

General Automatic Human Shape and Motion Capture Using Volumetric Contour Cues

Helge Rhodin¹ Nadia Robertini^{1, 2} Dan Casas¹
Christian Richardt^{1, 2} Hans-Peter Seidel¹ Christian Theobalt¹

¹MPI Informatik, ²Intel Visual Computing Institute

Abstract. Markerless motion capture algorithms require a 3D body with properly personalized skeleton dimension and/or body shape and appearance to successfully track a person. Unfortunately, many tracking methods consider model personalization a different problem and use manual or semi-automatic model initialization, which greatly reduces applicability. In this paper, we propose a fully automatic algorithm that jointly creates a rigged actor model commonly used for animation – skeleton, volumetric shape, appearance, and optionally a body surface – and estimates the actor’s motion from multi-view video input only. The approach is rigorously designed to work on footage of general outdoor scenes recorded with very few cameras and without background subtraction. Our method uses a new image formation model with analytic visibility and analytically differentiable alignment energy. For reconstruction, 3D body shape is approximated as a Gaussian density field. For pose and shape estimation, we minimize a new edge-based alignment energy inspired by volume ray casting in an absorbing medium. We further propose a new statistical human body model that represents the body surface, volumetric Gaussian density, and variability in skeleton shape. Given any multi-view sequence, our method jointly optimizes the pose and shape parameters of this model fully automatically in a spatiotemporal way.

1 Introduction

Markerless full-body motion capture techniques refrain from markers used in most commercial solutions, and promise to be an important enabling technique in computer animation and visual effects production, in sports and biomechanics research, and the growing fields of virtual and augmented reality. While early markerless methods were confined to indoor use in more controlled scenes and backgrounds recorded with eight or more cameras [1], recent methods succeed in general outdoor scenes with much fewer cameras [2, 3].

Before motion capture commences, the 3D body model for tracking needs to be personalized to the captured human. This includes personalization of the bone lengths, but often also of biomechanical shape and surface, including appearance. This essential initialization is, unfortunately, neglected by many methods and solved with an entirely different approach, or with specific and complex manual or semi-automatic initialization steps. For instance, some methods for motion




Fig. 1. Method overview. Pose is estimated from detections in Stage I, actor shape and pose is refined through contour alignment in Stage II by space-time optimization. Outputs are the actor skeleton, attached density, mesh and motion.

capture in studios with controlled backgrounds rely on static full-body scans [4–6], or personalization on manually segmented initialization poses [7]. Recent outdoor motion capture methods use entirely manual model initialization [3]. When using depth cameras, automatic model initialization was shown [8–12], but RGB-D cameras are less accessible and not usable outdoors. Simultaneous pose and shape estimation from in-studio multi-view footage with background subtraction was also shown [13–15], but not on footage of less constrained setups such as outdoor scenes filmed with very few cameras.

We therefore propose a fully-automatic space-time approach for simultaneous model initialization and motion capture. Our approach is specifically designed to solve this problem automatically for multi-view video footage recorded in general environments (moving background, no background subtraction) and filmed with as few as two cameras. Motions can be arbitrary and unchoreographed. It takes a further step towards making markerless motion capture practical in the aforementioned application areas, and enables motion capture from third-party video footage, where dedicated initialization pose images or the shape model altogether are unavailable. Our approach builds on the following contributions.

First, we introduce a body representation that extends a scene model inspired by light transport in absorbing transparent media [16]. We represent the volumetric body shape by Gaussian density functions attached to a kinematic skeleton. We further define a novel 2D contour-based energy that measures contour alignment with image gradients on the raw RGB images using a new volume raycasting image formation model. We define contour direction and magnitude for each image position, which form a ridge at the model outline, see Figure 1. No explicit background segmentation is needed. Importantly, our energy features analytic derivatives, including fully-differentiable visibility everywhere.

The second contribution is a new data-driven body model that represents human surface variation, the space of skeleton dimensions, and the space of volumetric density distributions optimally for reconstruction using a low-dimensional parameter space.

Finally, we propose a space-time optimization approach that fully automatically computes both the shape and the 3D skeletal pose of the actor using both contour and ConvNet-based joint detection cues. The final outputs are (1) a rigged character, as commonly used in animation, comprising a personalized skeleton and attached surface, along with the (optionally colored) volumetric human shape representation, and (2) the joint angles for each video frame. We tested our method on eleven sequences, indoor and outdoor, showing reconstructions with fewer cameras and less manual effort compared to the state of the art.

2 Related work

Our goal is to fully automatically capture a personalized, rigged surface model, as used in animation, *together with* its sequence of skeletal poses from sparse multi-view video of general scenes where background segmentation is hard. Many multi-view markerless motion capture approaches consider model initialization and tracking separate problems [2]. Even in recent methods working outdoors, shape and skeleton dimensions of the tracked model are either initialized manually prior to tracking [3], or estimated from manually segmented initialization poses [7]. In controlled studios, static shape [17, 18] or dimensions and pose of simple parametric human models [13, 14] can be optimized by matching against chroma-keyed multi-view image silhouettes. Many multi-view performance capture methods [19] deform a static full-body shape template obtained with a full-body scanner [4, 5, 20, 21], or through fitting against the visual hull [22–25] to match scene motion. Again, all these require controlled in-studio footage, an off-line scan, or both. Shape estimation of a naked parametric model in single images using shading and edge cues [26], or monocular pose and shape estimation from video is also feasible [27–29], but require substantial manual intervention (joint labeling, feature/pose correction, background subtraction etc.). For multi-view in-studio setups (3–4 views), where background subtraction works, Bălan et al. [15] estimate shape and pose of the SCAPE parametric body model. Optimization is independent for each frame and requires initialization by a coarse cylindrical shape model. Implicit surface representations yield beneficial properties for pose [30] and surface [31] reconstruction, but do not avoid the dependency on explicit silhouette input. In contrast to all previously mentioned methods, our approach requires no manual interaction, succeeds even with only two camera views, and on scenes recorded outdoors without any background segmentation.

Recently, several methods to capture both shape and pose of a parametric human body model with depth cameras were proposed [9, 11, 32]; these special cameras are not as easily available and often do not work outdoors. We also build up on the success of parametric body models for surface representation, e.g. [29, 33–35], but extend these models to represent the space of volumetric shape models needed for tracking, along with a rigged surface and skeleton.

Our approach is designed to work without explicit background subtraction. In outdoor settings with moving backgrounds and uncontrolled illumination, such segmentation is hard, but progress has been made by multi-view segmentation

[36–38], joint segmentation and reconstruction [39–42], and also aided by propagation of a manual initialization [43, 20]. However, the obtained segmentations are still noisy, enabling only rather coarse 3D reconstructions [42], and many methods would not work with only two cameras.

Edge cues have been widely used in human shape and motion estimation [1, 2, 44–47], but we provide a new formulation for their use and make edges in general scenes the primary cue. In contrast, existing shape estimation works use edges as supplemental information, for example to find self-occluding edges in silhouette-based methods and to correct rough silhouette borders [26]. Our new formulation is inspired by the work of Nagel et al., where model contours are directly matched to image edges for rigid object [48] and human pose tracking [49]. Contour edges on tracked meshes are found by a visibility test, and are convolved with a Gaussian kernel. This approach forms piecewise-smooth and differentiable model contours which are optimized to maximize overlap with image gradients. We advance this idea in several ways: our model is volumetric, analytic visibility is incorporated in the model and optimization, occlusion changes are differentiable, the human is represented as a deformable object, allowing for shape estimation, and contour direction is handled separately from contour magnitude.

Our approach follows the generative analysis-by-synthesis approach: contours are formed by a 3D volumetric model and image formation is an extension of the volumetric ray-tracing model proposed by Rhodin et al. [16]. Many discriminative methods for 2D pose estimation were proposed [50–52]; multi-view extensions were also investigated [46, 53, 54]. Their goal is different to ours, as they find single-shot 2D/3D joint locations, but no 3D rigged body shape and no temporally stable joint angles needed for animation. We thus use a discriminative detector only for initialization. Our work has links to non-rigid structure-from-motion that finds sparse 3D point trajectories (e.g. on the body) from single-view images of a non-rigidly moving scene [55]. Articulation constraints [56] can help to find the sparse scene structure, but the goal is different from our estimation of a fully dense, rigged 3D character and stable skeleton motion.

3 Notation and overview

Input to our algorithm are RGB image sequences $\mathcal{I}_{c,t}$, recorded with calibrated cameras $c = 1, \dots, C$ and synchronized to within a frame (see list of datasets in supplemental document). The output of our approach is the configuration of a virtual actor model $\mathcal{K}(\mathbf{p}_t, \mathbf{b}, \boldsymbol{\gamma})$ for each frame $t = 1, \dots, T$, comprising the per-frame joint angles \mathbf{p}_t , the personalized bone lengths \mathbf{b} , as well as the personalized volumetric Gaussian representation $\boldsymbol{\gamma}$, including color, of the actor.

In the following, we first explain the basis of our new image formation model, the Gaussian density scene representation, and our new parametric human shape model building on it (Section 4). Subsequently, we detail our space-time optimization approach (Section 5) in two stages: (I) using ConvNet-based joint detection constraints (Section 5.1); and (II) using a new ray-casting-based volumetric image formation model and a new contour-based alignment energy (Section 5.2).




Fig. 2. Registration process of the body shape model. Skeleton and Gaussians are once manually placed into the *reference* mesh, vertex correspondence transfers Gaussian- and joint-neighborhood weights (green and red respectively), to register reference bones and Gaussians to all instance meshes.

4 Volumetric statistical body shape model

To model the human in 3D for reconstruction, we build on sum-of-Gaussians representations [7, 16] and model the volumetric extent of the actor using a set of 91 isotropic Gaussian density functions distributed in 3D space. Each Gaussian G_q is parametrized by its standard deviation σ_q , mean location μ_q in 3D, and density c_q , which define the Gaussian shape parameters $\gamma = \{\mu_q, \sigma_q, c_q\}_q$. The combined density field of the Gaussians, $\sum_q c_q G_q$, smoothly describes the volumetric occupancy of the human in 3D space, see Figure 1. Each Gaussian is rigidly attached to one of the bones of an articulated skeleton with bone lengths \mathbf{b} and 16 joints, whose pose is parameterized with 43 twist pose parameters, i.e. the Gaussian position μ_q is relative to the attached bone. This representation allows us to formulate a new alignment energy tailored to pose fitting in general scenes, featuring analytic derivatives and fully-differentiable visibility (Section 5).

In their original work, Rhodin et al. [16] create 3D density models for tracked shapes by a semi-automatic placement of Gaussians in 3D. Since the shape of humans varies drastically, a different distribution of Gaussians and skeleton dimensions is needed for each individual to ensure optimal tracking. In this paper, we propose a method to automatically find such a skeleton and optimal attached Gaussian distribution, along with a good body surface. Rather than optimizing in the combined high-dimensional space of skeleton dimensions, the number of Gaussians and all their parameters, we build a new specialized, low-dimensional parametric body model.

Traditional statistical human body models represent variations in body surface only across individuals, as well as pose-dependent surface deformations using linear [57, 58] or non-linear [33, 28] subspaces of the mesh vertex positions. For our task, we build an enriched statistical body model that parameterizes, in addition to the body surface, the optimal volumetric Gaussian density distribution γ for tracking, and the space of skeleton dimensions \mathbf{b} , through linear functions $\gamma(\mathbf{s})$, $\mathbf{b}(\mathbf{s})$ of a low-dimensional shape vector \mathbf{s} . To build our model, we use an existing database of 228 registered scanned meshes of human bodies in neutral

pose [59]. We take one of the scans as *reference* mesh, and place the articulated skeleton inside. The 91 Gaussians are attached to bones, their position is set to uniformly fill the mesh volume, and their standard deviation and density is set such that a density gradient forms at the mesh surface, see Figure 2 (left). This manual step has to be done only once to obtain Gaussian parameters γ_{ref} for the database reference, and can also be automated by silhouette alignment [16].

The best positions $\{\mu_q\}_q$ and scales $\{\sigma_q\}_q$ of Gaussians γ_i for each remaining database *instance* mesh i are automatically derived by weighted Procrustes alignment. Each Gaussian G_q in the reference has a set of neighboring surface mesh vertices. The set is inferred by weighting vertices proportional to the density of G_q at their position in the reference mesh, see Figure 2 (right). For each Gaussian G_q , vertices are optimally translated, rotated and scaled to align to the corresponding instance mesh vertices. These similarity transforms are applied on γ_{ref} to obtain γ_i , where scaling multiplies σ_q and translation shifts μ_q .

To infer the adapted skeleton dimensions \mathbf{b}_i for each instance mesh, we follow a similar strategy: we place Gaussians of standard deviation 10 cm at each joint in the reference mesh, which are then scaled and repositioned to fit the target mesh using the same Procrustes strategy as before. This yields properly scaled bone lengths for each target mesh.

Having estimates of volume γ_i and bone lengths \mathbf{b}_i for each database entry i , we now learn a joint body model. We build a PCA model on the data matrix $[(\gamma_1; \mathbf{b}_1), (\gamma_2; \mathbf{b}_2), \dots]$, where each column vector $(\gamma_i; \mathbf{b}_i)$ is the stack of estimates for entry i . The mean is the average human shape $(\bar{\gamma}; \bar{\mathbf{b}})$, and the PCA basis vectors span the principal shape variations of the database. The PCA coefficients are the elements of our shape model \mathbf{s} , and hence define the volume $\gamma(\mathbf{s})$ and bone lengths $\mathbf{b}(\mathbf{s})$. Due to the joined model, bone length and Gaussian parameters are correlated, and optimizing \mathbf{s} for bone length during pose estimation (stage I) thus moves and scales the attached Gaussians accordingly. To reduce dimensionality, we use only the first 50 coefficients in our experiments.

To infer the actor body surface, we introduce a volumetric skinning approach. The reference surface mesh is deformed in a free-form manner along with the Gaussian set under new pose and shape parameters. Similar to linear blend skinning [60], each surface vertex is deformed with the set of 3D transforms of nearby Gaussians, weighted by the density weights used earlier for Procrustes alignment. This coupling of body surfaces to volumetric model is as computationally efficient as using a linear PCA space on mesh vertices [34], while yielding comparable shape generalization and extrapolation qualities to methods using more expensive non-linear reconstruction [33], see supplemental document. Iso-surface reconstruction using Marching Cubes would also be more costly [30].

5 Pose and shape estimation

We formulate the estimation of the time-independent 50 shape parameters \mathbf{s} and the time-dependent $43T$ pose parameters $\mathbf{P} = \{\mathbf{p}_1, \dots, \mathbf{p}_T\}$ as a combined space-time optimization problem over all frames $\mathcal{I}_{c,t}$ and camera viewpoints c of

the input sequence of length T :

$$E(\mathbf{P}, \mathbf{s}) = E_{\text{shape}}(\mathbf{s}) + \sum_t \left(E_{\text{smooth}}(\mathbf{P}, \mathbf{s}, t) + E_{\text{pose}}(\mathbf{p}_t) + \sum_c E_{\text{data}}(c, \mathbf{p}_t, \mathbf{s}) \right). \quad (1)$$

Our energy uses quadratic prior terms to regularize the solution: E_{shape} penalizes shape parameters that have larger absolute value than any of the database instances, E_{smooth} penalizes joint-angle accelerations to favor smooth motions, and E_{pose} penalizes violation of manually specified anatomical joint-angle limits. The data term E_{data} measures the alignment of the projected model to all video frames. To make the optimization of Equation 1 succeed in unconstrained scenes with few cameras, we solve in two subsequent stages. In Stage I (Section 5.1), we optimize for a coarse skeleton estimate and pose set without the volumetric distribution, but using 2D joint detections as primary constraints. In Stage II (Section 5.2), we refine this initial estimate and optimize for all shape and pose parameters using our new contour-based alignment energy. Consequently, the data terms used in the respective stages differ:

$$E_{\text{data}}(c, \mathbf{p}_t, \mathbf{s}) = \begin{cases} E_{\text{detection}}(c, \mathbf{p}_t, \mathbf{s}) & \text{for Stage I (Section 5.1)} \\ E_{\text{contour}}(c, \mathbf{p}_t, \mathbf{s}) & \text{for Stage II (Section 5.2)}. \end{cases} \quad (2)$$

The analytic form of all terms as well as the smoothness in all model parameters allows efficient optimization by gradient descent. In our experiments we apply the conditioned gradient descent method of Stoll et al. [7].

5.1 Stage I – Initial estimation

We employ the discriminative ConvNet-based body-part detector by Tompson et al. [50] to estimate the approximate 2D skeletal joint positions. The detector is independently applied to each input frame $\mathcal{I}_{c,t}$, and outputs heat maps of joint location probability $\mathcal{D}_{c,t,j}$ for each joint j in frame t seen from camera c . Importantly, the detector discriminates the joints on the left and right side of the body (see Figure 1). The detections exhibit noticeable spatial and temporal uncertainty, but are nonetheless a valuable cue for an initial space-time optimization solve. The output heat maps are in general multi-modal due to detection ambiguities, but also in the presence of multiple people, e.g. in the background.

To infer the poses \mathbf{P} and an initial guess for the body shape of the subject, we optimize Equation 1 with data term $E_{\text{detection}}$. It measures the overlap of the heat maps \mathcal{D} with the projected skeleton joints. Each joint in the model skeleton has an attached joint Gaussian (Section 4), and the overlap with the corresponding heat map is maximized using the visibility model of Rhodin et al. [16]. We use a hierarchical approach by first optimizing the torso joints, followed by optimizing the limbs; please see the supplemental document for details. The optimization is initialized with the average human shape $(\bar{\gamma}, \bar{\mathbf{b}})$ in T-pose, at the center of the capture volume. We assume a single person in the capture volume; people in the background are implicitly ignored, as they are typically not visible from all cameras and are dominated by the foreground actor.

Please note that bone lengths $\mathbf{b}(\mathbf{s})$ and volume $\gamma(\mathbf{s})$ are determined through \mathbf{s} , hence, Stage I yields a rough estimate of γ . In Stage II, we use more informative image constraints than pure joint locations to better estimate volumetric extent.

5.2 Stage II – Contour-based refinement

The pose \mathbf{P} and shape \mathbf{s} found in the previous stage are now refined by using a new density-based contour model in the alignment energy. This model explains the spatial image gradients formed at the edge of the projected model, between actor and background, and thus bypasses the need for silhouette extraction, which is difficult for general scenes. To this end, we extend the ray-casting image formation model of Rhodin et al. [16], as summarized in the following paragraph, and subsequently explain how to use it in the contour data term E_{contour} .

Ray-casting image formation model. Each image pixel spawns a ray that starts at the camera center \mathbf{o} and points in direction \mathbf{n} . The visibility of a particular model Gaussian G_q along the ray (\mathbf{o}, \mathbf{n}) is defined as

$$\mathcal{V}_q(\mathbf{o}, \mathbf{n}) = \int_0^\infty \exp\left(-\int_0^s \sum_i G_i(\mathbf{o} + t\mathbf{n}) dt\right) G_q(\mathbf{o} + s\mathbf{n}) ds. \quad (3)$$

This equation models light transport in a heterogeneous translucent medium [61], i.e. \mathcal{V}_q is the fraction of light along the ray that is absorbed by Gaussian G_q . The original paper [16] describes an analytic approximation to Equation 3 by sampling the outer integral.

Different to their work, we apply this ray casting model to infer the visibility of the background, $\mathcal{B}(\mathbf{o}, \mathbf{n}) = 1 - \sum_q \mathcal{V}_q(\mathbf{o}, \mathbf{n})$. Assuming that the background is infinitely distant, \mathcal{B} is the fraction of light not absorbed by the Gaussian model:

$$\mathcal{B}(\mathbf{o}, \mathbf{n}) = \exp\left(-\int_0^\infty \sum_q G_q(\mathbf{o} + t\mathbf{n}) dt\right) = \exp\left(-\sqrt{2\pi} \sum_q \bar{\sigma}_q \bar{c}_q\right). \quad (4)$$

This analytic form is obtained without sampling, but rather it stems from the Gaussian parametrization: the density along ray (\mathbf{o}, \mathbf{n}) though 3D Gaussian G_q is a 1D Gaussian with standard deviation $\bar{\sigma}_q = \sigma_q$ and density maximum $\bar{c}_q = c_q \cdot \exp\left(-\frac{(\boldsymbol{\mu}_q - \mathbf{o})^\top (\boldsymbol{\mu}_q - \mathbf{o}) - ((\boldsymbol{\mu}_q - \mathbf{o})^\top \mathbf{n})^2}{2\sigma_q^2}\right)$, and the integral over the Gaussian density evaluates to a constant (when the negligible density behind the camera is ignored). A model visibility example is shown in Figure 3 left.

To extract the contour of our model, we compute the gradient of the background visibility $\nabla \mathcal{B} = \left(\frac{\partial \mathcal{B}}{\partial u}, \frac{\partial \mathcal{B}}{\partial v}\right)^\top$ with respect to pixel location (u, v) :

$$\nabla \mathcal{B} = \mathcal{B} \sqrt{2\pi} \sum_q \frac{\bar{c}_q}{\bar{\sigma}_q} (\boldsymbol{\mu}_q - \mathbf{o})^\top \mathbf{n} (\boldsymbol{\mu}_q - \mathbf{o})^\top \nabla \mathbf{n}. \quad (5)$$

$\nabla \mathcal{B}$ forms a 2D vector field, where the gradient direction points outwards from the model, and the magnitude forms a ridge at the model contour, see Figure 3 center. In (calibrated pinhole) camera coordinates, the ray direction thus depends on the 2D pixel location (u, v) by $\mathbf{n} = \frac{(u, v, 1)^\top}{\|(u, v, 1)\|_2}$ and $\nabla \mathbf{n} = \left(\frac{\partial \mathbf{n}}{\partial u}, \frac{\partial \mathbf{n}}{\partial v}\right)^\top$.




Fig. 3. Contour refinement to image gradients through per-pixel similarity. Contour color indicates direction, green and red energy indicate agreement and disagreement between model and image gradients, respectively. Close-ups highlight the shape optimization: left arm and right leg pose are corrected in Stage II.

In contrast to Rhodin et al.’s visibility model [16], our model is specific to background visibility, but more accurate and efficient to evaluate. It does not require sampling along the ray to obtain a smooth analytic form, has linear complexity in the number of model Gaussians instead of their quadratic complexity, and improves execution time by an order of magnitude.

Contour energy. To refine the initial pose and shape estimates from Stage I (Section 5.1), we optimize Equation 1 with a new contour data term E_{contour} , to estimate the per-pixel similarity of model and image gradients:

$$E_{\text{contour}}(c, \mathbf{p}_t, \mathbf{s}) = \sum_{(u,v)} E_{\text{sim}}(c, \mathbf{p}_t, \mathbf{s}, u, v) + E_{\text{flat}}(c, \mathbf{p}_t, \mathbf{s}, u, v). \quad (6)$$

In the following, we omit the arguments $(c, \mathbf{p}_t, \mathbf{s}, u, v)$ for better readability. E_{sim} measures the similarity between the gradient magnitude $\|\nabla \mathcal{I}\|_2$ in the input image and the contour magnitude $\|\nabla \mathcal{B}\|_2$ of our model, and penalizes orientation misalignment (contours can be in opposite directions in model and image):

$$E_{\text{sim}} = -\|\nabla \mathcal{B}\|_2 \|\nabla \mathcal{I}\|_2 \cos(2\angle(\nabla \mathcal{B}, \nabla \mathcal{I})). \quad (7)$$

The term E_{flat} models contours forming in flat regions with gradient magnitude smaller than $\delta_{\text{low}} = 0.1$:

$$E_{\text{flat}} = \|\nabla \mathcal{B}\|_2 \max(0, \delta_{\text{low}} - \|\nabla \mathcal{I}\|_2). \quad (8)$$

We compute spatial image gradients $\nabla \mathcal{I} = (\frac{\partial \mathcal{I}}{\partial u}, \frac{\partial \mathcal{I}}{\partial v})^\top$ using the Sobel operator, smoothed with a Gaussian ($\sigma = 1.1$ px), summed over the RGB channels and clamped to a maximum of $\delta_{\text{high}} = 0.2$.

Appearance estimation. Our method is versatile: given the shape and pose estimates from Stage II, we can also estimate a color for each Gaussian. This is needed by earlier tracking methods that use similar volume models, but color




Fig. 4. Reconstruction of challenging outdoor sequences with complex motions from only 3-4 views, showing accurate shape and pose reconstruction.

appearance-based alignment energies [16, 7] – we compare against them in our experiments. Direct back-projection of the image color onto the model suffers from occasional reconstruction errors in Stages I and II. Instead, we compute the weighted mean color $\bar{a}_{q,c}$ over all pixels separately for each Gaussian G_q and view c , where the contribution of each pixel is weighted by the Gaussian’s visibility \mathcal{V}_q (Equation 3). Colors $\bar{a}_{q,c}$ are taken as candidates from which outliers are removed by iteratively computing the mean color and removing the largest outlier (in Euclidean distance). In our experiments, removing 50% of the candidates leads to consistently clean color estimates, as shown in Figures 4 and 9.

6 Evaluation

We evaluate our method on 11 sequences of publicly available datasets with large variety, both indoor and outdoor, and show comparisons to state-of-the-art methods (see supplementary document for dataset details). The quality of pose and shape reconstruction is best assessed in the supplemental video, where we also apply and compare our reconstructions to tracking with the volumetric Gaussian representations of Rhodin et al. [16] and Stoll et al. [7].

Robustness in general scenes. We validate the robustness of our method on three outdoor sequences. On the `walk` dataset [3], people move in the background, and background and foreground color are very similar. Our method is nevertheless able to accurately estimate shape and pose across 100 frames from 6 views, see Figure 1. We also qualitatively compare against the recent the model-free method of Mustafa et al. [42]. On the `Cathedral` sequence of Kim et al. [62], they achieve rough surface reconstruction using 8 cameras without the explicit need for silhouettes; in contrast, 4 views and 20 frames are sufficient for us to reconstruct shape and pose of a quick outdoor run, see Figure 4 (top) and supplementary material. Furthermore, we demonstrate reconstruction of complex motions on `Subject3` during a two-person volleyball play from only 3 views and 100 frames, see Figure 4 (bottom). The second player was segmented out during Stage I, but Stage II was executed automatically. Fully automatic model





Fig. 5. Visual comparison of estimated body shapes at the different stages. In each subfigure (from left to right): mean PCA ($\bar{\gamma}$, $\bar{\mathbf{b}}$), Stage I, Stage II and ground-truth shape, respectively.

Table 1. Quantitative evaluation of estimated shapes in different stages and comparison to Guan et al.’s results [26]. We use three body measures (chest, waist and hips, as shown on the right) to evaluate predicted body shapes against the ground truth (GT) captured using a laser scan.

	Chest size [cm]				Waist size [cm]				Hip size [cm]				Height [cm]				
	[26]	Stage I	Stage II	GT	[26]	Stage I	Stage II	GT	[26]	Stage I	Stage II	GT	[26]	Stage I	Stage II	GT	
Pose1	92.7	-	92.8	92.6	79.6	-	82.5	80.2	-	-	98.0	-	-	-	183.2	185.0	
Pose2	87.4	-	91.3	91.6	78.5	-	82.1	79.4	-	-	98.9	-	-	-	181.9	185.0	
Pose3	91.9	-	93.5	91.4	76.9	-	83.2	80.3	-	-	101.3	-	-	-	182.9	185.0	
Subject1	-	92.7	132.6	131.3	-	76.7	127.1	132.7	-	108.2	135.4	136.1	-	187.5	194.2	195.0	
Subject2	-	92.3	99.3	100.1	-	77.3	90.6	96.5	-	92.5	102.3	99.7	-	168.5	162.5	162.0	

and pose estimation are even possible from only two views as we demonstrate on the **Marker** sequence [3], see Figure 6.

Shape estimation accuracy. To assess the accuracy of the estimated actor models, we tested our method on a variety of subjects performing general motions such as walking, kicking and gymnastics. Evaluation of estimated shape is performed in two ways: (1) the estimated body shape is compared against ground-truth measurements, and (2) the 3D mesh derived from Stage II is projected from the captured camera viewpoints to compute the overlap with a manually segmented foreground. We introduce two datasets **Subject1** and **Subject2**, in addition to **Subject3**, with pronounced body proportions and ground-truth laser scans for quantitative evaluation. Please note that shape estimates are constant across the sequence and can be evaluated at sparse frames, while pose varies and is separately evaluated per frame in the sequel.

The shape accuracy is evaluated by measurements of chest, waist and hip circumference. **Subject1** and **Subject2** are captured indoor and are processed using 6 cameras and 40 frames uniformly sampled over 200 frames. **Subject3** is an outdoor sequence and only 3 camera views are used, see Figure 4. All subjects are reconstructed with high quality in shape, skeleton dimensions and color, despite inaccurately estimated poses in Stage I for some frames. We only observed little variation dependent on the performed motions, i.e. a simple walking motion is sufficient, but bone length estimation degrades if joints are not sufficiently articulated during performance. All estimates are presented quantitatively in Table 1 and qualitatively in Figure 5. In addition, we compare against




Fig. 6. We obtained accurate results even using only two views on the **Marker** sequence.




Fig. 7. Overlap of the estimated shape in Stages I and II for an input frame of **Box** (left) and **Walk** (right) sequences [63]. Note how the white area (correct estimated shape) significantly increases between Stage I and II, while blue (overestimation) and red (underestimation) areas decrease.

Table 2. Quantitative evaluation of Figure 7. See Section 6 for definitions of *Precision* and *Recall*.

	Precision	Recall
Walk Stage I	87.43%	87.25%
Walk Stage II	95.18%	86.89%
Box Stage I	93.26%	81.11%
Box Stage II	95.42%	85.28%


Guan et al. [26] on their single-camera and single-frame datasets **Pose1**, **Pose2** and **Pose3**. Stage I requires multi-view input and was not used; instead, we manually initialized the pose roughly, as shown in Figure 8, and body height is normalized to 185 cm [26]. Our reconstructions are within the same error range, demonstrating that Stage II is well suited even for monocular shape and pose refinement. Our reconstruction is accurate overall, with a mean error of only 2.3 ± 1.9 cm, measured across all sequences with known ground truth.

On top of these sparse measurements (chest, waist and hips), we also evaluate silhouette overlap for sequences **Walk** and **Box** of subject 1 of the publicly available **HumanEva-I** dataset [63], using only 3 cameras. We compute how much the predicted body shape overlaps the actual foreground (*precision*) and how much of the foreground is overlapped by the model (*recall*). Despite the low number of cameras, low-quality images, and without requiring background subtraction, our reconstructions are accurate with 95% precision and 85% recall, and improve slightly on the results of Bălan et al. [15]. Results are presented in Figure 7 and Table 2. Note that Stage II significantly improves shape estimation. The temporal consistency and benefit of the model components are shown in the supplemental video on multiple frames evenly spread along the *HumanEva-I* **Box** sequence.

Pose estimation accuracy. Pose estimation accuracy is quantitatively evaluated on the public **HumanEva-I** dataset, where ground-truth data is available, see Table 3. We tested the method on the designated validation sequences **Walk** and **Box** of subject **S1**. Reconstruction quality is measured as the average Euclidean distance of estimated and ground-truth joint locations, frames with ground truth inaccuracies are excluded by the provided scripts.

Table 3. Pose estimation accuracy measured in mm on the HumanEva-I dataset. The standard deviation is reported in parentheses.

Seq.	Trained on	Our	Amin et al. [53]	Sigal et al. [46]	Belagiannis et al. [54]	Elhayek et al. [3]
S1,	general	74.9 (21.9)	—	66	68.3	66.5
Walk	HumanEva	54.6 (24.2)	54.5	—	—	—
S2,	general	59.7 (15.0)	—	—	62.7	60.0
Box	HumanEva	35.1 (19.0)	47.7	—	—	—

**Fig. 8.** Monocular reconstruction experiment. Our reconstruction (right) shows high-quality contour alignment, and improved pose and shape estimates.

Our pose estimation results are on par with state-of-the-art methods with 6–7 cm average accuracy [46, 53, 54, 3]. In particular, we obtain comparable results to Elhayek et al. [3], which however requires a separately initialized actor model. Please note that Amin et al. [53] specifically trained their model on manually annotated sequences of the same subject in the same room. For best tracking performance, the ideal joint placement and bone lengths of the virtual skeleton may deviate from the real human anatomy, and may generally vary for different tracking approaches. To compensate differences in the skeleton structure, we also report results where the offset between ground truth and estimated joint locations is estimated in the first frame and compensated in the remaining frames, reducing the reconstruction error to 3–5 cm. Datasets without ground-truth data cannot be quantitatively evaluated; however, our shape overlap evaluation results suggest that pose estimation is generally accurate. In summary, pose estimation is reliable, with only occasional failures in Stage I, although the main focus of our work is on the combination with shape estimation.

Runtime. In our experiments, runtime scaled linearly with the number of cameras and frames. Contour-based shape optimization is efficient: it only takes 3 seconds per view, totaling 15 minutes for 50 frames and 6 views on a standard desktop machine. Skeleton pose estimation is not the main focus of this work and is not optimized; it takes 10 seconds per frame and view, totaling 50 minutes.

Limitations and discussion. Even though the body model was learned from tight clothing scans, our approach handles general apparel well, correctly reconstructing the overall shape and body dimensions. We demonstrate that even if not all assumptions are fulfilled, our method produces acceptable results, such as for the dance performance *Skirt* of Gall et al. [5] in Figure 9 (top left) that




Fig. 9. In-studio reconstruction of several subjects. Our estimates are accurate across diverse body shapes and robust to highly articulated poses.

features a skirt. However, our method was not designed to accurately reconstruct fine wrinkles, facial details, hand articulation, or highly non-rigid clothing.

We demonstrate fully automatic reconstructions from as few as two cameras and semi-automatic shape estimation using a single image. Fully automatic pose and shape estimates from a single image remains difficult.

7 Conclusion

We proposed a fully automatic approach for estimating the shape and pose of a rigged actor model from general multi-view video input with just a few cameras. It is the first approach that reasons about contours within sum-of-Gaussians representations and which transfers their beneficial properties, such as analytic form and smoothness, and differentiable visibility [16], to the domain of edge- and silhouette-based shape estimation. This results in an analytic volumetric contour alignment energy that efficiently and fully automatically optimizes the pose and shape parameters. Based on a new statistical body model, our approach reconstructs a personalized kinematic skeleton, a volumetric Gaussian density representation with appearance modeling, a surface mesh, and the time-varying poses of an actor. We demonstrated shape estimation and motion capture results on challenging datasets, indoors and outdoors, captured with very few cameras. This is an important step towards making motion capture more practical.

Acknowledgements

We thank PerceptiveCode, in particular Arjun Jain and Jonathan Tompson, for providing and installing the ConvNet detector, Ahmed Elhayek, Jürgen Gall, Peng Guan, Hansung Kim, Armin Mustafa and Leonid Sigal for providing their data and test sequences, The Foundry for license support, and all our actors. This research was funded by the ERC Starting Grant project CapReal (335545).

References

1. Moeslund, T.B., Hilton, A., Krüger, V.: A survey of advances in vision-based human motion capture and analysis. *Computer Vision and Image Understanding* **104**(2) (2006) 90–126
2. Holte, M.B., Tran, C., Trivedi, M.M., Moeslund, T.B.: Human pose estimation and activity recognition from multi-view videos: Comparative explorations of recent developments. *IEEE Journal of Selected Topics in Signal Processing* **6**(5) (Sept 2012) 538–552
3. Elhayek, A., de Aguiar, E., Jain, A., Tompson, J., Pishchulin, L., Andriluka, M., Bregler, C., Schiele, B., Theobalt, C.: Efficient ConvNet-based marker-less motion capture in general scenes with a low number of cameras. In: *CVPR*. (2015) 3810–3818
4. de Aguiar, E., Stoll, C., Theobalt, C., Ahmed, N., Seidel, H.P., Thrun, S.: Performance capture from sparse multi-view video. *ACM Transactions on Graphics* **27**(3) (2008) 98
5. Gall, J., Stoll, C., de Aguiar, E., Theobalt, C., Rosenhahn, B., Seidel, H.P.: Motion capture using joint skeleton tracking and surface estimation. In: *CVPR*. (2009) 1746–1753
6. Zollhöfer, M., Nießner, M., Izadi, S., Rehmann, C., Zach, C., Fisher, M., Wu, C., Fitzgibbon, A., Loop, C., Theobalt, C., Stamminger, M.: Real-time non-rigid reconstruction using an RGB-D camera. *ACM Transactions on Graphics* **33**(4) (2014) 156
7. Stoll, C., Hasler, N., Gall, J., Seidel, H.P., Theobalt, C.: Fast articulated motion tracking using a sums of Gaussians body model. In: *ICCV*. (2011) 951–958
8. Shotton, J., Sharp, T., Kipman, A., Fitzgibbon, A., Finocchio, M., Blake, A., Cook, M., Moore, R.: Real-time human pose recognition in parts from single depth images. *Communications of the ACM* **56**(1) (2013) 116–124
9. Bogo, F., Black, M.J., Loper, M., Romero, J.: Detailed full-body reconstructions of moving people from monocular RGB-D sequences. In: *ICCV*. (2015) 2300–2308
10. Tong, J., Zhou, J., Liu, L., Pan, Z., Yan, H.: Scanning 3D full human bodies using Kinects. *IEEE Transactions on Visualization and Computer Graphics* **18**(4) (2012) 643–650
11. Helten, T., Baak, A., Bharaj, G., Müller, M., Seidel, H.P., Theobalt, C.: Personalization and evaluation of a real-time depth-based full body tracker. In: *3DV*. (2013) 279–286
12. Newcombe, R.A., Fox, D., Seitz, S.M.: DynamicFusion: Reconstruction and tracking of non-rigid scenes in real-time. In: *CVPR*. (2015) 343–352
13. Kakadiaris, I.A., Metaxas, D.: Three-dimensional human body model acquisition from multiple views. *International Journal of Computer Vision* **30**(3) (1998) 191–218
14. Ahmed, N., de Aguiar, E., Theobalt, C., Magnor, M., Seidel, H.P.: Automatic generation of personalized human avatars from multi-view video. In: *ACM Symposium on Virtual Reality Software and Technology*. (2005) 257–260
15. Bălan, A.O., Sigal, L., Black, M.J., Davis, J.E., Haussecker, H.W.: Detailed human shape and pose from images. In: *CVPR*. (2007)
16. Rhodin, H., Robertini, N., Richardt, C., Seidel, H.P., Theobalt, C.: A versatile scene model with differentiable visibility applied to generative pose estimation. In: *ICCV*. (2015)

17. Hilton, A., Beresford, D., Gentils, T., Smith, R., Sun, W.: Virtual people: Capturing human models to populate virtual worlds. In: *Computer Animation*. (1999) 174–185
18. Bălan, A.O., Black, M.J.: The naked truth: Estimating body shape under clothing. In: *ECCV*. (2008) 15–29
19. Theobalt, C., de Aguiar, E., Stoll, C., Seidel, H.P., Thrun, S.: Performance capture from multi-view video. In Ronfard, R., Taubin, G., eds.: *Image and Geometry Processing for 3-D Cinematography*. Springer (2010) 127–149
20. Wu, C., Stoll, C., Valgaerts, L., Theobalt, C.: On-set performance capture of multiple actors with a stereo camera. *ACM Transactions on Graphics* **32**(6) (2013) 161
21. Wu, C., Varanasi, K., Theobalt, C.: Full body performance capture under uncontrolled and varying illumination: A shading-based approach. In: *ECCV*. (2012) 757–770
22. Vlastic, D., Baran, I., Matusik, W., Popović, J.: Articulated mesh animation from multi-view silhouettes. *ACM Transactions on Graphics* **27**(3) (2008) 97
23. Starck, J., Hilton, A.: Model-based multiple view reconstruction of people. In: *ICCV*. (2003) 915–922
24. Ballan, L., Cortelazzo, G.M.: Marker-less motion capture of skinned models in a four camera set-up using optical flow and silhouettes. In: *3DPVT*. (2008)
25. Allain, B., Franco, J.S., Boyer, E.: An efficient volumetric framework for shape tracking. In: *CVPR*. (2015) 268–276
26. Guan, P., Weiss, A., Bălan, A.O., Black, M.J.: Estimating human shape and pose from a single image. In: *ICCV*. (2009) 1381–1388
27. Guo, Y., Chen, X., Zhou, B., Zhao, Q.: Clothed and naked human shapes estimation from a single image. In: *Computational Visual Media*. (2012) 43–50
28. Hasler, N., Ackermann, H., Rosenhahn, B., Thormählen, T., Seidel, H.P.: Multilinear pose and body shape estimation of dressed subjects from image sets. In: *CVPR*. (2010) 1823–1830
29. Jain, A., Thormählen, T., Seidel, H.P., Theobalt, C.: MovieReshape: Tracking and reshaping of humans in videos. *ACM Transactions on Graphics* **29**(5) (2010)
30. Plankers, R., Fua, P.: Articulated soft objects for multi-view shape and motion capture. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **25**(9) (2003) 63–83
31. Ilic, S., Fua, P.: Implicit meshes for surface reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **28**(2) (2006) 328–333
32. Cui, Y., Chang, W., Nöll, T., Stricker, D.: KinectAvatar: Fully automatic body capture using a single Kinect. In: *ACCV Workshops*. (2012) 133–147
33. Anguelov, D., Srinivasan, P., Koller, D., Thrun, S., Rodgers, J., Davis, J.: SCAPE: shape completion and animation of people. *ACM Transactions on Graphics* **24**(3) (2005) 408–416
34. Pishchulin, L., Wuhler, S., Helten, T., Theobalt, C., Schiele, B.: Building statistical shape spaces for 3D human modeling. *arXiv:1503.05860* (2015)
35. Loper, M., Mahmood, N., Black, M.J.: MoSh: Motion and shape capture from sparse markers. *ACM Transactions on Graphics* **33**(6) (2014) 220
36. Campbell, N.D.F., Vogiatzis, G., Hernández, C., Cipolla, R.: Automatic 3D object segmentation in multiple views using volumetric graph-cuts. In: *BMVC*. (2007) 530–539
37. Wang, T., Collomosse, J., Hilton, A.: Wide baseline multi-view video matting using a hybrid Markov random field. In: *ICPR*. (2014) 136–141

38. Djelouah, A., Franco, J.S., Boyer, E., Le Clerc, F., Pérez, P.: Sparse multi-view consistency for object segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **37**(9) (2015) 1890–1903
39. Szeliski, R., Golland, P.: Stereo matching with transparency and matting. In: *ICCV*. (1998) 517–524
40. Guillemaut, J.Y., Hilton, A.: Joint multi-layer segmentation and reconstruction for free-viewpoint video applications. *International Journal of Computer Vision* **93**(1) (2011) 73–100
41. Bray, M., Kohli, P., Torr, P.H.S.: POSECUT: simultaneous segmentation and 3D pose estimation of humans using dynamic graph-cuts. In: *ECCV*. (2006) 642–655
42. Mustafa, A., Kim, H., Guillemaut, J.Y., Hilton, A.: General dynamic scene reconstruction from multiple view video. In: *ICCV*. (2015)
43. Hasler, N., Rosenhahn, B., Thormahlen, T., Wand, M., Gall, J., Seidel, H.P.: Markerless motion capture with unsynchronized moving cameras. In: *CVPR*. (2009) 224–231
44. Deutscher, J., Blake, A., Reid, I.: Articulated body motion capture by annealed particle filtering. In: *CVPR*. (2000) 126–133
45. Sidenbladh, H., Black, M.J.: Learning the statistics of people in images and video. *International Journal of Computer Vision* **54**(1–3) (2003) 183–209
46. Sigal, L., Isard, M., Haussecker, H., Black, M.J.: Loose-limbed people: Estimating 3D human pose and motion using non-parametric belief propagation. *International Journal of Computer Vision* **98**(1) (2012) 15–48
47. Kehl, R., Bray, M., Van Gool, L.: Markerless full body tracking by integrating multiple cues. In: *ICCV Workshop on Modeling People and Human Interaction*. (2005)
48. Kollnig, H., Nagel, H.H.: 3D pose estimation by fitting image gradients directly to polyhedral models. In: *ICCV*. (1995) 569–574
49. Wachter, S., Nagel, H.H.: Tracking of persons in monocular image sequences. In: *Nonrigid and Articulated Motion Workshop*. (1997) 2–9
50. Tompson, J.J., Jain, A., LeCun, Y., Bregler, C.: Joint training of a convolutional network and a graphical model for human pose estimation. In: *NIPS*. (2014) 1799–1807
51. Felzenszwalb, P.F., Huttenlocher, D.P.: Pictorial structures for object recognition. *International Journal of Computer Vision* **61**(1) (2005) 55–79
52. Andriluka, M., Roth, S., Schiele, B.: Pictorial structures revisited: People detection and articulated pose estimation. In: *CVPR*. (2009) 1014–1021
53. Amin, S., Andriluka, M., Rohrbach, M., Schiele, B.: Multi-view pictorial structures for 3D human pose estimation. In: *BMVC*. (2013)
54. Belagiannis, V., Amin, S., Andriluka, M., Schiele, B., Navab, N., Ilic, S.: 3D pictorial structures for multiple human pose estimation. In: *CVPR*. (2014) 1669–1676
55. Park, H.S., Shiratori, T., Matthews, I., Sheikh, Y.: 3D trajectory reconstruction under perspective projection. *International Journal of Computer Vision* **115**(2) (2015) 115–135
56. Fayad, J., Russell, C., Agapito, L.: Automated articulated structure and 3D shape recovery from point correspondences. In: *ICCV*. (2011) 431–438
57. Allen, B., Curless, B., Popović, Z.: The space of human body shapes: reconstruction and parameterization from range scans. *ACM Transactions on Graphics* **22**(3) (2003) 587–594
58. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: a skinned multi-person linear model. *ACM Transactions on Graphics* **34**(6) (2015) 248

59. Hasler, N., Stoll, C., Sunkel, M., Rosenhahn, B., Seidel, H.P.: A statistical model of human pose and body shape. *Computer Graphics Forum* **28**(2) (2009) 337–346
60. Lewis, J.P., Cordner, M., Fong, N.: Pose space deformation: A unified approach to shape interpolation and skeleton-driven deformation. In: *SIGGRAPH*. (2000) 165–172
61. Cerezo, E., Pérez, F., Pueyo, X., Seron, F.J., Sillion, F.X.: A survey on participating media rendering techniques. *The Visual Computer* **21**(5) (2005) 303–328
62. Kim, H., Hilton, A.: Influence of colour and feature geometry on multi-modal 3D point clouds data registration. In: *3DV*. (2014) 202–209
63. Sigal, L., Bălan, A.O., Black, M.J.: HumanEva: Synchronized video and motion capture dataset and baseline algorithm for evaluation of articulated human motion. *International Journal of Computer Vision* (87) (2010) 4–27