Person-Specific Subspace Analysis for Unconstrained Familiar Face Identification

Giovani Chiachia[1,2]
giovanichiachia@gmail.com
Nicolas Pinto[2,3]
pinto@mit.edu
William Robson Schwartz[4]
william@dcc.ufmg.br
Anderson Rocha[1]
anderson.rocha@ic.unicamp.br
Alexandre X. Falcão[1]
afalcao@ic.unicamp.br
David Cox[2]
davidcox@fas.harvard.edu

[1] Institute of Computing
University of Campinas
Campinas, Brazil

[2] Rowland Institute
Harvard University
Cambridge, USA

[3] McGovern Institute
Massachusetts Institute of Technology
Cambridge, USA

[4] Department of Computer Science
Universidade Federal de Minas Gerais
Belo Horizonte, Brazil

Present face recognition systems can surpass human performance in the task of matching unfamiliar faces from images acquired under relatively controlled conditions [2]. However, in settings where images are less controlled and where human subjects are familiar with the faces that are tested, the advantage of humans over machines is still substantial [3]. While the issue of uncontrolled variation has received increased attention in recent years, the notion of "familiarity" in face recognition systems has been relatively unexplored.

The recognition of "familiar" faces is increasingly relevant in an age where an ever-growing torrent of images of friends and family members is made available. While many current face recognition benchmark sets are organized around deciding whether two probe faces are the same or different, the problem is often instead that of recognizing which individual a given face image belongs to. This identification problem has a natural relationship to the notion of "familiarity" in human face recognition, in that a large number of past examples of a relatively small cohort of individuals are leveraged to recognize new examples.

A growing body of neuroscience and psychology research suggests that human face recognition with familiar and unfamiliar faces is substantially different, possibly even relying on qualitatively different internal representations. Indeed, while human performance with unfamiliar faces is generally poor, performance with familiar faces is excellent [1].

Inspired by the idea that humans may rely on enhanced representations for familiar individuals, in this work we explore the construction of per-individual subspaces for performing face identification. We build these subspaces from orthonormal projection vectors obtained using a person-specific configuration of partial least squares [5, 6], which we refer to as PS-PLS models. A key motivating idea for this work is that such person-specific subspaces, due to its supervised nature, can capture both those aspects of the face that are good for discriminating it from others, as well as natural variation in appearance that is present in the unconstrained images of that individual.

Partial least squares (PLS) is a class of methods primarily designed to model relations between sets $\mathbf{X}$ and $\mathbf{Y}$ of observed variables by means of latent vectors [5]. The projection vectors $\mathbf{w}$ that are the basis of our person-specific subspaces are determined iteratively by the NIPALS algorithm [6] such that

$$\max_{||\mathbf{w}||=||\mathbf{c}||=1} [cov(\mathbf{Xw}, \mathbf{Yc})]^2, \quad (1)$$

where $cov(\mathbf{Xw}, \mathbf{Yc})$ is the sample covariance between $\mathbf{Xw}$ and $\mathbf{Yc}$ and $\mathbf{c}$ is dependent on $\mathbf{w}$ in our case. For each person, we model $\mathbf{Y}$ as a set with a single indicator variable such that $\mathbf{Y}_{n\times 1} = \mathbf{y}_c$, and $y_{cs} = 1$ if sample s (out of n) belongs to class c or $y_{cs} = 0$ otherwise. In this case, obtaining projection vectors $\{\mathbf{w}\}_i$ is straightforward [5]. At each iteration i,

$$\mathbf{w}_i = \mathbf{X}i^T \mathbf{y}_c, \quad (2)$$

where $\mathbf{X}i$ is the matrix $\mathbf{X}$ deflated up to iteration i according to the NIPALS algorithm [5, 6].

We use the *Pubfig83* dataset, a subset of the *Pubfig* face dataset [3] reconfigured for the problem of unconstrained face identification [4]. We replicate the previous best results with this data set [4] and consider them as baselines. The baseline methods consist of binary linear support vector machines (SVMs) trained on different visual representations of faces in a one-versus-all setting. To compare these methods, we project feature descriptor vectors from this method into custom PS-PLS subspaces that we construct so that each binary linear SVM is trained in a different and person-specific space corresponding to the positive class (Fig. 1).

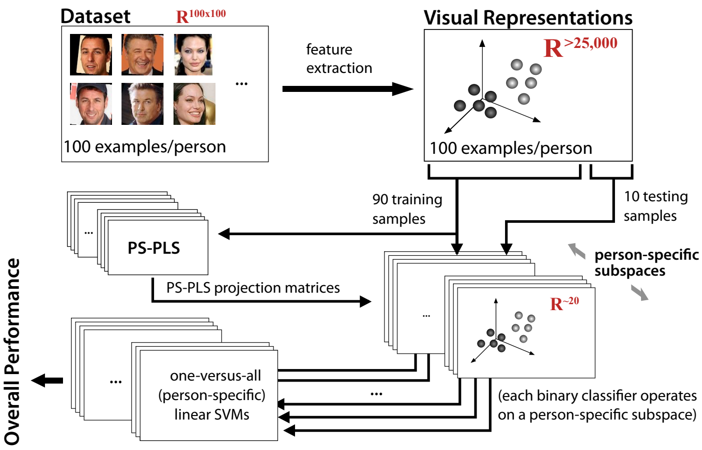

Figure 1: Our approach. From the training samples, PS-PLS creates a different face subspace for each individual. All training samples are then projected onto each subspace, so that a classification engine can be trained by considering the different representations of the samples over the subspaces. Given a test sample, an overall decision is made according to decisions made in each person-specific subspace.

We also compare our approach with subspaces built via other linear techniques. Among them, we consider person-specific principal component analysis (PCA), similar in spirit to the approach of Burton *et al.* [1], along with traditional non-person-specific PCA and linear discriminant analysis (LDA). Finally, as an additional test, we evaluate the approach on the *Facebook100* dataset [4], which is constructed from a large set of real-world face images taken from the Facebook social network.

With the use of the PS-PLS models, we could consistently get better results across the four different visual representations we consider. In the paper, we present details of the method and the results. In general, we argue that these subspaces are useful both for noise removal and for accentuating discriminative person-specific face aspects.

[1] A. M. Burton, R. Jenkins, and S. R. Schweinberger. Mental representations of familiar faces. *British Journal of Psychology*, 2011.

[2] P.J. Phillips et al. FRVT 2006 and ICE 2006 large-scale experimental results. *IEEE TPAMI*, 2010.

[3] N. Kumar, A. C. Berg, P. N. Belhumeur, and S. K. Nayar. Attribute and simile classifiers for face verification. *IEEE ICCV*, 2009.

[4] N. Pinto, Z. Stone, T. Zickler, and D. D. Cox. Scaling-up biologically-inspired computer vision: A case study in unconstrained face recognition on facebook. *IEEE CVPR*, 2011.

[5] R. Rosipal and N. Kramer. Overview and recent advances in partial least squares. *LNCS*, 2006.

[6] H. Wold. Partial least squares. *Encyclopedia of Statistical Sciences*, 1985.