Modeling the "Gorilla Arm" Effect with Reinforcement Learning

Noshaba Cheema
Max-Planck Institute for Informatics & DFKI
Saarland Informatics Campus
ncheema@mpi-inf.mpg.de

Laura A. Frey-Law
Carver College of Medicine,
University of Iowa Health Care
Iowa City, USA
laura-freylaw@uiowa.edu

Kourosh Naderi
Games Research Group,
Aalto University
Helsinki, Finland
kourosh.naderi@aalto.fi

Jaakko Lehtinen
Aalto University &
NVIDIA Research
Helsinki, Finland
jaakko.lehtinen@aalto.fi

Philipp Slusallek
Computer Graphics Lab
Saarland University & DFKI
Saarbrücken, Germany
philipp.slusallek@dfki.de

Perttu Hämäläinen
Games Research Group,
Aalto University
Helsinki, Finland
perttu.hamalainen@aalto.fi

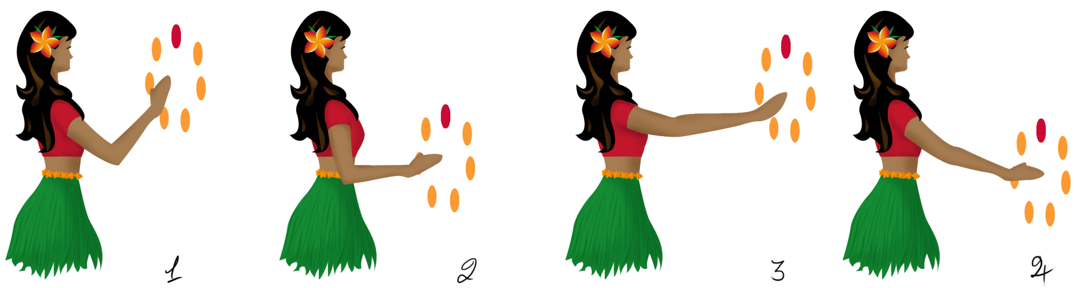

Figure 1: Four interaction zones used for determining the best model. 1) Target is shoulder height and arm is bent. 2) Target is waist height and arm is bent. 3) Target is shoulder height and arm is straight. 4) Target is waist height and arm is bent.

ABSTRACT

The "Gorilla arm" effect is a common problem of mid-air interaction which appears during excessive arm fatigue. To predict and prevent such problems at a low cost, we investigate user testing without real users, utilizing biomechanically simulated AI agents trained with Reinforcement Learning using a cumulative fatigue function from biomechanical literature. We show that the simulated fatigue data matches human perceived fatigue ratings based on the Borg CR10 scale. Our work demonstrates that deep RL combined with the fatigue model provides a viable tool for predicting both interaction movements and user experience *in silico*, without users.

ACM Reference Format:
Noshaba Cheema, Laura A. Frey-Law, Kourosh Naderi, Jaakko Lehtinen, Philipp Slusallek, and Perttu Hämäläinen. 2021. Modeling the "Gorilla Arm" Effect with Reinforcement Learning. In *MIG'21: Motion, Interaction and Games*. ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/1122445.1122456

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

1 INTRODUCTION

Touch screens or extended reality goggles enable a more intuitive and natural user experience with the use of mid-air gestures as an interaction tool. As such, physical ergonomics is an important design factor for mid-air interaction. In particular, arm fatigue, also known as "Gorilla arm effect" [Hincapié-Ramos et al. 2014; Jang et al. 2017], is a common problem that negatively effects user experience. A rising trend in design and human-computer interaction is to utilize computational models of users to predict the user experience [Biswas et al. 2012; Fischer 2001; Guckelsberger et al. 2017; Oulasvirta et al. 2018]. If this can be done with sufficient accuracy, one can rapidly evaluate alternative solutions to design problems *in silico*, without users, or at least preselect the most likely solutions to be tested in real life.

In this regard, we have contributed the first user modeling experiment that combines deep reinforcement learning (RL) with a biomechanical arm simulation model that allows both synthesizing mid-air interaction movements and predicting the associated embodied user experience, with a focus on subjective fatigue, as well as a biomechanical fatigue reward for reinforcement learning for improved motion quality [Cheema et al. 2020].

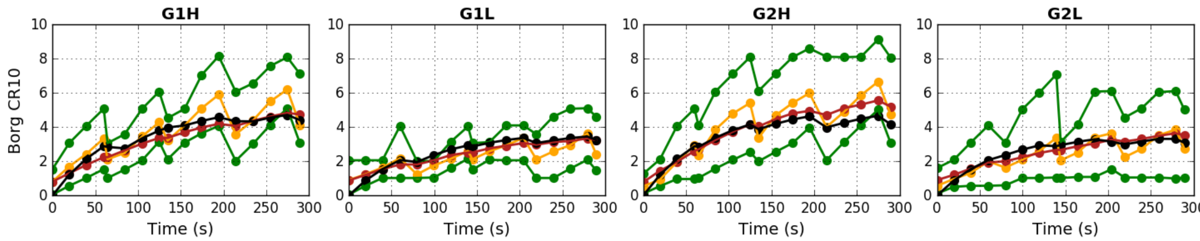

Figure 2: Results of predicting the Borg CR10 rating. Green: Upper/lower bound of ground truth. Yellow: Average of ground truth. Red: Average 3CC estimate of ground truth computed using motion capture data [Jang et al. 2017]. Black: Our simulation-based average 3CC-*r* estimate. Our simulation model yields similar modeling accuracy as [Jang et al. 2017], but does not require motion capture data.

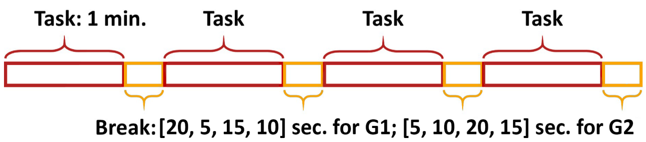

Figure 3: Experiment protocol of relaxedness vs. accuracy measure (G1), as well as during comparison against ground truth human data (G1, G2).

2 METHOD

We train 25 virtual agents consisting of shoulder, upper arm, and lower arm (5 for each interaction zone) to point at randomly highlighted targets in an ISO 9241-9 reciprocal pointing task with 7 targets (Fig. 1) using RL. The reward is based on the distance between hand and target. In addition we add a cumulative fatigue reward $\rho(t)_F$ based on the Three-Compartment Controller-r (3CC-r) model by Looft et al. [2018], which is extension of their previous 3CC-model [Xia and Law 2008] with an additional rest recovery factor r for dynamic load conditions. The reward is calculated using the difference between the active motor units M_A that can be currently used and the target load TL conditions needed for the task [Cheema et al. 2020]:

$$\rho(t)_F = \exp\left(-\frac{\left\|\frac{\vec{M}_A}{100} - \vec{T}L(\vec{T})\right\|^2}{\tau_F^2}\right)$$

With $\vec{T}L(\vec{T}) = \left[\frac{T_1^+}{T_{max}}, \frac{T_1^-}{-T_{max}}, \frac{T_2^+}{T_{max}}, \frac{T_2^-}{-T_{max}}, \frac{T_3^+}{T_{max}}, \frac{T_3^-}{-T_{max}}\right]^\top$ and $\vec{M}_A = \left[M_{A_1}^+, M_{A_1}^-, M_{A_2}^+, M_{A_2}^-, M_{A_3}^+, M_{A_3}^-\right]^\top$ (two values for each dimension of the shoulder joint roughly corresponding to opposing muscle groups [Cheema et al. 2020]). $\vec{T}$ is the current joint torque at the shoulder and T_{max} the maximum voluntary torque at a joint based on biomechanical literature [Hageman et al. 1989].

Additionally, the agents exhibit random pauses of varying length in between varying point periods and randomly switch their sex to generalize to different pointing conditions, as well as different maximum torque and weight conditions for each sex.

3 EVALUATION & RESULTS

We compare our method against ground truth human data of perceived fatigue ratings useing the Borg CR10 [Borg 1982] scale, which ranges from "nothing at all" (rating 0) to "extremely strenuous" (rating 10) obtained from Jang et al. [2017]. We futhermore compare our method against their method, which uses the 3CC [Xia and Law 2008] model to predict fatigue ratings based on torque measures estimated from motion capture data obtained from a Kinect [Zhang 2012] sensor. To compare our model we replicate the four conditions in [Jang et al. 2017]. For that we use the first two interaction zones shown in Fig. 1, and two groups with different rest periods in between the four 60 s pointing periods: [20s, 5s, 15s, 10s] for group 1 and [5s, 10s, 20s, 15s] for group 2 (Fig. 3). In Fig. 2 we refer to group 1 and 2 as G1 and G2, and the high and low interaction zones as H and L. Jang et al. [2017] use 24 participants in their study of which two were female. Since there was no ground truth data published of each participant's weight and their corresponding maximum torque estimate, we gauge their subjects in a virtual environment by using average torque and arm weight estimates found in literature [De Leva 1996; Hageman et al. 1989]. Similar to Jang et al. [2017] we assume a linear relationship between the fatigued motor units obtained from the 3CC-r model and the Borg CR10 scale with $\varphi(x) = 0.3 \cdot x$ denoting the linear mapping. An overview of our results is shown in Fig. 2. The average root mean squared error (RMSE) between the fatigue estimates from [Jang et al. 2017] (*red*) and the average Borg CR10 ground truth data (*yellow*) is 0.58, while ours (*black*) to ground truth is 0.66. However, despite using no ground truth human data for our calculations our fatigue estimates using virtual agents follows mostly the trend of Jang et al. [2017], as well as the ground truth average Borg CR10 data. In conclusion, we achieve a similar accuracy to Jang et al. [2017] just by fitting a single scaling parameter φ without using any human motion capture data.

Furthermore, we have shown that the cumulative fatigue function results in more relaxed and natural looking movements [Cheema et al. 2020], compared to existing continuous control methods in animation and RL which make use of instantaneous joint torques [Al Borno et al. 2012; Brockman et al. 2016; Peng et al. 2018; Wang et al. 2012] as a means to minimize effort.

ACKNOWLEDGMENTS

This work was funded by the Horizon 2020 EIC Pathfinder project Carousel+ (grant no. 101017779), the Academy of Finland (grant no. 299358), and an IMPRS-CS doctoral fellowship. The calculations presented in the paper were performed with the computational resources provided by the Aalto Science-IT project.

REFERENCES

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.

Pradipta Biswas, Peter Robinson, and Patrick Langdon. 2012. Designing inclusive interfaces through user modeling and simulation. *International Journal of Human-Computer Interaction* 28, 1 (2012), 1–33.

Gunnar A Borg. 1982. Psychophysical bases of perceived exertion. *Med sci sports exerc* 14, 5 (1982), 377–381.

Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. Openai gym. *arXiv preprint arXiv:1606.01540* (2016).

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. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.

Paolo De Leva. 1996. Adjustments to Zatsiorsky-Seluyanov's segment inertia parameters. *Journal of biomechanics* 29, 9 (1996), 1223–1230.

Gerhard Fischer. 2001. User modeling in human–computer interaction. *User modeling and user-adapted interaction* 11, 1-2 (2001), 65–86.

Christian Guckelsberger, Christoph Salge, Jeremy Gow, and Paul Cairns. 2017. Predicting player experience without the player.: An exploratory study. In *Proceedings of the Annual Symposium on Computer-Human Interaction in Play*. ACM, 305–315.

Patricia A Hageman, Debra K Mason, Kelly W Rydlund, and Scott A Humpal. 1989. Effects of position and speed on eccentric and concentric isokinetic testing of the shoulder rotators. *Journal of Orthopaedic & Sports Physical Therapy* 11, 2 (1989), 64–69.

Juan David Hincapié-Ramos, Xiang Guo, Paymahn Moghadasian, and Pourang Irani. 2014. Consumed endurance: a metric to quantify arm fatigue of mid-air interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1063–1072.

Sujin Jang, Wolfgang Stuerzlinger, Satyajit Ambike, and Karthik Ramani. 2017. Modeling cumulative arm fatigue in mid-air interaction based on perceived exertion and kinetics of arm motion. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 3328–3339.

John M Looft, Nicole Herkert, and Laura Frey-Law. 2018. Modification of a three-compartment muscle fatigue model to predict peak torque decline during intermittent tasks. *Journal of biomechanics* 77 (2018), 16–25.

Antti Oulasvirta, Xiaojun Bi, and Andrew Howes. 2018. *Computational interaction*. Oxford University Press.

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), 143.

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* 31, 4 (2012).

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.

Zhengyou Zhang. 2012. Microsoft kinect sensor and its effect. *IEEE multimedia* 19, 2 (2012), 4–10.