



Computational Modeling: Human Dynamic Model

Lijia Liu^{1*}, Joseph L. Cooper^{1,2} and Dana H. Ballard¹

¹ Department of Computer Science, The University of Texas at Austin, Austin, TX, United States, ² Google Inc., Mountain View, CA, United States

OPEN ACCESS

Edited by:

Jose De Jesus Rubio,
Instituto Politécnico Nacional (IPN),
Mexico

Reviewed by:

Stefano Tortora,
University of Padua, Italy
Giacinto Barresi,
Italian Institute of Technology (IIT), Italy
Shiyin Qiu,
Harbin Institute of Technology, China
Dante Mujica-Vargas,
Centro Nacional de Investigación y
Desarrollo Tecnológico, Mexico
Sheng Cao,
Kobe University, Japan

*Correspondence:

Lijia Liu
lijialiu@utexas.edu

Received: 11 June 2021

Accepted: 24 August 2021

Published: 24 September 2021

Citation:

Liu L, Cooper JL and Ballard DH
(2021) Computational Modeling:
Human Dynamic Model.
Front. Neurobot. 15:723428.
doi: 10.3389/fnbot.2021.723428

Improvements in quantitative measurements of human physical activity are proving extraordinarily useful for studying the underlying musculoskeletal system. Dynamic models of human movement support clinical efforts to analyze, rehabilitate injuries. They are also used in biomechanics to understand and diagnose motor pathologies, find new motor strategies that decrease the risk of injury, and predict potential problems from a particular procedure. In addition, they provide valuable constraints for understanding neural circuits. This paper describes a physics-based movement analysis method for analyzing and simulating bipedal humanoid movements. The model includes the major body segments and joints to report human movements' energetic components. Its 48 degrees of freedom strike a balance between very detailed models that include muscle models and straightforward two-dimensional models. It has sufficient accuracy to analyze and synthesize movements captured in real-time interactive applications, such as psychophysics experiments using virtual reality or human-in-the-loop teleoperation of a simulated robotic system. The dynamic model is fast and robust while still providing results sufficiently accurate to be used to animate a humanoid character. It can also estimate internal joint forces used during a movement to create effort-contingent stimuli and support controlled experiments to measure the dynamics generating human behaviors systematically. The paper describes the innovative features that allow the model to integrate its dynamic equations accurately and illustrates its performance and accuracy with demonstrations. The model has a two-foot stance ability, capable of generating results comparable with an experiment done with subjects, and illustrates the uncontrolled manifold concept. Additionally, the model's facility to capture large energetic databases opens new possibilities for theorizing as to human movement function. The model is freely available.

Keywords: dynamic modeling, motor control, kinematic representation, movement costs, human movement simulation

1. INTRODUCTION

The complexity of human motion was first dramatically captured via the Muybridge high-speed photographs (Muybridge, 1887; Andriacchi and Alexander, 2000; Wolpert and Landy, 2012) which spawned several different analysis techniques in different disciplines. Visualization first used keyframing techniques, but later sophisticated models were used in advanced rendering for computer graphics (e.g., Zordan and Hodgins, 2002). The early cognitive analyses of human behavior (Badler et al., 1993) focused on human motion in problem-solving, using an

essentially logical approach. In robotics, insights have been obtained by building physical systems directly (Ijspeert et al., 2007) that straddle the boundary between humans and robotics that have shed light on human design. However, these efforts are characteristically specialized. In another development, machine learning techniques have been introduced for use in analyzing animal-like motion (Schulman et al., 2016).

At the neural level, the brain's control of movement utilizes an array of specialized subsystems. One is the vast cortical memory of movement plans. Others are the sequencing via the Basal Ganglia and mediation of input and output via the Thalamus. The details of how they coordinate (Callahan et al., 2013) are being continually worked out, but there are many open problems (Loeb and Tsianos, 2015). Given this incompleteness, a research choice can be to search for descriptions at more abstract levels.

Since movements have to respect the dynamics of the musculoskeletal system, one approach is to take these systems for granted and start with measurements of motion and work backward. However, modeling all the complexity of the human musculoskeletal system can be challenging, with over 300 degrees of freedom and 650 muscles controlling a complex interconnected skeletal system. Nonetheless, there are many modeling approaches with a dynamics focus, which can be introduced approximately chronologically.

A straightforward approach is to build a humanoid dynamic model. This method is a valuable way to imitate human movements, and it is widely used in biomedical engineering due to its compliance with real-world physical rules. However, it is not without limitations: (1) it is too difficult to model and control a complex system like the whole human body. (2) it is complicated to represent "kinematic loops," such as postures with both feet on the ground. (3) for large systems, the equations of motion in nested, rotating reference frames become very complex, making them demanding to approximate well. As a result, due to the complexity of the direct modeling method for large complex systems, the significant portion of the studies choose the simplification of two-dimensional models that constrain body motions to the sagittal plane.

For example, some studies build a two-dimensional dynamic bipedal robot by modeling the whole body with a skeleton of rigid segments connected with joints. However, those method's simplifications to human bodies limit studies to simple single-behavior human movements. The simplest bipedal robot uses three links to represent the torso and two legs in the sagittal plane (Lee, 1988; Čelikovský and Anderle, 2018). Nonetheless, there have some extensions. Five-link biped robots use two links to represent each leg (Mu and Wu, 2003, 2004; Krishchenko et al., 2007), and seven-link biped robots further extend it by adding feet (Mousavi and Bagheri, 2007; Bajrami and Murturi, 2017). Furthermore, those methods also make many assumptions in studying locomotion. For example, most researchers assume an instantaneous exchange of biped support occurs when the swing leg contacts the ground. In this way, as robot manipulators, the biped foot support can be considered a successive open loop of kinematic chain from the support point to the free ends. There have been some recent three-dimensional model improvements to the two-dimensional biped robot models (Grizzle et al., 2010;

Khusainov et al., 2016; Bailly et al., 2021). However, they are still not sophisticated compared with a real human body.

To get beyond the disadvantages of these simple models, a major way forward is to incorporate more detailed musculoskeletal models (Durandau et al., 2017; Lee et al., 2019). Among them, the most advanced and sophisticated one is OpenSim (Delp et al., 2007; Seth et al., 2011; Rajagopal et al., 2016; Dembia et al., 2020), which allows modeling large systems including detail to the level of attached muscles. OpenSim is a significant advance by providing a modeling level that includes muscle contractions, which are essential in shaping movement dynamics. However, this level works best in fitting contraction data. It can be challenging to solve desired muscle co-contractions from a movement kinematic plan because those equations are non-linear and normally under-constrained.

To avoid this complexity, many studies that need such generating capability choose to eschew muscle components and focus on simpler dynamic models at the level of inertias and joint torques and model abstract versions of the human system that still use multiple degrees of freedom but summarize detailed dynamics via joint torques. For example, co-contraction can use a simple Hill model (Blümel et al., 2012). These alternative methods of dynamics computation of such multi-jointed systems have also experienced significant advances (OpenSim can also be used in this way). The foremost of these is to use a kinematic plan to integrate the dynamic equations directly. Several dynamic libraries were designed for this purpose, such as MuJoCo¹, Bullet², Havok³, Open Dynamic Engine(ODE)⁴, and PhysX⁵. An evaluation of these dynamic libraries by Erez et al. (2015) found them roughly comparable in capability. However, the focus of these systems is on a physics engine with the expectation that the users will program their own applications.

In contrast to previous studies, our Human Dynamic Model (HDM) system is a complete model that is readily available ^{6,7}. It is built on top of the physics engine ODE, the most commonly used dynamic library in the robotic area. Our 48 degree-of-freedom HDM focuses on individual human movement modeling and applies a direct dynamics integration method (Cooper and Ballard, 2012; Johnson and Ballard, 2014) to extract torques from motion data using a novel unifying spring constraint formalism.

One of the advantages of the HDM system is that the simulation's costs can be recovered. At each frame, the instantaneous power is computed from the net joint torque and joint angular velocity. The work performed at each joint is determined by numerically integrating the instantaneous powers over the entire tracing task. In this way, given motion capture data, we can compute the mechanical cost without building a humanoid biped robot with motion equations. Note that it is common to use mechanical measures of work to indicate cost

¹MuJoCo, <http://www.mujoco.org/>

²Bullet, <https://pybullet.org/>

³Havok, <https://www.havok.com/>

⁴OpenDE, <http://www.ode.org/>

⁵PhysX, <https://developer.nvidia.com/gameworks-physx-overview>

⁶HDM, https://github.com/EmbodiedCognition/HDM_UI

⁷HDM Demo, <https://youtu.be/ASs4Wo5PQeM>

instead of metabolic energy consumption (Burdett et al., 1983). The “energetic cost” mentioned in the following sections means the mechanical cost.

The focus of the paper is to describe the HDM simulator as a practical laboratory instrument and detailed demonstrations that illustrate the model’s capabilities. The paper is organized as follows: first, the methods section starts with an introduction to the model topology and then elaborates on the use of the motion markers to capture the model’s joint torques, which allows the estimation of energetic costs. Next, the results section describes several tests taken to measure the performance of the HDM system. After that, the discussion section highlights the features and issues of the HDM and points out one of the potential applications of the model, which is its use to capture large amounts of analyzed movement data. Finally, the **Supplementary Material** includes two appendices that represent the essential low-level implementation of the model, which is the integration steps of the dynamics equations.

2. METHODS—MODEL DESIGN AND OPERATION

This section provides a comprehensive exposition of building the human dynamic model from high level. Low level details of building the HDM and the derivation of the mathematics underlying the physics simulation are presented separately in the Appendix.

2.1. Body Structural Details

Our techniques use a simulated model of the human whose movement is analyzed. The first order of business is to build a physical model capable of representing human movements, of which the accuracy influences the outcome of the analysis. **Figure 1** shows the body segments and topology of the model. The humanoid model is a collection of rigid bodies connected by joints. Each joint connects two rigid bodies with anchor points (center of rotation), defined in the reference frame of both bodies. The body dimensions (bone lengths) of the character model were determined based on motion capture data.

2.2. Data Fitting

The technique for fitting a human model to motion capture data begins with a character model that serves as a template, **Figure 1** presenting the number of body segments and topology of the model. We further assign all labeled markers used in motion capture to specific model segments. It is straightforward to derive these using techniques such as in Kirk et al. (2005) and De Aguiar et al. (2006). However, manually assigning markers is also not complicated because the motion capture suit typically puts markers on the same body segments (**Figure 2**), even if they are in slightly different places or the body segments have different dimensions.

We present a method in S2 Appendix section, for using marker data to help determine the dimensions of the model segments and where markers attach to the model. Although this method could easily be automated, in practice, the research did not rely on very many different models, so the system uses

a mechanism for relaxing the marker attachment points and joint anchors with the click of a button in the graphical user interface (**Figure 5**). With a new data set, a handful of iterations proved sufficient to produce a reasonable model with marker attachments that fit the data well enough for further analysis. This algorithm does not address joint limits on a range of motion. These can also be learned (Tournier et al., 2009), but in our case, the range of motion for each joint is set *a priori*. After determining segment lengths, we set other segment dimensions as appropriate to fit against the markers. Mass properties for each segment assume uniform density by volume.

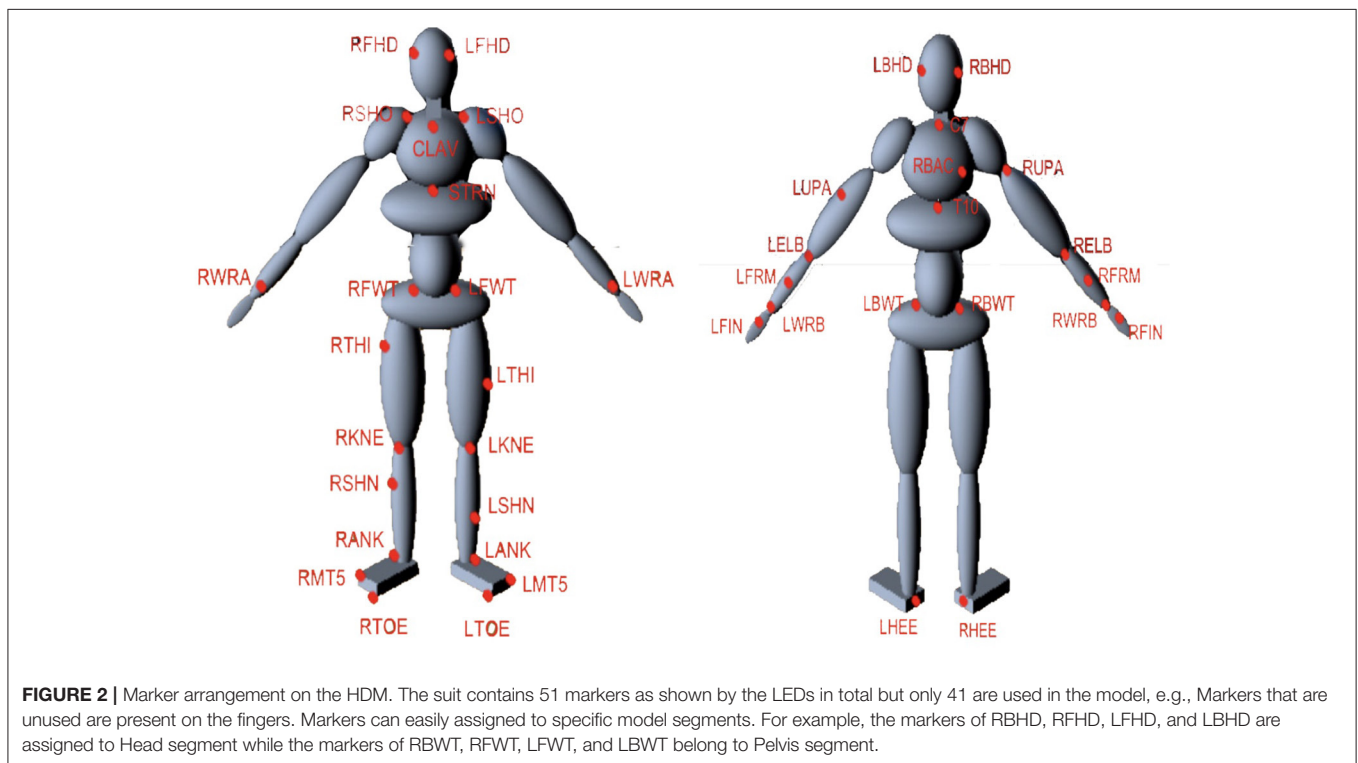
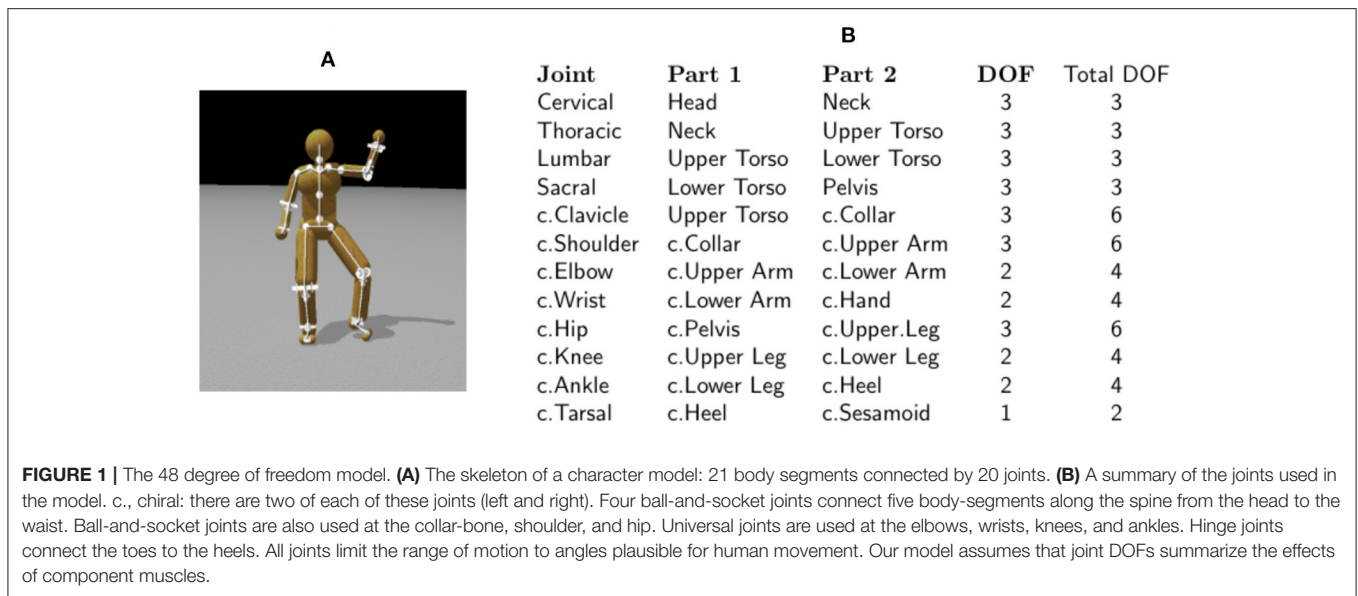
Given motion capture data of a subject, the model is fit to the subject’s dimensions, and joint-range-of-motion is constrained to approximate the subject’s flexibility. Additionally, the model segments have inertial matrix properties. The initial mass assignment to each segment assumes a uniform density of water ($1,000 \frac{\text{kg}}{\text{m}^3}$) for the volume associated with each rigid body. The mass assignment should be modified to match that of a specific subject roughly. The increased fidelity required for individual subjects in clinical biomechanics research would employ more sophisticated techniques to better approximate mass distribution in the model. However, interestingly the experimental results discussed above show that even this low fidelity model is sufficient to produce high-quality data that compares favorably with data gathered from independent sensors.

2.3. Pose Fitting

Having addressed the issues in attaching the model to motion capture data, we turn to the construction of its capability of representing human movements. Various commercial packages provide different methods for converting marker trajectories into sequences of body poses, but they can be time-consuming, expensive, or difficult to use. This section describes an approach related to Demircan et al. (2008) and Zordan and Van Der Horst (2003) that is free, fast, uses intuitive parameters, and allows the user to fit markers to whatever model they wish.

The method uses the physics engine to constrain a character model to fit marker data and other constraints. Markers are modeled as infinitely massed points attached to the character model. Given a frame of marker data, the position and orientation of all body segments can be found by balancing internal joint targets and external marker data. From the global position and orientation of the different body segments, it becomes simple to compute relative orientations (joint angles).

The internal degrees of freedom are limited by the range of motion constraints, e.g., the elbows and knees cannot bend backward. All other joints have similar range-of-motion limitations based on the subject’s flexibility. Furthermore, each joint is set to have a “target state,” a preferred relative orientation between its connected bodies. These preferences can be thought of as “muscle stiffnesses” and are modeled as weak constraints with limited force. Joint limits and stiffness serve as a prior over possible poses so that in the absence of any marker data, the model still takes on a pose. Consequently, every internal degree of freedom is constrained to some degree. These constraints hold the model together and give it a default pose. Next, in a pivotal step shown in **Figure 3**, marker data pull the model from the



default pose into a new pose. Each marker is connected to a body segment using a ball-and-socket joint constraint for a given frame of motion capture data. A total of 41 markers, which do not contribute any degrees of freedom because of their infinite mass, attach to the character model, adding additional $3 \times 41 = 123$ constraints.

Finally, collisions between the ground and the feet also influence the model pose. Each foot can form up to three contact points with the ground. Inequality constraints at these points

prevent penetration with the ground. When both feet are firmly on the ground, all markers are actively pulling the body into a pose, all joints are holding the body together, and joint limits and stiffness are biasing the relative orientation of the bodies. The experiments described above show that the model can simulate the ground force correctly.

This approach is intuitive: attach markers to the model with springs and drag the body along. The parameter, tunable for each constraint, known in ODE as the constraint force mixing

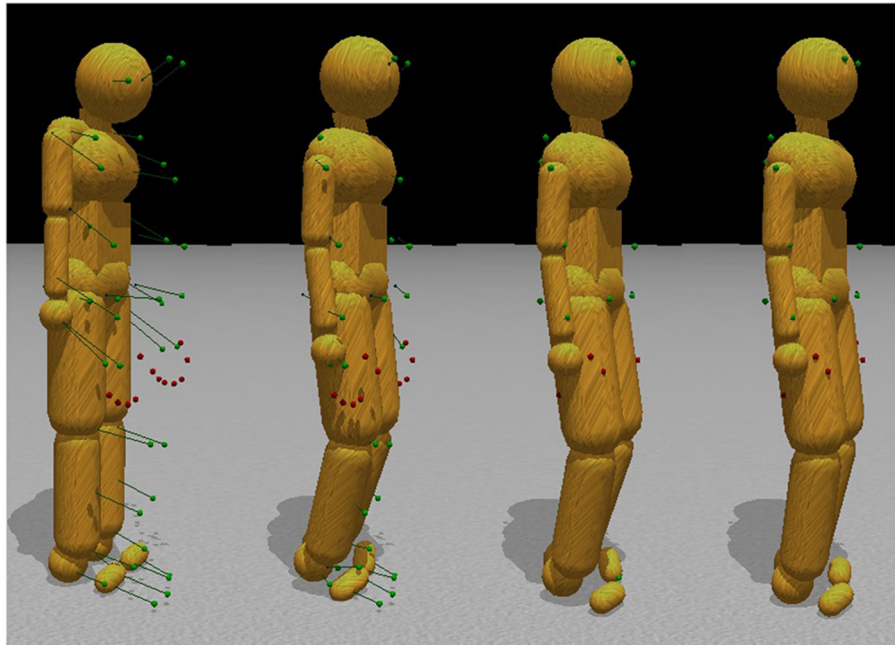


FIGURE 3 | Pose fitting. Initially, the motion capture data points are in a very different configuration than the initial stance of the model. To find the appropriate correspondences, simulated markers attach to the humanoid model through ball-and-socket joints and pull the body parts into place, subject to model joint constraints. The left to right sequence in the figure shows the body targets being gradually reconciled with the external markers.

parameter (CFM)⁸, allows a constraint to slip proportional to the amount of force required to maintain the constraint. We use a CFM value of 1×10^{-5} for the regular internal body joints and contact constraints while using 1×10^{-4} for the constraints between markers and body parts. Both of these values represent very stiff springs, although they are different by order of magnitude. This stiffness stabilizes the simulation by allowing the markers to stretch slightly from their mapped locations if the marker constraints are not compatible with the character model. **Figure 3** shows that when the markers move, the constraints drag the character along with them.

2.4. Inverse Dynamics

It can be useful to know the torques to apply at each joint or the required effort to accomplish a particular movement. Given a kinematic sequence of body poses, the physics engine ODE can archive the computation with minimal effort. Given constraints like each joint's angular velocity, it can correctly compute the desired torques/forces measurements.

The process is straightforward. Given the current joint angle and the desired joint angle for the next frame, the body parts' relative angular velocity is constrained to achieve the target orientation on the next frame. Contact constraints are necessary to prevent ground surface penetration as well. The ODE physics library handles the constraints and solves the torques and forces that satisfy each constraint in the process.

For computing inverse dynamics, the first step is to initialize the model to a starting dynamic state. The initial state can be found from the first and second frames of kinematic pose data. The model pose is set by using the second frame of data, and the initial linear and angular velocity of each joint is computed by taking the finite difference between the two frames (and dividing by the time-step). Computing velocity through finite differences is appropriate for a physics engine using first-order semi-implicit Euler integration. After that, continuously find the torques between two consecutive frames of pose data using the finite difference between poses to compute angular velocities that will move the model from the current to the next pose.

Differentiating again, this time between the current and future velocity gives a target acceleration that constrains the model. The primary difference between this step and the previously discussed method for finding pose from marker data is that no marker constraints are dragging the body into place. The internal dynamic constraints drive the model toward a target pose on each frame instead of toward a “default” pose. Because there are fewer constraints in play, stiffer muscle forces are used, but the absolute forces the muscles can apply are limited to prevent muscle forces from being unreasonably large. Again, we can use the relative spring stiffness to express confidence in the measurements in this case. We use very stiff springs (CFM = 10^{-10}) to keep the model segments together. We use looser constraints to keep the feet from penetrating the ground (CFM = 10^{-5}) and to constrain the model to adopt the appropriate pose (CFM = 10^{-8}).

⁸ODE CFM http://www.ode.org/ode-latest-userguide.html#sec_3_8_0

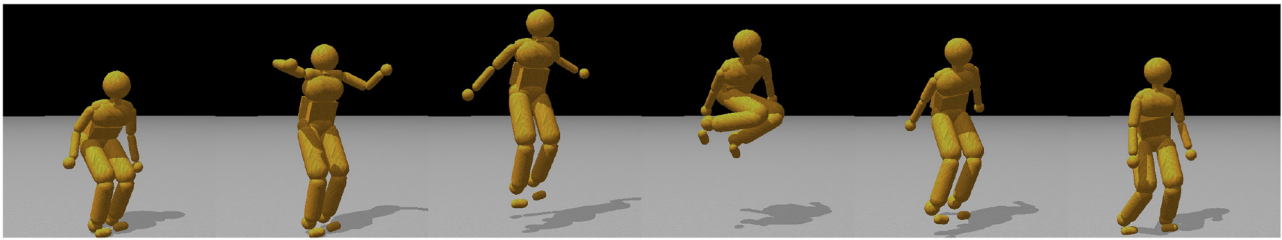


FIGURE 4 | Model capability illustration. A complex jump sequence reproduced with physics-engine-based inverse dynamics using recorded motion capture data from a human subject. The recreated jump height is achieved completely from ground forces, augmented with small residual torques (≤ 100 Nm), allowing the model to maintain balance.

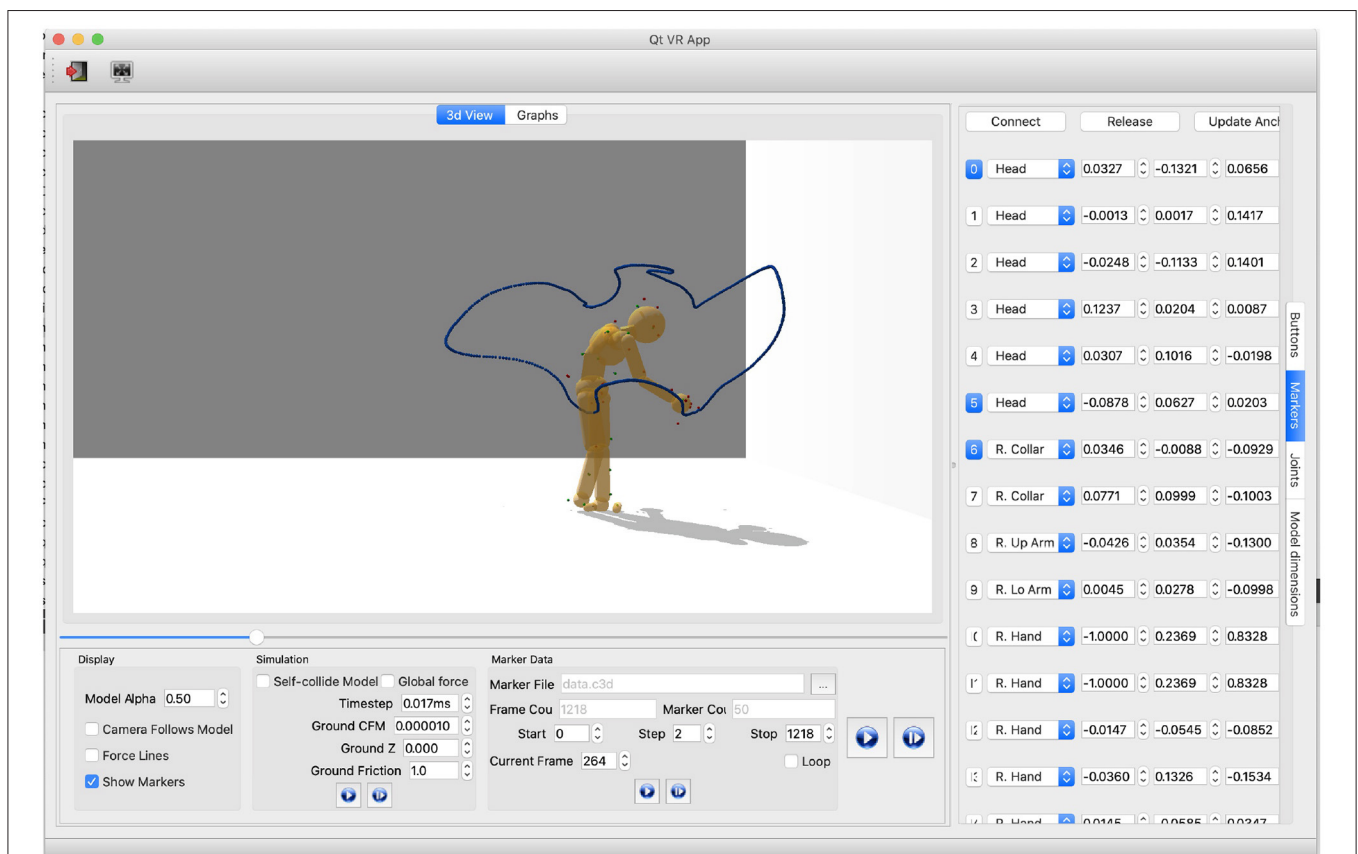


FIGURE 5 | Relevant parameters for analyzing and simulating physically-based movements must be tuned manually. Parameters of the model consist of physical world parameters, joints constraints, and the model's body-marker relative positions. In this depiction shows how users can get the current HMD configurations by clicking the buttons on the rightmost vertical menu. "Marker" is selected, meaning the marker information is shown: (1) The first column represents marker index buttons. Buttons in blue means the corresponding markers are attached to the HDM. Users can attach/detach markers by clicking index buttons. (2) The second column shows body segments where markers are attached. Each spin box is a collective item of all body segment names. Users can use it to change the body-marker attachment relationship. (3) The three-five columns present the marker-body relative positions. Users can modify the values directly using this interface. (4) The "Connect" button and "Release" button on the top are to attach or detach all the markers, respectively. The "Update Anchor" button automatically updates the marker-body relative positions based on the current motion posture.

2.4.1. Residual Torques/Forces

The torque calculation by the HDM is ideal in the sense of solving the inverse dynamic equations. The inverse dynamics uses measured kinematics and external forces to calculate net joint torques in a rigid body-linked segment model (van der Kooij

et al., 2005). However, discrepancies between the dynamic forces of the model and the kinematic of the reality make it so that the dynamic model falls over unless action is taken to stabilize it. Therefore, there needs a corrective system for unexpected errors in practice. In the human system, there are multiple corrective

systems based on vision, proprioception, and the vestibular system. Such corrective systems have been extensively studied (e.g., van der Kooij et al., 2005; Sentis et al., 2010; Welch and Ting, 2014).

In classical inverse dynamic areas, a common way to compensate for this problem is by introducing “residual forces and torques” (van der Kooij et al., 2005). In the HDM, a 6 degree of freedom joint between the waist segment and the global frame generates the external forces. A weak, limited spring constrains the waist segment to achieve its recorded pose relative to the global frame. The experiments show that attaching the external constraint to the head or the feet has little noticeable difference. The non-realistic external forces (residuals) account for noise as well as discrepancies between the model and the human generating the data. In particular, differences in how the feet interact with the ground cause errors in our analysis. In most cases, it is only necessary to constrain two of the six angular degrees of freedom (pitch and roll), leaving the other four external degrees of freedom disabled. The two angular constraints keep the body from falling over but allow it to move about through simulated ground interactions.

The stabilization system completes the model. It can be implemented in parallel, with the control used to stabilize the residual necessary to balance. With this included, The simulation can reproduce highly dynamic motions (e.g., see **Figure 4**), which shows a jumping sequence made originally by a human subject and recreated using the torques computed by the inverse dynamics model.

2.5. Method Summary

We construct a dynamic model for each human subject and force the model to follow the subject’s motion capture data, which leads directly to the recovery of joint angles. Our algorithm constrains the dynamic model to track these angles and consequently can estimate the correct joint torques. This concept was originally demonstrated in two dimensions for human walking by Faure et al. (1997). We have extended the method to the significantly more demanding case of 48 DOFs in three dimensions and arbitrary posture changes. **Figure 1** lists the body segments. The dimensions of each segment are matched those of an individual subject. The principal difficulty is that the constraints in the high DOF 3D model present many delicate numerical issues for the ODE solver that need to be addressed (Cooper and Ballard, 2012). Currently, the dynamic model does not attempt to model stiffness components, with the consequence that it can only directly recover the net torques at each DOF.

The calculation of mass properties is crucial for simulating rigid body collisions. Mass and inertial are computed using the volume of the body parts with a constant density of $1,000 \frac{Kg}{m^3}$. The articulations are designed to allow the dynamic model to simulate the majority of human movements. For instance, elbow joints have two DOFs to represent the hinge movements of the elbow as well as the twisting movements of the radius and ulna bones in the arm. Joint angles are also limited to avoid impossible movements, such as reverse bending of the elbows or knees.

For data capture, a subject wears the motion capture suit developed by PhaseSpace. Each LED marker on the suit is

mapped to a corresponding point on the model. The markers are then introduced into the physics simulation as kinematic bodies without collision geometry. As a heuristic, each marker is considered infinite mass. Thus, when another dynamic body is attached to the marker through a joint constraint, the dynamic body’s trajectory will follow the marker’s trajectory completely.

The PhaseSpace motion capture system records 3-dimension positions of specific human body locations over time. When the simulation steps forward, the constraint solver adapts the dynamic model to a state that satisfies the internal joint constraints, the external marker constraints, and other constraints such as ground forces and conservation of momentum. Knowing the kinematics allows the recovery of the dynamics since the joint velocities allow the equations of motion to be inverted. The retrieved forces can be used to generate feed-forward torque profiles for actuating the character.

Utilizing the human dynamic model to analyze human movements includes the following five steps:

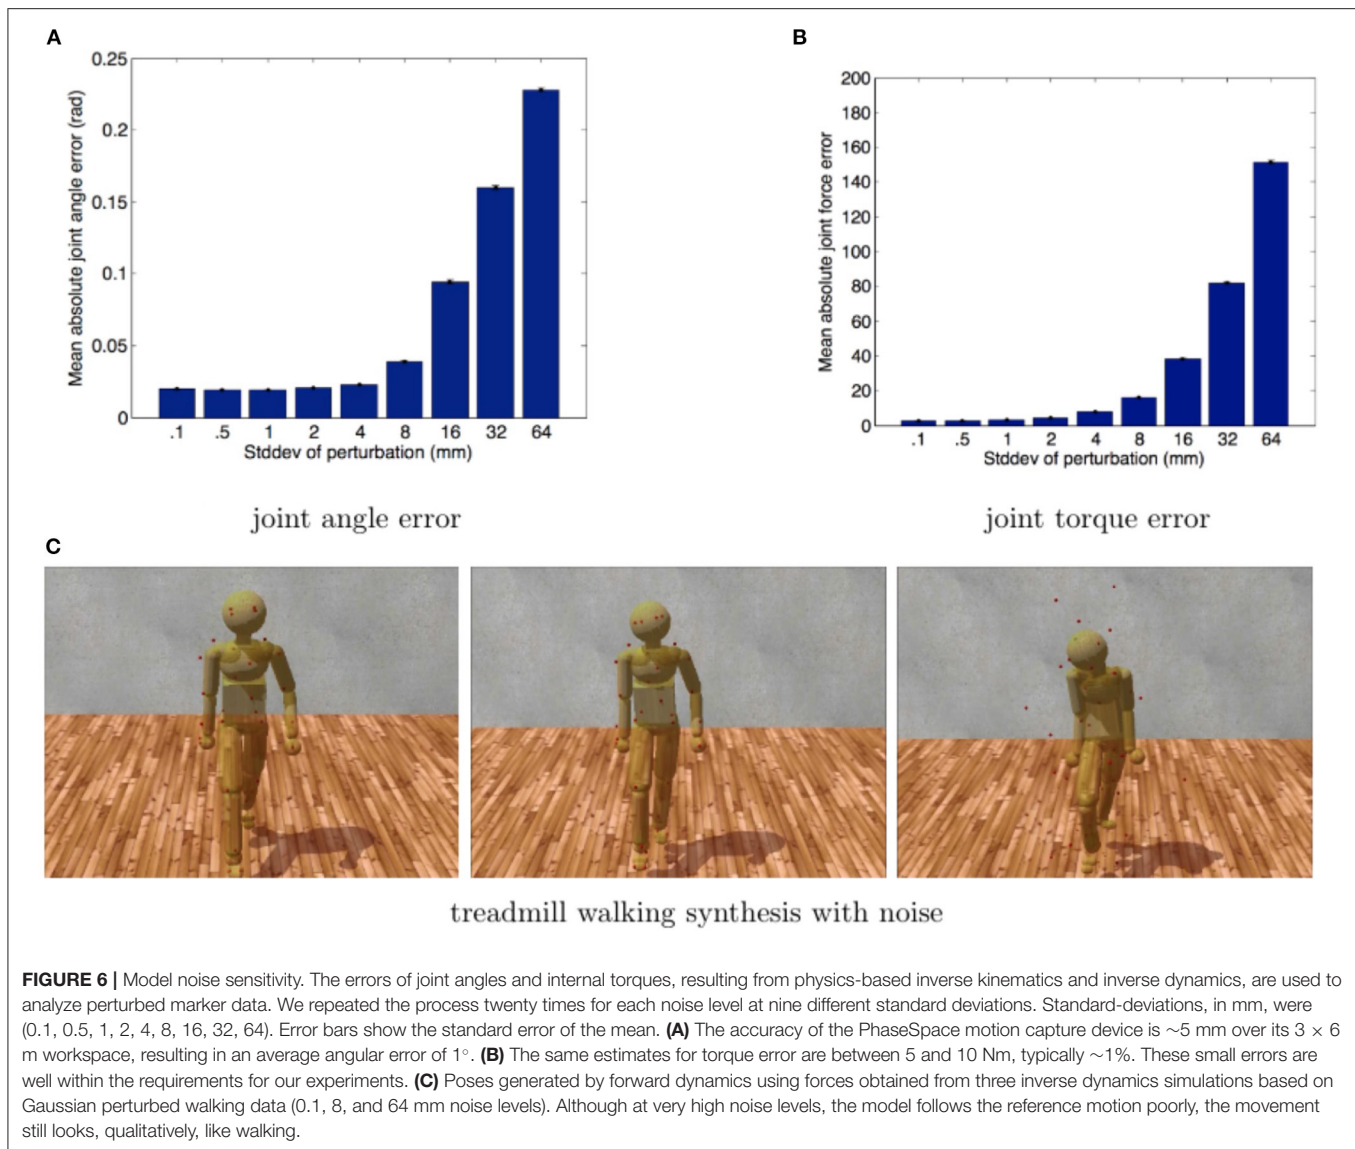
1. Motion synthesis step: it represents human motion in terms of motion capture data.
2. Inverse kinematics step: the ODE built-in functions are called to compute the joint angular velocities and joint angles at each frame.
3. Forward kinematics step: human motions are simulated based on the joint angles and angular velocities acquired from the previous step. This step is to check the correctness of recovered kinematic properties.
4. Inverse dynamics step: ODE built-in functions are called to calculate the required joint torques.
5. Forward dynamics step: human motions are simulated based on the computed torques and residual forces. This step is to check the recovered dynamic properties against the original motions.

At each frame, instantaneous power was computed from the product of net joint torque and joint angular velocity. The work performed at each joint was determined by numerically integrating the instantaneous powers over the entire tracing task. In this way, the energy cost of human motions can be computed given motion capture data. The overall idea behind the method for calculating joint torques/angles is straightforward using our implementation in ODE. The mathematics underlying the rigid body simulation software used in our work is explained in the Appendix section.

3. RESULTS—HDM BASIC CAPABILITIES

The HDM is a fast, robust, intuitive, and inexpensive multi-purpose tool for simulating, analyzing, and synthesizing humanoid movement. **Figure 5** shows a frame of tracing movements collected from the virtual reality experiment (Liu et al., 2019). Users can specify the configurations of the human dynamic model via a multi-purpose graphical interface⁹ for analyzing movement data captured through interaction with the virtual environment. With this tool, it is possible to manually fit

⁹HDM UI Demo, <https://youtu.be/ASs4W05PQcM>



a model to motion capture data, dynamically adjust parameters to test different effects, and visualize the results of kinematic and dynamic analysis. In the particular experiment shown, a subject's fitted model traces a virtual curve to generate a kinematics posture trajectory that allows the cost of the subject's dynamics model to be calculated.

This section focuses on describing the model's capabilities through a series of examples in different settings. Several test experiments provide qualitative and quantitative validation of the physics-based movement analysis techniques described here.

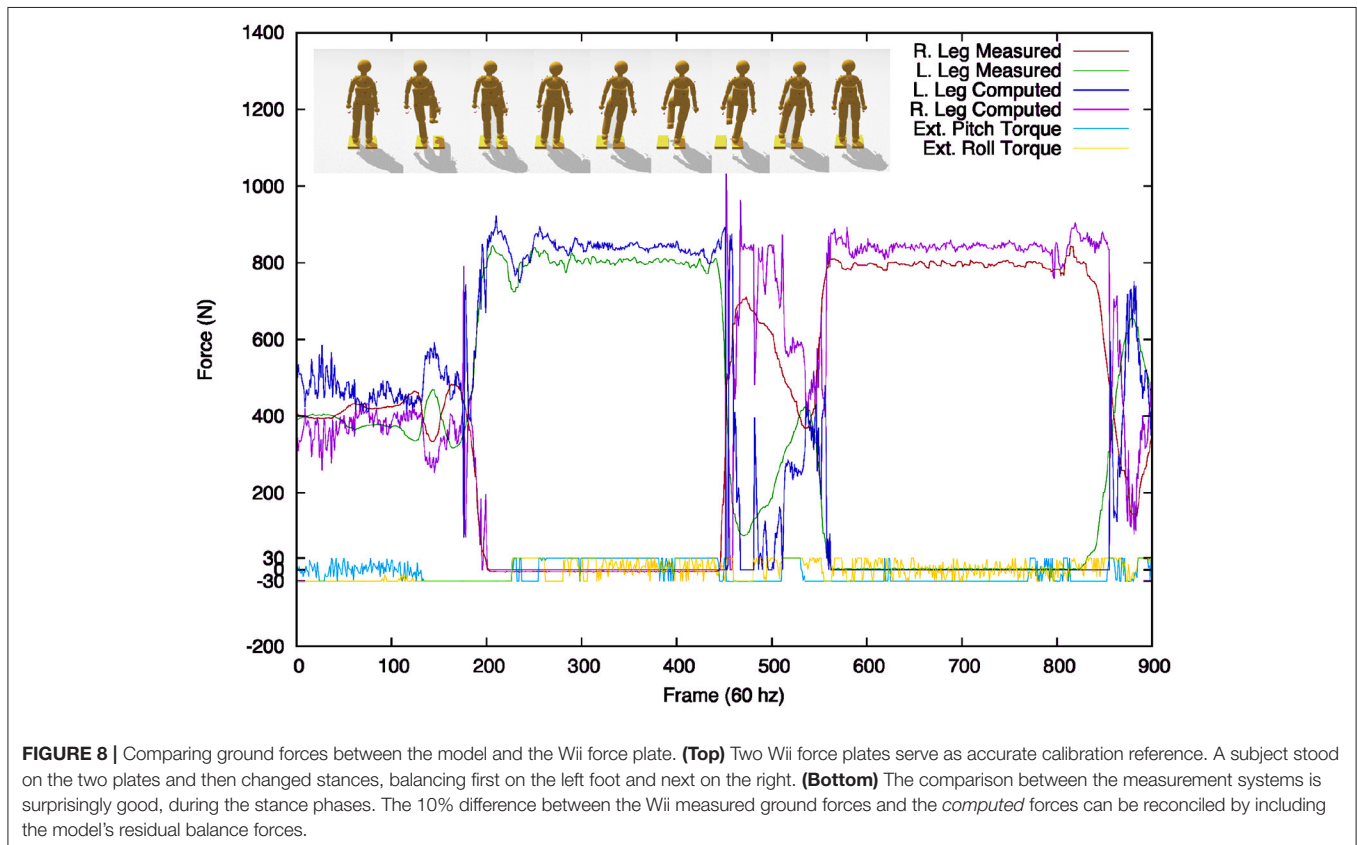
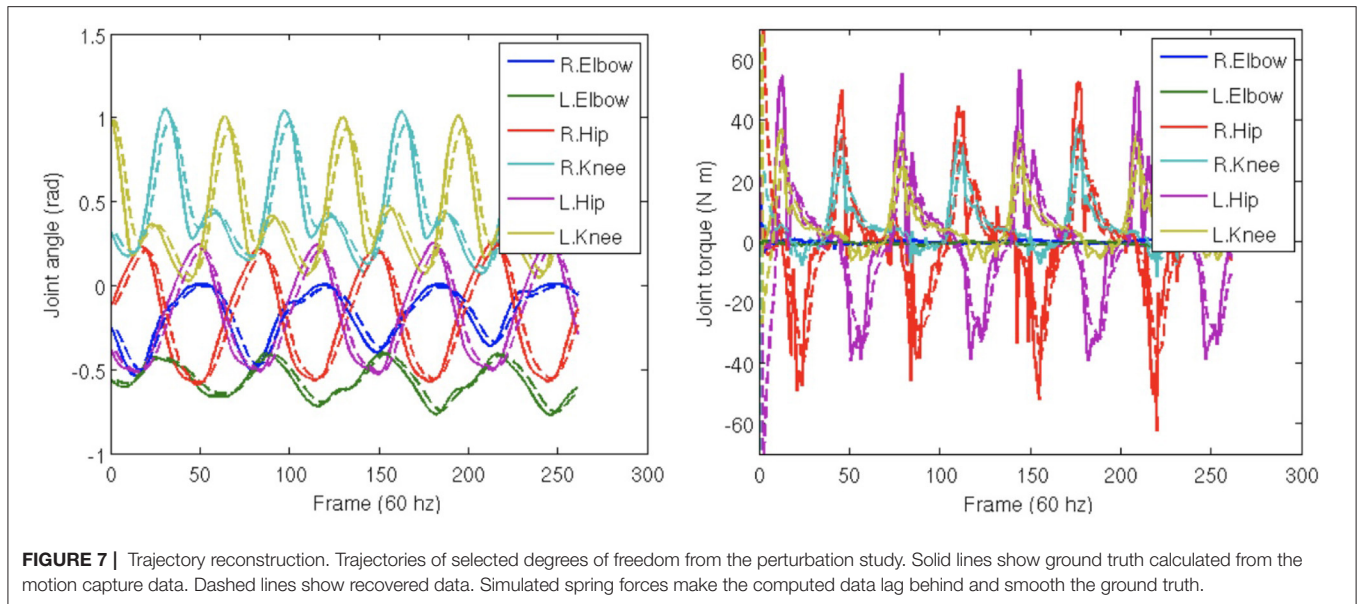
3.1. HDM Accuracy

Given that the torque recovery technique will be the basis for our experiments, it is essential to establish its accuracy in absolute terms. A straightforward way to do this is to use a *particular model* to generate joint torque data and then verify that these generating torques can be recovered with sufficient accuracy. To

test the model accuracy and noise sensitivity, we first use the PhaseSpace motion capture system to gather the walking data and then let the model simulate the walking motion. To simulate possible sensor errors in the PhaseSpace system, we introduce noise into the simulated marker positions and study recovery accuracy with increasing noise levels.

3.1.1. Model Data Sensitivity Tolerance

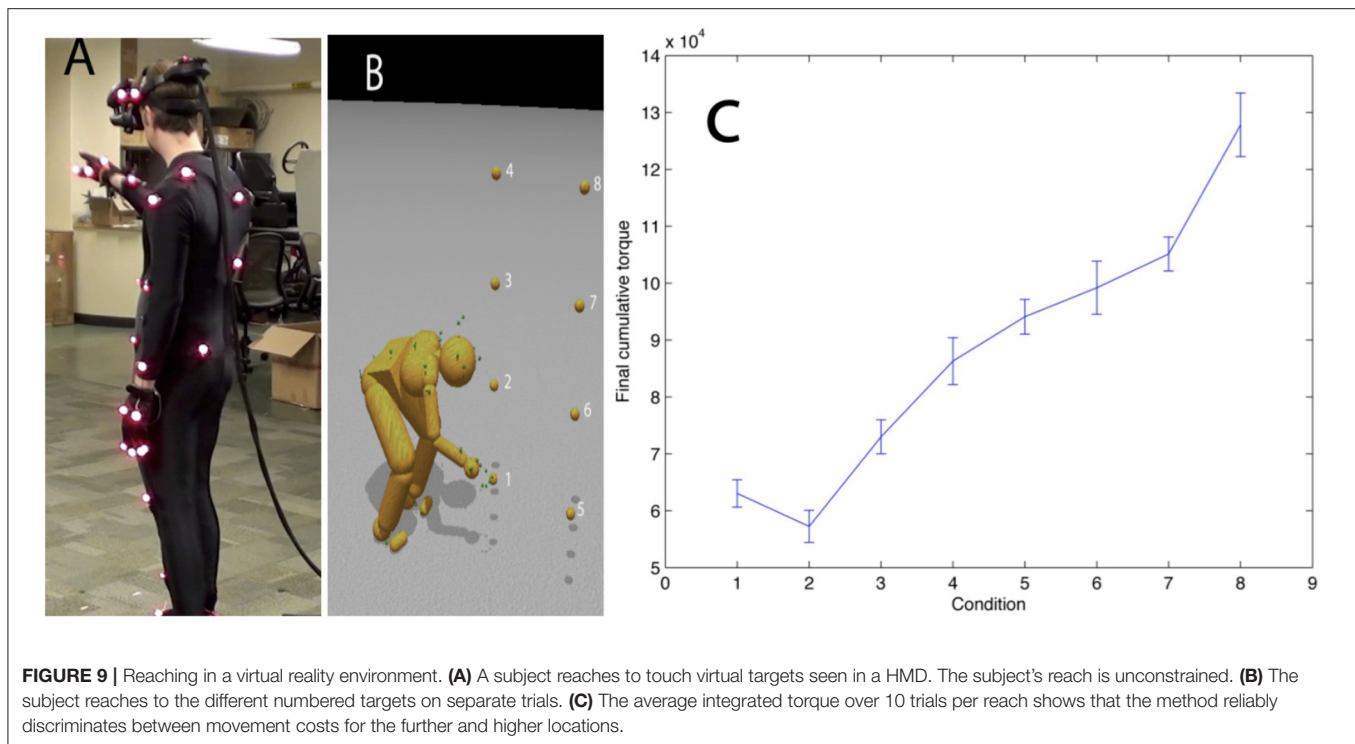
Inverse dynamics computations rely on first finding the model's pose. Therefore, given motion capture data, it is essential to synthesize the pose sequence precisely. We used the HDM to synthesize treadmill walking and then compute its accuracy. This study aimed to assess the effect of sensor perturbations on the results and compare the joint angles and torques found with our method to those used to generate marker data. We used an experimental process similar to that employed in Remy and Thelen (2009). In this experiment, both steps were tested by



studying eight steps of marker data captured from treadmill walking. The movement lasts a little longer than 4 s, giving us 260 frames of data. For this computation, we used data arbitrarily sampled at 60 Hz.

We used a preliminary pass through the motion capture data to generate synthesized “ground truth” marker, pose, and

torque. After using the physics-based inverse kinematics to compute joint angles, we constrained the body to use inverse dynamics to reproduce the joint angles with internal torques (and residual forces at the waist segment). As the model performed the movement, we recorded the global position of the marker attachment points. We also recorded the forces used and the



resulting joint angles. Thus we had synthetic “ground truth” data directly from the model. The model is capable of falling over, when it did not, we recorded the test as a success.

Using the synthetic marker data, we analyzed the process by perturbing all marker positions at each frame in time along all three axes with mean-centered Gaussian noise of a controlled standard deviation. Applying physics-based pose-fitting followed by inverse dynamics produced a new set of virtual marker positions, joint angles, and torques. The results are shown in **Figure 6**.

Gaussian perturbations render the marker data dynamically inconsistent. This dynamic inconsistency also pushes a constrained system toward singularity, making it more challenging to solve numerically. We included very high noise levels to see if they would slow the system down or prevent it from finding any solution. In all cases, the system analyzed the perturbed data and found poses to fit the marker data.

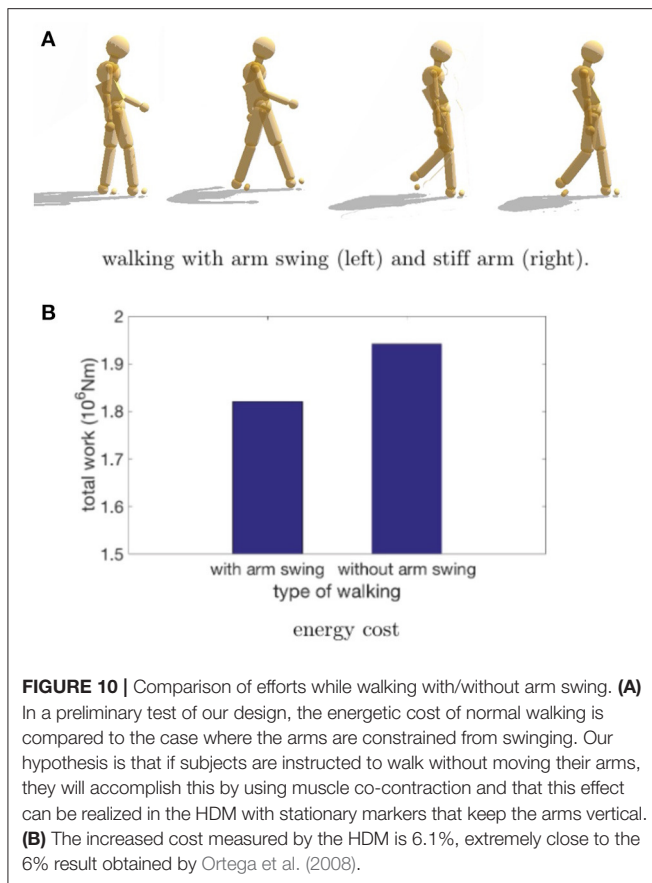
After running through an inverse kinematics pass, an inverse dynamics pass for each trial runs, we compared the marker attachment points, joint angles, and joint torques from the second pass to the synthetic ground truth data. **Figure 6** shows the mean error across all degrees of freedom from eight steps walking. Although the perturbations make the marker data dynamically inconsistent, small amounts of noise have minimal effect on the computed measurements. **Figures 6A,B** standard error deviations show that functional effects are minimal up to 8 mm. A ± 1 mm PhaseSpace marker position accuracy translates in our model into an average joint angle error of 0.02 radians and average force errors of 3 Newtons, which is acceptable for our experiments.

There is a systematic error in both the marker positions and joint angles caused by the fact that the constraints behave like springs. The spring-like behavior causes the marker positions and joint angles to lag behind their targets by a small amount and dampens the overall movement. This lag and damping are apparent in **Figure 7** comparing individual trajectories for selected dimensions of the joint angles and torques. As shown in **Figure 7**, the data follow ground truth very well under low noise conditions.

3.1.2. Residual Torques/Forces and Ground Forces

In the classical inverse dynamic area, discrepancies between the model and humans that created the data necessitate non-realistic “residual” forces to keep the model from falling over when dynamically reproducing most movements. The HDM includes a joint to the model’s waist to constrain it to reproduce orientation deviations found during the pose-fitting pass. To minimize the effect of these external forces, we used torque limits on the amount of stabilizing torque available.

The fully configured system could be tested against an objective set of measurements. We compared HDM data together with ground force data from a pair of balance boards. **Figure 8** shows the calibration of the ground force computed from our method compared to those taken from *Wii*TM force plates. A subject standing on two force plates varied their stance from one being supported exclusively by leg standing on one plate and then shifted their weight to the other leg supported by the other plate. For this simple movement of transitioning from standing on one foot or the other, residual angular torques of 30 Nm



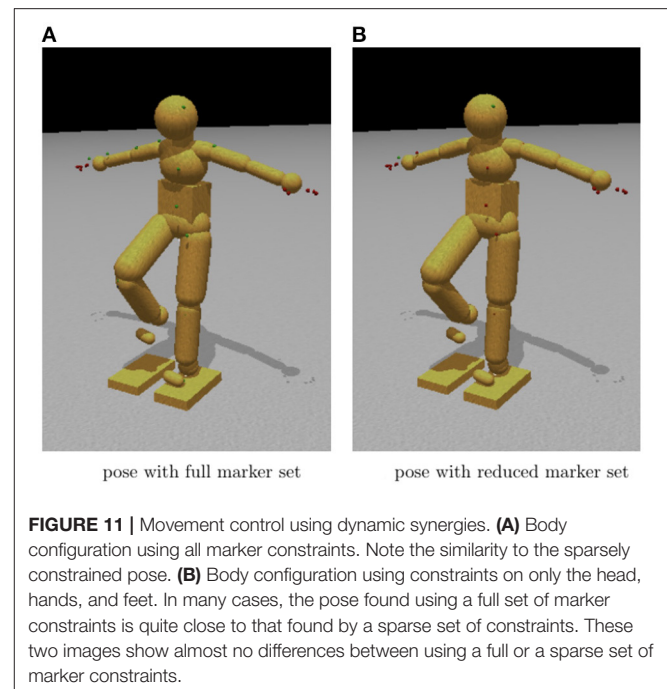
were sufficient to keep the dynamic model quite close to its target trajectory.

The residual torques are very modest, being within $\pm 5\%$ of the maximum excursion. The correspondence is a little better as the faux vestibular balance forces are not factored into the comparison. Note also that we cannot expect the correspondence to be exact during the phase between the two stances as there is no attempt in the model in this test to make the dynamics of the changing stance match that of the force plates. To generate independent movements, such as grasping might need additional accuracy (Sentis et al., 2010), but for estimating a subject's energetic cost, the accuracy is well within range.

Figure 8 also shows the comparison results between the sensor-measured ground forces for the right and left feet (red and green lines) with the computed ground forces found through physics-based inverse dynamics (blue and pink lines). During the bipedal stance phase, the forces come surprisingly close. The most significant discrepancies come during the transition from one foot to the other. These discrepancies can be blamed mainly on poor collision detection resulting from an abstract model of the foot.

3.2. HDM Experimental Evaluation

The previous demonstrations report on tests of the fundamental performance of the system. This section focuses on three tests



of the HDM's ability to fit experimental data. The first test uses a subject carrying out successively more difficult reaches in a virtual reality environment to test whether the model's estimate of movement costs correlates with increasing task difficulty. The second test simulates data from an issue facing movements in an aging population. Do aging subjects' reduced use of arm swing while walking incur a movement cost, and does the HDM's estimate correspond to laboratory treadmill data? The final test demonstrates an essential property of the model concerning its degrees of freedom. The critical observation is that virtues of their interconnections constrain the degree of freedom of the model; thus, the control of a posture can be achieved with a significantly reduced set of crucial marker positions. This property has implications for movement control programs.

3.2.1. Experiment 1: Whole Body Reaching

The movement accuracy test is encouraging, but the importance of the method depends on its usefulness to capture the energetic cost of whole-body movements in a complex experimental setting. One such venue is a three-dimensional Virtual Reality (VR) environment. The advantage of the VR environment for studying human movements is that the dimensions and the dynamic variations of the parametric quantities describing the setting can be varied with full experimental control.

In this experiment, we studied how a subject chose whole-body movements in reaching targets. **Figure 9** demonstrates the experimental setup. The subject wears the PhaseSpace motion capture suit and the nVisor head-mounted stereo display. In each trial, the subject starts from a particular starting position marked with tape on the lab floor, approaches a target suspended in 3D space, and finally touches it.

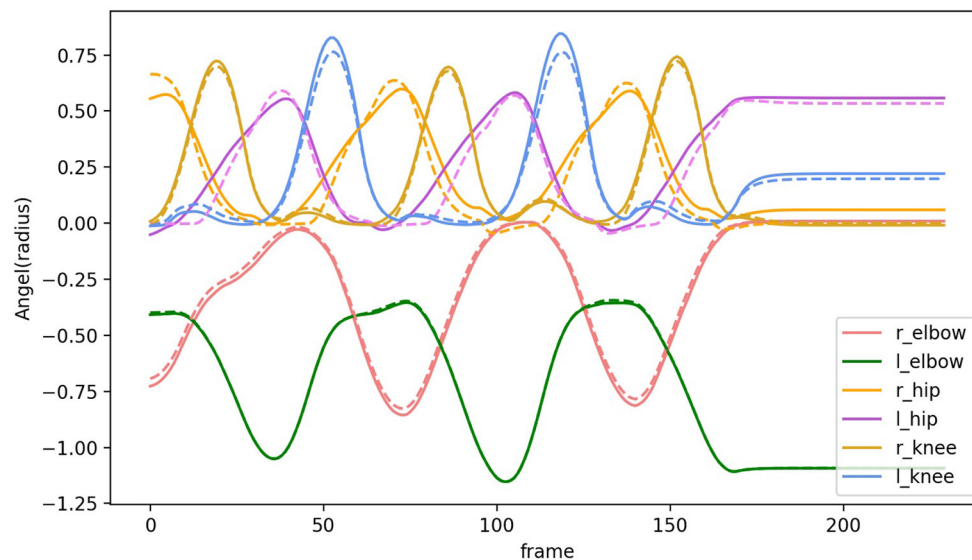


FIGURE 12 | Comparison of joint angles along the selected degree of freedom. Solid lines show joint angles recovered based on the full marker set. Dashed lines show joint angles recovered based on a reduced marker set.

Tests were able to establish that, just focusing on integrated net torque and avoiding stiffness, the relative cost of a movement recorded by our system reliably discriminates the energetic costs of the movement in the way hypothesized. The hypothesized cost of reaching for and touching each target was ranked based on distance and height relative to the subject. Note that target 2 is the least expensive as the subject does not have to crouch or extend significantly to touch it. Targets 5 through 8 are more costly than targets 1 through 4 as they require that the subject take a step to touch them. These results were expected, but the point was to show that the overall setting and model could produce reliable torque estimates.

This demonstration shows that the model can be used in any setting where the cost of a movement is hypothesized to be a constituent factor. We develop this technique further in the next demonstration.

3.2.2. Experiment 2: The Cost of Stiff Arm Walking

Once the joint stiffness parameters were adjusted appropriately, can it reproduce the results of a stiffness modulating experiment? The experiment we tried was to replicate that of Ortega et al. (2008). They showed that arresting the arm swing during treadmill walking incurred an increased metabolic cost of 6%. Our hypothesis was that to reproduce this result we could modify our walking data for the model so that the arms were clamped by the sides with stiff stationary markers.

To test this feasibility, we used one of our HDM walking data sets in a test situation. The cost of walking was computed with a modification designed to model the experimental protocol in Ortega et al. (2008). To simulate their experiment, we modified the model data so the arms could swing with the walking gait for the standard case, but for the restricted case, the arms were constrained by markers that move with the stride but are not

allowed to swing. Since the arms under restricted situations could not balance the leg movements, we expected the energetic cost to be higher. As shown in **Figure 10**, the result was that the constrained walk was about 6% more expensive than the standard walk, which was essentially the value obtained by the Farley lab (Ortega et al., 2008). The use of the HDM in imitating this experiment shows off the model's utility; no elaborate tuning was necessary to obtain the preliminary result other than restraining the arms. It should be appreciated, that the model experiment required only hours to program, whereas the human experiment takes many days to set up, and the corresponding measure of cost is delicate.

3.2.3. Experiment 3: Controlling Poses Using Reduced Marker Sets

Human pose sequences from simple single-behavior motions lie on a very low-dimensional linear subspace (Barbič et al., 2004). Previous research showed that for many movements, with suitable internal stiffness, it is only necessary to control the location of a reduced set consisting of the head, hands, and feet markers (Liu et al., 2006). This observation is the centerpiece of *uncontrolled manifold theory*, which restrict control to a subspace of the degrees of freedom, leaving the rest to the natural system dynamics (Scholz and Schönner, 1999; Torres-Oviedo and Ting, 2007). An obvious comes from ice skating. Pairs spinning on the ice would use one set of markers while jumping would use another set. Another example is using a subset of the markers to constrain the dynamics still produces reasonable walking gaits. This property could also have been expected from studies of muscle synergies, which show that muscle contractions coordinate in movement generation (Ting and Macpherson, 2005; Ting, 2007). As described above, the HDM uses 41 markers, which means a pose is represented

by a 123-dimension coordinate system. Our test compares the generated pose using a reduced marker set with the one using the full marker set. Tests of movement accuracy revealed that the dynamics engine was able to function with significant fewer numbers of markers. Another benefit is that fewer markers can make interpolating across drop-outs in the data easier.

Figure 11 shows a qualitative comparison between a pose found using the whole marker set (on the left) and one found using only head, hands, and feet (on the right). To achieve the reduced marker pose, we started the model in an upright stance with the arms by the side, and then the reduced set markers are moved slowly along trajectories that leave them in the final posture. The straight arms take advantage of the elbow joint angle limitation.

Joint limits on the knees and elbows and general joint stiffness naturally bias the physics engine to find a pose that is very close to the fully constrained pose. Body inertia and joint stiffness naturally clean up minor noise and occlusions in the captured marker data. The resulting joint angles in transit allow the specification of the complete set of dynamic torques. To test this feature of HDM quantitatively, the recovered joints angles while walking according to the reduced marker set were compared with those from the full marker set. **Figure 12** illustrates the recovered joints angles are quite similar with the original joints angles.

This result has important general implications. First of all, the finding suggests that the kinematic plan for movements can be compressed into a subset of formative trajectories, leaving the remaining degrees of freedom interpolated using the body's dynamic constraint. Another aspect of this observation is that the reduced set can be used to adjust movements to individual circumstances, again leaving the detailed interpolation to the dynamics.

4. CONCLUSION AND DISCUSSION

This paper provides a method of building a humanoid model on top of a physical engine that can analyze and synthesize human movements. Its 48 degrees of freedom and generalized spring constraints allow models of scale that are robust to disturbances. In addition to being an analytical tool for experiments, it can also generate movements from a kinematic plan. The system has several features:

1. It utilizes the realizations of constraints as implicit springs. The spring parameters exhibit many advantageous properties. They stabilize the simulation, pushing a constrained system away from singularities, and reducing constraint error.
2. It calculates the movements' energetic cost to provide the capability to compare different movement scenarios. Achieving this goal can be tricky, owing to the lack of systems that can provide independent cost measures. The HDM achieved an excellent correspondence with force plates, as shown in the experiment to measure the human's stance change.
3. It produces correlations with similar tests with human subjects, such as our research with stiff-arm walking. Once we have vetted the system in many such areas, it can be used as

a predictive tool, as in the experiment showing the different costs of reaching targets. We have developed a large-scale three-dimensional tracing experiment in virtual reality (Liu et al., 2019) to elicit natural whole-body movements under common goals. Our future work is to analyze the energetic cost using the HDM.

4. It shows that such a model can play a valuable role in studying the kinematic-plan model's consequences. In particular, the reduced degree of freedom control demonstration supports the *uncontrolled manifold* view wherein a subset of crucial degrees of freedom can direct a movement with the uncontrolled degrees of freedom interpolating the movement using the system's dynamics (Carpenter, 1968; Latash, 2008).

One way of illustrating the method's robustness is to combine a kinematic data set from the source with another set of dynamic parameters. So far, we have explored the HMD capability of using data from two other laboratories. One source is Carnegie Mellon University's graphics laboratory's motion capture database¹⁰. The other source is the motion data from the Hayhoe laboratory at the University of Texas psychology department, motion capture data from subjects traversing rough outdoor terrain. In tests, the data gathered with a different motion capture device is combined with the inertial data from another model to make a composite. We could use our dynamics calculation to compute joint torques for the hybrid system by adopting the imported database's marker conventions. Although the estimate is thus done for a synthetic pairing of kinematic data and dynamic parameters, this combination, the integration is stable and leads to identifiable torques.

Besides its use of a mechanism for interpreting experiments, the system can also serve as an adjunct for theorizing about the human system's organization concerning its space-time performance since many of these issues are open. While an enormous amount of research in human motor control has produced ever more refined subsystem components' elucidations, a comprehensive theory at the level of large-scale dynamics is still unsettled.

One question is a description of how the motor cortex can code information to drive the high temporal bandwidths of the spinal cord circuitry. Several possibilities were debated at the *Neural Control of Movement* conference in 2013 without definitive results. We have emphasized that the motor cortex communicates a coded kinematic plan together with stiffness settings that play an essential role in shaping the dynamics in muscle spindles (Blum et al., 2020).

Studies with kinematics coded with temporal basis functions have shown that a kinematic plan can be coded to reduce the bandwidth needed by a factor of $\sim 10^3$ Hz (Iyer and Ballard, 2011; Won et al., 2020). More recently, sparse coding has been used to solve this problem (Glanz et al., 2021), despite unnecessary caution of its applicability to cortical motor areas (Beyeler et al., 2019). A tack that is still open would be our large sets of kinematic data to attempt to model cell responses

¹⁰CMU Graphics Lab Motion Capture Database: <http://mocap.cs.cmu.edu/>

in cortical area M1 to see if they turn out to be correlated with Graziano's data. The expectation would be that cells might have an informational interpretation analogous to the one used for oriented calls in the striate cortex. The facility with which the HDM system can collect motor data would allow a ready exploration of motor coding strategies similar to the visual domain. Being able to store movement segments would obviate the need to use an unlikely option to generate dynamic codes in the cortex neurally. This observation finds support in neural recordings that show movements generated cortically in discrete phases (Zimnik and Churchland, 2021), and the Boston Dynamics Atlas robot also uses a look-up strategy (Feng et al., 2014).

Regarding the uncontrolled manifold concept, a critical insight was the use of reduced degrees of freedom constraints in computing the dynamics. If the limitations are near the number of DOFs of the system, the torque recovery can quickly become numerically unstable. However, between 20 and 41 markers in the HDM provide sufficient constraints to integrate the dynamic equations reliably by allowing the system's natural dynamics to interpolate the motion appropriately.

In summary, the method has several advantages over alternative methods:

1. It can be easily implemented in a single robust framework of the physics engine. Using the physics engine for multiple tasks allows a unique human model to be used from start to finish, rather than being forced to use the conventions built into a commercial package.
2. The method is fast. The simulation engine is designed for performance, making it possible to analyze movement captured in real-time and create interactive experiments with stimuli dependent on the feedback results.
3. The software is free. Freely accessible code, such as ODE, is useful because it facilitates comparison and collaboration in research. Fourth, the method handles multiple ground contacts and noisy data challenging to related approaches. Kinematic loops do not require any special treatment. The method is robust even to large perturbations making data dynamically inconsistent.
4. The tunable parameters (CFM), couched in the physics framework, are intuitive. It is more straightforward to specify the importance of a constraint in force and mass rather than arbitrary gains and weightings. We illustrate these advantages by using ODE to analyze and reproduce movement recorded from optical motion capture.

REFERENCES

- Andriacchi, T. P., and Alexander, E. J. (2000). Studies of human locomotion: past, present and future. *J. Biomech.* 33, 1217–1224. doi: 10.1016/S0021-9290(00)00061-0
- Atkeson, C. G., and Schaal, S. (1997). "Robot learning from demonstration," in *ICML, Vol. 97* (Vienna), 12–20. Available online at: https://link.springer.com/chapter/10.1007/978-3-7091-6874-5_4
- Awrejcewicz, J., Kudra, G., and Zagrodny, B. (2012). Nonlinearity of muscle stiffness. *Theor. Appl. Mech. Lett.* 2:053001. doi: 10.1063/2.1205301
- Badler, N. I., Phillips, C. B., and Webber, B. L. (1993). *Simulating Humans: Computer Graphics Animation and Control*. New York, NY: Oxford University Press. doi: 10.1093/oso/9780195073591.001.0001
- Bailly, F., Carpentier, J., and Souères, P. (2021). "Optimal estimation of the centroidal dynamics of legged robots," in *ICRA 2021-IEEE International Conference on Robotics and Automation* (Xi'an). Available online at: <https://hal.archives-ouvertes.fr/hal-03193940/document>
- Bajrami, X., and Murturi, I. (2017). Kinematic model of the seven link biped robot. *IJMET* 8, 454–462.

There are several ways to improve the system in the future, and the following two are the most important among them. One limitation of our method for computing torque is that it is insensitive to muscle stiffness, which is both passive and can be actively modulated (Atkeson and Schaal, 1997; Awrejcewicz et al., 2012). Increasing stiffness will increase the overall net movement energetic cost and needs to be taken into account. The observation somewhat ameliorates this issue that subjects will try to minimize energetic costs in most natural tasks and thus exploit natural dynamics whenever they can, reducing high levels of co-contraction (Carpenter, 1968; Shadmehr and Arbib, 1992; Sternad and Sternad, 2009). However, the ubiquitous use of spring as a constraint means opening up the possibility of adding springs to the joint degrees of freedom to model stiffness. These could also have parametric programmable spring constants to model muscle co-contraction. The second feature that could be added is an improved system to keep the human model upright. Any of the three human sources of this needed information—visual, vestibular, and proprioception—would be candidates for this practical constraint. At present, the HDM uses a faux system of rotational torques at the center of gravity, but these could easily be replaced with more appropriate ankle torques.

AUTHOR CONTRIBUTIONS

JC created the initial protocol of the dynamic model and did some primary validations. LL further developed the model, conducted the experiments, and wrote the paper. DB designed the experiments and worked with LL on the overall paper editing and scientific presentation.

FUNDING

This research was supported by the National Institutes of Health EY05729.

ACKNOWLEDGMENTS

Thanks to John Matthis and Leif Johnson for helpful discussions.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fnbot.2021.723428/full#supplementary-material>

- Barbič, J., Safonova, A., Pan, J.-Y., Faloutsos, C., Hodgins, J. K., and Pollard, N. S. (2004). "Segmenting motion capture data into distinct behaviors," in *Proceedings of Graphics Interface 2004* (Ontario, CA), 185–194. Available online at: https://www.researchgate.net/publication/221474944_Segmenting_Motion_Capture_Data_into_Distinct_Behaviors.
- Beyeler, M., Rounds, E. L., Carlson, K. D., Dutt, N., and Krichmar, J. L. (2019). Neural correlates of sparse coding and dimensionality reduction. *PLoS Comput. Biol.* 15:e1006908. doi: 10.1371/journal.pcbi.1006908
- Blum, K. P., Campbell, K. S., Horslen, B. C., Nardelli, P., Housley, S. N., Cope, T. C., et al. (2020). Diverse and complex muscle spindle afferent firing properties emerge from multiscale muscle mechanics. *Elife* 9:e55177. doi: 10.7554/eLife.55177
- Blümel, M., Hooper, S. L., Guschlbauer, C., White, W. E., and Büschges, A. (2012). Determining all parameters necessary to build hill-type muscle models from experiments on single muscles. *Biol. Cybern.* 106, 543–558. doi: 10.1007/s00422-012-0531-5
- Burdett, R. G., Skrinar, G. S., and Simon, S. R. (1983). Comparison of mechanical work and metabolic energy consumption during normal gait. *J. Orthop. Res.* 1, 63–72. doi: 10.1002/jor.1100010109
- Callahan, D. M., Umberger, B. R., and Kent-Braun, J. A. (2013). A computational model of torque generation: neural, contractile, metabolic and musculoskeletal components. *PLoS ONE* 8:e56013. doi: 10.1371/journal.pone.0056013
- Carpenter, M. (1968). *The Co-Ordination and Regulation of Movements*. Baltimore, MD: Waverly Press doi: 10.1097/00005072-196804000-00011
- Čelikovský, S., and Anderle, M. (2018). Stable walking gaits for a three-link planar biped robot with two actuators based on the collocated virtual holonomic constraints and the cyclic unactuated variable. *IFAC-PapersOnLine* 51, 378–385. doi: 10.1016/j.ifacol.2018.11.571
- Cooper, J. L., and Ballard, D. (2012). "Realtime physics-based marker following," in *International Conference on Motion in Games* (Berlin: Springer), 350–361. doi: 10.1007/978-3-642-34710-8_32
- De Aguiar, E., Theobalt, C., and Seidel, H.-P. (2006). "Automatic learning of articulated skeletons from 3d marker trajectories," in *International Symposium on Visual Computing* (Berlin: Springer), 485–494. doi: 10.1007/11919476_49
- Delp, S. L., Anderson, F. C., Arnold, A. S., Loan, P., Habib, A., John, C. T., et al. (2007). OpenSim: open-source software to create and analyze dynamic simulations of movement. *IEEE Trans. Biomed. Eng.* 54, 1940–1950. doi: 10.1109/TBME.2007.901024
- Dembia, C. L., Bianco, N. A., Falisse, A., Hicks, J. L., and Delp, S. L. (2020). OpenSim moco: musculoskeletal optimal control. *PLoS Comput. Biol.* 16:e1008493. doi: 10.1371/journal.pcbi.1008493
- Demircan, E., Sentis, L., De Sapio, V., and Khatib, O. (2008). "Human motion reconstruction by direct control of marker trajectories," in *Advances in Robot Kinematics: Analysis and Design*, eds J. Lenarčič and P. Wenger (Dordrecht: Springer), 263–272. doi: 10.1007/978-1-4020-8600-7_28
- Durandau, G., Farina, D., and Sartori, M. (2017). Robust real-time musculoskeletal modeling driven by electromyograms. *IEEE Trans. Biomed. Eng.* 65, 556–564. doi: 10.1109/TBME.2017.2704085
- Erez, T., Tassa, Y., and Todorov, E. (2015). "Simulation tools for model-based robotics: comparison of bullet, Havok, Mujoco, Ode and Physx," in *2015 IEEE international conference on robotics and automation (ICRA)* (Seattle), 4397–4404. doi: 10.1109/ICRA.2015.7139807
- Faure, F., Debonne, G., Cani-Gascuel, M.-P., and Multon, F. (1997). "Dynamic analysis of human walking," in *Computer Animation and Simulation'97* (Budapest: Springer), 53–65. doi: 10.1007/978-3-7091-6874-5_4
- Feng, S., Whitman, E., Xinjilefu, X., and Atkeson, C. G. (2014). "Optimization based full body control for the atlas robot," in *2014 IEEE-RAS International Conference on Humanoid Robots* (Boston, MA; Madrid), 120–127. doi: 10.1109/HUMANOIDS.2014.7041347
- Glanz, R. M., Dooley, J. C., Sokoloff, G., and Blumberg, M. S. (2021). Sensory coding of limb kinematics in motor cortex across a key developmental transition. *J. Neurosci.* 41, 6905–6918. doi: 10.1523/JNEUROSCI.0921-21.2021
- Grizzle, J. W., Chevallereau, C., Ames, A. D., and Sinnet, R. W. (2010). 3D bipedal robotic walking: models, feedback control, and open problems. *IFAC Proc.* 43, 505–532. doi: 10.3182/20100901-3-IT-2016.00302
- Ijspeert, A. J., Crespi, A., Ryczko, D., and Cabelguen, J.-M. (2007). From swimming to walking with a salamander robot driven by a spinal cord model. *Science* 315, 1416–1420. doi: 10.1126/science.1138353
- Iyer, R., and Ballard, D. (2011). "Humanoid muscle movement representation," in *2011 11th IEEE-RAS International Conference on Humanoid Robots* (Bled), 409–415. doi: 10.1109/Humanoids.2011.6100892
- Johnson, L., and Ballard, D. (2014). "Efficient codes for inverse dynamics during walking," in *Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 28* (Palo). Available online at: <https://www.aaai.org/Library/AAAI/aaai14contents.php>
- Khusainov, R., Shimchik, I., Afanasyev, I., and Magid, E. (2016). "3D modelling of biped robot locomotion with walking primitives approach in simulink environment," in *Informatics in Control, Automation and Robotics 12th International Conference, ICINCO 2015 Colmar* (Cham: Springer), 287–304. doi: 10.1007/978-3-319-31898-1_16
- Kirk, A. G., O'Brien, J. F., and Forsyth, D. A. (2005). "Skeletal parameter estimation from optical motion capture data," in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05), Vol. 2* (San Diego, CA), 782–788. doi: 10.1109/CVPR.2005.326
- Krishchenko, A., Tkachev, S., and Fetisov, D. (2007). Planar walking control for a five-link biped robot. *Comput. Math. Model.* 18, 176–191. doi: 10.1007/s10598-007-0018-8
- Latash, M. L. (2008). *Synergy*. New York: Oxford University Press. doi: 10.1093/acprof:oso/9780195333169.001.0001
- Lee, S., Park, M., Lee, K., and Lee, J. (2019). Scalable muscle-actuated human simulation and control. *ACM Trans. Graph.* 38, 1–13. doi: 10.1145/3306346.3322972
- Lee, T.-T. (1988). "Trajectory planning and control of a 3-link biped robot," in *IEEE International Conference on Robotics and Automation*, 820–823. doi: 10.1109/ROBOT.1988.12161
- Liu, G., Zhang, J., Wang, W., and McMillan, L. (2006). "Human motion estimation from a reduced marker set," in *Proceedings of the 2006 Symposium on Interactive 3D Graphics and Games*, 35–42. doi: 10.1145/1111411.1111418
- Liu, L., Johnson, L., Zohar, O., and Ballard, D. H. (2019). Humans use similar posture sequences in a whole-body tracing task. *IScience* 19, 860–871. doi: 10.1016/j.isci.2019.08.041
- Loeb, G. E., and Tsianos, G. A. (2015). Major remaining gaps in models of sensorimotor systems. *Front. Comput. Neurosci.* 9:70. doi: 10.3389/fncom.2015.00070
- Mousavi, P. N., and Bagheri, A. (2007). Mathematical simulation of a seven link biped robot on various surfaces and zmp considerations. *Appl. Math. Model.* 31, 18–37. doi: 10.1016/j.apm.2006.06.018
- Mu, X., and Wu, Q. (2003). Synthesis of a complete sagittal gait cycle for a five-link biped robot. *Robotica* 21, 581–587. doi: 10.1017/S0263574702004903
- Mu, X., and Wu, Q. (2004). "Sagittal gait synthesis for a five-link biped robot," in *Proceedings of the 2004 American Control Conference, Vol. 5*, Boston 4004–4009.
- Muybridge, E. (1887). *Animal Locomotion*. Oxford: Da Capo Press.
- Ortega, J. D., Fehلمان, L. A., and Farley, C. T. (2008). Effects of aging and arm swing on the metabolic cost of stability in human walking. *J. Biomech.* oxford, 41, 3303–3308. doi: 10.1016/j.jbiomech.2008.06.039
- Rajagopal, A., Dembia, C. L., DeMers, M. S., Delp, D. D., Hicks, J. L., and Delp, S. L. (2016). Full-body musculoskeletal model for muscle-driven simulation of human gait. *IEEE Trans. Biomed. Eng.* 63, 2068–2079. doi: 10.1109/TBME.2016.2586891
- Remy, C. D., and Thelen, D. G. (2009). Optimal estimation of dynamically consistent kinematics and kinetics for forward dynamic simulation of gait. *J. Biomech. Eng.* 131:031005. doi: 10.1115/1.3005148
- Scholz, J. P., and Schöner, G. (1999). The uncontrolled manifold concept: identifying control variables for a functional task. *Exp. Brain Res.* 126, 289–306. doi: 10.1007/s002210050738
- Schulman, J., Ho, J., Lee, C., and Abbeel, P. (2016). "Learning from demonstrations through the use of non-rigid registration," in *Robotics Research*, eds M. Inaba M. and P. Corke (Cham: Springer), 339–354. doi: 10.1007/978-3-319-28872-7_20
- Sentis, L., Park, J., and Khatib, O. (2010). Compliant control of multicontact and center-of-mass behaviors in humanoid robots. *IEEE Trans. Robot.* 26, 483–501. doi: 10.1109/TRO.2010.2043757

- Seth, A., Sherman, M., Reinbolt, J. A., and Delp, S. L. (2011). Opensim: a musculoskeletal modeling and simulation framework for *in silico* investigations and exchange. *Proc. Iutam* 2, 212–232. doi: 10.1016/j.piutam.2011.04.021
- Shadmehr, R., and Arbib, M. A. (1992). A mathematical analysis of the force-stiffness characteristics of muscles in control of a single joint system. *Biol. Cybern.* 66, 463–477. doi: 10.1007/BF00204111
- Sternad, D., and Sternad, D. (2009). *Progress in Motor Control*. Springer. doi: 10.1007/978-0-387-77064-2
- Ting, L. H. (2007). Dimensional reduction in sensorimotor systems: a framework for understanding muscle coordination of posture. *Prog. Brain Res.* 165, 299–321. doi: 10.1016/S0079-6123(06)65019-X
- Ting, L. H., and Macpherson, J. M. (2005). A limited set of muscle synergies for force control during a postural task. *J. Neurophysiol.* 93, 609–613. doi: 10.1152/jn.00681.2004
- Torres-Oviedo, G., and Ting, L. H. (2007). Muscle synergies characterizing human postural responses. *J. Neurophysiol.* 98, 2144–2156. doi: 10.1152/jn.01360.2006
- Tournier, M., Wu, X., Courty, N., Arnaud, E., and Reveret, L. (2009). “Motion compression using principal geodesics analysis,” in *Computer Graphics Forum, Vol. 28* (Wiley), 355–364. doi: 10.1111/j.1467-8659.2009.01375.x
- van der Kooij, H., van Asseldonk, E., and van der Helm, F. C. (2005). Comparison of different methods to identify and quantify balance control. *J. Neurosci. Methods* 145, 175–203. doi: 10.1016/j.jneumeth.2005.01.003
- Welch, T. D., and Ting, L. H. (2014). Mechanisms of motor adaptation in reactive balance control. *PLoS ONE* 9:e96440. doi: 10.1371/journal.pone.0096440
- Wolpert, D. M., and Landy, M. S. (2012). Motor control is decision-making. *Curr. Opin. Neurobiol.* 22, 996–1003. doi: 10.1016/j.conb.2012.05.003
- Won, J., Gopinath, D., and Hodgins, J. (2020). A scalable approach to control diverse behaviors for physically simulated characters. *ACM Trans. Graph.* 39:33. doi: 10.1145/3386569.3392381
- Zimnik, A. J., and Churchland, M. M. (2021). Independent generation of sequence elements by motor cortex. *Nat. Neurosci.* 24, 412–424. doi: 10.1038/s41593-021-00798-5
- Zordan, V. B., and Hodgins, J. K. (2002). “Motion capture-driven simulations that hit and react,” in *Proceedings of the 2002 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (New York, NY), 89–96. doi: 10.1145/545261.545276
- Zordan, V. B., and Van Der Horst, N. C. (2003). “Mapping optical motion capture data to skeletal motion using a physical model,” in *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (Goslar), 245–250. Available online at: <https://dl.acm.org/doi/proceedings/10.5555/846276>

Conflict of Interest: JC is employed by company Google Inc.

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher’s Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2021 Liu, Cooper and Ballard. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.