Biometrics of Cut Tree Faces

W. A. Barrett
billbarrett04@gmail.com
Department of Computer Engineering, retired
San Jose State University
San Jose, CA
October, 2007

Abstract

An issue of some interest to those in the lumber and timber industry is the rapid matching of a cut log face with its mate. For example, the U.S. Forest Service experiences a considerable loss of its valuable tree properties through poaching every year. They would like to have a tool that can rapidly scan a stack of cut timber faces, taken in a suspect lumber mill yard, and identify matches to a scanned photograph of poached stump faces. Such a tool would clearly fall into the category of a biometric identifier.

We have developed such a tool and have shown that it has usefully high biometric discrimination in the matching of a stump photograph to its cut face. It has certain limitations, described in this paper, but is otherwise eminently suitable for the task for which it was created.

Other applications may be found in the timber industry, as a way of tracking and controlling the movement of logs from the forest to the lumber mill, and of providing an independent means of inventory for timber stock.

Introduction

The science of biometrics probably began with a Chinese system of finger printing, in the 14th century, as reported by explorer Joao de Barros [1]. Nothing was made of this in the western world until the nineteenth century -- most police forces made use of human matching of live faces. As photography improved, matches using mugshots, also matched by a victim. Alphonse Bertillon developed a system of body measurements in the 1890s that reduced the problem of matching people to statistically comparing a list of numbers [1]. One must therefore credit Bertillon with creating the first automated biometric system, although his “automation”, prior to the computer, still required human calculation.

A *biometric measure* is some set of measurements made on an object that is intended to provide a unique or near-unique signature, typically a vector of numbers, of that object. A biometric measure is not designed to reproduce an exact image of the object, but rather to provide enough statistical variance that one such object can be distinguished from another through comparison of their signatures alone. A biometric measure may be considered a very lossy compression of an image, or other measure, into a small set of numbers. The loss is typically so large that the image cannot be reconstructed from the biometric measure, or, perhaps only in a very fuzzy way.

The main advantage of high compression is that a large original data set is reduced to a much smaller set, and that implies that a very large database of biometric subjects can be carried in a relatively small file. A smaller file can also be searched more rapidly than a large file, reducing the time required to find a match for a candidate in a large set of potential matches.

A set of biometric measures should also be *distinctive* and *orthogonal*. By *distinctive,* we mean that each such measure will vary significantly between different biometric samples. By *orthogonal,* we mean that the cross-correlation of the measures in the set will be reasonably low. A measure with low distinctiveness might as well be eliminated. A measure with poor orthogonality, *i.e.* one that is strongly correlated with some other measure, might also be removed from the set.

An important goal in obtaining a biometric measure is the *removal of extraneous information,* a process called *segmentation.* An original image may be of a face photographed in a background containing other people, desks, clocks, chairs, etc., all of which should be masked away before obtaining the biometric measure. The face itself might well be subjected to further masking to eliminate variable features such as the hair arrangement. One biometric uses the pattern of an eye's iris to obtain a biometric measure; it must clearly segment the iris out of the background of the rest of the eye, the face, and the pupil. Segmentation is achieved in a variety of ways for human recognition. In the case of digital fingerprints, segmentation is not even an issue, as the image itself encloses the print with essentially no extraneous background. Face recognition segmentation can be done reasonably well using motion video [2] or (for example) facial color, followed by an analysis of the eyes, nose and mouth that can lead to an efficient rejection of a non-facial image [3]. In our case, automated segmentation turned out to be a major problem, one that remains unsolved.

A third goal in biometrics is finding a suitable *comparison* algorithm. The problem this solves is reducing a pair of biometric measurement sets to a single value that estimates the difference between the two sets. Ideally, a large difference implies that the sets are of different objects, while a small difference implies the same object. But every such distance is subject to statistical variation. The variation of measures of the same object, taken under different circumstances and times, can be expressed in terms of an *authentics* distribution. The variation of measures of different objects is expressed as the *imposters* distribution. A suitable comparison algorithm, along with a suitable set of measures, should result in a reasonably clean separation between these two distributions. An overlap between the two means that a difference that is not clearly outside either of these distributions has a corresponding uncertainly in classification.

The classification uncertainty is typically expressed in a *receiver operating curve,* or *ROC* for short. An ROC is a plot of the integrated authentics distribution *vs.* the imposters distribution, for a representative set of biometric measures, with the biometric distance used as a parameter. One may consider this to be a plot of the false match rate, or *FMR,* against the false non-match rate, or *FNMR.* This curve resembles a plot of the inverse function $y= n/x$, for positive x. An *equal error rate,* or *EER,* is the distance parameter for which *FMR* = *FNMR.* When a biometric distance is less than the *EER,* we can say that the two measures represent the same object, though that conclusion has a much higher probability if the distance is considerable less than *EER.* Similarly, when a biometric distance is greater than the *EER,* we can say that the two measures represent different objects, again subject to an uncertainty that decreases rapidly with increasing distance. Figure 4 is an example plot of integrated authentics and imposters probabilities *vs.* a biometric distance.

Biometric Measure of a Log Face

In this work, we consider digital color photographs of the faces of cut timber, typically taken in daylight, with the camera axis essentially at right angles to the face. The image face should be wholly within the camera view, but does not have to be centered in the image. It turns out that the biometric measure that we have developed is reasonably independent of the camera angle, so that elaborate means of ensuring a correct view angle are not necessary.

The background in the images varies considerably, as does the number of faces photographed. The background may consist of partial side view of other logs, the sky, or other natural scenery. The background of a stump is typically a combination of twigs, grass, earth and other plants. Some images are partially obscured by other log faces or other growth.

Some upper limit on the number of faces is imposed by the camera resolution -- a large number of faces at low resolution will yield only a smear of a few pixels for each face, rendering any biometric measure useless.

The nature of the face image might be expected to be of the tree rings. Unfortunately, one also sees strong saw kerf patterns, which also present variations depending on the angle of the incident light and orientation of the log. There are often significant variations in coloration of the different parts of the faces, plus changes in coloration with age. A freshly cut pine log will have a strong white-yellow color, but this shifts to a darker yellow and later brown over a few weeks as the log ages. Some logs will also develop a split starting near the center toward the bark -- as the log dries out, it contracts more in the outer layers than the inner ones, causing a high peripheral stress to develop in the outer layers.

One major variation is in the *orientation* of the log face. After the tree is cut, its trunk is stripped of branches in the forest, then loaded into a carrier. The orientation relative to the tree stump is lost. The implication with respect to a biometric measure is that the measure must be orientation independent. That alone makes biometrics of log faces quite different from those used for human classification -- fingerprints are inherently oriented along the finger axis; faces are typically taken in an upright view, and can be oriented by finding the eyes; iris images can be oriented if necessary by examining the curve of the eye lids.

Orientation invariance requires an image transform that is orientation invariant. It also requires that we locate an origin in the segmented image that can be used as a rotation origin. We resolved the first problem by choosing a *pseudo-Zernike* polynomial moment [4, 5, 6]. The second problem is easily solved by first obtaining a good segmentation of the face's bark layer, then computing the center of mass of the face.

The moments of an image can be rapidly computed using pre-computed tables of coefficients, and methods developed by various authors, for example, [5]. It happens that all the orientation information appears in the first few low-order moments, of orders (0,0), (1,0), and (0,1). By dropping these, all the remaining moments are orientation invariants. It also happens that as the order increases, one obtains a measure of finer detail in the image. Moment theory assures that the moments are orthogonal -- each such measure contributes something useful to the biometric measure, with little or no interference from image features that contribute to other moments.

Circular moments, of which the pseudo-Zernike polynomials are one instance, are typically a combination of a radial component, in which the radius is considered to range from 0 to 1, combined with a circular component over the whole circle. The radial component is comprised of a polynomial whose degree increases with the order. The circular component has the general structure of a Fourier series expansion, as one might expect, with 0, 1, 2, 3, *etc.* cycles in each full circle.

One might expect that the origin would cause a singularity, but the form of the equations prohibit that, with all but one of the radial functions vanishing at the origin.

One problem with the use of pseudo-Zernike polynomials for biometric measures is that the complete set of orders permitted is much larger than a minimal orthogonal set. One must select a subset of all possible orders which is orthogonal. Belkasim *et al* [6] describe how to do that.

A second problem is that polynomial expansions typically require computing some high powers of complex numbers, together with additions and subtractions of these numbers. Powers of size 15 are needed for good resolution, yet such high powers cause significant loss of precision with 8-byte IEEE floating-point numbers. Power above 20 are essentially useless for the purpose, and merely generate quantization noise in the measures. Belkasim recognized this problem and proposed using fractional powers instead of integral powers, a measure that significantly extends the number of useful orders by reducing quantization errors. Fractional powers of floating point numbers requires a bit more calculation time, but are worth it in the improved quality of the results.

Our system therefore follows the Belkasim system. It permits the sophisticated user to select a good set of orders to be used in the biometric measurements. We have found that by just dropping a few of the low order terms, and keeping a set of some forty floating-point measures, provides a good biometric discrimination in the several hundred log faces examined with our software tools.

Software Platform

The Forest Service requires a tool that can be operated by any ranger with minimal training. At a minimum, this entails having the range take digital photographs of a stump and log faces, then sending these to an analysis center for matching purposes. The equipment required in the field is therefore simple and cheap. Almost any modern digital camera has sufficient resolution to carry this out.

We would also like to know the camera distance and focal length when an image was taken, as this would also provide an estimate of the physical size of each face, but this was not done. The face size would be a useful and powerful biometric, when combined with the other feature measurements.

At the analysis center, a technician is required to segment the faces using a Windows platform. Each image is brought up on a screen (figure 1). Each face can be accurately delimited with a few mouse movements and clicks, using a unique closed cubic-spline fitting tool (figure 2). Once a face is delimited, its enclosing polynomial is saved in a database. The polynomial is also used to compute the face's biometric measure, an operation that is typically completed in a fraction of a second per face.

Face matching can proceed when several images have been so entered and segmented.

If the operator is interested in finding those log faces that match a particular stump, the stump image is brought up, the stump selected, and then one key click locates a set of cut faces that most closely match the stump. These matching faces are ordered by closeness of match, *i.e.* by increasing biometric distance measure from the stump's biometric measure.

The matching faces are shown in a special window in which both the stump image and the face image can be visually compared (figure 3). A unique algorithm examines both of these images, and arrives at a canonical "rotation angle" for each of them. The face image is then rotated for display by the difference between these angles, such that it will typically appear to be orientationally aligned with respect to the stump. Needless to say, both images are scaled to exactly fit into their square image frame. The result of this automatic scaling and rotational alignment is to make a visual comparison of the two images very easy. After all, the biometric measurement process is not exact, and sometimes produces a false match. These can usually be eliminated through a visual comparison.

Although this logface tool has been designed to be as user-friendly as possible, with many safeguards against foolish activities, some training is nevertheless required to install and use it.

We were privileged to train one computer-astute person in its use at the San Dimas office of the Forest Service's research group. The main problem he faced was in getting the product installed. It turned out that his office Windows XP system was non-standard, owing to security and virus precautions enforced by their computer support staff. We were able to install it through a more primitive approach. Installation is otherwise automated. Our operator easily learned to use the tool for face segmentation and matching in less than an hour. Needless to say, the tool is richly documented for both the casual operator and a more sophisticated one interested in tuning its biometric parameters.

That experience, together with the need to segment each face manually through the tool's windows interface, suggested that a central bureau should be used to carry out the image analysis, rather than the field agents. A central bureau could also explore biometric quality and make appropriate adjustments. They would also be in a better position to identify matches across forest districts than could someone in the field. That effort has so far not been made.

Biometric Quality

In order to better gauge the quality of the biometrics, we incorporated two tools into our system. Both require a special “training set” of images, in which the matching faces have been manually identified.

One tool uses this set to construct the imposter-authentics ROC data in the form of an Excel spreadsheet data list -- see figure 4 for an example. This makes use of an Euclidean distance algorithm applied to the biometric data.

The other tool produces a very detailed report in Excel form of all the biometric measures, organized by image index and match index (not shown). This report, though voluminous, permits one to experiment with different distance measures, and also to examine the separate variances of each of the measure categories.

We have estimated the biometric quality of matching using a reasonably large set of log face images. Unfortunately, we do not have more than a few dozen images of known matching faces and stumps, too small a sample to yield a good ROC estimate. In fact, those pairs show a high quality of matching -- the system has consistently located the matching face to every stump in our collection, with the matching face at the top of the match list.

We have also compared log faces with other log faces, using, of course, different photographic images of the same face. This provides a larger set of matches than stumps to log faces, though this obviously can be criticized as comparing one image of an object to another image, and that should provide a close match. A tentative result using a random sampling of known image-pairs drawn from the known matches and non-matches yielded an ROC with a cross-over probability of 0.04 (figure 4). This essentially means that if 25 random samples are compared to a given candidate, the probability is even that the correct matching sample will be chosen from the set.

The ROC also provides a good estimate of an appropriate selection threshold, set to select as large a sample of authentics as possible, and as few imposters. This provides a good initial selection from a large sample of candidates, from which a visual inspection of a few top candidates is sufficient to find a matching face, or to decide that there are no matching faces.

Summary

A software tool designed to match a cut log face image with that of its mating face or stump, using biometric principles, has been designed and implemented. It accepts a set of digital images, then provides a means of segmenting the log faces in the images. Once segmented, the faces can be compared using a pseudo-Zernike polynomial moment comparison approach. Matching or nearly-matching faces are brought up in a matching tool for a final manual comparison by the operator.

Some unique properties of this tool include the use of orientation-invariant pseudo-Zernike polynomial moments, face segmentation using a cubic-spline fitting scheme, and a matching tool that automatically rotates a candidate image for final visual matching purposes.

Automatic segmentation of the faces has not been achieved. There are several possible avenues that could be explored in this regard, but this appears to be a difficult problem due to the varied background and the frequent resemblance of the background to the face.

Once segmented, the face matching appears to be excellent, and in line with reports of similar matching experiments using pseudo-Zernike moments.

The interested reader may download a version of this tool through the author’s web site, http://www.engr.sjsu.edu/wbarrett. Installation directions are provided in this link.

Acknowledgements

This work was supported through a private contract between the author and the U.S. Department of Agriculture, through the Forest Service Technology and Development group, San Dimas, CA.

I'd like to thank Mr. Ed Messerlie of the Forest Service for instigation of this contract, his close working relationship, and his patience, as well as providing us with numerous digital photographs. Thanks also to Andy Horcher of the San Dimas office, for his continued support.

References

[1] http://ctl.ncsc.dni.us/biomet%20web/BMHistory.html

[2] Motion detection uses a video camera on a fixed ount viewing a near stationary background. The image collection software continuously compares a frame with its predecessor, looking for significant changes. A person entering a room can therefore be detected and tracked, and a face entering the field of view can be segmented by comparing a known background image against the current one.

[3] Sobottka and Pitas, *Face Localization and Facial Feature Extraction Based on Shape and Color Information,* Proceedings, International Conference on Image Processing, Sept. 16-19, 1996, Lausanne, Switzerland.

[4] R. Mukundan, K. R. Ramakrishnan, *Moment Functions in Image Analysis,* World Scientific, 1998

[5] Chong, Mukundan, and Raveendran, *An Efficient Algorithm for Fast Computation of Pseudo-Zernike Moments,* Intl. Conf. on Image and Vision Computing, IVCNZ01 New Zealand (Nov. 2001), pp 237-242.

[6] S.O. Belkasim, M. Shridhar and M. Ahmadi, *Pattern Recognition with Moment Invariants: A Comparative Study and New Results,* in Pattern Recognition, Vol. 24, No. 12, pp. 1117-1138, 1991.

Figures

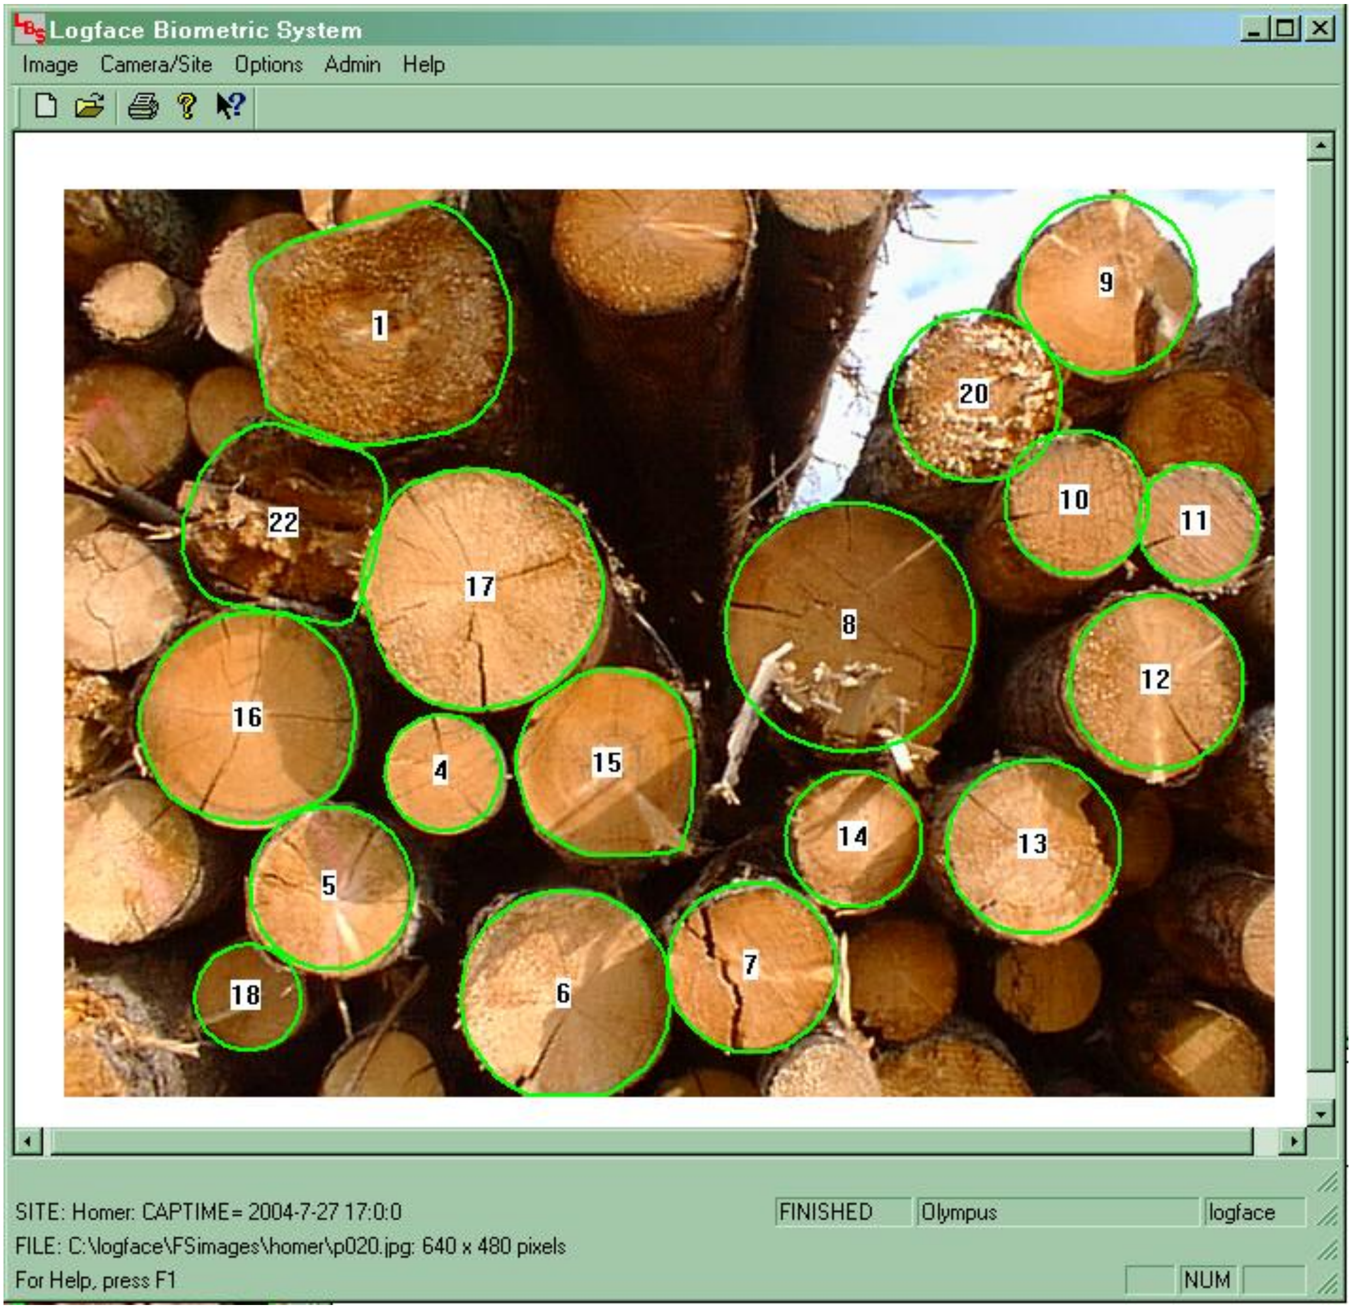

SITE: Homer: CAPTIME= 2004-7-27 17:0:0

FILE: C:\logface\FSimages\homer\p020.jpg: 640 x 480 pixels

For Help, press F1

FINISHED | Olympus | logface

NUM

Figure 1. Editing panel for the Logface Biometric System. Separate faces are manually delimited through a spline-curve fitting tool, as shown by the green outlines. Face numbers within this image are shown.

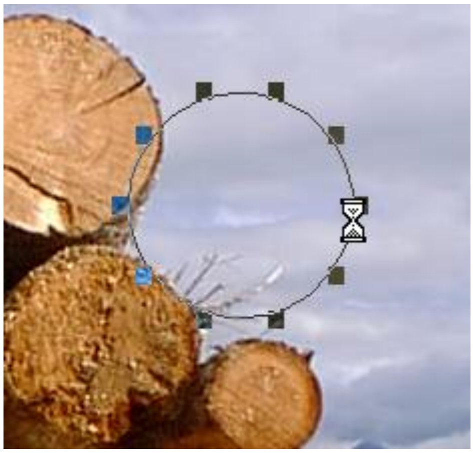

Figure 2. Appearance of an initial spline curve fitting tool. The spline control points can be moved about, merged or created, to fit a face. The whole pattern can also be dragged to a face.

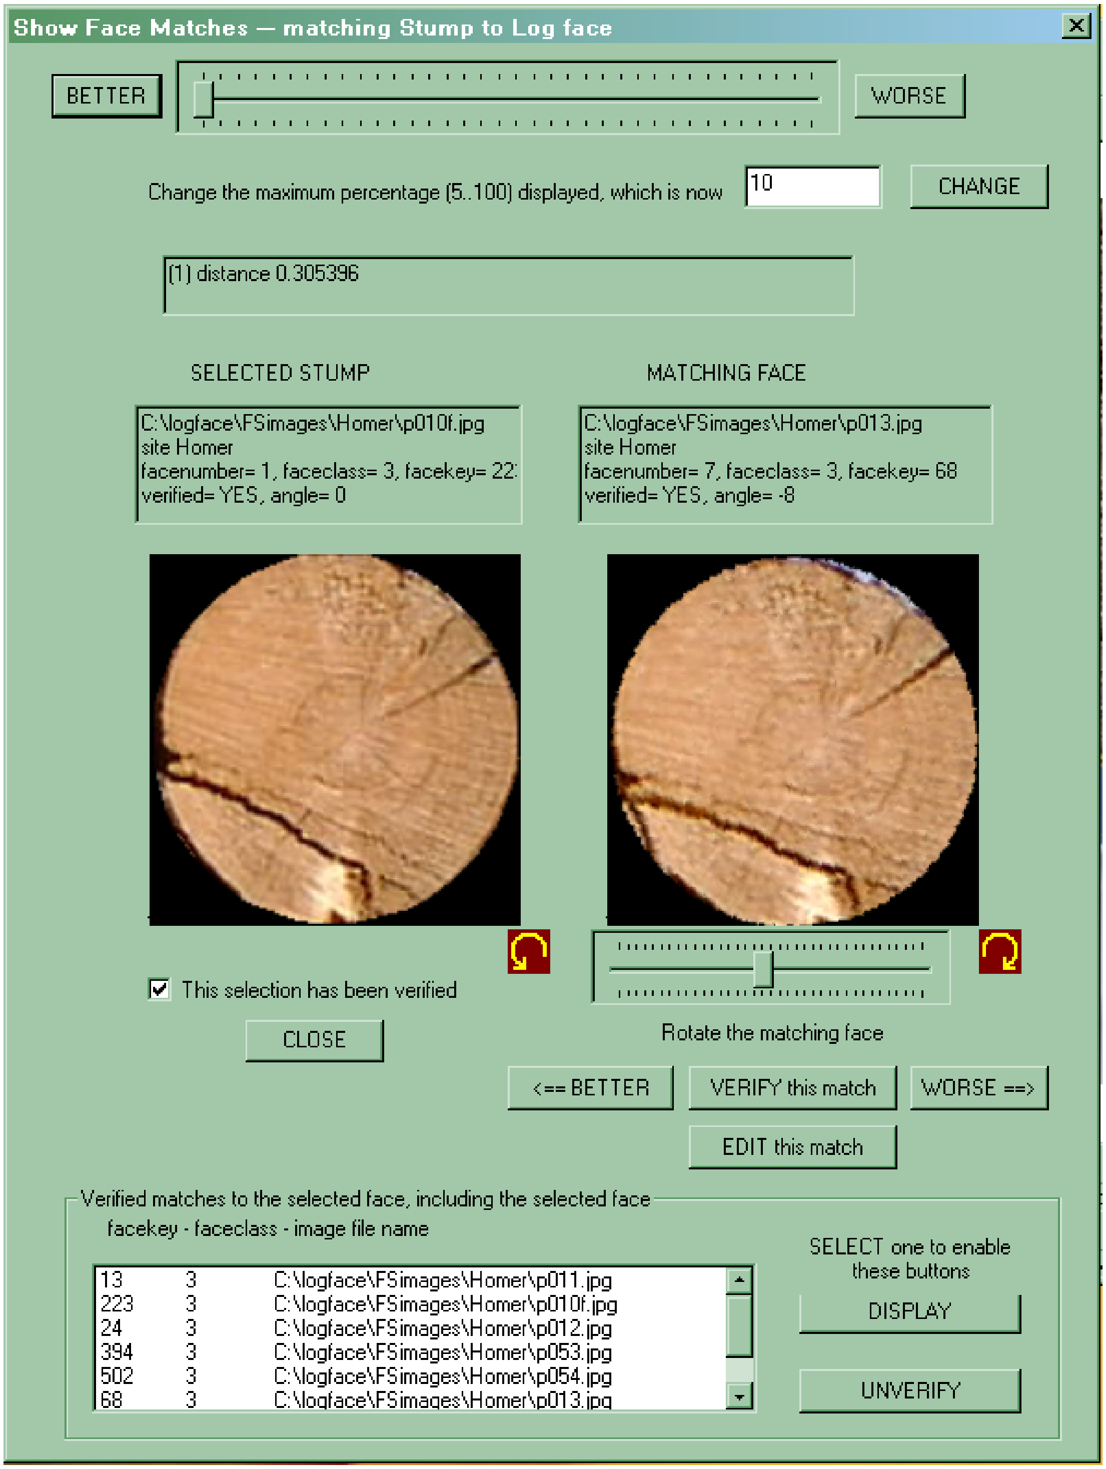

Show Face Matches — matching Stump to Log face

Change the maximum percentage (5..100) displayed, which is now 10

(1) distance 0.305396

SELECTED STUMP

C:\logface\FSimages\Homer\p010f.jpg
site Homer
facenumber= 1, faceclass= 3, facekey= 22
verified= YES, angle= 0

MATCHING FACE

C:\logface\FSimages\Homer\p013.jpg
site Homer
facenumber= 7, faceclass= 3, facekey= 68
verified= YES, angle= -8

☑ This selection has been verified

Rotate the matching face

Verified matches to the selected face, including the selected face

facekey - faceclass - image file name

13	3	C:\logface\FSimages\Homer\p011.jpg
223	3	C:\logface\FSimages\Homer\p010f.jpg
24	3	C:\logface\FSimages\Homer\p012.jpg
394	3	C:\logface\FSimages\Homer\p053.jpg
502	3	C:\logface\FSimages\Homer\p054.jpg
68	3	C:\logface\FSimages\Homer\p013.jpg

SELECT one to enable these buttons

Figure 3. A face-matching panel. The image on the left has been selected as the candidate for a matching search. That on the right looks identical, but has appeared in a separate image, scaled and rotated to best match the left image. The comparison can be verified by the operator, and marked as taken of the same face.

The slider at the top selects one of a set of biometrically matching faces, ordered by closeness of match. The table in the bottom lists the top match candidates, providing more details.

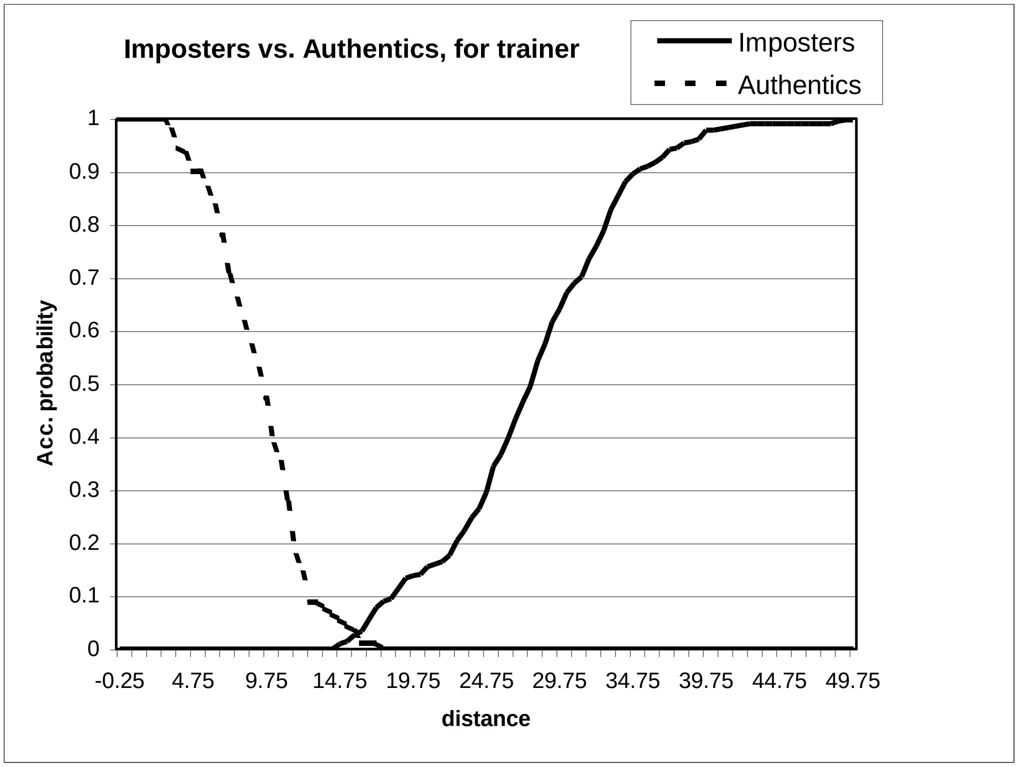

Figure 4. Integrated imposters vs. authentics probabilities. This is estimated from a sample of several hundred manually matched "trainer" images. Most of the matching images are of the same face, photographed at different distances and viewpoints.