

Hand Modeling and Simulation Using Stabilized Magnetic Resonance Imaging

BOHAN WANG, University of Southern California, USA
GEORGE MATCUK, University of Southern California, USA
JERNEJ BARBIČ, University of Southern California, USA




Fig. 1. **Opposition of the thumb:** Our method produces an accurate data-driven skeleton mesh kinematic “rig”. Here, we used our rig to reliably reproduce the well-known opposition of the thumb to all other 4 fingers. The solid soft tissue was computed using a FEM simulation attached to the articulated bone meshes. The fingers do not merely touch but also orient to be co-planar at the contact location, subject to biomechanical limits. The right-most four images show representative FEM frames of the opposition between the thumb and the pinky finger.

We demonstrate how to acquire complete human hand bone anatomy (meshes) in multiple poses using magnetic resonance imaging (MRI). Such acquisition was previously difficult because MRI scans must be long for high-precision results (over 10 minutes) and because humans cannot hold the hand perfectly still in non-trivial and badly supported poses. We invent a manufacturing process whereby we use lifecasting materials commonly employed in film special effects industry to generate hand molds, personalized to the subject, and to each pose. These molds are both ergonomic and encasing, and they stabilize the hand during scanning. We also demonstrate how to efficiently segment the MRI scans into individual bone meshes in all poses, and how to correspond each bone’s mesh to same mesh connectivity across all poses. Next, we interpolate and extrapolate the MRI-acquired bone meshes to the entire range of motion of the hand, producing an accurate data-driven animation-ready rig for bone meshes. We also demonstrate how to acquire not just bone geometry (using MRI) in each pose, but also a matching highly accurate surface geometry (using optical scanners) in each pose, modeling skin pores and wrinkles. We also give a soft tissue Finite Element Method simulation “rig”, consisting of novel tet meshing for stability at the joints, spatially varying geometric and material detail, and quality constraints to the acquired skeleton kinematic rig. Given an animation sequence of hand joint angles, our FEM soft tissue rig produces quality hand surface shapes in arbitrary poses in the hand range of motion. Our results qualitatively

Authors’ addresses: Bohan Wang, University of Southern California, Los Angeles, CA, USA, bohanwan@usc.edu; George Matcuk, University of Southern California, Los Angeles, CA, USA, matcuk@usc.edu; Jernej Barbič, University of Southern California, Los Angeles, CA, USA, jnb@usc.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.

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.
0730-0301/2019/7-ART115 \$15.00
<https://doi.org/10.1145/3306346.3322983>

reproduce important features seen in the photographs of the subject’s hand, such as similar overall organic shape and fold formation.

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

Additional Key Words and Phrases: hand, MRI, animation, simulation, modeling, anatomy, deformable objects, FEM, optical scanning

ACM Reference Format:

Bohan Wang, George Matcuk, and Jernej Barbič. 2019. Hand Modeling and Simulation Using Stabilized Magnetic Resonance Imaging. *ACM Trans. Graph.* 38, 4, Article 115 (July 2019), 14 pages. <https://doi.org/10.1145/3306346.3322983>

1 INTRODUCTION

Hands and their modeling and animation are of paramount importance in many applications. In computer games and film, clothing may occlude the body of the characters, but the hands are often exposed and important for the story. In engineering, hand anatomical models can be used to design tools and equipment. In computer vision, anatomically based modeling can improve the tracking of hands, because it provides better statistical priors on hand shapes. In healthcare, accurate hand shapes and motion can be used to design better finger and partial hand prosthetics, and better tools for surgeons. As is well-known and argued by medical authorities in the field [Kapandji 2009], hand’s dexterity stems from a thumb opposing four fingers (Figure 1), making the hand perfectly suited for precise grasping, lifting, climbing and daily manipulations of objects. To improve realism, virtual hands should be modeled similarly to biological hands, and this requires building precise anatomical and kinematic models of real human hands.

Unfortunately, complex hand biomechanics today is not modeled, measured or resolved in any quality way. While there are several computational models of the hand kinematics, few attempted to

model internal structures, let alone produce anatomically precise models that match real-world data. Existing detailed human hand anatomy books such as [Kapandji 2009] focus on applications in medicine, where the goal is to treat and repair hand injuries, not analyze the function of healthy hands. It is challenging to acquire the motion of the internal human hand anatomy. It must be performed *in vivo* (on a live person), as the hand motion is greatly affected by the tendons, fat tissue and active muscles. It cannot be reliably done using exterior scanning techniques (motion capture), due to tissue deformations and sliding. We give a method to acquire high-resolution geometry of the bones of the human hand (collectively referred to as the “skeleton”) in multiple hand poses, using MRI scanning. We scan two subjects (1 male and 1 female) in 12 poses. Our poses were chosen to reasonably sample the range of motion of each important hand joint. We note that in medicine, it is common to use cadavers for human anatomy research. We exclude such studies because we want to study the kinematics of healthy hands, as activated by a live person.

In modern medicine, there are two imaging techniques that are potentially suitable to acquire 3D hand anatomy. In Computed Tomography (CT), a 3D image is generated from a large number of two-dimensional X-ray images taken around a fixed axis of rotation. In computer animation, hand bone skeletons have been acquired using CT scans in the pioneering work of Kurihara and Miyata [2004]. Similarly, Marai et al. have also used CT scans to acquire the carpal bone motion of the human hand [Marai et al. 2003]. Unfortunately, CT emits ionizing radiation, which can cause cancer in high radiation doses. Building a quality skeleton animation model requires acquiring many poses for each individual, leading to unacceptably high radiation doses. We therefore decided to exclude CT from our work. Magnetic Resonance Imaging (MRI) does not use any ionizing radiation nor pose other healthy risk for the vast majority of people. Compared to Kurihara’s CT work, which produced CT scans for 1 individual in 5 poses, we produced MRI scans for 2 individuals in 12 poses without any health risks. We performed our MRI scans in a hospital, using clinical 3T scanners manufactured by General Electric. We obtained our institution’s review board (IRB) human subjects research approval for these experiments. An additional advantage of MRI over CT is that CT only generates sharp images of bones, whereas MRI produces all major hand organs, including bones, muscles and fat. For example, in the CT images in Kurihara’s work [Kurihara and Miyata 2004], one can barely observe any significant hand components other than the bones and the surface (Figure 2). Although we do not utilize muscles and fat in this work, their availability bodes well for further research in this area.

Although MRI scanners can provide great anatomical detail, there is a practical challenge that prior work has not addressed: the hand




Fig. 2. Left: Kurihara’s CT scan [Kurihara and Miyata 2004]. Reprinted with permission from Eurographics. Right: our MRI scan.

must be kept perfectly still in the scanner for 10-15 minutes. Long scanning times are needed to decrease the signal-to-noise ratio, which decreases with the square root of the scanning time [Stillfried 2015]. If the hand is not still, the scanning image quickly loses resolution and becomes useless. If the goal is to scan just one pose, this is doable (not easily comfortable, but doable) for most people: lie down on the scanner bed, firmly press the hand against the bed, and hold it still during the scan. However, general hand poses cannot be scanned in this way, because they position the palm and fingers in some general configuration, e.g., bend the fingers, close the hand into a fist, partially close the hand, and so on. Such poses do not have any physical support inside the scanner. Without support, they are impossible to scan sharply, as a human simply cannot keep an unsupported hand pose still for more than a few seconds.

Our solution is to generate molds that hold the hand in a known and fixed pose. We generate our molds by placing a hand, in a chosen pose, into a liquid lifecasting solution that solidifies into an elastic rubber-like solid, then retrieve the hand. We then fill the hole with liquid plastic, generating a plastic replica of the hand in each pose. This replica captures extremely high detail on the skin’s surface. We then scan the plastic replica using an optical 3D scanner (Artec Spider [Artec3D 2018]), and repeat the lifecasting process to generate a rubber mold (negative image of plastic hand). We then cut the mold in two parts, place the hand into the mold, cover it with the upper part, and place the hand into the MRI scanner. To the best of our knowledge (including that of our co-author Dr. Matcuk who is a medical doctor radiologist and active researcher in radiology), we give the first solution for stable long-time (>10 minutes) MRI scanning of the hand in multiple poses. Another key benefit of our approach is that we can obtain, for each pose, not just the anatomy in that pose, but also a high-resolution surface skin detail in that pose, which can be used for high-resolution hand rendering.

Next, we demonstrate how to segment the MRI scans into bone meshes, for each pose. We show how to register the bone meshes to have equal number of vertices and triangle connectivity. Although there is substantial literature on *body* MRI data segmentation, the literature on *hand* MRI segmentation is very sparse. We give new methods that improve the state-of-the-art (Section 3.1). We then demonstrate how to build a data-driven skeleton kinematic model (a skeleton animation “rig”). This rig is able to articulate the skeleton mesh into an arbitrary pose, using interpolation / extrapolation of the acquired MRI meshes. Our rig captures complex real-world rotations and translations of bones relative to their parent bone.

We also give a new Finite Element Method (FEM) simulation method for the hand’s soft tissue (“soft tissue rig”). Naive tet meshing produces unstable tet meshes that explode due to pinching of the elastic material at the joints, or over-constraining at the joints. We show how to stably constrain the soft tissue tetrahedral mesh to the skeleton mesh animated using our bone rig. We give a method to increase the spatial mesh resolution and adjust elastic material properties in selected regions, which can reliably and automatically simulate folds and creases, both on the palm and under fingers. We show experimentally that our bone rig greatly improves the stability of our soft tissue rig, compared to naive bone rigs. We also demonstrate that our FEM soft tissue rig outperforms skinning with standard Maya skinning weights.

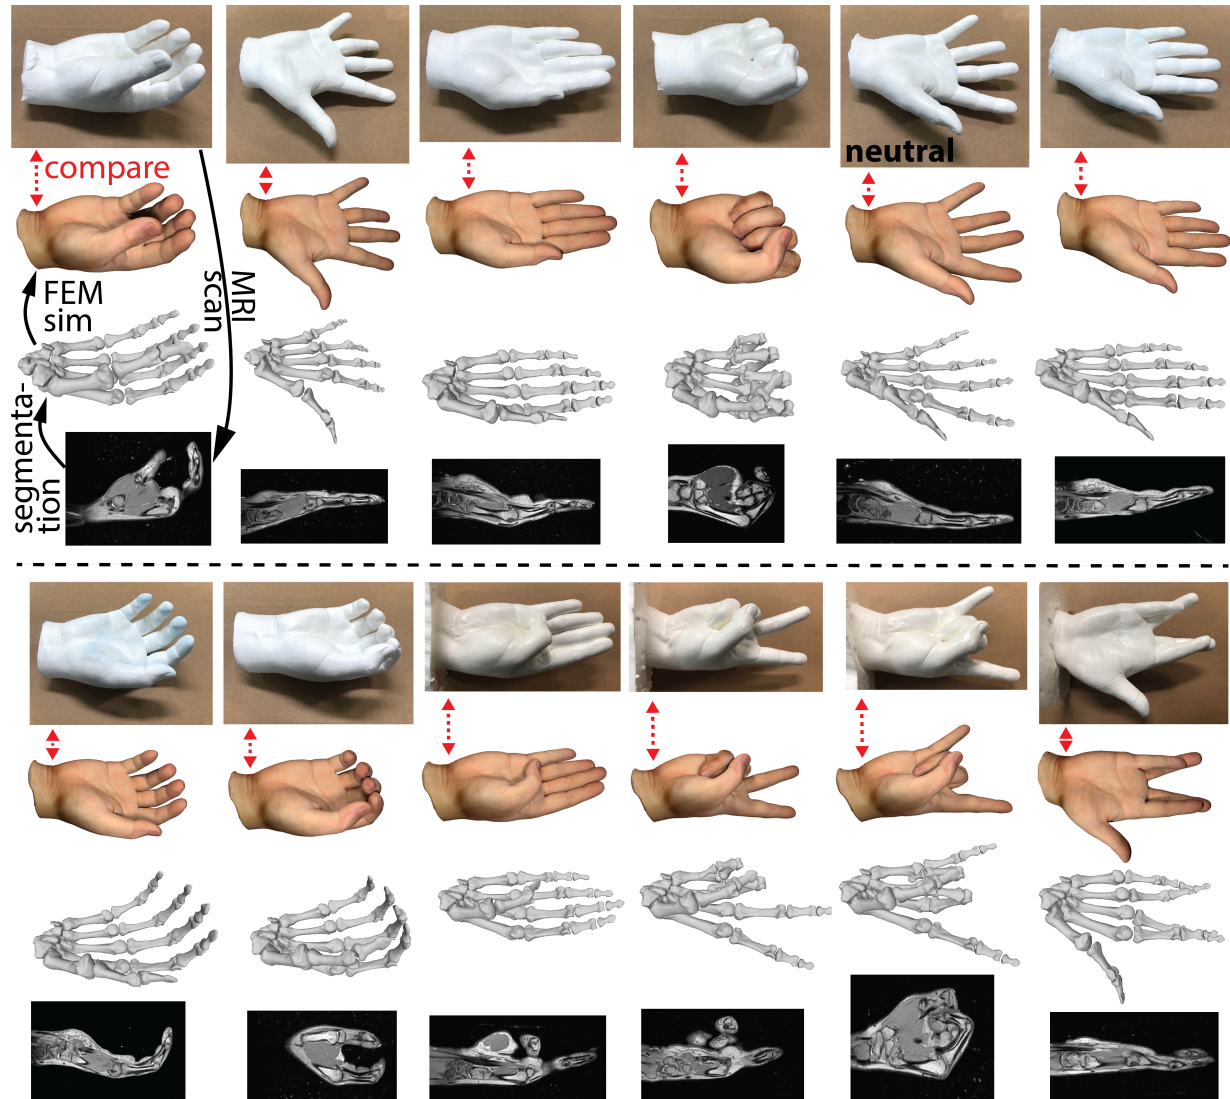


Fig. 3. 12 MRI-scanned poses (1 subject; male; late 20s). Four rows are the cloned plastic hand (the "ground truth" shapes), our FEM result, skeleton mesh obtained from MRI scan, and the MRI scan. The last 4 plastic poses have a plastic stand (seen on the left of each image) to easily stand up on a table. FEM closely matches the plastic hand in all poses, despite not actually using plastic shapes anywhere in our system (except the neutral shape: column 5 in the top set). The solid black arrows indicate the logical sequence of steps; the FEM and cloned hands are shown adjacent for easier comparison. Image is zoomable.

2 RELATED WORK

Hands are biomechanically complex and it is therefore natural to model their shapes using physically based simulation; such as, for example, simulating a solid mesh constrained to the underlying skeleton [Capell et al. 2005; Kim and Pollard 2011; Liu et al. 2013]. Lee et al. comprehensively modeled the upper human body using anatomically based simulation [Lee et al. 2009]. However, they excluded hands from simulation and modeled them kinematically. The skin and tendons of the human hand have been simulated to enhance the visual appearance or control of hand articulation [Li et al. 2013; Sachdeva et al. 2015; Sueda et al. 2008]. McAdams et al. [2011] and Smith et al. [2018] used FEM to simulate cartoon hands. For human hands, Kry et al. [2002] and Garre et al. [2011] modeled hand

deformations using a layer of FEM soft tissue around articulated bones. We give a hand FEM simulation model that is more elaborate than previous approaches. We model spatially varying geometric detail and material properties, automatically remove elastic material in between the heads of bones, and determine attachments to bones at the joints to avoid pinching. If such features are omitted, it is in our experience impossible to stably simulate hands that aspire to look like real hands. Our approach can, for example, model the fold formation on the entire hand.

To a first-degree approximation, hands can be animated using skinning techniques [Jacobson et al. 2014]. Skinning can be improved using several techniques, such as dual quaternions [Kavan et al. 2008], or implicit methods [Vaillant et al. 2014]. Pose-Space

Deformation (PSD) is a method that combines skeleton subspace deformation [Magenat-Thalmann et al. 1988] with artist-corrected pose shapes [Lewis et al. 2000]. It is widely used in industry due to its speed, simplicity and the ability to incorporate real-world scans and arbitrary artist corrections. Kurihara and Miyata [2004] presented a variant of PSD suitable for hand animation, and Rhee et al. [2006] demonstrated how to efficiently implement it on a GPU.

Recently, Romero et al. [2017] gave a surface scanning approach to animate human hands together with the rest of the body. Their model is purely data-driven and does not model interior anatomy. Different from our approach, their goal is to generate a statistical model to parameterize the variations of human hand shapes due to personalization and pose. In contrast to data-driven methods [Kurihara and Miyata 2004; Romero et al. 2017] and model-based methods [Garre et al. 2011; Kry et al. 2002], we demonstrate how to *simultaneously* acquire both a kinematic model of the internal anatomy and high-quality surface scans in matching poses. We are not aware of any prior work that has achieved this for hands.

Our hands are driven by standard and familiar skeleton joint angle animations. Such animations can be created manually by artists, or acquired from real performances, for example, using gloves [CyberGlove Systems 2017; Han et al. 2018; Wang and Popović 2009], or computer vision techniques [LeapMotion 2017; NimbleVR 2012; Wang et al. 2011]. In computer animation, there are many methods that compute kinematic hand skeleton models based on surface tracking methods, such as optical or magnetic motion capture; see [Wheatland et al. 2015] for a good survey. Unlike surface-based techniques, our imaging-based method can generate internal anatomy, which is required for anatomically based simulation. Because of their small motion and combined effect on the surface of the palm, surface-based techniques have difficulties tracking the bones in the palm. Our MRI method works equally well for all bones. We were able to extract the rigid motion of palmar bones without any difficulty.

Common medical imaging techniques that are in principle applicable to in-vivo measurements include X-rays, Computed Tomography (CT), Positron Emission Tomography (PET), Magnetic Resonance Imaging (MRI), and Ultrasound (for a good review, see [Duncan and Ayache 2000; Fenster and Downey 1996]). PET requires injecting the subjects with positron-emitting radionuclides, equivalent to 8-years of natural background radiation [RadiologyInfo 2018]. Although ultrasound is a safe imaging techniques and has been studied for decades [Fenster and Downey 1996; Solberg et al. 2007], it is not suitable for hands. This is because ultrasound signals are blocked by bones, and are thus unable to generate complex three-dimensional geometry due to sonic occlusions. MRI provides good contrast for all anatomical structures [Dempsey et al. 2002; Watson 2015], and can capture soft tissues such as muscles and fat. Commercial companies have used MRI to build complete human body 3D anatomy [Zygot 2016]. A few publications analyzed the hand bone geometry using MRI scans [Miyata et al. 2005; Rusu 2011; Stillfried 2015; van der Smagt and Stillfried 2008]. However, because the MRI scanning times need to be long to decrease the signal to noise ratio (10-15 minutes in our work), prior work has not addressed an important limitation. Namely, if the hand is not held perfectly still during the scan, the MRI image is blurred (useless). Of course, humans cannot

hold the hand still in arbitrary poses for 10 minutes. Researchers attempted to overcome this limitation by employing clay support, and by asking the subject to hold objects during the scan. Using such an approach, Miyata et al. [2005] succeeded in MRI-scanning and modeling 2 finger joints on 4 fingers, in 4 poses. In our work, we scanned the complete hand in 12 poses. We use ergonomic molds which, due to complete hand encasement, provide superior support to using clay or relying on grasping. To combat hand shake, Stillfried's work [2015] used a balancedSSFP MRI scanning sequence to shorten the scanning time to 2-3 minutes. In medicine, this sequence is usually employed to scan moving objects (e.g., a beating heart). While this reduces hand fatigue, it also decreases tissue detail and sacrifices the overall image quality. Furthermore, even if such procedures are used, the specific MRI-scanned poses in prior work do not correspond to any optical scan of the surface geometry. Note that the quality of any hand skin surface extracted from a MRI scan is far inferior to what can be achieved via surface optical scanning. In our work, we use the PD Cube MRI sequence, which is a common MRI sequence to image stationary objects. That said, the "fast MRI" (balancedSSFP) approaches are orthogonal to our mold stabilization approach and could be combined if so desired.

Most medical imaging segmentation techniques focus on the general body (and many on the hip area), and not on hands (see [McInerney and Terzopoulos 2008] for a good review). Although segmentation algorithms with prior knowledge have succeeded in automatic musculoskeletal segmentation of the hip [Baudin et al. 2012; Gilles and Magnenat-Thalmann 2010; Grady 2006; Schmid et al. 2011; Schmid and Magnenat-Thalmann 2008], they have not been demonstrated to work on hands. Note that hands have a much smaller and more complicated structure. Compared to the hip, a human hand (without the wrist) has 27 bones and 19 flexible joints, which dramatically increases the segmentation complexity. The most common approach (aimed at general human bodies, but also applicable to hands) in practice is to mostly perform a manual segmentation, with the help of a few basic computer vision techniques such as thresholding, filtering and water-shedding [Stillfried 2015; van der Smagt and Stillfried 2008]. Such capabilities are often integrated into commercial medical image analysis software, such as Amira [Amira 2018], VTK [VTK 2018] and ITK-SNAP [ITK-SNAP 2018]. We tried Amira, but found it too inefficient for hands and not automated enough for our application. Recently, with the rapid development of GPUs, deep neural network techniques (DNN) became a popular and powerful technique for image segmentation, including MRI segmentation [Deniz et al. 2017; Kayalibay et al. 2017]. However, these techniques are still in their infancy, e.g., [Kayalibay et al. 2017] does not solve the inhomogeneous bone problem. DNN approaches are currently inhibited due to a lack of hand medical imaging datasets and approaches to stably build such datasets.

3 ACQUIRING SKELETON MESHES IN MANY POSES

We now describe our procedure to acquire bone geometry **and** detailed external hand surface geometry and color texture in **multiple** hand poses, spanning the typical range of motion of the human hand (Figure 3). Our skeleton geometry is acquired using MRI, and the external geometry is acquired using laser scanner metrology (for

shape) and photogrammetry (for color texture). For each bone, we generate a mesh for each pose, with equal mesh connectivity. We then use the skeleton geometry to run FEM simulations to generate high-quality unscanned “in-between” poses, producing registered high-resolution surface geometry that can be articulated to any continuous pose of interest.

We stabilize the human hand inside the MRI scanners in known prescribed poses, by manufacturing rigid molds. Our technology “silver bullet” is the observation that human hand can be physically **cloned** into a shape made of a variety of materials (plastic, silicon, etc.) using lifecasting materials (Figure 5). Inspired by the film special and visual effects industry, we observe that it is possible to generate extremely precise replicas of human hands, using commonly available rubber-like materials such as the AljaSafe™ material from Smooth-On, Inc [AljaSafe 2018]. We note that lifecasting has been used in the visual effects industry to model human faces; we are not aware of prior usage to model hands in multiple poses. Lifecasting has been used for fabricating an accurate neutral shape of a human hand in robotics [King et al. 2018; Schlagenhaut et al. 2018], to perform hand control, such as for grasping.

We perform lifecasting using AljaSafe [AljaSafe 2018], an alginate skin-safe material naturally occurring in the cells of brown algae. The resulting hand replicas capture extremely detailed surface hand features, down to individual pores and tiny lines on the hand surface (Figure 5, g). In order to make the replicas, we prepare a liquid AljaSafe solution (mixture of powder and water) in a properly hand-sized bucket, and then position the hand into it, in a chosen pose. The solution solidifies in approximately 8-10 minutes into a rubbery solid. The shape of the hand and the micro-detail on the surface get imprinted into the AljaSafe surface. We note that AljaSafe gently encases the hand. It does not change volume when it solidifies, so there is minimal hand compression; only small water-like hydrostatic pressure due to gravity. Because AljaSafe is sturdy, but sufficiently flexible, it is easily possible to withdraw the hand, without damaging the imprinted surface detail. This process produces a bucket of solidified AljaSafe rubber with a hole in the shape of the hand pose. We then fill this hole with liquid plastic (Smooth-Cast™ 300Q), which solidifies in a few minutes. We then remove the AljaSafe rubber. The result is a high-quality plastic replica of the hand in the specific pose (Figure 5, g, h). This manufacturing process is also shown in our video. We note that the subject just has to approximately eyeball each pose during lifecasting. The created mold and the plastic hand are automatically consistent.

We then scan the plastic hand using a precise laser scanner, such as Artec Spider [Artec3D 2018]. We scan the plastic hand, as opposed to the real hand directly, because the plastic hand is perfectly still, and therefore such a scan can achieve high precision (Figure 5, h), on the order of 0.05 - 0.1mm with Artec Spider, producing a mesh with 11M triangles. For comparison, the number of hand triangles in the dataset on the MANO website [Romero et al. 2017] is 100K. The Artec Spider precision already is an order of magnitude better than the precision of the skin geometry acquired by our MRI scan (which is 0.5 - 1mm). We are not aware of any prior computer graphics work to use such a chemistry-based solution to “clone” hands. We observe that computer graphics and vision acquisition

research typically focuses on optical surface scanning technologies, e.g., reconstructing shapes from multiple camera views. Chemistry-based approaches offer clear advantages in their ability to capture high-precision surface detail on human body parts with **minimal** costs: we only spent about \$600 worth of chemical materials to make our replicas for 1 subject.

By repeating the above process for every pose, we obtain a high-quality plastic hand replica in each pose, as well as a high-precision surface geometry scan of each shape. Next, we use these plastic hands to stabilize the real hand during the MRI scan, by generating a **mold** into which the subject inserts the hand prior to MRI scanning (Figure 4). The mold keeps the hand still in a fixed, known, optically scanned pose. The presence of the mold does not corrupt the MRI signal in any way. We tested this in practice and discussed the issue with MRI technicians. Our mold generates almost no signal outside of the hand, and zero signal inside the hand. We generate the mold by repeating the AljaSafe casting process. We prepare the liquid AljaSafe solution, and then place the **plastic** hand into it. After AljaSafe solidifies, we cut the resulting mold in two parts along the hand’s frontal (i.e., coronal) plane, using a precision knife (Figure 4, a). For MRI scanning, we place the lower part of the mold onto the scanner bed. The subject inserts her/his hand into the mold, and then the hand is covered and secured by the top part of the mold, and additionally secured and fastened with tape (Figure 4, c). Because the top mold piece is largely supported by the bottom piece, and because our molds are precise hand negative images, there is only a minimal gravity and twisting during MRI scanning. Note that if one does not need the plastic hand, one can already use the first AljaSafe cast for the MRI mold. However, molds obtained from plastic hands can reorient the hand, for MRI ergonomics. The mold has to be sturdy to keep the hand still, but still somewhat flexible, so that the subject can “nudge” the hand position slightly before the scan, for ergonomics during the scan. This will inevitably result in a small mismatch between the plastic hand pose that was used to make the mold, and the actual scanned MRI pose. We address this in Section 5.2.

3.1 Segmenting MRI data into bone meshes

The MRI scan data consist of a regular volumetric grid of scalar values proportional to tissue mass density, and must be segmented into individual bone meshes. Although this task is well-understood for general full-body MRI scans, the literature on segmenting **hand** MRI scans is very sparse. We found only a single research paper (in any field) that discusses (in any substantial detail) how to segment bones from **hand** MRI images [Rusu 2011], and even this publication only showed a single example. We implemented their method and found it challenging to use: (1) there are more than four parameters to tune, and (2) the result cannot be incrementally enhanced by users. Moreover, in many cases it did not work well (for example, see Figure 6), so we abandoned it. We give a new technique to perform such a segmentation. In general, common issues in MRI images include (a) inhomogeneous bone tissue intensity: cortical bone (i.e. outer bone layer) is darker than cancellous bone (inner layer), and (b) unclear and fuzzy bone boundaries in articulation areas, as described in [Schmid et al. 2011].




Fig. 4. **Stabilized MRI scanning:** (a) The mold with a plastic hand before cutting. (b1,2,3) Mold has been cut into two parts. (c) Hand secured into the mold. (d) MRI scanning with the mold. Clinical MRI scanner manufactured by General Electric. Magnetic field strength of 3T, and resolution of $0.5 \times 0.5 \times 0.5 \text{ mm}^3$.




Fig. 5. We “cloned” this hand onto a plastic hand. (a) Mixing the AljaSafe. (b) The hand in a bucket filled with liquid AljaSafe. (c) Hand removed, leaving a hand-shaped hole. (d, e) Casting a liquid plastic into solidified AljaSafe. (f) Removed the AljaSafe to obtain the plastic hand. (g) Photo of plastic hand. (h) Rendered 3D-scanned mesh of plastic hand. Scanner: Artec Spider. Note the high-resolution detail in g, h (zoomable in PDF).




Fig. 6. **Our method outperforms prior hand bone segmentation methods:** It is challenging to segment the “metacarpal II” bone (red rectangle) for 2 reasons: (1) the MRI signal (shown in A) on the bone is not of uniform intensity, and (2) the neighboring bone “metacarpal III” (yellow rectangle; only head of the bone is visible in this 2D view) has very similar signal intensity to “metacarpal II”. Our method (shown in B) successfully segmented this challenging case. Prior method [Rusu 2011] uses a region-based active contour method and produces a suboptimal result: a significant part of II’s bone head is missing (C).

We give a method that segments the image with as few parameters as possible and is as automatic as possible. We use a variant of Laplacian-based segmentation [Grady 2006], but show how to extend it to 3D in a scalable manner. We also give an intuitive interface specifically for segmentation of hands. Our method has only




Fig. 7. **Our segmentation interface.** Showing the result of the first step of our 3D Laplacian-based segmentation. The segmented bone is red. Yellow is the local segmentation region selected by the user for this bone.

one parameter, and users are able to incrementally improve the result by delineating additional bone voxels which the segmentation classified incorrectly. The algorithm in [Grady 2006] is unacceptably slow when solving the Laplace equation in 3D on the entire volume. Our solution is to segment each bone separately and combine the results together. When merging segmentations into a global labeled volume, we resolve any labeling conflicts by treating conflicted voxels as unlabeled voxels again, and resolve the Laplacian problem again in a local region. Using our method, we were able to segment a $512 \times 512 \times 512$ volume of the cancellous bones of the entire hand in about 3 hours (5-10 minutes per bone). For cortical bones, there is no existing work on segmentation without using training data (which is unavailable for hands) [Schmid et al. 2011]. The Laplacian method does not work here because the cortical bone is dark in the MRI images and too similar to tendons. We segment it by the classical computer vision watershed method [Meyer 1992] in 3D. We deem the connected component closest to the cancellous bone as the cortical segmentation. Finally, we merge each cortical and cancellous pair into a single bone using dilation and/or random walk [Grady 2006]. After the medical image segmentation, it is not difficult to build the implicit function for each bone from the labeled images. With such implicit functions, we can easily generate the explicit 3D geometry using an isosurface meshing algorithm [Boissonnat and Oudot 2005].

We now describe how we register each bone’s pose-varying meshes to the same number of vertices and triangle connectivity, and how we compute a rigid transformation between the bone in the neutral pose and all other poses. For the neutral pose, we perform complete segmentation of the bone as described in the previous paragraph. This gives us a mesh of the cancellous part of the bone, and a complete bone mesh combining both the cancellous and cortical parts of the bone, in the neutral pose. For all other poses, we only perform cancellous segmentation. This saves a lot of time and is all that is needed. Unlike the cancellous bone which is white and has good contrast, the cortical bone is dark and intensive to segment in terms of manual work; and its segmentation in multiple poses is redundant. We bypass the need for segmenting cortical bones in non-neutral poses by executing a rigid Iterative Closest Point (ICP) algorithm [Besl and McKay 1992] to align the neutral cancellous bone mesh onto the cancellous mesh of each pose, extracting the translation $x \in \mathbb{R}^3$ and rotation matrix R for each pose. We then transform the entire neutral bone mesh using the rigid transformation (x, R) . This produces the output of our segmentation: a mesh for each bone in each pose, with the same connectivity across all poses. We smooth the bone meshes using 2 levels of the Laplacian smoothing in MeshLab [Cignoni et al. 2008]. Our final skeleton has 155,450 triangles total for 23 bones. Using these methods, we successfully segmented 12 hand poses of our two subjects (Figure 3). Figure 7 shows a screenshot of our segmentation user interface. Next, we demonstrate how we can animate this skeleton mesh, by building a “skeleton rig”.

4 SKELETON KINEMATIC MODEL

In computer animation, the motion of the human hand is typically modeled using a hierarchical joint structure; and we adopt this approach also in our work. Controlling motion by prescribing joint angles using commonly available inputs (mocap, IK, keyframing, etc.) is commonly done in practice because it is easier (and admittedly less principled) than optimizing muscle actuators [Lee et al. 2009]. Given the vector of joint angles θ , a typical pipeline in computer animation industry today is **not** to compute any bone geometry, but instead to bypass bones and apply a skinning algorithm that computes the positions of the skin vertices, based on θ and some suitably defined skinning weights, optionally with sculpted pose-varying corrections. Such modeling is imprecise because the fat layer and the skin are sliding relative to the bones, and because real bones undergo general 6-DOF rigid body motion relative to one another. The rigid transformations between bones depend nonlinearly on θ , and are dictated by the anatomical constraints between the bones and/or muscles, such as ligaments and tendons, and muscle activations. Accurate anatomical modeling requires computing the bone geometry (mesh vertex positions) in arbitrary hand poses. We therefore depart from previous work and give a new kinematic model (a “bone rig”) that computes the positions of all vertices in all bone meshes, based on θ . The input to our rig is θ , and the output is the anatomically based rigid transformation (translation and rotation) of each bone. We obtain our rig by fitting it to our segmented bone MRI meshes in the captured poses (Section 3). Our model can both interpolate and extrapolate the MRI-acquired poses. In later sections, we will

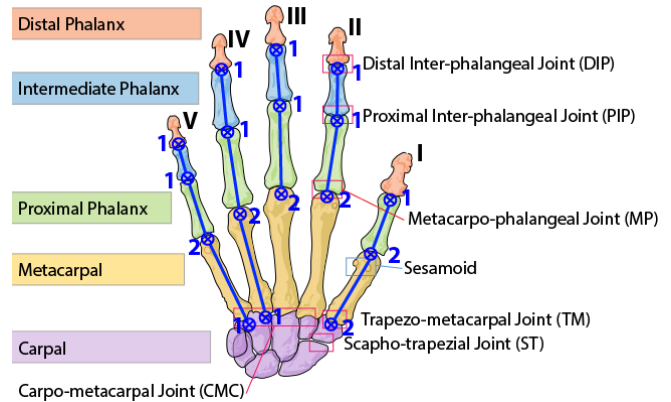


Fig. 8. **The bones and joints of a human hand.** Our joint hierarchy is shown superimposed in blue, with the number of joint DOFs indicated.

use our bone rig to compute FEM simulations of the soft tissue of the hand.

We note that Kurihara and Miyata [2004] optimized bone rotation centers to data, but did not optimize joint rotation axes. Instead, they specified the relative rotation of each joint via three Euler angles corresponding to unoptimized axes that in general do not match biomechanical axes. For example, the finger DIP and PIP joints rotate mostly around one axis. With unoptimized axes, none of the elemental Euler rotations correspond to the main biomechanical axis of rotation. Hence, the animator has to either accept inaccurate motion, or animate all three Euler angles simultaneously. In contrast, we optimize rotation axes to the MRI-acquired bone motion. We give a rig where 1-DOF and 2-DOF joints can be animated by specifying only 1 and 2 angles, respectively; namely the angles of rotations around our optimized joint axes (Figure 10). Our rig drives the biomechanical joint axes directly, and as such requires fewer parameters to animate.

Hand kinematics. Modeling the hand joints and bone motion is a complex problem. The human hand (see Figure 8) has 27 bones: 5 distal phalanges, 4 intermediate phalanges, 5 proximal phalanges, 5 metacarpals and 8 carpals, most of which are movable. The hand has 9 inter-phalangeal joints (DIPs and PIPs), 5 metacarpo-phalangeal joints (MP) and 5 carpo-metacarpal joints (CMC). The thumb is the only finger capable of opposing the other four fingers; therefore, it has a larger range of motion and greater complexity. The CMC joint of the thumb is usually called Trapezo-Metacarpal joint (TM). Researchers also study the inter-carpal joints, i.e., scapho-trapezoidal joint (ST) [Kapandji 2009]. The employed models vary largely according to the specific setup and application. The number of modeled joints can range from 5 to 36 [Yuan et al. 2017]. The number of DOFs of a single joint can range from 1 to 3. The motion of carpal bones is small and on (or slightly over) the boundary of our MRI scanning volume for our male subject; as commonly done [Stillfried 2015], we do not model this motion. The IP and CMC joint are usually modeled to have one rotation axis (or zero for CMC II and III), and MP and TM joints to have two rotation axes [Gustav and van der Smagt 2016; Kapandji 2009]. We adopt these conventions in our work. There are 17 joints in total in our human hand

Table 1. **Bone rig errors for each joint.** We list them separately for rotations (R) and translations (x). Here, FM_R and FM_x are the absolute errors using our Full Model (i.e., quadratic fitting of \hat{R} and \hat{x}), while PM_R and PM_x are the errors under a Partial Model where $\hat{R} = I$ and $\hat{x} = 0$. The error is computed as the average difference between the output of our bone rig, and the ground truth (segmented MRI data), across all 12 poses (male subject). It can be seen that the fitting of \hat{R} and \hat{x} decreases the error. The first and second half of the table give 1-DOF and 2-DOF joints, respectively. Observe that the pinky finger has the largest error; this is because it is the smallest. Consequently, it is resolved with fewer voxels in the MRI scan.

Joint	FM_R [deg]	PM_R [deg]	FM_x [mm]	PM_x [mm]
DIP I	2.40	2.98	0.52	0.67
DIP II	2.10	3.25	0.31	0.40
DIP III	1.93	2.87	0.22	0.36
DIP IV	2.41	2.78	0.30	0.33
DIP V	2.84	3.30	0.35	0.73
PIP II	1.90	2.08	0.32	0.36
PIP III	1.58	2.19	0.27	0.50
PIP IV	1.32	2.81	0.34	0.46
PIP V	2.77	3.44	0.26	0.31
CMC IV	2.06	2.07	0.69	0.70
CMC V	2.74	2.79	1.07	1.15
MP I	0.59	1.03	0.44	0.80
MP II	1.58	1.82	0.55	0.76
MP III	1.76	1.87	0.37	0.65
MP IV	0.87	2.57	0.24	0.51
MP V	1.67	1.81	0.46	0.79
CMC I	0.79	1.94	0.90	2.26
Average	1.84	2.45	0.45	0.69

model, of which 11 are single-axis and 6 are two-axis, for a total of 23 degrees of freedom (Figure 8). We assign no degrees of freedom to CMC II and III because it is well-known that these joints almost do no move [Austin 2005]; in our dataset, maximum motion was under 2 degrees. The configuration of the entire joint hierarchy is then specified via the vector θ of 23 joint angles. These angles can be driven using any suitable computer animation method, such as keyframe animation, motion capture or even motion control.

Our bone rig. We first describe our approach for 1-DOF joints. Consider one specific parent-child bone pair. During hand articulation, both the parent and child undergo rigid body motion. In this section, we conceptually undo the parent's transformation so that the parent is stationary in the world, and the child undergoes some rigid body motion relative to parent (Figure 9). All motions of the child bone in this section are to be understood in this way. Denote the angle of the 1-DOF rotation of the child bone by the scalar ϕ . Our MRI scan and its segmentation give us the rigid transformation (x_i, p_i) of the child bone in each pose i , for $i = 1, \dots, N$, where N is the number of acquired poses. Here, $x_i \in \mathbb{R}^3$ is the translation vector and $p_i \in \mathbb{R}^4$ is the unit quaternion. In pose i , the




Fig. 9. Transformation (x_i, R_i) is expressed in parent's coordinate system. It transforms the neutral child bone into pose i .

parent is stationary in the world, and the child undergoes some rigid body motion relative to parent (Figure 9). All motions of the child bone in this section are to be understood in this way. Denote the angle of the 1-DOF rotation of the child bone by the scalar ϕ . Our MRI scan and its segmentation give us the rigid transformation (x_i, p_i) of the child bone in each pose i , for $i = 1, \dots, N$, where N is the number of acquired poses. Here, $x_i \in \mathbb{R}^3$ is the translation vector and $p_i \in \mathbb{R}^4$ is the unit quaternion. In pose i , the

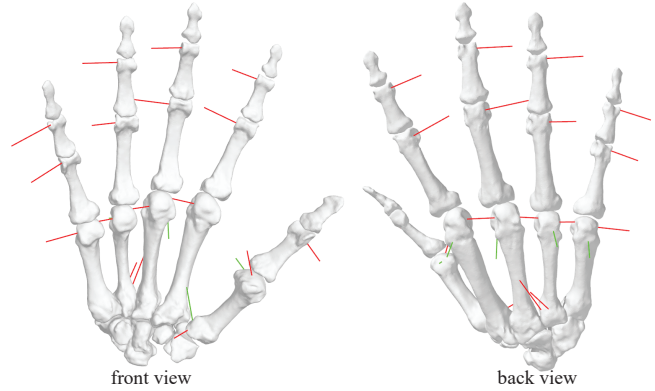


Fig. 10. **The rotation axes of our joints.** They are shown as red lines (for 1-DoF joints), and red and green lines (for 2-DoF joints).

child bone undergoes rigid transformation $X \mapsto R_i X + x_i$, where R_i is the rotation matrix corresponding to p_i , and $X \in \mathbb{R}^3$ is an arbitrary point on the child bone. Our task is to discover a unit rotation axis $a \in \mathbb{R}^3$, and angles $\phi_i \in \mathbb{R}$, such that p_i is as close as possible to a rotation by ϕ_i around axis a . This leads to a constrained optimization problem for a and angles ϕ_i :

$$\max_{a, \phi_1, \dots, \phi_N} \sum_{i=1}^N \text{dot}\left(p_i, \cos\left(\frac{\phi_i}{2}\right) + \sin\left(\frac{\phi_i}{2}\right)a\right). \quad (1)$$

s.t. $a^T a = 1$

Here $\text{dot}(p, q)$ denotes the dot product of quaternions p and q . We robustly solve this problem using the IPOPT constrained nonlinear optimizer [Wächter and Biegler 2006] with ADOL-C automatic differentiation [Walther and Griewank 2009].

After a and ϕ_i are known, we construct our rig as follows. Given an arbitrary angle ϕ , we model the rigid transformation of the child bone relative to its parent as

$$X \mapsto \hat{R}(\phi)R(\phi)(X - C) + C + \hat{x}(\phi), \quad (2)$$

where $R(\phi)$ is the rotation around axis a by angle ϕ , and \hat{R} and \hat{x} are the residual rotation and translation, respectively (to be determined). Here, $C \in \mathbb{R}^3$ is a center of rotation, which we determine using least squares so that $\hat{x}(\phi)$ is minimized, assuming $\hat{R} = I$,

$$\min_C \sum_{i=1}^N \|x_i + (R_i - I)C\|^2. \quad (3)$$

Finally, we compute the residual translations $\hat{x}(\phi_i) = x_i + (R_i - I)C$ and residual rotations $\hat{R}(\phi_i) = R_i R(\phi_i)^T$ at each sampled pose i . The only remaining task is to interpolate / extrapolate these residuals to arbitrary ϕ . We achieve this by representing \hat{R} using Euler angles (note that \hat{R} is a small rotation), and then interpolating / extrapolating each degree of freedom of translations and rotations by fitting a 1D quadratic function to samples $(\phi_i, \hat{x}_i(\phi_i))$ and $(\phi_i, \hat{R}_i(\phi_i))$, respectively. We initially employed natural cubic




Fig. 11. Fitting the residual \hat{x} .

splines, and then also polyharmonic splines; however, this resulted in a lot of spurious wiggles in the rig; hence, we settled for quadratic functions which gave smooth residuals. We extrapolate the functions using the tangents at each end of the samples. We experimentally observed that the residuals \hat{x}_i are small, on the order or smaller than the accuracy of our MRI scan (0.5 mm) (Figure 11), and the wiggles in \hat{x} are therefore mostly just experimental noise. Our conclusion is that bones in real hands, to a large extent, do rotate around parent bones around a fixed center C . However, the deviation away from the “fixed-rotation-center” model is still substantial. Table 1 analyzes the accuracy of our bone rig. It demonstrates that using our quadratically interpolated offsets \hat{R} and \hat{x} improves the accuracy of the rig. Even if one accepts fixed centers, one still needs our method to compute the centers and the rotation axes; otherwise suboptimal shapes occur (Figures 16, 20).

Our approach for two-dimensional joints is similar, except that we now have to find two orthogonal unit rotation axes a and b and a pair of joint angles ϕ_i and ψ_i . Our kinematic model is that we first perform a rotation around a by ϕ_i , followed by a rotation around the rotated axis b by ψ_i . Therefore, we maximize

$$\max_{a, b, \phi_1, \dots, \phi_N, \theta_1, \dots, \theta_N} \sum_{i=1}^N \text{dot}(p_i, r_i q_i), \text{ where} \quad (4)$$

$$\text{s.t. } a^T a = 1, b^T b = 1, a^T b = 0$$

$$q_i = \cos\left(\frac{\phi_i}{2}\right) + \sin\left(\frac{\phi_i}{2}\right)a, \quad r_i = \cos\left(\frac{\theta_i}{2}\right) + \sin\left(\frac{\theta_i}{2}\right)b'_i, \quad (5)$$

and b'_i is obtained by rotating vector b using q_i . As in the 1D case, this optimization is greatly facilitated by the ADOL-C automatic differentiation library. The computation of the entire rig took 1 minute total for all poses and the entire skeleton. The rotation axes are shown in Figure 10.

5 FEM SOFT TISSUE SIMULATION MODEL

Our soft tissue simulation model consists of bone geometry (Section 3), surrounded by an elastic soft tissue modeled as a tetrahedral mesh. The bone geometry is animated using our skeleton rig (Section 4). We now describe how we generate a stable tetrahedral mesh (Figure 12), assign its material properties, and how we constrain the tetrahedral mesh to the bones for FEM simulation.

5.1 Generation of the tet mesh outer surface

We start with the high-resolution (11M triangles) mesh of the hand in the neutral pose, obtained using optical scanning (Artec Spider) of the plastic neutral hand. Because the fingers are sufficiently spread in this pose, this mesh had no webbing or occlusion issues, and only tiny imperfections which we easily cleaned in Maya using smoothing. We then simplify this mesh to a smaller mesh (0.5M triangles) using MeshLab; this is done to make the rest of the pipeline more tractable. The average edge length in this mesh is 0.5mm, however the edge lengths are not uniform and some triangles are of bad quality. We then manually (using painting in Maya) separate the triangles in two groups: those on creases (where higher deformation precision is needed; both at the finger joints and on the palm), and those elsewhere (“non-crease”). We then constrain the crease triangles, and remesh the non-crease triangles using CGAL’s

isotropic_remeshing tool, with a 2.5mm target edge length. Next, we constrain the non-crease triangles and remesh the crease triangles using the same tool, at 0.5mm. We note that such a 2-stage remeshing process is necessary; otherwise, artefacts appear (Figure 13, bottom-left). This procedure gives us a good isotropic manifold triangle mesh $\hat{\mathcal{T}}_{\text{outer}}$ (31,986 triangles) of the hand’s neutral pose surface, with smaller edge lengths in the crease areas. We need to correct this mesh, as described next.

5.2 Resolving discrepancies due to subject’s small pose change in the mold prior to MRI scanning

Note that $\hat{\mathcal{T}}_{\text{outer}}$ encloses the surface of the plastic hand, which is the negative image of the MRI mold of the neutral pose. However, when the hand is placed into the rubber-like mold for MRI scanning, for ergonomics reasons, the hand slightly changes its pose as the subject settles it into the mold prior to scanning. Also, the rubber-like mold somewhat gently squeezes the hand, which slightly alters its volume. We resolve these discrepancies by constructing the hand surface geometry (the skin) from the MRI scan. This is feasible for the neutral shape because the fingers are well-separated. Note that due to limited MRI scanning resolution, this mesh has significantly lower quality than $\hat{\mathcal{T}}_{\text{outer}}$, and cannot serve directly as an input mesh for tetrahedral meshing. Instead, we align $\hat{\mathcal{T}}_{\text{outer}}$ onto the MRI-scanned surface mesh, using a nonlinear Iterative Closest Point (ICP) tool Wrap3 [Wrap3 2018]. Here, we set the parameters in Wrap3 so that the wrap is nearly rigid and only adjusts the pose, preserving the local surface detail in $\hat{\mathcal{T}}_{\text{outer}}$. This is a slight, but important adjustment, to compensate for subject’s ergonomic settling. The result is a triangle mesh $\mathcal{T}_{\text{outer}}$ of the surface of the neutral hand that has good quality and local surface detail, does not suffer from a loss of volume, and is positioned consistently with the MRI-scanned skeleton (see Figure 12, top). This mesh will serve as the outer mesh for our simulation tet mesh.

5.3 Generation of the tet mesh inner boundary

We now describe how we generate a triangle mesh $\mathcal{T}_{\text{inner}}$ which will serve as the inner boundary for our tet mesh. This mesh has to conform to the bones; however, it should *not* mesh the space in between the bones at the joints; otherwise, those tets will be immediately pinched as the joints articulate, leading to tet inversions and simulation instabilities. We start with the skeleton rig mesh \mathcal{B} (155,450 triangles; animated in Section 4), and use CGAL’s isotropic_remeshing function (at 2.5mm) to produce a mesh $\mathcal{B}_{\text{coarse}}$ with 10,408 triangles for the entire skeleton. This mesh cannot serve as $\mathcal{T}_{\text{inner}}$ because the mesh simplification and remeshing sometimes introduce collisions between the bone meshes. To remove collisions, we use the TetWild algorithm [Hu et al. 2018] to create a tetrahedral mesh of the union of the volumes enclosed by bones in $\mathcal{B}_{\text{coarse}}$. The surface of this tet mesh is our inner mesh $\mathcal{T}_{\text{inner}}$. It is a good-quality isotropic mesh (10,682 triangles) whose surface closely matches $\mathcal{B}_{\text{coarse}}$, except that $\mathcal{T}_{\text{inner}}$ joins (“welds”) bone meshes in the collision areas. Note that the vertices of $\mathcal{T}_{\text{inner}}$ are different to those of $\mathcal{B}_{\text{coarse}}$. Here, we considered an alternative: compute mesh boolean union of the bone meshes in $\mathcal{B}_{\text{coarse}}$, and then use CGAL’s isotropic_remeshing tool; however, this destroyed some




Fig. 12. **Simulation tet mesh.** Top row: outer and inner tet surface mesh. In $\mathcal{T}_{\text{outer}}$, note the greater resolution at the folds. Bottom: three cutaway views of the FEM mesh.

geometric detail and introduced new collisions, so we abandoned this approach.

Next, we execute constrained Delaunay tetrahedralization with refinement (we use TetGen [Hang Si 2011]) on input $\mathcal{T}_{\text{outer}} \cup \mathcal{T}_{\text{inner}}$. We set TetGen’s parameters so that $\mathcal{T}_{\text{outer}}$ and $\mathcal{T}_{\text{inner}}$ are not refined, but the interior tetrahedra are refined. The running time of TetGen was under 1 second. The resulting tet mesh is not our final tet mesh because we need to remove tetrahedra located at the joints between two bones. For each bone, we manually select the triangles of $\mathcal{B}_{\text{coarse}}$ on the proximal epiphysis (the head closer to the parent bone), roughly corresponding to the region of articular cartilage (smooth white tissue where bone forms a joint to its parent bone); call this set $\mathcal{S}_{\text{child}}$. Next, we execute a procedure that we call “sweeping”: for each vertex of a triangle from $\mathcal{S}_{\text{child}}$, find the closest triangle t on the parent bone. If the distance is less than a threshold (we use 2.5mm), and the normals are approximately opposite (dot product is smaller than -0.7), we add t to the swept set $\mathcal{S}_{\text{parent}}$. We then form the convex hull (using CGAL’s `convex_hull_3`) of $\mathcal{S}_{\text{child}} \cup \mathcal{S}_{\text{parent}}$. We then remove from the tet mesh every tet whose center is inside the convex hull (see Figure 13, top-middle). This gives us our simulation tet mesh (101,941 tets, 29,583 vertices). Note that this mesh has an external surface and an internal surface (i.e., there are internal voids corresponding to the bones).

5.4 Constraining the tet mesh to the bones

Our tet mesh is attached to the bones using soft constraints. We constrain a subset of the vertices of the internal tet mesh surface. Constrained vertices are determined by sweeping (in the sense of Section 5.3) the child bone against the parent bone, across the entire range of motion of the joint. Note that hand joints may have 1-DOF or 2-DOFs. Our sweep resolution is 1 degree for each joint degree

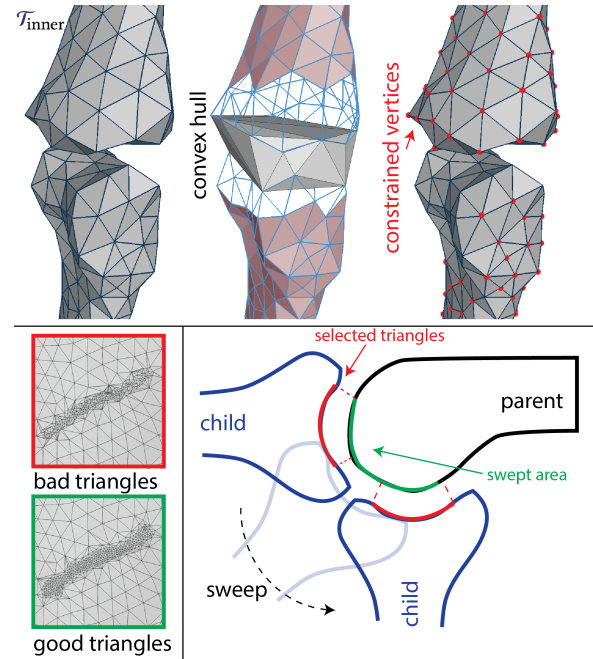


Fig. 13. **Tet mesh creation.** Top: mesh $\mathcal{T}_{\text{inner}}$; removing tets at the joints; tet mesh constrained vertices. Bottom-left: the two-step remeshing process at the folds produces good triangles, and is necessary: under a single-step process, bad triangles appear. Bottom-right: illustration of sweeping in 2D.

of freedom, both for 1-DOF and 2-DOF joints. For each value of the joint degrees of freedom, we compute $\mathcal{S}_{\text{parent}}$. We form the union \mathcal{U} of $\mathcal{S}_{\text{parent}}$ across the entire range of motion (Figure 13, bottom-right). The sweeping takes 10 minutes total for the entire skeleton. We then traverse all tet mesh vertices that are on the internal surface. If the closest triangle on $\mathcal{B}_{\text{coarse}}$ is further than 0.01mm, or is in \mathcal{U} , we do not constrain this vertex, otherwise we constrain it to the closest location on $\mathcal{B}_{\text{coarse}}$, using a spring of stiffness 10 N/mm. This avoids constraining \mathcal{U} , and the region on the tet mesh adjacent to the removed tets in between the bones. The constraint springs are integrated using implicit integration by computing their force gradients, for simulation stability. We compute the motion of the coarse bone mesh $\mathcal{B}_{\text{coarse}}$ by skinning it to the output mesh \mathcal{B} of our bone rig; each vertex of $\mathcal{B}_{\text{coarse}}$ is skinned to the closest bone in \mathcal{B} . Because nails are rigid, we need to also constrain nail vertices. We do this manually, by painting the nail triangles in Maya. The nail vertices are constrained to follow the transformation of the distal phalanx finger bones. In addition to the constraints, we also apply collision response between the bones and the soft tissue, and soft tissue self-collision response. For the former, we use the technique presented in [Barbič and James 2008], whereas for the latter, we detect pairs of non-neighboring overlapping triangles and push them apart using penalty springs. In one of our examples (connect-all-poses), we observed small spurious wrinkles on the skin at the dorsal (i.e., back) side of the PIP joint, when the PIP joint is bent to its extreme. We attribute this to the fact that the skin is pre-folded in that region in the neutral configuration, and “unfolds” as PIP articulates; pre-folding is not modeled by our system. We remedy




Fig. 14. **Young's modulus of various human musculoskeletal tissues.** Note the unit $kPa = 1000N/m^2$. From [Discher et al. 2009]. Reprinted with permission from AAAS.

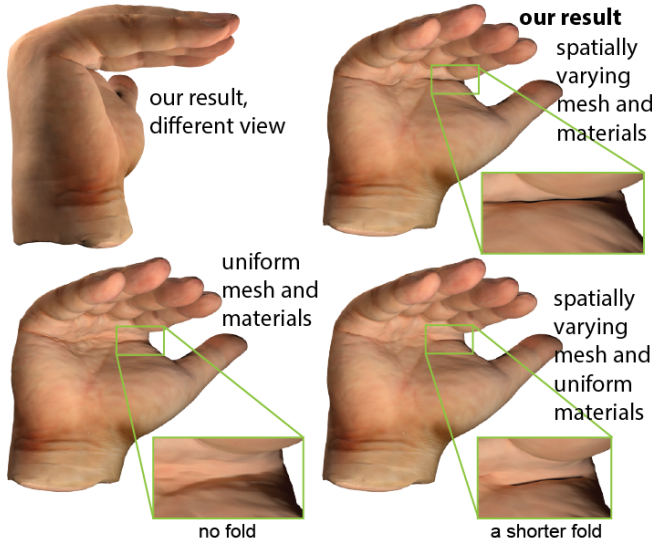


Fig. 15. **Spatially varying mesh resolution and materials.** Fold formation is improved by increasing the tet mesh resolution under the finger joints (via \mathcal{T}_{outer} ; Sections 5.1 and 5.2), and by making the material in the same region 6x softer (Section 5.5). Our method produces a quality horizontal fold across the entire length of the palm. Other methods only produce a partial fold, or no fold at all.

this issue by post-processing the small local region using the “Delta Mush” deformer [Mancewicz et al. 2014].

5.5 Material properties

For elastic material properties, we use the Saint-Venant Kirchhoff material. We add a volume preservation term which prevents large compressions and successfully guards against collapse. Specifically, we model volume preservation by adding a term $\lambda(J - 1)^2/2$ to the elastic strain energy density function. Here, λ is the first Lamé parameter, and J is the local volume growth factor (product of the three principal stretches [Irving et al. 2004]). The combination of StVK and volume preservation produced organic hand shapes that visually match photographs of the subject's hand (Figure 19). We model our tissue as a nearly incompressible isotropic hyperelastic solid, and therefore use a Poisson's ratio of 0.495. Although biological tissues are not completely incompressible, incompressibility and isotropy is appropriate for many biological materials [Wex et al. 2015], because human tissue largely consists of an incompressible liquid (water). Our choice of Young's modulus is motivated by experimental measurements available in literature, such as Figure 14 reproduced from [Discher et al. 2009], which the authors assembled from various experimental studies. As can be seen, fat ($2500-4000 N/m^2$) is softer than muscles ($9000-15000 N/m^2$). Because we




Fig. 16. **Anatomical bone rig stabilizes FEM soft tissue simulations.** Left: non-anatomical rig, created by pivoting bones around the center of the parent bone's head. Bones pinch elastic material, causing artefacts. Right: our rig produces stable soft tissue simulations. Both results use FEM.




Fig. 17. **Comparison of FEM simulation to skinning.** Skinning weights were computed using Maya's geodesic voxel method. Other Maya skinning weight methods give even worse results. Both methods use our bone rig.

do not model fat and muscles as separate objects, we set our Young's modulus to a single value of $12000 N/m^2$. In locations where there are visible hand lines, such as on the palm and at finger joints, we decrease Young's modulus to $2000 N/m^2$, and Poisson's ratio to 0.45, which helps with crease formation. Because there are no muscles on the dorsal side of the hand, we decrease Young's modulus in that region to $6000 N/m^2$, including on the back side of the fingers. For mass density, we choose the density of water ($1000 kg/m^3$). This is because the reported densities of muscles and fat are $1060 kg/m^3$ and $920 kg/m^3$, respectively [Farvid et al. 2005].

5.6 Embedding of surface triangle meshes into tet mesh

We deform any of our surface triangle meshes by embedding it into the tet simulation mesh. Sometimes, some vertices of these surface triangle meshes in the neutral pose may be slightly outside of the tet mesh due to coarsening or discretization errors; in which case we simply deform the vertex using the closest tetrahedron. There are two surface meshes that we deform in this way and that appear in our rendering results: the high-resolution mesh obtained by scanning the neutral plastic hand using Artec Spider (11M triangles), and a low-resolution version with 15,232 triangles, obtained by simplifying the high-resolution mesh using MeshLab.

6 RESULTS

We used a 3 Tesla GE MRI scanner with a PD CUBE (3D fast spin echo) sequence with the following parameters: NEX (number of excitations) 1; FA (flip angle) 90 degrees; TR (time to repetition) 1500 ms; TE (time to echo) 25.93 ms; slice thickness 1 mm; slice spacing 0.5 mm; matrix 256×256 pixels; and FOV (field of view) 28.3×25.6 cm. We scanned, processed and simulated two subjects: a male (age: late 20s; Figure 3), and a female (age: late 40s; Figure 18).




Fig. 18. **The opposition of the thumb for our female subject (late 40s).** We show the bone and FEM geometry of 3 animation frames.

We tested our bone rig and FEM simulation on 4 motion sequences. Two sequences were recorded using a LeapMotion system [LeapMotion 2017], one was created using Inverse Kinematics (IK) to perform the opposition of the thumb to the other fingers, and one connects all of our scanned poses into one continuous sequence, also using IK. We implemented a simple inverse kinematics engine, using the ADOL-C symbolic differentiation library. In our IK, we can constrain both positions and orientations of the end effectors; and we did so for the fingers subject to biological limits. The running time of our FEM simulations is approximately 1 sec / timestep, on an Intel i7 6950X PC (manufactured in 2016, 1 processors x 10 cores at 3.00 GHz), 128GB of RAM. We are performing 60 timesteps per frame, at 25 FPS; and therefore, the computational cost is 1 minute per frame, or 25 minutes per second of motion. We do not attempt to simulate accurate tissue dynamics, only static shapes under the animated bone meshes. Hence mass, damping and timestep properties only matter for simulation stability. The four sequences (recorded motions 1, 2; thumb opposition, connect-all-poses) are 13, 16, 10 and 16 seconds long, and took 5.4, 6.7, 4.2 and 6.7 hours to compute, respectively.

Our pipeline produces a quality un-colored surface mesh consistent with internal anatomy. In order to add a color texture, we took photos of the subject’s hand from approximately 60 different angles using a DSLR camera (Canon EOS 7). The photos were taken rapidly while the subject held the hand firmly against a table. The subject then reversed the hand and another set of 60 photos was taken. We then constructed two separate textured surfaces from these photographs using the Photoscan software [Agisoft 2018]. We then separately wrapped each of these two surfaces onto our un-colored low-resolution surface mesh, using Wrap3. At the seams, the surfaces and colors overlap because each of the two half-scans covers slightly more than one half of the hand. We adjusted and smoothed the seams using Photoshop and Maya. We then transferred the texture onto the high-resolution mesh, using Maya. This

produced a good quality color texture map for both low-resolution and high-resolution surface meshes, which we then use for rendering (Figure 20). In this project, we discovered that it is actually not easy to obtain a complete hand static model with a high-resolution geometry and texture. Our Photoscan approach benefited from the fact that we already obtained high-precision geometry by scanning the (perfectly still) plastic hand using Artec Spider, and only needed to add color texture. Before using Photoscan, we consulted a local professional photogrammetry studio who attempted to create both geometry and color texture. They failed to produce a good scan, despite taking over 60 simultaneous photographs of the subject’s hand with high-quality cameras from multiple angles. Because a live hand is moving, Artec Spider also failed to simultaneously create geometry and color texture, despite repeated scanning attempts.

We compare our anatomy-driven bone rig to a non-anatomical bone rig created in Maya, by placing the joint pivots at the end of each bone, and then running FEM soft tissue simulation. Our rig clearly outperforms the non-anatomical rig, where bone collisions artificially squeeze the elastic material, causing instabilities in FEM simulations (Figure 16). We also compare our FEM method to skinning (Figure 17). Both FEM and skinning used our bone rig. The comparison shows that FEM produces higher quality shapes. Figure 15 demonstrates that our spatially varying materials enable the generation of quality folds near the finger joints. In Figure 19, we compare our FEM result to photographs of the hand of the same subject. Figure 20 demonstrates that our bone rig also improves the quality of skinning (i.e., without using FEM). Table 1 analyzed the accuracy of our bone rig.

7 CONCLUSION

We demonstrated how to acquire highly accurate human skeleton geometry in multiple poses using MRI. We achieved this using a novel molding procedure to stabilize hands during MRI scanning. We registered all poses to the same bones mesh connectivity, and built a skeleton mesh kinematic model (“skeleton rig”) to interpolate or extrapolate the acquired pose to the entire hand range of motion. We demonstrated that our skeleton rig can be used to drive soft-tissue FEM simulation to produce anatomically plausible organic hand shapes that qualitatively match photographs. Our accurate hand model can potentially benefit virtual hands in games / film / VR, robotic hands, grasping, and medical education, such as visualizations of internal hand anatomy motion.

We performed 12 scans for each subject, which is sufficient to demonstrate stable and precise common hand motions, including opposition of the thumb to all the other four fingers. Accuracy would be improved with more scans. An interesting research questions is how to automatically select the next best pose to scan, to maximize coverage of the hand’s range of motion. Some complex hand poses are challenging for our technique, for example, all fingertips touching at a single point. This is because we cut our molds in two pieces manually using a knife. It would be interesting to explore how to cut the molds into 3 or more pieces, to improve the ergonomics of the insertion of the hand into the mold. Although we produced and scanned 12 plastic shapes, we only used the neutral plastic shape in our simulation pipeline; the other 11 serve to evaluate the accuracy of our FEM rig (Figure 3). It is challenging to register the scanned

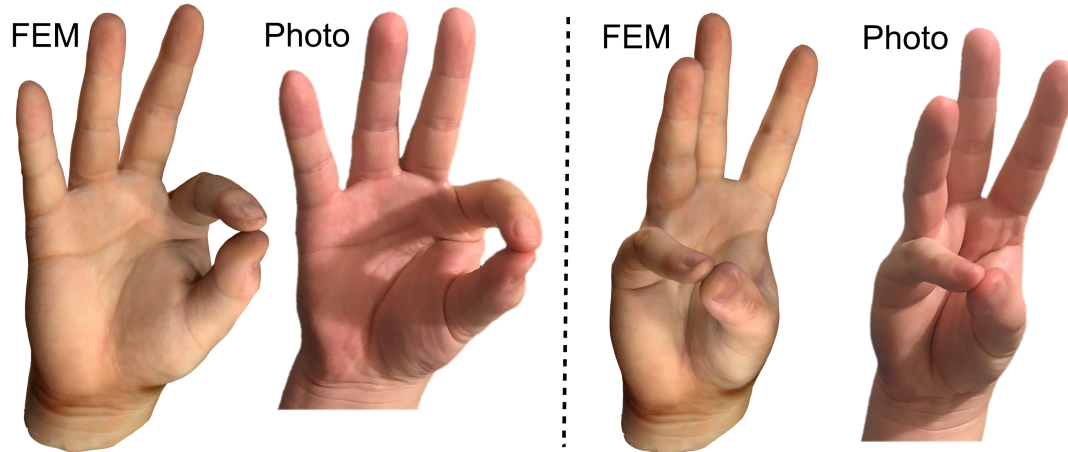


Fig. 19. **Comparison of our FEM results to photographs**, on two “in-between” poses (i.e., not a part of the acquired 12 pose-dataset). The FEM results are rendered using our high-resolution mesh (11M triangles), with subsurface scattering; individual pores are visible as geometric features. Image is zoomable. Generally, our FEM method produces a good qualitative visual match to the photographs: finger and palm look “organic”, and the two main palmar lines are in the correct location, and are produced by FEM as geometric mesh features (folds), as opposed to only a texture map illusion. This figure also demonstrates the limitations of our method. We do not model the wrinkling of the skin, which is absent in the FEM result, but visible in photographs. We also do not simulate muscle activation; and hence the soft tissue at the root of the thumb is flatter in FEM vs in photographs.



Fig. 20. **Our bone rig also improves skinning**. Left: skinning with a non-anatomical bone rig, created manually by placing the joints at the center of the parent bone’s bonehead. Right: skinning with our anatomical rig. No FEM simulation was used in this example, only skinning. It can be seen that skinning works better when it uses our anatomical bone rig (see the dent at the joint). Both results use Maya’s “geodesic voxel” skinning weights.

plastic meshes to consistent complete mesh connectivity because many of these optical scans are incomplete due to palm / finger occlusions, and webbing of fingers in close proximity.

We only acquired bone geometry; muscles, tendons and subcutaneous fat are left for future work. While our FEM simulations (driven by our skeleton rig) produce anatomically plausible shapes, accuracy would be improved by modeling muscles and subcutaneous fat. These structures are available in our MRI scans and could be in principle segmented and animated; new research techniques will likely be required to do so in a stable and accurate way. We use spatially-varying tissue elastic properties to improve the generation of folds at the joints. Results could be further improved by optimizing spatially-varying material properties. Physically based grasping would also benefit from good material properties. We do not compute the joint hierarchy motion; but instead rely on the animator to provide us with the animations of the bone’s joints. Such an approach is standard in industry today [Tissue 2013], but requires the animator to be careful not to cause self-collisions. A better, more principled approach would be for motion to originate from hand muscle activations, with bones and the rest of the anatomy animating based on a physically based simulation. This poses interesting neuromechanical control problems, due to the well-known muscle redundancy,

and the specific nature of hands which are mostly controlled by long muscle tendons originating in the human arm. The understanding of motion of internal hand anatomy as proposed here represents one step closer to lifelike robots with hands that mimic biological hands. Researchers at Columbia Univ. recently demonstrated how to make artificial muscles [Miriyeve et al. 2017]. With recent advances in 3D printing technology, our work may contribute to the physical replication of anatomically realistic hands and improved hand prosthetics, in the not so distant future.

ACKNOWLEDGMENTS

This research was sponsored in part by NSF (IIS-1422869), USC Annenberg Fellowship to Bohan Wang, Bosch Research and Adobe Research. We would like to thank the anonymous shepherd for outstanding shepherding work.

REFERENCES

- Agisoft. 2018. Photoscan, <http://www.agisoft.com>.
- AljaSafe. 2018. SmoothOn Inc. www.smooth-on.com.
- Amira. 2018. Amira for Life Sciences. <https://www.fei.com/software/amira-3d-for-life-sciences/>.
- Artec3D. 2018. Spider Scanner, <http://www.artec3d.com>.
- Noelle M. Austin. 2005. *Chapter 9: The Wrist and Hand Complex*. In *Levangie, Pamela K.; Norkin, Cynthia C. Joint Structure and Function: A Comprehensive Analysis* (4th ed.). F. A. Davis Company.
- J. Barbič and D. L. James. 2008. Six-DoF haptic rendering of contact between geometrically complex reduced deformable models. *IEEE Trans. on Haptics* 1, 1 (2008), 39–52.
- Pierre-Yves Baudin, Noura Azzabou, Pierre G Carlier, and Nikos Paragios. 2012. Automatic skeletal muscle segmentation through random walks and graph-based seed placement. In *IEEE Int. Symp. on Biomedical Imaging (ISBI)*. 1036–1039.
- Paul J. Besl and N.D. McKay. 1992. A Method for Registration of 3-D Shapes. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 14, 2 (1992), 239–256.
- Jean-Daniel Boissonnat and Steve Oudot. 2005. Provably good sampling and meshing of surfaces. *Graphical Models* 67, 5 (2005), 405 – 451.
- S. Capell, M. Burkhart, B. Curless, T. Duchamp, and Z. Popović. 2005. Physically Based Rigging for Deformable Characters. In *Symp. on Computer Animation (SCA)*. 301–310.
- Paolo Cignoni, Marco Callieri, Massimiliano Corsini, Matteo Dellepiane, Fabio Ganovelli, and Guido Ranzuglia. 2008. MeshLab: an Open-Source Mesh Processing Tool.

- In *Eurographics Italian Chapter Conference*.
 CyberGlove Systems. 2017. CyberGrasp. <http://www.cyberglovesystems.com/cyberg rasp>.
 Mary F Dempsey, Barrie Condon, and Donald M Hadley. 2002. MRI safety review. In *Seminars in Ultrasound, CT and MRI*, Vol. 23. Elsevier, 392–401.
 C. M. Deniz, S. Xiang, S. Hallyburton, A. Welbeck, S. Honig, K. Cho, and G. Chang. 2017. Segmentation of the Proximal Femur from MR Images using Deep Convolutional Neural Networks. *arXiv preprint arXiv:1704.06176* (2017).
 D. E. Discher, D. J. Mooney, and P. W. Zandstra. 2009. Growth factors, matrices, and forces combine and control stem cells. *Science* 324, 5935 (2009), 1673–1677.
 James S Duncan and Nicholas Ayache. 2000. Medical image analysis: Progress over two decades and the challenges ahead. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 22, 1 (2000), 85–106.
 M. S. Farvid, T. W. K. Ng, D. C. Chan, P. H. R. Barrett, and G. F. Watts. 2005. Association of adiponectin and resistin with adipose tissue compartments, insulin resistance and dyslipidaemia. *Diabetes, Obesity and Metabolism* 7, 4 (2005), 406–413.
 A. Fenster and D. B. Downey. 1996. 3-D ultrasound imaging: a review. *IEEE Engineering in Medicine and Biology Magazine* 15, 6 (1996), 41–51.
 Carlos Garre, Fernando Hernández, Antonio Gracia, and Miguel A Otaduy. 2011. Interactive simulation of a deformable hand for haptic rendering. In *IEEE World Haptics Conference (WHC)*. IEEE, 239–244.
 Benjamin Gilles and Nadia Magnenat-Thalmann. 2010. Musculoskeletal MRI segmentation using multi-resolution simplex meshes with medial representations. *Medical image analysis* 14, 3 (2010), 291–302.
 Leo Grady. 2006. Random walks for image segmentation. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 28, 11 (2006), 1768–1783.
 Agneta Gustus and Patrick van der Smagt. 2016. Evaluation of joint type modelling in the human hand. *Journal of Biomechanics* 49, 13 (2016), 3097 – 3100.
 Shangchen Han, Beibei Liu, Robert Wang, Yuting Ye, Christopher D Twigg, and Kenrick Kin. 2018. Online optical marker-based hand tracking with deep labels. *ACM Transactions on Graphics (SIGGRAPH 2018)* 37, 4 (2018), 166.
 Hang Si. 2011. TetGen: A Quality Tetrahedral Mesh Generator and a 3D Delaunay Triangulator.
 Yixin Hu, Qingnan Zhou, Xifeng Gao, Alec Jacobson, Denis Zorin, and Daniele Panozzo. 2018. Tetrahedral Meshing in the Wild. *ACM Trans. on Graphics (SIGGRAPH 2018)* 37, 4 (2018), 60:1–60:14.
 G. Irving, J. Teran, and R. Fedkiw. 2004. Invertible Finite Elements for Robust Simulation of Large Deformation. In *Symp. on Computer Animation (SCA)*. 131–140.
 ITK-SNAP. 2018. ITK-SNAP. <http://www.itk-snap.org/pmwiki/pmwiki.php>.
 Alec Jacobson, Zhigang Deng, Ladislav Kavan, and JP Lewis. 2014. Skinning: Real-time Shape Deformation. In *ACM SIGGRAPH 2014 Courses*.
 A.I. Kapandji. 2009. *The physiology of the joints, 6th Edition, Vol. 1: The Upper Limb*. Elsevier Exclusive.
 L. Kavan, S. Collins, J. Zara, and C. O’Sullivan. 2008. Geometric Skinning with Approximate Dual Quaternions Blending. *ACM Trans. on Graphics* 27, 4 (2008).
 Baris Kayalibay, Grady Jensen, and Patrick van der Smagt. 2017. CNN-based segmentation of medical imaging data. *arXiv preprint arXiv:1701.03056* (2017).
 Junggon Kim and Nancy S Pollard. 2011. Fast simulation of skeleton-driven deformable body characters. *ACM Trans. on Graphics (TOG)* 30, 5 (2011), 121.
 Jonathan P King, Dominik Bauer, Cornelia Schlangenau, Kai-Hung Chang, Daniele Moro, Nancy Pollard, and Stelian Coros. 2018. Design, Fabrication, and Evaluation of Tendon-Driven Multi-Fingered Foam Hands. In *IEEE-RAS Int. Conf. on Humanoid Robots (Humanoids)*. 1–9.
 Paul G. Kry, Doug L. James, and Dinesh K. Pai. 2002. EigenSkin: Real Time Large Deformation Character Skinning in Hardware. In *Proc. of the Symp. on Comp. Animation 2002*. 153–160.
 Tsuneya Kurihara and Natsuki Miyata. 2004. Modeling deformable human hands from medical images. In *Symp. on Computer Animation (SCA)*. 355–363.
 LeapMotion. 2017. <https://www.leapmotion.com>.
 S. H. Lee, E. Sifakis, and D. Terzopoulos. 2009. Comprehensive Biomechanical Modeling and Simulation of the Upper Body. *ACM Trans. on Graphics* 28, 4 (2009), 99:1–99:17.
 J. P. Lewis, Matt Corder, and Nickson Fong. 2000. Pose Space Deformations: A Unified Approach to Shape Interpolation and Skeleton-Driven Deformation. In *Proc. of ACM SIGGRAPH 2000*. 165–172.
 Duo Li, Shinjiro Sueda, Debanga R Neog, and Dinesh K Pai. 2013. Thin Skin Elastodynamics. *ACM Trans. Graph. (Proc. SIGGRAPH)* 32, 4 (2013), 49:1–49:9.
 Libin Liu, KangKang Yin, Bin Wang, and Baining Guo. 2013. Simulation and control of skeleton-driven soft body characters. *ACM Trans. on Graphics (SIGGRAPH Asia 2013)* 32, 6 (2013), 215.
 N. Magnenat-Thalmann, R. Laperrère, and D. Thalmann. 1988. Joint-dependent local deformations for hand animation and object grasping. In *Proc. of Graphics Interface*.
 Joe Mancewicz, Matt L. Derksen, Hans Rijpkema, and Cyrus A. Wilson. 2014. Delta Mush: Smoothing Deformations While Preserving Detail. In *Proceedings of the Fourth Symposium on Digital Production (DigiPro ’14)*. 7–11.
 G. Marai, D. Laidlaw, J. Coburn, M. Upal, and J. Crisco. 2003. A 3D method for segmenting and registering carpal bones from CT volume images. In *Proc. of Annual Meeting of the American Society of Biomechanics*.
 A. McAdams, Y. Zhu, A. Selle, M. Empey, R. Tamstorf, J. Teran, and E. Sifakis. 2011. Efficient elasticity for character skinning with contact and collisions. *ACM Trans. on Graphics (SIGGRAPH 2011)* 30, 4 (2011).
 Tim McInerney and Demetri Terzopoulos. 2008. Deformable models in medical image analysis: a survey. *Medical Image Analysis* 1, 2 (2008), 91–108.
 Fernand Meyer. 1992. Color image segmentation. In *International Conf. on Image Processing and its Applications*. IET, 303–306.
 Aslan Miriyev, Kenneth Stack, and Hod Lipson. 2017. Soft material for soft actuators. *Nature Communications* 8, 596 (2017).
 N. Miyata, M. Kouch, M. Mochimaru, and T. Kurihara. 2005. Finger joint kinematics from MR images. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*. 2750–2755.
 NimbleVR. 2012. <http://nimblevr.com>.
 RadiologyInfo. 2018. Radiation Dose in X-Ray and CT Exams. <https://www.radiologyinfo.org/en/pdf/safety-xray.pdf>.
 Taehyun Rhee, J.P. Lewis, and Ulrich Neumann. 2006. Real-Time Weighted Pose-Space Deformation on the GPU. In *Proc. of Eurographics 2006*, Vol. 25.
 Javier Romero, Dimitrios Tzionas, and Michael J. Black. 2017. Embodied Hands: Modeling and Capturing Hands and Bodies Together. *ACM Trans. on Graphics (SIGGRAPH Asia 2017)* 36, 6 (2017), 245:1–245:17.
 Alexandru Rusu. 2011. *Segmentation of bone structures in Magnetic Resonance Images (MRI) for human hand skeletal kinematics modelling*. Master’s thesis. German Aerospace Center.
 Prashant Sachdeva, Shinjiro Sueda, Susanne Bradley, Mikhail Fain, and Dinesh K. Pai. 2015. Biomechanical Simulation and Control of Hands and Tendinous Systems. *ACM Trans. Graph.* 34, 4 (2015), 42:1–42:10.
 Cornelia Schlangenau, Dominik Bauer, Kai-Hung Chang, Jonathan P King, Daniele Moro, Stelian Coros, and Nancy Pollard. 2018. Control of tendon-driven soft foam robot hands. In *IEEE-RAS Int. Conf. on Humanoid Robots (Humanoids)*. 1–7.
 Jérôme Schmid, Jinman Kim, and Nadia Magnenat-Thalmann. 2011. Robust statistical shape models for MRI bone segmentation in presence of small field of view. *Medical image analysis* 15, 1 (2011), 155–168.
 Jérôme Schmid and Nadia Magnenat-Thalmann. 2008. MRI bone segmentation using deformable models and shape priors. In *Int. Conf. on Medical Image Computing and Computer-Assisted Intervention*. 119–126.
 Breannan Smith, Fernando De Goes, and Theodore Kim. 2018. Stable Neo-Hookean Flesh Simulation. *ACM Trans. Graph.* 37, 2 (2018), 12:1–12:15.
 Ole Vegard Solberg, Frank Lindseth, Hans Torp, Richard E. Blake, and Toril A. Nagelhus Hernes. 2007. Freehand 3D Ultrasound Reconstruction Algorithms: A Review. *Ultrasound in Medicine and Biology* 33, 7 (2007), 991 – 1009.
 Georg Stillfried. 2015. *Kinematic modelling of the human hand for robotics*. Ph.D. Dissertation. Technische Universität München.
 Shinjiro Sueda, Andrew Kaufman, and Dinesh K. Pai. 2008. Musculotendon Simulation for Hand Animation. *ACM Trans. Graph. (Proc. SIGGRAPH)* 27, 3 (2008).
 Tissue. 2013. Weta Digital: Tissue Muscle and Fat Simulation System.
 Rodolphe Vaillant, G ael Guennebaud, Loic Barthe, Brian Wyvill, and Marie-Paule Cani. 2014. Robust Iso-surface Tracking for Interactive Character Skinning. *ACM Trans. on Graphics (SIGGRAPH Asia 2014)* 33, 6 (2014), 189:1–189:11.
 Patrick van der Smagt and Georg Stillfried. 2008. Using MRI data to compute a hand kinematic model. In *Conf. on Motion and Vibration Control (MOVIC)*.
 VTK. 2018. VTK. <https://www.vtk.org/>.
 Andreas W achter and Lorenz T. Biegler. 2006. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming* 106, 1 (01 Mar 2006), 25–57.
 Andrea Walther and Andreas Griewank. 2009. Getting Started with ADOL-C. In *Combinatorial scientific computing*. 181–202.
 R. Y. Wang, S. Paris, and J. Popovi c. 2011. 6D Hands: Markerless Hand Tracking for Computer Aided Design. In *ACM User Interface Software and Technology (UIST)*.
 Robert Y. Wang and Jovan Popovi c. 2009. Real-time hand-tracking with a color glove. *ACM Trans. on Graphics (SIGGRAPH 2009)* 28, 3 (2009), 63:1–63:8.
 Robert E. Watson. 2015. Lessons Learned from MRI Safety Events. *Current Radiology Reports* 3, 10 (12 Aug 2015), 37.
 C. Wex, S. Arndt, A. Stoll, C. Bruns, and Y. Kupriyanova. 2015. Isotropic incompressible hyperelastic models for modelling the mechanical behaviour of biological tissues: a review. *Biomedical Engineering / Biomedizinische Technik* 60, 6 (2015), 577–592.
 Nkenge Wheatland, Yingying Wang, Huaiguang Song, Michael Neff, Victor Zordan, and Sophie J org. 2015. State of the art in hand and finger modeling and animation. In *Computer Graphics Forum*, Vol. 34. 735–760.
 Wrap3. 2018. Nonlinear Iterative Closest Point mesh registration software. <https://www.russian3dscanner.com>.
 Shanxin Yuan, Qi Ye, Bj orn Stenger, Siddhant Jain, and Tae-Kyun Kim. 2017. Bighand2. 2M benchmark: Hand pose dataset and state of the art analysis. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*. 2605–2613.
 Zygote. 2016. Zygote body. <http://www.zygotebody.com>.