

Predict Your Child: a system to suggest the facial appearance of children

Charlie D. Frowd

University of Central Lancashire, Preston, UK

Email: cfrowd@uclan.ac.uk

Vicki Bruce

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

Email: Head.CHSS@ed.ac.uk

Helen Y. Chang, Yvonne Plenderleith, Alex H. McIntyre, and Peter J.B. Hancock

Psychology Department, Stirling University, Stirling, UK

Email: helento2005@gmail.com, yvonneplenderleith@hotmail.com, ahm5@stir.ac.uk, and pjbh1@stir.ac.uk

Abstract— We have developed a novel software program called ‘Predict Your Child’ that, given photographs of potential parent faces, generates plausible looking children. The parent photographs are imported into a PCA-based model of facial appearance to give a set of face parameters that can be mixed together to produce offspring faces. The program is intended for entertainment and has been used commercially, with customers emailing photographs of parents to be bred. In this paper, we describe the system, developed from an evolutionary facial composite system called EvoFIT, outline some of the problems encountered and present some performance data.

Index Terms— face perception; evolution; breeding; facial composite system

I. INTRODUCTION

How might we predict what our children might look like with a given partner? An old fairground method is to sit the potential parents either side of a half silvered mirror. By adjusting their head position appropriately, their faces can be roughly aligned and each sees a mix of the two faces. However, unless they are already unusually similar, the alignment will not be very good and a blurry set of features will be produced. Advances in computer graphics technology have allowed a more sophisticated approach where photographs of the two faces can be morphed together [1]. Key locations around each feature are marked, allowing an average of the two face shapes to be calculated, and the facial appearances can then be mapped to this average shape and themselves averaged to give a blend of the two faces. This will indeed resemble both parents, but fails to be a plausible predictor of children. Firstly, there is only one mid point, while real children all differ in appearance: even ‘identical’ twins are distinguishable. Secondly, everything is simply averaged, while real children may at least appear to inherit one or more features from each parent, such as a father’s nose and a mother’s eyes. Finally, and crucially,

the average will be of indeterminate sex, while most children are either one sex or the other.

The Predict Your Child system was designed to take a pair of photographs of potential parents and generate a set of plausible-looking children. Its purpose is entertainment and we see no particular scientific application, other than a possible role in the assessment of paternity (described below). In particular, it does not seek to emulate the true genetic and developmental processes that produce someone’s face. It works by simply importing pairs of photographs of ‘parents’ into a mixed gender model of facial appearance, thereby reducing each to a set of parameters, and producing children by taking parameters at random from each parent. As each child inherits a different set of parameters, they all look different, but share some characteristics of the parents. The system can still generate children of ambiguous sex, but this was addressed by shifting the parameters responsible for the appearance of sex in one direction or another. However, to avoid minor artefacts, such as slight facial hair for a female offspring face, we projected all offspring into a single sex face model. Note that the designed process produces children as they might be when grown up to a young adult age: for a variety of technical and practical reasons, our face model does not include faces of true children.

The system is derived from EvoFIT, a computer program designed to help witnesses of crime generate a facial likeness of a perpetrator [2][3]. Traditional facial composite systems allow the independent selection of facial features – such as eyes, nose and mouth – but this is a very difficult task and the likenesses produced are typically very poor [4][5][6]. EvoFIT contains a model constructed using Principal Components Analysis (PCA) of a set of photographs of faces. PCA is a standard statistical technique to store, manipulate and compress data, and has been used widely to build models of facial appearance [7][8][9][10][11]. In brief, it identifies the major modes of variation in facial appearance – animated

examples have been put on line for researchers at www.psychology.stir.ac.uk/staff/phancock/ – and handles shape and image parameters separately, though it is also possible to combine them into a single ‘appearance’ model [8]. Each face is described by a set of values along each Principal Component and novel faces may be produced by generating a random set of component values. In use, EvoFIT initially presents a set of such random faces to a user who selects those that most resemble a target face. The parameter sets for the selected faces are then recombined to give a new set of faces which share some characteristics of the chosen ‘parent’ faces. The process is iterated to gradually move closer to a likeness of a target by the act of repeated selection and breeding. The identification rate of ‘EvoFITs’ is significantly better than that of composites produced from traditional systems [3] and is in use in the UK by Lancashire police. We are aware of two other similar PCA-based composite systems in development [11][12].

In EvoFIT, we already have a system capable of producing plausible-looking offspring of parent faces. Note, however, that the faces are currently male, to reflect the serious crime statistics, and also are modelled in monochrome, as colour does not appear to enhance identification [13][14]. In order to produce a Predict Your Child system, we needed to add female faces to the database and make use of colour information. We also had to solve the problem of how to import new photographs into the model. This is non-trivial, as PCA face models are very sensitive to changes in lighting. This is addressed within EvoFIT by ensuring that the original photographs are taken under controlled conditions. There will be no such control in typical photographs of people submitted to Predict Your Child. In what follows, we describe the PCA-based face model, indicate how we addressed the problem of unconstrained lighting and present two evaluations.

II. SYSTEM DESIGN

A. Face model

The basic model was derived from a set of front-face photographs taken under consistent lighting. We initially used photographs of 150 Caucasian university students, half of each sex, mainly in the age range 17 to 21 years, and none wore spectacles or beards; later, an additional 50 were included, as described in IIB, to provide a better import of photographs of faces taken under different lighting conditions. Subjects were asked to face the camera but this still left some variations in head pose, which we did not seek to eliminate as it would produce useful generality in this application, since imported pictures will inevitably also vary in pose. Images were cropped to 180x240 pixels. A total of 223 landmark points were precisely located on each face around the outline of the face and each facial feature: eyes, nose, mouth, etc. This was achieved by an initial fit using an active shape model [8], as implemented in PsychoMorph [15], but final positioning was done manually. We next carried out a basic face normalisation procedure to

remove irrelevant scaling and positional variation that otherwise would have been encoded in the PCA models. This involved the eight landmarks that identified the inner and outer acanthus of the eyes, the centres of the eyes, the base of the nose and the tip of the chin. Across the image set, landmark files were rotated, linearly rescaled and repositioned – and images morphed to reflect these changes – so that these eight co-ordinates within each landmark file gave a best fit (minimal error) across the entire set.

The normalised sets of landmark points have two roles: firstly, they can be combined to give an average face shape to which each individual face is morphed. Secondly, they can be subjected to a Principal Components Analysis, or PCA, which extracts the principal modes of variation in both pose and face shape. This analysis, along with the main face breeding software, was carried out in the Matlab matrix manipulation language [24]. The PCA provides a set of reference faces, generically called eigenvectors, but eigenshapes in the current application, and a set of shape coefficient values, one for each data item (facial shape co-ordinates). These coefficient values allow the eigenvectors to be recombined in specific proportions for the original data items, in this case facial shapes, to be regenerated. Combining the eigenshapes in novel combinations, however, allows random shapes to be generated. For EvoFIT, this is part of the initial procedure for constructing a composite: the generation of random facial shapes for witnesses to choose. In the current application, the shape coefficients were obtained from photographs of parent faces and those were mixed together to generate novel shapes. PCA applied to the 150 item face set generates 149 eigenshapes.

PCA is often used for image compression since the original dataset, or any interpolated data point therein, can be represented using these 149 floating point numbers. So long as the eigenvectors are available when decoding, the original or interpolated data can be recovered. The PCA also produces eigenvectors that are holistic in nature: they provide a coding for the face as a whole. However, the eigenvectors tend to have quite complex functions. For example, one eigenshape may code the face width, making a face appear long and thin or wide and fat; another, the head pose; a third may exaggerate the width of the forehead while at the same time giving the appearance of a smaller chin. In spite of care with photography, as mentioned above, head pose tends to be a component which accounts for most variation in a shape PCA.

The next phase was a second PCA, this time on the pixel intensity values for the internal features of the face – the region containing the 6,500 pixels encompassing the eyes, brows, nose and mouth. Morphed images were used, those where the outline of the face and the facial features were in register, and therefore the resulting ‘texture’ model provided a coding for these features as well as the overall skin tone. For EvoFIT, a single texture model was constructed, but here, three models, one each for the red, green and blue pixel intensity values (or

channels) in the images. Thus, in addition to coefficients for facial shape, each face within the system is described by three sets of coefficients, one for each of these models. Thus, 149 eigentextures were generated for the red, green and blue pixel intensities. The first 50% of generated coefficients were used for breeding shape and texture, 75 for each, which are sufficient to account for the vast majority of variance in these models [10]; also, as described below, to limit the time taken to compute the best fit of texture in the models. Note that with the controlled lighting used for photography, early texture components are associated with the biggest source of variation in appearance, which is the difference between male and female faces. These, together with equivalent components in the shape model, were used to control the sex of generated faces.



Figure 1. Removing residual masculinity effects when generating a female face. Left image is a randomly generated female face from the original mixed sex model; right, the face after a global best fit of the facial shape and facial texture in the female only model.

In more detail, we first determined which coefficients of shape and texture coded for sex in the model. While this may be achieved manually, by visually inspecting the eigenshapes, a better way is to carry out a (point-biserial or Pearson) correlation across PCA coefficients for the image set and the sex of the face. This produced a strong correlation for one of the facial shape components, and similarly, two of facial texture. This also indicated the sense of the coefficient: whether positive values increased either masculinity or femininity of the face. The general approach was problematic, however, as female faces would sometimes be produced with male characteristics such as light facial stubble (either dark or abnormally pale, a kind of inverse stubble). The situation arises as PCA is concerned with modelling variance, rather than psychological variables, and so may contain parameters that contain complex and not necessary ‘clean’ variations. Since not all men display stubble, it is not simply a male-female difference. The solution used was to build a similar model containing only the female faces into which a female face could be fitted, or projected. Thus, a global best-fit approximation was carried out, a well-established procedure, for example [16]. The Matlab matrix division instruction ‘`mldivide`’ was used to divide the eigenshapes into the shape co-ordinates; this provided a least squares solution in the form of a set of shape coefficients. We then repeated the process for the eigentextures and facial textures, to give the texture coefficients. This is a computationally intensive process for facial texture, involving 6,500 pixels for each channel,

so only 50% of the eigentextures were used, to speed up the process, which provided an acceptable fit; total import time was 10 seconds per face on an Intel Pentium IV PC running at 2.0GHz. An example of the process is presented in Figure 1. Note that this projection (best fit) was carried out for both facial texture and shape. We also generalised the procedure by building a male only face model to allow male faces to be similarly projected.

To generate a face, a facial shape and a facial texture was first produced by combining eigenshapes and eigentextures (respectively) in variable proportions. Hair does not model well within a PCA texture model, since combinations of ‘hair’ eigenvectors tend not to produce clear, unambiguous styles. Instead, hair, along with ears and neck, referred to collectively as the ‘external facial features’, were taken from the original photographs. Then, the facial texture – which encompasses the eyes, brows, nose and mouth – was blended into the external facial features. The final stage was an image morph to align the facial features to those defined by the facial shape.

B. Importing photographs of ‘parent’ faces

The first stage in the face breeding process is to import a pair of parent faces from photographs. This was achieved, as for the model’s faces, by the identification of 250 facial landmarks on each photograph. Each image was then pre-processed using the basic face normalisation procedure described in Section IIA. We developed a user interface to allow a best fit of the landmark data in the shape model and the pixel intensities into the texture model. The best-fit procedure was the same as that used to control sex, a least-squares minimisation of the facial landmarks in the shape model, to produce a set of PCA shape components, and facial texture into texture models, to provide the texture coefficients.

Import into the shape model was successful for faces photographed in front view, and also into the texture model for faces taken from the same photoshoot, but poor under fluorescent lighting, or when lit by flash or from the side, as illustrated in Figure 2. The original texture models specifically coded variations between faces, under specific lighting, but attempts to import a face with different lighting results in the system attempting to match the lighting differences by variations in perceived identity, with poor results. We addressed this issue by including 50 faces in the PCA models that were taken under a range of lighting conditions, as implemented elsewhere [7]. Specifically, an additional 10 people were each filmed under a variety of lighting conditions: under fluorescent lighting, flash light and strong incandescent lighting from the left, right and above. Adding the additional photographs introduced a source of variation where identity was constant but lighting changes. PCA was thus carried out on 200 item face set. When a new face was imported into the system, these new components tended to code for any variations in lighting; by setting them to zero, the face possesses a good approximation of the standard lighting.

It is perhaps worth emphasising that the face import is a global process. In the fit to a texture model, for

example, there is a tendency for smoothing of facial lines, wrinkles, moles, etc. This is because such areas occupy a small number of pixels in the image relative to the total but the current algorithm gives equal weighting to each pixel. The effect is illustrated in Figure 2 and is rather like photography using a soft focus. In this case, the effect is exacerbated by the properties of the image set, which were mainly of young adults posing in a neutral expression. Thus, the model was somewhat invariant to the effect of aging – e.g. bags under the eyes and laughter lines – and produced a best fit for individuals in their twenties. This was appropriate in the current application as our intention was to produce faces within this age range. Similarly, for import into the shape model, as broad smiles were not included into the original face set, imported shapes tended to have a somewhat neutral expression, although minor changes in subjects' relaxed expressions allowed for some variation.




Figure 2. Building lighting invariance in the texture model. Image (a) illustrates a poor import into the texture model without invariance to fluorescent lighting, (b) import in the improved texture model, and (c) the target face. Image (d) was imported from a photograph of a female face with strong side lighting and (e) after correction.

After obtaining the shape and texture coefficients from photographs of both parent faces, the production of an offspring involved combining the parents' coefficients first for facial shape then facial texture. A standard uniform cross-over procedure was used, with each of the offspring's shape coefficients taken randomly from one parent or the other; the same for the three texture models. As described above, these coefficients were used to generate a facial shape (by multiplication of the eigenshapes) and a facial texture (by multiplication by the eigentextures). The resulting texture was blended into a set of external facial features followed by an image morph to the given facial shape. Note that the current implementation does not apply mutation to the face parameters; to do so would reduce similarity between parent and offspring faces, whose apparent familial likeness is already rather hard to detect (see section on System Evaluation).

In summary, a system has been developed to enable photographs of two 'parent' faces to be imported and bred together. Developments have been made to avoid generating female faces with residual male characteristic and to reduce the effects of variations in lighting. In general, the faces produced are of high quality, as can be seen in Figures 3 to 5, although we do note that occasionally faces appear 'slightly strange', apparently due to an over-expression of face parameters during breeding – see General Discussion. Thus, the faces generated from the system are given an additional visual check and, perhaps 10% of the time, are discarded as appearing inappropriate.




Figure 3. Face breeding example given two novel photographs. Images (a) and (b) are parent faces, (c) and (d) are the male and female offspring faces produced. Photograph (e) was taken from one of the four 'real' daughters and illustrates a close likeness to synthetic face (d).

SYSTEM EVALUATION

Natural family likeness can be striking but in general may not be as strong as we think. It turns out that humans are more inclined to see a resemblance where a relationship is already believed to exist. For example, [17] recorded stronger rated similarity between pictures of parents and children when told that the faces were related. In general, however, blind testing has demonstrated that we are able to reliably detect family resemblance [18][19][20], although only about 10-15% above the level of chance.

In spite of human observers being fairly poor at detecting family resemblance, the face breeding system was really only of practical value if the offspring faces produced could be reliably matched to their parents in a blind test. Therefore, a human factors type of evaluation was carried out; this was not a qualitative analysis, but rather an objective test using identity matching of photographs of parents and images of offspring generated by the system. We expected reliable matching (above the level of chance) given the successful nature of EvoFIT – e.g. [3][5][12][25]. However, the success of EvoFIT is generally measured at the end of a long series of face

selection and breeding; here, we are essentially exploring people's ability to detect faces bred from a single breeding attempt. Two carefully designed psychological experiments were administered. In the first, offspring matching was carried out using faces that were generated randomly, to explore the functionality of the basic face breeding software. The second was similar, but based on pairs of photographs of familiar faces that were imported and then bred together.

III. EXPERIMENT 1: RANDOMLY-GENERATED FACES

The first investigation was a test of whether offspring faces generated from within the system could be identified. This did not involve the importing of photographs; instead, pairs of parent faces were randomly generated, one male and one female for each pair, that were bred together to produce a single offspring face, and then testing whether people were able to identify the offspring from among other randomly-generated faces. For this test, the original face model was used, the one containing 150 identities, and not the lighting invariant version, to avoid generating parents and children that could be matched in part by lighting cues.




Figure 4. One of the face breeding examples used to test the system. The female and male parent faces are along the top row, the associated array underneath. All faces were randomly generated, except face 1, which was bred from the parent faces.

A six alternative-forced choice method for the evaluation was employed, a standard format in psychology, e.g. [26][27], with offspring faces accompanied by five alternatives. For simplicity, and to ensure attention was focused on the important inner features of the face, the same hairstyle was used

throughout. The style chosen was fairly short and could be considered to be either male or female so as to limit cues to gender based on hair alone.

A. Stimuli

The stimulus set comprised 12 pairs of male/female parent faces, each with an associated array of six 'child' faces. All faces had the same texture for the hair, ears, and neck, and were generated randomly, except for a single offspring face for each array that was produced by breeding the parent faces. Within the set of six, therefore, one was genuinely related to the two parents, the other five were randomly generated. The design also ensured that sex was split evenly, both within each array and in the number of offspring faces. The 'target' offspring was placed in a random position that occurred equally across the array; otherwise, the order of the array members was random.

For presentation to participants, the faces were printed using a high quality printer in colour at 5.5cm (width) x 7.3cm (high): parent faces were printed on one page of A4, the offspring faces on another page in two rows of three. One of the arrays used is shown in Figure 4.

B. Participants

Thirty-six staff and students at the University of Stirling volunteered. There were 17 males and 19 females with an age range from 20 to 80 years ($M = 30.0$, $SD = 12.9$).

C. Procedure

Participants were tested individually and told that they would be evaluating a new face breeding system by trying to identify the child faces in the presence of the parents. They were presented with each pair of parent faces and, in their own time, selected the most likely offspring face from the associated array. The order of presentation for each breeding example was randomised for each person.

D. Results

Participants achieved a mean score of 22.0% ($SD = 13.7$), with a maximum of 50%. Within-subjects t-tests indicated that performance was significantly above chance, $t(35) = 2.34$, $p = 0.025$, by-subjects, $t(11) = 2.41$, $p = 0.034$, by-items; chance was calculated as $1 / 6$ or 16.7%.

E. Discussion

Participants were presented with pairs of parent faces and were required to select the most likely offspring face from an associated array. All faces used were randomly generated except for one of the faces in the array, which was generated from a random mix of facial shape and texture parameters of the two parents. The data supports the notion that participants could reliably select the correct offspring faces. Performance was 5% above chance, similar to that found with real photographs of faces [18][19][20]. In the next experiment, the full system was tested by breeding pairs of faces imported from photographs.

IV. EXPERIMENT 2: FACE BREEDING FROM PHOTOGRAPHS

In this evaluation, photographs of famous celebrity couples were used as parents. This allowed the system to be used more as designed, with the celebrity faces imported into the system and then bred together. The design was similar to Experiment 1, with performance measured by how well the offspring faces could be identified among other faces. The full lighting-invariant model was used, to allow a good import of the famous faces. In order to prevent matching based on residual lighting or pose cues, different photographs of the parent celebrities were presented to participants for matching.

A. Stimuli

The stimuli were six well-known UK celebrity couples that may be considered related in some way (real or fictitious) plus an associated array of six faces. The celebrities were model Jordan and singer Peter Andre; actors Colin Firth and Renee Zellweger (Darcy and Bridget Jones in the film of her ‘diary’); actor Hugh Grant and model Jemima Khan; footballer Wayne Rooney and partner Colleen McLaughlan; singer Rod Stewart and model Penny Lancaster; and royals Camilla Parker-Bowles and Prince Charles. Two photographs of each celebrity were used, each taken in a largely front face pose and a neutral expression. Photographs of the male celebrities had little or no facial hair and none of the celebrities wore spectacles. One photograph from each celebrity pair was imported into the face model and bred together to produce three male and three female offspring faces (the other celebrity photographs were presented to participants, see below). This time, male offspring were given a short brown hairstyle and female offspring, a shoulder-length blond one. Next, the offspring faces were distributed evenly among the six arrays such that each array was balanced by sex. Therefore, all faces seen in the arrays were the result of breeding, although only one of them was bred from the presented parent photographs. As for Experiment 1, the order of the array members was random except that the ‘target’ offspring was positioned in a different location. The faces were printed also as for Experiment 1, with parent faces on one page and face arrays on the other. The parent faces shown were different to those used for breeding. Figure 5 illustrates offspring bred from the system from one of the celebrity couples.

B. Participants

Six male and 18 female staff and students at Stirling University volunteered. They were aged between 22 and 67 years with a mean age of 38.0 years ($SD = 12.4$).

C. Procedure

The procedure for administering the face arrays was the same as in Experiment 1.

D. Results

Participants correctly identified 34.0% ($SD = 20.0$) of the offspring faces, which is 17.4% above chance. Within-subjects t-test confirmed this, $t(23) = 4.26$, $p < 0.001$, by-subjects, $t(5) = 2.67$, $p = 0.044$, by-items.




Figure 5. Suggested offspring of the system given photographs of the pop singer Rod Stewart and model Penny Lancaster.

E. Discussion

In the current test, photos of celebrity couples were used as parents. Results suggested that people were still able to identify faces bred from the system. Performance was also significantly better than when using the random faces in Experiment 1, $t(60) = 2.77$, $p = 0.007$.

V. GENERAL DISCUSSION

We have described a face breeding system that is able to breed together photographs of potential parent faces. The software used was developed from the EvoFIT composite system by designing a mixed sex, colour face model and an interface for importing photos of parents. Work was necessary to permit good import of faces photographed under a variety of lighting conditions and to remove residual mixed-sex effects. The result of two evaluations suggests that the faces bred from this system are identifiable to third persons.

The two tasks used to evaluate the system presented parent faces and required participants to identify the offspring from among five other faces. In Experiment 1, the basic breeding ability of the system was evaluated using random faces. While the offspring faces were able to be identified, performance was still only about 5% above chance. A likely reason for the low scores is the general blandness of random faces generated from EvoFIT [3], which will tend to look very similar to each other in the arrays and thus can be difficult to discriminate. In Experiment 2, the system was used more as intended, by importing photographs of ‘parent’ faces, and performed much better. This may be partly because the celebrity faces are more distinctive than those generated from within the system but may also reflect some engagement of familiar face processing (which may naturally promote better performance). It is relatively difficult even to match different pictures of the same person in an array task [21]. When celebrities are used, there is the possibility of recognising a likeness directly from the offspring, rather than attempting a match to the parent photographs. As mentioned previously, people in general are only able to correctly detect family

resemblance from photographs about 10-15% above chance [17][18][20][22]. The figure of 17% measured in Experiment 2 would appear to be sensible. Both experiments thus suggest that the breeding system was behaving appropriately and was demonstrably valuable. The work also contributes to the body of knowledge which suggests that human observers are able to detect offspring faces of children, albeit not very well, and in this case from synthesised faces, but nevertheless above the level expected by chance alone.

We make no claim that the Predict Your Child system is mimicking real life breeding. While the random crossover procedure used might be considered broadly similar to real life, and PCA-type systems and humans have been shown to be perceptually similar [23], the processes by which real faces develop are entirely different. Of course, we could compare real children with those synthesised here. While there is no expectation that an actual child would resemble one generated by this system, any more than it might resemble its true siblings, it might be interesting to present true child photographs in amongst the alternatives in the array task, to see whether it or the generated image is chosen more often, although there would be problems in equalising the image quality. Anecdotally, we have received some feedback, after breeding about 50 couples for members of the public, that sometimes a good likeness to an existing child is produced, as Figure 3 illustrates.

One semi-serious application of the system might be assessing paternity, for example of historical figures where photographs are available of all parties concerned. The system would generate multiple potential offspring from the mother and both potential fathers. Naïve viewers would be presented with the true child photograph and one putative child from each pair of parents and asked to judge which is most similar. By repeating this for multiple generated images, it should become apparent whether one couple's children are rated as more similar more often, which would be suggestive evidence that they are the true parents of the actual child.

A potential enhancement could model eye colour. It is known that the colour of the iris emerges from a complex interaction between several gene pairs and may be estimated by an analysis of a person's extended family. The breeding system currently takes no account of this and merely combines face parameters, but it should be possible to create a vector within the face model that describes a progression from blue to brown eyes, for example, and then progresses faces along this dimension to produce children with an appropriate eye colour.

Ultimately, we would like to make the system available via a web interface. Unfortunately, two problems remain. Firstly, we do not know of a sufficiently robust fully automatic mechanism for alignment of the facial landmarks. As a result, an autonomous system would require users to manually check and adjust coordinates as necessary. Secondly, the system occasionally produces an unrealistic-looking face, presumably by an over-expression of face parameters, as mentioned earlier. This would appear to be the result of

multiple parameters within the face models having overlapping functions, such as face width, and all being selected at the same time during the breeding process (resulting in an implausibly wide face). The problem occurs as the Principal Components do not generally map cleanly with psychological perceptions, as already found with face sex, and we are planning to examine whether this may be overcome by applying a limiting function to multiple parameters that code common functions.

In summary, we have described a face breeding system that imports photographs of faces into a model and breeds them together. The result of two short studies suggests that the system is successful in producing a likeness of the parents that other people can detect. While development is required before a fully-automated system is possible, we can report that a version has been trialled, with users emailing photographs to us, and the breeding process being carried out manually.

ACKNOWLEDGEMENT

The work was funded by grants from the Engineering and Physical Sciences Research Council (EPSRC), EP/C522893/1, and Stirling University Research and Enterprise (SURE).

REFERENCES

- [1] P.J. Benson, and D.I. Perrett, "Perception and recognition of photographic quality caricatures: Implications for the recognition of natural images", *European Journal of Cognitive Psychology*, 3, 1991, pp. 105-135.
- [2] P.J.B. Hancock, "Evolving faces from principal components", *Behavior Research Methods, Instruments and Computers*, 32, 2000, pp. 327-333.
- [3] C.D. Frowd, P.J.B. Hancock, and D. Carson, "EvoFIT: A holistic, evolutionary facial imaging technique for creating composites", *Association for Computer Machinery (ACM) Transactions on Applied Psychology (TAP)*, 1, 2004, pp. 1-21.
- [4] G.M. Davies, "Face recognition: issues and theories". In M.M. Gruneberg, P.E. Morris, and R.N. Sykes (Eds.) *Practical aspects of memory*, New York: Academic Press, 1978.
- [5] 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 and Criminological Psychology*, 10, 2005, pp. 63-81.
- [6] C.E. Koehn, and R.P. Fisher, "Constructing facial composites with the Mac-a-Mug Pro system", *Psychology, Crime and Law*, 3, 1997, pp. 215-224.
- [7] V. Blanz, and T. Vetter, "A morphable model for the synthesis of 3D faces", *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, Addison-Wesley: New York, 1999, pp. 187-194.
- [8] T.F. Cootes, K.N. Walker, and C.J. Taylor, "View-Based Active Appearance Models", *Proceedings of the International Conference on Face and Gesture Recognition*, 2000, pp. 227-232.
- [9] I. Craw, and P. Cameron, "Parameterising images for recognition and reconstruction", *Proceedings of the British Machine Vision Conference*, Springer Verlag: London, 1991, pp-367-370.

- [10] 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.
- [11] C. Tredoux, and Y. Rosenthal, "Face reconstruction using a configural, eigenface-based composite system". *Presented at SARMAC III*, Boulder, Colorado, July 10, 1999.
- [12] P. Marks, "New photofit 'evolves' a suspect's face", *NewScientist.com*, 2005, retrieved on 19 March 2005 at <http://www.newscientist.com/article.ns?id=dn7143>.
- [13] R. Kemp, G. Pike, P. White, and A. Musselman, "Perception and recognition of normal and negative faces: the role of shape from shading and pigmentation cues", *Perception*, 25, 1999, pp. 37-52.
- [14] D. Perrett, P.J. Benson, J.K. Hietanen, M.W. Oram, and W.H. Dittrich, "When is a face not a face?" In R. Gregory, J. Harris, P. Heard and D. Rose (Eds.). *The Artful Eye*, Oxford University Press, 1995, pp. 95-124.
- [15] B. Tiddeman, D.M. Burt, and D. Perrett, "Prototyping and Transforming Facial Textures for Perception Research", *IEEE Computers Graphics and Applications*, 21, 2001, pp. 42-50.
- [16] B.W. Hwang, V. Blanz, T. Vetter, and S.W. Lee, "Face Reconstruction from a Small Number of Feature Points", *Proceedings of the 15th International Conference on Pattern Recognition*, 2000, pp. 842-845.
- [17] P. Bressan, and M.F. Dal Martello, "Talis pater, talis filius: perceived resemblance and the belief in genetic relatedness", *Psychological Science*, 13, 2002, pp. 213-218.
- [18] P. Bressan, and M. Grassi, "Parental resemblance in one-year-olds and the Gaussian curve", *Evolution and Human Behavior*, 25, 2004, pp. 133-141.
- [19] R.M. Nesse, A. Silverman, and A. Bortz, "Sex differences in ability to recognize family resemblance", *Ethology and Sociobiology*, 11, 1990, pp. 11-21.
- [20] R.H. Porter, J.M. Cernoch, and R.D. Balogh, "Recognition of neonates by facial-visual characteristics", *Pediatrics*, 74, 1984, pp. 501-504.
- [21] V. Bruce, Z. Henderson, K. Greenwood, P.J.B. Hancock, A.M. Burton, and P. Miller, "Verification of face identities from images captured on video", *Journal of Experimental Psychology: Applied*, 5, 1999, pp. 339-360.
- [22] S. Bredart, and R.M. French, "Do babies resemble their fathers more than their mothers? A failure to replicate Christenfeld and Hill (1995)", *Evolution and Human Behavior*, 20, 1999, pp. 129-135.
- [23] P.J. Hancock, A.M. Burton, and V. Bruce, "Face processing: human perception and principal components analysis", *Memory and Cognition*, 24, 1996, pp. 26-40.
- [24] *Using Matlab*. Mathworks Inc.
- [25] C.D. Frowd, V. Bruce, Y. Plenderleith, & 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.
- [26] G.M. Davies, & D. Christie. Face recall: an examination of some factors limiting composite production accuracy. *Journal of Applied Psychology*, 67, 1982, pp. 103-109.
- [27] V. Bruce, H. Ness, P.J.B. Hancock, C. Newman & J. Rarity. Four heads are better than one. Combining face composites yields improvements in face likeness. *Journal of Applied Psychology*, 87, 2002, pp. 894-902.

Charlie Frowd is based at the University of Central Lancashire. He developed EvoFIT as part of his Ph.D. at the University of Stirling and has been improving it ever since. His research involves improving the quality of

facial composites from composite systems. He is also interested in how we perceive familiar and unfamiliar faces. He has been an investigator on 2 EPSRC grants, one to develop EvoFIT, from which the current paper is produced, and one to develop a version of EvoFIT as an exhibit to the Sensation Science Centre, Dundee, as part of a Partnerships for Public Engagement (PPE) project.

Vicki Bruce is Vice-Principal and Head of the College of Humanities and Social Science at the University of Edinburgh and holds a personal chair in Psychology there. As Head of College she is in charge of academic disciplines across the arts, humanities and social sciences including law, education and management. She is an experimental psychologist who has researched and published extensively in the areas of visual perception and cognition, particularly focussing on human face recognition and face perception. She continues to collaborate with colleagues at Stirling.

Helen Chang completed her PhD at the University of Stirling, where she studied the effects of smiling and motion on perceived facial attractiveness. She was also involved in the commercial application of Predict Your Child. She currently works at the University of Toronto as an institutional researcher in graduate education.

Yvonne Plenderleith is a Clinical Psychology Assistant Practitioner for the Child and Adolescent Mental Health Service, NHS Forth Valley. As an undergraduate student of the University of Stirling, she developed an interest in the effect of Autism on perception and cognition. Since achieving her degree, Yvonne has contributed to university research investigating face recognition and facial composite systems. She hopes to continue both experimental and clinical work.

Alex McIntyre graduated from the University of Stirling with a first class honours degree in Psychology. She is currently employed at Stirling as a Research Assistant and is working on her Ph.D. Her approach is to apply theoretical insights to forensic problems, her thesis concerns factors affecting memory, identification and matching of faces in forensic applications.

Peter Hancock is a Professor in Psychology at the University of Stirling. His PhD in Computing Science looked at the application of Genetic Algorithms to the design of Neural Nets. He moved on to study human face perception and the extent to which it might be explained by principal components analysis of face images. The combination of face PCA and Genetic Algorithms led to the idea behind EvoFIT. He has subsequently researched ways to further improve recall of face information by witnesses. He is also interested in evolutionary approaches to the understanding of human behaviour and in the role of temporal synchrony in the binding problem in early vision.