Volume xx (200y), Number z, pp. 1–7

The Empathic Visualisation Algorithm (EVA) – An Automatic Mapping from Abstract Data to Naturalistic Visual Structure

A. Loizides and M.Slater

Department of Computer Science, University College London (UCL), London, UK

Abstract

This paper demonstrates a technique for visualising multidimensional data sets holistically. The EVA algorithm provides an automatic mapping *from semantically important features of the data to emotionally or perceptually significant features of the corresponding visual structure (such as a face). In other words a single glance at the visual structure informs the observer of the global state of the data, since the visual structure has an emotional impact on the observer. The visualisation is designed to correspond to the impact that would have been generated had the observer been able to access the relevant semantics of the underlying data.*

Categories: Information Visualization, Multi-Variate Visualization

1. Introduction

The 4 pictures of faces shown in Figure 1 represent the financial state of 4 different companies. In which one would you invest? Without knowing any details of what the facial expressions represent, it is very likely that you would have chosen the company represented by the face to the bottom right as your target of investment. In fact, more information can be gleaned from the faces by knowing that the degree of happiness represents profitability, the degree of fear represents lack of liquidity and so on.

The Empathic Visualisation Algorithm (EVA) is concerned with multivariate data visualiation [4,5,19], except that instead of using abstract diagrams to represent the data [23,12] a naturalistic visual structure is generated instead. By ‘naturalistic’ we imply one that is so familiar that people do not have to learn how to understand its meaning, and where this meaning is conformant with relevant meaning in the underlying data. An excellent example of such a visual structure is the apparent view of a human face. No one needs to be an expert in order to understand if a face is ‘sad’ or ‘happy’ - this understanding follows from being human. Complex emotional expressions (e.g., generally looking happy, but with a tinge of anxiety) can be easily recognised by human observers, and lead to their questioning the underlying data for the cause of such a visualisation. The EVA provides a global overview of what the data is ‘saying’ relative to the underlying values or interests of the data consumer.

Figure 1: *Examples of Faces Produced by the Method*

A basic problem lies in finding a mapping from the multivariate data set to the features of the visual structure that is designed to produce the appropriate emotional responses in the human observer. If the mapping is arbitrary (as in the case of Chernoff faces [7] for example) then this method has no intrinsic value - we require that the emotional response of an observer to the visual structure is an appropriate one corresponding to the goals that the observer has in understanding the data and the actions that should flow from this. The 'happy' face should lead the observer to think that (e.g.) the financial state of the companies that it represents is a good one, the 'sad' face indicating that action to avert a problem may be needed. We have devised an automatic mapping from the data to the characteristics of the visual structure that preserves this correspondence. This is done using Genetic Programming techniques [1].

There are two main problems in such data visualisation:

- Choosing a suitable mapping from the data to the chosen visual structure. The major difference is between arbitrary, handcrafted visualisations (constructed manually by a designer or end-user using appropriate tools) and automatic mapping (visualisations that are composed on the fly by the underlying system according to a set of pre-supplied layout rules and components). The problem of taking into account the impact the visualisation has on the emotional response of users in an automatically generated visualisation is fundamental to this work.
- The choice of a suitable representation of the data - whether abstract or more naturalistic. Use of naturalistic visual structures, based on entities encountered in everyday life that need no special knowledge for understanding by an average human observer, is at the heart of this research.

In this research we use an automatic mapping paradigm on a naturalistic visual representation, which up to now we have based on human facial expressions. This takes into account psychological research regarding characteristics of the face (e.g. universality of facial expressions). Paul Ekman spent more than 20 years studying faces, and derived a system called the Facial Action Coding System (FACS) which measures individual facial expressions [9]. This is the basis of the software system [21] that we have used to create faces with different emotional expressions. We call this method the 'Empathic Visualisation Algorithm (EVA)'.

2. Related Work

Information Visualisation is the visual presentation of information spaces and structures to facilitate their rapid assimilation and understanding.

The techniques described in the literature describe several different approaches to visualise multidimensional data sets of quantitative variables. There are numerous techniques that use an arbitrary mapping and abstract representations [8,12,16,10], some techniques that use arbitrary mapping and naturalistic representations [7,24], a few systems that use automatic mapping and abstract representations [14,17] and the proposed EVA scheme that uses automatic mapping combined with naturalistic visual structures [3,13].

Techniques falling in the first category (arbitrary mapping, abstract representation) present a number of problems. As the dimensionality increases, the complexity of those systems increases as well, overloading the display with excessive information. Such techniques require user-learning time to understand and be able to manipulate such interactive displays effectively. It becomes non-trivial to visualise relationships of variables beyond quadratic and in particular to get a holistic view of the data set.

The main example of the second category (arbitrary mapping, naturalistic representations) is that of Chernoff faces [7]. Hernan Chernoff was the first to realise the potential of using a human face as a representation for data. His idea capitalises on two important principles. Firstly, is our familiarity with human faces, and our ability to rapidly process even the smallest nuances and changes in a human face.Secondly, the fact that a human face evokes an emotional response in us and can therefore affect the way in which we behave. Since humans are optimised in some sense, for face recognition, it was hoped that using faces would aid in the grouping of the data and also to illustrate 'trends' in multi-dimensional data. The method involves assigning to each column of the data a facial feature such as width of the eye, position of the mouth, and for each row of the data constructing the face associated with the assignment. It is believed to be able to represent a total of 20 different dimensions.

Many researchers have built upon Chernoff's initial idea [11,15,18,20,22]. However, there are a number of limitations when using Chernoff faces. The most important of all is that the variables in the representation are treated equally, but their emotional impact depends on which part of the face they affect. Variables assigned to an area near the mouth or eyes have higher emotional significance compared to variables assigned to a different facial feature. Hence, this arbitrary mapping from data variables to facial features does not take into account the impact on the emotions of the observer. Moreover, it is necessary to spend some time training test subjects as to which features apply to which variables. It can be argued that the visualisation loses its effectiveness when we have extreme values, since these will produce unrealistic faces.

In the category of automatically generated visualisations to abstract representations the visualisations are composed on the fly by the underlying system according to a set of pre-supplied layout rules and components. These pre-supplied rules and components restrict both the problem domain and the visual interface. Such techniques require user-learning time while they widen the gap between expert and non-expert users.

The problems with existing methods described above led to the original motivation to build a system using naturalistic representations with an automatically generated visualisation. The system, EVA, has been built and initial experiments with the method [13] show the following advantages over existing methods:

1. Interrelationships between the data variables are incorporated into one visual structure.
2. The visual structure gives an overview of the whole data set.
3. The visual structure is simple and thus easy to understand.
4. Both experts and non-experts are able to give a quick overview of what is going on.
5. The complexity of the system does not increase with dimensionality of the data.

3. Method

3.1. Basic Approach

Throughout this paper we are considering the representation of multivariate data in an $n \times k$ data matrix X, consisting of n cases on k quantitative variables $x_1, x_2, \ldots, x_k$. Each row in the data matrix typically represents an individual and there are k observations made for each individual. The objective is to construct a visualisation of the data matrix where the salient features of the data can be intuitively recognised by an observer, such that the representation gives an overall view of the data set. Within this overall objective there are two fundamental goals:

1. Naturalistic visual representation. The visual structure should be something encountered in everyday life, something that does not require special knowledge for interpretation by a normal human observer.
2. Automatic mapping. The mapping should be such that semantically important features in the data are mapped to perceptually or emotionally important features of the visual structure.

Examples of (1.) include faces, buildings, towns and others. No human needs to be an expert to recognise the emotional content of another human face - it is recognisably ‘happy’, ‘sad’, ‘angry’, ‘relaxed’, ‘fearful’, and various possibilities of these basic emotions, including neutrality. We frequently use, throughout this research, an example of a ‘face’ because it is the epitome of a naturalistic visual structure in the sense we mean here.

Given a set of data, and given a visual structure, it is trivial to construct a mapping from the data to a visual structure; for example map variable x_i to the ith facial feature. However, such a mapping is arbitrary; it does not take into account the impact of the face on the emotions of the observer. Now the data is of interest to the observer for some reason: associating with the data there will be some ‘value system’ (metric) reflecting the interest or importance or consequences of aspects of the data for the situation of the observer. A fundamental goal, reflected in (2.) is that the perceptually or emotionally significant features of the visual structure directly reflect the value system over the data. We call this *visual homomorphism*. Hence the mapping from data to visual structure is not arbitrary, but constructed in such a way that this visual homomorphism is realised.

What follows is a description of how such a mapping is constructed in the situation where each row (i.e. individual) in the data is to be represented by a different instance of the same type of visual structure. That is to say, by using faces as the visual structure, each row is mapped to a different face.

Let $\nu_s(X)$, $s = 1, 2, \ldots, p$ be p given functions over the data representing specific utilities or values over the data matrix. Suppose for example that the data matrix represents a set of customers of a telephone company, and the variables are quantities such as age, gender, marital status, income, number of years with the company, number of telephone calls made per week, monthly phone bill, and so on. One utility or value might be a function of the overall age distribution of the population - such as the mean age, or the percentage over 65, or the percentage under 20. Another value might be the ’flatness’ of the data - for example the ratio of the variance of the first principle component to the total variation in the data. Another value might be the correlation coefficient between age and monthly bill, and so on, for many other quantities that characterise the interests or value system of the observer. That is, the user of the data must be able to specify a set of utility functions, relationships over the data, that are important - that essentially define the reasons why the data is of interest at all.

Now consider a visual structure (Ω). We will use *p importance metrics*, that are characteristics of Ω which are measurable and significant to human perception or emotions, $e_s(\Omega)$, $s = 1, 2, \ldots, p$. In the example of the face these might be the degree of anger, happiness, anxiety, - or characteristics such as age, beauty, gender and so on.

The fundamental goal, in terms of any X and any Ω , is to produce a mapping $\mu(X) \longrightarrow \Omega$ such that the values over the data matrix are reflected in characteristics of the visual structure. In particular that $e_s(\Omega)$ is a monotonically increasing function of $\nu_s(X)$. If happiness reflects the utility function for a required relationship (financial stability) between profitability, investment, turnover and season, and anger represents the utility function for the rate of staff leaving a company, then a face with combined characteristics of happiness and anger will immediately present to the observer a situation where there may be high financial stability, but also too high a rate of staff turnover. Of course many such combinations can be considered in parallel, not simply two at a time.

A *characteristic* is a measurement of some aspect of Ω as a whole (such as the emotions of a face) rather than some individual *feature* (such as the shape of the mouth). A characteristic is a measure representing the totality of the face such as the degree of 'happiness'. The manifestation of a characteristic (such as happiness) depends on many different individual *features* of the face - specific configurations of muscle tensions. Similarly, the appearance of beauty or age or gender is derived from many different features - such as size of the eyes, inter-ocular distance, shape of the mouth, symmetry, and so on. In other words, *features* are the individual components that make up a face - such as the specific configuration of muscle tensions for a specific face, or the geometric and material properties of the actual features (eyes, colour of the eye, mouth, nose, lips) that make up a face. Therefore, we can *render* a face knowing its features. Once rendered the face will then have a set of measurable *characteristics*.

For instance, suppose that there are r features of the visual structure: $\phi_t(\Omega), t = 1, 2, ..., r$. Knowing these features of Ω results in being able to render it. Once rendered, we can measure it to determine its characteristics $e_s(\Omega), s = 1, 2, ..., p$.

Finally, we introduce *feature functions* over the data matrix: $f_t(X), t = 1, 2, ..., r$. These functions completely determine the features of the visual structure, in fact

$$\phi_t(\Omega) = f_t(X), t = 1, 2, ..., r \quad (1)$$

In other words, the values of these functions are *interpreted* as the values for the features of the visual structure. Our goal is to select these functions f in order to attain an accurate correspondence between the value system over the data and the characteristics of the visual structure. Therefore, attaining the visual homomorphism.

Let $\nu = (\nu_1, \nu_2, ..., \nu_p)$ and $e = (e_1, e_2, ..., e_p)$. Suppose that, $\| \nu - e\|$ is a measure of the 'distance' between these two *vectors*. Then the specific goal is to choose $f_t(X), t = 1, 2, ..., r$ such that $\| \nu - e\|$ is minimised.

3.2. Genetic Programming

EVA tackles the minimisation problem introduced above (the selection of appropriate functions) using Genetic Programming (GP) techniques.

Let $F_i = (f_{1i}, f_{2i}, ..., f_{ri})$ be a specific set of feature functions which when applied to X (the data set) produce the visual structure features. We choose at random a large collection of such sets of functions F_i, $i = 1, 2..., N$. This collection defines a *population* of sets of feature functions and is ready for 'evolution'. The ith member of the population produces a specific visual structure Ω_i. This visual structure has characteristics $e_s(\Omega_i), s = 1, 2, ..., p$. These characteristics can be used to produce the distance measurement $\| \nu - e\|$. The distance measurement can be used to compute a 'fitness' for the ith member of the population. Hence, each member of the population has an associated fitness, which can be expressed as a probability. These probabilities determine survival into the next generation and selection for mating - thus producing a second generation. The process continues until (possible) convergence. The most fit member (i.e. set of feature functions) of all populations is chosen for the required mapping. Figure 2, shows a diagrammatic review of the method.

3.3. Measuring Emotional Expressions

Values returned by the feature functions $f_t(X), t = 1, 2, ..., r$ were assigned to be contractions of individual muscles in the face [21]. These values allow an individual face to be rendered. The muscles are shown as red lines in bottom right face of Figure 3, and areas affected by the muscles are shown in the face next to it. The lighter the colour the more the muscles affect that particular area.

An experimental set-up was created to quantify emotional expressions on the face, for example, the scale of happiness-sadness. This was achieved using a number of landmark points on the face that are significantly influenced by muscle contractions expressing the various emotions measured. Two sets of data (randomly created facial expressions) were created, the first consisting of 250 faces and the second 150 faces, and positions of each of the landmarks, for each face, were recorded. Each face was subjectively assessed for emotional state by at least 3 different people (from a pool of about 30 people) for both data sets. The answers of the subjects for the *first data set* only were used to create symbolic regression equations with response variable y_i representing the individual emotional expression and explanatory variables x_i representing each of the landmarks on the face. The *second data set* was then used to verify the equations formed, and high positive correlations (all above 0.7) were found between the results from the equation produced and the subjective evaluations.

We use these estimated regression equations [2] to measure the emotional expressions for the method we describe here. Note that this data analysis on generated faces only had to be done once, and the results can be used in any number of different visualisations.

4. Results

Figure 3, shows some sample results of the method applied to financial data. The data set used to construct these faces, includes Balance-Sheet and Profit & Loss accounts for 250 companies. Each row represents an individual company for a specific year. In the data there are a number of highly correlated variables and it is very hard if not impossible to get an overall view of the performance of the company.

The goal is to get an overall understanding on how well these companies were performing on a number of utility functions. For these we used 3 financial ratios that cover important features of the data. These are the $\nu_s(X), s =$

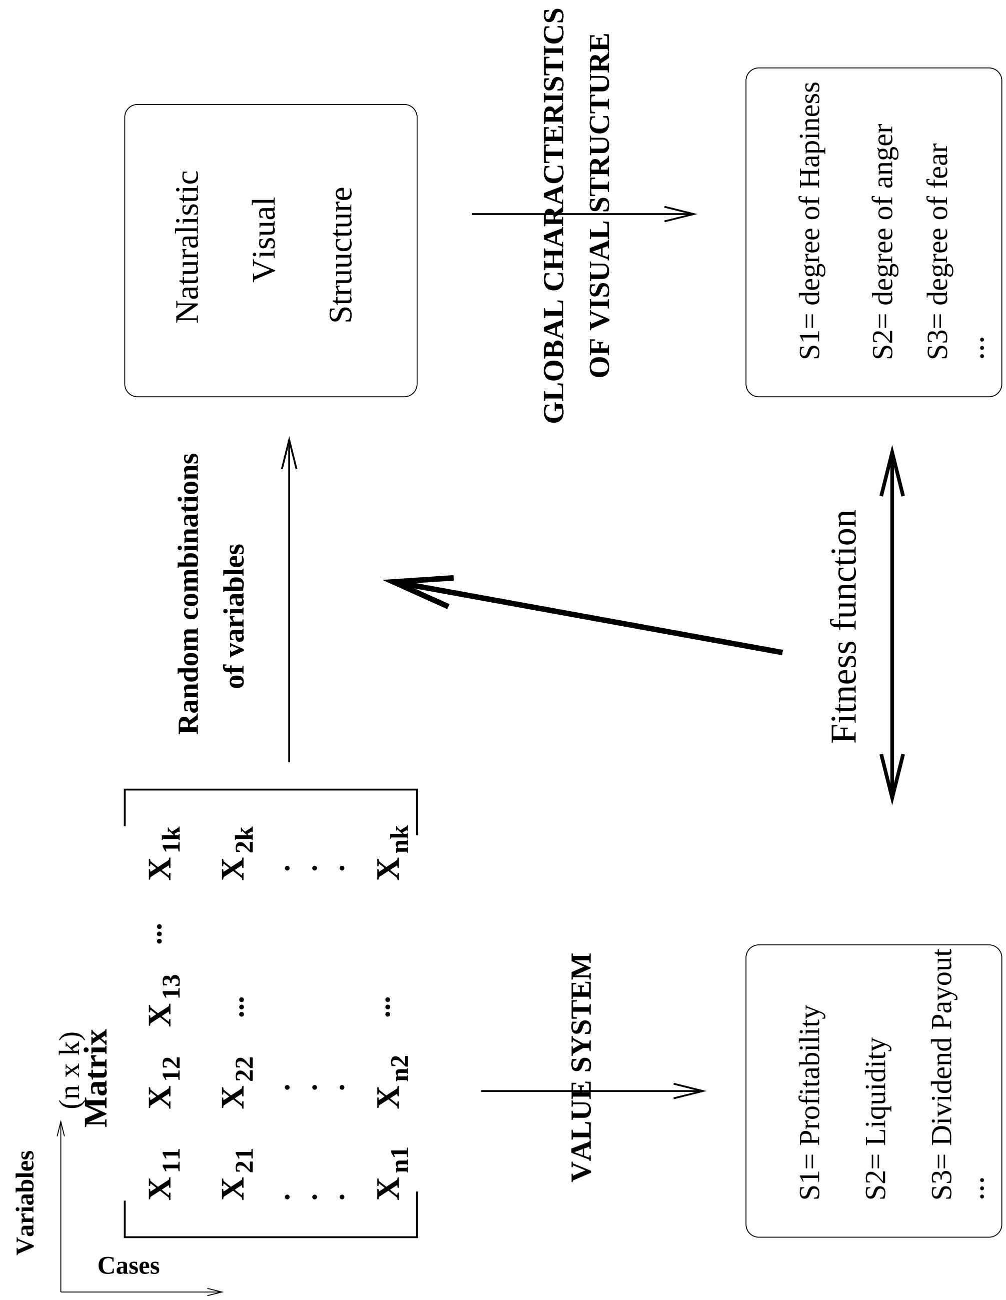

Figure 2: *From Data to Visual Structure: An automatic mapping*

$1, 2, \ldots, p$ functions where $p = 3$. The financial ratios were: profitability and effectiveness in the management of the capital of the firm (**Return on Capital Employed**), ability of organisation to cover its immediate liabilities (**Current Liabilities over Total Assets**) and risk ness and long-term solvency (**Gearing Ratio**).

The corresponding emotions on the face used are:

1. Happiness-Sadness scale representing Profitability. The happier the face the more profitable the organisation.
2. Degree of Anger representing Liabilities. Angry faces representing companies that are "technically insolvent".

3. Degree of Fear representing Gearing. Risky organisations are mapped to fearful faces.

When using the EVA method we can produce a visual structure that represents the company as a whole, and therefore quickly form a view regarding overall performance of the company.

When looking at the ‘face matrix’ in Figure 3 , it is easy to spot happy faces (second face in the first row, first face in last row, second face in the third row) that may mean a good overall performance according to profitability. It is also easy to spot sad, angry and fearful faces that can denote differing types of poor performance. The mixing of these emotions in one naturalistic visual structure, results in a very powerful visualisation that is easy to interpret. This could become even more powerful, if we could actually interact with and question the visual structure to get more information about the data set - in other words perform backtracking from the visualisation to the data. This is the subject of further research and out of the scope of this paper.

5. Conclusions

We have introduced a new method for constructing an automatic mapping from data to visual structure, which enforces a homomorphism between important characteristics of the data and the emotional or perceptual impact of the visual structure. Salient global aspects of the data (the utility or value functions) are mapped to emotionally or perceptually significant aspects of a visual structure. The features that allow rendering of the visual structure are determined by a genetic program that breeds generations of visual structures such that in each successive generation there is a greater match between the utility functions and the visual structure features.

The type of visual structure produced by this method are meant to be informative ‘at a glance’, and can also reveal important detailed information or unusual characteristics present in the data (e.g., a happy face with a hint of anxiety). The method is not put forward as an alternative to other types of visualisation, but rather it provides a ‘first-pass’ visualisation that may, in particular applications, raise interesting features that then may be explored in detail through traditional visualisation techniques, or indeed statistical analysis.

Acknowledgements

Special thanks goes to the Knowledge Lab of NCR Financial Services who financially supports this work.

References

1. John R. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, 1992.
2. A. Loizides, M.Slater, W.B. Langdon, “Emotional Expressions Using Genetic Programming, Springer Engineering Books”, *6th World Conference on Soft Computing, WSC6*, 2001.
3. A. Loizides, M.Slater, “The Empathic Visualisation Algorithm (EVA): Chernoff faces re-vistited”, *Conference Abstracts and Applications, Technical Sketch SIGGRAPH 2001*, pg 179, 2001.
4. Stuart K. Card, Jock D. Mackinlay, and Ben Schneiderman. *Readings in Information Visualisation: Using Vision to Think, San Francisco, CA*. Morgan Kaufman, 1999.
5. M. Chalmers. Tutorial on Information Visualisation. *VLDB Conference*, 1999.
6. C. Chen and M. Czerwinski, “Spatial ability and visual navigation: An empirical study”, *News Reviews of Hypermedia and Multimedia* **3**, pp. 67–89 (1997).
7. H. Chernoff. The use of faces to represent points in n-dimensional space graphically, Department of Statistics, Stanford University. *Research Note NR-042-993*, December 1971.
8. Stephen G. Eick, Joseph L. Steffen and Eric E. Summer Jr, “A tool for visualising line oriented software statistics”, *IEEE Transactions on Software Engineering*, pp. 957–968 (1992).
9. P. Ekman and W. V. Friesen. Facial Action Coding System (Investigator’s Guide). *Consulting Psychologists Press, Inc.*, Palo Alto, California, USA, 1978.
10. S. Feiner, C.Beshers, “Worlds within worlds: metaphors for exploring n-dimensional virtual worlds. In ACM, editor, UIST. Third Annual Symposium on User Interface Software and Technology”, *Proceedings of the ACM SIGGRAPH Symposium, Snowbird, Utah, USA*, pp. 76–83 (1990).
11. H. Homa and Schwartz, “Perceptibility of schematic face stimuli: evidence of a perceptual gestalt”, *Journal of Memory and Cognition* **4**, (1976).
12. A. Inselberg, “Parallel coordinates for multidimensional displays. In Spatial Information Technologies for Remote Sensing Today and Tomorrow”, *The Ninth William T. Pecora Memorial Remote Sensing Symposium*, pp. 312–324 (1984).
13. A. Loizides, M.Slater. An application of the Empathic Visualisation Algorithm (EVA), Department of Computer Science, University College London. *Research Note RN /99/50*, 1999.
14. J. Mackinlay, “Automating the design of graphical presentations of relational information”, *ACM Transactions on Graphics* **5**(2), pp. 110–141 (April 1986).

15. Moriarity, "Communicating financial information through multidimensional graphics", *In Journal of Accounting Research* **17**, pp. 205–224 (1979).

16. R. Rao and S. Card, "Exploring large tables with the table lens", *In Proceedings of ACM CHI'95 Conference on Human Factors in Computing Systems, Videos* **2**, pp. 403–404 (1995).

17. S. Roth, J. Kolojejchick, J. Mattis, M. C. Chuah, J. Goldstein and O. Juarez, "SAGE tools: A knowledge-based environment for designing and perusing data visualisations", *In Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systems* **2**, pp. 27–28 (1994).

18. G. M. Smith and R. Taffler, "Improving the communication of accounting information through cartoon graphics", *In Accounting, Auditing and Accountability Journal* **9**, pp. 68–85 (1996).

19. R. Spence. *nformation Visualization*. Addison-Wesley, May 2000.

20. J. Walker, J. Sproull and R. Subramani, "Using a human face in an interface", *In Human Factors in Computer Systems*, pp. 85–91 (1994).

21. K. Waters and F. Parke, "A muscle model for animating 3- dimensional facial expression", *In M. C. Stone, editor, SIGGRAPH '87 Conference Proceedings, Computer Graphics* **21**(4), pp. 17–24 (1987).

22. R. Wilkinson, "An experimental evaluation of multivariate graphical point representations", *In Human Factors in Computer Systems, ACM*, pp. 202–209 (1982).

23. C. Ahlberg and B. Shneiderman, "Visual Information Seeking using the FilmFinder, Conference on Human Factors in Computing Systems", *In Proceedings of ACM CHI' 94* **2**, pp. 433 (1994).

24. M. Alexa, W. Muller, "Visualisation by examples: Mapping data to visual representations using few correspondences. In Proceedings of the Joint EUROGRAPHICS and IEEE TCVG", *Symposium on Visualisation in Vienna, Austria* **,** pp. 23–32 (May 1999).

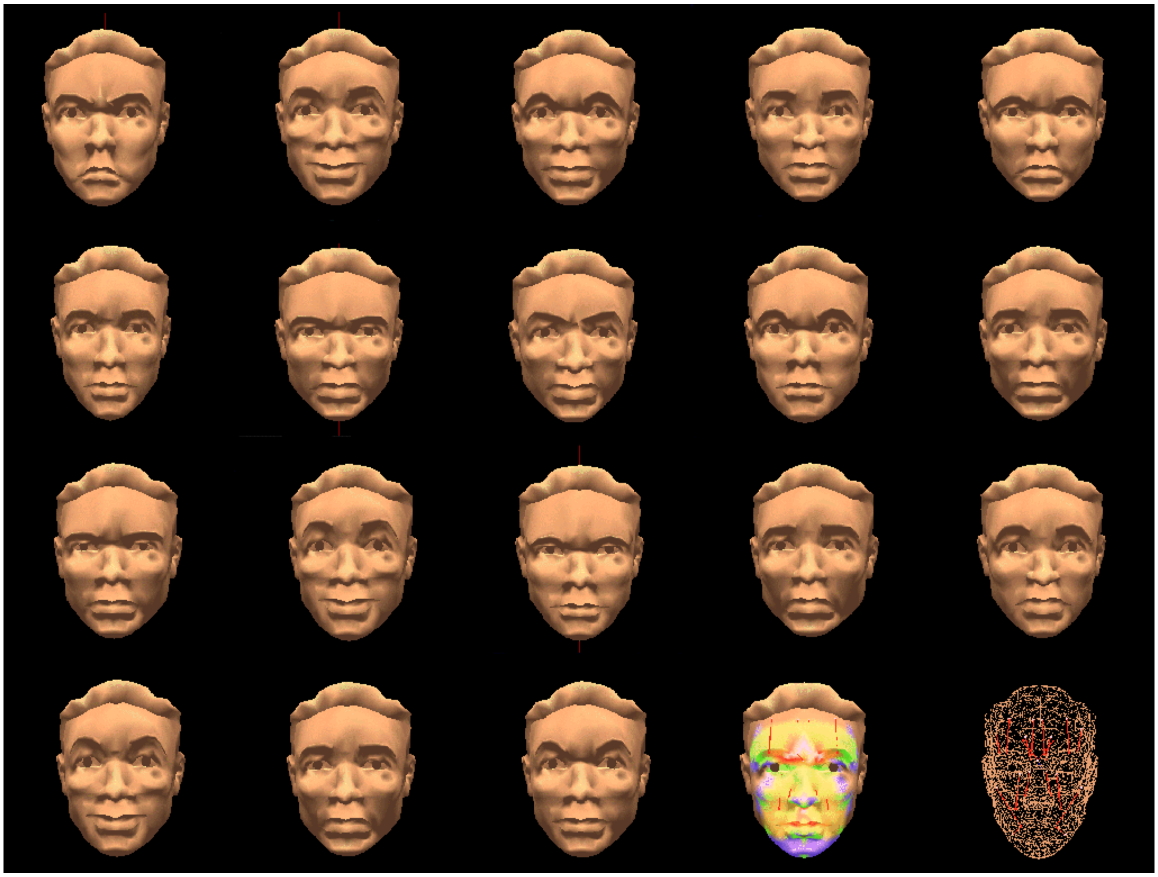

Figure 3: *More faces produce by the method*