Reverse Engineering Garments

Nils Hasler, Bodo Rosenhahn, and Hans-Peter Seidel

Max-Planck-Institut für Informatik, 66123 Saarbrücken, Germany,
hasler@mpi-inf.mpg.de

Abstract. Segmenting garments from humanoid meshes or point clouds is a challenging problem with applications in the textile industry and in model based motion capturing. In this work we present a physically based template-matching technique for the automatic extraction of garment dimensions from 3D meshes or point clouds of dressed humans. The successfull identification of garment dimensions also allows the semantic segmentation of the mesh into naked and dressed parts.

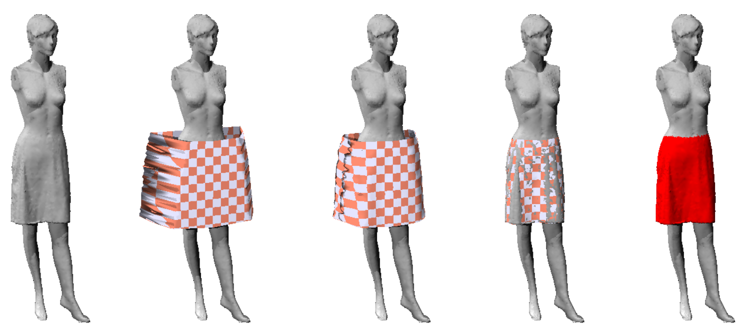

Fig. 1. The laser scan of a dressed person is segmented into dressed and naked parts by fitting the parameterised template of a garment to the mesh.

1 Introduction

Segmenting point clouds or meshes of dressed humans into covered and naked parts is a problem with applications e. g. in the textile and movie industry. A number of general mesh segmentation algorithms have been published [1] but they cannot be applied to the given problem because they provide a segmentation but are unable to supply labels for the segmented parts. So a different approach is suggested here. Like Jojic and Huang [2] we fit a simulated cloth to the 3D data to recreate the observed configuration. Segmenting the mesh or point cloud is then straight forward.

Another application of the approach is to provide initialisation informations for a model-based motion capture algorithm of loosely dressed humans such as [3]. For this approach it is essential that a good model of both the tracked person and of the loose garments worn during the motion capture session is available. Since the model of the tracked person can easily be acquired using a full-body laser scanner, we here deal with the more challenging problem of extracting the dimensions of the attire from the 3D scan of a person. Additionally, extracting the dimensions of the clothing allows the semantic segmentation of the 3D data into covered and unclad parts.

The approach proposed here is a template based analysis-by-synthesis technique [4]. That is we assume a priori knowledge of the general type of garments worn but not their exact dimensions. These dimensions of the parameterised garments are chosen by a simulated annealing algorithm [5] and the resulting attire is draped over the mesh or point cloud. The evaluation function of the procedure is based on the distance from the simulated cloth to the measured 3D data, the stretch of the springs, and the silhouettes of model and simulated garment.

The remainder of this work is structured as follows: Section 2 introduces publications related to the present work. Section 3 presents details about the employed draping procedure. Results are presented in Section 4 and a summary is given in Section 5.

2 Previous Work

Even though the algorithm presented here is aimed at segmenting a mesh into parts describing the garments the scanned human is wearing and those that belong to the person itself, the approach cannot readily be compared with general mesh segmentation algorithms such as [1]. These approaches are after all only aimed at segmenting arbitrary shapes into smaller parts that do not necessarily have any intuitive semantics attached to them. Even if they do these cannot be assigned automatically to the segments which defeats their applicability to the given problem.

Other approaches were specially adapted to segmenting human shapes. Nurre [6] or Xiao et al. [7] are able to reliably segment a human body scan into torso and limbs and can label the parts automatically. However, the approaches were not designed to work with dressed humans.

Allen et al. [8] introduced a different class of algorithms that can be used to segment the scan of a human. They proposed to fit a generic human model to the observed mesh by deforming it locally. That is an affine transformation is applied to every point of the mesh. Local smoothness of the transformations, the distance between the two models, and the distance of manually selected landmarks are used to evaluate a configuration. They also show that reasonable results can be obtained if no landmarks or even if landmarks but no further surface information is available. Since the template model may contain arbitrary additional information, segmentation or animation of the model becomes trivial.

Seo and Magnenat-Thalmann [9] presented a similar approach, which operates in two phases. They first roughly find the pose and scale of a skeleton based template model and then iteratively fit the surface of the template model to the measured 3D data. Unfortunately, their rough fitting stage requires the manual localisation of landmarks.

A recent survey of these and related human model generation and animation techniques was presented by Magnenat-Thalmann et al. [10].

Unfortunately, all the above techniques work well only for tightly or scantily dressed models. Going the other way, Oh et al. [11] propose a technique for creating animatable models of dressed humans that rely on simulating the garments only during model generation. They proceed by first draping the garments over the undressed model. Then they remove hidden surfaces and transfer skinning weights from the underlying model to the cloth surface. During animation the model can then be deformed based on simple skinning weights attached to a skeleton.

With a completely different goal in mind but arguably the work most closely related to our own Jojic and Huang [2] presented an approach for estimating the static parameters of a square piece of fabric from range data. Their algorithm follows a two-phased pattern. In the first phase a simulated piece of cloth is attracted to the triangulated range data of an object which is covered by a square textile. They do this to ensure that the simulation converges to the same minimum energy configuration as the real-world fabric did when it was draped over the real object. In the second phase the attraction forces are turned off and the simulation is again run until it comes to rest. The mean distance between the simulated cloth and the range data is used to tune the static cloth properties of their simulation.

Since the approach that is proposed here involves the simulation of a piece of fabric, a brief overview of the particle based cloth simulation literature is given in the following. Physically based cloth simulation was pioneered by Terzopoulos et al. [12]. In their work a number of techniques that are still common today such as semi-implicit integration, hierarchical bounding boxes, and adaptive time-step control were proposed. Until Baraff and Witkin reintroduced semi-implicit integration [13], decreasing the computational cost of cloth simulation significantly, explicit integration techniques were common.

In the last few years two major strands of development can be made out in the cloth simulation community. One, aiming for real-time simulation, focusses on computation speed alone, sacrificing expressiveness and accuracy of the employed model if necessary. Desbrun et al. simplified the equation system that needs to be solved every step by precomputing parts of it [14]. Kang and Choi used a coarse mass-spring discretisation and added wrinkles in a post-processing step by interpolating with a cubic spline [15].

The other strand attempts to simulate cloth as realistically as possible. The use of nonlinear cloth properties has been introduced by Eberhardt et al. [16]. Simplified nonlinearities have since been integrated into a number of systems such as [17, 18]. Impressive results have been presented by Volino and Magnenat-Thalmann [19]. The fabric properties employed in their system are not only nonlinear but even exhibit hysteretic behaviour.

3 Algorithm

In this section the template deformation approach for estimating garment dimensions is detailed. We first give a rough overview of the approach and then provide details for the different components of the system.

The method for segmenting a 3D scan into dressed and naked parts proposed here can be seen as a physically based template deformation approach. That is we assume a priori knowledge of the type of garment the scanned person is wearing but not the exact measurements of the attire. A physically based cloth simulation system is used to drape the garment template over the captured mesh. The resulting resting position of the garment is evaluated using the distance to the captured mesh, the stretch of the springs employed in the cloth simulation, and the silhouettes of 3D scan and simulated clothing. A simulated annealing algorithm then modifies the dimensions of the garments and updates the drape until convergence is reached.

3.1 Garment Templates

Two generic garment templates are shown in Figure 2. The skirt pattern has four degrees of freedom one vertical measurement (the length of the skirt) and three horizontal ones: One at the waistline which in this case actually sits on the upper hip, one at the hem and an additional measurement at the lower hip. Similarly the T-shirt template has a length, a width which is normally the same at all places for a simple T-shirt, sleeve length, and sleeve circumferences at the seam and at the hem. Unfortunately, we found that the T-shirt frequently rests so closely on the skin that it is impossible to extract the neckline from range data alone. So parameters defining the shape of the neckline were omitted from the optimisation procedure. As the parameters chosen by the optimisation procedure define the complete circumference of the garment the lengths have to be distributed to the different parts the garment is made of. The distribution of the parameters to the different parts of the pattern is also marked in Figure 2.

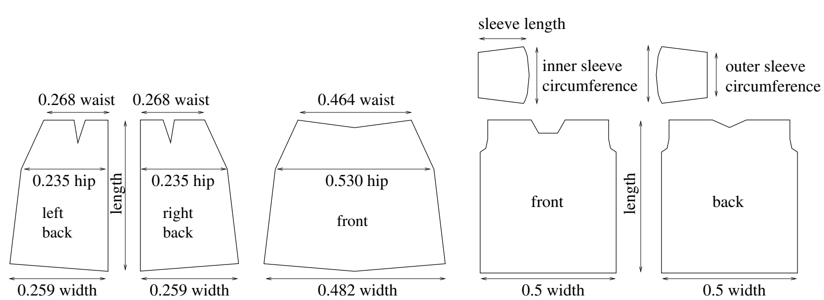

Fig. 2. Generic patterns of a skirt and a T-shirt

During optimisation the garment dimensions are modified in a way similar to the technique described by Volino et al. [20]. Only the restlengths of affected springs, the areas of deformed triangles and masses of involved particles are modified. Their 3D coordinates are left unchanged. That way invalid configurations, such as parts of the cloth extending into the body, are avoided and abrupt parameter changes are handled smoothly. The general procedure for applying dimension changes is shown in Figure 3.

The length of a T-shirt, for example, is changed by modifying the two-dimensional coordinates of affected points of the unrefined Delaunay triangulation. Particle positions of the refined triangulation are derived from their barycentric coordinates within this original triangulation.

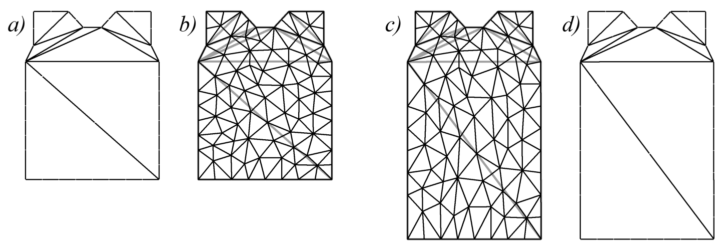

Fig. 3. When modifying the length of a T-shirt first the original underlying Delaunay triangulation gets changed (a, d). The coordinates of the refined triangles are then derived from their barycentric coordinates within this triangulation (b, c).

3.2 Cloth Simulation

The cloth simulation employed at the heart of the optimisation procedure is in essence a damped mass-spring system as used in most cloth simulations. The force $\mathbf{f}$ on two connected mass-points is given by

$$\mathbf{f} = \pm \left(k_s \frac{|\mathbf{\Delta x}| - l_0}{|\mathbf{\Delta x}|} + k_d \frac{\mathbf{\Delta x} \cdot \mathbf{\Delta v}}{|\mathbf{\Delta x}|^2} \right) \mathbf{\Delta x}$$

where $\mathbf{\Delta x}$ and $\mathbf{\Delta v}$ denote the differences between the two involved particles in position and speed respectively, while l_0 signifies the restlength of the connecting spring. The constants k_s and k_d specify stiffness and damping of the cloth. Additionally, bending resistance is implemented based on neighbouring triangles as proposed by Bridson et al. [21]. As we are only interested in obtaining the dimensions of the clothing and finding a segmentation of the model, the properties defining stretch and bending do not have to represent the properties of the captured fabric exactly. They are consequently not included in the optimisation.

One of the greatest problems encountered when employing physically based cloth simulations to real-world modelling such as tracking or in this case draping of garments is that cloth simulations are inherently instable [17]. That is slight deviations of the initial configuration or the simulation physics can cause significantly different outcomes. The only solution to the problem is to introduce attraction forces between the simulation and the observed configuration. So additional linear forces $\mathbf{f_a}$ act on all particles of the simulation.

$$\mathbf{f_a} = k_a \cdot d(\mathbf{p}, \mathbf{M})$$

Here k_a is the stiffness of the attraction force and $d(\mathbf{p}, \mathbf{M})$ is a function returning a vector pointing from particle $\mathbf{p}$ to the closest point of the model $\mathbf{M}$. If the model is a point cloud d always returns a vector pointing to one of $\mathbf{M}$'s points. If it is, however, a mesh any point on the surface of the mesh may be its target.

To speed up the computation of the closest point the k-DOP bounding box hierarchy, also used for collision detection, was engaged. Collisions are, however, completely disabled in our experiments because the attraction forces antagonise the collision resolution forces slowing the simulation significantly. In addition it was found that garment dimensions tended to get overestimated if collisions were enabled.

3.3 Optimisation

The optimisation procedure we use is a simulated annealing variant of the downhill simplex method as described by Press et al. [5]. The error $E(\mathbf{x})$ is a function of the parameter set $\mathbf{x}$.

$$E(\mathbf{x}) = \alpha \cdot E_{dist}(\mathbf{x}) + \beta \cdot \sqrt{|1 - E_{stretch}(\mathbf{x})|} + \gamma \cdot E_{sil}(\mathbf{x}) \tag{1}$$

It is a weighted sum of three evaluation functions. The first, $E_{dist}(\mathbf{x})$, describes the average distance of garment particles to the model.

$$E_{dist}(\mathbf{x}) = \frac{1}{N_p} \sum_{i=1}^{N_p} |d(\mathbf{p_i}, \mathbf{M})|$$

Here N_p is the number of particles of the garment template and $\mathbf{p_i}$ stands for the position of a particle in space. $E_{stretch}(\mathbf{x})$ computes the average stretch of the springs employed in the simulation.

$$E_{stretch}(\mathbf{x}) = \frac{1}{N_s} \sum_{i=1}^{N_p} \sum_{j=1}^{N_p} \frac{|\mathbf{p_i} - \mathbf{p_j}|}{l_0(i,j)} \cdot c(i,j), \quad i \neq j$$

with

$$c(i,j) = \begin{cases} 1 & \text{if } \mathbf{p_i} \text{ and } \mathbf{p_j} \text{ are connected} \\ 0 & \text{otherwise} \end{cases}$$

and N_s representing the number of springs in the cloth. In Equation (1) $|1 - \ldots|$ is employed because the average deviation from the restlength (stretch or compression) of all springs is the property we are really interested in. The square root of the whole term emphasises the valleys in the otherwise fairly shallow error surface. The last component of Equation (1), $E_{sil}(\mathbf{x})$, is a measure for the difference in the silhouettes. It is computed as follows

$$E_{sil}(\mathbf{x}) = \sum_{i=1}^{4} \frac{\sum |\mathbf{S_m}(i) - \mathbf{S_g}(i)|}{\sum \mathbf{S_g}(i)},$$

where $\mathbf{S_m}(i)$ denotes the silhouette of the model from side i and $\mathbf{S_g}(i)$ the silhouette of the garment. The $\sum$ operator applied to one of these matrices computes the sum of all elements.

The factor α is always set to 1 while β, and γ are chosen such that the three terms of Equation (1) range in the same order of magnitude.

3.4 Mesh

If a mesh is used instead of a point cloud additional information is available. Namely, the orientation of the triangles conveys the notion of inside and outside. Assuming that the garment's triangles are oriented in the same way it then becomes possible to define more selective attraction forces. By attracting cloth particles only to faces that have similar normals the convergence of the procedure can be improved significantly. This is visualised in Figure 4. On the left the initial position of the garment is shown. In the middle cloth particles were attracted to the closest point on the surface of the model. The sleeve crumples on the upper surface of the arm. On the right the sleeve's lower half gets attracted by the arm's lower side because only their normals are sufficiently similar. The sleeve expands as expected.

Fig. 4. On the left the initial configuration of a T-shirt template is shown. In the middle cloth particles are attracted to the closest point on the surface. On the right attraction forces are restricted to faces with similar normals.

Another advantage of using meshes instead of point clouds is that attraction forces are not limited to discrete particle positions but can be defined to arbitrary points on the mesh's triangular surfaces. As a consequence the cloth slides more easily over the surface of the mesh.

A significant disadvantage of using meshes, however, results from the fact that a closed surface is normally generated from a point cloud. Consider for example the lower edge of the skirt shown in Figure 5. The triangulation algorithm adds an additional surface at the bottom connecting the skirt to the legs. If a skirt that is slightly too long or was initially placed slightly too low is draped on this model, its hem will simply be attracted by this lower surface folding towards the legs. This behaviour is first of all unnatural and secondly defeats our silhouette based evaluation function. If the surface is missing, the hem of the simulated skirt will be attracted by the hem of the model pushing the whole skirt up.

3.5 Segmentation

After the dimensions of the garments have been found and the draped configuration of the garments on the scan has been calculated it is trivial to segment the mesh into parts belonging to the clothing and those that are unclad or tightly dressed. The shortest distance from a point of the scan to the garment can simply be used to classify the point as belonging to the garment or not employing a simple, hard threshold.

Fig. 5. On the left a closed model was used to drape the skirt on. The cloth clearly folds at the bottom instead of pushing the fabric up. On the right the plane connecting legs and the hem of the skirt was manually cut out. The skirt hangs straight as would be expected of a real skirt.

3.6 Layered Garments

In order to segment persons wearing layered garments an iterative approach is used. That is the dimensions of only one garment are estimated at a time, starting at the outermost layer working towards the inside. The first iteration is identical to the above procedure when only the outermost article of clothing is considered. After the point cloud has been labelled the outermost garment is removed from the simulation and the next piece of clothing is inspected. For the second and all following garments the procedure is slightly altered. If a point of the simulated cloth is attracted to a mesh point that has previously been labelled as belonging to another garment the attraction force is attenuated. This prevents inner garments from being deformed more than necessary by points not actually belonging to the garment. The attraction forces cannot, however, be omitted completely because the hidden particles would be unsupported causing the garment to crumple inside the body, dragging neighbouring particles along.

4 Results

In this section we present results using on the one hand, manual measurements of the employed garments and renderings of finished segmentations on the other. For all estimation experiments the optimisation was run three times and the run with the smallest error was selected for presentation.

In Table 1 three different scans of the same skirt are used to estimate its dimensions. While the first scan works almost perfectly the other two overestimate the length of the skirt. This happens because it is almost impossible to distinguish the upper edge of the skirt in the scan data alone. The only cue a human can perceive is that the belly button is visible. So the skirt must end below it.

Table 2 shows the same evaluation for two different T-shirts. The T-shirts' dimensions were estimated using the same template and the same optimiser settings.

Table 3 and Figure 6 show a very another advantage of the algorithm. Even in the presence of severe data corruption it is possible to estimate the dimensions of the skirt well and by substituting the missing points with the simulated cloth the holes can be filled nicely.

experiment 1		experiment 2		experiment 3	
GT	est.	GT	est.	GT	est.
54	56	54	59	54	62
110	107	110	109	110	114
78	79	78	76	78	74
94	94	94	93	94	90

Table 1. Top: Comparison between ground truth (GT) and estimated (est.) measurements in cm. The rows show length, width, waist, and hip measurements. Bottom: Segmentation results

Layered garments are estimated one at a time. While estimation of the outermost garment is equivalent to estimating just one garment. Inner garments are harder to estimate. Some dimensions such as the length of the skirt or the waist circumference are impossible to estimate if the upper half of the skirt is covered by a T-shirt. The results for these measurements presented in Table 4 consequentially deviate more severely.

5 Summary

We have presented a physically based template-deformation approach for segmenting full body laser scans of dressed humans into dressed and naked parts. The method at the same time estimates the dimensions of the garments the scanned person is wearing providing a model of the garments which can be used for example by the cloth simulation at the heart of a model based motion capture algorithm.

Acknowledgements

We gratefully acknowledge funding by the Max-Planck Center for Visual Computing and Communication.

experiment 1		experiment 2		experiment 3		experiment 4	
GT	est.	GT	est.	GT	est.	GT	est.
62	62	62	63	76	77	76	78
112	117	112	109	125	127	125	178
17	6.2	17	13	17	46	17	35
40	43	40	38	42	46	42	37
44	70	44	68	50	47	50	41

Table 2. Top: Comparison between ground truth (GT) and estimated (est.) measurements in cm. The rows show length, width, sleeve length, and sleeve circumference at the opening and at the seam. Bottom: Segmentation results

References

1. Attene, M., Katz, S., Mortara, M., Patane, G., Spagnuolo, M., Tal, A.: Mesh segmentation - a comparative study. In: SMI '06: Proceedings of the IEEE International Conference on Shape Modeling and Applications 2006 (SMI'06), Washington, DC, USA, IEEE Computer Society (2006)
2. Jojic, N., Huang, T.S.: Estimating cloth draping parameters from range data. In: Proceedings of the International Workshop on Synthetic-Natural Hybrid Coding and Three Dimensional Imaging, Rhodes, Greece (1997) 73–76
3. Rosenhahn, B., Kersting, U., Powell, K., Seidel, H.P.: Cloth x-ray: Mocap of people wearing textiles. In: Pattern Recognition (DAGM). (2006) 495–504
4. Bhat, K., Twigg, C., Hodgins, J., Khosla, P., Popović, Z., Seitz, S.: Estimating cloth simulation parameters from video. In: Proceedings of ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA 2003), ACM Press (2003) 37–51
5. Press, W., Vetterling, W., Teukolsky, S., Flannery, B.: Numerical Recipes in C++: the art of scientific computing. 2nd edn. Cambridge University Press (2002)
6. Nurre, J.H.: Locating landmarks on human body scan data. In: NRC '97: Proceedings of the International Conference on Recent Advances in 3-D Digital Imaging and Modeling, Washington, DC, USA, IEEE Computer Society (1997) 289
7. Xiao, Y., Siebert, P., Werghi, N.: A discrete reeb graph approach for the segmentation of human body scans. In: The 4th International Conference on 3-D Digital Imaging and Modeling, Banff, Alberta, Canada, IEEE Computer Society (2003)
8. Allen, B., Curless, B., Popović, Z.: Articulated body deformation from range scan data. In: SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques, New York, NY, USA, ACM Press (2002) 612–619

experiment 1

GT	est.
54	55
110	120
78	81
94	104

experiment 2

GT	est.
54	65
110	109
78	76
94	98

Table 3. Tables: Comparison between ground truth (GT) and estimated (est.) measurements in cm. The rows show length, width, waist, and hip measurements. Images: Segmentation results in the presence of severe data corruption.

9. Seo, H., Magnenat-Thalmann, N.: An automatic modeling of human bodies from sizing parameters. In: SI3D '03: Proceedings of the 2003 symposium on Interactive 3D graphics, New York, NY, USA, ACM Press (2003) 19–26
10. Magnenat-Thalmann, N., Seo, H., Cordier, F.: Automatic modeling of animatable virtual humans - a survey. In: 4th International Conference on 3D Digital Imaging and Modeling (3DIM 2003). (2003) 2–11
11. Oh, S., Kim, H., Magnenat-Thalmann, N., Wohn, K.: Generating unified model for dressed virtual humans. The Visual Computer **21** (2005) 522–531
12. Terzopoulos, D., Platt, J., Barr, A., Fleischer, K.: Elastically deformable models. In: Computer Graphics (Proceedings of ACM SIGGRAPH 87), ACM Press (1987) 205–214
13. Baraff, D., Witkin, A.: Large steps in cloth simulation. In: Proceedings of ACM SIGGRAPH 98, ACM Press (1998) 43–54
14. Desbrun, M., Schröder, P., Barr, A.: Interactive animation of structured deformable objects. In: Proceedings of Graphics Interface (GI 1999), Canadian Computer-Human Communications Society (1999) 1–8
15. Kang, Y.M., Cho, H.G.: Bilayered approximate integration for rapid and plausible animation of virtual cloth with realistic wrinkles. In: Proceedings of Computer Animation, IEEE Computer Society (2002) 203–214
16. Eberhardt, B., Weber, A., Straßer, W.: A fast, flexible, particle-system model for cloth draping. IEEE Computer Graphics and Applications **16** (1996) 52–59
17. Choi, K.J., Ko, H.S.: Stable but responsive cloth. ACM Transactions on Graphics (ACM SIGGRAPH 2002) **21** (2002) 604–611
18. Bridson, R.: Computational aspects of dynamic surfaces. PhD thesis, Stanford University (2003)
19. Volino, P., Magnenat-Thalmann, N.: Accurate garment prototyping and simulation. Computer-Aided Design Applications **2** (2005) 645–654
20. Volino, P., Cordier, F., Magnenat-Thalmann, N.: From early virtual garment simulation to interactive fashion design. Computer-Aided Design **37** (2005) 593–608
21. Bridson, R., Marino, S., Fedkiw, R.: Simulation of clothing with folds and wrinkles. In: Proceedings of ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA 2003), ACM Press (2003) 28–36

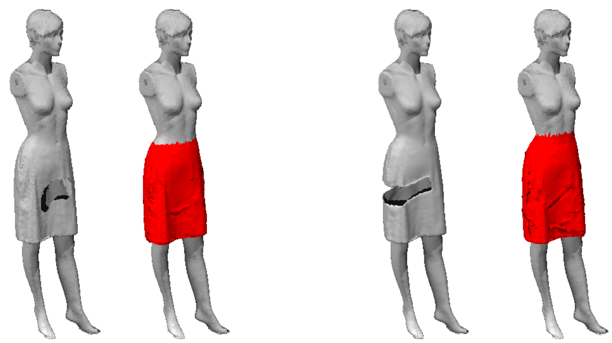

Fig. 6. Even in the presence of severe data corruption garment fitting allows the reasonable filling of holes.

experiment 1

GT	est.
54	63
110	111
78	92
94	119

experiment 2

GT	est.
54	91
110	161
78	67
94	101

Table 4. Tables: Comparison between ground truth (GT) and estimated (est.) measurements in cm. The rows show length, width, waist, and hip measurements. Images: Segmentation results of a skirt which is partially covered by a T-shirt.