

Using A Generalized Linear Mixed Model to Study the Configuration Space of a PCA+LDA Human Face Recognition Algorithm

Geof H. Givens¹, J. Ross Beveridge², Bruce A. Draper², and David Bolme²

¹ Statistics Department

Colorado State University

Fort Collins, CO, 80523 U.S.A

² Department of Computer Science

Colorado State University

Fort Collins, CO, 80523 U.S.A

Abstract. A generalized linear mixed model is used to estimate how rank 1 recognition of human faces with a PCA+LDA algorithm is affected by the choice of distance metric, image size, PCA space dimensionality, supplemental training, and the inclusion of test subjects in the training data. Random effects for replicated training sets and for repeated measures on people were included in the model. Results indicate that variation among subjects was a dominant source of variability, and that there was moderate correlation within people. Statistically significant effects and interactions were found for all configuration factors except image size. The most significant interaction is that changes to the PCA+LDA configuration only improved recognition for test subjects who were included in the training data. For subjects not included in training, no configuration changes were helpful.

This study is a model for how to evaluate algorithms with large numbers of parameters. For example, by accounting for subject variation as a random effect and explicitly looking for interaction effects, we are able to discern effects that might otherwise have been masked by subject variation and interaction effects. This study is instructive for what it reveals about PCA+LDA.

1 Introduction

Proper configuration or tuning is a critical factor in making effective use of most vision algorithms. Virtually all algorithms have a configuration space, and it is well understood performance varies with different configuration choices. There are a variety of approaches to configuration. Perhaps the most obvious and most common is semi-structured exploration: the developers of the algorithm run the algorithm, examine results, make changes, and repeat until they are satisfied. Somewhat more rigorous is iterative refinement with an objective performance standard as illustrated by the work of Bowyer [1].

We advocate designing experiments, collecting performance data, and applying statistical tools to determine how configuration choices influence performance. We apply this approach here to understand the configuration space of an interesting human face

recognition algorithm. The statistical tool is a generalized linear mixed model. The face recognition algorithm uses principal components analysis (PCA) followed by linear discriminant analysis (LDA) and is based closely on the work of Zhao [2].

The use of the generalized linear mixed model in this context provides several very important features lacking from simpler approaches. First, it permits us to model what are called fixed effects that are attributable to various algorithm configuration choices. These include factors such as image size and the selection (or generation) of training data. Second, it also permits us to model random effects that influence performance but are not under our direct control. A good example of a random effect is subject identity.

Accounting for random effects due to subject identity is particularly important in domains such as human face recognition. Some individuals are harder to recognize than others, but we are seldom afforded the luxury of choosing to only recognize the easy ones. Thus, a large source of variation in observed performance is attributable to a recorded variable that is not under our control. In simpler models that do not explicitly account for random effect, variation among subjects can easily obscure other effects.

2 Image Pre-processing

Data pre-processing is arguably as important as any other aspect of an algorithm. In the experiments presented below, the face imagery is pre-processed in essentially the same fashion as in the FERET tests [3, 4]. Specifically, the following pre-processing steps are applied to images:

1. Integer to float conversion.
2. Geometric normalization to align hand-picked eye coordinates.
3. Cropping to an elliptical mask centered on the face.
4. Histogram equalization of non-cropped pixels.
5. Pixel normalization so that every image has a mean of zero and a standard deviation of one.

Optional pre-processing steps considered as part of the configuration space include:

- Adding zero mean Gaussian noise to each pixel.
- Reflecting the image around the vertical axis. This creates a mirror-image of the original face.
- Shifting the image left, right, up or down.
- Reducing the image size by a factor of two.

3 The PCA+LDA Algorithm

The PCA+LDA human face recognition algorithm considered in this study is based upon the work of Zhao [2]. It processes a set of training images and builds a PCA subspace as in the standard Eigenfaces algorithm [5]. The training images are projected into the PCA subspace and grouped according to subject identity. Each subject is treated as a distinct class and the LDA basis vectors are computed. In the experiments that

follow, training is performed on 64 randomly selected subjects with 3 images each. Thus, the maximum possible dimensionality of the PCA space is 192, and one of the configuration choices involves how many of these to retain. Since there are 64 subjects, LDA generates 63 Fisher basis vectors. The two linear projections are then composed, yielding a single projection from image space into the PCA+LDA space.

In the testing phase, the PCA+LDA algorithm reads the subspace basis vectors generated during training, projects test imagery into this subspace, and measures distances between image pairs. Two distance measures are considered in the experiments below. The first is the standard L2 norm (\mathcal{L}_2). The second is a heuristic distance measure defined by Zhao [2] called Soft L2 (\mathcal{L}_{Soft}). Formally, these two distance measures are defined as:

$$\mathcal{L}_2(A, B) = \sum_{i=1}^N (A_i - B_i)^2 \quad \mathcal{L}_{Soft}(A, B) = \sum_{i=1}^N (\lambda_i)^{0.2} (A_i - B_i)^2 \quad (1)$$

where λ_i is the i th eigenvalue from the the solution of the generalized eigenvector problem associated with finding the fisher basis vectors [6].

4 Experimental Design

This study used 1,024 pre-processed images of 256 FERET subjects. The choice of subjects is based upon the data available: there are 256 FERET subjects with at least 4 frontal images. These images are divided into *training*, *gallery* and *probe* sets. Training images are used to define the PCA+LDA space. Probe images are test images to be identified. Gallery images are the stored images that probe images are compared to. For every subject, one image is selected as the probe image and one of the other three images (selected at random on each trial) is the gallery image. When a subject is used for training, all three non-probe images are training images. Note that whether the sets of gallery and training subjects are disjoint is a configuration parameter.

The response variable is the recognition rank for a given subject and configuration of the algorithm. The PCA+LDA algorithm sorts the gallery images according to their distance from the probe image; recognition rank is the position of the subject's gallery image in this sorted list. Thus a recognition rank of 1 indicates that a nearest neighbor classifier will correctly match the subject to their gallery image. A recognition rank of 2 indicates that one imposter appears more similar to the probe image than the subject's gallery image does. Any recognition rank larger than 1 suggests an algorithm will do a poor job of recognizing the subject.

One of the most important aspects of algorithm configuration is training. In these experiments, training images were drawn from the 3 non-probe images of each subjects. For each test, the PCA+LDA algorithm was trained on $\frac{1}{4}$ of the subjects. Thus, the algorithm was trained on 192 images, 3 images for each of the 64 randomly selected subjects. An additional 192 images were added to the training in some trials to see if supplementing the training with altered versions of training images improves performance. Three specific sources of supplemental imagery were used (note that in all three cases one supplemental image is generated for each training image):

1. Reflect each image horizontally about its mid-line.
2. Add independent mean zero Gaussian noise (with a standard deviation of 2% of the dynamic range of the image) to the grey-scale value of each pixel in each image.
3. Shift each image left by one pixel.

The entire training procedure using 64 randomly selected subjects was itself replicated 8 times in order to determine how much the exact choice of training subjects matters.

The change in the size of the training set from 192 images without supplemental training to 384 with supplemental training is particularly important if it affects the dimensionality of the PCA space used for discrimination. Therefore, this dimensionality was explicitly controlled by setting it to either 114 eigenvectors (60% of 191 or 30% of 383) or the number of eigenvectors that comprise 90% of the energy, denoted $E_{90\%}$ [7]. Typical values of k using the $E_{90\%}$ rule to truncate the eigenspace ranged from 64 to 107.

To test the effect of gallery image choice on algorithm performance, each probe was tested on 30 different gallery sets. These galleries were generated by randomly selecting 1 of the remaining 3 images per subject. The same 30 randomly generated gallery sets were used throughout the entire experiment.

The tests in this experiment were run using the standard L_2 norm and the L_{Soft} measure proposed by Zhao [2] and defined above in equation 1. It was also brought to our attention in discussions with Zhao that performance may benefit from reducing the image size prior to recognition. Thus, two different image sizes were tested: the FERET standard size (130×150) and smaller images with $\frac{1}{4}$ as many pixels (65×75).

One additional variable was observed in this experiment: whether a particular probe had been included in the training set used in each test. Since there were 256 probes and 64 subjects in the training set, only 25% of probes were included in any training set.

The experimental design was a balanced complete factorial; in other words, we ran each choice in every combination. This results in $8 \times 30 \times 4 \times 2 \times 2 \times 2 = 7680$ attempts to recognize each probe image under different circumstances. A summary of the configuration and replication factors used in this experimental design is given in Table 1.

Factor	Index	Values
Training Set	t	$0, \dots, 7$
Gallery Set	g	$0, \dots, 29$
Probe Image	p	$0, \dots, 255$
Training Set Supplement	u	None, Shift, Reflection, or Noise
Metric	m	Standard or Soft L2, i.e. L_2 or L_{Soft}
PCA Space Dimension	d	114 or $E_{90\%}$
Image Size	s	Standard (130×150) or Small (65×75)
Training Set Inclusion Flag	f	No or Yes

Table 1. Summary of the configuration and replication factors used in the experimental design.

5 The Model

Let $R_{tpgusmdf}$ represent the recognition rank when the p^{th} probe image of the s^{th} size and f^{th} value of the training set inclusion flag is matched to the g^{th} gallery, using an algorithm with the m^{th} metric limiting the PCA discrimination space in the d^{th} manner, when the algorithm is trained on the t^{th} baseline training set with the u^{th} type of training set supplementation. Then let $Y_{tpusmdf}$ be the number of gallery sets for which $R_{tpgusmdf} = 1$. Note that $0 \leq Y_{tpusmdf} \leq 30$.

We model $Y_{tpusmdf}$ as a binomial random variable, namely $Y_{tpusmdf} \sim \text{Bin}(30, p_{tpusmdf})$. The probability of rank 1 recognition (namely $p_{tpusmdf}$) is modeled as:

$$\log \frac{p_{tpusmdf}}{1 - p_{tpusmdf}} = L + U_u + M_m + D_d + zI_t(p) + \tau_t + \pi_p + (UM)_{um} + (UF)_{uf} + (UD)_{ud} + (MF)_{mf} + (MD)_{md} + (DF)_{df} + (MDF)_{mdf} \quad (2)$$

where

- L is the grand mean
- U_u is the effect of including the u^{th} supplement to the training set
- M_m is the effect of using the m^{th} distance metric
- D_d is the effect of the d^{th} strategy for limiting PCA dimensionality
- F_f is the effect of the f^{th} option for including a probe in the training set
- $(UM)_{um}, \dots$ are 2- and 3-way interactions between main effects
- τ_t is the random effect of the t^{th} set of people used for training, and
- π_p is the random effect of the p^{th} probe image.

Commonly $\log(p/(1-p))$ is termed the logit transformation.

We are interested in the amount of variability in rank 1 recognition rate attributable to the training set, but the 8 training sets used in our experiment are merely a random sample from the variety of sets that might have been used. Therefore, we fit $\tau_t \sim N(0, \sigma_\tau^2)$ as independent random effects. Using the same rationale, we modeled the π_p as mean zero normal random effects with constant variance (σ_π^2), but with a more complex covariance structure. In particular, results from matching the same probe image in different circumstances are probably more correlated than matching different probe images. We assumed a compound symmetric covariance structure, which allows correlation ρ between any pair of $Y_{tpusmdf}$ corresponding to the same probe image. Results for different probe images were modeled as independent. Thus, the model accounts for the possibilities that some training sets were more effective than others and that some probes were ‘easier’ than others. It also allows for correlation induced by repeated measurement of the same probes under different experimental conditions.

This model is a generalized linear mixed model [8], or more precisely a mixed effects logistic regression with repeated measures on subjects. Such models can be fit routinely using restricted pseudo-likelihood or the MIVQEU0 method [9–12].

6 Results

We originally fit a larger model including the main effect for image size and all its two-way interactions. A major finding of our study was that image size did not significantly

affect rank 1 recognition rate, either overall or in any interaction. Through a sequence of model fitting steps, we removed all image size terms from the original model (and added the three-way *MDF* interaction) to obtain the final fitted model given in (2).

On the logit scale, the random training sets accounted for virtually no variability in rank 1 recognition rate. Any randomly selected training set of this size is as good as any other. Also on the logit scale, the correlation between probes of the same person under different conditions was about 0.70. Thus our findings strongly support the conclusion that some people are harder to recognize at rank 1 than others, and that this variation accounts for the majority of the total variability in rank 1 recognition rate on the logit scale.

The best configuration (among those tested) of the PCA+LDA algorithm used the \mathcal{L}_{Soft} distance metric with PCA dimensionality set at 60% of the number of eigenvectors and supplementing the training set with shifted images. Estimated probability of rank 1 recognition for this configuration was .87 for subjects included in the training set and .75 for those who were not. In comparison, the estimated probabilities were .77 and .74, respectively, for the ‘baseline’ PCA+LDA configuration that use the standard \mathcal{L}_2 distance metric and no training supplementation.

The main focus of our study is the estimation of the main effects and interactions included in (2). These allow us to quantify how changes to the PCA+LDA configuration affect estimated probability of rank 1 recognition. Table 2 shows tests of the statistical significance of the various fixed effects. Figures 1 and 2 show the various estimated effects, as described in more detail in the rest of this section. The threshold for statistical significance in these figures is substantially smaller than the magnitude of difference that would be scientifically important in the context of human identification problems. Therefore it suffices to examine these figures for interesting effects with the understanding that any interesting effect would be statistically significant according to a simple F-test. The significance threshold is so small because the large sample size provides a lot of discriminatory power.

Effect	Label (Num.)	d.f.	F value	p-value
Training Supplement	<i>U</i>	3	139.1	<.0001
Distance Metric	<i>M</i>	1	1087.8	<.0001
PCA Dimensionality	<i>D</i>	1	936.6	<.0001
Training Inclusion Flag	<i>F</i>	1	333.8	<.0001
<i>U</i> × <i>M</i> Interaction	<i>UM</i>	3	14.2	<.0001
<i>U</i> × <i>F</i> Interaction	<i>UF</i>	3	67.0	<.0001
<i>U</i> × <i>D</i> Interaction	<i>UD</i>	3	41.5	<.0001
<i>M</i> × <i>F</i> Interaction	<i>MF</i>	1	717.6	<.0001
<i>M</i> × <i>D</i> Interaction	<i>MD</i>	1	59.1	<.0001
<i>D</i> × <i>F</i> Interaction	<i>DF</i>	1	703.7	<.0001
<i>M</i> × <i>D</i> × <i>F</i> Interaction	<i>MDF</i>	1	5.04	.0248

Table 2. Table of significance of model effects and interactions based on approximate F-tests.




Fig. 1. Effect of training set supplementation on predicted probability of rank 1 recognition. See the text for codings for color, shape, fill, and line type.

Figures 1 and 2 employ a consistent labeling to illustrate the results. Color indicates training supplement (orange for none, green for reflections, red for noise, and blue for shifts). Shape indicates distance metric (square for standard \mathcal{L}_2 and circle for \mathcal{L}_{Soft}). Fill indicates whether a probe is included in the training set (solid fill for yes and hollow for no). Finally, PCA dimensionality is indicated by line type (solid lines for 114 and dashed for $E_{90\%}$).

Figure 1 shows the effect of training set supplementation on predicted probability of rank 1 recognition. Overall, adding shifted images was best and adding reflected images was actually worse than no supplementation. This refutes a hypothesis we originally held that reflections were effectively equivalent to shifts due to the rough horizontal symmetry of human faces. This figure also illustrates the interaction between the training set inclusion variable and supplementation type: the improvement or interference provided by training set supplementation was substantial only for recognition of probes that were included in the training set.




Fig. 2. Effects of distance metric, PCA dimensionality, and inclusion of the probe in the training set on predicted probability of rank 1 recognition. See the text for codings for color, shape, fill, and line type. There is no line type variation in the right panel because PCA dimensionality is labeled on the horizontal axis.

The left panel of Figure 2 shows the effect of distance metric on predicted probability of rank 1 recognition. Overall, the soft L2, \mathcal{L}_{Soft} , metric was superior, but this difference was mostly limited to instances when the probe was included in the training set. When training did not include the probe, very little was gained by switching from \mathcal{L}_2 to \mathcal{L}_{Soft} . This is the significant interaction.

The effect of including the probe in the training set is further shown in the center panel of Figure 2. Somewhat surprisingly, the overall effect of training set inclusion on predicted probability of rank 1 recognition was only slightly positive. There was an important three-way interaction with distance metric and PCA dimensionality which can be described as follows. The effect of training set inclusion was always positive when the PCA dimensionality was limited to 114. When the PCA dimensionality was set at $E_{90\%}$, the effect of training set inclusion depended on the choice of metric: positive effect for \mathcal{L}_{Soft} and negative for standard \mathcal{L}_2 . A definitive explanation for why inclusion

of the probe in the training set actually impeded recognition when using the \mathcal{L}_{Soft} with $E_{90\%}$ PCA eigenvectors eludes us.

The right panel of Figure 2 shows the effect of PCA dimensionality on predicted probability of rank 1 recognition. Overall, the use of $E_{90\%}$ PCA eigenvectors was inferior to fixing 114 eigenvectors. Since $E_{90\%}$ was virtually always less than 114, it is perhaps not surprising that this configuration was inferior. Again a significant interaction can be seen: the negative effect was mostly confined to probes that were included in the training set.

An important broader conclusion can be distilled from the details above. Variations in distance metric, PCA dimensionality, and training set supplementation exemplify the many sorts of strategies employed by researchers to improve the performance of the basic PCA+LDA algorithm. Such strategies appear to span a conceptual dimension with opposing approaches at each extreme. Direct manipulation of the training lies at one extreme (eg. training supplementation) and changes in discriminatory strategy of the algorithm lies at the other extreme (eg. changing the distance metric). Both ends of the spectrum are motivated by the hope that such changes will improve recognition for probes not used in training. Our results firmly establish that none of these strategies succeed in this goal. On the other hand, our results also show that either end of the strategic spectrum can yield improved results on probes that are included in the training set. Changes to the algorithm provide about twice the improvement as changes to the training.

7 Discussion

Since all the tested variations in the baseline configuration of the PCA+LDA algorithm only improved rank 1 recognition for probes that were included in the training set, our results raise several important questions. First, the earlier work of Zhao indicated the soft-L2 measure would improve performance for subjects not in the training set. The discrepancy between our findings and Zhao’s suggests more work may be needed before this matter is fully understood. Second, and more broadly, how might one best configure a PCA+LDA algorithm to enhance performance for such subjects.

Another avenue of future work might consider alternate and more generous definitions of recognition success. For example, do our results hold when a ‘success’ is defined as ranking the correct gallery image at rank 10 or higher? Although our statistical methodology could be applied equally well to such a problem, the FERRET data are not very suitable for such an analysis. About 47% of the 65,536 attempts at recognition under different configurations and training resulted in a rank 1 match for all 30 galleries; the overall frequency of rank 1 matches was 76%. If the response variable in our analysis were changed to, say, rank 10 recognition, the outcomes would be skewed even more toward successful matches and the reliability of the inference would be decreased, all else being unchanged.

The randomness inherent in our outcome—rank 1 recognition—is generated by the repeated testing of probes under 30 different galleries. In a real application, similar randomness applies to probe images as well. The advantage of the experimental design used here is that the binomial sample size parameter can be set quite large (we chose 30).

Consequently, the $Y_{t_{pusmdf}}$ are very informative about the recognition rate parameter, $p_{t_{pusmdf}}$. An alternative experimental design would have been to fix a single gallery and to test different probe images of the same person. In the FERRET data, however, there are only a small number of frontal images for most subjects. Thus, the observed recognition rate for each individual subject is less informative. When we have access to a large data set with many replicates per subject, we plan to run such an analysis varying probe presentation.

8 Conclusion

The results presented here represent one of the largest tuning studies of a significant face recognition algorithm on record. As an exercise in advancing the state of experiment design, it is great success. For example, the existence of strong interactions between whether a subject is in the training set and whether the algorithm does better or worse is expected, but characterizing this dependence in clear quantitative terms is significant. As to what this study tells us about our ability to significantly alter, in particular improve, algorithm performance, the results are somewhat disappointing. Clearly, there is a need for additional study of the PCA+LDA algorithm in order to truly demarcate those factors of algorithm configuration that may, in future, prove to significantly improve performance. At a minimum, such studies should adhere to the standards established here.

Acknowledgments

This work supported by the Defense Advanced Research Projects Agency under contract DABT63-00-1-0007.

References

1. Heath, M., Sarkar, S., Sanocki, T., Bowyer, K.: A robust visual method for assessing the relative performance of edge detection algorithms. *PAMI* **19** (1997) 1338–1359
2. Zhao, W., Chellappa, R., Krishnaswamy, A.: Discriminant analysis of principal components for face recognition. In: In Wechsler, Philips, Bruce, Fogelman-Soulie, and Huang, editors, *Face Recognition: From Theory to Applications*. (1998) 73–85
3. Phillips, P., Moon, H., Rizvi, S., Rauss, P.: The FERET Evaluation Methodology for Face-Recognition Algorithms. *T-PAMI* **22** (2000) 1090–1104
4. FERET Database: <http://www.itl.nist.gov/iad/humanid/feret/>. NIST (2001)
5. Turk, M.A., Pentland, A.P.: Face Recognition Using Eigenfaces. In: *Proc. of IEEE Conference on Computer Vision and Pattern Recognition*. (1991) 586 – 591
6. Beveridge, J.R.: The Geometry of LDA and PCA Classifiers Illustrated with 3D Examples. Technical Report CS-01-101, Computer Science, Colorado State University (2001)
7. Kirby, M.: *Dimensionality Reduction and Pattern Analysis: An Empirical Approach*. Wiley (2000)
8. Breslow, N.E., Clayton, D.G.: Approximate inference in generalized linear mixed models. *J. Amer. Statist. Assoc.* **8** (1993) 9–25
9. Rao, C.R.: Estimation of variance and covariance components in linear models. *J. Amer. Statist. Assoc.* **67** (1972) 112–115

10. LaMotte, L.R.: Quadratic estimation of variance components. *Biometrics* **29** (1973) 311–330
11. Wolfinger, R., O’Connell, M.: Generalized linear models: a pseudo-likelihood approach. *Journal of Statistical Computation and Simulation* **48** (1993) 233–243
12. Wolfinger, R., Tobias, R., Sall, J.: Computing gaussian likelihoods and their derivatives for general linear mixed models. *SIAM Journal of Scientific Computing* **15** (1994) 1294–1310