Towards Fatigue Modeling for Locomotion Tasks

RUI XU*, DFKI and Saarland Informatics Campus (SIC), Germany
NOSHABA CHEEMA*, Max-Planck Institute for Informatics, DFKI, and SIC, Germany
ERIK HERRMANN, DFKI and SIC, Germany
PERTTU HÄMÄLÄINEN, Aalto University, Finland
PHILIPP SLUSALLEK, DFKI, Germany and SIC, Germany

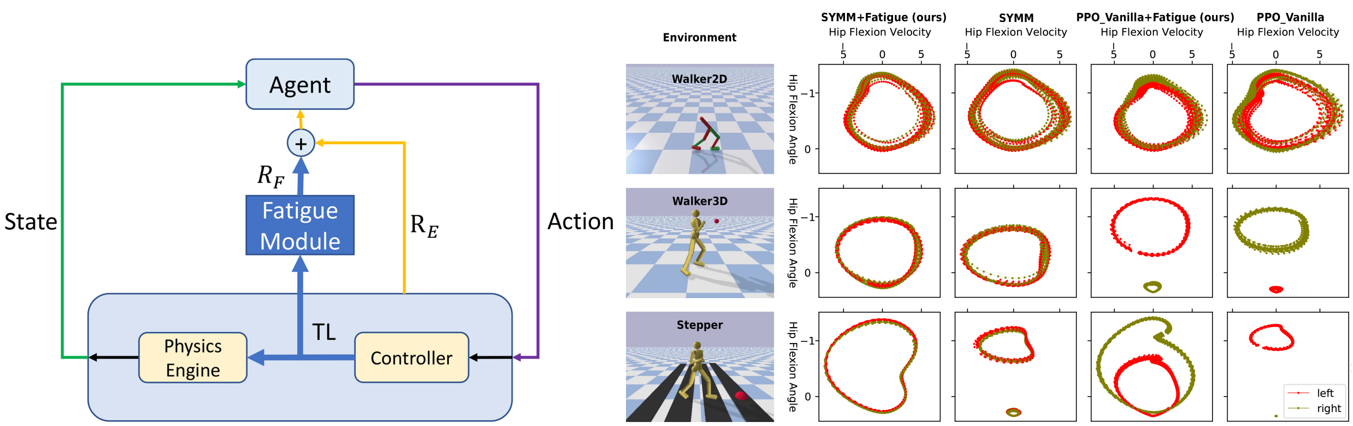

Fig. 1. *Left:* System overview. The cumulative fatigue reward R_F is calculated using the target load (TL) and is added to the environment rewards R_E. TL is the current torque used. *Middle:* Environments. *Right:* Phase-portraits.

Modern deep reinforcement learning (RL) methods allow simulated characters to learn complex skills such as locomotion from scratch. To generate realistic and smooth movements, domain-specific knowledge, such as motion capture data, finite state machines or morphology-specific attributes are needed to guide the motion generation algorithms. Here we investigate biomechanical fatigue to improve symmetry and periodicity of the generated locomotion movements compared to methods found in previous literature.

CCS Concepts: • **Computing methodologies** → **Animation**; *Physical simulation*; *Bio-inspired approaches*; *Neural networks*; • **Theory of computation** → **Reinforcement learning**.

Additional Key Words and Phrases: reinforcement learning, continuous control, character control, locomotion, animation, fatigue, biomechanics

ACM Reference Format:
Rui Xu, Noshaba Cheema, Erik Herrmann, Perttu Hämäläinen, and Philipp Slusallek. 2021. Towards Fatigue Modeling for Locomotion Tasks. In *MIG'21: Motion, Interaction and Games.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/1122445.1122456

*co-firstauthorship

MIG'21, November 10–12, 2021, Lausanne, Switzerland

ACM ISBN 978-1-4503-XXXX-X/18/06…$15.00
https://doi.org/10.1145/1122445.1122456

1 INTRODUCTION

It is a long standing task in computer animation to make characters walk on their own. In this context, Deep Reinfocement Learning (DRL) has become a promising method for automatic generation of movement controls for interactive, physics-based characters. However, in many cases the resulting motions are still not perceived as natural [Schulman et al. 2017]. A common approach to mitigate this is to use motion capture or animation data [Bergamin et al. 2019; Peng et al. 2018, 2021; Won et al. 2020]. Nevertheless, such approaches are limited to characters and movements to which data is readily available. Furthermore, obtaining qualitatively good data is oftentimes expensive, and many biomechanical constrains that are implicit in captured motions are not preserved during editing and retargeting – which is often required when data is limited. Another method for improving motion quality is to optimize for movement characteristics that shape the motion such as symmetric gait properties [Abdolhosseini et al. 2019; Yu et al. 2018] or minimal energy consumption and task goals. While such methods overcome the need of motion capture data, the absence of biomechanical constraints still may lead to unwanted behaviour and unnatural torque patterns. Another group of methods that have emerged come from bio-mechanical literature, which include musculoskeletal models and other forms of biological constraints. Previous works [Geijtenbeek et al. 2013; Lee et al. 2014; Wang et al. 2012] in this direction have explored biomimetic muscles and tendons to simulate a variety of human and animal motions. However, such muscle-based methods are usually computationally expensive, especially under a reinforcement learning framework [Kidziński et al. 2018]. In this research we work towards developing a cumulative fatigue reward, akin to [Cheema et al. 2020], based on biomechanical literature to account for a computationally efficient way to include motion constraints that are implicit in articulated figures driven by musculotendon units, in the context of locomotion. To improve on quality we further incorporate movement characteristics, such as gait symmetry [Abdolhosseini et al. 2019; Yu et al. 2018].

2 PROPOSED METHOD

Previous work in computer animation, robotics and standard RL [Peng et al. 2018; Yu et al. 2018] uses instantaneous squared joint torques as a simple measurement to minimize the effort of a given task. However, such a measure is not very biologically accurate as it does not take the duration of the given task into account and the increasing amount of perceived fatigue the longer a task is sustained. We therefore propose a reward built on cumulative fatigue based on biomechanical literature [Xia and Frey Law 2008] to mitigate this discrepancy. An overview of our proposed system can be seen in Fig. 1 *(left)*, where the target load (TL) in torque is used to compute the cumulative fatigue reward R_F, which is added to the environment reward(s) R_E. For the *Fatigue Model* the Three-Compartment Model (3CC) by [Xia and Frey Law 2008] is used. The model overview can be seen in Fig. 2. Resting motor units (M_R) become activate units (M_A) during contraction. Active motor units become fatigued (M_F) over time with the rate of F. When stopping contraction, they become rested units with the rate of R. The behaviour of the 3CC-model can be seen in Fig. 3. The rate of these motor units is computed by % of *Maximum Voluntary Contraction* ($\%MVC$), which can be the percentage of maximum torque or forced used. Since locomotion can be sustained for a long time, we assume that the target load (TL) is always $TL \leq \frac{R}{F+R}$ (see Fig. 3 *right*) and normalize the fatigue function in such a way. We then use the fatigued motor units (2 per DoF for each opposing direction) as a reward by negating their rate. We further add the mirror symmetry loss function first proposed by [Yu et al. 2018] into our model for improved movement quality by assuring a symmetric gait.

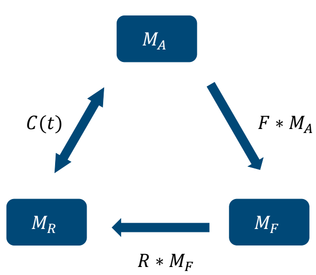

Fig. 2. Three Compartment Controller Model

3 EXPERIMENTAL RESULTS

We test our method on the *Walker2D*, *Walker3D* and *Stepper* environments by [Abdolhosseini et al. 2019] (Fig. 1 *middle*) by comparing phase-portraits to investigate gait symmetry. Our results can be seen in Fig. 1 *(right)*. We compare our method against [Abdolhosseini et al. 2019; Yu et al. 2018] who only use the symmetry loss, and against [Schulman et al. 2017] who does neither and found ours to produce more symmetric and qualitatively better results.

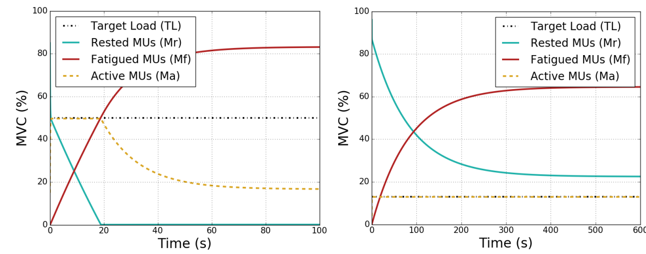

(a) $TL = 0.5$ (b) $TL = 0.13 < \frac{R}{F+R}$

Fig. 3. Behavior of the 3CC model at (a) 50% Maximum Voluntary Contraction (MVC) and at (b) 13% MVC $< \frac{R}{F+R}$. Note how the full load cannot be held any longer after 20s in (a) (yellow dashed line), while the load in (b) can be held indefinitely.

ACKNOWLEDGMENTS

This research was funded by the EU Horizon 2020 grant Carousel+ (101017779), an ITEA3 grant (01|S18060C) and an IMPRS-CS Fellowship. Computational resources were provided by the Aalto Science-IT project, as well as the BMWi under the grants 01MK20004D and 01MD19001B.

REFERENCES

Farzad Abdolhosseini, Hung Yu Ling, Zhaoming Xie, Xue Bin Peng, and Michiel Van de Panne. 2019. On learning symmetric locomotion. *Proceedings - MIG 2019: ACM Conference on Motion, Interaction, and Games* (2019). https://doi.org/10.1145/3359566.3360070

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.

Noshaba Cheema, Laura A. Frey-Law, Kourosh Naderi, Jaakko Lehtinen, Philipp Slusallek, and Perttu Hämäläinen. 2020. Predicting Mid-Air Interaction Movements and Fatigue Using Deep Reinforcement Learning. (2020), 1–13. https://doi.org/10.1145/3313831.3376701

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.

Łukasz Kidziński, Sharada Prasanna Mohanty, Carmichael F Ong, Zhewei Huang, Shuchang Zhou, Anton Pechenko, Adam Stelmaszczyk, 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.

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).

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, Ze Ma, Pieter Abbeel, Sergey Levine, and Angjoo Kanazawa. 2021. AMP: Adversarial Motion Priors for Stylized Physics-Based Character Control. *arXiv preprint arXiv:2104.02180* (2021).

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).

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.

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.

Ting Xia and Laura A. Frey Law. 2008. A theoretical approach for modeling peripheral muscle fatigue and recovery. *Journal of Biomechanics* 41, 14 (2008), 3046–3052. https://doi.org/10.1016/j.jbiomech.2008.07.013

Wenhao Yu, Greg Turk, and C. Karen Liu. 2018. Learning symmetric and low-energy locomotion. *ACM Transactions on Graphics* 37, 4 (2018). https://doi.org/10.1145/3197517.3201397