



HAL
open science

Learning and Tracking the 3D Body Shape of Freely Moving Infants from RGB-D sequences

Nikolas Hesse, Sergi Pujades, Michael J Black, Michael Arens, Ulrich Hofmann, Sebastian Schroeder

► **To cite this version:**

Nikolas Hesse, Sergi Pujades, Michael J Black, Michael Arens, Ulrich Hofmann, et al.. Learning and Tracking the 3D Body Shape of Freely Moving Infants from RGB-D sequences. IEEE Transactions on Pattern Analysis and Machine Intelligence, 2020, 42 (10), pp.2540-2551. 10.1109/TPAMI.2019.2917908 . hal-02162171

HAL Id: hal-02162171

<https://inria.hal.science/hal-02162171>

Submitted on 21 Jun 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Learning and Tracking the 3D Body Shape of Freely Moving Infants from RGB-D sequences

Nikolas Hesse, Sergi Pujades, Michael J. Black, Michael Arens, Ulrich G. Hofmann,
and A. Sebastian Schroeder

Abstract—Statistical models of the human body surface are generally learned from thousands of high-quality 3D scans in predefined poses to cover the wide variety of human body shapes and articulations. Acquisition of such data requires expensive equipment, calibration procedures, and is limited to cooperative subjects who can understand and follow instructions, such as adults. We present a method for learning a statistical 3D Skinned Multi-Infant Linear body model (SMIL) from incomplete, low-quality RGB-D sequences of freely moving infants. Quantitative experiments show that SMIL faithfully represents the RGB-D data and properly factorizes the shape and pose of the infants. To demonstrate the applicability of SMIL, we fit the model to RGB-D sequences of freely moving infants and show, with a case study, that our method captures enough motion detail for General Movements Assessment (GMA), a method used in clinical practice for early detection of neurodevelopmental disorders in infants. SMIL provides a new tool for analyzing infant shape and movement and is a step towards an automated system for GMA.

Index Terms—body models, data-driven, RGB-D, infants, motion analysis.

1 INTRODUCTION

STATISTICAL parametric models of the human body, such as SCAPE [1] or SMPL [2] describe the geometry of the body surface of an observed population in a low-dimensional space. They are usually learned from dense high quality scans of the surface of the human body. Typically, subjects are instructed to stand in the same pose to simplify the problem of modeling body shape.

Since the pioneering work of Blanz and Vetter [3] on a morphable face model, parametric shape models have been evolved and have found a wide range of applications in computer vision and computer graphics. For example, the low-dimensional representation of the human body surface has played a key role in enabling i) the precise capture of shape and pose of humans in motion from low quality RGB-D sequences [4]; ii) the temporal registration of highly dynamic motions of the human body surface [5]; iii) the prediction of human shape and pose from single RGB images based on deep neural networks [6], [7], [8]; and iv) learning detailed avatars from monocular video [9].

Human movements contain key information allowing to infer internal properties of the observed subject [10]. In our work, we consider the application of *motion analysis*, i.e., the acquisition and quantification of poses an observed subject strikes. Human motion analysis is used in medicine

for patient monitoring [11], quantifying therapy or disease progression [12], or performance assessment [13]. Most interestingly, it can be applied to the early detection of neurodevelopmental disorders like cerebral palsy (CP) in infants at a very early age. The General Movements Assessment (GMA) enables trained experts to detect CP at an age of 2 to 4 months, based on assessing the movement quality of infants from video recordings [14]. Infants with abnormal movement quality have very high risk of developing CP or minor neurological dysfunction [15]. While GMA is the most accurate clinical tool for early detection of CP, it is dependent on trained experts and is consequently subject to human perceptual variability. GMA experts require regular practice and re-calibration to assure accurate ratings. Automation of this analysis could reduce this variability and dependence on human judgment. To allow GMA automation, a practical system must first demonstrate that it is capable of capturing the relevant information needed for GMA. Moreover, to allow its widespread use, the solution needs to be seamlessly integrated into the clinical routine. Approaches aimed at GMA automation have relied on wearable sensors or vision-based systems for capturing infant motion. For a review of existing methods, we refer the reader to [16] and [17].

Inspired by previous work on capturing motion from RGB-D data using a body model [4], we follow this direction in our work. Two main problems arise on the way to capture *infant* motion using a body model. The first problem is that there is no realistic infant body model. While parametric body models like SMPL [2] cover a wide variety of adult body shapes, the shape space does not generalize to the new domain of infant bodies (see Fig. 1a). As the body part dimensions between infants and adults vary significantly, the goal of this work is to learn an infant body model that faithfully captures the shape of infants (see Fig. 1b).

However, most statistical models are learned from high-

- N. Hesse and M. Arens are with the Fraunhofer Institute of Optronics, System Technologies and Image Exploitation, Ettlingen, Germany. E-mail: nikolas.hesse@iosb.fraunhofer.de
- S. Pujades is with Univ. Grenoble Alpes, Inria, CNRS, GrenobleINP, LJK, Grenoble, France. The work was done at Max Planck Institute for Intelligent Systems. E-mail: sergi.pujades-rocamora@inria.fr.
- M. J. Black is with Max Planck Institute for Intelligent Systems, Tübingen, Germany.
- U. G. Hofmann is with University Medical Center Freiburg, Faculty of Medicine, University of Freiburg, Germany.
- A. S. Schroeder is with Ludwig Maximilian University, Hauner Children's Hospital, Munich, Germany.

Manuscript revised March 18, 2019




Fig. 1. (a) Simply scaling the SMPL adult body model and fitting it to an infant does not work as body proportions significantly differ. (b) The proposed SMIL model properly captures the infants' shape and pose.

quality scans, which are expensive, and demand cooperative subjects willing to follow instructions. This is the second problem we face: there is no repository of high quality infant 3D body scans from which we could learn the statistics of infant body shape. Acquiring infant shape data is not straightforward, as one needs to comply with strict ethics rules as well as an adequate environment for the infants. Therefore, we acquire sequences of moving infants in a children's hospital. We use one RGB-D sensor and a laptop as this provides a low-cost, easy-to-use alternative to bulky and expensive 3D scanners. Our proposed system produces minimal overhead to the standard examination protocol, and does not affect the behavior of the infants. We use the captured RGB-D sequences to learn an infant body model.

Infant RGB-D data poses several challenges. We have to deal with incomplete data, i.e., partial views, where large parts of the body are occluded most of the time. The data is of low quality and noisy, and captured subjects are not able to follow instructions and take predefined poses.

Contributions. We present the first work on 3D shape and 3D pose estimation of infants, as well as the first work on learning a statistical 3D body model from low-quality, incomplete RGB-D data of freely moving humans. We contribute (i) a new statistical *Skinned Multi-Infant Linear* model (SMIL), learned from 37 RGB-D low-quality sequences of freely moving infants, and (ii) a method to register the SMIL model to the RGB-D sequences, capable of handling severe occlusions and fast movements. Quantitative experiments show how SMIL properly factorizes the pose and the shape of the infants, and allows the captured data to be accurately represented in a low-dimensional space. With a case-study involving a high-risk former preterm study population, we demonstrate that the amount of motion detail captured by SMIL is sufficient to enable accurate GMA ratings by humans. Thus, SMIL provides a fundamental tool that can form a component in an automated system for the assessment of GMs. We make SMIL available to the community for research purposes¹. This article is an extended version of [18].

1. Model available at <http://s.fhg.de/smil>

2 RELATED WORK

We review two main areas: the creation of statistical models of the human body surface and the estimation of shape and pose of humans in movement from RGB-D sequences.

2.1 Human Surface Models

Statistical parametric models of the human body surface are usually based on the *Morphable Model* idea [3], stating that a single surface representation can morph and explain the different samples in a population. These models can be intuitively viewed as a mathematical function taking shape and pose parameters as input and returning a surface mesh as output. The *shape space* and the *pose prior*, i.e., the statistics of most plausible poses, are learned by registering the single surface representation, i.e., a template surface, to real-world data. The shape space and the pose prior allow a compact representation of the human body surface describing the geometry of the human body surface of an observed population in a low-dimensional space. Existing models have been learned from different real-world data. For example, to model the variation of faces De-Carlo et al. [19] learn a model from a cohort of anthropometric measurements. Blanz and Vetter [3], use dense geometry and color data to learn their face model. Allen et al. [20] use the CAESAR dataset to create *the space of human body shapes* by using the geometry information as well as sparse landmarks identified on the bodies. Similarly, Seo et Magnenat-Thalmann [21] learn a shape space from high quality range data. Angelov et al. propose SCAPE [1], a statistical body model learned from high quality scan data, which does not only contain the shape space, but also accounts for the pose dependent deformations, i.e., the surface deformations that a body undergoes when different poses are taken. SCAPE and all successive statistical models of the human body surface [2], [22], [23], [24], [25] or their soft tissue dynamics [26], [27] have been learned from a relatively large number of high quality range scans of adult subjects. Adults, in contrast to infants, are typically cooperative and can be instructed to strike specific poses during scanning.

Animal shape modeling methods face a similar difficulty as ours: live animals are generally difficult to instruct and their motions make them difficult to scan. Thus, they provide a source of inspiration to create models without a large cohort of high quality scans. Cashman and Fitzgibbon [28] learn a deformable model of dolphins by using manually annotated 2D images. Kanazawa et al. [29] learn the deformations and the stiffness of the parts of a 3D mesh template of an animal from manually annotated 2D images. To create the animal model SMAL [30], Zuffi et al. circumvent the difficulty to instruct and scan real animals by using high quality scans of toy figurines. They fit their model to new animals using a set of multi-view images with landmarks and silhouette annotations.

In this work, we learn our statistical model of the shape of infants from low-quality RGB-D sequences. While there are several methods that fit body shape to RGB-D data, we do not know of any that estimates a statistical model from such input. Our method does not rely on manually annotated landmarks and leverages the ones that are automatically extracted from RGB images [31], [32], [33].

2.2 Capturing motion from RGB-D using a body model

The existing body models have proven to be successful in capturing the pose and shape of a subject from RGB-D sequences. The model is parameterized with the shape, giving information about the joint locations, and the pose, defining the angles between limbs. Once a model is registered to the input data, one can obtain the desired motion information. As the joint locations depend on the shape, the closer the estimated body shape is to the actual subject’s shape, the better, i.e., more accurate, the tracking of motions will be.

Parametric body models generally model the space of human bodies without clothing or hair. Some approaches propose to overcome the discrepancy between such a shape space and the real world by creating personalized avatars from the input data and then registering them to the dynamic sequences by keeping the personalized shape fixed. This creates an additional step and usually requires cooperative subjects to take predefined poses. Personalized avatars have been created from multiple Kinect scans [34], [35], Kinect fusion scans [36], [37] or laser scanners [38] and the obtained avatars are then registered to different RGB-D scans of the same person in different poses.

Other methods use a parametric body model to capture the pose without the preliminary step of the personalized shape creation. Ganapathi et al. [39] use a simplistic body model for real-time pose tracking from range data. Ye and Yang [40] introduce a method for real-time shape and pose tracking from depth. They register an articulated deformation model to point clouds within a probabilistic framework. Yu et al. [41] capture non-rigid surface motion by combining a parametric body model to represent the inner body shape with a freely deformable outer surface layer capturing surface details.

Bogo et al. [4] fit a multi-resolution body model to Kinect sequences. This work is the closest to ours, which is why we give a brief summary and identify similarities and differences. The authors aim at creating highly realistic textured avatars from RGB-D sequences. They capture shape, pose and appearance of freely moving humans based on a parametric body model, which is learned from 1800 high-quality 3D scans of 60 adults. They create a personalized shape for each sequence by accumulating shape information over the sequence in a “fusion cloud”. The captured subjects wear tight clothing and take a predefined pose at the beginning of each sequence, as is common in scanning scenarios. They use their body model at different resolutions in a coarse-to-fine manner to increasingly capture more details. They use a displacement map to represent fine details that lie beyond the resolution of the model mesh. The model contains a head-specific shape space for retrieving high-resolution face appearance.

In our work, we also capture shape and pose of sequences containing unconstrained movements. To create a personalized shape, we also merge all temporal information into fusion clouds. In the gradient-based optimization, some of our energy terms are similar to the ones from [4]. In contrast to their work, an initial infant body model is not available and we must create it. We adapt an existing adult body model to a different domain, namely infants, to use as an initial model for registering our infant sequences.

The fact that infants lie in supine position in our scenario presents two different constraints. First, it means that very few backs are visible and we have to deal with large areas of missing data in our fusion clouds. Second, as infants are in contact with the background, i.e., the examination table, we can not rely on a background shot to segment the relevant pointcloud. When the infants move they wrinkle the towel they are lying on with their hands and feet. However, we can take advantage of the planar geometry to fit a plane to the table data to segment it. Moreover, we can (and do) use the fitted plane as a geometric constraint, as we know the back of the infants can not be inside the examination table. Also, in contrast to [4], we can not rely on predefined poses for initialization since the infants are too young to strike poses on demand. We contribute a new automatic method for choosing the best poses for initialization. Moreover, the clothing in our setting is not constrained: we have to deal with diapers, onesies and tights. In particular, diapers pose a challenge since their shape largely deviates from the human body. We handle the unconstrained cloth condition by segmenting the points corresponding to clothes, and by introducing different energy terms for clothing and skin parts. Finally, in our work we do not use the appearance of the surface, but rather use the RGB information to extract 2D landmark estimates to have individual constraints on the face and hand rotations.

3 LEARNING THE SKINNED MULTI-INFANT LINEAR MODEL FROM LOW QUALITY RGB-D DATA

Learning a body model from data is a chicken-and-egg problem. We need a model to register the data to a common topology, and we need registered, or aligned, meshes to learn a model. Since no infant body model is available, we first create an initial infant model by adapting the adult SMPL model [2] (see Sec. 3.2). We then register this initial model to RGB-D sequences of moving infants (Sec. 3.3). To mitigate the incompleteness of data due to the monocular setup, we accumulate shape information from each sequence into one “personalized shape” (Sec. 3.6). Finally, we learn a new infant shape space from all personalized shapes, as well as a new prior over plausible infant poses from our registrations (Sec. 3.7). An overview of the complete learning pipeline is given in Fig. 2.

3.1 Data

There are multiple reasons why no public repository of infant 3D scans exists. Protection of privacy of infants is more strict as compared to adults. The high cost of 3D scanners prevents them from being widespread. Creating a scanning environment that takes into consideration the special care required by infants, like warmth and hygiene, requires additional effort. Finally, infants can not be instructed to strike poses on demand, which is usually required in standard body scanning protocols. RGB-D sensors offer a cheap and lightweight solution for scanning infants, only requiring the sensor and a connected laptop. The data used to learn SMIL was obtained by setting up recording stations at a children’s hospital where infants and parents regularly visit for examinations. The acquisition protocol was integrated in the doctor’s medical routine in order to minimize overhead.

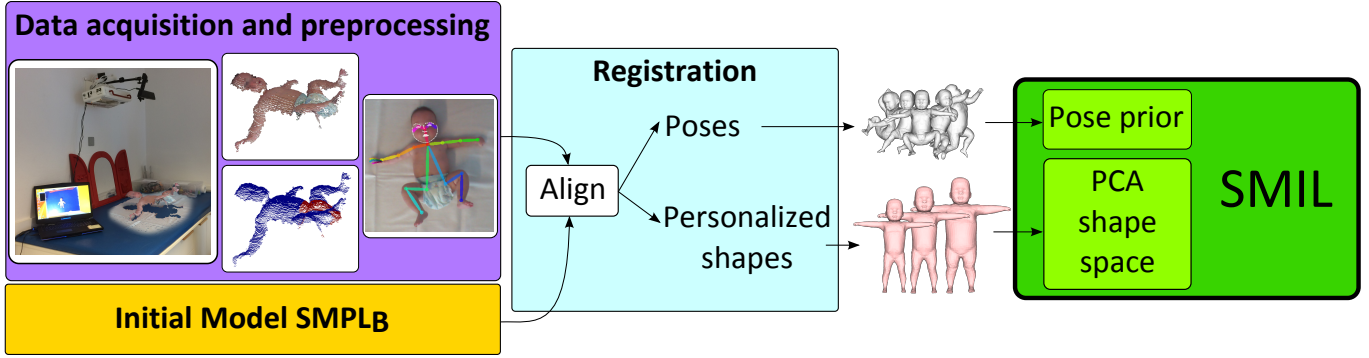


Fig. 2. Skinned Multi-Infant Linear (SMIL) model creation pipeline. We create an initial infant model based on SMPL. We perform background and clothing segmentation of the recorded sequences in a preprocessing step, and estimate body, face, and hand landmarks in RGB images. We register the initial model, $SMPL_B$, to the RGB-D data, and create one personalized shape for each sequence, capturing infant shape details outside the SMPL shape space. We learn a new infant specific shape space by performing PCA on all personalized shapes, with the mean shape forming our base template. We also learn a prior over plausible poses from a sampled subset of all poses.

Preprocessing. We transform depth images to 3D point clouds using the camera calibration. To segment the infant from the scene, we fit a plane to the background table of the 3D point cloud using RANSAC [42] and remove all points close to or below the table plane and apply a simple cluster-based filtering. Further processing steps operate on this segmented cloud, in which only points belonging to the infant remain. Plane-based segmentation is not always perfect, e.g., in case of a wrinkled towel very close to the infant body, some noise may remain. However, the registration methods have proven to be robust to outliers of this kind. The estimated table plane will be reused for constraining the infants’ backs in the registration stage (Sec. 3.3).

In order to avoid modeling diapers and clothing wrinkles in the infant shape space, we segment the input point clouds into *clothing* and *skin* using the color information by adapting the method from Pons-Moll et al. [43]. The clothing type determines the number of color modes k and the distribution of clothing and skin probabilities on the model surface. We perform unsupervised k-means clustering to obtain the dominant modes in RGB, and one of k labels is assigned to each scan point. We register the initial model, $SMPL_B$ (Sec. 3.2), to the scan in order to transfer the scan point labels to the model vertices, and solve a minimization problem on a Markov random field defined by the model topology. This gives us a clean segmentation of the scan points belonging to clothing (or diaper) and the ones belonging to the skin. An example of the segmentation result can be seen in the *Data acquisition and preprocessing* box of Fig. 2; blue is skin and red is diaper. To avoid registering all scans twice, i.e., a first rough registration to segment the sequence and a second to obtain the final registration, we transfer the clothing labels from the registration at frame $t-1$ to the point cloud at frame t . In practice this works well, since most of the clothing is on, or close to the trunk, which does not move much in consecutive frames. In sequences with clothing types longarm onesie and tights we do not observe severe labeling failures due to fast motion.

Scanning of adults typically relies on them striking a simple pose to facilitate model fitting and registrations. The scanned infants can not take a predefined pose to facilitate an initial estimate of model parameters. However, existing

approaches on 2D pose estimation from RGB images (for adults) have achieved impressive results. Most interestingly, experiments show that applying these methods to images of infants produces accurate estimates of 2D pose [17]. In order to choose a “good” candidate frame to initialize the model parameters (see Sec. 3.5), we leverage the 2D body landmarks together with their confidence values. From the RGB images we extract body pose [32] as well as face [31] and hand [33] landmarks. We experimentally verify that they provide key information on the head and hand rotations to the registration process, which is complementary to the noisy point clouds.

3.2 Initial model

We manually create an initial model, $SMPL_B$, by adapting the *Skinned Multi-Person Linear model* (SMPL) [2], which we briefly recap. SMPL is a linear model of shape and pose. It represents shape deformations as a combination of identity specific shape (shape blend shapes) and pose-dependent shape (pose blend shapes). The pose blend shapes are learned from 1786 registered scans of adults in predefined poses, while the shape blend shapes are learned from registered scans of 1700 males and 2100 females. The SMPL shape space is represented by a mean template shape and principal shape directions that are created by performing principal component analysis (PCA) on pose-normalized registered meshes. The shape of a body is described by a vector of linear coefficients, β , that multiply the principal component displacements. The model can be viewed as a mapping from shape and pose parameters to a shaped and posed mesh. Shape and pose blend shapes are modeled as vertex offsets, which are added to the mean template, and the result is then transformed by a standard blend skinning function to form the output mesh. The SMPL template consists of 6890 vertices and 23 body joints. Each body joint has 3 degrees of freedom (DoF) resulting, with 3 DoF for global rotation, in 72 pose parameters, θ . SMPL contains a learned joint regressor for computing joint locations from surface vertices.

Adaptation to infants. We manually create an initial infant mesh using makeHuman, an open source software for creating 3D characters. We wish to register SMPL to this

mesh to use this base shape in the SMPL model. Directly registering SMPL to the infant mesh fails due to differences in size and proportions. We make use of the fact that meshes exported from makeHuman share the same topology, independent of shape parameters. We register SMPL to an adult makeHuman mesh, and describe makeHuman vertices as linear combinations of SMPL vertices. This allows us to apply this mapping to the infant mesh and transfer it to the SMPL topology. We then replace the SMPL base adult-shape template with the registered infant mesh.

We further scale the SMPL pose blend shapes, which correct skinning artifacts and pose-dependent shape deformations, to infant size. Specifically, we divided infant height by average adult height and multiply the blend shapes by this factor. We keep the SMPL joint regressor untouched, since it worked well for infants in our experiments.

The initial pose priors, i.e., prior probabilities of plausible poses, are learned from data of adults in upright positions. We used the same pose prior as [4] which is computed from more than 700 sequences of the CMU MoCap data set. These priors can not be directly transferred to lying infants. We experienced that the (adult) shape space of SMPL_B does not properly explain round infant bellies, which is why the spine was bent to create a more infant-like belly. To avoid explaining shape deformations with pose parameters, we penalize the bending of the spine, which is supposed to be relatively straight in supine position. Infant feet are not conditioned on standing poses, which is why we further adjust the prior to allow the ankle to extend more.

3.3 Registration

We register the initial model to the segmented point cloud using gradient-based optimization. The main energy being optimized w.r.t. shape β and pose θ parameters is

$$E(\beta, \theta) = E_{\text{data}} + E_{\text{lm}} + E_{\text{table}} + E_{\text{sm}} + E_{\text{sc}} + E_{\beta} + E_{\theta}. \quad (1)$$

We omit the weight factors λ_{term} associated with E_{term} , $\text{term} \in \{\text{data}, \text{lm}, \text{table}, \text{sm}, \text{sc}, \beta, \theta\}$, for compactness. In the following, we explain each term of the energy in detail.

Data term. The data term E_{data} consists of two terms:

$$E_{\text{data}} = \lambda_{\text{m2s}} E_{\text{m2s}} + E_{\text{s2m}}. \quad (2)$$

E_{m2s} accounts for the distance of the visible model vertices to the scan points and E_{s2m} accounts for the distance of the scan points to the model surface

E_{m2s} can be written as

$$E_{\text{m2s}}(M, P) = \sum_{m \in \text{vis}(M)} \rho(\min_{v \in P} \|(m, v)\|), \quad (3)$$

where M denotes the model surface and P the scan points. $\min_{v \in P} \|(m, v)\|$ is the Euclidean distance between each visible model point and its closest scan point, and ρ is the robust Geman-McClure function [44]. In the preprocessing stage, P is segmented into the scan points belonging to the skin (P_{skin}) and the ones belonging to clothing (P_{cloth}). The function $\text{vis}(M)$ selects the visible model vertices. The visibility is computed using the Kinect V1 camera calibration and the OpenDR renderer [45].

E_{s2m} consists of two terms,

$$E_{\text{s2m}} = \lambda_{\text{skin}} E_{\text{skin}} + \lambda_{\text{cloth}} E_{\text{cloth}}. \quad (4)$$




Fig. 3. a) Original RGB image. b) Skin weights used for weighted PCA. White denotes high weight, red denotes low weight.

E_{skin} enforces the skin points to be close to the model mesh and E_{cloth} enforces the cloth points to be outside the model mesh. The skin term can be written as

$$E_{\text{skin}}(M, P_{\text{skin}}, W) = \sum_{v_i \in P_{\text{skin}}} W_i \rho(\text{dist}(v_i, M)), \quad (5)$$

where $\text{dist}(v_i, M)$ gives the minimal distance of vertex v_i to the model mesh surface M , and W_i are the skin weights associated with v_i , which are computed using the same method as in [46]. These are used to avoid learning clothing wrinkles in the shape space (cf. Sec. 3.7). We assign low weights to clothing points and high weights to skin points with smooth transition weights in between (see Fig. 3) to avoid discontinuities at the clothing borders. The cloth term is divided into two more terms, depending on cloth points lying inside or outside the model mesh:

$$E_{\text{cloth}} = E_{\text{outside}} + E_{\text{inside}}, \quad (6)$$

with E_{outside} applying a Geman-McClure penalty to cloth points lying *outside* the body. This is to keep the model from matching clothing points that are far away from the body. At the same time, nearby points are closely matched.

$$E_{\text{outside}}(M, P_{\text{cloth}}, W) = \sum_{v_i \in P_{\text{cloth}}} \delta(v_i)^{\text{out}} W_i \rho(\text{dist}(v_i, M)), \quad (7)$$

with $\delta(v_i)^{\text{out}}$ an indicator function, returning 1 if v_i lies outside the model mesh, else 0.

The term E_{inside} “pushes” the model *inside* the cloth points. We apply a quadratic penalty to cloth points lying *inside* the body mesh, since clothing is not supposed to penetrate the body at all.

$$E_{\text{inside}}(M, P_{\text{cloth}}, W) = \sum_{v_i \in P_{\text{cloth}}} \delta(v_i)^{\text{in}} W_i \text{dist}(v_i, M)^2, \quad (8)$$

with $\delta(v_i)^{\text{in}}$ an indicator function, returning 1 if v_i lies inside the model mesh, else 0. Only using E_{inside} would favor shrinking the model too much.

Landmark term. Due to the low-quality of depth data, depth-only methods can not reliably capture details like head or hand rotations. However, we can estimate 2D landmark positions from the RGB images and use them as additional constraints in the optimization energy. Body landmarks [32] are used for initialization (Sec. 3.5), whereas face [31] and hand [33] landmarks are used in the registration energy of every frame. In cases where the face

detection [31] fails, mostly profile faces, we use the ear and eye landmarks from the body pose estimation method [32]. These help to guide the head rotation in such extreme cases.

The landmark term E_{lm} is similar to Eq. 2 from [47], where the distances between the 2D landmarks estimated from RGB and the corresponding projections of the 3D model landmarks are measured. Instead of using the body joints, which we use only for guiding the initial fitting (Sec. 3.5), we only use the estimated 2D face landmarks (nose, eyes outlines, mouth outline and ears) as well as the hand landmarks (knuckles). We note the set of all markers as L . The 3D model points corresponding to the above landmarks were manually selected through visual inspection. They are projected into the image domain using the camera calibration matrix in order to compute the final 2D distances to the estimated landmarks. The landmark term is then

$$E_{lm} = \lambda_{lm} \sum_{l \in L} c_l \rho(\|l_M - l_{est}\|), \quad (9)$$

where c_l denotes the confidence of an estimated landmark 2D location l_{est} , and l_M is the projected model landmark location. All confidences from the different methods are in the interval $[0, 1]$, making them comparable in terms of magnitudes.

Table term. The recorded infants are too young to roll over, which is why the back is rarely seen by the camera. However, the table on which the infants lie, lets us infer shape information of the back. We assume that the body can not be inside the table, and that a large part of the back will be in contact with it. The table energy has two terms: E_{in} prevents the model vertices M from lying inside the table, i.e., behind the estimated table plane, by applying a quadratic error term on points lying inside the table. E_{close} acts as a gravity term, pulling the model vertices M towards the table by applying a robust Geman-McClure penalty function to vertices that are close to the table and facing away from the camera. We write the table energy term as

$$E_{table} = \lambda_{in} E_{in} + \lambda_{close} E_{close}, \quad (10)$$

with

$$E_{in}(M) = \sum_{x \in M} \delta^{in}(x) \text{dist}(x, \Pi)^2, \quad (11)$$

and

$$E_{close}(M) = \sum_{x \in M} \delta^{close}(x) \rho(\text{dist}(x, \Pi)), \quad (12)$$

with Π denoting the table. $\delta^{in}(x)$ is an indicator function, returning 1 if x lies inside the table (behind the estimated table plane), or 0 otherwise and $\delta^{close}(x)$ is an indicator function, returning 1 if x is close to the table (distance less than 3 cm) and faces away from the camera, or 0 otherwise.

To account for soft tissue deformations of the back, which are not modeled, we allow the model to virtually penetrate the table. We effectively enforce this by translating the table plane by 0.5 cm, i.e., pushing the virtual table back.

Other terms. Depth data contains noise, especially around the borders. To avoid jitter in the model caused by that noise, we add a temporal pose smoothness term. It avoids important changes in pose unless one of the other terms has strong evidence. The temporal pose smoothness term E_{sm} is the same as in Eq. 21 in [48] and penalizes

large differences between the current pose θ and the pose from the last processed frame θ' . The penalty for model self intersections E_{sc} and the shape prior term E_{β} are the same as in Eq. 6 and Eq. 7 in [47] respectively. Bending the model in unnatural ways might decrease the data term error, which is why the pose prior term keeps the pose parameters in a realistic range. The SMIL pose prior consists of a mean and covariance matrix that were learned from 37,000 sample training poses; these are not used during testing. E_{θ} penalizes the squared Mahalanobis distance between θ and the pose prior, as described in [4].

3.4 Registration Optimization

To compute the registrations of a sequence we start by computing an initial shape using 5 frames. In this first step, we optimize for the shape and pose parameters, β and θ , as well as the global translation t . The average shape parameters from these 5 frames will be kept fixed and used later on as a shape regularizer. Experiments showed that otherwise the shape excessively deforms in order to explain occlusions in the optimization process.

With the initial shape fixed, we compute the poses for all frames in the sequence, i.e., we optimize the following energy w.r.t. pose parameters θ and the global translation t :

$$E(\theta, t) = E_{data} + E_{lm} + E_{table} + E_{sm} + E_{sc} + E_{\theta}. \quad (13)$$

Notice that this energy is equal to Eq. 1 without the shape term E_{beta} , as shape is kept fixed. We denote S_f the computed posed shape at frame f .

In the last step, we compute the registration meshes R_f and allow the model vertices $v \in R_f$ to freely deform to best explain the input data. We optimize w.r.t. v the energy

$$E(v) = E_{data} + E_{lm} + E_{table} + E_{cpl}, \quad (14)$$

where E_{cpl} is a ‘‘coupling’’ term, used to keep the registration edges close to the edges of the initial shape. We use a similar energy term as Eq. 8 in [4]:

$$E_{cpl}(R_f, S_f) = \lambda_{cpl} \sum_{e \in V'} \| (AR)_e - (AS)_e \|_F^2, \quad (15)$$

with V' denoting the edges in the model mesh. AR and AS are edge vectors of the triangles of R_f and S_f , and e indexes the edges. The results of these optimizations are the final registrations.

All energies are minimized using a gradient-based dogleg minimization method [49] with OpenDR [45] and Chumpy [50]. We display registration samples in Fig. 4.

3.5 Initialization

In order to find the global minimum, the optimization needs a good initial estimate. In adult settings, subjects are usually asked to take an easy pose, e.g., T-pose (extended arms and legs), at the start of the recording. Infants are not able to strike poses on demand, which is why we can not rely on a predefined pose.

We automatically find an initialization frame containing an ‘‘easy’’ pose by relying on 2D landmark estimates acquired in the preprocessing stage. We make the assumption that a body segment is most visible if it has maximum 2D



Fig. 4. SMIL registration results. From left to right: RGB, point cloud, point cloud (other view), point cloud with registered SMIL, rendered registration.

length over the complete sequence. Perspective projection would decrease 2D body segment length and therefore visibility. The initialization frame is chosen as

$$f_{\text{init}} = \operatorname{argmax}_f \sum_{s \in S} \operatorname{len}(s, f) c(s, f), \quad (16)$$

where S is the set of body segments, $\operatorname{len}(s, f)$ is the 2D length of the segment s in frame f , and $c(s, f)$ is the estimated confidence of joints belonging to s in frame f .

For f_{init} , we optimize a simplified version of Eq. 1, i.e., the initialization energy

$$E_{\text{init}} = \lambda_{j2d} E_{j2d} + \lambda_{\theta} E_{\theta} + \lambda_{\beta} E_{\beta} + \lambda_a E_a + \lambda_{s2m} E_{s2m} \quad (17)$$

where E_{j2d} is similar to E_{lm} with landmarks being 2D body joint positions. E_{j2d} is needed to guide the initial pose towards the target pose without getting stuck in local minima (due to other error terms). In the sequential registration (Eq. 1), the initial pose is taken as the result from the previous frame and therefore is already close to the target pose, rendering E_{j2d} obsolete. E_{θ} and E_{β} are pose and shape priors as in Eq. 1 with strong weights λ_{θ} and λ_{β} . $E_a(\theta) = \sum_i \exp(\theta_i)$ is an angle limit term for knees and elbows. The minimum of E_{init} provides a coarse estimation of shape and pose, which is refined afterwards. In contrast to [47], we omit the self intersection term, and add a scan-to-mesh distance term E_{s2m} , containing 3D information, while [47] solely relies on 2D information.

3.6 Personalized Shape

To capture the subject specific shape details, we create one personalized shape from each sequence, which we do not restrict to the shape space of the model. We *unpose* a randomly

selected subset of 1000 frames per sequence. The process of unposing changes the model pose to a normalized pose (T-pose) in order to remove variance related to body articulation. For each scan point, we calculate the offset normal to the closest model point. After unposing the model, we add these offsets to create the unposed point cloud for each of the 1000 frames. Since the recorded infants lie on their backs most of the time, the unposed clouds have missing areas on the back side. To take advantage of the table constraint in each frame and sparsely fill the missing areas, we add *virtual points*, i.e., points from model vertices that belong to faces oriented away from the camera, to the unposed cloud. We retain the clothing segmentation labels for all unposed scan points. We call the union of all unposed point clouds including virtual points the *fusion cloud*; cf. [4].

To compute the personalized shape, we uniformly sample 1 million points at random from the fusion cloud and proceed in two stages. First, we optimize $E = E_{\text{data}} + E_{\beta}$ w.r.t. the shape parameters β , and keep the pose θ fixed in the zero pose of the model (T-pose with legs and arms extended). We obtain an initial shape estimate that lies in the shape space of the initial model SMPL_B. Second, we allow the model vertices to deviate from the shape space, but tie them to the shape from the first stage with a coupling term. We optimize $E = E_{\text{data}} + E_{\text{cpl}}$ w.r.t. the vertices.

The clothing segmentation is also transformed to the unposed cloud and therefore, the fusion cloud is labeled into clothing and skin parts. These are used in the data term to enforce clothing points to lie outside the model surface and to avoid learning clothing artifacts in the shape space.

3.7 Learning SMIL shape space and pose prior

We compute the new infant shape space by doing weighted principal component analysis (WPCA) on personalized shapes of all sequences. Despite including the clothing segmentation in the creation of personalized shapes, clothing deformations can not be completely removed and diapers typically tend to produce body shapes with an over-long trunk. The recorded sequences contain infants with long-arm onesies, short-arm onesies, tights, diapers and without clothing. These different clothing types cover different parts of the body. We use a weight value of 1 for clothing points P_{cloth} and a weight value of 3 for skin points P_{skin} in the PCA. We compute weighted PCA with an iterative expectation-maximization approach² and retain the first 20 shape components.

In Sec. 4, we describe the data set from which SMIL is learned, consisting of more than 200K frames. We create a pose data set by looping over all poses of all sequences and only add poses to the set if the dissimilarity to any pose in the set is larger than a threshold. The final set contains 47K poses and is used to learn the new pose prior. As the Gaussian pose prior can not penalize illegal poses, e.g., unnatural bending of knees, we manually add penalties to avoid such poses. The final SMIL model is composed of the shape space, the pose prior, and the base template, which is the mean of all personalized shapes.

2. <https://github.com/jakevdp/wPCA>

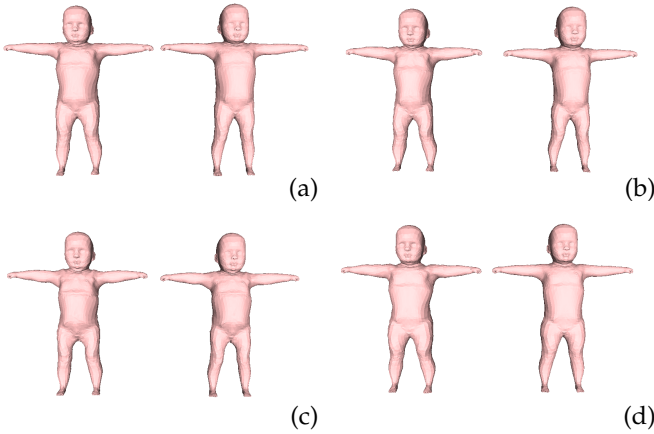


Fig. 5. Visual comparison of four SMPL_B (left sides) and SMIL (right sides) registration samples with all 20 shape components. Most visually significant differences occur at waist, chest, shoulders and head.

3.8 Manual Intervention

In our method we use manual intervention three times: i) to decide which type of clothing the infant is wearing (see Sec 3.1); ii) to generate the initial model SMPL_B (see Sec 3.2) and iii) to define illegal poses in the pose prior. The illegal poses are only defined once and the initial model is no longer used once SMIL is learned. However, given a new sequence, one still needs to manually define the type of clothing: short onesie, long onesie, tights, naked or diapers. Each cloth type defines the corresponding number of color modes and priors to be used. We believe that a classifier predicting the clothing type from RGB images could be learned, making our method fully automatic.

3.9 Method Parameters

The values of the weights in the energy functions were empirically adjusted to keep the different terms balanced. For optimization of the main energy w.r.t. shape and pose parameters (Eq. 1) and the modified energy w.r.t. pose parameters (Eq. 13) we use the weight values: $\lambda_{skin} = 800$, $\lambda_{cloth} = 300$, $\lambda_{m2s} = 400$, $\lambda_{lm} = 1$, $\lambda_{table} = 10000$, $\lambda_{sm} = 800$, $\lambda_{sc} = 1$, $\lambda_{\beta} = 1$ and $\lambda_{\theta} = 0.15$. For optimization of the energy w.r.t. the model vertices (Eq. 14) we use: $\lambda_{skin} = 1000$, $\lambda_{cloth} = 500$, $\lambda_{m2s} = 1000$, $\lambda_{lm} = 0.03$, $\lambda_{table} = 10000$ and $\lambda_{cpl} = 1$. For the creation of the personalized shape (Sec. 3.6), we use: $\lambda_{skin} = 100$, $\lambda_{cloth} = 100$, $\lambda_{\beta} = 0.5$ and $\lambda_{cpl} = 0.4$. Finally, for the initialization energy (Eq. 17), we use: $\lambda_{j2d} = 6$, $\lambda_{\theta} = 10$, $\lambda_a = 30$, $\lambda_{\beta} = 1000$, $\lambda_{s2m} = 30000$. We keep the chosen weights constant for all experiments.

4 EXPERIMENTS

As elaborated in the introduction, gathering high quality 3D scans of infants is highly unpractical, which is why we quantitatively evaluate SMIL and our initial model SMPL_B on the 37 acquired RGB-D sequences of infants. We record the infants using a Microsoft Kinect V1, which is mounted 1 meter above an examination table, facing downwards. All parents gave written informed consent for their child to participate in this study, which was approved by the

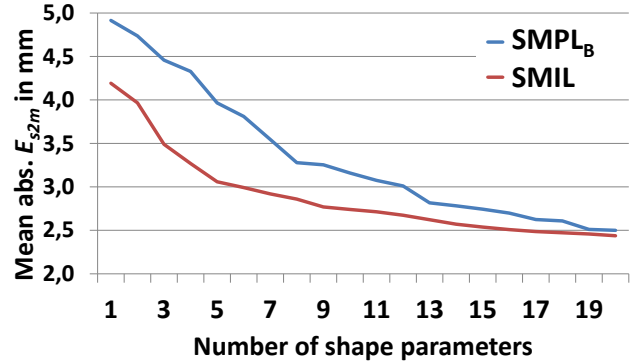


Fig. 6. Average scan-to-mesh error E_{s2m} in mm w.r.t. the number of shape parameters for the two models registered to all fusion scans.

ethics committee of Ludwig Maximilian University Munich (LMU). The infants lie in supine position for three to five minutes without external stimulation, i.e., there is no interaction with caregivers or toys. The recorded infants are between 9 and 18 weeks of corrected age (post term), and their size range is 42 to 59 cm, with an average of 53.5 cm. They wear different types of clothing: none, diaper, onesie shortarm / longarm, or tights. All sequences together sum up to roughly 200K frames, and have an overall duration of over two hours. We evaluate SMIL with a 9-fold cross-validation, using 33 sequences for training the shape space and the pose prior, and 4 for testing. We distribute different clothing styles across all training sets. We measure the distance E_{s2m} (cf. Eq. 4) of the scan to the model mesh by computing the Euclidean distance of each scan point to the mesh surface. For evaluation, we consider all scan points to be labeled as skin, which reduces Eq. 4 to Eq. 5. Note that we do not use the Geman-McClure function ρ here, as we are interested in the actual Euclidean distances.

To compare the SMPL_B shape space to the SMIL shape space, we register both models to each of the 37 fusion scans, using different numbers of shape components. Results are displayed in Fig. 6. We plot average error heatmaps for using the first 1, 3, 5 and all 20 shape components for the registrations (Fig. 7). We observe lower error for SMIL for smaller numbers of shape parameters, and a nearly identical error when using all 20 parameters. The small difference between SMPL_B³ and SMIL poses the question if there is a significant difference between the two. A visual comparison of the personalized shape registration with all 20 shape components shows that SMPL_B shapes display more adult-like features, e.g., a more pronounced waisting, a more “muscular” chest and shoulder region, as well as a smaller (fore)head area (Fig. 5).

To evaluate how well the computed personalized shapes and poses explain the input sequences, we calculate E_{s2m} for all 200K frames. SMIL achieves an average scan-to-mesh distance of 2.51 mm (SD 0.21 mm), SMPL_B has an average E_{s2m} of 2.67 mm (SD 0.22 mm), which is slightly higher than the error of the fusion cloud registrations (cf. Fig. 6). This is in the range of Bogo et al. [4], who report an average

3. Please note that SMPL_B is not the SMPL model [2], but our initial infant model, registered to the data using our method.




Fig. 7. Average error heatmaps for SMIL and SMPL_B on fusion clouds for different numbers of shape components (sc). Top: SMIL. Bottom: SMPL_B. Blue means 0 mm, red means ≥ 10 mm.




Fig. 8. Registration errors. Top: Unnatural pose sample. Bottom: Failure case sample. From left to right: RGB image, 3D point cloud (rotated for improved visibility), overlay of point cloud and registration result, registration result, rendered result from same viewpoint as RGB image.

error between 2.45mm (static scans) and 3.23 mm (arbitrary motions) on adults captured with a Kinect V2.

Due to the lack of ground truth data for evaluation of infant pose correctness, we perform a manual inspection of all sequences to reveal pose errors. We distinguish between “unnatural poses” and “failure cases”. Unnatural poses contain errors in pose, like implausible rotations of a leg (cf. Fig. 8 top row), while the overall registration is plausible, i.e., the 3D joint positions are still at roughly the correct position. Failure cases denote situations in which the optimization gets stuck in a local minimum with a clearly wrong pose, i.e., one model body part registered to a scan part which it does not belong to (cf. Fig. 8 bottom row). We count 16 unnatural leg/foot rotations lasting 41 seconds (= 0.54% of roughly 2 hours) and 18 failure cases (in 7 sequences) lasting 49 seconds (= 0.66% of roughly 2 hours).

To evaluate how well SMIL generalizes to older infants, we register the model to 25 sequences of infants at the age between 21 and 36 weeks, at an average of 26 weeks. The resulting average scan to mesh distance is 2.83 mm (SD: 0.31 mm). With increasing age, infants learn to perform




Fig. 9. Older infant in very challenging pose. (a) RGB image, (b) 3D point cloud (rotated for improved visibility), (c) overlay of point cloud and SMIL registration result, (d) rendered SMIL registration result.

directed movements, like touching their hands, face, or feet, as displayed in Fig. 9. This makes motion capture even more challenging, as standard marker-based methods would not be recommended because of the risk of infants grabbing (and possibly swallowing) markers.

Failure cases. The most common failure is a mixup of feet, i.e., left foot of the model registered to the right foot of the scan and vice versa. Despite our energy having the interpenetration penalty E_{sc} , we observe a few cases where the legs interpenetrate, as in the bottom row in Fig. 8. The registration of all sequences is time consuming (between 10 and 30 seconds per frame) and rerunning the full 200K registrations many times to optimize the parameters is not feasible. The energy term weights are manually selected in order to balance the different terms, and by visually inspecting the results. Further manual adjustment of the E_{sc} weight could fix these rare cases. In the example in the top row of Fig. 8, the right knee is twisted in an unnatural way after the right foot was completely occluded. When the foot is visible again, the pose recovers (knee twisted for 5-6 seconds). Similar to the first failure case, a higher weight on the pose prior would prevent such cases, but finding the perfect weight which completely forbids all illegal poses while allowing all legal poses would require a significant engineering effort.

Motion analysis case study. To show that SMIL captures enough motion information for medical assessment we conduct a case study on GMA. Two trained and certified GMA-experts perform GMA in different videos. We use five stimuli: i) the original RGB videos (denoted by V_{rgb}), and ii) the synthetic registration videos (V_{reg}). For the next three stimuli we use the acquired poses of infants, but we animate a body with a different shape, namely iii) a randomly selected shape of another infant (V_{other}), iv) an extreme shape producing a very thick and large baby (V_{large}), and v) the mean shape (V_{mean}). We exclude three of the 37 sequences, as two are too short and one has non-nutritive sucking, making it non suitable for GMA. As the number of videos to rate is high (34×5), for iv) and v) we only use 50% of the sequences, resulting in 136 videos. For a finer evaluation, we augment GMA classes *definitely abnormal* (DA), *mildly abnormal* (MA), *normal suboptimal* (NS), and *normal optimal* (NO) of [15] into a one to ten scale. Scores 1-3 correspond to DA, 4-5 to MA, 6-7 to NS, and 8-10 to NO. We consider two ratings with an absolute difference ≤ 1 to *agree*, and otherwise to *disagree*.

Rater R_1 is a long-time GMA teacher and has worked on GMA for over 25 years, R_2 has 15 years experience in




Fig. 10. Results of GMA case study. Percentage of ratings of synthetic sequences, generated using SMIL, that *agree* with the reference ratings $R_1 V_{rgb}$ (left) and $R_2 V_{rgb}$ (right), respectively. $V_{\{reg,other,large,mean\}}$ denotes different stimuli.

GMA. Average rating score (and standard deviation) for R_1 is 4.7 (1.4), for R_2 4.0 (1.9). The agreement on original RGB ratings V_{rgb} between R_1 and R_2 is 65%. This further stresses that GMA is challenging and its automation important. In Fig. 10, we present rating differences between synthetic and reference sequences. Each rater is compared to her own V_{rgb} ratings as a reference. $R_1 V_{reg}$ ratings *agree* on 91% of the reference ratings, whereas R_2 achieves an agreement rate of 79%. The agreement decreases more (R_2) or less (R_1) when the motions are presented with a different body shape. We intend to conduct further studies to elucidate the biases introduced by the variation of the infants’ shape.

Generation of realistic data. Human body models have been used to create training data for deep neural networks [51]. We used SMIL to create a realistic (but privacy preserving) data set of *Moving Infants In RGB-D⁴* (MINI-RGBD) [17]. To create the data set, we captured shape and pose of infants from RGB-D sequences as described in Sec. 3.3, but additionally captured texture. We selected random subsets of shapes and textures, and averaged them to create new, synthetic, but realistic shapes and textures. We mapped the real captured poses to the new synthetic infants and extracted ground truth 3D joint positions. We used OpenDR [45] for rendering RGB and depth images to resemble commodity RGB-D sensors. We created the data set with the intention to provide an evaluation set for pose estimation in medical infant motion analysis scenarios. We plan to create a larger data set to provide enough data to train neural networks for infant shape and pose estimation.

5 DISCUSSION / CONCLUSION

Why does it work? Even though each RGB-D frame is a partial observation of the shape, the motions in the different frames of a sequence reveal previously hidden body parts. Moreover, even though the backs of the infants are rarely visible, we can still faithfully restrict the shape by taking into account the background table constraints. This lets us accumulate shape information over sequences. In addition, we leverage 2D pose estimation from RGB in two ways: i) it allows us to add landmark constraints to guide the model where depth-only approaches fail and ii) allows

us to find a good initialization frame and circumvent the need of predefined poses. Key to handling fast motion is a careful balancing of the different error terms. It has to be ensured that large pose changes are allowed, but at the same time, we have to make sure that body parts are not moved to distant scan points in the event of occlusions of corresponding scan parts or noisy segmentation.

Why has it not been done before? Most previous work has been done on adults, who can be instructed and 3D scanned much more easily. Recording infants poses more challenges. As our setup only relies on a low-cost RGB-D camera and a laptop, we can capture infants with great flexibility, for instance at a children’s hospital that parents and children are visiting anyway. The creation of an initial model was not straightforward. Thanks to the flexibility of the SMPL model, we were able to adapt it to infants and get a starting point for capturing real infant shapes.

Limitations. Due to the data resolution, fine structures like finger and toe movements are not captured. Moreover, the model is not yet able to capture facial expressions. In our registration results, we found few cases of unusual neck twists. These occur in frames with side views when there is a mismatch between the computed initial head/neck shape and the head/neck shape in the scan. The model neck tries to match the scan shape by applying a small twist, making the neck look unnatural. By applying a stronger pose prior this could be forbidden, but might at the same time restrict the neck too much in other cases where strong rotations should be allowed. We will rather handle this by creating better initial shapes by sampling from the complete sequence instead of using 5 consecutive frames.

Conclusion. We contribute a method for learning a body model from RGB-D sequences of freely moving infants. We show that our learned Skinned Multi-Infant Linear model factorizes pose and shape and achieves metric accuracy of 2.5 mm. Although the metric difference to our initial model SMPL_B is small, SMIL creates visually more realistic shapes. SMPL_B is a miniaturized adult model that mimics infant appearance while SMIL is characteristic of a real infant population.

We applied our method to the task of medical infant motion analysis. Two expert GMA raters achieve a scoring agreement of 91% and 79% when comparing the GMA ratings from RGB video and from rendered SMIL registration results of the same sequence. Our method is a step towards a fully automated system for GMA, to help early-detect neurodevelopmental disorders like cerebral palsy.

Future work. We have not learned pose-dependent shape deformations for infants, and we reused the scaled down pose blend shapes of SMPL. While these provide sufficient numerical accuracy, we will learn infant pose blend shapes to further increase the realism of SMIL.

In an ongoing study, we are collecting more RGB-D data by taking advantage of the lightweight recording setup: we capture the infants’ motions in their homes to minimize stress and effort for both infants and parents. We will further pursue the clinical goal to create SMIL: the automation of GMA, i.e., learn how to infer GMA ratings from captured motions. We are currently applying the system to infants affected by spinal muscular atrophy, with the goal to quantify the disease progress as well as the impact of therapy.

4. Data set available at <http://s.fhg.de/mini-rgbd>

ACKNOWLEDGMENTS

Authors thank Javier Romero for helpful discussions, Mijna Hadders-Algra and Uta Tacke for GMA ratings and medical insights, R. Weinberger for data acquisition and help with evaluation and Paul Soubiran for the video voice over.

Disclosure: MJB has received research gift funds from Intel, Nvidia, Adobe, Facebook, and Amazon. While MJB is a part-time employee of Amazon, his contribution was performed solely at, and funded solely by, MPI.

REFERENCES

- [1] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, and J. Davis, "SCAPE: Shape completion and animation of people," *ACM Trans. Graph.*, vol. 24, no. 3, pp. 408–416, Jul. 2005.
- [2] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black, "SMPL: A skinned multi-person linear model," *ACM Trans. Graph.*, vol. 34, no. 6, pp. 248:1–248:16, Oct. 2015.
- [3] V. Blanz and T. Vetter, "A morphable model for the synthesis of 3D faces," in *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '99. ACM Press/Addison-Wesley Publishing Co., 1999, pp. 187–194.
- [4] F. Bogo, M. J. Black, M. Loper, and J. Romero, "Detailed full-body reconstructions of moving people from monocular RGB-D sequences," in *2015 IEEE International Conference on Computer Vision (ICCV)*, Dec 2015, pp. 2300–2308.
- [5] F. Bogo, J. Romero, G. Pons-Moll, and M. Black, "Dynamic FAUST: Registering human bodies in motion," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [6] H.-Y. Tung, H.-W. Tung, E. Yumer, and K. Fragkiadaki, "Self-supervised learning of motion capture," in *Advances in Neural Information Processing Systems*, 2017, pp. 5236–5246.
- [7] M. Omran, C. Lassner, G. Pons-Moll, P. Gehler, and B. Schiele, "Neural Body Fitting: Unifying deep learning and model based human pose and shape estimation," in *2018 International Conference on 3D Vision (3DV)*, Sep. 2018, pp. 484–494.
- [8] G. Pavlakos, L. Zhu, X. Zhou, and K. Daniilidis, "Learning to estimate 3D human pose and shape from a single color image," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018, pp. 459–468.
- [9] T. Alldieck, M. Magnor, W. Xu, C. Theobalt, and G. Pons-Moll, "Detailed human avatars from monocular video," in *2018 International Conference on 3D Vision (3DV)*, Sep. 2018, pp. 98–109.
- [10] J. Lu, G. Wang, and P. Moulin, "Human identity and gender recognition from gait sequences with arbitrary walking directions," *IEEE Transactions on Information Forensics and Security*, vol. 9, no. 1, pp. 51–61, 2014.
- [11] F. Achilles, A.-E. Ichim, H. Coskun, F. Tombari, S. Noachtar, and N. Navab, "Patient MoCap: Human pose estimation under blanket occlusion for hospital monitoring applications," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016*. Springer International Publishing, 2016, pp. 491–499.
- [12] P. Kotschieder, J. F. Dorn, C. Morrison, R. Corish, D. Zikic, A. Sellen, M. D'Souza, C. P. Kamm, J. Burggraaff, P. Tewarie *et al.*, "Quantifying progression of multiple sclerosis via classification of depth videos," in *MICCAI*. Springer, 2014, pp. 429–437.
- [13] R. A. Clark, Y.-H. Pua, K. Fortin, C. Ritchie, K. E. Webster, L. Denehy, and A. L. Bryant, "Validity of the Microsoft Kinect for assessment of postural control," *Gait & posture*, vol. 36, no. 3, pp. 372–377, 2012.
- [14] H. Prechtl, "Qualitative changes of spontaneous movements in fetus and preterm infant are a marker of neurological dysfunction," *Early human development*, vol. 23, no. 3, pp. 151–158, 1990.
- [15] M. Hadders-Algra, "General movements: a window for early identification of children at high risk for developmental disorders," *The Journal of pediatrics*, vol. 145, no. 2, pp. S12–S18, 2004.
- [16] C. Marcroft, A. Khan, N. D. Embleton, M. Trenell, and T. Pltz, "Movement recognition technology as a method of assessing spontaneous general movements in high risk infants," *Frontiers in Neurology*, vol. 5, p. 284, 2015.
- [17] N. Hesse, C. Bodensteiner, M. Arens, U. G. Hofmann, R. Weinberger, and A. Sebastian Schroeder, "Computer vision for medical infant motion analysis: State of the art and rgb-d data set," in *Computer Vision – ECCV 2018 Workshops*. Springer International Publishing, 2019, pp. 32–49.
- [18] N. Hesse, S. Pujades, J. Romero, M. J. Black, C. Bodensteiner, M. Arens, U. G. Hofmann, U. Tacke, M. Hadders-Algra, R. Weinberger, W. Müller-Felber, and A. Sebastian Schroeder, "Learning an infant body model from RGB-D data for accurate full body motion analysis," in *Medical Image Computing and Computer Assisted Intervention*. Springer International Publishing, 2018, pp. 792–800.
- [19] D. DeCarlo, D. DeCarlo, D. Metaxas, and M. Stone, "An anthropometric face model using variational techniques," in *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '98. ACM, 1998, pp. 67–74.
- [20] B. Allen, B. Curless, and Z. Popović, "The space of human body shapes: Reconstruction and parameterization from range scans," *ACM Trans. Graph.*, vol. 22, no. 3, pp. 587–594, Jul. 2003.
- [21] H. Seo and N. Magnenat-Thalmann, "An automatic modeling of human bodies from sizing parameters," in *Proceedings of the 2003 symposium on Interactive 3D graphics*. ACM, 2003, pp. 19–26.
- [22] N. Hasler, C. Stoll, M. Sunkel, B. Rosenhahn, and H.-P. Seidel, "A statistical model of human pose and body shape," in *Computer Graphics Forum*, vol. 28, no. 2. Wiley Online Library, 2009, pp. 337–346.
- [23] D. A. Hirshberg, M. Loper, E. Rachlin, and M. J. Black, "Coregistration: Simultaneous alignment and modeling of articulated 3D shape," in *Computer Vision – ECCV 2012*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 242–255.
- [24] Y. Chen, Z. Liu, and Z. Zhang, "Tensor-based human body modeling," in *2013 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2013, pp. 105–112.
- [25] L. Pishchulin, S. Wuhrer, T. Helten, C. Theobalt, and B. Schiele, "Building statistical shape spaces for 3D human modeling," *Pattern Recognition*, vol. 67, pp. 276–286, 2017.
- [26] G. Pons-Moll, J. Romero, N. Mahmood, and M. J. Black, "Dyna: A model of dynamic human shape in motion," *ACM Trans. Graph.*, vol. 34, no. 4, pp. 120:1–120:14, Jul. 2015.
- [27] M. Kim, G. Pons-Moll, S. Pujades, S. Bang, J. Kim, M. J. Black, and S.-H. Lee, "Data-driven physics for human soft tissue animation," *ACM Trans. Graph.*, vol. 36, no. 4, pp. 54:1–54:12, Jul. 2017.
- [28] T. J. Cashman and A. W. Fitzgibbon, "What shape are dolphins? Building 3D morphable models from 2D images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 1, pp. 232–244, 2013.
- [29] A. Kanazawa, S. Kovalsky, R. Basri, and D. Jacobs, "Learning 3D deformation of animals from 2D images," *Computer Graphics Forum*, vol. 35, no. 2, pp. 365–374, 2016.
- [30] S. Zuffi, A. Kanazawa, D. W. Jacobs, and M. J. Black, "3D Menagerie: Modeling the 3D shape and pose of animals," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017, pp. 5524–5532.
- [31] S. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh, "Convolutional pose machines," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 00, June 2016, pp. 4724–4732.
- [32] Z. Cao, T. Simon, S. Wei, and Y. Sheikh, "Realtime multi-person 2D pose estimation using part affinity fields," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017, pp. 1302–1310.
- [33] T. Simon, H. Joo, I. Matthews, and Y. Sheikh, "Hand keypoint detection in single images using multiview bootstrapping," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017, pp. 4645–4653.
- [34] A. Weiss, D. Hirshberg, and M. J. Black, "Home 3D body scans from noisy image and range data," in *2011 IEEE International Conference on Computer Vision (ICCV)*. IEEE, 2011, pp. 1951–1958.
- [35] T. Helten, A. Baak, G. Bharaj, M. Miller, H. Seidel, and C. Theobalt, "Personalization and evaluation of a real-time depth-based full body tracker," in *2013 International Conference on 3D Vision - 3DV 2013*, June 2013, pp. 279–286.
- [36] Q. Zhang, B. Fu, M. Ye, and R. Yang, "Quality dynamic human body modeling using a single low-cost depth camera," in *2014 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014, pp. 676–683.
- [37] Y. Chen, Z. Cheng, C. Lai, R. R. Martin, and G. Dang, "Realtime reconstruction of an animating human body from a single depth camera," *IEEE Transactions on Visualization and Computer Graphics*, vol. 22, no. 8, pp. 2000–2011, Aug 2016.
- [38] G. Ye, Y. Liu, N. Hasler, X. Ji, Q. Dai, and C. Theobalt, "Performance capture of interacting characters with handheld Kinects," in *Computer Vision–ECCV 2012*. Springer, 2012, pp. 828–841.

- [39] V. Ganapathi, C. Plagemann, D. Koller, and S. Thrun, "Real-time human pose tracking from range data," in *Computer Vision – ECCV 2012*. Springer Berlin Heidelberg, 2012, pp. 738–751.
- [40] M. Ye and R. Yang, "Real-time simultaneous pose and shape estimation for articulated objects using a single depth camera," in *2014 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014, pp. 2353–2360.
- [41] T. Yu, Z. Zheng, K. Guo, J. Zhao, Q. Dai, H. Li, G. Pons-Moll, and Y. Liu, "DoubleFusion: Real-time capture of human performances with inner body shapes from a single depth sensor," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018, pp. 7287–7296.
- [42] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.
- [43] G. Pons-Moll, S. Pujades, S. Hu, and M. J. Black, "ClothCap: Seamless 4D clothing capture and retargeting," *ACM Trans. Graph.*, vol. 36, no. 4, pp. 73:1–73:15, Jul. 2017.
- [44] S. Geman and D. E. McClure, "Statistical methods for tomographic image reconstruction," in *Proceedings of the 46th Session of the International Statistical Institute, Bulletin of the ISI*, vol. 52, 1987.
- [45] M. M. Loper and M. J. Black, "OpenDR: An approximate differentiable renderer," in *Computer Vision – ECCV 2014*. Springer International Publishing, 2014, pp. 154–169.
- [46] C. Zhang, S. Pujades, M. Black, and G. Pons-Moll, "Detailed, accurate, human shape estimation from clothed 3D scan sequences," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017, pp. 5484–5493.
- [47] F. Bogo, A. Kanazawa, C. Lassner, P. Gehler, J. Romero, and M. J. Black, "Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image," in *Computer Vision – ECCV 2016*. Springer International Publishing, Oct. 2016.
- [48] J. Romero, D. Tzionas, and M. J. Black, "Embodied Hands: Modeling and capturing hands and bodies together," *ACM Trans. Graph.*, vol. 36, no. 6, pp. 245:1–245:17, Nov. 2017.
- [49] J. Nocedal and S. J. Wright, "Numerical optimization," 2006.
- [50] M. Loper, "Chumpy." [Online]. Available: <http://chumpy.org>
- [51] G. Varol, J. Romero, X. Martin, N. Mahmood, M. J. Black, I. Laptev, and C. Schmid, "Learning from synthetic humans," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, Jul. 2017.



Nikolas Hesse has received his Diploma in computer science from Karlsruhe Institute of Technology (KIT) in 2011. From 2011 to 2013 he worked as a software developer. Since 2013, he has been a research associate at Fraunhofer Institute of Optronics, System Technologies and Image Exploitation (IOSB), and a Ph.D. student at University of Freiburg. His research interests include human motion capture, body modeling, RGB-D vision and the intersection of computer vision and medicine.



Sergi Pujades graduated both in Applied Mathematics at FME-UPC in Barcelona, Spain and in Computer Science at ENSIMAG, Grenoble, France in 2005. After working as a research engineer at INRIA Grenoble he joined Binocle3D (2007) where he managed the R&D Software Department (2009 -2011). In 2015 he received his PhD from University Grenoble Alpes for his work on camera models and algorithms for the creation of video content. Between 2016 and 2018 he was a Post-Doctoral Researcher at Max

Planck Institute for Intelligent Systems in Tübingen, Germany. Since December 2018 he is an associate professor at University Grenoble Alpes, and a researcher at the Morpheo team, a joint research team of Inria Grenoble and Laboratoire Jean Kuntzmann (LJK), in France.



Michael J. Black (Senior Member) received his B.Sc. from the University of British Columbia (1985), his M.S. from Stanford (1989), and his Ph.D. in computer science from Yale University (1992). After post-doctoral research at the University of Toronto, he joined the Xerox Palo Alto Research Center in 1993 where he managed the Image Understanding Area and founded the Digital Video Analysis group. From 2000 to 2010 he was on the faculty of Brown University in the Department of Computer Science (Assoc. Prof. 2000-2004, Prof. 2004-2010). Since 2011 he is a founding director at the Max Planck Institute for Intelligent Systems in Tübingen, Germany, where he leads the Perceiving Systems department. He is also a Distinguished Amazon Scholar, an Honorarprofessor at the University of Tübingen, and Adjunct Professor at Brown University. He is a member of the Royal Swedish Academy of Science and a recipient of the Koenderink Prize and the Helmholtz Prize. He co-founded Body Labs Inc., which was acquired by Amazon.



Michael Arens received his diploma in Computer Science and his PhD (Dr.rer.nat.) from the University of Karlsruhe in 2001 and 2004, respectively. Since 2006, he works at the Fraunhofer Institute of Optronics, System Technologies and Image Exploitation in various positions. In 2011 Dr. Arens was appointed Head of the Department of Object Recognition. His research interests include image and video analysis with an emphasis on the automatic conceptualization and reasoning on results obtained from such an analysis. Dr. Arens is a member of the German Association of Pattern Recognition (DAGM) and member of the IEEE.



Ulrich G. Hofmann received the Ph.D. degree in technical physics and physics from the Technical University of Munich. After post-doctoral periods in Turku, Finland, and the Caltech, Pasadena, he was habilitated at the University of Lübeck at the Institute for Signal Processing and held the Endowed Peter-Osyпка Professorate at the Medical Center University of Freiburg. He is currently a Professor and the Head of the Section for Neuroelectronic Systems. He uses image processing methods to improve outcome in biomedical engineering experiments. His research interest has long centered on brain-machine-interfaces in rodent models. He is now shifting his focus to diverse physical means to stimulate the nervous system. He is on the editorial board of *Frontiers in Neuroscience: Neural Technology* and has chaired the inaugural Gordon Research Conference on Neuroelectronic Interfaces in 2018.



A. Sebastian Schroeder studied medicine in Greifswald, Tübingen, Freiburg, Dresden and Munich. Since then, he has been working as a Pediatrician at the Dr. von Hauner Children's Hospital, University of Munich. Dr. Schroeder has treated children with spasticity and dystonia including botulinum toxin since 2003. Since 2014 he is head of the movement disorder outpatient clinic at the iSPZ Hauner. His research foci are: (i) musculoskeletal ultrasound, (ii) focal effects of botulinum toxin in healthy and spastic adult and pediatric skeletal muscle, (iii) robotic gait-therapy in children, and (iv) movement analysis in infancy. He is a member of the American Academy of Cerebral Palsy and Developmental Medicine, the German Society of Pediatric Neurology and lecturer of the University of Munich.