Volume 0 (1981), Number 0 pp. 1–12 COMPUTER GRAPHICS *forum*

Multi-Segment Foot for Human Modeling and Simulation

Hwangpil Park, Ri Yu, and Jehee Lee

Seoul National University, South Korea

Abstract

Realistic modeling of a human-like character is one of the main topics in computer graphics to simulate human motion physically and also look realistically. Of the body parts, a human foot interacts with the ground, and plays an essential role in weight transmission, balancing posture, and assisting ambulation. However, in the previous researches, the foot model was often simplified into one or two rigid bodies connected by a revolute joint. We propose a new foot model consisting of multiple segments to reproduce human foot shape and its functionality accurately. Based on the new model, we develop a foot pose controller that can reproduce foot postures that are generally not obtained in motion capture data. We demonstrate the validity of our foot model and the effectiveness of our foot controller with a variety of foot motions in a physics-based simulation.

CCS Concepts

• ***Computing methodologies*** → ***Physical simulation;***

1. Introduction

Reproducing realistic movement of a character in a physics-based simulation is a primary issue in computer graphics and robotics societies. Realistic modeling of characters not only makes their appearance look realistic but also helps simulated behavior to work more physically feasible.

The foot is the body part that directly interacts with the ground when standing, walking, or running. The human foot contains 26 bones and 33 joints, and more than a hundred muscles and ligaments. These components form a complicated structure of the foot and move organically to perform various foot functions. For example, the foot can alleviate the shock felt when walking on the ground, provide balance by changing the way it supports the body when standing, and gain momentum by pushing the ground to increase speed.

Despite these functions of the foot when controlling a biped character, the importance of foot modeling has been overlooked, and foot models have not been developed well. As a result, the foot models of biped characters used in physics simulations have mostly consisted of one or two boxes. These types of models cannot reproduce a human foot shape and motions enough because they lack sufficient degrees of freedom (DoFs).

In this paper, we propose a new human foot model based on real human foot anatomy, and an algorithm that changes the shape of the foot model depending on human motion. By using our new foot model and the foot shape controlling method, we can represent the foot shapes and motions better than the existing foot models.

The main contribution of our work is that we created a new foot model that consists of several segments. By designing a foot model with several pieces, we can reproduce the shape of the various poses that a human foot can achieve. Moreover, because there are many contact points between the foot and the ground, it is possible to enhance the stability of the foot. Another contribution is that we can create foot motion without extra foot motion data. Previously, because the foot was treated as a body that has simple geometry, there has been little motion data information available for the foot. In particular, motion capture data deal only with information about the ankle joint. Using our foot model and control method, we can generate the foot motion without detailed motion data of the foot.

In the following sections, we explain our multi-segment model and foot pose control method. Our foot model is based on a human foot and contains several bones and joints. The model imitates a human foot structure and the characteristics of foot-to-ground collisions without becoming overly complex. We introduce our multi-segment foot model in detail and explain how to construct the foot model in section 3. A human foot can change its shape appropriately to perform functions according to the situation or the environment. We used a simple rule to allow the foot to take an appropriate pose depending on the movements of the character. We describe the foot pose control method in section 4 and its simulation in section 5.

We demonstrated the expressiveness and the robustness of our newly designed foot model by conducting several experiments, which are discussed in section 6. We reproduced several complicated foot motions such as foot circling or tiptoeing motion to show the expressiveness of our model. Also, we demonstrate the walking motion by using deep reinforcement learning. To prove the robustness of the model against the environment and the perturbations, we changed the terrain or pushed the character using two different

foot models (two-segment foot and ours) and compared the results. The results prove the robustness of our foot model.

2. Related work

Foot modeling. The foot is an important component that directly interacts with the ground for standing and locomotion. Therefore, progressive models of the foot were proposed, especially in the area of biomechanics [CHT*01]. Meglan and Berme [MB93] suggested the first foot model that was isolated from the ground. They designed the foot as an one-segment model, and the model features a viscoelastic heel made of a single sphere. A two-segment foot model was also proposed by Gilchrist and Winter [GW96]. They placed one hinge joint between the metatarsal bones and the phalanges to achieve a smooth transition between the swing phase and the stance phase. Studies were conducted to design the foot model to three segments or more [MCN03, LNAS15]. However, their major goal was to reproduce and analyze kinetic or kinematic information rather than to achieve realistic foot movement reproduction.

Further, as the number of foot segments has increased, there have been attempts to use various segment shapes instead of a box for computational efficiency. These shapes include a sphere [MMK09], an ellipsoid [KQW13, LNAS15], and a cylinder [Kec11]. We designed each foot segment as a capsule to accelerate the computation time for calculating the foot-to-ground collisions and to reflect the structure of foot bones and joints.

In the computer graphics community, contact handling has been executed in the usual way, such as non-sliding or non-penetrating. However, most studies on controlling physically simulated characters have treated the foot as a simple structure such as a box or capsule consisting of one or two bodies. Wang et al. [WFH09] used a two-segment foot model, which provides more flexibility when the heel strike or toe-off occurs. Jain and Liu [JL11] proposed a soft foot model using a deformable surface on the foot. Using their soft foot model, they achieved more robust walking control in simulation. Those foot models improved the robustness of controllers and the quality of the resultant motion, but the shape of the foot was not their main concern. Thus, the purpose of modeling the foot was limited to functional areas. We designed the foot as a multi-segment model enough to represent the several important characteristic foot shapes.

Physics-based simulation with a humanoid character. Developing an accurate walking control scheme for an articulated humanoid character in a physics-based simulation is a major problem in computer graphics. Many researchers have tried to reproduce realistic human walking motions with a physically simulated character by using pre-designed finite state machines [YLvdP07] or motion capture data [SKL07, dSAP08, LKL10a, LPKL14, LPY16]. In their works, they used human motion data as a reference for the controller. By using the real human data, they can show more realistic results for controlling humanoid characters. On the other hand, optimization-based controls, which set several pre-defined high-level objectives, were used to find the optimal actuator values to represent various motions [WFH09, HYL12, MWTK13, GvdPvdS13]. The objectives were usually defined by selected features considered to represent the main principles of human motion such as a joint torques minimization [dLMH10], a foot landing position [BdLH13], or metabolic energy expenditure [WHDK12]. Recently, deep reinforcement learning (DRL) has been successfully adapted to control biped locomotion [PBYVDP17, PALvdP18, YTL18]. We used the method proposed by [PALvdP18] to generate a walking motion.

Meanwhile, researches about balancing a biped character have been conducted in computer animation and robotics. Most of works have controlled the representative physical characteristics, such as the center of mass (CoM), the center of pressure (CoP), and the zero-moment point (ZMP). These characteristics should lie on a base of support that is a convex hull composed of the projection of the contacted bodies (often foot). To control CoM, CoP, and ZMP, hip and ankle strategies [Ste07] inspired by human balance control were used. Linear and angular momentum [KKK*03, AdSP07, MZS09] were also controlled to regulate CoM and CoP. We used the system of Macchietto et al. [MZS09], which formulated a quadratic problem to attain desirable linear and angular momentum changes to balance the character.

3. Multi-segment foot model

The human foot has a complex structure consisting of 26 bones, 33 joints, and more than 100 muscles and ligaments. However, in computer graphics, most of the foot models used in previous studies focusing on physically simulating biped characters often consisted of one or two bodies [LKL10b, WHDK12, HL14]. These foot models can hardly reproduce the actual foot shapes that a human can make. For example, foot models with one or two rigid bodies cannot perform foot rotations, such as pronation and supination. Further, it was difficult for them to control the interaction between the foot and the ground during standing because of the lack of DoFs. Consequently, they cannot reproduce the fine foot movements. We intend to design a multi-segment foot model that can describe various foot movements realistically.

To achieve this goal, we wondered if we should apply the real anatomy of human equally to our model. However, as the foot anatomy is quite complex, controlling such a complex structure would be inefficient, so we agreed that there should be a trade-off between accuracy and efficiency. In addition to efficiency, we have considered the tendency of actual foot movement caused by the foot anatomy. Although the foot is composed of 26 bones, the bones do not move separately because they are strongly connected with ligaments. Looking at the anatomy of the foot, the tarsal bones near the ankle are joined by dozens of ligaments. However, it is not the case with the toes located distal to the foot, so the toe bones have high DoFs, relatively. Based on this feature, we tried to design our foot model more realistically by increasing the number of segments and DoFs.

In order to develop the model for the foot while considering the above mentioned concerns, we referred to biomechanics researches. Several studies on foot modeling and segmenting in that field have been reported [DSR*11]. Researchers in biomechanics area were able to analyze foot anatomy and create a model to meet their needs. Among these studies, we found that some divide the foot along the coronal plane (hind foot, forefoot, and hallux) to measure foot kinematics during gait [KAHJ96, CHT*01].

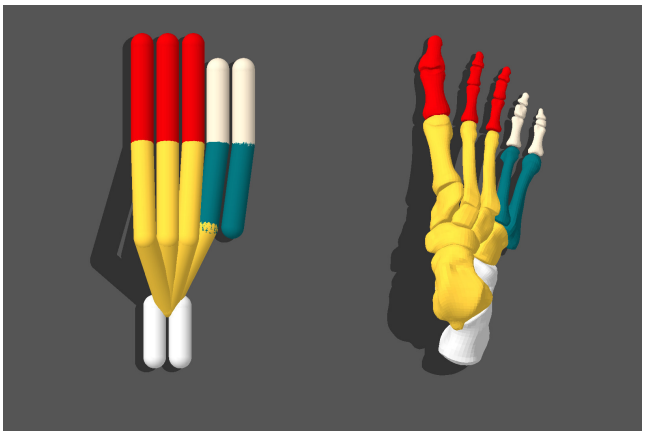

Figure 1: *Our multi-segment foot model consists of 16 artificial bones and five segments. Each bone has a capsule shape that is set to have two contact points at most. To help understand, each segment was painted in a different color. red: medial phalanges, ivory: lateral phalanges, yellow: medial metatarsal, green: lateral metatarsal, white: heel.*

However, these models are not sufficient to represent fine movements, such as inversion/eversion, which appear along the sagittal plane. To obtain more delicate foot information, such as kinematics and kinetics data, Macwilliams et al. [MCN03] designed a nine-segment foot model: hallux, medial toes, lateral toes, medial forefoot, lateral forefoot, calcaneus, cuboid, talus/navicular/cuneiform, and tibia/fibula. However, they segmented human foot model for analytic and diagnostic purposes. Particularly, placing the cuboid and talus/navicular/cuneiform as separate segments in the midfoot is aimed to identify specific movements of the tarsalmetatarsal and midtarsal joints caused by pathologies such as equinus or mid-foot breaks. We decided to use a slightly modified version of this model. Cuboid, talus, navicular, and cuneiform are bones that hardly move and have no contact with the ground, hence we placed these four bones in one segment. Also, tibia and fibula are not included in our foot segment definition. We noticed that toes except hallux are ignored in most studies [DSR*11]. To reproduce fine foot movements such as inversion/eversion, we divided toes into two parts, medial and lateral phalanges.

We first selected six segments that are in direct contact with the ground (heel, medial and lateral forefoot, first toe, second and third toes, and fourth and fifth toes) and then incorporated the first toe and the second and third toes into one segment for simplicity. Furthermore, we made other segments (talus, navicular and cuneiform) that do not have contact with the ground belong to the medial phalanges, which form an arch of the foot.

To summarize, our foot model consists of five segments and each segment is a control unit. To move each segment properly, we planted four joints (3 DoFs per each joint) in our foot model. A segment is composed of several artificial bones that represent phalanges, metatarsal bones, or the tarsal bones. In the following subsection, we describe the artificial foot bone first, followed by foot segmentation and joint placement in detail.

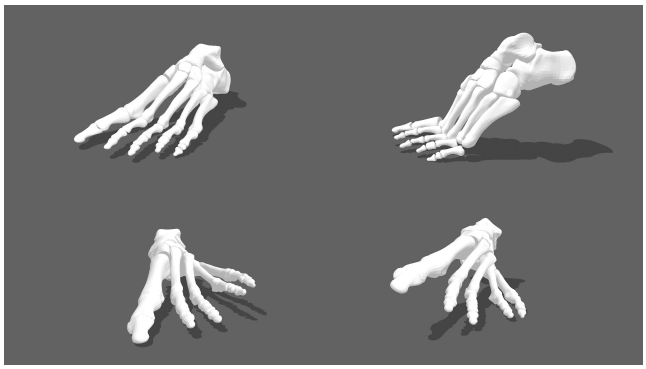

Figure 2: *Four basic foot poses. Top left: rest pose, top right: tiptoe, bottom left: inside tilt, and bottom right: outside tilt.*

3.1. Artificial foot bones

Human foot bones are largely divided into three parts: tarsal bones, metatarsal bones, and phalanges. Tarsal bones are seven bones that make up the midfoot and the hindfoot. Metatarsal bones are located in the midfoot and phalanges are in the forefoot. Our foot model consists of 16 capsule-shaped artificial bones that correspond to the human foot bones. Among the 16 bones, the five bones on the forefoot correspond to phalanges, the other five at the middle are metatarsal bones, and the other two form the calcaneus. We modeled the tarsal bones (colored with yellow in Figure 1) but decided not to directly control them for two reasons. One reason is that they rarely come into contact with the ground, so there is no interaction. The second reason is that they experience little movement because they are strongly bonded to the surrounding bones with ligaments.

Metatarsal bones and phalanges have cylindrical shapes. Wedge-shaped metatarsal bones and phalanges can be treated as a cylinder [Pat15], so we designed artificial bones with capsule primitives with the same radius. A capsule is a cylinder with hemispheres on both ends. This is used often for collision detection. We assume that each artificial bone can collide with the ground. We have made the collision occur only at both ends of the capsule-shaped bone, which are hemispheres.

3.2. Foot segmentation

We divided the 16 artificial bones into five segments (Figure 1). After observing various foot motions, we found that the individual movements of foot bones are constrained because they are connected to each other with ligaments. Therefore, we grouped bones that are moving and contact the ground together into same segment, and we control our foot model on a segment basis.

We defined five segments with the following names: medial phalanges, lateral phalanges, medial metatarsal, lateral metatarsal, and heel. For stable contact, we selected the number and the shape of the bones that belong to the segment so that each segment has three or more contact points when contacting with the ground. There are four basic foot poses: rest pose, a tiptoe, an inside tilt, and an outside tilt (Figure 2). By making the artificial bones that belong to

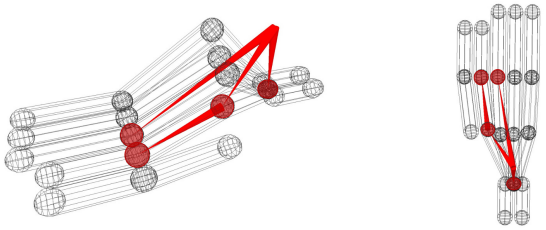

Figure 3: *We put four joints (marked with the red circle) on our foot model to move each segment properly. The locations of joints are between the third metatarsal and the third phalanges, between the fourth metatarsal and the fourth the phalanges, between the fourth metatarsal and the cuboid, and on the heel.*

some segments contact the ground, we can make the foot model achieve one of the basic poses.

3.3. Joint placement

We put four joints on our foot model to move each segment properly. Each joint is located based on the joint location on a real human foot. We use 3-DoF ball joints for foot joints so the foot model has 12 DoFs (Figure 3). The location of each joint is between the third metatarsal and the third phalanges, the fourth metatarsal and the fourth the phalanges, the fourth metatarsal and the cuboid, and on the heel.

The place where the foot is bent most is between the toes and the metatarsals. Therefore, we put two joints between the metatarsal bones and the phalanges. Looking at the midfoot during the rest pose from the front, the lateral side of the foot is in contact with the ground while the medial side is in arched form without touching the ground. When the foot tilts to the outside, the lateral metatarsals and heel are in contact with the ground. To control the outside metatarsals, we put a joint between the fourth metatarsal and the cuboid. However, when inside tilt occurs, medial metatarsals are still arched and have no contact or any notable movement, so we decided not to place a joint there. Instead, we made this part subordinate to the ankle joint. The heel joint does not exist on a human foot but we made one for the following reason. The heel is the part that can contact with the ground at any pose except the tiptoe pose because of the characteristic shape of the calcaneus bone. The bone can make contact in various directions. To allow our foot model to perform the function of the calcaneus bone, we put a joint on the heel.

All the capsule-shaped artificial bones we designed have the same radius and each end of the bones overlaps to form a sphere. We set the joint location to the center of the sphere, and the intersection of two bones is connected by that joint. Because two bones sharing a joint have the same contact point, one calculation for contact per joint is reduced.

4. Foot pose control

When a person is standing on the ground, a foot can be in contact with the ground in various ways (Figure 2). To reproduce these various shapes and motions of the foot naturally, the foot model must have many DoFs. We made a multi-segment foot model that has 12 DoFs based on human foot shape and have discussed the model in detail in the previous section.

Although our foot model has the power to represent various shapes and motions of the foot, the joint angles in the foot should be determined to make a specific pose. However, because there was not much concern about foot modeling until now, when capturing human motion, the foot was not the subject of consideration. In other words, motion capture data describing delicate foot motion are difficult to obtain. Therefore, we need a methodology that can control the foot depending on the environment or desired behavior without reference data of the foot.

We propose the foot pose control method for our multi-segment foot model. By using the foot pose controller, we reproduce an appropriate foot pose for a given ankle position and segments that are specified by the user to be in contact with the ground.

We take an ankle position and the contacting segments as a input of the foot pose controller to decide the foot shape at each moment because we consider it was sufficient to express the foot pose. For example, the tiptoe pose can be made by setting a high ankle position and making the toe segments attached to the ground and vice versa. Our foot model is embedded in the existing biped character model. The ankle joint position is obtained from the motion data, and the segments that must have contact with the ground for the task are provided by a user. Because the inputs of our controller can be obtained without detailed information about foot motion, we can reconstruct the foot pose together with the full body pose in common motion capture data.

Foot pose control proceeds in the following order (Figure 4). First, place the ankle joint at the ankle position of the motion capture data. Next, adjust the orientation of the ankle joint to place the joint belonging to the segment entered as the control input on the ground. Lastly, rotate the joint belonging to the segment in order to make the bones in input-segment be parallel with the ground.

The joint orientations can be obtained by the following calculation. Let $\vec{p}_a$ be ankle joint position, and $\vec{p}_s$ be joint position of marked segment. Now we let $\vec{v}_{sa} = \vec{p}_s - \vec{p}_a$. To change the original ankle orientation as little as possible in the motion capture data, we try to rotate the ankle joint to the smallest angle so that the segment joint could touch the ground. A rotation axis of the smallest rotation to make the segment joint touched the ground is perpendicular to the plane which contain $\vec{v}_{sa}$ and $\hat{n}$, where the rotation axis is perpendicular to $\vec{v}_{sa}$ and $\hat{n}$. The rotation axis $\hat{\omega}$ can be obtained as follows.

$$\hat{\omega} = \frac{\hat{n} \times \vec{v}_{sa}}{\|\hat{n} \times \vec{v}_{sa}\|}. \tag{1}$$

The rotation angle that causes $\vec{v}_{sa}$ to touch the ground is equal to the value obtained by subtracting the angle between $\vec{v}_{sa}$ and $-\hat{n}$ from the angle between $\vec{v}'_{sa}$ and $-\hat{n}$ where $\vec{p}_a + \vec{v}'_{sa} - r_c\hat{n}$ is on the ground. When ankle joint is rotated at angle ϕ, segment joint can

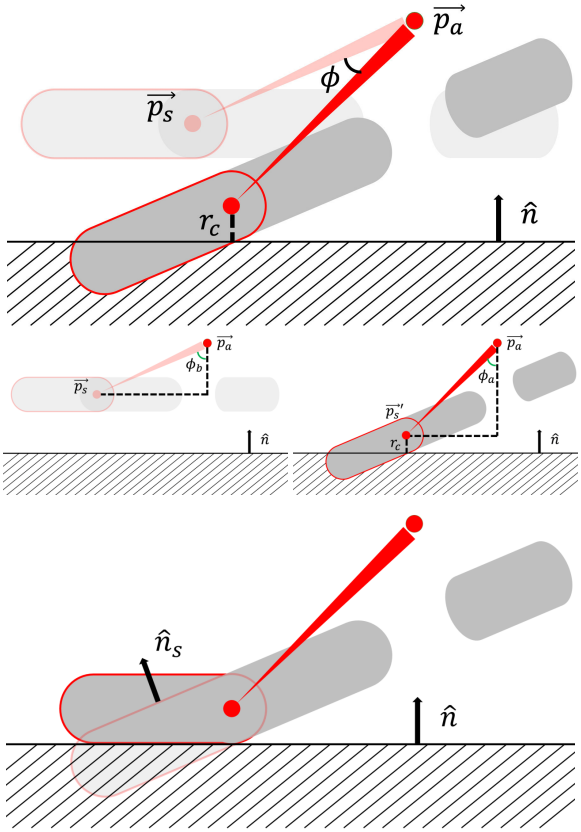

Figure 4: *Foot pose control process. For clarity, we draw our foot model in 2D. An artificial bone marked with the red border is the one belonging to the selected segment to contact with the ground. Top: We rotate the ankle joint to make the selected segment joint touch the ground. Middle: Representations of ϕ_b and ϕ_a. Bottom: We rotate the selected segment joint to make the selected bone parallel to the ground.*

touch the ground.

$$\begin{aligned}\phi &= \phi_b - \phi_a \\ &= \arctan(\|\vec{v}_{sa} - (\vec{v}_{sa}\cdot\hat{n})\hat{n}\|,\ -\vec{v}_{sa}\cdot\hat{n}) - \arccos(\frac{\vec{p}_a\cdot\hat{n} - r_c}{\|\vec{v}_{sa}\|}),\end{aligned} \tag{2}$$

where r_c is a radius of a capsule of each bone, ϕ_b is the angle between $\vec{v}_{sa}$ and $-\hat{n}$, and ϕ_a is the angle between $\vec{v}'_{sa}$ and $-\hat{n}$, respectively (see Figure 4).

After the joint of the segment touches the ground, we should make the bones in the segment be parallel to the ground, not to penetrate the ground. We rotated the segment joint so that the segment normal $\hat{n}_s$ matches the ground normal $\hat{n}$ for that. After that, we need to determine the remaining DoF of the segment. We reduced the unnaturalness by making the direction of the segment coincide with the target direction vector $\hat{t}$, which is the projection

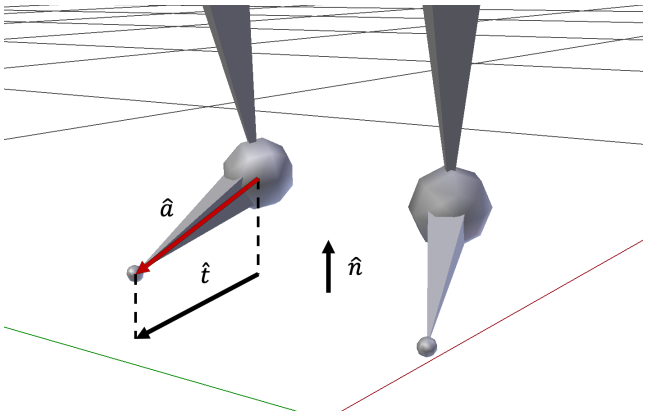

Figure 5: *A representation of $\hat{a}$ and $\hat{t}$ calculated from reference motion. $\hat{a}$ is the unit direction vector from the ankle joint to the toe joint, and $\hat{t}$ is the unit direction vector which is the projection of $\hat{a}$ to the ground.*

of the foot direction vector $\hat{a}$ on the ground plane. We set $\hat{a}$ to be a unit direction vector from an ankle to a toe in the original motion data (Figure 5).

$$\hat{t} = \frac{\hat{a} - (\hat{a}\cdot\hat{n})\hat{n}}{\|\hat{a} - (\hat{a}\cdot\hat{n})\hat{n}\|}, \tag{3}$$

$$R_s = R(\hat{s},\hat{t})R(\hat{n}_s,\hat{n}), \tag{4}$$

where R_s is rotation that segment joint should be rotated with. $\hat{t}$ is the segment target direction that is calculated from the ankle direction, $\hat{s}$ is the segment direction, and $\hat{n}_s$ is the segment normal unit vector. $R(\hat{c},\hat{d}) \in R^{3\times3}$ is the rotation matrix rotating $\hat{c}$ to $\hat{d}$ with the smallest rotation angle, where

$$R(\hat{c},\hat{d}) = \exp(\arccos(\hat{c}\cdot\hat{d})\frac{\hat{c}\times\hat{d}}{\|\hat{c}\times\hat{d}\|}). \tag{5}$$

The calculation of the pose control of the foot model is performed every time step and it is performed very quickly. Therefore, our controller works in real time.

5. Simulation with multi-segment foot model

In this section, we perform physics-based simulations considering a full-body character equipped with our foot model and foot pose control. First, we simulate assorted balancing motions in place and then simulate the walking motion as well. To control a full-body character model equipped with our multi-segment foot, we use existing controllers [MZS09, PALvdP18]. By using these controllers, we successfully reproduce diverse human motions. This simulation confirms that our newly designed foot model can be applied well to existing biped controllers.

5.1. Balancing

We use the balancing system of Macchietto et al. [MZS09] to keep the humanoid character balanced in the standing position. To achieve this goal, the linear and angular momentum of the humanoid character are adjusted. The character maintains its balance if its CoM and CoP are inside the base of support (BoS) which is the convex hull of the bones in contact with the ground. After deciding the desired CoM and CoP according to the positions of segments in contact with the ground, the quadratic optimization problem is defined to move the current CoM and CoP to the desired position. The output of this optimization is joint angle accelerations. Next, joint torques are obtained by using these accelerations as input for the dynamics solver and solving the floating-base hybrid dynamics [Fea07]. The simulation is performed by applying the obtained joint torques to each joint of the humanoid character. The optimization problem is defined as follwing:

$$\begin{aligned} \ddot{\theta}^* &= \underset{\ddot{\theta}}{\operatorname{argmin}}\ w_t C_t + w_l C_l + w_h C_h \\ \text{subject to: } & a_{sup} = J_{sup}\ddot{\theta} + \dot{J}_{sup}\dot{\theta}, \end{aligned} \tag{6}$$

where $\dot{\theta}$ and $\ddot{\theta}$ are the joint angle velocities and accelerations, $C_t(\ddot{\theta})$ is a quadratic objective to track the reference motion, $C_l(\ddot{\theta})$ and $C_h(\ddot{\theta})$ are also quadratic objectives to achieve the desired linear and angular momentum changes. J_{sup} is the Jacobian of the supporting bodies, and $\dot{J}_{sup}$ is the time derivative of J_{sup}. w_t, w_l, and w_h are the weights for tracking, linear momentum, and angular momentum objectives, respectively.

Non-slip condition constraints in Eq. (6) require determined supporting bodies. Because we obtained which segments should be in contact with the ground from the user-input when deciding the foot pose described in the previous section, non-slip constraints can be calculated using the user-input directly. Those user-input selected segments are corresponding to the supporting bodies in the optimization and using these information, the optimization is processed together with the foot pose control. J_{sup} and $\dot{J}_{sup}$ are calculated for those segments. We set a_{sup} to keep the segments in contact and parallel with the ground in a proportional-derivative (PD) control manner for both linear and angular quantities of the segments.

$$a_{sup} = \begin{pmatrix} k_s(p_{s,ref} - p_s) - d_s v_s \\ k_s \text{diff}(q_{s,ref},\ q_s) - d_s \omega_s \end{pmatrix}, \tag{7}$$

where k_s is a proportional gain, $d_s = 2\sqrt{k_s}$ is a derivative gain, $p_{s,ref}$ and p_s are joint positions in Cartesian coordinates of selected segments for the reference motion and the simulated character, $q_{s,ref}$ and q_s are joint orientations in SO(3) of selected segments for the reference motion and the simulated character, and v_s and ω_s are joint linear and angular velocities in Cartesian coordinates of selected segments for the simulated character. We used $k_s = 28$.

We set the values of θ and $\dot{\theta}$ as a current configuration of the character at the moment of optimization. Because J_{sup} and $\dot{J}_{sup}$ can be calculated from θ and $\dot{\theta}$, the constraints of Equation (6) are linear with respect to $\ddot{\theta}$. Because the objective of Equation (6) is quadratic with respect to $\ddot{\theta}$, the results of the optimization can be obtained by solving a set of linear equations.

5.2. Walking

Recently, the character control method based on reinforcement learning (RL) has been implemented successfully. A deep reinforcement learning (DRL), which introduces deep neural network models into RL, is suitable for calculation of high dimensional input and output such as a joint angle coordinate vector of an articulated character. A DRL approach allows the physics simulation of various movements including biped locomotion without much effort to design a delicate, hand-tuned controller. We equip a humanoid character with our foot model based on the control system proposed by Peng et al. [PALvdP18] to verify that our foot model works well using the existing controller.

Basically, we follow a state, action, and reward design defined in a previous work [PALvdP18], except for a modification in the action design for the system to adapt our foot model. The state vector is defined by $s = [p, v, q, \omega, \psi]$, where p, v, q and ω are aggregated body positions, velocities, orientations, and angular velocities represented in the Cartesian coordinate, respectively. $\psi \in [0, 1]$ is a phase variable that indicates the current stage of the character in a cycle of motion.

We designed an action vector as $a = [\theta_{disp}, \alpha_{pd}]$, where θ_{disp} is the displacement of motion and used to calculate the target motion θ_{des}, or $\theta_{des} = \theta_{ref} + \theta_{disp}$, where θ_{ref} is the reference motion from the motion capture data. We used a proportional-derivative (PD) control to track θ_{des} with PD gains k_p and k_d. Our foot model has many segments, with each having a small mass. This segmented design causes instability during the simulation when the joint torque is too strong. To make the simulation stable with this segmented design, we used a stable PD controller [TLT11] at every time step in the simulation.

α_{pd}, which is the second part of the action vector, is the normalized PD gain vector of foot joints. We took α_{pd} into the action vector so that the PD gain on each foot joint can be changed in every frame according to the situation. The PD gain adjustment gives flexibility to the joints on the foot. A significant amount of torque is needed when the foot stance to maintain its shape or when the foot needs to push the ground strongly. Conversely, at the moment of landing the foot, joint torque must be reduced to change the shape of the foot to fit into the ground. By adding PD gains to the action, we can expect that the deep neural network will learn the proper PD gain in each of these cases. PD gains are adjusted in the training process so that when a relatively large torque is required on the joint, it can be obtained by increasing the gain of the joint. PD gains of foot joints are calculated as follows:

$$k_p^{foot,i} = \beta \cdot \exp(\ln\beta \cdot \alpha_{pd}^i), \tag{8}$$

$$k_d^{foot,i} = \gamma \cdot \exp(\ln\gamma \cdot \frac{\alpha_{pd}^i}{2}), \tag{9}$$

where $k_p^{foot,i}$, $k_d^{foot,i}$ are a proportional gain and a derivative gain of the i-th foot joint, respectively. α_{pd}^i is the i-th element of the α_{pd}, $\beta = k_p$, and $\gamma = k_d$.

We use the proximal policy optimization algorithm [SWD*17] to train our control policy. A three-depth deep neural network is used

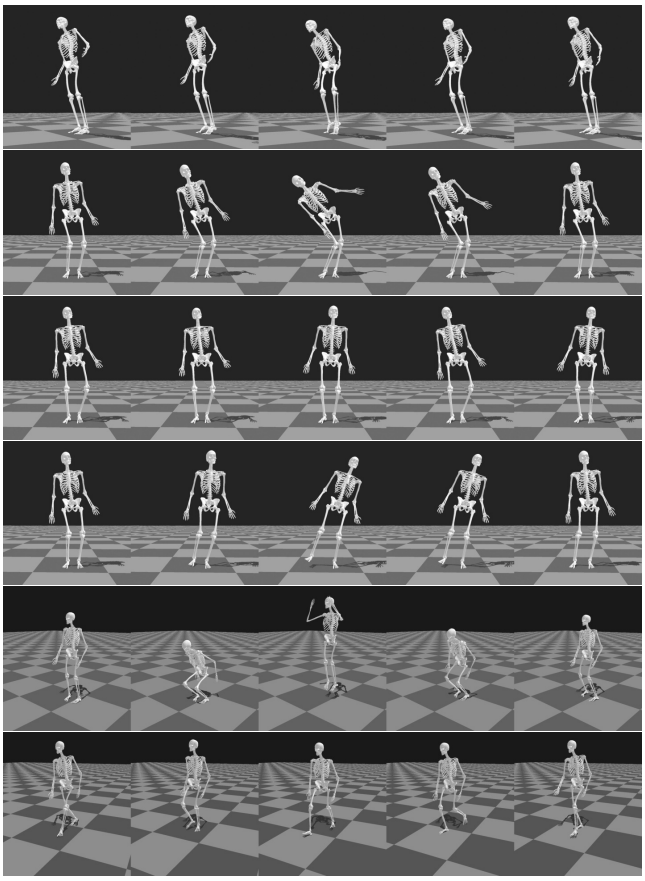

Figure 6: *We can simulate various foot motions with our multi-segment foot model: tiptoe, tilt, one foot circling, standing on one foot, in-place jump, and walk.*

for the value and action networks, and each layer is fully connected. The first layer has 256 neurons and the second and the third layers have 128 hidden units each.

The reward we use to train the neural network is as follows:

$$r = w_p r_p + w_v r_v + w_c r_c + w_e r_e, \tag{10}$$

where r_p, r_v, r_c, and r_e are the pose, velocity, center-of-mass, and the end-effector rewards. These rewards are the same as set in a previous work [PALvdP18].

6. Experimental results

Our human model has 66 DoFs (including 12 DoFs for each foot), is 150 cm tall, and weighs 48.5 kg. We use the dynamic system based on the Lie Group theory [PBP95] to solve the forward and the inverse dynamics for balancing simulation. For walking simulation, we use DART [LGH*18] simulator to solve the forward dynamics. Our simulation rate is 1800 Hz to simulate small segments on the foot model robustly. Foot pose control, optimization for maintaining the balance of the character, and calculation of actions for walking control are performed at the rate of 30 Hz. We set PD gains of $k_p = 400$, $k_d = 40$ all joints for walking motion except the joints on foots. To compare robustness to simple foot model, we used two-segment foot model. We set the coefficient of friction to be 1.0 in all of our experiments.

6.1. Reproducing various foot motions

Using our multi-segment foot model, we can simulate natural, sophisticated foot motions such as tiptoeing, tilting, circling on one foot, and standing on one foot (Figure 6). Given the ankle position and the foot segments that should be in contact with the ground as inputs, the foot pose is decided by our foot pose controller. In the figures and video, the foot bones are shown in red to indicate that the bones are in contact with the ground. Please watch the accompanying video for detailed motion.

Tiptoe. Tiptoeing is an unstable pose because it uses only the front, narrow area of the foot to balance while lifting the heel. The foot models that consist of a single body or two cannot produce the tiptoeing motion properly because they have a small number of contact points. In our experiment, 10 contact points exist in the toe area, and these points are activated by turns according to the COM of the body to balance like a real human does. Therefore, the tiptoe motion looks natural.

Tilt. When a human leans to the side like a skiing motion, the side of the foot comes into contact with the ground. We demonstrate this as foot tilt using our foot model. When we make the foot model tilt to the left side, the upper body is bent to the opposite side for balance. There are some internal collisions between body parts, and we will discuss this issue later.

One foot circling. With our multi-segment foot model, delicate foot movement such as foot circling can be produced. When stretching, a person turns their one ankle to release their ankle joint. At this moment, with the heel lifted, the toes touch the ground in turn.

Standing on one foot. Our model can balance with one foot. We let the character lift the right foot and balance with the left foot only. The character tilts its upper body toward the stance foot and shifts the CoM position to balance.

In-place jump. We can simulate the in-place jump. By demonstrating the jump motion, we wanted to prove that our foot model can adjust its shape adequately for more dynamic motion. At the landing phase, after the jump, the character tries to stand upright. At this moment, the toes and heel alternately contact the ground to maintain balance.

6.2. Walking motion

When humans walk, they place their foot on the ground, and the foot goes through heel-strike, mid-stance, and toe-off in turn. During these foot stance phases, humans appropriately change the shape of the foot to function well at each moment. For example, in the heel-strike phase, the foot absorbs the impact from the ground. In the toe-off phase, toes push the ground to gain the momentum

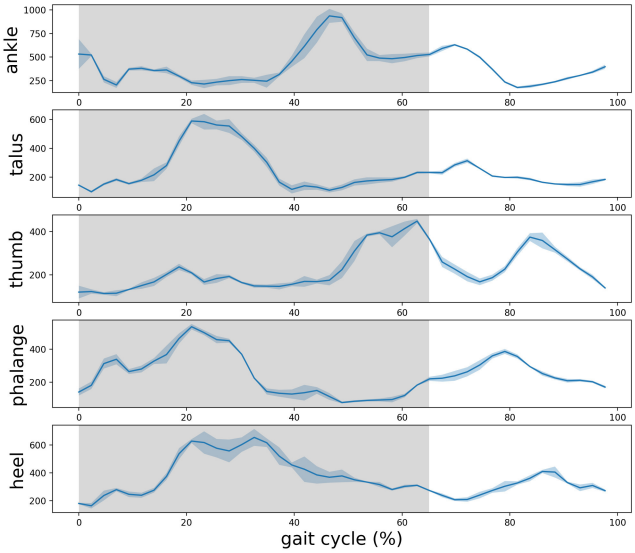

Figure 7: *Proportional gain values of joints in the right foot of walking motion. The values are averaged over 5 gait cycles. Gray zone is the stance phase of right foot. Blue shaded areas represent one standard deviation.*

to go forward. These changes in the shape of the foot not only give the naturalness by themselves but also improve the stability of the character on simulation. As a result, it helps to track the reference motion better. We applied the proposed segmented foot model to the humanoid character to confirm that the changes in the shape of the foot and functions adjust well during walking.

The heel is the first part that touches the ground when the foot lands. However, in case of using the box foot model in the simulation, we observed the motion that the entire sole of the foot, not just the heel, touches the ground during the landing motion. This action is caused by the structural nature of this box foot model because there are only two contact points in the back and two in the front in this foot model. Therefore, in order to land the box foot on the ground in a stable manner, a base of support (BoS) should be a surface, which is the sole. On the other hand, in the case of using our multi-segment foot model, the natural motion, in which the heel touches the ground first, can be achieved because the heel is composed of two capsules and it has 4 contact points. In the simulation result using segmented foot, we confirmed that the character touches the ground with the heel first when landing the foot. This result shows that our multi-segment foot model can mimic the walking motion of the human foot during walking better than the simpler foot model.

As mentioned in section 5.2, we added the proportional gains of foot joints to the action vector so that the foot joints are able to adjust their stiffness as necessary. We confirmed that the suitable PD gain for the actual walking pattern of the person is obtained. We plotted the proportional gain k_p^{foot} of each joint in the foot model to show the benefit of this scheme clearly (Figure 7). If the k_p^{foot} value is low, the corresponding joint is relatively soft, and it can

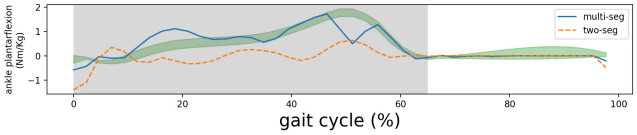

Figure 8: *Comparison of normalized ankle torques to human walking data. Green shaded area represents torques from human walking data [PB10]. Simulated result of the multi-segment foot model matches human data better than the result of the two-segment foot model.*

adapt to the ground. On the other hand, if the k_p^{foot} value is high, the corresponding joint becomes stiffer and can generate stronger power.

For the ankle joint, the proportional gain slightly increases at the heel-strike, decreases at mid-stance, and significantly increases again at toe-off phase. This is similar to the gait pattern of a human, which moves the center of mass forward by softening the ankle in the stance foot in the mid-stance, and then pushes the ground strongly when toe-off. This tendency is also seen in the third phalanges joint (the thumb joint), which is the joint performing a toe-off. The heel joint has the smallest proportional gain value during the heel-strike, because the faster the heel is in parallel with the ground and the more contact points touch the ground, the faster the character lands stably on the ground. The proportional gains of the fourth metatarsal joint and the fourth phalanges joint increase significantly in the mid-stance, which seems to maintain the foot shape so that the character does not fall over during the mid-stance.

We plotted ankle joint torques when walking for comparison (Figure 8). We can observe that the ankle torque values using the multi-segment foot model are better matched with human data than those using two-segment foot model. This implies that detailed modeling of the human musculoskeletal system would help to generate more plausible simulation results.

6.3. Robustness

A human can maintain balance while standing on flat ground, a gravelly field, or rocky terrain. In addition, a human is able to withstand the force of a push. These actions can be achieved by adjusting the foot pose to suit the situation. We conducted experiments to verify the robustness of our foot model on the curved terrain and in a situation where an external force is acting. The results were compared with those obtained using the two-segment foot model.

Rough terrain. In the real world, the ground is not always flat. At times, it is uneven and bumpy. Our proposed foot model can modify its shape to be adapted adequately to the environment. We conducted two experiments using the model to test its adaptability: one to stand on curved ground surface and the other to walk in the rough terrain without falling.

First, we let the two foot models (two-segment foot and multi-segment foot) stand on the sphere-shaped ground and compared the results (Figure 9). In the case of the two-segment foot model, the simulated character could not balance on the spherical terrain,

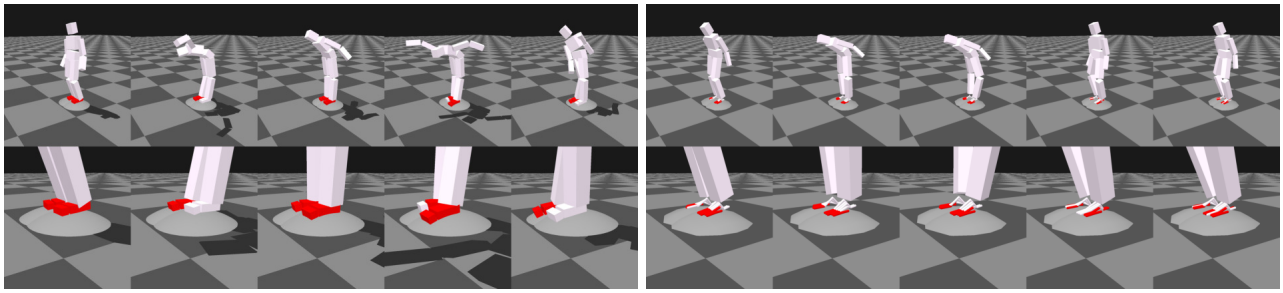

Figure 9: *We created bumpy terrain and placed our character on top. The two-segment model (left column) cannot stand on the bumpy terrain stably while our model (right column) can stand stably by properly changing the foot shape according to the terrain.*

and it slipped gradually. The character equipped with our multi-segment foot model could balance well on this terrain by changing the shape of the foot model. The result of this experiment shows that our foot model adapts well to the environment in a manner similar to a human. Next, we prepared an uneven terrain with an unevenness factor of $\pm$ 3 cm and make the character walk on it (Figure 10). The character was given no additional training for this environment. Our proposed foot model showed adaptation to the bumpy terrain by transforming its shape to fit the ground without falling.

Moving slope. We simulated a character standing on moving slopes swaying back and forth, and side to side (Figure 11). The joints between the phalanges and the metatarsals were flexed when the slope tilted back and forth, while the joints between metatarsals bent when the slope swaying side to side. In both cases, our foot model adjusted its shape to stand on the inclined floor and balance as the angle of the slope increased. This is the similar to how a human balances themselves on a slope.

External perturbation. To show the robustness of our controller, we applied a perturbation while the model was standing. We applied the external force to the torso of the character from three different directions (front, side, and back) for 0.4 sec. The character with the multi-segment foot model endured a maximum external force of 65 N from the front. When the character was pushed from the side, it could withstand a force of 75 N. When it was pushed from the back, it could withstand a force of 50 N.

Table 1: *We measured the maximum forces that each foot model could withstand. Our multi-segment foot model is more robust than the two-segment foot model. The multi-segment model endured a maximum external force of 75 N.*

Force(N)	Two-segment	Multi-segment
Front	45	65
Side	50	75
Back	40	50

Further, we compared the resilience of the multi-segment foot model with that of the two-segment foot model (Figure 12). Because the two-segment foot has few DoFs, when an external force is applied, the foot slips off the ground, becomes unstable, and eventually falls. However, our foot model is an articulated body consisting of segments and joints, so it can function much like a human foot. Therefore, when an external force is applied, only some segments of it come off the ground. The segments that remain on the ground make the overall model more stable. The experimental results confirm that our multi-segment foot model is more robust than the two-segment foot when force is applied from any direction (Table 1). We also added more contact points to the edges of the box in a two-segment foot model and repeated the pushing experiment, but the result did not change.

We conducted another push-recovery experiment while the character was walking. We applied an external force, starting with 50 N and going up to 100 N, on the right of the character for 0.4 sec. The character equipped with our foot model withstood force up to 80 N; it fell when the applied force reached 100 N.

Virtual contact plate. We display the virtual contact plate to show the changes in the foot contact during the simulation. Through the robustness experiments, we proved the robustness of our multi-segment foot model because our foot model can change its shape and contact points. By visualizing the contact points with the virtual contact plate, we can show the changes of the foot state more clearly (Figure 13).

7. Discussion

In this paper, we proposed a new foot model that consists of 16 artificial bones and five segments, and we proposed a controller that can successfully manage the interactions between the foot model and the ground. Previous studies on computer graphics tried several control methods to reproduce realistic human movements, but they did not pay much attention to foot modeling. We constructed the multi-segment foot model to resemble a human foot and simulate delicate foot motions.

For design convenience and computational efficiency, we used the capsule-shaped primitive for every foot bone, and each joint on the foot is a 3-DoF ball joint. Our foot model has more segments and DoFs than the foot models used in the previous works; therefore, we can generate delicate foot motion such as foot circling. In addition, our foot model is robust enough to adapt to rough terrain

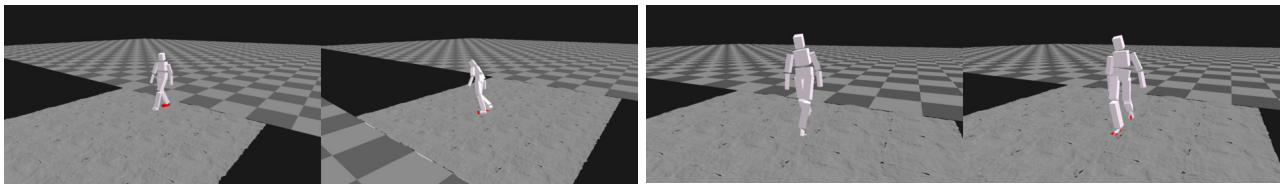

Figure 10: *Walking motion on a rough terrain. The two-segment model (left) tripped and fell on the bumpy terrain while our model (right) can walk through the uneven terrain.*

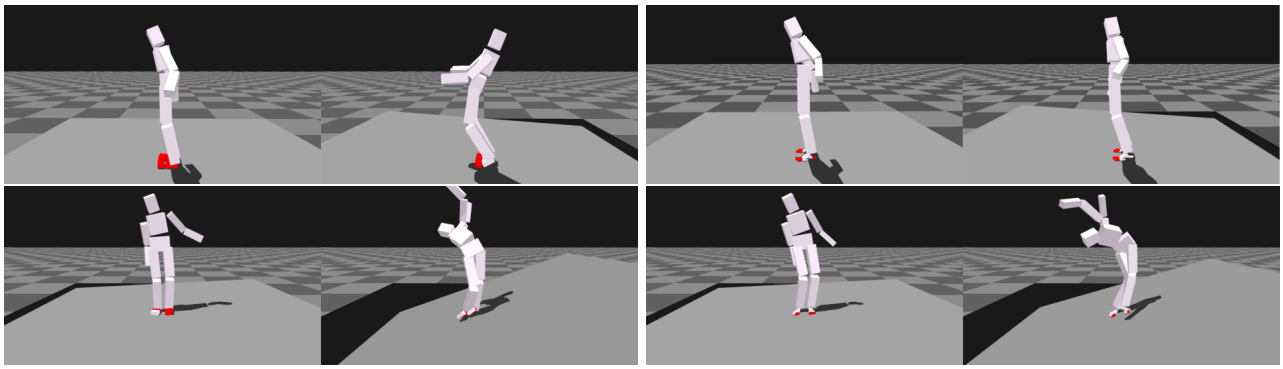

Figure 11: *(Top) Standing motion on a moving slope swaying back and forth. Although both the two-segment foot (top-left) and the multi-segment foot model (top-right) can balance on the moving slope, the case using the two segment foot model (top-left) has more wobbling motion on torso when the slope goes down. (Bottom) Standing motion on a moving slope swaying laterally. The case using the two segment foot model (bottom-left) failed to maintain balance while the multi-segment foot model (bottom-right) can balance.*

or endure external perturbation. We believe that by using our multi-segment foot model, a biped character simulation can be more expressive and robust.

Although the soft foot model [JL11] using a soft body simulation can achieve robust simulation results, we did not apply softness to our foot model. We applied our foot model to the existing rigid body character model to exploit the advantages of an articulated rigid body simulation, which include simple calculations and precise articulation.

Our foot modeling approach is meaningful in that we paid attention to the structure of a human foot when creating our model; the structure had previously been regarded as a simple form. Nevertheless, our foot model and the control method also have several limitations. Our controller determines the foot pose for the given inputs. The foot is the major body part that takes the contact force from the ground, so we expect that controlling not only the foot pose but also the ground reaction forces makes system more stable. We will build a controller that can directly control the ground reaction force.

There are some issues regarding foot modeling. We divided a foot into small pieces for expressiveness, but this can cause a problem. Because each bone in the foot model has a relatively small mass, the simulation will be unstable when a strong ground reaction force is generated. Therefore, we used a small time step for dynamics integration.

To control the foot model conveniently, we grouped the bones according to the movement of the foot. However, the foot movement of the grouped bones may seem unnatural in some cases. For example, in "rough terrain" experiments, it is natural for the toes to touch the ground all the way to grab the rounded ground. However, we grouped three toes (marked in red in Figure 1) together and treated them as one rigid body. All the bones in this segment, called medial phalanges, rotate together such that some of the bones in that segment cannot reach the ground. We decided to control the foot model on a segmented basis because the foot bones are strongly connected by ligaments. However, because the phalanges have a high DoF, it is better to control them individually rather than to group them together. Inevitably, the foot model needs to be further subdivided and controlled for more natural foot motion.

In some examples, especially the "tilt" example, self collisions occurred. For computational convenience, we did not consider the joint limit and self collisions. Because setting the joint limit and avoiding self collisions can make the simulated character motions more realistic, it would be desirable to consider these issues in the next study.

An actual human foot has a very complex structure composed of dozens of bones and hundreds of muscles and ligaments. In particular, the bones that comprise the foot are not separated from each other; they are joined together by ligaments. Hence, they can interact flexibly without major changes in shape. Recently, studies controlling human character using muscles have been successful [LPKL14, LYP*18]. Foot modeling with muscles is an interesting research topic. Using muscles for foot modeling can explain

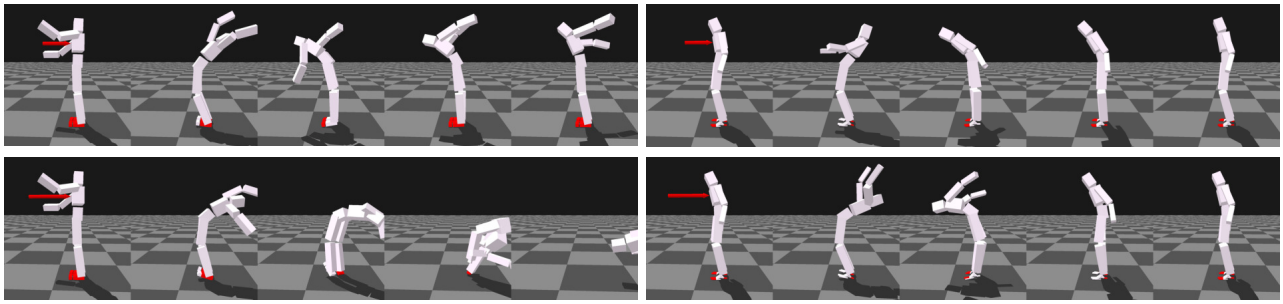

Figure 12: *We applied an external force to the character to show the robustness of the controller. We used two foot models, one is the two-segment foot model (left column) and the other is our proposed multi-segment model (right column). An external force of 40 N is applied to the experiment in the first row and 60 N in the second row. The two-segment model cannot endure 60 N force, while our multi-segment model can. Our model can withstand external force through proper adjustment of the foot pose.*

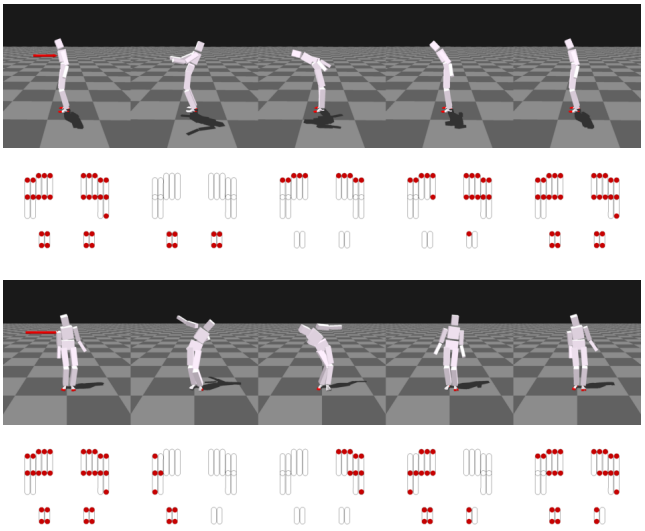

Figure 13: *We displayed the contact points on the virtual contact plate to visualize the change in contact while the human character is pushed by the external force. Note that the medial metatarsal segment was not displayed on the contact plate. This is because the segment is assumed to have no contact with the ground.*

many things about foot movement. Furthermore, as Pai pointed out [Pai10], the effective inertia of distal joints like interphalangeal joints can be large than the inertia calculated by their bone mass because of connections to proximal muscles. Therefore, if a musculoskeletal model is used for simulation, we might get more robust simulation results for stronger torques to distal joints. In addition, we expect that studies about the human hand will be a good reference for further research about foot structure and movement [SSB*15].

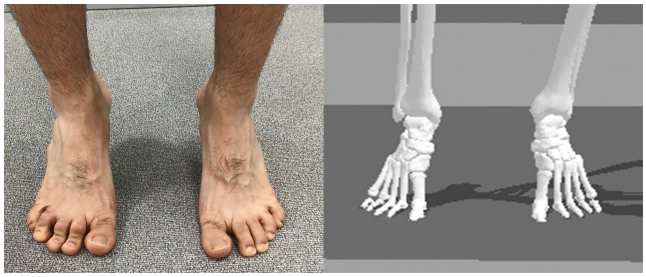

Figure 14: *Our foot model can reproduce human foot shapes and movements naturally. The left picture is the human foot tiptoeing and the right picture is tiptoe simulation result. Our controller can generate natural foot motion like human motion.*

Acknowledgements

This research was supported by the MSIT(Ministry of Science and ICT), Korea, under the SW Starlab support program(IITP-2017-0-00878) supervised by the IITP(Institute for Information & communications Technology Promotion), and the SNU-Samsung Smart Campus Research Center at Seoul National University provides research facilities for this study.

References

[AdSP07] ABE Y., DA SILVA M., POPOVIĆ J.: Multiobjective control with frictional contacts. In *Proceedings of the 2007 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2007), SCA '07, pp. 249–258. 2

[BdLH13] BORNO M. A., DE LASA M., HERTZMANN A.: Trajectory optimization for full-body movements with complex contacts. *IEEE Transactions on Visualization and Computer Graphics 19*, 8 (Aug 2013), 1405–1414. doi:10.1109/TVCG.2012.325. 2

[CHT*01] CARSON M., HARRINGTON M., THOMPSON N., O'CONNOR J., THEOLOGIS T.: Kinematic analysis of a multi-segment foot model for research and clinical applications: a repeatability analysis. *Journal of biomechanics 34*, 10 (2001), 1299–1307. 2

[dLMH10] DE LASA M., MORDATCH I., HERTZMANN A.: Feature-based locomotion controllers. *ACM Trans. Graph. 29*, 4 (2010), 131:1–131:10. 2

[dSAP08] DA SILVA M., ABE Y., POPOVIĆ J.: Interactive simulation of stylized human locomotion. *ACM Trans. Graph. 27*, 3 (2008), 82:1–82:10. 2

[DSR*11] DESCHAMPS K., STAES F., ROOSEN P., NOBELS F., DESLOOVERE K., BRUYNINCKX H., MATRICALI G. A.: Body of evidence supporting the clinical use of 3d multisegment foot models: a systematic review. *Gait & posture 33*, 3 (2011), 338–349. 2, 3

[Fea07] FEATHERSTONE R.: *Rigid Body Dynamics Algorithms*. Springer, Nov 2007. 6

[GvdPvdS13] GEIJTENBEEK T., VAN DE PANNE M., VAN DER STAPPEN A. F.: Flexible muscle-based locomotion for bipedal creatures. *ACM Trans. Graph. 32*, 6 (2013), 206:1–206:11. 2

[GW96] GILCHRIST L., WINTER D.: A two-part, viscoelastic foot model for use in gait simulations. *Journal of biomechanics 29* (07 1996), 795–8. doi:10.1016/0021-9290(95)00141-7. 2

[HL14] HA S., LIU C. K.: Iterative training of dynamic skills inspired by human coaching techniques. *ACM Transactions on Graphics (TOG) 34*, 1 (2014), 1. 2

[HYL12] HA S., YE Y., LIU C. K.: Falling and landing motion control for character animation. *ACM Trans. Graph. 31*, 6 (2012), 155:1–155:9. 2

[JL11] JAIN S., LIU C. K.: Controlling physics-based characters using soft contacts. *ACM Trans. Graph. 30*, 6 (2011), 163:1–163:10. 2, 10

[KAHJ96] KIDDER S. M., ABUZZAHAB F. S., HARRIS G. F., JOHNSON J. E.: A system for the analysis of foot and ankle kinematics during gait. *IEEE Transactions on Rehabilitation Engineering 4*, 1 (1996), 25–32. 2

[Kec11] KECSKEMÉTHY A.: A novel cylinder-plane foot contact model for human gait motion reproduction. In *Proc. ECCOMAS Thematic Conf. on Multibody Dynamics, Brussels, Belgium* (2011). 2

[KKK*03] KAJITA S., KANEHIRO F., KANEKO K., FUJIWARA K., HARADA K., YOKOI K., HIRUKAWA H.: Resolved momentum control: humanoid motion planning based on the linear and angular momentum. In *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003) (Cat. No.03CH37453)* (2003), vol. 2, pp. 1644–1650 vol.2. 2

[KQW13] KOOP D., Q. WU C.: Passive dynamic biped walking-part i: Development and validation of an advanced model. *Journal of Computational and Nonlinear Dynamics 8* (10 2013), 041007. doi:10.1115/1.4023934. 2

[LGH*18] LEE J., GREY M. X., HA S., KUNZ T., JAIN S., YE Y., SRINIVASA S. S., STILMAN M., LIU C. K.: Dart: Dynamic animation and robotics toolkit. *Journal of Open Source Software 3*, 22 (2018), 500. doi:10.21105/joss.00500. 7

[LKL10a] LEE Y., KIM S., LEE J.: Data-driven biped control. *ACM Trans. Graph. 29*, 4 (July 2010), 129:1–129:8. 2

[LKL10b] LEE Y., KIM S., LEE J.: Data-driven biped control. *ACM Transactions on Graphics (TOG) 29*, 4 (2010), 129. 2

[LNAS15] LOPES D., NEPTUNE R., AMBRÓSIO J., SILVA M.: A superellipsoid-plane model for simulating foot-ground contact during human gait. *Computer methods in biomechanics and biomedical engineering 19* (2015), 1–10. doi:10.1080/10255842.2015.1081181. 2

[LPKL14] LEE Y., PARK M. S., KWON T., LEE J.: Locomotion control for many-muscle humanoids. *ACM Trans. Graph. 33*, 6 (2014), 218:1–218:11. 2, 10

[LPY16] LIU L., PANNE M. V. D., YIN K.: Guided learning of control graphs for physics-based characters. *ACM Trans. Graph. 35*, 3 (2016), 29:1–29:14. 2

[LYP*18] LEE S., YU R., PARK J., AANJANEYA M., SIFAKIS E., LEE J.: Dexterous manipulation and control with volumetric muscles. *ACM Trans. Graph. 37*, 4 (2018). 10

[MB93] MEGLAN D., BERME N.: A 3d passive mechanical model of the human foot for use in locomotion synthesis. *Journal of Biomechanics 26*, 3 (1993), 331. 2

[MCN03] MACWILLIAMS B. A., COWLEY M., NICHOLSON D. E.: Foot kinematics and kinetics during adolescent gait. *Gait & Posture 17*, 3 (2003), 214–224. 2, 3

[MMK09] MILLARD M., MCPHEE J., KUBICA E.: *Multi-Step Forward Dynamic Gait Simulation*. springer, 2009, pp. 25–43. 2

[MWTK13] MORDATCH I., WANG J. M., TODOROV E., KOLTUN V.: Animating human lower limbs using contact-invariant optimization. *ACM Trans. Graph. 32*, 6 (2013), 203:1–203:8. 2

[MZS09] MACCHIETTO A., ZORDAN V., SHELTON C. R.: Momentum control for balance. *ACM Trans. Graph. 28*, 3 (2009), 80:1–80:8. 2, 5, 6

[Pai10] PAI D. K.: Muscle mass in musculoskeletal models. *Journal of biomechanics 43*, 11 (2010), 2093–2098. 11

[PALvdP18] PENG X. B., ABBEEL P., LEVINE S., VAN DE PANNE M.: Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Trans. Graph. 37*, 4 (July 2018), 143:1–143:14. URL: http://doi.acm.org/10.1145/3197517.3201311, doi:10.1145/3197517.3201311. 2, 5, 6, 7

[Pat15] PATTON K. T.: *Anatomy & Physiology*. Elsevier, jan 2015. 3

[PB10] PERRY J., BURNFIELD J.: *Gait Analysis: Normal and Pathological Function*. Slack Incorporated, Feb 2010. 8

[PBP95] PARK F. C., BOBROW J. E., PLOEN S. R.: A lie group formulation of robot dynamics. *Int. J. Rob. Res. 14*, 6 (1995), 609–618. 7

[PBYVDP17] PENG X. B., BERSETH G., YIN K., VAN DE PANNE M.: Deeploco: Dynamic locomotion skills using hierarchical deep reinforcement learning. *ACM Trans. Graph. 36*, 4 (July 2017), 41:1–41:13. URL: http://doi.acm.org/10.1145/3072959.3073602, doi:10.1145/3072959.3073602. 2

[SKL07] SOK K. W., KIM M., LEE J.: Simulating biped behaviors from human motion data. *ACM Trans. Graph. 26*, 3 (2007). 2

[SSB*15] SACHDEVA P., SUEDA S., BRADLEY S., FAIN M., PAI D. K.: Biomechanical simulation and control of hands and tendinous systems. *ACM Trans. Graph. 34*, 4 (2015), 42:1–42:10. 11

[Ste07] STEPHENS B.: Integral control of humanoid balance. In *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems* (2007), pp. 4020–4027. 2

[SWD*17] SCHULMAN J., WOLSKI F., DHARIWAL P., RADFORD A., KLIMOV O.: Proximal policy optimization algorithms. *CoRR abs/1707.06347* (2017). URL: http://arxiv.org/abs/1707.06347, arXiv:1707.06347. 6

[TLT11] TAN J., LIU K., TURK G.: Stable proportional-derivative controllers. *IEEE Comput. Graph. Appl. 31*, 4 (July 2011), 34–44. URL: http://dx.doi.org/10.1109/MCG.2011.30, doi:10.1109/MCG.2011.30. 6

[WFH09] WANG J. M., FLEET D. J., HERTZMANN A.: Optimizing walking controllers. *ACM Trans. Graph. 28*, 5 (2009), 168:1–168:8. 2

[WHDK12] WANG J. M., HAMNER S. R., DELP S. L., KOLTUN V.: Optimizing locomotion controllers using biologically-based actuators and objectives. *ACM Trans. Graph. 31*, 4 (2012), 25:1–25:11. 2

[YLvdP07] YIN K., LOKEN K., VAN DE PANNE M.: Simbicon: Simple biped locomotion control. *ACM Trans. Graph. 26*, 3 (2007). 2

[YTL18] YU W., TURK G., LIU C. K.: Learning symmetric and low-energy locomotion. *ACM Trans. Graph. 37*, 4 (July 2018), 144:1–144:12. URL: http://doi.acm.org/10.1145/3197517.3201397, doi:10.1145/3197517.3201397. 2