

Modeling and Estimation of Nonlinear Skin Mechanics for Animated Avatars

Cristian Romero

Miguel A. Otaduy

Dan Casas

Jesus Perez

Universidad Rey Juan Carlos, Madrid, Spain



Figure 1: Our soft-tissue avatar combines a statistical model and an FEM simulation. Our custom constitutive material produces highly dynamic effects and realistic external interactions. We characterize material parameters to match 4D captures using numerical optimization.

Abstract

Data-driven models of human avatars have shown very accurate representations of static poses with soft-tissue deformations. However they are not yet capable of precisely representing very nonlinear deformations and highly dynamic effects. Nonlinear skin mechanics are essential for a realistic depiction of animated avatars interacting with the environment, but controlling physics-only solutions often results in a very complex parameterization task. In this work, we propose a hybrid model in which the soft-tissue deformation of animated avatars is built as a combination of a data-driven statistical model, which kinematically drives the animation, an FEM mechanical simulation. Our key contribution is the definition of deformation mechanics in a reference pose space by inverse skinning of the statistical model. This way, we retain as much as possible of the accurate static data-driven deformation and use a custom anisotropic nonlinear material to accurately represent skin dynamics. Model parameters including the heterogeneous distribution of skin thickness and material properties are automatically optimized from 4D captures of humans showing soft-tissue deformations.

CCS Concepts

• Computing methodologies → Physical simulation;

1. Introduction

Soft-tissue deformation is essential for a realistic depiction of animated characters. The human body deforms due to its own move-

ment and its interaction with the environment, creating rich and expressive effects. Not only are soft-tissue deformations unique for every individual, but they also produce highly nonlinear forces in

response to the interaction with surrounding objects. A jiggling belly, the quick transition of skin from soft to stiff when we pull from it, or the bulging induced by tight apparel, are familiar examples to all of us. Finding accurate and inexpensive methods for the animation of soft-tissue characters in highly dynamic and contact-intensive scenarios has been a long-term goal in the Computer Graphics community, with numerous applications in VFX, video games or garment design, among others. In this paper, we present a model to animate personalized characters with rich soft-tissue dynamics produced by both skeletal motion and external interactions.

There are two distinct approaches to produce expressive animated characters: physics-based and data-driven. Physics-based solutions model the mechanical response of flesh, and simulate deformations by solving the equations of elasticity [LYWG13, KIL*16]. Their advantage is the ability to respond in a plausible way to arbitrary interactions; they are fully generative by construction. Their challenge, on the other hand, is to design materials that mimic the properties of skin, and to parameterize them to match the behavior of each individual. Data-driven solutions learn nonlinear mappings from skeletal motion to soft-tissue deformation [LMR*15, PMRMB15, CO18]. Their advantage is the ability to reproduce observed data with high accuracy, thanks to high-dimensional parameterizations and rich nonlinear building blocks, not constrained by physics laws. Their challenge, on the other hand, is the lack of response to external interactions or, more generally, the need for large training data and the difficulty to generalize to unseen situations.

We propose a model for avatar animation that combines the advantages of physics-based and data-driven approaches. The model responds to external interactions, yet it reproduces accurately observed deformations. The high potential of the combined approach is evidenced by the work of Kim et al. [KPMP*17]. However, our deformation model enjoys major contributions that enhance notably the response to external interactions and the ability to reproduce observed deformations (See Fig. 1).

We design a deformation model that reproduces with very high accuracy static deformations, matches well dynamic deformations, and exhibits plausible and robust response to unseen interactions. We build the deformation model by extending the SMPL parametric human model [LMR*15] within the volume of the body, and adding a parametric smooth soft-tissue layer in neutral shape and pose. We achieve effective personalization of the soft-tissue layer through shape-and-pose-dependent transformation plus adjustment of its thickness parameters.

The key to match static deformations with very high accuracy is a formulation of mechanics in neutral pose, through inverse skinning of the full motion. With this formulation, soft-tissue mechanics ignore static deformations already considered by the parametric model, and focus instead on accurately capturing highly dynamic deformations. We introduce a sound and simple derivation of this mechanical formulation, and we show how animated characters converge exactly to the output of the parametric model when they reach a static pose.

For plausible and robust interaction, we design a custom nonlinear skin material. Furthermore, by observing human performance data, we have identified strong anisotropy in the motion of skin. Therefore, we augment the custom material with anisotropy, and

define anisotropic error metrics for the estimation of material parameters from observed motions.

Following the definition of the deformation model and the skin material, we personalize them by estimating soft-tissue thickness and heterogeneous material properties from 4D human performance data. We execute the estimation as a numerical optimization, and to optimize efficiently we design acceleration strategies for gradient-based optimization methods. As a result of this optimization, our method accurately reproduces observed deformations. While our model does not intend to match real material stress, it shows realistic behaviors under external interactions.

We structure our paper as follows. After a discussion of related work, we describe the volumetric parametric human model (Section 3), we introduce our deformation model (Section 4), and we discuss our method for model estimation (Section 5).

2. Related Work

The modeling of deformable human characters has been approached with a large variety of methods, ranging from fully physics-based methods to purely data-driven models that ignore physics altogether.

Works on physics-based models aim at modeling an anatomically-inspired representation of the human body that can be deformed in a physics-based simulation framework to reproduce real-world human body behavior. Anatomical methods work inside out to represent the musculoskeletal elements of the body [LST09], and then place a soft layer of skin, which provides the final appearance. Some early works focus on specific body parts, including the modeling of head [KHYS02, SNF05], neck [LT06], hands [AHS03], torso [ZCCD04], or upper-body [TSB*05, LST09].

Other simulation methods approximate the skeletal structure using an articulated body, and model the soft flesh as a continuum, coupled to the skeleton [LYWG13]. Some works have placed emphasis on the properties of the flesh models, including robustness [SGK18] or anisotropy [KDG19]. Other have placed emphasis on the efficiency of the model, as it provides a cost-effective solution for responsive animated characters [KB18, LLK19].

Setting up a physics-based model of a character provides interaction capabilities and a plausible skin behavior. But personalization requires, in addition, the estimation of geometric and material properties from observed data. Some works have addressed the estimation of full-body models [KIL*16], specific models for particularly complex areas such as the face [KK19], or soft-tissue layers from local measurements [PRWH*18], and some have even addressed learning of motor control [NZC*18]. As discussed in the introduction, the work of Kim et al. [KPMP*17] is closest to ours, as they also estimate a soft-tissue model to augment a parametric human model. Throughout the paper we discuss in detail the differences in our model, but there are two that are most significant. Our formulation of the mechanics model is fundamentally novel, and it allows highly accurate matching of static deformations by construction. The model of Kim et al. relies instead on optimization accuracy, which is strongly limited. Our material model and the estimation

metrics support important properties of skin, such as nonlinearity and anisotropy. Kim et al. use a simple linear material that may produce decent matches to deformations produced by skeletal motion, but fails to produce realistic behaviors under external interactions.

At the other end of the spectrum, leaving physics behind, there are the purely data-driven models. These methods aim at finding a model to represent the surface deformations of the human body directly from data. Initial works leverage static 3D scans to build linear models that represent upper torso [ACP02] and full body [ACCP03] static deformations as a function of body shape. The seminal work of SCAPE [ASK*05] went one step further and learned an articulated human body model parameterized by shape and pose. Dynamic deformations caused by soft tissue and muscles have been also attempted from a data-driven perspective. Pioneering works used a marker-based tracking system to capture the trajectory of a few hundred markers to reconstruct [PH06] and model [PH08] soft-tissue deformations.

With 3D scanning technologies becoming more accessible and precise, it is nowadays possible to reconstruct full-body sequences exhibiting highly nonrigid deformations [DTF*15, PMPHB17, RC-DAT17, BRPMB17]. Such detailed and dense reconstructions of human performances have been very recently leveraged to build data-driven models capable of learning soft-tissue dynamics [PM-RMB15, CO18], but with no interaction capabilities.

3. Construction of a Soft-Tissue Avatar

Our method creates soft-tissue deformations as a combination of both data-driven and physics-based components. Our deformable avatars are composed of a soft-tissue layer with varying thickness throughout the body that is simulated using a nonlinear FEM. The data-driven model is used to define the overall shape of the avatar, to kinematically drive the inner vertices of the soft-tissue layer and to formulate deformation mechanics in an unposed reference shape. This provides us with a personalized simulation model that retains as much as possible of the data-driven deformation for static poses and only uses mechanics to precisely model the response to external interactions as well as highly dynamic effects. In this section, we first briefly review SMPL, a statistical model, which is capable of creating static deformations of the body surface for some input shape and pose parameters. Then, we first describe how we create a volumetric mesh of the soft-tissue layer with varying thickness, and adapt SMPL to also account for volumetric deformations. This pipeline results in a parametric soft-tissue avatar ready to be used in simulation.

3.1. Body Surface Model

Our model relies on SMPL [LMR*15], a statistical model of body surface static deformations that modifies a rigged template mesh \mathcal{T} , with $N = 6890$ vertices and $K = 24$ skeletal joints. Vertex positions are adapted depending on two sets of parameters: i) the *pose*, $\theta, |\theta| = 75$, encoding the translation and rotations of the skeletal chain; and ii) the *shape*, $\beta, |\beta| = 10$, representing a series of identity-dependent features that model aspects such as height, slenderness or muscularity. Given these parameters, SMPL produces




Figure 2: We generate a volumetric discretization of the soft-tissue layer with smoothly varying thickness throughout the body. This avoids the negative effect that an irregular discretization of the inner surface might have on the outer surface deformation.

shape and pose dependent surface deformations following the expression:

$$\begin{aligned} S(\beta, \theta) &= W(\bar{S}, J(\beta), \theta, \mathbf{W}), \\ \bar{S}(\beta, \theta) &= \mathcal{T} + S_s(\beta) + S_p(\theta). \end{aligned} \quad (1)$$

Here, $W(\bar{S}, J(\beta), \theta, \mathbf{W})$ is a linear blend skinning function [MTLT88] that takes as input the surface vertices in a reference T-pose $\bar{S} \in \mathbb{R}^{3N}$, the joint locations $J(\beta) \in \mathbb{R}^{3K}$, the pose parameters $\theta \in \mathbb{R}^{75}$ and some blend weights $\mathbf{W} \in \mathbb{R}^{3N \times K}$. To create the reference T-pose \bar{S} , SMPL modifies the template in an additive way through per-vertex 3D offset functions called *blend shapes*. While shape blend shapes, $S_s(\beta)$, model changes due to identity body features, pose blend shapes, $S_p(\theta)$, deform the template to compensate for skinning artifacts and features such as muscle bulging. Both blend shapes and joint location functions are learnt from 3D human captures in different poses. We refer to the original work [LMR*15] for a more detailed explanation.

3.2. Soft-Tissue Mesh

Simulating the soft-tissue layer requires creating a volume mesh with varying thickness throughout the body. One straightforward solution used in [KPMP*17] would be to compute a volumetric discretization of the template mesh, \mathcal{T} , and then select those elements within a given distance of the surface. However, this leads to irregular and jaggy discretizations of the inner surface. This might potentially affect the robustness of the simulation against element inversion and the smoothness of the outer surface deformation due to the highly nonlinear response to external interactions. For this reason, we propose an alternative method that first creates smooth outer and inner boundaries of the flesh layer and then generates the volumetric discretization using a standard meshing software.

The outer surface of the soft-tissue layer corresponds to the template mesh \mathcal{T} . To define the inner surface \mathcal{B} , each vertex of the template mesh $\mathbf{t}_i \in \mathcal{T}, i = 1, \dots, N$ is initially projected to the closest bone of the kinematic chain leading $\hat{\mathbf{t}}_i$. This mapping often results




Figure 3: These are three examples of posed avatars (inset) subject to various soft-tissue deformations expressed in unposed space through a color map. For a static pose (left), the unposed shape is undeformed and identical to the reference shape. For highly dynamic (center) or external interaction (right) scenarios, the unposed shape shows high-frequency local deformations.

on failure cases where neighboring vertices in the surface correspond to points on the skeleton which are far from each other. To correct this issue, we refine this initial solution by iteratively applying a Laplacian smoothing of vertex projections, which slide along the skeleton until convergence. This yields a point distribution where each pair of neighboring surface vertices also have neighboring skeleton projections. In practice, this allows us to effectively shrink the body surface moving each vertex along the projection direction $\mathbf{u}_i = (\hat{\mathbf{t}}_i - \mathbf{t}_i) / |\hat{\mathbf{t}}_i - \mathbf{t}_i|$, while preventing face inversion. Consequently, given some thickness distribution defined at each vertex of the template mesh, $\mathbf{h} = \{h_1, \dots, h_N\}$, the inner surface $\mathcal{B}(\mathbf{h})$ can be easily generated by setting each vertex $\mathbf{b}_i \in \mathcal{B}, i = 1, \dots, N$ as $\mathbf{b}_i = \mathbf{t}_i + \mathbf{u}_i h_i$. Here, thickness values are kept in a feasible range $h_i \in [h_i^{\min}, h_i^{\max}]$. The lower bound h_i^{\min} is equal for all surface vertices and corresponds to a minimum soft-tissue layer thickness. The upper bound h_i^{\max} varies for each of the vertices and can be computed as $h_i^{\max} = |\hat{\mathbf{t}}_i - \mathbf{t}_i| - b$, for some common minimum bone thickness b .

Once both the outer and inner meshes are created, we generate the volumetric mesh using the Tetgen software package [Si15], making sure we retain the original topology for the boundaries. This leads to a parameterized volumetric mesh $\mathcal{V}(\mathbf{h})$ with M vertices and D elements. Notice that none of the soft-tissue deformations affects extremities (i.e., hands, feet and head). In practice, we remove those parts from the volumetric mesh and enforce thickness to slowly decay to zero at the boundaries. As it can be seen in Fig. 2, our way of generating the volumetric mesh produces smooth discretizations that are better suited for simulation.

3.3. Volumetric SMPL

Our deformation method relies on the statistical model to constrain the kinematics of the inner surface of the soft-tissue layer, as well as to formulate deformation mechanics in an unposed reference state. However, the model defined in Equation 1 only handles surfaces and thus we must adapt the SMPL method to account for volumetric deformations. This is done by interpolating blend shape functions and blending weights from the surface points to the volume mesh through Laplacian interpolation. More precisely, as the outer boundary of \mathcal{V} is conformal with the template mesh \mathcal{T} , we simply

constrain the values at the outer vertices and compute a Laplacian interpolation matrix $\mathbf{L} \in \mathbb{R}^{3M \times 3N}$ using the cotangent approximation [Sor05]. Thus, our volumetric SMPL function simply results

$$\begin{aligned} V(\beta, \theta) &= Z(\tilde{V}, \mathcal{J}(\beta), \theta, \mathbf{Z}), \\ \tilde{V}(\beta, \theta) &= \mathcal{V}(\mathbf{h}) + V_s(\beta) + V_p(\theta), \end{aligned} \quad (2)$$

where $V_s = \mathbf{L} \cdot S_s, V_p = \mathbf{L} \cdot S_p, \mathbf{Z} = \mathbf{L} \cdot \mathbf{W}$. We notice that this approximation is not exactly equal to effectively applying Laplacian interpolation to blend shape functions as done in [KPMP*17]. However, we have not detected any significant difference in the performance of the method.

3.4. Parametric Soft-Tissue Avatar

The statistical method described in the previous section provides us with a parametric model of a soft-tissue animated avatar ready to be simulated. The model has a personalized shape β and varying skin thickness \mathbf{h} and can be statically deformed given some time-dependent pose θ . The separation between fixed and time-varying properties defines two clearly different parametric shapes that are used in combination with the FEM simulation:

- The *reference shape* $\bar{\mathbf{x}}(\mathbf{h}, \beta) = \mathcal{V}(\mathbf{h}) + V_s(\beta), \bar{\mathbf{x}} \in \mathbb{R}^{3M}$, which depends only on character-based properties that are fixed through the animation and define a natural rest configuration.
- The *skinned shape* $\mathbf{p}(\bar{\mathbf{x}}, \theta) = Z(\bar{\mathbf{x}} + V_p(\theta), \mathcal{J}(\beta), \theta, \mathbf{Z}), \mathbf{p} \in \mathbb{R}^{3M}$, which conversely depends on the time-changing pose and provides an accurate approximation to soft-tissue static deformation. Note this function can be expressed in a compact way as an affine transformation $\mathbf{P}(\theta)$ of the reference shape $\mathbf{p} = \mathbf{P}(\theta) \cdot \bar{\mathbf{x}}$.

The first expression defines the rest shape of the FEM simulation. In Section 4, we leverage the second expression to drive the kinematics of the inner surface vertices of the soft-tissue layer, as well as formulating mechanics in an unposed deformed shape.

4. Soft-Tissue Avatar Mechanics

The volumetric SMPL method introduced in the previous section provides us with a tool to parametrically deform a volumetric soft-



Figure 4: This picture shows the resulting deformation for two static poses when defining the deformation gradient in world space as in [KPMP*17]. The color map highlights where the result differs from the data-driven static deformation of SMPL.

tissue model given some thickness \mathbf{h} , shape β , and pose θ parameters. We aim to build a hybrid data-driven and physically based model such that it retains as much as possible of the data-driven pose-dependent static deformation, while taking advantage of the capabilities of FEM simulation to capture nonlinear skin dynamics. Overall, this is done by using the volumetric SMPL to impose kinematic boundary conditions on the internal surface of the soft-tissue layer, while letting the rest of the vertices deform freely. However, two design decisions are key for an accurate depiction of the skin behavior. First, we formulate deformation mechanics in an unposed reference space by inverse skinning the SMPL transformation to ignore all pose-dependent deformations already captured by the SMPL model. And second, we design a custom constitutive material that is capable of representing two essential features of the skin: anisotropy and nonlinearity.

In this section, we first formalize the dynamics equation that is solved to simulate the deformation of the soft-tissue layer. Then, we explain how we formulate deformation mechanics and discuss its advantages with respect to other alternatives presented in the related work. Finally, we describe in detail our custom material.

4.1. Dynamics Problem

Simulating the deformation of the FEM model due to skeletal-driven animation and external interactions requires solving the system of nonlinear differential equations defined by discrete Newtonian dynamics. Acceleration and external forces add a deformation offset to the skinned shape, $\mathbf{p}(\bar{\mathbf{x}}, \theta)$, defined in Section 3.4. We denote the resulting configuration of the volumetric mesh $\mathbf{x} \in \mathbb{R}^{3M}$ as *deformed shape*. Then, the dynamics problem is formulated as:

$$\begin{aligned} \mathbf{M}\ddot{\mathbf{x}} - \mathbf{f} - \mathbf{f}_e &= 0 \\ \mathbf{C}_B \cdot \dot{\mathbf{x}} &= \mathbf{C}_B \cdot \mathbf{p}(\bar{\mathbf{x}}, \theta). \end{aligned} \quad (3)$$

Here, \mathbf{C}_B is a matrix selecting only those vertices in the inner boundary of the soft-tissue layer, which are kinematically constrained to the skinned configuration defined by the statistical model $\mathbf{p}(\bar{\mathbf{x}}, \theta) \in \mathbb{R}^{3M}$ for a given reference shape $\bar{\mathbf{x}} \in \mathbb{R}^{3M}$ and time-dependent pose θ . Additionally, \mathbf{M} is the mass matrix, \mathbf{f}_e groups all external forces produced by environmental interactions and \mathbf{f} is the sum of all internal forces generated by skin elastic response. In-

ternal forces can be computed from the derivative of a scalar elastic potential $U(\mathbf{x})$ as $\mathbf{f} = -\nabla_{\mathbf{x}}U$. In the following section, we will focus on how we formulate such elastic potential to ignore pose-dependent static deformations and maximize its response to local dynamic interactions.

4.2. Skin Mechanics Formulation

One key concept for our deformation mechanics formulation is that of *inverse skinning*. The deformed shape \mathbf{x} used in Equation 3 is expressed in world coordinates. However, we can also express deformation in an unposed space by applying the inverse of the skinning affine transformation defined in Section 3.4, $\mathbf{P}^{-1}(\theta)$. This results in a configuration that we denote *unposed deformed shape* $\mathbf{u} = \mathbf{P}^{-1}(\theta) \cdot \mathbf{x}$, which has some desirable properties. In absence of acceleration or external forces, it is identical to the reference shape $\mathbf{u} \equiv \bar{\mathbf{x}}$. In other cases, the resulting shape differs from the reference only in local high-frequency deformations due to external interactions and dynamic effects; big global deformations are undone by the inverse skinning. This opens the possibility to formulate soft-tissue mechanics so that only deformations different from the static pose defined by the skinning transformation produce an actual elastic response. This is similar to the well-known co-rotational formulation [MG04], where the rotational part of the deformation is removed. The co-rotational formulation bears a cost as the estimated rotation depends on the deformed shape, hence affecting derivative computations. This is not the case of our method as the skinning transform depends only on kinematically-defined pose parameters.

Continuum mechanics define the elastic response to deformation in terms of a strain energy density $\Psi(\mathbf{F})$ that depends on the deformation gradient \mathbf{F} , a tensor field defined at each point of the volume as the partial derivative of the deformed position w.r.t. the rest configuration. The total elastic potential is then computed as the volumetric integral $U = \int \Psi(\mathbf{F}) \cdot dV$. FEM provides an interpolation framework to define \mathbf{F} at each element of the volumetric discretization and approximate the total potential integral as a summation [ITF06].

The definition of the deformation gradient \mathbf{F} affects the response to elastic deformation. We propose an alternative definition as the derivative of the unposed deformed shape w.r.t. the rest configuration, $\mathbf{F} = \nabla_{\bar{\mathbf{x}}}(\mathbf{P}^{-1}\mathbf{x})$. This way, our local strain metric ignores large global deformations that are accounted for by the pose-dependent SMPL model and focuses solely on high-frequency local deformations that arise from inertial effects and external interactions. With this formulation, it is not necessary to consider gravity forces because their effect is already included in the statistical model. Inertia effects appear naturally, since the method does not modify the definition of kinetic energy. Finally, note that our deformation gradient \mathbf{F} measures deviation w.r.t. the unposed reference state and this might affect the expected stress response for a given strain. However, we do not aim to model anatomically accurate skin and so strain magnitudes are not exact anyway. Instead, we seek to produce realistic behaviors through a simplified model and, to this end, we optimize material parameters to best fit captured data.

Alternatively, [KPMP*17] use the standard definition of the deformation gradient, i.e., the derivative of deformed shape expressed



Figure 5: This picture shows the deformation of our material (top) in comparison with the nonlinear flesh material defined in [SGK18] (bottom), for two different weights under gravity: 0.5Kg (left) and 2.0Kg (right). It can be clearly seen that our material provides a higher nonlinear response.

in world coordinates $\mathbf{F} = \nabla_{\bar{\mathbf{x}}}\mathbf{x}$. However, this solution does not take advantage of the deformation data provided by the statistical model and does not preserve pose-dependent static deformations in absence of external interactions. As it can be seen in Fig. 4, deformations resulting from the definition of the deformation gradient in world space significantly differ from SMPL parametric shapes even for static poses. There is an additional way of defining a deformation metric such that static deformations are preserved, the so-called *rest-state retargetting* [KDG19]. In that case, the deformation gradient is computed in world coordinates but the reference shape is incrementally updated to match the static deformation produced by the statistical model, i.e., $\mathbf{F} = \nabla_{\mathbf{p}}\mathbf{x}$, with $\mathbf{p} = \mathbf{P}(\theta) \cdot \bar{\mathbf{x}}$. However, this alternative comes with the burden of having to compute all rest-state-dependent magnitudes that remain constant in other cases. Plus, generating the volumetric mesh might produce meshing errors for some extreme poses and affect the solution of the dynamics problem.

Having defined our local measure of strain based on the unposed reference shape $\mathbf{F} = \nabla_{\bar{\mathbf{x}}}(\mathbf{P}^{-1}\mathbf{x})$, the response of the material to deformation is fully governed by the properties of the elastic energy density $\Psi(\mathbf{F})$. In the next section, we will describe in detail how we define such elastic potential to capture deformation properties that are key for an accurate depiction of skin deformation.

4.3. A Nonlinear Skin Material

We aim to design an elastic energy function that accurately captures the complexity of skin dynamics. More precisely, we want to support two deformation properties: *anisotropy* and *nonlinearity*. Skin is a complex structure composed of several layers of het-

erogeneous materials. It has been shown experimentally that this structure results in a highly nonlinear anisotropic response to tangential and normal components of the deformation. After extensive experimentation, we designed a custom elastic potential to support each of these properties as mixture of an orthotropic *StVK* material [LB14], combined with a *Fung* type exponential saturation to increase the nonlinear response [SGK18]. The elastic energy density $\Psi(\mathbf{F})$ results

$$\Psi(\mathbf{F}) = \mu(\exp(\eta \cdot \sigma) - 1)/\eta + \frac{\lambda}{2}(J - 1)^2, \quad (4)$$

$$\sigma = \sum_{i=1}^3 \sum_{j=1}^3 \tau_{ij} \cdot E_{ij}^2, \quad J = \det(\mathbf{F}).$$

Here, E_{ij} are the coefficients of the Green strain tensor $\bar{\mathbf{E}}$ expressed in an orthonormal basis \mathbf{B} that aligns the axes in a meaningful direction for modeling skin anisotropic response, $\bar{\mathbf{E}} = \mathbf{B}^T \cdot \mathbf{E} \cdot \mathbf{B}$, with \mathbf{E} the Green strain tensor evaluated at unposed deformed shape. The first term of Equation 4 models material response to stretch and compression and is controlled through the first Lamè parameter $\mu \in [0, \text{inf})$, together with a saturation parameter $\eta \in (0, \text{inf})$ that regulates the nonlinearity of the elastic response. Anisotropy is achieved by individually weighting the contribution of each quadratic component in σ using weights $\tau_{ij} \in [0, 1]$. The second term depends on the relative volume change $J = \det(\mathbf{F})$ and modulates material response to incompressibility through the second Lamè parameter λ .

Fig. 5 shows a comparison between our material and the nonlinear flesh material defined in [SGK18]. In both cases, material parameters have been estimated from data to best fit target 4D human captures. The saturation component in Equation 4 produces a highly nonlinear response preventing very large deformation from happening even under big external loads. To prove the capabilities of our method to handle external interactions, we have integrated the *ArcSim* cloth simulator [NSO12] in our pipeline as well as incorporated collision handling. Fig. 6 shows realistic soft-tissue deformations as a result of the avatar wearing tight cloths. In addition, handling self-contacts improves the default SMPL solution even in the absence of other external interactions.

4.4. Mechanics Parameterization

Our custom constitutive material requires the characterization of Lamè parameters and several weights, $\{\mu, \lambda, \tau_{ij}\}$. In practice, we estimate the elastic modulus for the normal and tangential directions $Y_N, Y_T \in [0, \text{inf})$, the Poisson's ratio $\nu \in [0, 0.5)$ and the saturation parameter $\eta \in (0, \text{inf})$. We compute Lamè parameters as

$$\mu = \frac{Y_{max}}{2(1 + \nu)}, \quad \lambda = \frac{Y_{max}\nu}{(1 + \nu)(1 - 2\nu)}, \quad (5)$$

where $Y_{max} = \max(Y_N, Y_T)$. In a similar way as done in [LB14], we set the coefficients corresponding to the main stretch directions relative to elastic moduli, $\tau_{00} = Y_N/Y_{max}$ and $\tau_{11} = \tau_{22} = Y_T/Y_{max}$, and automatically compute the value of shear direction coefficients as $\tau_{ij} = \max(\tau_{ii}, \tau_{jj}), i \neq j$ to ensure a reasonable behavior.

Overall, the mechanical behavior of our animated avatar is completely determined by the thickness of the soft-tissue layer at each point of the surface mesh \mathbf{h} ; and the material parameters $\mathbf{m}_i =$

$\{Y_N, Y_T, \nu, \eta\}$ and anisotropy basis \mathbf{B}_i at each element of the volumetric mesh. In practice, we fix Poisson's ratio $\nu = 0.2$ and the saturation parameter $\eta = 4$ to constant values throughout the volume that were experimentally found to produce reasonable volume preservation and nonlinear effects. Anisotropy basis \mathbf{B}_i are defined in the surface mesh to be aligned with the surface normals of the reference shape $\bar{\mathbf{x}}$ defined in Section 3.4, and interpolated through the volume using Laplacian interpolation.

In the next section, we will present how we characterize the thickness, \mathbf{h} , and material, $\{Y_N, Y_T\}_i$, parameters from 4D human captures using numerical optimization. To reduce the computational burden of solving this problem, we consider the material distribution is symmetric along the longitudinal axis, thus we only care about half of the body. Plus, parameters are defined at a reduced number $C = 42$ of control points distributed throughout the surface and interpolated to surface vertices using biharmonic interpolation [JBPS11]. Finally, material parameters are interpolated from surface points to volume vertices using again Laplacian interpolation. To ensure that these properties do not decay with the distance to the surface, values at the vertices of the inner surface are set equal to the value at vertices in the outer surface. This yields a reduced parameter space of $\mathbf{d} \in \mathbb{R}^{3C}$.




Figure 6: This picture shows four instances of soft-tissue deformation due to the effect of tight cloth, contact and friction. We refer to the supplementary video for more detailed animations.

5. Data-Driven Estimation

In the previous sections, we have introduced our hybrid method for the deformation of soft-tissue animated characters. The performance of this method heavily depends on having a good characterization of all the parameters affecting the deformation, which can be an arbitrarily complex task. This includes shape and pose parameters used by the statistical model in Equation 1 but, more importantly, the geometry and material constituting the soft-tissue layer. In this section, we will describe in detail how these parameters are automatically estimated from 4D captures of people using numerical optimization. Our key contribution is the definition of a novel error metric based on the anisotropic motion variance of surface vertices. This is required to accurately estimate the value of the anisotropic material parameters. In addition, we use sparsity acceleration to make gradient-based optimization feasible.

The estimation of SMPL shape β and pose θ parameters that best approximate an input 4D capture can be formulated as a least-squares problem and has been previously described. We refer to e.g. [CO18] for a more detailed explanation. In the following sections, we will focus on the estimation of the parameters affecting the mechanical model. First, we will describe the input data and introduce metrics that can be used to characterize the avatars. And then, we will pose the optimization problem and describe in detail the methods used for solving it.

5.1. Input Data

Our input data consists of several sequences of animated 4D meshes captured from 3 different subjects performing various tasks [PMRMB15]. For a given subject with V sequences, the i -th sequence \mathcal{Z}_i is composed by a series of frames $\mathcal{Z}_i = \{\mathcal{Z}_i^1, \dots, \mathcal{Z}_i^{K_i}\}$ captured at 60 FPS. Each frame is a triangular mesh with $N = 6890$ vertices and exactly the same topology as the template mesh \mathcal{T} described in Section 3.1. Sequences are initially grouped by subject and preprocessed to estimate SMPL parameters as described in [CO18]. This results in a unique shape parameter vector β per subject and as many pose parameter vectors as the total number of frames across all sequences $\{\theta_1^1 \dots \theta_1^{K_1}, \theta_2^1 \dots \theta_2^{K_2}, \dots, \theta_V^1 \dots \theta_V^{K_V}\}$.

5.2. Motion Variance Metrics

We aim to minimize some distance metric between the soft-tissue deformation produced by our hybrid method and the 4D captures under the same input animation. One straightforward choice would be to use the L2 distance between the captured and simulated data, but we are not interested in matching exact trajectories. Instead, we seek to approximate the characteristic movement of the captured data and, for that purpose, L2 distance between trajectories has some undesirable properties. In particular, for static deformations, the error is not zero because SMPL does not exactly match captured data. In situations with low dynamic behavior, the base error caused by the static deformation might hide the contribution of dynamic effects. For these reasons, we look for a simplified descriptive metric of the surface motion in order to approximately match the overall deformation behavior of the skin. We explore the use of the statistical variance of vertex trajectories.

Through early observation of captured data we noticed a skewed




Figure 7: Picture showing the trajectory of a vertex for different configurations of material and error metric (from left to right): captured data, isotropic-isotropic, isotropic-anisotropic, anisotropic-anisotropic. It can be seen that the last configuration clearly outperforms the rest.

distribution of vertex trajectories, where the resulting point cloud takes the approximated shape of an ellipsoid aligned with the surface normal as shown in Fig. 7. This is caused by the highly nonlinear and anisotropic nature of skin dynamics. As a consequence, in our initial estimation experiments, the isotropic motion variance of the deformed mesh approximately matched that of the input captures but still the resulting trajectories differed significantly. This anomaly motivated the creation of a custom anisotropic material as well as exposed the need for a different motion variance metric that takes into account anisotropy.

For notation simplicity, in the following we will drop the subscripts corresponding to subjects and sequences. The trajectory of the i -th vertex of a mesh within a sequence \mathcal{Z} follows a discrete path $\pi_{\mathcal{Z},i} = \{\mathbf{z}_i^1, \dots, \mathbf{z}_i^K\}$, where the subscript and superscript indicate respectively the index of the vertex and the frame. This trajectory includes large displacements due to the large differences in input poses. This might result in a very high dispersion that would not correctly represent the local behavior of the skin. For this reason, we first transform vertex positions to unposed reference space by applying the inverse skinning operation $\mathbf{P}^{-1}(\theta)$ as described in Section 4.2. Therefore, we define the anisotropic motion variance of a vertex as

$$\begin{aligned} M_N(\pi_{\mathcal{Z},i}) &= \frac{1}{K} \sum_{j=1}^K |\mathbf{N}_i(\mathbf{P}_i^{-1}(\theta^j)) \cdot \mathbf{z}_i^j - \bar{\mathbf{z}}_i|^2, \\ M_T(\pi_{\mathcal{Z},i}) &= \frac{1}{K} \sum_{j=1}^K |(\mathbf{I} - \mathbf{N}_i)(\mathbf{P}_i^{-1}(\theta^j)) \cdot \mathbf{z}_i^j - \bar{\mathbf{z}}_i|^2, \\ \bar{\mathbf{z}}_i &= \frac{1}{K} \sum_{j=1}^K \mathbf{P}_i^{-1}(\theta^j) \cdot \mathbf{z}_i^j, \end{aligned} \quad (6)$$

where \mathbf{N}_i is the projection matrix to the subspace defined by the normal of the surface at the i -th vertex, $\mathbf{P}_i^{-1}(\theta^j)$ is the inverse skinning transformation corresponding to the i -th vertex at the j -th frame, and $\bar{\mathbf{z}}_i$ is the average vertex position across all frames. This allows us to define an anisotropic variance residual between the simulated and captured sequences:

$$D(\mathcal{X}, \mathcal{Z}) = M(\mathcal{X}) - M(\mathcal{Z}), \quad (7)$$

where $M(\mathcal{X}) = \{M_N(\pi_{\mathcal{X},i}), M_T(\pi_{\mathcal{X},i})\}, i = 1, \dots, N$, is a vector containing the normal and tangential motion variances corresponding to all the vertex trajectories of a sequence. Note that here we only consider those vertices in the outer surface of the soft-tissue layer. In the following sections, we will pose an optimization prob-

lem to minimize an error based on this residual and show that our anisotropic distance metric leads to a more precise description of skin behavior and ultimately better results.

5.3. Optimization Procedure

For a given subject with V sequences, we pose the optimization a nonlinear least squares problem based on the residual defined in Equation 7

$$\min_{\mathbf{d}} \sum_{i=1}^V D(\mathcal{X}_i(\mathbf{d}), \mathcal{Z}_i)^T D(\mathcal{X}_i(\mathbf{d}), \mathcal{Z}_i), \quad (8)$$

where $\mathbf{d} \in \mathbb{R}^{3C}$ is the vector that concatenates all free parameters described in Section 4.4. For notation simplicity and without loss of generality we will assume we are optimizing just one sequence. We aim to solve this problem using a gradient-based method to take advantage of all the available machinery in standard optimization packages. This requires the computation of the gradient as $\nabla_{\mathbf{d}} D \cdot D$ which is highly nontrivial. Each evaluation of the residual D requires the simulation of the complete animated sequence, which means solving a large number of nonlinear systems of equations defined in Section 4.1. Moreover, the solution of each dynamic problem depends on all the previous frames of the sequence which complicates the computation of the analytic Jacobian $\nabla_{\mathbf{d}} D$. Consequently, we opt to estimate the gradient using finite differences, which requires one full evaluation of the simulation per optimized parameter.

To alleviate this large computational cost, we take advantage of the fact that the Jacobian $\nabla_{\mathbf{d}} D$ is sparse in practice. The trajectory of each of the deformed vertices depends only on the thickness and material distribution in a close neighborhood. For instance, it is very unlikely that the material of the leg and the mechanical behavior of the chest are related in any way. We solve the optimization problem using the off-the-shelf trust-region nonlinear least squares solution offered in *SciPy* optimization package that takes advantage of a sparsity pattern to estimate the Jacobian using finite-differences. The sparsity is defined through a matrix $\mathbf{S} \in \mathbb{R}^{N \times C}$ such that $S_{ij} \neq 0$ if and only if the mechanical behavior of the i -th vertex is affected by the parameters of the j -th control point. We build such matrix heuristically:

1. We first compute the biharmonic interpolation matrix from control points to surface vertices $\mathbf{H} \in \mathbb{R}^{N \times C}$. Then, we build the p -Influence matrix \mathbf{H}_p by taking the p higher coefficients of

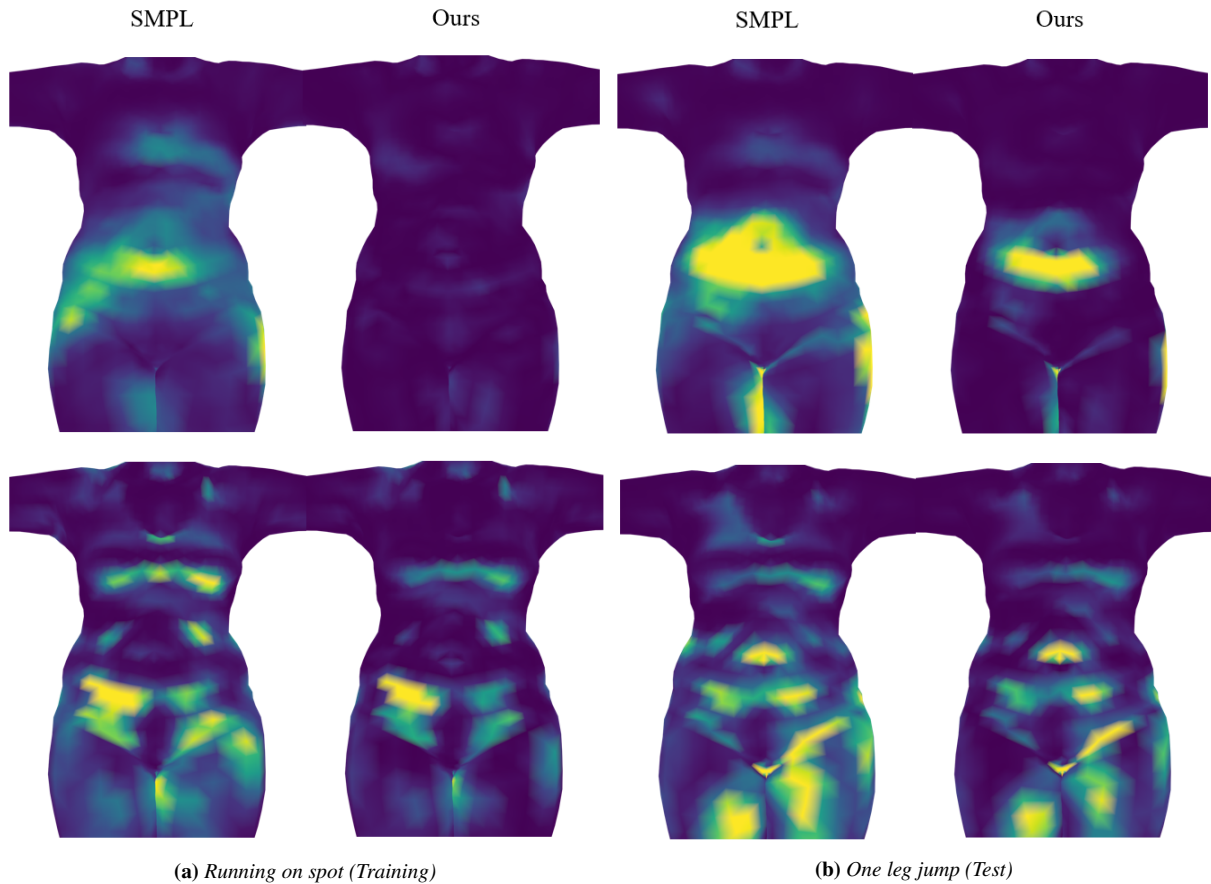



Figure 8: This figure shows a comparison of the tangential (top) and normal (bottom) variance errors for subject S_B , between SMPL and our soft-tissue model. It can be seen that our method approximates the overall dynamic behavior better for both training and test sequences.

each row and rounding them up to 1, leaving the rest of them to 0. This relates each surface vertex to its p most influential control points.

2. The adjacency matrix $\mathbf{A}_p = \mathbf{H}_p^T \cdot \mathbf{H}_p$ defines neighborhoods of control points sharing surface vertices for which such control points are among the p most influential. This adjacency criterion can be finally used to build k -ring expansions of the influence matrix $\mathbf{S}(p, k) = \mathbf{H}_p \cdot \mathbf{A}_p^k$.

In our experiments, the values $p = 2$ and $k = 2$ result in a sufficiently sparse pattern and a good trade-off between the quality of the estimation and computation time. The inset figure shows the dependency area corresponding to one of the points in the chest and its corresponding symmetric point. To further alleviate the cost of the optimization and avoid redundant effects of material stiffness and layer thickness we separate both sets of parameters and follow and alternating optimization scheme. First, we optimize material parameters to convergence using a simple distribution of $C = 12$



points. Then we alternate optimization steps of the thickness and material parameters until no further improvement is possible. Finally, we take this initial result and refine it using the total number of control points $C = 42$, directly on an alternating scheme. In the following section, we will show the results of fitting material and thickness parameters from data for 3 different subjects with various input sequences and discuss the performance of the optimization procedure.




Figure 9: The three test subjects: S_A , S_B and S_C .

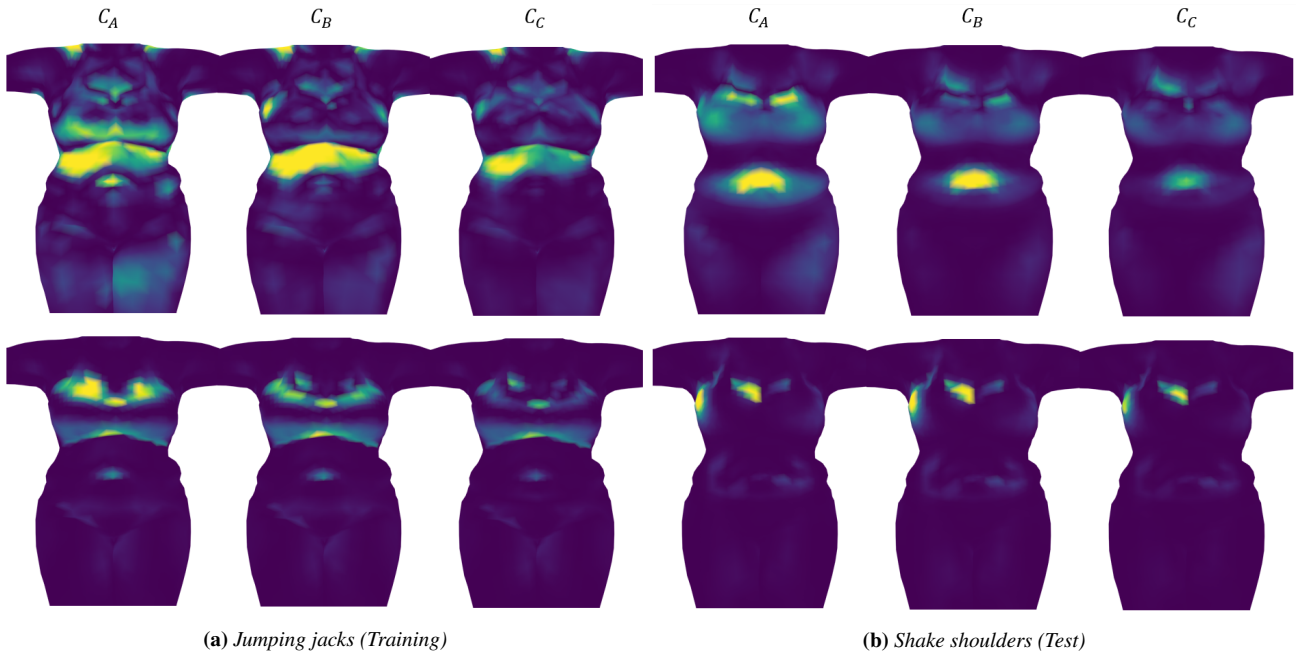


Figure 10: This figure shows a comparison of the tangential (top) and normal (bottom) variance errors for subject S_A , using the three different configurations of the optimization. It can be seen that for C_C the error w.r.t. the captured data is lower for both training and test sequences.

6. Experiments and Results

The optimization procedure described in the previous section is capable of characterizing the parameters of our mechanical model. As it can be seen in Fig. 1 and the supplementary video, this produces soft-tissue characters that realistically deform under highly dynamic animations and external interactions. In this section, we




Figure 11: This graph shows a comparison of the anisotropic variance error for each of the five sequences, averaged across the three different subjects, between SMPL (pink) and our hybrid data-driven and physics based method (purple). It can be seen that our soft-tissue avatar clearly outperforms the SMPL solution.

describe in more detail our parameter estimations experiments and discuss our optimization scheme.

We estimated the parameters of our soft-tissue avatar model for three different subjects that we denote S_A , S_B and S_C (Fig. 9). For each subject, we use two animation sequences of 1.5s as the training set and three animation sequences of [3,4]s as the test set. We focus our main analysis in this section on subjects S_A and S_B , as those show stronger dynamic behavior. Fig. 8 shows the distribution of the normal and tangential variance error defined in Equation 8, for subject S_B in two different animation sequences. It can be seen that our method significantly reduces the error of the SMPL model and approximates the dynamic behavior of the captures from just two training sequences. Thanks to the formulation of the error metric in terms of the anisotropic motion variance, we notice that most of it is concentrated in the normal part of the residual. We believe that this is caused by an existing trade-off in our material definition between compliance in the normal direction and volume preservation (which is a desired property when modeling skin mechanics). This is expected as the skin has inherently limited stretch in the tangential direction, thus constraining the deformation in the normal direction if volume must be preserved. We found that using a Poisson's Ratio of $\mu = 0.2$ was a reasonable compromise solution. For future work, we would like to explore more complex soft-tissue models (e.g., multi-layer) that might alleviate this limitation. Overall, our method reduces the error of the statistical model by 56.8% in the case of training sequences and 31.4% in the case of test sequences (Fig. 11). For some particular tests that notably differ from the training set, the optimization is not fully capable of generalizing the behavior of the soft-tissue. In such cases, we have observed




Figure 12: This graph shows the convergence of the optimization considering the anisotropic material and error metric for each of the three subjects tested. The dots indicate the iteration separating the three stages of the optimization: i) 12 points, material parameters; ii) 12 points, alternating material and thickness; and iii) 42 points, alternating material and thickness.

that the solution provided by SMPL produces smaller variance errors. We found that this happens because the anisotropic metric is particularly sensitive to the misalignment between the captured and simulated vertex trajectories. However, as it can be seen in the supplementary video, perceptually, our method provides qualitatively better results.

We repeated the optimization of subject S_A several times to compare the performance of our custom constitutive material when considering isotropic and anisotropic elastic moduli. Additionally, we explore the effect that our anisotropic error metric has on the results of the optimization. Overall, this result in three configurations:

- C_A : isotropic material with isotropic error metric.
- C_B : isotropic material with anisotropic error metric.
- C_C : anisotropic material with anisotropic error metric.

Fig. 7 shows a comparison between the trajectories obtained for one of the training sequences, for each of the mentioned configurations. It can be seen that the C_C produce vertex trajectories that fit much better with the captured one than C_A . While the anisotropic error metric used in C_B helps to improve the overall quality of the fit, the isotropic constitutive material cannot capture skin behavior accurately. Fig. 10 shows in detail the spatial distribution of the anisotropic variance for each of the considered configurations. It can be seen that the configuration C_C clearly outperforms the other two in approximating the captured data.

Overall, the complexity of the problem complicates convergence and suffers risk of getting stuck in local minima. Our optimization scheme notably alleviates the burden and makes solving the problem feasible. As shown in Fig. 12, it is effective to provide an initialization for the full optimization using the solution obtained with the reduced parameter space of 12 control points. However, the big computational cost remains the main limitation and prevents us from further exploring other options that might lead to

better results. All our simulations and optimizations were run in a desktop machine Intel Core i7-8700 @ 3.2GHz, with 32GiB RAM. We use a time step of 1/120s. Each time step converges in 1 or 2 Newton iterations, and the cost per iteration is around 0.7s. We estimate the parameters of each subject using 2 sequences of 1.5s each. Altogether, simulating the two sequences takes about 6min. In each optimization step, we estimate the Jacobian of error terms using finite differences. With 42 control points, this implies running 43 simulations per thickness iteration and 85 simulations per material iteration. Our sparsity acceleration approximately divides the number of simulations by 2. Overall, this leads to a total optimization time of around 30-40 hours, depending on the subject and the chosen training set. In future work, we would like to focus on improving the efficiency of the optimization.

7. Limitations and Future Work

We have presented a hybrid model of soft-tissue animated avatar that computes deformation as a combination of a data-driven statistical model and a FEM simulation. In addition, we have presented an optimization procedure to estimate the parameters of the model from 4D human captures. Unlike previous work, we ensure that two important properties are achieved. We reformulate strain computations such that, in absence of external interactions and dynamics effects, the deformation generated by our method matches exactly that of the data-driven statistical model. Plus, we model soft-tissue mechanics using a custom constitutive material that captures the nonlinearity and anisotropy of the skin, which are essential for representing complex interactions. To best explore the parameter space of our material model, we have formulated the estimation problem in terms of a novel anisotropic motion variance metric.

We have shown that our method is capable of realistically modeling soft-tissue avatars under highly dynamic effects and extreme external interactions. However, our work is not free from limitations. During our experiments we detected that there is a trade-off between the volume preservation and the anisotropic capabilities of the material, which limits the quality of the fitting to captured data. This might be solved using a more complex model of the soft-tissue that would consider several layers with different mechanical properties. In addition, characterizing material parameters requires solving a complex optimization problem with large computational cost. Even though our sparsity accelerated gradient alleviates part of the burden, using finite-differences for the computation of the Jacobian quickly becomes hardly feasible for a large number of control points. It would be interesting to find an appropriate reformulation of the dynamics problem such that the gradient can be analytically approximated. Solving both these problems would constitute interesting lines of future work.

Acknowledgements

The authors wish to thank Igor Santesteban for his support with SMPL and data acquisition, as well as Rosa Sanchez and Hector Barreiro for their help with results preparation. We also thank other members at URJC for their support. The work was funded in part by the European Research Council (ERC Consolidator Grant no. 772738 TouchDesign) and Spanish Ministry of Science (RTI2018-098694-B-I00 VizLearning).

References

- [ACCP03] ALLEN B., CURLESS B., CURLESS B., POPOVIĆ Z.: The space of human body shapes: Reconstruction and parameterization from range scans. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 22, 3 (July 2003), 587–594. 3
- [ACP02] ALLEN B., CURLESS B., POPOVIĆ Z.: Articulated Body Deformation from Range Scan Data. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 21, 3 (2002), 612–619. 3
- [AHS03] ALBRECHT I., HABER J., SEIDEL H.-P.: Construction and animation of anatomically based human hand models. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation* (2003), Eurographics Association, pp. 98–109. 2
- [ASK*05] ANGUELOV D., SRINIVASAN P., KOLLER D., THRUN S., RODGERS J., DAVIS J.: SCAPE: Shape Completion and Animation of People. *ACM Trans. Graph.* 24, 3 (July 2005), 408–416. 3
- [BRPMB17] BOGO F., ROMERO J., PONS-MOLL G., BLACK M. J.: Dynamic FAUST: Registering human bodies in motion. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017* (Piscataway, NJ, USA, July 2017), IEEE. 3
- [CO18] CASAS D., OTADUY M. A.: Learning nonlinear soft-tissue dynamics for interactive avatars. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 1, 1 (2018), 10. 2, 3, 7
- [DTF*15] DOU M., TAYLOR J., FUCHS H., FITZGIBBON A., IZADI S.: 3D scanning deformable objects with a single RGBD sensor. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2015), pp. 493–501. 3
- [ITF06] IRVING G., TERAN J., FEDKIW R.: Tetrahedral and hexahedral invertible finite elements. *Graphical Models* 68, 2 (2006), 66 – 89. Special Issue on SCA 2004. 5
- [JBPS11] JACOBSON A., BARAN I., POPOVIĆ J., SORKINE O.: Bounded biharmonic weights for real-time deformation. *ACM Trans. Graph.* 30, 4 (July 2011), 78:1–78:8. 7
- [KB18] KOMARITZAN M., BOTSCH M.: Projective skinning. *Proc. ACM Comput. Graph. Interact. Tech.* 1, 1 (2018), 12:1–12:19. 2
- [KDG19] KIM T., DE GOES F., IBEN H.: Anisotropic elasticity for inversion-safety and element rehabilitation. *ACM Trans. Graph.* 38, 4 (July 2019), 69:1–69:15. 2, 6
- [KHYS02] KÄHLER K., HABER J., YAMAUCHI H., SEIDEL H.-P.: Head shop: Generating animated head models with anatomical structure. In *Proceedings of the 2002 ACM SIGGRAPH/Eurographics symposium on Computer animation* (2002), ACM, pp. 55–63. 2
- [KIL*16] KADLEČEK P., ICHIM A.-E., LIU T., KRIVÁNEK J., KAVAN L.: Reconstructing personalized anatomical models for physics-based body animation. *ACM Trans. Graph.* 35, 6 (2016), 213:1–213:13. 2
- [KK19] KADLEČEK P., KAVAN L.: Building accurate physics-based face models from data. *Proc. ACM Comput. Graph. Interact. Tech.* 2, 2 (2019), 15:1–15:16. 2
- [KPMP*17] KIM M., PONS-MOLL G., PUJADES S., BANG S., KIM J., BLACK M. J., LEE S.-H.: Data-driven physics for human soft tissue animation. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 54. 2, 3, 4, 5
- [LB14] LI Y., BARBIĆ J.: Stable orthotropic materials. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (Aire-la-Ville, Switzerland, Switzerland, 2014), SCA '14, Eurographics Association, pp. 41–46. 6
- [LLK19] LI J., LIU T., KAVAN L.: Fast simulation of deformable characters with articulated skeletons in projective dynamics. In *Proceedings of the 18th Annual ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2019), SCA '19, pp. 1:1–1:10. 2
- [LMR*15] LOPER M., MAHMOOD N., ROMERO J., PONS-MOLL G., BLACK M. J.: Smpl: A skinned multi-person linear model. *ACM Trans. Graph.* 34, 6 (Oct. 2015), 248:1–248:16. 2, 3
- [LST09] LEE S.-H., SIFAKIS E., TERZOPOULOS D.: Comprehensive biomechanical modeling and simulation of the upper body. *ACM Trans. Graph.* 28, 4 (2009), 99:1–99:17. 2
- [LT06] LEE S.-H., TERZOPOULOS D.: Heads up!: Biomechanical modeling and neuromuscular control of the neck. *ACM Trans. Graph.* 25, 3 (July 2006), 1188–1198. 2
- [LYWG13] LIU L., YIN K., WANG B., GUO B.: Simulation and control of skeleton-driven soft body characters. *ACM Trans. Graph.* 32, 6 (Nov. 2013), 215:1–215:8. 2
- [MG04] MÜLLER M., GROSS M.: Interactive virtual materials. In *Proceedings of Graphics Interface 2004* (Waterloo, CAN, 2004), GI '04, Canadian Human-Computer Communications Society, p. 239–246. 5
- [MTLT88] MAGNENAT-THALMANN N., LAPERRIÈRE R., THALMANN D.: Joint-dependent local deformations for hand animation and object grasping. In *Proceedings on Graphics Interface '88* (Toronto, Ont., Canada, Canada, 1988), Canadian Information Processing Society, pp. 26–33. 3
- [NSO12] NARAIN R., SAMII A., O'BRIEN J. F.: Adaptive anisotropic remeshing for cloth simulation. *ACM Trans. Graph.* 31, 6 (Nov. 2012). 6
- [NZC*18] NAKADA M., ZHOU T., CHEN H., WEISS T., TERZOPOULOS D.: Deep learning of biomimetic sensorimotor control for biomechanical human animation. *ACM Trans. Graph.* 37, 4 (2018), 56:1–56:15. 2
- [PH06] PARK S. I., HODGINS J. K.: Capturing and animating skin deformation in human motion. *ACM Transactions on Graphics (TOG)* 25, 3 (2006), 881–889. 3
- [PH08] PARK S. I., HODGINS J. K.: Data-driven modeling of skin and muscle deformation. *ACM Trans. Graph.* 27, 3 (2008), 96:1–96:6. 3
- [PMPHB17] PONS-MOLL G., PUJADES S., HU S., BLACK M. J.: Clothcap: Seamless 4d clothing capture and retargeting. *ACM Transactions on Graphics (Proc. of SIGGRAPH)* 36, 4 (2017), 73:1–73:15. 3
- [PMRMB15] PONS-MOLL G., ROMERO J., MAHMOOD N., BLACK M. J.: Dyna: A model of dynamic human shape in motion. *ACM Trans. Graph.* 34, 4 (July 2015), 120:1–120:14. 2, 3, 7
- [PRWH*18] PAI D. K., ROTHWELL A., WYDER-HODGE P., WICK A., FAN Y., LARIONOV E., HARRISON D., NEOG D. R., SHING C.: The human touch: Measuring contact with real human soft tissues. *ACM Trans. Graph.* 37, 4 (2018), 58:1–58:12. 2
- [RCDAT17] ROBERTINI N., CASAS D., DE AGUIAR E., THEOBALT C.: Multi-view performance capture of surface details. *International Journal of Computer Vision (IJCV)* 124, 1 (Aug 2017), 96–113. 3
- [SGK18] SMITH B., GOES F. D., KIM T.: Stable neo-hookean flesh simulation. *ACM Trans. Graph.* 37, 2 (Mar. 2018), 12:1–12:15. 2, 6
- [Si15] SI H.: Tetgen, a delaunay-based quality tetrahedral mesh generator. *ACM Trans. Math. Softw.* 41, 2 (Feb. 2015), 11:1–11:36. 4
- [SNF05] SIFAKIS E., NEVEROV I., FEDKIW R.: Automatic determination of facial muscle activations from sparse motion capture marker data. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 24, 3 (2005), 417–425. 2
- [Sor05] SORKINE O.: Laplacian Mesh Processing. In *Eurographics 2005 - State of the Art Reports* (2005), Chrysanthou Y., Magnor M., (Eds.), The Eurographics Association. 4
- [TSB*05] TERAN J., SIFAKIS E., BLEMKER S. S., NG-THOW-HING V., LAU C., FEDKIW R.: Creating and simulating skeletal muscle from the visible human data set. *IEEE Transactions on Visualization and Computer Graphics* 11, 3 (2005), 317–328. 2
- [ZCCD04] ZORDAN V. B., CELLY B., CHIU B., DILORENZO P. C.: Breathe easy: model and control of simulated respiration for animation. In *Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on Computer animation* (2004), Eurographics Association, pp. 29–37. 2