



Kent Academic Repository

visinmetric ltd (2006) *Generation of Facial Composites*. WO 2006/008485 A1.

Downloaded from

<https://kar.kent.ac.uk/85412/> The University of Kent's Academic Repository KAR

The version of record is available from

<https://doi.org/WO 2006/008485 A1>

This document version

Other

DOI for this version

Licence for this version

UNSPECIFIED

Additional information

Versions of research works

Versions of Record

If this version is the version of record, it is the same as the published version available on the publisher's web site. Cite as the published version.

Author Accepted Manuscripts

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding. Cite as Surname, Initial. (Year) 'Title of article'. To be published in *Title of Journal*, Volume and issue numbers [peer-reviewed accepted version]. Available at: DOI or URL (Accessed: date).

Enquiries

If you have questions about this document contact ResearchSupport@kent.ac.uk. Please include the URL of the record in KAR. If you believe that your, or a third party's rights have been compromised through this document please see our [Take Down policy](https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies) (available from <https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies>).

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
26 January 2006 (26.01.2006)

PCT

(10) International Publication Number
WO 2006/008485 A1

(51) International Patent Classification⁷: **G06T 11/00**

(74) Agent: **ELKINGTON AND FIFE LLP**; Prospect House,
8 Pembroke Road, Sevenoaks, Kent TN13 1XR (GB).

(21) International Application Number:
PCT/GB2005/002780

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(22) International Filing Date: 14 July 2005 (14.07.2005)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
0416032.1 16 July 2004 (16.07.2004) GB

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

(71) Applicant (for all designated States except US): **VISION-METRIC LTD.** [GB/GB]; 1 Oast Cottages, Broom Street, Graveney, Kent ME13 9DN (GB).

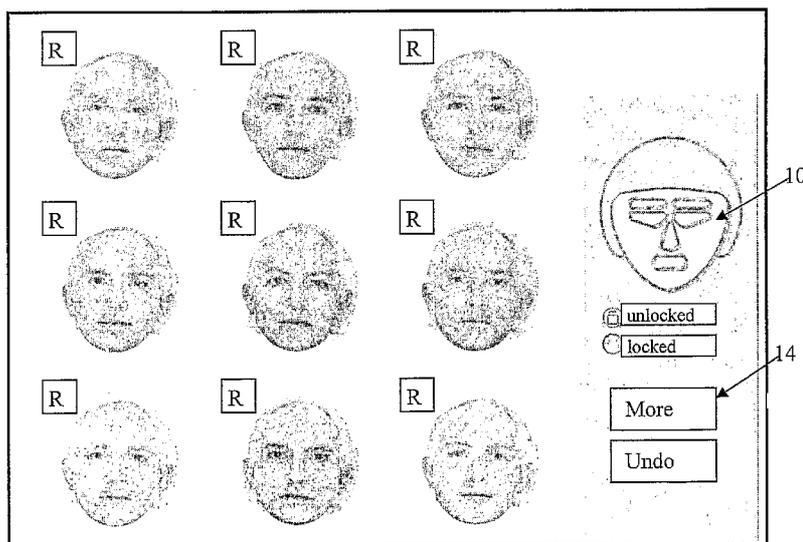
(72) Inventors; and

(75) Inventors/Applicants (for US only): **SOLOMON, Christopher, John** [GB/GB]; University of Kent at Canterbury, Enterprise Unit, Canterbury, Kent CT2 7NZ (GB). **GIBSON, Stuart, James** [GB/GB]; University of Kent at Canterbury, Enterprise Unit, Canterbury, Kent CT2 7NZ (GB). **PALLARES-BEJARANO, Alvaro** [GB/GB]; University of Kent at Canterbury, Enterprise Unit, Canterbury, Kent CT2 7NZ (GB). **MAYLIN, Matthew, Ian, Scott** [GB/GB]; University of Kent at Canterbury, Enterprise Unit, Canterbury, Kent CT2 7NZ (GB).

Published:
— with international search report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: GENERATION OF FACIAL COMPOSITES



(57) Abstract: A system for generating facial composites, comprises a processor for processing facial composite data and a display for displaying images constructed from facial composite data. The processor is adapted to implement an interface in which a plurality of facial images are presented to a user, and in response to user input, including the selection of one best match facial image of the plurality of faces, a further plurality of facial images is presented to the user. A mutation algorithm is used for generating the further plurality of facial images, and which generates facial images corresponding to facial composites which vary from the facial composite of the best match facial image in dependence on a random control parameter. The facial composite comprises a vector having a plurality of terms, and each term is assigned a probability of mutating. The random control parameter determines which vector terms are altered. This provides an efficient representation of multiple choice facial images.

WO 2006/008485 A1

GENERATION OF FACIAL COMPOSITES

Field of the invention

This invention relates to a method and apparatus for the production and display of facial
5 composites.

Background of the invention

The use of facial composites in situations in which a target facial appearance is sought is
known. A target facial appearance will in many cases exist only in the memory or
10 imagination of an operator of the system and thus may not be viewed directly or is not
accessible in concrete form as a photograph or visual image on a computer monitor or any
other visual form.

There are various applications for the generation of such facial composites, and examples
15 are given below:

- 20 i) For purposes of criminal and police investigation in which a user (typically a victim or witness) generates a facial composite designed to match the facial appearance of an individual associated with a crime.
- ii) For purposes of advertising and for the beauty and cosmetics industry in which the goal is to evolve a desired facial appearance i.e. one possessing attractive, interesting or other characteristics.
- 25 iii) To assist in the planning of reconstructive and cosmetic surgery in which the goal is to generate a desired facial appearance which is constrained by surgically achievable outcomes.

Facial composites are widely used in criminal investigations as a means to generate a
30 likeness to a suspected perpetrator of a crime. Current commercial systems for producing composites such as EFit (Trade Mark) and ProFit (Trade Mark) rely on the ability of a witness to recall individual facial features, select the best choice from a large sample of examples and then place them in the appropriate spatial configuration.

There are two major drawbacks to this approach. Firstly, many authors working in the field of psychology as early as the late 70s demonstrated the shortcomings of recall as a means of identification and it has been suggested that the requirement for the witness to recall (as distinct from recognise) the face is the weakest link in the composite process. Secondly, a considerable body of evidence now suggests that the task of face recognition and synthesis does not lend itself to simple decomposition into features, and is partly a global process relying as much on the inherent spatial/textural relations between all the features in the face.

10

It has been demonstrated that a principal components analysis (PCA) on a suitably normalized set of faces produces a highly efficient representation of a human face as a linear superposition of global principal components or "eigenfaces". There has been a significant amount of research into this technique such that the PCA technique is now a standard paradigm in face recognition and both 2-D and 3-D face modelling research.

15

Principal component analysis is a standard statistical technique in which a sample of data vectors are analysed to extract component vectors which successively embody the maximum amount of variation in the sample.

20

No commercial composite systems currently exist which use principal components to achieve facial synthesis/composite production. However, some experimental systems have now been developed, for example by P. Hancock (Hancock, P.J.B. Evolving faces from principal components. Behaviour Research Methods, Instruments and Computers, 32-2, 327-333, 2000), and which have used global principal components as the basic building blocks of composite faces.

25

The first use of evolutionary/genetic algorithms in the production of facial composites was in the FacePrints system developed by Johnston and disclosed in U.S. 5,375,195. In the system proposed by Johnston, an evolutionary procedure is applied in which selection, crossover and mutation are applied to interchange individual feature components and their relative positions within the facial region. Principal components are not used in either global or local form.

30

The combination of an evolutionary algorithm with global principal components as a means to produce composite faces has been described by Hancock (referenced above) and Gibson, Solomon and Pallares (Stuart J. Gibson, Christopher J. Solomon, Alvaro Pallares-Bejarano "Synthesis of photographic quality facial composites using evolutionary algorithms", Proceedings of the British Machine Vision Conference 2003).

It is a standard result of principal components analysis that any data vector in the original training sample may be expressed as a linear combination of the derived principal components, thus the coefficients describing a data vector provide a parametric representation of that data vector. The essence of the method is that a parametric representation of both the shape characteristics and the texture characteristics of a representative population sample is obtained so that the probability density functions of both the shape and texture parameters over the population may be estimated from the sample data - this enables quite new but plausible examples of faces to be generated using standard random number generation techniques. The major difference between the approaches described by Hancock and Gibson et al is that in the former case, facial characteristics of shape and texture are modelled independently using PCA. Gibson et al obtain a more compact representation by calculating a single set of parameters which combines both shape and texture models into a single unified statistical appearance model.

A detailed mathematical treatment of statistical appearance models is given in the paper by T.F.Cootes, G.J. Edwards and C.J.Taylor entitled "Active Appearance Models", IEEE PAMI, Vol.23, No.6, pp.681-685, 2001.

For completeness, the main steps in the generation of the model and the means by which a face is described in parametric form are described below, although this theory will be known to those skilled in the art.

The construction of a statistical appearance model is essentially a three-stage process:

Stage 1: Training – the generation of the facial appearance model

The faces in a training set are first hand marked at a number of control points to form a set of shape model vectors S_i . The Procrustes aligned mean of the shape vectors, \bar{S} is calculated. This is referred to as the "prototype" shape.

5 A principal component analysis PCA is carried out on the ensemble of aligned shape vectors - that is, a linear combination of the shape vectors $P_s = (S - \bar{S})B_s$ is found which satisfies the required orthogonality relationship $P_s^T P_s = \Lambda_s$ where Λ_s is a diagonal matrix and P_s is the matrix containing the principal components. The required diagonalising matrix B_s can be found by standard eigenvector analysis.

10

The corresponding texture map vectors T_s are warped using a linear, piecewise affine transformation to the prototype shape. The resulting texture values are referred to as the "shape-free" texture maps.

15 A PCA is carried out on the shape-free texture maps. That is to say, a diagonalising matrix B_T is found such that $P_T = (T - \bar{T})B_T$ with $P_T^T P_T = \Lambda_T$.

It is important to recognise that the shape and texture in a human face are correlated. In the final stage, separate linear models are combined by decorrelating the shape and texture.

20 A block matrix, B is formed:

$$B = \begin{bmatrix} WB_s \\ B_T \end{bmatrix}$$

where the upper element of the block contains the eigenvectors which diagonalise the shape covariance and the lower element comprises the eigenvectors which diagonalise the texture (shape-normalised) covariance. The matrix W is a diagonal matrix of weights which is required to make the shape and texture parameters, which have different units, commensurate. A further PCA is applied on the columns of B, namely an orthogonal matrix C is obtained such that:

25

$$C = Q^T B$$

where the columns of Q are the eigenvectors and C is the matrix of appearance parameters for the training sample. The key result is that each column of C provides a parametric description of the corresponding face in the training sample which is optimally compact in the linear, least-squares sense.

5

Stage 2: Decomposition of a face into appearance parameters

Decomposition of a given face into its appearance parameters proceeds by the following stages:

10 -The facial landmarks are placed and the Procrustes aligned shape vector S of the face is calculated.

- S is projected onto the shape principal axes P_S to yield the "decoupled" shape parameter vector, b_S .

-The face texture is warped to the prototype or "shape-free" configuration.

15 -The shape-free texture map is projected onto the texture principal axes P_T to yield the decoupled texture appearance parameters.

The appearance parameters are calculated using the eigenvector matrix Q -

$$c = Q^T b = \begin{bmatrix} Q_S^T & Q_T^T \end{bmatrix} \begin{bmatrix} W b_S \\ b_T \end{bmatrix}$$

Stage 3: Synthesis of face from appearance parameters

20 The reconstruction of the separate shape and (shape-free) texture vectors of a sample face from its appearance parameters c is calculated through the linearity of the model according to the equations:

$$S = \bar{S} + P_S W_S^{-1} Q_S c \quad T = \bar{T} + P_T Q_T c \quad (1)$$

25

where \bar{S} and \bar{T} are the mean shape and shape-free textures, P_S and P_T are the shape and texture principal components and Q is the eigenvector matrix, separable into shape and texture block form as $Q = \begin{bmatrix} Q_S \\ Q_T \end{bmatrix}$. The decoupled shape and texture appearance parameters

are given by -

30

$$\mathbf{b}_S = \mathbf{W}_S^{-1} \mathbf{Q}_S \mathbf{c} \qquad \mathbf{b}_T = \mathbf{Q}_T \mathbf{c} \quad (2)$$

Warping the shape-free texture to the required shape completes the facial synthesis.

5 This invention relates to a facial composite system in which composites are generated by an evolutionary procedure, in which the evaluation of an array of visually displayed facial composites by an operator/witness drives the evolutionary process towards a final composite of desired appearance.

10 In particular, the invention relates to a system which combines a statistical appearance model of human facial appearance with an interactive evolutionary/genetic algorithm. These two core elements are established generic techniques which have found application in a variety of scientific and engineering applications. For example, evolutionary and genetic algorithms are widely employed in scientific, engineering and other fields.

15

Despite the existence of the techniques outlined and reference above, there remains a need for an efficient and reliable process for the generation of accurate facial composites and for a user interface which interacts most efficiently with a user to prompt the most accurate and simplest creation of a recognised facial image.

20

Summary of the invention

According to a first aspect of the invention, there is provided a system for generating facial composites, comprising:

a processor for processing facial composite data;

25

a display for displaying images constructed from facial composite data,

wherein the processor is adapted to implement an interface in which a plurality of facial images are presented to a user, and in response to user input, including the selection of one best match facial image of the plurality of faces, a further plurality of facial images is presented to the user,

30

wherein the processor further comprises a mutation algorithm for generating the further plurality of facial images, and which generates facial images corresponding to facial composites which vary from the facial composite of the best match facial image in dependence on a random control parameter, and wherein the facial composite comprises a

vector having a plurality of terms, wherein each term is assigned a probability of mutating, and wherein the random control parameter determines which vector terms are altered.

5 The use of a randomly primed mutation algorithm provides an efficient representation of multiple choice facial images. The basic operational procedure is characterised by extreme simplicity. Namely, a randomly generated face is selected as the preferred member from an array of displayed faces, this face then forms the seed for evolution of the next array of faces, another face is selected as the preferred member which forms the seed for the next array and so on. The selection procedure requires a simple input from the user – a single
10 click of a mouse, a single touch on a touch-sensitive screen, or even a voice prompt are all possible implementations.

The facial composite preferably comprises a vector having a plurality of terms, and each term is assigned a probability of mutating. The random control parameter then determines
15 which vector terms are altered. The new values of the vector terms that are altered can be derived from a random sampling of a normal distribution of values.

Preferably, the probability of mutating can be varied in dependence on the number of further pluralities of facial images that have been generated. In this way, an automatic
20 mutation control algorithm is implemented in which the variation in the array of faces displayed to the user is monitored and controlled by an algorithm which dynamically adjusts the mutation rate. For example, convergence can be provided by decreasing the probability of mutating as the number of further pluralities of facial images that have been generated increases. The mutation rate is thus determined by the step number (the number
25 of arrays of faces displayed to date) but can also be determined by the frequency with which the selected best match has changed over the number of displayed steps to that point. The advantage of this automatic control is that it relieves the user of the need to specify the degree of variation that he/she wants in the array of displayed faces. Of course, sufficient flexibility can be provided so that the individual operator may override this
30 automated control if desired. This step takes away from the operator the need to provide complex input to the system if the system is to be used to rapidly evolve a composite. This approach provides analysis of the system and the specific evolutionary history of a given composite to automatically control the mutation so as to achieve fast convergence.

The probability of mutating can be decreased as a scaled exponential of the number of further pluralities of facial images that have been generated or as a function including the number of further pluralities of facial images that have been generated raised to a negative
5 power. The probability of mutation can be returned to a previous value in response to user input - for example so as to demand increased variation.

The user interface can be adapted to allow the selection of one or more of the facial images which are considered by the user to be the worst match or matches. This provides a multi-
10 touch algorithm which effects a faster convergence. In this case, the operator is required to identify both the worst matching face(s) in the displayed array and the best. Knowing the worst as well as the best gives information on what structures to suppress in the subsequent evolutionary procedure.

For example, a further plurality of facial images can be generated from a best match facial
15 image and a second sub-set of facial composites derived from the worst matching facial images. This is generated from a composite vector comprising the subtraction from the best match facial image composite vector, a composite vector derived as a weighted combination of the worst matching facial images.

20 The system can also allow one or more facial features of the selected best match to be fixed for the next plurality of images.

Each further plurality of images preferably includes the best match facial image from the
25 previous plurality of facial images, so that the best match is always retained.

The system can allow facial images to be produced as a weighted combination of the plurality of images to be selected for subsequent use as a new best match facial image. This provides a face-blending tool which enables new candidate faces to be produced as a
30 weighted combination of the faces currently displayed in the array. The weightings can be selected by means of sliders in a graphical interface. Thus two or more faces which display an appearance deemed in some way similar to the target face by the operator can be combined.

The system can allow facial features of the plurality of facial images to be manually altered in scale and/or position. This provides a means for adjusting the appearance of local features in the face (e.g. nose, mouth, eyes, eyebrows, chin) in a controlled and seamless way through use of a graphical/touch sensitive control. The way in which features can change through this mechanism is constrained to be a priori statistically plausible – i.e. noses can only assume nose-like shapes as determined by training on a population sample. The blending of the altered feature with the rest of the face is seamless because the local feature model alters in such a way as to be consistent with the global facial appearance model. Thus, only statistically plausible facial appearances are permitted.

The system can also allow facial attributes of the plurality of facial images to be manually altered. Sliders (or any other suitable verbal or graphical computer control such as drop-down menus) can be used as a means of controlling semantically labelled attributes in the facial appearance. Examples of such facial attributes are Masculinity/Femininity, Kindness/unkindness, honesty/dishonesty, placidity/aggressiveness (the list is far from exhaustive) and these perceived attributes or properties of human faces can be adjusted according to the subjective impression of the witness.

The first plurality of facial images can be derived from a random selection of facial composites satisfying an initial set of criteria. This provides an initialisation procedure for evolving a satisfactory composite in which examples of specific sub-groups are generated within a general population, immediately narrowing the search and promoting faster convergence. By way of simple illustration, if the witness says that the suspect/perpetrator was a Caucasian female aged 30-40, the system can be initiated by randomly generating examples of faces which are constrained to this class only. This is clearly much better than producing unconstrained examples of faces from different age groups, gender or racial origin. Specifically, this is achieved by building a single appearance model incorporating all sub-groups but using statistical transformation of random number generation techniques to effectively sample faces from a chosen sub-group.

According to a second aspect of the invention, there is provided a system for generating facial composites, comprising:

a processor for processing facial composite data;
a display for displaying images constructed from facial composite data,
wherein the processor is adapted to implement an interface in which a plurality of
facial images are presented to a user, and in response to user input, a further plurality of
5 facial images is presented to the user, wherein the user interface is adapted to allow the
selection of one or more of the facial images which are considered by the user to be the
worst match or matches.

The use of worst match selection (preferably in addition to the selection of one best match)
10 enables a more rapid conversion to the target image.

According to a third aspect of the invention, there is provided a system for generating
facial composites, comprising:

a processor for processing facial composite data;
15 a display for displaying images constructed from facial composite data,
wherein the processor is adapted to implement an interface in which a plurality of
facial images are presented to a user, and in response to user input, a further plurality of
facial images is presented to the user, wherein the user interface is adapted to allow facial
images from the plurality of images to be selected for combination as a weighted
20 combination for use as a new best match facial image.

According to a fourth aspect of the invention, there is provided a system for generating
facial composites, comprising:

a processor for processing facial composite data; and
25 a display for displaying images constructed from facial composite data,
wherein the processor is adapted to implement an interface in which a plurality of
facial images are presented to a user, and in response to user input, a further plurality of
facial images is presented to the user, wherein the user interface is adapted to allow
selection of one best match facial image, and also to allow one or more facial features of a
30 selected facial image or images to be fixed for the next plurality of images.

This feature freezing process enables additional user control, for more rapid convergence.
The user interface is preferably adapted to allow unfixing of a previously fixed feature, and

the processor then reintroduces variation into the previously fixed feature in subsequent pluralities of facial images.

According to a fifth aspect of the invention, there is provided a system for generating facial composites, comprising:

a processor for processing facial composite data;

a display for displaying images constructed from facial composite data,

wherein the processor is adapted to implement an interface in which a plurality of facial images are presented to a user, and in response to user input, including the selection of one best match facial image of the plurality of faces, a further plurality of facial images is presented to the user,

wherein the processor further comprises a mutation algorithm for generating the further plurality of facial images, and which generates facial images corresponding to facial composites which vary from the facial composite of the best match facial image in dependence on a random control parameter, and wherein the interface is adapted to allow the user to alter the degree of variation.

The ability to influence the mutation rate enables the user to backtrack if there is an impression that they have converged too rapidly in the wrong direction. The degree of variation is preferably varied automatically by the processor in dependence on the number of mutations, and the user alteration of the degree of variation overrides the automatic degree of variation. The processor will try to converge by reducing variation progressively, and the user input allows an override function to be implemented.

The system of the invention can be operated on a hand-held computer /device. This changes the paradigm of facial composite generation which has always previously been conducted in the police station some days after the event.

Furthermore, in accordance with a sixth aspect of the invention, there is provided a system for generating facial composites, comprising:

a first processor for processing facial composite data and including a wireless transmission and reception system;

a second processor, implemented as a portable wireless device, including a display for displaying images constructed from facial composite data and comprising a wireless transmission and reception system for communicating with the first processor,

wherein the second processor is adapted to implement an interface in which a plurality of facial images are presented to a user, and in response to user input, a further plurality of facial images is presented to the user, and

wherein the first processor is adapted to process the user input to generate the further plurality of facial images.

The invention also provides methods of generating facial composites, and corresponding to the different aspects of the invention outlined above. The invention also provides the software for implementing these methods.

Brief description of the drawings

Examples of the invention will now be described in detail with reference to the accompanying drawings, in which:

Figure 1 shows the cyclic operation of the system of the invention;

Figures 2 to 4 show examples of user interface provided by the system of the invention; and

Figure 5 shows a system of the invention.

Detailed Description

The system of the invention is an evolutionary facial composite system using a statistical appearance model. This type of system has been described in general terms in Stuart J. Gibson, Christopher J. Solomon, Alvaro Pallares-Bejarano "Synthesis of photographic quality facial composites using evolutionary algorithms", Proceedings of the British Machine Vision Conference 2003). This invention concerns the detailed operation of the system and new aspects of the user interface.

The first step in the implementation of the system is to build a statistical appearance model using a suitable sample of faces in digital image form. These faces constitute a set of training data which, if sufficient in number and judiciously chosen, enable subsequent generation of new, artificial examples of faces which are entirely plausible in appearance.

The precise method for calculating such an appearance model is described in detail in T.F.Cootes, G.J. Edwards and C.J.Taylor entitled "Active Appearance Models", IEEE PAMI, Vol.23, No.6, pp.681-685, 2001, and so is not repeated here.

- 5 As far as this application is concerned, the central result of constructing such an appearance model is that both the original training sample of faces and new examples can be parametrically encoded as a compact vector of numerical parameters which retains all the important shape and textural information in the facial appearance. The general assumption is made that N (typically N~50) such parameters are sufficient to encode the
10 facial appearance of the faces to the required accuracy.

Thus, in general, such a set of numerical parameters is referred to as an appearance vector and denoted by $\vec{c} = [c_1, c_2, \dots, c_N]$.

- 15 Each appearance vector, consisting of N numerical parameters, may be considered to occupy a certain location in an abstract parameter space of N dimensions, the magnitude of the kth component c_k thereby corresponding to the extension along the kth axis of this abstract space. Altering any of the components in an appearance vector thus moves to a different position in this abstract space and alters the facial appearance of the individual.
20 The parameter space effectively defines a very large number of different facial appearances each of which corresponds to a specific point location encoded by the vector $\vec{c} = [c_1, c_2, \dots, c_N]$.

- The appearance vector $\vec{c} = [c_1, c_2, \dots, c_N]$ encodes information about both the shape of the
25 main features in the face and of their colour and intensity characteristics (hereafter referred to as the texture). Related to the appearance vector $\vec{c} = [c_1, c_2, \dots, c_N]$ in an entirely deterministic way are the "shape" vector of the face $\vec{s} = [x_1, x_2, \dots, x_p; y_1, y_2, \dots, y_p]$ which encodes the geometric shape variation of the face and the "texture" vector of the face
 $\vec{T} = [g_1, g_2, \dots, g_M]$ which encodes the textural appearance of the face.

Because of the specific deterministic connection between \bar{c} , \bar{s} and \bar{T} (outlined above), the following causal relations hold:

-Alteration of any of the components of \bar{c} will correspondingly alter both the shape and texture vectors \bar{s} and \bar{T} .

5 -Alteration of the shape vector \bar{s} will alter the appearance vector \bar{c} (but not the texture vector \bar{T})

-Alteration of the texture vector \bar{T} will alter the appearance vector \bar{c} (but not the shape vector \bar{s})

10 In the envisaged applications, a method is required for searching within this parameter space to find a target location or indeed target locations within the space which correspond to a facial appearance which most closely resembles the desired one. The concept of the desired appearance is a general one – in certain applications (such as the attempt to create a composite of a criminal offender by a witness or police officer) the desired appearance will
15 be an accurate and/or recognisable likeness of a living individual. In other applications, the target appearance may simply be one possessing subjective characteristics (e.g. beauty, honesty, masculinity etc) in appropriate combinations as required by the operator.

Thus, in general, there is no direct mathematical means by which the precise coordinates of
20 the target location in the parameter space can be found. This is because the required facial appearance is not available in digital form but only in the memory or imagination of the operator/witness. For this reason, the basic approach is to search for the target location by presenting examples of facial appearance to an operator/witness and using their response to such examples to guide a search within the space. Such an approach is fundamentally
25 based on cognitive processes in which the perception of similarity (or other relevant cognitive concept such as beauty, masculinity etc) between the presented face and the target face forms the basis for the fitness evaluation in the evolutionary search procedure.

The basic iterative procedure is illustrated in Figure 1. The evolutionary process produces
30 appearance vectors (genotypes) whose corresponding facial appearance (the phenotypes) are evaluated by the operator. This evaluation gives rise to an iterative cycle as shown. In this cycle, the user evaluation guides the generation of the new genotypes.

One aspect of the invention relates to a mutation algorithm/procedure used in the evolutionary algorithm, and this will now be described in detail.

- 5 The primary means by which random but controlled variation in facial appearance is produced in is through the action of a dynamically adjustable "*mutation operator*" which is denoted in the following description by M . The mutation operator M is a stochastic operator which acts on each element of an appearance vector with a certain probability of changing its value. Thus, the action of M on the elements of an arbitrary appearance vector
- 10 $\vec{c} = [c_1, c_2, \dots, c_N]$ is that

$$\begin{array}{ll} c_k \rightarrow c'_k & \text{if } r < M_k(t) \\ \text{else} & c_k \rightarrow c_k \end{array}$$

- where $0 < M_k(t) < 1$ is the value of the k th element of M at time t . In other words,
- 15 $M_k(t)$ determines the probability of the k th element in the appearance vector mutating at time t , r is a random number sampled from the canonical density function and c'_k indicates the new mutated value for the k th element of the appearance vector. In general, the values of M_k can be different for all elements from $k=1, 2, \dots, N$, thereby defining different probabilities of mutation for each element in the appearance vector. Note that, in this
- 20 context, time t does not correspond to clock time but is rather a discrete quantity corresponding to the number of generations produced since the start of the evolutionary procedure.

- The statistical distribution of the elements of an appearance vector is known to be
- 25 independent, multivariate normal. In the construction of the statistical appearance model of which the considered appearance vectors are members, a suitable scaling of the principal components can ensure that all elements of the appearance vector are distributed identically as independent normal with similar mean and variance. Where the action of the mutation operator M requires a new value to be produced in the appearance vector (i.e.
- 30 mutation occurs), this is therefore achieved easily by sampling a normal distribution of

zero mean and unit variance (for which standard computational routines exist) and then scaling appropriately. Thus:

$$c'_k = \alpha S\{N(0,1)\}$$

5

where α is a scaling parameter and $S\{N(0,1)\}$ denotes a random sampling of a normal distribution of zero mean and unit variance.

The mutation operator $M(t)$ is dynamic – i.e. changes with generation number t . As a
10 general principle, it is desirable that

$M(t)$ is relatively large when t is small. In the early stages of the evolutionary process, we cannot expect to be very close to the target appearance and presentation of significant variation in facial appearance to the witness/operator is desirable.

15

$M(t)$ should smoothly decrease to small values as t becomes large. The reason for this is that as the evolutionary procedure progresses, the operator will progress closer to the target face and it is desirable to correspondingly reduce the scope of the search to modest variations about the currently favoured appearance.

20

This behaviour for $M(t)$ can be achieved by a variety of pre-determined functional forms. Two such forms are proposed. The first form is a scaled exponential:

$$M(t) = \gamma_{\max} \exp(-t/\beta)$$

25

where γ_{\max} is the maximum mutation rate and β controls the decay rate. Best values for these parameters have been established as approximately 0.8 and 2.5 respectively, though a range from 2-5 for the latter parameter is also effective and the precise numerical value is not critical. Thus the mutation function has a time (generation) dependent decay.

30

A second functional form has been determined by a regression analysis on a large number of studies in which a computer program was made to act like an "ideal witness", rating faces according to their actual distance from the known target in the parametric appearance space. This form is:

5

$$M(t) = \alpha + \beta t^\gamma$$

Effective values for these parameters have been found to be $\alpha = 0.1, \beta = 0.15, \gamma = -0.18$.

10 However, the use of a mutation function $M(t)$ whose value is always strictly determined by the actual time/generation number t is not, in general, sufficiently flexible. In certain instances, the mutation rate may drop to unacceptably low values before a reasonably accurate facial appearance has been achieved. If the mutation function is left unaltered, the search is then, broadly speaking, confined to small variations about a facial appearance
15 which is still far from the desired one, so that there is premature convergence. A mechanism is therefore required which can increase the mutation rate.

A means of avoiding premature convergence is provided by the invention by a method of moving back up the mutation curve (i.e. increasing the mutation rate and thereby
20 increasing the variation in the population) based on operator response. This can be alternatively viewed as moving the mutation operator back to a value it possessed at some earlier point in time. This operates under two conditions as follows.

25 Firstly, if the operator/witness should elect for more variation in the sample of faces, the current stallion is retained and the mutation rate increased prior to producing the next generation.

30 Secondly, if the image selected by the operator/witness over a specified sequence of N generations is the same face, the mutation rate can also be increased. This is because repeated selection of the same face tends to suggest that insufficient variation is being generated. The counter argument, namely that repeated selection of the same face indicates that a good composite has been already been achieved may, in certain instances, be true. However, since in the evolutionary process used, the selected image stallion

member is never lost, increasing the mutation rate for this scenario does not have any significant detrimental effect (the best face is retained no matter what).

In both the cases described, the mutation rate can be increased by restoring the value of the
5 mutation function as it was at two previous time steps:

$$M(t) \rightarrow M(t - 2)$$

Having described the basic principle of the mutation operator, the overall evolutionary
10 procedure will now be described.

Interactive Evolutionary procedure

Precisely because the fitness evaluation is an interactive one in which the human operator must effect the fitness evaluation himself, it is essential for practical purposes that the
15 evaluation process be a simple one that does not make untenable cognitive demands on the operator/witness.

At the same time, and for the same reasons of practicality, the search procedure cannot require excessive numbers of evaluations and a long time to achieve a satisfactory result.
20 Accordingly, the invention provides an interactive evolutionary search procedure which is both cognitively simple and which provides satisfactory results in a short time.

The procedure is based on the repeated generation of an array of M examples of facial appearance (for example an array of 3 x 3 faces so that M=9 although other numbers and
25 combinations are possible). All examples conform to witness/operator selected categories of gender, racial origin and approximate age.

Initially, the user is thus given the opportunity to select gender, age and race, and also hairstyle. This information can be input using a purely graphical interface, where a set of
30 normalised faces are presented to the user, each time representing different races, genders, ages and hair styles.

The system then generates the first batch of randomly mutated variations. The operator/witness can respond to this set of examples in a number of ways as follows:

a) Reject all 9 examples as unsatisfactory and proceed to the next generation.

5

b) Reject specific selected examples from the 9 as being poor likenesses to the required facial appearance and proceed to the next generation.

c) Reject specific selected examples from the 9 as being poor likenesses to the required facial appearance, Accept/select one example as being the best likeness to the required facial appearance and proceed to the next generation.

10

d) Accept/select one example as being the best likeness to the required facial appearance and proceed to the next generation.

15

e) To fix/freeze one or more selected features in the stallion so that this feature is preserved in subsequent generations. (The option also exists to "unfix" features currently in the frozen state).

When the operator responds in one of these ways a) to e), the system will be described as operating in "*EasyFIT*" mode. The operator has a further option, which is to invoke a mode of operation which will be termed the "*ExpertFIT*" mode of operation.

20

Considering only the *EasyFIT* mode, and choices a) to e), the precise basis from which the next evolutionary step proceeds depends on which of these specific choices is made by the operator/witness. A face which has been selected as a best likeness to a required facial appearance is an elite member which is termed a "stallion". The use of this term lies in the fact that the stallion is the encoded facial appearance which effectively seeds the evolutionary process, the next generation of faces being produced by random mutations effected on the stallion member. Each generation of faces is thus produced by application of a dynamically changing mutation operator M on the current stallion, thereby producing variation in facial appearance about that appearance encoded by the stallion. Thus new appearance vectors are created as -

25

30

$$\bar{\mathbf{c}}_{NEW} = M \left\{ \bar{\mathbf{c}}_{ST} \right\}$$

where $\bar{\mathbf{c}}_{ST}$ is the current stallion.

- 5 From the beginning of the evolutionary search procedure, a stallion member only exists from the first time at which either option c) or d) is taken. Thereafter, a stallion always exists as the best likeness selected to date. Thus, until a stallion member has been selected, the procedure is effectively just a random sampling of the search space to find an appropriate starting point from which to attempt convergence towards the target.

10

The specific procedure which is followed upon selection of options a) to e) is now described.

Selection of option a)

- 15 If there is an existing stallion, the stallion is cloned and exists in the new generation whereas the other 8 members in the next generation are randomly generated by application of the mutation operator on the stallion. If there is no existing stallion, the mutation rate is modified and then all 9 members are randomly generated.

20 Selection of option b)

- If there is an existing stallion, this stallion is cloned and exists in the new generation. A first fraction of the remaining members of the new generation is created by applying the mutation operator on the stallion member whilst a second fraction is created from a shifted version of the stallion. Specifically, the shifted stallion is given by subtracting a scalar multiple of the average appearance vector of the rejected faces. Denoting the appearance vector of the stallion by $\bar{\mathbf{x}}_{st}$ we thus apply the mutation operator M on the vector $\bar{\mathbf{x}}'$. Thus

$$\bar{\mathbf{x}}' = \bar{\mathbf{x}}_{st} - \alpha \langle \mathbf{x}_{rej} \rangle$$

and the new generation is produced as

30

$$M \left\{ \bar{\mathbf{x}}' \right\} = M \left\{ \bar{\mathbf{x}}_{st} - \alpha \langle \mathbf{x}_{rej} \rangle \right\}.$$

$\alpha = 0$ thus corresponds to the creation of new members from the stallion itself.

If there is no existing stallion, the new generation is produced in the same way as described for option a) when there is no existing stallion. An effective value for the shift parameter α has been determined as 0.1. In the preferred implementation of this process, two new members are created by mutating on the shifted stallion, two by mutating on the current stallion with an artificially increased mutation rate by a factor of 1.3 and the remainder (5 in the preferred implementation) by mutating on the current stallion at the current mutation rate.

Selection of option c)

The new selected stallion is cloned and exists in the next generation. The remaining members of the next generation are created by applying the mutation operator on a shifted version of this stallion. The shifted stallion is calculated by an identical formula as described in option b) by subtracting a scalar multiple of the average appearance vector of the rejected faces from the stallion. Denoting the appearance vector of the stallion by $\bar{\mathbf{x}}_{st}$ we thus apply the mutation operator M on the vector $\bar{\mathbf{x}}^i$. Thus

$$\bar{\mathbf{x}}^i = \bar{\mathbf{x}}_{st} - \alpha \langle \mathbf{x}_{rej} \rangle$$

and the new generation is produced as

$$M\{\bar{\mathbf{x}}^i\} = M\{\bar{\mathbf{x}}_{st} - \alpha \langle \mathbf{x}_{rej} \rangle\}.$$

Selection of option d)

The new selected stallion is cloned and exists in the next generation. The remaining members of the next generation are created by applying the mutation operator on this stallion. Denoting the appearance vector of the stallion by $\bar{\mathbf{x}}_{st}$ we thus apply the mutation operator M on the vector $\bar{\mathbf{x}}_{st}$.

30

and the new generation is produced as

$$M\{\overline{\mathbf{x}}_{st}\}.$$

Selection of option e)

5 Selection of this option does not exclude options a) to d), rather its selection may be considered as a constraint on the results of taking the other options. Thus, “frozen” features will be propagated throughout the generations until they are unfixed by the operator.

10 A major advantage of operation in the EasyFIT mode is the simplicity of operation and cognitive task. One embodiment of this operational mode is shown in the Figure 2 which shows the user interface.

15 The process of selection in the EasyFIT interface can be effected either by clicking with a mouse on the appropriate area of the screen or by direct contact with a touch sensitive input device. The process of selection is referred to as "touching" the corresponding graphic or area of the screen. Referring to Figure 2, the following functionality is employed:

20 -Touching the reject symbol "R" at the upper left of the corresponding face image will remove that face from view.

-Touching a given face will elect that face as the stallion member and automatically invoke calculation of the next generation of faces.

25 -Touching an unlocked region (colour coded grey for example) of the icon face on the right hand side of the interface will freeze the given feature in future generations. The corresponding feature region is then differently colour coded (for example in blue). The freeze can be released by touching the blue region.

-Touching the pushbutton 14 entitled “Generate More” randomly produces a new generation of faces.

30 Options a) to e) described previously are thus effected by the following:

a) Touch the “Generate More” button 14.

- b) Touch one or more reject icons R and then touch the “Generate More” button 14.
- c) Touch one or more reject icons R and then touch the chosen facial image.
- d) Touch the chosen facial image.
- e) Touch the corresponding region in the face icon 10 (toggle grey-blue to fix feature;
- 5 toggle blue-grey to release feature).

The EasyFIT mode of operation enables changes in both specific features and the overall facial appearance to be effected. However, it does not allow an operator to produce any specific changes by direct intervention. By invoking the ExpertFIT mode of operation, the
10 operator/witness can produce changes in facial appearance using a set of additional manipulation tools. Specifically, the ExpertFIT mode provides three deterministic ways in which the facial appearance can be altered:

- E1 – Individual features (eyes, nose, mouth, eyebrows and face shape) can be altered.
 - 15 E2 – New faces can be produced which are weighted combinations of faces existing in the current generation.
 - E3 – Faces can be altered by adding attribute-based components to enhance or decrease certain attributes (e.g. masculinity, perceived age, honesty etc) of the face.
- 20 Invoking the ExpertFit mode produces three new function tabs in the interface entitled Local Feature (E1), Blend Faces (E2) and Attributes (E3).

Selection of the Local Feature function produces the interface shown in Figure 3. The interface displays the face icon 10, together with movement arrows 20 and scaling arrows
25 22 for operating on a selected part of the facial image.

Manipulation of a feature is achieved by touching the corresponding region in the face icon (thereby making this region active) and scaling and moving the feature using the function buttons as indicated.

30

A preferred implementation of the invention enables features to be controlled with little or no guidance from a third party. Once a facial feature has been locked, it appears highlighted in the schematic image to inform the user that no further shape deformation of

the selected feature will occur during subsequent generations. In terms of the facial composite system, a snap shot of the stallion at certain instances in time is effectively taken and the shape of one or more chosen features are fixed. Subsequent generations nevertheless cause variations in texture and shape changes in the features that remain
5 unlocked.

This concept can be expressed in terms of a vector addition comprising the current stallion St and a snap shot of a previous stallion St_0 captured at time t_0 . The term "time" is used to refer to a particular generation number.

10

$$St' = St [I - Wf] + St_0.Wf \dots (a)$$

15

Where I is the identity matrix and Wf is a diagonal matrix with elements equal to one or zero. Wf can be considered to be a feature extractor as it effectively extracts all of the coordinates of St_0 corresponding to the fixed feature.

The equation above can be extended to include multiple features locked at different times:

$$St' = St [I - Wf_1 - Wf_2 - \dots - Wf_n] + St_1.Wf_1 + St_2.Wf_2 + \dots + St_k.Wf_k \dots (b)$$

20

Wf_1, Wf_2 etc. are the feature selectors for the 1st, 2nd etc features respectively (i.e. nose, mouth... etc). St_1, St_2 and St_k are snap shots of the stallion taken at times t_1, t_2 and t_k . Hence one or more features may be locked at once. If the user wishes to evolve a single feature in isolation, all other features can be locked.

25

When a feature is locked, the new facial composites are still generated using the same evolutionary algorithm, but the selected shape features of the new facial composites are effectively inhibited, i.e. not displayed. Thus, although a given facial feature will appear the same in all 9 images, the underlying facial composite will have different data relating
30 to that facial feature. The reason for this is explained below.

The stallion is defined by a vector of parameters which control both shape and texture (colouring) of the whole face. From this underlying vector of parameters, actual stallion

face shape is derived and stallion face texture. The stallion face texture is unaffected by the freeze feature function, and for this reason the mutation algorithm needs to operate on the full composite data even when a feature is frozen, so that the face texture mutation can take place. For this reason, it is not possible to lock a facial feature by freezing one or
 5 more of the underlying parameters without interfering with the appearance of the face as a whole. Hence, in order to freeze a facial feature, the facial shape is locked (as viewed on screen) rather than the underlying parameters that define the stallion.

When a chosen feature is unlocked, shape variation needs to be re-introduced so that the
 10 feature may evolve as it did prior to locking. This could be achieved by simply reverting to the current stallion such that $S = S_t$ (i.e. acting to reverse inhibition of the display of that feature). However, this would cause an abrupt subsequent change in shape of the unlocked facial feature which is both counter intuitive (as the feature had previously been locked) and visually displeasing. Instead, a continuous shape transition is required by gradually re-
 15 introducing variation into the unlocked feature. To accomplish this smooth transition, a decay function is established, $a(t)$ and equation (a) above becomes:

$$S' = S_t [I - a(t).Wf] + S_{t_0}.a(t).Wf \dots (c)$$

20 A linear decay function can be used, $a(t) = -0.1t$, in the range $0 < t < 10$. This gives an aesthetically pleasing transition between the fixed feature and stallion. An exponential decay may be preferable when a smoother decay is required. For a fixed feature a remains constant ($a = 1$) and once $a(t)$ has decayed to zero it remains at that value until the feature is fixed again.

25

Equation (c) has a useful limiting behaviour. When a feature is unlocked the shape set ebbs away over time and facial shape as seen by the user reverts to the underlying stallion (St).

30 The decay method outlined above allows the facial feature to return towards the stallion facial feature, and also ensures that when features are locked and unlocked, the underlying facial composite remains plausible.

Selection of the Blend Faces function produces the interface shown in Figure 4.

5 A blended face 30 is produced in the bottom right-hand corner, the relative weights assigned to each face in the 3 x 3 array on the left being indicated by the slider controls 32 above it. If the blended face is considered a better face than any of those existing in the 3 x 3 array, the operator touches it (thereby making it the current stallion) and a new generation of faces is produced.

10 Selection of the Attributes function generates an interface showing word descriptions of facial attributes. The specified attribute is selected by touching the corresponding word and its contribution to the face adjusted by means of the slider control. The resulting transformed appearance is then displayed on the bottom right. If the transformed face is considered a better face than any of those existing in the 3 x 3 array, the operator touches it
 15 (thereby making it the current stallion) and a new generation of faces is produced. This interface operates in similar manner to that of Figure 4.

The functionality included in the ExpertFIT mode of operation and denoted by E1, E2 and E3 is achieved through the following techniques.

20

Local Feature manipulation E1

The appearance vector of an arbitrary face in the system, denoted by \vec{c} , directly determines the values of an associated shape vector $\vec{s} = [x_1, x_2, \dots, x_p; y_1, y_2, \dots, y_p]$ which encodes the shape of the face and an associated vector $\vec{x} = [g_1, g_2, \dots, g_M]$ which encodes the textural
 25 appearance of the face.

Local feature manipulation is achieved by considering only those coordinates within the shape vector $\vec{s} = [x_1, x_2, \dots, x_p; y_1, y_2, \dots, y_p]$ which delineate the particular feature under consideration. It operates in two basic steps:

30

(i) Application of the controls in the interface such as moving the feature in the vertical or horizontal directions, scaling the vertical and horizontal coordinates

independently and other more sophisticated operations operate only on the selected coordinates within the shape vector. Thus the manipulation has produced a local shape offset vector $\overline{\mathbf{s}}_{local} = [0, 0, \dots, \Delta_{x1}, \dots, \Delta_{xL}, \Delta_{y1}, \dots, \Delta_{yL}, \dots, 0, 0]$ between the global shape vector $\overline{\mathbf{s}}$ and the desired, manipulated shape.

5

(ii) For purposes of display, a copy of the shape vector $\overline{\mathbf{s}}$ is modified such that the current values for the selected feature are replaced by the modified feature coordinates (i.e. $\overline{\mathbf{s}}' = \overline{\mathbf{s}} + \overline{\mathbf{s}}_{local}$) and the current texture vector $\overline{\mathbf{x}} = [g_1, g_2, \dots, g_M]$ then warped through a standard piecewise affine transformation to the new shape $\overline{\mathbf{s}}'$. The resulting image is then displayed.

10

At subsequent stages in the evolutionary procedure, $\overline{\mathbf{s}}_{local}$ is always considered and added to the global, evolving shape vector $\overline{\mathbf{s}}$ so that the texture vector is warped onto the desired shape as determined by the operator.

15

Face Blending E2

Face blending is achieved simply by a linear, weighted combination of the appearance vectors of the chosen faces. Thus the blended face is given as:

$$\overline{\mathbf{c}}_{blend} = \frac{\sum_{i=1}^L w_i \overline{\mathbf{c}}_i}{\sum_{i=1}^L w_i}$$

20

where w_i is the weight given to the i th face in the array using the slider or other suitable control in the interface and $\overline{\mathbf{c}}_i$ is corresponding appearance vector. The reconstruction of the actual blended image from the appearance vector is a standard procedure.

25

Attribute manipulation E3

A general statistical approach can be taken to attribute manipulation in which arbitrary facial attributes of an individual face can be accurately manipulated by moving along predetermined directions in the abstract vector space defined by the appearance model

parameters. The basis of the method of the invention is to combine a set of facial attribute ratings obtained from a group of observers on a sample of faces which form the training data (or part thereof) for the statistical appearance model.

- 5 The transformed face \bar{c}' , in which a given attribute is decreased or enhanced from the original \bar{c} , is calculated by adding a scalar multiple of the given attribute vector \bar{v} to the original appearance vector of the face:

$$\bar{c}' = \bar{c} + \beta \bar{v}$$

10

Details of the calculation of the attribute vectors will now be described.

An attribute vector defines a direction in the complete appearance space along which a specified attribute exhibits maximum variation over the sample population. The gradual transformation of the facial appearance is then achieved by adding some scalar multiple of the given attribute vector to the veridical appearance vector of the subject. To identify an attribute vector, consider that we have a sample of M faces in our appearance model training sample. After decomposition, each of these may be described by a vector consisting of N appearance parameters. Thus, the jth face in a sample with shape vector S_j and texture vector T_j is represented in parametric form as the (mean-subtracted) appearance vector $c_j = [c_j(1) c_j(2) \dots c_j(N)]$. The appearance vectors can be written as the columns of the $N \times M$ appearance matrix C:

15

20

$$C = \begin{bmatrix} \uparrow & \uparrow & \dots & \uparrow \\ \mathbf{c}_1 & \mathbf{c}_2 & \dots & \mathbf{c}_M \\ \downarrow & \downarrow & \dots & \downarrow \end{bmatrix}$$

Supposing that a group of P observers each numerically rates the M faces for the degree to which the given face exhibits a single chosen attribute. Considering the scale used for the ratings to be arbitrary, and that the kth observer assigns an attribute score of w_{kj} to the jth face, c_j . Denoting all M scores of the kth observer by the (mean-subtracted) score vector w_k and writing all such vectors as the columns of the $M \times P$ rating matrix W:

25

$$W = \begin{bmatrix} \uparrow & \uparrow & \dots & \uparrow \\ w_1 & w_2 & \dots & w_p \\ \downarrow & \downarrow & \dots & \downarrow \end{bmatrix}$$

The attribute matrix is formed as the product $D=CW$. Expressing this in tableau form reveals the essential structure -

$$5 \quad D = CW = \begin{bmatrix} \uparrow & \uparrow & \dots & \uparrow \\ Cw_1 & Cw_2 & \dots & Cw_p \\ \downarrow & \downarrow & \dots & \downarrow \end{bmatrix} = \begin{bmatrix} \uparrow & \uparrow & \dots & \uparrow \\ d_1 & d_2 & \dots & d_p \\ \downarrow & \downarrow & \dots & \downarrow \end{bmatrix} \quad (6)$$

The columns of matrix D, are thus given by weighted combinations of the appearance vectors, the weights corresponding to the attribute scores assigned by the observers. Specifically, the columns of matrix D are $\{d_j\}$ where:

$$d_j = \sum_{k=1}^M w_{kj} c_k \quad (7)$$

10 It can be seen that each column of D is a weighted average of the appearance vectors effectively defining each individual observer's estimate of the mean vector in appearance space along which the attribute in question varies. Note that for an entirely objective attribute (such as actual age in years, actual gender or some other attribute in which all observers agree exactly on their "scores") then $w_k = w$ for all k, the columns of matrix D
 15 are identical, and there is a rank 1 matrix and there is a single direction which accounts for the sample variance and in which the attribute changes, given by Aw . Considering the other extreme, namely a quite "imaginary" facial attribute on which the observers do not show any consistent relationship in their scores, the ratings matrix W would then be characterised by having completely uncorrelated columns. The covariance matrix

20 $\frac{1}{N-1} D^T D$ would then tend (in the limit of a large number of faces) to diagonal form

indicating that no particular direction in appearance space can be associated with the attribute in question. In general, we assume that there will be some but not unanimous agreement amongst the observers. In this case, the $\{w_k\}$ will differ and so, therefore, will the $\{d_j\}$.

One approach to the task is to find linear combinations of the basis vectors in D which are orthogonal and which successively account for the directions in appearance space in which the given attribute exhibits most variance. This is easily accomplished by a discrete PCA or Karhunen-Loeve analysis. In matrix form:

5

$$P=DU \quad (8)$$

Where the columns of P are the new desired basis vectors (the attribute axes) and U is the required matrix which combines the columns of D to produce P . Enforcing the
10 orthogonality condition on the attribute axes, U must be found such that:

$$P^T P = U^T \left\{ \frac{D^T D}{M-1} \right\} U = U^T C U = \Lambda \quad (9)$$

where Λ is a diagonal matrix and C is the covariance matrix of the data. Thus the problem
15 reduces to an eigenvector/eigenvalue decomposition of the covariance which yields the diagonalising matrix U and the eigenvalues Λ through standard numerical procedures.

This simple linear approach finds both the dominant directions in appearance space associated with the given attribute and the amount of variance (through the eigenvalues)
20 that is associated with each of them. The simplest scenario is thus to take the first principal component as defining the attribute vector, effectively defining a “single control” for altering the attribute. Naturally, this will be more or less satisfactory depending on the fraction of the total variance over the attribute which is associated with it and more subjective attributes may exhibit substantial variations along two or more orthogonal
25 directions. The eigenvalues in the matrix Λ conveniently describe the degree of objectivity or subjectivity of the given attribute as they specify directly the level of agreement amongst the sample of observers.

Hairstyles can be selected independently, and these can be blended with the given facial
30 appearance so as to ensure correct position, scaling and colour appearance at the seams between the hair and face. This can be implemented through the use of multi-resolution spline techniques.

The system of the invention can be implemented as a hand held portable device, as shown in Figure 5, which shows the device 50 with an output touch sensitive screen 52. The touch sensitive screen can allow the interface to be implemented fully with touch input, mainly involving the selection between presented images.

The hand held device may comprise a wireless display and with only minimum processing power implemented in the device, for example sufficient to render received data as a graphical output. A separate processing device 60 can be provided in a different location, with a wireless link between the two. For example, a processing unit 60 may be located in the boot of a police vehicle, and the police witness interviewer can then simply carry the portable device 50. This allows the processing power required of the portable device to be kept to a minimum, and allows any sensitive information, such as stored images, to be stored more securely immediately.

The processing of facial composite data within the processing unit 60 is computationally intensive. Reconstruction of a given facial composite requires a shape-normalised texture map to be warped to its associated shape vector of landmark coordinates. This warping process is accomplished through a linear piecewise affine transformation based on a Delaunay triangulation of the facial region of interest.

The piecewise affine transformation is a standard procedure described in many texts. The basic principle is to define corresponding triangles between an *input* image (in this case, the shape-normalised face) and a target or *output* image (the face with its corresponding shape) where three corresponding landmarks define the matching triangles. The texture values lying within a given triangle in the input image are then mapped to their corresponding locations in the output image. Repeating this over all such corresponding triangles produces the final composite appearance.

When the calculation described above is carried out on the central processing unit (CPU) this task can take a substantial amount of the overall CPU time and constitutes the major part of the computational overhead in the production of each new generation of faces in the

system. The amount of CPU time required also scales approximately linearly with the image size in pixels.

To circumvent this problem, the reconstruction of the composite face for visual display can be divided into two distinct steps. The first step, namely the reconstruction of the shape-normalised texture map and the associated shape vector of landmark coordinates from the appearance vector of the facial composite proceeds by execution of code on the main CPU of the processing device.

The second step, namely the transformation of the shape-normalised texture to the true shape vector of coordinates, proceeds by execution of code on dedicated hardware (typically the graphics processing unit) of the computer. The affine transformation of textures from an input image to a target image is effected much more efficiently on the CPU of dedicated hardware such as typical graphics processing units. This division of the two distinct steps into tasks executed on two distinct processing units in this way drastically reduces the overall processing time required to produce a new generation of faces. The overall effect of this procedure is to make the system response to user input considerably quicker, enhancing its overall usability and effectiveness. This division of processing tasks can provide an overall speed increase of between 400% and 1200% compared to the implementation on the main CPU of the computer alone.

There are a number of different aspects to the system described. Although the system described has all of the features of the invention, relating to the different aspects of the invention, the various features can be used in any combination. The most basic system design would use an interface by which a single image is selected as a best match from a plurality of images, and in response to this a further plurality of images is generated using a randomly seeded algorithm. The additional features such as selecting worst match, freezing features, altering and moving individual features, attribute selection, mutation rate flexibility could be applied in any combination. The invention is intended to cover any such combinations of features.

Various modifications to the detailed implementation of the invention described above will be apparent to those skilled in the art.

CLAIMS

1. A system for generating facial composites, comprising:
a processor for processing facial composite data;
5 a display for displaying images constructed from facial composite data,
wherein the processor is adapted to implement an interface in which a plurality of facial images are presented to a user, and in response to user input, including the selection of one best match facial image of the plurality of faces, a further plurality of facial images is presented to the user,
10 wherein the processor further comprises a mutation algorithm for generating the further plurality of facial images, and which generates facial images corresponding to facial composites which vary from the facial composite of the best match facial image in dependence on a random control parameter, and wherein the facial composite comprises a vector having a plurality of terms, wherein each term is assigned a probability of mutating,
15 and wherein the random control parameter determines which vector terms are altered.
2. A system as claimed in claim 1, wherein the vector terms that are altered are altered to a random sample of a normal distribution of vector terms.
- 20 3. A system as claimed in claim 1 or 2, wherein the probability of mutating is varied in dependence on the number of further pluralities of facial images that have been generated.
4. A system as claimed in claim 3, wherein the probability of mutating is decreased as
25 the number of further pluralities of facial images that have been generated increases.
5. A system as claimed in claim 4, wherein the probability of mutating is decreased as a scaled exponential of the number of further pluralities of facial images that have been generated.
- 30 6. A system as claimed in claim 4, wherein the probability of mutating is decreased as a function including the number of further pluralities of facial images that have been generated raised to a negative power.

7. A system as claimed in any one of claims 3 to 6, wherein the probability of mutating is returned to a previous value in response to user input.

5 8. A system as claimed in claim 7, wherein the probability of mutating is returned to a previous value in response to a user request for more image diversity and/or repeated selection of the same facial image.

9. A system as claimed in any preceding claim, wherein the user interface is adapted
10 to allow the selection of one or more of the facial images which are considered by the user to be the worst match or matches.

10. A system as claimed in claim 9, wherein a first sub-set of facial composites for the further plurality of facial images are generated from a best match facial image and a
15 second sub-set of facial composites for the further plurality of facial images are generated from a composite vector comprising the subtraction from a best match facial image composite vector a composite vector derived from the worst match or matches facial images.

20 11. A system as claimed in any preceding claim, wherein the user interface is adapted to allow one or more facial features of the selected best match to be fixed for the next plurality of images.

12. A system as claimed in any preceding claim, wherein the user interface is adapted
25 initially to allow no best match to be selected, in response to which a further plurality of facial images are presented to the user which are randomly generated.

13. A system as claimed in any preceding claim, wherein each further plurality of images includes the best match facial image from the previous plurality of facial images.

30

14. A system as claimed in any preceding claim, wherein the user interface is adapted to allow facial images from the plurality of images to be selected for combination as a weighted combination for use as a new best match facial image.

15. A system as claimed in any preceding claim, wherein the user interface is adapted to allow facial features of the plurality of facial images to be manually altered in scale and/or position.

5

16. A system as claimed in claim 15, wherein each facial composite comprises a vector which encodes a shape vector and a textural vector, and wherein the facial feature alterations are implemented by altering the shape vector.

10

17. A system as claimed in any preceding claim, wherein the user interface is adapted to allow facial attributes of the plurality of facial images to be manually altered.

15

18. A system as claimed in claim 17, wherein the attributes comprise masculinity/femininity and/or kindness/unkindness and/or honesty/dishonesty and/or placidity/aggressiveness.

20

19. A system as claimed in any preceding claim, wherein the first plurality of facial images is derived from a random selection of facial composites satisfying an initial set of criteria.

20. A system as claimed in claim 19, wherein the initial set of criteria relate to age and/or racial origin and/or gender.

25

21. A system for generating facial composites, comprising:

a processor for processing facial composite data;

a display for displaying images constructed from facial composite data,

wherein the processor is adapted to implement an interface in which a plurality of facial images are presented to a user, and in response to user input, a further plurality of facial images is presented to the user, wherein the user interface is adapted to allow the selection of one or more of the facial images which are considered by the user to be the worst match or matches.

30

22. A system as claimed in claim 21, wherein the user interface is further adapted to allow a user input comprising the selection of one best match facial image of the plurality of faces

5 23. A system as claimed in claim 22, wherein a first sub-set of facial composites for the further plurality of facial images are generated from a best match facial image and a second sub-set of facial composites for the further plurality of facial images are generated from a composite vector comprising the subtraction from a best match facial image composite vector a composite vector derived from the worst match or matches facial
10 images.

24. A system for generating facial composites, comprising:
a processor for processing facial composite data;
a display for displaying images constructed from facial composite data,
15 wherein the processor is adapted to implement an interface in which a plurality of facial images are presented to a user, and in response to user input, a further plurality of facial images is presented to the user, wherein the user interface is adapted to allow facial images from the plurality of images to be selected for combination as a weighted combination for use as a new best match facial image.

20

25. A system for generating facial composites, comprising:
a processor for processing facial composite data; and
a display for displaying images constructed from facial composite data,
wherein the processor is adapted to implement an interface in which a plurality of
25 facial images are presented to a user, and in response to user input, a further plurality of facial images is presented to the user, wherein the user interface is adapted to allow selection of one best match facial image, and also to allow one or more facial features of a selected facial image or images to be fixed for the next plurality of images.

30 26. A system as claimed in claim 25, wherein the user interface is adapted to allow unfixing of a previously fixed feature, and wherein the processor then reintroduces variation into the previously fixed feature in subsequent pluralities of facial images.

27. A system as claimed in claim 26, the processor then reintroduces variation into the previously fixed feature in subsequent pluralities of facial images using a decay function.

28. A system for generating facial composites, comprising:

5 a processor for processing facial composite data;

a display for displaying images constructed from facial composite data,

wherein the processor is adapted to implement an interface in which a plurality of facial images are presented to a user, and in response to user input, including the selection of one best match facial image of the plurality of faces, a further plurality of facial images
10 is presented to the user,

wherein the processor further comprises a mutation algorithm for generating the further plurality of facial images, and which generates facial images corresponding to facial composites which vary from the facial composite of the best match facial image in dependence on a random control parameter, and wherein the interface is adapted to allow
15 the user to alter the degree of variation.

29. A system as claimed in claim 28, wherein the degree of variation is varied automatically by the processor in dependence on the number of mutations, and wherein the user alteration of the degree of variation overrides the automatic degree of variation.

30. A system as claimed in any one of claims 1 to 29 implemented as a hand held portable device.

31. A system for generating facial composites, comprising:

25 a first processor for processing facial composite data and including a wireless transmission and reception system;

a second processor, implemented as a portable wireless device, including a display for displaying images constructed from facial composite data and comprising a wireless transmission and reception system for communicating with the first processor,

30 wherein the second processor is adapted to implement an interface in which a plurality of facial images are presented to a user, and in response to user input, a further plurality of facial images is presented to the user, and

wherein the first processor is adapted to process the user input to generate the further plurality of facial images.

32. A method of generating facial composites each comprising a vector having a
5 plurality of terms, the method comprising:

assigning a probability of mutating to each term;

presenting a plurality of facial images to a user;

receiving user input including at least the selection of one best match facial image
of the plurality of faces;

10 generating a further plurality of facial images using a mutation algorithm which
generates facial images corresponding to facial composites which vary from the facial
composite of the best match facial image in dependence on a random control parameter,
wherein the vector terms to be altered are determined using the random control parameter;
and

15 presenting the further plurality of facial images to the user.

33. A method as claimed in claim 32, wherein altering a vector term comprises
selecting a random sample of a normal distribution of vector terms.

20 34. A method as claimed in claim 32 or 33, wherein assigning a probability comprises
assigning a probability in dependence on the number of further pluralities of facial images
that have been generated.

25 35. A method as claimed in claim 34, wherein the probability of mutating is decreased
as the number of further pluralities of facial images that have been generated increases.

36. A method as claimed in claim 35, wherein the probability of mutating is decreased
as a scaled exponential of the number of further pluralities of facial images that have been
generated.

30

37. A method as claimed in claim 35, wherein the probability of mutating is decreased
as a function including the number of further pluralities of facial images that have been
generated raised to a negative power.

38. A method as claimed in any one of claims 34 to 37, further comprising returning the probability of mutating to a previous value in response to user input.

5 39. A method as claimed in claim 38, wherein the user input comprises a user request for more image diversity and/or detection of repeated selection of the same facial image.

40. A method as claimed in any one of claims 31 to 39, wherein receiving user input comprises receiving the selection of one or more of the facial images which are considered
10 by the user to be the worst match or matches.

41. A method as claimed in claim 40, further comprising:

generating a first sub-set of facial composites for the further plurality of facial images from a best match facial image; and

15 generating a second sub-set of facial composites for the further plurality of facial images from a composite vector comprising the subtraction from a best match facial image composite vector a composite vector derived from the worst match or matches facial images.

20 42. A method as claimed in any one of claims 31 to 41, wherein receiving user input comprises receiving a selection of one or more facial features of the selected best match to be fixed for the next plurality of images.

25 43. A method as claimed in any one of claims 31 to 42, wherein, in response the first plurality of images, a user input providing no best match can be received, and wherein the method comprises generating a further plurality of facial images which are randomly generated.

30 44. A method as claimed in any one of claims 31 to 43, wherein each further plurality of images includes the best match facial image from the previous plurality of facial images.

45. A method as claimed in any one of claims 31 to 44, wherein receiving user input comprises receiving a selection of facial images from the plurality of images for combination as a weighted combination.

5 46. A method as claimed in any one of claims 31 to 45, wherein receiving user input comprises receiving facial features of the plurality of facial images to be manually altered in scale and/or position.

47. A method as claimed in claim 46, wherein each facial composite comprises a vector
10 which encodes a shape vector and a textural vector, and wherein the method further comprises implementing facial feature alterations by altering the shape vector.

48 A method as claimed in any one of claims 31 to 47, wherein receiving user input
15 comprises receiving a selection facial attributes of the plurality of facial images to be manually altered.

49. A method as claimed in claim 48, wherein the attributes comprise
20 masculinity/femininity and/or kindness/unkindness and/or honesty/dishonesty and/or placidity/aggressiveness.

50. A method as claimed in any one of claims 31 to 49, further comprising:
generating a first plurality of facial images, derived from a random selection of
facial composites satisfying an initial set of criteria.

25 51. A method as claimed in claim 50, wherein the initial set of criteria relate to age and/or racial origin and/or gender.

52. A method of generating facial composites, comprising:
displaying a plurality of images constructed from facial composite data;
30 receiving user input comprising at least the selection of one or more of the facial images which are considered by the user to be the worst match or matches; and
generating a further plurality of images taking into account the user input.

53. A method as claimed in claim 52, wherein generating a further plurality of facial images comprises using a randomly-seeded mutation algorithm.

54. A method as claimed in claim 53, wherein receiving user input comprises receiving
5 the selection of one best match facial image of the plurality of faces.

55. A method as claimed in claim 54, wherein generating a further plurality of facial images comprises generating a first sub-set of facial composites from a best match facial image and generating a second sub-set of facial composites from a composite vector
10 comprising the subtraction from a best match facial image composite vector a composite vector derived from the worst match or matches facial images.

56. A method for generating facial composites, comprising:
displaying a plurality of images constructed from facial composite data;
15 receiving user input comprising at least the selection of a plurality of facial images from the plurality of images;
generating a weighted combination of the selected plurality of facial images for use as a new best match facial image; and
generating a further plurality of images based on the new best match facial image.
20

57. A method for generating facial composites, comprising:
displaying a plurality of images constructed from facial composite data;
receiving user input comprising the selection of one best match facial image, and also the selection of one or more facial features of a selected facial image or images;
25 generating a further plurality of images based on the best match facial image and also in which the appearance of the one or more selected facial features are substantially fixed.

58. A method as claimed in claim 57, further comprising receiving user input to unfix
30 of a previously fixed feature, and reintroducing variation into the previously fixed feature in subsequent pluralities of facial images.

59. A method as claimed in claim 58, wherein the processor reintroduces variation into the previously fixed feature in subsequent pluralities of facial images using a decay function.

5 60. A method for generating facial composites, comprising:
displaying a plurality of images constructed from facial composite data;
receiving user input comprising at least the selection of one best match facial image; and

generating a further plurality of images based on the best match facial image using
10 a mutation algorithm which generates facial images corresponding to facial composites which vary from the facial composite of the best match facial image in dependence on a random control parameter, and wherein receiving user input comprises receiving a request for increased variation, in response to which the degree of variation is altered.

15 61. A method as claimed in claim 60, wherein the degree of variation is varied automatically by the processor in dependence on the number of mutations, and wherein the user alteration of the degree of variation overrides the automatic degree of variation.

62. A method for generating facial composites, comprising:

20 in a first processor, processing facial composite data to generate a first plurality of facial composites for presentation to a user;

using a wireless transmission and reception system to send the facial composite data to a second processor, implemented as a portable wireless device having a display for displaying images constructed from facial composite data; and

25 on the second processor:

constructing facial images from the composite data;

presenting the plurality of facial images to a user;

receiving user input in response to presented images; and

30 transmitting the user input to the first processor, for subsequent processing to generate a further plurality of facial composites.

63. A computer program comprising computer program code means adapted to perform all the steps of any one of claims 32 to 62 when said program is run on a computer.

64. A computer program as claimed in claim 63 embodied on a computer readable medium.

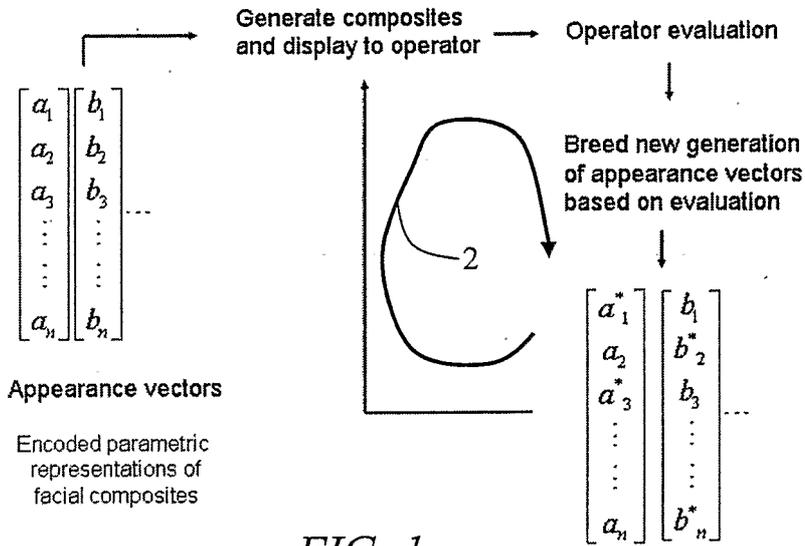


FIG. 1

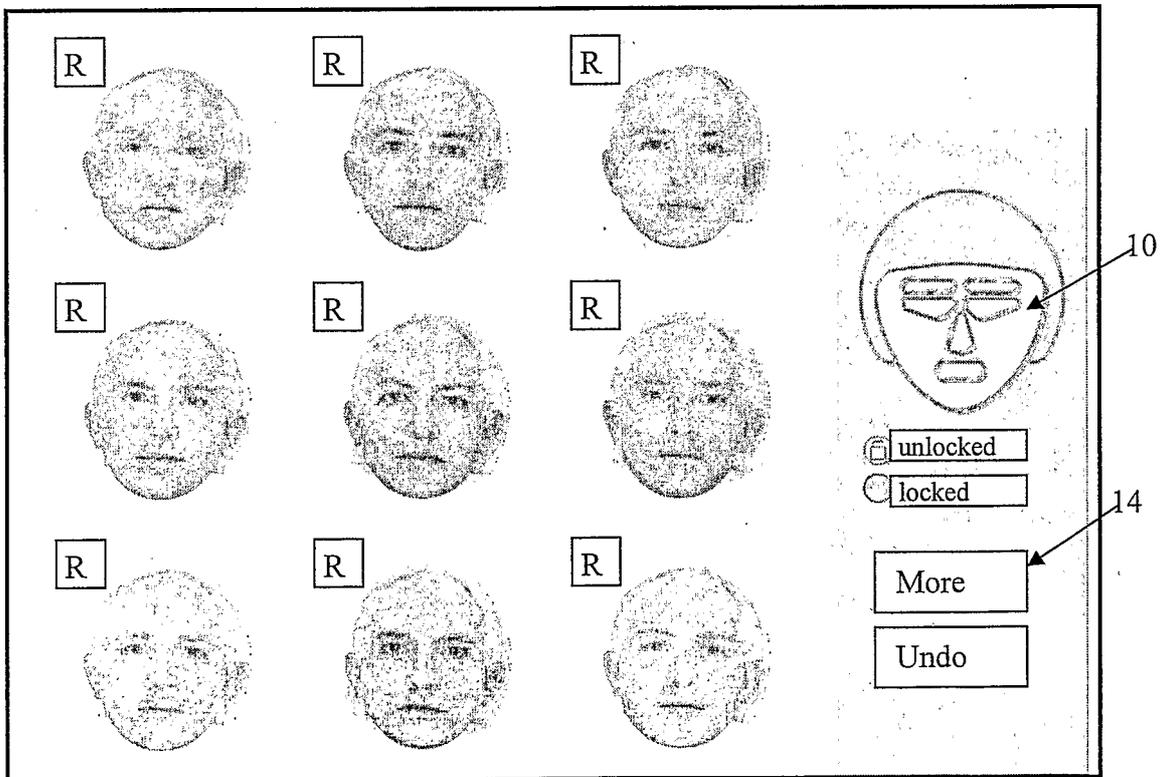


FIG. 2

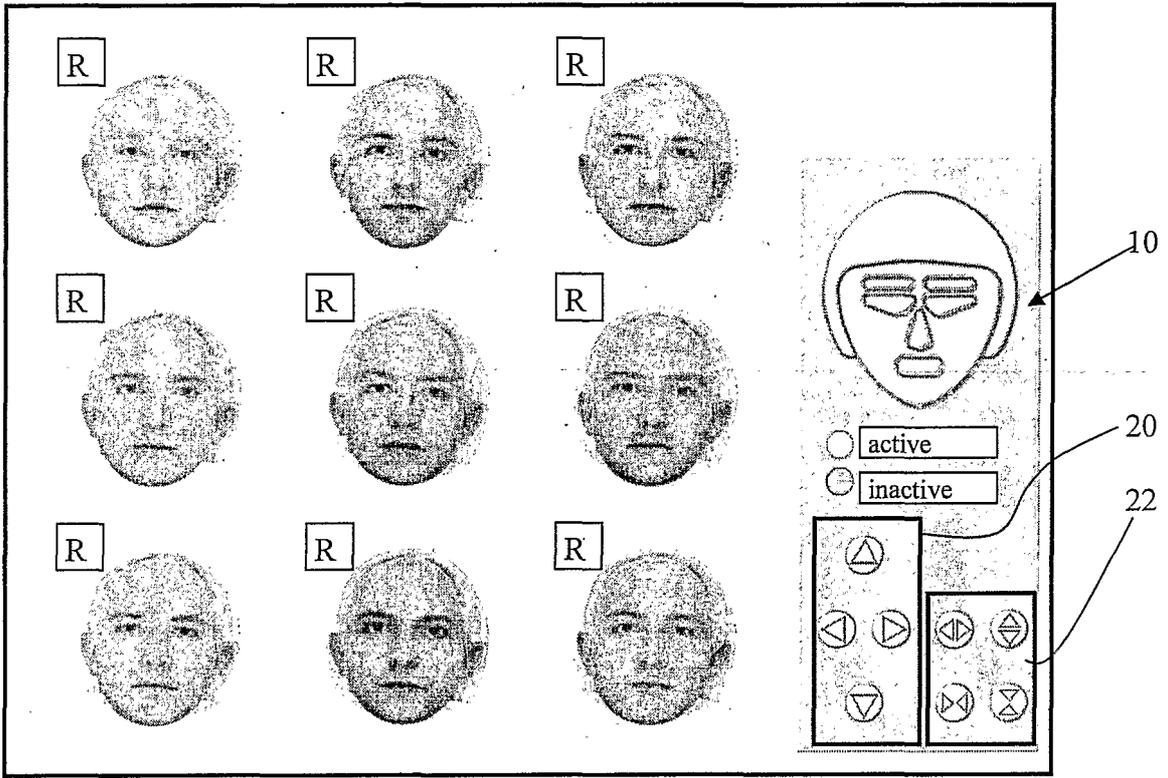


FIG. 3

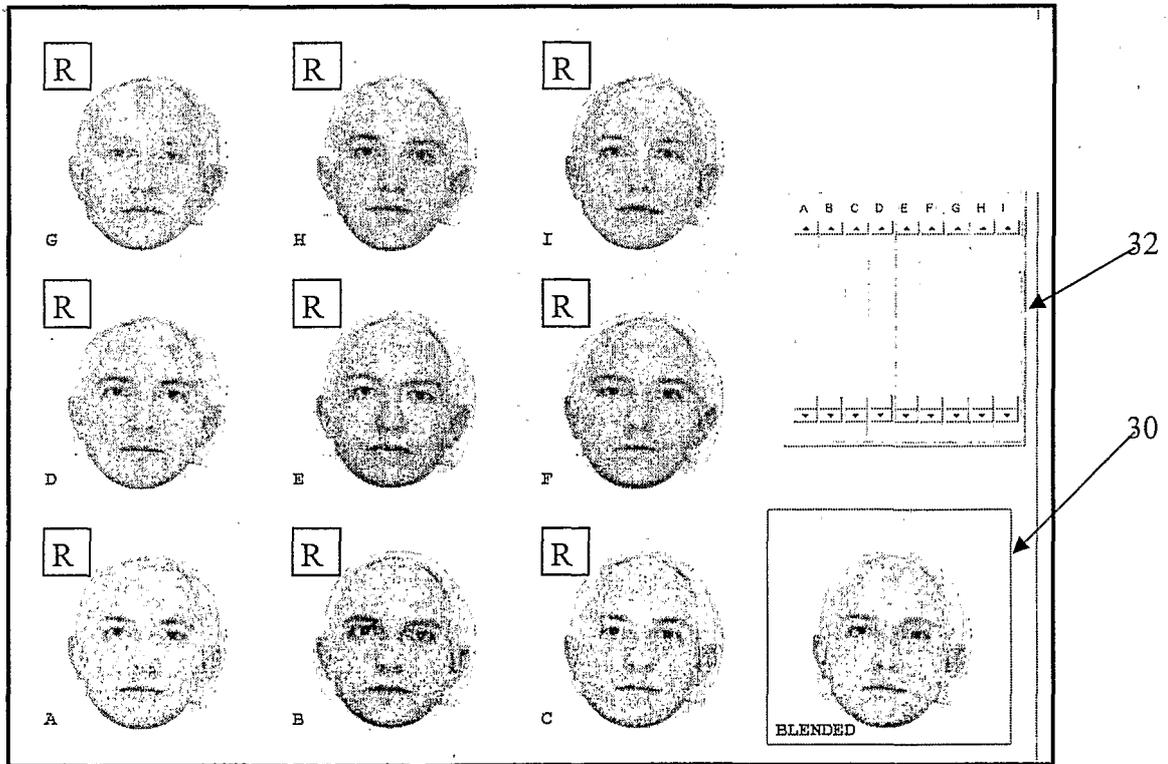


FIG. 4

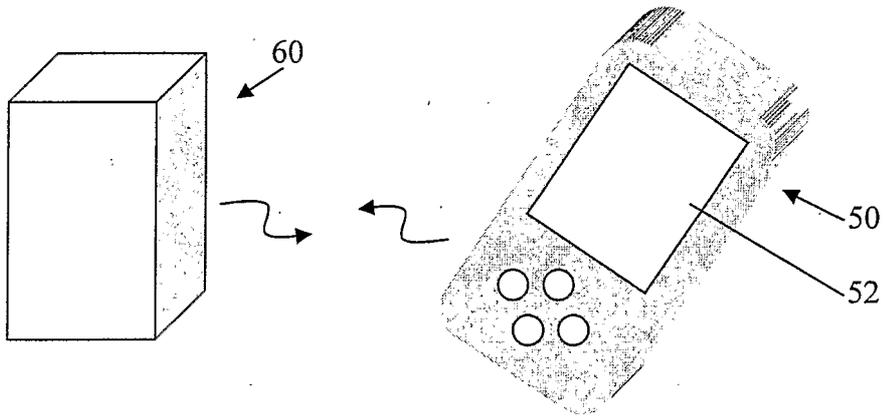


FIG. 5

INTERNATIONAL SEARCH REPORT

PCT/GB2005/002780

A. CLASSIFICATION OF SUBJECT MATTER IPC 7 G06T11/00		
According to International Patent Classification (IPC) or to both national classification and IPC		
B. FIELDS SEARCHED		
Minimum documentation searched (classification system followed by classification symbols) IPC 7 G06T		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic data base consulted during the international search (name of data base and, where practical, search terms used) EPO-Internal		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	GIBSON S ET AL: "SYNTHESIS OF PHOTOGRAFIC QUALITY FACIAL COMPOSITES USING EVOLUTIONARY ALGORITHMS" PROCEEDINGS OF THE BRITISH MACHINE VISION CONFERENCE, XX, XX, 2003, pages 1-10, XP009053617 cited in the application page 3, last paragraph - page 4, paragraph 3; figure 2 <div style="text-align: center;">----- -/--</div>	1-64
<input checked="" type="checkbox"/> Further documents are listed in the continuation of box C.		
<input checked="" type="checkbox"/> Patent family members are listed in annex.		
° Special categories of cited documents :		
A document defining the general state of the art which is not considered to be of particular relevance	*T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention	
E earlier document but published on or after the international filing date	*X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone	
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.	
O document referring to an oral disclosure, use, exhibition or other means	*&* document member of the same patent family	
P document published prior to the International filing date but later than the priority date claimed		
Date of the actual completion of the international search <div style="text-align: center; font-size: 1.2em;">22 September 2005</div>	Date of mailing of the international search report <div style="text-align: center; font-size: 1.2em;">06/10/2005</div>	
Name and mailing address of the ISA European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016	Authorized officer <div style="text-align: center; font-size: 1.2em;">Pierfederici, A</div>	

INTERNATIONAL SEARCH REPORT

PCT/GB2005/002780

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>HANCOCK P J B: "EVOLVING FACES FROM PRINCIPAL COMPONENTS" BEHAVIOR RESEARCH METHODS, INSTRUMENTS AND COMPUTERS, PSYCHONOMIC SOCIETY, US, vol. 32, no. 2, 18 May 1999 (1999-05-18), pages 327-333, XP001040938 ISSN: 0743-3808 cited in the application page 331, right-hand column, paragraph 2 - page 332, right-hand column, paragraph 2</p>	1-64
X	<p>US 5 375 195 A (JOHNSTON ET AL) 20 December 1994 (1994-12-20) cited in the application abstract; figure 2 column 3, line 47 - column 4, line 19 column 10, line 25 - line 31; claim 1</p>	1-64
X	<p>CALDWELL C ET AL INTERNATIONAL SOCIETY FOR GENETIC ALGORITHMS: "TRACKING A CRIMINAL SUSPECT THROUGH FACE-SPACE WITH A GENETIC ALGORITHM" PROCEEDINGS OF THE INTERNATIONAL CONFERENCE ON GENETIC ALGORITHMS. SAN DIEGO, JULY 13 - 16, 1991, SAN MATEO, MORGAN KAUFMANN, US, vol. CONF. 4, 13 July 1991 (1991-07-13), pages 416-421, XP000260130 page 417, left-hand column, last paragraph - right-hand column, paragraph 2</p>	1-64
X	<p>BRUNELLI R ET AL: "SpotIt An Interactive Identikit System" CVGIP GRAPHICAL MODELS AND IMAGE PROCESSING, ACADEMIC PRESS, DULUTH, MA, US, vol. 58, no. 5, September 1996 (1996-09), pages 399-404, XP004418978 ISSN: 1077-3169 page 401, left-hand column, paragraph 1 page 402, left-hand column, paragraph 1 - paragraph 2 page 404, left-hand column, paragraph 2 - right-hand column, paragraph 1</p>	1-64
A	<p>WO 01/16882 A (BIOADAPTIVE SYSTEMS PTY LTD; LIDDLE, ALAN, LESLIE; STANNARD, MARTIN, J) 8 March 2001 (2001-03-08) abstract; claim 1</p>	1-64

-/--

INTERNATIONAL SEARCH REPORT

PCT/GB2005/002780

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>COOTES T F ET AL: "ACTIVE APPREARANCE MODELS" IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, IEEE INC. NEW YORK, US, vol. 23, no. 6, June 2001 (2001-06), pages 681-685, XP001110809 ISSN: 0162-8828 cited in the application -----</p>	

INTERNATIONAL SEARCH REPORT

PCT/GB2005/002780

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 5375195	A	20-12-1994	NONE	
WO 0116882	A	08-03-2001	NONE	