

Laughing Out Loud: Control for modeling anatomically inspired laughter using audio

Paul C. DiLorenzo*
University of California, Riverside
DreamWorks Animation SKG

Victor B. Zordan†
University of California, Riverside

Benjamin L. Sanders‡
University of California, Riverside

Abstract

We present a novel technique for generating animation of laughter for a character. Our approach utilizes an anatomically inspired, physics-based model of a human torso that includes a mix of rigid-body and deformable components and is driven by Hill-type muscles. We propose a hierarchical control method which synthesizes laughter from a simple set of input signals. In addition, we present a method for automatically creating an animation from a soundtrack of an individual laughing. We show examples of laugh animations generated by hand-selected input parameters and by our audio-driven optimization approach. We also include results for other behaviors, such as coughing and a sneeze, created using the same model. These animations demonstrate the range of possible motions that can be generated using the proposed system. We compare our technique with both data-driven and procedural animations of laughter.

CR Categories: I.3.1 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation; I.6.8 [Simulation and Modeling]: Types of Simulation—Animation

Keywords: human simulation, human animation, laughter

1 Introduction

Given how entertaining and contagious laughter is, it is surprising that there has been virtually no attention placed on the animation of such a common and important behavior. In our survey of several animated films, the effect of laughter is often created with exaggerated secondary actions, such as throwing the head back and slapping one’s leg. However, the motion of the trunk, where laughter is actually generated, is often simplified, and largely symbolic. We present a novel technique to generate and control laughter using a physically based, anatomically inspired torso simulation. We use this system to synthesize motion of the trunk and to create secondary effects that propagate to other parts of the body, e.g. the arms and the head. A hybrid set of rigid and flexible components comprise the model of the torso: spine, ribs, diaphragm, and abdomen as well as neck and shoulders. We employ hierarchical, Hill-type muscles to actuate laughter movement. The result is a

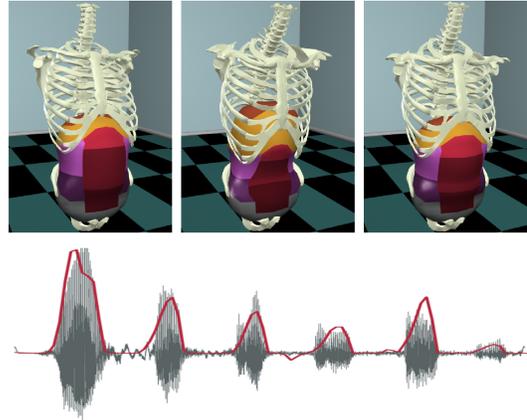


Figure 1: Sample output from our laughter simulation and a plot of an input audio and the simulated pressure for a corresponding laugh animation found with our optimization.

rich, controlled motion derived from a simple, intuitive set of control parameters which we can use to demonstrate a range of laughing animations.

Further, by taking advantage of biomedical research in laughter, we offer a second, innovative method for controlling the described system automatically starting from a source audio track of laughing. Our key insight, which is described in more detail in the body of the paper, is that lung pressure is a direct reflection of the strength and phase of the laughter which can be assessed from the time varying amplitude of an audio signal. From this observation, we infer that we can determine the pressure required for a desired recording of laughter. By adding a pressure-based model of the lungs to our torso simulation, we formulate an optimization problem which determines activations to produce the desired pressure profiles. We use this technique to generate animation that syncs well with a given audio track. Notably, this approach is aligned with common practices used in production animation where movement is often made to match a pre-recorded soundtrack.

Our approach differs in important ways from pure data-driven techniques, such as the impressive skin capture examples shown in [Park and Hodgins 2006] which attempt to precisely recreate a given example motion. Foremost, we argue that our approach offers a strong advantage in control over the motion. It is difficult to modify a surface model of a giggle to be more uproarious, or even to be somewhat longer, without a reasonable model of laughter. In addition, with our emphasis placed on controlling laughter with an audio clip, we support the accepted practice of separating the motion generation (or capture) from the sound recording. The matching of a pre-recorded animation with a pre-recorded soundtrack is a dizzyingly complex proposition. In addition, our anatomically inspired, physically-based torso is superior to a procedural model, such as one which uses a weighted blending technique. During laughter, there is rich interplay between the subsystems of the torso

*e-mail: pdiloren@cs.ucr.edu

†e-mail: vbz@cs.ucr.edu

‡e-mail: bsanders@cs.ucr.edu

(i.e. abdominal cavity, ribcage, clavicles, and spine) which changes over the course of a single bout of laughter and across the range of possible laughing behaviors (a little giggle to a deep-belly uproar.) The degree of effects depends on the intensity of the contraction and the pressure in the gut as well as the contraction of the other components. Propagation is intensified if the contraction speed exceeds the damping effects of any ‘link’ in the chain. Furthermore, our method could be made to react to external forces in a physically plausible manner. A procedural approach would not be a suitable method for this type of interaction.

The primary contributions of this paper are: 1) we introduce an anatomically motivated model for the generation of laughter motion; 2) we propose a straightforward method for controlling the activation of the model using simple, intuitive input signals; and 3) we provide an approach for generating inputs to control the system automatically using an audio soundtrack of laughing.

2 Background

Anatomical modeling is not a new concept in animation and researchers have identified a wide array of benefits from adding biomechanical realism to kinematic and dynamic systems for motion generation. Several researchers have focused on deriving the shapes of muscles based on procedural rules [Scheepers et al. 1997; Wilhelms and Gelder 1997; Nedel and Thalmann 2000; Dong et al. 2002; Pratscher et al. 2005], while others have investigated various physics-based models for generating muscle deformations [Chen and Zeltzer 1992; Teran et al. 2005; Lemos et al. 2005]. The shape of muscles, particularly skeletal muscles, changes the outward appearance of the character and gives the impression of a rich, layered bodily physique. Another approach is to exploit biomechanical structure to generate movement. For example, synthesis of anatomically based hand motion has been derived from the interaction of muscles, tendons, and bone [Albrecht et al. 2003; Tsang et al. 2005]. Additionally, several researchers have described such methods for generating muscle-based systems for the face [Platt and Badler 1981; Waters 1987; Lee et al. 1995; Kähler et al. 2002; Sifakis et al. 2005].

Our goal of animating laughter centers on the movement of the trunk. In the human body, laughter erupts as a connection of movements throughout the thoracic and abdominal regions which demands that we generate a complete torso model. Researchers have focused on individual components, such as the spine [Monheit and Badler 1991], the neck [Lee and Terzopoulos 2006], the abdominal area [Promayon et al. 1997], and the lungs [Kaye et al. 1997]. However, our interest requires that these elements be compiled into a single system that captures the interplay of the subsystems. One example of a more complete model is described for breathing simulation [Zordan et al. 2004], but this system lacks sufficient articulation and control to achieve the range of movements required for a laughing simulation. To this end, according to our investigations, the anatomically based simulation we present for laughter is the most complete animated version of the human torso to date.

Our torso includes pseudo- muscle fibers that control the ribcage, abdomen, and shoulders. The *manual* activation of several hundred independent degrees of freedom is not feasible, but simplification of the activation that exploits symmetry and redundancy provides one option for managing complexity. A nice example of this technique is the simulation control of swimming fish [Tu and Terzopoulos 1994] where the authors reduce the control of approximately ninety spring elements to a manageable set of four control parameters and then tune these by hand. We show results using a similar technique where basic motions can be generated using four intuitive parameters. However, for our soundtrack-driven laughter simulations, we

employ an optimization to compute the precise inputs needed to generate the desired animation. This addition is analogous to the follow-on work on learned behaviors for fish and other marine animals seen in [Grzeszczuk and Terzopoulos 1995]. Another notable likeness is found in the work of Sifakis et al. [2005] on automatic determination of muscle activation from motion capture. In their research, a pre-recorded (human) data signal provides a sparse input about the desired action. And the anatomical model, of the face in their case, is employed to synthesize the final motion. In our case, the soundtrack is treated as the sparse signal for optimization.

A potentially fruitful alternative would be to record (sparse) data from motion capture data, i.e. of an individual laughing, and to use this motion to inform the simulation, much like Sifakis et al. This approach could likely leverage off of the various techniques set forth for full body capture and animation [Allen et al. 2002; Sand et al. 2003; Seo and Magnenat-Thalmann 2003; Angelov et al. 2005; Park and Hodgins 2006]. However, a strength of our system is that we can apply motion to any skeleton, not only the actor being recorded. Thus our technique is well-suited for animating a wide variety of characters while avoiding the potentially difficult task of re-targeting a recorded laughter embedded in a high-resolution surface model.

3 Laughing Mechanics

Laughter is a highly dynamic action that stems from involuntary contractions in the abdominal and respiratory muscles. During bouts of laughter, regular breathing is interrupted and several subsystems of the body are affected including vocal function and bodily and facial activities. Our focus is on the respiration and abdomen dynamics which drive laughter. In this section, we describe the mechanics of laughter in two parts, first based on the muscle activity which leads to the visible movement of the torso, and second based on the pressure and air flow which yield the sounds associated with laughter. For more details on basic respiratory function we refer readers to [Mines 1993; West 2004]. For a more in-depth overview of laughter mechanics, see [Filippelli et al. 2001].

Muscle Perspective. Laughter can be initiated at any point during the respiratory cycle. When laughter begins it is marked by a quick drop in lung volume. This drop stems from a severe contraction in the abdominal muscles which push on the diaphragm and in the inner intercostals which act to collapse the ribcage. Both sets of these so-called “expiratory” muscles contribute to reduction in the chest wall volume which drastically increases the (pleural) pressure surrounding the lungs. The resulting increase in lung pressure causes increased expiratory flow. Along with this drop, quick pulses of expiratory flow are also generated, appearing around the frequency of 5 Hz. Expiratory muscles act in concert to create these small bursts as they continue to contract. Filippelli and colleagues propose that the diaphragm remains active to behave like a load balancer, protecting the lungs from extreme forces generated by the abdominal muscle contraction [Filippelli et al. 2001]. Laughter ends at, or close to, Residual Volume (RV), at which point the expiratory muscles cease their contractions.

Pressure Perspective. There is a tight link between the lung activity derived from the expiratory muscle contractions and the sounds produced during laughter. This connection is driven by lung pressure and controlled via air flow. Expiratory air flow has the effect of dissipating lung pressure, while restriction of air flow (i.e., turbulence in the trachea) has the effect of building lung pressure. In related fields of study, laughter between inhales is defined as a *bout* of laughter. Inside a bout are multiple *calls* of laughter, each defined as a single ‘ha’ (or a single expiratory flow pulse in the case of silent laughter.) Studies report synchronized modulation of air

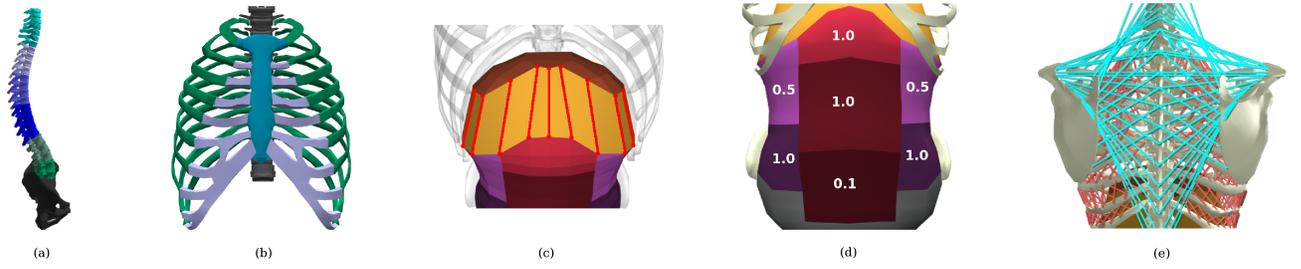


Figure 2: Our torso simulation is broken down into five main components: spine (a), ribcage (b), diaphragm (c), abdominal cavity (d), and clavicles (e). We use a mixed combination of Hill-type muscles for the muscle elements in the ribcage, diaphragm, abdomen, and clavicles and torque actuators, applied to the spine.

flow over the duration of a single call - which matches the 5 Hz frequency pulses we see in the expiratory muscles. We exploit this relationship in our audio-driven laughter and provide more detail in this regard in Section 6.

4 Torso Simulation

Our system follows the design and implementation described in Zordan et al. [2004] with a few notable exceptions. We highlight differences in this section. For a complete description of the torso simulation, we refer interested readers to [DiLorenzo 2008].

As seen in the coloring of Figure 2(a), we break up the lumbar into two sections (L1-L3 and L4-L5,) the thoracic into three sections (T1-T4, T5-T8, T9-T12,) and the cervical into two sections (C1-C3, C4-C7.) The spine is rooted at the pelvic girdle and each joint is treated as a ball joint. In Figure 2(b), we attach the cartilage to the sternum using universal joints. We explicitly model the central tendon (shown in rust color at the top of the diaphragm in Figure 2(c).) The diaphragm, which is shaped like an inverted bowl, contracts largely along the ‘sides’ of the bowl while the central tendon maintains its shape. Our diaphragm is modeled so that the muscle elements are aligned in the direction of contraction making the diaphragm more effective as a “pump” for the lungs. The bright red line segments in Figure 2(c) highlight the contractile muscles of the diaphragm. According to [Hoit et al. 1988], the pattern of activation during laughter is not uniform. We use their findings to identify five regions in the abdomen (shown in multiple colors in Figure 2(d)) that we activate proportionately, as shown, based on their findings. To capture the interplay of the clavicle, ribcage, and spine we attach a representative set of muscles that have proper origin and insertions (see Figure 2(e)). These muscles propagate activity between the components.

5 Hierarchical Control System

To control the torso simulation, we propose a three-tier hierarchical controller. At the lowest level, we compute forces using Hill muscle elements for the respiratory muscles, the abdomen, and the muscles attached to the clavicles. At the mid-level, we group muscle components and activate them in unison. For example, all the muscle elements in the diaphragm receive the same inputs, even though each is applying a different force as appropriate to its specific conditions. At the highest level, we generate activation signals based on the high-level characteristics of the motion.

5.1 Hill muscle model

We apply force through the muscle elements using a Hill-type muscle model. Our formulation is derived from insights drawn from

previous research on Hill-type muscle modeling, particularly [Zajac 1989; Ng-Thow-Hing 2001; Buchanan et al. 2004]. The basic structure of the muscle model includes a passive element (PE) which contributes forces when the muscle is stretched and the contractile element (CE) which contributes force when the muscle is activated. The total force can be written as $F_M = F_{PE} + F_{CE}$.

The terms of the passive and contractile forces vary from model to model. For our passive element, we define

$$F_{PE} = \max(0, k_{PE} * \frac{e^{10(\tilde{\ell}-1)}}{e^5} + b_{PE}\dot{s}),$$

where k_{PE} is the stiffness coefficient, b_{PE} is the damping coefficient, $\tilde{\ell} = \ell/\ell_o$ is the normalized muscle fiber length with ℓ and ℓ_o the muscle element length and the rest muscle length, respectively. $\dot{s} = v/\ell_o$ is the strain rate of the muscle with v being the muscle contraction velocity. Our stiffness term follows the function described by [Schutte 1992] (as reported by [Buchanan et al. 2004]) and the damping term follows the suggestions of [Hatze 1981] (as reported by [Ng-Thow-Hing 2001]).

The contractile force is modeled as $F_{CE} = a(t)\tilde{F}_\ell(\tilde{\ell})\tilde{F}_v(\tilde{v})F_{max}$, where $a(t)$ is the time-varying activation level, $\tilde{F}_\ell(\tilde{\ell})$ is the dimensionless force-length relationship of the muscle, and $\tilde{F}_v(\tilde{v})$ is the dimensionless force-velocity relationship with $\tilde{v} = v/v_{max}$. Each of these terms is forced to stay between the values of 0 and 1. As such, F_{max} , the maximum contractile force, controls the strength of the element. Following the curve suggested by [van den Bogert and Nigg 1999] (as reported by [Ng-Thow-Hing 2001]),

$$\tilde{F}_\ell(\tilde{\ell}) = 1 - \left(\frac{\tilde{\ell} - 1}{w}\right)^2.$$

Per [Zajac 1989], the range of force generation is 0.5 to 1.5, thus we set w to be 0.5. We adapt the function by [Hill 1970] (as reported by [Buchanan et al. 2004]) to a dimensionless term for $\tilde{F}_v(\tilde{v})$:

$$\tilde{F}_v(\tilde{v}) = \frac{b - a|\tilde{v}|}{b + |\tilde{v}|}$$

and set a and b to be 0.25 based on discussion by Buchanan et al. Plots of $\tilde{F}_\ell(\tilde{\ell})$ and $\tilde{F}_v(\tilde{v})$ appear in Figure 3. In total, our model yields a muscle element that is tunable with four parameters k_{PE}, b_{PE}, v_{max} , and F_{max} - values for these terms appear in the table below. More details about our Hill-type muscle model are described in [DiLorenzo 2008].

5.2 Activation for Laughter

According to Filippelli [2001], during a bout of laughter, “all expiratory muscles are well coordinated” and (they suggest) the di-

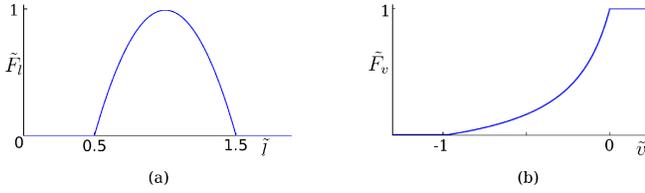


Figure 3: Dimensionless multipliers for our Hill model for (a) Force-length relationship and (b) Force-velocity relationship.

	k_{PE}	b_{PE}	v_{max}	F_{max}
abdominals	16.0	0.16	8.0	8.0
diaphragm	16.0	0.16	5.0	5.0
inner intercostals	1.5	0.05	1.0	1.5
outer intercostals	1.5	0.05	1.0	1.5

aphragm remains active to behave as a load balancer. We treat this anecdotal commentary as a recipe for controlling our model. We accomplish this goal by organizing the activation of the expiratory muscles and the diaphragm into a manageable number of inputs and then generate intuitive control signals for these parameters. We combine the activation of the muscles of each: the diaphragm; the inner intercostals; and the outer intercostals as well as for each section of the abdomen. Further, we reduce the activation of the different sections of the abdomen to a single value based on the findings of [Hoit et al. 1988]. We control the activation levels of the different groupings by generating smooth sinusoidal time-varying input signals. In total we generate four input signals, one each for: 1) the abdomen; 2) the diaphragm; 3) the inner intercostals, and 4) the outer intercostals. We build these functions using a simple scripting tool and choose period, offset, and amplitude by drawing from the biomechanical descriptions for laughter summarized in Section 3 as well as our own intuition and visual aesthetic.

5.3 Spine and Clavicle Control

We animate the physically based spine and clavicle to add visual appeal while capturing the interplay of the whole system. This adds two activation signals for the clavicle (up and down) and a set of joint angles for the spine. for flexion/extension both front/back and left/right. We combine the degrees of freedom of the spine by distributing the overall desired bend along the length of the spine and employ PD-servos to account for the muscles that maintain the stability of the back. We found the spine and clavicle motion simple to control manually using the described method.

6 Audio driven control

The sound of laughter is derived from two factors, the time-varying air flow which passes over the vocal cords and the contraction of the laryngeal muscles which tighten the cords and create sound. In turn, air flow is largely dependent on the lung pressure, though other factors also play a role. Here, our goal is to develop a model of the connections between the pressure, air flow, and sound so that we can use audio to drive our laughter simulation. In [Luschei et al. 2006], Luschei and colleagues study the relationships between laryngeal activity, sound production, tracheal pressure, and the respiratory system in laughter. Their findings reveal a strong correlation between the tracheal pressure and the sound produced during laughter. Note, tracheal and lung pressure are very close in humans and can be treated as the same for our purposes (see [Finnegan et al. 1999] for more details), though we match the vocabulary here for

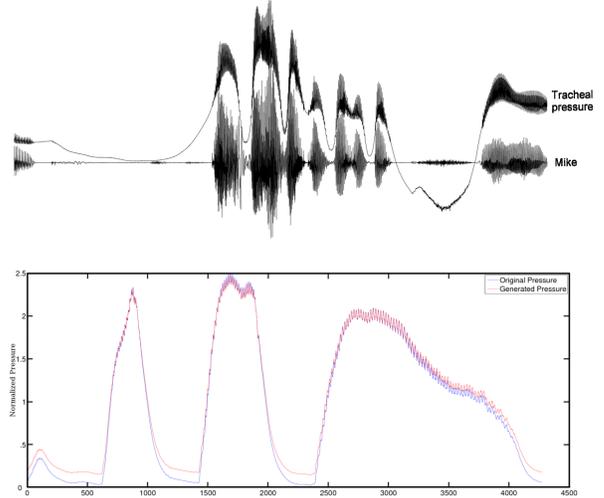


Figure 4: (Top) Reprint of figure from [Luschei et al. 2006] with the mike and pressure channels overlaid. (Bottom) Taking audio from another person in the study, we used the linear fit to generate a new pressure profile. This figure shows the similarity between our estimated pressure and the recorded pressure.

the discussion of Luschei et al.’s work. We use these findings in our pressure model, described in Section 6.1.

Furthering our investigation, we requested and received the data from Luschei’s study on laughter [Luschei et al. 2006]. We reprint (with permission) a portion of the results in Figure 4. In the top row, we can visually see a good correlation between the audio track (labelled Mike) and tracheal pressure. An aggregate correlation analysis of all trials of a single subject revealed a significant relationship between audio and tracheal pressure ($r = 0.837$, $p < 0.001$). This analysis was performed on data taken during active bouts with the pressure and absolute value of the audio filtered using a low-pass, first-order Butterworth filter. We compute a simple linear fit using the data from one subject and input new audio from a different individual in the study to generate a predicted tracheal pressure for that soundtrack. Since we know the measured tracheal pressure for that audio file, we can compare the results - they appear in the bottom row of Figure 4. Correlation analysis reveals a significant relationship between the original tracheal pressure and our generated tracheal pressure ($r = 0.924$, $p < 0.001$). Note, we did observe overall differences between individual subjects, e.g. Subject A was consistently more ‘soft-spoken’ than Subject B. To account for this difference, we normalize pressure based on average peak height for this comparison. This normalization also accounts for differences in recording levels and our assumption is that once normalized an *intensity* scale factor can be used as an intuitive animator knob to exaggerate or tone down motion derived from a given audio file.

6.1 Pressure Model

To use these findings for animation, we develop a model to approximate the pressure of the lungs. From the lung’s pressure, we can compare the simulation to the desired audio. Our pressure model consists of the interplay between three components: the lungs which are represented only by their volume and their pressure; the air flow represented as its volume flow rate; and the chest wall cavity which we treat as the volume enclosed by the ribcage and the diaphragm.

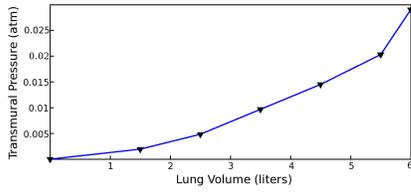


Figure 5: Linearized version of pressure/volume relationship. Samples shown taken from [Mines 1993].

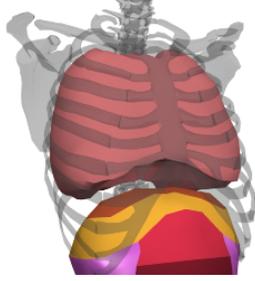


Figure 6: Chest wall cavity shown in pink.

Following established methods from basic respiratory mechanics (for example, see [West 2004]), the pressure of the lungs can be derived from

$$P_{lung} = P_{pl} + P_{tm} \quad (1)$$

where P_{pl} is the *pleural* pressure, or the pressure between the chest wall and the lung, and P_{tm} is the *transmural* pressure. P_{tm} is defined to be an all encompassing term which accounts for several physical components such as the elastic recoil of the stretched lungs. This term is well-studied in the respiratory field and a linearized approximation from Mines appears in Figure 5 for the P_{tm} as a function of lung volume, V_{lung} . Given this relationship we derive P_{lung} from running estimates of P_{pl} and V_{lung} .

From the torso simulation, we can determine the volume of the chest wall cavity, V_{cw} . A geometric enclosure for the chest wall cavity is created inside the volume formed by the ribcage and diaphragm, as in Figure 6. Each control point on the geometric enclosure is bound to the four closest neighbors (i.e., bones or vertices) of the torso simulation and are weighted equally when determining the final position. Every time step, the volume of the chest wall is calculated using a summation of elemental volumes, a similar approach is described by Zordan et al. [Zordan et al. 2004]. To update P_{pl} based on the movement of the chest wall, we use Boyle's law:

$$P_{pl,i} = \frac{V_{pl,i-1}}{V_{pl,i}} \cdot P_{pl,i-1} \quad (2)$$

where we define $V_{pl} = V_{cw} - V_{lung}$.

In order to solve Equations 1 and 2, we need an approximation of V_{lung} . Change in volume for the lungs is accounted for by two factors, the change in pressure of the lungs, and the air flowing in or out of the lungs. We can compute an approximate update due to flow as follows:

$$V_{lung,i} \approx V_{lung,i-1} + Q \cdot dt \quad (3)$$

where Q is the volume flow rate. Based on the results of [Slutsky et al. 1981] we approximate this turbulent flow as

$$Q \approx \sqrt{\frac{\Delta P}{k}} \quad (4)$$

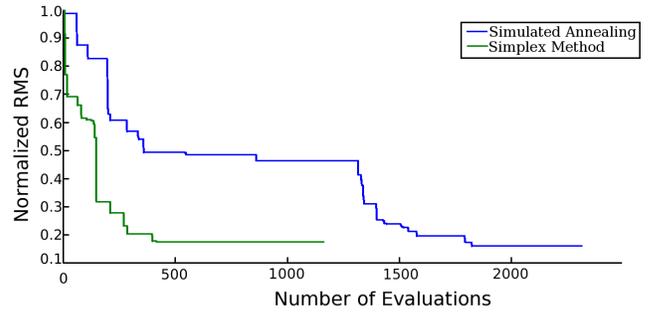


Figure 7: Error for a single call comparing simulated annealing versus simplex method.

where $\Delta P = P_{lung} - P_{atm}$ to account for atmospheric pressure and k is set to $8.2 \times 10^{-4} \text{ atm} \cdot \text{s}^2/\text{liter}^2$ to match the pressure-flow curve described by Slutsky et al.

Because of the dependency of Equation 4 on P_{lung} , we solve the system with two update loops. First, we solve for P_{pl} (Equation 2) using the updated chest volume from the simulation and the previous lung volume. Next, we solve Equation 1, using the previous value for P_{tm} to get an intermediate value of P_{lung} . In the second loop, we use the temporary value for P_{lung} to determine the flow rate Q (Equation 4) and update V_{lung} in Equation 3. Finally, we re-evaluate Equation 1 using updated values for P_{pl} and P_{tm} to determine the final value of P_{lung} .

To initialize this system, the volume of the lungs is estimated to be eighty percent of the rest volume of the chest wall. We assume that at the start of the simulation the pressure inside the lung is equal to the pressure at the mouth, therefore, P_{lung} is set to one atmosphere of pressure.

6.2 Optimization

Since pressure cannot be directly controlled in the simulation, we determine the activation levels required to achieve the desired pressure in the lungs using an optimization. Rather than optimizing for the entire laugh simultaneously, we break the complete laugh into separate calls and optimize over each call in turn. This trade-off allows us to compute a small set of parameters in exchange for a larger number of optimizations. We choose a straightforward cost function to be the root mean squared error (RMS) between the desired and simulated pressure and normalize the error by using a constant zero pressure as the maximum error.

The optimizer chooses activation for the three main muscle groups that contribute to a laugh's call: the inner intercostals; the diaphragm; and the abdomen. The outer intercostals are used for inhalation and we choose not to activate them since laughing is an expiratory behavior. Although the abdomen is broken up into five sections, as described in Section 4, we determine a single value for their collective activation and use it to compute the specific activation levels of each section. We define a basic activation primitive: a single period of a sine function with three parameters, amplitude, start time, and frequency. We use this primitive for the contraction levels of the ribcage, abdomen, and diaphragm. We link the timing parameters of the ribcage and abdomen based on the remarks made by [Filippelli et al. 2001], totalling seven optimized parameters. We experimented with other activation inputs such as a step function and a linear interpolation of control points, but found that the sinusoid produced the most visually pleasing results as well as came closest to the input pressures.



Figure 8: Deep belly laugh skinned, shown at 0.6 sec intervals.

Our system computes an optimized laugh call (0.2 sec) in about one hour. We use OPT++’s [Meza et al. 2007] parallel direct search method, which implements the algorithm of [Dennis Jr and Torczon 1991]. This algorithm uses a simplex to generate the direction of the search. One advantage of this method is that the optimization does not require expensive function evaluations to compute the gradient. Also, with parallelism, we can divide the time by the number of processors used for the optimization. We compare the effectiveness of simulated annealing versus this method. In Figure 7, we show the running best result from an exemplary single-call optimization. Simulated annealing produces a final result of 0.16 after approximately 1800 evaluations while the simplex method produces a final result of 0.18 after approximately 400 evaluations. Simulated annealing produces a quantitatively better result, but only after almost 5x more evaluation. Given the cost of the simulation evaluations and the observation that the visual quality of the laughter did not improve between these examples, we chose to use the simplex method for our optimization results.

7 Results

In the accompanying video, we include several hand-animated laughing examples as well as a set of optimized, audio-driven laugh animations. We also demonstrate the breadth of the system with a series of non-laughter behaviors. For rendering, we developed a proof-of-concept rigging tool within DreamWorks Animation’s production pipeline. Since our technique reproduces detailed sections of the human torso, the skinning process is straightforward and requires minimal effort to change DreamWork’s existing rig system. We bind the skin to the torso with the same approach described for the geometric lung enclosure outlined in Section 6.1. An example of our rendered output can be seen in Figure 8. A comparison with the torso simulated geometry appears in Figure 9.

Hand-Animated Laughter. We animate three types of laughter: a giggle, an average laugh, and a deep belly laugh. The skinned version of the deep belly laugh can be seen in Figure 8. For each, we activate the abdomen using a sine wave with a frequency of 5 Hz, following the respiratory dynamics literature. Variations in the activation levels are seen in the activation table below. The diaphragm has a constant activation throughout the laughter to create resistance against the abdomen contraction. The inner intercostals are activated in phase with the abdomen while the outer intercostals are out of phase. In addition, we include activations for secondary movement of the spine and clavicle. While some spine and clavicle motion occurs on its own based on their interplay with the abdomen and ribcage, we exaggerate this movement with additional flexion and extension of the spine and clavicle.

Optimized Laughter. Our optimizer was able to find animations that closely match the approximated pressure curves for many examples, several appear in the video. An example of the fit for a sin-

gle call appears in Figure 10. To test the strength and range of our system, we ran the audio-driven laughter optimization over many different types of laughter, some natural, some computer generated. We observed that more natural laughs produced better results than computer generated laughs. We added aggressive filtering to audio tracks that were ‘muddy’ between calls to create more distinct boundaries. This process allowed the optimizer to find more visually pleasing results. In addition, it found better solutions when we segmented the calls in the center of the valley between the peaks rather than at the beginning of the steep incline. The optimizer is unable to find solutions when competitive audio sounds appeared in the audio, for example when the recorded person was speaking while laughing. Similar to the hand-animated laughter, we added additional spine and clavicle movement to the optimized results shown in the video. For multiple bout laughs, we also manually add an inhalation between bouts.

Other Behaviors. We generate a set of hand-crafted (non-optimized) motions including breathing, coughing, and sneezing. Control inputs for these motions are included in the activation table. Breathing is simpler than the other motions, such as laughter, because the abdomen is passive. For the cough, we control the abdomen using a function that quickly activates, holds the contraction for 250 ms and quickly releases. We use a smooth wave for the inner intercostals for exhales as well as for the outer intercostals for inhales between coughs. The sneeze motion is broken up into two separate stages: the inspiration (“ah-”) and the violent expiration (“-choo”). The inspiration is created with short quick inhales that increase in intensity. The exhalation is a single severe, extended exhale.

Technique Comparison. In the video, we compare our technique to a pure data-driven animation (motion capture) and a procedural animation. For the pure data-driven motion, we recorded a male subject with 35 markers on the front of his torso. We provoked laughter by showing him a comedy skit and generated a surface based on the trajectories of the markers. For the procedural motion, we use the same geometric models of the ribcage, spine, and abdominal cavity. We apply periodic, time-varying rotations to the ribs, sternum and cartilage. For the abdominal cavity we scale each vertex using the ratios from Figure 2(d). To generate the motion for an audio track, as seen in the video, we carefully hand-select sine

	diaphragm	abdomen	intercostals
giggle	0.05	0.125	0.2
normal laugh	0.05	0.5	0.5
deep laugh	0.05	1.0	0.5
breath	0.1	-	0.3
cough	0.05	1.0	1.0
sneeze: ah -	0.1	0	0.2 (outer only)
sneeze: - choo	0.0	1.0	0.8 (inner only)

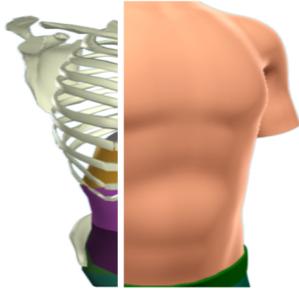


Figure 9: Comparison of rendering skin and torso simulation components. Note, detail features for the pectoral muscles, abdominal definition, and shoulders come from the artist’s skin geometry existing in the original character rig.

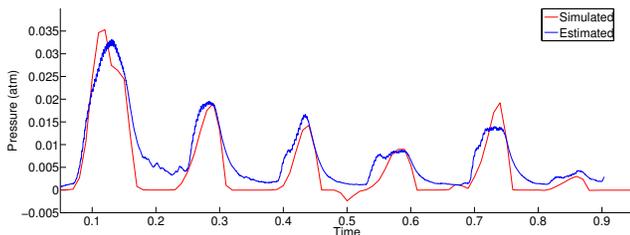


Figure 10: Comparison of pressure for generated animation (red) and input motion (blue) for a six call bout of laughter.

waves for both the ribcage and abdomen with different periods and amplitudes that match the volume and speed of the input audio. We find this process quite tedious and time-consuming and the quality of the outcome can be assessed in the video.

8 Conclusion

We have presented an innovative approach for animating laughter by controlling a physical model. Our system can be driven in two fashions: first using hand-crafted input signals and second using an audio clip as input. This animation technique is novel in several ways, one of which is in using audio to drive motion - an approach seldom used other than for lip syncing in facial motion. Also, we present a much more expressive torso than has been seen previously which is capable of a variety of behaviors. The torso is a difficult system to animate without a physical model because of the subtle interplay of rigid and deforming components. We capture many of these in our model, including effects such as a hard contraction of the abdomen pulling the spine and neck forward in a deep laugh or a sneeze. We also demonstrate that this system can be controlled simply with an audio input thereby making it such that an unskilled animator can generate motions with ease.

Our model is still limited in its capacity to model the real system. For example, we do not model the larynx. Also, Filippelli and colleagues [Filippelli et al. 2001] suggest that laughter may excite maximum expiratory flow, a phenomenon where the air flow is capped based on the volume of air in the lungs. This phenomenon may contribute to the visual effect of the motion. In addition, our system does not include a fat layer which would add realism for a wider range of body types. Finally, our conversion from audio to pressure has specific downfalls, e.g., it cannot capture silent laughter or deal well with noise.

We hope this work inspires other researchers to consider an anatomically inspired torso. The physiology of the trunk is one of the most complicated systems in the human body and it is sorely overlooked by the animation community. Likewise, laughter is an amazing behavior and we hope our efforts excite others to investigate new means for generating such motion using related or novel approaches. Finally, with our idea for driving laughing animations with audio we look forward to provoking consideration about other applications where sound can drive interesting animation.

Acknowledgments

The first author was supported by the University of California, Riverside Research Assistantship/Mentorship Program and Chancellor’s Dissertation Fellowships. The authors would like to thank Erich Luschei and Eileen Finnegan for allowing us use of their data and Kim Prisk for his help on respiratory physiology. We are thankful for the encouragement and support of Jeff Wike and Bill Ballew as well as the many fruitful conversations with experts at DreamWorks Animation. We acknowledge Vicon Motion Systems for their equipment donation. Thanks to Jose Medina for his help with graphics and rendering; Marc Soriano and Chun-Chih Wu for support with motion capture; Melissa DiLorenzo for her assistance with statistical analysis; and Eamonn Keogh for allowing us to record his laughter while everyone watched. Finally, we thank each of our paper reviewers for their time and input.

References

- ALBRECHT, I., HABER, J., AND SEIDEL, H. 2003. Construction and animation of anatomically based human hand models. In *ACM SIGGRAPH Symposium on Computer Animation*, 98–109.
- ALLEN, B., CURLLESS, B., AND POPOVIĆ, Z. 2002. Articulated body deformation from range scan data. *ACM Transactions on Graphics* 21, 3, 612–619.
- ANGUELOV, D., SRINIVASAN, P., KOLLER, D., THRUN, S., RODGERS, J., AND DAVIS, J. 2005. Scape: shape completion and animation of people. *ACM Transactions on Graphics* 24, 3 (Aug.), 408–416.
- BUCHANAN, T., LLOYD, D., MANAL, K., AND BESIÈR, T. 2004. Neuromusculoskeletal Modeling: Estimation of Muscle Forces and Joint Moments and Movements From Measurements of Neural Command. *Journal of applied biomechanics* 20, 4.
- CHEN, D. T., AND ZELTZER, D. 1992. Pump it up: computer animation of a biomechanically based model of muscle using the finite element method. *SIGGRAPH 1992* 26, 2, 89–98.
- DENNIS JR, J., AND TORCZON, V. 1991. Direct Search Methods on Parallel Machines. *SIAM Journal on Optimization* 1, 448.
- DILORENZO, P. C. 2008. *Breathing, Laughing, Sneezing, Coughing: Model and Control of an Anatomically Inspired, Physically-Based Human Torso Simulation*. PhD thesis, University of California, Riverside.
- DONG, F., CLAPWORTHY, G. J., KROKOS, M. A., AND YAO, J. 2002. An anatomy-based approach to human muscle modeling and deformation. *IEEE Transactions on Visualization and Computer Graphics* 8, 2 (Apr./June), 154–170.
- FILIPPELLI, M., PELLEGRINO, R., IANDELLI, I., MISURI, G., RODARTE, J., DURANTI, R., BRUSASCO, V., AND SCANO, G. 2001. Respiratory dynamics during laughter. *Journal of Applied Physiology* 90, 4, 1441–1446.

- FINNEGAN, E. M., LUSCHEI, E. S., AND HOFFMAN, H. T. 1999. Estimation of Alveolar Pressure During Speech Using Direct Measures of Tracheal Pressure. *Journal of Speech, Language and Hearing Research* 42, 5, 1136–1147.
- GRZESZCZUK, R., AND TERZOPOULOS, D. 1995. Automated learning of muscle-actuated locomotion through control abstraction. In *SIGGRAPH 95*, 63–70.
- HATZE, H. 1981. Myocybernetic control models of skeletal muscle: Characteristics and applications. *University of South Africa*.
- HILL, A. 1970. The force-velocity relation in shortening muscle. In *First and Last Experiments in Muscle Mechanics*, 23–41.
- HOIT, J. D., PLASSMAN, B. L., LANSING, R. W., AND HIXON, T. J. 1988. Abdominal muscle activity during speech production. *Journal of Applied Physiology* 65, 6, 2656–2664.
- KÄHLER, K., HABER, J., YAMAUCHI, H., AND SEIDEL, H.-P. 2002. Head shop: Generating animated head models with anatomical structure. In *ACM SIGGRAPH Symposium on Computer Animation*, 55–64.
- KAYE, J., METAXAS, D. N., AND JR., F. P. P. 1997. A 3d virtual environment for modeling mechanical cardiopulmonary interactions. In *CVRMed*, 389–398.
- LEE, S.-H., AND TERZOPOULOS, D. 2006. Heads up!: biomechanical modeling and neuromuscular control of the neck. *ACM Transactions on Graphics* 25, 3 (July), 1188–1198.
- LEE, Y., TERZOPOULOS, D., AND WATERS, K. 1995. Realistic modeling for facial animation. *SIGGRAPH '95*, 55–62.
- LEMOIS, R. R., ROKNE, J., BARANOSKI, G. V. G., KAWAKAMI, Y., AND KURIHARA, T. 2005. Modeling and simulating the deformation of human skeletal muscle based on anatomy and physiology. *Computer Animation and Virtual Worlds* 16, 3, 319–330.
- LUSCHEI, E., RAMIG, L., FINNEGAN, E., BAKER, K., AND SMITH, M. 2006. Patterns of Laryngeal Electromyography and the Activity of the Respiratory System During Spontaneous Laughter. *Journal of Neurophysiology* 96, 1, 442.
- MEZA, J., OLIVA, R., HOUGH, P., AND WILLIAMS, P. 2007. OPT++: An object oriented toolkit for nonlinear optimization. *ACM Transactions on Mathematical Software* 33, 2 (June), 12. Article 12, 27 pages.
- MINES, A. H. 1993. *Respiratory Physiology*. Raven Press.
- MONHEIT, G., AND BADLER, N. I. 1991. A kinematic model of the human spine and torso. *IEEE Computer Graphics & Applications* 11, 2 (Mar.), 29–38.
- NEDEL, L. P., AND THALMANN, D. 2000. Anatomic modeling of deformable human bodies. *The Visual Computer* 16, 6, 306–321.
- NG-THOW-HING, V. 2001. *Anatomically-Based Models for Physical and Geometric Reconstruction of Humans and Other Animals*. PhD thesis, University of Toronto.
- PARK, S. I., AND HODGINS, J. K. 2006. Capturing and animating skin deformation in human motion. *ACM Transactions on Graphics* 25, 3 (July), 881–889.
- PLATT, S., AND BADLER, N. 1981. Animating facial expression. *Computer Graphics*, 245–252.
- PRATSCHER, M., COLEMAN, P., LASZLO, J., AND SINGH, K. 2005. Outside-in anatomy based character rigging. In *2005 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, 329–338.
- PROMAYON, E., BACONNIER, P., AND PUECH, C. 1997. Physically-based model for simulating the human trunk respiration movements. In *Lecture Notes in Computer Science 1205*, 379–388. Springer Verlag CVRMed II-MRCAS III first joint conference.
- SAND, P., MCMILLAN, L., AND POPOVIĆ, J. 2003. Continuous capture of skin deformation. *ACM Transactions on Graphics* 22, 3, 578–586.
- SCHEEPERS, F., PARENT, R. E., CARLSON, W. E., AND MAY, S. F. 1997. Anatomy-based modeling of the human musculature. *SIGGRAPH 1997*, 163–172.
- SCHUTTE, L. 1992. *Using musculoskeletal models to explore strategies for improving performance in electrical stimulation-induced leg cycle ergometry*. PhD thesis, Stanford University.
- SEO, H., AND MAGNENAT-THALMANN, N. 2003. An automatic modeling of human bodies from sizing parameters. *Symposium on Interactive 3D graphics*, 19–26.
- SIFAKIS, E., NEVEROV, I., AND FEDKIW, R. 2005. Automatic determination of facial muscle activations from sparse motion capture marker data. *ACM Transactions on Graphics* 24, 3 (Aug.), 417–425.
- SLUTSKY, A., DRAZEN, J., O'CAIN, C., AND INGRAM, R. 1981. Alveolar pressure-airflow characteristics in humans breathing air, He-O₂, and SF₆-O₂. *Journal of Applied Physiology* 51, 4, 1033–1037.
- TERAN, J., SIFAKIS, E., BLEMKER, S. S., NG-THOW-HING, V., LAU, C., AND FEDKIW, R. 2005. Creating and simulating skeletal muscle from the visible human data set. *IEEE Transactions on Visualization and Computer Graphics* 11, 3, 317–328.
- TSANG, W., SINGH, K., AND FIUME, E. 2005. Helping hand: An anatomically accurate inverse dynamics solution for unconstrained hand motion. In *2005 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, 319–328.
- TU, X., AND TERZOPOULOS, D. 1994. Artificial fishes: Physics, locomotion, perception, behavior. In *SIGGRAPH 94*, 43–50.
- VAN DEN BOGERT, A. J., AND NIGG, B. M. 1999. Simulation. In: *Biomechanics of the Musculo-skeletal System*, 594–617.
- WATERS, K. 1987. A muscle model for animating three-dimensional facial expressions. *SIGGRAPH 1987*, 17–24.
- WEST, J. 2004. *Respiratory Physiology: the essentials*. Lippincott Williams & Wilkins.
- WILHELMS, J., AND GELDER, A. V. 1997. Anatomically based modeling. *SIGGRAPH 1997*, 173–180.
- ZAJAC, F. E. 1989. Muscle and tendon: Properties, models, scaling, and application to biomechanics and motor control. *CRC Critical reviews in biomedical engineering* 17, 359–411.
- ZORDAN, V. B., CELLY, B., CHIU, B., AND DILORENZO, P. C. 2004. Breathe easy: model and control of simulated respiration for animation. *ACM SIGGRAPH/Eurographics symposium on Computer animation*, 29–37.