

Evolving the Face of a Criminal: How to Search a Face Space More Effectively

Charlie D. Frowd*¹, Vicki Bruce², Carol Gannon¹, Mark Robinson³, Colin Tredoux⁴, Jo Park³, Alex McIntyre³ and Peter J.B. Hancock³

¹*Department of Psychology, University of Central Lancashire, Preston, UK*

²*College of Humanities and Social Science, University of Edinburgh, Edinburgh, UK*

³*Department of Psychology, University of Stirling, Stirling, UK*

⁴*Department of Psychology, University of Cape Town, Cape Town, SA*

*Corresponding author: Dr Charlie Frowd, Department of Psychology, University of Central Lancashire, Preston PR1 2HE. Email: cfrowd@uclan.ac.uk. Phone: (01772) 893439.

Abstract

Witnesses and victims of serious crime are often required to construct a facial composite, a visual likeness of a suspect's face. The traditional method is for them to select individual facial features to build a face, but often these images are of poor quality. We have developed a new method whereby witnesses repeatedly select instances from an array of complete faces and a composite is evolved over time by searching a face model built using PCA. While past research suggests that the new approach is superior, performance is far from ideal. In the current research, face models are built which match a witness's description of a target. It is found that such 'tailored' models promote better quality composites, presumably due to a more effective search, and also that smaller models may be even better. The work has implications for researchers who are using statistical modelling techniques for recognising faces.

1. Introduction

For people we know well, recognition normally occurs accurately and effortlessly. For less familiar faces, recognition involves greater error and uncertainty (e.g. [1]). Likewise, computer recognition systems involve error and appear to perform similarly to human observers perceiving unfamiliar faces (e.g. [2,3]). One approach to improve performance is to converge information from several sources. For example, verbal descriptions have been used for many years to help locate visual exemplars in semi-automatic systems (e.g. [4]). A recent trend has been to combine different biometric modalities – face, iris and fingerprints (e.g. [5]) – or different types of descriptive

information [6]. Other approaches may combine different viewpoints or different images of the same person [7,8]. Here, we evaluate an enhancement to a face production system that allows witnesses to construct faces of criminals. The heart of this technology is a model built using Principal Components Analysis, or PCA, a typical component of many face recognition systems (e.g. [9,10]). The work combines visual and verbal information provided by a witness to both fine-tune the face model and evolve a better likeness of a target.

Witnesses and victims of crime are often asked to describe the appearance of an unfamiliar face – a suspect – and then to construct a facial composite. These images are sometimes shown in the newspapers and on TV crime programmes in an attempt to locate a criminal. The traditional procedure is for witnesses and victims to select individual facial features from a kit of 'face parts'. However, this is not an optimal procedure since we do not perceive faces as a set of parts – rather, more as a complete or *holistic* image (e.g. [11]). Unsurprisingly then, composites constructed using the 'feature' method are not recognised very well (e.g. [12-16]). Following a realistic delay of several days, recognition tends to be very poor indeed (e.g. [17,18]).

An alternative is to present sets of complete faces for a witness to select, a more natural procedure than viewing individual facial features; it is also somewhat similar to a witness selecting faces from a police line-up or mugshot album. While there are several such whole face systems (e.g. ID [19]; EigenFIT [20]), the focus here will be on the EvoFIT composite system [13,18,21-23] developed jointly by the Universities of Stirling and Central Lancashire. With EvoFIT, witnesses select from a set of complete faces, and the selected faces are bred together using a Genetic Algorithm to produce another set for selection. While

the faces contain random characteristics at the start, repeating the selection and breeding process a few times enables a specific face to be ‘evolved’ – hence the name Evolving Facial Identification Technique, or EvoFIT. In practice, witnesses’ first select facial *shape*, corresponding to the size and position of features, then facial colouring or *texture*, the colour of the eyes, brows, mouth and overall skin tone. To model hair, a specific hairstyle is chosen from about 500 alternatives and presented on each face (which is subsequently blurred, see below). Tools are available to manipulate the shape and position of features on demand as well as the more overall or *holistic* properties of the face (e.g. age and masculinity).

Central to the current EvoFIT are a set of face models that are able to generate a large number of synthetic faces. The models are constructed from PCA, a statistical technique that extracts the dimensions of variation – the eigenvectors – from a set of items, in this case faces. These ‘standard’ models are built from 72 front-view, white male faces for a given age range. In tests, EvoFIT produces composites that are correctly named on average about 10-15% of the time after a realistic delay of two days. While this figure is still rather low, it is at least twice that of a current ‘feature’ system [18,24]; as such, it is now available for UK police use. More recent developments have improved performance (e.g. [16,23]).

A problem with EvoFIT has been to ensure reliable convergence on a target face. One of the problems with PCA face models is their complexity: in being able to generate a wide range of faces, they contain a great deal of information, which is difficult to search, even with genetic algorithms. One successful approach is to run the system more than once with a witness, each time using a new set of random faces [24], capitalizing on the fact that random starting points in a complex search space can produce quite different results. An alternative is to be more selective about the faces used to build the model in the first place. For example, if a witness remembers a suspect’s face to be thin with small eyes, a model built containing faces that match this description is likely to be valuable. Such a model would not generate wide faces, nor faces with large eyes (since it is based on the statistics of the reference faces). A ‘tailored’ model of this type could be searched more effectively and better likenesses evolved more often.

The thrust of the current work is to explore the effectiveness of evolving faces using standard face models that are matched on age, the approach used with EvoFIT to date, versus tailored face models, which are built specifically to match a witness’s memory of a target. Two experiments are presented. The first explores the quality of composites constructed

from standard and tailored models built using the same number of faces; the second compares a range of tailored models of different sizes. It was expected that a tailored model would outperform a standard one, and that reducing the complexity of the model still further, by reducing its size, would be even more effective.

1.1. Detailed background to EvoFIT

The procedure used to build the current face models follows that of Sirovich and Kirby [25] and Troje and Vetter [26]. The initial stage was to carefully photograph 200 white male faces in a front-view pose (the system currently works for Caucasian male faces). These subjects possessed a wide age range, from about 18 to 65, as shown in Table 1, and allowed 4 overlapping face models to be constructed, each containing 72 faces centred at 10 year intervals: 20, 30, 40 and 50 years. Models were constructed in greyscale, since there is presently no evidence that colour improves results [24].

Table 1. Age distribution of reference faces used to build the PCA face models.

Age	15-20	20-29	30-39	40-49	50-75
Count	13	63	51	45	28

The next stage was to identify the position of individual landmarks on the 200 reference faces. This is largely a manual procedure that involves the careful identification of about 300 standard locations on each face (see Figure 1, far right). For the standard models, these shape co-ordinate files are then subjected to a Principal Components Analysis (PCA) to provide a set of 72 reference shapes, or eigenshapes. While the eigenshapes can be added together in different amounts to reconstruct the original shapes, adding them in random proportions allows a novel shape to be generated. Note that the face model constructed by PCA is itself holistic in nature, since the eigenshapes change the overall appearance of the face [27]. While the behaviour of these eigenshapes tends to be fairly complex, normally one of the components models face length and width.

The features of the reference faces are then morphed to a standard size and position and a second, texture PCA is carried out on the image pixel values. This produces a set of reference eigentextures that allow random textures to be generated. To produce a random face, a random texture is generated that is morphed to a random shape. Hair is selected via the PRO-fit ‘feature’ system and added to the facial texture prior to the shape morph. The result is a good quality synthetic face, as shown in Figure 1.

Witnesses are presented with 72 facial shapes, on 4 screens of 18 faces, the maximum number that can sensibly be displayed on a computer monitor, and select about 6; in the same way, 72 facial textures are presented and they and similarly select about 6. Witnesses go on to select a single face that contains the shape and texture with the best likeness: known as the ‘best’ face. The selected faces are then subjected to a Genetic Algorithm (refer to [28], for an introduction to GAs). To do this, faces are selected in pairs, first for shape, then texture. Each time, an offspring is produced containing a random mix of coefficients from both faces (uniform crossover). All faces are selected as parents for breeding with equal probability, except the best, which is given twice the number of breeding opportunities. The best is also copied directly to the next generation, an elitist strategy which avoids damage through breeding. For all other faces, a small amount of mutation, a probability of 0.1, is applied that replaces each coefficient with a random value. The breeding process is repeated to create a new set of 72 shapes and 72 textures for selection. Witnesses normally inspect 3 generations of faces before selecting one that is saved to disk as the composite image. They are also given the opportunity to enhance the face artistically: adding shading, stubble or eye bags, for example, as required.



Figure 1. From left to right: random shape, random texture, combined shape and texture, location of shape landmarks

A small utility called the Feature Shift is available to change the shape and position of features in the best face. Such changes might include moving the eyes together or making the mouth larger. To do this, the points of the face are moved and a best fit is carried out in the shape model to maintain a holistic code for the face. Two additional enhancements have been found to be of benefit. Firstly, once a hairstyle has been selected, the set of so-called *external facial features* – the region comprising the hair, ears and neck – are blurred until the end of the evolution. This procedure allows a witness to focus on and thereby improve the internal part of the composite, which is important for recognising the face later (e.g. [15,29]). Secondly, a set of ‘holistic’ tools has been developed that allow an evolved face to be re-worked holistically [23]. These

make a face appear older, more masculine, or even more threatening; eight such dimensions have been implemented.

1.2. Building a ‘tailored’ face model

To build a face model that matches a witness’s description of a target, referred to here as a ‘tailored’ model, descriptive labels were first assigned to each of the 200 reference faces. The descriptive labels used were taken mainly from the ‘Aberdeen’ Index [30], which are also used to classify features in the UK PRO-fit and E-FIT composite systems. Examples include brow thickness, eye colour and mouth shape.

A total of 70 facial features were classified by giving each a whole number value (rating). These features were themselves grouped into 7 categories: face shape, hair, eyes, nose, mouth, brows and holistic attributes. Some features were rated along an interval scale – e.g. brow colour, size and thickness. Ratings were typically given a value between 0 and 2; brow thickness for example: 0=thin / 1=average / 2 = thick. Some features were categorical and were classified along dimensions that were broadly ordered by feature similarity. For example, for chin shape: 0=square / 1=oval / 2=round / 3=angular / 4=triangular. Some features were clearly dichotomous, and were given a value of 0 or 1. Examples include: broken nose, slanting eyes and crow’s feet. Faces were also given a value along each of the eight dimensions used in the above holistic tools; these were the mean participant scores obtained from [23].

In use, a witness first recalls about 3 distinctive features of their target face. An error score (the absolute numerical difference) is then calculated between each specified feature and the relevant rating for each face in the database. Those faces with the lowest overall error score are used to build the model. To avoid building models that are too old or too young, an additional classification is made that excludes all faces that are not within about 15 years of the target face, also specified by the witness.

2. Experiment 1 – Standard vs tailored models

The first evaluation explored the effectiveness of a standard face model, built using faces of a similar age to a target, compared to a tailored model, built from faces that broadly matched on age as well as other facial characteristics. The basic procedure involved recruiting 24 participants to serve as ‘witnesses’, with each constructing a single composite using EvoFIT. Half of these participants constructed a face using a standard EvoFIT model, and the other half provided 3

to 4 distinctive features from which a tailored model was built; all models contained 72 faces. The targets in the experiment were six members of the Psychology staff at Stirling University.

Witnesses who were unfamiliar with the target faces were recruited: in order to effect this important parallel to police practice, witnesses were drawn from another university department who did not know the targets in the study. Composites constructed by these witnesses were given to psychology staff and final year psychology students, all of whom were familiar with the targets, and who were asked to name the person represented by the composite. Thus, composite quality was assessed by composite naming, and can be considered an analogue to real life usage of composites. An additional task was administered that required further participants, also familiar with the targets, to identify each composite from just the inner part of the face, the so-called internal facial features, which are important for naming (e.g. [29]) – see Figure 2 (far right). This was carried out to check that composite naming was not being driven exclusively by the presence of hair, which can be an important cue when the number of potential targets is fairly small – in the present study this totalled 24 academic and support staff.

2.1. Procedure

Laboratory witnesses were tested individually and asked to watch a short video of an unfamiliar member of staff. This was carried out in the knowledge that a composite would be required the following day. There were 6 target videos, each of a different member of staff giving directions to a local train station, and lasted for about a minute. Each video was shown to 4 people, two of whom went on to construct a composite of the target using a standard model, the other two, a composite using a tailored model.

Witnesses returned 24 hours later and were met by an experimenter, a person experienced in the use of EvoFIT, who helped them construct a composite. The experimenter was blind to the identity of the targets until all composites had been constructed. Each person first described the appearance of their target face with the assistance of a Cognitive Interview (e.g. [31]). This involved describing the face freely (free recall) and then attempting to recall further details about each feature (cued recall). Those witnesses in the tailored model condition were additionally asked to identify three to four distinctive features of the face and a bespoke model was then built; those in the other condition used the standard model that matched on age.

Both groups went on to evolve a composite with EvoFIT, as described above. Thus, they first selected

an appropriate hairstyle that was displayed in blurred form, along with the ears, neck, etc. They then selected 6 facial shapes from a set of 72 shapes, then 6 facial textures from a set of 72 textures. Next, the best face was selected from the best combination of shape and texture. The faces were then bred together and this procedure repeated until a good likeness was evolved, whereupon the blurring filter was switched off. The tool for manipulating specific facial features, the Feature Shift, was offered for use on the best face from the second cycle onwards. The final face was reworked using the holistic tools, to manipulate the face's perceived age, masculinity etc. Lastly, witnesses were given the opportunity to enhance the face using the GIMP, an artistic package available at no cost (<http://www.gimp.org/>). The experimenter improved the likeness of the hair, mainly by adding or deleting textured areas, or adding stubble to the chin, as directed by the witness. Composites took about an hour to construct.

The 24 composites were given to a group of 18 participants to name (as described earlier). Twelve additional composites of unfamiliar faces were added to this set to make the task more challenging (and little more life-like). Participants were told that many of the composites were constructed of members of staff and were to try to name them. They were also told to expect more than one composite of the same member of staff. Each person was tested individually and shown the composites sequentially. Afterwards, as a check to verify that the targets were known, they were shown a static photograph of the targets and similarly asked to name them. The order of presentation of composites and target faces was randomised for each person. A further group of 17 participants were similarly shown just the inner part of the face, the so-called internal facial features, and for each were asked to select the most likely person from a list of 12 written names, a list containing the original six names mixed in with a further six staff.

2.2. Results and Discussion

Participants were very familiar with the target set, correctly naming them 97.2% of the time. Composites constructed using the standard face model were correctly named on average quite well, at 35.2%; those using a tailored model were much better, at 54.2%. Example composites are presented in Figure 2. A two-tailed paired samples t-test applied to the participant data confirmed that the tailored models performed significantly better than the standard variety, $t(17) = 6.03$, $p < .001$; the relatively weaker items analysis approached significance, $t(5) = 4.43$, $p = .06$. An analysis of the incorrect names mentioned by

participants was included, which provides an indication of guessing, and was very similar across both conditions ($M \approx 14\%$). For internal features composites, identification was significantly better for the tailored model ($M = 37.3\%$) compared with the standard ($M = 24.5\%$), $t(16) = 2.85$, $p = .012$, by-subjects. While the items analysis was not significant, a more powerful by-item test was run with the type of task (complete composite / internal features) as a between-subjects factor and the type of model (standard / tailored) as within-subjects. The ANOVA indicated that tailored models were better, $F(1,5) = 7.10$, $p = .045$; all other $F_s < 2.34$, $p > .1$.



Figure 2. (From left to right) A composite face evolved using a standard model; a face evolved of the same target using a tailored model built from the witness's description of 'spiky hair', 'chiselled jaw' and 'large nose bridge'; a photograph of this target face; an example internal features composite used in the evaluation.

It would appear valuable then to build a model from faces that matched a target on a few distinctive features mentioned by a witness. The improvement in naming was sizeable, at almost 20%. Note that only five of the six target faces were better named overall when a tailored model was used. It is likely that the sixth target was problematic to construct as this person is in his 60s and there were relatively few faces from which to produce a model (refer to Table 1). As such, the tailored model may have behaved much the same as the standard one. Better performance would be expected for a tailored model had a greater number of older aged reference faces been available.

3. Experiment 2 – Different sized tailored models

The previous experiment demonstrated benefit for a tailored model, one built from a subset of faces whose features matched a witness's memory of a target face. In Experiment 2, we explored whether further benefit could be obtained using a tailored model built from fewer faces. One would expect smaller models to contain reference faces that better match a description, thus promoting an even better composite. Clearly, there

is a lower limit in the number of faces that should be used to build a model, since it needs to generalise well for faces that match the description. Early work on the development of the system used about 30 faces [22], which appeared to work fairly well, and so this was taken as the smallest model size. The number of faces used in a standard model was taken as the largest, and rounded down to 70 faces for convenience, and compared with an intermediate one of 50 faces. Thus, 3 tailored models were built from 30, 50 and 70 faces and then evaluated.

A procedure similar to the previous one was employed to compare the ability of these models to construct recognisable composites. A somewhat realistic procedure was employed in Experiment 1, with witnesses recruited to construct a single composite of an unfamiliar face. This time, a more powerful within-subjects design was used whereby an experienced EvoFIT operator evolved all the composites from her memory alone. To allow the resulting faces to be evaluated by adults in general, thus facilitating ease of participant recruitment, the targets were well-known UK celebrities, and included Gordon Brown (politician), David Tennant (actor), Declan Donnelly (TV presenter), Simon Cowell (TV celebrity), Daniel Craig (James Bond) and Wayne Rooney (footballer). Each of these identities was a familiar face to the operator and was evolved for each size of model (30 / 50 / 70), a total of 18 composites.

3.1. Procedure

The EvoFIT operator looked at one of the celebrity photographs for one minute, to refresh her memory of the face, constructed a tailored face model as before based on 3 distinctive features of the face, and then evolved a composite. This procedure (including looking again at the photograph) was repeated for all three model sizes and then for the remaining five celebrity targets. The same hairstyle was selected for each target. Each composite took about an hour to construct and, working full-time with sensible breaks, the operator produced the set of 18 within 3 days. The order of construction using the three sizes of model was randomised and based on a Latin Square design that allowed all possible combinations to be used.

The composites were printed at 8cm wide by 10cm high on A4 paper, one per page. A group of 17 participants volunteered to identify the composites, comprising adult visitors to a newsagents and health club in Wigan, UK. They were presented with each composite in sequence and asked to select the most likely candidate from a list of 6 written names. Participants were tested individually and were self

paced. The order of presentation of the composites was randomised for each person.

3.2 Results and Discussion

Example composites produced are presented in Figure 3. Identification was lowest from composites evolved from the largest model ($M = 63.7\%$, $SD = 32.3\%$); it was about 10% better from the smallest face model ($M = 71.6\%$, $SD = 19.8\%$) and better again by a similar amount from the intermediate one ($M = 81.4\%$, $SD = 13.6\%$). A repeated-measures ANOVA of these subject data was significant for model size, $F(2,32) = 5.9$, $p = .006$, and simple contrasts of the ANOVA confirmed that the 50 face model was superior to the other two, $p < .05$; a two-tailed paired t-test provided weak evidence for a benefit of the 30 over the 70 face model, $t(16) = 1.73$, $p = .104$. While the by-items ANOVA was not significant, $F < 2$, relative to the largest model, the intermediate one produced composites that were on average better identified on 5 of the 6 celebrities; a one-tailed t-test also suggested a non-significant trend between these two model sizes, $t(5) = 1.57$, $p = .089$.



Figure 3. Example composites of the British footballer, Wayne Rooney. They were evolved from tailored models of size 30, 50 and 70 faces (from left to right).

In summary, the intermediate size face model was found to evolve a better quality composite than either the larger or the smaller versions. The expectation was for an inverse relationship between model size and composite quality, but this was only found to be partially true: model performance improved considerably when reducing the build from 70 to 50 faces; this trend did not continue to 30 faces, though there was weak evidence of benefit relative to the largest. It would appear therefore that a somewhat smaller sized tailored model is of value for evolving composites. Note that the variability in composite quality from the intermediate model was considerably less than the largest: the standard deviation was about 50% less and approaches significance on an F-test, $p = .081$. Thus, the evidence is that the intermediate sized

model not only evolves an overall better composite than the largest, it is also more consistent.

4. General Discussion

EvoFIT allows a witness to produce a likeness by the selection and breeding of complete faces. While the general approach is more appealing than selecting individual facial features, overall performance is by no means ideal. In this paper, we explored whether tailoring a face model to more closely match a witness's description of a target might help to produce a better quality composite. In the first experiment, a fixed face model containing 72 faces for a given age, a 'standard' face model, was compared with a tailored model. The tailored model was better. In the second experiment, different sized tailored models were evaluated. An intermediate sized model built from 50 faces produced the better quality composites than ones built from 30 or 70 faces.

The results from Experiments 1 and 2 clearly indicate value in tailoring a face model to match a target. The evolving process includes a Genetic Algorithm (GA) that searches face space iteratively, identifying a sub-space likely to contain the target face. Witnesses select faces that are bred together to produce another set of solutions. If a model is not constrained, it will tend to generate examples with potentially irrelevant features and so produce less accurate solutions: If a target has a thin appearance, for example, then generating wide faces is ineffectual. As Figure 2 illustrates, building a model from thin faces will preclude wide faces from being generated. The other advantage of this approach concerns the nature of the average face. While the initial faces have random characteristics, they are solutions that radiate from the model's average. With a tailored model, the average should be closer to the desired region of face shape and therefore the initial solutions should similarly be closer.

We have previously evaluated a face model built on the basis of a very detailed description, but this did not appear to work very well. This is perhaps due to a limitation on the number of reference faces available – i.e. 200. Our original approach may have worked better given a larger set of references and thus more candidates from which to select. This is difficult to achieve currently given the appreciable time necessary to locate features in each face and classify them. Using a few distinctive features appears to overcome this limitation.

One surprising result was that our smallest tailored model did not perform the best, despite the likelihood of it containing the most accurate set of reference faces

(the best ranked faces). PCA is a statistical technique that, when applied to face stimuli, captures variations in facial shape and texture. However, to be effective, the model must generalise well to faces that match the description. One possibility then is that the model of 30 faces was too small to effectively generate a sufficient number of faces for that description: a larger model may generalise better. An alternative explanation is that people's descriptions of faces are not consistent: One person's idea of 'thin' may be different to another's. A somewhat more general model, like that built with 50 faces, might give a more consistent match on average.

One possible way to investigate these possibilities might be to evolve faces from standard models of a specific age (like the model used in Experiment 1) – i.e. containing 30, 50 and 70 faces of a given age. In this case, a feature description is not necessary and one would merely explore the generalisation ability of the model. If it is the case that an intermediate size model still comes out the best, this would also suggest that a smaller model might be preferable to larger one for witnesses who are unable to produce a description of a suspect (as is sometimes the case).

There are a number of potential extensions to this project. Arguably the most pressing is to explore the generalisation ability of the intermediate sized tailored model. Both experiments employed six targets, a fairly small number, and so a sensible next step would be a replication with a larger set of targets. Such a study would ideally involve the recruitment of laboratory-witnesses, rather than a single observer (the EvoFIT Operator), with each person constructing a single composite from memory. It would also be interesting to further investigate model size. While 50 faces was the optimum number here, what about 40 or 60? However, as human evaluations are time-consuming, a better approach might be to use computer simulations initially. A third possibility might be to further refine face selection. At present, the same faces are used to build both shape and texture models. Selecting a face on the basis of a wide mouth, which, while perhaps appropriate for its shape, may not be for texture if other properties are inappropriate (e.g. wrong eye or brow colour). More careful face selection is likely to help.

Our work would appear to resonate with other researchers who use face recognition applications premised on PCA. These typically involve a PCA shape model where the reference faces are themselves the targets (e.g. [9,10]). 'Recognition' occurs when the (Euclidean) distance is minimal between the correct reference and the projected probe in shape space. However, the distance metric for similarity tends to be noisy when the number of principal components is large and therefore a smaller space may yield better

results; it is also difficult to know, due to their complex behaviour, which components are the most important. It should be possible to rank the reference faces according to an incoming probe and re-build the space on the fly to compute similarity. Such a procedure would need to be fairly rapid – within a second or two for good performance – but would appear to be conceivable for a manageable face set (e.g. up to 100 candidates) and shape-only matching, the norm.

In summary, the current work sought to improve the underlying model for a face evolving system. It found that a more identifiable composite was produced from a model built with faces that matched on key aspects of a target rather than from a more generic model. The work also found that reducing the model size by about 30% was also valuable, but a further similar sized reduction was less effective. While further research is prudent before suggesting that the police use a version of EvoFIT with a smaller face model, it does suggest that a tailored model of 72 faces is of benefit in situations where a witness is able to report a few distinctive features of a suspect's face.

5. Acknowledgements

The work was funded by a grant from the Engineering and Physical Sciences Research Council, EP/C522893/1. We would like to thank Martin's Newsagent, Winstanley, Wigan, for allowing participant recruitment for Experiment 2.

6. References

- [1] P.J.B. Hancock, V. Bruce, and A.M. Burton, "Recognition of unfamiliar faces", *Trends in Cognitive Sciences*, 4-9, 2000, pp. 330-337.
- [2] P.J. Hancock, A.M. Burton, and V. Bruce, "Face processing: human perception and principal components analysis", *Memory & Cognition*, 24, 1996, pp. 26-40.
- [3] P.J. Phillips, P. Grother, R.J. Michaelis, D.M. Blackburn, E. Tabassi, and J.M. Bone, (). "Face recognition vendor test 2002". *Evaluation report NISTIR6965*, 2003.
- [4] J.W. Shepherd, An interactive computer system for retrieving faces. In Ellis, H.D., Jeeves, M.A., Newcombe, F., & Young, A. (Eds.). *Aspects of face processing*. Dordrecht: Martinus Nijhoff, 1986, pp. 398-409.
- [5] A.K. Jain, S.C. Dass, and K. Nandakumar, , "Soft Biometric Traits for Personal Recognition Systems", *Proceedings of International Conference on Biometric Authentication*, Hong Kong, July (2004).
- [6] J. Annesley, V.L. Leung, A. Colombo, J. Orwell, and S.A. Velastin, "Fusion of multiple features for identity

- estimation”, *IEE Conference on Crime and Security*, London:IET, 2006, pp. 534-539.
- [7] A.M. Burton, R. Jenkins, P.J.B. Hancock, and D. White, “Robust representations for face recognition: the power of averages”, *Cognitive Psychology*, 51(3), 2005, pp. 256-284.
- [8] S.J.D. Prince, J.H. Elder, Y. Hou and Y. Oleviskiy, “Towards face recognition at a distance”, *IEE Conference on Crime and Security*. London:IET, 2006, pp. 570-575
- [9] M. Kirby, and L. Sirovich, “Application of the karhunen-loeve procedure for characterization of human faces”, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1990, 12, pp. 103-108.
- [10] A.J. O’Toole, H. Abdi, K.A. Deffenbacher, and D. Valentin, “Low dimensional representation of faces in high dimensions of the space”, *Journal of the Optical Society of America, A*, 10, 1993, pp. 405-410.
- [11] J.W. Tanaka, and J.A. Sengco, “Features and their configuration in face recognition”, *Memory & Cognition*, 25, 5, 1997, pp. 583-592.
- [12] G.M. Davies, P. van der Willik, and L.J. Morrison, “Facial Composite Production: A Comparison of Mechanical and Computer-Driven Systems”, *Journal of Applied Psychology*, 85, 1, 2000, pp. 119-124.
- [13] C.D. Frowd, P.J.B. Hancock, and D. Carson, “EvoFIT: A holistic, evolutionary facial imaging technique for creating composites”, *ACM Transactions on Applied Psychology (TAP)*, 1, 2004, pp. 1-21.
- [14] C.D. Frowd, D. Carson, H. Ness, J. Richardson, L. Morrison, S. McLanaghan, and P.J.B. Hancock, “A forensically valid comparison of facial composite systems”, *Psychology, Crime & Law*, 11 (1), 2005b, pp. 33-52.
- [15] C.D. Frowd, V. Bruce, A. McIntyre, and P.J.B. Hancock, “The relative importance of external and internal features of facial composites”. To appear in *British Journal of Psychology*, 98 (1), 2007, pp. 61-77
- [16] C.D. Frowd, V. Bruce, D. Ross, A. McIntyre, and P.J.B. Hancock, “An application of caricature: how to improve the recognition of facial composites”, *Visual Cognition*, 15 (8), 2007, pp. 1-31.
- [17] C.D. Frowd, D. Carson, H. Ness, D. McQuiston, J. Richardson, H. Baldwin, and P.J.B. Hancock, “Contemporary Composite Techniques: the impact of a forensically-relevant target delay”. *Legal & Criminological Psychology*, 10, 2005, pp. 63-81.
- [18] C.D. Frowd, V. Bruce, H. Ness, L. Bowie, C. Thomson-Bogner, J. Paterson, A. McIntyre, and P.J.B. Hancock, “Parallel approaches to composite production”. *Ergonomics*, in press, 2007, 50 (4), pp. 562-585.
- [19] C. Tredoux, D.T. Nunez, L. da Costa, and Y. Rosenthal, “Face reconstruction using a configural, eigenface-based composite system”. *Presented at SARMAC III*, Boulder, Colorado, July 10, 1999.
- [20] S.J. Gibson, C.J. Solomon, and A. Pallares-Bejarano, “Synthesis of photographic quality facial composites using evolutionary algorithms”, In R. Harvey and J.A. Bangham, Ed., *Proceedings of the British Machine Vision Conference*, 2003, pp. 221-230.
- [21] P.J.B. Hancock, “Evolving faces from principal components”. *Behavior Research Methods, Instruments and Computers*, 32-2, 2000, pp. 327-333.
- [22] C.D. Frowd, “EvoFIT: A Holistic, Evolutionary Facial Imaging System”. *Unpublished PhD thesis*, University of Stirling, 2002.
- [23] C.D. Frowd, V. Bruce, A. McIntyre, D. Ross, and P.J.B. Hancock, “Adding Holistic Dimensions to a Facial Composite System”, *Proceedings of the Seventh International Conference on Automatic Face and Gesture Recognition*, Los Alamitos: Ca. 2006, pp. 183-188
- [24] C.D. Frowd, V. Bruce, Y. Plenderleith, and P.J.B. Hancock, “Improving target identification using pairs of composite faces constructed by the same person”. *IEE Conference on Crime and Security*, London:IET., 2006, pp. 386-395.
- [25] L. Sirovich, and M. Kirby, “Low-dimensional procedure for the characterization of human faces”, *Journal of the Optical Society of America*, 4, 1987, pp. 519-524.
- [26] N.F. Troje, and T. Vetter, “Representation of human faces”. *Technical report*, Max-Planck-Institut, Tubingen, Germany, 1996.
- [27] P.J.B. Hancock, V. Bruce, and A.M. Burton, “Testing principal component representations for faces”. In J.A. Bullinaria, D.W. Glasspool and G. Houghton (Eds.). *Proceedings of 4th Neural Computation and Psychology Workshop*, London: Springer-Verlag, 1997, pp. 84-97.
- [28] M. Mitchell, *An introduction to genetic algorithms*. MIT, 1996.
- [29] H. Ellis, J. Shepherd, and G.M. Davies, “Identification of familiar and unfamiliar faces from internal and external features: some implications for theories of face recognition”. *Perception*, 8, 1979, pp. 431-439.
- [30] G.M. Davies, J. Shepherd, J. Shepherd, R. Flin, and H. Ellis, “Training skills in police photofit operators”, *Policing*, 2, 1986, pp. 35-46.
- [31] R.E. Geiselman, R.P. Fisher, D.P. MacKinnon, and H.L. Holland, “Eyewitness memory enhancement with the cognitive interview”, *American Journal of Psychology*, 99, 1986, pp. 385-401.