Controlling Human Head Kinematics under External Loads Using Reinforcement Learning

Sayak Mukherjee, Daniel Perez-Rapela, Jason Forman, Kelley Virgilio, Matthew B. Panzer

I. INTRODUCTION

Human Body Models (HBMs) have become the tool of choice to study human response under varied external loads. One of the major drawbacks of HBMs is the absence of a robust active muscle recruitment strategy. Active muscle response has been found to influence injury risks by altering the body kinematics and changing the effective mass as a result of muscle recruitment [1-2]. Current HBM muscle controllers are either limited to a prescribed muscle activation history or joint torque, or use a closed-loop controller to actuate specified groups of muscles [3-4]. Both methods rely on a pre-defined strategy for the complex coordination of muscle actuation schemes that are tuned to specific loading cases, and may not be applicable for the general loading scenarios that are associated with HBMs. In this study, we integrate an alternative active muscle control based on reinforcement learning (RL) into an HBM. The RL approach allows the HBM to learn how muscle activation affects kinematics like how humans learn to use their muscles, thereby achieving a muscle activation model without a predefined recruitment strategy. We demonstrate this technique using a multibody (MB) head-neck model with 23 independent muscle pairs.

II. METHODS

Multibody Head-neck Model

The MB head-neck model was based on a similar finite element model developed by [5]. The model consists of eight vertebrae (T1 – C1) and the head represented by a Hybrid-III headform with the same mass and inertial properties as in the original model (Figure 1a). Vertebrae T1 – C2 were modelled as separate bodies with C2 and C1 combined as one. The head and vertebrae were connected using a 6-DOF joint with non-linear stiffness and damping responses [6-8]. The muscles were modelled as Hill-type muscle with an active contractile element [9] and passive parallel element with suitable origin and insertion points. Each muscle segment was split into several equal length sections along the length and the ends of each section were attached to the nearest vertebra considering the changing load direction during simulation (Figure 1b).

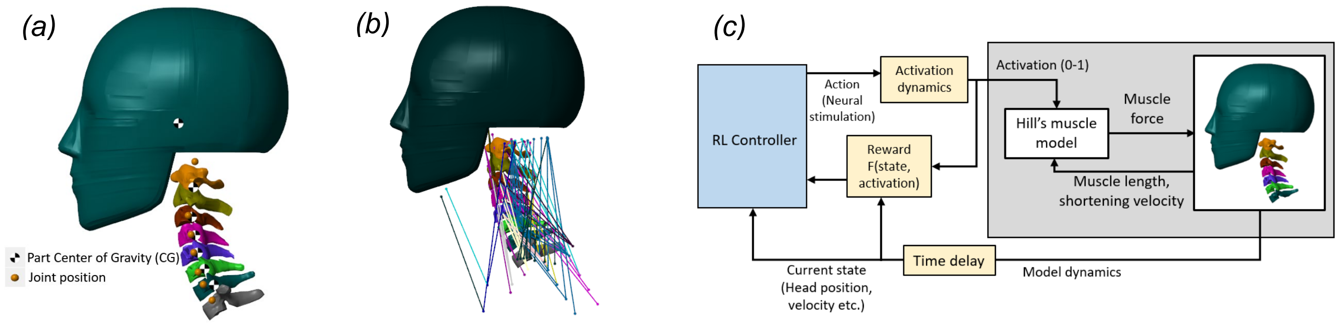

Fig. 1. (a) Multibody head-neck model; (b) Head-neck model with muscle elements; (c) RL control framework.

Muscle Recruitment Model

Reinforcement learning (RL) is a branch of machine learning which uses an iterative approach to train the controller to perform best sequence of actions in a given state. Each state is associated with a reward level that helps the system learn to maximise the use of states associated with high rewards, in order to achieve a desired objective. In this study we developed the RL muscle controller based on the actor-critic algorithm [10]. The actor-critic algorithm utilises two deep neural networks to accomplish the desired goal, actor network which takes the state parameters as inputs and determines the best action, and critic network which evaluates the

S. Mukherjee (e-mail: sm8ma@virginia.edu; tel: +1 434.297.8042) is a PhD Student, D. Perez-Rapela is a Research Associate, J. Forman is a Principal Scientist, and M. Panzer is an Associate Professor at the Center for Applied Biomechanics in the Department of Mechanical and Aerospace Engineering, University of Virginia, USA. K. Virgilio is a Research Associate at Luna Innovations Inc., Charlottesville, Virginia, USA.

action. The current states were defined using HBM kinematics (position, velocity and intervertebral angles). The actor network comprised of three layers, with the outputs of the first two layers activated using rectified linear (ReLu) function. The output of the final layer was activated with sigmoid function to limit the output values between 0 and 1. The output layer has 23 nodes, one for each muscle pair which provides the neural excitation response based on the current state of the model (Figure 1c). The critic network was modelled with one hidden layer between the input and final layer with ReLu transfer function after the input and hidden layer. The input includes the state parameter values as well as the actions from the actor network, and the network output estimates the significance of the action taken in the present state.

For this preliminary study, we trained the RL model to keep the head and neck in the neutral position under gravity. Training was accomplished using Matlab by applying the initial velocity of 1 m/s to the head centre of gravity (CG) in random orientations in the sagittal plane. The reward function penalised the model for distance of the head CG from the initial position, head velocity in all 6 DOFs, the rotation between two successive vertebrae, and sum of the muscle activations. Each iteration ran for maximum 750 ms.

III. INITIAL FINDINGS

With adequate training (up to 1050 simulations to convergence), the controller learned to generate a muscle activation scheme that maintained head-neck stability under gravity, in the presence of moderate head perturbation. Testing the RL muscle recruitment strategy with anterior and posterior initial velocities to the head demonstrated that the muscle scheme could stabilise the head within 300 ms for both the initial condition (Figure 2) compared to the purely passive case which resulted in unstable biomechanics (as expected).

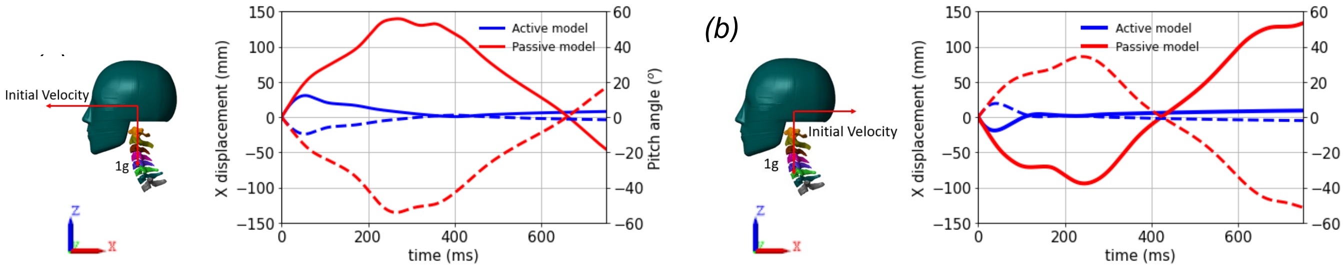

Fig. 2. Head CG displacement (solid) and pitch angle (dashed) for between the active and passive head-neck models for two perturbation conditions defined by an initial velocity of the head: a) anterior and b) posterior. The head pitch angle is measured with T1 CG as origin.

IV. DISCUSSION

The current study is a preliminary investigation into the use of machine learning for the training and control of muscle response with HBMs. The major advantage of the RL control framework is that, apart from the reward function, it requires negligible user input or assumptions on how the muscles should work together to behave like a living human. The training of a RL model is an ongoing and evolving process, and the existing model will continue to improve with additional training scenarios. The biofidelity of the controller will be determined by calibrating the components of the reward function based on the many different sources of human volunteer data. We anticipate that the control framework developed in this study will eventually result in a general-use muscle control model that can produce realistic head kinematics under omni-directional automotive or sport loading environments. The current control framework can also be generalised to other body regions as well, but is best suited for complex, multi-DOF systems with many muscles like the head and neck.

V. REFERENCES

[1] Brolin et al., Traffic Inj Prev, 2005.
[2] Choi HY et al., ESV, 2005.
[3] Bayer A et al., J Neurophysiol, 2006.
[4] Martynenko OV et al., ESV , 2019.
[5] Panzer MB et al., Med eng & Phy, 2011.
[6] Shea M et al., J Biomech, 1991.
[7] Liu YK et al., AFAMRL, 1982.
[8] Dibb et al., J Biomech, 2009.
[9] Zazak FE et al., Crit rev Biomed eng, 1989.
[10] Lillicrap TP, ICLR, 2016.