Face Fitting using a Genetic Algorithm.

D. W. Hunter[1], B. P. Tiddeman[2] and D. I. Perrett[1]

[1]School of Psychology, University of St Andrews, St. Andrews, Scotland, UK
[2]Department of Computer Science, University of Aberystwyth, Wales, UK

Abstract

The ability to estimate the three-dimensional structure of a human face from a single image has many useful applications, including but not limited to the fields of face recognition, modelling and actors' head replacement. Although a lot of promising techniques have been developed in this area a number of significant technical hurdles exist preventing the development of a reliable method to extract three-dimensional structure without human intervention.

In this poster we present work in progress *on a new technique that should overcome some of the problems associated with the current state-of-the-art methods. We intend to use* Genetic Algorithms *to search the face space to find a match with a given input image. We hope to show that this method will avoid local-minima problems associated with traditional gradient descent techniques as well as provide a method that locates the position, pose and shape of the face without any user interaction.*

Categories and Subject Descriptors (according to ACM CCS): I.3.6 [Computer Graphics]: Methodology and Techniques—I.4.8 [Computer Vision]: Scene Analysis—Surface fitting

1. Introduction

Most current methods involve minimizing a cost function based on the L^2-norm between a rendered face model with a particular set of parameters and a target image. As the derivatives of this function can be approximated, many previous authors have used gradient descent methods. However these methods are prone to local-minima problems. Also the derivatives are only approximated and are only valid if the face is already closely aligned with the target image. These derivatives introduce a new source of error to the fitting that is most pronounced when the gradient is shallow as well as a windowing effect that makes it difficult to detect shape updates that differ significantly in scale. We plan to use an alternative minimization approach that will avoid many of the problems associated with gradient descent, called a *Genetic Algorithm*. This algorithm uses the 'best' results from the previous iteration to seed a new set of trial parameters. This allows a greater proportion of the parameters space to be analysed and so is less likely to become 'stuck' in a local-minima and so find a globally optimum.

In their original paper on Morphable Models, Blanz and Vetter used a stochastic gradient descent method to minimize the L^2-norm between a synthesized face image and a target image [BV99]. Romdhani also developed a multi-feature fitting strategy that combined, in a Baysian fashion, a set of different differentiable cost functions designed to extract different aspects of the image. For example edges, and particular illumination artefacts, such as specular reflection [Rom05]. Like previous methods these functions were differentiable and required a good initial estimate of parameters. Xiao et al used a 2D to 3D method whereby an Active Appearance Model was constructed from a 3DMM. Thus methods developed to fit and track AAMs can be used with 3D models. However the combined model also spans a large set of parameter values that result in invalid 3D shape models [XBMK04]. These methods all suffer from both the local-minima and windowing problems described above. Moghaddam et al. [MLPM03] used an XOR based function to fit a model to a silhouette of the boundary between the face and the image background. As this function is not differentiable they used a simplex minimization method.

2. Genetic Algorithm

In order to extract 3D facial features we minimize the L^2-norm of the difference between a rendered 3D face model and the target image.

The algorithm begins by generating a set of samples from a distribution believed to contain the global minima of an error function. The error function is applied this set of samples and the *m* best selected as 'parents.' A new set of samples are generated from the parent set by selecting random pairs of parents and combining their parameters. The combined pairs are randomly mutated to introduce variation into the population. The error function is applied to these child samples and the *m* best become the 'parents' of the next generation. The algorithm repeats until either a sample is found that sufficiently minimizes the error function or, no further improvement is made over a pre-defined number of generations.

Selection Each subject in the current generation is evaluated using the fitness function and a subset with the 'best,' i.e. lowest, scores selected as parents for the next generation. These 'parents' survive into the next generation and are randomly paired to produce offspring.

Cross-Over In order to create a set of new 'child' subjects from a pair of 'parents' we select, randomly, one of the two parents and copy that parent's parameter. Thus creating a new subject that retains the desirable attributes of its parents.

Mutation In order that the whole parameter space is potentially explored by the algorithm, each parameter has a small chance of being mutated, the probability of mutation is known as the *mutation factor*.

Elitism The best result from the previous generation is preserved in the new generation. This makes the search similar to a 'down-hill' search.

3. Current Progress

As a proof-of-principle we have trialled the algorithm on synthetic images, created by rendering the Morphable-Model with a number of different shapes and poses. Figure 3 shows the results of fitting to 3 different synthetic images. The images are generated from random shape and pose co-efficients, constrained such that the shape co-efficients are with in the range of the Morphable-Model and the pose co-efficients are in-frame and facing between 45° of face-forward around all three principal axes.

4. Further Work

Fitting a 3D model to a synthetic image with a black background is clearly significantly easier than to a real-world image. Clearly this is the next step for our research. Also Genetic Algorithms provide a frame work for the investigation of non-continuous and non-differentiable functions. A investigation of alternative cost functions such as the XOR function suggested by Moghaddan et al [MLPM03] is one possible function to investigate.

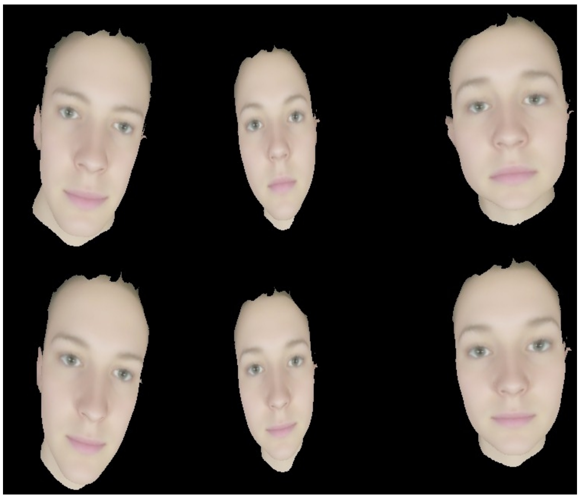

Figure 1: *The results of applying the fitting algorithm to three synthetic images (top row). The resulting models are show in the bottom row.*

5. Conclusion

Although preliminary we believe that the results show that using Genetic Algorithms to solve the face-fitting problem is viable, and that research in the area is worth pursuing.

Acknowledgement

The authors would like to thank Vinet Coetzee, Janek Lobmaier and Ian D. Stephen for capturing the 3D heads used in this paper and Davis Buls for manually placing landmarks on each of the 3D models. This work has been funded by Unilever plc.

References

[BV99] BLANZ V., VETTER T.: A morphable model for the synthesis of 3d faces. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1999), ACM Press/Addison-Wesley Publishing Co., pp. 187–194.

[MLPM03] MOGHADDAM B., LEE J., PFISTER H., MACHIRAJU R.: Model-based 3d face capture with shape-from-silhouettes. In *In IEEE International Workshop on Analysis and Modeling of Faces and Gestures* (2003), p. pages.

[Rom05] ROMDHANI S.: *Face Image Analysis using a Multiple Feature Fitting Strategy*. PhD thesis, University of Basel, 2005.

[XBMK04] XIAO J., BAKER S., MATTHEWS I., KANADE T.: Real-time combined 2d+3d active appearance models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (June 2004), vol. 2, pp. 535 – 542.