



(86) Date de dépôt PCT/PCT Filing Date: 2011/06/08

(87) Date publication PCT/PCT Publication Date: 2011/12/15

(85) Entrée phase nationale/National Entry: 2012/12/03

(86) N° demande PCT/PCT Application No.: US 2011/039605

(87) N° publication PCT/PCT Publication No.: 2011/156474

(30) Priorité/Priority: 2010/06/10 (US61/353,407)

(51) Cl.Int./Int.Cl. *G06K 9/00* (2006.01)

(71) Demandeur/Applicant:
BROWN UNIVERSITY, US

(72) Inventeurs/Inventors:
BLACK, MICHAEL J., US;
FREIFELD, OREN, US;
WEISS, ALEXANDER W., US;
LOPER, MATTHEW M., US;
GUAN, PENG, US

(74) Agent: BORDEN LADNER GERVAIS LLP

(54) Titre : MODELE PARAMETRE DE FORME HUMAINE ARTICULEE EN 2 DIMENSIONS

(54) Title: PARAMETERIZED MODEL OF 2D ARTICULATED HUMAN SHAPE





Figure 1

(57) Abrégé/Abstract:

A novel "contour person" (CP) model of the human body is proposed that has the expressive power of a detailed 3D model and the computational benefits of a simple 2D part-based model. The CP model is learned from a 3D model of the human body that

(57) **Abrégé(suite)/Abstract(continued):**

captures natural shape and pose variations; the projected contours of this model, along with their segmentation into parts forms the training set. The CP model factors deformations of the body into three components: shape variation, viewpoint change and pose variation. The CP model can be "dressed" with a low-dimensional clothing model, referred to as "dressed contour person" (DCP) model. The clothing is represented as a deformation from the underlying CP representation. This deformation is learned from training examples using principal component analysis to produce so-called eigen-clothing. The coefficients of the eigen-clothing can be used to recognize different categories of clothing on dressed people. The parameters of the estimated 2D body can be used to discriminatively predict 3D body shape using a learned mapping approach. The prediction framework can be used to estimate/predict the 3D shape of a person from a cluttered video sequence and/or from several snapshots taken with a digital camera or a cell phone.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
15 December 2011 (15.12.2011)(10) International Publication Number
WO 2011/156474 A3(51) International Patent Classification:
G06K 9/00 (2006.01)(21) International Application Number:
PCT/US2011/039605(22) International Filing Date:
8 June 2011 (08.06.2011)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/353,407 10 June 2010 (10.06.2010) US(71) Applicant (for all designated States except US):
BROWN UNIVERSITY [US/US]; 1 Prospect Street,
Providence, RI 02912 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **BLACK, Michael, J.** [US/US]; 134 Laurel Avenue, Providence, RI 02906 (US). **FREIFELD, Oren** [US/US]; 405 Kipling Street, Apt. #3, Palo Alto, CA 94301 (US). **WEISS, Alexander, W.** [US/US]; 30 Groton Road, Shirley, MA 01464 (US). **LOPER, Matthew, M.** [US/US]; 34 East George Street, Providence, RI 02906 (US). **GUAN, Peng** [CN/US]; 81 Hampden Drive, Apt. 2A, Norwood, MA 02062 (US).(74) Agents: **SAUNDERS, Thomas, M.** et al.; **RISSMAN HENDRICKS & OLIVERIO, LLP**, 100 Cambridge Street, Suite 2101, Boston, MA 02114 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

[Continued on next page]

(54) Title: PARAMETERIZED MODEL OF 2D ARTICULATED HUMAN SHAPE





Figure 1

(57) **Abstract:** A novel "contour person" (CP) model of the human body is proposed that has the expressive power of a detailed 3D model and the computational benefits of a simple 20 part-based model. The CP model is learned from a 3D model of the human body that captures natural shape and pose variations. The CP model factors deformations of the body into three components: shape variation, viewpoint change and pose variation. The CP model can be "dressed" with a low-dimensional clothing model. The clothing is represented as a deformation from the underlying CP representation. This deformation is learned from training examples using principal component analysis to produce so-called eigen-clothing. The coefficients of the eigen-clothing can be used to recognize different categories of clothing on dressed people. The parameters of the estimated 20 body can be used to discriminatively predict 3D body shape using a learned mapping approach.

WO 2011/156474 A3

(88) Date of publication of the international search report:
19 April 2012

PARAMETERIZED MODEL OF 2D ARTICULATED HUMAN SHAPE

CROSS-REFERENCES TO RELATED APPLICATIONS

This application claims the benefit of prior filed U.S. provisional Application No. 61/353,407, filed June 10, 2010, pursuant to 35 U.S.C. 119(e).

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH OR DEVELOPMENT

This invention was made with support from Grant No. NIH EUREKA 1R01NS066311-01 from the National Institute of Health, and Grant No. NSF IIS-0812364 from the National Science Foundation. The United States Government has certain rights in the invention.

BACKGROUND OF THE INVENTION

[0001] The present invention relates to estimating human shape and pose from images where the shape and pose is in two dimensions. The present invention also relates to modeling clothing on 2D models of the human body, and more particularly to accurately estimate body shape underneath clothing and to recognize the type of clothing worn by a person.

[0002] The following discussion of related art is provided to assist the reader in understanding the advantages of the invention, and is not to be construed as an admission that this related art is prior art to this invention.

[0003] Three dimensional (3D) models of the human body have been widely reported and have become sophisticated and highly detailed, with the ability to accurately model human shapes and poses. Disadvantageously, however, they are complex and computationally intensive. Additionally, the

estimation of such models from a single, monocular, image is ambiguous and may require sophisticated optical scanning methods for data acquisition.

[0004] Conversely, two-dimensional (2D) models of the human body are popular due to their representational and computational simplicity. Existing models include articulated pictorial structures models, active shape models (or point distribution models), parameterized non-rigid templates, and silhouette models. However, most 2D articulated person models have focused on estimating human pose and have ignored body shape.

[0005] It would therefore be desirable and advantageous to address this problem and to obviate other prior art shortcomings by developing models that explicitly represent human shape with contours and, furthermore, represent non-rigid human shape and pose. Such a representation enables addressing the issue of body shape recognition, a task that is beyond the scope of traditional 2D models. It would also be desirable and advantageous to explicitly model how clothing influences human shape based the 2D human shape model. It would also be desirable and advantageous to be able to predict 3D body shape from silhouettes in one or several 2D uncalibrated images.

SUMMARY OF THE INVENTION

[0006] Common models such as pictorial structures are typically fairly crude, lacking realistic detail which limits their expressive power within a generative framework. The 2D representation of a human body according to the present invention is based on a “contour person” (CP) model of the human body that has the expressive power of a detailed 3D model and the computational benefits of a simple 2D part-based model. The CP model is learned from a 3D model of the human body that captures natural shape and pose variations; the projected contours of this model, along with their segmentation into parts, form the training set. The CP model factors deformations of the body into three

components: shape variation, viewpoint change and pose variation.

[0007] To model clothing, the CP model is “dressed” with a low-dimensional clothing model, to be referred to as “dressed contour person” (DCP) model. The clothing is represented as a deformation from the underlying CP representation. This deformation is learned from training examples using principal component analysis to produce eigen-clothing.

[0008] The 3D body shape is discriminatively predicted from parameters of the estimated 2D body using learned mapping (e.g. a mixture-of-experts model). For example, the method can be used to predict the 3D shape of a person from a cluttered video sequence. Alternatively, the 3D shape of a person can be estimated from several snapshots taken with a digital camera or cell phone.

[0009] One aspect of the present invention relates to a method for modeling two-dimensional (2D) contours of a human body by capturing natural shapes/poses of a three-dimensional (3D) model of a human body, projecting the 3D model into a camera plane to derive 2D contours, segmenting the 2D contours based on a predetermined segmentation of the 3D model, and creating a factored representation of non-rigid deformations of the 2D contours, wherein the deformations are represented by line segments.

[0010] According to one advantageous feature of the present invention, the deformations may be factored into at least two components selected from shape, pose and camera viewpoint. Edges of a contour proximate to a joint between segmented 2D contours may be corrected by applying non-rigid deformations to the edges of the contour proximate to the joint.

[0011] According to one advantageous feature of the present invention, a training set of clothing outlines and an underlying naked human body represented by the 2D contours may be created, and a displacement between

contour points of the clothing outlines and the underlying naked human body may be determined to define a clothing deformation. The clothing deformation may represent a deformation (through scaling and rotation) of the line segments of 2D contours of the human body.

[0012] According to one advantageous feature of the present invention, the clothing deformation obtained from the training set may represent “eigen-clothing” which may be separated into “eigen-separates”, with each “eigen-separate” covering a part of the 2D contours of the human body.

[0013] According to one advantageous feature of the present invention, different types of clothing may be classified based on clothing deformation coefficients.

[0014] According to one advantageous feature of the present invention, the “eigen-clothing” may be subtracted from a clothed human body, and the underlying naked human body may be estimated therefrom.

[0015] According to one advantageous feature of the present invention, the 3D model may be projected into different camera planes to derive 2D contours in at least two different views to obtain a plurality of body silhouettes, the 2D contour model may be fitted to the plurality of body silhouettes, 2D shape parameters may be determined from the factored representation of the non-rigid deformations of each of the plurality of body silhouettes, and 3D shape parameters may be predicted from the determined 2D shape parameters. The body silhouettes obtained for different shapes, poses and views of a person of the projected 3D model may be combined into an estimate of the 3D shape of the person.

[0016] According to one advantageous feature of the present invention, after subtracting “eigen-clothing” from the 2D contour model of a clothed human

body to obtain the 2D contour model of an unclothed human body, 3D shape parameters of the unclothed human body may be predicted from the determined 2D shape parameters of the clothed human body.

[0017] According to another aspect of the invention, a two-dimensional (2D) contour person (CP) model for modeling an unclothed human body includes a shape variation component, a viewpoint change component, and a pose variation component, wherein the contour person (CP) model is learned from a 3D model of the human body that captures natural shape and pose variations, with projected contours of this model, along with their segmentation into parts, form a training set.

[0018] According to another aspect of the invention, a dressed contour person (DCP) model for modeling a clothed human body, includes the 2D CP model of an unclothed human body, wherein deformations are computed by aligning the contour of the clothed human body with the contour of the unclothed human body, and wherein an “eigen-clothing” model is learned using principal component analysis (PCA) applied to these deformations.

BRIEF DESCRIPTION OF THE DRAWING

[0019] Other features and advantages of the present invention will be more readily apparent upon reading the following description of currently preferred exemplified embodiments of the invention with reference to the accompanying drawing, in which:

FIG. 1 shows a conventional 2D body models (left) in form of simple articulated collections of geometric primitives and a contour person (CP) model (right) according to the present invention;

FIG. 2 shows the CP model with a range of articulation of different body

parts;

- FIGS. 3a–d illustrate non-rigid deformation of the left arm (heavy line) of the contour person, showing (a) template; (b) rigid transformation of upper arm; (c) same as (b) but with parts which should be non-rigidly deformed due to the rigid motion marked in heavy lines; and (d) final deformed contour with the non-rigidly deformed parts marked in heavy lines.
- FIG. 4 shows 2D contour people sampled from the model: Row 1: variations in body shape; Row 2: variations in pose; Row 3: variations in camera view; and Row 4: a combination of all variations;
- FIG. 5 shows the dressed contour model of the invention in different body shapes and poses, dressed in different types of eigen-clothing;
- FIGS. 6a-g show for the same naked shape the mean clothing contour (a) and the clothing contour ± 3 std from the mean for several principal components (b)-(d). The associated statistics of the clothing deformations are shown in (e)-(g);
- FIG. 7 shows synthetic data results, which each pair showing the DCP result is on the left and NM (naïve method) result is on the right;
- FIG. 8 shows sample DCP results of estimated underlying bodies overlaid on clothing. Results are shown for a variety of poses (left to right) and viewing directions (top to bottom);
- FIGS. 9a-e show clothing types for three types of upper clothing: long sleeves (a and e; top), short sleeves (b and d; top) and sleeveless tops (c;

top) and four types of lower clothing: short pants (b and e; bottom), long pants (d; bottom), short skirt (c; bottom), and long skirts (a; bottom);

FIG. 10 shows classification results for the DCP model for the 7 clothing types of FIG. 9 and all 8 poses of FIG. 8, as compared to “Chance”;

FIG. 11 shows results from an ICP (Iterative Closest Points) model with points mapped to their respecting closest observed points (heavy lines);

FIGS. 12a-l show a 3D Shape Prediction aspect of the invention from 2D Contour People. (a-d) and (g-j) show respective CP representations of a particular person in several views and poses. (e) and (k) are the corresponding predicted 3D shapes, while (f) and (i) are the true 3D shapes; and

FIGS. 13a-e shows exemplary fits of the 2D CP model for multiple views and a predicted 3D mesh.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

[0020] Throughout all the figures, same or corresponding elements may generally be indicated by same reference numerals. These depicted embodiments are to be understood as illustrative of the invention and not as limiting in any way. It should also be understood that the figures are not necessarily to scale and that the embodiments are sometimes illustrated by graphic symbols, phantom lines, diagrammatic representations and fragmentary views. In certain instances, details which are not necessary for an understanding of the present invention or which render other details difficult to perceive may have been omitted.

[0021] The following abbreviation will be used throughout the specification:

CP	Contour Person
DCP	Dressed Contour Person
PCA	Principal Component Analysis
ICP	Iterative Closest Point
PS	Pictorial Structures
SPM	Scaled Prismatic Model
ASM	Active Shape Model
PC	Principal Component
NM	Naïve method
NP3D	Naked People Estimation in 3D
SCAPE	Shape Completion and Animation of People
GT	Ground truth

[0022] The detection of people and the analysis of their pose in images or video have many applications and have drawn significant attention. In the case of uncalibrated monocular images and video, 2D models dominate while in calibrated or multi-camera settings, 3D models are popular. In recent years, 3D models of the human body have become sophisticated and highly detailed, with the ability to accurately model human shape and pose. In contrast, 2D models typically treat the body as a collection of polygonal regions that only crudely capture body shape (Fig. 1(a)). Two-dimensional models are popular because they are relatively low dimensional, do not require camera calibration, and admit computationally attractive inference methods. For many problems, such as pedestrian detection, full 3D interpretation may not be needed.

[0023] In this invention, a novel 2D model of the body is described that has many of the benefits of the more sophisticated 3D models while retaining the computational advantages of 2D. This 2D Contour Person (CP) model (Fig. 1(b)) provides a detailed 2D representation of natural body shape and captures

[0024] The CP model builds on a person detector that uses a conventional pictorial structure (PS) model. However, the CP model according to the invention increases the realism by modeling shape variation across bodies as well as non-rigid deformation due to articulated pose changes. Moreover, it provides several types of pose parameterizations based on the PS model, thus making use of existing PS inference algorithms. Importantly, the CP model according to the invention models deformations of 2D contours which is important for explicitly modeling articulation and for factoring different types of deformations.

[0025] Referring now to FIG. 2, the 2D body shape is factored into: 1) a linear model characterizing shape change across the population; 2) a linear approximation to distortions caused by local camera view changes; 3) an articulation of the body parts represented by a rotation and length scaling; and 4) a non-rigid deformation associated with the articulation of the parts. An example of the full model with a range of articulations of different body parts is shown in Fig. 2.

[0026] The CP model is built from training data generated from a 3D SCAPE (Shape Completion and Animation of PEople) body model capturing realistic body shape variation and non-rigid pose variation. Each training body for the CP model is generated by randomly sampling a body shape, in a random pose and viewed from a random camera. The bounding contour of the 3D body is

projected onto the camera plane to produce a training contour. The known segmentation of the 3D model into parts induces a similar 2D contour segmentation (Fig. 2).

[0027] As described in more detail in Provisional US Patent Application Serial No. 61/353,407, the contents of which are incorporated by reference in their entirety as if fully set forth herein, the CP is a pattern-deformable template model whose basic deformation unit is a scaled rotation matrix acting on a line segment connecting two contour points.

$$\mathbf{T} = (\mathbf{R}_1 \ \mathbf{R}_2 \ \mathbf{R}_3 \ \mathbf{R}_4 \ \dots \ \mathbf{R}_N) \quad (1)$$

and

$$\mathbf{C} = (x_1 \ y_1 \ x_2 \ y_2 \ \dots \ x_N \ y_N)^T \quad (2)$$

represent a template and a deformed contour respectively, where N is the number of contour points. For now, assume for simplicity that the contour is closed and linear. In effect, every contour point is connected exactly to two points: its predecessor and successor in a (for example) clockwise orientation. This graph connectivity is expressed by a 2N by 2N sparse matrix E:

$$E = \begin{pmatrix} -1 & 0 & 1 & 0 & \dots & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 1 & \dots & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & \dots & 0 & 0 & 0 & 0 \\ & & & \vdots & & & & & \\ 0 & 0 & 0 & 0 & & -1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & & 0 & -1 & 0 & 1 \\ 1 & 0 & 0 & 0 & \dots & 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 & \dots & 0 & 0 & 0 & -1 \end{pmatrix}. \quad (3)$$

[0028] The CP deformation is encoded in $\mathbf{D}$ ($\mathbf{D}$ is parameterized by Θ , to be described later), a block diagonal matrix, whose 2 by 2 blocks are scaled rotation matrices. In effect, let (S_i, θ_i) denote the scale and angle of the i^{th}

directed line segment between two points in T and let $D_{2 \times 2}^i$ represent the associated scaled rotation matrix acting on this line segment:

$$D_{2 \times 2}^i = S^i \begin{pmatrix} \cos \theta_i & -\sin \theta_i \\ \sin \theta_i & \cos \theta_i \end{pmatrix} = \exp \left(s^i \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} + \theta^i \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \right) \quad (4)$$

where $s^i = \log S^i$ is the log of the scale, and $\exp$ stands for the matrix exponential. Thus, the matrix product

$$EC = DET \quad (5)$$

defines the directed line segments of C as a deformed version of those of T . Note that left multiplication of ET by D can be viewed as an action of a Lie group.

[0029] However, for an arbitrary $D = D(\Theta)$ this matrix product may not result in a meaningful contour (e.g., not closed). A possible solution is imposing a global constraint on the local deformations, with the disadvantage of losing the group structure of the deformations. Instead, a different approach is employed. Specifically, given a prescribed deformation matrix D , we seek a contour C such that its line segments (denoted by e_i) are close to the desired deformed line segments of T (denoted by $D_{2 \times 2}^i l_i$) in a least squares sense. In effect, we seek to minimize

$$\sum_{i=1}^N \|D_{2 \times 2}^i l_i - e_i\|^2 = \|DET - EC\|^2. \quad (6)$$

The minimizer yields the *contour synthesis* equation:

$$C = E^\dagger DET \quad (7)$$

where $E^\dagger$, the Moore-Penrose pseudoinverse of the constant matrix E , is computed offline. The connectivity of G ensures the closure of C .

Eq. 7 shows how to synthesize C from D and the template. Conversely, given known l_i and e_i , we compute D_i^1 by solving the invertible ($\|l_i\| > 0$) linear system

$$e_i = s_i R_{\theta_i} l_i = \begin{pmatrix} l_i^{(1)} & -l_i^{(2)} \\ l_i^{(2)} & l_i^{(1)} \end{pmatrix} \begin{pmatrix} s_i \cos \theta_i \\ s_i \sin \theta_i \end{pmatrix}. \quad (8)$$

Deformations of the template contour are factored into several constituent parts: pose, shape, and camera, which are then composed to derive the full model.

[0030] Starting from the 3D SCAPE model, numerous realistic body shapes are generated in a canonical pose and their contours are projected into the image. The segmentation of contour points in 2D is known from the segmentation of the body parts in 3D, which is used to evenly space points along a training part. The known segmentation prevents points from “sliding” between parts. The result is 2D training contours with known alignment of the contour points.

[0031] The deformation for each contour is computed from a single template contour. A matrix of all these training deformations is then formed (subtracting the mean) and PCA (Principal Component Analysis) is performed, resulting in a linear approximation to contour deformations caused by body shape variation parameterized by the PCA coefficients.

[0032] Global deformations, which are used in linear pose parametrization, can be accommodated like any other deformation in the PCA analysis, for example, by adding a sparse row to E , with 1/2 and -1/2 at suitable places in the matrix.

[0033] The deformations due to camera variation are also well captured by PCA, with 6 components accounting for more than 90% of the variance.

[0034] In the 3D SCAPE model, deformations due to body articulation are modeled by a two-step process. First, a rigid rotation is applied to the entire limb or body part, and then local non-rigid deformations are applied according to a learned linear model. A similar approach is employed here in 2D.

[0035] Turning now to FIG. 3, as seen for example in Fig. 3(b), a rigid motion of the upper arm does not account for non-rigid deformations of the shoulder. This is corrected by applying a learned non-rigid deformation to the edges of the contour in the vicinity of the joint (Fig. 3(d)). Note that a rigid motion of the upper arm affects the non-rigid deformation of the upper arm as well as those of the lower arm and the shoulder. The residual is the deformation of the contour that is not accounted for by part-rotation and part-scaling.

[0036] The CP model according to the invention utilizes 3D information when it is constructed; this is quite different from standard 2D models and allows it to handle self-occlusions as well as out of the plane rotations. In a standard 2D contour model, the ordering of the points would be poorly defined. However, since the contours of the CP model according to the invention are generated from a 3D mesh, the correspondence between 2D contour points and the respective points and body parts on the 3D mesh is known, so that the contour is correctly connected even in the event of a crossover in 2D.

[0037] The full model is derived by first training each deformation model independently and then composing them by way of a simple matrix multiplication. Since 2D rotation matrices commute, the composition order is immaterial. Given parameters for shape, pose, and camera view, the overall deformation is given by the deformation synthesis equation:

$$D(\Theta) = D_{pose} D_{shape} D_{camera} \quad (9)$$

where $\Theta = \{\Theta_{pose}, \Theta_{shape}, \Theta_{camera}\}$. $D(\Theta)$ can be substituted into the contour synthesis equation $C = E^\dagger DET$ to produce a new contour C . In the example

illustrated in FIG. 4, 24 pose parameters (12 joints $\times$ 2), 10 shape coefficients and 6 camera coefficients are used, for a total of 40 parameters.

[0038] An exemplary application of the CP model is related to the problem of segmenting images of humans. The CP model provides a strong prior over human body shape that can be used to constrain more general segmentation algorithms. Specifically, one searches over the CP parameters that optimally segment the image into two regions (person and non-person) using a cost function that 1) compares image statistics inside the contour with those outside; 2) favors contours that align with image edges; and 3) enforces our prior model over shape, pose and camera parameters. A region term of the segmentation objective compares intensity and color histograms inside and outside the body contour. Because the segmented contour should also follow image edges, image edges are detected using a standard edge detector and a thresholded distance transform is applied to define an edge cost map normalized to [0, 1].

[0039] The model is not clothed and consequently will produce segmentations that tend to ignore clothing. While the optimization could be made explicitly robust to clothing, it might be advantageous for segmenting clothed people to explicitly model clothing which will be described in detail below.

[0040] Conventional 2D models which are widely used in computer vision tasks, such as pose estimation, segmentation, pedestrian detection and tracking, fail to explicitly model how clothing influences human shape. In the following, an exemplary embodiment of a fully generative 2D model is described that decomposes human body shape into two components: 1) the shape of the underlying naked body and 2) the shape of clothing relative to the underlying body. The naked body shape is represented by the Contour Person (CP) model as described in the previous section. Given training examples of people in clothing with known 2D body shape, we compute how clothing deviates from the naked body to learn a low-dimensional model of this deformation. The resulting

[0041] The key idea is to separate the modeling of the underlying body from its clothed appearance. The most likely naked body shape can be inferred from images of clothed people by explicitly modeling clothing. The model also supports new applications such as the recognition of different types of clothing.

[0042] There are several novel properties of the DCP model. First, so-called eigen-clothing is defined to model deformation from an underlying 2D body contour. Given training samples of clothed body contours, where the naked shape of the person is known, the naked contour is then aligned with the clothing contour to compute the deformation. The eigen-clothing model is learned using principal component analysis (PCA) applied to these deformations. A given CP model is then “clothed” by defining a set of linear coefficients that produce a deformation from the naked contour. This is illustrated, for example, in Fig. 5.

[0043] The estimation of a person’s 2D body shape under clothing from a single image is demonstrated, showing clearly the advantages over a principled statistical model of clothing.

[0044] Finally the problem of clothing category recognition is introduced. It can be shown that the eigen coefficients of clothing deformations are distinctive and can be used to recognize different categories of clothing such as long pants, skirts, short pants, sleeveless tops, etc. Clothing category recognition could be useful for person identification, image search and various retail clothing applications.

[0045] In summary, the key contributions of this invention include: 1) the first model of 2D eigen-clothing; 2) a full generative 2D model of dressed body shape that is based on an underlying naked model with clothing deformation; 3)

the inference of 2D body shape under clothing that uses an explicit model of clothing; and 4) a shape-based recognition of clothing categories on dressed humans.

[0046] We directly model the deformation from a naked body to a clothed body by virtually “dressing” the naked contour with clothing. We start with a training set (described below) of clothing outlines and corresponding naked body outlines underneath. The CP model is first fit to the naked body outline to obtain a CP representation with a fixed number of contour points. For each point on the CP, we find the corresponding point on the clothing outline and learn a point displacement model using PCA. We further learn a prior over the PCA coefficients using a Beta distribution to prevent infeasible deformations (i.e. “negative clothing” that causes the naked body to appear smaller than it is). Finally we define a two layer deformation model in which the first layer generates a naked body deformation from a template body and the second layer deforms the naked body to a clothed body. The parameters controlling the pose and shape of the body can be changed independently of the parameters controlling clothing type. These method requires training contours of people in clothing for which we know the true underlying naked body shape. We use two such training sets.

[0047] Synthetic data provides GT (Ground Truth) body shapes that enable accurate quantitative evaluation. We use 3D body meshes generated from the CAESAR database (SAE International) of laser range scans and then dress these bodies in simulated clothing. We use 60 male and 100 female bodies spanning a variety of heights and weights and use commercial software (OptiTex International, Israel) to generate realistic virtual clothing. The clothing simulation produces a secondary 3D mesh that lies outside the underlying body mesh by construction. Given a particular camera view, we project the body mesh into the image to extract the body outline and then do the same for the combined body and clothing meshes. This provides a pair of training outlines. We restrict the

clothing to a single type (Army Physical Training Uniforms) but in sizes appropriate to the body model. While narrow, this dataset provides perfect training data and perfect ground truth for evaluation.

[0048] For training data of real people in real clothing we used a dataset with a set of images of 6 subjects (3 male and 3 female) captured by 4 cameras in two conditions: 1) A “naked condition” where the subjects wore tight fitting clothing; 2) A “clothed condition” in which they wore various different “street” clothing. Each subject was captured in each condition in a fixed set of 11 postures. Each posture was performed with 6-10 different sets of clothing (trials) provided by the subjects. Overall there are 47 trials with a total of 235 unique combinations of people, clothing and pose. For each clothed image we used standard background subtraction to estimate the clothed body silhouette and extracted the outline. To obtain the underlying naked body outlines of those clothed images, we obtained the 3D parametric body model fit using the 4 camera views of the naked condition. We consider the resulting 3D body shape to be the true body shape. For each subject in each posture the pose of the 3D body has been optimized using the 4 camera views while holding the shape fixed. The resulting 3D body outline is then projected into a certain camera view and paired with the segmented clothed body in that view. Note that the fitting of the 3D body to the image data is not perfect and, in some cases, the body contour actually lies outside the clothing contour. This does not cause significant problems and this dataset provides a level of realism and variability not found in the perfect synthetic dataset.

[0049] Given the naked and clothed outlines, we need to know the correspondence between them. Defining correspondence is nontrivial and how it is done is important. Incorrect correspondence (i.e. sliding of points along the contour) results in eigen shapes that are not representative of the true deformations of the contours.

[0050] Two contours $B_T(\psi)$ for the CP and G for the clothing contour, both of which have N points, are computed. For each body/clothing pair, x and y , coordinates of the body and clothing contours are stacked to get $\mathbf{u} = (b_{x_1}, b_{y_1}, b_{x_2}, b_{y_2}, \dots, b_{x_N}, b_{y_N})^T$ for the body and $\mathbf{v} = (c_{x_1}, c_{y_1}, c_{x_2}, c_{y_2}, \dots, c_{x_N}, c_{y_N})^T$ for the clothing. The clothing deformation is then defined by $\delta = \mathbf{u} - \mathbf{v}$ on which PCA is performed. It can be shown that the first 8 principal components account for around 90% of the variance to define the eigen-clothing model. FIGS. 6a-d show the mean clothing base and the first few clothing bases. This illustrates how the bases can account for things like long pants, skirts, baggy shirts, etc. The statistics of clothing deformations are shown in FIGS. 6e-g, including exemplary histograms and Beta distribution fits to linear eigen-clothing coefficients. Note the skew that results from the fact that clothing generally makes the body to appear larger

[0051] With this model, new body shapes in new types of clothing can be generated by first sampling CP parameters ψ to create a naked body contour $B_T(\psi)$ and generating a clothed body using $C(B_T(\psi), \eta) = B_T(\psi) + C_{mean} + \sum_{i=1}^{N_\eta} \eta_i \cdot C_i$, where N_η is the number of eigen vectors used, the η_i 's are coefficients, C_{mean} is the mean clothing deformation, and C_i is the i^{th} eigen-clothing vector.

[0052] Based on the observation of natural clothing statistics, a prior on the clothing deformation coefficients can be learned to penalize infeasible clothing deformations. The clothing contour $C(B_T(\psi), \eta)$ is the result of two layer deformations, with the first one, $B_T(\psi)$, from a body template to a particular body and the second one, $C(B_T(\psi), \eta)$, from the body to a clothing contour. The inference problem is to estimate variables ψ and η with a single clothing view. A likelihood function is defined in terms of silhouette overlap. The final energy

function $E(\psi, \eta) = E_{data}(\psi, \eta) + \lambda E_{prior}(\eta)$ is minimized, wherein λ indicates the importance of the prior. Problems with “negative clothing” and clothing that is unusually large are avoided due to the prior. Optimization is performed using MATLAB’s fminsearch function, although other optimization methods may be used, such as gradient descent or stochastic search.

[0053] Two novel applications of the proposed method will now be described. One is to estimate the underlying 2D body given a clothing observation. The other is clothing category recognition by classifying the estimated clothing deformation parameters. The model will be examined for: (1) naked body estimation for synthetic data, (2) naked body estimation for real data, and (3) clothing type classification for real data. The results of the first two tasks will be compared with approaches that do not explicitly model clothing deformation.

[0054] The results of DCP are also compared for synthetic data with a naive method (NM) in which the CP model is simply fitted to the clothing observation. The NM takes clothing into account by assessing a greater penalty if the estimated silhouette falls outside of the clothing observation and less if it does not fully explain the clothing observation. The average estimation errors obtained with NM for males and females are 4.56% and 4.72% respectively while DCP achieves 3.16% and 3.08%. The DCP model thus improves the relative accuracies over NM by 30% (for male) and 35% (for female). While the synthetic dataset has only one clothing type, the bodies span a wide range of shapes. The results show the advantage of modeling clothing deformation compared with ignoring clothing. FIG. 7 shows some representative results from the test set, with the DCP result on the left and NM result on the right. The first pair shows an estimated body silhouette overlaid on clothing silhouette; overlapped regions are shown with light texture. The second pair shows the estimated body overlaid on the ground truth (GT) body. While the first and second pairs are for a male, the third and fourth pairs show body silhouettes for a female. NM typically

overestimates the underlying body while still keeping the body inside the clothing silhouette.

[0055] In a body estimation under clothing for real data, eight different poses (arranged in each row of FIG. 8) are evaluated, each pose with 47 instances having unique combinations of subject and clothing. Since the number of body/clothing pairs is limited, in each pose, a leave-one-out strategy is used where the body of instance i is estimated using the eigen-clothing space learned from all the other 46 instances, excluding i . The underlying naked body DCP is estimated for a total of $47 \times 8 = 376$ instances (Fig. 8) and the results are compared with two other methods: 1) NM described in the previous experiment; and 2) “Naked People estimation in 3D”(NP3D). Since DCP and NM are 2D methods using a 2D CP model, they only use one camera view. NP3D, however, uses a 3D model with multiple views and the 3D body was computed with all 4 camera views. To compare with NP3D the estimated body is projected from NP3D into the image using the camera view used by the method of the invention. Table 1 shows the comparison of Average Estimation Error (AEE) and standard deviation (std) on 47 instances for each pose. By modeling clothing deformations the 2D method of the invention even outperforms the conventional multi-camera method.

Table 1: Comparison on real data: DCP, NM, and NP3D methods

	Pose 1	Pose 2	Pose 3	Pose 4	Pose 5	Pose 6	Pose 7	Pose 8	Average
DCP	3.72	5.25	5.08	4.37	4.33	4.51	5.03	6.68	4.87%
(AEE)	%	%	%	%	%	%	%	%	
DCP (std)	0.019	0.028	0.029	0.031	0.022	0.028	0.026	0.038	0.022
NP3D	4.11	6.28	5.62	4.84	4.94	4.60	4.72	7.23	5.29%
(AEE)	%	%	%	%	%	%	%	%	
NP3D	0.027	0.032	0.034	0.028	0.036	0.026	0.036	0.051	0.023
(std)									
NM (AEE)	8.56	9.12	8.46	8.35	8.77	9.21	9.02	11.84	9.18%
	%	%	%	%	%	%	%	%	
NM (std)	0.023	0.026	0.031	0.029	0.028	0.035	0.031	0.043	0.025

[0056] Turning now to Fig. 9, different types of clothing can be classified from the estimated clothing deformation coefficients. Upper clothing and lower clothing are separated; in this example, 7 types of clothing are shown: three types of upper clothing: long sleeves (a and e; top), short sleeves (b and d; top) and sleeveless tops (c; top) and four types of lower clothing: short pants (b and e; bottom), long pants (d; bottom), short skirt (c; bottom), and long skirts (a; bottom).

[0057] Fig. 10 shows the classification results for the DCP model for the 7 aforementioned clothing types and the 8 different poses shown in FIG. 8, as compared to “Chance”. We use a simple nearest neighbor (NN) classifier with Euclidean distances computed from the first 8 principal components. Other classifiers, such as Support Vector Machines, Bayesian classifiers, etc, may also be used. Since the number of clothing instances (47) for each pose is limited, a leave-one-out strategy can be used by assuming that the categories of all the instances, except for the one being tested, are known. Each instance is then assigned a category for both upper clothing and lower clothing based on its nearest neighbor.

[0058] While clothing deformation models have been shown for static poses, it should be understood that this model can be applied to a range of poses. In particular, given training data of the CP model in different poses and the associated clothing deformation for each pose, the model can be extended to learn a deformation as a function of the pose. This function can be represented by a linear, multi-linear, or non-linear model. Given the body in a particular pose, the appropriate clothing deformation can then be synthesized based on the pose.

[0059] Here we have described clothing deformation as a point deformation from the underlying CP model. The CP model itself is defined in terms of edges. It should be understood that the clothing model can also adopt

this representation. In this case clothing becomes a deformation (scale and rotation) of line segments from the underlying CP model.

[0060] Although all the clothing on the body was modeled at once, it should be clear that different pieces of clothing can be modeled individually. These models will be referred to as “eigen-separates”, with separate eigen models for tops and bottoms for example, as described with reference to FIG. 9. A body is then clothed by combining the separate deformation from each piece of clothing to produce a dressed body. This same idea can be used to model hair, hat and shoe deformations. In fact it can also model things like backpacks, purses or other items worn or carried by humans.

[0061] The estimation of 2D body shape under clothing has numerous potential applications, especially when multiple cameras are not available. Consider forensic video from a single camera in which the anthropometric measurements of the suspect are to be identified while the body’s shape is obscured by clothing. In the following section we show how to go directly from 2D body shape to 3D measurements.

[0062] In computer vision, 3D human detailed shape estimation from 2D marker-less image data is a problem of great interest with many possible applications. The present invention extends the conventional 3D human shape estimation from 2D image data by parameterizing the 2D body through representation of pose deformations using a low dimensional linear model, supporting limb scaling and rotations. In addition, the CP model is extended to views other than frontal. This in turn has the desired side effect of not only making the CP model suitable for additional single-view settings, but also renders the model capable of modeling the human body in a multi-view setup or in a sequential data (such as the estimation of TV character from a full episode).

[0063] In addition, the body shape, pose and camera deformation parameters can be optimized, in particular, based various gradient-based optimization methods, for example, Levenberg-Marquardt or Newton's methods.

[0064] Lastly, a framework is described to predict 3D body shape from silhouettes in one or several 2D uncalibrated images. The framework is based on the invariance property of the hidden shape parameter across different views or poses, and relies on the CP factorization. By factoring shape deformation from pose and camera view, we show how to directly predict 3D shape parameters from 2D shape parameters using a discriminative method. In addition to predicting 3D shape from one or more images, the framework allows additional information (such as given height or stature) to be optionally used in predicting 3D body shape.

[0065] The 2D shape model parameters described above in relation to the CP model are used directly to estimate 3D shape. This approach has an important advantage. It allows multiple images of a person taken at different times and in different poses to be combined to estimate a consistent 3D shape. Body shape from pose and camera view are taken in consideration for each 2D view of a person. Each 2D view then provides information about 3D shape that should be consistent. The 2D views are then combined into a single 3D shape prediction.

[0066] This approach allows body shape to be recovered from a sequence of images (e.g. a video sequence) in which a person is moving. In another application, a person might snap several pictures of a subject with a digital or cell-phone camera. Between each picture, the subject may move or change pose. By factoring 2D body shape from pose and camera changes, the 3D shape can still be estimated accurately.

[0067] The first step in estimating the 3D shape from the 2D shape model parameters is to describe how to fit the CP model to body silhouettes. Given an outline of a silhouette, we would like to find the best fit of the CP model. We employ a bi-directional Iterative Closest Point (ICP) approach. Given a current value of the model parameter β , for each point of the outline of the model we find the closest point among the observed points (we also impose a threshold on the difference between the normals). This is illustrated in FIG. 11. Likewise, for every observed point we find the closest point among the points of the outline of the model. The method then computes a cost function using an iterative process.

[0068] If additional information is available, such as height, weight or other measurements, it is possible to condition the shape model accordingly. This is done by appending the (properly scaled) measurement (or measurements) for the training instances to the vector deformations when learning the PCA model. Then, given a nominal value of height, a conditional PCA shape model can be used. Our experiments show that this type of information, while optional, can improve both fitting and prediction results.

[0069] Alternatively, a separate function mapping measurements to 2D shape coefficients can be learned. This could be linear, multi-linear, non-linear or probabilistic. Then given some input measurements, this provides a prediction of 3D shape that can be incorporated as a regularizing prior in the objective function. In this way the known measurements are incorporated during optimization.

[0070] Another inherent advantage of the CP model is the fact that every contour point is affiliated with a specific contour part (e.g., the outer side of the left calf, the right side of the torso, etc.). This enables us to reduce the weight of particular body parts during model fitting. For example, by reducing the weight of points that correspond to the head, we increase robustness to errors that are caused by presence of hair. In a similar approach, downweighting the weight in

the presence of occlusions enables us to fit the model to data from a cluttered video sequence (e.g., a TV episode during which parts of the body keep appearing and disappearing). In such a case, only some parts of the body may be seen at a given instant. We are thus able to combine many partial views taken at different times, in different poses, and from different camera views into a coherent estimate of 3D shape.

[0071] Turning now to FIG. 12, this section will explain how to use the generative 2D CP model to discriminatively predict a 3D human shape. To illustrate the approach we use 4 different views: T-pose in a frontal view (a, g); T-pose in a 3/4 view (b, h); a side view with arms behind the body (c, i); and a side view with arms in front of the body (d, j). While other poses and views may also be possible, these poses were found to be both informative enough for purposes of shape prediction, as well as simple for subjects to understand.

[0072] We evaluate our 3D prediction framework on 4000 synthetic examples generated from SCAPE. Visual results showing the corresponding predicted 3D shapes are shown in FIGS. 12e and k, with FIGS. 12f and l showing the true 3D shapes. Numerical results for different biometric measurements are listed in Table 2, demonstrating the high accuracy of the results.

Table 2: Error summary		
Measurement	Average Relative Error	RMS
Stature (training)	0.64%	18.5 [mm]
Stature (test)	0.54%	11.5 [mm]
Knee height (training)	0.87%	7.0 [mm]
Knee height (test)	0.82%	5.3 [mm]
Thigh circumference (training)	1.86%	17.4 [mm]
Thigh circumference (test)	1.74%	13.18 [mm]
Crotch height (training)	1.0%	12.5 [mm]
Crotch height (test)	1.0%	9.1 [mm]

[0073] A realistic fitting result can be seen in FIG 13e based on the 2D CP model and the given segmentation. The experiments show that results improved

when the square error is replaced with a robust error function as described in the optimization section. This may be important when silhouettes contain outliers due to hair or loose clothing or other errors.

[0074] In summary, a novel and complete solution for representing 2D body shape has been described, estimating the 2D body shape from images, and using the 2D body shape to compute 3D body shape. The key property of the 2D contour person model is its ability to factor changes in shape due to identity, pose and camera view. This property allows images of a person taken in different poses and different views to be combined in a 3D estimate of their shape. Moreover we have described how to estimate 2D body shape under clothing using a novel eigen-clothing representation.

[0075] There are many application of this invention to person detection, segmentation and tracking; 3D shape estimation for clothing applications, fitness, entertainment and medicine; clothing detection and recognition; shape analysis in video sequences; surveillance and biometrics; etc.

[0076] While the invention has been illustrated and described in connection with currently preferred embodiments shown and described in detail, it is not intended to be limited to the details shown since various modifications and structural changes may be made without departing in any way from the spirit and scope of the present invention. The embodiments were chosen and described in order to explain the principles of the invention and practical application to thereby enable a person skilled in the art to best utilize the invention and various embodiments with various modifications as are suited to the particular use contemplated.

[0077] What is claimed as new and desired to be protected by Letters Patent is set forth in the appended claims and includes equivalents of the elements recited therein:

Claims:

What is claimed is:

1. A method for modeling two-dimensional (2D) contours of a human body, comprising the steps of:
capturing natural shapes/poses of a three-dimensional (3D) model of a human body,
projecting the 3D model into a camera plane to derive 2D contours,
segmenting the 2D contours based on a predetermined segmentation of the 3D model, and
creating a factored representation of non-rigid deformations of the 2D contours, wherein the deformations are represented by line segments.
2. The method of claim 1, wherein the deformations are factored into at least two components selected from shape, pose and camera viewpoint.
3. The method of claim 2, further comprising the step of training shape deformation by generating realistic 3D body shapes in canonical pose and projecting the contours of the 3D body shapes into the camera plane.
4. The method of claim 3, further comprising the step of correcting for edges of a contour proximate to a joint between segmented 2D contours by applying non-rigid deformations to the edges of the contour proximate to the joint.
5. The method of claim 1, wherein proper connectivity of the segmented 2D contours is created from a known correspondence between contour point of the 2D model and respective points and body parts of a mesh applied to the 3D model.

6. The method of claim 1, further comprising the steps of:
creating a training set of clothing outlines and an underlying naked human body represented by the 2D contours, and
determine a displacement between contour points of the clothing outlines and the underlying naked human body to define a clothing deformation.
7. The method of claim 6, wherein the clothing deformation represents a deformation (scale and rotation) of the line segments of 2D contours of the human body.
8. The method of claim 6, wherein the clothing deformation obtained from the training set represents “eigen-clothing”.
9. The method of claim 8, wherein “eigen-clothing” is separated into “eigen-separates”, with each “eigen-separate” covering a part of the 2D contours of the human body.
10. The method of claim 6, further comprising the steps of:
estimating clothing deformation coefficients, and
classifying different types of clothing based on the clothing deformation coefficients.
11. The method of claim 8, further comprising the steps of:
subtracting the “eigen-clothing” from a clothed human body, and
estimating therefrom the underlying naked human body.
12. The method of claim 1, further comprising the steps of:
projecting the 3D model into different camera planes to derive 2D contours in at least two different views to obtain a plurality of body silhouettes,
fitting the 2D contour model to each of the plurality of body silhouettes,
determining 2D shape parameters from the factored representation of the non-

rigid deformations of each of the plurality of body silhouettes, and predicting 3D shape parameters from the determined 2D shape parameters.

13. The method of claim 12, wherein the body silhouettes are obtained for different shapes and poses of the projected 3D model.

14. The method of claim 12, further comprising combining the body silhouettes of a person taken in different poses and different views into an estimate of the 3D shape of the person.

15. The method of claim 12, further comprising the steps of: subtracting "eigen-clothing" from the 2D contour model of a clothed human body to obtain the 2D contour model of an unclothed human body, and predicting 3D shape parameters of the unclothed human body from the determined 2D shape parameters of the clothed human body.

16. A two-dimensional (2D) contour person (CP) model for modeling an unclothed human body, comprising
a shape variation component,
a viewpoint change component, and
a pose variation component,
wherein the contour person (CP) model is learned from a 3D model of the human body that captures natural shape and pose variations, with projected contours of this model, along with their segmentation into parts, form a training set.

17. The CP model of claim 16, wherein the model visualizes frontal views and non-frontal views of the human body.

18. The CP model of claim 16, wherein the pose variation component includes body parts rotation and foreshortening.

19. The CP model of claim 16, wherein the shape variation component or the pose variation component, or both, capture the non-rigid deformation of the body that occurs with articulation.

20. The CP model of claim 16, wherein the 2D body is factored into

- a linear model characterizing shape change across the population;
- a linear approximation to distortions caused by local camera view changes;
- an articulation of the body parts represented by a rotation and length scaling; and
- a non-rigid deformation associated with the articulation of the parts.

21. A dressed contour person (DCP) model for modeling a clothed human body, comprising
a 2D contour person (CP) model of an unclothed human body according to one of claims 16 to 20,
wherein deformations are computed by aligning the contour of the clothed human body with the contour of the unclothed human body, and
wherein an “eigen-clothing” model is learned using principal component analysis (PCA) applied to these deformations.

22. The dressed contour person (DCP) model of claim 21, wherein the CP model is “clothed” by defining a set of linear coefficients that produce from the naked contour a deformation associated with a clothing type and therefrom model of the clothed human body.

23. The dressed contour person (DCP) model of claim 21, wherein the “eigen-clothing” is used to classify different types of clothing.

24. The dressed contour person (DCP) model of claim 21, and further deforming the “eigen-clothing” commensurate with shape and pose changes of

the 2D CP model.

25. A method for predicting a 3D body shape of a human body from uncalibrated 2D images, comprising the steps of:
taking multiple 2D images of a person at different times and in different poses;
combining the multiple images to estimate a consistent 3D shape by including for each 2D view of a person body shape, body pose and camera view.

26. The method of claim 25, wherein the 2D images of a person taken include at least one of a single image, multiple images taken simultaneously and several non-simultaneous images of subject standing in different poses, with possible different occlusions.

27. The method of claim 25, wherein the 2D images are a sequence of cluttered video images.

28. The method of claim 25, wherein the 2D images are silhouettes.

29. A method for estimating an underlying 2D naked body (CP model) according to claims 16-20 from a DCP model according to claims 21-24, comprising
associating a type of clothing with the clothed body,
modeling clothing deformations,
removing the deformed "eigen-clothing" from the clothed body, and
computing the CP model of the underlying 2D naked body.




Figure 1




Figure 2




Figure 3




Figure 4






Figure 5



(a) mean (b) PC1 (c) PC2 (d) PC4 (e) PC1 (f) PC2 (g) PC8

Figure 6



(a) male (b) compared to GT (c) female (d) compared to GT

Figure 7



Figure 8




Figure 9




Figure 10




Figure 11




Figure 12




Figure 13





Figure 1