# CHIMPANZEE IDENTIFICATION USING GLOBAL AND LOCAL FEATURES

*Alexander Loos*

Audio-Visual Systems
Fraunhofer IDMT
98693 Ilmenau, Germany

**ABSTRACT**

Because of the ongoing biodiversity crisis many species like chimpanzees or gorillas for example are threatened and need to be protected. To overcome this agitating issue, biologist recently started to use remote camera devices for wildlife monitoring and estimation of remaining population sizes. Unfortunately, the huge amount of data makes the necessary manual analysis extremely tedious and highly cost intensive. To reduce the burden of time consuming routine work, we have recently started to develop computer vision algorithms to identify individuals. In this paper we extend our previous work using both global and local information for identification. To combine the results of the two approaches we apply a decision based parallel fusion scheme where we take the confidences of both classifiers into account. We show that the proposed approach outperforms our previous work for full-frontal faces while at the same time being more robust against pose variations. We evaluate our algorithm on two datasets of captive and free-living chimpanzees. The outcome of this paper builds the basis of a semi-automatic identification system for African Great Apes which will help biologists to develop new and innovative protection strategies.

***Index Terms*—** Primates, Face Recognition, Decision Fusion

## 1. MOTIVATION

In recent years the world's biodiversity is declining on an unprecedented scale. Many species, including African Great Apes such as chimpanzees for instance, are endangered and remaining populations need to be protected. According to Walsh *et al.* , ape populations in western equatorial Africa between 1983 and 2000 declined by more than half [1]. Another study by Campbell *et al.* even observed a 90% decrease in the number of chimpanzee nests over a 17 year period between 1990 and 2007 in Côte d'Ivoire [2]. Those agitating results demonstrate the urgent need to intensify close surveillance of this threatened species. Individual identification of animals is not only a prerequisite for measuring the success of implemented protection schemes but also for lots of other biological questions, e.g. wildlife epidemiology and social network analysis. An essential part of effective biodiversity conservation management is non-invasive population monitoring using remote camera devices. Unfortunately, the manual analysis of the large amount of data is extremely cost and labor intensive. Consequently, there is a high demand for automated processing of remotely gathered video recordings. Especially so-called capture-mark-recapture methods, commonly used in ecology, could benefit from an automated system for identification of Great Apes.

In this paper we extend the approaches from [3] and [4] to improve the system's invariance against pose variations, difficult lighting conditions, and partial occlusion. While global descriptors represent the whole appearance of the chimpanzee's face, local features around certain facial fiducial points are more robust against local changes because they only encode the detailed traits of the corresponding point of interest. It is well known that from psychophysics and neuroscience that both holistic and local information are crucial for perception and recognition of faces. Starting from the assumption that a combination of global and local descriptors should improve the performance and robustness of the system, we use a decision fusion scheme to combine their results. We show that global feature vectors obtained by Gabor features in combination with Speeded-Up Robust Features (SURF) [5] as local face representation achieve promising results in the new field of face recognition of Great Apes and clearly outperforms the system presented in [4]. For evaluation we use two realistic real-world datasets of chimpanzees, gathered in the zoo and in the field. The outcome of this paper builds the basis of a semi-automatic system for primate identification in photos and videos, which open up new venues in effective wildlife monitoring and biodiversity conservation management.

The remainder of the paper is organized as follows: In the subsequent section we give a short recap of existing work in the field of animal identification and our own previous work. A detailed description of the proposed system is presented in section 3 and we thoroughly evaluate our system on two publicly available datasets of free-living and captive chimpanzees in section 4. Finally, in section 5 we conclude this paper and give further ideas of improvement.

## 2. RELATED WORK

The field of computer vision and pattern recognition has been an active research field for years. Even though automatic image and video processing techniques become more and more important for the detection and identification of animals, yet only few publications do exist dealing with that topic. Ardovini *et al.* [6] for instance proposed a system for semi-automatic recognition of elephants from photos based on shape comparison of the nicks characterizing the elephants ears. Also Burghardt *et al.* [7] presented a full automatic system for penguin identification. After a penguin has been detected, unique individual-specific spot patterns on the penguin's coat are used for identification. More recently a method called StripeCodes for zebra identification was published by Lahiri *et al.* [8]. The authors claim that their algorithm efficiently extracts simple image features used for the comparison of zebra images to determine whether the animal has been observed before or not.

The aforementioned approaches use characteristic coat patterns or other individually unique biometrics like the pattern of fur and skin as well as unique nicks in ears to distinguish between individuals. Unfortunately, such an approach is often infeasible for the identification of Great Apes since unique coat markings are not existent or cannot be used because of the limited resolution of video recordings. Based on the assumption that humans and our closest relatives share similar properties of the face, we suggested to use and adapt face recognition techniques, originally developed to identify humans, for the identification of chimpanzees and gorillas [3, 4]. Although the results of [4] are very promising, the accuracy of the system drops significantly if non-frontal face images are used for testing. In this paper we try to overcome this limitation by using a combination of global and local features.

## 3. PROPOSED APPROACH

### 3.1. Feature Extraction

Since global features gather holistic information of the face and local descriptors around facial points represent intrinsic factors, both should be used for classification. Additionally, it has been reported in the literature that different representations misclassify different patterns. Therefore, different features offer complementary information which can be used to improve the recognition results. As global features we propose to use Gabor features, which are known to perform well in pattern recognition tasks. The complimentary local descriptor is SURF, a powerful visual descriptor of interest points in an image.

#### *3.1.1. Gabor Descriptor*

Gabor Features are extracted by convolving of the graylevel input image $I(z)$ with a set of Gabor kernels $\psi_{\mu,\nu}(z)$.

$$G_{\mu,\nu}(z) = I(z) * \psi_{\mu,\nu}(z), \tag{1}$$

where $G_{\mu,\nu}(z)$ is the output image at orientation $\mu$ and scale $\nu$ at pixel $z = (x, y)$. Complex Gabor kernels are defined as

$$\psi_{\mu,\nu}(z) = \frac{\|k_{\mu,\nu}\|^2}{\sigma^2} e^{\frac{-\|k_{\mu,\nu}\|^2\|z\|^2}{2\sigma^2}} [e^{ik_{\mu,\nu}z} - e^{\frac{-\sigma^2}{2}}], \tag{2}$$

where the wave vector $k_{\mu,\nu}$ is defined as $k_{\mu,\nu} = k_\nu e^{i\theta_\mu}$ with $k_\nu = \frac{k_{max}}{f^\nu}$ and $\theta_\mu = \frac{\pi\mu}{8}$. The maximum frequency is denoted as $k_{max}$ and $f$ is the spacing between kernels in the frequency domain. Additionally, $\sigma$ represents the ratio of the Gaussian window to the wavelength.

In general, $G_{\mu,\nu}(z)$ is complex and can be rewritten as $G_{\mu,\nu}(z) = M_{\mu,\nu}(z)e^{i\theta_{\mu,\nu}(z)}$, where $M_{\mu,\nu}(z)$ denotes the magnitude and $\theta_{\mu,\nu}(z)$ the phase at pixel location $z$. Since the magnitude contains the local energy variation in the facial image, $M_{\mu,\nu}$ is used as feature. Finally, the overall feature vector is constructed as

$$x_{GABOR} = \left(m^{(\rho)}_{0,0}, m^{(\rho)}_{0,1}, \cdots, m^{(\rho)}_{1,0}, \cdots, m^{(\rho)}_{K,L}\right), \tag{3}$$

where $m^{(\rho)}_{\mu,\nu}$ is a column vector representing the normalized and vectorized version of the magnitude matrix $M_{\mu,\nu}$, which was downsampled by factor $\rho$.

#### *3.1.2. SURF Descriptor*

Speeded Up Robust Features (SURF) is a fast and robust scale- and rotation-invariant interest point detector and descriptor. Because in this task we already know the position of the interest points, we only refer to the descriptor part of SURF in this paper. In the following we briefly describe the main ideas of SURF. A more detailed description including the detection part can be found in [5]. As claimed by the authors the standard version of SURF is several times faster, more compact and at the same time more robust against certain image transformations than comparable local descriptors like Scale Invariant Feature transform (SIFT) [9] for instance. Similar to SIFT and its variants, SURF describes the distribution of intensity content within a certain neighborhood around the interest point. However, instead of using gradient information directly, SURF uses first order Haar wavelet responses in $x$ and $y$-direction. For efficiency SURF exploits integral images, which drastically reduces both processing time and robustness of the resulting descriptors. In order to increase the robustness against rotation, usually the first step of feature extraction is to identify a reproducible orientation for the interest point. The dominant orientation can be found by calculating the sum of the Gaussian weighted Haar wavelet responses using a sliding window around a circular region around the interest point. The next step is to construct a square region with correct orientation symmetrically around the interest point. This region is then split into $4 \times 4$ subregions. Finally, the feature vector can be calculated by again using Haar wavelet responses weighted with a Gaussian kernel, which is centered at the particular interest point. The horizontal and vertical wavelet responses, $dx$ and $dy$, as well as their absolute values are summed up over each sub-region to construct the final feature vector

$$x_{SURF} = \left(\sum dx, \sum dy, \sum |dx|, \sum |dy|\right). \tag{4}$$

### 3.2. Feature Space Transformation

The goal of many feature space transformation techniques is to project the $N$ high dimensional vectorized feature vectors $\{x_1, \cdots, x_N\}$ of size $n$ into a smaller dimensional subspace of size $m$ using a unitary projection matrix $W \in \mathbb{R}^{n\times m}$.

$$y_k = W^T x_k \tag{5}$$

The resulting feature vectors $y_k \in \mathbb{R}^m$, with $k = 1, \cdots, N$, can then be used for classification. In this paper we propose to use Locality Preserving Projections (LPP) [10] for feature space transformation. This approach assumes that the feature vectors reside on a nonlinear sub-manifold hidden in the original feature-space. LPP tries to find an embedding that preserves local information and obtains a subspace that best detects the essential manifold structure of the feature-space. To preserve the local structure of the feature-space, this manifold structure is modeled by a nearest-neighbor graph. First, an adjacency graph $G$ with $m$ nodes is defined. An edge is put between two nodes $k$ and $j$ if they are within an $\epsilon$ neighborhood, i.e. if $\|x_k - x_j\|^2 < \epsilon$. LPP will try to optimally preserve this graph in choosing projections. After constructing the graph, weights have to be assigned to the edges. Therefore a sparse symmetric matrix $S$ of size $N \times N$ is created with $S_{k,j}$ having the weight of the edge joining vertices $k$ and $j$, and 0 otherwise. The weights are calculated as follows:

$$S_{k,j} = \begin{cases} e^{\frac{\|x_k - x_j\|^2}{t}}, & \text{if } \|x_k - x_j\|^2 < \epsilon \\ 0, & \text{otherwise.} \end{cases} \tag{6}$$

The constant values $t$ and $\epsilon > 0$ have to be chosen adaptively. Here, $\epsilon$ defines the radius of the local neighborhood. Therefore, the objective function of LPP is defined as

$$w_{opt} = \min \sum_{kj} (y_k - y_j)^2 S_{k,j}. \tag{7}$$

Following some simple algebraic steps, it is possible to show that Eq. (7) finally results in a generalized eigenvalue problem:

$$XLX^T w = \lambda XDX^T w, \tag{8}$$

where $D$ is a diagonal matrix whose entries are column sums of $S$ and $L = D - S$ is the so called Laplacian matrix. The $k$-th column of matrix $X$ is $x_k$.

The projection matrix $W$ is constructed by concatenating the solution to the above equation, *i.e.* the column vectors of $W_{LPP} = [w_1, \cdots, w_m]$ are ordered ascendingly according to their eigenvalues. Usually, the original features are first projected into the PCA subspace before applying LPP by deleting the smallest principle components. Thus, the final embedding is as follows:

$$W_{final} = W_{PCA} W_{LPP}. \tag{9}$$

Details about the algorithm and the underlying theory can be found in [10].

### 3.3. Classification

#### *3.3.1. Sparse Representation Classification*

Sparse Representation Classification (SRC) was developed by Wright *et al.* and is known to perform well for face recognition [11, 12].

Let $\tilde{A}$ be the normalized matrix of training samples transformed into the feature space and $\tilde{t}$ be the normalized vectorized test image in the feature domain, then classification can be done by first solving a convex optimization problem via $l_1$-norm minimization:

$$\hat{p} = \arg\min_{p} \|p\|_1 \quad \text{subject to} \quad \tilde{t} = \tilde{A}p, \tag{10}$$

where $p$ is a sparse coefficient vector whose entries only associated with the $i$-th class are 1 and the rest is 0. We then assign $\tilde{t}$ to the object class that minimizes the residual $r_i(\tilde{t})$ between $\tilde{t}$ and $\tilde{A}\delta_i(\hat{p})$, such that

$$\min_{i} r_i(\tilde{t}) = \|\tilde{t} - \tilde{A}\delta_i(\hat{x})\|_2, \tag{11}$$

where $\delta_i$ is the characteristic function of class $i$, which is 1 for all training samples of class $i$ and 0 elsewhere. A detailed description of SRC can be found in [11].

#### *3.3.2. Support Vector Machines*

A Support Vector Machine (SVM) is a discriminative classifier, attempting to generate an optimal decision plane between feature vectors of the training classes. Oftentimes, classification with linear separation planes is not possible in the original feature space for real-world applications. Using a so called kernel trick, the feature vectors are transformed to a higher dimensional space in which they can be linearly separated.

### 3.4. Decision Fusion

The decision fusion paradigm we propose in this paper was influenced by ideas of [13]. A parallel ensemble classifier which fuses the rank-outputs of different classifiers is used to combine the results of local and global features. In contrast to the parallel fusion scheme in [13], where only a single function $f(x) = x^n$ is used for the non-linear rank-sum method, we propose to weight the results of both classifiers using different weighting functions. Additionally, the confidences of each classifier can be taken into account when generating the weighting function $f(x) = e^c$, where $c$ represents the confidence of SRC or SVM, respectively. For SRC we use the minimal residual $r_{min}$ from equation 11 as confidence measure, while for SVM the probability estimates of LibSVM [14] can be utilized. Details on the estimation of probabilities can be found in [15]. Figure 1 illustrates the proposed parallel fusion scheme. Note that for every of the six facial interest point we transform the resulting SURF descriptor separately into a smaller dimensional subspace before calculating the final feature vector. The position of the interest points are calculated based on the annotated coordinates of eyes and mouth and are depicted in Figure 1.

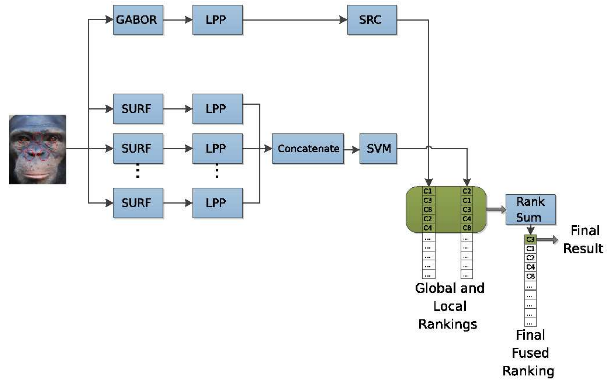


**Fig. 1**. The proposed parallel fusion paradigm used in this paper.

## 4. EXPERIMENTS AND RESULTS

In this section we give a detailed description of the experiments we conducted and present the results on two realistic databases of free-living as well as captured chimpanzee individuals. We compare the performance of our baseline algorithm from [4], where we suggested to use global Gabor features in combination with LPP and SRC for recognition, with the system proposed in this paper. We show that a decision fusion technique of holistic global features and local information gathered around facial interest points, outperforms the algorithm of previous work while at the same time being more robust against pose variations. Throughout the whole evaluation, we use 5 scales and 8 orientations for the generation of Gabor kernels. After convolving an image with the resulting 40 Gabor wavelets, we downsample the magnitude-matrix $M_{\mu,\nu}$ by a factor of 8. For LPP we decided to have 160 features after feature space transformation. For the local SURF descriptors we transform the resulting 64 dimensional feature features separately for every facial fiducial point into a feature space of size 50 and combine them by concatenating the resulting feature vectors.

### 4.1. Datasets

Due to the lack of publicly available benchmark databases for primates, we assembled two different sets of facial images for chimpanzees and made them available on our project website *http://www.saisbeco.com*. Table 1 shows an overview of the datasets we used in our experiments. The two datasets consist of different chimpanzee individuals, one of captured individuals from the zoo of Leipzig, Germany (*ChimpZoo*) and one of free-living primates from the Taï National Park, Africa (*ChimpTaï*). All images were

| Dataset | Origin | Images | Individuals |
|---|---|---|---|
| ChimpZoo | Zoo Leipzig | 1839 | 24 |
| ChimpTaï | Tai NP | 3193 | 71 |

**Table 1**. Overview of the three datasets we used in our experiments.

annotated by marking the region of the apes face and setting marker points for eyes and the mouth. We also assigned meta-information such as the name of the individual, species, gender and age to every facial image. To evaluate the system's robustness to pose variation, we also annotated the pose of each face. By using this information we can generate pose-specific subsets for every dataset, such as *Front*, *SemiLeft* and *SemiRight*. The subset *Front* then only contains full-frontal face images of every individual, while the subsets *SemiLeft* and *SemiRight* contain images of full-frontal and semi-left as well as full-frontal and semi-right faces, respectively. Example images of one individual per dataset with different poses can be seen in Figure 2. Note that because the datasets were gathered in uncontrolled environments, unlike most of the human face databases, not only pose-variations but also different lighting conditions, expressions and even partial occlusion are present in the images, which makes both datasets very challenging for face recognition tasks. For both datasets and their subsets we only considered face images

(a)

(b)

**Fig. 2**. Two individuals with three different poses: *SemiLeft* (left), *Front* (middle) and *SemiRight* (right). Images were taken from the datasets (a) ChimpZoo (b) ChimpTaï.

that have a minimal size of $64 \times 64$ pixels. Furthermore, we only focused on individuals with at least $5$ images to get an appropriate amount of training data for every individual. These criteria lead to $24$ individuals for the ChimpZoo datasets and $44$ individuals for the ChimTaï datasets. After the face images were converted to gray scale and rotated into an upright position, we used a projective transformation based on the annotated eye and mouth coordinates to align all the faces. For lighting normalization we applied a simple histogram equalization. For all the experiments described in the following sections we used a stratified $10$-fold crossvalidation to get valid results. For testing the *SemiLeft* and *SemiRight* subsets, respectively, only full-frontal faces are used for training.

### 4.2. Results

The results of our experiments can be seen in the Tables 2a and 2b. We evaluated four different algorithms against each other. The first method is the one we proposed in [4], where we use Gabor features as global descriptors in combination with LPP for feature space transformation and SRC for classification. Secondly, we evaluated how the local features, extracted by SURF, perform without any additional holistic information. Here, we also used LPP for dimensionality reduction, but applied it separately for the descriptors of every interest point, resulting in $50$ features per descriptor. For classification we used a SVM with RBF kernel. It is obvious that the results of GABOR and SURF features are comparable if only full frontal faces are contained in the dataset. However, for non-frontal faces SURF performs much better for both databases. Therefore, the combination of GABOR and SURF using fusion techniques boosts the performance for frontal faces while at the same time increase the robustness against pose variation. Note, that our proposed fusion scheme using the confidences of both classifiers for weighting the ranked results outperforms the fusion paradigm by [13] significantly for both datasets and all pose subsets.

(a)

| **Acc. (Std.) [%]** | SemiLeft | Front | SemiRight |
|---|---|---|---|
| GABOR | 85.81 (3.77) | 91.43 (3.85) | 87.86 (4.06) |
| SURF | 87.96 (3.11) | 90.41 (2.54) | 88.66 (4.27) |
| Fusion [13] | 88.45 (4.27) | 91.11 (2.89) | 89.88 (3.83) |
| Own | **91.28 (2.01)** | **94.26 (1.53)** | **91.53 (4.21)** |

(b)

| **Acc. (Std.) [%]** | SemiLeft | Front | SemiRight |
|---|---|---|---|
| GABOR | 74.55 (4.38) | 77.29 (4.50) | 75.02 (3.51) |
| SURF | 78.76 (3.55) | 80.51 (4.29) | 78.89 (2.68) |
| Fusion [13] | 80.29 (4.79) | 82.19 (4.32) | 80.40 (2.55) |
| Own | **83.69 (4.81)** | **84.29 (4.27)** | **82.27 (2.35)** |

**Table 2**. Rank-1 accuracy and standard deviation across the $10$ folds of global features (GABOR), local features (SURF), their combination using the fusion method proposed in [13] and our own fusion scheme (Own) for the different pose subsets (a) ChimpZoo (b) ChimpTaï.

## 5. CONCLUSION AND FUTURE WORK

In this paper we extended our approach from [4] for identification of captive and wild-living chimpanzees. We significantly improved the performance and the invariance against pose variations of the current system by fusing holistic and local information in a decision based manner. Moreover, we improved the parallel fusion scheme by [13] by taking the confidence of both classifiers into account. We thoroughly evaluated the algorithms on two self-established datasets of captive chimpanzees from the zoo of Leipzig, Germany (ChimpZoo), and free-living chimpanzee individuals from the Taï National Park, Africa (ChimpTaï). Both datasets were annotated by experts using a provided annotation tool. Based on the achieved results, the outcome of this paper builds the basis of a semi-automatic tool for the identification of Great Apes. Such a software can assist biologists with tedious annotation work of gathered video material and therefore has the potential to open up new venues in effective biodiversity conservation management. In future works we want to develop a complete identification system including automatic face detection, face alignment and face recognition.

**Acknowledgments**

We thank the Ivorian authorities for long term support, especially the Ivorian Ministére de l'Environnement, des Eaux et Forêts and the Ministére de l'Enseignement Supérieur et de la Recherche Scientifique, the directorship of the Taï National Park, the OIPR and the CSRS in Abidjan. Financial support is gratefully acknowledged from the Swiss Science Foundation. We would like to thank especially Dr. Tobias Deschner for collecting videos and pictures over the last years and for providing invaluable assistance during the data collection. We thank all the numerous field assistants and students for their work on the Taï Chimpanzee Project. We thank the Zoo Leipzig and the Wolfgang Köhler Primate Research Center (WKPRC), especially Josep Call and all the numerous research assistants, zoo-keepers and Josefine Kalbitz for support and collaboration. We also thank Laura Aporius for providing videos and pictures in 2010 and annotation of data. This work was supported by the Max Planck Society.

## 6. REFERENCES


[1] P. D. Walsh, K. A. A. Abernethy, and M Bermejo, "Catastrophic Ape Decline in Western Equatorial Africa," in *Nature 422:611 - 614*, 2003.

[2] Genevieve Campbell, Hjalmar Kuehl, Paul N'Goran Kouame, and Christophe Boesch, "Alarming Decline of West African Chimpanzees in Côte d'Ivoire," in *Current Biology, Vol. 18, No. 19, R904 - 905*, 2008.

[3] Alexander Loos, Martin Pfitzer, and Laura Aporius, "Identification of Great Apes Using Face Recognition," in *19th European Signal Processing Conference (EUSIPCO)*, 2011.

[4] Alexander Loos, "Identification of Great Apes Using Gabor Features and Locality Preserving Projections," in *1st ACM International Workshop on Multimedia Analysis for Ecological Data (MAED) in conjunction with ACM Multimedia*, 2012.

[5] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool, "Speeded-Up Robust Features (SURF)," *Computer Vision and Image Understanding*, vol. 110, no. 3, pp. 346 – 359, 2008.

[6] A. Ardovini, L. Cinque, and E. Sangineto, "Identifying elephant photos by multi-curve matching," *Pattern Recognition*, vol. 41, no. 6, pp. 1867–1877, June 2008.

[7] Tilo Burghardt and Neill Campbell, "Individual animal identification using visual biometrics on deformable coat patterns," in *5th International Conference on Computer Vision Systems (ICVS)*, 2007.

[8] Mayank Lahiri, Rosemary Warungu, Daniel I. Rubenstein, Tanya Y. Berger-Wolf, and Chayant Tantipathananandh, "Biometric Animal Databases from Field Photographs: Identification of Individual Zebra in the Wild," in *ACM International Conference on Multimedia Retrieval (ICMR)*, 2011.

[9] D.G. Lowe, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.

[10] Xiaofei He, Shuicheng Yan, Yuxiao Hu, Partha Niyogi, and Hong-Jiang Zhang, "Face recognition using laplacianfaces.," *IEEE transactions on pattern analysis and machine intelligence*, vol. 27, no. 3, pp. 328–40, Mar. 2005.

[11] John Wright, Allen Y Yang, Arvind Ganesh, S Shankar Sastry, and Yi Ma, "Robust face recognition via sparse representation.," *IEEE transactions on pattern analysis and machine intelligence*, vol. 31, no. 2, pp. 210–27, Feb. 2009.

[12] Meng Yang and Lei Zhang, "Gabor feature based sparse representation for face recognition with gabor occlusion dictionary," in *European Conference on Computer Vision (ECCV)*, 2010.

[13] Berk Gökberk, Albert Ali Salah, and Lale Akarun, "Rank-based decision fusion for 3D shape-based face recognition," in *International Conference on Audio-and Video-Based Biometric Person Identification (AVBPA)*, 2005, pp. 1019–1028.

[14] Chih-Chung Chang and Chih-Jen Lin, "LIBSVM: A library for support vector machines," *ACM Transactions on Intelligent Systems and Technology*, vol. 2, pp. 27:1–27:27, 2011.

[15] T.-F. Wu, C.-J. Lin, and R.C. Weng, "Probability estimates for multi-class classification by pairwise coupling," *Journal of Machine Learning Research*, vol. 5, pp. 975 – 1005, 2004.