Breathing Life into Shape:
Capturing, Modeling and Animating 3D Human Breathing

Aggeliki Tsoli* Naureen Mahmood† Michael J. Black‡
Max Planck Institute for Intelligent Systems, Tübingen, Germany

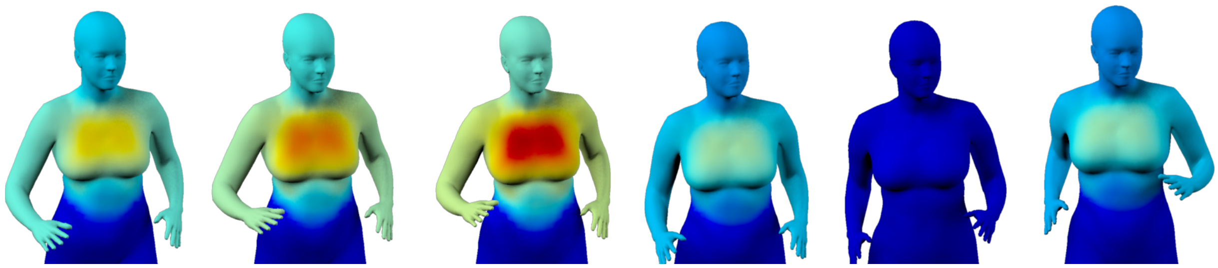

Figure 1: Animating breathing. *Here we add realistic breathing shape deformations to a 3D model of a person running. The Euclidean distances between vertices of the breathing and non-breathing model are color coded (red is large, blue is small). Here the runner is breathing mostly with the chest and the temporal pattern of breathing was animated by a "breath actor."*

Abstract

Modeling how the human body deforms during breathing is important for the realistic animation of lifelike 3D avatars. We learn a model of body shape deformations due to breathing for different breathing types and provide simple animation controls to render lifelike breathing regardless of body shape. We capture and align high-resolution 3D scans of 58 human subjects. We compute deviations from each subject's mean shape during breathing, and study the statistics of such shape changes for different genders, body shapes, and breathing types. We use the volume of the registered scans as a proxy for lung volume and learn a novel non-linear model relating volume and breathing type to 3D shape deformations and pose changes. We then augment a SCAPE body model so that body shape is determined by identity, pose, and the parameters of the breathing model. These parameters provide an intuitive interface with which animators can synthesize 3D human avatars with realistic breathing motions. We also develop a novel interface for animating breathing using a *spirometer*, which measures the changes in breathing volume of a "breath actor."

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation I.6.8 [Simulation and Modeling]: Types of Simulation—Animation

Keywords: Breathing animation, human body modeling, statistical model, respiration, 3D shape, learning

Links: DL PDF

*e-mail:aggeliki.tsoli@tue.mpg.de
†e-mail:nmahmood@tue.mpg.de
‡e-mail:black@tue.mpg.de

1 Introduction

We describe a method to animate realistic breathing in virtual humans with a simple intuitive interface. Realistic human models and avatars are common in movies and video games. While 3D body scanning technology produces realistic looking 3D body meshes, making them look "alive" requires that they breathe. Moreover, breathing is part of body language and is important for conveying emotions. Apart from visually pleasing animations in the film or video game industry, realistic animation of breathing is also essential in the medical domain (e.g. for planning radiation therapy). Given the importance of breathing, there are surprisingly few techniques that produce realistic breathing motions, across a range of body shapes, without extensive animation by hand.

Modeling breathing in a realistic, lifelike way is a hard problem. First, it entails modeling subtle yet complex deformations of the human body that vary across time and context. Second, breathing has a time-varying global effect on the human body; it induces shape change mainly in the torso, but also posture changes over the whole body. Previous work on animating breathing 3D avatars has been either limited in realism or does not generalize easily to new shapes and breathing types [Park and Hodgins 2006; Promayon et al. 1997; Sanders et al. 2009; Zordan et al. 2004; Veltkamp and Piest 2009].

We propose a new approach for modeling body deformations due to breathing using high-resolution 3D human scans, a statistical model of the human body, and physiological parameters related to respiration. An example of animating the breathing of a running character is shown in Fig. 1. To capture the subtle and complex deformations of the human body shape due to breathing, we scanned 58 human subjects at multiple time instants during natural breathing. High resolution triangulated meshes were captured using 22 pairs of stereo cameras and a random projected texture pattern together with 22 color cameras and a white-light flash system; this system gives high quality 3D meshes with registered texture. To elicit a range of deformations, subjects were asked to breathe normally, with the chest, and with the stomach. To separate breathing-induced shape change in our data from pose-induced shape change, we register all scans to a statistical model of body shape and pose variation. We compute a mean shape for each subject and the deviations from the mean due primarily to breathing.

We perform principal component analysis (PCA) on the estimated breathing deformations to produce a low-dimensional model of

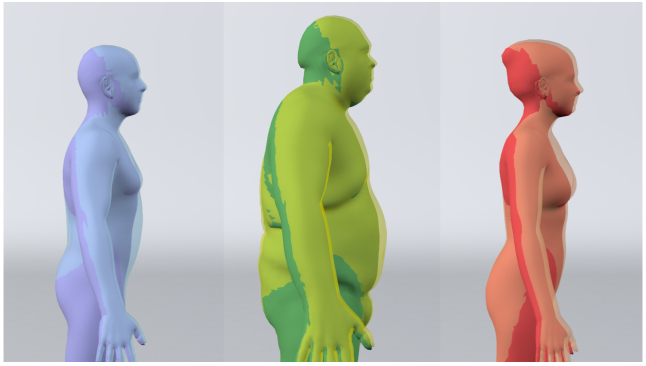

Figure 2: Animating breathing types. *Respiration induces changes in torso shape and posture. We learn a model of how 3D breathing deformations relate to lung volume and breathing type and use it to animate bodies of varying shape and pose. Here we show the maximal inhale and exhale shapes overlaid for three different bodies breathing mainly with the stomach (left), mainly with the chest (right), or using a combination of chest and stomach (middle).*

breathing variation. The PCA space has distinct components for "chest breathing" and "stomach breathing". For instance, the shape change during chest breathing is as much up and down as in and out; this is quite different from the shape changes used in simple animations. Maybe surprisingly, we found little significant difference in breathing shape change between men and women. We did find postural changes that were significantly correlated with breathing and that differed between men and women. We also found that the dominant breathing shape deformations were only weakly correlated with body shape but that body shape is correlated with fine-scale differences of shape change during breathing.

To animate breathing we need natural controls that are related to the statistics of pose and shape deformations. To that end, we compute the difference in volume between each 3D body and the mean shape of the subject. We take the change in volume as a proxy for change in lung volume. This allows us to model breathing deformations as a function of volume and to use volume as a simple, and physiologically relevant, control for animation. We also define different *types* of breathing as illustrated in Fig. 2. For a given breathing type, we find that body shape varies linearly with volume. This linear relationship, however varies non-linearly with breathing type. We learn a novel mathematical model of body shape deformation and pose change as a function of volume and type. We also extend the SCAPE body model [Anguelov et al. 2005] to include body shape deformations predicted by our breathing model. These deformations are combined with identity and pose deformations to produce realistic breathing for bodies of any shape and any pose (Figs. 1 and 2).

We describe an intuitive interface for creating temporal breathing patterns and for changing breathing types. To more easily capture realistic lung volume sequences for animation we use a device called a spirometer. This device makes it possible for a human "breath actor" to "act out" a particular breathing sequence to correspond with the desired action or emotional content. The recorded changes in lung volume drive the animated character using the learned shape deformation model, providing an easy and novel way to achieve realistic breathing animation.

While our shape model is built from subjects in a standing pose, we show that the learned model applies to other poses. We animate a 3D body model and use breath acting to recover the corresponding breathing sequence. The animated sequences with breathing look more natural than sequences without breathing.

2 Related Work

Animation by hand. Breathing is a strong indication of life and realistic characters in feature films often have many parameters for hand animation of breathing; the animation is labor intensive. For simpler characters (e.g. in video games) fairly primitive models may be used that capture the gist of breathing through changes in posture (rocking back and forth) or simple cyclic expansion of the chest. Basic breathing controls like these are sometimes used for idle motion generation [Egges et al. 2004; Egges et al. 2006]. In this case breathing is seen as a cause of idling motion, rather than something to model on its own. What is missing is a realistic model of breathing, with simple animation controls, that can be applied to many body shapes in motion.

Anatomy- and physics-based modeling. There is extensive work on anatomy- and physics-based modeling of the human body; see [Lee et al. 2012; Magnenat-Thalmann et al. 2009] for reviews. For breathing, prior work focuses on modeling the torso [Promayon et al. 1997; Zordan et al. 2004; Veltkamp and Piest 2009]. Zordan et al. [2004] propose an anatomically motivated model of the human torso that consists of rigid parts (bones) and deformable parts (muscles). Animation requires physical simulation. Veltkamp et al. [2009] introduce a similar model that combines better control over abdominal and chest breathing using two independent breathing systems. Lee et al. [2009] present a comprehensive biomechanical model of the upper human body with a proof-of-concept demonstration of synthesized breathing motions. In the medical domain, breathing models focus primarily on representing lung shape [Moreno et al. 2007; Cavagnaro et al. 2013; Santhanam et al. 2003].

Although anatomy- and physics-based body models offer the potential for high detail, they do not generalize easily to new subjects. Synthesizing new human bodies as well as tuning the parameters to generate specific types of breathing is not straightforward. Since these breathing models focus on the torso, they do not model whole-body posture variation during breathing. These issues, combined with the computational expense of physics simulation, mean that such methods are difficult to use in practice. In contrast, our model is learned from data, generalizes easily to new subjects, models whole-body posture variation during breathing, and provides intuitive controls for synthesizing breathing animations. In previous anatomy-based models, volume change over time is the observed outcome that is used to evaluate whether an animation is realistic. In our case, breathing volume is the input that drives the animation and we can animate arbitrary breathing sequences using volume.

Statistical human body models. Previous work on synthesizing breathing in a data-driven way is limited to replaying recorded breathing motions for 3D shapes similar to the shape of the recorded subjects [Park and Hodgins 2006; Sanders et al. 2009]. However, statistical body models have been used successfully in the past to model the observed body shape across the human population [Anguelov et al. 2005; Allen et al. 2006; Hasler et al. 2009; Chen et al. 2013]. In these models the observed shape is conceptually decomposed to the intrinsic shape of the subject and deformations that change based on the pose of the subject, such as muscle bulging etc. They do not model breathing deformations or, in fact, other deformations not due to identity or pose. In this work, we extend the SCAPE model [Anguelov et al. 2005] to include breathing deformations and define controls to easily animate these deformations.

Dynamic shape capture and modeling. The modeling of breathing shape deformations has been limited by a lack of data. High-resolution 3D body scanners typically require several seconds be-

Figure 3: Example scans. *58 subjects were scanned in an "A" pose while breathing. Subjects wore tight fitting clothing so that shape changes during breathing were evident. They were asked to perform different "types" of breathing: normal, breathing with the chest, and breathing with the stomach. The full dataset consists of 2807 3D meshes with associated texture.*

tween consecutive scans meaning that fine temporal resolution is lost. Depth maps with high temporal resolution are available from range sensors [Penne et al. 2008] but these are noisy and have low spatial resolution. Despite progress on tracking complex surfaces such as human clothing in video sequences [Stoll et al. 2010], capturing accurate subtle deformations of the human body remains a challenge.

Low spatial resolution and high temporal resolution is available from tracked markers but, with standard marker sets, breathing is not readily visible. Larger marker sets can capture breathing motions of individuals [Park and Hodgins 2006]. Park and Hodgins capture shape changes of an actor breathing using approximately 350 markers and replay this motion on a similar body model by deforming the body mesh. The markers provide a high-dimensional control signal, with high temporal resolution, that can be used to reproduce highly realistic and nuanced animations. In contrast to our approach, however, they do not provide a model of breathing that can be easily controlled and animated to produce different effects. We learn a low-dimensional model of breathing deformations from examples that can be applied to different body shapes, poses, and motion sequences. Our model is parameterized by lung volume and breathing type using concepts from the physiology of respiration [West 2012; Mines 1981].

Learning a breathing model requires capturing the breathing shape of different people and types of breathing. To analyze breathing across a population, we need breathing deformations that are in correspondence across people and over time. Marker data could be hard to align because, on areas like the torso, it is difficult to place markers in the same location on different people. Instead Park and Hodgins [2008] use large marker sets to learn a deformation model for static and dynamic motions. They then control the deformations using smaller marker sets (40-50 makers); they do not address breathing. Small marker sets, with appropriately placed markers, could be used to control breathing animations. In contrast, here we develop a model of breathing that can be animated with simpler controls.

Computer vision methods with texture painted on the skin provide a possible alternative to marker-based systems [Neumann et al. 2013] but have not been used to model breathing. In contrast to previous approaches, we acquire a dense reconstruction of the human subject's shape using a high-resolution 3D scanner. We acquire multiple scans of each subject at unknown time instants in the breathing cycle and register them by taking into account both the geometry and appearance of the 3D scans. We find that high-resolution meshes facilitate the computation of correspondence across subjects and across time.

Controls for animation. The motion of markers on the chest has been used to drive an anatomy-inspired model [Sanders et al. 2009]. Other controls for breathing animation include audio [DiLorenzo et al. 2008; Cosker and Edge 2009] and parameters related to human physiology. Animation from physiology-related input (including a stretch sensor on the chest, EKG, pulse, skin temperature) has been limited to the anatomically-based models described above and lacks visual realism. Our approach is more similar in concept to [Kider et al. 2011], where human body surface deformations are correlated with recorded physiology data related to the level of fatigue. In our case, we link the observed surface deformations with the lung volume during breathing. In addition we animate 3D human characters using spirometer data (lung volume measurements) recorded by "breath actors."

3 Breath Taking (Data Capture)

To model deformations of the human body due to breathing as realistically as possible we capture high-resolution 3D full-body scans of 58 subjects (28 men and 30 women) of various shapes; Fig. 3 shows a few representative scans. Subjects were a mix of professional models (with a modeling contract), paid participants (8 Euro/hour) and volunteers. Before a capture session, each subject gave their informed written consent for the analysis and publication of their 3D scan data including images and scans of their faces. The non-professional subjects agreed that their data could be displayed, provided that their face is blurred or masked. The scans were captured with a custom multi-camera stereo-based system (3dMD LLC, Atlanta, GA), posing no risk to the participants, using flashed texture patterns (for stereo) and white light flashes (for texture capture). Shape capture happens in about 34ms and, since it is flash-based, there is no motion blur. There is a recovery time between captures meaning that we can only capture discrete instants during breathing. Subjects wore minimal tight-fitting clothing (bike-shorts style bottoms for both men and women and a sports-bra style top for women) as shown in Fig. 3; this clothing makes shape changes during respiration readily apparent. To make later registration of scans with a common template more accurate, some of the subjects were painted in a multi-colored pattern using a water-based paint [Bogo et al. 2014].

We focus on normal breathing of the upright body in an "A pose" (Fig. 3); that is, we do not consider different activities or pose-

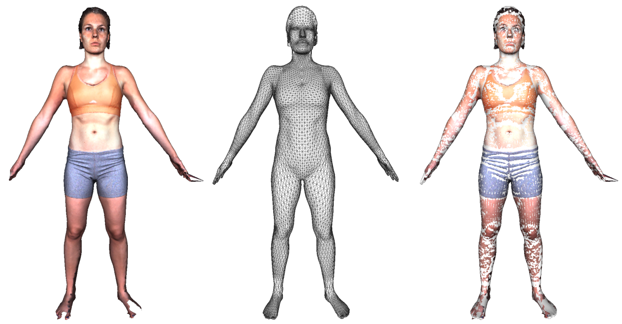

(a) *(b)* *(c)*

Figure 4: Example of 3D scan registration. *A template mesh is warped to match a high-resolution 3D scan (a). The warped template mesh at the end of the registration procedure is shown in (b). Overlaying the warped template on the scan (c), we see that the two surfaces are very close to each other (they interleave with high frequency). We followed the registration procedure described in [Bogo et al. 2014]. No landmarks were used for initialization.*

dependent changes in breathing. Subjects were informed that the study was about breathing and were instructed to breathe at what they considered a normal pace. The physiology of respiration [West 2012] leads to two main types of breathing: chest and abdominal breathing which correspond to different motions of the diaphragm. We initially asked the subjects to breathe normally. Then we explicitly asked them to focus on breathing with the chest or with the stomach. Additionally, to be able to represent the extremes of the breathing deformation, we recorded the subject shape during complete inhale and complete exhale. In total we captured and analyzed 2807 full body scans.

Although breathing is naturally a time evolving process, current high-resolution 3D body scanning systems can give us only sparse samples of this temporal process. Thus, our data consist of static 3D scans that were taken at unknown time instants of the subject's breathing activity. We address this limitation below.

3.1 Data processing

Our first step is to bring all the 3D scans into correspondence by registering (aligning) them to a 3D body template represented as a triangulated mesh (10,777 vertices, 21,550 triangles) as illustrated in Fig. 4. The detailed process is described elsewhere [Bogo et al. 2014]; the result is that all 2807 meshes are in correspondence with the template. Shapes are represented as triangle deformations from a template shape. Behind this process is a 3D parametric shape model similar to SCAPE [Anguelov et al. 2005] in that it factors body shape changes due to identity from those due to pose. We normalize all registered scans to a common pose and save the pose parameters. For each subject we compute the mean shape and, for each scan, we then compute the residual shape deformation from the mean. This constitutes our shape training data. Additionally we have the pose of each aligned scan and this is used as pose training data.

The shape and pose change during respiration is directly related to the volume of air in the lungs and the motion of the diaphragm. Consequently lung volume and diaphragm motion would provide natural controls for breathing animation. Unfortunately, neither is directly observable from the scans. What is observable, however, is mesh volume, which is easily computed from the aligned meshes using signed volumes of tetrahedra as described in [Sánta and Kato 2013]. We assume that mesh volume changes result exclusively

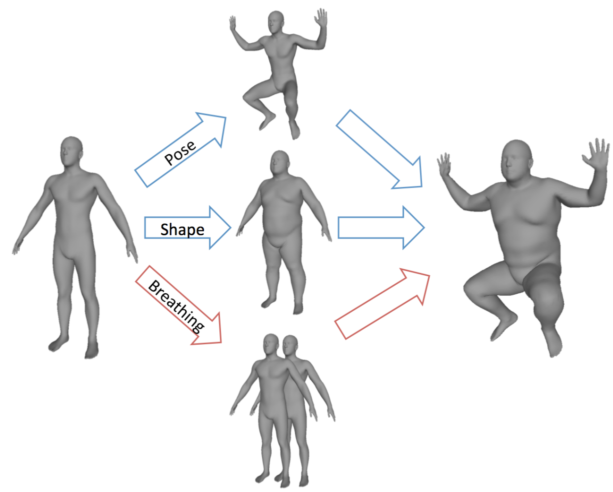

Figure 5: BreathSCAPE. *The standard SCAPE model [Anguelov et al. 2005] factors body shape into intrinsic shape and pose-dependent shape (blue). We add a new type of shape deformation for breathing and combine all three into a model with separate controls for breathing (red).*

from changes in lung volume and consequently take mesh volume (and change in volume) as a proxy for actual lung volume.

As we saw before, according to the physiology of respiration there are two main types of breathing: chest and abdominal breathing. In practice, however, people breathe in a variety of ways with varying amounts of chest and stomach deformation. While we cannot observe the diaphragm's motion, we can observe its effect on body shape. To define the type of breathing we segment the torso into an upper and lower segment of roughly equal volume. At maximal inhale we compute the difference in volume of each segment from that of the mean segment volume. The ratio of chest volume change over the total volume change defines the percentage of "chest breathing", which we refer to as the "type" of breathing.

4 Breathing Space (Shape Model)

Given a single 3D scan of a subject it is not well defined what part of the observed shape is due to breathing and what is due to the intrinsic shape of the person; e.g. do they have a large chest or are they inhaling deeply? However, given multiple scans of the same subject at different time instants in the breathing cycle, we can extract the shape and pose variations due to breathing. After registering the initial 3D scans (above), our data consist of aligned 3D meshes of multiple subjects at unknown time instants in their breathing cycle. Given a set of K 3D meshes in correspondence, $X_{ij}, i = 1, \ldots, K$, for a subject j, we extract their intrinsic shape, D^j, as well as the shape deformations due to breathing B^{ij} for each X_{ij} by extending a SCAPE body model [Anguelov et al. 2005].

SCAPE represents body shape as a deformation from a template mesh to an instance mesh using deformation gradients [Sumner and Popović 2004]. The basic idea is summarized in Fig. 5 and the reader is referred to [Anguelov et al. 2005] for details. The deformation gradients in SCAPE are represented as 3×3 deformation matrices that transform triangles, t, in a template mesh, T, into corresponding triangles in an instance mesh, X_{ij}. Since we have aligned the template with all the scans, T and X_{ij} have the same topology, and the transformation matrices are given. To reconstruct a mesh X_{ij} using the SCAPE model, three types of de-

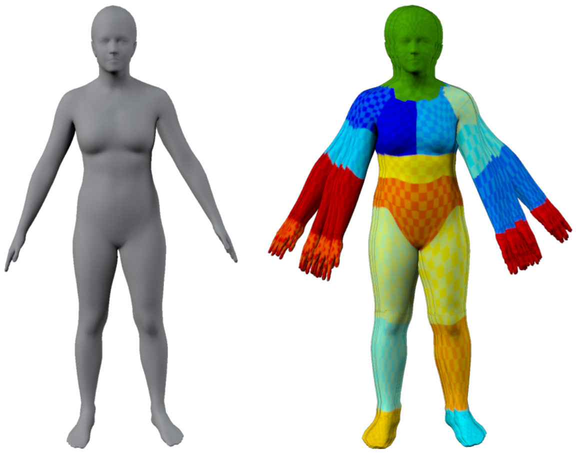

Figure 6: Pose change during breathing. *Left: Mean pose of a subject (A-pose). Right: Posture variation of the same subject during the scanning session. Some pose variation is due to breathing and some is not. Each part is color coded to show the body segmentation.*

formation gradients are applied to the triangles t of a template mesh T: pose-dependent transformations, Q_t^{ij}, identity-dependent transformations, S_t^{ij}, and rigid part rotations $R_{l[t]}^{ij}$. More specifically, given the edges $\hat{v}_{t,e}, e = 0, 1$ of each triangle t on the template, we compute the edges $v_{t,e}^{ij}, e = 0, 1$ of triangle t belonging to the i-th mesh of subject j as

$$v_{t,e}^{ij} = R_{l[t]}^{ij} S_t^{ij} Q_t^{ij} \hat{v}_{t,e} \tag{1}$$

where $l[t]$ denotes the body part to which triangle t belongs. The template mesh is segmented into distinct parts and all the triangles of the part undergo the same rotation $R_{l[t]}$; the part segmentation is illustrated in Fig. 6.

There is one extra step to SCAPE. The above equation acts on every triangle in the mesh independently, resulting in a collection of triangles that do not necessarily form a valid mesh. SCAPE adds an extra step of solving for the valid mesh with triangle deformations that best match those above; see [Anguelov et al. 2005].

Additionally, the identity dependent deformations for a population of people can be approximated as a linear combination of basis deformations learned using principal component analysis (PCA). Here we use approximately 4000 laser scans of men and women from the US and European CAESAR datasets [Robinette et al. 2002]. After registering the scans with the template, we take the deformations describing each scan and stack them in a vector. We then perform PCA on the matrix of vectors for all the registered scans. The principal components capture the directions of shape variation in the population and we can approximate any body shape deformation as a linear combination of a relatively small number of these components.

To model pose-dependent deformations we captured an registered approximately 1800 scans of people in a wide variety of poses. The registration process and our BlendSCAPE formulation is described in [Bogo et al. 2014; Hirshberg et al. 2012]. The template mesh is segmented into parts and, from the registered template meshes, we compute the rotation of each part. Using this dataset, we learn the non-rigid, pose-dependent, deformations, Q_t^{ij}, which are a function of the part rotations (see [Anguelov et al. 2005; Hirshberg et al. 2012]).

4.1 Adding Breathing

We now define the deformation matrices R, S, Q mentioned above as functions of either pose parameters, $\mathbf{r}$, or shape parameters, $\mathbf{u}$, corresponding to linear coefficients in the PCA space; that is,

$$v_{t,e}^{ij} = R_{l[t]}(\mathbf{r}^{ij}) S_t(\mathbf{u}^{ij}) Q_t(\mathbf{r}^{ij}) \hat{v}_{t,e}. \tag{2}$$

These parameters provide the animator controls to create a body shape $\mathbf{u}$ in pose $\mathbf{r}$.

One of our key contributions is to extend SCAPE by separating the identity-dependent deformations S into two parts: one due to the intrinsic shape of the person, D, and one due to breathing, B (Fig. 5). The functions D and B depend on intrinsic shape parameters, $\mathbf{d}$, and the shape parameters related to breathing, $\mathbf{b}$, respectively. Additionally, we separate the pose into static pose, $\mathbf{a}$, and, optionally, pose due to breathing, $\mathbf{c}$. Our new model, BreathSCAPE, takes the following form:

$$v_{t,e}^{ij} = R_{l[t]}(\mathbf{a}^{ij} + \mathbf{c}^{ij})(D_t(\mathbf{d}^j) + B_t(\mathbf{b}^{ij})) Q_t(\mathbf{a}^{ij} + \mathbf{c}^{ij}) \hat{v}_{t,e}. \tag{3}$$

To describe pose we use an axis-angle representation. In this representation it is meaningful to add pose parameters as long as self-intersection contraints and joint limits are not violated. Previous SCAPE models (and related models) ignore breathing deformations. Here we make them explicit. Below we show how to learn and then parameterize these by breathing type s, volume v, and gender g. We end up with a model of the following form:

$$\begin{aligned} v_{t,e} = R_{l[t]}(\mathbf{a} + \mathsf{E}(g, \mathsf{v}))(D_t(\mathbf{d}) + B_t(\mathsf{F}(s, \mathsf{v}))) \\ Q_t(\mathbf{a} + \mathsf{E}(g, \mathsf{v})) \hat{v}_{t,e} \end{aligned} \tag{4}$$

where $\mathsf{E}(g, \mathsf{v})$ returns an estimate of the breathing-dependent pose, $\mathbf{c}$, and $\mathsf{F}(s, \mathsf{v})$ returns an estimate of the breathing shape parameters, $\mathbf{b}$.

4.2 Extracting the Breathing Deformations and Pose

Given multiple scans from the subjects in our training set, our goal is to extract the intrinsic shape, D^j, of each subject as well as the shape deformations due to breathing, B^{ij}. Recall that all scans are in correspondence with the template (and hence the SCAPE model).

Consider one subject, j, with K aligned meshes X_{ij}; we seek to extract the breathing-related deformations B^{ij}. This means we want to effectively factor out pose, pose-dependent deformations, and identity to focus on what is left. The remainder should be due to breathing.

To recover the deformations for X_{ij}, we first solve for the shape deformations S^{ij} by minimizing (see [Anguelov et al. 2005])

$$\begin{aligned} \underset{S^{ij}}{\operatorname{argmin}} \sum_t \sum_{e=0,1} \| R_{l[t]}^{ij} S^{ij} Q_t^{ij} \hat{v}_{t,e} - v_{t,e}^{ij} \|_F^2 \\ + \beta \sum_{t_1, t_2 adj} \| S_{t_1}^{ij} - S_{t_2}^{ij} \|_F^2. \end{aligned} \tag{5}$$

The first term minimizes the reconstruction error between the vertices of the captured meshes and their mesh representation based on deformation gradients. The second term enforces smooth deformations between adjacent triangles. Note that the smoothness term here is applied only to the shape deformations.

Given that our meshes are in correspondence and segmented, it is easy to estimate the rigid rotation matrices $R_{l[t]}^{ij}$ between corresponding body parts in the aligned mesh. We convert the rotation

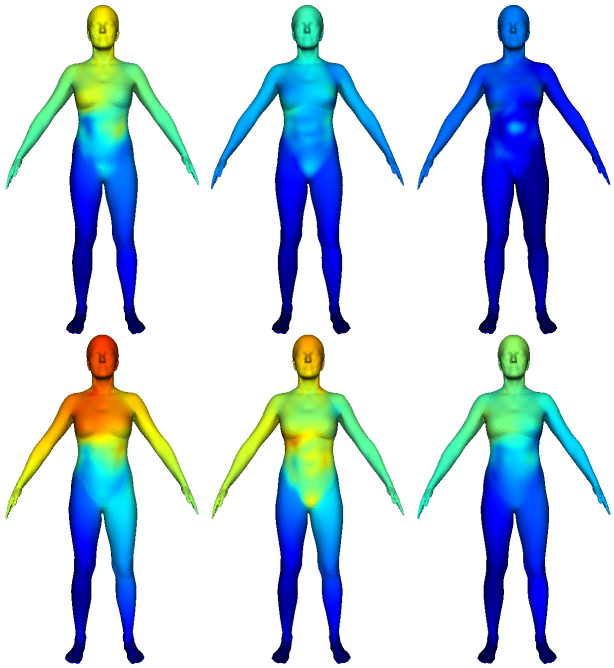

Figure 7: Shape change during breathing. *Several examples of registered meshes and how they deviate from the mean shape of the subject. Here we have pose-normalized the meshes to the A-pose. Hot colors indicate greater distance from the mean (red approximately equals 1 cm and blue 0 cm).*

matrix per body part to an axis-angle representation of pose relative to the template mesh consisting of 3 parameters. That amounts to a vector, $\mathbf{r}^{ij}$, of 57 pose parameters per mesh (3 parameters, 19 body parts). We approximate the static pose with the average pose parameters over all meshes per subject, $\mathbf{a}^j = \frac{1}{K}\sum_i \mathbf{r}^{ij}$, and the dynamic pose with the residual pose parameters $\mathbf{c}^{ij} = \mathbf{r}^{ij} - \mathbf{a}^j$. Figure 6 shows the mean pose and the variation in pose for one subject. Below we will correlate these variations with breathing to factor out pose changes that are not breathing-related.

Each subject was scanned multiple times at unknown time instants in their breathing cycle. After we estimate $S^{ij}, i = 1, \ldots, K$, we approximate the intrinsic shape of the subject as the average of the deformations, $D^j = \frac{1}{K}\sum_i S^{ij}$. Figure 7 shows example meshes and how they deviate from the mean. We found that as few as $K = 20$ scans were sufficient to extract a reasonable representation of a subject's intrinsic shape. The residual shape deformation due to breathing is $B^{ij} = S^{ij} - D^j$. We do this for all subjects in our dataset and use the residual deformations below to learn a shape model of breathing.

5 Statistics of Breathing

Respiration induces change in body shape and pose. In this section, we study the statistics of body deformations and posture variation due to breathing. In addition, we examine correlations with intrinsic attributes of humans, such as gender and intrinsic shape.

5.1 Breathing Shape Statistics

After estimating S^{ij} over all subjects, we end up with a very high-dimensional representation of the shape of each mesh. The dimensionality of S^{ij} is $9 \times F$, where F is the number of mesh triangles and 9 is the number of parameters of the 3×3 deformation gradient per triangle. Intuitively, the shape deformations due to breathing can be expressed with a much smaller number of parameters. Similar in concept to SCAPE, we learn a low-dimensional representation of shape change during breathing expressed as a linear combination of basis vectors, $G^m \in \mathbb{R}^{9F}, m = 1, \ldots, M, M \ll F$. We learn the basis vectors of breathing by computing the principal components (PCs) of the breathing deformations B^{ij} of all subjects. We then represent a breathing shape deformation using a small number of linear coefficients, $\mathbf{b}^{ij}$; these can provide breathing animation controls. Breathing deformations are approximated using the basis vectors and the linear coefficients as

$$\hat{B}^{ij} = \sum_m b_m^{ij} G^m. \tag{6}$$

Figure 8 illustrates the principal components of shape variation during breathing. Conceptually, the first two components correspond mostly to motion of the chest and the stomach, respectively. The remaining components represent higher-frequency variation of shape in the torso area. In our experiments, we have used $N = 20$ PCs which account for 76% of the variance in the data. The number of components was selected empirically; using more components does not visually improve the realism of the synthesized breathing animations. Note that there are other subject-specific shape variations and noise in the registered meshes, resulting in a fairly low signal to noise ratio in the data. Below we relate the shape variations to breathing to model how the coefficients change with breathing.

We evaluated whether breathing deformations were linearly correlated with body shape. In general the correlation coefficients are below 0.5. In particular the first few principal components of body shape are not strongly correlated with breathing shape deformations. For higher-order shape components, capturing finer details of the body (e.g. rolls of fat), we find stronger correlation with breathing deformations. Intuitively, we would expect dependence between intrinsic shape and breathing shape in areas of the body where there are prominent skin folds and fat. Here we do not model body-shape-specific breathing deformations but this is an interesting direction for future work.

5.2 Breathing Pose Statistics

As with breathing shape, we extract a low-dimensional representation, $\mathbf{p}^{ij}$ of breathing pose variation, $\mathbf{c}^{ij}$, using PCA. The low-dimensional pose representation can be expressed as $p_n^{ij} = P^{nT}\mathbf{c}^{ij}, n = 1, \ldots, N$ where $P^n \in \mathbb{R}^{57}$ are the principal components of pose variation; here we use $N = 4$ components. This results in a low-dimensional description of breathing pose, $\mathbf{p}^{ij}$. We found, however, that not all components were correlated with breathing. The subjects were allowed to "relax" between consecutive scans (20 sec) and adjusted their pose and moved their feet slightly. Consequently we discarded pose components that were not strongly correlated with breathing (i.e. volume). Figure 9 shows the three most informative principal components of pose change during breathing. As expected, they are related to spine and shoulder/neck motion during breathing. Examining the low-dimensional pose space, we did not find strong correlations with the intrinsic shape of the subject, but we did find correlation with gender. In particular, women show a more pronounced forward/backward rocking of the upper body during breathing. Consequently, we build a separate model of pose variation for men and women.

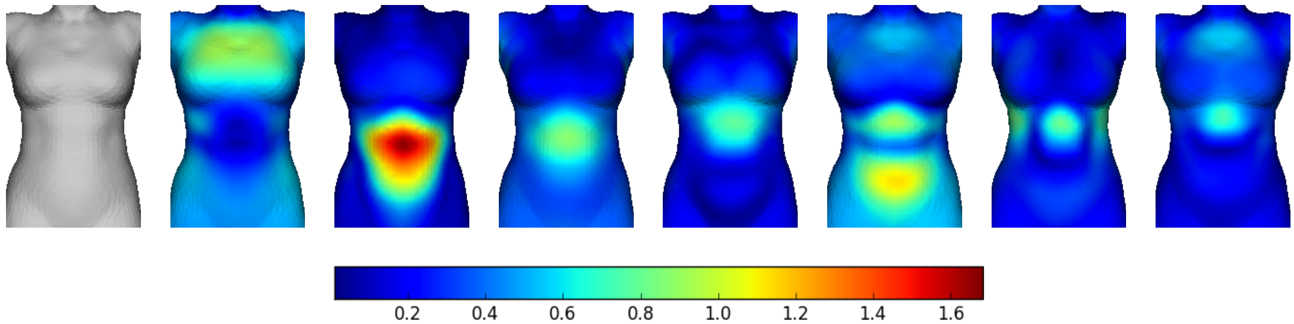

Figure 8: Principal components of shape variation during breathing. *(Gray) Mean female body. (Color) Ordered principal components (PCs) of breathing deformation shown at +5 standard deviations. Each body is color coded based on the Euclidean distance (in cm) between corresponding vertices of the mean shape and the mean deformed along the PC direction.*

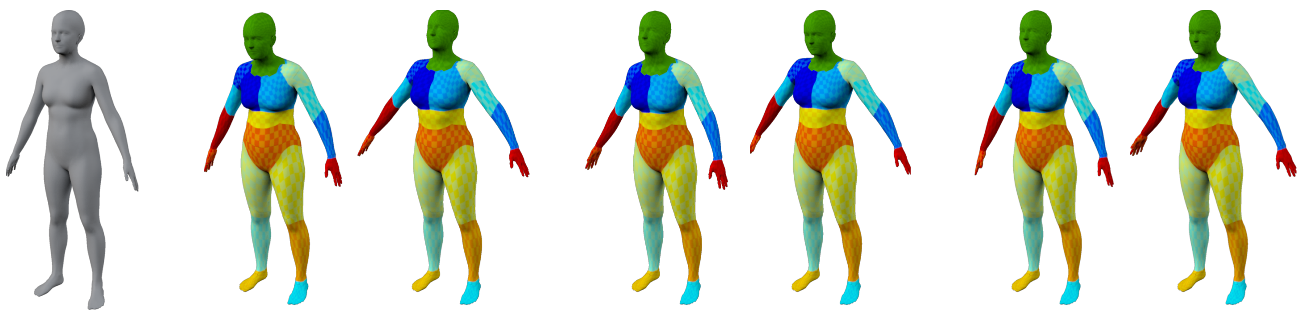

Figure 9: Principal components of pose variation during breathing. *(Gray) Mean female body. (Color) The three pose principal components most correlated with breathing (volume) displayed at +/-6 standard deviations from the mean pose.*

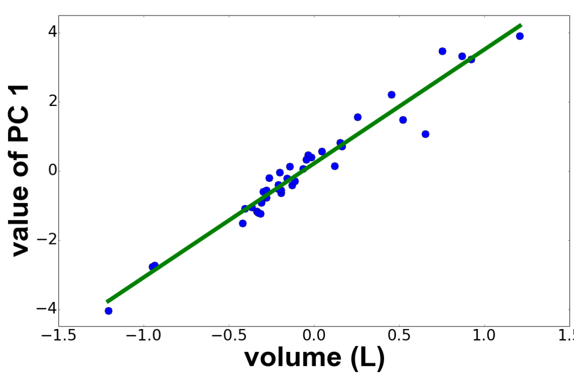

Figure 10: Volume change versus shape change. *For a specific subject, with a particular type of breathing, we find there is a linear relationship between the breathing shape coefficients and changes in mesh volume. Here we see projections to the first principal component of breathing shape for various values of mesh volume. Volume is expressed in litres (L).*

6 Breathing Model

The statistics of breathing shape and pose change do not provide a *model* for animation. What we need is a model that relates these changes in pose and shape to physiological parameters like lung volume over time. We develop our model in stages.

6.1 Shape change during breathing

Subjects were instructed to breathe in three different ways: normally, with the chest, and with the stomach. Scans from each of these conditions were treated as separate trials. Using mesh volume as a proxy for lung volume, we express shape change of a subject within a trial as a function of changes in mesh volume from the mean subject mesh. We find a largely linear relationship between the coefficients of breathing shape and mesh volume change (Fig. 10). Let $\mathbf{Z}_j \in \mathbb{R}^{K \times 2}$ be a matrix containing a column with ones and a column with the volume differences, v_{ij}, between the i^{th} mesh of subject j and the mesh corresponding to their mean shape; K is the number of meshes in the trial. Let $\mathbf{Y}_j \in \mathbb{R}^{K \times M}$ be a matrix containing the low-dimensional breathing shape coefficients, $\mathbf{b}^{ij}$, representing the breathing shape deformations of the training meshes (Sec. 5.1). For each trial, we learn a subject-specific linear model, $\mathbf{W}_j$, relating changes in breathing volume to shape deformation coefficients

$$\underset{\mathbf{W}_j \in \mathbb{R}^{2 \times M}}{\operatorname{argmin}} \|\mathbf{Z}_j \mathbf{W}_j - \mathbf{Y}_j\|_F^2. \tag{7}$$

6.2 Breathing types

In the linear model above we assume that the subject performs the same type of breathing throughout each of the three scanning sessions (normal, chest, stomach breathing). The type of breathing plays an important role in animation. However, the trial classification above provides only a crude classification of the type of breathing. To more precisely classify the type of breathing performed in a trial we use the linear function and the maximum inhale volume to predict the shape of the body at maximum volume. Using the segmentation of the torso into chest and stomach regions (Sec. 3.1),

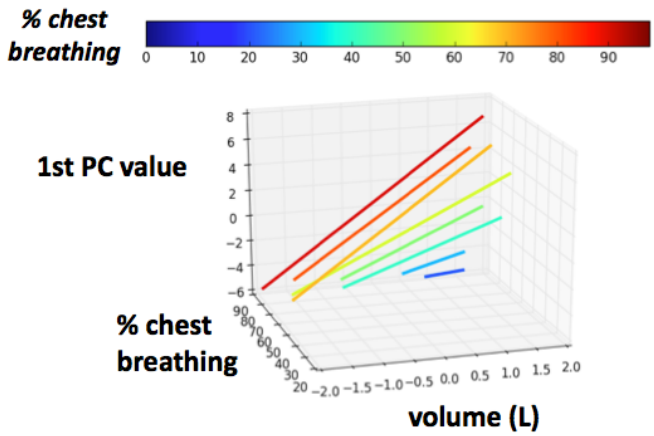

Figure 11: Shape as a function of volume and type of breathing. *Linear models of shape change during breathing for various breathing types (percentage of chest breathing) considering only the* 1^{st} *PC. Color coding is based on breathing type.*

we compute the ratio of chest volume change of this mesh from the mean to total volume change of chest and stomach. This gives a value, s_j, for each trial, indicating the percentage of chest breathing present in that trial.

6.3 Breathing shape model

Finally, we have what we need to learn a function, $\mathsf{F}(s, \mathsf{v})$, that takes as input the breathing type s and volume difference v and returns the corresponding linear shape deformation coefficients. Given the classification of breathing type above, we divide the trials into 10 categories corresponding to 0%-10%, 10%-20%, ..., and 90%-100% chest breathing. Within each category we combine all the individual linear models into an aggregate linear model relating each shape coefficient to change in volume. This aggregate model can be thought of as the average linear relationship predicting shape change from volume change.

Figure 11 shows what this looks like for the first principal component. Each colored line is an aggregate linear model for a specific value of breathing type, s. Note that the slope of each line is different. Recall that the first principal component captures mostly chest deformation (Fig. 8). The higher the value of s, the more the chest is involved, and the greater the correlation of the first component with changes in volume. Note further that this results in a function that is non-linear in s and v.

We want a model of breathing that is continuous in s and v and we achieve this by fitting a surface to the changing regression functions using cubic interpolation. Figure 12 shows some examples of the resulting functions $w_m(s, \mathsf{v})$. As we saw before, the first two principal components are highly correlated with chest and stomach breathing respectively. This is evident in the corresponding weight functions (left two subplots in Fig. 12).

In the final breathing shape model then, we weight the principal components of breathing, G^m, by a non-linear function $w_m(s, \mathsf{v})$ which predicts the breathing coefficients

$$\mathsf{F}(s, \mathsf{v}) = \sum_{m=1}^{M} w_m(s, \mathsf{v}) G^m. \tag{8}$$

Figure 13 shows two synthesized meshes at maximum inhale: a female breathing with the stomach and a male breathing with the chest.

6.4 Breathing pose model

Based on the insights from Sec. 5.2, we derive a generic model of pose change per gender g, $\mathsf{E}(g, \mathsf{v})$, parameterized additionally by the breathing volume v. Let $\mathbf{O}_g \in \mathbb{R}^{K_g \times 2}$ be a matrix containing a column with ones and a column with the volume differences, v_{ij}, over all subjects j of gender $g = \{\text{male, female}\}$. Let $\mathbf{H}_g \in \mathbb{R}^{K_g \times N}$ be a matrix containing the PCA projections of pose, $\mathbf{p}^{ij}$, corresponding to the training meshes as described in Sec. 5.2. For each gender, we define a linear model for predicting breathing induced pose deformations using linear least squares regression:

$$\operatorname*{argmin}_{\mathbf{L}_g \in \mathbb{R}^{2 \times N}} \|\mathbf{O}_g \mathbf{L}_g - \mathbf{H}_g\|_F^2. \tag{9}$$

We then define $\mathbf{h}(g, \mathsf{v}) = [1, \mathsf{v}]\mathbf{L}_g$ and

$$\mathsf{E}(g, \mathsf{v}) = \sum_{n=1}^{N} h_n(g, \mathsf{v}) P^n. \tag{10}$$

Here we do not model any dependence on breathing type. Note also that the pose model does not need to be used for animation; for example, when animating the breathing of a moving character, we do not use the pose model.

7 Breathing Animation

Respiration is time varying. In particular, as air moves in and out of the lungs, their volume changes. To animate breathing using the model defined above, we need a way to vary lung volume over time.

7.1 Trajectory editing

We developed a Maya tool (not shown here) to create and edit realistic 3D body shapes that is similar to previous work on body shape modeling [MPI IS 2011; Allen et al. 2003; Anguelov et al. 2005; Hasler et al. 2009; Jain et al. 2010]. The tool also allows an animator to edit the temporal pattern of breathing.

Our breathing model takes two inputs: the breathing type and volume difference. Our interface includes a slider with which the animator selects the percentage of chest breathing enabling them to achieve different “styles.” A common assumption in the physiology of respiration [West 2012; Mines 1981] is that air flow in lungs during breathing at rest pose is a sinusoidal function of time. Thus we provide an interface for controlling the parameters of a sinusoid function of volume over time. The intensity of the pose change can be adjusted separately from shape deformation. The amplitude and frequency of the sinusoid can be varied using sliders.

7.2 Breath acting

Breathing in real life does not always follow a pure sinusoidal function. It varies with activity and emotion and plays a role in telling a story. We use a device call a spirometer (ndd Medizintechnik AG, Zurich), which measures change in air flow, to capture the breathing pattern of a “breath actor” (Fig. 14). Like a voice actor, the breath actor observes an animation and acts out the breathing that fits the action. We then use the recorded changes in volume to produce deformations (and possibly pose changes) and to animate a 3D avatar. We manually adjust the breathing type (chest or stomach) based on the action and emotions in a scene. This approach provides a simple and intuitive interface to produce realistic and compelling breathing animations.

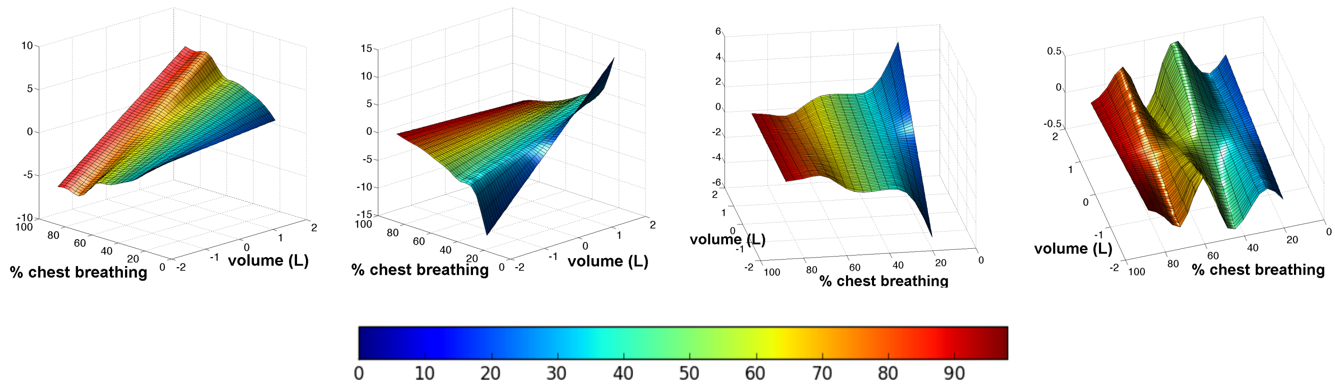

Figure 12: Shape change during breathing as a function of breathing type and volume. *The breathing type is expressed as the percentage of chest relative to stomach breathing. The first two PCA coefficients of shape change (first two plots from the left) are correlated with "chest-breathing" and "abdominal-breathing" respectively. The third and fourth PCA coefficients (two rightmost plots) reveal a more complex relationship. Color coding is based on breathing type (% of chest breathing).*

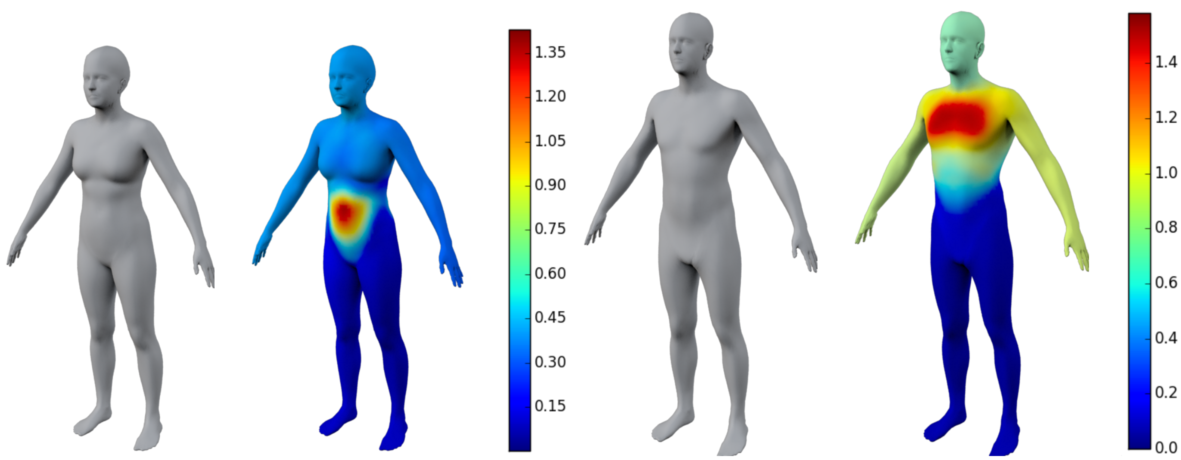

Figure 13: Examples of synthesized breathing. *"Abdominal-breathing" (left) and "chest-breathing" (right). The gray bodies represent the mean shape of two subjects. Shape change due to breathing is color coded based on the Euclidean distance (in cm) between every mesh vertex and the surface of the mean shape.*

8 Results

To evaluate the realism of our model, we compare with a reference video of a subject breathing. We had a breath actor (different from the subject) watch the video and imitate the breathing using the spirometer. We used a 3D scan of the subject to create an avatar for their body shape, selected the amount of chest breathing manually, and then animated the body in the style of the subject. Note the we did not capture the pose of the actor during breathing and did not attempt to match the pose. Focusing on the breathing deformations, however, we find a good qualitative match between reference and animation (Fig. 15).

8.1 Breathing in Action: Poses and Motion

Our model of respiration is trained using body scans of people in a standing "A" pose. While the pose variation model may be quite specific to this pose, the shape deformation model can be easily applied to other poses using BreathSCAPE with realistic results. Figure 16 shows a body in a seated pose and a standing pose with the same breathing model applied. Notice that the breathing deformations are, in fact, different because the mesh is in a different pose. The animation of breathing looks realistic.

We also animate the breathing of characters in motion. Figure 1 shows frames from a running sequence with the breathing deformation color coded in terms of distance from the animated average shape. In this case a breath actor observed the animated body without breathing and simulated the breathing to go with it. While the running motion makes it harder to see the breathing animation, one can readily tell the difference between sequences animated with and without breathing.

9 The Last Breath (Conclusions)

We describe a model for realistic breathing animation. A key novelty of our approach is the use of high-resolution 3D scans in combination with a human body model to capture pose and shape change during breathing. By analyzing the statistics of breathing shape changes we found that: 1) there are statistically significant changes in whole-body posture and shape during breathing, 2) the differences in breathing shape between men and women are not great but that there are some significant postural differences; 3) the dominant breathing shape changes are somewhat independent of body shape but more detailed changes are correlated with body shape; 4) people can perform different types of breathing (chest

Figure 15: Reference video. *Example frames from a reference video of a subject breathing with different styles. On the left of each image pair we see the recorded motion. On the right, we show a roughly corresponding frame from our animation.*

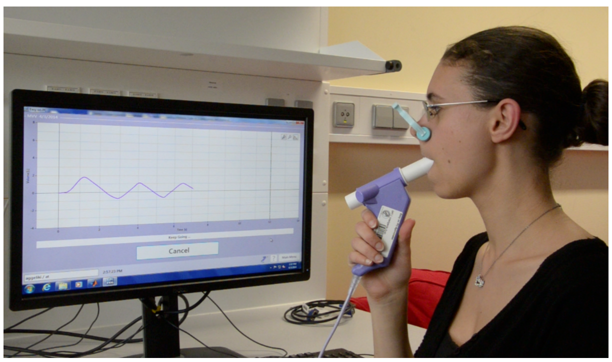

Figure 14: Breath acting. *An actor breathes into a spirometer to convey the action and emotional content of a character. The changing volume of the lungs is recorded and used to animate breathing.*

and stomach) and these are clearly reflected in the principal components of breathing shape; 5) for a particular type, lung volume is linearly related to these principal components; 6) this linear relationship varies with type, resulting in a non-linear model. Based on this analysis, we learn a low-dimensional model of breathing shape and pose variation that is parameterized by breathing volume and the type of breathing.

We capture breathing in a fixed pose but clearly shape changes will be influenced by posture (e.g. lying down). Future work should study how pose affects breathing deformations. Breathing shape is likely also correlated with activity and it would be good to build a temporal model of breathing *dynamics* as it relates to pose changes during activity. We used mesh volume to measure lung volume but it would be interesting to synchronize the output of a spirometer directly with the 3D scanning process. We focused on the two dominant types of breathing described in the literature but we would like to capture a much wider range of scenarios including other actions like puff, pant, blow, gasp, wheeze, sigh, huff. Our methods could be use to give an animator the ability to select among these styles. We have focused on the body but it would be interesting to simultaneously analyze facial motions, which are also influenced by breathing. In contrast to our model-based approach, one could also explore example-based methods that are popular with motion capture data [Hsu et al. 2004; Park and Shin 2004] but have not been applied to 3D breathing shape. As mentioned, the relationship between body shape and breathing deformations deserves further study. Animation of our model from marker data could also be done. Finally, our scan data does not reveal the detailed temporal evolution of breathing; either new scanning methods are needed or possibly one could use a combination of scanning and marker-based motion capture.

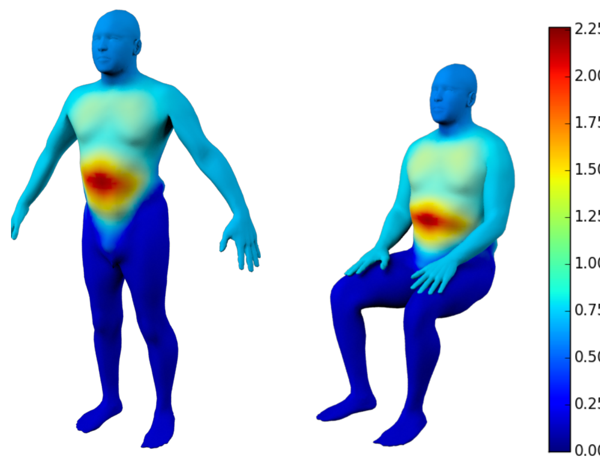

Figure 16: Pose and motion. *We apply the breathing model with 40% chest breathing to a standing and seated pose. The body corresponds to a person in the CAESAR dataset. Color coding represents the distance (in cm) between the vertices of the meshes at full inhale and the original shape of the person.*

Acknowledgments. We thank F. Bogo, J. Romero, M. Loper for sharing their code for mesh registration and 3D human modeling, E. Holderness for data collection, S. Lupas for breath acting, J. Anning for video recordings, 3dMD, and the anonymous reviewers.

References

ALLEN, B., CURLESS, B., AND POPOVIĆ, Z. 2003. The space of human body shapes: Reconstruction and parameterization from range scans. *ACM Trans. on Graphics (TOG) 22*, 3, 587–594.

ALLEN, B., CURLESS, B., POPOVIĆ, Z., AND HERTZMANN, A. 2006. Learning a correlated model of identity and pose-dependent body shape variation for real-time synthesis. In *ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA)*, Eurographics Association, 147–156.

ANGUELOV, D., SRINIVASAN, P., KOLLER, D., THRUN, S., RODGERS, J., AND DAVIS, J. 2005. SCAPE: Shape Completion and Animation of PEople. *ACM Trans. on Graphics (TOG) 24*, 3, 408–416.

BOGO, F., ROMERO, J., AND BLACK, M. J. 2014. FAUST: Dataset and evaluation for 3D mesh registration. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE.

CAVAGNARO, M., PITTELLA, E., AND PISA, S. 2013. Anatomical models of breathing subjects for absorption and scattering analysis. In *International Symposium on Electromagnetic Compatibility (EMC EUROPE)*, IEEE, 149–153.

CHEN, Y., LIU, Z., AND ZHANG, Z. 2013. Tensor-based human body modeling. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, 105–112.

COSKER, D., AND EDGE, J. 2009. Laughing, crying, sneezing and yawning: Automatic voice driven animation of non-speech articulations. In *Computer Animation and Social Agents (CASA)*.

DILORENZO, P. C., ZORDAN, V. B., AND SANDERS, B. L. 2008. Laughing out loud: Control for modeling anatomically inspired laughter using audio. *ACM Trans. on Graphics (TOG) 27*, 5, 125–133.

EGGES, A., MOLET, T., AND MAGNENAT-THALMANN, N. 2004. Personalised real-time idle motion synthesis. In *Pacific Conf. on Comp. Graphics and Applications*, IEEE, 121–130.

EGGES, A., GIACOMO, T., AND THALMANN, N. 2006. Synthesis of realistic idle motion for interactive characters. *Game Programming Gems 6*, 3–2.

HASLER, N., STOLL, C., SUNKEL, M., ROSENHAHN, B., AND SEIDEL, H. 2009. A statistical model of human pose and body shape. *Computer Graphics Forum 28*, 2, 337–346.

HIRSHBERG, D., LOPER, M., RACHLIN, E., AND BLACK, M. 2012. Coregistration: Simultaneous alignment and modeling of articulated 3D shape. In *European Conf. on Computer Vision (ECCV)*, Springer-Verlag, LNCS 7577, Part IV, 242–255.

HSU, E., GENTRY, S., AND POPOVIĆ, J. 2004. Example-based control of human motion. In *ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA)*, Eurographics Association, 69–77.

JAIN, A., THORMÄHLEN, T., SEIDEL, H.-P., AND THEOBALT, C. 2010. MovieReshape: Tracking and reshaping of humans in videos. *ACM Trans. on Graphics (TOG) 29*, 6, 148:1–148:10.

KIDER, J., POLLOCK, K., AND SAFONOVA, A. 2011. A data-driven appearance model for human fatigue. In *ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA)*, ACM, 119–128.

LEE, S.-H., SIFAKIS, E., AND TERZOPOULOS, D. 2009. Comprehensive biomechanical modeling and simulation of the upper body. *ACM Trans. on Graphics (TOG) 28*, 4, 99:1–99:17.

LEE, D., GLUECK, M., KHAN, A., FIUME, E., AND JACKSON, K. 2012. Modeling and simulation of skeletal muscle for computer graphics: A survey. *Foundations and Trends in Computer Graphics and Vision 7*, 4, 229–276.

MAGNENAT-THALMANN, N., ZHANG, J. J., AND FENG, D. D. 2009. *Recent Advances in the 3D Physiological Human*. Springer.

MINES, A. H. 1981. *Respiratory Physiology*. Raven Press.

MORENO, A., CHAMBON, S., SANTHANAM, A. P., BROCARDO, R., KUPELIAN, P., ROLLAND, J. P., ANGELINI, E., AND BLOCH, I. 2007. Thoracic CT-PET registration using a 3D breathing model. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer, 626–633.

MPI IS, 2011. `http://bodyvisualizer.com/`.

NEUMANN, T., VARANASI, K., HASLER, N., WACKER, M., MAGNOR, M., AND THEOBALT, C. 2013. Capture and statistical modeling of arm-muscle deformations. *Computer Graphics Forum 32*, 285–294.

PARK, S. I., AND HODGINS, J. K. 2006. Capturing and animating skin deformation in human motion. *ACM Trans. on Graphics (TOG) 25*, 3, 881–889.

PARK, S. I., AND HODGINS, J. K. 2008. Data-driven modeling of skin and muscle deformation. *ACM Trans. on Graphics (TOG) 27*, 3, 96:1–96:6.

PARK, M. J., AND SHIN, S. Y. 2004. Example-based motion cloning. *Comp. Anim. and Virtual Worlds 15*, 3-4, 245–257.

PENNE, J., SCHALLER, C., HORNEGGER, J., AND KUWERT, T. 2008. Robust real-time 3D respiratory motion detection using Time-of-Flight cameras. *International Journal of Computer Assisted Radiology and Surgery 3*, 5, 427–431.

PROMAYON, E., BACONNIER, P., AND PUECH, C. 1997. Physically-based model for simulating the human trunk respiration movements. In *Computer Vision, Virtual Reality and Robotics in Medicine and Medical Robotics and Computer-Assisted Surgery (CVRMed-MRCAS)*, Springer, 379–388.

ROBINETTE, K., BLACKWELL, S., DAANEN, H., BOEHMER, M., FLEMING, S., BRILL, T., HOEFERLIN, D., AND BURNSIDES, D. 2002. Civilian American and European Surface Anthropometry Resource (CAESAR) final report. Tech. Rep. AFRL-HE-WP-TR-2002-0169, US Air Force Research Laboratory.

SANDERS, B., DILORENZO, P., ZORDAN, V., AND BAKAL, D. 2009. Toward anatomical simulation for breath training in mind/body medicine. In *Recent Advances in the 3D Physiological Human*. Springer, 105–119.

SÁNTA, Z., AND KATO, Z. 2013. Correspondence-less non-rigid registration of triangular surface meshes. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2275–2282.

SANTHANAM, A., PATTANAIK, S., ROLLAND, J., IMIELINSKA, C., AND NORFLEET, J. 2003. Physiologically-based modeling and visualization of deformable lungs. In *Pacific Conf. on Comp. Graphics and Applications*, IEEE, 507–511.

STOLL, C., GALL, J., DE AGUIAR, E., THRUN, S., AND THEOBALT, C. 2010. Video-based reconstruction of animatable human characters. *ACM Trans. on Graphics (TOG) 29*, 6, 139:1–139:10.

SUMNER, R. W., AND POPOVIĆ, J. 2004. Deformation transfer for triangle meshes. *ACM Trans. on Graphics (TOG) 23*, 3, 399–405.

VELTKAMP, R. C., AND PIEST, B. 2009. A physiological torso model for realistic breathing simulation. In *Modelling the Physiological Human*. Springer, 84–94.

WEST, J. 2012. *Respiratory Physiology: The Essentials*. Lippincott Williams & Wilkins.

ZORDAN, V. B., CELLY, B., CHIU, B., AND DILORENZO, P. C. 2004. Breathe Easy: Model and control of simulated respiration for animation. In *ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA)*, 29–37.