

Anatomically Detailed Simulation of Human Torso

SEUNGHWAN LEE, Stanford University, USA

YIFENG JIANG, Stanford University, USA

C. KAREN LIU, Stanford University, USA




Fig. 1. (Left) Our method recovers underlying spine and shoulder bone movements unobservable by standard motion capture systems. (Right) Our torso model includes vertebrae bones, ligaments, facet joints and discs, scapula, clavicles, and a rib cage to generate physically and physiologically plausible movements.

Many existing digital human models approximate the human skeletal system using rigid bodies connected by rotational joints. While the simplification is considered acceptable for legs and arms, it significantly lacks fidelity to model rich torso movements in common activities such as dancing, Yoga, and various sports. Research from biomechanics provides more detailed modeling for parts of the torso, but their models often operate in isolation and are not fast and robust enough to support computationally heavy applications and large-scale data generation for full-body digital humans. This paper proposes a new torso model that aims to achieve high fidelity both in perception and in functionality, while being computationally feasible for simulation and optimal control tasks. We build a detailed human torso model consisting of various anatomical components, including facets, ligaments, and intervertebral discs, by coupling efficient finite-element and rigid-body simulations. Given an existing motion capture sequence without dense markers placed on the torso, our new model is able to recover the underlying torso bone movements. Our method is remarkably robust that it

can be used to automatically “retrofit” the entire Mixamo motion database of highly diverse human motions without user intervention. We also show that our model is computationally efficient for solving trajectory optimization of highly dynamic full-body movements, without relying any reference motion. Physiological validity of the model is validated against established literature.

CCS Concepts: • **Computing methodologies** → **Animation; Physical simulation.**

Additional Key Words and Phrases: Biomechanics, Human anatomy, Finite elements, Spine modeling, Trajectory optimization

ACM Reference Format:

Seunghwan Lee, Yifeng Jiang, and C. Karen Liu. 2023. Anatomically Detailed Simulation of Human Torso. *ACM Trans. Graph.* 42, 4, Article 0 (August 2023), 11 pages. <https://doi.org/10.1145/3592425>

1 INTRODUCTION

Realistic modeling and simulation of human body movements have a wide range of applications from entertainments to robotics to medicine. Conventionally, human models used in computer graphics focus on visual realism while those used in biomechanics focus on anatomical realism. With recent advances in modeling, simulation and generative models, digital humans that achieve both high fidelity in perception and in functionality, while still robust and computationally efficient, are closer to the horizon.

Human skeletal system is typically approximated as articulated rigid bodies, with hinge or ball-and-socket joints connecting the body segments. While the simplification can be acceptable for legs

Authors' addresses: Seunghwan Lee, Stanford University, 450 Serra Mall, Stanford, CA, 94305, USA, ls9021@gmail.com; Yifeng Jiang, Stanford University, 450 Serra Mall, Stanford, CA, 94305, USA, yifengj@stanford.edu; C. Karen Liu, Stanford University, 450 Serra Mall, Stanford, CA, 94305, USA, karenliu@cs.stanford.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

0730-0301/2023/8-ART0 \$15.00

<https://doi.org/10.1145/3592425>

and arms, the human torso is significantly more complex. For example, the vertebrae bones can rotate to provide spine flexibility and can also translate to absorb compressive impact along the spinal column. Similarly, the shoulder complex manifests complicated range of motion (RoM) governed by multiple non-linear inter-bone constraints. For visualizing and synthesizing upper-body-rich human activities, such as gymnastics, dancing, and baseball pitching, using only rotational joints to model the human torso often leads to laborious and ad-hoc tuning in order to achieve visual plausibility.

This paper takes a step toward building a biomechanically realistic torso model that goes beyond the chain-of-rigid-bodies structure. Instead of connected by joints, the vertebral bones of our spine model are connected and regulated by discs, ligaments and facets, where we simulate the elastic discs with finite-element method based on Projective Dynamics [Bouaziz et al. 2014], two-way coupled with the bones, and simulate the ligaments and facets as non-linear inequality constraints in an LCP (Linear Complementary Problem) framework. Similarly, the movements of shoulder complex are also subject to their own non-linear constraints. The interplay of all elements collectively determines the RoM of the skeleton, as well as its dynamic characteristics. The model parameters are systematically adjusted such that the simulated torso can match the real human's RoM.

Our torso skeleton is kinematically and dynamically accurate and easily swappable with the simplified torso of existing human models. We show that existing motion capture (moCap) data can be "retrofit" to obtain realistic and detailed torso and shoulder bone movements that match the movements of limbs in the moCap data. Because this problem is inherently under determined, the dynamic constraints imposed by our model are necessary for solving physically and anatomically plausible torso motions. As such, our method can serve as a principled tool to enrich motion capture data by adding unobserved skeletal movements. We also show that our model is computationally fast enough for optimal control with physics simulation. We optimize a sequence of control force from the torso to produce plausible motions, without using any reference data. In comparison to simplified torso, our model can generate more momentum without anatomically implausible artifacts, achieving more human-like movements.

We evaluate our model by reaffirming findings from biomechanics literature, demonstrating the biological validity of our model and its potentials to be used for biomedical analysis. The detailed analysis shows that the RoM of our spine model is similar to those reported in biomechanics literature. We also found that an emergent behavior of our torso model consistent with the well-known Scapulohumeral rhythm, a key signature movement of human shoulder complex.

2 RELATED WORK

Existing human models for character animation often depend on arbitrary parameterization and manual rigging. Artists usually design human models to have sufficient degrees of freedom (Dofs) for reconstructing captured human motions [Holden et al. 2017; Kovar et al. 2008], but anatomically realistic range of motion and physical properties are rarely top priorities. Automatic rigging methods have been proposed to accelerate the manual process by learning

rig parameters from a collection of motion data [Hong et al. 2010; Le and Deng 2014; Xu et al. 2020]. Our model can also be used as an animation rig, but we differ from existing methods in that we utilize biomechanics knowledge and physics simulation to synthesize additional details of the motion that cannot be observed by the moCap device, rather than simply fitting to the observed signals using arbitrarily parameterized rig.

Physics-based modeling of human body parts has been broadly investigated in computer graphics, such as fingers [Sachdeva et al. 2015], face [Sifakis et al. 2005], hair [Selle et al. 2008], muscles [Angles et al. 2019], eyes [Nakada et al. 2018], jaws [Daumas et al. 2005], skin [McAdams et al. 2011], and feet [Camacho et al. 2002]. For full human skeletons, however, most existing methods adopt the structure of a rigid-body tree connected by standard hinge or ball-and-socket joints, despite that rotation joints are known to be error-prone for modeling certain human parts, such as the shoulder complex [Maurel and Thalmann 2000]. The most closely related work is by Lee and Terzopoulos [2006], who built a neck model connected with vertebrae bones using ball-and-socket joints. In their work, pivot points of each joint are carefully selected following biomechanics literature [Kapandji 1974]. Lee and Terzopoulos [2008] also introduce spline joints to model complex biomechanical joints such as the scapulohumeral joints and the knee. Rotational spring-dampers are designed to mimic the effect of intervertebrae discs. The same chain-of-rigid-bodies framework is later extended to a full upper-body model in [Lee et al. 2009], combined with finite-element-based muscle and soft tissue modeling. Our model does not consider muscles. We focus on detailed and realistic modeling of skeletal connections between vertebrae bones, the rib cage, and the shoulder complex, allowing more anatomically correct range of motion.

While our work focuses on the modeling of human, control is also an inseparable component in physics-based animation [Peng et al. 2018; Witkin and Kass 1988; Yin et al. 2007]. Multiple previous studies have revealed that better modeling can lead to more effective control. For instance, more humanlike behaviors were observed when simulated muscles [Komura et al. 2000; Wang et al. 2012], learned actuators [Jiang et al. 2019], or passive dynamics of joints [Liu et al. 2005] were used. Our work also shows that a detailed spine model enables natural torso motion to emerge with little tweaking on the optimization cost function.

A rich body of work in biomechanics and medicine has been focusing on modeling human spines. Gray and Norman [1990] propose a kinematic model of the human spine with all vertebral joint movements having three axes of rotations. Panjabi and colleagues [1976] modeled the discs as bushing elements between vertebrae bones, parameterized as 6 by 6 stiffness matrices. Such bushing elements have been extended to have non-linear stiffness to better reflect the complex dynamics of intervertebrae discs, with parameters identified from in-vitro experiments [Wang et al. 2020]. Dicko and colleagues [2015] modelled discs using finite-element and ligaments as non-linear springs. However, their model was only used for static spine pose analysis, possibly due to less stable dynamics caused by springs. In general, robustness and computational efficiency are the main reasons why most existing spine models were rarely used beyond static pose equilibrium analysis [Ghezbash et al. 2018; Wang et al. 2020]. Our model adopts a hybrid FEM




Fig. 2. Our torso model include vertebrae bones, discs, ligaments, facet joints, a rib cage, and a clavicle and a scapula on each side of the shoulder complex.

plus rigid approach, utilizing an efficient finite-element simulation framework, coupled with efficient rigid-body simulation with robust impulse-based constraint handling through LCP solvers, enabling stable simulation of full-body motion even with large foot-ground impact.

Human shoulder complex has also been studied extensively in biomechanics, Seth and colleagues [2016] showed that analytic modeling of the scapulothoracic joint improved the accuracy of shoulder movement reconstruction. They used measured statistics of the scapulothoracic joint movements gathered from in-vivo experiments [McClure et al. 2001]. More realistic skeleton modeling facilitated the construction of shoulder muscle models in their later study [Seth et al. 2019]. Biomechanics researchers have also combined spine and shoulder models in biomechanics simulators, e.g. OpenSim [Delp et al. 2007], to study various activities. For example, [Bruno et al. 2015; Senteler et al. 2016] studied spine load and reaction forces during primitive motions such as bending and lifting with weight, and [Cazzola et al. 2017] studied shoulder injury of rugby players.

3 TORSO MODELING

We present a physics-based model for human torso that simulates detailed human anatomy. Our torso model consists of bones, joints, discs, ligaments, and facets. We include 23 vertebrae bones (from sacrum to C3) assembling a spinal column, a rib cage, and a clavicle and a scapula on each side of shoulder complex (See Figure 2). 20 ligaments and 44 facets are attached to the spine, and 23 vertebrae discs are sandwiched between adjacent vertebrae bones. There are 28 additional costovertebral and costotransverse facet joints that connect the rib cage to the thoracic bones. The collaboration of the




Fig. 3. We simulate intervertebral discs as deformable materials using finite element method.

ligaments, the discs and the joints defines the passive dynamics of our model as well as kinematic characteristics such as the range of motion. Muscles are not included in our current model.

All the bones are modeled as rigid bodies with detailed geometries used by anatomical modeling [Zygot [n. d.]]. The rib cage is assumed to be a single rigid body as the movements of the ribs are mostly associated with breathing. We use ball-and-socket joints to connect the rib cage and the clavicle, and the clavicle and the scapula. Between the vertebrae bones and between the rib cage and the thoracic vertebrae, we connect them with placeholder 6-DoF free joints, and the relative movements are instead constrained by finite-element simulation of the discs and non-linear constraints rising from ligaments and facets. The use of 6-DoF joints allows us to still treat the torso as an articulated rigid-body tree with q being the state of the rigid-body in generalized coordinates, facilitating incorporating existing frameworks for physics simulation. Two-way weak coupling between the finite-element discs and the rigid-body skeleton is discussed in Section 3.1 and 3.5.

3.1 Intervertebral Disc

Intervertebral discs are deformable materials attached in-between adjacent vertebrae bones (See Figure 3). The elasticity of the discs prevents the bones from colliding and rubbing against each other, and absorbs and dissipates large compressive impacts on the spinal column. Such a deformable body is often simulated using Finite Element Method (FEM) which discretizes a body into a finite number of elements such as tetrahedrons. The states of the disc at time t can be represented by its nodal positions $x_t \in \mathbb{R}^{3k}$, where k is the number of discretized vertices.

Following previous work, we adopt the standard implicit Euler integration in its variational form [Gast et al. 2015; Martin et al. 2011], where at each time step t , we solve a non-linear optimization problem for the current state x_t :

$$x_t = \operatorname{argmin} \frac{1}{2h^2} \left\| M^{\frac{1}{2}} (x_t - x_{t-1} - hv_{t-1}) \right\|_F^2 + \sum_i k_i E_i(x_t). \quad (1)$$

Here M is the mass matrix and h is timestep size. v_t is obtained through backward finite differencing of x_t . By minimizing Equation 1, we essentially solve for x_t that balances the change of momentum and the forces expressed as the gradients of the potential energies E_i ,

with scalar weights k_i . For our case, E_i include isotropic elasticity [Chao et al. 2010] and volume preservation [Bouaziz et al. 2014] of the disc.

The minimization problem can be solved by modified Newton's Method, but it is usually computationally expensive. This is because Hessian $\nabla^2 E_i(\mathbf{x})$ obtained through linearization changes at each time step, preventing us from utilizing matrix factorization to accelerate. Facing this issue, Projective Dynamics [Bouaziz et al. 2014] showed that certain types of energy functions can be solved much more efficiently in an alternating "local-global" manner. Following their framework, we define each E_i as a minimization process:

$$E_i(\mathbf{x}) = \min_{\mathbf{m}_i \in C_i} \frac{1}{2} \|\mathbf{A}_i \mathbf{x} - \mathbf{m}_i\|^2, \quad (2)$$

where the constraint manifold C_i and coefficient matrix \mathbf{A}_i depend on the specific energy term. Equation 2 solves for \mathbf{m}_i , the projection of $\mathbf{A}_i \mathbf{x}$ onto C_i . The projection error is then defined as the energy E_i . Since each E_i is independently computed without considering other E_i 's, we call this step a "local" solve. A large class of FEM energies can be written into this projection form. For instance, we can cast the isotropic elasticity term as the distance from the current deformation gradient to its closest projection on the $SO(3)$ manifold (isometry rotation with zero deformation), such that a rigid rotation of the rest-pose tetrahedron will lead to zero elastic energy.

Once all \mathbf{m}_i 's are obtained through local solves, we proceed with the standard Newton steps ("global solve"). We plug in the solved \mathbf{m}_i in Equation 2 and express E_i as a quadratic function in \mathbf{x} , where the constant \mathbf{A}_i enables acceleration using Cholesky decomposition. In practice, since the discs are tightly attached to the bones with negligible inertial effects comparing to the bones, we adopt a quasi-static simplification, effectively forgoing the momentum term ($h \rightarrow \infty$) to improve simulation stability, similar to previous works [Lee et al. 2018b]. As such, Equation 1 simplifies to: $\mathbf{L}\mathbf{x}_t = \mathbf{d}$, where $\mathbf{L} = \sum_i k_i \mathbf{A}_i^T \mathbf{A}_i$ and $\mathbf{d} = \sum_i \mathbf{A}_i^T \mathbf{m}_i$.

An additional energy term E_{pos} is needed to couple the vertebrae bones and the discs. We compute a set of attached points between the bone surface and the disc surface at the rest pose and enforce the attachment via minimizing $E_{\text{pos}} = \|\mathbf{x}_B - \mathbf{P}(\mathbf{q})\|^2$, where $\mathbf{P}(\mathbf{q})$ evaluates attach points on the bone surface in the world frame given the current state \mathbf{q} , and \mathbf{x}_B are the boundary disc vertices in the world frame. Note that we use the bone configuration \mathbf{q} from the previous time step, resulting in a weak coupling between the bones and the discs. $\mathbf{f}_{\text{pos}} := -\nabla E_{\text{pos}}$ is therefore the coupling force between finite-elements and rigid bodies (See Figure 3).

3.2 Ligament

A ligament is a spring-like fibrous connective tissue attached between bones. It regulates the relative movements among the bones, keeping them stable. The anatomical functionality of each ligament varies by attachment sites (termed *origin* and *insertion*). For example, the intertransverse ligament is routed on the transverse process of the vertebrae bone, and it mainly limits the opposite-side lateral flexion of the spine (See Figure 4). We use piece-wise linear approximations to represent the ligament and its geometry. Starting from the origin of a ligament, we divide the ligament into multiple line




Fig. 4. Ligaments (red) and facet joints (blue). \mathbf{p}^+ and \mathbf{p}^- being waypoints on the ligament's routing from its origin on sacrum to insertion on neck. The location on cartilage A where the facet resides is termed "inferior articular process" and the counterpart on cartilage B is termed "superior articular process".

segments, where each segment is expressed by \mathbf{p}^+ and \mathbf{p}^- defined in the local coordinate of the associated bones respectively.

Each ligament is represented by a unilateral spring; it generates contractile forces when the ligament is extended and zero force when relaxed, a characteristic described by the Hill's muscle model [Zajac 1989]. According to the model, the force-length curve of ligaments has two phases: tendon elongations below the thresholds create negligible forces, while tendon elongations above a threshold create exponentially growing forces. In practice, it is desirable to avoid a large amount of contractile forces as it decreases simulation stability. Hence we add inequality constraints (to be solved by a stable LCP time-stepping scheme) in our simulation where they preemptively bound the length of the ligament within a moderate range:

$$C_{\text{ligament}}(\mathbf{q}) = \|\mathbf{p}^+(\mathbf{q}) - \mathbf{p}^-(\mathbf{q})\|^2 \leq l_{\text{max}}, \quad (3)$$

where l_{max} is the maximum length that varies by each ligament (see Section 3.6). Comparing with the Hill's model, this modification to inequality constraints (zero force below threshold and infinite-as-needed beyond threshold) doesn't change the biomechanical property much. Note that one ligament will create multiple $C_{\text{ligament}}(\mathbf{q})$ constraints with different l_{max} for each segment. Section 3.6 will describe how the ligament parameters are identified systematically.

3.3 Facet

The facets are located on the posterior side of the vertebrae bones, as well as between the rib cage and thoracic vertebrae bones. They are synovial joints consisting of two cartilages facing each other, wrapped up by the joint capsules. One cartilage (A) slides on the surface of the other cartilage (B) with near-frictionless movement due to highly smooth surface and the existence of fluids (See Figure 4). Facet forms a closed-chain kinematics with the disc which increases the stability and robustness of the spinal column against excessive external perturbations. The surface geometry of the cartilage B determines the movements of the facet [Drake et al. 2009; Williams et al. 2010]. As the cartilage is flat and oval-shaped, we represent the surface geometry of cartilage B as a flat ellipsoid:

$$C_{\text{facet}}(\mathbf{p}) = \|\mathbf{S}^{-1} \mathbf{R}^T (\mathbf{p} - \mathbf{p}_{\text{center}})\|^2 - 1, \quad (4)$$




Fig. 5. The anatomy of shoulder joint. We use ball-and-socket joints for the sternoclavicular joint, acromioclavicular joint, and glenohumeral joint (green). We utilize facet constraints to simulate the scapulothoracic joint where the ellipsoid (blue) defines the surface of the rib cage whereas points attached to the scapula (red) are moving on the ellipsoid.

where $S \in \mathbb{R}^{3 \times 3}$ is the diagonal matrix representing each axis scale, $R \in \text{SO}(3)$ is the rotation matrix where each column denotes the axis of the ellipsoid, and $\mathbf{p}_{\text{center}}$ is the center of the ellipsoid. Here, the level-set $C_{\text{facet}}(\mathbf{p}) = 0$ represents the surface where we would like points on cartilage A to stay on. In other words, we need to select three non-colinear points on A that are also on $C_{\text{facet}}(\mathbf{p}) = 0$ at the rest configuration, and the constraint solver will then keep these three "wheels" on B's ellipsoid surface during the entire simulation. In practice, we select four points instead of three to ensure numerical robustness, and relax $C_{\text{facet}}(\mathbf{p}) = 0$ to $|C_{\text{facet}}(\mathbf{p})| \leq 0.2$.

3.4 Shoulder

We also utilize the ellipsoid constraint used for facets to simulate the shoulder movements (See Figure 5). The clavicles and the scapula together determine the location and the orientation of the shoulder. The joint of the clavicle is on the sternum (sternoclavicle joint) and can be modeled by a rotational joint. The movement of the scapula is much more complex and determined by both the acromioclavicle joint and the scapulothoracic joint. The scapulothoracic joint represents the sliding of the scapula on the posterior of the rib cage. It has been revealed in many studies (e.g. [Seth et al. 2016]) that the rib cage can be approximated by an ellipsoid on which the scapula slide on. This movement is identical to the movement of the facet joint so we add constraints between the rib cage and the scapula to mimic the scapulothoracic joint anatomy. Unlike facets joint, however, the ellipsoid constraints for the scapula will create a close-chain kinematics.

3.5 Torso Dynamics

With finite-element simulation presented in Section 3.1, the equations of motions of our rigid skeleton model can be described by the Lagrangian dynamics, where we solve for $\ddot{\mathbf{q}}$ to drive the simulation

forward:

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{c}(\mathbf{q}, \dot{\mathbf{q}}) = \boldsymbol{\tau}_{\text{disc}}(\mathbf{x}_{\mathcal{B}}, \mathbf{P}(\mathbf{q})) + \boldsymbol{\tau}_{\text{ext}} \quad \text{subject to} \quad (5)$$

$$C_{\text{ligament}}(\mathbf{q}) \leq \mathbf{l}_{\text{max}} \quad (6)$$

$$|C_{\text{facet}}(\mathbf{p}(\mathbf{q}))| \leq 0.2, \quad (7)$$

where \mathbf{q} is the generalized positions, $\mathbf{M}(\mathbf{q})$ is the mass matrix, $\mathbf{c}(\mathbf{q}, \dot{\mathbf{q}})$ is the coriolis and gravitational forces, $\boldsymbol{\tau}_{\text{ext}}$ is the external forces, and $\boldsymbol{\tau}_{\text{disc}}(\mathbf{x}_{\mathcal{B}}, \mathbf{P}(\mathbf{q})) = -\mathbf{J}^T \mathbf{f}_{\text{pos}}$ is the sum of forces generated by the disc attached to the bones via positional constraints $E_{\text{pos}}(\mathbf{x}_{\mathcal{B}}, \mathbf{P}(\mathbf{q}))$. \mathbf{J} is the Jacobian matrix that maps the generalized coordinate to the attachment (bone surface) points. Equation 6 and 7 are the ligament constraints and facet joint constraints, respectively.

Algorithm 1 shows the process of forward dynamics of our model for each time step t , where we sequentially solve the FEM simulation and the articulated rigid body simulation. For FEM, we iterate between solving the local projection for each constraint (Equation 2) and linear solves for the Newton steps. Featherstone's algorithm [Featherstone 2014] is used to solve Equation 5 while the constraints 6 and 7 are solved via impulse-space simulation which yields a Linear Complementary Problem (LCP) [Anitescu and Potra 1997]. We use Danzig's LCP solver [Cottle et al. 2009] to compute the lagrangian multipliers of the constraints which are eventually used to modify the generalized velocities $\dot{\mathbf{q}}_{t+1}$ such that they will not violate the constraints. Finally, the solved \mathbf{q}_{t+1} will be used for the FEM simulation at next time step.

ALGORITHM 1: Torso Forward Dynamics

```

Input: Disc Positions  $\mathbf{x}_t$ , Joint Positions, Joint Velocities
 $\mathbf{q}_t, \dot{\mathbf{q}}_t$ , Bone Surface Locations  $\mathbf{P}(\mathbf{q}_t)$   $\mathbf{L} = \sum_i k_i \mathbf{A}_i^T \mathbf{A}_i$ 
for  $l = 0, 1, \dots, \text{numFEMIterations}$  do
     $\mathbf{d} = 0$ 
    for  $i = 0, 1, \dots, \text{numFEMEnergies}$  do
         $\mathbf{m}_i = \text{project}(\mathbf{A}_i \mathbf{x}_t, C_i)$ 
         $\mathbf{d} = \mathbf{d} + \mathbf{A}_i^T \mathbf{m}_i$ 
    end
     $\mathbf{x}_t = \mathbf{L}^{-1} \mathbf{d}$ 
end
 $\boldsymbol{\tau}_{\text{disc}} = -\mathbf{J}^T \mathbf{f}_{\text{pos}}(\mathbf{x}_t, \mathbf{P}(\mathbf{q}_t))$ 
 $\dot{\mathbf{q}}_{t+1} = \text{ForwardDynamics}(\boldsymbol{\tau}_{\text{ext}}, \boldsymbol{\tau}_{\text{disc}}, \mathbf{q}_t, \dot{\mathbf{q}}_t)$ 
 $\dot{\mathbf{q}}_{t+1} = \text{integrateVelocities}(\dot{\mathbf{q}}_{t+1}, \dot{\mathbf{q}}_t)$ 
 $\mathbf{q}_{t+1} =$ 
    ResolveConstraintsByLCP( $\dot{\mathbf{q}}_{t+1}, C_{\text{ligament}}(\mathbf{q}_t), C_{\text{facet}}(\mathbf{q}_t)$ )
 $\mathbf{q}_{t+1} = \text{integratePositions}(\dot{\mathbf{q}}_{t+1}, \mathbf{q}_t)$ 
 $\mathbf{x}_{t+1} = \mathbf{x}_t$ 

```

3.6 Parameters Identification

Ligaments. The maximum lengths l_{max} of the the ligaments are determined from biomechanics data of spine range of motion. For each adjacent pairs of vertebrae bones, existing literature [Šavlovskis and Raits [n.d.]] measured a few extreme relative orientations (shown as cross symbols in Figure 6) that an average subject is able to achieve. As shown in Figure 6, we gradually shorten l_{max} until it is no longer enough to realize all documented extreme poses. At each shortening step, this check is done by running the simulation for the bone pair with random actions. After all ligaments




Fig. 6. The maximum ligament lengths are adjusted sequentially to match the human's range of motion. The cross symbols represented the extreme relative orientations between two vertebrae bones known to biomechanics literature. Ligament 1 is shortened until one of extreme poses is no longer feasible, followed by the repeated operation for Ligament 2. Axial rotation DoF not shown here for simplicity.

are "tightened", we will obtain a non-linear range of motion (the yellow and red lines) that is strictly more flexible than the simplex (green lines) of measured extreme poses. In practice, some of the ligaments only affects one degree of freedom, and we shorten these independent ligaments first. Though different orders of shortening can lead to slightly different non-linear shapes of the RoM, they are equally valid in the sense of matching all available measurement data.

Facets. The location $\mathbf{p}_{\text{center}}$, orientation \mathbf{R} and dimension \mathbf{S} of each facet ellipsoid are determined by annotating the bone geometries. For each pair of the adjacent bones, we visually inspect and manually annotate the positions of *superior articular process* and *inferior articular process* with knowledge of anatomy. We determine the parameters such that the ellipsoid contains all the vertices on the *superior articular process*. We randomly select four non-colinear points \mathbf{p} among the vertices on the *inferior articular process*. Note that these four points are irrelevant to the shape of the facet ellipsoid and thus do not affect RoM; they are only used to enable adjacent bones to slide with each other. Therefore, randomly choosing four non-colinear points is sufficient to guarantee sliding.

Discs. FEM parameters such as the Young's modulus are tuned manually for stable and visually plausible results. We leave matching disc properties to the non-linear stiffness observed in measurement data [Wang et al. 2020] to future work.

4 EXPERIMENTAL RESULTS

We evaluate our model based on three aspects. *Realism*: Can our model generate detailed torso motions that are visually realistic and consistent with the realistic human range of motion? *Robustness*: Can we use our model as an automatic tool to robustly process large-scale motion datasets? *Efficiency*: Is our model computationally efficient enough for physics simulation and optimal control applications? We describe a few experiments to validate these aspects below and provide a supplemental video for further evaluation.

4.1 Experiment setup

Our model includes 24 place-holder free joints and 17 ball-and-socket joints, totaling 195 DoFs. We model an adult human with a

height of 170cm and a mass of 72kg. The mass of a bone is computed by the bone density and the bounding box of the mesh while the mass of some bones with concave shapes such as ribs, scapula, and clavicle, is manually assigned. We apply joint limits to the sternoclavicular joint, the acromioclavicular joint, and the glenohumeral joint [Magee 2013] while the other joints have no limits. We also use passive spring-damper for all joints which increase simulation stability as well as the robustness of retrofitting. We use a PD controller for each DoF of our torso model.

Our model consists of 20 ligaments and 74 facets, including 44 facet joints, 28 costovertebral and costotransverse joints, and 2 scapulothoracic joints. Each ligament is modelled by a sequence of line segments, each of which corresponds to each individual constraint equation (6). This yields 440 tendon constraints on the spinal column. We also sample four non-colinear points for each facet. As a result, our rigid-body world is simulated by 680 constraints. Our rigid-body simulation is written in C++ building upon the open-source library DART [Lee et al. 2018a]. We directly use the LCP solver in the DART engine for body-ground contact. The simulation runs at 600Hz for retrofitting and 1800Hz for trajectory optimizations. Simulation runs 8 to 40 times slower than real-time depending on configurations. We tetrahedralize the disc meshes using TetWild [Hu et al. 2018]. Each disc consists of 127 to 285 nodal vertices and 307 to 1090 tetrahedrons depending on its shape.

4.2 Motion Retrofitting

Our method provides a tool to estimate detailed movements of torso bones unobserved by standard moCap devices. We apply our tool to a large existing motion capture dataset, Mixamo, without user intervention to demonstrate its robustness. We will host a large-scale public dataset of motion capture data with enriched torso movements.

Our tool takes as input a target motion on any rigged skeleton and first retarget it to a temporary skeleton, followed by simulating the detailed torso movements on our detailed skeleton. In the retargeting step, we utilize the IK-based retargeting algorithm provided by *MotionBuilder*TM to match the limb and head movements of the target motion on our temporary skeleton. The limbs and head of the temporary skeleton are identical to those of the detailed skeleton, but the torso is simplified to three ball joints, similar to many existing human rigs used in computer animation. The temporary skeleton has shoulders that include sternoclavicular and acromioclavicular joints, as it needs to allow movements such as shoulder shrugging and straightening. We set a wider range of motion for the temporary torso so that the arms and the head can match the target motion precisely.

For each frame, using the current positions of the temporary skeleton, we set the positions of humerus, femur, and head bones as positional constraints and run our torso dynamics (Section 3.5) until convergence to a static equilibrium. The dynamic simulation is necessary because there are far more degrees of freedom on the torso than the constraints. In our experiments, we simulate the torso dynamics for 0.5s for most examples, while for some motions containing excessive bending we simulate the dynamics for 1s.

We retrofit 1293 motion clips from *Mixamo* dataset including assorted dance moves, martial arts, baseball, jumping, walking, running etc. The dataset contains motion clips ranging from 1 to 44 seconds long, totaling an hour and 18 minutes of high quality motion data. Motion retrofitting takes 20 hours with AMD Ryzen 9 5950X. We accelerate the process with 16 parallel cores. A few representative motions can be found in the supplemental video.

4.3 Trajectory Optimization

We evaluate our torso model by solving trajectory optimization problems to simulate highly dynamic motions, i.e. jumping and swinging. The optimizer solves for a control trajectory to optimize the performance of a physically simulated motion sequence. Since the main goal of this experiment is to validate that the large number of DoFs does not prohibit physics simulation or optimal control in terms of efficiency, and motion naturalness is secondary, we use a very simple cost function, such as jumping as high as possible, with little regularization or tuning and without any use of human motion data.

We adopt a standard stochastic trajectory optimization framework MPPI [Williams et al. 2017], a shooting-based gradient-free algorithm for solving open-loop control sequences. Given a constant initial state $\mathbf{q}_0, \dot{\mathbf{q}}_0$, and a transition function defined by a physics simulator $\mathbf{q}_t = F(\mathbf{q}_{t-1}, \mathbf{a}_{t-1})$ with any action space \mathbf{a} , MPPI solves for $(\mathbf{a}_0, \dots, \mathbf{a}_{N-1})$ such that the total optimal control cost $\sum_{t=0}^N c(\mathbf{q}_t, \dot{\mathbf{q}}_t)$ of the simulated trajectory $(\mathbf{q}_0, \mathbf{q}_1, \dots, \mathbf{q}_N)$ is minimized. We additionally compare motions solved by our model to those from a fixed-torso model. Optimization takes around 30 MPPI iterations for both our model and the fixed-torso model, though our model has a much larger action space. For our model MPPI optimizations take around 6 hours each with 6 CPU cores, and the rigid-torso model is around 60x faster.

4.3.1 Falling Test. Before running optimization, we make sure that when combining our model with an existing full-body skeleton, the combined model will be robust enough to interact with the environment and execute random control sequences commanded by MPPI during optimization. For this we drop our character from 2 meter high to the floor with small random controls commanded. As shown in the supplemental video, our model can handle discontinuous, large contact impact without issues. Interestingly, even without any optimization, our model already generates a more humanlike falling motion upon touching the ground, comparing with the fixed-torso model.

4.3.2 Swing. In this task, the character attempts to swing on a horizontal bar from zero velocity. The objective is to maximize the product of horizontal and vertical speeds at the final frame of the trajectory, such that the character would travel the longest distance after releasing the bar:

$$-\dot{q}_{y,N} \cdot \dot{q}_{z,N} + \sum_{t=0}^N 0.1 q_{x,t}^2, \quad (8)$$

where $q_{x,t}, q_{y,t}, q_{z,t}$ are respectively the x, y, z components of the center of mass of the character at time step t . The second term penalizes deviation from the sagittal plane:




Fig. 7. The ranges of motion of the vertebrae bones. (red) Real human RoM of each bone. (blue) RoM measured by our model.

The action space at each control step is the change of PD targets of relative bone orientations. While this action space is largely redundant compared with the real human musculoskeletal system which uses about 30 muscles to actuate the spine, we did not find MPPI having practical difficulty converging to a solution, nor did we find unnatural spine poses exhibited in the solution (except the task-irrelevant neck bones tends to be moderately jittery). We hypothesize that the coupling of discs, ligaments and facets in our spine model effectively restrict the search space of MPPI to a biomechanically realistic state-action space. On the other hand, the fixed spine character which uses only torso DoFs to swing demonstrates unsmooth and suboptimal behaviors, potentially due to its inability to utilize the spine to store and release potential energy.

4.3.3 Jumping. The cost function tasks the character to jump as high as possible, with a small regularization term discouraging deviation from the vertical (y) direction:

$$-\text{sign}(\dot{q}_{y,N}) \frac{\dot{q}_{y,N}^2}{2g} - q_{y,N} + \sum_{t=0}^N 0.001(q_{x,t}^2 + q_{z,t}^2). \quad (9)$$

For the jumping task, the action space for both our model and the fixed-torso model additionally include the knee and ankle DoFs. The difference between motions solved with two models are more subtle for this task, as excepted since spine mainly performs a balancing and buffering role in jumping. Nevertheless, the motion simulated by our torso model exhibits a propagation of acceleration due to the ground reaction force from the lower limbs through the spine, while the motion simulated by the fixed-torso model only uses ankles to exert large torque, resulting in more robot-like motion.

4.4 Comparisons with Real Human Data

4.4.1 Range of Motion. An effective validation of our model is to compare the range of motion (RoM) of each vertebrae bone in the simulated motion against real human RoM recorded in biomechanics literature. Because we do not explicitly “set” joint limits for our torso model, the resulting RoM is an emergent quantity from the complex interplay of physical and geometry properties of discs, ligaments, and facet joints.




Fig. 8. Scapulohumeral rhythm. A unique phenomenon observed when humans rotate the upper arm laterally. The scapula elevation and humerus abduction maintains a fixed ratio after threshold configuration is reached. Our model (movement trajectory in black curve) is able to reproduce this phenomenon (shown in red curve).

We compare the range of motion of each bone with established biomechanics literature [Anderst et al. 2015; Bhalla and Simmons 1969; Fujii et al. 2007]. The RoM of existing real human data is highly variant by subject's age, sex, and physical capability. We compare the RoM of our model with the average of existing data shown in red line in Figure 7). For each vertebrae bone, we measure the maximum and the minimum angle along anatomical axes (Flexion-Extension, Axial rotation and Lateral bending) to identify the RoM. The ranges of axial rotation and lateral bending are divided by half as they are symmetric. The RoM of our model is measured by the minimum and maximum angles (5% outliers are removed) among all retrofitted sequences of the Mixamo dataset.

We observe that most vertebrae bones have similar RoMs compared to the biomechanics literature. There are still a few noticeable differences. For flexion-extension, the real human has a wider RoM on L1 and L2 while our model has a wider RoM on upper thoracic region. We suspect that the maximum length of ligaments we identified might be larger than the real human statistics due to our overly conservative algorithm. As a result, bending upper thoracic region is sufficient to match the positional constraints. We also observe a peak on C7 especially in axial rotation. We suspect that the locations and the shapes of the costovertebral joints and the costotransverse joints overly constrain the thoracic bones, and thus C7 has to be more flexible to compensate the narrowed RoM on the thoracic bones.

4.4.2 Scapulohumeral Rhythm. Scapulohumeral rhythm is one of the signature features observed from human shoulder movements [Rockwood 2009]. It describes the relation between scapula elevation and humerus abduction. When a healthy subject laterally rotates the upper arm on the coronal plane, the glenohumeral joint (GH) and the scapulothoracic joint (ST) both contribute to the rotation (See Figure 8). The lateral rotation begins with the first phase called setting phase where the rotation is almost purely done by the GH joint while the ST joint remains relatively unchanged. If the rotation continues beyond a threshold (around 30 degrees), both GH and




Fig. 9. Pressure Load on L5-L4 Disc.

ST joints start to rotate at a 1-to-2 ratio. This unique phenomenon emerges in our simulated motion from using sliding facet constraints between the scapula and the rib cage.

4.4.3 Pressure Load on Disc. Existing biomechanics literature primarily focuses on spine and shoulder geometry and kinematics, such as lordosis (spine curvature), RoM, and Scapulohumeral Rhythm, while studies examining spine dynamics are less common due to the challenges of conducting in-vivo experiments. Among studies that measure dynamical properties, researchers often perform quasi-static analyses, such as assessing pressure loads on intervertebral discs. In [Wilke et al. 1999], the authors implanted a transducer with a pressure sensor into the L4-L5 disc, measuring loads during everyday tasks like standing, sitting, and holding objects. One experiment revealed a 0.5 MPa pressure during standing, which increased to 1.1 MPa when holding a 20 kg box and further to 1.8 MPa when holding the box 60 cm from the chest (See Figure 9).

To benchmark our model against their data, we compute pressure load by calculating elastic forces at the disc surface and dividing by the area attached to the bone (14.4 cm²). Using a PD controller, we emulate a pose holding a box with feet welded to the ground, allowing the character to hold the box without learning or optimizing motion control. Our experiment demonstrates similar pressure for the standing task and an increase when holding a box, albeit less than the clinical data, possibly due to the PD controller's high gain. Holding a box should cause spinal bending and both rotational and translational deformation of discs, but the high gain results in negligible rotational disc deformation. Translational DoFs between bones are not controlled by PD (to prevent the spine from moving freely with magic forces) and are solely influenced by the translational disc deformation, which correctly increases during box carrying.

4.5 Ablations

In this section, we conduct ablation studies to demonstrate the roles of intervertebral discs and ligaments in our torso simulation model.

4.5.1 Disc. The intervertebral discs play a critical role in preventing vertebral discs from colliding and rubbing against each other. When we remove the discs during motion retrofitting, we observe a significant increase in bone-to-bone collisions. To quantify the impact of the discs, we measure bone-to-bone displacements relative to




Fig. 10. Bone-to-bone displacements and collision incidence rate during motion retrofitting.




Fig. 11. (left) We visualize distributions of flexion-extension angle of L5 during retrofitting the motions in the Mixamo dataset. (right) We accumulate the flexion-extension angles among lumbar bones during 65° flexion.

rest-pose displacements and count the frames with collisions during retro-fitting, with and without the discs (See Figure 10). We approximate each vertebral bone as a cylinder to prevent the posterior side of bone (linked with the facets) from dominantly affecting distance measurements. We use mesh-to-cylinder distances to measure the displacements and compute the average and standard deviation (error bars) of the displacements across all motion frames in the dataset.

We observe that without the discs, the displacements at the lumbar bones are below 0 cm due to gravity, meaning that the distances between the bones are narrow. As a result, the bones collide with each other frequently at the incidence rate of 8.0% and 14.1% for L5 and L4 respectively during motion retrofitting. We also observe high collision rates among the upper thoracic bones, which are closer to each other at rest comparing to bones in other sections of the spine. With discs modeled (red in figure), collision rates are close to zero.

4.5.2 Ligament. In addition, we also study the impact of ligaments. Removing the ligaments results in a unrealistically wider range of motion (RoM) than that of a real human spine (See Figure 11). Without the ligaments, the L5 exhibits a large flexion-extension angle ranging from -30° to 60° . When retrofitting the spine to the existing motions, our method tends to find extreme solution where only L5 rotates to fit the motion while other vertebral bones remain in the rest pose (See Figure 11 Right). This experiment shows that the ligaments with identified parameters provide accurate RoM, significantly improving our simulation model in producing more human-like results.

5 DISCUSSION

This paper introduces a novel human torso model featuring key anatomical elements from the spine and the shoulder complex. We show that our model is robust enough to retrofit a large and diverse motion database, and to be readily incorporated into existing full-body skeletons to simulate highly dynamic movements. We further partially validate that the statistics of our model is consistent with the human data from biomechanics literature. While there has been some attempts from the biomechanics community to build similarly detailed spine or shoulder models, we demonstrate in this work that tools, algorithms, and implementation know-hows that continue to mature in Computer Graphics can make anatomically realistic human models robust and efficient for computationally-heavy applications, accelerating wide adoption of these models.

Our current implementation has a few limitations. First, we have yet to model muscles on the torso which drive the spine and the shoulder. Our detailed skeletal model should help the development of an accurate torso muscle model. Second, passive dynamics of the joints and the discs, such as spring-damping coefficients and Poisson ratio, are tuned manually. While measurement of dynamic properties of human skeleton is much less available compared with kinematics, our future work would include matching passive dynamics of our model to experiment data. Third, our current retrofitting is solved frame-by-frame and can lack temporal smoothness in some cases. Optimizing over a time window, or better initialization scheme would improve the quality of retrofitting.

Perhaps the most exciting application domain enabled by our model lies in medicine and healthcare. The proposed computer graphics techniques have the potential to contribute to the development of noninvasive diagnostic tools for analyzing and predicting internal spinal conditions during daily activities. For instance, back pain, a prevalent medical issue affecting 80% of individuals at some point in their lives, is often caused by slipped (herniated) discs that impinge on the spinal nerve cord. The depth of this impingement affects the severity of pain and varies with different poses. Our model could simulate such protrusions and anticipate how various full-body poses during a gait cycle influence the depth of impingement. Furthermore, incorrect squatting posture during strength training, leading to an improperly bent spine, can increase pressure on the spinal column. Our method may be capable of reproducing these phenomena and predicting the optimal squat position to minimize spinal load. Considering the wide variability in torso RoM due to factors like age, sex, and physical ability, personalizing our model for medical applications presents another promising future direction.

Our framework also holds promise for enhancing computer animation applications. Existing skinning techniques such as linear blend skinning and dual quaternion skinning can be readily integrate to our model. In addition, by accounting for detailed bone meshes, our spine model can enable novel skinning techniques to visualize the outlines of scapular and vertebral protrusions commonly seen in actual human skin. Our model can thus serve as the foundation for a novel skinning pipeline. Additionally, adapting our spine model to quadruped animals, such as cats, may lead to improvements in the realism of animated quadruped motion.

ACKNOWLEDGMENTS

This work was supported by Wu Tsai Human Performance Alliance at Stanford University, the Stanford Institute for Human-Centered AI (HAI), and NSF:FRR 2153854.

REFERENCES

- William J Anderst, William F Donaldson III, Joon Y Lee, and James D Kang. 2015. Three-dimensional intervertebral kinematics in the healthy young adult cervical spine during dynamic functional loading. *Journal of biomechanics* 48, 7 (2015), 1286–1293.
- Baptiste Angles, Daniel Rebain, Miles Macklin, Brian Wyvill, Loic Barthe, JP Lewis, Javier Von Der Pahlen, Shahram Izadi, Julien Valentin, Sofien Bouaziz, et al. 2019. Viper: Volume invariant position-based elastic rods. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 2, 2 (2019), 1–26.
- Mihai Anitescu and Florian A Potra. 1997. Formulating dynamic multi-rigid-body contact problems with friction as solvable linear complementarity problems. *Nonlinear Dynamics* 14, 3 (1997), 231–247.
- SK Bhalla and EH Simmons. 1969. Normal ranges of intervertebral-joint motion of the cervical spine. *Canadian Journal of Surgery* 12, 2 (1969), 181–187.
- Sofien Bouaziz, Sebastian Martin, Tiantian Liu, Ladislav Kavan, and Mark Pauly. 2014. Projective dynamics: Fusing constraint projections for fast simulation. *ACM transactions on graphics (TOG)* 33, 4 (2014), 1–11.
- Alexander G. Bruno, Mary L. Bouxsein, and Dennis E. Anderson. 2015. Development and Validation of a Musculoskeletal Model of the Fully Articulated Thoracolumbar Spine and Rib Cage. *Journal of Biomechanical Engineering* 137, 8 (06 2015). <https://doi.org/10.1115/1.4030408> arXiv:https://asmedigitalcollection.asme.org/biomechanical/article-pdf/137/8/081003/6092282/bio_137_08_081003.pdf 081003.
- Daniel LA Camacho, William R Ledoux, Eric S Rohr, Bruce J Sangeorzan, and Randal P Ching. 2002. A three-dimensional, anatomically detailed foot model: a foundation for a finite element simulation and means of quantifying foot-bone position. *Journal of rehabilitation research and development* 39, 3 (2002), 401–410.
- Dario Cazzola, Timothy P Holsgrove, Ezio Preatoni, Harinderjit S Gill, and Grant Trewartha. 2017. Cervical spine injuries: a whole-body musculoskeletal model for the analysis of spinal loading. *PLoS one* 12, 1 (2017), e0169329.
- Isaac Chao, Ulrich Pinkall, Patrick Sanan, and Peter Schröder. 2010. A simple geometric model for elastic deformations. *ACM transactions on graphics (TOG)* 29, 4 (2010), 1–6.
- Richard W Cottle, Jong-Shi Pang, and Richard E Stone. 2009. *The linear complementarity problem*. SIAM.
- B Daumas, WL Xu, and J Bronlund. 2005. Jaw mechanism modeling and simulation. *Mechanism and Machine Theory* 40, 7 (2005), 821–833.
- Scott L Delp, Frank C Anderson, Allison S Arnold, Peter Loan, Ayman Habib, Chand T John, Eran Guendelman, and Darryl G Thelen. 2007. OpenSim: open-source software to create and analyze dynamic simulations of movement. *IEEE transactions on biomedical engineering* 54, 11 (2007), 1940–1950.
- Ali Hamadi Dicko, Nicolas Tong-Yette, Benjamin Gilles, François Faure, and Olivier Palombi. 2015. Construction and Validation of a Hybrid Lumbar Spine Model For the Fast Evaluation of Intradiscal Pressure and Mobility. *International Science Index, Medical and Health Science* 9, 2 (Feb. 2015), 134–145. <https://doi.org/10.5281/zenodo.1099356>
- Richard Drake, A Wayne Vogl, and Adam WM Mitchell. 2009. *Gray's anatomy for students E-book*. Elsevier Health Sciences.
- Roy Featherstone. 2014. *Rigid body dynamics algorithms*. Springer.
- Ryutaro Fujii, Hironobu Sakaura, Yoshihiro Mukai, Noboru Hosono, Takahiro Ishii, Motoki Iwasaki, Hideki Yoshikawa, and Kazuomi Sugamoto. 2007. Kinematics of the lumbar spine in trunk rotation: in vivo three-dimensional analysis using magnetic resonance imaging. *European spine journal* 16, 11 (2007), 1867–1874.
- Theodore F Gast, Craig Schroeder, Alexey Stomakhin, Chenfanfu Jiang, and Joseph M Teran. 2015. Optimization integrator for large time steps. *IEEE transactions on visualization and computer graphics* 21, 10 (2015), 1103–1115.
- F. Ghezelbash, A.H. Eskandari, A. Shirazi-Adl, N. Arjmand, Z. El-Ouaaid, and A. Plamondon. 2018. Effects of motion segment simulation and joint positioning on spinal loads in trunk musculoskeletal models. *Journal of Biomechanics* 70 (2018), 149–156. <https://doi.org/10.1016/j.jbiomech.2017.07.014> 2nd International Workshop on Spine Loading and Deformation.
- Daniel Holden, Taku Komura, and Jun Saito. 2017. Phase-functioned neural networks for character control. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–13.
- Q Youn Hong, Sang Il Park, and Jessica K Hodgins. 2010. A Data-driven Segmentation for the Shoulder Complex. In *Computer Graphics Forum*, Vol. 29. Wiley Online Library, 537–544.
- Yixin Hu, Qingnan Zhou, Xifeng Gao, Alec Jacobson, Denis Zorin, and Daniele Panozzo. 2018. Tetrahedral meshing in the wild. *ACM Trans. Graph.* 37, 4 (2018), 60–1.
- Yifeng Jiang, Tom Van Wouwe, Friedl De Groote, and C Karen Liu. 2019. Synthesis of biologically realistic human motion using joint torque actuation. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–12.
- IA Kapandji. 1974. *The Physiology of the Joints*. vol III. *Trunk and Vertebral Column, 1st ed.* New York: Churchill Livingstone 10 (1974), 88–90.
- Taku Komura, Yoshihisa Shinagawa, and Toshiyasu L Kunii. 2000. Creating and retargeting motion by the musculoskeletal human body model. *The visual computer* 16, 5 (2000), 254–270.
- Lucas Kovar, Michael Gleicher, and Frédéric Pighin. 2008. Motion graphs. In *ACM SIGGRAPH 2008 classes*. 1–10.
- Binh Huy Le and Zhigang Deng. 2014. Robust and accurate skeletal rigging from mesh sequences. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–10.
- Jeongseok Lee, Michael X. Grey, Sehoon Ha, Tobias Kunz, Sumit Jain, Yuting Ye, Siddhartha S. Srinivasa, Mike Stilman, and C. Karen Liu. 2018a. DART: Dynamic Animation and Robotics Toolkit. *The Journal of Open Source Software* 3, 22 (2018), 500.
- Seunghwan Lee, Ri Yu, Jungnam Park, Mridul Aanjaneya, Eftychios Sifakis, and Jehee Lee. 2018b. Dexterous manipulation and control with volumetric muscles. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–13.
- Sung-Hee Lee, Eftychios Sifakis, and Demetri Terzopoulos. 2009. Comprehensive Biomechanical Modeling and Simulation of the Upper Body. *ACM Trans. Graph.* 28, 4, Article 99 (sep 2009), 17 pages. <https://doi.org/10.1145/1559755.1559756>
- Sung-Hee Lee and Demetri Terzopoulos. 2006. Heads up! Biomechanical Modeling and Neuromuscular Control of the Neck. *ACM Trans. Graph.* 25, 3 (jul 2006), 1188–1198. <https://doi.org/10.1145/1141911.1142013>
- Sung-Hee Lee and Demetri Terzopoulos. 2008. Spline joints for multibody dynamics. In *Karen SIGGRAPH 2008 papers*. 1–8.
- C Karen Liu, Aaron Hertzmann, and Zoran Popović. 2005. Learning physics-based motion style with nonlinear inverse optimization. *ACM Transactions on Graphics (TOG)* 24, 3 (2005), 1071–1081.
- David J Magee. 2013. *Orthopedic Physical Assessment-E-Book*. Elsevier Health Sciences.
- Sebastian Martin, Bernhard Thomaszewski, Eitan Grinspun, and Markus Gross. 2011. Example-Based Elastic Materials. *ACM Transactions on Graphics* 30, 4 (2011), 72.
- Walter Maurel and Daniel Thalmann. 2000. Human shoulder modeling including scapulo-thoracic constraint and joint sinus cones. *Computers & Graphics* 24, 2 (2000), 203–218.
- Aleka McAdams, Yongning Zhu, Andrew Selle, Mark Empey, Rasmus Tamstorf, Joseph Teran, and Eftychios Sifakis. 2011. Efficient elasticity for character skinning with contact and collisions. In *ACM SIGGRAPH 2011 papers*. 1–12.
- Philip W McClure, Lori A Michener, Brian J Sennett, and Andrew R Karduna. 2001. Direct 3-dimensional measurement of scapular kinematics during dynamic movements in vivo. *Journal of shoulder and elbow surgery* 10, 3 (2001), 269–277.
- Gary Monheit and Norman I Badler. 1990. A kinematic model of the human spine and torso. (1990).
- Masaki Nakada, Tao Zhou, Honglin Chen, Tomer Weiss, and Demetri Terzopoulos. 2018. Deep learning of biomimetic sensorimotor control for biomechanical human animation. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–15.
- Manohar M. Panjabi, Richard A. Brand, and Augustus A. White. 1976. Three-dimensional flexibility and stiffness properties of the human thoracic spine. *Journal of Biomechanics* 9, 4 (1976), 185–192. [https://doi.org/10.1016/0021-9290\(76\)90003-8](https://doi.org/10.1016/0021-9290(76)90003-8)
- Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel Van de Panne. 2018. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions On Graphics (TOG)* 37, 4 (2018), 1–14.
- Charles A Rockwood. 2009. *The shoulder*. Vol. 1. Elsevier Health Sciences.
- Prashant Sachdeva, Shinjiro Sueda, Susanne Bradley, Mikhail Fain, and Dinesh K Pai. 2015. Biomechanical simulation and control of hands and tendinous systems. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–10.
- Andrew Selle, Michael Lentine, and Ronald Fedkiw. 2008. A mass spring model for hair simulation. In *ACM SIGGRAPH 2008 papers*. 1–11.
- Marco Senteler, Bernhard Weiss, Dominique A. Rothenfluh, and Jess G. Snedeker. 2016. Intervertebral reaction force prediction using an enhanced assembly of OpenSim models. *Computer Methods in Biomechanics and Biomedical Engineering* 19, 5 (2016), 538–548. <https://doi.org/10.1080/10255842.2015.1043906>
- Ajay Seth, Meilin Dong, Ricardo Matias, and Scott Delp. 2019. Muscle contributions to upper-extremity movement and work from a musculoskeletal model of the human shoulder. *Frontiers in neurorobotics* 13 (2019), 90.
- Ajay Seth, Ricardo Matias, António P Veloso, and Scott L Delp. 2016. A biomechanical model of the scapulothoracic joint to accurately capture scapular kinematics during shoulder movements. *PLoS one* 11, 1 (2016), e0141028.
- Eftychios Sifakis, Igor Neverov, and Ronald Fedkiw. 2005. Automatic determination of facial muscle activations from sparse motion capture marker data. In *ACM SIGGRAPH 2005 Papers*. 417–425.
- Jack M Wang, Samuel R Hamner, Scott L Delp, and Vladlen Koltun. 2012. Optimizing locomotion controllers using biologically-based actuators and objectives. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 1–11.
- Wei Wang, Dongmei Wang, Friedl De Groote, Lennart Scheys, and Ilse Jonkers. 2020. Implementation of physiological functional spinal units in a rigid-body model of the thoracolumbar spine. *Journal of Biomechanics* 98 (2020), 109437.

- Hans-Joachim Wilke, Peter Neef, Marco Caimi, Thomas Hoogland, and Lutz E Claes. 1999. New in vivo measurements of pressures in the intervertebral disc in daily life. *Spine* 24, 8 (1999), 755–762.
- Grady Williams, Andrew Aldrich, and Evangelos A Theodorou. 2017. Model predictive path integral control: From theory to parallel computation. *Journal of Guidance, Control, and Dynamics* 40, 2 (2017), 344–357.
- Gregory M Williams, Elaine F Chan, Michele M Temple-Wong, Won C Bae, Koichi Masuda, William D Bugbee, and Robert L Sah. 2010. Shape, loading, and motion in the bioengineering design, fabrication, and testing of personalized synovial joints. *Journal of biomechanics* 43, 1 (2010), 156–165.
- Andrew Witkin and Michael Kass. 1988. Spacetime constraints. *ACM Siggraph Computer Graphics* 22, 4 (1988), 159–168.
- Zhan Xu, Yang Zhou, Evangelos Kalogerakis, Chris Landreth, and Karan Singh. 2020. Rignet: Neural rigging for articulated characters. *arXiv preprint arXiv:2005.00559* (2020).
- KangKang Yin, Kevin Loken, and Michiel Van de Panne. 2007. Simbicon: Simple biped locomotion control. *ACM Transactions on Graphics (TOG)* 26, 3 (2007), 105–es.
- Felix E Zajac. 1989. Muscle and tendon: properties, models, scaling, and application to biomechanics and motor control. *Critical reviews in biomedical engineering* 17, 4 (1989), 359–411.
- Zygote. [n. d.]. Zygote 3D model. <http://https://www.zygote.com>. [Online;].
- Jānis Šavlovskis and Kristaps Raits. [n. d.]. Biomechanics of the Spine. <https://www.anatomystandard.com>