A Parameterized Family of Anatomically Accurate Human Upper-Body Musculoskeletal Models for Dynamic Simulation & Control

Samir Menon, Toki Migimatsu, and Oussama Khatib

***Abstract*— Modeling human motion requires an accurate specification of musculoskeletal physiology, yet there exists no method to quantify the modeling accuracy required, or to predict the effect of modeling errors on subsequent analyses. Quantifying how inaccuracies in physiology, kinematics, or dynamics affect the study of human motion is a challenge that must be solved before we can construct robust generative models of human motor control at appropriate levels of detail. In this paper, we overcome two fundamental problems in characterizing the effect of model accuracy: the lack of ground truth about the arm's musculoskeletal kinematics, and the inability to systematically vary modeling accuracy. To do so, we developed a family of upper-body musculoskeletal models for a live human individual where modeled musculature was parameterized by decomposing volumetric muscles into fiber-groups of varying diameter and geometric complexity. The family of models thus obtained offer an unprecedented level of detail, and enable empirical comparisons of human motion analysis results across varying levels of anatomical accuracy and geometry. This sets the stage for large-scale studies of human motion that connect high level behavior to low level musculoskeletal dynamics, with applications in robotics, biomechanics, and human motor control.**

I. Introduction

Efforts to model and characterize human motion are presently limited by a lack of automated methods to create subject-specific models, and a lack of systematic methods that can quantify a given musculoskeletal model's robustness with respect to a given analysis. Moreover, human-centered analyses are diverse and include a variety of robotics research areas, including human-robot interaction [1], ergonomics [2], robot design [3], motion reconstruction [4], [5], computer vision [6], [7] and human-robot safety [8], [9], [10]. This large set of use cases has led to fragmentation in efforts across groups, which have very different goals. A set of openly available models parameterized by variables of shared interest promises to synchronize disparate efforts. Methods and tools to automate model creation and compare model performance thus promise to impact a large and rapidly growing section of research.

One strategy adopted by human modelers, in the face of uncertainty about the required level of modeling detail, is to pursue an ever-increasing level of detail. The underlying assumption is that the most detailed model that can be feasibly constructed will be sufficiently robust to errors in model parameter estimates and motion measurements. While intuitive, this is an unverifiable assumption in the absence of ground truth about the modeled system—an assumption that warrants scrutiny. Error checks, when included, primarily focus on high-level physical performance metrics [11], or robustness to parameter tuning and numerical errors in analysis methods [12], [13] while keeping the model fixed. We propose, in addition, that analyses should identify the gradient of errors with respect to modeling detail.

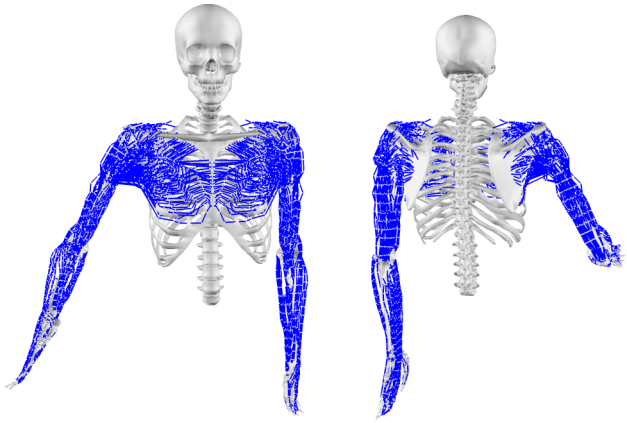

Fig. 1. **Subject-specific human musculoskeletal modeling.** An anatomically precise subject-specific musculoskeletal arm model extracted from magnetic resonance imaging data.

Here, we construct an anatomically accurate upper-body musculoskeletal model using magnetic resonance imaging (MRI) "*ground truth*" data (Fig. 1), vary it parametrically to create a family of models, and propose an analysis framework to empirically analyze how model accuracy influences human musculoskeletal simulation and control (Fig. 2). The primary challenges we overcame while doing so were to: *(i)* obtain high resolution MRI data for limbs in the upper-body; *(ii)* co-register the MRI scans of different limbs into a composite arm model; *(iii)* segment MRI scans into volumetric muscles and skeletons; and *(iv)* automatically generate families of musculoskeletal systems from the segmented MRI scans. Comparing our models with existing human musculoskeletal models, we found a large gain in anatomical accuracy at the shoulder and elbow joints. We also found that commonly used affine scaling methods have the potential to introduce large errors in human motion analyses when the musculoskeletal system of the canonical model does not match an individual subject. This makes a strong case for developing detailed subject-specific models, or at least developing a family of models large enough that it is feasible

*This work was supported by National Science Foundation National Robotics Initiative grant (IIS-1427396, O. Khatib and R. Bajcsy), and by a Stanford Center for Neurobiological Imaging Seed Grant (S. Menon).

Stanford Robotics Laboratory, Department of Computer Science, Stanford University, Stanford, CA 94305, USA {smenon, takatoki, ok}@cs.stanford.edu

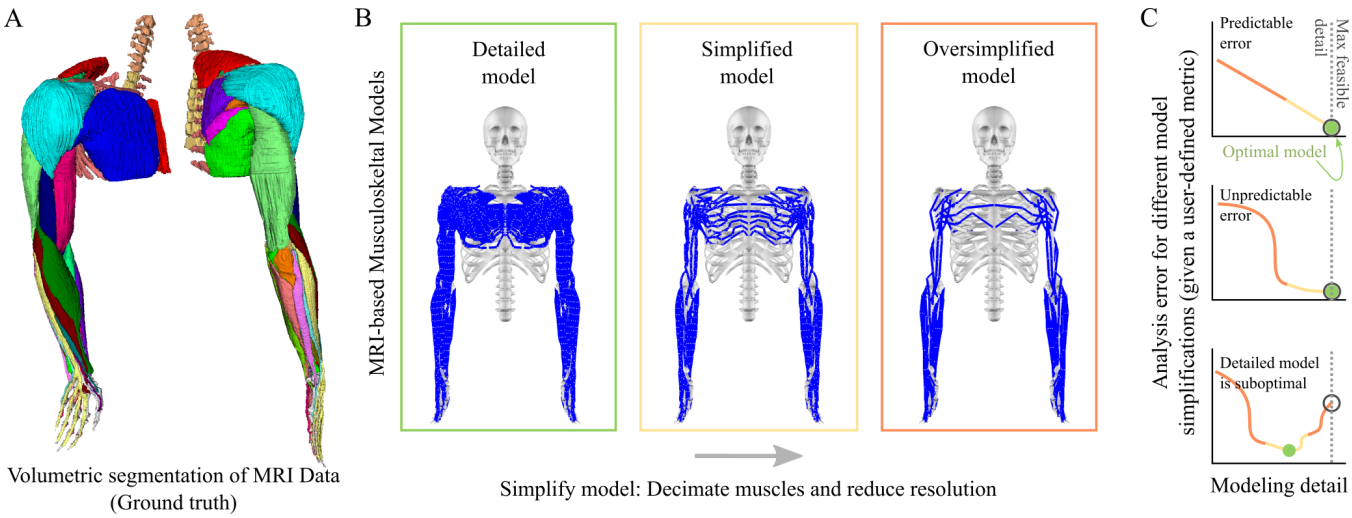

Fig. 2. **Comparing Model Accuracy and Analysis Error. A.** Volumetric rendering of bones and muscles extracted from a subject's anatomical MRI data. **B.** A family of models generated from the volumetric data. Skeletons are identical. The muscle model on the left very accurately captures muscle volumes (2.5 mm radius and 2 cm length fiber-group segments). The other two models are parametrically decimated by reducing the number of fiber groups per unit area, without dropping muscles. The musculature in the lower arm is better preserved since the muscles are more numerous and thinner, and thus lose less detail. **C.** Analyzing a family of MRI-based models with varying accuracy provides insights into the level of detail required for a given biomechanical analysis. Note that these graphs show hypothetical scenarios and should be computed by users of the model with an error metric suitable for their application. Ideal models have predictable error trends, such as the one in the top graph. Unpredictable errors make it difficult to determine the model detail required for a given application. Increasing detail may not necessarily decrease error, as shown in the bottom graph.

to map one of them to a given individual human.

II. Related Work

A. Musculoskeletal Simulation

Human musculoskeletal models and dynamic simulation have been widely used to reconstruct and study quasi-static [13], [14], [15], [16], cyclical [17], and high performance motions [18], [19]. One approach to dynamic simulations uses motion capture data to solve an inverse dynamics problem for the simulated musculoskeletal model [20]. An alternative approach uses control algorithms to move simulated models in a manner that synthesizes the captured human motion [21], [22], [23], [5]. Both methods require large computational resources for complex models, and are sensitive to errors introduced by kinematic and physiological accuracies.

B. Musculoskeletal Modeling

Realizing the importance of subject-specific modeling and parametric variations, researchers have made progress in using anatomical imaging measurements to improve human models (see [24], [25] for an extensive review). A notable effort is the parametric human project (*www.parametrichuman.org*), which has led to dramatic progress in modeling humans. Related efforts include ribcage skeleton generation [26], MRI-based knee moment arm testing [13], [27], [28], hip muscle modeling [29], bicep tendon modeling [30], pennation angle modeling with ultrasound [31], and a cadaver-based forearm model [32]. While a comprehensive review of the associated literature is beyond the scope of this paper, we have attempted to sample recent as well as long-standing advances. Notably, such efforts are synergistic with our own. An upper-body model that was developed recently [33], for instance, forms a benchmark with which to compare our work.

III. Proposed Empirical Analysis to Ensure Results are Robust to Model Errors

The primary motivation for our work was to be able to determine how analysis error (for arbitrary analyses) varies with modeling detail. As such, our remaining results work towards the goal of simplifying the following steps for any upper-body human musculoskeletal analysis:

1) Metric specification: A given musculoskeletal analysis method must clearly state metrics that map aspects of performance to real numbers.

2) Model variation: The identical analysis should be carried out while varying the model used across the parametric range that is expected to span all possible subjects with a generous margin of error. The dataset of recorded physical measurements used to instantiate modeled motions must be split into testing and training sets. Only the training set should be used at this stage, and a set of optimal and near-optimal models may be selected.

3) Model testing: Having selected a suitable set of models, the final analysis should be performed on them with the remaining test data. Results should be similar across the set of models.

This process, which we call "*model cross-validation*", is common in statistical analysis and helps avoid under- or over-fitting to the data. For musculoskeletal modeling, it will also build confidence that obtained results are not sensitive to model perturbation. In the remaining paper, we will outline our efforts to build a pipeline to obtain the family of models required to perform this method. We present results

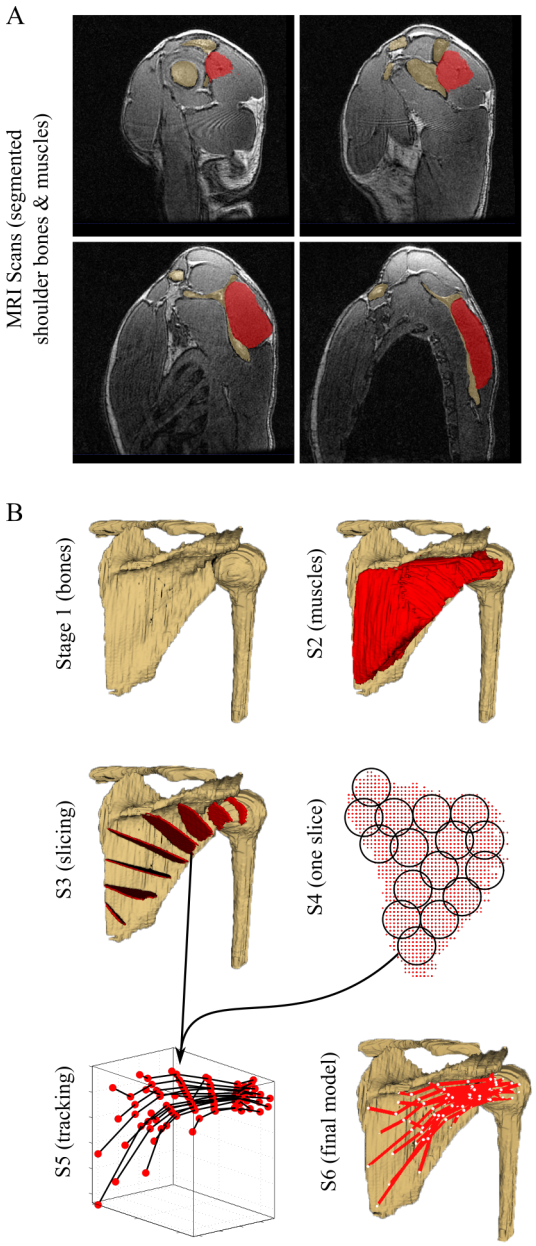

Fig. 3. **Model Generation. A.** MRI-based musculoskeletal models were obtained by segmenting high-resolution anatomical scans. Exemplar sagittal cross-sections for the shoulder are shown, matching the volumetric reconstruction below. **B.** The model generation pipeline consists of six stages. Stages 1 and 2 involved extracting three dimensional volumes for bones and muscles. Stage 3 involved slicing muscles normal to their direction of force. Stage 4 involved packing fiber-group actuator cross-sections into the muscle slices. Stage 5 involved associating actuator intersection circles across slices using Hungarian assignment [34]. And, finally, stage 6 involved connecting actuators to create piece-wise muscle approximations. The muscle may terminate in the bone across a large area, so some fibers may be shorter than others. Stages 3 and 4 may be parameterized to create families of models.

for parameterized models obtained from a specific human subject. Expanding our dataset to include many subjects of varying physical stature is essential for this proposed method to be widely applicable and successful.

IV. Subject-Specific Musculoskeletal Modeling

Human models used in the study of motion are either scaled from canonical models, where a single model (say, cadaver-based) is mapped to different subjects, or are subject-specific, where a unique model is generated for each human. The complexity associated with creating accurate musculoskeletal models for individual subjects has led human modelers to develop scaling methods that map canonical models to different humans [35]. The accuracy of such methods is limited by the fact that human physiology is diverse, scaling measurements are taken at the skin's surface, and the underlying skeletal structure of a candidate human may differ dramatically from the canonical model. A library of musculoskeletal models, however, might provide sufficient bases for scaling methods to work. Since no such library exists for upper-body musculoskeletal models, we decided to develop a subject-specific model using magnetic resonance imaging (MRI) data for our decimation analysis.

The methods described in this paper are only concerned with subject-specific musculoskeletal geometry and rigid body dynamics, not muscular dynamics. This model may be paired with different muscle actuator models for full dynamic simulations. It is important to note that dynamic parameters such as peak force, slack length, and activation dynamics cannot be obtained from MRI data, and thus alternative methods are needed to obtain this data for individual subjects.

We will now outline our results obtained for an MRI-based subject-specific upper-body (arms and chest) model. We also discuss how to parameterize the model creation process to allow creating a family of models of varying complexity.

A. MRI Measurements for Upper-Body Musculoskeletal Models

MRI is a versatile technique that requires carefully selecting scan parameters that trade-off spatial resolution and time taken. The parameter trade-off is primarily driven by the fact that scan signal decreases as a cubic of the spatial resolution, and by the physical properties of human tissue when subjected to magnetization. Engineering advances have helped improve MRI scanners and modern scanners can support sub-mm resolution using multi-channel (8–32) receiver coils. However, receiver coils are often small and cannot scan the whole arm. Moreover, the quality of scan measurements varies dramatically with the part of the body being scanned, and regions near cavities (like the lungs) tend to produce noisy data.

For our MRI measurements—keeping in mind scanner resolution constraints—we used multiple scans (2–5 scans at $0.47 \times 0.47 \times 0.8\,\mathrm{mm}^3$ resolution), made sure each scan differed from the others by a small offset (1–5 mm), and finally merged the multiple scans into a higher resolution scan (see Appendix for more details). Increasing voxel size beyond the bare minimum provided a cubic increase in signal, and averaging across multiple scans provided a reduction in noise. It is noteworthy that the up-sampling process effectively smooths the data. As such, our approach might not be ideal where sharpness is critical, such as when

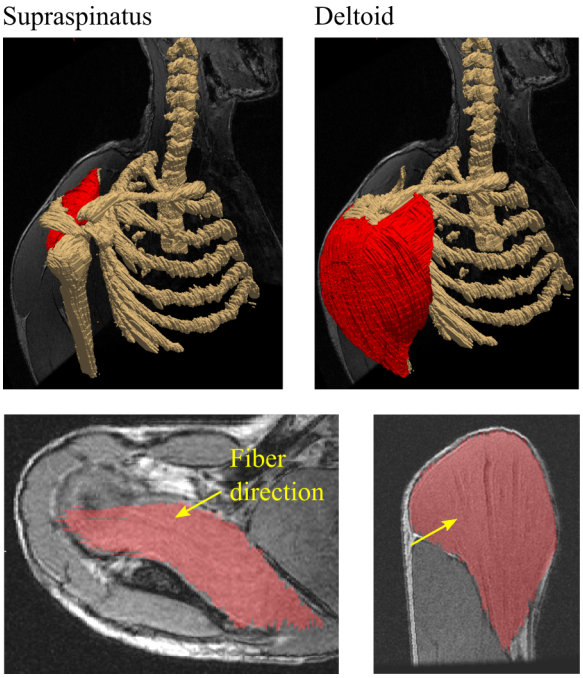

Fig. 4. **Muscle Volumes and Fiber Directions.** High resolution MRI scans allow accurate estimates of muscle volume (top). Cross-sectional slices reveal muscle fiber directions and attachment points for muscles, allowing accurate muscular kinematic modeling (bottom). Two muscle exemplars are shown. The MRI images (below) have a translucent red overlay marking the segmented muscle cross section.

mapping the attachment points of tendon-like structures in the hand or muscles near the wrist. For such cases, we at times reverted to the unaveraged (unsmoothed) scans.

Small receiver coils are best for producing high resolution images, but they are limited to scanning small regions of the body. Thus, we scanned different body parts separately and merge them to obtain a full upper body scan. To facilitate the registration process, we placed fish oil capsules, which are highly visible in MRI scans, on the subjects in the overlapping regions of the scans. We then used affine transformations defined by three oil capsules shared between any two scans to combine the scans.

B. Volumetric Segmentation of Skeleton and Muscles

Once we obtained MRI scans of acceptable quality, we proceeded to segment bones and extract volumetric skeletons. Since our focus was whole-arm motion, we treated the hand as a rigid body. It is noteworthy that the MRI scans do provide sufficient information for finger bones, which we hope to model in future work. For the whole arm, we thus categorized bones into the following groups: torso, upper arm, lower arm (radius), lower arm (ulna), and hand. The skeletal model for one arm had seven degrees of freedom, three at the torso-upper arm joint, one each at the lower arm joints, and two at the hand (Table I). Future work includes integrating more sophisticated kinematic models to deal with complex shoulder motion [36].

After segmenting the bones, we segmented the MRI scans to obtain precise volumetric estimates of the muscles. We identified thirty-three muscles (Table II).

C. Automatic MRI-Based Musculoskeletal Model Generation

Since our goal was to develop a family of MRI-based musculoskeletal models, we decided to deconstruct muscle volumes into actuator segments (fiber-groups) parameterized by varying length and diameter. Allowing parametric model variations required automating the muscle model generation process (Fig. 3). As such, we developed a pipeline that began with segmented bones and muscles in MRI scans, reconstructed bone and muscle volumes, sliced the muscles at planes normal to their fiber direction, packed circles in each slice, connected circles in consecutive slices with cylindrical actuator segments, and, finally, terminated actuator segments by attaching their ends to bones (see Appendix for details about methods).

Our first parameter of interest for creating a family of muscle models was the mean actuator segment length, which is determined by the inter-slice distance (see Fig. 3. S3). Our second parameter, the muscle fiber radius, determined the number of segments (circles) packed into a slice (see Fig. 3. S4). Note that the muscle fiber directions are visible in our MRI scans (Fig. 4).

To summarize, our pipeline allowed us to create families of models parameterized by: *(i)* the length of segments with which a piece-wise linear simulated muscle approximates a real muscle's curvature, and *(ii)* the radius of individual actuators used to approximate each muscle's volume. Varying the two parameters, we constructed musculoskeletal models from MRI data that spanned a large range in actuator segments (50–$200\,\mathrm{segments/m}$) and diameters ($5\,\mathrm{mm}$–$2\,\mathrm{cm}$; produces 1–$47\,\mathrm{actuators/muscle}$). The parameter range was selected in a manner that the most detailed settings very closely approximated the actual curvature of the volumetric muscles. Moreover, focusing on actuator diameter instead of actuators per muscle helped keep the volumetric sampling uniform across large and small muscles.

V. Comparing Subject-Specific and Canonical Models

While anatomically accurate subject-specific models are attractive, it is presently challenging to create a new model for every conceivable individual. It is inevitable—at least in the near future—to rely on scaled canonical models. Moreover, as the number of subject-specific models increases, it may become feasible to accurately scale one out of the many existing models to a new individual. As such, we decided to compare subject-specific models to scaled canonical models.

While comparisons between models are best interpreted with respect to a specific desired analysis metric, it remains feasible to simply compare anatomical accuracy after scaling. Since our detailed MRI-based upper-body model serves as ground truth, we used it to visualize the degree to which commonly used affine scaling methods can approximate skeletal structure.

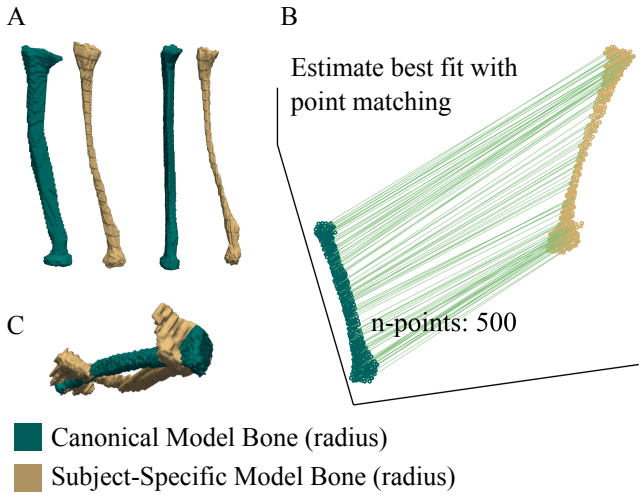

Fig. 5. **Model Scaling Errors. A.** A canonical model's radius bone side-by-side with an MRI-based subject-specific model's radius bone; front and side view shown. The subject-specific model is accurate to $< 1\,\mathrm{mm}$, and was considered to be the ground truth. **B.** The canonical model was scaled to the subject's radius with an affine transformation that optimized the distance between five hundred corresponding points between the two bones. **C.** The scaled canonical model was unable to match the geometry of the subject-specific model. Moreover, affine fits can be expected to be substantially worse when ground truth is unavailable.

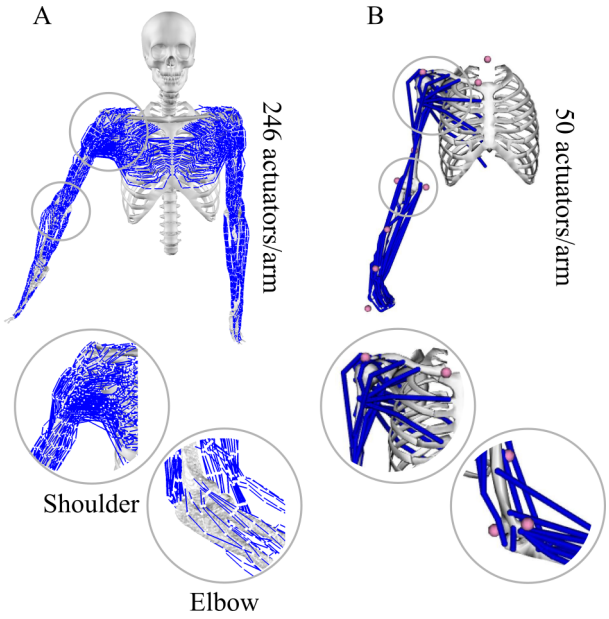

Fig. 6. **Comparison with a state-of-the-art canonical model. A.** The MRI-based subject-specific model. The detailed musculature at the shoulder and elbow is highlighted. **B.** An existing state-of-the art model. While the lower arm has a lot of detail, the upper arm does not match the detail provided by the MRI-based model.

A. Bone-to-Bone Scaling

To simplify the interpretation of scaling method performance, we focused on bone-to-bone scaling error as a metric (Fig. 5). We selected a canonical radius bone model provided with OpenSim [35], as well as a radius bone model from our MRI-based model. We then selected a finite number of control points (500) in each bone using a volume filling algorithm. Each successive control point was picked to maximize its distance from the closest among the previously selected control points. Next, we estimated point-correspondence between the subject-specific and canonical models using the Hungarian algorithm [37]. Finally, we fit an affine transformation to collectively warp all the points in a manner that minimized the aggregate point-to-point Euclidean distance. Increasing the number of points did not change the resulting transformation substantially.

Our bone-to-bone scaling method's use of point-correspondence between the bones allows non-linear scaling methods. We chose not to use nonlinear scaling since it was our goal to simply identify a close-to-optimal affine transformation between the two bones. We also note that bone-to-bone scaling differs from other scaling methods in that it assumes access to a ground truth bone, which, if available, would eliminate the need for scaling itself. As such, it is not useful as a method for mapping a canonical skeleton to a given subject. It merely serves as a testbed to identify how accurate a given scaling method can possibly be.

The scaled canonical bone differed in geometry from the subject's bone (see Fig. 5.C). Whether this difference is consequential or not will require comparing the performance of the scaled model to the subject-specific model for a given analysis metric. As such, we chose not to presume any analysis metric and instead simply visualized the differences.

B. Comparison with Existing Canonical Models

In addition to a bone-to-bone scaling comparison, we also compared a recently developed arm model [33] with our subject-specific model (Fig. 6). The canonical model has thirty-three muscles and fifty actuator segments. Most acuators, however, are concentrated at the lower arm and hand. As such, the model does not accurately capture muscular kinematics at the elbow, shoulder, or chest.

We must note that the shoulder joint is complex [36]. While our MRI-based model has highly accurate kinematics, it requires additional data to ensure that the accuracy is sustained with shoulder motion. As such, its advantage in kinematic detail over the canonical model is confined to small motions that do not cause substantial movement in the scapulothoracic joint.

VI. Dynamics & Control for Motion Reconstruction

We demonstrated that our models can be used for realtime dynamics and control analyses using motion reconstruction as an example application. To do so, we first modeled the skeletal structure as an articulated body mechanism with seven degrees-of-freedom (see Appendix for details) whose dynamics may be modeled using the dynamics equation:

$$\Gamma = A(q)\ddot{q} + b(q, \dot{q}) + g(q), \tag{1}$$

where $q, \dot{q}, \ddot{q}$ are the generalized positions, velocities, and accelerations, respectively, Γ is the vector of generalized forces, $A(q)$ is the generalized articulated body inertia, $b(q, \dot{q})$ is

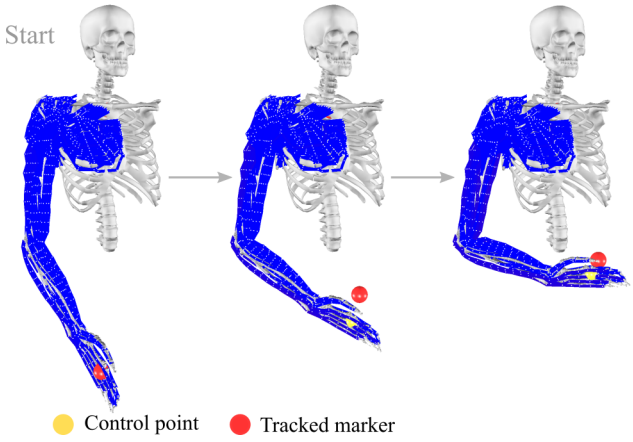

Fig. 7. **Actuating MRI-Based Musculoskeletal Models.** To demonstrate our model in a dynamic simulation, we simulated a specific model with two hundred and forty-six muscle segments. The model performed an end-effector position control task. A control point (yellow) attached to the simulated model's hand tracked an optical marker attached to a human's hand (red). The human hand marker's motion was recorded using optical motion capture (see Appendix for details). A single arm model was used for clarity.

the vector of centrifugal and coriolis forces, and $g(q)$ is the vector of generalized gravity. We modeled muscles by assuming that sections of muscle were attached to their nearest bone, which merged the inertias of the muscles into the links. The piece-wise segments of muscle thus moved as if they were attached to the closest bone. Sections of the muscle that crossed over bone-to-bone joints changed shape as the arm moved. Integrating the dynamics with the recursive articulated body algorithm [38] allowed us to develop a dynamic simulation of the full biomechanical model.

Having developed a dynamic simulation of our model, we then used the model to track the motion of a human by implementing a direct marker-space control formulation [5] using the Standard Control Library (SCL) [39] (Fig. 7). To briefly review the control strategy, we used a marker space task controller:

$$F^*_{marker} = k_p(x_{des} - x) + k_v(\dot{x}_{des} - \dot{x}) \quad (2)$$

$$\Gamma = \Lambda(q)F^*_{marker} + \mu(q, \dot{q}) + p(q), \quad (3)$$

where F^*_{marker} is the marker control force (or task acceleration), k_p and k_v are proportional-derivative control gains, $x, x_{des}, \dot{x}$, and $\dot{x}_{des}$ are current and desired marker positions and velocities, $\Lambda(q)$ is the marker space inertia matrix, $\mu(q, \dot{q})$ is the marker space centrifugal and coriolis force, and $p(q)$ the marker space gravity force. Marker positions were measured using an OptiTrack motion capture system (see Appendix for details), and control gains were hand-tuned to ensure a stable simulation.

We note that since our model is fitted to musculoskeletal kinematics at the base configuration, it works best near that configuration. Notably, muscular kinematics are affected by changes in muscle volume caused by contraction, and by muscle displacement due to sliding over other muscles and bones [29]. In addition, skeletal kinematics are affected by bones that move with respect to each other [36]. Quantifying the effect of these phenomena is an active research area. As such, we focused our dynamic simulation to a small motion range where the kinematic assumptions embedded in our model were satisfied: to motions that do not displace the shoulder-bone complex, and that do not deform muscles to a point where our model's fiber routing algorithm produces fiber overlap. Future research in developing articulation models will generalize our model to a wider range of motion.

VII. Conclusions

To summarize, we present a novel subject-specific upper extremity human musculoskeletal model that was extracted from high-resolution MRI data. The model's bones and musculature can be independently parameterized to create a large family of models. The family of models can serve as a testbed to cross-validate results obtained with arbitrary analysis methods and can help identify whether the analysis methods (or associated data measurements) are sensitive to variations in modeling accuracy and detail. Finally, our models can be used in realtime dynamics and control simulations. Future work includes developing a much larger set of subject-specific models, which, potentially, will provide a large enough basis set to make model scaling a reasonable approach.

References

[1] A. Steinfeld, T. Fong, D. Kaber, M. Lewis, J. Scholtz, A. Schultz, and M. Goodrich, "Common metrics for human-robot interaction," in *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-robot Interaction*, ser. HRI '06. New York, NY, USA: ACM, 2006, pp. 33–40.

[2] S.-Y. Baek and K. Lee, "Parametric human body shape modeling framework for human-centered product design," *Computer-Aided Design*, vol. 44, no. 1, pp. 56 – 67, 2012.

[3] M. Grebenstein, M. Chalon, W. Friedl, S. Haddadin, T. Wimböck, G. Hirzinger, and R. Siegwart, "The hand of the dlr hand arm system: Designed for interaction," *The International Journal of Robotics Research*, vol. 31, no. 13, pp. 1531–1555, 2012.

[4] J. M. Wang, D. J. Fleet, and A. Hertzmann, "Gaussian process dynamical models for human motion," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 30, no. 2, pp. 283–298, Feb 2008.

[5] E. Demircan, T. Besier, S. Menon, and O. Khatib, "Human motion reconstruction and synthesis of human skills," in *Advances in Robot Kinematics*, J. Lenarčič and M. Stanisic, Eds. Springer, Berlin, Heidelberg, Germany, 2010, pp. 283–292.

[6] C. Hu, Q. Yu, Y. Li, and S. Ma, "Extraction of parametric human model for posture recognition using genetic algorithm," in *Automatic Face and Gesture Recognition, 2000. Proceedings. Fourth IEEE International Conference on*, 2000, pp. 518–523.

[7] S. Liu, X. Liang, L. Liu, X. Shen, J. Yang, C. Xu, L. Lin, X. Cao, and S. Yan, "Matching-cnn meets knn: Quasi-parametric human parsing," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.

[8] A. D. Santis, B. Siciliano, A. D. Luca, and A. Bicchi, "An atlas of physical humanrobot interaction," *Mechanism and Machine Theory*, vol. 43, no. 3, pp. 253 – 270, 2008.

[9] A. Albu-Schäffer, S. Haddadin, C. Ott, A. Stemmer, T. Wimböck, and G. Hirzinger, "The dlr lightweight robot: design and control concepts for robots in human environments," *Industrial Robot: An International Journal*, vol. 34, no. 5, pp. 376–385, 2007.

[10] S. Haddadin, A. Albu-Schäffer, and G. Hirzinger, "Requirements for safe robots: Measurements, analysis and new insights," *The International Journal of Robotics Research*, vol. 28, no. 11-12, pp. 1507–1527, 2009.

[11] D. A. Padua, S. W. Marshall, M. C. Boling, C. A. Thigpen, W. E. Garrett, and A. I. Beutler, "The landing error scoring system (less) is a valid and reliable clinical assessment tool of jump-landing biomechanics: The jump-acl study," *The American Journal of Sports Medicine*, vol. 37, no. 10, pp. 1996–2002, 2009.

[12] H. B. Henninger, S. P. Reese, A. E. Anderson, and J. A. Weiss, "Validation of computational models in biomechanics," *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, vol. 224, no. 7, pp. 801–812, 2010.

[13] D. J. A. Allison S. Arnold, Silvia Salinas and S. L. Delp, "Accuracy of muscle moment arms estimated from mri-based musculoskeletal models of the lower extremity," *Computer Aided Surgery*, 2000.

[14] M. Viceconti, D. Testi, F. Taddei, S. Martelli, G. Clapworthy, and S. Jan, "Biomechanics modeling of the musculoskeletal apparatus: Status and key issues," *Proceedings of the IEEE*, vol. 94, no. 4, pp. 725–739, April 2006.

[15] G. E. G. Silvia S. Blemker, Deanna S. Asakawa and S. L. Delp, "Image-based musculoskeletal modeling: Applications, advances, and future opportunities," *Journal of Magnetic Resonance Imaging*, 2007.

[16] M. Krokos, D. Podgorelec, G. Clapworthy, R. Liang, D. Testi, and M. Viceconti, "Patient-specific muscle models for surgical planning," in *Medical Information Visualisation - Biomedical Visualisation, 2005. (MediVis 2005). Proceedings. Third International COnference on*, July 2005, pp. 3–8.

[17] D. G. Thelen, F. C. Anderson, and S. L. Delp, "Generating dynamic simulations of movement using computed muscle control," *Journal of Biomechanics*, vol. 36, no. 3, pp. 321 – 328, 2003.

[18] M. Hayashibe, G. Venture, K. Ayusawa, and Y. Nakamura, "Muscle strength and mass distribution identification toward subject-specific musculoskeletal modeling," in *IROS'11*, 2011, pp. 3701–3707.

[19] J. Teran, E. Sifakis, S. S. Blemker, V. Ng-Thow-Hing, C. Lau, and R. Fedkiw, "Creating and simulating skeletal muscle from the visible human data set," *Visualization and Computer Graphics, IEEE Transactions on*, vol. 11, no. 3, pp. 317–328, 2005.

[20] S. L. Delp, F. C. Anderson, A. S. Arnold, P. Loan, A. Habib, T. John, E. Guendelman, and D. G. Thelen, "Opensim: Open-source software to create and analyze dynamic simulations of movement."

[21] O. Khatib, J. Warren, V. D. Sapio, and L. Sentis, "Human-like motion from physiologically-based potential energies," *Advances in Robot Kinematics*, pp. 149–163, 2004.

[22] V. D. Sapio, J. Warren, O. Khatib, and S. Delp, "Simulating the Task-level Control of Human Motion: A Methodology and Framework for Implementation," *The Visual Computer*, vol. 25, no. 1, pp. 289–302, June 2006.

[23] O. Khatib, E. Demircan, V. DeSapio, L. Sentis, T. Besier, and S. Delp, "Robotics-based synthesis of human motion," *Journal of Physiology - Paris*, vol. 103(3-5), pp. 211–219, 2009.

[24] D. Lee, M. Glueck, A. Khan, E. Fiume, and K. Jackson, "A survey of modeling and simulation of skeletal muscle," *ACM Transactions on Graphics*, vol. 28, no. 4, pp. 1–13, 2010.

[25] D. Lee, M. Glueck, A. Khan, E. Fiume, K. Jackson, *et al.*, *Modeling and simulation of skeletal muscle for computer graphics: A survey*. Now Publishers, 2012.

[26] B. Aubert, C. Vergari, B. Ilharreborde, A. Courvoisier, and W. Skalli, "3d reconstruction of rib cage geometry from biplanar radiographs using a statistical parametric model approach," *Computer Methods in Biomechanics and Biomedical Engineering: Imaging and Visualization*, vol. 4, no. 5, p. 281295, Sep 2016.

[27] D. S. Asakawa, K. S. Nayak, S. S. Blemker, S. L. Delp, J. M. Pauly, D. G. Nishimura, and G. E. Gold, "Real-time imaging of skeletal muscle velocity," *Journal of Magnetic Resonance Imaging*, vol. 18, no. 6, pp. 734–739, 2003.

[28] S. S. Blemker, D. S. Asakawa, G. E. Gold, and S. L. Delp, "Image-based musculoskeletal modeling: Applications, advances, and future opportunities," *Journal of Magnetic Resonance Imaging*, vol. 25, no. 2, pp. 441–451, 2007.

[29] S. Blemker and S. Delp, "Three-dimensional representation of complex muscle architectures and geometries," *Annals of Biomedical Engineering*, vol. 33, no. 5, pp. 661–673, 2005.

[30] C. Walton, Z. Li, A. Pennings, A. Agur, and A. Elmaraghy, "A 3-dimensional anatomic study of the distal biceps tendon: Implications for surgical repair and reconstruction," *Orthopaedic Journal of Sports Medicine*, vol. 3, no. 6, 2015.

[31] D. Lee, Z. Li, Q. Z. Sohail, K. Jackson, E. Fiume, and A. Agur, "A three-dimensional approach to pennation angle estimation for human skeletal muscle," *Computer Methods in Biomechanics and Biomedical Engineering*, vol. 18, no. 13, pp. 1474–1484, 2015.

[32] Z. Li, J. P. Mogk, D. Lee, J. Bibliowicz, and A. M. Agur, "Development of an architecturally comprehensive database of forearm flexors and extensors from a single cadaveric specimen," *Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization*, vol. 3, no. 1, pp. 3–12, 2015.

[33] K. R. Saul, X. Hu, C. M. Goehler, M. E. Vidt, M. Daly, A. Velisar, and W. M. Murray, "Benchmarking of dynamic simulation predictions in two software platforms using an upper limb musculoskeletal model," *Computer Methods in Biomechanics and Biomedical Engineering*, vol. 18, no. 13, pp. 1445–1458, 2015.

[34] J. Munkres, "Algorithms for the assignment and transportation problems," *Journal of the Society for Industrial and Applied Mathematics*, vol. 5, no. 1, pp. 32–38, 1957.

[35] D. S. et. al., "Opensim: Open-source software to create and analyze dynamic simulations of movement," *IEEE Transactions on Biomedical Engineering*, 2007.

[36] A. Seth, R. Matias, A. P. Veloso, and S. L. Delp, "A biomechanical model of the scapulothoracic joint to accurately capture scapular kinematics during shoulder movements," *PloS one*, vol. 11, no. 1, p. e0141028, 2016.

[37] H. W. Kuhn, "The hungarian method for the assignment problem," *Naval research logistics quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955.

[38] R. Featherstone, *Rigid body dynamics algorithms*. Springer, 2014.

[39] S. Menon, "The Standard Control Library Version 1.00," 2016, http://samirmenon.org/scl.html.

[40] G. Wu, F. C. van der Helm, H. D. Veeger, M. Makhsous, P. V. Roy, C. Anglin, J. Nagels, A. R. Karduna, K. McQuade, X. Wang, F. W. Werner, and B. Buchholz, "Isb recommendation on definitions of joint coordinate systems of various joints for the reporting of human joint motionpart ii: shoulder, elbow, wrist and hand," *Journal of Biomechanics*, vol. 38, no. 5, pp. 981 – 992, 2005.

ACKNOWLEDGMENTS

We acknowledge Kaidi Yan and Anupama Rajan for contributing to segmenting bones and muscles in the MRI voxel data, Hari Ganti and Paul Quigley for assisting with MRI data collection, and Fanny Chen for rendering images. We also thank Kaidi Yan for helping prototype an early version of the muscle model generation pipeline. Finally, we thank Bob Dougherty and Laima Baltusis for their assistance with developing MRI scanning protocols, and Scott Delp and Samuel Hamner for discussions about modeling humans and insights into developing musculoskeletal models.

APPENDIX

MRI Protocol:

All MRI scans were conducted at Stanford University's Center for Cognitive and Neurobiological Imaging on a GE Discovery MR750 3T MRI scanner. Three different coils were used: *(i)* a thirty-two channel Nova Medical head coil for scanning hands and forearms, *(ii)* a sixteen channel coil for scanning elbows and upper arms, and *(iii)* the body coil to scan shoulders. Multiple T1 anatomical scans (2–5 scans at $0.47 \times 0.47 \times 0.8\,\mathrm{mm}^3$ resolution) were performed and averaged to obtain high signal to noise. Scans of separate body parts were taken such that they overlapped with other scans and could be merged using fish oil capsule markers to create a full body image.

TABLE I

SKELETAL ARTICULATION MODEL

Joint Number (q_i)	*Rotational Axis (Z_i)*
Ground	N/A
1	Axis perpendicular to Z_3 in the anterior direction
2	Axis perpendicular to Z_3 and Z_1
3	Axis between $\mathcal{O}_1$ and $\mathcal{O}_4$
4	Axis between medial and lateral epicondyles
5	Axis between $\mathcal{O}_4$ and ulnar styloid process
6	Axis between radial and ulnar styloid processes
7	Axis perpendicular to Z_6 and axis connecting $\mathcal{O}_6$ to third metacarpal head

Joint Number (q_i)	*Frame Origin ($\mathcal{O}_i$)*	*Attached Link*
Ground	Incisura Jugularis	Torso
1	Glenohumeral rotation center	
2	$\mathcal{O}_1$ *(same as joint 1)*	
3	$\mathcal{O}_1$	Humerus
4	Midpoint between medial and lateral epicondyles	
5	$\mathcal{O}_4$	Radius and ulna
6	Midpoint between radial and ulnar styloid processes	
7	$\mathcal{O}_6$	Hand

TABLE II

trapezius	latissimus dorsi
pectoralis major	pectoralis minor
deltoid	teres major
teres minor	supraspinatus
infraspinatus	subscapularis
coracobrachialis	biceps brachii
brachialis	triceps brachii
anconeus	pronator teres
flexor carpi radialis	flexor carpi ulnaris
palmaris longus	flexor digitorum superficialis
flexor digitorum profundus	pronator quadratus
flexor pollicis longus	extensor digitorum
extensor digiti minimi	extensor carpi ulnaris
extensor carpi radialis	brachio radialis
supinator	extensor indicis
abductor pollicis longus	extensor pollicis brevis
extensor pollicis longus	

List of Arm Muscles: We segmented thirty three arm muscle volumes from MRI data. By modeling muscles with groups of fiber paths, the fiber generation process naturally divides muscles with multiple segments (such as the deltoid, with anterior, lateral, and posterior heads) into their functional components. The family of upper body models created using these muscles had 66–492 actuator segments.

Human Subject:

The human subject was a healthy right handed male, 21 years old, $160\,\mathrm{lb}$, and $5\,\mathrm{ft}\ 9\,\mathrm{in}$. Informed consent was obtained in advance on a protocol approved by the Institutional Review Board (IRB) at Stanford University.

Motion Capture:

We used an OptiTrack Prime 13 system operating at $200\,\mathrm{Hz}$ to record human motion for motion reconstruction. An Optitrack rigid body marker tracked the head of the third metacarpal. In the dataset used for motion reconstruction, subjects flexed their right wrists, shoulders, and elbows in a workspace of about $0.7\,\mathrm{m}^3$ in front of them.

Skeleton Articulation Model:

We modeled simplified articulation kinematics that aim to match present standards [40] (Table I). As with the standard, the humerus-shoulder connection was modeled as a spherical joint, and the elbow and the radio-ulnar complex were modeled as hinge joints. In contrast with the standard specification, which uses external landmarks, we specified joint rotation axes by directly estimating them from our MRI scans (see Table I for details).

Muscle Fiber Generation Method:

1) Slicing: The first step in generating piece-wise linear muscle fibers is slicing the muscle segmentation from MRI data normal to the direction of fibers. These slices define the vertices of the piece-wise linear approximations.

We assume muscle fibers follow the curvature of the muscle segmentation volume from its origin points to its insertion points. To slice the muscle normal to this curvature, we first slice the volume linearly along the axis from the centroid of the muscle origin to the centroid of the insertion, with the slices spaced according to the specified muscle segment length. We then connect the centroids of each slice to create a piecewise-linear curve that approximates a single muscle fiber going through the center of the muscle. Finally, we re-slice the muscle normal to this central fiber.

2) Circle Packing: To obtain cylindrical fiber approximations, we pack the largest muscle slice with circles of the specified fiber radius. To distribute the circles across the slice, we first shrink the slice by performing a morphological erosion with a circle of the specified fiber radius, allowing us to treat circles as points on the reduced slice. We choose the first sample point randomly and sample a succeeding point at a distance given by the fiber radius from the first sample. We continue sampling in this manner until no more valid samples exist.

3) Point Matching: Circle packing gives us the centers of cylindrical muscle fibers on the largest slice. To track the fibers through all the slices, we sample a large number of points on each slice (greater than the number of fibers), and connect samples between slices using the Hungarian method [34]. Samples connected to the original fiber centers from circle packing form the complete fibers.

4) Post Processing: All slices within a muscle initially contain the same number of circles, which means that every actuator segment intersects every muscle slice. However, some actuator segments are shorter and thus should intersect fewer slices. To solve this, we go through each fiber and project the intersection points onto the axis connecting the origin and insertion points. We then remove all intersection points that fall beyond the attachment points.