

Generative GaitNet

Jungnam Park
Department of Computer Science and
Engineering
Seoul National University
Republic of Korea
jungnam04@mrl.snu.ac.kr

Sehee Min
Department of Computer Science and
Engineering
Seoul National University
Republic of Korea
sehee@mrl.snu.ac.kr

Phil Sik Chang
Department of Computer Science and
Engineering
Seoul National University
Republic of Korea
phil@mrl.snu.ac.kr

Jaedong Lee
Department of Computer Science and
Engineering
Seoul National University
Republic of Korea
Jaedong@mrl.snu.ac.kr

Moon Seok Park
Department of Orthopaedic Surgery
Seoul National University Bundang
Hospital
Republic of Korea
pmsmed@gmail.com

Jehee Lee
NCsoft
Republic of Korea
Department of Computer Science and
Engineering
Seoul National University
Republic of Korea
jehee@mrl.snu.ac.kr

ABSTRACT

Understanding the relation between anatomy and gait is key to successful predictive gait simulation. In this paper, we present Generative GaitNet, which is a novel network architecture based on deep reinforcement learning for controlling a comprehensive, full-body, musculoskeletal model with 304 Hill-type musculotendons. The Generative GaitNet is a pre-trained, integrated system of artificial neural networks learned in a 618-dimensional continuous domain of anatomy conditions (e.g., mass distribution, body proportion, bone deformity, and muscle deficits) and gait conditions (e.g., stride and cadence). The pre-trained GaitNet takes anatomy and gait conditions as input and generates a series of gait cycles appropriate to the conditions through physics-based simulation. We will demonstrate the efficacy and expressive power of Generative GaitNet to generate a variety of healthy and pathological human gaits in real-time physics-based simulation.

CCS CONCEPTS

• **Computing methodologies** → *Animation; Physical simulation; Reinforcement learning.*

KEYWORDS

GaitNet, Musculoskeletal Simulation, Predictive Gait Simulation, Clinical Gait Analysis

ACM Reference Format:

Jungnam Park, Sehee Min, Phil Sik Chang, Jaedong Lee, Moon Seok Park, and Jehee Lee. 2022. Generative GaitNet. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Proceedings (SIGGRAPH '22 Conference Proceedings, August 7–11, 2022, Vancouver, BC, Canada)*

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

SIGGRAPH '22 Conference Proceedings, August 7–11, 2022, Vancouver, BC, Canada

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9337-9/22/08...\$15.00

<https://doi.org/10.1145/3528233.3530717>

'22 Conference Proceedings), August 7–11, 2022, Vancouver, BC, Canada. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3528233.3530717>

1 INTRODUCTION

A long-standing goal of human animation is a simulation algorithm that predicts gait in specific conditions. Specifically, we consider predictive gait simulation that takes anatomical conditions (e.g., mass distribution, body proportion, bone deformity, and muscle weakness/contracture) and gait conditions (e.g., stride and cadence) as input and generates a series of periodic gait cycles appropriate to the given conditions via physics-based simulation. This goal with a musculoskeletal simulation model is challenging owing to its large control space and the difficulty of evaluating the combined dynamics of muscle contraction and bone articulations. Feedback/feedforward control, linear/nonlinear control theory, various forms of trajectory optimization, anatomical modeling and simulation have been studied to achieve robust gait simulation for several decades [Geijtenbeek et al. 2013; Lee et al. 2014; Wang et al. 2012]. Recently, rapid progress toward this goal has been made using DRL (deep reinforcement learning) [Kidziński et al. 2018; Lee et al. 2019].

The core of predictive gait simulation is establishing a conceptual mapping between human anatomy and its gait. Ideally, we wish that appropriate gait emerges from biological principles for any given body condition. This approach often requires dynamical models much simpler than human anatomy [Kajita et al. 2003; Wang et al. 2012] or results in aberrant gaits that do not look human-like [Heess et al. 2017; Waterval et al. 2021]. Alternatively, many gait simulation algorithms leverage motion capture data as reference to formulate the problem as trajectory tracking [Fussell et al. 2021; Lee et al. 2014; Peng et al. 2018]. Given a reference trajectory to track, gait simulation algorithms can faithfully reproduce every nuance of human gaits captured in the reference data and the simulation can also be more robust and stable. Despite the advantages of reference-tracking approaches, a reference trajectory is considered harmful for predictive gait simulation because the simulated gait is constrained by the reference trajectory and thus may not fully adapt to new conditions as desired.

Although the state-of-the-art algorithms based DRL have demonstrated highly-detailed, anatomically-valid gaits, the process is computationally demanding. Given a reference trajectory, muscle-actuated gaits can be reproduced for a specific anatomy condition through policy learning and physics-based simulation. Policy learning takes several hours to several days depending on the complexity of the anatomical model and the level of dynamic skills. Whenever new conditions are given, a new policy must be learned from scratch.

In this paper, we present a Generative GaitNet (or simply GaitNet), which is a pre-trained, integrated system of artificial neural networks that learns the control policy of human gaits eligible for a high-dimensional continuous domain of anatomical conditions. The control policy is massively parameterized by a large number of bone and muscle conditions. The GaitNet produces a broad spectrum of human gaits starting from a single motion clip of the reference gait cycle, which serves as gentle guidance of cyclic patterns. Accurately tracking the reference motion is not the primary goal of the GaitNet. With the GaitNet, simulated gaits can deviate from the reference motion by using biologically-inspired minimal rewards instead of imitation rewards and employing hierarchical motion displacement mapping and variable phase stepping. The key to the success of Generative GaitNet is our novel CSN (Cascaded Subsumption Network) architecture that captures the correlations among muscle activations and their influence on the simulated gait in a hierarchical manner. The CSN allows us to learn the large integrated policies progressively while alleviating the risks of catastrophic forgetting.

We will demonstrate the efficacy and expressive power of Generative GaitNet to generate a broad variety of healthy and pathological human gaits. The pre-trained GaitNet allows gait simulation to be interactive and responsive. The change of anatomical condition is immediately reflected to the real-time simulation. Code and data for this paper are available at <https://github.com/namjohn10/GenerativeGaitNet>.

2 RELATED WORK

Reproducing human locomotion in physics-based simulation has been studied in graphics, robotics, and biomechanics for decades. Early approaches leveraged finite state machines and hand-crafted feedback rules to control dynamics systems [Hodgins et al. 1995; Yin et al. 2007]. Conceptual (often Simplified) dynamics models, such as an inverted pendulum and its variants, have been comprehensively explored to simulate the dynamics of biped locomotion [Coros et al. 2010; Kajita et al. 2003; Mordatch et al. 2010]. The use of motion capture data significantly improved the quality and stylistic variation of locomotion simulation [Lee et al. 2010; Sok et al. 2007]. Statistical optimal control theory provides a solid theoretical foundation for understanding the relation among energy efficiency, biological plausibility, and task achievement [Wang et al. 2009; Yin et al. 2008]. Many trajectory optimization methods that minimize either total joint torques or metabolic energy expenditure have been explored [Al Borno et al. 2012; Anderson and Pandy 2001; Kwon and Hodgins 2017; Kwon et al. 2020; Mordatch and Todorov 2014; Tassa et al. 2012].

Recently, remarkable progresses have been made in biped locomotion control using DRL (Deep Reinforcement Learning), which

allows robust control policies to be learned by deep artificial neural networks. It has been demonstrated that deep control policies can deal with uneven terrain and various motor skills [Peng et al. 2017]. Deep control policies can imitate human motion closely if reference motion capture data are provided [Bergamin et al. 2019; Liu and Hodgins 2018; Park et al. 2019; Peng et al. 2018; Won et al. 2020]. Although the use of reference motion data is often preferred for best-looking results, the reference motion is not mandatory. Stylistic, physically-plausible motions can be produced without any reference data [Merel et al. 2020; Yu et al. 2018]. The usability and flexibility of DRL are significantly improved if the control policies can be learned in continuously parameterized environments [Lee et al. 2021; Portelas et al. 2020; Won and Lee 2019]. Unfortunately, previous studies reported that brute-force curriculum learning can deal with only a few parameters (up to four dimensions as reported by Lee et al. [2021]). Learning control policies in higher-dimensional parametric domain is a technical challenge we address in this work.

The dynamical model of the human body in aforementioned studies consists of rigid links connected by idealized motors. The anatomical simulation model consisting of bones, musculotendons, and their contraction dynamics achieves better accuracy, flexibility, and applicability in gait simulation. The combination of muscle actuation models and nonlinear optimization has successfully demonstrated the simulation of healthy and pathological human gaits [Ong et al. 2019; Peng and van de Panne 2017; Song and Geyer 2018; Thatte and Geyer 2015; Wang et al. 2012; Waterval et al. 2021]. Since control optimization is notoriously time-consuming, early studies used simplified musculoskeletal models with only one or two dozens of musculotendon units. There have been continuing effort to develop efficient trajectory optimization and muscle approximation methods that can deal with scalable anatomical models with more muscles [Dembia et al. 2020; Falisse et al. 2019; Jiang et al. 2019; Lee et al. 2014]. In addition to humans, the locomotion of various creatures and underwater animals has also been studied [Geijtenbeek et al. 2013; Min et al. 2019].

DRL has also brought great progress in muscle-actuated simulation and control [Anand et al. 2019; Song et al. 2020; Wang et al. 2019]. Many results of series of the NeurIPS AI for prosthetics challenge produced muscle-actuated gaits of interactively-controllable speed [Kidziński et al. 2018, 2020]. The state-of-the-art, DRL-based algorithm presented by Lee and his colleagues [2019] uses a comprehensive anatomical model with 346 muscles. Their algorithm is capable of learning robust control policies that produce various human movements, such as running, jumping, and cartwheel, by mimicking reference motion data. Our Generative GaitNet is also based on DRL focusing on bipedal locomotion under various body sizes, proportions, and muscle deficits.

3 PARAMETERIZED ANATOMICAL MODEL

Our anatomical model has 23 rigid bones and 304 muscles. Each muscle is attached to bones at its origin and insertion points. The geometric route of the muscle is represented by a series of way-points, which transmit force between the origin and insertion points. The muscle generates contraction force according to Hill-type muscle model.

$$f_{\text{muscle}}(l, \dot{l}, a) = f_{\text{max}}(a \cdot g_c(l, \dot{l}) + g_p(l)), \quad (1)$$

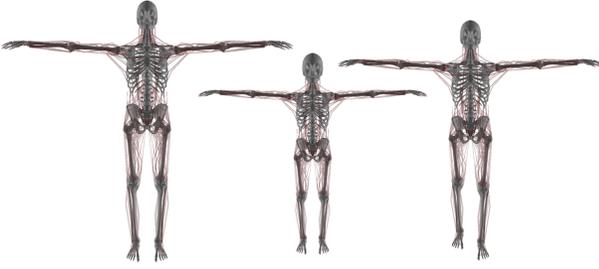


Figure 1: Parameterized musculoskeletal models. (Left) Reference Model. (Middle) Child. (Right) Uneven leg length.

where f_{\max} is the maximal isometric force. The muscle can generate its maximal force when its length is optimal l_{optimal} . l and \dot{l} , respectively, denote the muscle length and the rate of length change normalized by the optimal length. $a \in [0, 1]$ is its activation level. g_c is the contractile force, which depends on the force-length and force-velocity curves. g_p is the passive force. The active force occurs when the muscle is active, while the passive force occurs when the muscle is stretched. We determined the parameters of the Hill-type model based on data sheets provided in previous studies [Delp et al. 2007; Lee et al. 2019].

Our anatomical model is morphable with adjustable body and muscle conditions (see Figure 1). The reference model represents a healthy individual who is 168.7 cm tall and weighs 72.9 kg. The skeleton is parameterized by ten parameters

$$C_{\text{body}} = (c_{\text{head}}, c_{\text{torso}}, c_1, \dots, c_8) \quad (2)$$

that generates a continuous spectrum of body sizes and proportions. c_{head} and c_{torso} are scale factors of the head and the torso, respectively, and the others are scale factors of four (upper and lower) limbs. The size, length, and weight of individual body parts are scaled accordingly. We consider two types of muscle deficits: weakness and contracture. The maximal isometric force indicates how strong the muscle is. The weakness parameter $c_{\text{weakness}} \in [0, 1]$ of a muscle is the scaling factor of its maximal isometric force with respect to the reference model. Contracture refers to permanent shortening of muscles and tendons. The contracture parameter $c_{\text{contracture}} \in [0, 1]$ of a muscle is the scaling factor of its optimal length. The anatomical condition of a parametrically-varied model is represented by a vector $C_{\text{anatomy}} = (C_{\text{body}}, C_{\text{weakness}}, C_{\text{contracture}}) \in \mathbb{R}^{10+2 \times 304}$ concatenating the body scale, weakness, and contracture factors over all body parts and muscles.

Given the size and proportion of the skeleton altered, the geometric route of each muscle and its parameters must be altered as well to match the change of the skeleton. We use the Musculature Retargeting algorithm by Ryu et al. [2021], which systematically updates Hill-type muscle parameters and waypoints such that the muscle functionality and the range of joint motion are preserved for parameterically-varied models.

4 GENERATIVE GAITNET

The Generative GaitNet is an integrated system of artificial neural networks that learns the relation between anatomy and gait

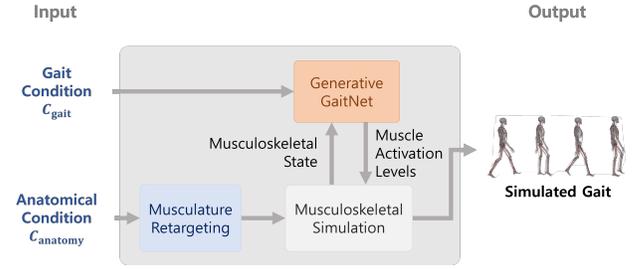


Figure 2: System Overview.

(see Figure 2). The GaitNet takes gait condition C_{gait} and anatomical condition C_{anatomy} as input and generates the corresponding physically- and anatomically-plausible gait through physics-based musculoskeletal simulation. The gait condition include stride and cadence. High-dimensionality of the anatomy condition explained in the previous section makes the learning of GaitNet very challenging. In this section, we will explain a baseline algorithm, which works well with a moderate subset of the anatomical condition. We will explain how to learn the GaitNet with the full set of 608-dimensional anatomical condition through cascading and subsumption in the next section.

In deep reinforcement learning, an agent in state s receives reward r from the environment through action u . Through the action sequence of the agent in the environment, the cumulative reward $r_0 + \gamma r_1 + \gamma^2 r_2 + \dots + \gamma^n r_n$ is calculated, where γ is a discount factor. The goal of reinforcement learning is to learn the policy $\pi_\theta(u|s)$ that maximizes the cumulative reward. Deep reinforcement learning showed remarkable performance in torque control of 3D characters in physics-based simulation [Park et al. 2019; Peng et al. 2018; Won et al. 2020]. Recently, DRL showed its potential in muscle-driven control using two-level hierarchical learning in the work of Lee et al [2019]. We adopted the two-level network architecture, though our work is fundamentally different from their trajectory-tracking approach. Our anatomically-inspired rewards and variable phase-stepping method achieves realistic human gait without imitation rewards. Although we use a reference gait to guide the progression of the simulation, the simulation is not hindered by the guidance but is free to deviate from the reference gait to satisfy required conditions.

In the two-level network architecture, the torque-based controller $\pi_\theta(u|s)$ is a stochastic policy that produces PD target poses as output. The PD servos compute desired joint torques τ from the target poses. The muscle coordinator $a = \pi_\psi(\tau, s)$ is a regression network that computes muscle activation a that drives the musculoskeletal system in physics-based simulator. The policy network π_θ and the regression network π_ψ are learned jointly in the DRL framework.

4.1 State and Action

The state of the musculoskeletal system has four terms.

$$s = (s_{\text{skeleton}}, s_{\text{muscle}}, s_{\text{joint}}, s_{\text{gait}}), \quad (3)$$

where s_{skeleton} includes the center of mass (COM) velocity, the relative position and linear velocity of body parts with respect to

COM, and their orientation and angular velocity. s_{muscle} includes muscle weakness and contracture parameters.

Multiple muscles around a joint are responsible for actuating the joint. This relation is denoted by

$$\tau = \sum_m j_m f_m = \sum_m j_m (f_m^c a + f_m^p) = J^c A + F^p, \quad (4)$$

where $A = (a_0, a_1, \dots, a_m)$ is a vector of muscle activation. The torque τ around a joint is the sum of muscle forces f_m multiplied by its Jacobian j_m . The muscle force has two components: contractile f_m^c and passive f_m^p . The contractile force is linearly proportional to the activation level a , while the passive force is produced when the muscle is stretched beyond a threshold. The state s_{joint} of muscles around a joint is represented by its force capacity $(F_{\text{min}}^c, F_{\text{max}}^c, F^p)$, where F_{min}^c and F_{max}^c , respectively, are the minimum and maximum contractile torque the muscles around a joint can produce and F^p is the aggregate passive force at the joint. We compute the minimum and maximum contractile torques approximately from the row-wise minimum and maximum values of the Jacobian matrix J^c .

The gait state s_{gait} includes six terms $(\phi, \hat{t}, g_s, g_c, g_v, h)$, where $\phi \in [0, 1]$ is a phase of the gait cycle. The cycle begins with left heel strike at $\phi = 0$, passes the halfway with right heel strike at $\phi = 0.5$, and ends with the next left heel strike. The phase is different from the normalized time $\hat{t} \in [0, 1]$ in the cycle, although they both parameterize the gait cycle. The normalized time increases at the fixed rate from zero to one and resets to zero repeatedly in the simulation, whereas the phase has variable step sizes to simulate (possibly asymmetric) gait variations. g_s , g_c , and g_v are the desired stride, cadence, and walking speed, respectively. The foothold location h of the next heel strike can be derived from the desired stride.

Action $u = (\Delta M, \Delta\phi, \beta)$ includes pose displacement ΔM , phase increment $\Delta\phi$, and threshold β . ΔM and $\Delta\phi$ describe how the reference gait $M_o(\phi)$ should be altered spatially and temporally to properly guide the simulation at state s . Specifically, the PD target

$$\hat{M} = M_o(\phi + \alpha\Delta\phi) \oplus \alpha\Delta M \quad (5)$$

is a full-body pose of the skeletal system. Here, we use the motion displacement notation (see [Lee 2008] for detail). $\alpha \in [0, 1]$ is the confidence level of the policy. With high confidence, the action is fully reflected in gait control. We will explain how to compute the confidence α and the threshold β in the next section. The PD target thus obtained is fed into stable PD controller to convert it into the desired joint torque.

The phase is updated by $\Delta\phi$ at every simulation time step. The phase should be synchronized with the normalized time at the end of each gait cycle. To do so, we use a clipping method during reinforcement learning. If the phase does not reach one yet at the end of the gait cycle ($\phi < 1$ and $\hat{t} = 1$), the phase jumps to one to synchronize. If the phase reaches one prematurely before the gait cycle ends ($\phi = 1$ and $\hat{t} < 1$), the phase is fixed to one until the end of the cycle. This simple clipping method is sufficient because the control policy gradually learns to synchronize smoothly in the process of reinforcement learning.

4.2 Reward

The reward r includes four terms:

$$r = r_{\text{head}} r_{\text{stride}} r_{\text{vel}} + w_{\text{energy}} r_{\text{energy}} \quad (6)$$

It is well-known that humans tend to stabilize their head during walking, because visual sensory and vestibular systems are in the head. r_{head} regularizes changes in linear velocity and rotation of the head:

$$r_{\text{head}} = \exp\left(-\frac{\|\Delta v_{\text{head}}\|^2}{\sigma_v} - \frac{\|\theta_{\text{head}}\|^2}{\sigma_r}\right), \quad (7)$$

where Δv_{head} is the rate of change in linear velocity and θ_{head} is the head orientation with respect to the frame aligned to the progression direction and the global up direction. r_{stride} and r_{vel} , respectively, encourage to maintain the desired stride and walking speed.

$$r_{\text{stride}} = \exp\left(-\frac{\|h - h_{\text{desired}}\|^2}{\sigma_{\text{stride}}}\right), \quad (8)$$

$$r_{\text{vel}} = \exp\left(-\frac{\|v - v_{\text{desired}}\|^2}{\sigma_{\text{vel}}}\right).$$

r_{energy} regularizes the energy expenditure during gait.

$$r_{\text{energy}} = \exp\left(-\frac{\|a\|^2}{\sigma_{\text{energy}}}\right), \quad (9)$$

where a is the vector of muscle activation. Note that we multiply the first three terms, but the last term is added to the others. This reward design is intended to enforce head stability, desired stride, and desired velocity simultaneously, since the multiplication gets rewarded only when all three terms get rewarded. The last term is intended to gently regularize excessive energy consumption by seeking a compromise between energy efficiency and gait requirements.

4.3 Muscle Coordination

Muscle coordination is the process of determining activation levels for all muscles that achieve desired joint torques. We employed the QP formulation for muscle coordination and its conversion to regression-by-supervised-learning presented by Lee et al [2019]. Deriving from Eq (4), the QP formulation is

$$\arg \min_A \|\tau_{\text{desired}} - (J^c A + F^p)\|^2 + w_{\text{reg}} \|A\|^2 \quad (10)$$

subject to $0 \leq a_i \leq 1$ for $0 \leq i \leq m$,

The regression network π_{ψ} learns the solution of the QP by supervised learning. Since DRL is episodic, it generates a large collection of simulation roll-outs in the learning phase. The regression network is jointly learned in the DRL framework by sampling tuples (J^c, F^p, τ) from the simulation roll-outs. Since J^c is a large, sparse matrix with n columns and m rows, where n is the degrees of freedom of the skeleton and m is the number of muscles, we vectorized the matrix for regression by compactly packing non-zero values. Then, the regression network learns the relation $\tau = \pi_{\psi}(\text{vec}(J^c), F^p)$ from the tuples. The last layer of the regression network uses a sigmoid activation function to enforce the inequality condition $0 \leq a_i \leq 1$.

4.4 Simulation and Learning

In the learning phase, DRL picks random values for anatomy and gait conditions within a certain domain and runs physics-based simulation to evaluate its outcome repeatedly while gradually updating the policy and value networks. It is important to establish a stable initial configuration for each simulation. In particular, severe muscle contracture entails higher passive muscle force and a narrower joint ROM (Range of Motion). With low values for contracture conditions, skeletal poses selected from the reference gait may be outside the joint ROM. This violation can result in numerical instability of the simulation. To alleviate this risk, we set the initial configuration through pose optimization that minimizes

$$\min_q (FP)^T W_{\text{joint}} FP + w_{\text{com}} \|p_{\text{com}} - p_{\text{StanceFoot}}\|^2, \quad (11)$$

where q is a full-body pose of the skeletal system at the beginning of the simulation episode and W_{joint} is a diagonal weight matrix. The joints in the lower body are weighed more than the joints in the upper body since the lower body joints have a stronger influence on the gait. The first term penalizes the violation of the joint ROM, while the second term penalizes off-balance poses.

Although human gaits are not necessarily symmetric, the gait policy has an inherent symmetry because the control policy for left leg swing is also eligible for right leg swing if states and actions are properly reflected. We employed Phase-Based Mirroring [Abdolhosseini et al. 2019] that allows us to learn the policy only for the half of the gait cycle $\phi \in [0, 0.5]$. The policy of the other half $\phi \in [0.5, 1]$ is the mirror reflection. This simple technique cuts the computation time in half and helps to produce symmetrical gait. Note that this symmetric policy can generate asymmetric gait without any difficulty.

5 CASCADED SUBSUMPTION LEARNING

The DRL algorithm in the previous section can successfully learn a parameterized control policy with a moderate number of adjustable parameters. This algorithm does not generalize easily to deal with more (100+) parameters even with carefully planned curriculum because the policy network easily forgets what it learned while learning a subsequent batch of curriculum. In this section, we present a new idea, cascaded subsumption, to learn a control policy with 618 adjustable parameters that include 10 body parameters C_{body} and 608 muscle parameters $C_{\text{muscle}} = C_{\text{weakness}} \cup C_{\text{contracture}}$. The key to the scalability of Cascaded Subsumption Network (CSN) is its ability to learn new knowledge progressively while keeping previously learned knowledge unaltered by over-learning.

Let $\{c_i^j\}$ be a partition of the muscle parameters at level $0 < j < N$.

$$C_{\text{muscle}} = \bigcup_i c_i^j \text{ for any } j. \quad (12)$$

The multiple subsets at level $j - 1$ are subsumed into a parent set at level j . The child sets are a partition of their parent set. Each subset can have only one parent except for the ones at the base and top levels.

The CSN consists of N network levels $\{\pi_i^j\}$ for $0 \leq j < N$. There is only one network in the base level 0 and the top level $N - 1$. Each network is associated with a subset of muscle parameters except for π^0 at the base level, which does not take any adjustable

muscle parameters. The networks are learned progressively level-by-level. The base network π^0 learns a control policy of healthy (all weakness and contracture parameters are fixed to 1) individuals parameterized only by body and gait parameters. The network π_i^1 at the next level learns how muscle disorders associated with c_i^1 affect the gait. The new knowledge π_i^1 learned over the base network π^0 is reflected to gait control by overlaying actions. Specifically, we add the spatial and temporal displacements of the two networks to compute the PD target for π_i^1 . Eq (5) is replaced by

$$\hat{M} = M_o(\phi + \Delta\phi^0 + \alpha_i^1 \Delta\phi_i^1) \oplus \Delta M^0 \oplus \alpha_i^1 \Delta M_i^1, \quad (13)$$

where $(\Delta\phi^0, \Delta M^0)$ and $(\Delta\phi_i^1, \Delta M_i^1)$ are part of action of π^0 and π_i^1 , respectively. While π^0 is always confident about its action, the confidence of π_i^1 depends on state s . For example, if the muscles in state s are all healthy and normal, π^0 already knows how to control and simulate the gait. The intervention of π_i^1 risks over-learning and forgetting, so the confidence α_i^1 should be low. If some muscles in c_i^1 has weakness or contracture, π_i^1 should have higher confidence to exploit new knowledge it gained.

Algorithm 1 Confidence calculation

π^n : a network at level n
 $\pi_1^{n-1}, \pi_2^{n-1}, \dots, \pi_m^{n-1}$: Child networks of π^n .
 β : threshold value learned by $\pi^n(u|s)$
 $\alpha(u, s)$: confidence of taking action u given state s
 W_{joint} : a diagonal weight matrix of joints

```

for  $i = 1, \dots, m$  do
   $\Delta c_i \leftarrow c_i^n \setminus c_i^{n-1}$ 
   $\Delta s_{\text{joint}} \leftarrow$  Changes in joint state  $s_{\text{joint}}$  induced by  $\Delta c_i$ 
   $d \leftarrow \|(\Delta s_{\text{joint}})^T W_{\text{joint}} \Delta s_{\text{joint}}\|$ 
  if  $d_{\text{min}} < d$  then
     $d_{\text{min}} \leftarrow d$ 
  end if
end for
 $\alpha \leftarrow \text{sigmoid}(d_{\text{min}} - \beta)$ 

```

The general procedure of computing the confidence is as follows (see Algorithm 1). Let π^n be any network at level n and $\pi_1^{n-1}, \pi_2^{n-1}, \dots, \pi_m^{n-1}$ be its child networks at level $n - 1$. π^n is a super-set of any of its children. We evaluate the representation power of each network based on its force capacity around joints, which is described by $s_{\text{joint}} = (F_{\text{min}}^c, F_{\text{max}}^c, FP)$. If inclusion of new muscle parameters Δc_i makes impact on its force capacity, π^n has to learn new knowledge over its child networks and thus has high confidence. The threshold β decides how big the impact should be to have confidence. β is part of policy action learned in RL.

Muscle coordination also requires cascaded subsumption. Composition of muscle regression networks are different from composition of policy networks since muscle activation is bounded between zero and one. There is no guarantee that the sum of muscle activation outputs weighed by confidence is within $[0, 1]$. To alleviate this problem, we take muscle activation output $A_{\text{unnormalized}}$ before it passes through the sigmoid function, weighted sum the outputs

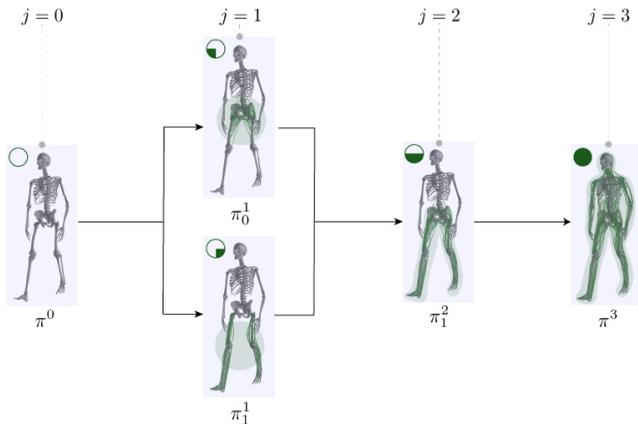


Figure 3: An example of cascaded subsumption.

of the overlaying networks, and feed it back to the sigmoid activation function. This simple technique works well with cascaded regression networks.

The first version of our Generative GaitNet has four CSN network levels (see Figure 3). The base network learned the influence of body and gait parameters. There are two networks at level 1. One of them learned the influence of muscle disorders around hip joints, while the other network learned the muscle parameters around knee and ankle joints. The network at level 2 subsumes its two child networks to learn control policies with an entire set of lower body muscles. Since the influences of upper body muscles are relatively weaker, we did not explicitly construct networks associated with the upper body at level 1 and 2. The last network at level 3 managed to learn the co-occurrence relation between upper and lower body muscles without upper body child networks.

6 EXPERIMENTS AND RESULT

Our simulation and learning system is implemented in C++ and Python. The skeletal system and Hill-type muscle contracture dynamics are implemented on the DART physics engine [2018a]. The skeleton consists of 23 rigid bones connected by 8 revolute joints at elbows, knees, and toes and 14 ball-and-socket joints. In total, the skeletal system has 50 degrees of freedom. The physics simulation time step is 480 Hz, the PD control time step is 30 Hz, and the muscle activation time step is 480 Hz. Deep reinforcement learning uses a PPO (Proximal Policy Optimization) algorithm with adaptive KL penalty and clipping implemented in Ray RLlib [2018]. Each policy/value network consists of three fully connected layers of 512 nodes and each value function, while each muscle network consists of three fully connected layers of 256 nodes. The learning parameters are summarized in Table 1.

The parameter domain is spanned by body, muscle, gait parameters. We can adjust the body and head size by 15% and the limb length by 10% in the domain. Stride and cadence ranges from -25% to $+25\%$. Ideally, both weakness and contracture parameters range from zero to one. In practice, low values for weakness parameters close to zero can have singularity issues. Therefore, their lower bound are set greater than zero. We determined the lower bound of contracture parameters based on the magnitude of passive force

Table 1: learning parameters

Policy/Value and Muscle learning rate	10^{-4}
Discount factor (γ)	0.99
GAE and TD (λ)	0.99
# of tuples per policy update	65536
Batch size for policy update	1024
Iteration for policy/value update	4
Iteration for muscle update	4-10
Maximum time horizon (sec)	10
Clip parameter (ϵ)	0.2
KL target (ϵ)	0.01
Minimum tuples per selected parameter	460

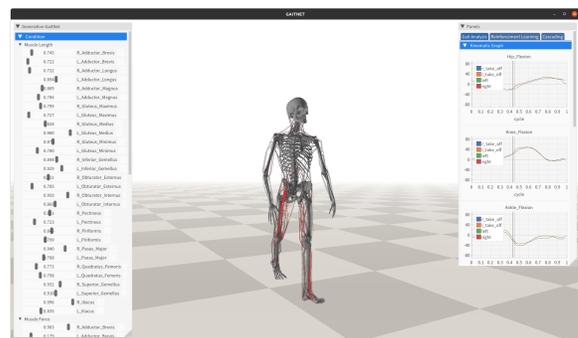


Figure 4: Interactive user interface for real-time gait simulation.

incurred by contraction at a neutral pose. The lower bound ranges from 0.5 to 0.8.

We used a cluster server with four Intel Xeon 6242 CPU 2.8GHz and a Nvidia RTX 3090 to learn the GaitNet. Learning a single network in CSN requires approximately 100 million simulation roll-outs and takes 6 hours to 24 hours in the cluster. Learning all networks of CSN takes 4 to 5 days. Once the GaitNet is learned, gait simulation runs in real-time. Through our interactive user interface system, the user can change the body size and proportion, limb length, stride, cadence, and any of muscle parameters interactively (see Figure 4). The change is reflected immediately in real-time simulation.

6.1 Pathological Gaits

We reproduced many pathological gaits with Generative GaitNet to demonstrate its expressive power. We manually tuned body and muscle parameters to match the gaits in video clips. The simulation results are best viewed in the supplementary video.

6.1.1 Foot Drop. Foot drop is a condition in which someone has difficulty lifting the front part of the foot. To simulate foot drop, we weakened muscles around the ankle (with weakness parameter 0.05). The ankle dorsiflexor muscles include tibialis anterior, extensor hallucis longus, and extensor digitorum longus. The MASS system is probably the current state-of-the-art in musculoskeletal

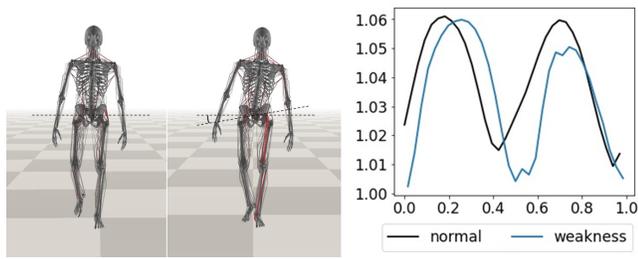


Figure 5: Normal gait and Trendelenburg gait. Pelvic obliquity plots in the gait cycle.

simulation [Lee et al. 2019]. The MASS system equipped with imitation rewards demonstrated excellent quality and performance in simulating pathological gaits induced by muscle contraction, which is enforced by inequality constraints in simulation. However, the reference-tracking method is not adept at dealing with muscle weakness probably because DRL is too smart. Even with weak muscles, it often manages to find a clever way to imitate the reference trajectory even though the result does not look human-like. Our GaitNet is not hindered by imitation requirements and thus has a better chance to simulate both weakness and contracture accurately.

6.1.2 Equinus. Equinus is a condition in which the upward bending motion (dorsiflexion) of the ankle is limited. Someone with equinus tends to walk on their toes. We tightened ankle plantarflexor muscles (gastrocnemius and soleus) to simulate toe walking (contracture 0.8).

6.1.3 Lumbar Hyperlordosis. Psoas major muscles are powerful hip flexors. Contracture (0.65) in psoas major muscles lifts the upper leg towards the body and consequently makes someone with this condition bend forward during walking. There is a secondary effect of compensating for forward bending, called lumbar hyperlordosis. Lordosis is the inward curve of the lumbar spine. Hyperlordosis often occurs to compensate for tight psoas muscles. The MASS system can only simulate the first effect (forward bending), while our GaitNet predicts the subsequent effect more accurately to produce lumbar hyperlordosis.

6.1.4 Stiff Knee. We reproduced stiff knee gaits by setting contracture in vastus muscles (0.87), tibialis posterior (0.78), peroneus tertius (0.71), gluteus muscles (0.85), and extensor digiti mini (0.87).

6.1.5 Crouch Gait. Crouch gait is characterized by knee flexion during the stance phase. This gait disorder is common among patients with cerebral palsy. We reproduced the crouch gait by setting contracture in adductor brevis (0.7), psoas major (0.65), bicep femoris short (0.65), extensor digitorum longus (0.8) extensor hallucis longus (0.75), flexor digiti minimi brevis (0.78), tibialis anterior (0.8), and tibialis posterior (0.78).

6.1.6 Trendelenburg and Waddling Gait. Trendelenburg gait is an abnormal gait caused by weakness of gluteal muscles. In normal gait, the pelvis is balanced by the gluteal muscles. The weakness of gluteal muscles on one side results in pelvic drop on the opposite

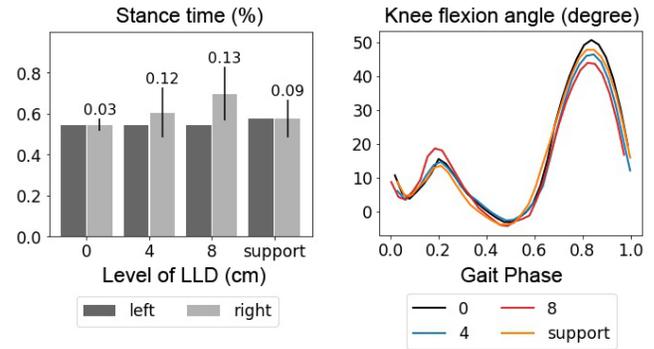


Figure 6: Leg length discrepancy.

side and the upper body leans toward the weakened side to compensate for the pelvic drop (see Figure 5). The GaitNet successfully produces both pelvic drop and upper body leaning. Gluteal weakness on both sides causes waddling gait, which is characterized by upper body swaying.

6.1.7 Leg Length Discrepancy. Leg length discrepancy is a condition in which one leg is shorter than the other (see Figure 6). It has been reported that the longer leg has a longer stance duration than the shorter leg. Bhavne et al [1999] showed this trend in their experiments with human subjects. We reproduced their experiments in physics-based simulation with GaitNet. We generated four cases. The first case has an even leg length. The second case has one leg 4cm shorter than the other. The third case has one leg 8cm shorter than the other. The fourth case uses a thick insole to compensate for 8cm discrepancy in leg length. We found that the simulation results closely matched the human experiments.

6.2 Comparison to Curriculum Learning

The biggest advantage of CSN is its ability to keep the previously learned policy unaltered while learning policies in a new parameter domain. The normal gait in a healthy condition generated by the base network π^0 is almost identical to the gait in the same condition generated by the trained CSN which consists of $\pi^0, \pi_0^1, \dots, \pi^3$. It is because the knowledge learned by the base network remains intact while learning networks in the subsequent levels. Brute-force curriculum learning with a single (non-cascaded) network does not guarantee this property. If the network learns the normal gait first and then explores in the parameter domain according to curriculum, the normal gait the network remembers would change as the learning proceeds because the new knowledge interferes the previous learned knowledge.

7 DISCUSSION

Generative GaitNet successfully learned the relation between anatomy and gait of the human musculoskeletal system. With the pre-trained GaitNet, predictive gait simulation can be a practical tool for physicians, clinical practitioners, and biomechanics/ergonomics/sports researchers seeking to understand how anatomical conditions affect human gait. We anticipate that Generative GaitNet will be used to

design wearable walking assist devices and predict the outcome of orthopedic surgical procedures for improving pathological gaits.

Cascaded Subsumption Networks are remarkably adept at dealing with high-dimensional parameter domains. With 618 adjustable parameters, simply examining all corners of the domain would require a tremendously large number ($2^{618} \approx 10^{186}$) of simulation roll-outs. We learned the first version of our GaitNet with only 340 million simulation roll-outs, which is computationally demanding but manageable in modern CPU/GPU clusters. This result shows the scalability and efficacy of our approach.

Generative GaitNet also has numerous limitations. Our musculoskeletal model lacks many anatomical features such as nervous system, ligaments, skin and soft tissue. The nervous system is particularly important because it is an integral part of the sensory-motor system. Proper computational modeling of the nervous system would allow the simulation of Parkinson gait and spasticity in cerebral palsy patients. Volumetric FEM modeling of muscles offers promising opportunities to achieve more accurate simulations of muscle contraction dynamics [Abdrashitov et al. 2021; Lee et al. 2018b, 2009].

Although Generative GaitNet currently focuses on simulating bipedal walking, there are numerous possibilities to generalize its scope because the learning procedure is based on fundamental principles, such as energy minimization and head stabilization. These principles are also applicable to any type of locomotion such as bipedal running, quadrupedal locomotion, swimming and even flapping flying [Ju et al. 2013; Yu et al. 2018].

ACKNOWLEDGMENTS

This work was supported by *Samsung Research Funding Center* under Project Number SRFC-IT1801-01.

REFERENCES

- Farzad Abdolhosseini, Hung Yu Ling, Zhaoming Xie, Xue Bin Peng, and Michiel van de Panne. 2019. On learning symmetric locomotion. In *Motion, Interaction and Games*. 1–10.
- Rinat Abdrashitov, Seungbae Bang, David IW Levin, Karan Singh, and Alec Jacobson. 2021. Interactive Modelling of Volumetric Musculoskeletal Anatomy. *arXiv preprint arXiv:2106.05161* (2021).
- Mazen Al Borno, Martin De Lasa, and Aaron Hertzmann. 2012. Trajectory optimization for full-body movements with complex contacts. *IEEE transactions on visualization and computer graphics* 19, 8 (2012), 1405–1414.
- Akhil S Anand, Guoping Zhao, Hubert Roth, and Andre Seyfarth. 2019. A deep reinforcement learning based approach towards generating human walking behavior with a neuromuscular model. In *2019 IEEE-RAS 19th International Conference on Humanoid Robots (Humanoids)*. IEEE, 537–543.
- Frank C Anderson and Marcus G Pandy. 2001. Dynamic optimization of human walking. *J. Biomech. Eng.* 123, 5 (2001), 381–390.
- Kevin Bergamin, Simon Clavet, Daniel Holden, and James Richard Forbes. 2019. DReCon: data-driven responsive control of physics-based characters. *ACM Transactions On Graphics (TOG)* 38, 6 (2019), 1–11.
- Anil Bhawe, Dror Paley, and John E Herzenberg. 1999. Improvement in gait parameters after lengthening for the treatment of limb-length discrepancy. *JBJS* 81, 4 (1999), 529–34.
- Stelian Coros, Philippe Beaudoin, and Michiel Van de Panne. 2010. Generalized biped walking control. *ACM Transactions On Graphics (TOG)* 29, 4 (2010), 1–9.
- Scott L Delp, Frank C Anderson, Allison S Arnold, Peter Loan, Ayman Habib, Chand T John, Eran Guendelman, and Darryl G Thelen. 2007. OpenSim: open-source software to create and analyze dynamic simulations of movement. *IEEE transactions on biomedical engineering* 54, 11 (2007), 1940–1950.
- Christopher L Dembia, Nicholas A Bianco, Antoine Falisse, Jennifer L Hicks, and Scott L Delp. 2020. Opensim moco: musculoskeletal optimal control. *PLOS Computational Biology* 16, 12 (2020), e1008493.
- Antoine Falisse, Gil Serranocoli, Christopher L Dembia, Joris Gillis, Ilse Jonkers, and Friedl De Groot. 2019. Rapid predictive simulations with complex musculoskeletal models suggest that diverse healthy and pathological human gaits can emerge from similar control strategies. *Journal of the Royal Society Interface* 16, 157 (2019), 20190402.
- Levi Fussell, Kevin Bergamin, and Daniel Holden. 2021. SuperTrack: motion tracking for physically simulated characters using supervised learning. *ACM Transactions on Graphics (TOG)* 40, 6 (2021), 1–13.
- Thomas Geijtenbeek, Michiel Van De Panne, and A Frank Van Der Stappen. 2013. Flexible muscle-based locomotion for bipedal creatures. *ACM Transactions on Graphics (TOG)* 32, 6 (2013), 1–11.
- Nicolas Heess, Dhruva TB, Srinivasan Sriram, Jay Lemmon, Josh Merel, Greg Wayne, Yuval Tassa, Tom Erez, Ziyu Wang, SM Eslami, et al. 2017. Emergence of locomotion behaviours in rich environments. *arXiv preprint arXiv:1707.02286* (2017).
- Jessica K Hodgins, Wayne L Wooten, David C Brogan, and James F O'Brien. 1995. Animating human athletics. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*. 71–78.
- Yifeng Jiang, Tom Van Wouwe, Friedl De Groot, and C Karen Liu. 2019. Synthesis of biologically realistic human motion using joint torque actuation. *ACM Transactions On Graphics (TOG)* 38, 4 (2019), 1–12.
- Eunjung Ju, Jungdam Won, Jehee Lee, Byungkuk Choi, Junyong Noh, and Min Gyu Choi. 2013. Data-driven control of flapping flight. *ACM Transactions on Graphics (TOG)* 32, 5 (2013), 1–12.
- Shuuji Kajita, Fumio Kanehiro, Kenji Kaneko, Kiyoshi Fujiwara, Kensuke Harada, Kazuhito Yokoi, and Hirohisa Hirukawa. 2003. Biped walking pattern generation by using preview control of zero-moment point. In *2003 IEEE International Conference on Robotics and Automation (Cat. No. 03CH37422)*, Vol. 2. IEEE, 1620–1626.
- Lukasz Kidziński, Sharada Prasanna Mohanty, Carmichael F Ong, Zhewei Huang, Shuchang Zhou, Anton Pechenko, Adam Stelmasczyk, Piotr Jarosik, Mikhail Pavlov, Sergey Kolesnikov, et al. 2018. Learning to run challenge solutions: Adapting reinforcement learning methods for neuromusculoskeletal environments. In *The NIPS'17 Competition: Building Intelligent Systems*. Springer, 121–153.
- Lukasz Kidziński, Carmichael Ong, Sharada Prasanna Mohanty, Jennifer Hicks, Sean Carroll, Bo Zhou, Hongsheng Zeng, Fan Wang, Rongzhong Lian, Hao Tian, et al. 2020. Artificial intelligence for prosthetics: Challenge solutions. In *The NeurIPS'18 Competition*. Springer, 69–128.
- Taesoo Kwon and Jessica K Hodgins. 2017. Momentum-mapped inverted pendulum models for controlling dynamic human motions. *ACM Transactions on Graphics (TOG)* 36, 1 (2017), 1–14.
- Taesoo Kwon, Yoonsang Lee, and Michiel Van De Panne. 2020. Fast and flexible multilegged locomotion using learned centroidal dynamics. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 46–1.
- Jehee Lee. 2008. Representing Rotations and Orientations in Geometric Computing. *IEEE Computer Graphics and Applications* 28, 2 (2008), 75–83.
- Jeongseok Lee, Michael X Grey, Sehoon Ha, Tobias Kunz, Sumit Jain, Yuting Ye, Siddhartha S Srinivasa, Mike Stilman, and C Karen Liu. 2018a. DART: Dynamic animation and robotics toolkit. *The Journal of Open Source Software* 3, 22 (2018), 500.
- Seyoung Lee, Sunmin Lee, Yongwoo Lee, and Jehee Lee. 2021. Learning a family of motor skills from a single motion clip. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–13.
- Seunghwan Lee, Moonseok Park, Kyoungmin Lee, and Jehee Lee. 2019. Scalable muscle-actuated human simulation and control. *ACM Transactions On Graphics (TOG)* 38, 4 (2019), 1–13.
- Seunghwan Lee, Ri Yu, Jungnam Park, Mridul Aanjaneya, Eftychios Sifakis, and Jehee Lee. 2018b. Dexterous manipulation and control with volumetric muscles. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–13.
- Sung-Hee Lee, Eftychios Sifakis, and Demetri Terzopoulos. 2009. Comprehensive biomechanical modeling and simulation of the upper body. *ACM Transactions on Graphics (TOG)* 28, 4 (2009), 1–17.
- Yoonsang Lee, Sungeun Kim, and Jehee Lee. 2010. Data-driven biped control. In *ACM SIGGRAPH 2010 papers*. 1–8.
- Yoonsang Lee, Moon Seok Park, Taesoo Kwon, and Jehee Lee. 2014. Locomotion control for many-muscle humanoids. *ACM Transactions on Graphics (TOG)* 33, 6 (2014), 1–11.
- Eric Liang, Richard Liaw, Robert Nishihara, Philipp Moritz, Roy Fox, Ken Goldberg, Joseph Gonzalez, Michael Jordan, and Ion Stoica. 2018. RLlib: Abstractions for Distributed Reinforcement Learning. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, 3053–3062.
- Libin Liu and Jessica Hodgins. 2018. Learning basketball dribbling skills using trajectory optimization and deep reinforcement learning. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–14.
- Josh Merel, Saran Tunyasuvunakool, Arun Ahuja, Yuval Tassa, Leonard Hasenclever, Vu Pham, Tom Erez, Greg Wayne, and Nicolas Heess. 2020. Catch & Carry: reusable neural controllers for vision-guided whole-body tasks. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 39–1.
- Sehee Min, Jungdam Won, Seunghwan Lee, Jungnam Park, and Jehee Lee. 2019. SoftCon: Simulation and control of soft-bodied animals with biomimetic actuators. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–12.

- Igor Mordatch, Martin De Lasa, and Aaron Hertzmann. 2010. Robust physics-based locomotion using low-dimensional planning. In *ACM SIGGRAPH 2010 papers*. 1–8.
- Igor Mordatch and Emo Todorov. 2014. Combining the benefits of function approximation and trajectory optimization. In *Robotics: Science and Systems*, Vol. 4.
- Carmichael F Ong, Thomas Geijtenbeek, Jennifer L Hicks, and Scott L Delp. 2019. Predicting gait adaptations due to ankle plantarflexor muscle weakness and contracture using physics-based musculoskeletal simulations. *PLoS computational biology* 15, 10 (2019), e1006993.
- Soohwan Park, Hoseok Ryu, Seyoung Lee, Sunmin Lee, and Jehee Lee. 2019. Learning predict-and-simulate policies from unorganized human motion data. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–11.
- Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. 2018. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–14.
- Xue Bin Peng, Glen Berseth, KangKang Yin, and Michiel Van De Panne. 2017. Deeploco: Dynamic locomotion skills using hierarchical deep reinforcement learning. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–13.
- Xue Bin Peng and Michiel van de Panne. 2017. Learning locomotion skills using deeprl: Does the choice of action space matter?. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. 1–13.
- Rémy Portelas, Cédric Colas, Katja Hofmann, and Pierre-Yves Oudeyer. 2020. Teacher algorithms for curriculum learning of deep rl in continuously parameterized environments. In *Conference on Robot Learning*. PMLR, 835–853.
- Hoseok Ryu, Minseok Kim, Seungwhan Lee, Moon Seok Park, Kyoungmin Lee, and Jehee Lee. 2021. Functionality-Driven Musculature Retargeting. In *Computer Graphics Forum*, Vol. 40. Wiley Online Library, 341–356.
- Kwang Won Sok, Manmyung Kim, and Jehee Lee. 2007. Simulating biped behaviors from human motion data. In *ACM SIGGRAPH 2007 papers*. 107–es.
- Seungmoon Song and Hartmut Geyer. 2018. Predictive neuromechanical simulations indicate why walking performance declines with ageing. *The Journal of physiology* 596, 7 (2018), 1199–1210.
- Seungmoon Song, Lukasz Kidziński, Xue Bin Peng, Carmichael Ong, Jennifer L Hicks, Serge Levine, Christopher Atkeson, and Scot Delp. 2020. Deep reinforcement learning for modeling human locomotion control in neuromechanical simulation. *bioRxiv* (2020).
- Yuval Tassa, Tom Erez, and Emanuel Todorov. 2012. Synthesis and stabilization of complex behaviors through online trajectory optimization. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 4906–4913.
- Nitish Thattai and Hartmut Geyer. 2015. Toward balance recovery with leg prostheses using neuromuscular model control. *IEEE Transactions on Biomedical Engineering* 63, 5 (2015), 904–913.
- Jianpeng Wang, Wenhui Qin, and Libo Sun. 2019. Terrain adaptive walking of biped neuromuscular virtual human using deep reinforcement learning. *IEEE Access* 7 (2019), 92465–92475.
- Jack M Wang, David J Fleet, and Aaron Hertzmann. 2009. Optimizing walking controllers. In *ACM SIGGRAPH Asia 2009 papers*. 1–8.
- Jack M Wang, Samuel R Hamner, Scott L Delp, and Vladlen Koltun. 2012. Optimizing locomotion controllers using biologically-based actuators and objectives. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 1–11.
- NFJ Waterval, K Veerkamp, T Geijtenbeek, J Harlaar, F Nollet, MA Brehm, and MM van der Krogt. 2021. Validation of forward simulations to predict the effects of bilateral plantarflexor weakness on gait. *Gait & Posture* 87 (2021), 33–42.
- Jungdam Won, Deepak Gopinath, and Jessica Hodgins. 2020. A scalable approach to control diverse behaviors for physically simulated characters. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 33–1.
- Jungdam Won and Jehee Lee. 2019. Learning body shape variation in physics-based characters. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–12.
- KangKang Yin, Stelian Coros, Philippe Beaudoin, and Michiel Van de Panne. 2008. Continuation methods for adapting simulated skills. In *ACM SIGGRAPH 2008 papers*. 1–7.
- KangKang Yin, Kevin Loken, and Michiel Van de Panne. 2007. Simbicon: Simple biped locomotion control. *ACM Transactions on Graphics (TOG)* 26, 3 (2007), 105–es.
- Wenhao Yu, Greg Turk, and C Karen Liu. 2018. Learning symmetric and low-energy locomotion. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–12.