (rLRRml for short). Based on the fused affinity matrix, we additionally propose a new iterative method by formulating the face naming problem as an integer programming problem with linear constraints, where the constraints are related to the feasible label set of each image.

Our main contributions are summarized as follows.

- 1) Based on the caption-based weak supervision, we propose a new method rLRR by introducing a new regularizer into LRR and we can calculate the first affinity matrix using the resultant reconstruction coefficient matrix (Section III-C).
- 2) We also propose a new distance metric learning approach ASML to learn a discriminative distance metric by effectively coping with the ambiguous labels of faces. The similarity matrix (i.e., the kernel matrix) based on the Mahalanobis distances between all faces is used as the second affinity matrix (Section III-D).
- 3) With the fused affinity matrix by combining the two affinity matrices from rLRR and ASML, we propose an efficient scheme to infer the names of faces (Section IV).
- 4) Comprehensive experiments are conducted on one synthetic dataset and two real-world datasets, and the results demonstrate the effectiveness of our approaches (Section V).

II. RELATED WORK

Recently, there is an increasing research interest in developing automatic techniques for face naming in images [3]–[9] as well as in videos [10]–[13]. To tag faces in news photos, Berg *et al.* [3] proposed to cluster the faces in the news images. Ozkan and Duygulu [4] developed a graph-based method by constructing the similarity graph of faces and finding the densest component. Guillaumin *et al.* [6] proposed the multiple-instance logistic discriminant metric learning (MildML) method. Luo and Orabona [7] proposed a structural support vector machine (SVM)-like algorithm called maximum margin set (MMS) to solve the face naming problem. Recently, Zeng *et al.* [9] proposed the low-rank SVM (LR-SVM) approach to deal with this problem,

based on the assumption that the feature matrix formed by faces from the same subject is low rank. In the following, we compare our proposed approaches with several related existing methods.

Our rLRR method is related to LRR [2] and LR-SVM [9]. LRR is an unsupervised approach for exploring multiple subspace structures of data. In contrast to LRR, our rLRR utilizes the weak supervision from image captions and also considers the image-level constraints when solving the weakly supervised face naming problem. Moreover, our rLRR differs from LR-SVM [9] in the following two aspects. 1) To utilize the weak supervision, LR-SVM considers weak supervision information in the partial permutation matrices, while rLRR uses our proposed regularizer to penalize the corresponding reconstruction coefficients. 2) LR-SVM is based on robust principal component analysis (RPCA) [14]. Similarly to [15], LR-SVM does not reconstruct the data by using itself as the dictionary. In contrast, our rLRR is related to the reconstruction based approach LRR.

Moreover, our ASML is related to the traditional metric learning works, such as large-margin nearest neighbors (LMNN) [16], Frobnmetric [17], and metric learning to rank (MLR) [18]. LMNN and Frobnmetric are based on accurate supervision without ambiguity (i.e., the triplets of training samples are explicitly given), and they both use the hinge loss in their formulation. In contrast, our ASML is based on the ambiguous supervision, and we use a max margin loss to handle the ambiguity of the structural output, by enforcing the distance based on the best label assignment matrix in the feasible label set to be larger than the distance based on the best label assignment matrix in the infeasible label set by a margin. Although a similar loss that deals with structural output is also used in MLR, it is used to model the ranking orders of training samples, and there is no uncertainty regarding supervision information in MLR (i.e., the groundtruth ordering for each query is given).

Our ASML is also related to two recently proposed approaches for the face naming problem using weak supervision, MildML [6], and MMS [7]. MildML follows the multi-instance learning (MIL) assumption, which assumes that each image should contain a face corresponding to each name in the caption. However, it may not hold for our face naming problem as the captions are not accurate. In contrast, our ASML employs a maximum margin loss to handle the structural output without using such an assumption. While MMS also uses a maximum margin loss to handle the structural output, MMS aims to learn the classifiers and it was designed for the classification problem. Our ASML learns a distance metric that can be readily used to generate an affinity matrix and can be combined with the affinity matrix from our rLRR method to further improve the face naming performance.

Finally, we compare our face naming problem with MIL [19], multi-instance multilabel learning (MIML) [20], and the face naming problem in [21]. In the existing MIL and MIML works, a few instances are grouped into bags, in which the bag labels are assumed to be correct. Moreover, the common assumption in MIL is that one positive bag contains at least one positive instance. A straightforward

way to apply MIL and MIML methods for solving the face naming problem is to treat each image as a bag, the faces in the image as the instances, and the names in the caption as the bag labels. However, the bag labels (based on candidate name sets) may be even incorrect in our problem because the faces corresponding to the mentioned names in the caption may be absent in the image. Besides, one common assumption in face naming is that any two faces in the same image cannot be annotated using the same name, which indicates that each positive bag contains no more than one positive instance rather than at least one positive instance. Moreover, in [21], each image only contains one face. In contrast, we may have multiple faces in one image, which are related to a set of candidate names in our problem.

III. LEARNING DISCRIMINATIVE AFFINITY MATRICES FOR AUTOMATIC FACE NAMING

In this section, we propose a new approach for automatic face naming with caption-based supervision. In Sections III-A and III-B, we formally introduce the problem and definitions, followed by the introduction of our proposed approach. Specifically, we learn two discriminative affinity matrices by effectively utilizing the ambiguous labels, and perform face naming based on the fused affinity matrix. In Sections III-C and III-D, we introduce our proposed approaches rLRR and ASML for obtaining the two affinity matrices respectively.

In the remainder of this paper, we use lowercase/uppercase letters in boldface to denote a vector/matrix (e.g., \mathbf{a} denotes a vector and \mathbf{A} denotes a matrix). The corresponding nonbold letter with a subscript denotes the entry in a vector/matrix (e.g., a_i denotes the i th entry of the vector \mathbf{a} , and $A_{i,j}$ denotes an entry at the i th row and j th column of the matrix \mathbf{A}). The superscript $'$ denotes the transpose of a vector or a matrix. We define \mathbf{I}_n as the $n \times n$ identity matrix, and $\mathbf{0}_n, \mathbf{1}_n \in \mathbb{R}^n$ as the $n \times 1$ column vectors of all zeros and all ones, respectively. For simplicity, we also use $\mathbf{I}, \mathbf{0}$ and $\mathbf{1}$ instead of $\mathbf{I}_n, \mathbf{0}_n$, and $\mathbf{1}_n$ when the dimensionality is obvious. Moreover, we use $\mathbf{A} \circ \mathbf{B}$ (resp., $\mathbf{a} \circ \mathbf{b}$) to denote the element-wise product between two matrices \mathbf{A} and \mathbf{B} (resp., two vectors \mathbf{a} and \mathbf{b}). $\text{tr}(\mathbf{A})$ denotes the trace of \mathbf{A} (i.e., $\text{tr}(\mathbf{A}) = \sum_i A_{i,i}$), and $\langle \mathbf{A}, \mathbf{B} \rangle$ denotes the inner product of two matrices (i.e., $\langle \mathbf{A}, \mathbf{B} \rangle = \text{tr}(\mathbf{A}'\mathbf{B})$). The inequality $\mathbf{a} \leq \mathbf{b}$ means that $a_i \leq b_i \ \forall i = 1, \dots, n$ and $\mathbf{A} \geq \mathbf{0}$ means that \mathbf{A} is a positive semidefinite (PSD) matrix. $\|\mathbf{A}\|_F = (\sum_{i,j} A_{i,j}^2)^{1/2}$ denotes the Frobenious norm of a matrix \mathbf{A} . $\|\mathbf{A}\|_\infty$ denotes the largest absolute value of all elements in \mathbf{A} .

A. Problem Statement

Given a collection of images, each of which contains several faces and is associated with multiple names, our goal is to annotate each face in these images with these names.

Formally, let us assume we have m images, each of which contains n_i faces and r_i names, $i = 1, \dots, m$. Let $\mathbf{x} \in \mathbb{R}^d$ denote a face, where d is the feature dimension. Moreover, let $q \in \{1, \dots, p\}$ denote a name, where p is the total number of names in all the captions. Then, each image can be represented as a pair $(\mathbf{X}^i, \mathcal{N}^i)$, where $\mathbf{X}^i = [\mathbf{x}_1^i, \dots, \mathbf{x}_{n_i}^i] \in \mathbb{R}^{d \times n_i}$

is the data matrix for faces in the i th image with each \mathbf{x}_f^i being the f th face in this image ($f = 1, \dots, n_i$), and $\mathcal{N}^i = \{q_1^i, \dots, q_{r_i}^i\}$ is the corresponding set of candidate names with each $q_j^i \in \{1, \dots, p\}$ being the j th name ($j = 1, \dots, r_i$). Moreover, let $\mathbf{X} = [\mathbf{X}^1, \dots, \mathbf{X}^m] \in \mathbb{R}^{d \times n}$ denote the data matrix of the faces from all m images, where $n = \sum_{i=1}^m n_i$.

By defining a binary label matrix $\mathbf{Y} = [\mathbf{Y}^1, \dots, \mathbf{Y}^m] \in \{0, 1\}^{(p+1) \times n}$ with each $\mathbf{Y}^i \in \{0, 1\}^{(p+1) \times n_i}$ being the label matrix for each image \mathbf{X}^i , then the task is to infer the label matrix \mathbf{Y} based on the candidate name sets $\{\mathcal{N}^i\}_{i=1}^m$. Considering the situation where the ground-truth name of a face does not appear in the associated candidate name set \mathcal{N}^i , we use the $(p+1)$ th name to denote the null class, so that the face should be assigned to the $(p+1)$ th name in this situation. Moreover, the label matrix \mathbf{Y}^i for each image should satisfy the following three image-level constraints [9].

- 1) *Feasibility*: the faces in the i th image should be annotated using the names from the set $\tilde{\mathcal{N}}^i = \mathcal{N}^i \cup \{(p+1)\}$, i.e., $Y_{j,f}^i = 0, \forall f = 1, \dots, n_i$ and $j \notin \tilde{\mathcal{N}}^i$.
- 2) *Nonredundancy*: each face in the i th image should be annotated using exactly one name from $\tilde{\mathcal{N}}^i$, i.e., $\sum_j Y_{j,f}^i = 1, \forall f = 1, \dots, n_i$.
- 3) *Uniqueness*: two faces in the same image cannot be annotated with the same name except the $(p+1)$ th name (i.e., the null class), i.e., $\sum_{f=1}^{n_i} Y_{j,f}^i \leq 1, \forall j = 1, \dots, p$.

B. Face Naming Using a Discriminative Affinity Matrix

First, based on the image-level constraints, we define the feasible set of \mathbf{Y}^i for the i th image as follows:

$$\mathcal{Y}^i = \left\{ \mathbf{Y}^i \in \{0, 1\}^{(p+1) \times n_i} \left| \begin{array}{l} \mathbf{1}_{(p+1)}'(\mathbf{Y}^i \circ \mathbf{T}^i)\mathbf{1}_{n_i} = 0, \\ \mathbf{1}_{(p+1)}'\mathbf{Y}^i = \mathbf{1}_{n_i}', \\ \mathbf{Y}^i\mathbf{1}_{n_i} \leq [\mathbf{1}_p', n_i]' \end{array} \right. \right\} \quad (1)$$

where $\mathbf{T}^i \in \{0, 1\}^{(p+1) \times n_i}$ is a matrix in which the rows related to the indices of the names in $\tilde{\mathcal{N}}^i$ are all zeros and the other rows are all ones.

Accordingly, the feasible set for the label matrix on all images can be represented as

$$\mathcal{Y} = \{\mathbf{Y} = [\mathbf{Y}^1, \dots, \mathbf{Y}^m] \mid \mathbf{Y}^i \in \mathcal{Y}^i \quad \forall i = 1, \dots, m\}.$$

Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be an affinity matrix, which satisfies that $\mathbf{A} = \mathbf{A}'$ and $A_{i,j} \geq 0, \forall i, j$. Each $A_{i,j}$ describes the pairwise affinity/similarity between the i th face and the j th face [2]. We aim to learn a proper \mathbf{A} such that $A_{i,j}$ is large if and only if the i th face and the j th face share the same groundtruth name. Then, one can solve the face naming problem based on the obtained affinity matrix \mathbf{A} . To infer the names of faces, we aim to solve the following:

$$\max_{\mathbf{Y} \in \mathcal{Y}} \sum_{c=1}^p \frac{\mathbf{y}_c' \mathbf{A} \mathbf{y}_c}{\mathbf{1}' \mathbf{y}_c} \quad \text{s.t. } \mathbf{Y} = [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_{(p+1)}]' \quad (2)$$

where $\mathbf{y}_c \in \{0, 1\}^n$ corresponds to the c th row in \mathbf{Y} . The intuitive idea is that we cluster the faces with the same inferred label as one group, and we maximize the sum of the average

affinities for each group. The solution of this problem will be introduced in Section IV. According to (2), a good affinity matrix is crucial in our proposed face naming scheme, because it directly determines the face naming performance.

In this paper, we consider two methods to obtain two affinity matrices, respectively. Specifically, to obtain the first affinity matrix, we propose the rLRR method to learn the low-rank reconstruction coefficient matrix while considering the weak supervision. To obtain the second affinity matrix, we propose the ambiguously supervised structural metric learning (ASML) method to learn the discriminative distance metric by effectively using weakly supervised information.

C. Learning Discriminative Affinity Matrix With Regularized Low-Rank Representation (rLRR)

We first give a brief review of LRR, and then present the proposed method that introduces a discriminative regularizer into the objective of LRR.

1) *Brief Review of LRR*: LRR [2] was originally proposed to solve the *subspace clustering* problem, which aims to explore the subspace structure in the given data $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n] \in \mathbb{R}^{d \times n}$. Based on the assumption that the subspaces are linearly independent, LRR [2] seeks a reconstruction matrix $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_n] \in \mathbb{R}^{n \times n}$, where each \mathbf{w}_i denotes the representation of \mathbf{x}_i using \mathbf{X} (i.e., the data matrix itself) as the dictionary. Since \mathbf{X} is used as the dictionary to reconstruct itself, the optimal solution \mathbf{W}^* of LRR encodes the pairwise affinities between the data samples. As discussed in [2, Th. 3.1], in the noise-free case, \mathbf{W}^* should be ideally block diagonal, where $W_{i,j}^* \neq 0$ if the i th sample and the j th sample are in the same subspace.

Specifically, the optimization problem of LRR is as follows:

$$\min_{\mathbf{W}, \mathbf{E}} \|\mathbf{W}\|_* + \lambda \|\mathbf{E}\|_{2,1} \quad \text{s.t. } \mathbf{X} = \mathbf{X}\mathbf{W} + \mathbf{E} \quad (3)$$

where $\lambda > 0$ is a tradeoff parameter, $\mathbf{E} \in \mathbb{R}^{d \times n}$ is the reconstruction error, the nuclear norm $\|\mathbf{W}\|_*$ (i.e., the sum of all singular values of \mathbf{W}) is adopted to replace $\text{rank}(\mathbf{W})$ as commonly used in the rank minimization problems, and $\|\mathbf{E}\|_{2,1} = \sum_{j=1}^n (\sum_{i=1}^d (E_{i,j})^2)^{1/2}$ is a regularizer to encourage the reconstruction error \mathbf{E} to be column-wise sparse. As mentioned in [2], compared with the sparse representation (SR) method that encourages the sparsity using the ℓ_1 norm, LRR is better at handling the global structures and correcting the corruptions in data automatically. Mathematically, the nuclear norm is nonseparable with respect to the columns, which is different from the ℓ_1 norm. This good property of the nuclear norm is helpful for grasping the global structure and making the model more robust. The toy experiments in [2, Sec. 4. 1] also clearly demonstrate that LRR outperforms SR (which adopts the ℓ_1 norm). Similarly in many real-world applications such as face clustering, LRR usually achieves better results than the sparse subspace clustering [22] method (see [2], [23], and [24] for more details).

2) *LRR With a Discriminative Regularization*: In (3), LRR learns the coefficient matrix \mathbf{W} in an unsupervised way. In our face naming problem, although the names from captions are ambiguous and noisy, they still provide us with the weak

supervision information that is useful for improving the performance of face naming. For example, if two faces do not share any common name in their related candidate name sets, it is unlikely that they are from the same subject, so we should enforce the corresponding entries in \mathbf{W} to be zeros or close to zeros.

Based on this motivation, we introduce a new regularization term $\|\mathbf{W} \circ \mathbf{H}\|_F^2$ by incorporating the weak supervised information, where $\mathbf{H} \in \{0, 1\}^{n \times n}$ is defined based on the candidate name sets $\{\mathcal{N}_i^m\}_{i=1}^m$. Specifically, the entry $H_{i,j} = 0$ if the following two conditions are both satisfied: 1) the i th face and the j th face share at least one common name in the corresponding candidate name sets and 2) $i \neq j$. Otherwise, $H_{i,j} = 1$. In this way, we penalize the nonzero entries in \mathbf{W} , where the corresponding pair of faces do not share any common names in their candidate name sets, and meanwhile, we penalize the entries corresponding to the situations where a face is reconstructed by itself.

As a result, with weak supervision information encoded in \mathbf{H} , the resultant coefficient matrix \mathbf{W} is expected to be more discriminative. By introducing the new regularizer $\|\mathbf{W} \circ \mathbf{H}\|_F^2$ into LRR, we arrive at a new optimization problem as follows:

$$\begin{aligned} \min_{\mathbf{W}, \mathbf{E}} \quad & \|\mathbf{W}\|_* + \lambda \|\mathbf{E}\|_{2,1} + \frac{\gamma}{2} \|\mathbf{W} \circ \mathbf{H}\|_F^2 \\ \text{s.t.} \quad & \mathbf{X} = \mathbf{XW} + \mathbf{E} \end{aligned} \quad (4)$$

where $\gamma \geq 0$ is a parameter to balance the new regularizer with the other terms. We refer to the above problem as rLRR. The rLRR problem in (4) can reduce to the LRR problem in (3) by setting the parameter γ to zero. The visual results for the resultant \mathbf{W} from rLRR and the one from LRR can be found in Fig. 2 (Section V-A).

Once we obtain the optimum solution \mathbf{W}^* after solving (5), the affinity matrix \mathbf{A}_W can be computed as $\mathbf{A}_W = \frac{1}{2}(\mathbf{W}^* + \mathbf{W}^{*'})$, similarly as in [2], and \mathbf{A}_W is further normalized to be within the range of [0, 1].

3) *Optimization*: The optimization problem in (4) can be solved similarly as in LRR [2]. Specifically, we introduce an intermediate variable \mathbf{J} to convert the problem in (4) into the following equivalent problem:

$$\begin{aligned} \min_{\mathbf{W}, \mathbf{E}, \mathbf{J}} \quad & \|\mathbf{J}\|_* + \lambda \|\mathbf{E}\|_{2,1} + \frac{\gamma}{2} \|\mathbf{W} \circ \mathbf{H}\|_F^2 \\ \text{s.t.} \quad & \mathbf{X} = \mathbf{XW} + \mathbf{E}, \quad \mathbf{W} = \mathbf{J}. \end{aligned} \quad (5)$$

Using the augmented Lagrange multiplier (ALM) method, we consider the following augmented Lagrangian function:

$$\begin{aligned} \mathcal{L} = & \|\mathbf{J}\|_* + \lambda \|\mathbf{E}\|_{2,1} + \frac{\gamma}{2} \|\mathbf{W} \circ \mathbf{H}\|_F^2 + \langle \mathbf{U}, \mathbf{X} - \mathbf{XW} - \mathbf{E} \rangle \\ & + \langle \mathbf{V}, \mathbf{W} - \mathbf{J} \rangle + \frac{\rho}{2} (\|\mathbf{X} - \mathbf{XW} - \mathbf{E}\|_F^2 + \|\mathbf{W} - \mathbf{J}\|_F^2) \end{aligned} \quad (6)$$

where $\mathbf{U} \in \mathbb{R}^{d \times n}$ and $\mathbf{V} \in \mathbb{R}^{n \times n}$ are the Lagrange multipliers, and ρ is a positive penalty parameter. Following [2], we solve this problem using inexact ALM [25], which iteratively update the variables, the Lagrange multipliers, and the penalty parameter until convergence is achieved. Specifically, we set $\mathbf{W}_0 = (1/n)(\mathbf{1}_n \mathbf{1}_n' - \mathbf{H})$, $\mathbf{E}_0 = \mathbf{X} - \mathbf{XW}_0$, and $\mathbf{J}_0 = \mathbf{W}_0$, and we set $\mathbf{U}_0, \mathbf{V}_0$ as zero matrices. Then at the t th iteration,

the following steps are performed until convergence is achieved.

- 1) Fix the others and update \mathbf{J}_{t+1} by

$$\min_{\mathbf{J}_{t+1}} \|\mathbf{J}_{t+1}\|_* + \frac{\rho_t}{2} \left\| \mathbf{J}_{t+1} - \left(\mathbf{W}_t + \frac{\mathbf{V}_t}{\rho_t} \right) \right\|_F^2$$

which can be solved in closed form using the singular value thresholding method in [26].

- 2) Fix the others and update \mathbf{W}_{t+1} by

$$\begin{aligned} \min_{\mathbf{W}_{t+1}} \quad & \frac{\gamma}{2} \|\mathbf{W}_{t+1} \circ \mathbf{H}\|_F^2 + \langle \mathbf{U}_t, \mathbf{X} - \mathbf{XW}_{t+1} - \mathbf{E}_t \rangle \\ & + \langle \mathbf{V}_t, \mathbf{W}_{t+1} - \mathbf{J}_{t+1} \rangle + \frac{\rho_t}{2} \|\mathbf{X} - \mathbf{XW}_{t+1} - \mathbf{E}_t\|_F^2 \\ & + \frac{\rho_t}{2} \|\mathbf{W}_{t+1} - \mathbf{J}_{t+1}\|_F^2. \end{aligned} \quad (7)$$

Due to the new regularizer $\|\mathbf{W} \circ \mathbf{H}\|_F^2$, this problem cannot be solved as in [2] by using precomputed SVD. We use the gradient descent method to efficiently solve (7), where the gradient with respect to \mathbf{W}_{t+1} is

$$\begin{aligned} & \gamma (\mathbf{H} \circ \mathbf{H}) \circ \mathbf{W}_{t+1} + \rho_t (\mathbf{X}'\mathbf{X} + \mathbf{I}) \mathbf{W}_{t+1} \\ & + \mathbf{V}_t - \rho_t \mathbf{J}_{t+1} - \mathbf{X}'(\rho_t (\mathbf{X} - \mathbf{E}_t) + \mathbf{U}_t). \end{aligned}$$

- 3) Fix the others and update \mathbf{E}_{t+1} by

$$\begin{aligned} \min_{\mathbf{E}_{t+1}} \quad & \frac{\lambda}{\rho_t} \|\mathbf{E}_{t+1}\|_{2,1} + \frac{1}{2} \left\| \mathbf{E}_{t+1} - \left(\mathbf{X} - \mathbf{XW}_{t+1} + \frac{\mathbf{U}_t}{\rho_t} \right) \right\|_F^2 \\ & \text{which can be solved in closed form based on [27, Lemma 4.1].} \end{aligned}$$

- 4) Update \mathbf{U}_{t+1} and \mathbf{V}_{t+1} by respectively using

$$\begin{aligned} \mathbf{U}_{t+1} &= \mathbf{U}_t + \rho_t (\mathbf{X} - \mathbf{XW}_{t+1} - \mathbf{E}_{t+1}) \\ \mathbf{V}_{t+1} &= \mathbf{V}_t + \rho_t (\mathbf{W}_{t+1} - \mathbf{J}_{t+1}). \end{aligned}$$

- 5) Update ρ_{t+1} using

$$\rho_{t+1} = \min(\rho_t(1 + \Delta\rho), \rho_{\max})$$

where $\Delta\rho$ and ρ_{\max} are the constant parameters.

- 6) The iterative algorithm stops if the two convergence conditions are both satisfied

$$\begin{aligned} \|\mathbf{X} - \mathbf{XW}_{t+1} - \mathbf{E}_{t+1}\|_\infty &\leq \epsilon \\ \|\mathbf{W}_{t+1} - \mathbf{J}_{t+1}\|_\infty &\leq \epsilon \end{aligned}$$

where ϵ is a small constant parameter.

D. Learning Discriminative Affinity Matrix by Ambiguously Supervised Structural Metric Learning (ASML)

Besides obtaining the affinity matrix from the coefficient matrix \mathbf{W}^* from rLRR (or LRR), we believe the similarity matrix (i.e., the kernel matrix) among the faces is also an appropriate choice for the affinity matrix. Instead of straightforwardly using the Euclidean distances, we seek a discriminative Mahalanobis distance metric \mathbf{M} so that Mahalanobis distances can be calculated based on the learnt metric, and the similarity matrix can be obtained based on the Mahalanobis distances. In the following, we first briefly review the LMNN method, which deals with fully-supervised problems with the ground-truth labels of samples provided, and then introduce our proposed ASML method that extends LMNN for face naming from weakly labeled images.

1) *Brief Review of LMNN*: Most existing metric learning methods deal with the supervised learning problems [16], [28] where the ground-truth labels of the training samples are given. Weinberger and Saul [16] proposed the LMNN method to learn a distance metric \mathbf{M} that encourages the squared Mahalanobis distances between each training sample and its target neighbors (e.g., the k nearest neighbors) to be smaller than those between this training sample and training samples from other classes. Let $\{(\mathbf{x}_i, y_i)|_{i=1}^n\}$ be the n labeled samples, where $\mathbf{x}_i \in \mathbb{R}^d$ denotes the i th sample, with d being the feature dimension, and $y_i \in \{1, \dots, z\}$ denotes the label of this sample, with z being the total number of classes. $\eta_{i,j} \in \{0, 1\}$ indicates whether \mathbf{x}_j is a target neighbor of \mathbf{x}_i , namely, $\eta_{i,j} = 1$ if \mathbf{x}_j is a target neighbor of \mathbf{x}_i , and $\eta_{i,j} = 0$ otherwise, $\forall i, j \in \{1, \dots, n\}$. $v_{i,l} \in \{0, 1\}$ indicates whether \mathbf{x}_l and \mathbf{x}_i are from different classes, namely, $v_{i,l} = 1$ if $y_l \neq y_i$, and $v_{i,l} = 0$ otherwise, $\forall i, l \in \{1, \dots, n\}$. The squared Mahalanobis distance between two samples \mathbf{x}_i and \mathbf{x}_j is defined as

$$d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_j) = (\mathbf{x}_i - \mathbf{x}_j)' \mathbf{M} (\mathbf{x}_i - \mathbf{x}_j).$$

LMNN minimizes the following optimization problem:

$$\begin{aligned} \min_{\mathbf{M} \geq 0} \quad & \sum_{(i,j): \eta_{i,j}=1} d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_j) + \mu \sum_{(i,j,l) \in \mathcal{S}} \xi_{i,j,l} \\ \text{s.t.} \quad & d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_l) - d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_j) \geq 1 - \xi_{i,j,l} \quad \forall (i, j, l) \in \mathcal{S} \\ & \xi_{i,j,l} \geq 0 \quad \forall (i, j, l) \in \mathcal{S} \end{aligned} \quad (8)$$

where μ is a tradeoff parameter, $\xi_{i,j,l}$ is a slack variable, and $\mathcal{S} = \{(i, j, l) | \eta_{i,j} = 1, v_{i,l} = 1, \forall i, j, l \in \{1, \dots, n\}\}$. Therefore, $d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_j)$ is the squared Mahalanobis distance between \mathbf{x}_i and its target neighbor \mathbf{x}_j , and $d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_l)$ is the squared Mahalanobis distance between \mathbf{x}_i and \mathbf{x}_l that belong to different classes. The difference between $d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_l)$ and $d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_j)$ is expected to be no less than one in the ideal case. The introduction of the slack variable $\xi_{i,j,l}$ can also tolerate the cases when $d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_l) - d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_j)$ is slightly smaller than one, which is similar to the one in soft margin SVM for tolerating the classification error. The LMNN problem in (8) can be equivalently rewritten as the following optimization problem:

$$\begin{aligned} \min_{\mathbf{M} \geq 0} \quad & \sum_{(i,j): \eta_{i,j}=1} d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_j) \\ & + \mu \sum_{(i,j,l) \in \mathcal{S}} |1 - d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_l) + d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_j)|_+ \end{aligned}$$

with $|\cdot|_+$ being the truncation function, i.e., $|x|_+ = \max(0, x)$.

2) *Ambiguously Supervised Structural Metric Learning*: In the face naming problem, the ground-truth names of the faces are not available, so LMNN cannot be applied to solve the problem. Fortunately, weak supervision information is available in the captions along with each image; hence, we propose a new distance metric learning method called ASML to utilize such weakly supervised information.

Recall that we should consider the image-level constraints when inferring the names of faces in the same image. Therefore, we design the losses with respect to each image, by considering the image-level constraints in the feasible label sets $\{\mathcal{Y}^i|_{i=1}^m\}$ defined in (1).

Let us take the i th image for example. The faces in the i th image are $\{\mathbf{x}_f^i|_{f=1}^{n_i}\}$. Let \mathbf{Y}_*^i be the ground-truth label matrix for the faces in the i th image, which is in the feasible label sets \mathcal{Y}^i . Let $\bar{\mathbf{Y}}^i$ be an infeasible label matrix for the faces in the i th image, which is contained in the infeasible label set $\bar{\mathcal{Y}}^i$. Note the infeasible label set $\bar{\mathcal{Y}}^i$ is the set of label matrices that is excluded in \mathcal{Y}^i and, meanwhile, satisfies the nonredundancy constraint

$$\bar{\mathcal{Y}}^i = \left\{ \bar{\mathbf{Y}}^i \in \{0, 1\}^{(p+1) \times n_i} \middle| \begin{array}{l} \bar{\mathbf{Y}}^i \notin \mathcal{Y}^i, \\ \mathbf{1}_{(p+1)}' \bar{\mathbf{Y}}^i = \mathbf{1}_{n_i}' \end{array} \right\}.$$

Assume that the face \mathbf{x}_f^i is labeled as the name q according to a label matrix, we define face to name (F2N) distance $D_{F2N}(\mathbf{x}_f^i, q, \mathbf{M})$ to measure the disagreement between the face \mathbf{x}_f^i and the name q . Specifically, $D_{F2N}(\mathbf{x}_f^i, q, \mathbf{M})$ is defined as follows:

$$D_{F2N}(\mathbf{x}_f^i, q, \mathbf{M}) = \frac{1}{|\mathcal{X}_q|} \sum_{\tilde{\mathbf{x}} \in \mathcal{X}_q} d_{\mathbf{M}}^2(\mathbf{x}_f^i, \tilde{\mathbf{x}})$$

where $d_{\mathbf{M}}^2(\mathbf{x}_f^i, \tilde{\mathbf{x}})$ is the squared Mahalanobis distance between \mathbf{x}_f^i and $\tilde{\mathbf{x}}$, \mathcal{X}_q is the set of all the faces from the images with each image associated with the name q , and $|\mathcal{X}_q|$ is the cardinality of \mathcal{X}_q . Intuitively, $D_{F2N}(\mathbf{x}, q, \mathbf{M})$ should be small if q is the ground-truth name of the face \mathbf{x} , and $D_{F2N}(\mathbf{x}, q, \mathbf{M})$ should be large otherwise. Recall that in LMNN, we expect $d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_j)$ (i.e., the squared Mahalanobis distance between \mathbf{x}_i and its target neighbor \mathbf{x}_j) to be somehow smaller than $d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_l)$ (i.e., the squared Mahalanobis distance between \mathbf{x}_i and \mathbf{x}_l that belong to different classes). Similarly, we expect that $D_{F2N}(\mathbf{x}_f^i, q, \mathbf{M})$ should be smaller than $D_{F2N}(\mathbf{x}_f^i, \bar{q}, \mathbf{M})$ to some extent, where q is the assigned name of \mathbf{x}_f^i according to the ground-truth label matrix \mathbf{Y}_*^i , and \bar{q} is the assigned name of \mathbf{x}_f^i according to an infeasible label matrix $\bar{\mathbf{Y}}^i$. For all the faces in the i th image and a label matrix \mathbf{Y}^i , we define the I2A distance $D(\mathbf{X}^i, \mathbf{Y}^i, \mathbf{M})$ to be the sum of F2N distances between every face and its assigned names. Mathematically, $D(\mathbf{X}^i, \mathbf{Y}^i, \mathbf{M})$ is defined as

$$D(\mathbf{X}^i, \mathbf{Y}^i, \mathbf{M}) = \sum_{f=1}^{n_i} \sum_{q: \mathbf{Y}_{q,f}^i=1} D_{F2N}(\mathbf{x}_f^i, q, \mathbf{M}).$$

In the ideal case, we expect that $D(\mathbf{X}^i, \mathbf{Y}_*^i, \mathbf{M})$ should be smaller than $D(\mathbf{X}^i, \bar{\mathbf{Y}}^i, \mathbf{M})$ by at least $h(\bar{\mathbf{Y}}^i, \mathbf{Y}_*^i)$, where $h(\bar{\mathbf{Y}}^i, \mathbf{Y}_*^i)$ is the number of faces that are assigned with different names based on two label matrices $\bar{\mathbf{Y}}^i$ and \mathbf{Y}_*^i . To tolerate the cases where $D(\mathbf{X}^i, \bar{\mathbf{Y}}^i, \mathbf{M}) - D(\mathbf{X}^i, \mathbf{Y}_*^i, \mathbf{M})$ is slightly smaller than $h(\bar{\mathbf{Y}}^i, \mathbf{Y}_*^i)$, we introduce a *nonnegative* slack variable ξ_i for the i th image and have the following constraint for any $\bar{\mathbf{Y}}^i \in \bar{\mathcal{Y}}^i$:

$$D(\mathbf{X}^i, \bar{\mathbf{Y}}^i, \mathbf{M}) - D(\mathbf{X}^i, \mathbf{Y}_*^i, \mathbf{M}) \geq h(\bar{\mathbf{Y}}^i, \mathbf{Y}_*^i) - \xi_i. \quad (9)$$

However, the groundtruth label matrix \mathbf{Y}_*^i is unknown, so $h(\bar{\mathbf{Y}}^i, \mathbf{Y}_*^i)$ and $D(\mathbf{X}^i, \mathbf{Y}_*^i, \mathbf{M})$ in (9) are not available. Although \mathbf{Y}_*^i is unknown, it should be a label matrix in the feasible label set \mathcal{Y}^i . In this paper, we use $\ell(\bar{\mathbf{Y}}^i, \mathcal{Y}^i)$ to approximate $h(\bar{\mathbf{Y}}^i, \mathbf{Y}_*^i)$, where $\ell(\bar{\mathbf{Y}}^i, \mathcal{Y}^i)$ measures the difference between

an infeasible label matrix $\bar{\mathbf{Y}}^i$ and the most similar label matrix in \mathcal{Y}_i . Similarly as in [7], we define $\ell(\bar{\mathbf{Y}}^i, \mathcal{Y}^i)$ as follows:

$$\ell(\bar{\mathbf{Y}}^i, \mathcal{Y}^i) = \min_{\mathbf{Y}^i \in \mathcal{Y}^i} h(\bar{\mathbf{Y}}^i, \mathbf{Y}^i).$$

On the other hand, since \mathbf{Y}_*^i is in the feasible label set \mathcal{Y}^i and we expect the corresponding I2A distance should be small, we use $\min_{\mathbf{Y}^i \in \mathcal{Y}^i} D(\mathbf{X}^i, \mathbf{Y}^i, \mathbf{M})$ to replace $D(\mathbf{X}^i, \mathbf{Y}_*^i, \mathbf{M})$, where $\min_{\mathbf{Y}^i \in \mathcal{Y}^i} D(\mathbf{X}^i, \mathbf{Y}^i, \mathbf{M})$ is the smallest I2A distance based on the feasible label matrix inside \mathcal{Y}_i . In summary, by replacing $h(\bar{\mathbf{Y}}^i, \mathbf{Y}_*^i)$ and $D(\mathbf{X}^i, \mathbf{Y}_*^i, \mathbf{M})$ with $\ell(\bar{\mathbf{Y}}^i, \mathcal{Y}^i)$ and $\min_{\mathbf{Y}^i \in \mathcal{Y}^i} D(\mathbf{X}^i, \mathbf{Y}^i, \mathbf{M})$, respectively, the constraint in (9) becomes the following one for any $\bar{\mathbf{Y}}^i \in \bar{\mathcal{Y}}^i$:

$$\xi_i \geq \ell(\bar{\mathbf{Y}}^i, \mathcal{Y}^i) - D(\mathbf{X}^i, \bar{\mathbf{Y}}^i, \mathbf{M}) + \min_{\mathbf{Y}^i \in \mathcal{Y}^i} D(\mathbf{X}^i, \mathbf{Y}^i, \mathbf{M}). \quad (10)$$

Instead of enforcing ξ_i to be no less than every $\ell(\bar{\mathbf{Y}}^i, \mathcal{Y}^i) - D(\mathbf{X}^i, \bar{\mathbf{Y}}^i, \mathbf{M}) + \min_{\mathbf{Y}^i \in \mathcal{Y}^i} D(\mathbf{X}^i, \mathbf{Y}^i, \mathbf{M})$ (each based on an infeasible label matrix $\bar{\mathbf{Y}}^i$ in $\bar{\mathcal{Y}}^i$) as in (10), we can equivalently enforce ξ_i to be no less than the largest one of them. Note that the term $\min_{\mathbf{Y}^i \in \mathcal{Y}^i} D(\mathbf{X}^i, \mathbf{Y}^i, \mathbf{M})$ is irrelevant to $\bar{\mathbf{Y}}^i$. Accordingly, we rewrite (10) with respect to the nonnegative slack variable ξ_i in the following equivalent form:

$$\xi_i \geq \max_{\bar{\mathbf{Y}}^i \in \bar{\mathcal{Y}}^i} [\ell(\bar{\mathbf{Y}}^i, \mathcal{Y}^i) - D(\mathbf{X}^i, \bar{\mathbf{Y}}^i, \mathbf{M})] + \min_{\mathbf{Y}^i \in \mathcal{Y}^i} D(\mathbf{X}^i, \mathbf{Y}^i, \mathbf{M}).$$

Hence, we propose a new method called ASML to learn a discriminative Mahalanobis distance metric \mathbf{M} by solving the following problem:

$$\min_{\mathbf{M} \geq 0} \frac{\sigma}{2} \|\mathbf{M} - \mathbf{I}\|_F^2 + \frac{1}{m} \sum_{i=1}^m |\max_{\bar{\mathbf{Y}}^i \in \bar{\mathcal{Y}}^i} [\ell(\bar{\mathbf{Y}}^i, \mathcal{Y}^i) - D(\mathbf{X}^i, \bar{\mathbf{Y}}^i, \mathbf{M})] + \min_{\mathbf{Y}^i \in \mathcal{Y}^i} D(\mathbf{X}^i, \mathbf{Y}^i, \mathbf{M})|_+. \quad (11)$$

where $\sigma > 0$ is a tradeoff parameter and the regularizer $\|\mathbf{M} - \mathbf{I}\|_F^2$ is used to enforce \mathbf{M} to be not too far away from the identity matrix \mathbf{I} , and we also rewrite ξ_i as $|\max_{\bar{\mathbf{Y}}^i \in \bar{\mathcal{Y}}^i} [\ell(\bar{\mathbf{Y}}^i, \mathcal{Y}^i) - D(\mathbf{X}^i, \bar{\mathbf{Y}}^i, \mathbf{M})] + \min_{\mathbf{Y}^i \in \mathcal{Y}^i} D(\mathbf{X}^i, \mathbf{Y}^i, \mathbf{M})|_+$, similarly to that in LMNN. Note that we have incorporated weak supervision information in the max margin loss in (11). A nice property of such max margin loss is the robustness to label noise.

Optimization: Since $\min_{\mathbf{Y}^i \in \mathcal{Y}^i} D(\mathbf{X}^i, \mathbf{Y}^i, \mathbf{M})$ in (11) is concave, the objective function in (11) is nonconvex with respect to \mathbf{M} . For convenience, we define two convex functions $f_i(\mathbf{M}) = \max_{\bar{\mathbf{Y}}^i \in \bar{\mathcal{Y}}^i} [\ell(\bar{\mathbf{Y}}^i, \mathcal{Y}^i) - D(\mathbf{X}^i, \bar{\mathbf{Y}}^i, \mathbf{M})]$ and $g_i(\mathbf{M}) = -\min_{\mathbf{Y}^i \in \mathcal{Y}^i} D(\mathbf{X}^i, \mathbf{Y}^i, \mathbf{M})$, $\forall i = 1, \dots, m$. Inspired by the concave-convex procedure (CCCP) method [29], we equivalently rewrite (11) as follows:

$$\min_{\mathbf{M} \geq 0} \frac{\sigma}{2} \|\mathbf{M} - \mathbf{I}\|_F^2 + \frac{1}{m} \sum_{i=1}^m |f_i(\mathbf{M}) - g_i(\mathbf{M})|_+. \quad (12)$$

We solve the problem in (12) in an iterative fashion. Let us denote \mathbf{M} at the s th iteration as $\mathbf{M}_{(s)}$. Similarly as in CCCP, at the $(s+1)$ th iteration, we replace the nonconvex term $|f_i(\mathbf{M}) - g_i(\mathbf{M})|_+$ with a convex term $|f_i(\mathbf{M}) - \langle \mathbf{M}, \tilde{g}_i(\mathbf{M}_{(s)}) \rangle|_+$, where $\tilde{g}_i(\cdot)$ is the subgradient [7] of $g_i(\cdot)$. Hence, at the $(s+1)$ th iteration, we solve the following relaxed version of the problem in (12):

$$\min_{\mathbf{M} \geq 0} \frac{\sigma}{2} \|\mathbf{M} - \mathbf{I}\|_F^2 + \frac{1}{m} \sum_{i=1}^m |f_i(\mathbf{M}) - \langle \mathbf{M}, \tilde{g}_i(\mathbf{M}_{(s)}) \rangle|_+ \quad (13)$$

Algorithm 1 ASML Algorithm

Input: The training images $\{\mathbf{X}^i\}_{i=1}^m$, the feasible label sets $\{\mathcal{Y}^i\}_{i=1}^m$, the parameters σ , N_{iter} and ε .

- 1: Initialize¹ $\mathbf{M}_{(0)} = \mathbf{I}$.
- 2: **for** $s = 1 : N_{iter}$ **do**
- 3: Calculate $\mathbf{Q}_{(s)}$ as $\mathbf{Q}_{(s)} = \mathbf{M}_{(s)} - \mathbf{I}$.
- 4: Obtain $\mathbf{Q}_{(s+1)}$ by solving the convex problem in (14) via the stochastic subgradient descent method.
- 5: Calculate $\mathbf{M}_{(s+1)}$ as $\mathbf{M}_{(s+1)} = \mathbf{Q}_{(s+1)} + \mathbf{I}$.
- 6: break if $\|\mathbf{M}_{(s+1)} - \mathbf{M}_{(s)}\|_F \leq \varepsilon$.
- 7: **end for**

Output: the Mahalanobis distance metric $\mathbf{M}_{(s+1)}$.

which is now convex with respect to \mathbf{M} . To solve (13), we define $\mathbf{Q} = \mathbf{M} - \mathbf{I}$ and $\mathbf{Q}_{(s)} = \mathbf{M}_{(s)} - \mathbf{I}$, and equivalently rewrite (13) as the following convex optimization problem:

$$\begin{aligned} \min_{\mathbf{Q}, \tilde{\xi}_i} \quad & \frac{\sigma}{2} \|\mathbf{Q}\|_F^2 + \frac{1}{m} \sum_{i=1}^m \tilde{\xi}_i \\ \text{s.t.} \quad & f_i(\mathbf{Q} + \mathbf{I}) - \langle \mathbf{Q} + \mathbf{I}, \tilde{g}_i(\mathbf{Q}_{(s)} + \mathbf{I}) \rangle \leq \tilde{\xi}_i, \quad \tilde{\xi}_i \geq 0 \quad \forall i \\ & \mathbf{Q} + \mathbf{I} \succeq \mathbf{0}. \end{aligned} \quad (14)$$

Although the optimization problem in (14) is convex, it may contain many constraints. To efficiently solve it, we adopt the stochastic subgradient descent method similarly as in Pegasos [30]. Moreover, to handle the PSD constraint on $\mathbf{Q} + \mathbf{I}$ in (14), at each iteration when using the stochastic subgradient descent method, we additionally project the solution onto the PSD cone by thresholding the negative eigenvalues to be zeros, similarly as in [31]. The ASML algorithm is summarized in Algorithm 1.

IV. INFERRING NAMES OF FACES

With the coefficient matrix \mathbf{W}^* learned from rLRR, we can calculate the first affinity matrix as $\mathbf{A}_W = \frac{1}{2}(\mathbf{W}^* + \mathbf{W}^{*'})$ and normalize \mathbf{A}_W to the range $[0, 1]$. Furthermore, with the learnt distance metric \mathbf{M} from ASML, we can calculate the second affinity matrix as $\mathbf{A}_K = \mathbf{K}$, where \mathbf{K} is a kernel matrix based on the Mahalanobis distances between the faces. Since the two affinity matrices explore weak supervision information in different ways, they contain complementary information and both of them are beneficial for face naming. For better face naming performance, we combine these two affinity matrices and perform face naming based on the fused affinity matrix. Specifically, we obtain a fused affinity matrix \mathbf{A} as the linear combination of the two affinity matrices, i.e., $\mathbf{A} = (1 - \alpha)\mathbf{A}_W + \alpha\mathbf{A}_K$, where α is a parameter in the range $[0, 1]$. Finally, we perform face naming based on \mathbf{A} . Since the fused affinity matrix is obtained based on rLRR and ASML, we name our proposed method as rLRRml. As mentioned in Section III-B, given this affinity matrix \mathbf{A} , we perform face naming by solving the following optimization problem:

$$\max_{\mathbf{Y} \in \mathcal{Y}} \sum_{c=1}^p \frac{\mathbf{y}_c' \mathbf{A} \mathbf{y}_c}{\mathbf{1}' \mathbf{y}_c}, \quad \text{s.t. } \mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_{(p+1)}]'. \quad (15)$$

¹Our experiments show that the results using this initialization are comparable with those using random initialization.

However, the above problem is an integer programming problem, which is computationally expensive to solve. In this paper, we propose an iterative approach to solve a relaxed version of (15). Specifically, at each iteration, we approximate the objective function by using $\tilde{\mathbf{y}}_c' \mathbf{A} \mathbf{y}_c / \mathbf{1}' \tilde{\mathbf{y}}_c$ to replace $\mathbf{y}_c' \mathbf{A} \mathbf{y}_c / \mathbf{1}' \mathbf{y}_c$, where $\tilde{\mathbf{y}}_c$ is the solution of \mathbf{y}_c obtained from the previous iteration. Hence, at each iteration, we only need to solve a linear programming problem as follows:

$$\max_{\mathbf{Y} \in \mathcal{Y}} \sum_{c=1}^p \mathbf{b}_c' \mathbf{y}_c, \quad \text{s.t. } \mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_{(p+1)}]' \quad (16)$$

where $\mathbf{b}_c = \mathbf{A} \tilde{\mathbf{y}}_c / \mathbf{1}' \tilde{\mathbf{y}}_c$, $\forall c = 1, \dots, p$. Moreover, the candidate name set \mathcal{N}^i may be incomplete, so some faces in the image \mathbf{X}^i may not have the corresponding ground-truth names in the candidate name set \mathcal{N}^i . Therefore, similarly as in [32], we additionally define a vector $\mathbf{b}_{p+1} = \theta \mathbf{1}$ to allow some faces to be assigned to the null class, where θ is a predefined parameter. Intuitively, the number of faces assigned to null changes when we set θ with different values. In the experiments, to fairly compare the proposed methods and other methods, we report the performances of all methods when each algorithm annotates the same number of faces using real names rather than null, which can be achieved by tuning the parameter θ (see Section V-C for more details).

By defining $\mathbf{B} \in \mathbb{R}^{(p+1) \times n}$ as $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_{p+1}]'$, we can reformulate the problem in (16) as follows:

$$\max_{\mathbf{Y} \in \mathcal{Y}} \langle \mathbf{B}, \mathbf{Y} \rangle. \quad (17)$$

Recall that the feasible set for \mathbf{Y} is defined as $\mathcal{Y} = \{\mathbf{Y} = [\mathbf{Y}^1, \dots, \mathbf{Y}^m] | \mathbf{Y}^i \in \mathcal{Y}^i, \forall i = 1, \dots, m\}$, which means the constraints on \mathbf{Y}^i s are separable. Let us decompose the matrix \mathbf{B} as $\mathbf{B} = [\mathbf{B}^1, \dots, \mathbf{B}^m]$ with each $\mathbf{B}^i \in \mathbb{R}^{(p+1) \times n_i}$ corresponding to \mathbf{Y}^i , then the objective function in (17) can be expressed as $\langle \mathbf{B}, \mathbf{Y} \rangle = \sum_{i=1}^m \langle \mathbf{B}^i, \mathbf{Y}^i \rangle$, which is also separable with respect to \mathbf{Y}^i s. Hence, we optimize (17) by solving m subproblems, with each subproblem related to one image in the following form:

$$\max_{\mathbf{Y}^i \in \mathcal{Y}^i} \langle \mathbf{B}^i, \mathbf{Y}^i \rangle \quad (18)$$

$\forall i = 1, \dots, m$. In particular, the i th problem in (18) can equivalently be rewritten as a *minimization* problem with detailed constraints as follows:

$$\begin{aligned} \min_{Y_{q,f}^i \in \{0,1\}} & \sum_{q \in \tilde{\mathcal{N}}^i} \sum_{f=1}^{n_i} -B_{q,f}^i Y_{q,f}^i \\ \text{s.t.} & \sum_{q \in \tilde{\mathcal{N}}^i} Y_{q,f}^i = 1 \quad \forall f = 1, \dots, n_i \\ & \sum_{f=1}^{n_i} Y_{q,f}^i \leq 1 \quad \forall q \in \mathcal{N}^i \\ & \sum_{f=1}^{n_i} Y_{(p+1),f}^i \leq n_i \end{aligned} \quad (19)$$

in which we have dropped the elements $\{Y_{q,f}^i | q \notin \tilde{\mathcal{N}}^i\}$, because these elements are zeros according to the feasibility

Algorithm 2 Face Naming Algorithm

Input: The feasible label sets $\{\mathcal{Y}^i |_{i=1}^m\}$, the affinity matrix \mathbf{A} , the initial label matrix $\mathbf{Y}(1)$ and the parameters \tilde{N}_{iter} , θ .

- 1: **for** $t = 1 : \tilde{N}_{\text{iter}}$ **do**
- 2: Update \mathbf{B} by using $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_{p+1}]'$, where $\mathbf{b}_c = \frac{\mathbf{A} \tilde{\mathbf{y}}_c}{\mathbf{1}' \tilde{\mathbf{y}}_c}$, $\forall c = 1, \dots, p$ with $\tilde{\mathbf{y}}_c$ being the c -th column of $\mathbf{Y}(t)$, and $\mathbf{b}_{p+1} = \theta \mathbf{1}$.
- 3: Update $\mathbf{Y}(t+1)$ by solving m subproblems in (19).
- 4: **break** if $\mathbf{Y}(t+1) = \mathbf{Y}(t)$.
- 5: **end for**

Output: the label matrix $\mathbf{Y}(t+1)$.

constraint in (1). Similarly as in [32], we solve the problem in (19) by converting it to a minimum cost bipartite graph matching problem, for which the objective is the sum of the costs for assigning faces to names. In this paper, we adopt the Hungarian algorithm to efficiently solve it. Specifically, for the i th image, the cost $c(f, q)$ for assigning a face \mathbf{x}_f^i to a real name q is set to $-B_{q,f}^i$, and the cost $c(f, p+1)$ for assigning a face \mathbf{x}_f^i to the corresponding null name is set to $-B_{(p+1),f}^i$.

In summary, to infer the label matrix \mathbf{Y} for all faces, we iteratively solve the linear programming problem in (17), which can be efficiently addressed by solving m subproblems as in (19) with the Hungarian algorithm. Let $\mathbf{Y}(t)$ be the label matrix at the t th iteration. The initial label matrix $\mathbf{Y}(1)$ is set to the label matrix that assigns each face to all names in the caption associated with the corresponding image that contains this face. The iterative process continues until the convergence condition is satisfied. In practice, this iterative process always converges in about 10–15 iterations, so we empirically set \tilde{N}_{iter} as 15. The iterative algorithm for face naming is summarized in Algorithm 2.

V. EXPERIMENTS

In this section, we compare our proposed methods rLRR, ASML, and rLRRml with four state-of-the-art algorithms for face naming, as well as two special cases of our proposed methods using a synthetic dataset and two real-world datasets.

A. Introduction of the Datasets

One synthetic dataset and two real-world benchmark datasets are used in the experiments. The synthetic dataset is collected from the Faces and Poses dataset in [33]. We first find out the top 10 popular names and then for each name, we randomly sample 50 images where this name appears in the image tags. In total, the synthetic dataset contains 602 faces in 500 images, with a total number of 20 names appearing in the corresponding tags, which include these top 10 popular names and other names associated with these 500 images.

Other than the synthetic dataset, the experiments are also conducted on the following two real-world datasets.

TABLE I
DETAILS OF THE DATASETS. THE COLUMNS IN TURN ARE THE TOTAL NUMBER OF IMAGES, FACES AND NAMES, THE AVERAGE
NUMBER OF DETECTED FACES PER IMAGE, AND THE AVERAGE NUMBER OF DETECTED NAMES PER
CAPTION AND THE GROUNDTRUTH RATIO, RESPECTIVELY

Dataset	#images	#faces	#names	#faces/image	#names/caption	groundtruth ratio
Synthetic	500	602	20	1.20	1.07	0.89
Soccer player	8640	17472	170	2.02	1.74	0.51
Labeled Yahoo! News	10128	15868	214	1.57	1.37	0.56

1) *Soccer Player Dataset*: This dataset was used in [9], with the images of soccer players from famous European clubs and names mentioned in the captions. The detected faces are manually annotated using names from the captions or as null. Following [9], we retain 170 names that occur at least 20 times in the captions and treat others as the null class. The images without containing any of these 170 names are discarded.

2) *Labeled Yahoo! News Dataset*: This dataset was collected in [34] and further processed in [6]. It contains news images as well as the names in the captions. Following [7] and [9], we retain the 214 names occurred at least 20 times in the captions and treat others as the null class. The images that do not contain any of the 214 names are removed.

The detailed information about the synthetic and real-world datasets is shown in Table I, where the *ground-truth real name ratio* (or *ground-truth ratio* in short) is the percentage of faces whose groundtruth names are real names (rather than null) among all the faces in the dataset. In the Soccer player dataset, there are more images with multiple faces and multiple names in the captions when compared with the Labeled Yahoo! News dataset, which indicates that the Soccer player dataset is more challenging. For the synthetic dataset and the two real-world datasets, we extract the feature vectors to represent the faces in the same way as in [10]. For each face, 13 interest points (facial landmarks) are located. For each interest point, a simple pixel-wised descriptor is formed using the gray-level intensity values of pixels in the elliptical region based on each interest point, which is further normalized to achieve local photometric invariance [10]. Finally, a 1937-D descriptor for each face is obtained by concatenating the descriptors from the 13 interest points.

B. Baseline Methods and Two Special Cases

The following four state-of-the-art methods are used as baselines.

- 1) MMS learning algorithm [7] that solves the face naming problem by learning SVM classifiers for each name.
- 2) MildML [6] that learns a Mahalanobis distance metric such that the bags (images) with common labels (names in captions) are pulled closer, while the bags that do not share any common label are pushed apart.
- 3) Constrained Gaussian mixture model (cGMM) [32], [35]. For this Gaussian mixture model based approach, each name is associated with a Gaussian density function in the feature space with the parameters estimated from the data, and each face is assumed to be independently generated from the associated

Gaussian function. The overall assignments are chosen to achieve the maximum log likelihood.

- 4) LR-SVM [9] that simultaneously learns the partial permutation matrices for grouping the faces and minimize the rank of the data matrices from each group. SVM classifiers are also trained for each name to deal with the out-of-sample cases.

More details of these methods can be found in Section II. For detailed analysis of the proposed rLRRml, we also report the results of the following two special cases.

- 1) *Low Rank Representation With Metric Learning (LRRml for Short)*: rLRRml reduces to LRRml if we do not introduce the proposed regularizer on \mathbf{W} . In other words, we set the parameter γ in rLRR to 0 when learning \mathbf{W} .
- 2) *LRR*: rLRRml reduces to LRR if we neither consider the affinity matrix \mathbf{A}_K nor pose the proposed regularizer on \mathbf{W} . In other words, we set the parameter γ in rLRR to 0 when learning the coefficient matrix \mathbf{W} , and we use \mathbf{A}_W as the input affinity matrix \mathbf{A} in Algorithm 2.

On the synthetic dataset, we empirically set γ to 100 for our rLRR, and we empirically set λ to 0.01 for both LRR and rLRR.² On the real-world datasets, for MMS, we tune the parameter C in the range of $\{1, 10, \dots, 10^4\}$ and report the best results from the optimal C . For MildML, we tune the parameter about the metric rank in the range of $\{2^2, 2^3, \dots, 2^7\}$ and report the best results. For cGMM, there are no parameters to be set. For the parameters λ and C in LR-SVM, instead of fixing $\lambda = 0.3$ and choosing C in the range of $\{0.1, 1, 10\}$, as in [9], we tune these parameters in larger ranges. Specifically, we tune λ in the range of $\{1, 0.3, 0.1, 0.01\}$ and C in the range of $\{10^{-2}, 10^{-1}, \dots, 10^2\}$, and report the best results from the optimal λ and C . The parameter α for fusing the two affinity matrices in rLRRml and LRRml is empirically fixed as 0.1 on both real-world datasets, namely, we calculate \mathbf{A} as $\mathbf{A} = 0.9\mathbf{A}_W + 0.1\mathbf{A}_K$. On the two real-world datasets, after tuning λ in LRR in the range of $\{1, 0.1, 0.01, 0.001\}$, we observe that LRR achieves the best results when setting λ to 0.01 on both datasets, so we fix the parameter λ for LRR, rLRR, LRRml, and rLRRml to 0.01 on both datasets. The parameter γ for rLRR and rLRRml is empirically set to 100, and the tradeoff parameter σ for ASML, LRRml, and rLRRml is empirically fixed to one. For the kernel matrix \mathbf{K} in ASML, LRRml, and rLRRml, we use the kernel matrix based on the Mahalanobis dis-

²We set the parameters ρ_{max} , $\Delta\rho$, and ϵ to the default values in the code from <http://www.columbia.edu/~js4038/software.html>. We set the number of iterations to 20 since the result becomes stable after about 20 iterations.

tances, namely we have $K_{i,j} = \exp(-\sqrt{\nu} D_{\mathbf{M}}(\mathbf{x}_i, \mathbf{x}_j))$, where $D_{\mathbf{M}}(\mathbf{x}_i, \mathbf{x}_j) = ((\mathbf{x}_i - \mathbf{x}_j)' \mathbf{M} (\mathbf{x}_i - \mathbf{x}_j))^{1/2}$ is the Mahalanobis distance between \mathbf{x}_i and \mathbf{x}_j , and ν is the bandwidth parameter set as the default value $1/\beta$, with β being the mean of squared Mahalanobis distances between all samples [36].

C. Experimental Results

1) *Results on the Synthetic Dataset:* First, to validate the effectiveness of our proposed method rLRR for recovering subspace information, we compare the coefficient matrices obtained from LRR and rLRR with the ideal coefficient matrix \mathbf{W}^* according to the groundtruth, as shown in Fig. 2.

Fig. 2(a) shows the ideal coefficient matrix \mathbf{W}^* according to the groundtruth. For better viewing, the faces are reordered by grouping the faces belonging to the same name at contiguous positions. Note the white points indicate that the corresponding faces belong to the same subject (i.e., with the same name), and the bottom-right part corresponds to the faces from the null class. The diagonal entries are set to be zeros since we expect self-reconstruction can be avoided.

Fig. 2(b) shows the coefficient matrix \mathbf{W}^* obtained from LRR. While there exists block-wise diagonal structure to some extent, we also observe the following.

- 1) The diagonal elements are large, meaning that a face is reconstructed mainly by itself. It should be avoided.
- 2) In general, the coefficients between faces from the same subject are not significantly larger than the ones between faces from different subjects.

Fig. 2(c) shows the coefficient matrix \mathbf{W}^* obtained from our rLRR. It has smaller values for the diagonal elements. In general, the coefficients between faces from the same subject become larger, while the ones between faces from different subjects become smaller. Compared with Fig. 2(b), Fig. 2(c) is more similar to the ideal coefficient matrix in Fig. 2(a), because the reconstruction coefficients exhibit more obvious block-wise diagonal structure.

2) *Results on the Real-World Datasets:* For performance evaluation, we follow [37] to take the accuracy and precision as two criteria. The accuracy is the percentage of correctly annotated faces (also including the correctly annotated faces whose ground-truth name is the null name) over all faces, while the precision is the percentage of correctly annotated faces over the faces that are annotated as real names (i.e., we do not consider the faces annotated as the null class by a face naming method). Since all methods aim at inferring names based on the faces in the images with ambiguous captions, we use all the images in each dataset for both learning and testing. To fairly compare all methods, we define the *real name ratio* as the percentage of faces that are annotated as real names using each method over all the faces in the dataset, and we report the performances at the same real name ratio.

To achieve the same real name ratio for all methods, we use the minimum cost bipartite graph matching method (introduced in Section IV) to infer the names of the faces, and vary the hyperparameter θ to tune the real name ratio, as suggested in [37]. Specifically, the costs $c(f, q)$ and $c(f, p + 1)$ are set as follows. For MildML, we set

TABLE II
PERFORMANCES (AT GROUND-TRUTH RATIOS) OF DIFFERENT METHODS ON TWO REAL-WORLD DATASETS. THE BEST RESULTS ARE IN BOLD

Method	Soccer player		Labeled Yahoo! News	
	Accuracy	Precision	Accuracy	Precision
MMS	0.613	0.583	0.784	0.797
LR-SVM	0.574	0.534	0.697	0.690
cGMM	0.611	0.553	0.750	0.755
MildML	0.630	0.579	0.684	0.658
ASML	0.646	0.594	0.710	0.715
LRR	0.664	0.632	0.797	0.797
rLRR	0.725	0.694	0.830	0.836
LRRml	0.708	0.671	0.810	0.812
rLRRml	0.736	0.703	0.832	0.839

$c(f, q) = -\sum_{\mathbf{x} \in \mathcal{S}_q} w(\mathbf{x}_f^i, \mathbf{x})$ and $c(f, p + 1) = \theta$, as in [6], where $w(\mathbf{x}_f^i, \mathbf{x})$ is the similarity between \mathbf{x}_f^i and \mathbf{x} and \mathcal{S}_q contains all faces assigned to the name q while inferring the names of the faces. For cGMM, we set $c(f, q) = -\ln \mathcal{N}(\mathbf{x}_f^i; \boldsymbol{\mu}_q, \Sigma_q)$, and $c(f, p + 1) = -\ln \mathcal{N}(\mathbf{x}_f^i; \boldsymbol{\mu}_{(p+1)}, \Sigma_{(p+1)}) + \theta$, as in [32], where $\boldsymbol{\mu}_q$ and Σ_q (resp. $\boldsymbol{\mu}_{(p+1)}$ and $\Sigma_{(p+1)}$) are the mean and covariance of the faces assigned to the q th class (resp., the null class) in cGMM. Similarly for MMS and LR-SVM, we consider the decision values from the SVM classifiers of the n th name and the null class by setting the cost as $c(f, q) = -\text{dec}_q(\mathbf{x}_f^i)$ and $c(f, p + 1) = -\text{dec}_{\text{null}}(\mathbf{x}_f^i) + \theta$, where $\text{dec}_q(\mathbf{x}_f^i)$ and $\text{dec}_{\text{null}}(\mathbf{x}_f^i)$ are the decision values of SVM classifiers from the q th name and the null class, respectively. The accuracies and precisions of different methods on the real-world datasets are shown in Table II, where the real name ratio for each method is set to be close to the ground-truth ratio using a suitable hyperparameter θ , as suggested in [37]. For a more comprehensive comparison, we also plot the accuracies and precisions on these two real-world datasets when using different real name ratios for all methods, by varying the value of the parameter θ . In Fig. 3, we compare the performances of our proposed methods ASML and rLRRml with the baseline methods MMS, cGMM, LR-SVM, and MildML on these two real-world datasets, respectively. In Fig. 4, we compare the performances of our proposed methods rLRRml, rLRR with the special cases LRRml, and LRR on these two real-world datasets, respectively. According to these results, we have the following observations.

- 1) Among the four baseline algorithms MMS, cGMM, LR-SVM, and MildML, there is no consistent winner on both datasets in terms of the accuracies and precisions in Table. II. On the Labeled Yahoo! News dataset, MMS achieves the best accuracy and precision among four methods. On the Soccer player dataset, MMS still achieves the best precision, but MildML achieves the best accuracy.
- 2) We also compare ASML with MildML, because both methods use captions-based weak supervision for distance metric learning. According to Table II, ASML outperforms MildML on both datasets in terms of both accuracy and precision. From Fig. 3, we observe that ASML consistently outperforms MildML

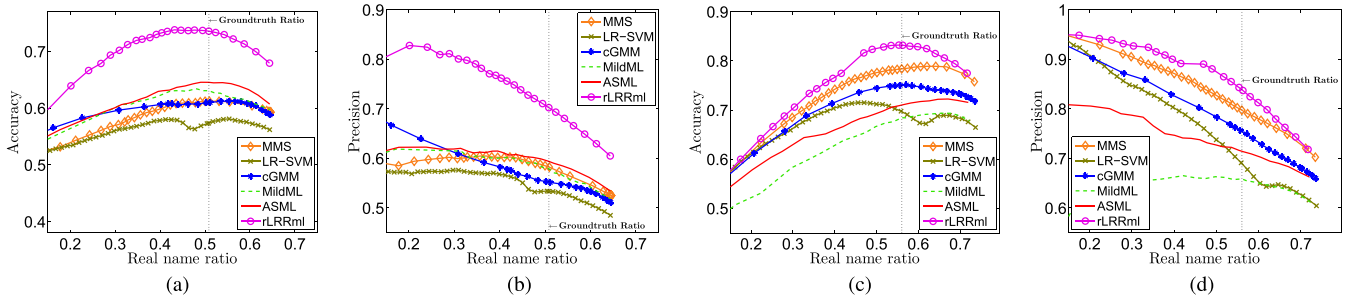


Fig. 3. Accuracy and precision curves for the proposed methods rLRRml and ASML, as well as the baseline methods MMS, cGMM, LR-SVM, and MildML, on the Soccer player dataset and the Labeled Yahoo! News dataset, respectively. (a) Accuracy versus real name ratio on the Soccer player dataset. (b) Precision versus real name ratio on the Soccer player dataset. (c) Accuracy versus real name ratio on the Labeled Yahoo! News dataset. (d) Precision versus real name ratio on the Labeled Yahoo! News dataset.

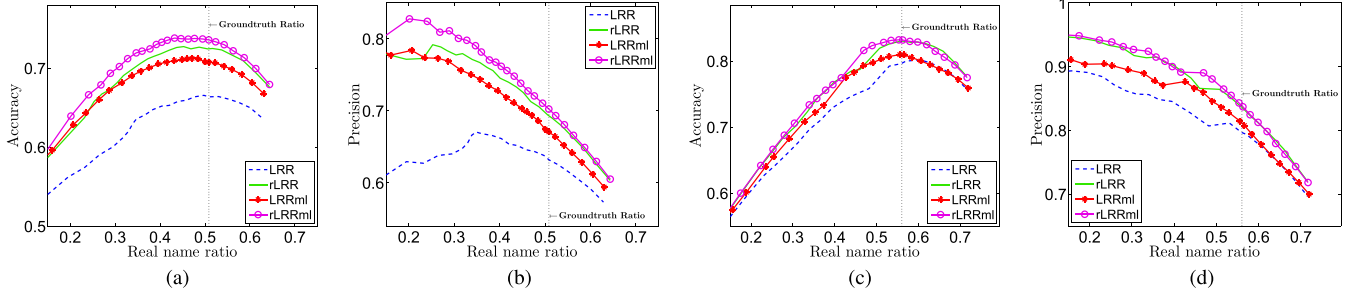


Fig. 4. Accuracy and precision curves for the proposed methods rLRRml and rLRR, as well as the special cases LRRml and LRR, on the Soccer player dataset and the Labeled Yahoo! News dataset. (a) Accuracy versus real name ratio on the Soccer player dataset. (b) Precision versus real name ratio on the Soccer player dataset. (c) Accuracy versus real name ratio on the Labeled Yahoo! News dataset. (d) Precision versus real name ratio on the Labeled Yahoo! News dataset.

on the Labeled Yahoo! News dataset, and generally outperforms MildML on the Soccer player dataset. These results indicate that ASML can learn a more discriminative distance metric by better utilizing ambiguous supervision information.

- 3) LRR performs well on both datasets, which indicates that our assumption that faces in a common subspace should belong to the same subject/name is generally satisfied on both real-world datasets. Moreover, rLRR consistently achieves much better performance compared with the original LRR algorithm on both datasets (Table II and Fig. 4), which demonstrates that it is beneficial to additionally consider weak supervision information by introducing the new regularizer into LRR while exploring the *subspace structures among faces*.
- 4) According to Table II, rLRRml is better than rLRR, and LRRml also outperforms LRR on both datasets in terms of accuracy and precision. On the Soccer player dataset [Fig. 4(a) and (b)], rLRRml (resp., LRRml) consistently outperforms rLRR (resp., LRR). On the Labeled Yahoo! News dataset [Fig. 4(c) and (d)], rLRRml (resp., LRRml) generally outperforms rLRR (resp., LRR). One possible explanation is that these two affinity matrices contain complementary information to some extent, because they explore weak supervision information in different ways. Hence, the fused affinity matrix is more discriminative for face naming. Note that the performance of rLRR on the Labeled Yahoo! News dataset is already high, so the improvement of rLRRml over rLRR on this dataset is not as significant as that on the Soccer player dataset.

- 5) Compared with all other algorithms, the proposed rLRRml algorithm achieves the best results in terms of both accuracy and precision on both datasets (Table II). It can be observed that rLRRml consistently outperforms all other methods on the Soccer player dataset [Fig. 3(a) and (b) and Fig. 4(a) and (b)], and rLRRml generally achieves the best performance on the Labeled Yahoo! News dataset [Fig. 3(c) and (d) and Fig. 4(c) and (d)]. These results demonstrate the effectiveness of our rLRRml for face naming.
- 6) For all methods, the results on the Soccer player dataset are worse than those on the Labeled Yahoo! News dataset. One possible explanation is that the Soccer player dataset is a more challenging dataset because there are more faces in each image, more names in each caption, and relatively more faces from the null class in the Soccer player dataset (Table I).

More Discussions on \mathbf{H} in Our rLRR: In our rLRR, we penalize the following two cases using the specially designed \mathbf{H} : 1) a face is reconstructed by the irrelevant faces that do not share any common names with this face according to their candidate name sets and 2) a face is reconstructed by using itself. If we only consider one case when designing \mathbf{H} in our rLRR, the corresponding results will be worse than the current results in Table II. Taking the Soccer player dataset as an example, we redefine \mathbf{H} by only considering the first (resp., the second) case, the accuracy and precision of our rLRR method become 0.714 and 0.682 (resp., 0.694 and 0.664), respectively. These results are worse than the results (i.e., the accuracy is 0.725 and the precision is 0.694) of our rLRR in Table II that considers both cases when designing \mathbf{H} ,

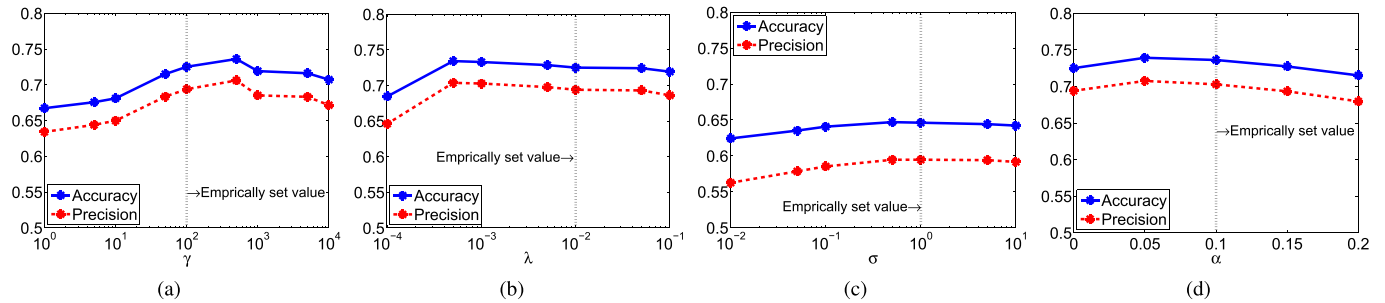


Fig. 5. Performances (accuracies and precisions) of our methods on the Soccer player dataset when using different parameters. The black dotted line indicates the empirically set value (i.e., the default value) of each parameter. (a) Performances of rLRR with respect to γ . (b) Performances of rLRR with respect to λ . (c) Performances of ASML with respect to σ . (d) Performances of rLRRml with respect to α .

which experimentally validates the effectiveness of penalizing both cases.

3) Performance Variations of Our Methods Using Different Parameters: We take the Soccer player dataset as an example to study the performances (i.e., accuracies and precisions) of our methods using different parameters.

We first study the performances of our rLRR when using different values of the parameters γ and λ , and the results are shown in Fig. 5(a) and (b), respectively. Note that we vary one parameter and set the other parameter as its default value (i.e., $\gamma = 100$ and $\lambda = 0.01$). In (4), γ is the tradeoff parameter for balancing the new regularizer $\|\mathbf{W} \circ \mathbf{H}\|_F^2$ (which incorporates weakly supervised information) and other terms. Recall that our rLRR reduces to LRR when γ is set to zero. When setting γ in the range of (1, 500), the performances of rLRR become better as γ increases and rLRR consistently outperforms LRR, which again shows that it is beneficial to utilize weakly supervised information. We also observe that the performances of rLRR are relatively stable when setting γ in the range of (50, 5000). The parameter λ is used in both LRR and our rLRR. We observe that our rLRR is relatively robust to the parameter λ when setting λ in the range of $(5 \times 10^{-4}, 10^{-1})$.

In Fig. 5(c), we show the results of our new metric learning method ASML when using different values of the parameter σ in (11). It can be observed that our ASML is relatively stable to the parameter σ when σ is in the range of (0.1, 10).

Finally, we study the performance variations of our rLRRml when setting the parameter α to different values, as shown in Fig. 5(d). When setting $\alpha = 0$ and $\alpha = 1$, rLRRml reduces to rLRR and ASML, respectively. As shown in Table II, rLRR is better than ASML in both cases in terms of accuracy and precision. Therefore, we empirically set α as a smaller value such that the affinity matrix from rLRR contributes more in the fused affinity matrix. When setting α in the range of (0.05, 0.15), we observe that our rLRRml is relatively robust to the parameter α and the results are consistently better than rLRR and ASML, which demonstrates that the two affinity matrices from rLRR and ASML contain complementary information to some extent.

VI. CONCLUSION

In this paper, we have proposed a new scheme for face naming with caption-based supervision, in which one image

that may contain multiple faces is associated with a caption specifying only who is in the image. To effectively utilize the caption-based weak supervision, we propose an LRR based method, called rLRR by introducing a new regularizer to utilize such weak supervision information. We also develop a new distance metric learning method ASML using weak supervision information to seek a discriminant Mahalanobis distance metric. Two affinity matrices can be obtained from rLRR and ASML, respectively. Moreover, we further fuse the two affinity matrices and additionally propose an iterative scheme for face naming based on the fused affinity matrix. The experiments conducted on a synthetic dataset clearly demonstrate the effectiveness of the new regularizer in rLRR. In the experiments on two challenging real-world datasets (i.e., the Soccer player dataset and the Labeled Yahoo! News dataset), our rLRR outperforms LRR, and our ASML is better than the existing distance metric learning method MildML. Moreover, our proposed rLRRml outperforms rLRR and ASML, as well as several state-of-the-art baseline algorithms.

To further improve the face naming performances, we plan to extend our rLRR in the future by additionally incorporating the ℓ_1 -norm-based regularizer and using other losses when designing new regularizers. We will also study how to automatically determine the optimal parameters for our methods in the future.

REFERENCES

- [1] P. Viola and M. J. Jones, "Robust real-time face detection," *Int. J. Comput. Vis.*, vol. 57, no. 2, pp. 137–154, 2004.
- [2] G. Liu, Z. Lin, and Y. Yu, "Robust subspace segmentation by low-rank representation," in *Proc. 27th Int. Conf. Mach. Learn.*, Haifa, Israel, Jun. 2010, pp. 663–670.
- [3] T. L. Berg *et al.*, "Names and faces in the news," in *Proc. 17th IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, Washington, DC, USA, Jun./Jul. 2004, pp. II-848–II-854.
- [4] D. Ozkan and P. Duygulu, "A graph based approach for naming faces in news photos," in *Proc. 19th IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, New York, NY, USA, Jun. 2006, pp. 1477–1482.
- [5] P. T. Pham, M. Moens, and T. Tuytelaars, "Cross-media alignment of names and faces," *IEEE Trans. Multimedia*, vol. 12, no. 1, pp. 13–27, Jan. 2010.
- [6] M. Guillaumin, J. Verbeek, and C. Schmid, "Multiple instance metric learning from automatically labeled bags of faces," in *Proc. 11th Eur. Conf. Comput. Vis.*, Heraklion, Crete, Sep. 2010, pp. 634–647.
- [7] J. Luo and F. Orabona, "Learning from candidate labeling sets," in *Proc. 23rd Annu. Conf. Adv. Neural Inf. Process. Syst.*, Vancouver, BC, Canada, Dec. 2010, pp. 1504–1512.

- [8] X. Zhang, L. Zhang, X.-J. Wang, and H.-Y. Shum, "Finding celebrities in billions of web images," *IEEE Trans. Multimedia*, vol. 14, no. 4, pp. 995–1007, Aug. 2012.
- [9] Z. Zeng *et al.*, "Learning by associating ambiguously labeled images," in *Proc. 26th IEEE Conf. Comput. Vis. Pattern Recognit.*, Portland, OR, USA, Jun. 2013, pp. 708–715.
- [10] M. Everingham, J. Sivic, and A. Zisserman, "Hello! My name is... Buffy—Automatic naming of characters in TV video," in *Proc. 17th Brit. Mach. Vis. Conf.*, Edinburgh, U.K., Sep. 2006, pp. 899–908.
- [11] J. Sang and C. Xu, "Robust face-name graph matching for movie character identification," *IEEE Trans. Multimedia*, vol. 14, no. 3, pp. 586–596, Jun. 2012.
- [12] Y.-F. Zhang, C. Xu, H. Lu, and Y.-M. Huang, "Character identification in feature-length films using global face-name matching," *IEEE Trans. Multimedia*, vol. 11, no. 7, pp. 1276–1288, Nov. 2009.
- [13] M. Tapaswi, M. Bäumel, and R. Stiefelhagen, "'Knock! Knock! Who is it?' Probabilistic person identification in TV series," in *Proc. 25th IEEE Conf. Comput. Vis. Pattern Recognit.*, Providence, RI, USA, Jun. 2012, pp. 2658–2665.
- [14] E. J. Candès, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?" *J. ACM*, vol. 58, no. 3, pp. 1–37, 2011, Art. ID 11.
- [15] Y. Deng, Q. Dai, R. Liu, Z. Zhang, and S. Hu, "Low-rank structure learning via nonconvex heuristic recovery," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 24, no. 3, pp. 383–396, Mar. 2013.
- [16] K. Q. Weinberger and L. K. Saul, "Distance metric learning for large margin nearest neighbor classification," *J. Mach. Learn. Res.*, vol. 10, pp. 207–244, Feb. 2009.
- [17] C. Shen, J. Kim, and L. Wang, "A scalable dual approach to semidefinite metric learning," in *Proc. 24th IEEE Conf. Comput. Vis. Pattern Recognit.*, Colorado Springs, CO, USA, Jun. 2011, pp. 2601–2608.
- [18] B. McFee and G. Lanckriet, "Metric learning to rank," in *Proc. 27th Int. Conf. Mach. Learn.*, Haifa, Israel, Jun. 2010, pp. 775–782.
- [19] S. Andrews, I. Tsochantaridis, and T. Hofmann, "Support vector machines for multiple-instance learning," in *Proc. 16th Annu. Conf. Neural Inf. Process. Syst.*, Vancouver, BC, Canada, Dec. 2003, pp. 65–72.
- [20] M.-L. Zhang and Z.-H. Zhou, "M³MIML: A maximum margin method for multi-instance multi-label learning," in *Proc. 8th IEEE Int. Conf. Data Mining*, Pisa, Italy, Dec. 2008, pp. 688–697.
- [21] T. Cour, B. Sapp, C. Jordan, and B. Taskar, "Learning from ambiguously labeled images," in *Proc. 22nd IEEE Conf. Comput. Vis. Pattern Recognit.*, Miami, FL, USA, Jun. 2009, pp. 919–926.
- [22] E. Elhamifar and R. Vidal, "Sparse subspace clustering," in *Proc. 22nd IEEE Conf. Comput. Vis. Pattern Recognit.*, Miami, FL, USA, Jun. 2009, pp. 2790–2797.
- [23] C. Lu, J. Feng, Z. Lin, and S. Yan, "Correlation adaptive subspace segmentation by trace Lasso," in *Proc. 12th IEEE Int. Conf. Comput. Vis.*, Sydney, VIC, Australia, Dec. 2013, pp. 1345–1352.
- [24] S. Xiao, M. Tan, and D. Xu, "Weighted block-sparse low rank representation for face clustering in videos," in *Proc. 13th Eur. Conf. Comput. Vis.*, Zürich, Switzerland, Sep. 2014, pp. 123–138.
- [25] Z. Lin, R. Liu, and Z. Su, "Linearized alternating direction method with adaptive penalty for low-rank representation," in *Proc. 24th Annu. Conf. Neural Inf. Process. Syst.*, Granada, Spain, Dec. 2011, pp. 612–620.
- [26] J.-F. Cai, E. J. Candès, and Z. Shen, "A singular value thresholding algorithm for matrix completion," *SIAM J. Optim.*, vol. 20, no. 4, pp. 1956–1982, 2010.
- [27] J. Yang, W. Yin, Y. Zhang, and Y. Wang, "A fast algorithm for edge-preserving variational multichannel image restoration," *SIAM J. Imag. Sci.*, vol. 2, no. 2, pp. 569–592, 2009.
- [28] C. Shen, J. Kim, and L. Wang, "Scalable large-margin Mahalanobis distance metric learning," *IEEE Trans. Neural Netw.*, vol. 21, no. 9, pp. 1524–1530, Sep. 2010.
- [29] A. L. Yuille and A. Rangarajan, "The concave-convex procedure," *Neural Comput.*, vol. 15, no. 4, pp. 915–936, 2003.
- [30] S. Shalev-Shwartz, Y. Singer, N. Srebro, and A. Cotter, "Pegasos: Primal estimated sub-gradient solver for SVM," *Math. Program.*, vol. 127, no. 1, pp. 3–30, 2011.
- [31] K. Q. Weinberger and L. K. Saul, "Fast solvers and efficient implementations for distance metric learning," in *Proc. 25th IEEE Int. Conf. Mach. Learn.*, Helsinki, Finland, Jun. 2008, pp. 1160–1167.
- [32] M. Guillaumin, T. Mensink, J. Verbeek, and C. Schmid, "Face recognition from caption-based supervision," *Int. J. Comput. Vis.*, vol. 96, no. 1, pp. 64–82, 2012.
- [33] J. Luo, B. Caputo, and V. Ferrari, "Who's doing what: Joint modeling of names and verbs for simultaneous face and pose annotation," in *Proc. 22nd Annu. Conf. Neural Inf. Process. Syst.*, Vancouver, BC, Canada, Dec. 2009, pp. 1168–1176.
- [34] T. L. Berg, E. C. Berg, J. Edwards, and D. A. Forsyth, "Who's in the picture?" in *Proc. 19th Annu. Conf. Neural Inf. Process. Syst.*, Vancouver, BC, Canada, Dec. 2006, pp. 137–144.
- [35] M. Guillaumin, T. Mensink, J. Verbeek, and C. Schmid, "Automatic face naming with caption-based supervision," in *Proc. 21st IEEE Conf. Comput. Vis. Pattern Recognit.*, Anchorage, AL, USA, Jun. 2008, pp. 1–8.
- [36] X. Xu, I. W. Tsang, and D. Xu, "Soft margin multiple kernel learning," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 24, no. 5, pp. 749–761, May 2013.
- [37] M. Özcan, L. Jie, V. Ferrari, and B. Caputo, "A large-scale database of images and captions for automatic face naming," in *Proc. 22nd Brit. Mach. Vis. Conf.*, Dundee, U.K., Sep. 2011, pp. 1–11.



Shijie Xiao received the B.E. degree from the Harbin Institute of Technology, Harbin, China, in 2011. He is currently pursuing the Ph.D. degree with the School of Computer Engineering, Nanyang Technological University, Singapore.


His current research interests include machine learning and computer vision.



Dong Xu (M'07–SM'13) received the B.E. and Ph.D. degrees from the University of Science and Technology of China, Hefei, China, in 2001 and 2005, respectively.

He was with Microsoft Research Asia, Beijing, China, and the Chinese University of Hong Kong, Hong Kong, for over two years, while pursuing the Ph.D. degree. He was a Post-Doctoral Research Scientist with Columbia University, New York, NY, USA, for one year. He also worked as a faculty member in the School of Computer Engineering, Nanyang Technological University, Singapore. He is currently a faculty member in the School of Electrical and Information Engineering, The University of Sydney, Australia. His current research interests include computer vision, statistical learning, and multimedia content analysis.

Dr. Xu co-authored a paper that received the Best Student Paper Award in the IEEE International Conference on Computer Vision and Pattern Recognition in 2010. His another coauthored paper also won the IEEE Transactions on Multimedia (T-MM) prize paper award in 2014.



Jianxin Wu (M'09) received the B.S. and M.S. degrees in computer science from Nanjing University, Nanjing, China, and the Ph.D. degree in computer science from the Georgia Institute of Technology, Atlanta, GA, USA.

He was an Assistant Professor with Nanyang Technological University, Singapore. He is currently a Professor with the Department of Computer Science and Technology, Nanjing University. His current research interests include computer vision and machine learning.