Efficient Clothing Fitting from Data

Marco Gillies
University College London
Ross Building pp1
Adastral Park
Ipswich, IP5 3RE, UK

m.gillies@ucl.ac.uk

Daniel Ballin
BT Exact
Ross Building pp4
Adastral Park
Ipswich, IP5 3RE, UK

daniel.ballin@bt.com

Balázs Csanád Csáji
Department of General Computer Science

Faculty of Science

Eötvös Loránd University

1117, Budapest, Pázmány Péter sétány 1, Hungary

csaji@sztaki.hu

ABSTRACT

A major drawback of shopping for clothes on-line is that the customer cannot try on clothes and see if they fit or suit them. One solution is to display clothing on an avatar, a 3D graphical model of the customer. However the normal technique for modeling clothing in computer graphics, cloth dynamics, suffers from being too processor intensive and is not practical for real time applications. Hence, retailers normally rely on a fixed set of body models to which clothes are pre-fitted. As the customer has to choose from this limited set the fit is typically not very representative of how the real clothes will fit. We propose a method that uses a compromise between these two methods. We generate a set of example avatars by performing Principal Component Analysis on a dataset of avatars. Clothes are pre-fitted to these examples off-line. Instead of asking the customer to choose from the set of examples we are able to represent the users avatar as a weighted sum of the examples, we then fit clothes as the same weighted sum over the clothes fitted to the examples.

Keywords
Virtual Clothing, E-Commerce, Data driven techniques

1.INTRODUCTION

On-line shopping is becoming increasingly popular both with customers and retailers. Customers have a convenient time saving experience while retailers can save costs with a largely automated sales procedure. Clothes retailing is one of the primary retail areas (worth about £30billion (€42billion) in the UK) but there is a significant disadvantage to selling clothes on-line. It is very important for customers to be able to try on clothes to get an idea of how the clothes will fit and suit them. This discourages many people from shopping for clothes on-line. In order for more people to adopt e-commerce for clothing some sort of alternative to trying clothes on in person is needed. One alternative proposed is to provide users with a 3D graphical model of themselves (called an avatar)

and display the selected clothes on this avatar. If it is an accurate enough a portrayal of the customer, the graphical fitting should give a reasonable idea of how the clothes would look in real life. This use of avatars for clothing retail needs a method for fitting clothing to avatars. The method must produce an accurate representation of how the clothes will look on the customer and do so without excessive processing requirements that would overload the servers or slow down the customer experience.

There are a number of existing methods for fitting clothing to human body models, which we will now discuss. Each of them has advantages and disadvantages in terms of creating personalized clothing in on-line applications.

Hand fitting
For the highest quality a skilled 3D modeler may fit the clothing to the avatar by hand. This is likely to produce a high quality result that makes the clothes look particularly flattering. Of course this method requires extensive human intervention and is not possible for an on-line system.

Cloth dynamics

The leading method for modeling clothes in computer graphics applications is by using physical modeling or cloth dynamics. The physical properties of the clothes are simulated numerically to generate the shape of the clothes. This has traditionally been the main area of research into clothes modeling and there has been a very large body of work of which we can only give a brief outline (Ng and Grimsdale [Ng96] give a fuller, though slightly dated, overview). Terzopoulos et al [Ter87] pioneered work in cloth dynamics. Other notable contributions include Breen, House and Wozny [Bre94], Provot [Pro95], Eischen, Deng and Clapp [Eis96], Baraff and Witkin [Bar98] and Vassilev and Spanlang [Vas00]. MIRALab has a particularly sustained and long lasting research programme starting with pioneering work by Carignan, Yang, Magnenat-Thalmann and Thalmann [Car92] to complete systems such as that by Volani and Magnenat-Thalmann [Vol00], among many others. Though the details of these various approaches differ, the essence consists of creating a discreet model of the clothing, commonly as a set of particles connected by springs as per Breen et al [Bre94] or Provot [Pro95]. This discrete model is simulated under a number of forces such as gravity, wind, and contact forces to determine the final positions of the elements. These positions are then used to determine the final shape of the cloth. Major research problems include the numerical simulation of the clothing properties including the work of Baraff and Witkin [Bar98] and Choi and Ko [Cho02] and efficiently determining and dealing with collisions between cloth and objects including the wearer and itself for example, Vassilev, Spanlang, and Chrysanthou [Vas01], Bridson, Fedkiw and Anderson [Bri02], and Baraff, Witkin and Kass [Bar03].

There has also been work that focuses specifically on fitting clothing to an avatar or body model, Vassilev and Spanlang [Vas00] and Volani and Magnenat-Thalmann [Vol00] deal with this issue. They both use flat pieces of cloth (similar to those used to make up real clothing) that are attracted to each other by seaming forces. Cordier, Seo and Magnenat-Thalmann present a system for on-line clothes shopping [Cor03].

Cloth dynamics methods are able to produce good results and realistic clothing, normally totally automatically. Their main disadvantage is that they are computationally intensive and thus unsuitable for on-line, interactive applications.

Pre-fitted examples

The above methods are not suitable for on-line use so clothing retailers must use alternatives. One possibility is to provide the user with a choice from a limited set of avatar body shapes and to then have a pre-fitted clothing model for each of these body shapes. It is difficult to know the exact method used by commercial systems, but "my virtual model" (www.myvirtualmodel.com) seems to take this approach. As the fitting is done off-line it can be of high quality but at the cost of providing very little customization to the individual user.

Texturing and geometric methods

A simple and cheap approach is to texture map the clothes onto the avatar model. A slightly more complex method is to do a number of scaling operations onto an existing clothing model to produce a rough fit. These methods are cheap and simple. The texturing method seems to be in use in most on-line multi-user environments and games. These methods do not take into account any properties of cloth and therefore the results are often poor and do not give a good impression of how clothes will fit in real life; they generally only work with tight fitting clothes. There have also been various geometric techniques that are used to help designers fit clothing to body models in graphical editors. These include Hinds and McCartney [Hin90], Ng, Grimsdale and Allen [Ng95] and Igarashi and Hughes [Iga03].

There are therefore difficulties with each of the existing methods. This paper describes a method that promises to provide a good compromise between the high quality results of cloth dynamics with lower computational cost. It approximates either a cloth dynamics solution or hand fitting at low on-line cost. It is in some way similar to pre-fitted methods in that it relies on clothing fitted off-line to a limited set of avatars. However, rather than using these examples directly they are used to produce an individual fitted item of clothing for each user avatar.

2. AN APPROXIMATION METHOD

As described in the previous section, a number of methods exist to accurately fit clothing to an avatar. Unfortunately these methods are not suitable for real-time clothing fitting. We propose a new method that generates a fitted garment based on two pre-generated datasets, a dataset of clothing that fits a corresponding dataset of avatars. The clothing dataset is calculated off-line using one of high quality techniques discussed previously, such as hand fitting or cloth dynamics.

Our method is composed of three fundamental steps:

1. Generating the dataset of example avatars

2. Generating a dataset of clothes by fitting a cloth model to the dataset of avatars.

3. Using both of these datasets to generate an item of clothing that fits a new avatar.
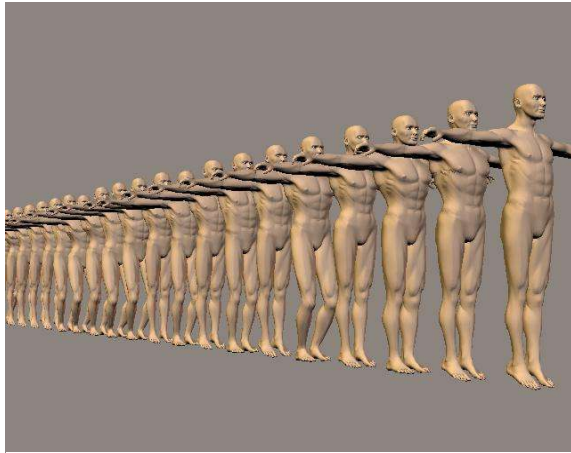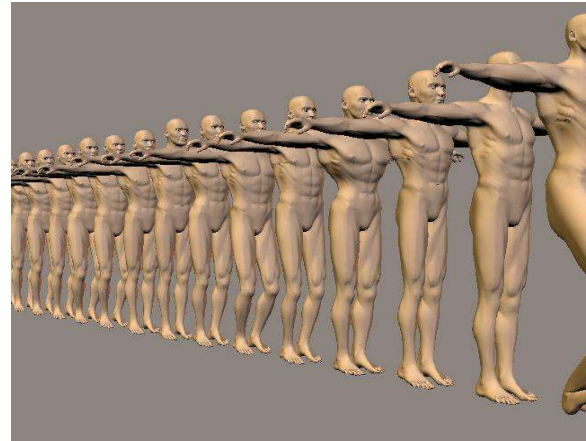
Figure 1. A data set of example avatars

Figure 2. Principal components of the data set

Generating a set of examples

The choice of an appropriate set of example avatars and a method of combining these examples to produce the new clothes is critical to the approximation method. For the method to work efficiently there are a number of properties that the set of examples should have:

- The examples should be representative of the real people that need clothes fitting. To achieve this we derive our examples from a dataset of avatars.

- The examples should provide a good coverage of the space of possible avatars.

- The operation to derive a new avatar from the examples, and therefore new clothes from the clothing examples should be computationally efficient. Ideally a new avatar or set of clothes should be simply a linear combination of the examples.

- Given a new avatar it should be simple to express it in terms of a combination of the examples. If the examples form an orthonormal basis, this becomes a simple matter of projecting the new avatar onto the basis.

- The number of examples should be as small as possible as the off-line fitting operation is relatively expensive.

Thus we need to produce a minimal, orthonormal basis from a dataset. The standard method for doing this is Principal Component Analysis (PCA) [Jol86]. Blanz and Vetter [Bla99] used PCA to form an orthonormal basis to represent a large dataset of heads. Allen Curless and Popović [All03] have used similar methods for human bodies.

Our method consists of two off-line preparation steps and two on-line fitting steps:

1. Perform a principal component analysis on a dataset of avatars

2. Fit the clothing items to each of the avatar dataset to form a parallel set of clothing.

3. Form a representation of the user avatar as weights over the principal components

4. Use the same weights to reconstruct clothes from the set of clothing items

We will now describe the process in detail.

Principal component analysis

To perform PCA the avatars must be represented as vectors. This is simple enough as the avatars are simply meshes consisting of a set of vertices each of which consists of an X, Y and Z component. The mesh can be represented as the vector $\{X_0, Y_0, Z_0, X_1, Y_1, Z_1, \ldots X_n, Y_n, Z_n\}$. These vectors must all be the same length, and for the results to be meaningful each element must be equivalent across the avatars. This means that the avatars in the dataset must all have the same mesh topology and than the i^{th} vertex must correspond to the same part of the body in all the avatars. For some datasets this constraint holds, for example, if a dataset is obtained by scanning real people using a method that conforms a single original mesh the constraint is likely to hold. To simplify our method we have used a set for which the constraint already holds. However, most sets do not have a fixed topology so a method is needed to convert all the avatars in the set to have the same topology. Luckily methods exist to do this alignment, for example those described in Dryden and Mardia [Dry98] and Allen, Curless and Popović [All03].

For our implementation we did not have access to a real dataset and therefore used an artificial dataset. We used Curious Labs Poser® to randomly transform a base avatar using a set of transforms designed to maintain the realism of the avatar. Since all the avatars are transformations of the same original mesh the alignment property automatically holds. The dataset of transformed avatars is shown in figure 1.
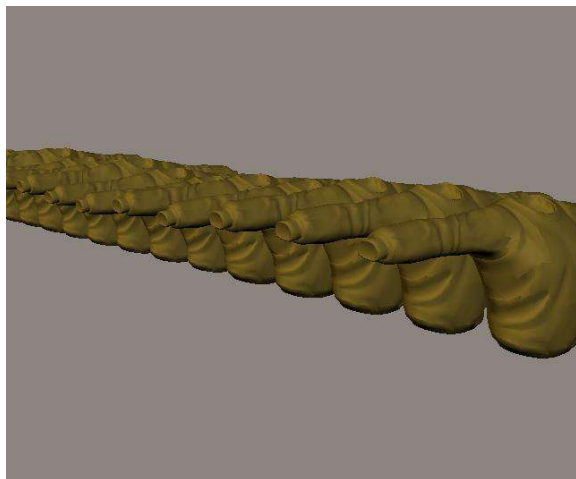
Figure 3. A set of clothes corresponding to the Principal components

Principal component analysis itself is described in detail in Jolliffe [Jol86]; we will give a brief overview. Firstly the mean of the set is calculated and subtracting this from the original set produces a set of differences:

$$\mathbf{x}' = \mathbf{x} - \bar{\mathbf{x}}$$

The covariance matrix of these differences is then calculated. The eigenvectors of the covariance matrix are then calculated. The eigenvectors (or principal components) form an orthonormal basis that can exactly represent any of the original avatars (to be exact they can represent the set of differences, the mean avatar must be added to generate the original avatars). The eigenvectors are ordered in terms of the amount of variation they account for in the data set. The first eigenvector is the axis with the greatest variation in the dataset; the second is the largest proportion of the remaining variation, etc. This ordering is given by the eigenvalues, each eigenvalue is proportional to the percentage variation represented by the eigenvector, and therefore the numerical ordering of the eigenvalues gives the ordering of the eigenvectors. The sum of a subset of the eigenvalues divided by the total sum of all eigenvalues can also be used to give the percentage of variation accounted for by the corresponding sub-set of the eigenvectors. Thus if a subset $\mathbf{X}$ of the eigenvalues sums to 95% of the total then the corresponding subset of principal components can reproduce the original data set to within 95% accuracy. This ordering of the eigenvalues can be used to discard the least important components, leaving a subset that account for a certain percentage of variation e.g. 95%. Thus Principal Component Analysis is a useful technique for generating an orthonormal basis from a dataset and for reducing its dimensionality.

Performing a weighted sum over the principal components and adding the mean avatar can construct new avatars. The avatar corresponding to a principal component is created by adding that component to the mean avatar.

Thus we can use the set of principal components we have generated as examples on which to base our approximate clothes fitting. These are shown in figure 2.

Fitting clothes to the examples

Once an orthogonal space of example avatars has been generated we must clothe these avatars so that they can be used in our approximation process. Once the mean avatar has been added to a principal component it is itself a valid avatar and so can be clothed by any method that can be used for a normal avatar. This part of the process is performed off-line and so can use an expensive method that gives good quality results, such as a cloth dynamics method or hand fitting. In our implementation we fitted the items of clothing by hand.

Fitting clothes to each principal component gives a new set of clothes. We take the mean of these clothes (again represented as a vector), subtract it from each item and normalize the result. This gives a new, normalized basis that is parallel to the original basis of avatars in the sense that each dimension on the first corresponds to a dimension of the second. The set is not orthogonal but this is not an important property as there is no need to project onto this space. Figure 3 gives an example set of clothing bases. If an avatar is created from a principal component by adding it to the mean; then the clothing created from the corresponding clothing base fits the avatar.

Real time cloth fitting

The preceding method gives a set of example avatars and examples of clothes fitted to those avatars. These examples are used to fit clothes to new avatars. These two databases are used to form an approximation of the original fitting method for the new avatar. The essence of the method is shown in figure 5. The set of clothing and the set of avatar principal components are both linear set, so a new avatar or item of clothing can be produced by a weighted sum over the set (and the addition of the mean). Thus applying a new set of weights to the principal components can generate a new avatar. More importantly the items in the clothing set each correspond to one of the avatar principal components. This means that we can apply the same set of weights that were used to create the avatar to the clothing set and obtain an item of clothing that approximately fits the new avatar. If the set of avatars is a good span over the set of possible avatars then the approximation will be a good one.

Of course we do not want to fit clothes to a new avatar generated from the principal components, we
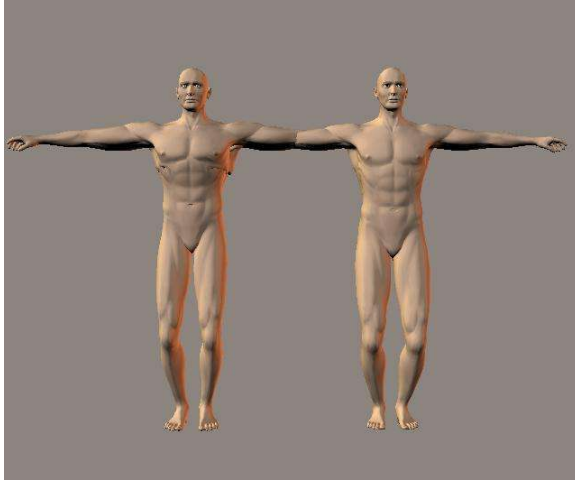
Figure 4. An example of a new avatar (left) reconstructed using principal components
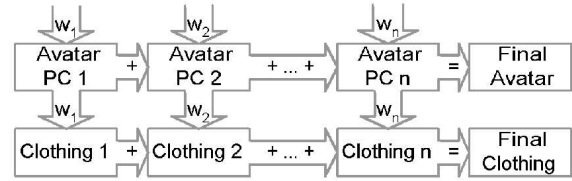

Figure 5. Using parallel sets of clothing and avatar. Applying the same weights to the set of avatars and to the set of clothes results in a new avatar and an item of clothing that fits that avatar.

want to fit it to an avatar of a real user. This means we must express the user avatar as a weighted sum over the principal components. As the principal components are an orthonormal base this is simply a matter of subtracting the mean avatar from the user avatar and performing a dot product of the result with each principal component. How well the principal components approximate the original avatar depends on how well they span the space of avatars and therefore on how representative the original dataset of avatars was. Figure 4 shows an example of the reconstruction of an avatar from principal components. Once the avatar has been represented as a set of weights over the principal components, applying the same weights to the clothing set generates an item of clothing for that avatar.

If there are N principal components each with M vertices, projecting the avatar onto the PCs consists of N dot products which result in 3MN multiplications and (3M-1)N additions. The projection only has to be calculated once for each new user avatar, and many types of clothing can then be fitted to it. The cloth fitting itself consists of 3M dot products of length N vectors so again takes 3MN multiplications and 3M(N-1) additions. This is fairly efficient compared to other cloth fitting methods and can clearly be improved by reducing N or M. Reducing N (the number of principal components) also has the benefit that there are fewer examples to fit clothes to in the off-line preprocessing step. To reduce N you can decrease the degrees of freedom of the dataset, for example, avatar models normally have a large number of vertices (and therefore degrees of freedom) in the head but this is largely irrelevant to the fitting process so they can be removed from models before the PCA process. Performing the fitting using a reduced mesh and then mapping the changes onto a higher resolution final

mesh could reduce M. We are yet to implement these possible improvements.

Figure 6 gives examples of clothes fitted to assorted avatars.

3. THE ONLINE CLOTH FITTING PROCESS

This section describes how our method might be used in a real application. As our method is based on approximation from data it is vital to base an application on good datasets. The primary development of this work should be obtaining data. There are two main datasets that are needed, avatars and clothes.

There are two ways in which to create an avatar and therefore a dataset of avatars. The first is to scan a users' body using either a laser scanning or photometric technique. This method potentially provides a way of obtaining highly accurate models of an individual (though many current methods are less accurate). Unfortunately the process of creating an avatar is complex and costly. It involves the user coming to a special, probably expensive scanning device, thus creating a barrier to entry for new users of an avatar clothes-shopping system. It is also a challenge to create a representative dataset of avatars, as a large number of people must be scanned, though existing databases such as the Civilian American and European Surface Anthropometry Resource Project (CAESAR) could be used. The other approach is to use a standard avatar that is deformed based on measurements of the user. This is analogous to traditional tailors' techniques where a small number of measurements of a user are taken and used to accurately fit garments. In fact the standard tailors' measurements can be used as the basis of this technique. This method is less accurate and it does not capture other aspects of appearance like face, hairstyle or skin colour. However, it is much simpler for users, they can measure themselves at home and enter the measurements in a web page to generate an avatar. If a facial photograph is also provided other aspects of appearance can also be mapped onto the avatar. Obtaining a dataset from this method involves measuring a large number of users but requires less equipment. This method can also re-use existing

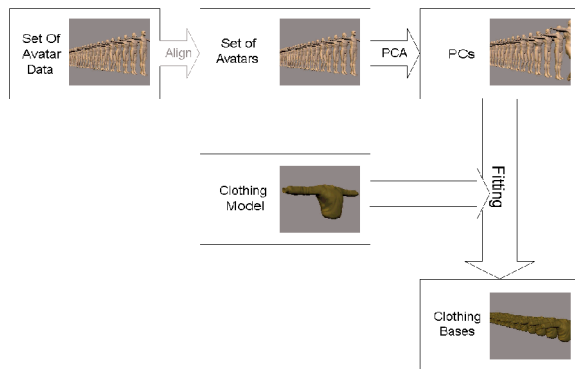**Figure 6. Examples of clothes fitted to various avatars**

Figure 7. The of-line preparation component of the process

The clothing models should clearly be based on the retailer's catalogue, ideally based on the original design data. The models can be fitted to the examples using an automated cloth dynamics system. This has the advantage over hand fitting the clothes that the process can be performed without human intervention, which would be a significant saving if the cloth fitting needs to be done for a large number of example avatars. However, retailers might prefer hand fitting if it gives a more flattering representation of their clothes.

Figure 6 shows the steps that need to be performed as an offline preparation for the cloth application. The dataset of avatars needs to be obtained as described above. Depending on how it was generated, an alignment step might be needed to ensure point-to-point correspondences between the avatars. Principal component analysis is then performed to obtain a set of example avatars. The clothing model is then fitted to each of the examples to obtain a basis set of clothes.

Once the above datasets have been created they can be used to allow customers to virtually try on the clothes in the e-commerce system. Figure 7 gives an overview of how such a system might work. The customer must first create an avatar of themselves, this is presumably done using the same method as for creating the original dataset of avatars, either by scanning with a system such as the AvatarBT scanning booth [Bal00], or by entering their measurements into a web page. The retailer could provide a number of different options. Ideally a common avatar format would be created and shared between retailers. In order to fit clothes to the avatar it needs to be represented as a set of weights over the principal components. This is done via the projection step described in section 2. The projection step only needs to be performed once for each customer and can be done when the customer first registers with the retailer. The most likely method is for the user to send the avatar to a server that then returns the weights to the user's computer (this is likely to be the
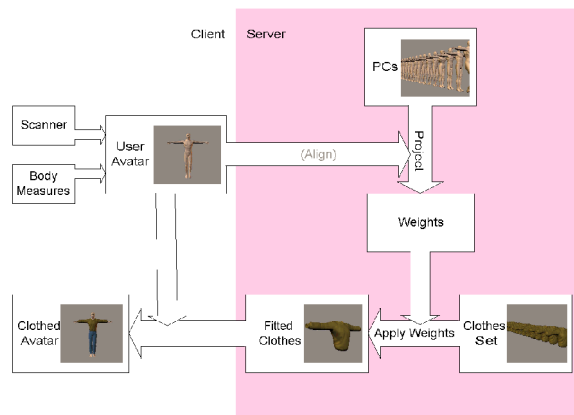

Figure 8. The online component of the process, fitting clothes to a new user avatar

most expensive step that must be performed while the user is waiting). Once the user has their avatar represented in terms of the principal components they can then try on any item of clothing in the retailer's database. The weights are transmitted to the server where they are applied to the clothing dataset. This generates a new clothing mesh specific to the user. This can then be transmitted back to the user and displayed on the user's avatar, giving a good impression of how the clothes will fit the user in real life. The fitting process is less computationally expensive than the multiple force integrations required for cloth dynamics and therefore providing relatively low load on the server with far more personalized results than a method that uses only a standard set of clothing models. In terms of user experience the fitting process is unlikely to take a noticeable time over that needed to transmit the final clothing model (which is required whatever method is used). Thus our method provides an efficient way to provide personalized clothing for a customer in an e-commerce system.

4. CONCLUSION AND FURTHER WORK

We have presented an alternative to current cloth fitting methods that promises quality similar to cloth dynamics methods (to which it is an approximation) with less of a computational overhead. We have presented the fundamental algorithm but there remains considerable work on the e-commerce process. The most important problem is obtaining datasets of avatars and clothes, in particular converting the clothing models used by manufacturers into a 3D mesh format that we require. A client/server web application must also be created in order to use this method. We would like to form partnerships with clothing retailers to investigate these practical issues.

5.ACKNOWLEDGMENTS

6.REFERENCES

[All03] Allen, B., Curless, B., and Popović, Z. The space of human body shapes: reconstruction and parameterization from range scans. proceedings of SIGGRAPH pp 587-594 2003

[Bal01] Ballin, D., Lawson, M., Crampton, S., Child, T., and Hilton, A. Practical, Reliable, Repeatable: Scanning Avatars in the Millennium Dome. Proceedings of SCANNING 2001 Congress, Paris, April 2001

[Bar98] Baraff, D., and Witkin, A. Large steps in cloth simulation, proceedings of SIGGRAPH pp 43-54 1998.

[Bar03] Baraff, D., Witkin, A., and Kass, M. Untangling cloth, proceedings of SIGGRAPH pp 862-870 2003.

[Bla99] Blanz, V. , and Vetter, T. A morphable model for the synthesis of 3D faces. proceedings SIGGRAPH pp187-194 1999

[Bre94] Breen, D. E., House, D. H., and Wozhny, M.J. Predicting the drape of woven cloth using interactive particles. proceedings of SIGGRAPH pp23-34 1994.

[Bri02] Bridson, R., Fedkiw, R., and Anderson, J. Robust treatment of collision, contact and friction for cloth animation. proceedings of SIGGRAPH pp.594-603 2003

[Car92] Carignan, M., Yang, Y., Magnenat-Thalmann, N., and Thalmann, D. Dressing Animated Synthetic Actors with Complex Deformable Clothes. proceedings of SIGGRAPH pp99-104 1992

[Cho02] Choi, K-J., and Ko H-S. Stable but responsive cloth. proceedings of SIGGRAPH pp604-611 2002

[Cor03] Cordier, F., Seo, H., and Magnenat-Thalmann, N. Made-to-measure technologies for online clothing store. IEEE Computer Graphics and Applications, pp38-48 2003

[Dry98] Dryden, I. L. and Mardia, K. V. Statistical Shape Analysis. Wiley, 1998

[Eis96] Eischen, J.W., Deng, S., and Clapp T.G., Finite-element modelling and control of flexible fabric parts, Computer Graphics in Textiles and Apparel (IEEE Computer Graphics and Applications) pp71-80 1996

[Hin90] Hinds, B.K., ,and McCartney, J. Interactive garment design. The Visual Computer, Springer-Verlag Vol. 6, pp.53-61 1990

[Iga03] Igarashi, T., and Hughes, J.F. Clothing manipulation. proceedings of the 15th Annual symposium on User Interface Software and Technology, Paris, France. pp. 91-100 2002

[Jol86] Jolliffe, I.T. Principal Component Analysis. Springer-Verlag 1986

[Ng95] Ng, H.N., Grimsdale, R.L., and Allen, W.G. A system for modeling and visualization of cloth materials. Computers and Graphics. Pergamon Press/Elsevier Sicence, Vol. 19 no. 3 pp. 423-430 1995

[Ng96] Ng, H.N., and Grimsdate, R. L. Computer graphics techniques for modeling cloth. IEEE Computer Graphics and Applications, Vol 16, pp28-41, 1996.

[Pro95] Provot, X. Deformation constraints in a mass spring model to describe rigid cloth behaviour. Proceedings of Graphics Interface, pp141-155 1995

[Rui02] Ruiz, M.C. Buxton, B. Douros, Y. Treleaven, P. Web Based Software Tools for 3D Body Database Access and Shape Analysis , Proceedings of Scanning 2002, Paris, April 2002.

[Siz] UK National Sizing Survey (SizeUK Project): http://www.sizeuk.org

[Ter87] Terzopoulos, D., Platt, J., Barr, A., and Fleischer, K. Elastically Deformable Bodies, proceedings of SIGGRAPH, pp205-214 1987.

[Vas00] Vassilev, T., and Spanlang, B. Efficient cloth model for dressing animated virtual people proceedings of Learning to Behave Workshop, Enschede, the Netherlands pp 89-100 2000

[Vas01] Vassilev, T., Spanlang, B., and Chrysanthou, Y. Efficient cloth model and collision detection for dressing virtual people proceedings of GeTech, Hong Kong 2001

[Vol00] Volani, P., and Magnenat-Thalmann, N. Virtual Clothing: Theory and Practice, Springer-Verlag 2000.