

Automatically Generating Virtual Humans using Evolutionary Algorithms

Adrian Albin-Clark and Toby Howard

School of Computer Science, The University of Manchester

Abstract

Virtual Humans are used in many applications either as an embodiment of a real person (an "avatar"), or under the control of a computer program (an "agent" or "non-player character"). The automatic generation of Virtual Humans is a challenging problem if they are to look both plausible and unique within a population. We present an approach which exploits the power of Evolutionary Algorithms (EAs), and provide illustrative examples of how our methods may be realised within the context of surface-based model geometry.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods, and Search - Heuristic methods I.6.5 [Simulation and Modeling]: Model Development - Modeling methodologies

1. Introduction

Applications for Virtual Humans cover the whole range of human activity, including simulation, training and games. Virtual humans can play many useful roles, some of the more common examples being patient, actor, teacher (as proposed by Jefferey [Jef08]), learner, soldier, worker, individual in a crowd, and conference delegate [MTT04].

A choice of starting avatars is often given when a person uses a multi-user virtual world for the first time, as seen in Figure 1, which shows avatars from Kaneva [Kan]. If the starting avatar is not customised, it effectively becomes a clone in the world. In some applications, an avatar's ap-



Figure 1: Kaneva's starting avatars. Left to right, for both female and male: Asian; Goth; Mixed; Irish; African; Blond.

pearance can be progressively customised such that it is no longer plausibly human.

For agents, there is often no reason for them to look like any living, identifiable person in the real world, although they should look like a *possible* person. The stochastic generation of virtual humans becomes even more challenging if visual diversity is a prerequisite.

A 3D model for an avatar or an agent may be created by a variety of standard methods, such as modelling software, 2D images which are processed for feature extraction, 3D range scans of real humans, and automated processes that attempt to constrain the final result.

The study of Virtual Humans is concerned with issues such as morphology, textures, animation, and behaviours. We focus primarily upon the morphology of Virtual Human models, and for clarity of presentation, restrict our graphical examples to human heads, but include discussion of full-body models in the approaches we present.

In this paper we present Evolutionary Algorithms (EAs) as an approach to the challenging problem of generating plausible Virtual Humans within a population.

2. Related Work

Early work by Fujiwara and Sawai [FS99] demonstrates a method that breeds geometry in order to approximate a 3D surface with simple facial features, starting from a 2D surface. For the generation of entire Virtual Humans, Ventrone [Ven00] refers to a genetic approach to avatar design and customisation and how it is possible, with this in place, to breed using sexual reproduction. He is also the chief author of a patent [VHW03] which sets out the client-server details for managing genetically-based avatars, which is used in *Second Life* [Sec]. The genotypes are non-geometric, floating-point representations of an avatar's anatomy and behaviour. Although there are some references to a model blender, and to the generation of many unique avatars using random numbers for genotypes, there is no mention of using an EA. In the literature, the interactive genetic algorithm paradigm is mostly employed, which Dawkins [Daw06] proposes, where users' aesthetic choices drive evolution. FaceLift [Facb], a tool to navigate through face space is developed by DiPaola [Dip02] for *The Sims* [Sim], to create game characters, offering slight mutations of faces for selection, until a favourite is found. In a similar spirit, Frowd et al. [FHC04] describe a system called EvoFIT, similar to the non-evolutionary application by Albrecht et al. [ABHS05], using a face model database, combined with an evolutionary algorithm, to converge towards a final face based upon user selections.

3. Experiments with Evolutionary Algorithms

Evolutionary Algorithms belong to the field of research known as Natural Computing, and are based on a combination of Darwin's theory of natural selection and genetic inheritance.

The general approach to EAs is that individuals in a population with higher fitness scores for an optimisation problem, are selected for reproduction. Over successive generations the population evolves better solutions to the problem.

In the following sections we shall look in more detail at some of the key aspects of an EA.

3.1. Representation Methods

A genotype is the encoding of an individual within a population. A phenotype is a decoded genotype, and a solution to some problem. Goldberg [Gol89] provides a simple example in which a genotype with binary string 01000 represents a phenotype, which when decoded, has a decimal value 8. Whilst this mapping from genotype to phenotype is both obvious and natural, in the general case, it should be possible for all solutions to be represented. In the building block approach common to EAs, which Mitchell [Mit98] describes, a chromosome should be represented by a short, low order string of values. The genotype-to-phenotype mapping usually involves some processing. There are three numerical




Figure 2: One of the 237 face groups from the *MakeHuman 0.9.1 unisex body mesh* (11222 vertices).

representations: floating-point; integer; and binary. We discuss examples of representations using just one type.

If the genotype-to-phenotype mapping is defined as the rendering process, a rendered model could be said to be a phenotype derived from an unrendered model, the genotype, comprising a string of floating-point coordinates. However, the use of such a lengthy chromosome holding absolute or relative (to a reference model) geometric data clearly defies the building block approach. A shorter chromosome holding blend weights, where each position in the string corresponds to a model, or part of a model, is to be preferred. The *Second Life* application uses a simple avatar mesh; for example, there are no toes. The "gene pool" for the body shape is defined by 76 floating-point numbers, each one encoding an avatar characteristic. Some examples of these measures are: "Arm Length"; "Body Fat"; and "Torso Muscles". This non-geometric, lightweight avatar representation is designed to be transmitted across networks, with the base mesh deformation being performed on the client machine, based upon the genetic value, or α , applied to morph targets. The *MakeHuman* project [BRM08] [Mak] uses a similar approach to control ethnicity, gender, age, muscle tone and body mass, though not explicitly in a genetic context.

Integers can be used as a reference to a face group from a larger model (as shown in Figure 2), or a complete model, where each is enumerated. The main advantage of face groups is that they provide semantics about the structure of a model, allowing recombination of subsets of models. Alternatively, a face group, or combination of them, might be deformed in some predefined way, similar to the *Second Life* and *MakeHuman* approaches.

A binary string representation can be used to represent the inclusion of a model in a morph. If there are 100 models, and the string length is 100, each position indicates whether

or not that model is included as part of the morph, with the weights being computed elsewhere.

Parke’s approach to face modelling calculates the geometry of one half of a face from the other by reflecting about the plane of symmetry [Par72], on the basis that a face is almost symmetrical. If this method is used, the chromosome length is effectively halved, though at the expense of a more natural-looking, asymmetrical human.

Roberto et al. [RVCN08] present a minimal chromosomal representation of the characteristics of a person using anthropometric head measurements. However, Kleinberg et al. [KKV*07] show the limitations of anthropometry as a facial identification technique, when comparing a crime suspect with that of a potential offender, since people can share the same anthropometric data. Caliper-based measurements may be good for finding the distance between landmarks, but they do not completely describe the complex topology of a face.

3.2. Evaluation of Candidates

The fitness of an individual is determined by a function that gives a score according to how well the individual solves the problem at hand. Our main criterion is that we desire unique models in the population. If they are too similar to another model they may be perceived as a clone, although what constitutes similarity when measured objectively may not necessarily be the same as that perceived by humans as Santini and Jain [SJ99] observe.

In order to be unique within a population there must be some form of distance measure between population members, and a way to place a model meaningfully in feature space. Ip and Wong [IW02] present a normals-based profile for 3D head models, which is used for database storage and retrieval.

To illustrate the scenario of similar models being produced, we generate ten using FaceGen Modeller 3.3 [Faca] to represent a mating pool. We randomly generate groups of five integers, e.g. [9,5,2,2,1] where each number is a model from the mating pool. Blend weights for each model are fixed at: 0.3; 0.15; 0.1; 0.15; and 0.3 respectively. Figure 3 shows three models that have been generated, using what is effectively a multi-parent (five) approach. To a human, at first glance, the models may look quite similar, and perhaps be considered clones.

A more subtle difference can be observed when we view the same models in profile, as seen in Figure 4. We consider, as do Bailenson and Beall [BB06], the question: “How different does a human (or Virtual Human) have to be before being considered unique?” Our current response to this, is to use an arbitrary threshold, based on the empirical observations of the authors, for the maximum percentage of normals permissible for two models to share, so as to not be considered too similar.




Figure 3: *Left: previously unique population member. Middle: shares 60% normals with Left. Right: shares 71% normals with Left.*



Figure 4: *Lateral view of models in Figure 3, showing subtle differences.*

4. Conclusions and Further Work

We have presented Evolutionary Algorithms as a powerful approach to optimisation and search, and suggested that they have a contribution to make in the generation of Virtual Humans, provided that an appropriate representation scheme is chosen. Our method differs from previous reported work in that we seek to use EAs to automatically generate unique and plausible Virtual Humans within a population.

We continue to explore the chromosome representation, and the genotype-to-phenotype mapping, which is so important for a successful EA. In addition, there are many other options when setting up the various parameters of an EA, such as the selection mechanisms, and the variation operators of recombination and mutation. This offers considerable opportunities for exploring which settings give the most effective results.

References

- [ABHS05] ALBRECHT I., BLANZ V., HABER J., SEIDEL H. P.: Creating face models from vague mental images. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Sketches* (New York, NY, USA, 2005), ACM.
- [BB06] BAIENSON J., BEALL A.: Transformed social interaction: Exploring the digital plasticity of avatars. In *Avatars at Work and Play*. Springer, 2006, pp. 1–16.
- [BRM08] BASTIONI M., RE S., MISRA S.: Ideas and methods for modeling 3d human figures: the principal algorithms used by makehuman and their implementation in a new approach to parametric modeling. In *Compute '08: Proceedings of the 1st Bangalore annual Compute conference* (New York, NY, USA, 2008), ACM, pp. 1–6.
- [Daw06] DAWKINS R.: *The Blind Watchmaker*. Penguin, April 2006.
- [Dip02] DIPAOLA S.: Facespace: a facial spatial-domain toolkit. In *Information Visualisation, 2002. Proceedings. Sixth International Conference on* (2002), pp. 105–109.
- [Faca] Facegen modeller, <http://facegen.com>.
- [Facb] Facelift, <http://www.dipaola.org/steve/facelift.html>.
- [FHC04] FROWD C. D., HANCOCK P. J. B., CARSON D.: Evofit: A holistic, evolutionary facial imaging technique for creating composites. *ACM Trans. Appl. Percept.* 1, 1 (2004), 19–39.
- [FS99] FUJIWARA Y., SAWAI H.: Evolutionary computation applied to mesh optimization of a 3-d facial image. *Evolutionary Computation, IEEE Transactions on* 3, 2 (1999), 113–123.
- [Gol89] GOLDBERG D. E.: *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley Professional, January 1989.
- [IW02] IP H. H., WONG W. Y.: 3d head models retrieval based on hierarchical facial region similarity. In *In Proceedings of the 15th International Conference on Vision Interface* (2002), pp. 314–319.
- [Jef08] JEFFERY C.: Using non-player characters as tutors in virtual environments. In *Researching Learning in Virtual Environments* (November 2008), pp. 181–188.
- [Kan] Kaneva, <http://kaneva.com>.
- [KKV*07] KLEINBERG, KRISTA F., VANEZIS, PETER, BURTON, MIKE A.: Failure of anthropometry as a facial identification technique using high-quality photographs. *Journal of Forensic Sciences* 52, 4 (July 2007), 779–783.
- [Mak] Makehuman, <http://makehuman.org>.
- [Mit98] MITCHELL M.: *An Introduction to Genetic Algorithms (Complex Adaptive Systems)*. The MIT Press, February 1998.
- [MTT04] MAGNENAT-THALMANN N., THALMANN D. (Eds.): *Handbook of Virtual Humans*. Wiley, August 2004.
- [Par72] PARKE F. I.: Computer generated animation of faces. In *ACM'72: Proceedings of the ACM annual conference* (New York, NY, USA, 1972), ACM, pp. 451–457.
- [RVCN08] ROBERTO, VIDAL C. A., CAVALCANTE-NETO J. B.: A biologically inspired generation of virtual characters. In *SAC '08: Proceedings of the 2008 ACM symposium on Applied computing* (New York, NY, USA, 2008), ACM, pp. 1218–1224.
- [Sec] Second life, <http://secondlife.com/>.
- [Sim] The sims, <http://thesims.ea.com/>.
- [SJ99] SANTINI S., JAIN R.: Similarity measures. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 21, 9 (1999), 871–883.
- [Ven00] VENTRELLA J.: Avatar physics and genetics. In *VW '00: Proceedings of the Second International Conference on Virtual Worlds* (London, UK, 2000), Springer-Verlag, pp. 107–118.
- [VHW03] VENTRELLA J. J., HARVEY W. D., WERTHER B. M.: Us patent 6545682: Method and apparatus for creating and customizing avatars using genetic paradigm, April 2003.