A

Presented to
the faculty of the School of Engineering and Applied Science
University of Virginia

in partial fulfillment
of the requirements for the degree

by

APPROVAL SHEET

This

is submitted in partial fulfillment of the requirements
for the degree of

Author:

Advisor:

Advisor:

Committee Member:

Committee Member:

Committee Member:

Committee Member:

Committee Member:

Committee Member:

Accepted for the School of Engineering and Applied Science:

Jennifer L. West, School of Engineering and Applied Science

ABSTRACT

Computational human body models (HBMs) are important tools used in biomechanics research to predict human responses under external loads. To improve the predictive capabilities of HBMs in some loading scenarios, it becomes important to incorporate a motor control mechanism that affects the human response by generating active muscle forces around skeletal joints. Current efforts to integrate muscle control into HBMs rely on feedback-based controllers which have been precisely tuned for specific load cases and may not generalize to cases beyond which it has been tuned. The dissertation proposes a novel approach to muscle control based on deep reinforcement learning (RL).

Reinforcement learning algorithms are recent advancements in the field of machine learning and allow for a complex system to learn how to work successfully to achieve the desired outcome and is analogous to how a child learns to walk by trying over and over again by constantly interacting with the surrounding environment. The proposed research presents a comprehensive study on the design of RL Muscle Activation Control (RLMAC) using detailed musculoskeletal multibody (MB) models and volunteer testing data. The central goal of this dissertation was to evaluate the utility of RL-based algorithms in muscle control for generating voluntary kinematics with eventual application to complex external loadings. The dissertation also examines the application of the trained muscle controller to changes in anthropometry, the addition of external mass to the body (such as helmets), and changes to the external loading environments.

An initial proof-of-concept study on the use of RLMAC was performed using a multibody model of the human arm incorporating muscles responsible for motion about the elbow. The human arm model provided a simple model setup with a revolute joint at the elbow which made it convenient for the preliminary analysis. The RLMAC was trained to perform extension and flexion movement

of the lower arm by activating the muscles, and the trained controller could generate goal-directed arm movements, synthesize the same motion in the presence of an external force field, and the trained controller could also maintain the stability of the elbow joint to high magnitude impulse loads. Following the initial investigation, the RLMAC was integrated into the head and neck body region MB model with the anatomy of a 50^{th} percentile male. The presence of multiple non-linear joints and the complex muscle orientation of the cervical spine make the head-neck complex a suitable body region to evaluate the ability of RL-based controls to generate coordinated muscle forces for joint control. At first, the RLMAC was trained assuming symmetry about the sagittal plane (i.e., left and right muscles were assigned identical activations). With the symmetrical control model, the RLMAC architecture was developed to maintain stability under gravity and synthesize the extension-flexion motion of the neck. The same control approach was then extended to all DOF models where each muscle was activated individually to maintain the desired posture.

A series of volunteer tests were performed to finetune and calibrate the architecture of the RL controller. The test subjects were asked to perform fast goal-directed rotations of the head in the sagittal plane (extension-flexion) and transverse plane (axial rotation), and the data gathered from the volunteers were used as datasets for the model validation. The trained RLMAC could replicate the desired head movements with both the symmetry model and the all-DOF model.

Finally, range of applicability studies were performed to gauge the ability of the RL controller to adapt to novel scenarios and develop responses to external loads, for which it has not been explicitly trained. The trained RLMAC was able to adapt to changes in anthropometry and was also able to maintain stability with an added mass representing a helmet. The trained model could also react to impact loads which provide evidence of its potential for controlling HBMs under novel loading environments for which it has not been previously trained.

2

The dissertation provides a detailed insight into the development of general use HBM muscle controller with the capability to simulate commonly encountered chaotic scenarios. The proposed approach can be extended to other body regions as well with the eventual application at the whole-body level. Active HBMs will serve as important tools for the development of improved injury mitigation devices by accurately predicting the response and thus, the injury risks. Furthermore, the RLMAC framework can also be used in biomechanics applications such as gait and occupational health research.

Contents

LIST OF FIGURES

11

LIST OF TABLES

LIST OF ABBREVIATIONS

HBM	Human body model
RL	Reinforcement learning
RLMAC	Reinforcement learned muscle actuation controller
MB	Multibody
BC	Boundary conditions
FE	Finite element
CNS	Central nervous system
IAMR	Individual Activated Muscle Recruitment
GAMR	Group Activated Muscle Recruitment
DDPG	Deep deterministic policy gradient
MTU	Muscle-tendon unit
CE	Contractile element
PE	Passive element
SEE	Series elastic element
SDE	Series damping element
a_t	Activation
u_t	Neural stimulation
F_{max}	Maximum muscle force
L_{opt}	Optimum muscle length
τ_{act}	Activation time constant
τ_{deact}	Deactivation time constant
ReLu	Rectified linear activation unit
NN	Neural Network
OU	Ornstein-Uhlenbeck
TD3	Twin-delayed deep deterministic policy gradient
VCR	Vestibulocollic reflex
CCR	Cervicocollic reflex

14

Chapter 1 – Introduction and background

1.1 Statement of Problem

Neuromuscular motor control mechanisms in humans are responsible for voluntary body motions as well as involuntary reactions to external forces and loads. The desired motion responses are achieved when the central nervous system (CNS) generates signals that actively contract skeletal muscles, resulting in forces that will move or stiffen up the body joints. However, human neuromuscular control is complex, even for basic motor functions such as postural stability. The musculoskeletal system has non-linear joints with multiple kinematic degrees of freedom (DoF), and when active muscle function is integrated into this structure, it results in an indeterminate problem with an infinite number of solutions of muscle activation to produce a similar motion outcome. This poses a substantial mathematical challenge when developing computational human body models (HBMs) that incorporate active muscle behavior and muscle control schemes.

Previous control studies have widely used two different approaches to model the active muscle forces. The first approach, known as open-loop control, applies a predetermined activation history to different muscles or muscle groups for maintaining posture under gravity and executing the desired motion of a joint or body segment. The open-loop control method has been effective in demonstrating the overall effect of muscle forces on model response or injury risks, however, the utility of such a simplified method is limited to cases where HBM posture or muscle activity is known in advance for the specific simulated load case. Thus, open-loop muscle control schemes need to be extrapolated to the general case, which may not represent real behavior.

The second approach, known as the closed-loop system, incorporates a feedback controller in the HBM environment which formulates the muscle activity levels depending on any physical signal error that can be measured in the model. The model signal is compared with a reference signal

15

describing the intended position of the joint, and the error is fed to the controller that generates the muscle activation levels. The closed-loop controller is widely used in physiological studies and for predicting responses under external force fields. However, one of the major limitations of a closed-loop controller is that they require tuning of the controller parameters based on achieving a desired output to a specific load case. Like open-loop controllers, they may not generalize beyond the scenarios for which they have been tuned or prescribed. Due to the inherent challenges of tuning PID controllers for multiple muscles simultaneously, researchers often simplify the muscle control scheme by grouping muscles around a joint by their agonist-antagonist function to reduce the complexity associated with joint movements. While such simplification can work reasonably well for simple joints such as elbows and knees, the redundant nature of more complex joints such as the neck vertebras makes it unsuitable for wider application.

Over the past decade, reinforcement learning (RL) has emerged as a strong alternative to traditional control loops. RL algorithms are advancements in the field of machine learning, where the controller learns to output the best sequence of control signals based on the control objective by constantly interacting with the system that it seeks to control. In the Reinforcement learning routine, the controller, also known as an agent, receives feedback on the value of the action concerning the overall goal which enables it to improve the control output over multiple iterations. The drawbacks of current active muscle control in HBMs can potentially be addressed using RL approaches as the algorithms are adaptive in nature and can generate multiple coordinated control outputs for non-linear systems. The major advantage of RL over conventional control systems is that the RL agent requires minimal input of the mechanics of the system. RL can also leverage the developments in supervised learning such as deep neural networks (NN) to control complex systems with limited domain information, which is useful for complicated body regions such as

16

the head and neck, where the muscles cannot be associated with any isolated kinematic DoF. Modeling the kinematics of skeletal joints by controlling the muscle actuations which can predict the human response under chaotic scenarios (such as automotive or sports impacts) is a complex problem. This dissertation seeks to explore the use of a control framework based on reinforcement learning to develop the active control mechanisms that are associated with humans for voluntary and autonomous control.

1.2 Human Musculoskeletal System

Specific movements in the human body happen when skeletal muscles apply forces to bones. Forces generated by skeletal muscles are responsible for maintaining posture and generating voluntary movements about joints. The magnitude of forces in muscles required to bring about a particular joint movement depends on the external environment and is controlled by the central nervous system (CNS). The skeletal muscles have varying shapes, depending upon the location of the muscles and the arrangement of muscle fibers. The forces produced by the skeletal muscles are proportional to the number of fibers across the muscle length (Figure 1-1).
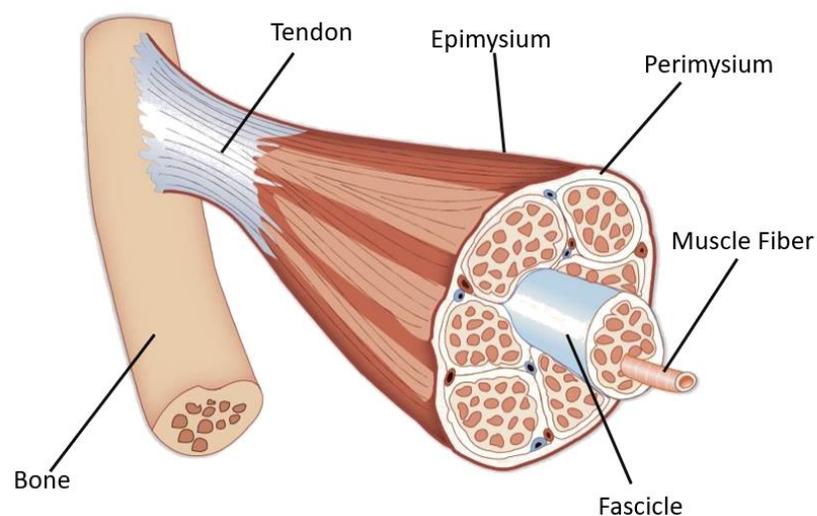
Figure 1-1: Anatomy of skeletal muscles. Adapted from google images (slideteam.net).

Skeletal muscles can only apply active tensile forces to shorten the muscle, i.e., when humans actuate a muscle, the muscle force attempts to pull the origin and insertion points toward each other. The muscle fibers are composed of long cylindrical structures called myofibrils. Myofibrils can further be divided into myofilaments, which are made up of repeated bands of proteins called actin and myosin, which results in the striated appearance of the skeletal muscles. Each band is known as the sarcomere, which is the primary contractile unit of the skeletal muscles. When a muscle is activated by the nervous system, the actin filaments slide over the myosin filaments, shortening the sarcomeres, which causes contraction of the muscle unit (Figure 1-2).
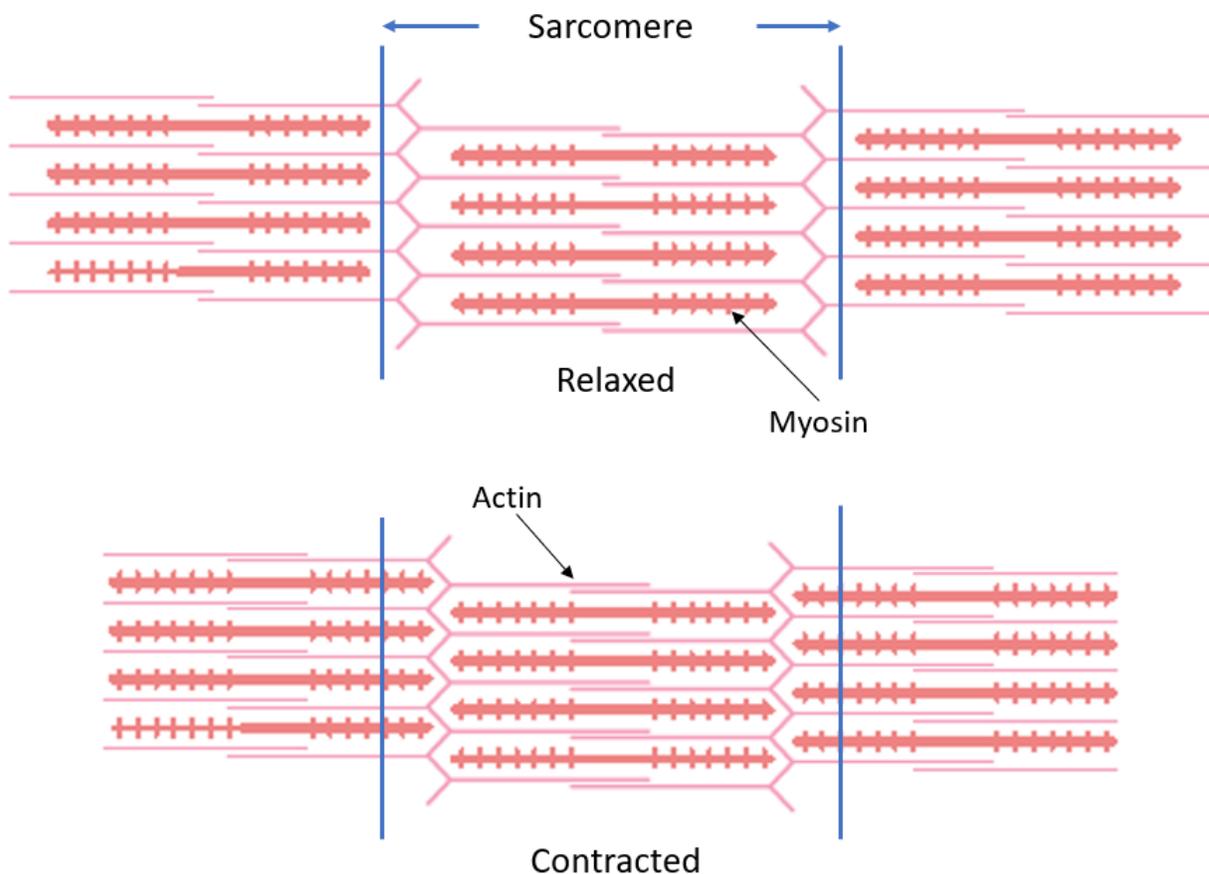
Figure 1-2: Sarcomere in relaxed (top) and contracted (bottom) state. The actin filaments slide over the myosin which leads to the contraction of the muscles. Adapted from google images (brainkart.com).

18

In numerical modeling, the Hill muscle model is used to represent the mechanical characteristics of the skeletal muscles (Hill, 1938). Hill proposed a two-component model to numerically represent the active forces generated by the muscles, the first being the contractile element that represents the active muscle forces, and a non-damping elastic element in series that represents the compliance of the tendons at the muscle ends. Hill's muscle model was later modified to include a passive element in parallel with the active element, which represents the tensile forces generated by stretching the inactivated muscle (Winters, 1995; Zajac, 1989).
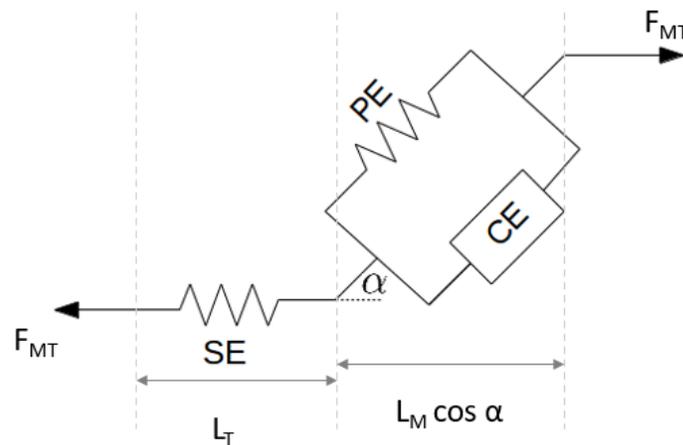
Figure 1-3: Hills muscle model with the active contractile element (CE), a passive element (PE), and a series element (SE).

Figure 1-3 displays the three-element Hill-type muscle model (Millard et al., 2013; Zajac, 1989). The pennation angle (α) is the angle at which the muscle fibers attach to the tendons (Figure 1-3). The muscle forces (F_M) are non-linear in nature and depend on the muscle physiology (length, area of cross-section, etc.) and the velocity of contraction. The total muscle forces are the sum of forces produced by the contractile element (F_{CE}), and the passive element (F_{PE}) (Equation 1-1) and are equal to the forces generated by the serial element (F_{SE}) (Equation 1-2).

$$F_M \ = \ F_{CE} \ + \ F_{PE} \qquad\qquad \text{Equation 1-1}$$

$$F_{MT} \ = \ F_M \ Cos \ \alpha \ = \ F_{SE} \qquad\qquad \text{Equation 1-2}$$

In some previous studies, the serial element and the pennation angle were ignored (Buhrmann and Di Paolo, 2014; Östh et al., 2012b; Panzer et al., 2011) while considering joint stabilization and voluntary motion. In the muscles of the upper extremities and the neck region, where the tendon slack lengths are comparable, such an assumption was not found to affect the overall response (Gribble et al., 1998; Lemay and Crago, 1996). While simulating fast arm movements, Bayer et al. (2017) found that the series element only contributes around 7.6 % towards the arm velocity using a muscle model with a serial damping element. However, in lower extremities where the tendon slack lengths are much longer, the effect the of series element was found to be significant (Bobbert, 2001; Scovil and Ronsky, 2006).

The active muscle forces are managed by the central nervous system (CNS), which directly controls the level of actuation (a_t) of each muscle. In Hill's model, the activation level varies between 0 (passive) and 1 (fully active). Physically, the activation level for the muscle represents the amount of relative effort of a muscle and will depend on the desired joint movement (voluntary movements, stability, etc.) and external environment (gravity, added inertia like helmets, etc.). The total forces generated by active element is represented by Equation 1-3 (Bahler et al., 1967; Winters, 1995; Zajac, 1989).

$$F_{CE} = a_t \text{ x } F_{max} \text{ x } F_l(l) \text{ x } F_v(V) \qquad \text{Equation 1-3}$$

The f_l and f_v terms in Equation 1-3 represent the normalized force-length and force-velocity relationship respectively (Figure 1-4). F_{max} is the maximum isometric force a muscle can generate and is dependent on the muscle cross-section area which is the indication of the amount of muscle fibers along the muscle length.

(a)   *(b)*
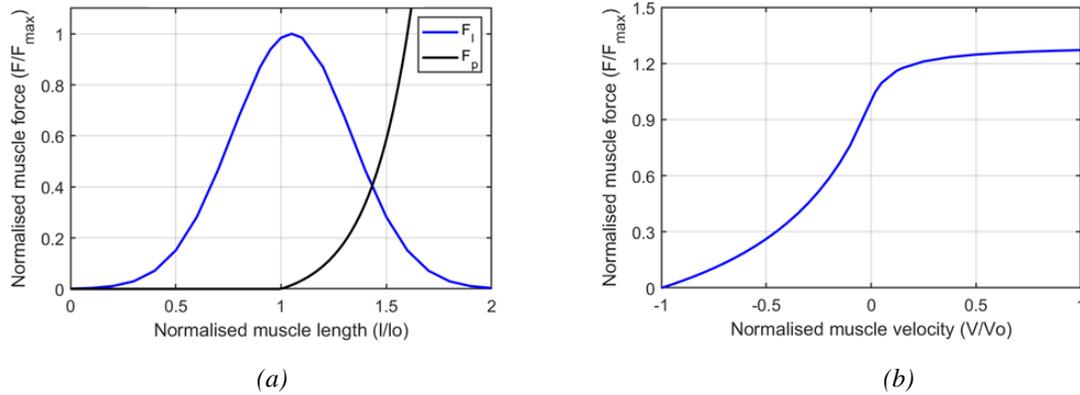
Figure 1-4: Hill-type muscle parameters (a) Force-length relationship of contractile (F_l) and passive (F_v) elements (b) Force-velocity (F_V) relationship of the contractile element.

The maximum force F_{max} in the active element develops when the muscle length is equal to a characteristic length or optimum length l_{opt}, and the muscle is neither shortening nor lengthening (isometric). The F_v curve describes the force-velocity relationship, which shows that when muscle elongates, the force asymptotes close to F_{max} (Figure 1-4b).

The active muscle forces only act to shorten or contract the muscles, i.e., the active muscle forces pull the origin and insertion points towards each other. For this reason, muscles are always arranged as agonist and antagonist groups for a specific DoF of a joint, and the kinetics of that joint depends on the balance of forces between the antagonistic pairs. When the agonist group is activated by the CNS, the antagonist muscles are stretched and provide passive stability. Sometimes under external loads, both the agonist and antagonist groups may co-contract, thus increasing the static stiffness (and internal load) of the joint, without undergoing any joint displacement.

1.3 Human Motor Control

Motor control in humans can be classified into two broad categories. The first category is called postural control, in which the nervous system tries to maintain a specific joint position. In some cases, like in maintaining the neutral head position under gravity, the passive stiffness of the

21

cervical spine and muscles is not enough to counteract the external force field of gravity and requires stabilizing muscle activations (Forbes et al., 2013; Goldberg and Cullen, 2011; Keshner et al., 1999). In most cases, postural maintenance does not require conscious effort and can be achieved through the peripheral nervous system (PNS), visual and vestibular input, and proprioception (Morningstar et al., 2005; Winters, 1995). Under some conditions where reflexive control is not sufficient, muscles can also co-contract to provide additional stability (Choi, 2003; de Vlugt et al., 2006; van Drunen et al., 2013).

The motor control mechanism is also responsible for the second category, which is synthesizing voluntary, target-directed movements around the joints. While performing a motion, the specific joint target position is determined by the CNS, and accordingly, the muscle activations are adjusted such that the net forces and moments about the joint, including those caused by external force fields (such as gravity), are zero at the target position. In theory, there can be multiple muscle activation patterns that can cause the identical desired joint orientation due to the inherent redundancy of the musculoskeletal system (Bernshteĭn, 1967; Wolpert, 1997). To determine a single activation pattern, several cost functions have been suggested in previous control studies to minimize redundancy. Cost functions have included kinematics-based parameters such as linear (Hoff and Arbib, 1993; Nelson, 1983) and angular acceleration (Ben-Itzhak and Karniel, 2008), or kinetic parameters like joint torques (Nakano et al., 1999; Uno et al., 1989), joint energy (Berret et al., 2008a) or muscle forces and energy (Koelewijn et al., 2019; Umberger et al., 2003).

The ability of humans to perform motor tasks also adapts to changes in external loads. Happee (1993) found that on repeatedly performing goal-directed movements with added inertia, the trajectory of the motion would adapt to the variation in the inertia. A similar trend was also observed by Oh et al, (2021) when stiffness and damping of the motion were also varied along

22

with the inertia. Smeets et al., (1990) argued that such complex adaptation to changes in the external environment cannot be easily explained by linear closed-loop control mechanisms.

The human musculoskeletal system is also redundant in nature, i.e., there are more muscles than possible degrees of freedom for a joint. The complexity of isolating a group of muscles based on the preferred direction of motion is exacerbated by the fact that a muscle can control the motion in two different DoFs for the same joint and across multiple joints in some body regions like the head and neck. The complicated orientation of the muscles along with various non-linearities in joint stiffness and delays associated with the neuromuscular system makes modeling of the motor control process in the HBMs challenging. The ambiguity and non-linearity associated with motor control and planning in humans require the evaluation of alternative control strategies to accurately represent such mechanisms in computer models.

1.4 Active Muscle Control in HBMs

Hill-type muscle model with suitable insertion points is most commonly used for muscle control studies in HBMs. The contractile element of Hill's model generates the active forces (Equation 1-3) and requires the activation level (a_t) as input to accurately scale the force magnitude for a given muscle length and velocity.

In an open-loop system, a predefined activation-time pulse is applied to each Hill muscle in the model. During the course of the simulation, the activation is not informed from the state of the model at any time step. This simplified approach was employed by de Jager et al. (1996) in a multibody (MB) model of head and neck to study the effects of muscle forces on impacts. The model was simulated in frontal and lateral impact conditions and compared with the passive model, and it was found that the active model compared better with the volunteer test data. The model

was then improved by van der Horst et al. (1997) by modeling accurate neck muscle trajectories. The accuracy in muscle representation improved the rotational response of the head and neck under high-severity impact conditions (15g frontal). However, the model did not consider the gravitational field which affected its response to low-severity impacts. Brolin et al. (2005) developed a finite element (FE) model of the head and neck region to evaluate the effects of muscle activations on spinal responses to frontal and side impacts. In the study, it was found that in frontal impact simulations, muscle activation reduced the loads on upper spinal ligaments whereas in the lateral impacts, reduced the injury risk throughout the spine. Dibb et al. (2013) performed FE simulations to study the effect of muscle activation on the response of models of different age groups (6-year-old, 10-year-old, and midsize adult) against low-speed frontal impact. The study found that when tensed, neck kinematics correlates well when compared with the volunteer corridors. Iwamoto et al. (2011) developed a FE model at the full body level to study the effect of pre-impact muscle activity on injury outcomes. The musculature in the model was represented as a combination of solid elements for passive stiffness and bar elements for the active forces. The study concluded that bracing may constrain the thoracic region and cause more bony fractures at the upper and lower extremities compared to the passive condition where more rib injuries were observed. In a FE study of SUV-to-pedestrian impacts, Iwamoto and Nakahira, (2014) developed a FE model of a mid-size male with 282 muscles of the trunk, upper extremities, and lower extremities. The simulations demonstrated that active musculature affects the location of the head impact and head injury risks. While the open-loop modeling approach has been able to explain the differences in response between the passive models and real-world data, the methodology cannot be extended for general use where the exact kinematics of the HBM are unknown during the simulation.

Closed-loop controllers have become the state of the art in motor control modeling studies and utilize a proportional-integral-derivative (PID) controller to output the muscle activation levels based on the difference between the current value and the target value of a physical signal associated with the model. Kistemaker et al. (2006) developed a multi-body model of the upper arm for generating rotational motion about the elbow. The developed arm model had three segments and two revolute joints representing the shoulder and the elbow, however, in the computational study the shoulder joint was fixed. The arm musculoskeletal model was simplified by simplifying the muscles into two extensors and two flexors. The muscles were activated by PID controllers with muscle length feedback that dragged the lower arm to the desired joint angle. Östh et al. (2012) implemented a closed-loop control in the FE model of a human arm with a revolute elbow joint. The FE model incorporated seven individual Hill muscles which were then grouped into extensors and flexors. The PID controller was tuned with elbow angle error to rotate the lower arm to a target angle and stabilize it in the presence of gravity. An active MADYMO model with spinal torque actuators was developed by Cappon et al. (2007) with PID controllers to stabilize the spine under perturbations. Optimization runs were carried out to tune the control parameters. The model was found to correlate with the pre-rollover phase. Happee et al. (2017) developed a cervical spine model with reflex control representing vestibular (VCR) and muscle (CCR) afferent feedback along with co-contraction. The controller was fed the sagittal angle and angular velocity, head linear acceleration, muscle length, and muscle velocity. A control parameter was defined or co-contraction along with a baseline activation for postural stability. The active cervical model was validated with experimental head translation data.

Prior studies have also implemented the closed-loop approach at the full body level to accurately predict response under impact scenarios or for physiological simulations. Meijer et al. (2008)

modeled an active occupant model with PID controllers at the elbow and spine for far-side impacts. The head kinematics with the active spine model was found to differ from the passive data. Östh et al. (2012) integrated three PID controllers in a Total Human Model for Safety AM50 v 3.0 (Toyota Motor Corporation 2008) FE model to evaluate the response of a vehicle occupant to autonomous braking. The PID controllers were used to actuate the muscles of the head, neck, and lumbar regions, and it was found that the cervical muscle activities are important to capture the response of occupants. Walter et al. (2021) developed a control architecture to simulate an upright stance and squatting under gravity, using hierarchical controllers based on joint angles and muscle lengths. The hierarchical control model was implemented in a simplified MB model of the skeletal system with 20 angular DoFs with only one antagonist unit per DoF. Devane et al. (2022) modified the Global Human Body Model Consortiums (GHBMC) simplified model for midsize males and small females by incorporating skeletal muscles as beam elements. The active model consisted of 32 PID controllers with joint angle error and 210 controllers with muscle length feedback. The active models were validated under pre-crash braking and low-speed impacts and were found to correlate with the volunteer data better.

Another approach often employed in active control studies is the use of an optimization routine to determine the muscle recruitment scheme for desired postural control. Chancey et al. (2003) developed an MB model of the 50th-percentile male cervical spine to stabilize it under gravity. The model was simulated under two different optimization objectives while maintaining the head stability within 5° - minimizing muscle fatigue (relaxed) or maximizing muscle forces (tensed). A similar optimization scheme was implemented by Dibb et al. (2013) to stabilize the cervical spine in the FE model of 6-year-old, 10-year-old, and midsize adults constrained at T1. Two separate activation routines were performed corresponding to relaxed and tensed states to counter the effect

of gravity. At the full body level, Bose et al. (2010, 2008) performed sensitivity studies using MB models representing 50th percentile drivers with 1-D Hill's muscles in the lower extremity. Optimizations were performed to determine the muscle actuation levels to maintain the stability of the joints by varying the driver anthropometries and driving positions. The study found that occupant statures such as posture and pre-impact bracing are important factors influencing collision-induced injuries.

While closed-loop control and dynamic optimization are improvements over simple open-loop control architecture and have influenced the development of injury mitigation systems, both these methodologies have major limitations. Optimum muscle activation levels (or PID parameters) have to be determined using a specific load case that the system is trying to match. For instance, separate controllers are used in a single HBM depending on whether the user is trying to simulate response in frontal, lateral, or rear impact cases. The closed-loop PID models also have been tuned under a particular loading environment and any major change in the external force field will require the retuning of the model. Tuning is not a straightforward process, so the PIDs for muscles around a joint are often grouped into extensors and flexors to reduce the complexity associated with tuning multiple parameters at once. So, not only are the PID parameters tuned to specific cases, but the muscle recruitment scheme is also pre-defined to facilitate this process. This simplification renders the closed-loop control effective under a narrow band of loads and may not extend to novel loading scenarios. Optimizing the grouped controller parameters for adaptation to a novel load set will be challenging as the activations of individual muscles may need to be adjusted.

Reinforcement learning algorithms can overcome some of the challenges commonly associated with closed-loop controls. The learning algorithms have been used in control problems with similar

27

dimensionality, and the coupling of RL algorithms with neural net frameworks make it suitable for approximating non-linear behavior observed in human motor control mechanisms.

1.5 Reinforcement learning for control

Reinforcement learning (RL) is a branch of machine learning in which the controller learns the control sequence for a system by interacting with the control system and exploring the action space over multiple interactions while receiving constant feedback. With multiple iterations, the controller develops an optimal control policy to achieve the desired objective, without being explicitly programmed to do so. RL algorithms are biologically inspired and are analogous to the process of how humans perfect a motor task through repeated practice and constant evaluation of the task which results in incremental improvement and learning over time.

While RL algorithms are not limited to control problems, the ability to evaluate and correlate controller output to the system response makes them suited for control applications. A reinforcement learning control problem can be expressed in terms of agent, environment, state, action, reward, and q-value (Sutton and Barto, 2015):

Agent – Agent refers to the RL algorithm used for a specific problem. The agent reads the environment parameters, outputs the control signal also called action, and while training receives feedback on the appropriateness of the action.

Actions (a) – Action is the control signal that the agent outputs for a given state of the control system. Depending on the nature of the system, the action space can be discrete or continuous.

Environment – The control system with which the agent interacts. The agent reads the state from the environment and based on the state performs actions.

State (s) – The state refers to the set of parameters that are read by the agent while interacting with the environment. Depending on a given state, the agent outputs the actions to the environment, and the environment returns the new state (s') as input to the agent.

Reward (r)– Reward is the feedback received by the agent after each action which quantifies the relative goodness of the action for a given state. The reward is generally calculated from a reward function that relates the different states (s') and actions with the desired objective.

Policy (π) – The control strategy that an agent formulates after multiple iterations which maps the actions to the state.

Q-value (Q) – Q-value quantifies the state-action pair based on the reward function. Q value for a state-action pair, Q(s, a) is the expected cumulative reward that an agent receives during the episode when it takes an action 'a' starting from current state 's'.

RL agents are functions that read the state (s_t) and output a specific action based on the policy, whereas the environment on receiving the action changes the current state s_t to the next state s_{t+1} and outputs the reward r_{t+1}. The training phase of RL algorithms are iterative in nature during which the agent is swept through multiple state-action pairs over numerous iterations or episodes calculating the reward and updating the policy. Over multiple iterations, the agent learns to output actions at each timestep that maximizes the cumulative reward for the iteration. This process of simulating the agent under different states and actions until it can determine the best actions for a

given state is called training. A trained agent can output the sequence of actions based on the training, and do not require the reward function once it has been trained.
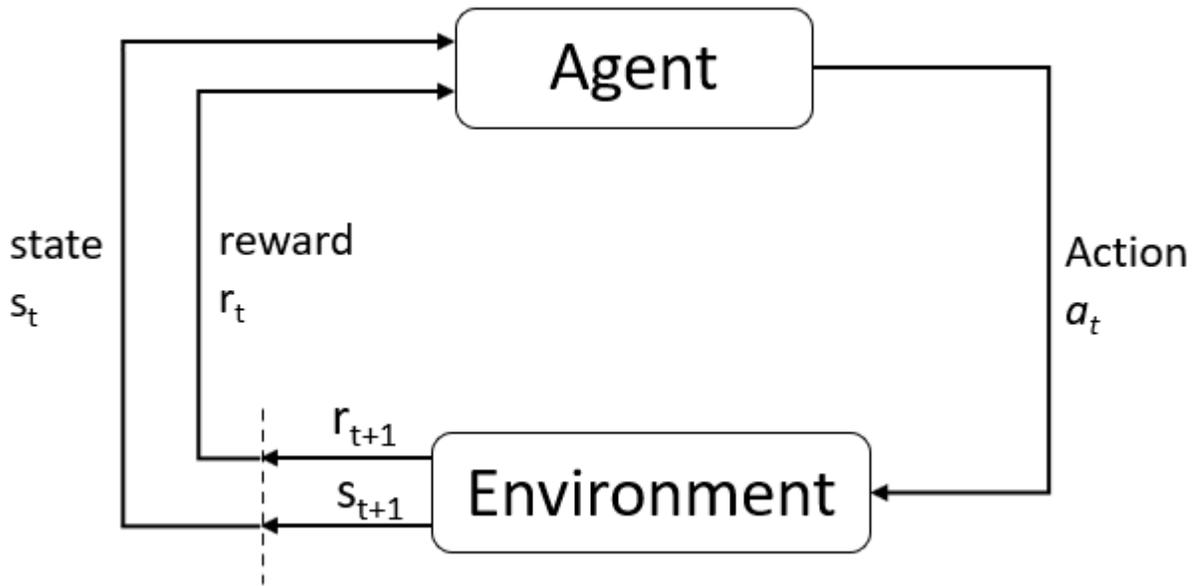
Figure 1-5: Reinforcement learning training cycle.

A standard RL algorithm is shown in Figure 1-5. The subscripts t and t+1 refer to the current timestep and the immediate next timestep. RL evaluates the actions for a given state (S_t), and the goal of the algorithms is to learn the best sequence of actions (a_t, a_{t+1}, …) to achieve the desired control objective which is defined by the reward function (Kakade and Langford, 2002; Sutton and Barto, 2015).

Unlike supervised learning algorithms which aim to learn the behavior of a system, the focus of reinforcement learning is to make sequential decisions for a target-oriented system. In supervised learning, the model learns from a labeled dataset by minimizing the error between the model output and the actual outcome, whereas in RL, the agent interacts with the environment over and over and learns from experience using an exploration and exploitation approach.

1.5.1 Deep reinforcement learning

Reinforcement learning when paired to neural networks with one or more hidden layers to improve the efficiency of the learning process is referred to as deep reinforcement learning. Neural networks (NN) are function approximations that are designed to map inputs to labeled outputs. Both the network inputs and the predicted outputs are in the form of numerical tensors.

Neural networks are broadly used for regression and classification tasks (LeCun et al., 1989; Werbos and John, 1974). Compared to the linear learning algorithms where the inputs are processed once, neural networks may consist of multiple layers to process the data before generating the outputs. The layers between the input layer and output layer are called hidden layers, and the input of one layer is the output of the previous layer. The architecture of the networks primarily depends on the objective. The simplest network called a feedforward network has sequential input, hidden, and output layers and processes information in one direction (Figure 1-6). Output from each layer may further be processed before feeding it to the following layer to improve the computational efficiency during the training (Glorot et al., 2011). Depending on the size of the input and output tensors, multiple hidden layers can be added.
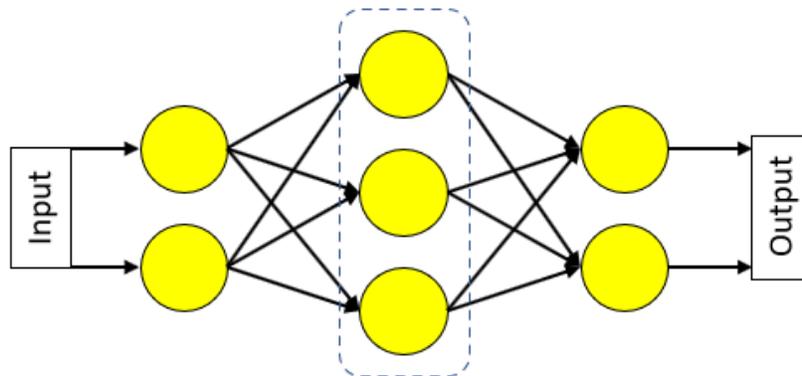
Figure 1-6: Feedforward network with one hidden layer.

Convolutional Neural Networks (CNNs) are mostly used in image classification and pattern recognition tasks (LeCun et al., 1989) due to improved performance with visual imagery data. CNN can be fed image or audio data in form of numerical tensors, and the network architecture is designed to extract features from the input to learn valuable information. Recurrent neural networks (RNN) have a cyclic structure that is developed for time series data (Hochreiter and Schmidhuber, 1997; Werbos and John, 1974). The architecture of the layer is designed in a way that the output of a layer is saved and fed back as input which makes it suitable for predicting temporal patterns. Deep learning with various NN architectures has enhanced state-of-the-art technologies in fields such as object recognition, natural language processing, genomics, and drug discovery (LeCun et al., 2015).

Neural Networks are coupled with RL algorithms when the control environment is complex in nature. Neural nets are particularly useful in RL problems with a large number of states and actions (Bengio et al., 2013). Rather than cycling through all the state and action pairs like in conventional RL algorithms, the NNs can sample a subset of state and action pairs to create functional approximations of either the Q-value or the policy. Mnih et al. (2015) introduced a deep Q-network (DQN) model which achieved performances comparable to professional players while playing 49 games in a classic Atari 2600 environment. The DQN used deep convolution layers to extract raw pixel information which conveyed the game situation based on which the action (game control) was determined by the agent.

Actor-critic methods are advancements in deep RL algorithms that can approximate Q-values and policy independently of each other (Peters et al., 2005; Sutton and Barto, 2015). The actor-critic algorithm utilizes two neural networks – the actor-network and the critic network. The actor network approximates the policy mapping states to the action whereas the critic network evaluates

the state-action pair by computing the Q-value. The critic uses the current estimated Q-value to guide the actor into giving better actions while learning how the Q-value varies with the state-action combinations.

Deep RL with further advancements in machine learning and computation has been able to solve a wide range of complex and decision-making problems. Silver et al., 2015 developed a novel algorithm based on actor-critic methods by partially differentiating the Q-values concerning actions called deterministic policy gradient (DPG). DRL was found to perform well for high dimensional control problems with continuous action space. Actor-critic networks combined with advanced tree search algorithms were used to develop the AlphaGo program which defeated professional players in the game of Go (Silver et al., 2016). AlphaGo was further modified to enable the RL algorithm to learn to play Go from scratch without any human interventions (Silver et al., 2017). Recently, a deep RL algorithm has been used to develop AlphaTensor, which discovered a matrix multiplication algorithm that improves upon the current state-of-the-art algorithms (Fawzi et al., 2022).

The ability of deep RL algorithms to simulate physics-based control tasks (Lillicrap et al., 2019; Mnih et al., 2015; Silver et al., 2015) makes them attractive for human motor control studies. Deep RL can also scale to high-dimensional state and action spaces requiring continuous control which is suited to deal with the inherent redundancy present in the musculoskeletal system.

1.6 Reinforcement learning architecture

An advantage of RL over closed-loop control is that the algorithms are not limited to the few error signals as inputs. This property becomes more important for musculoskeletal systems as several kinematics (translation, rotation, velocity etc.) as well as dynamic (joint force, torque, muscle forces, fatigue etc.) need to be considered for activating the muscles. Gradient based RL can be

appropriate for such kind of control inputs as the agents can handle multiple state parameters simultaneously, which can be a combination of body kinematics, joint responses, and muscle behaviors. While motor control studies involving closed-loop controllers assume muscle grouping as extensors and flexors to reduce the requirement of tuning the gains for multiple controllers, deep RL can generate multiple control outputs by reading the states and thus, do not require any grouping or assumptions for generating the muscle actuations as outputs. The output generated by the deep RL agent depends on the states and the reward function used for training.

During each training iteration, the RL agent reads the states and generates outputs based on the current policy. The objective of the training is for the agent to learn a policy that maximizes the cumulative reward of the system. The reward function can be defined as a function of parameters that can be measured from the HBM and the different components of the reward functions can be arranged such that control objective can be defined in terms of maximizing the reward function. Heess et al. (2017) showed that RL agents can learn complex behavior with a simple reward function, by constantly interacting with a rich control environment. Apart from the reward function, the RL agents require no other user defined input on how the system should behave and can generate the control outputs without any assumptions. The weights and biases of the actor network are adjusted during the training to map the HBM measurements (both kinematics and dynamics) to the control actions. Additionally, the gradient based RL algorithms can output continuous actions and have been found to be stable for large continuous input space (Silver et al., 2015, Lilicrap et al., 2016). As such, utilization of the developments in gradient based deep RL framework is ideal for muscle control studies as the framework can simultaneously handle multiple HBM parameters as inputs and does not require information of muscle coordination apriori.

1.7 Reinforcement learning for motor control in humans

Davoodi and Andrews (1998) implemented an RL control with a sagittal plane segmental model of the human body with the goal of arm-assisted standing. The agent seeks to control shoulder forces for standing while trying to limit the upper limb forces and knee joint velocity. Micera et al. (1999) developed a simplified elbow model with antagonistic muscle pair to simulate the extension-flexion motion. Izawa et al. (2008) developed a two-joint arm model with six muscles to generate arm-reaching motion using RL. Thomas et al. (2008) used actor-critic networks for control of functional electrical stimulation (FES) of the human arm under varying system dynamics. Iwamoto et al. (2012) used RL with an FE model of 50^{th}-percentile male to control head-neck motion in rear and frontal impacts. The agent was trained to control the head motion in the sagittal plane, and the activations produced by the trained agent were transferred to the FE model while simulating impact conditions. The active model had a better correlation to the volunteer data than the passive simulations, however, the agent overestimated the muscle forces making the head-neck region stiffer after the initial phase of the simulations. Heess et al. (2017) synthesized complex locomotion behavior in simplified torque-actuated bodies stimulating $6-21$ joints. In the study, it was shown that along with a suitable reward function, a rich environment can improve the learning behavior of an agent.

In the Neural Information Processing Systems (NIPS) 2017 conference, a 'Learning to Run Challenge' was conducted in which participants were provided a musculoskeletal multibody model of the human body (Delp et al., 2007) with 18 lower extremity muscles (Kidziński et al., 2018). The participants were asked to develop an RL-based controller to actuate the leg muscles and make the human model move forward. A total of nine teams were able to implement a controller that enabled the model to move 15 m in 10s. The winner of the competition, Jaśkowski et al. (2018)

35

used proximal policy optimization (PPO) (Schulman et al., 2017) to design the controller. In the following NIPS conferences (2018, 2019), contests were held to develop controllers for amputated leg musculoskeletal models and models to meet target velocity vectors respectively (Song et al., 2021). Various innovative solutions were developed for RL agent architectures and reward functions to meet the challenges respectively (Kolesnikov and Khrulkov, 2020; Zhou et al., 2019). RL has also been used to develop enhanced assistive devices and exoskeletons to aid in rehabilitation and lead to improvement in locomotion (Luo et al., 2021; Tu et al., 2021).

With the development of RL algorithms with higher efficiency and scalability, there is an increasing amount of work that has been done to combine RL and musculoskeletal modeling, however, some gaps remain in the field which needs to be addressed. In some locomotion studies, RL agents have been shown to reach the desired final position, but the physiology of the way the models moved was not accurate or biofidelic (Akimov, 2020; Kolesnikov and Khrulkov, 2020). The agent that Jaśkowski et al. (2018) trained for the 2017 Learning to Run Challenge used discrete action values of 0 or 1, also called the bang-bang strategy. Although the agent produced the intended kinematics and joint mechanics, the action space may not be feasible and lead to high energy expenditure and fatigue (Koelewijn et al., 2019; Umberger et al., 2003).

Also, most of the previous work using RL for muscle control has been limited to arm reaching or locomotion tasks, and the ability of the RL controller to synthesize response in a more chaotic environment like those experienced in sports collisions or motor vehicle crashes has not been evaluated in detail. The challenge in simulating active motor behavior lies in developing a control mechanism that can generate acceptable responses under a range of external loads. It also remains to be studied whether an agent trained for kinematics and goal-directed movement tasks can adapt

to dynamic forces. This is especially necessary as generating volunteer data for validation under such dynamic forces is not possible due to the injurious nature of the loads.

To summarize, the advancements in RL have not been utilized to their potential in the musculoskeletal modeling of active muscle behavior. The dissertation seeks to develop a comprehensive control framework for use of RL in active muscle control. The dissertation will utilize a detailed study on the design of RL based controller with a series of computations using detailed musculoskeletal models, volunteer testing, and validation.

1.8 Scope of research

The dissertation focuses on evaluating the use of RL algorithms for muscle control in HBMs. The simulations presented in the dissertation aim to replicate the most common scenarios generally experienced by humans. The research performed takes a step-by-step approach to determine the feasibility of RL based framework for general-use muscle controllers.

1.8.1 Research objectives and specific aims

The central objective of the dissertation is to employ gradient-based RL algorithms with deep RL to develop a muscle control framework that can synthesize physiologically accurate voluntary kinematics and respond to external perturbations. The dissertation is organized into four major tasks for incremental development of the competency, know-how, and data required to accomplish the underlying objectives.

The first specific aim of the dissertation is aimed at a feasibility study to gauge the applicability of RL algorithms for the control of joint kinematics. In specific aim 1, the framework for reinforcement learned muscle actuation controller (RLMAC) is formulated for the control of

motion around the elbow. The RLMAC is tested for its capability to adapt to changes in external forces and maintain stability under dynamic loading scenarios.

In the second specific aim, the RLMAC is coupled with a multibody model of a head-neck complex to control the head motion. The RLMAC is trained to maintain the head stability under a gravitational field and control the head position in the sagittal plane. The various parameters for the development of the reward function will also be evaluated.

The third specific aim involves performing volunteer studies to generate voluntary kinematics data on head extension, flexion, and axial rotation. The volunteer data will be used to calibrate the components of the reward function in the sagittal model. The head-neck control model will also be extended to all degrees of freedom to control the head kinematics in coronal and transverse planes as well.

In the fourth specific aim, the RLMAC will be evaluated for its range of applicability. The RLMAC trained for generating goal-directed kinematics will be assessed in its ability to synthesize complex head movements along different axes. The head-neck model will also be simulated under added inertia, low-severity automotive cases, and impact loads. The RLMAC will be tested for different anatomies by scaling the originally developed head-neck model. Specific aim 4 will illustrate the feasibility of the control framework in generating a generalized response to a wide range of intrinsic and external conditions. The future scope of the current methodology will also be discussed in the final chapter.

Specific Aim 1: *Demonstrate that RL can be used as the basis for a muscle control scheme that could simulate active anatomical kinematics.*

Specific Aim 2: *Demonstrate the ability of an RL framework for maintaining the desired head and neck posture under gravity.*

Specific Aim 3: *Validate the RL muscle control scheme using human voluntary kinematics data.*

Specific Aim 4: *Quantify the range of applicability of the RL muscle control model for use in head and neck kinematics.*

1.8.2 Dissertation overview

The specific aims mentioned in section 1.7.1 are organized into eight chapters in the dissertation. Chapter 1 of the dissertation provides a background and basic literature review of research in the field of musculoskeletal modeling and reinforcement learning. A more detailed literature review will be done in the following chapters which are relevant to the specific aims. Chapter 2 provides information on the development of the RLMAC and its integration with the multibody model of the arm. Details on the training of the RLMAC to simulate desired arm kinematics are also provided in chapter 2. Chapter 3 gives details on the development and validation of the head-neck multibody (MB) model which would be used for muscle control in later chapters. Chapter 4 will focus on the integration of RL-MAC with the head-neck multibody model to control kinematics in the sagittal plane. Components of the reward function will be evaluated relative to the goals of muscle control. Chapter 5 will provide information on the volunteer study and the protocols involved. Steps undertaken to measure the goal-directed head kinematics has been described in detail. Validation of the responses of the RLMAC trained in chapter 4 will also be performed and documented in this chapter. Chapter 6 focuses on extending the RLMAC controller for control in

all DoF for the head-neck model. Details on the training and validation of the RL agent for omnidirectional head kinematics control are provided in chapter 6. Chapter 7 focuses on evaluating the applicability of the agents trained for multi-DoF control for different anthropometries. Chapter 8 analyses the ability of the trained controller to adapt to changes in inertias and impact loads. The different criteria on which the agents will be simulated together with the details of the simulation results will also be documented.

All the multibody models that are mentioned in the dissertation are developed in MATLAB v20b Simscape Multibody toolbox. The MB models are integrated with MATLAB reinforcement learning toolbox v20b to develop and train the RLMAC. The RLMAC is trained in the parallel cores of UVa HPC cluster "Rivanna".

1.8.3 Contributions

The dissertation proposes a novel framework to develop and tune a muscle controller for HBMs using the latest developments in reinforcement learning. The control framework developed in the dissertation addresses some of the issues that are experienced in traditional closed-loop muscle controllers. The major contributions of the dissertation are summarized below.

- The work done in the dissertation shows that RLMAC can be trained for generating biofidelic joint motions without any prior assumptions on grouping or organization of muscles.

- Development of a computationally efficient multibody model of the head and neck region that can be used in control and optimization studies that are iterative in nature.

- The present study develops a deep RL control framework to generate goal-directed rotations of a human head-neck multibody model along the three anatomical planes. To the

best of my knowledge, no previous study has simulated targeted head kinematics in all three directions.

- The dissertation also generates volunteer kinematics data of goal-directed head rotations for validation of computational models.

The RLMAC developed in the dissertation shows the capability of RL for application in human motor control. With constant improvements in the field of machine learning and reinforcement learning, the efficiency and accuracy of the muscle controller based on RL will further improve. The work done in the dissertation is also an advancement towards the development of full-body active motor control which can generate accurate responses under loads associated with HBMs.

Chapter 2 – Reinforcement learning as a means of muscle control

In this chapter, a primary evaluation of the use of reinforcement learning (RL) algorithms for muscle control is performed. The ability of RL algorithms to synthesize biofidelic joint movements has been demonstrated using a multibody (MB) human arm model. The MB model of the human arm provides a simplified environment for the development and integration of the RL control framework. The RL muscle activation control (RLMAC) developed in this chapter is used to control the movement of the lower arm in extension and flexion, with other directions being constrained for simplicity. For the extension-flexion motion of the arm, two different muscle activation schemes were studied – the muscles implemented in the arm MB model were recruited as antagonistic groups or individually.

The RLMAC integrated with the arm MB model was trained to perform the extension-flexion motion of the arm. Simulations with the trained RLMAC show that the arm can perform the desired kinematics, even in presence of an external force field. The arm kinematics along with the muscle activations developed for the different ranges of extension flexion motions were measured and documented in this chapter. The response of the arm was compared with a previous experimental dataset to evaluate the biofidelity of the arm movements.

The trained controller was also subjected to pulse loads representing automotive collisions to evaluate the range of applicability of the RLMAC. The RLMAC, even though trained under constant loads, was able to maintain arm stability under dynamic loads implying the robustness of the motor control approach. Results from the present chapter provide an insight into the use of reinforcement learning for motor control applications and the adaptation of the control framework to more complex body regions.

42

2.1 Introduction

Performing a motor task involves coordinated force generation by the muscles which are activated by the central nervous system (CNS). The muscle forces maintain the stability of joints under external forces or synthesize the desired motions of the joints. Similar actuation of muscles in human body models (HBMs) is challenging due to the inherent complexity associated with muscle coordination or synergy as seen in humans. The present chapter is an attempt to apply reinforcement learning (RL) algorithms to the motor control problem. For this purpose, an MB model of the human arm was created with the anatomy of a 50[th] percentile male and integrated with the RL control framework with the intention of generating extension-flexion motion about the elbow joint.

The upper extremities with the elbow joint have been the subject of many previous control studies due to the simple nature of the joint. Kistemaker et al. (2006) developed a control strategy based on PID control and implemented it into a musculoskeletal model of the human arm. The elbow joint was modeled as a revolute joint and was actuated by muscles represented as four lumped Hill-type muscles. The muscles were grouped as extensors and flexors and actuated by error signals based on the difference between the current and target length of the contractile element of muscles. Budziszewski et al. (2008) implemented a PID controller in the MB model of the arm to control the extension-flexion and pronation-supination motion. The PID controllers operated with joint angles as feedback signals. The muscles were divided into agonist and antagonist groups and the same activations were developed for muscles belonging to the same group. Östh et al. (2012) incorporated the feedback-based control loop into a finite element (FE) model of a human arm. The FE model included nine individual muscles, and the muscles were grouped as extensors and flexors. The PID controller was driven by the elbow angle error, which activated the extensors and

43

flexors for arm-reaching movements. Martynenko et al. (2019) used a similar approach while modeling the activations of the arm for goal-directed motion. In this study, the PID controllers actuated the muscles based on the error of muscle length. A pre-simulation with the passive arm was carried out to obtain the muscle lengths at the initial and target positions which were then used to actuate the muscles as groups. Bayer et al. (2017) developed an arm musculoskeletal model with four-element hill-type muscles (Haeufle et al., 2014). In this study, a PID controller was used to activate the muscles, and the effects of the parameters of the muscle model and the activation functions on the kinematic responses of the arm model were evaluated.

Although the PID controllers used in previous studies were successful in synthesizing the required movements, some simplifications and assumptions were made while developing the control architectures. Muscles were grouped into agonist-antagonist pairs and identical activations were applied to each muscle group. Activating the muscles individually will require tuning multiple controllers at once which can be difficult. The controllers were also precisely tuned for a limited range of loads and the gains may need to be modified for a change in the loading environments (Happee et al., 2015, Zheng et al., 2021). Reinforcement learning can provide an alternative as it has been found to output multiple control actions simultaneously which can be both continuous (Heess et al., 2017; Silver et al., 2015) and discrete (Mnih et al., 2013) in nature.

Iwamoto and Kato (2021) used actor-critic RL agents to stabilize a multibody model of the human arm under gravity. Each muscle was stimulated individually using two reward functions – one based on joint angle and the other based on muscle length. Min et al. (2018) integrated an RL agent into a simplified finite element (FE) model of the human arm. The agent was trained to maintain joint equilibrium while carrying a weight of mass 1 kg. The reward function was developed using the position difference of the end effector and the agent was also penalized for high muscle

activities. Jagodnik et al. (2016) used actor-critic RL algorithms to train planar arm reaching movements for development of functional electrical stimulation (FES) system. The study demonstrated that sparse rewards at the final timestep can be useful for training the control system.

Although the RL-based control mechanisms have been used in the past, there have been some gaps in the research that this chapter and the dissertation in general seek to address. In the previous arm reaching tasks that have been modeled using RL, the human models were simulated under a similar environment for which the system was trained (Crowder et al., 2021; Driess et al., 2018; Fischer et al., 2021; Jagodnik et al., 2016). How the overall biomechanical response is affected by adding variability to the system, such as additional inertia or changing the external environment parameters, have not been studied. This is an important aspect for the application of RL in human body models designed for scenarios that people are not commonly exposed to, such as the event leading to an automotive crash. In computational control studies, such scenarios need to be defined in terms of the loading environment and control objectives to evaluate the utility of the controller. Human beings can adapt to changes in external loads (Happee, 1993; Smeets et al., 1990), and the ability of RL agents to replicate such adaptive response, during reaching tasks in this case, has also not been evaluated. Furthermore, the previous RL musculoskeletal control studies have been limited to kinematic tasks, and the ability of RL control mechanisms to extend to dynamic events such as automotive or sports environments, where the response time is relatively fast has also not been studied. Finally, the biofidelity of muscle activation patterns corresponding to a given joint movement, and the effect of change in the external environment on the muscle synergy also need to be investigated.

The present chapter aims to develop a RL-based control framework and integrate it into a musculoskeletal model of a human arm to evaluate the range of applicability of RL algorithms in

muscle control problems and also study the possible applications of the RL control methodology in more complex body regions.

2.2 Methodology

A simplified MB model of the arm was created with the geometry of the scapula, humerus, radius, and ulna. Muscles responsible for the extension and flexion motion were also included in the model. The MB model was integrated with an RL agent to develop the RLMAC which was then trained to generate arm kinematics.

2.2.1 Development of the arm multi-body model

The upper extremity bones were represented as rigid bodies with radius and ulna combined as one body (Figure 2-1a). Mass and inertia of the lower arm were applied at the radius-ulna body. As no motion at the shoulder is being considered in the present chapter, the glenohumeral joint was constrained effectively constraining the humerus in the multibody model. The elbow joint between the humerus and radius was represented as a revolute joint with a stiffness of 0.6 Nm/rad (Martynenko et al., 2019; Östh et al., 2012b) and damping of 0.4 Nms/rad (Cannon and Zahalak, 1982; Rack, 2011). The rotation space of the elbow revolute joint was limited between 0^o (full extension) and 160^o (full flexion).

(a)   *(b)*
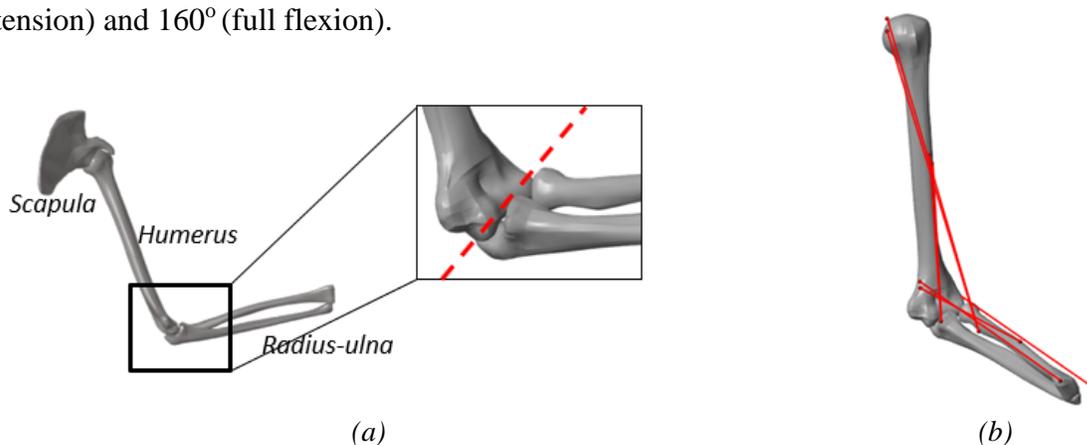
Figure 2-1: MB model of the human arm (a) Rigid bones with elbow joint (b) MB model with muscles.

Muscles were incorporated into the arm MB model were represented as tensile force elements acting between suitable origin and insertion points (Figure 2-1a) (Lieber et al., 1992; Moore and Dalley, 2009). The force magnitudes acting between the insertion points were calculated from simplified Hill-type muscle with an active contractile element and a passive element in parallel (Figure 2-2).
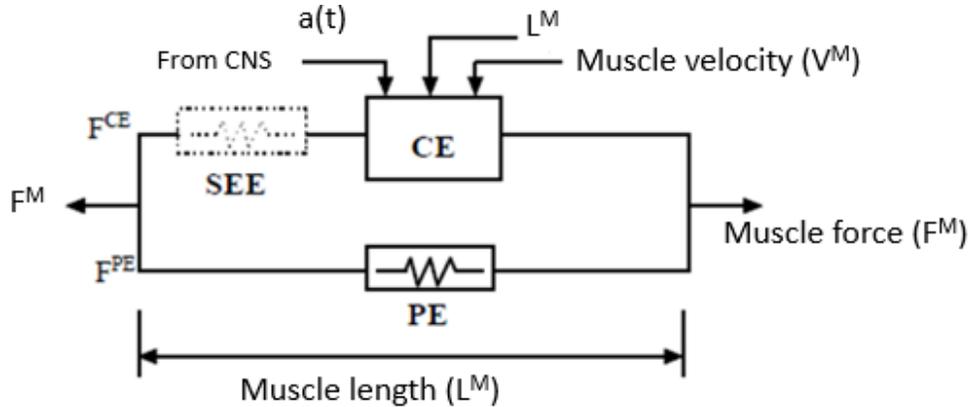
Figure 2-2: Hill-type muscle model with contractile element and passive element (Panzer, 2006).

The active forces (F_{CE}) are non-linear in nature and depend on the length and velocity of the muscles (Hill and Sec, 1938; Zajac, 1989). For this chapter, the normalized relationship between the force-length (F_l) and force-velocity (F_v) has been approximated as curve inputs and is considered the same for all the muscles (Panzer et al., 2011). The passive muscle force (F_{PE}) was modeled as an exponential function of length (Equation 2-1), and only start acting when the muscles are stretched beyond the optimum length (L_{opt}) (Winters, 1995). K_{sh} in Equation 2-1 is a dimensionless parameter influencing the rise of the passive force with length.

$$F_{PE} = \frac{1}{\exp(K_{sh}) - 1} \left\{ \exp\left[\frac{K_{sh}}{L_{max}} (L - 1) \right] - 1 \right\} \ for \ L > 1 \qquad \text{Equation 2-1}$$

The muscles in the MB arm model cause extension and flexion motion at the elbow joint. The muscles included in the MB model along with the muscle properties are listed in Table 2-1.

Table 2-1: Properties of muscles in the MB model

Muscles	F_{max} (N)	No. of strands in model	L_{opt} (mm)
Biceps brachii long head	360	1	336
Biceps brachii short head	248	1	327
Brachialis	568	2	166, 156
Brachioradialis	152	1	283
Pronator teres	320	1	157
Extensor Carpi Rad Longus	176	1	312
Triceps long head	456	1	324
Triceps lateral head	360	1	296
Triceps medial head	360	3	239, 211, 170

The optimum length (L_{opt}) of the muscles is obtained at the neutral position (90º flexion angle) (Hayes and Hatze, 1977)

The attachment points of the muscles to the bones have been approximated from different sources in the literature (Lieber et al., 1992; Moore et al., 2009). The muscles having large cross-section areas and long attachment points were split into multiple strands capturing the width of the muscles to properly distribute the muscle forces (Table 2-1).

2.2.2 Muscle control framework

For a movement around a joint, the central nervous system (CNS) actuates a set pattern of muscles (muscle synergy). In this chapter, the feasibility of using RLMAC for muscle control under a range of external environments is explored. Figure 2-3 shows the RL-based framework for muscle control.

Deep deterministic policy gradient (DDPG) is used to develop the controller for the RLMAC framework in the present chapter (Lillicrap et al., 2019). DDPG belongs to the class of gradient-based actor-critic agents (Silver et al., 2015). DDPG consists of an actor-network and a critic network and is suitable for muscle control because of its ability to output continuous actions (Lillicrap et al., 2019; Silver et al., 2015).
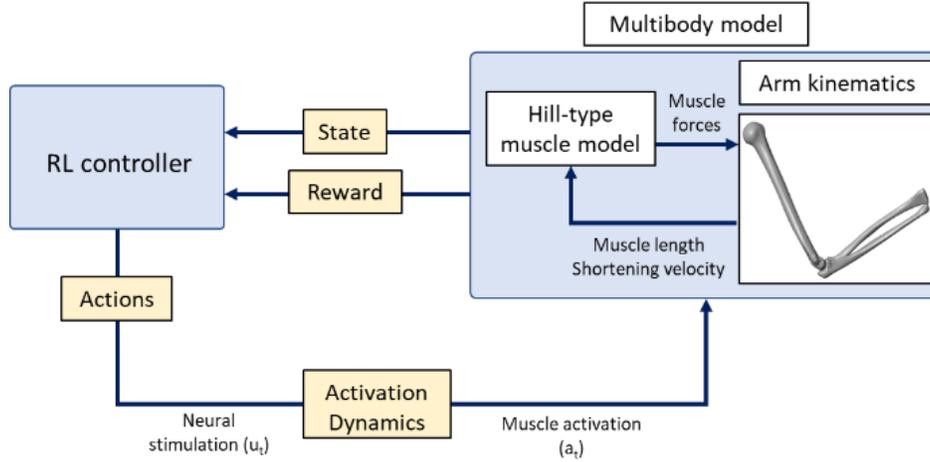
Figure 2-3: RLMAC framework for the arm MB model.

The RL controller (Figure 2-3) reads the state from the MB model of the upper extremity and outputs neural stimulations (u_t) that get converted into muscle activations (a_t) for input to the Hill's muscles in the MB model. The muscle length and velocity at a particular instance of the simulation are obtained by measuring the Euclidian distance between the insertion points. The forces generated by the muscles rotate the radius-ulna towards the target position. Based on how close the lower arm is moved toward the target angle, the RLMAC receives a scalar reward signal.

$$\frac{da}{dt} = \frac{1}{\tau_{act}}[u - (1 - \delta)au - \delta a] \qquad \text{Equation 2-2}$$

The activation build-up in the muscles was calculated using the activation dynamics relationship proposed by Zajac et al. (1989) (Equation 2-2). The τ_{act} is the activation time constant which relates to the increase in muscle activities as a result of u_t output by the RL agent. $\boldsymbol{\delta}$ is the ratio τ_{act}/τ_{deact}, where τ_{deact} is the depletion time constant controlling the drop in activation level (a_t) after a decrease in u_t output by the agent (Table 2-2).

Table 2-2: Parameters of the activation function

Parameters	Values (Bayer et al., 2017; Zajac, 1989)
τ_{act}	0.02
τ_{deact}	0.06
Minimum activation (a_o)	0.005

The architecture of the DDPG agent is shown in Figure 2-4. The DDPG agent incorporated in the current chapter is similar to that originally proposed by Lilicrap et al. (2016). The actor-network was composed of a feedforward network with one hidden layer. The inputs to the hidden layer and the final layer were activated with a rectified linear (ReLu) function (fx = max(0,x)) (Hara et al., 2015). The output of the actor-network was activated using a sigmoid function to limit the action space between 0 and 1 (Han and Moraga, 1995). The input layer of the actor-network had nodes equaling the number of states whereas the final layer nodes equal the number of muscles required to be activated. The critic layer is comprised of a feedforward network with one hidden layer. The output of the hidden layer was activated by the ReLu function before passing it on to the final layer. The state parameters were inputs to the first layer and were processed with the ReLU function before connecting to the hidden layer. The actions from the actor-network were passed on directly to the hidden layer without being activated *(Lilicrap et al., 2016)*. The critic layer outputs a scalar Q-value for the corresponding state action pair and thus has a single node at the final layer. The states can consist of any parameters that can be measured in the control environment and the values are passed on to the RL agent. The reward function is usually defined as a function of the environment parameters. Since the RL algorithms aim to maximize the rewards, the various parameters in the reward function can be calibrated to define the control objective.
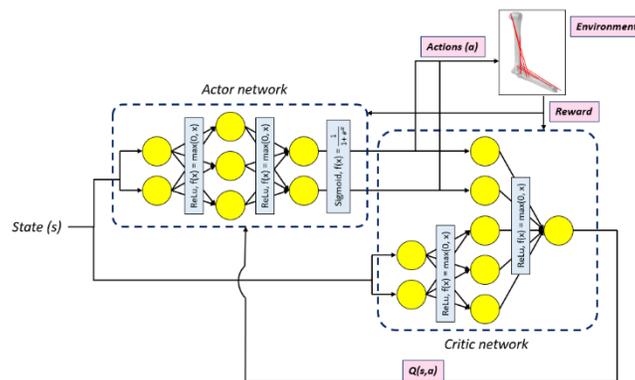
Figure 2-4: Architecture of DDPG agent used in the current chapter.

The DDPG agent is an actor-critic reinforcement learning agent that searches for an optimal state-action value (policy) that can maximize the expected long-term reward for a training iteration. The DDPG agent updates the actor and critic network at each time-step during the training. At the start of the training, the actor and critic networks are initialized with random weights and biases. For a state (s_t) that the agent observes, the actor network takes an action (a) according to the policy (initially randomized). Sometimes, a noise is also added to the action to enable the agent to explore the possible action space better. After the action, the agent receives a reward (r) according to the reward function, and the current state changes to the next state (s_{t+1}). The DDPG agent keeps a track of the sequence (s_t, a, r, s_{t+1}) and this set is known as experiences. The Q-value is calculated as the current reward (r), and the expected future reward from the state s_{t+1}. The gradient of the Q-value with respect to the action is used to update the actor network. The critic network is updated by minimizing the difference between the calculated Q-value and target Q-value (Lilicrap et al., 2016). At the culmination of the training, the weights and biases of the actor network are finalized, which is then used to output the actions in use-case scenarios, without requirement of the critic network or the reward function.

The control objective in this chapter was to move the lower arm (radius-ulna body) from an initial given position to a target position and stabilize the arm at that position. The states of the RLMAC were defined by the kinematics parameters – the current elbow angle, the angular error (difference between the target angle to the current angle), the elbow rotational velocity, and muscle activations from the previous timestep. The reward function was developed to minimize the angular error during the training (Equation 2-3). The agent was penalized proportional to the magnitude of the error and rewarded if it managed to maintain the angular error within 0.1 radians ($<5.7^{\circ}$). Due to the redundancy of the musculoskeletal system, there are many combinations of muscle activations

that may lead to similar arm kinematics. To overcome the redundancy problem, minimizing muscle activation (Chancey et al., 2003; Pedotti et al., 1978) is considered in this chapter, however, other costs of motion criteria are also used in biomechanics research (Koelewijn et al., 2019).

$$\text{Reward} = -\alpha \, \text{Error} - \beta \sum a(t) + \gamma \, (|\text{Error}| < 0.1 \text{ rad})$$
<div align="right">Equation 2-3</div>

The reward is calculated using Equation 2-3 at each timestep of the simulation. During the training, the agent tries to maximize the cumulative reward over the simulation time (Sutton and Barto, 2015). α, β, and γ are constant scalar terms that scale the contributions of the components of the reward function corresponding to the overall control objective.

2.2.3 Training and validation

The passive response of the arm MB model was evaluated before training. The moment was applied at the elbow joint with a magnitude varying between 1 and -1 Nm. The humerus was constrained, and the resulting rotation of the radius-ulna body was measured to obtain the passive stiffness of the joint which was also validated with published data.

The training of the arm MB model integrated with the RLMAC framework was carried out with the reward function mentioned in Equation 2-3. Noise was added with Ornstein-Uhlenbeck (OU) process with a standard deviation of 0.09 for adequate exploration of the action space (Lillicrap et al., 2019). The RLMAC was trained for desired arm kinematics in two different scenarios – bare arm motion control and arm motion control under an external force field (Table 2-3).

(a)   *(b)*   *(c)*
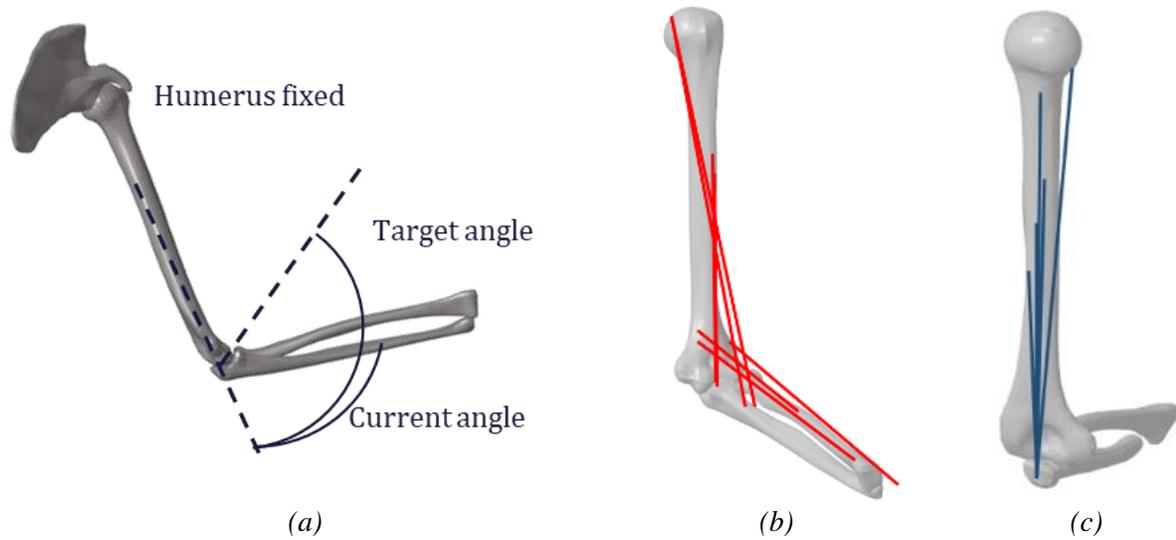
Figure 2-5: (a) Arm parameters for training (b) Flexor muscle group (c) Extensor muscle group

Training scenario 1 - Bare arm motion control

The MB arm model integrated with the RL control framework (Figure 2-3) was trained to carry out fast-goal directed rotations of the radius-ulna about the elbow joint, in absence of any added mass or inertial loads to the lower arm. For each iteration of the training, the lower arm was positioned randomly between 0^o (extension limit) and 160^o (flexion limit) at the start of the simulation (T = 0 s) (Figure 2-5a). The target position was also randomly varied between 0^o-160^o to prevent overfitting to any particular motion.

Two different muscle activation schemes were analyzed during the targeted rotations of the arm. In the first activation strategy, called Group Activated Muscle Recruitment (GAMR), the muscles in the MB arm model were split into extensors and flexors, and the muscles of the same group were activated identically. Biceps brachii long head and short head, the brachialis, the brachioradialis, the pronator teres, and the extensor carpi were grouped as flexors (Figure 2-5b) and the triceps long head, lateral head, and medial head were grouped as extensors (Figure 2-5c). For this muscle activation scheme, the RLMAC outputs two actions corresponding to the stimulation of each group. The second activation scheme was named Individual Activated Muscle

53

Recruitment (IAMR) where each muscle was actuated independent of the activations of other muscles. In IAMR, no assumptions were made as to how muscle will coordinate and respond to a target, and the RLMAC generated nine actions for stimulation of each muscle.

The response of the trained RLMAC in both activation schemes was compared with human volunteer data gathered from 6 subjects (Kistemaker et al., 2006). The simplified nature of the test setup where the volunteers were asked to perform fast goal directed rotations of the arm in the horizontal plane which minimized the effect of gravity on the arm motions and made it suitable to replicate the arm motions in the MB model. The arm MB model was initially stabilized in the initial position for 100 – 125 ms before a target angle was prescribed as a step function. The angle time history was calculated in the goal-directed simulations and compared with the test data.

Training scenario 2: Arm motion control under external load

The arm model was trained to perform fast goal-directed motions under control signals similar to scenario 1, but in this scenario, a point mass was added to the forearm and the magnitude was randomly varied between 0-5 kg (Figure 2-6) in each iteration. Gravity was also implemented in the MB model vertically (along Z) and the magnitude was randomly flipped (Figure 2-6) to train the extension-flexion motion along and against gravity. The RLMAC was trained to perform the goal-directed motions with the same states and reward function in scenario 1, no extra information regarding the added mass or the gravity vector was provided to the RLMAC. The RLMAC was trained to formulate a response based on the kinematics data inputs from the states.
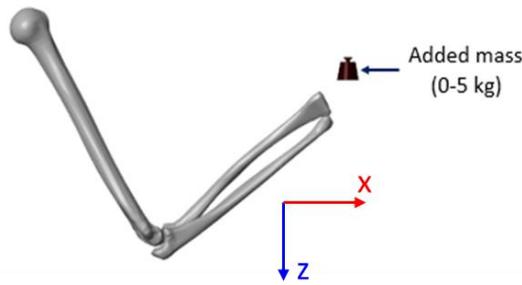
Figure 2-6: Arm MB model with attached point mass. The gravity was applied along Z axes and the magnitude was flipped in each iteration of the training.

In this scenario, training was performed only with the IAMR muscle recruitment scheme. Results from this training scenario will give an insight into the use of RLMAC to represent the motor adaptation to external loads as seen in humans that cannot be replicated by the use of linear feedback gains (Happee, 1993, 1992).

Testing scenarios: Response to novel environments

The robustness of the RLMAC depends on its ability to generate desired responses for novel loading environments that have not been encountered during the training. Thus, the RLMAC trained in the presence of a constant force field in scenario 2 was evaluated under a crash pulse (Shaw et al., 2009) representative of the inertial load experienced in the upper body during a motor vehicle collision applied horizontally to the forearm (Figure 2-7).
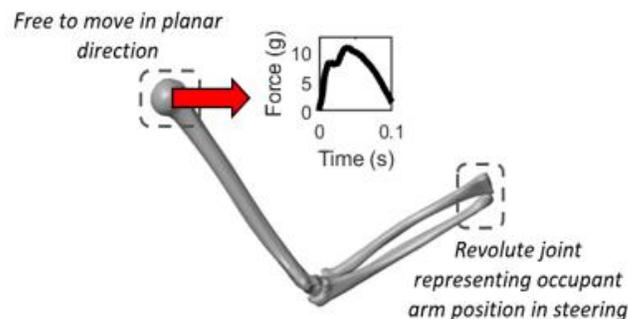
Figure 2-7: The trained arm model subjected to load case representing simplified automotive impact.

The radius-ulna lower end was constrained with a revolute joint limiting any planar movement similar to the interaction between an occupant's hand and the steering wheel. The scapula was free

to move in the planar (X-Z) direction and a stiffness of 1000 N/m was applied vertically (along Z) at the scapula proximal end to simulate the lower body weight. The crash pulse was applied to the scapula in the horizontal direction. At the start of the simulation, the arm was maintained at the neutral position (elbow angle 90^{o}) and the impact force was applied at the humerus end in the horizontal (X) direction (Figure 2-7). Two conditions were evaluated in this chapter – in the first condition, the force pulse was applied in the forward direction representing a frontal collision, trying to flex the elbow. In the second case, the force pulse was applied in the backward direction representing a rear collision and trying to extend the elbow. The results from the simulation with trained RLMAC were compared with the passive arm model to visualize the effects of active muscle forces on joint stiffness under external loads.

Table 2-3: Training and simulation scenarios

Training/ Simulation case	Load applied	Model Response
Evaluation of passive structural response	Torque at the elbow revolute joint	Moment-Angle response of the joint
Training scenario 1 – Goal-directed motion of the forearm with RL-MAC integrated MB model	No external loads	Angle-time response of the elbow joint
Training scenario 2 – Goal-directed motion of the forearm under external loads	Point mass attached to the radius and gravity field	Angle-time response of the elbow joint
Testing scenario: Response to novel loads	A crash pulse applied to the humerus proximal end	Angle-time response of the elbow joint

2.3 Results

The passive arm MB model was simulated by applying a moment varying between 1 and -1 to the revolute joint. The arm was initially placed at a neutral (90º elbow) position and the change in angle due to the application of the moment was evaluated. The passive stiffness of the muscle was found to be 0.955 Nm/ rad. The passive stiffness of the joint is the combination of the stiffness of 0.6 Nm/ rad prescribed at the elbow joint and passive muscle tensile forces. The stiffness magnitude was found to agree with previously published stiffness values (Hayes and Hatze, 1977; Howell et al., 1993; Wiegner and Watts, 1986) determined in experiments as well as measured computationally (Östh et al., 2012b).

The arm model integrated with the RLMAC was trained to perform the desired goal-directed movement. During each iteration of the training, the RLMAC was simulated for 600 ms and the target angle was provided to the elbow as a step function. The DDPG agent used in the RLMAC was updated every 5 ms during the training. In each training iteration, the initial angle and the target angle were varied randomly within the range of motion of the elbow.

Training scenario 1 - Bare arm motion control

The arm model was trained to perform goal-directed motion in the absence of any external loads. In this control scenario, the RLMAC was trained with both the GAMR and IAMR schemes. The training was assumed to converge when the average reward over 250 iterations converged, and the response of the trained agent at convergence was acceptable. The variation of average reward during the training for the GAMR and IAMR control schemes is shown in Figure 2-8. The RLMAC with the IAMR scheme took 3000 iterations to converge whereas the GAMR scheme took around 6000 iterations. The average reward can vary each time the arm model with RLMAC is trained

57

due to the inherent randomness in RL algorithms, but the nature of training is likely to be the same for all cases as the agent tries to maximize the reward based on reward function.
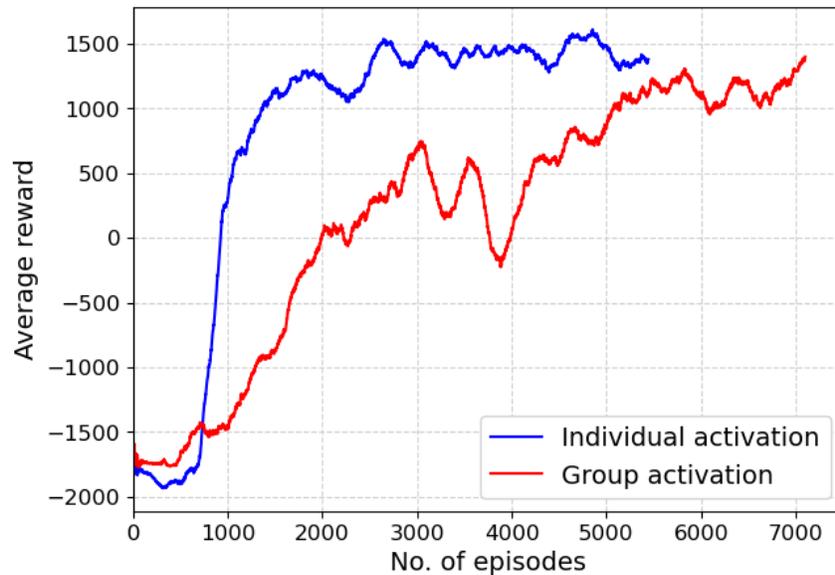
Figure 2-8: Plot showing the variation of average reward for individual activation (IAMR) and group activation (GAMR) with each episode during training. (The reward averaged over 250 episodes)

The trained RLMAC in scenario 1 was repositioned to perform a series of goal-directed kinematics to evaluate the response with volunteer data. The validation case selected for scenario 1 was simple in nature and could easily be replicated in the MB arm model (Kistemaker et al., 2006). The elbow joint was rotated in the range of 45° to 145° in the goal-directed simulations and the comparison with the volunteer data is provided in Figure 2-9. Both the activation schemes, the IAMR and the GAMR, show excellent agreement with the volunteer data as the response is dependent on the training where objective was to minimize the elbow angular error and the stiffness of the elbow joint. The trained RLMAC was able to stabilize the arm at the initial position and then move the arm once the target signal was applied as a step function. The RLMAC was able to maintain the arm within 5° of the target angle in all the simulation cases for both GAMR and IAMR. The minor position difference at the target can be due to the reward function used, which allows for an error of ±0.1 rad.
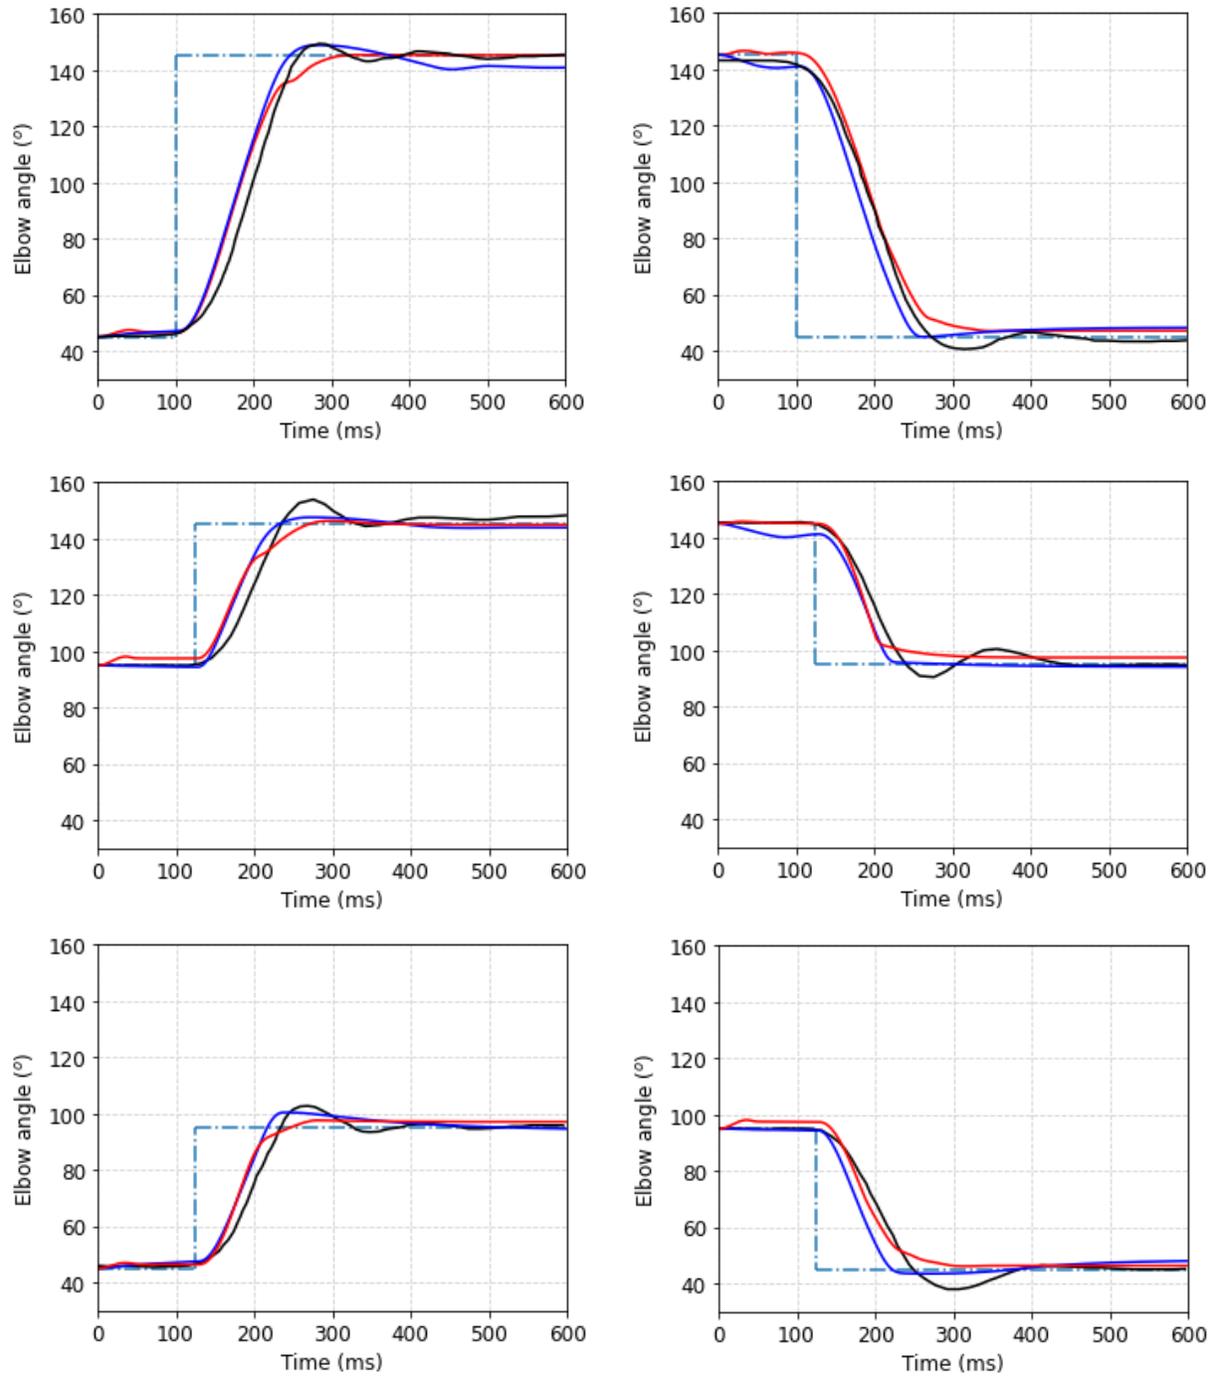
(a)   *(b)*

— Individual Activation   — Group Activation   — Volunteer data   – · – Target Angle

Figure 2-9: Elbow rotation angles by trained RL controllers in response to prescribed target angles and comparison with volunteer data with the individual (IAMR) and group (GAMR) activation scheme (a) Arm in flexion (b) Arm in extension.

The muscle activation levels of the extensors and flexors were measured during the simulation and are plotted in Figure 2-10. As the RL algorithms are randomized in nature and the DDPG agent in this chapter uses O-U noise to explore the action space, it is likely that a different agent may lead to a different set of muscle activations to achieve similar arm kinematics. However, the agents will have similar arm rotations as the RLMAC was trained with a reward function to minimize the elbow angular error.

(a)   (b)
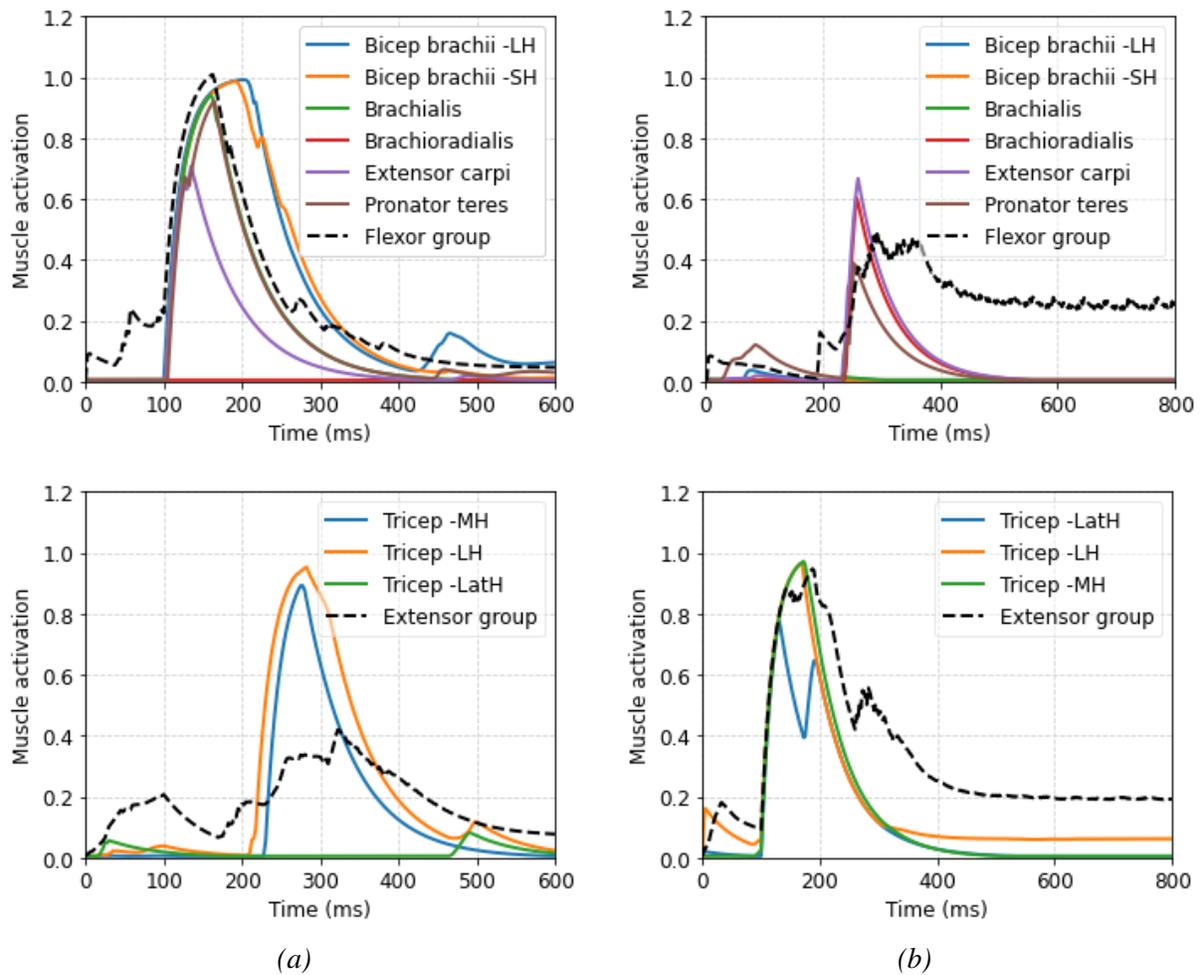
Figure 2-10: Muscle activation plots for extensor and flexor muscles (a) Arm in flexion - 45° -145° (b) Arm in extension - 145° - 45°. The individual curves in the figure are from the IAMR scheme and the extensor and flexor group curves are from the GAMR scheme.

The IAMR scheme outputs the activations of individual muscles as can be seen in Figure 2-10. Each muscle in the IAMR scheme has different peaks and a few muscles also remain inactivated

or with low activation during the entire simulation. The muscle activations also dropped after the target angle was reached. In the GAMR scheme, all the muscles have the same level of activation, and in the arm extension simulation, the muscle groups remain activated even after the radius-ulna reached the target and stabilized.

The muscle activations follow a biphasic or triphasic pattern which has also been observed in previous volunteer studies (Flament et al., 1984; Marsden et al., 1983). When the target signal is prescribed, the agonist muscles are activated initially which declines near the target joint position (Wadman et al., 1979). This decline in agonist activity is not enough to stop and stabilize the joint motion, thus a second burst of activation is seen, this time in the antagonist muscles to dampen the joint movement (Hannaford and Stark, 1983). In some cases, there is a third phase of agonist muscle activations when the antagonist muscle activities have not declined once the target position is reached (Happee, 1992). In the simulations performed in this chapter, the third phase was more prominent when the initial prescribed error was small. The fall in the activations after the elbow reached the target angle was due to a small penalty applied to the muscle activations in the reward function. Without the penalty factor, the extensors and flexors were found to co-contract at the final position.

Training scenario 2: Arm motion control under external loads

The RLMAC with added attached mass and gravity (scenario 2) took around 9000 iterations to converge. The trained RLMAC was evaluated under the extension-flexion motion with varied added mass. The trained agent was able to perform the desired kinematics in both extension and flexion motion of the elbow (Figure 2-11). In both the simulations, the arm was stable at the target position for different magnitude point masses.
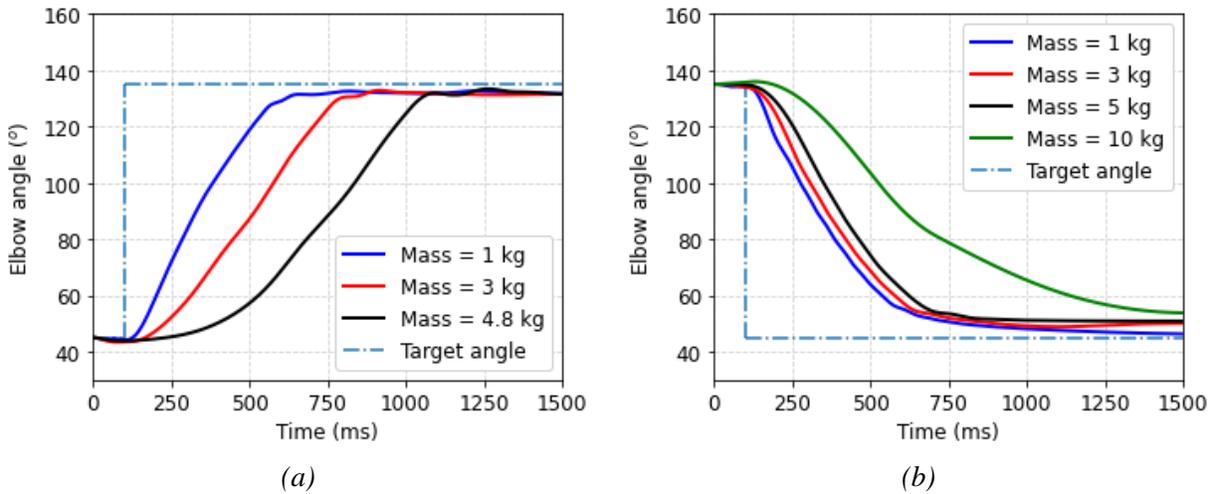
Figure 2-11: Rotation time history of the arm in the presence of external load against gravity (a) Arm in flexion - 45° -135° (b) Arm in extension - 135° - 45°.

In the flexion motion, the arm was able to carry a maximum mass of 4.8 kg against gravity. Figure 2-12a shows the activations of the bicep brachii and the brachialis at different magnitudes of point mass during the flexion motion. The muscle activations in scenario 2 had a similar biphasic or triphasic pattern. The flexor muscles had a maximum activation at 4.8 kg weight which eventually decreased at the target position due to the decrease in the moment arm. The trained RLMAC could carry up to 10 kg of mass against gravity in extension, even though the mass was limited to 5 kg during the training. This shows the robustness of the trained RLMAC that can operate under loads different than what is experienced during the training. During the extension with 10 kg, the agonist muscles had high activations throughout the simulation.

(a)   *(b)*
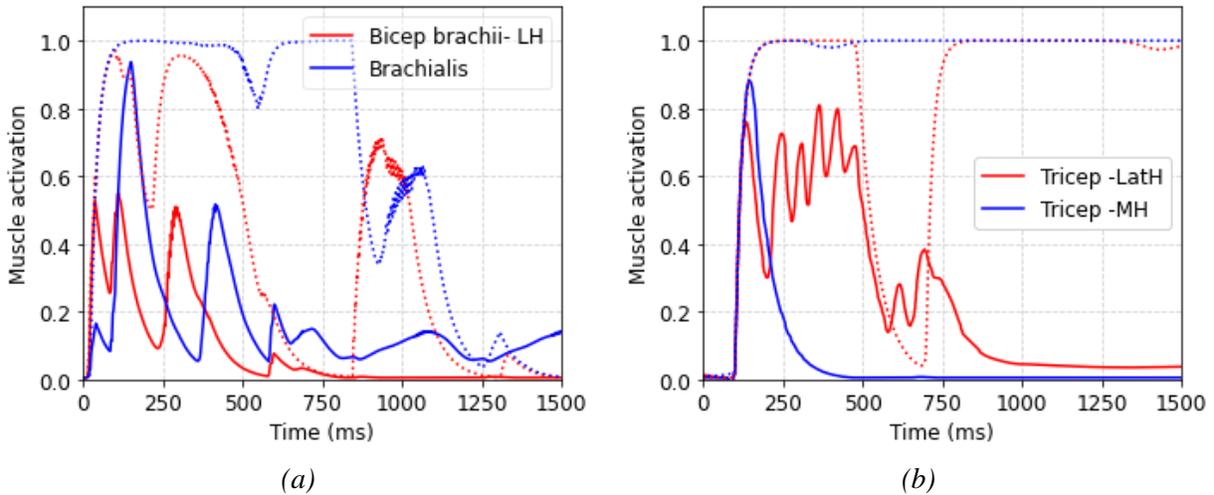
Figure 2-12: Activation time history of the agonist muscles in presence of external load against gravity (a) Arm in flexion - 45º -135º (solid – 1 kg, dashed – 4.8 kg) (b) Arm in extension - 135º - 45º (solid – 1 kg, dashed – 10 kg).

The RLMAC trained in scenario 2 was simulated for goal-directed motion in absence of any mass or gravity and the response was compared with the volunteer tests used for validation in scenario 1. The radius-ulna body undershot the experiment data initially which suggests that the peak velocity of the arm rotation was lower than in the experiments. However, the arm stabilized at the target angle within 300 ms after the target signal was prescribed to the RLMAC.

(a)   *(b)*
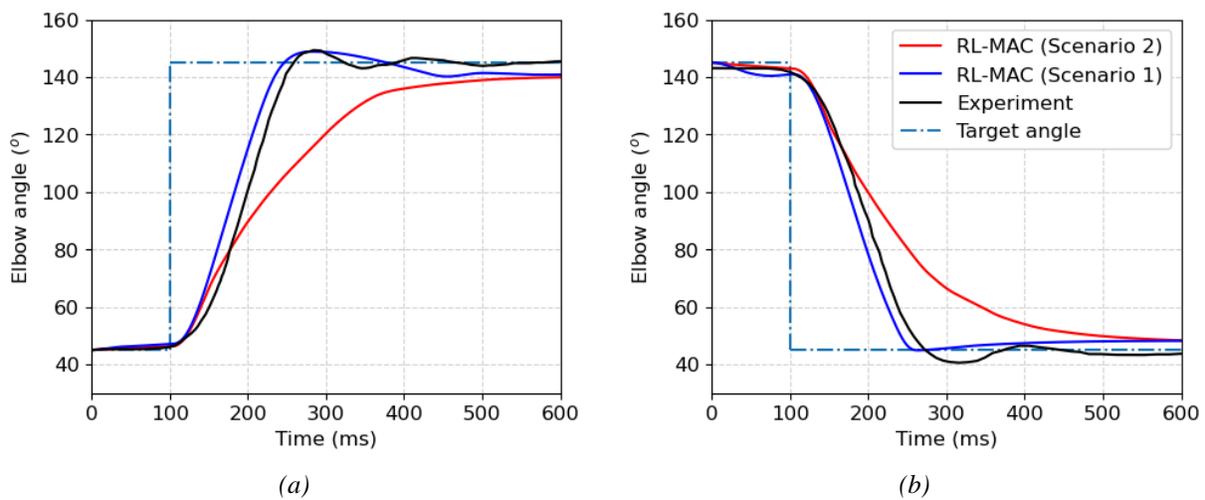
Figure 2-13: Comparison of RL-MAC from training scenario 1 (IAMR) and scenario 2 (a) Arm in flexion - 45º -145º (b) Arm in extension - 145º - 45º.

Testing scenario: Response to novel loads

RLMAC trained under external force fields (scenario 2) was subjected to a simplified crash pulse representing a driver bracing against the steering wheel during a low-speed motor vehicle collision. The active arm with RLMAC was stabilized at the neutral position before the crash pulse was applied horizontally (Figure 2-7). During the application of the load, the RLMAC was prescribed to maintain the initial position of the elbow joint to simulate the stiffening of the arm by an occupant. The objective of stiffening the elbow was used to evaluate the RLMAC response under extreme conditions and may not accurately represent the actual intent of the occupants during such impact cases.

(a)   (b)
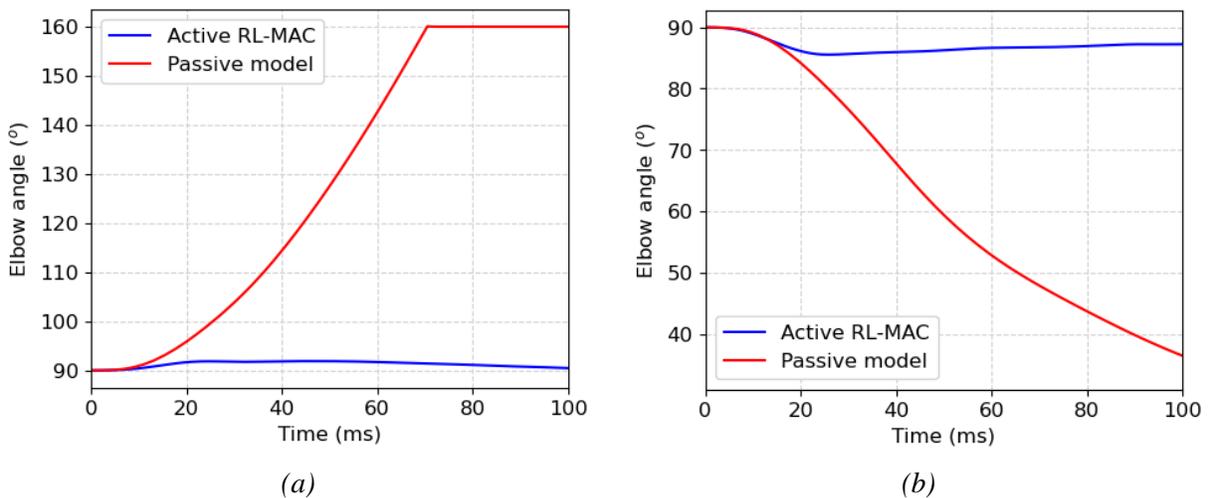
Figure 2-14: Rotation time history comparison of the arm in the presence of an external force pulse (a) Load applied backward (rear-end collision) (b) Load applied forward (frontal collision).

The response of the trained RLMAC has been compared with that of the arm passive model with baseline muscle activations in Figure 2-14. In both impact cases, the RLMAC was able to maintain the target angle at the elbow joint after the initial disturbance due to the application of load. The passive model however was unstable during the application of the impact pulse suggesting that the passive stiffness alone is insufficient to correctly predict the response of an occupant under such dynamic cases.
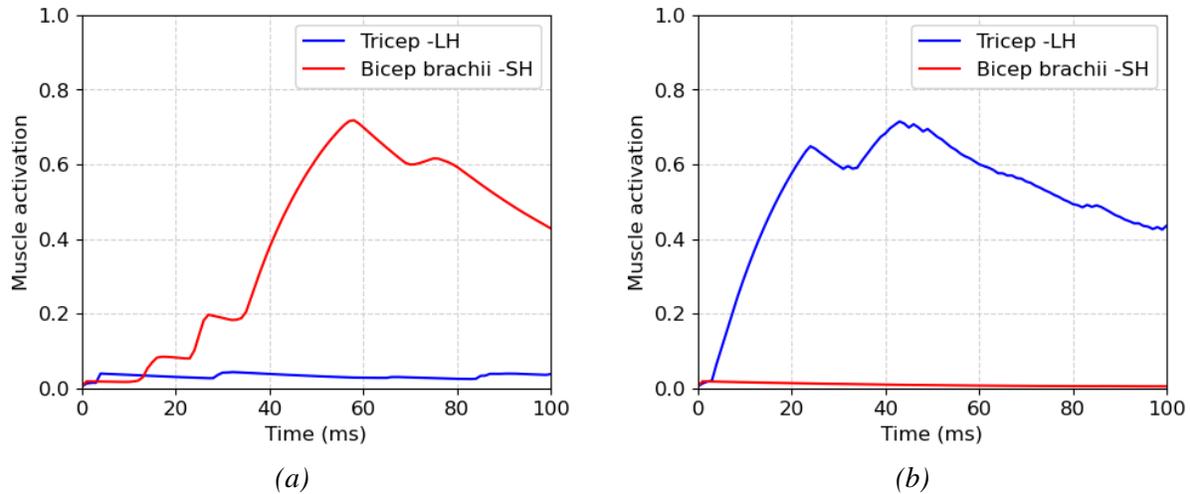
Figure 2-15: Muscle activation of tricep long head (extensor) and bicep brachii short head (a) Load applied backward (rear-end collision) (b) Load applied forward (frontal collision).

Figure 2-15 shows the activation histories of extensor and flexor muscles for rear and front impact cases. Initially at the onset of the force pulse, the agonist muscles begin to activate, and the activations drop once the force subsides. The antagonist muscles had minimal activations throughout, which increased slightly at the end of the run to damp the arm at the neutral position.

2.4 Discussions

The objective of the current chapter was to demonstrate the feasibility of the reinforcement learning (RL) framework for use in muscle control. The RL muscle activation control (RLMAC) framework implemented in this chapter was based on the DDPG agent which can output continuous output signals (Lillicrap et al., 2019). The RLMAC was integrated with a MB model of a human arm to generate goal-directed movements about the elbow.

Before integrating with the RLMAC, the passive stiffness of the arm model was evaluated and was determined to be within the ranges (0.96-1.2 Nm/rad) as previously measured in experiments (Hayes and Hatze, 1977; Howell et al., 1993; Wiegner and Watts, 1986). The stiffness of the elbow joint was due to the combination of joint stiffness and the passive muscle forces. As no damping

was considered in the Hill's muscle equations, the non-linearities at the joint were due to the constant joint damping assigned. Constant damping is a reasonable assumption as Popescu et al. (2003) found that the damping value was almost constant during the elbow rotations.

The RLMAC was trained to generate responses under two different muscle recruitment schemes – the GAMR in which the muscles were grouped into extensors and flexors and the IAMR scheme in which the muscles were activated individually. The objective of the training was to move the radius-ulna from an initial position to a target position by considering the angular error. Few previous studies have also considered the error in muscle contractile length as the control objective. Kistemaker et al. (2006) developed a PID-based muscle controller with the CE element length as the target signal (λ control) to generate similar goal-directed movements as performed in scenario 1 in this chapter. Martynenko et al. (2019) used a similar control approach to model the goal-directed arm rotations in an arm FE model. In the simulations studies using the CE error as the control criteria, pre-simulations with the passive model were performed to map the CE length with the elbow angles. While the approach was feasible for simple body regions like the arm model having 1 DoF, in a more complex region like the head and neck the muscle lengths corresponding to the target head orientation would be difficult to determine.

After the training, the RLMAC could generate the desired extension-flexion movements and the response of the RLMAC showed excellent agreement with the volunteer data (Kistemaker et al., 2006) for both the GAMR and IAMR scheme (Figure 2-9). The GAMR scheme had identical activations for all the muscles belonging to the same group, and the extensors and flexor muscles had a small activation level when the target orientation was achieved (Figure 2-10) which may lead to higher muscle fatigue. Also, grouping muscles corresponding to a preferred direction was possible due to the simple nature of the elbow joint, in a more complex body region isolating the

muscles corresponding to a preferred direction would be difficult. The IAMR scheme actuates only those muscles which are required for arm movement while keeping the other muscles inactivated which reduces muscle fatigue. In the following chapters of this dissertation, the IAMR scheme will be used for muscle control.

The RLMAC was also trained to synthesize the goal-directed motions in presence of a randomized point mass and gravity loads. Any information regarding the external force field was not included as a state parameter and the agent learned to generate the arm movement by the kinematics and muscle activation feedback based on how well the agent was performing relative to the control objective. Adding information on inertia and gravity to the state may improve the training process and the overall response of the RLMAC as the parameters are important in formulating the muscle response (Lacquaniti et al., 2015). However, these parameters will also result in over-fitting to the particular scenario (scenario 2) and would not extend well to other load cases. This kind of muscle response to changing loads is difficult to model using linear feedback gains as multiple muscles need to be modified simultaneously and the gains of all the muscles may not scale similarly (Happee, 1993; Smeets et al., 1990). In the RL training process, the external environment can be constantly changed to tune the neural networks in the agent to prevent any overfitting. The RLMAC in absence of any added inertia or gravity was found to undershoot the volunteer response (Figure 2-13) as the RLMAC has to depend on the kinematics state to formulate the response. Such adaptations to external loads is also seen in volunteers (Happee, 1993; Oh et al., 2021; Smeets et al., 1990; Wadman et al., 1979).

The RLMAC trained with a constant force field was able to maintain the stability of the joint in presence of dynamic loads implying the robustness of the controller. The dynamic loads applied to the arm model were substantially different from the training as a higher magnitude load was

applied for a shorter duration (Shaw et al., 2009). The testing scenario demonstrates the ability of the RLMAC to respond to novel loads which were not applied explicitly during the training. This feature is critical for muscle control in HBMs as there is a shortage of volunteer data under dynamic cases for validation and tuning of controllers. Along with the controller, evaluating the response of passive HBM also becomes important to ensure that forces developed in the active muscle and the range of motion of the joints in the active model are within the physiological limits. The robustness of RLMAC makes for a more general-use controller than the traditional feedback-based schemes and eliminates the requirement of constant calibrating and tuning of the controller based on the loading environments.

The objective of this chapter was to evaluate the feasibility of reinforcement learning for use in muscle control hence some simplifications were made in the human body model. A revolute joint was used to model the extension-flexion motion about the elbow and the motions in other DoFs were not considered. The goal-directed arm kinematics simulated in this chapter was mainly to compare the responses with volunteer data (Kistemaker et al., 2006) and other similar control studies (Kistemaker et al., 2006; Martynenko et al., 2019; Östh et al., 2012b). The series element representing tendon stiffness was also neglected while developing the MB model of the human arm. Few previous studies have included the series element while simulating goal-directed arm motions (Kistemaker et al., 2006; Wochner et al., 2019). However, a study conducted by Bayer et al. (2017) concluded that the series element contributes only 7-8% of the total forces during arm movements. The stiffness of the passive arm model was verified in the study ensuring that the assumption of a rigid tendon did not affect the response of the arm MB model or the overall objective of the study.

68

In this chapter, the reward function penalized the muscle activity as a means of solving the redundancy problem thus reducing muscle fatigue (Pedotti et al., 1978) during the simulations. However, other cost functions associated with the human musculoskeletal system such as muscle work (Margaria, 1968; Umberger et al., 2003) or joint energy expenditure (Berret et al., 2008b) can also be used.

2.5 Conclusions

The present chapter has presented a methodology to integrate a robust muscle control framework based on reinforcement learning for control of rotations of the elbow joint. The RLMAC developed in this study could generate goal-directed motion in an arm MB model, control the arm motion in the presence of an external force field, and also stabilize the elbow joint in presence of dynamic loads. Such a robust control mechanism is important for loading environments that can be injurious in nature and thus difficult to gather data for tuning the controller.

The present study is a preliminary investigation into the use of RL algorithms for muscle control. The potential for deep RL controllers can be better understood in more complex body regions like the head and neck with multiple non-linear joints and random muscle orientations which make the control of such a body region difficult. The development of RL based control framework for the head and neck region will be performed in the following chapters.

Chapter 3 – Development of a head-neck multibody model for postural control

The head-neck region is the most complex body region for control due to multiple skeletal joints and the complex arrangement of muscles. Neck physiology provides high flexibility, mobility, and stabilization to the head. However, due to the enormous complexity of neck geometry and muscle orientation, it is difficult to develop an efficient control strategy for the neck muscles. For the implementation of the RLMAC control framework for the head, a multibody model of the head-neck region of the 50[th] percentile male in a seated position has been developed. The neck model consists of T1-C1 vertebrae, the head modeled as Hybrid III headform, and 46 neck muscles. The neck model with passive muscles was validated under the rear impact loading scenario.

3.1 Introduction

The development of an accurate musculoskeletal model of the head and neck is necessary for modeling the active control behavior. Injury to the head and neck region being one of the leading causes of deaths and chronic disabilities has resulted in a lot of focus on the development of computational models of the head and neck region. Some of the earlier developed computational models of the head and neck used simplified lumped mass connected with deformable links (Bowman et al., 1981; Goldsmith et al., 1984; Williams and Belytschko, 1983). In these studies, the lumped masses were connected using non-linear springs and dampers modeled with empirical expressions (Huston et al., 1978; Tien and Huston, 1987).

de Jager et al. (1996) developed an anatomically detailed model of the head and neck region with rigid vertebrae and nonlinear joints and representation of muscles as the line of action of forces along the insertion points. Van der Horst et al. (1997) modified the rigid model to enhance the representation of muscles for improved rotational response under automotive loading. Stemper et al. (2004) developed a multibody model with detailed validation of passive kinematics for

70

whiplash-type loading. Vasavada et al. (1998) developed a multibody model with detailed Hill-type muscle representation to study the moment arm and moment generating capacity of the neck muscles in extension, flexion, and axial rotation. The model by Vasavada et al. was later improved by Mortensen et al. (2018) by adding hyoid muscles for a realistic representation of flexion strength. de Bruijn et al. (2016) developed a musculoskeletal model of the head-neck region with the non-linear representation of the intervertebral joints (T1-C1) in all 6 DoFs.

Brolin et al.(2005, 2004) developed a finite element (FE) model of the upper cervical spine with 1-D muscle representations to study the effect of the muscle activations on the occupant response. Hedenstierna et al. (2009, 2006) extended the FE model to include a continuum representation of the passive muscle properties for simulations under impacts. Panzer et al. (2011, 2009) developed a FE model of the neck region with a detailed representation of intervertebral disks and ligaments to predict occupant response under front crashes.

More recently, efforts have also been made to develop head and neck models for different anthropometries. Dibb et al. (2013) developed hybrid FE models of pediatric (6 years and 10 years) along with adult neck regions for automotive simulations. A detailed FE model of the head and neck region with a fifth percentile female anatomy was created by Davis et al. (2016) along with the whole body model for use in motor vehicle environments. Östh et al. (2017) developed a FE model for the head and neck with a detailed representation of ligaments and soft tissues and a rigid head for investigating whiplash-related injuries in females.

The development of a musculoskeletal model for integration with RLMAC poses a unique challenge. Due to the inherently iterative nature of the RL algorithms, the model should be fast running. However, the model should also represent anatomy and physiology with minimum simplifications for outputting acceptable responses. This chapter outlines the development process

71

of a multibody model of the 50th percentile male head and neck for use with RLMAC. The neck model developed in this study will also be validated under passive rear impact conditions.

3.2 Methodology

For this study, a previously developed 50th percentile (M50) head-neck model (Fice et al., 2011; Panzer and Cronin, 2009) was used for the anatomy of the neck model. The neck model consists of forces whose magnitude was determined from a Hill-type muscle model with an added damper in series (Bayer et al., 2017; Haeufle et al., 2014).

3.2.1 Anatomy of the cervical spine

The neck model incorporates seven cervical vertebrae (C1-C7), whose anatomy is based on 3D dataset by Viewpoint Datalabs (Orem, Utah) and implemented previously in a FE model (Fice et al., 2011; Panzer et al., 2011; Panzer and Cronin, 2009). The neck model also includes the first thoracic vertebra (T1) as a reference for applying the boundary conditions of the trunk. The skull was modeled using the shape of a Hybrid III head form and assigned the mass and inertia properties of the M50 skull by Panzer et al. (2011) (Figure 3-1 a). The skull was considered rigid, and the muscles attached to the skull applied forces to the skull as a whole.

The vertebrae T1-C3 were modeled separately and were connected with non-linear joints representing the stiffness imparted by the intervertebral discs, the ligaments, and the facet joints. The C2 and C1 vertebrae were lumped into one body as the atlanto-axial joint limits motion in the sagittal plane (Nightingale et al., 2002) and further connected to the C3 and the skull.

72

<center>(a)   (b)</center>
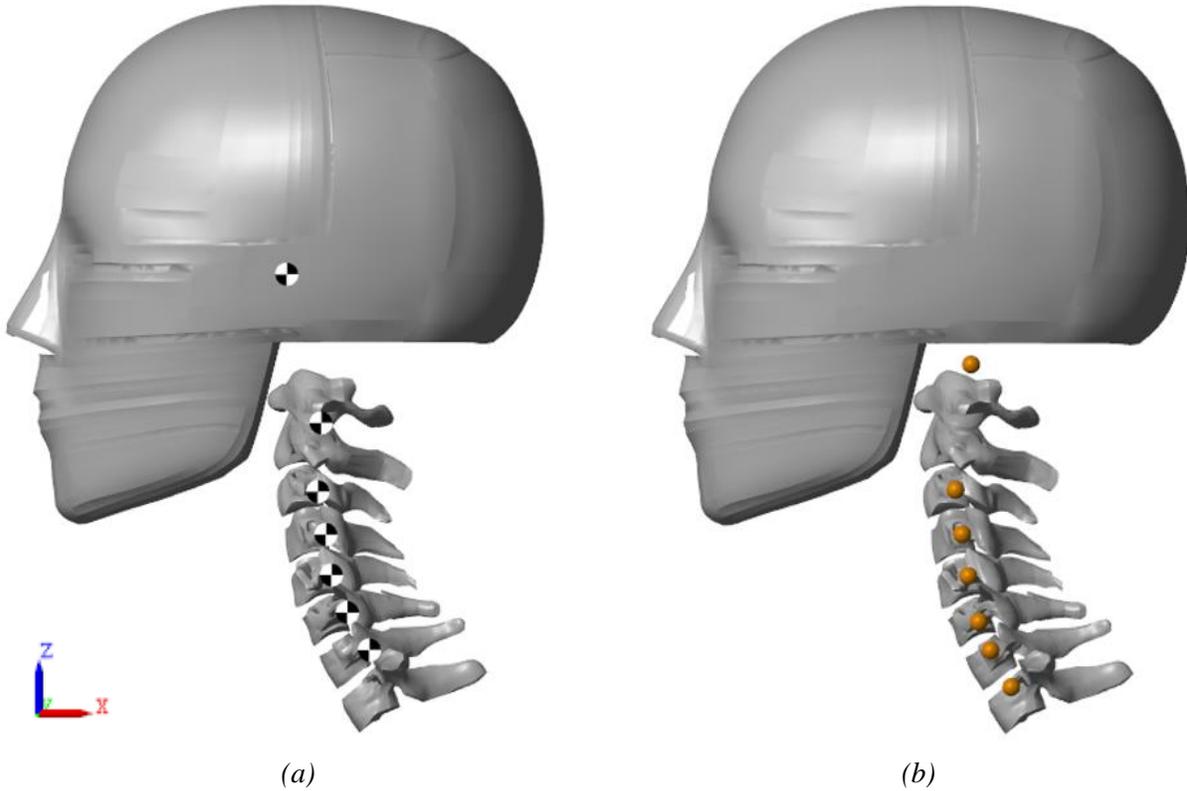
Figure 3-1: Development of the UVa neck model (a) Representation of the bony anatomies with each vertebra's center of masses (CoM). The combined CoM is displayed for C2-C1 combined mass. (b) Location of 6 DoF joints between the vertebrae and the skull.

The mass of the skull and vertebrae and their CoM locations were adapted from the source FE model (Panzer et al., 2011). The head mass and moment of inertia implemented in the neck model are within the ranges of values reported in the literature (Walker Jr et al., 1973; Yoganandan et al., 2009) and have been provided in Table 3-1.

<center>Table 3-1: Mass and inertia properties of the skull and cervical spine</center>

Part	Mass (kg)	Moment of Inertias (Kg-mm^2)				Center of Mass (mm)	
		I_{xx}	I_{yy}	I_{zz}	I_{xy}	X	Z
Skull	4.376	21060	23300	15200		-31.047	181.824
C1	0.0226	9.606	3.759	12.259	0.4085	-16.018	128.258
C2	0.0254	5.3329	5.3724	7.3772	1.1312	-17.328	112.225
C3	0.0162	2.8499	2.1659	4.2767	0.3067	-18.088	91.172
C4	0.017	3.5055	2.183	4.807	0.1967	-14.747	73.265
C5	0.0188	3.853	2.559	5.558	0.146	-12.441	56.381
C6	0.0191	4.1054	3.2598	6.441	0.34	-5.796	40.739

Note: The CoM of each vertebra is measured from the T1 CoM along the coordinate system in Figure 3-1a.

The intervertebral joints were allowed translation and rotation in all 6 DOFs. The location of the joints of the model was approximated from the literature (Chancey et al., 2007; Dibb et al., 2013; van Mameren et al., 1992).

Table 3-2: Location of the center of rotations of the intervertebral joints measured from T1 CoM

Joint	X (mm)	Z (mm)
Skull – C1	15.1	144.7
C1- C3	21.55	92.45
C3-C4	18.81	73.94
C4-C5	16.27	56.55
C5-C6	11.85	37.44
C6-C7	7.00	25.54
C7-T1	1.6	10.24

The location of the intervertebral joints has been shown in Figure 3-1 b and provided in Table 3-2. The skull-C1 joint connects the head (represented by HIII headform in the model) with the lumped C1-C2 vertebral mass. The seven joints in the neck model provide stiffness in all the 6 DoFs (3 translation and 3 rotational). The stiffnesses of the joints are represented by spring-damper combinations with non-linear force-deflection relation. Figure 3-1 shows the characterization of the joint stiffness. Bilinear curves are used to represent both the linear and rotational stiffness at the joints to capture the non-linear nature of the joints without affecting the simulation time.

(a)   *(b)*
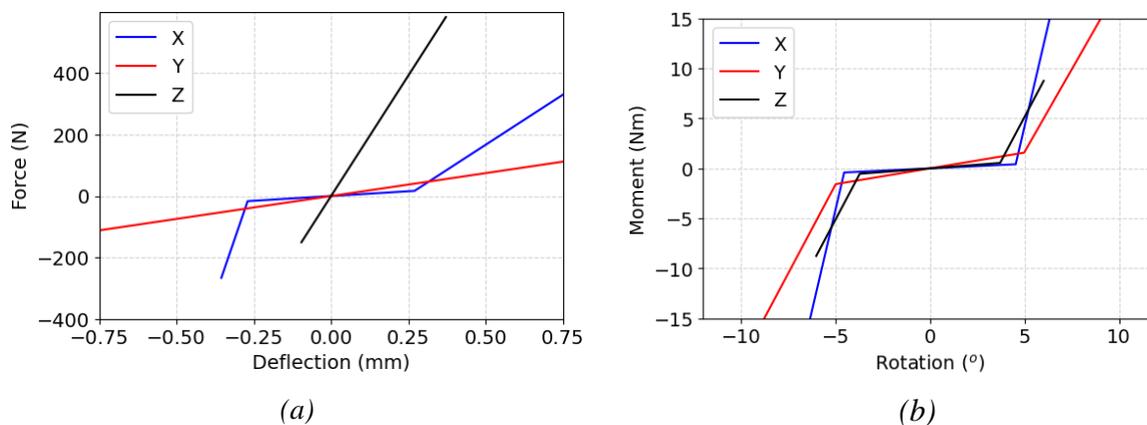
Figure 3-2: Stiffness of the joint between C4 and C5 (a) Linear stiffness (b) Rotational stiffness.

The joint stiffness curves in the neck model were adapted from many sources of data available in the literature (Table 3-1). The stiffness and damping prescribed at the joints approximate the stiffness imparted by intervertebral discs, the facet joint, ligaments, and other soft tissues present between the vertebrae. The joints do not consider passive muscle stiffness.

Table 3-3: Stiffness and damping of the intervertebral joints

Linear stiffness		
Direction	Stiffness	Damping
Anterior-Posterior (X)	Bilinear curve (Shea et al., 1991)	834 Ns/m
Lateral shear (Y)	138 – 232 N/mm (Liu et al., 1982)	352 Ns/m
Axial (Z)	800 – 2446 N/mm (Dibb et al., 2009; Shea et al., 1991)	614 Ns/m
Rotational stiffness		
Lateral bending (X)	Bilinear curve (Yoganandan et al., 2007)	0.78 – 1.5 Nms/ rad
Extension – Flexion (Y)	Bilinear curve (Nightingale et al., 2007)	0.78 – 1.5 Nms/ rad
Axial rotation (Z)	Bilinear curve (Chang et al., 1992; Dibb et al., 2013; Panjabi et al., 2001)	0.39 – 1.5 Nms/ rad

The damping coefficient values were adjusted based on previous studies to improve the model behavior in impact conditions and remove noises when stabilized (de Jager et al., 1996; Dibb et al., 2013; Happee et al., 2017; Mortensen et al., 2018). The damping coefficient considered in the joint stiffness allows for reproducing the viscoelasticity of the soft tissues as well as critically damping any vibrations between the vertebrae and skull during simulations.

3.2.2 Development of the cervical muscles

The neck muscles provide stability and strength to the vertebrae and head and support voluntary movements of the head. The neck model consists of 46 muscles that are symmetrical about the sagittal plane (Figure 3-3). The muscles are represented as force vectors between two points throughout the muscle length representing the active and passive stiffness.

(a)   *(b)*
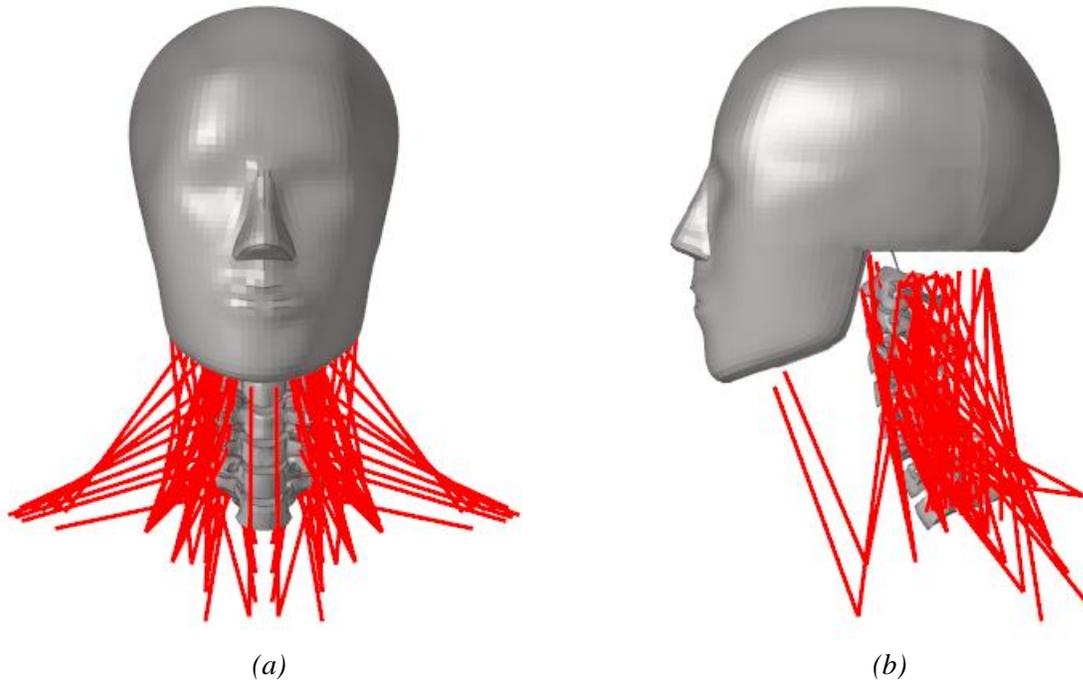
Figure 3-3: Neck model with muscle representation (a) Front view (b) Side view. The neck muscles are symmetrically located about the sagittal plane.

The muscle forces are modeled using a modified Hill-type muscle with a serial damping element (Bayer et al., 2017; Haeufle et al., 2014). The muscle-tendon unit (MTU) consists of four elements – the contractile (CE) and passive (PE) elements representing the active and passive muscle forces, and the series elastic element (SEE) and damping element (SDE) which represents the compliance of the tendon at the attachment to the bones (Figure 3-4). The total length of the MTU (l_{mtu}) is the summation of the length of the muscle (l_m) and the length of the tendon unit (l_t).
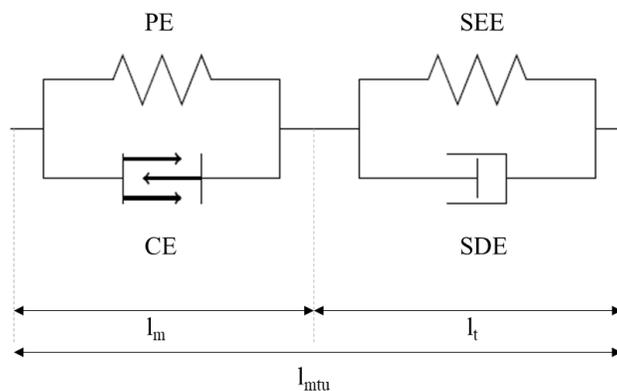
Figure 3-4: Muscle tendon unit of modified Hill-type model with serial elastic and damping elements (Haeufle et al., 2014).

The contractile element (CE) is the active part of the muscle and represents the force generated when the muscle is actuated by the CNS. The CE force depends on the length (F_l) and velocity (F_v) of the contraction of the muscle fibers.

$$F_l = exp\left\{ -\left| \frac{l_{CE}/l_{opt} - 1}{\Delta W} \right|^{k_{sh}} \right\}$$

Equation 3-1

Equation 3-1 describes the force-length relationship of the contractile element. ΔW controls the shape of the F_l curve while K_{sh} is the exponential parameter influencing the rise of force with the length of the CE.

The relationship between the active forces and the contractile velocity (F_v) is modeled as

$$F_v(V_{CE} \leq 0) = F_{max}\left\{ \frac{a_t F_{isom} + A_{rel}}{1 - \frac{V_{CE}}{B_{rel}l_{opt}}} - A_{rel} \right\}$$

Equation 3-2

Equation 3-2 shows the force-velocity relationship during the concentric motion of the muscle. For the eccentric part

$$F_v(V_{CE} \geq 0) = F_{max}\left\{ \frac{a_t F_{isom} + A_{rel,e}}{1 - \frac{V_{CE}}{B_{rel,e}l_{opt}}} - A_{rel,e} \right\}$$

Equation 3-3

The concentric and eccentric parts of the F_v are both continuous and differentiable at $V_{CE} = 0$. The Fv curve described by Equation 3-2 and Equation 3-3 gives a realistic force-velocity relationship where the force increases with increase in velocity while contracting and asymptotes at a maximum value when the CE elongates (van Soest and Bobbert, 1993). A_{rel} and B_{rel} are normalized hills parameter with peak values of $A_{rel,0}$ and $B_{rel,0}$ (Haeufle et al., 2014).

77

The muscle passive stiffness is also dependent on the muscle length (l_{ce}). The passive muscle element (PE) exerts tensile forces when extended beyond its rest length ($l_{PE,0}$). The magnitude of the forces generated by the PE (F_{PE}) is given by

$$F_{PE}\ (l_{CE}) = \begin{cases} 0 & l_{CE} < l_{PE,0} \\ K_{PE}\ (l_{CE} - l_{PE,0})^{v_{PE}} & l_{CE} < l_{PE,0} \end{cases}$$

Equation 3-4

K_{PE} is the non-linear coefficient relating the F_{PE} to elongation of the passive element from rest length (Günther et al., 2007).

The serial elastic element in the Hill-type muscle model is the representation of tendon forces. The force exerted by the series element (F_{SEE}) begins to increase exponentially which transitions into a linear relationship with the series element length (l_{SEE}).

$$F_{SEE} = \begin{cases} 0 & l_{SEE} < l_{SEE,0} \\ K_{SEE,nl}\ (l_{SEE} - l_{SEE,0})^{v_{SEE}} & l_{SEE} < l_{SEE.nl} \\ \Delta F_{SEE,0} + K_{SEE,l}\ (l_{SEE} - l_{SEE,nl}) & l_{SEE} \geq l_{SEE.nl} \end{cases}$$

Equation 3-5

The parameters in Equation 3-5 are derived below (Günther et al., 2007; Haeufle et al., 2014)

$$l_{SEE,nl} = (1 + \Delta U_{SEE,nl})l_{SEE,0}$$

Equation 3-6

$$v_{SEE} = \frac{\Delta U_{SEE,nl}}{\Delta U_{SEE,l}}$$

Equation 3-7

$$K_{SEE,nl} = \frac{\Delta F_{SEE,0}}{(\Delta U_{SEE,nl}\ l_{SEE,0})^{v_{SEE}}}$$

Equation 3-8

$$K_{SEE,l} = \frac{\Delta F_{SEE,0}}{\Delta U_{SEE,l}\ l_{SEE,0}}$$

Equation 3-9

$l_{SEE,0}$ is the rest length of SEE, $\Delta U_{SEE,nl}$ is the relative stretch of SEE at the transition between non-linear and linear regions, $\Delta F_{SEE,0}$ is the force at transition, and $\Delta U_{SEE,l}$ is the additional linear stretch causing a force increase of $\Delta F_{SEE,0}$.

78

The Hill-type muscle model used in the study also incorporates a serial damping element (SDE) modeled similarly to a viscous dashpot (Mörl et al., 2012). The force of the damping element (F_{SDE}) is characterized by

$$F_{SDE} = D_{max}\left((1 - R_{SDE})\frac{F_{CE} + F_{PE}}{F_{max}} + R_{SDE}\right).(V_{MTC} - V_{CE}) \qquad \text{Equation 3-10}$$

R_{SDE} is the damping at 0 muscle force and its maximum value is 1. D_{max} is the maximum damping coefficient when the force in the muscle equals the maximum force F_{max}.

The cervical muscle origin and insertion points were adapted from the source FE model by Panzer et al. (2011, 2009). The muscles with large cross-sectional areas and wide origins or insertions were separated into multiple strands with suitable vertebral insertion points for each strand. The forces developed by the muscles were divided among the muscle strands (Figure 3-5 a).

(a)   *(b)*
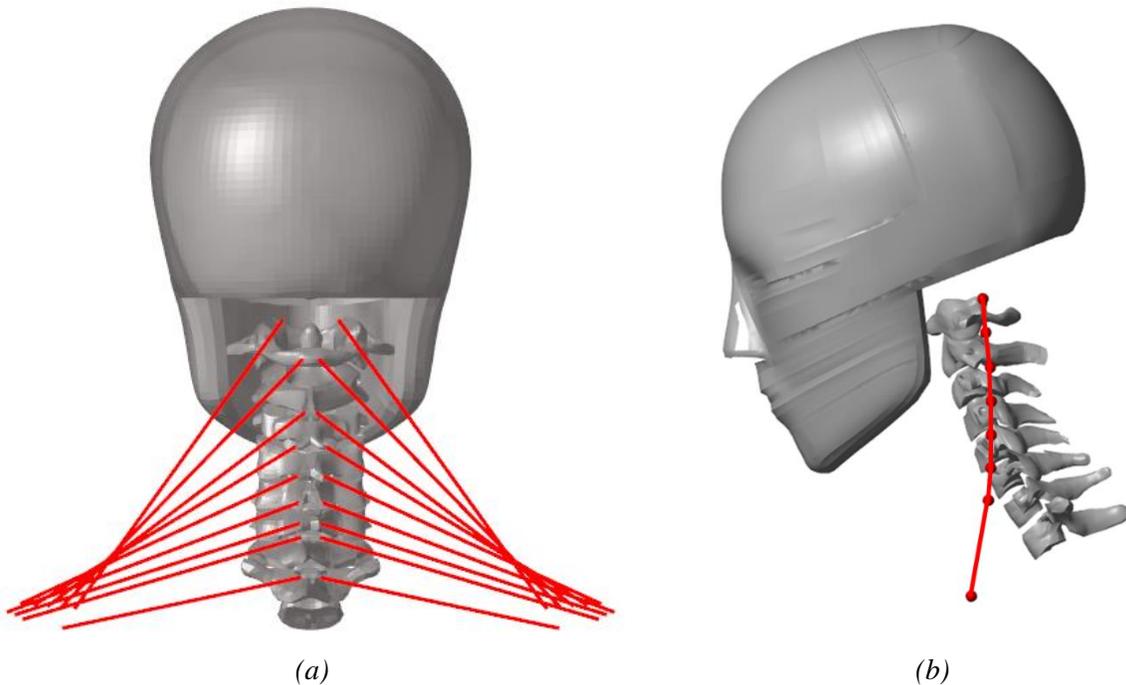
Figure 3-5: Muscles in the UVa neck model (a) The trapezius muscles are divided into multiple strands (b) The intermediate vertebral attachment points of the left sternocleidomastoid to model the changing load directions.

For maintaining a biofidelic force path between the endpoints, the muscle strands were split into segments in series, and the intermediate points were attached to the nearest vertebra (Figure 3-5 b). The splitting of the muscle strand across length accounted for maintaining the physiologically accurate muscle forces that follow the neck curvature during different neck motions (Dibb et al., 2013; Fice et al., 2011; van der Horst et al., 1997). Without considering the routing of the Hill's muscles along vertebrae, the muscle forces will act along a straight path between the origin and the insertion which may not be biofidelic (van der Horst et al., 1997). The muscle mass was divided between the origin and insertion points, as well as between the bony attachment points (Figure 3-6)

(a)   *(b)*
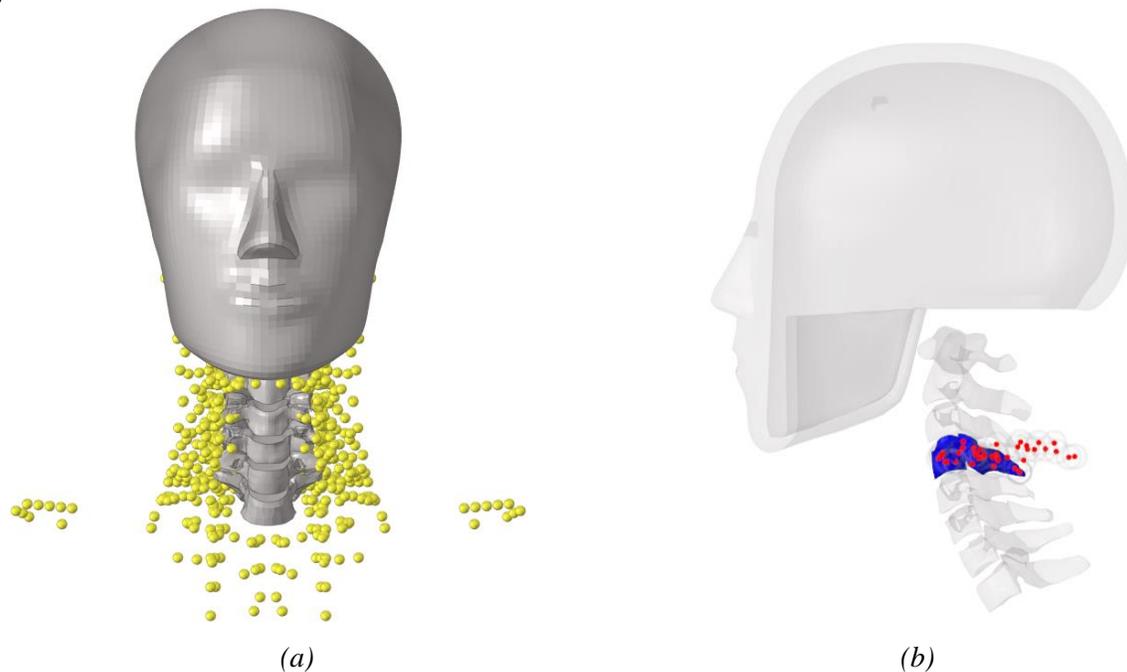
Figure 3-6: Muscle mass elements (a) Distribution of muscle mass in the model (b) Attachment points at the C4 vertebra.

The muscles in the neck model, which attach to the anatomical landmarks not included in the model, like the thoracic vertebrae other than T1, ribs, or the clavicle, were attached at T1. The cervical muscles along with the different muscle anatomical and physiological parameters used in the neck model are summarized in Table 3-4.

Table 3-4: Muscle parameters used in the neck model

Muscle (with strands)		Origin	Insertion	F_{max} (N)	L_{MTU} (mm)	L_{opt} (mm)	Pennation angle – α (°)
Oblique Capitis Superior		C1	Skull	43.98	21.879	12.0335	10
Rectus Capitis Major		C2	Skull	83.977	49.224	30.0266	5
Rectus Capitis Minor		C1	Skull	46.025	14.477	7.7	5
Rectus Capitis Anterior		C1	Skull	64.955	22.742	12.1	0
Rectus Capitis Lateral		C1	Skull	64.955	11.877	6.3	0
Longus Capitis	A	C3	Skull	17.167	62.042	31.021	5
	B	C4			86.302	43.151	
	C	C5			103.678	51.839	
	D	C6			121.179	60.589	
Longus Colli superior	A	C3	C1	11.50	40.849	20.424	5
	B	C4			64.550	32.275	
	C	C5			81.104	40.552	
Longus Colli inferior	A	T1	C5	17.259	63.742	31.871	5
	B		C6		61.050	30.525	
Longus Colli vert	A	C5	C2	22.827	76.368	38.184	5
	B	C6	C3		73.237	36.628	
	C	C7	C4		76.933	38.466	
Anterior Scalene	A	C3	T1	23.476	128.878	56.706	10
	B	C4			104.727	46.079	
	C	C5			88.200	38.808	
	D	C6			72.043	31.698	
Middle Scalene	A	C2	T1	11.321	122.847	55.035	10
	B	C3			116.027	51.980	
	C	C4			97.067	43.486	
	D	C5			80.895	36.241	
	E	C6			69.288	31.041	
	F	C7			54.142	24.255	
Posterior Scalene	A	C5	T1	17.538	84.746	38.983	10
	B	C6			73.107	33.629	
	C	C7			57.230	26.326	
Sternocleido Mastoid	A	T1	Skull	123.04	170	103.7	10
	B				202.335	123.424	
Iliocostalis Cervicis	A	T1	C3	12.991	120.073	91.736	0
	B		C4		102.708	78.469	
	C		C5		91.457	69.873	
	D		C6		79.826	60.987	
Longissimus Capitis	A	C4	Skull	9.836	67.146	52.378	0
	B	C5			84.165	65.648	
	C	C6			97.635	76.155	
	D	C7			116.596	90.945	
	E	T1			118.281	92.259	
Longissimus Cervicis	A	T1	C2	14.94	85.844	64.984	0
	B		C3		99.044	74.976	
	C		C4		98.128	74.283	
	D		C5		103.286	78.187	
	E		C6		108.995	82.509	

Muscle	Strand	Origin	Insertion	L_{MTU}			α
Multifidus	A	C4	C2		42.102	31.8712	
	B	C5	C3		33.992	25.732	
	C	C6	C4	10.02	33.954	25.703	0
	D	C7	C5		44.449	33.648	
	E	T1	C6		38.750	29.334	
	F		C7		45.557	34.487	
Semispinalis Capitis	A	C4			65.937	39.562	
	B	C5			82.398	49.438	
	C	C6			97.014	58.208	
	D	C7	Skull		114.356	68.613	
	E			30.622	118.095	70.857	10
	F				135.879	81.527	
	G	T1			153.812	92.287	
	H				188.814	113.288	
	I				207.966	124.779	
Semispinalis Cervicis	A		C2		80.538	57.423	
	B	T1	C3	38.23	79.908	56.974	5
	C		C4		83.750	59.713	
	D		C5		88.646	63.204	
Splenius Capitis	A	C7			119.273	64.407	
	B		Skull	38.602	139.859	75.524	0
	C	T1			156.124	84.307	
	D				169.640	91.606	
Splenius Cervicis	A		C1		169.979	91.788	
	B	T1	C2	23.848	169.472	91.515	0
	C		C3		184.296	99.512	
Levator Scapula	A	C1			135.657	73.255	
	B	C2	T1	38.973	120.238	64.928	0
	C	C3			116.244	62.772	
	D	C4			99.601	53.784	
Trapezius	A	Skull			176.061	95.073	
	B	C1			168.366	90.917	
	C	C2			156.787	84.665	
	D	C3	T1		148.039	79.941	
	E	C4		76.27	147.822	79.824	20
	F	C5			151.457	81.787	
	G	C6			144.849	78.218	
	H	C7			140.500	75.870	
Omohyoid		T1	Skull	58.738	124.066	86.846	0
Sternohyoid		T1	Skull	58.738	141.174	98.822	0

The L_{MTU} in Table 3-4 refers to the length of the muscle-tendon unit, measured as the Euclidian distance between the original and insertion points of the muscles at the initial position. F_{max} is the maximum force that can be exerted by each strand of the muscle. The pennation angles (α) of each muscle have been adapted from Kamibayashi and Richmond, (1998) and Vasavada et al. (1998).

As the mass of the muscles was distributed at the bony attachment points of the muscles (Figure 3-6), the segmental mass and the inertia got modified from the prescribed magnitudes. The updated mass of each segment is provided in Table 3-5.

Table 3-5: Mass at each segment with muscles

Segment	Mass (kg)
Head	4.449
C2-C1	0.207
C3	0.101
C4	0.112
C5	0.119
C6	0.124
C7	0.123
Total mass	5.235

The mass of each muscle was calculated by Panzer et al. (2011), considering a density of 1.06 g/cm^3 (Ward and Lieber, 2005). The total mass of the neck model, including the mass of the vertebra and the muscles was 5.235 kg, which was near to the average head and neck mass of 5.1 kg as reported by Dempster and Gaughran (1967).

3.2.3 Passive model validation

The fundamental objective of the neck model is to integrate it with RLMAC for active muscle control under voluntary kinematics and impact scenarios. However, to utilize the model for such purposes, it is necessary to validate the response of the passive model.

The neck model was evaluated under a load case resembling a rear impact scenario performed by Stemper et al. (2004). In the tests, the head-neck complexes from the cadavers were segmented and potted at T1. The extracted head-neck complexes were installed in a mini sled apparatus to impart the rear impact loading to accelerate the T1 in the anterior direction. Photo reflective markers were attached to cervical vertebras from C2 to C7 to track the relative angles between the vertebrae throughout the testing. The rear displacement of the head (retraction) and the pitch angle

83

with respect to T1 were also measured during the forward motion of the sled. Validation corridors were calculated and reported in the study for the sagittal angle between the adjacent vertebrae as well as for the head retraction and the pitch angle.

To impart a similar motion at the T1 in the neck model, the vertical and lateral translations, and rotations in all the DoFs were constrained. Sled acceleration corresponding to an impact velocity of 2.6 m/s was applied at the T1 in the anterior direction (Figure 3-7).

Application of the acceleration pulse

(a)   *(b)*
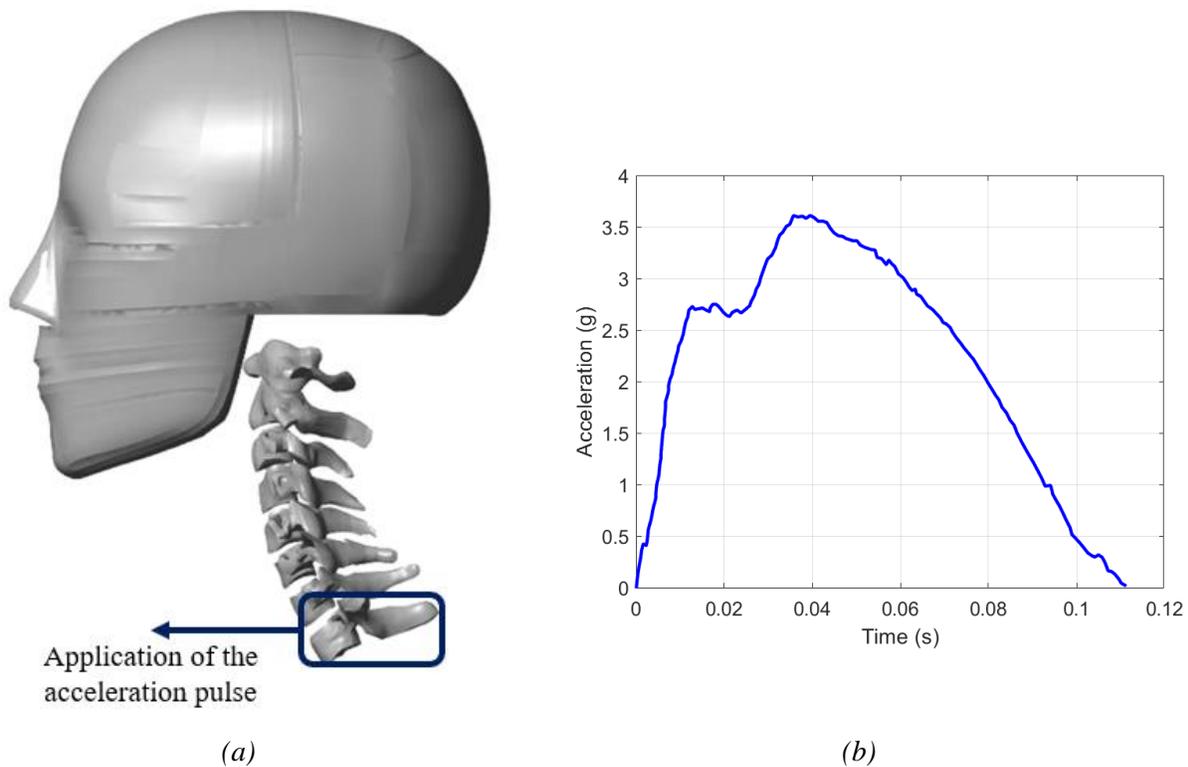
Figure 3-7: Application of rear impact load to the head model (a) The acceleration pulse applied to T1 horizontally in anterior direction (b) Acceleration pulse corresponding to 2.6 m/s impact velocity applied to T1.

The neck model was instrumented to output the sagittal angles at the intervertebral joints as well as the head angle and pitch rotations during the simulations. The overall kinematics of the head as well as the segmental angulations were compared with the test data to evaluate the biofidelity of the passive response of the head model.

84

3.3 Results

The neck model with acceleration boundary condition at the T1 was simulated for 110 ms, tracking the motion of the CoM of the skull and the vertebrae, as well as the joint angles. Figure 3-8 shows the response of the neck model under the 2.6 m/s acceleration pulse.

T = 0 ms   *T = 30 ms*

T = 60 ms   *T = 100 ms*
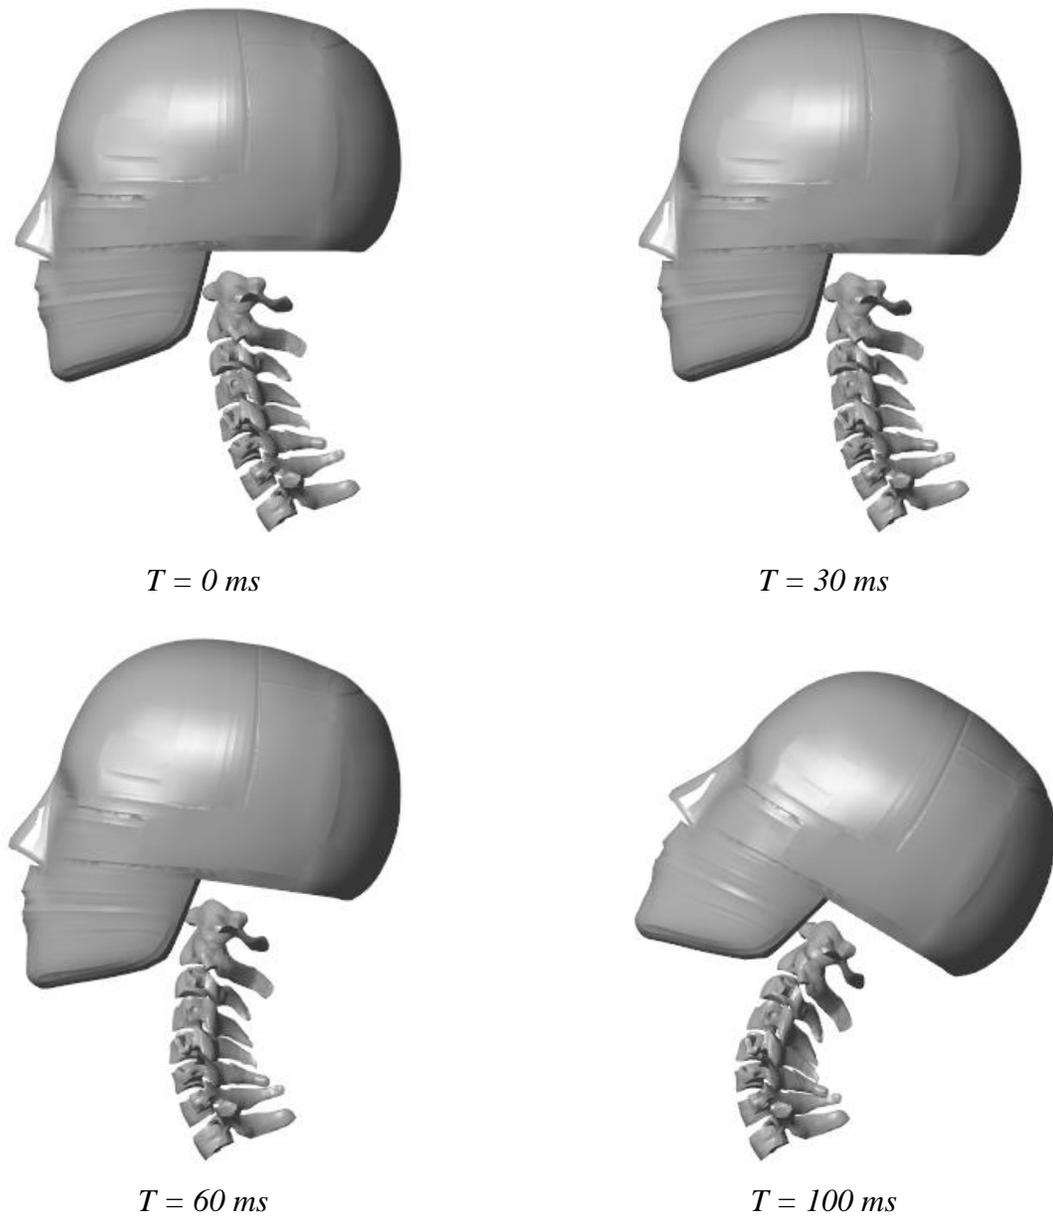
Figure 3-8: Behavior of the neck model during 2.6 m/s rear impact simulation at different points of time.

Similar to the tests, the spine in the simulation developed a S curvature – extension in the lower cervical joints followed by flexion in the upper intervertebral segments (Figure 3-8 T = 60 ms). The S pattern was followed by gross extension of the spinal column (Figure 3-8 T = 100 ms). At no point of the simulation, there was any lateral motion of any of the structure in the model.

3.3.1 Validation of the head kinematics

The kinematics of the head was measured in the simulation and the kinematics data of the head was compared with the corridors developed in the tests. The head trajectory measured from T1 has been plotted along with the test corridors in Figure 3-9.

(a)   (b)
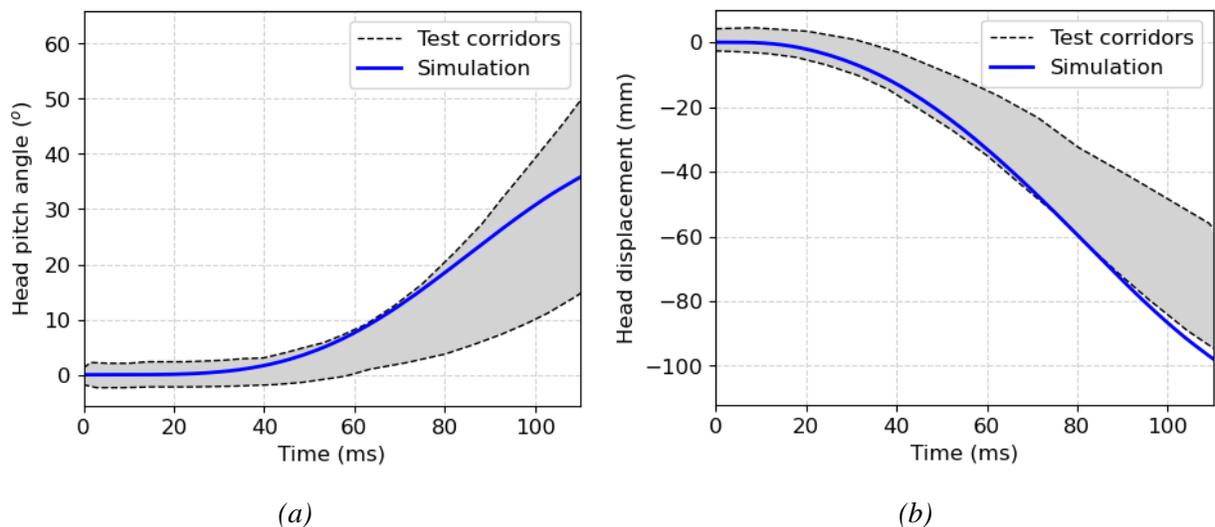
Figure 3-9: Head kinematics measured relative to T1 (a) Head angle in the sagittal plane (b) Head posterior displacement.

In general, the head kinematics time history in the simulation agreed well with the experimental response. The head pitch angle was within the test corridors as demonstrated in Figure 3-9 a. The head displacement also displayed realistic response compared to the experiment data (Figure 3-9 b). Despite the simulation curve being close to the lower bound of the test corridor, no major differences can be seen between the experiment and model response with regards to the posterior displacement of the head relative to the T1.

86

3.3.2 Validation of the intervertebral segment motion

The relative angle between the adjacent vertebrae was tracked in the simulation and compared with the test corridors in Figure 3-10. The overall trajectory of the sagittal angulation of the intervertebral angle were found to be within the test corridors.
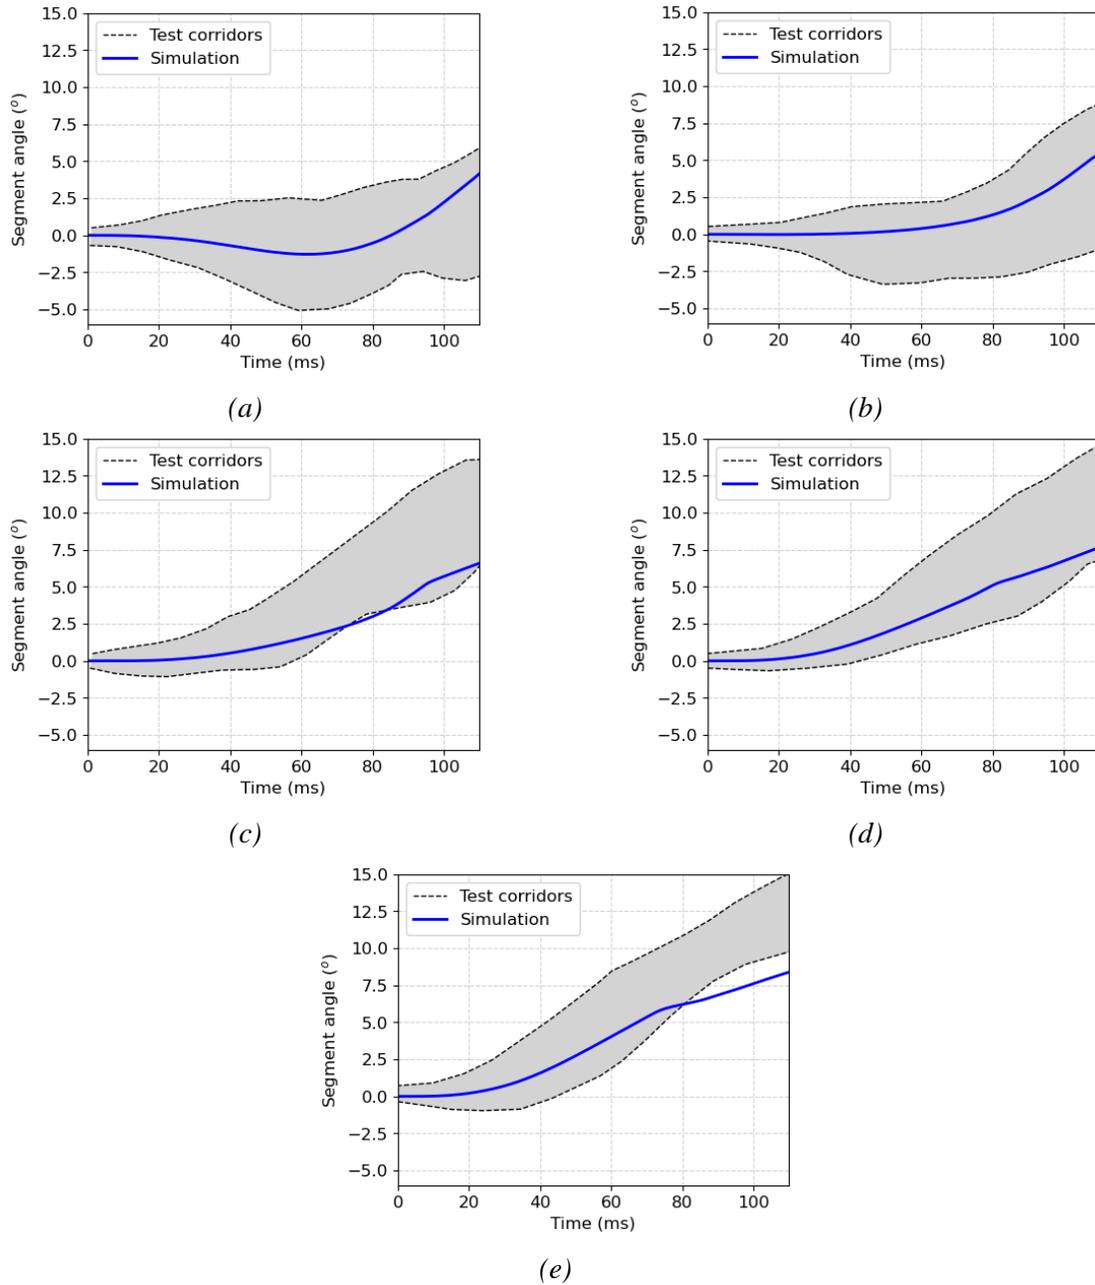
Figure 3-10: Intervertebral joint angle in the sagittal plane (a) C2-C3 angle (b) C3-C4 angle (c) C4-C5 angle (d) C5-C6 angle (e) C6-C7 angle.

The nature of the angle-time history of the intervertebral joints also followed the trends of the experiments. The C4-C5, C5-C6, and C6-C7 joints were all in extension throughout the simulation. The C2-C3 joint was in flexion for close to 90 ms which was followed by extension of the joint, causing the spine of the neck model to form an S curvature. Although the test corridor showed that in some specimens, C3-C4 segment was in flexion for a portion of the test, the simulation showed no flexion of the C3-C4 joint. In the later part of the simulation, the predicted C6-C7 joint angle was marginally outside of the lower test bound. However, the angle-time responses of the vertebral joints were within the test corridors for most of the simulation time.

3.3 Discussion

This chapter summarizes the development of the neck model which is a MB model of the head and neck region of a 50^{th} percentile male. The neck model incorporates rigid vertebrae from T1 to C1, and the skull. The rigid bony structures are attached to each other with joints 6 DoF joints. The muscles of the neck are modeled with Hill-type force characterization, with the line of action of the forces following the curvature of the spine during movements. The main intention behind development of the neck model was to integrate it with the RLMAC framework to generate head motions under various loading environments.

The neck model does not include soft tissues like the ligaments, facet joints, cartilage, and the intervertebral discs. The stiffness of these components has been accounted for while assigning the stiffness-damping parameters at the intervertebral joints. The joint stiffnesses were approximated from the values that have been reported in the literature. The Y and Z translation force-displacement relationship could be approximated with a linear stiffness value, however the X translation stiffness and the rotational stiffness in all DoFs were modeled with a bilinear relationship to effectively capture the non-linear force-displacement relationship (Figure 3-2). The

88

bilinear approximation of the stiffness curves allows for faster simulation times and symmetrical responses in absence of any lateral loads or moments.

The 6 DoF joints have been modeled between all the rigid structures in the model except C1-C2 which are combined into one part as not much relative movements happen between them (Dibb et al., 2013; Nightingale et al., 2002). The seven intervertebral joints effectively provide 42 DoFs to the spine. Some of the recent neck models developed for control studies have not considered any motion between the C2 vertebra and the skull (de Bruijn et al., 2016; Happee et al., 2017b) or have neglected translational motion at the joints (Mortensen et al., 2018). However, incorporating the joint compliance is necessary to accurately model the spine behavior.

The muscles have been modeled using a modified Hill-type muscle which considers a damping element for calculating the muscle forces (Haeufle et al., 2014). The Hill-type muscle was found to improve the joint kinematics by considering the eccentric force-velocity relationship and damping which reduce disturbances and vibrations at a stable joint position (Martynenko et al., 2018; Wochner et al., 2019). Muscles with wide physiological cross-sections were divided into multiple strands to improve the physiological representation and distribute the forces at the rigid vertebrae more accurately. Muscle wrapping around the bones were considered by splitting the strands and attaching the intermediate strands to the nearest vertebrae. For the force calculation using the Hill's muscle equations, the summation of change in length of each segment in a muscle strand is considered.

The passive responses of the neck model were validated under a rear impact loading scenario. The head kinematics as well as the joint angles were found to be within the test corridors. Under the impact scenario, the response of the model depends on the inertia of the rigid structures, the joint stiffness as well as the passive forces developed in the muscles. During the rear impact simulations

considerable passive forces developed in the sternocleidomastoid, the omohyoid, and the sternohyoid muscles, which is expected as the head angle respect to T1 was in extension. There was no lateral displacement of any structure during the simulation, which demonstrates that the neck model with non-linear joints and the muscles is symmetrical about sagittal plane.

The major limitation of the neck model is that the model does not consider the contacts between the different muscles included in the model. Contacts between the neighboring vertebrae are not explicitly modeled but the interaction between the cartilages and facets have been incorporated as the increase in stiffness in the intervertebral joints. The model responds sufficiently well under the validation case (Figure 3-9 and Figure 3-10), implying the simplifications do not affect the overall functionality and utility.

3.5 Conclusions

The neck model developed in this chapter contains the required stiffness and muscle properties for utilization in the control study using reinforcement learning. The developed neck model is representative of the cervical spine of a 50^{th} percentile male and capable of generating biofidelic responses under impact loads. The fast-running passive model will be integrated with reinforcement learning muscle controller (RLMAC) in the following chapters to simulate head kinematics under voluntary and impact cases.

Chapter 4 – Neck muscle control for postural stability and voluntary kinematics

The human head-neck region is a highly redundant control structure with several joints along the cervical spine and the skull. The head is stabilized at its upright position under gravity by the cervical muscles, which generate active forces to counteract the effects of gravity. The neck muscles also stabilize the head in the presence of external dynamic perturbations and are responsible for carrying out the desired voluntary motions of the head and cervical spine. The CNS generates the coordination of the muscle forces required for the desired orientation of the spine and head based on the external loading environments and sensory and proprioceptive feedback.

The current chapter outlines the integration of reinforcement learned muscle actuation controller (RLMAC) with the neck model developed in chapter 3 to simulate the control of neck muscles in humans. The RLMAC was trained to actuate the muscles in the neck model individually to perform voluntary head motions in the sagittal plane. The head motions performed include stabilization under gravity and goal-directed extension and flexion. The components of the reward function for the control objective were also evaluated and confirmed, aiming to eventually extend the methodology of the head control to all DoFs.

The RLMAC post-training could move the head and stabilize it at the target position. The muscle activation patterns developed by the controller for the head movements have been measured during the simulation and presented in the chapter. The ability of the controller to respond to novel target signals which were not included in the training process was also evaluated. Results from this chapter will provide the required information for the development and integration of neck muscle controllers into the neck model for omnidirectional control.

91

4.1 Introduction

Maintaining a stable head position requires constant activation of the neck muscles by the CNS (Crowninshield and Brand, 1981). Continuous control of the neck musculature happens due to various sensory and reflexive loops present in the nervous system (Bove et al., 2009; Keshner, 2009; Keshner et al., 1989). The CNS modulates between the different feedback control loops to cause voluntary motion and maintain head stability under external perturbations (McCrea et al., 1999). The neck muscles may also co-contract to increase the spinal stiffness as a reaction to changes in head inertia like wearing helmets (Bowman et al., 1981; Kumar et al., 2000). The neck muscles have a complex orientation about the spine and as such, it is difficult to associate the muscles to the corresponding head motion direction as a single muscle may be responsible for motion along different directions at different levels of activations (Lee et al., 1990; Vasavada et al., 2002). Fice et al. (2018) found that electrically activating muscles produced moment directions inconsistent with that produced during voluntary isometric contractions.

Active neural control of the head is possible due to the synchronized actuation of the neck muscles by the CNS; however, the phenomenon is not well understood for modeling purposes. Apart from the general redundancy of the system, the non-linearity of the structure and various delays associated with the muscle force generation add to the complexity of modeling the neuro-muscular control. Some previous computational studies used open-loop control for stiffening the neck under dynamic external loads (Brolin et al., 2005; de Jager et al., 1996; van Lopik and Acar, 2004). Chancey et al. (2003) performed an optimization study for the development of active muscle forces in a computational model of the head and neck of an average male to stabilize it under gravity. A similar optimization approach was extended to the anthropometries of six and ten-year-old children by Dibb et al. (2013). Peng et al. (1996, 1999) developed a lumped parameter model to

incorporate neural reflex loops through feedforward and feedback control for head kinematics. Cappon et al. (2007) used PID controllers with a MB human model to control the moment at spinal joints for roll-over automotive impacts. Brolin et al., (2015) used the moment control methodology with a MB model of a 6 year old child. Initially, an open-loop controller was used to stabilize the spine under gravity. PID controllers were used to stiffen the spine by generating torque to counter any external disturbances.

Fraga et al. (2009) used three PID controllers to control head kinematics for pitch and yaw in motorcycle riders. The controllers were tuned to maintain the head's vertical stability and found that the control approach could maintain the target head posture in simplified loading scenarios. Nemirovsky and van Rooij (2010) used three PID controllers for the roll, pitch, and yaw motions of the head with respect to T1. The control strategy involved grouping of muscles in a previously developed model of the head-neck complex (van der Horst et al., 1997) and activating each group with an individual PID to obtain the pure rotations of the head. The control model was validated in extension-flexion DoF at different levels of muscle co-contractions to simulate different levels of bracing by the occupants.

Östh et al. (2015) implemented PID control for head stability along with control of the upper extremities and shoulder to study restraints under autonomous braking conditions. Two separate controllers were used for the head and neck to maintain a vertical position relative to the vehicle coordinates. Feller et al. (2016) implemented a muscle strain-based controller for GHBMC male 50[th] percentile HBM. The response of the control sub-routine was validated with volunteer head fall tests in supine position under gravity (Feller et al., 2016), and the final head displacements were found to agree with the test data.

Happee et al. (2017) implemented a feedback control mechanism in a MB model of a head-neck complex with a deformable spine (de Bruijn et al., 2016) to stabilize the head in the sagittal plane. In the study, the authors found that while the head kinematics feedback alone can stabilize the head, it may result in a non-biofidelic spinal shape. To prevent the buckling of the neck, Happee et al. considered the feedback of muscle length and velocity in the controller. A similar approach was employed by Zheng et al. (2021) for postural control using an MB model of the head and neck with a detailed implementation of muscles (Mortensen et al., 2018). After the control loop was implemented, the model could maintain the stability of the head under gravity in the neutral sitting posture and when placed in prone and supine posture. Apart from maintaining stability, the control model could also respond to impact loads.

Correia et al. (2021) implemented the kinematics and muscle stretch-based feedback control in the GHBMC v5.1 50[th] percentile male model. The reflex responses by the controller could successfully simulate impact conditions at different severities. Putra et al. (2019) individually implemented the closed-loop feedback for head kinematics and muscle lengths in a FE model of 50[th] percentile female. It was observed in the study that the head and neck response to rear impact loading is different for the two control schemes. While the head kinematics controller could output head and neck motion similar to tests, buckling of the cervical spine was observed. The muscle length controller could not accurately capture the neck kinematics but reduced cervical spine buckling was predicted by the model. Ólafsdóttir et al. (2019) modified the THUMS v3.0 50[th] percentile male model with updated skin and muscle representation to simulate neck stability under omnidirectional loading. The control model incorporated feedback on the head kinematics and muscle spindle length. The active THUMS neck was simulated under gravity in the transverse plane in five different directions. When compared with the passive model, the active model was

found to reduce head translation and rotation for all the applied loading directions. Table 4-1 summarizes the various control strategies implemented for the cervical spine.

Table 4-1: Summary of neck control studies

Study	Model	Actuators	Control type	Control signal
de Jager et al. (1996)	Multibody (MADYMO)	1-D Hill-type muscles	Open-loop	NA
van der Horst et al. (1997)	Multibody (MADYMO)	1-D Hill-type muscles	Open-loop	NA
Brolin et al. (2005)	FE (LS Dyna)	1-D (spring) Hill-type muscle	Open-loop	NA
Brolin et al. (2008)	FE (LS Dyna)	1-D (spring) Hill-type muscle	Open-loop	NA
Panzer et al. (2011)	FE (LS Dyna)	1-D (beam) Hill-type muscle	Open-loop	NA
Chancey et al. (2003)	FE (LS Dyna)	1-D (spring) muscle	Optimization	Head motion less than 5° and 10 mm under relaxed and tensed activation states
Dibb et al. (2013)	FE (LS Dyna)	1-D (spring) muscle	Optimization	Head CG displacement less than 0.3° and 0.5 mm under relaxed and tensed activation states
Correia et al. (2020)	FE (LS Dyna)	1-D Hill-type active muscles with solid passive elements	Optimization	Head CG kinematics
Nemirovsky and Van Rooij (2010)	Multibody (MADYMO)	1-D Hill-type muscles	Closed-loop	Head roll, pitch, and yaw
van Rooij (2011)	Multibody (MADYMO)	Joint torques/ 1-D Hill-type muscles	Closed-loop	Thoracic vertebrae angles/ Neck pitch angle
Östh et al. (2012)	FE (LS Dyna)	1-D Hill-type muscles	Closed-loop	Head and neck pitch angle
Östh et al. (2015)	FE (LS Dyna)	1-D Hill-type muscles	Closed-loop	Head and neck pitch angle
Happee et al. (2017)	Multibody (MADYMO)	1-D Hill-type muscles	Closed-loop	Head angle, angular velocity, linear acceleration, muscle length
Ólafsdóttir et al. (2019)	FE (LS Dyna)	1-D Hill-type muscles	Closed-loop	Head angle, angular velocity, muscle length
Putra et al. (2019)	FE (LS Dyna)	1-D Hill-type muscles	Closed-loop	Head angle, angular velocity, muscle length

Study	Model	Actuators	Control type	Control signal
Putra et al., (2021)	FE (LS Dyna)	1-D Hill-type muscles	Closed-loop	Head angle, angular velocity
Correia et al. (2021)	FE (LS Dyna)	1-D Hill-type active muscles with solid passive elements	Closed-loop	Head angle, angular velocity, muscle length
Zheng et al. (2021)	Multibody (OpenSim)	1-D Hill-type muscles	Closed-loop	Head angle, angular velocity, linear acceleration, muscle length

Most of these studies have limited the development of the control architecture for postural control in the sagittal plane. Nemirovsky and Van Rooij (2010) developed a controller for rotation along all three anatomical planes, however, the model was only validated for sagittal motion. Other models which have considered omnidirectional control of the spine have grouped muscles according to response in different directions (Correia et al., 2021; Ólafsdóttir et al., 2019). The control parameters in the feedback models also need to be adjusted with external loads which are not suitable for dynamic events as loads acting on the model may change during the course of the simulation (Happee et al., 2017; Zheng et al., 2021).

In the previous neck control studies, control models were developed to maintain the neutral posture of the spine under external loads. To the best of my knowledge, only one optimization study (Silvestros et al., 2021) has tried to perform voluntary extension-flexion and lateral bending with a multibody head and neck model. In the study, two different optimization schemes were used for extension and flexion with two different muscle groupings. Grouping muscles may be suitable for motion about isolated axes, but for a model to perform omnidirectional voluntary motion, the muscles need to be activated individually as a single muscle may cause motion about two planes (Keshner et al., 1989).

In chapter 2, we demonstrated that reinforcement learning (RL) can generate voluntary kinematics by individually activating the muscles in a simplified arm model. The applicability of the trained RLMAC was also evaluated under static and dynamic forces. In the current chapter, the RL muscle activation controller (RLMAC) will be extended to maintain stability and synthesize voluntary kinematics in the neck model.

The nature of control required for the neck model, however, poses additional challenges. In the arm model, a revolute joint was used for the extension-flexion motion, neglecting any translation or supination and pronation motion at the elbow. The physiology and anatomy of the elbow muscles make it easier for the RLMAC to arrive at a functional activation scheme. The neck model however has seven 6-DoF joints sequentially stacked on top of each other. Moreover, for the neck model, both the state parameters and actions will be more in number due to the omnidirectional feedback and numerous neck muscles incorporated into the model. The complex nature of the neck model will make the RL controller computationally expensive as a large action space needs to be explored before forming an effective activation scheme.

In this chapter, the RLMAC control architecture will be explored for implementation in the neck model. As RLMAC will require high computational effort, the neck model was made symmetrical by mirroring the muscle activations about the sagittal plane for the initial study. The symmetrical model will allow for the fine-tuning of control parameters as well as the reward function before the RLMAC is eventually implemented for all-DOF control of the head and neck. RLMAC will be trained to maintain the stability of the spine under gravity and synthesize goal-directed voluntary extension-flexion of the neck. The activation pattern for the neck muscles generated for the voluntary motion will also be analyzed and compared with the direction preferences of the neck muscles that have been observed in the literature.

97

4.2 Methodology

This section details the implementation of a reinforcement learning controller (RLMAC) to the neck model. The neck model consists of a rigid representation of the vertebrae and the skull. 6 DoF joints were inserted between adjacent vertebrae to allow limited relative movements between the structures. The neck model contains 46 neck muscles represented as Hill-type muscles. The RLMAC aims to develop an activation pattern for the muscles to generate the desired kinematics of the head. In this chapter, detailed studies are performed to develop the control parameters for the neck model. To tune the state and reward function for the RLMAC, simulations are performed by considering symmetry along the sagittal plane. Simulations with different control objectives were performed to evaluate the effects of the control parameters on the trained response of the control model.

4.2.1 Muscle control framework

The head is stabilized under gravity by neck muscles which are activated by the CNS. To develop a biofidelic muscle response to an external force field in the neck model, a gradient-based actor-critic network is used for the RLMAC which is then coupled with the MB model. The RLMAC can read the force, energy, and kinematics data from the MB model. Based on the input data, the RLMAC outputs the muscle neural stimulation (u_t) for each muscle. As symmetry is being considered along the sagittal plane, equal stimulation was developed for the corresponding muscles on either side of the mid-sagittal plane. The neural stimulation was transformed into muscle activation (a_t) using the activation dynamics equation by Zajac (1989) (Equation 2-2). The activation was fed to the muscles in the neck model along with the length and velocity magnitudes that were used to generate the muscle active forces. Identical activations were applied to the strands belonging to the same muscle, which limited the controller action to 23 stimulations.

98
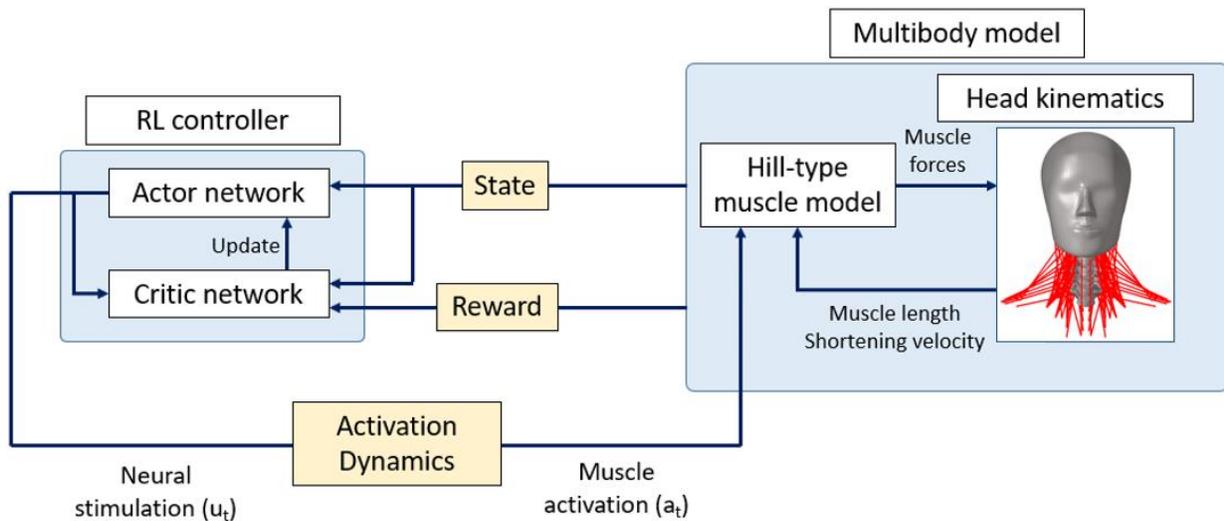
Figure 4-1: RLMAC framework integrated with UVa neck model.

Figure 4-1 displays the RLMAC framework for control of head motion. A twin-delayed deep deterministic policy gradient (TD3) is used as the agent for RLMAC in this case (Fujimoto et al., 2018). The TD3 agent is a development over the DDPG agent that was used in chapter 2 for the arm rotation control. It was found that the DDPG agent overestimates the long-term efficiency (Q-value) of the actions for a given state, a problem which was decreased in the TD3 architecture (Fujimoto et al., 2018). In an over-actuated system such as the head-neck region, overestimation of the Q-value may lead to the learning of sub-optimal policies. TD3 agent implemented in the chapter tries to reduce value overestimation by using two critic networks. Each critic network individually evaluates and quantifies the long-term efficiency of the actions, and the lower Q-value is used for updating the agent networks. TD3 agent was found to outperform the DDPG agent in control tasks that require continuous outputs (Fujimoto et al., 2018).

The actor network in the RLMAC consists of an input layer, an output layer, and a hidden layer. Inputs to the hidden layer and the final layer were activated by the rectified linear unit (ReLU)

function (Hara et al., 2015). The output of the actor network was bound between 0 and 1 representing the neural stimulations using a tanh and a scaling function. The state values constitute the inputs of the actor network. The final layer has 23 nodes, one for stimulation of each muscle. Both the critic networks have identical architectures with one hidden layer between the input and the final layer. The state parameters are read by the input layer and processed using the ReLU function before passing to the hidden layer while the actions are directly read by the hidden layer (Lillicrap et al., 2019). Output from the hidden layer is again activated by ReLU before being read by the final layer. The final layer of the critic outputs the Q-value and the lesser of the two Q-values is used for updating the agent's policy.

The policy is updated every two timesteps of the simulation, i.e., the policy of the agent is updated once for every two outputs of the Q-value by the critic networks. Ornstein–Uhlenbeck (OU) process was used to add noise with a standard deviation of 0.09 for adequate exploration of the action space (La Barbera et al., 2022; Lillicrap et al., 2019). The schematic of the TD3 agent used in the study is shown in Figure 4-2.
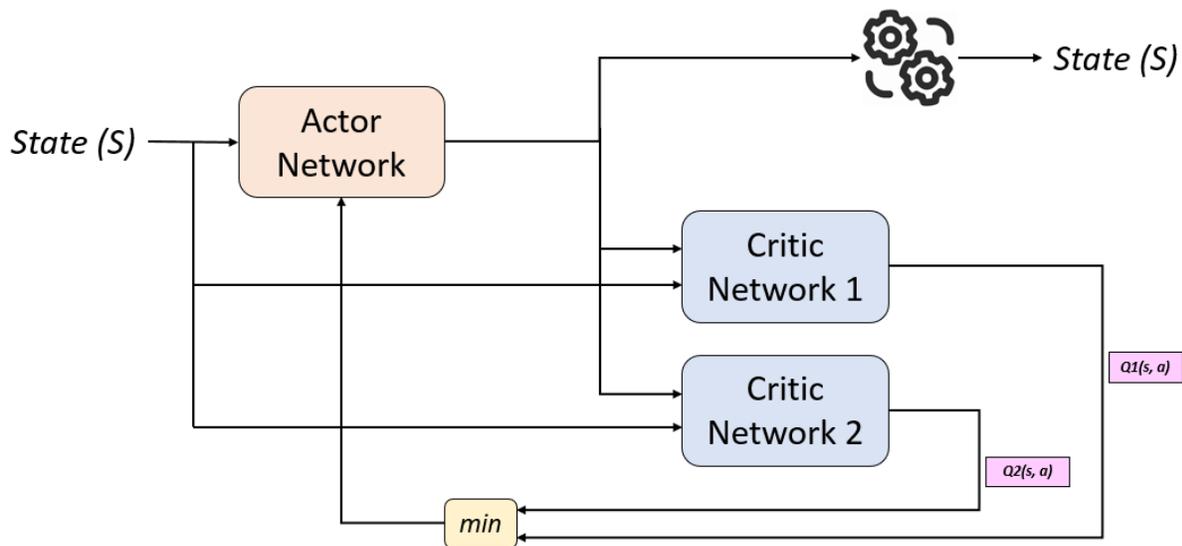
Figure 4-2: Schematic of the TD3 agent.

4.2.2 State and reward

The neck and head are stabilized by the voluntary contraction of the muscles as well as the various reflexive feedbacks and involuntary co-contraction. The central nervous system controls the contribution of reflexes to activate the muscles in synergy to stabilize the head. The postural reflexes are attributed to maintaining the orientation of the spine under gravity and external perturbations. Two major postural reflex loops are hypothesized to contribute towards head stability – the vestibulocollic reflex (VCR) and cervicocollic reflex (CCR) (Cullen, 2012; Goldberg and Cullen, 2011; Keshner, 2009).

The vestibular system consists of the semicircular canals and otolith organs which are located in the inner ear. The vestibular system senses the rotation and translation kinematics of the head and generates compensatory activations to stabilize the head. The cervicocollic reflex stabilized the spine in relation to the trunk and responds to the neck proprioceptive signals. The states for the RLMAC were defined considering both the VCR and CCR inputs. The agent receives input regarding the translation and rotation kinematics of the head for the vestibular feedback. In the symmetrical model, the kinematics were limited to the sagittal plane, i.e., displacement and velocity along X and Z, and rotation and angular velocity along Y. Due to the redundant nature of the head-neck region, kinematics-only feedback may not yield the desired result as there is no feedback regarding the nature of the spine and muscles to the agent, and this may result in stabilization of the head but a non-optimal and non-biofidelic position of the spine. Thus, CCR inputs are also included in the state. The agent was trained under two different state parameters concerning the CCR. In the first scenario, the length and velocity of the muscles were included in the states corresponding to the feedback by the muscle spindles. For the second condition, the joint kinematics feedback was added to the state corresponding to the combination of the muscle

101

afferent feedback and joint proprioceptors in the cervical spine. In both the CCR scenarios, muscle activations were included in the state for information regarding the muscle tension which is sensed by the golgi tendons.

In RL, the reward function quantifies the efficacy of an action relative to the overall control objective. In this chapter, the intention for the RLMAC is to perform voluntary head motions and maintain the stability of the head at the final position in presence of external forces. Thus, the reward function was formulated based on the stated objective. The agent was awarded for moving the head towards the target head angle and reducing the head angle error. The reward function can be described using the voluntary kinematics and contributions from VCR and CCR (Equation 4-1).

Reward function = Voluntary kinematics + VCR parameters + CCR parameters   Equation 4-1

The voluntary kinematics factor in the reward function is correlated to the error between the current angle and the target angle of the head (ε). The VCR parameters are included to stabilize the head kinematics in space, whereas the CCR parameters stabilize the spine and ensure a biofidelic movement of the head relative to the trunk.

For the voluntary kinematics along the sagittal plane, the agent received a reward proportional to the square of sine of twice the angular error $(1 - \sin(2\varepsilon))^2$. Higher order was considered rather than a linear relationship to minimize the error when the head is near the target angle. To maintain the stability of the head at the target position, and to attenuate any vibration or disturbances to the skull at the target position, the agent was penalized proportional to the linear velocities (V_x and V_z) and angular velocity (ω_y). The CCR parameters in the reward function were varied for the two state parameters considered for training. In the scenario in which muscle length is considered as a state, the agent was penalized proportional to the muscle energy expenditure for motion. When the

joint motion was accounted for in the state, the agent was penalized for cumulative joint energy expenditure for the motion.

Table 4-2: State and reward considerations for training

	State	Reward
VCR parameters	• Skull CoM translation (S_x and S_z) • Skull CoM linear velocity (V_x and V_z) • Skull CoM angle (θ_y) • Skull CoM angle error (ϵ_y) • Skull CoM rotational velocity (ω_y)	$(1 - \mathrm{Sin}(2\epsilon_y))^2 - \alpha\,(V_x + V_z) - \beta\,\omega_y$ Equation 4-2
CCR scenario 1	• VCR parameters • Muscle length • Muscle velocity • Muscle activation	$(1 - \mathrm{Sin}(2\epsilon_y))^2 - \alpha\,(V_x + V_z) - \beta\,\omega_y - \gamma\,\sum\text{Muscle energy}$ Equation 4-3
CCR scenario 2	• VCR parameters • Joint displacement • Joint velocity • Muscle activation	$(1 - \mathrm{Sin}(2\epsilon_y))^2 - \alpha\,(V_x + V_z) - \beta\,\omega_y - \eta\,\sum\text{Joint energy}$ Equation 4-4

Table 4-2 summarizes the state and reward parameters used for training the agent. α, β, γ, and η are non-negative scalar values that are used to scale the VCR and CCR parameters relative to the voluntary parameter in the reward function. A simple relationship between the components of the reward function is used as it is generally considered that agents can learn complex behavior from simple reward functions by constantly interacting with the environment (Heess et al., 2017). For the purpose of training, the head kinematics is calculated with the T1 CoM as reference.

The muscle energy (E_M) for state parameters is calculated using the metabolic energy costs from Margaria (1968) that states that muscles while shortening are only 25% efficient compared to 120% efficient while lengthening (Equation 4-5).

$$Muscle\ energy\ (E_M) = \begin{cases} \dfrac{F_M\ x\ \Delta l_M}{0.25} & V_{CE} < 0 \\ \dfrac{F_M\ x\ \Delta l_M}{1.2} & V_{CE} > 0 \end{cases} \qquad \text{Equation 4-5}$$

The joint energy (E_j) is calculated as the summation of the product of joint forces (F_j) and joint displacements (s_j) from the neutral position in X and Z translation and Y rotation (Equation 4-6).

$$Joint\ energy\ (E_j) = \sum F_j\ x\ s_j \qquad \text{Equation 4-6}$$

With the rewards described in Equation 4-2, Equation 4-3, and Equation 4-4, the neck model initially was trained to maintain stability (0^o sagittal angle relative to T1) under gravity. In the stabilization simulation, the relative effects of different parameters of the states and the reward function on the overall kinematics of the model were evaluated. Following the stabilization simulation, the RLMAC was trained to perform voluntary goal directed head motions in the sagittal plane. During the training, the target angle was randomly varied for each iteration between 25 degrees in flexion to 35 degrees in extension in increments of 5. The target angle was maintained at 0^o to train the RLMAC to maintain the head stability under gravity for 40% of the iterations. In the remaining iterations, the target angle was varied to train the RLMAC to perform the goal directed motions of the head. Each iteration ran for 500 ms with an agent time-step of 1.0 ms, and the agent was updated every two time-step. The control system was considered converged when the average reward plateaued and the response of the trained agent at the plateau was reasonable.

4.3 Results

At first, the neck model was simulated under gravity without muscle control to determine the kinematics of the head in the passive condition. A baseline activation of 0.005 was applied to each muscle during the passive simulation, and the head angular displacement was measured from the initial head orientation at which the model was developed. The simulation was run for 1 s until the

104

head was stabilized at the final position under gravity. The simulation shows that the passive muscle and joint properties are not enough to maintain the upright posture in the neck model under gravity and require active muscle forces to maintain the neutral head and spine position of 0^o head sagittal angle (Figure 4-3).

<table>
<tr><td>(a)</td><td>(b)</td></tr>
</table>
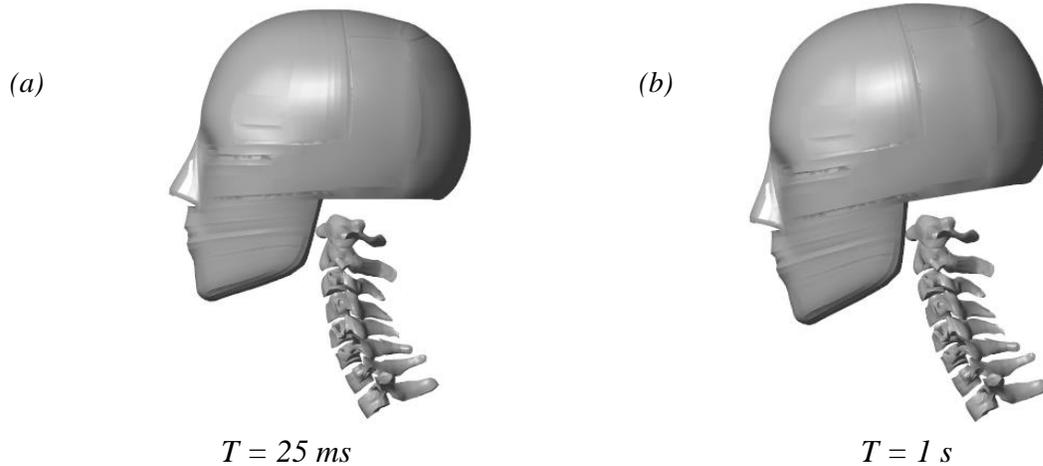
$T = 25 \text{ ms}$ $T = 1 \text{ s}$

Figure 4-3: UVa neck model under gravity (a) At 25 ms the head compressed the spine under gravity (b) Shape of the spine at the end of the simulation.

At the start of the simulation, the head compressed against the spinal column (Figure 4-3a) with the T1 fixed. After the initial compression, both the head and the spine deformed to flexion under the effect of gravity (Figure 4-3b, Figure 4-4a). During the bending of the spine, passive forces developed in the splenius capitis, splenius cervicis, semisplenius capitis, semisplenius cervicis, levator scapula, and longissimus capitis (Figure 4-4b).

(a)   (b)
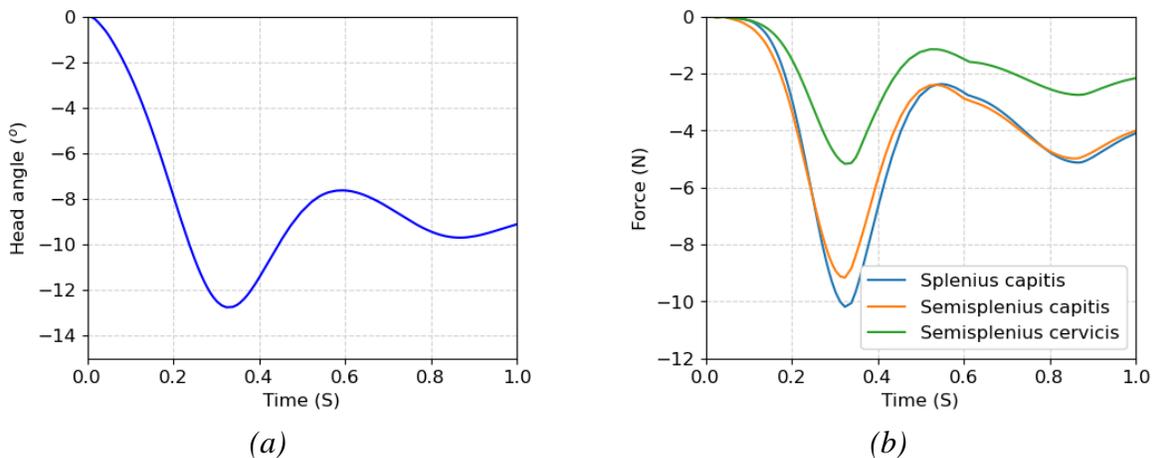
Figure 4-4: (a) Angle time history of Head CoM under gravity (b) Passive forces developed in the muscles during the head motion.

105

4.3.1 Contribution of VCR toward head stabilization

Following the passive simulation under gravity, the neck model was integrated with RLMAC architecture for active muscle control. First, RLMAC was trained only with VCR feedback to gauge the contribution of the vestibular system towards the head stabilization under gravity with state and rewards mentioned in VCR parameters in Table 4-2. The simulation was run for 4600 iterations before convergence.

(a)

(b)

(c)
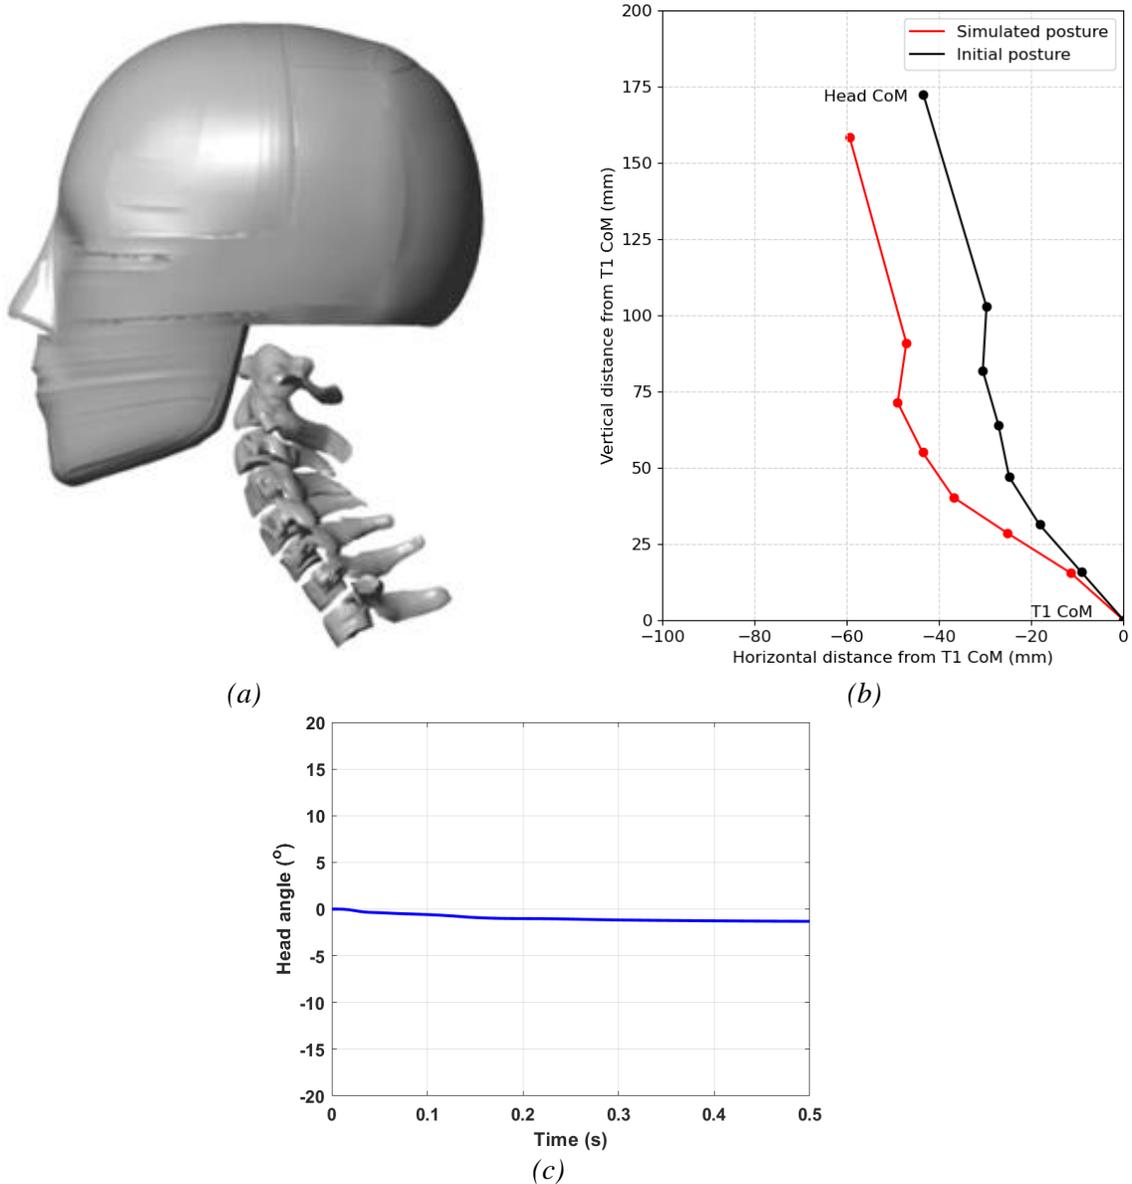
Figure 4-5: Control response with vestibular system (a) The final stabilized configuration of the head and spine (b) Head and vertebrae CoM position for initial and stabilized posture from T1 CoM. (c) Head angle with respect to T1 in the sagittal plane.

The neck model was simulated with the trained agent to obtain the response of the RLMAC with vestibular control. Figure 4-5a shows that although the RLMAC could maintain the desired head angle in presence of an external force field, the posture of the spinal column was non-biofidelic. The location of the vertebrae and the head CoM has been compared for the initial posture and the stabilized posture in Figure 4-5b. Under the presence of gravity, the head and neck were stabilized but the muscle forces developed by the RLMAC caused buckling of the spinal column. The head was stable at the final position with a head sagittal angle of around 1.2° in flexion (Figure 4-5c) which displays the influence of the linear and angular velocity terms in the state and the reward function for head motion control. As no term regarding the shape of the spinal column was included in the state or reward function, buckling of the spinal column with compression of the upper vertebrae was observed.

4.3.2 Contribution of CCR toward head stabilization

VCR alone was unable to elicit the stability of the spinal column even though head stability was achieved (Figure 4-5). To stabilize the cervical spine and the vertebral joints, the cervicocollic response needs to be considered while developing the architecture of RLMAC (Keshner, 2009). In this chapter, two separate pieces of training were performed with different considerations of CCR parameters for the state of the RLMAC. In the first scenario, the muscle spindle length was taken into account for the state and in the second scenario, the joint displacements of the cervical spine were included (Table 4-2) as a state parameter. Under both scenarios, the RLMAC was trained to maintain the neutral posture of the head CoM (0° head angle) in the sagittal plane and the effect of the CCR parameters on the response of the head and the cervical spine was evaluated. Under both CCR scenarios the RLMAC took around 5200 iterations for convergence for stability.

<div style="text-align:center">(a)   (b)</div>
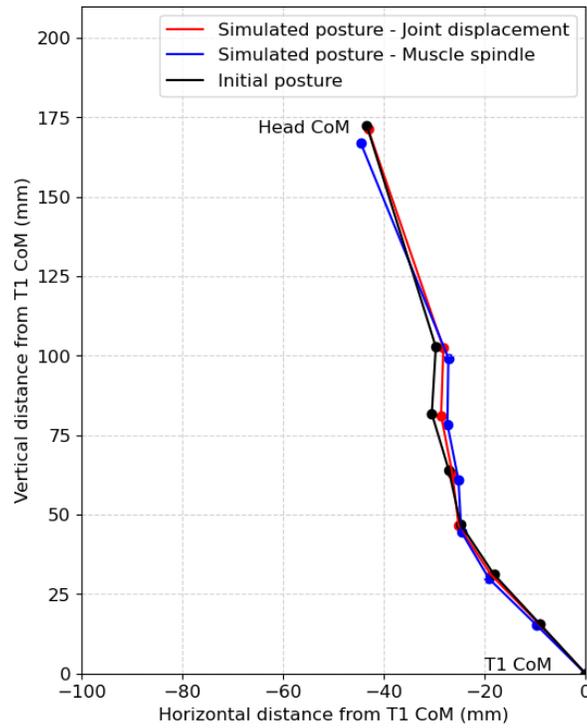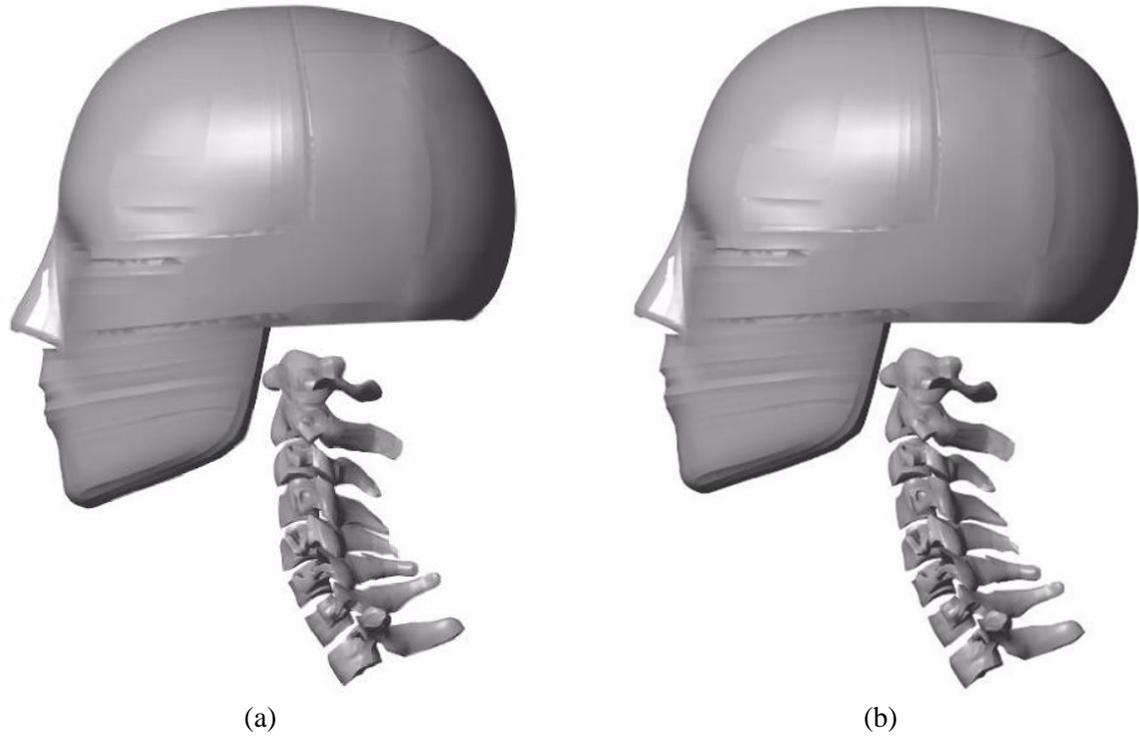
<div style="text-align:center">(c)</div>

Figure 4-6: Final posture of neck model under gravity with CCR consideration at 500 ms of the simulation (a) Muscle spindle length (b) Joint displacement as a state (c) Head and vertebrae CoM position for initial and stabilized posture from T1 CoM.

Figure 4-6 shows the stabilized posture of the head and spinal column under gravity for the two CCR cases. In the simulation of RLMAC trained with CCR scenario 1 (muscle length), the upper vertebrae moved anteriorly towards the T1 reducing the eccentricity, thus reducing the effective moment due to gravity. In this scenario, the head was stabilized by a combination of active muscle forces and passive stiffness of the spinal column, while compressing the upper spinal joints in this process (Figure 4-6a). When the joint displacements were considered for the CCR parameter, the head achieved a stable position maintaining a biofidelic posture of the spinal column (Figure 4-6b).

——— CCR as muscle spindle length   ——— CCR as joint displacement
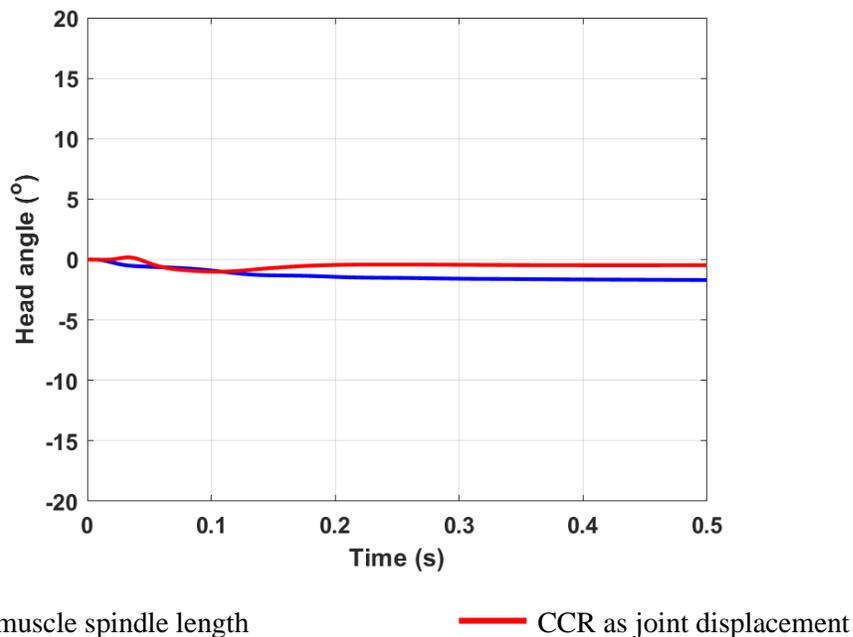
Figure 4-7: Head sagittal angle for the two CCR scenarios.

Figure 4-7 shows the time history of the head sagittal angle for the CCR scenarios including the muscle spindle length and joint displacements. In both the cases, the head was stabilized below 1.5° angle, even though the shape of the spine at the final time-step in both simulation conditions vary. Under both CCR conditions, there was no buckling or collapsing of the spinal column, however, in the CCR scenario 1, compression of the upper spine was observed. In CCR scenario 2, there was no major deformation of the cervical spine or the vertebral joints.
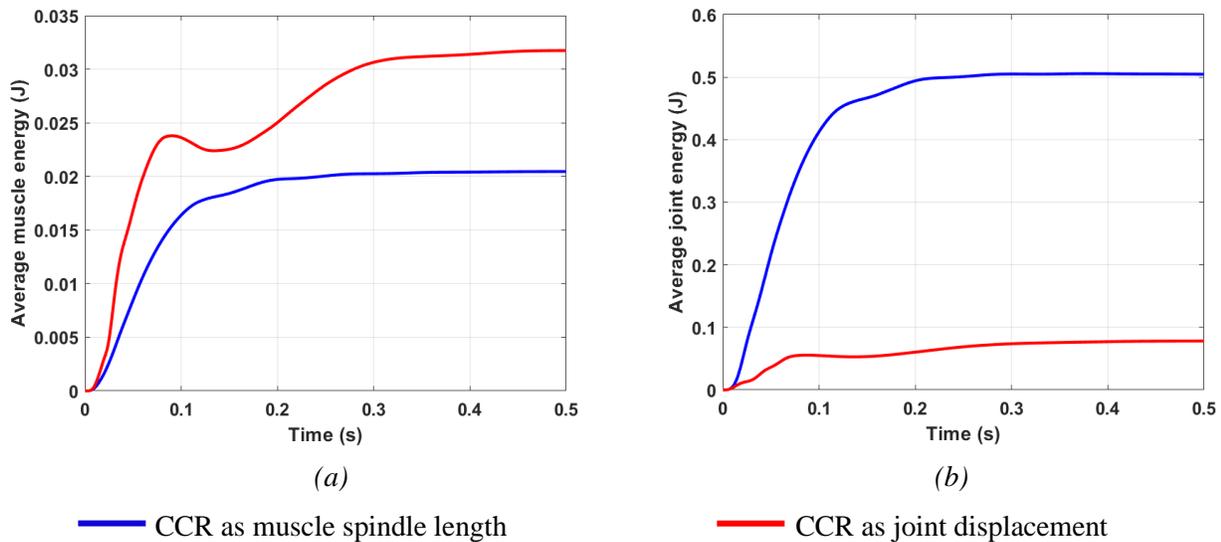
Figure 4-8: Response of the neck model with CCR parameters (a) Average muscle energy (b) Average joint energy.

Figure 4-8 displays the average energy response of the CCR control models while maintaining neck stability. Average muscle energy (Figure 4-8a) is the mean of the energy cost of motion by the muscles (Equation 4-5) over all the muscle strands in the model (86 on each side). The joint energy in Figure 4-8b is averaged for the seven intervertebral joints in the neck model. Due to the nature of the reward function associated with each CCR scenario, the average muscle energy was lower in scenario 1 while the average joint energy was lower in scenario 2.

The joint displacement proprioception as the cervicocollic parameter (CCR scenario 2) achieved a more stable shape of the spine with less compression of the spinal joints at the neutral stance, as such it was used in further control studies to synthesize voluntary kinematics in the sagittal plane. Each iteration with the CCR scenario 2 also took lesser time compared to muscle spindle feedback (CCR scenario 1) as scenario 1 uses muscle length and velocity feedback for 86 individual strands which increase the input layer nodes of both the actor and critic network and as a result, increasing the processing and computation time.

4.3.3 Voluntary head kinematics in the sagittal plane

The CCR scenario 2 with joint displacement feedback (Table 4-2) was used to train the RLMAC for generating voluntary extension and flexion motion of the head. During the training, the target angle was varied to train the RLMAC over a range of possible head motions. The starting position was always the neutral posture of the head, i.e., 0^o head angle. The training was distributed over 40 CPUs in the Rivanna cluster using the parallel computing toolbox. The training was performed in 6500 iterations. Figure 4-9 shows the variation of the average reward over the training period. The cumulative reward over 250 iterations were averaged to obtain the average reward.
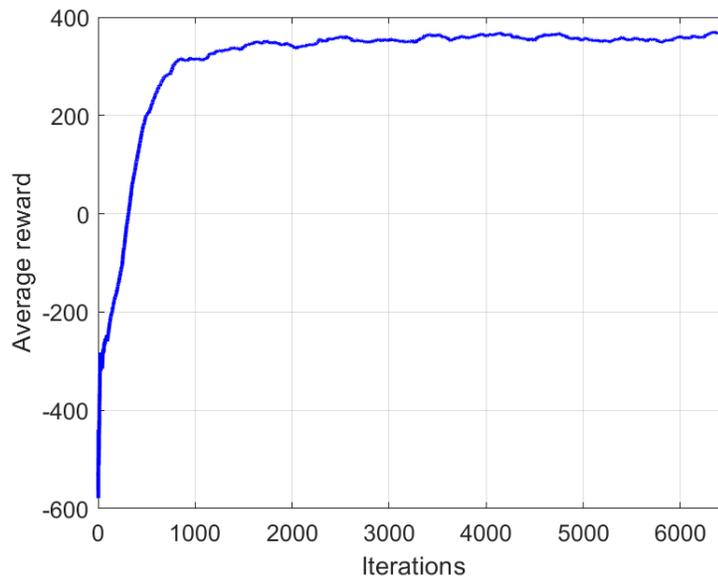
Figure 4-9: Average reward during the training.

The neck model was initially stabilized under gravity for 200 ms before the voluntary goal-directed motion was simulated. The simulation results show that the trained RLMAC could stabilize the head under gravity and generate both extension and flexion motion of the head (Figure 4-10).

111

| (a) | (b) |
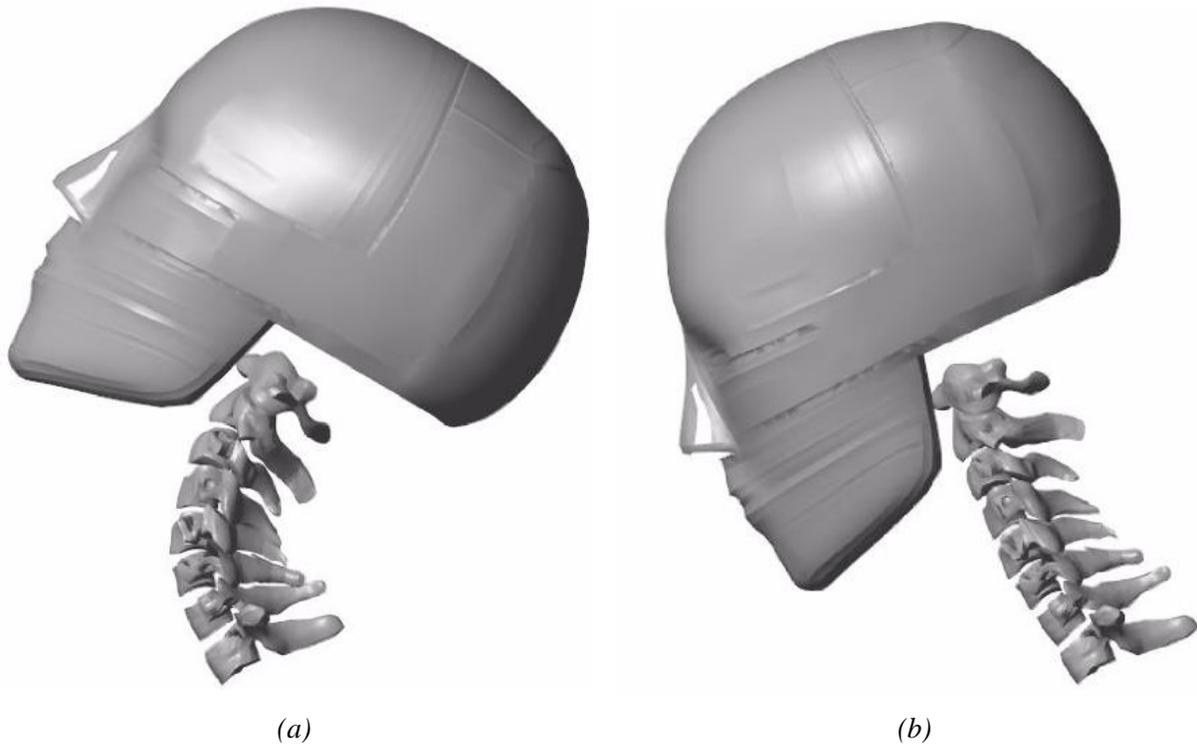
Figure 4-10: Head goal-directed motion along the sagittal plane (a) Extension (b) Flexion.

Figure 4-11 shows the angle time histories of the head during the voluntary rotation of the head. The final head angles are within 1.5° of the target values while simulating both the extension and flexion motion from the neutral position.

| (a) | (b) |

Head rotation angle — · — Target angle
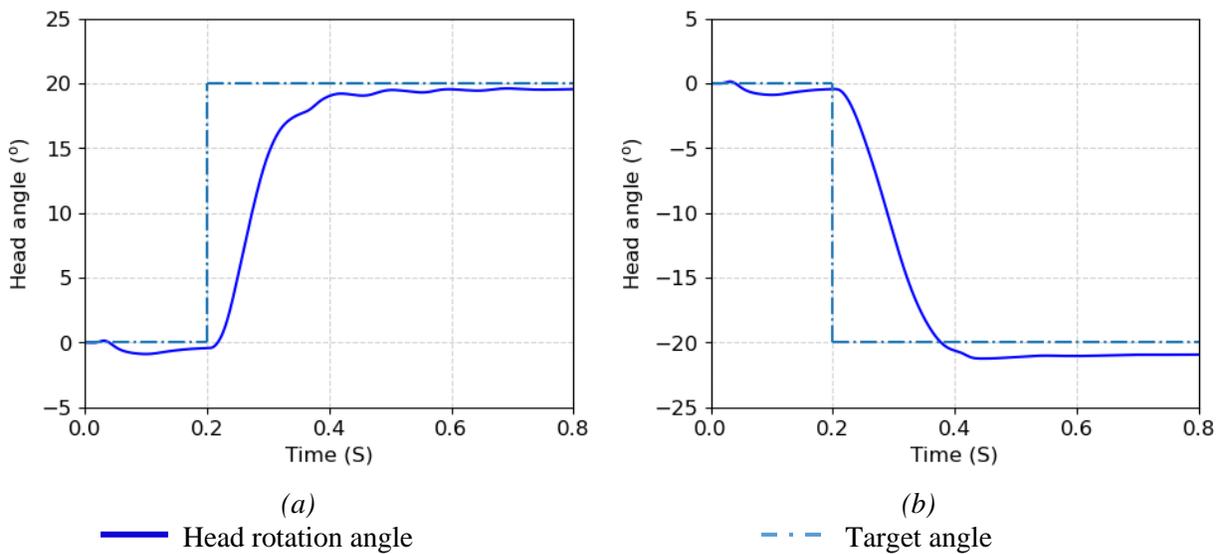
Figure 4-11: Head CoM angle measured during the goal-directed simulation (a) Extension (b) Flexion.

The activations of the neck muscles were measured for the head stabilization and the sagittal rotation cases. All the muscles in the neck model were individually activated without any consideration of predefined strategy or grouping based on the direction preference. The neck muscles were considered symmetrical about the mid-sagittal plane leading to 23 muscle activation outputs from the RLMAC.

Semispinalis cervicis   Sternohyoid
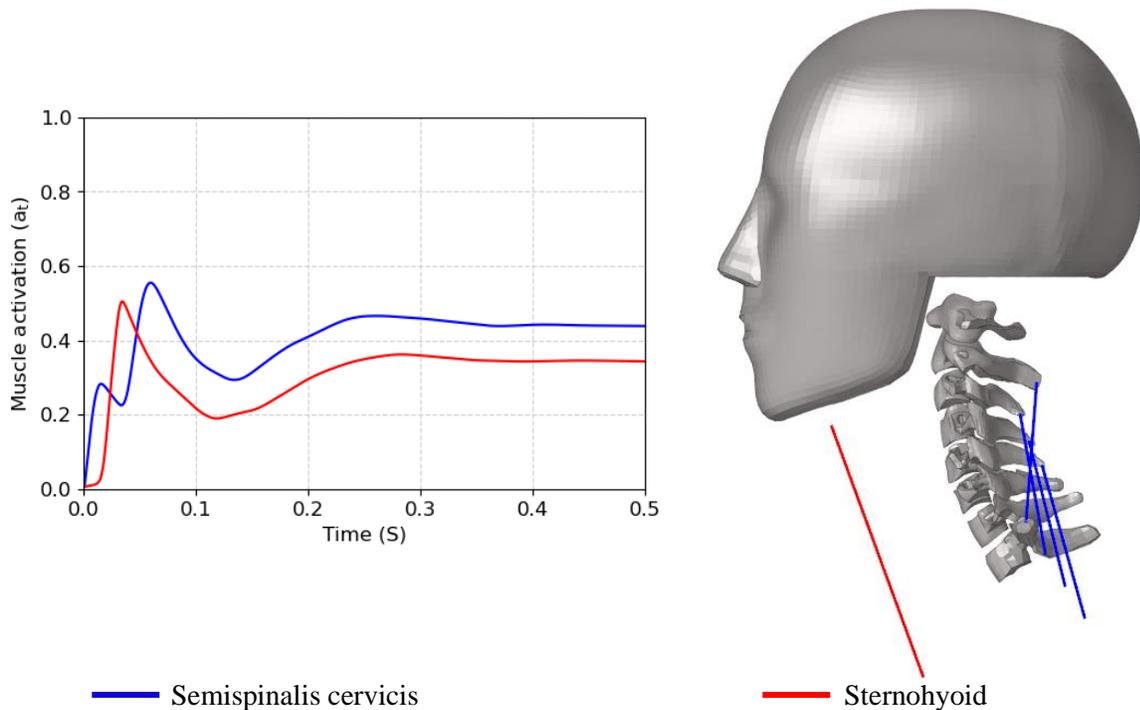
Figure 4-12: Co-contraction of the neck muscles in the stabilization simulation.

During the head stabilization run, the semispinalis cervicis and the sternohyoid muscles co-contract maintaining the stability of the spine and the head (Figure 4-12). The initial increase in the activations is due to the sudden application of gravity at the start of the simulation, requiring the muscles to apply forces. After 200 ms, the activations maintained a constant magnitude after securing the stability of the head and neck. Apart from the two muscles plotted in Figure 4-12, there were also minor activations in splenius capitis, longissimus capitis, omohyoid, and sternocleidomastoid, however, activations in these muscles were less than 0.1.
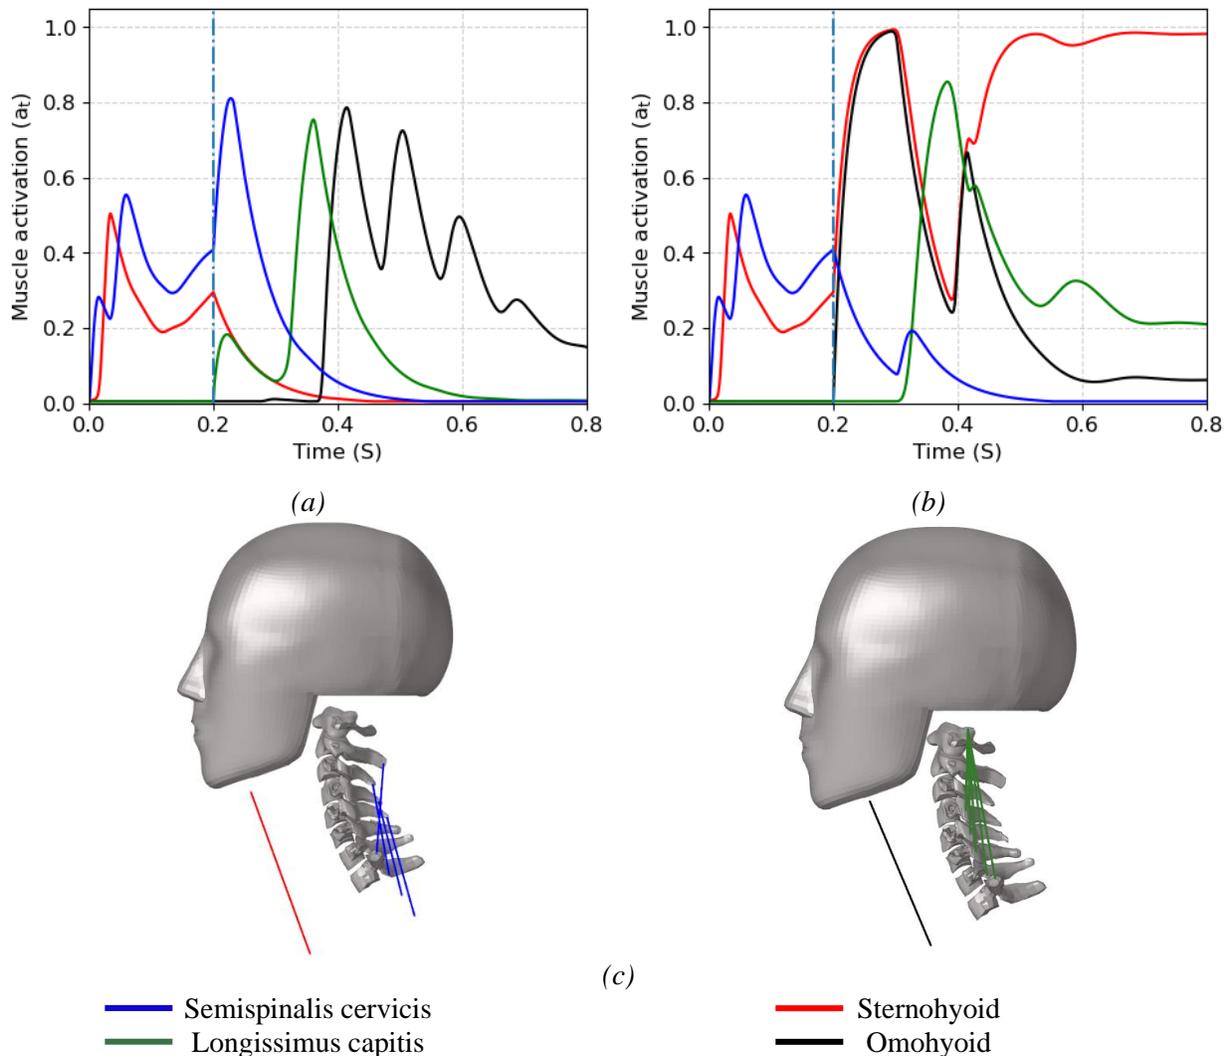
Figure 4-13: Activations of the neck muscles during head rotation (a) 0-20° Extension (b) 0-20° Flexion. The dotted line indicates the time at which the goal-directed motion starts (c) Muscles activated during extension-flexion of the head.

The muscle activations for the head extension-flexion motion are shown in Figure 4-13. The head rotations follow an initial period of stabilization which is indicated by the dotted line in the plots in Figure 4-13. The muscle activations during the goal-directed head motion follow a biphasic or triphasic pattern which is characterized by successive bursts of agonist and antagonist activities to reach a target and subsequently damp the joints at that position (Hannaford and Stark, 1983; Happee, 1992; Marsden et al., 1983). Similar tri-phasic activation patterns were also observed in the arm movement study in chapter 2.

In the extension motion, the activity of all the muscles reduced to the minimum value except the Omohyoid, whose activity level was around 0.2, possibly to support the head at the 20^o extension position (Figure 4-13a). In the flexion motion, the sternohyoid had high activations near the target position. The 20^o flexion angle is beyond the range of the passive model (Figure 4-4a) therefore extra force from the flexor, in this case, the sternohyoid is required to settle the head at the flexion target (Figure 4-13b). The sternocleidomastoid was activated during the flexion and splenius capitis and trapezius were activated during the extension at different points, however, their activations remain low throughout the simulation.

4.3.4 Response to novel targets

To evaluate the ability of the RLMAC to synthesize motions under novel targets, the neck model was first simulated to 20^o flexion and from the flexion position, the target angle was set to 20^o extension. Similarly, in a separate simulation, the head was first simulated to 20^o extension, and from that position, the head was made to undergo a target posture of 20^o flexion.
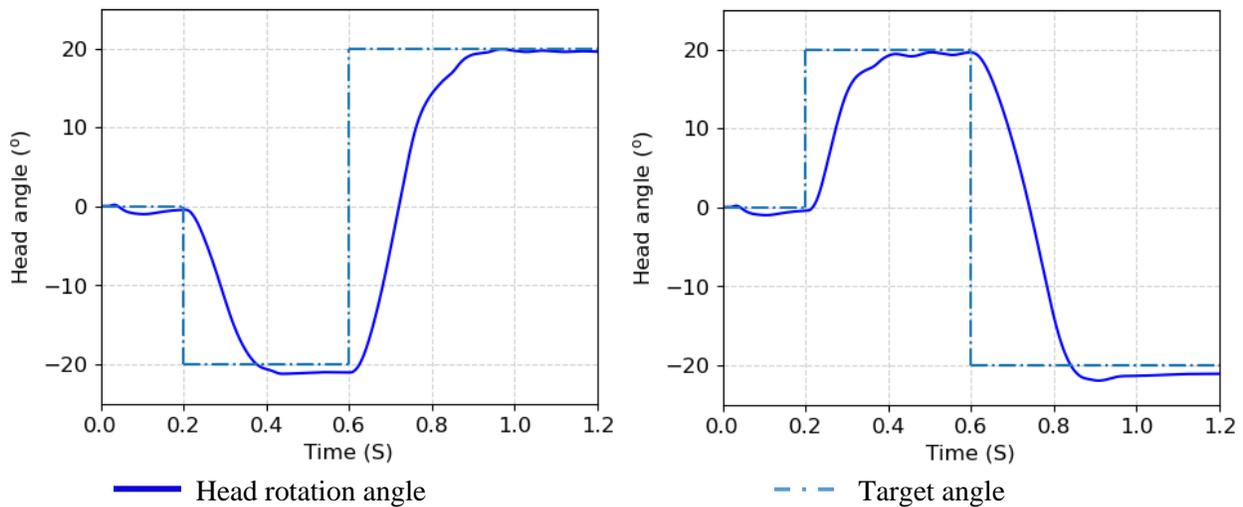
Figure 4-14: Head motion under novel target signals (a) 20^o flexion to 20^o extension (b) 20^o extension to 20^o flexion.

The trained RLMAC could generate the desired head kinematics from extension to flexion and flexion to extension (Figure 4-14) and the head remained in a stable posture at the end of the simulations. While performing the end-to-end kinematics, the head moved through the same eccentricity (horizontal distance between T1 and head CoM), however, at the same eccentricity, the head angle was 11.28^o in extension while performing the flexion to extension motion and 7.16^o in flexion while performing the extension to flexion motion (Figure 4-15). The shape of the spine at the neutral eccentricity was also different for both motions. The results show that the trained RLMAC can generate activations for head kinematics based on the target angle signal even if similar motions were not included in the training.

<div align="center">(a)   (b)</div>
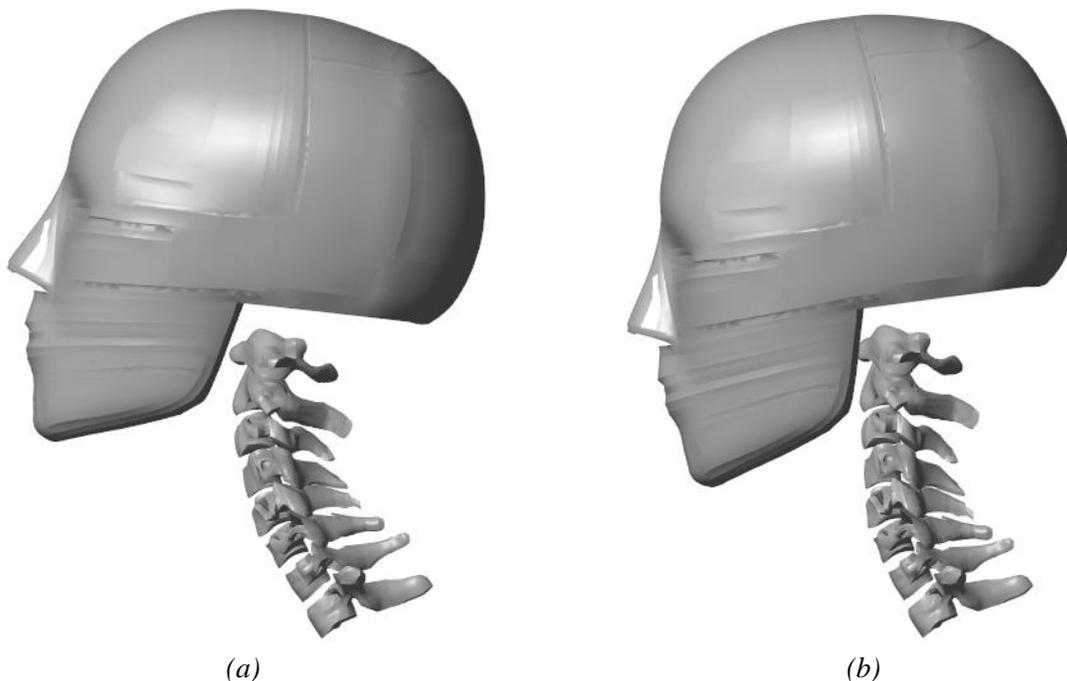
Figure 4-15: The head and spine alignment of the UVa neck model near the rest position (a) 20^o flexion to 20^o extension (b) 20^o extension to 20^o flexion.

The RLMAC was also evaluated under ramped target signals. During the training, a step signal was used to define the target head orientations, thus, in this chapter the ability of RLMAC to follow a ramped target was also evaluated. The head model was initially stabilized for 200 ms, following

which a ramped target was prescribed to move the head to 20° extension in 400 ms. The head was stabilized in the rotated position for 200 ms, and a target signal was again applied to bring the head back to the neutral position. 400 ms was used as the duration of the ramped signal as it was twice the time the RLMAC takes to rotate the head to 20° extension from neutral under step signal (Figure 4-11a). A similar process was also followed for moving the head to 20° flexion and back to neutral using ramp signals.

Figure 4-16 shows the head CoM angle-time history following a ramped target. The head was moved and was able to follow the target signal to 20° angle and back under both the extension and flexion cases.

(a)   *(b)*

——— Head rotation angle   — · — Target angle
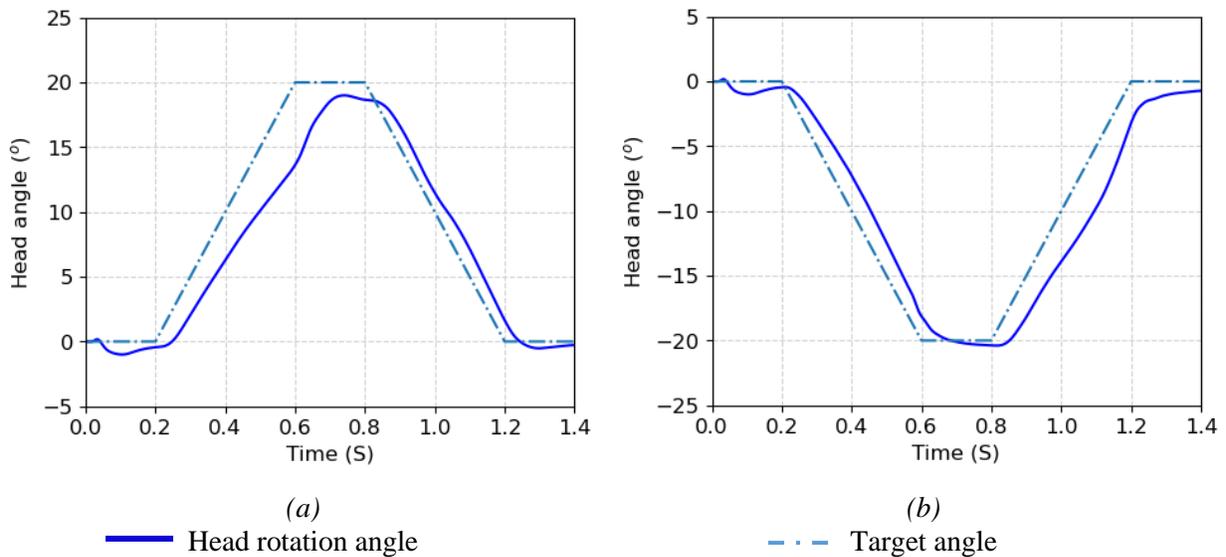
Figure 4-16: The head CoM angle following a ramped target signal (a) Extension (b) Flexion.

Figure 4-17 shoes the response of the neck model when the ramped flexion target signal is applied to the RLMAC (Figure 4-16b).

T = 0 ms *T = 200 ms*

T = 600 ms *T = 800 ms*

T = 1.2 s *T = 1.4 s*
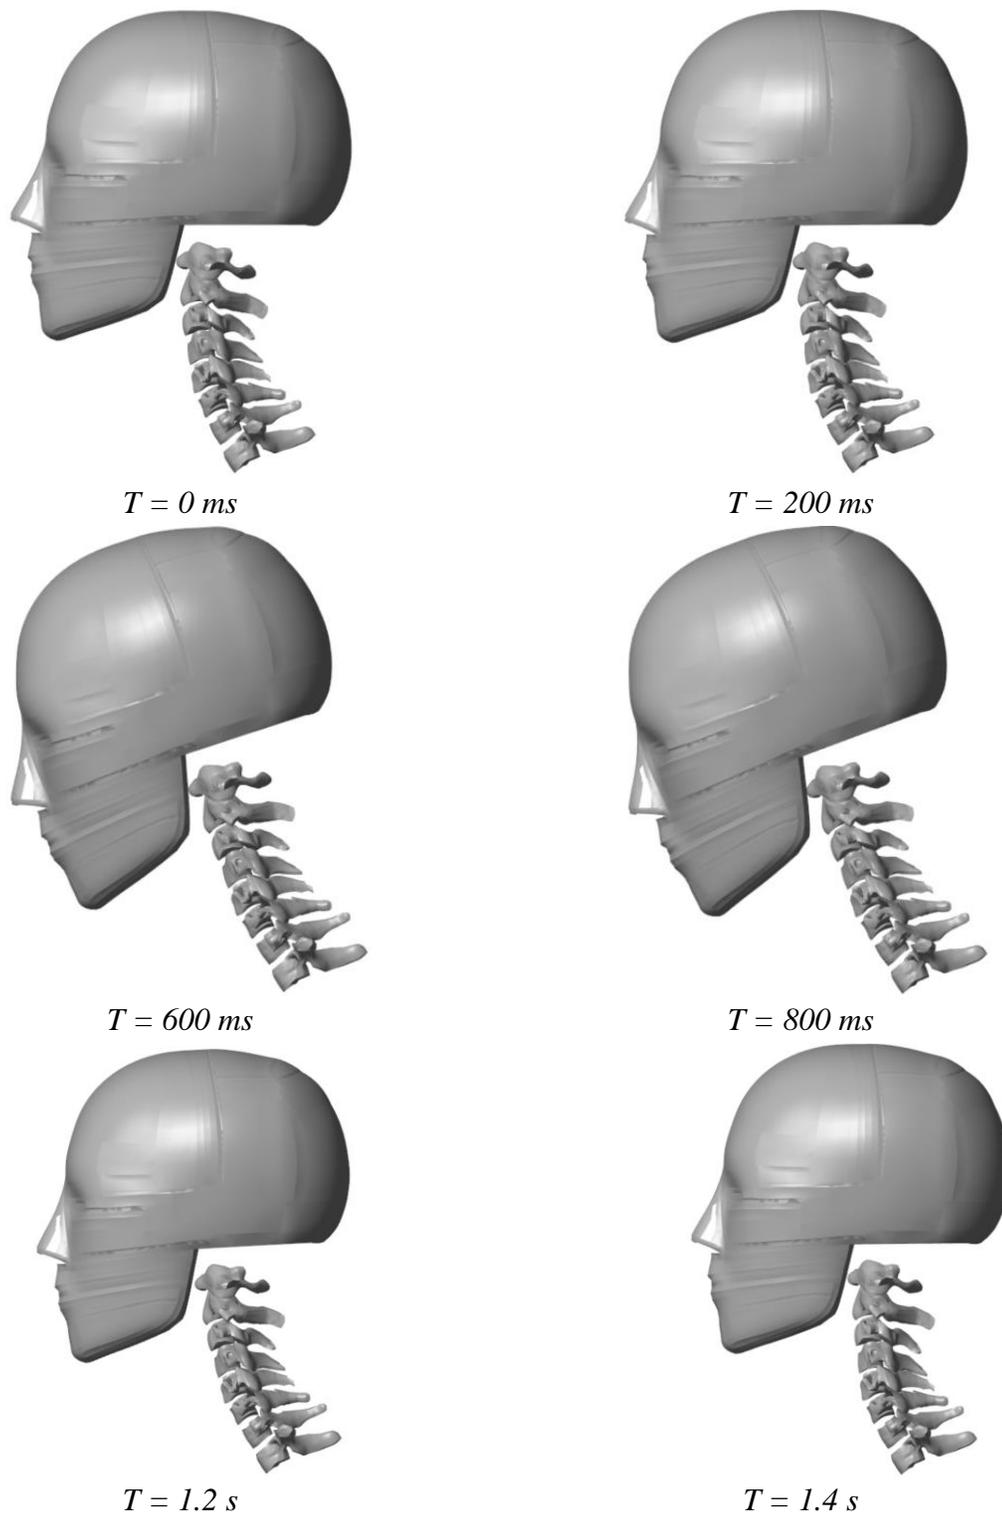
Figure 4-17: Simulation of the neck model in flexion under the ramped target signal

4.4 Discussion

The present chapter provides an overview of the incorporation of RLMAC with the head model. As the neck region is an over-constrained system, a twin-delayed deep deterministic policy gradient (TD3) was used to develop the RLMAC for muscle control of the neck (Fujimoto et al., 2018). As a first step towards the control of the head and neck motion, the muscle activities were made symmetrical about the mid-sagittal plane and only the motion about the plane was modeled.

The control objective of the RLMAC was to stabilize the head and spine under gravity as well as to produce goal-directed motion of the head. The state of the RLMAC included parameters that were representative of the sensory inputs, vestibulocollic reflex (VCR), and cervicocollic reflex (CCR). When the RLMAC was trained for the stability of the head in absence of the CCR terms, it was found that the agent could maintain the head stability, but in the process lead to the buckling of the spine (Figure 4-5). A similar trend was reported by Happee et al. (2017), who found that maintaining static stability with only VCR feedback leads to an unstable posture of the spine.

CCR provides feedback on the relative position and movement of the head with respect to the trunk (Goldberg and Peterson, 1986). The neck muscle spindles have been shown to be the major contributors towards neck proprioception (Richmond and Abrahams, 1979). Joint afferent neurons also affect the proprioceptive feedback when rotated beyond the ranges of motion (Grigg, 1994). Previous neck control studies have considered the neck muscle lengths as the CCR parameters corresponding to the neck proprioceptors (Happee et al., 2017; Ólafsdóttir et al., 2019, Zheng et al., 2021). In the present study, the position and the movement of the joints have also been evaluated as the CCR parameters (Table 4-1). This training and simulations performed in the study show that the RLMAC with joint displacements as CCR was able to stabilize the head and spine at the initial neutral position better (Figure 4-6). With the muscle length error as the CCR, the head

119

had to rely on the passive stiffness of the spine along with the muscle forces to be stabilized, as the reward function associated with the muscle length feedback tries to minimize the muscle energy (Table 4-1). In this study, a simplified muscle energy formulation is used (Margaria, 1968). There exist other more complex calculations of the muscle energy expenditure (Umberger et al., 2003), however, due to computational constraints, the simple formulation was used in the reward function of CCR scenario 1. The CCR parameters used in the present study are simplified assumptions of the various proprioceptive loops of the human reflex controls and have been used to inform the RLMAC of the relative position of the head and vertebrae compared to the trunk (T1 in the model).

The posture of the spinal column under the final stable state was different for the two different CCR scenarios that have been considered for the study. Although it is not uncommon for individuals to have different alignments of the cervical spine (Newell et al., 2018), the CCR scenario 2 (Table 4-2) was employed for further study as it maintained an upright, non-compressed shape of the spine. The reward function in the two CCR scenarios includes penalty terms for muscle energy and joint energy respectively (Table 4-2). Previous studies that have incorporated CCR parameters for muscle control in the head and neck region have used changes in muscle length and velocity as feedback control gains (Correia et al., 2021; Happee et al., 2017b; Olafsdottir et al., 2013; I. Putu A. Putra et al., 2019). As the objectives in the previous studies were to maintain the head and spine posture, muscle length error could be used. Since the focus of the present study was also to generate the goal-directed motion of the head, muscle length error or joint deflections couldn't be used in the sensory (target) reward as it is very difficult to correlate the muscle lengths or joint displacements to the target head angle. Therefore, we hypothesized that

the CNS tries to move the head to a target head angle and stabilize it such that the muscle energy or the joint displacements at the target state is minimum.

In the CCR scenario 1, muscle length feedback of individual strands was considered for the state, thus increasing the input size of the actor-network and the critic networks of the TD3 agent which also increased the training time. On the other hand, CCR scenario 2 considered the displacements of 7 intervertebral joints in 3 DoFs, reducing the CCR parameter size from 86 in scenario 1 to 21 in CCR scenario 2. Attempt was also made to combine the CCR parameters (muscle lengths and joint displacements) for training purposes but there was no major change in the output of the trained model from CCR scenario 2, and the computational cost of training increased with an increase in the size of the CCR parameters.

Next, RLMAC with the CCR scenario 2 was used to train the neck model to generate goal-directed head motions in the sagittal plane. The CCR scenario 2 (Table 4-2) state and reward were used for training the RLMAC and the target angle was varied in each iteration.

Post training the RLMAC was able to formulate a muscle synergy that allows for the goal-directed head movements, and the ability of the trained agent to move the head was verified in both flexion and extension (Figure 4-10), as well as in neutral position. Only one previous computational study has tried to generate head kinematics in the sagittal plane (Silvestros et al., 2021), however, in that study, a predefined control strategy was used based on electromyography data. Even in some of the previous stability simulations that include both VCR and CCR, assumptions were made on the muscle synergy. While maintaining stability under impact cases, Correia et al. (2021) split the neck muscles into four groups and maintained a predetermined ratio between the activations of muscles grouped as extensors and flexors. Olafsdottir et al. (2013) grouped the neck muscles into eight units and assumed a pre-defined contribution of each group towards neck stability. No such

assumptions are made regarding flexors or extensors in the present chapter, and the RLMAC actuates the 23 muscle pairs individually.

The muscle activation patterns generated by RLMAC have also been analyzed for the sagittal rotations. For the stability simulation, most of the muscles have low activity, while the semispinalis cervicis and the sternohyoid co-contract. This minimum level of co-contraction is required to maintain the head and spine posture under gravity (Choi, 2003; Vasavada et al., 2002). Due to the redundant nature of the musculoskeletal system, the trained controller at different levels of training may actuate the muscles differently for the identical target signals as no assumptions on the muscle activation levels were made in the reward function. In the arm control study in chapter 2, the muscle activations were minimized to reduce the muscle fatigue, which resulted in a drop in activation levels at the target elbow angle. Unlike the elbow model which was simplified as a revolute joint with constant stiffness, the vertebral joints in the neck model were modeled as 6 DoF joints with non-linear stiffnesses. High activation levels would lead to compression of the spine and thus high joint energy, thus penalizing joint energy reduces the muscle co-contraction at the neutral position to some extent. Penalizing the joint energy was also seen to reduce the muscle activations at the extension and flexion positions as can be observed from Figures 4-12 and 4-13.

During the flexion or extension motion of the spine, the neck muscle activations showed a biphasic or triphasic behavior. For a redundant system like the head and neck region, there can be many activation patterns that can lead to similar head motion. An experiment performed on cats showed that the voluntary kinematics muscle patterns differed between individual cats (Peterson et al., 1989). Keshner et al. (1989) also found that during voluntary kinematics in the sagittal plane in humans, the splenius muscles were activated during extension for half of the volunteers and during flexion in the other half. Siegmund et al. (2006) also found that the response of posterior muscles

in voluntary or sled movements varied among volunteers. With separate training hyperparameters, the stimulation output from the agent may be different for the same target kinematics, but the head kinematics synthesized would be similar as that depends on the state and the reward functions.

During the training, the head was always maintained at the neutral position (0^o) sagittal angle at the start of the iteration (T = 0 ms), and it was moved to a target angle from the neutral position. However, when the target angle was varied to move the head from initial extension or flexion position, the head could follow the target signals from flexion to extension or from extension to flexion (Figure 4-14). This demonstrates the robustness of the RLMAC, which can formulate a muscle activation pattern based on the current angle and the angular error, even though those values for the parameters were not explicitly used in the training.

It was observed in the simulations that when the head CoM approached the initial position from either direction, the spine shape, and the head angle were both different from that in the neutral posture (Figure 4-15). This phenomenon was tested by Newell et al. (2018), who found that the spinal alignment and the head angle differ for dynamic and static conditions. It was also found in the experiments that the shape of the spinal column during the motion near the stable rest position depends on the movement direction of the head. Similar trends of head and spinal orientation were seen in the end-to-end simulations.

The RLMAC was also able to move the head reacting to a ramped target angle (Figure 4-16). The head CoM followed the target angle under both the extension and flexion motion. There was a time delay between the onset of the target signal and the actual head motion, and the delay was present during head movement up to stabilization. The latencies observed in following the target signals result from the time required by the RLMAC to update the outputs based on the state as

well as delays associated with activating the muscles after receiving the neural stimulations from the activation dynamics (Zajac, 1989).

4.5 Conclusion

The training and simulations performed in this chapter show that muscle control frameworks based on deep reinforcement learning are effective for the control of head and neck motion. The RLMAC trained for control of head motion in the sagittal plane was able to stabilize the head and generate flexion and extension motion of the head and the spine. The trained RLMAC could also produce head kinematics under target signals that were not used during the training, verifying the robustness of the RLMAC. In the following chapters, the outputs from the RLMAC will be validated with testing data and the RLMAC will be extended for omnidirectional head motion control.

Chapter 5 – Measurement of goal-directed head motions in humans

The previous chapter outlined a framework for using reinforcement learning agents for synthesizing and control of head kinematics by selectively activating the neck muscles. While it was demonstrated in chapter 4 that the RLMAC can generate the desired head kinematics, it is important to evaluate the biofidelity of the response of RLMAC with human head kinematics data. Thus, a limited volunteer study has been performed to characterize target-specific head kinematics in humans and the protocol followed for the testing has been described in detail in this chapter.

The present volunteer study aimed to measure fast goal-directed head kinematics in eight study subjects. The volunteers had to perform ten targeted head motions in extension, flexion, and axial rotation. The head targets were set such that the targets are below the physiological limits previously measured and the trunk, as well as the T1 have minimum motion while the subjects perform the head rotations. The head was instrumented with a sensor to measure the angular velocity of the head movements. The head kinematics data were used to validate the response of the trained RLMAC in its ability to synthesize biofidelic head kinematics.

The simulation results of the neck model with the RLMAC trained in chapter 4 all lie within the ranges measured in the volunteer tests. The results suggest that the control architecture for RLMAC that was used for generating a response in the sagittal plane is biofidelic and can be extended for omnidirectional control of head kinematics.

5.1 Introduction

Gathering head movement data is important for validating a computer model because it allows for the comparison of the model's predictions to real-world observations. This helps to determine the accuracy and biofidelity of the model. Head movement data is particularly important for computational models that are designed to simulate human voluntary head motions, as it enables the assessment of how well the model can capture the complexity and variability of human head movements.

Many previous volunteer studies have analyzed the stabilization behavior of the head and neck under external loads. Keshner et al. (1989) performed human subject tests to observe the contribution of four neck muscles to the overall stability of the head under omnidirectional external loads. The authors found that the splenius muscles have different direction preferences among the test subjects. In another study, Keshner and Peterson (1995) measured the response while the human subjects tried to stabilize their heads under the presence and absence of visual inputs. Kuramochi et al. (2004) assessed the head kinematics when impacted by a 4 kg mass in the anterior-posterior direction at the mid-sagittal level. During the perturbations, the EMG data of the sternocleidomastoid and trapezius were measured, and it was observed that the activity of the sternocleidomastoid was higher in eyes closed condition than in eyes open. In a volunteer study by Siegmund et al. (2006), a 50 N sweeping force was applied to the head of the human subjects who maintained the head stability against the force. The muscle activities of the three subjects were measured during the application of the load and compared for the subjects. Eckner et al. (2014) applied impulsive loads along the extension, flexion, lateral bending, and axial rotation in passive and anticipatory conditions to evaluate the effect of neck strength and muscle activation level on the neck kinematics. Reynier et al. (2020) performed padded lateral impacts on human

126

volunteers and measured the head kinematics in three conditions – Passive, co-contraction, and unilateral activation. The head angular velocity and linear acceleration were measured during the impacts and the magnitudes were highest for the passive case followed by the unilateral case. Homayounpour et al. (2021) applied impulsive impacts on human subjects in anterior, posterior, and lateral directions under three muscle conditions – passive, directional, and co-contraction. During the impacts, the head kinematics and muscle activations were measured and reported for each condition.

Past studies have also focused on characterizing the head motion and muscle activities during voluntary head and neck rotation. Foust et al. (1973) measured the range of motion in the sagittal plane in 180 human subjects. Zangemeister et al. (1982, 1981) measured the head kinematics and muscle activities of human subjects while attempting rotations of the head in the horizontal plane. Mayoux-Benhamou et al. (1997) measured the activations of the dorsal muscles in 19 volunteers for head rotations in the sagittal plane. Margulies et al. (1998) compared the neck response of 5 male volunteers under voluntary and forced flexion of the head and found that the maximum neck flexion did not differ significantly under the two conditions. However, the head retractions and rotational velocity were higher in the forced flexion of the head. Lantz et al. (2003) measured and quantified the range of motion for rotations in all three anatomical planes. Siegmund et al. (2001) evaluated 20 human subjects in their ability to respond to two different auditory stimuli. The authors found that the muscle response time in presence of a startling signal was half of that observed in a standard "go" stimulus. Cheng et al. (2008) quantified the effect of muscle co-contraction during fast voluntary movements of the head in extension, flexion, and lateral bending. The muscle co-contraction patterns were different for flexion and extension motion, however, they were similar during the lateral bending in both directions. Siegmund et al. (2006), performed a

127

study that characterized the muscle activations during voluntary flexion-extension movements along with isometric stability and seated forward acceleration scenarios. Hernandez and Camarillo (2019) measured the maximum velocity at which humans can rotate their necks in the three anatomical planes.

Many volunteer studies have also been performed to investigate the effects of muscle forces in automotive impact scenarios. Matsushita et al. (1994) conducted volunteer tests in conditions imitating frontal, rear-end, and side impact scenarios while performing X-ray studies of the neck motions. Geigl et al. (1995) performed sled tests to analyze the response of human subjects, Hybrid III dummies, and cadavers in a rear-impact environment. Ono et al. (1997) performed a test series to obtain the response of human volunteers under four different conditions evaluating the effects of muscle tension, neck alignment, and seat stiffness. van den Kroonenberg et al. (1998) measured the head and neck kinematics for 19 subjects under low-speed impacts with ΔV ranging between 6.5 and 9.5 km/h. Siegmund et al. (2003) replicated whiplash-like conditions to compare the response of aware and surprised volunteers. The response of aware subjects was found to differ from that of surprised subjects. The studies also found that the muscle activities and the corresponding head kinematics changed with multiple whiplash-like perturbations. Siegmund et al. (2004) also evaluated the head-neck kinematics and muscle forces to changes in impact conditions. The study found that the neck muscle response correlated with the sled acceleration and change in velocity. Ejima et al. (2008, 2007) analyzed the response of volunteers under frontal impacts and emergency braking conditions. The volunteers were tested under relaxed and tensed muscle conditions and it was found that the reflex time for the head, neck, and torso region was between 70 ms – 200 ms for volunteers. Arbogast et al. (2009) conducted sled tests under frontal impact conditions to compare the response of children under such conditions to adults. The main

factor affecting the kinematics of the head and spine kinematics was found to be the anatomy of the head and neck with increasing age. Beeman et al. (2011) conducted human volunteer studies to evaluate the effect of bracing in low and medium-severity impacts. The authors found that the effect of bracing was more significant in low severity cases ($\Delta V = 4.8$ km/h) where the forward excursions of different body parts were reduced by 35-70 %, compared to 18-26 % reduction in the case of medium severity impacts ($\Delta V = 9.7$ km/h). Carlsson and Davidsson (2011) compared the response of 50[th] percentile male and female volunteers under rear impact scenarios. The study found that at impact velocities of 4 and 8 km/h, females had lesser rear translation and angular motions of the head and T1. Beeman et al. (2016) compared and quantified the neck forces and moments of relaxed and braced occupants against cadavers in low-speed frontal sled conditions. The study found that bracing did not significantly affect the magnitudes for peak neck forces or moments, however, the muscle activation did affect the timing at which the peak forces were obtained. Fice et al. (2021) evaluated the effect of bracing against the steering wheel on the head kinematics and neck muscle activity. The study found that bracing of arms against the steering wheel doesn't lead to an increase in neck muscle activations during the impacts.

Few studies have measured the neck response to automotive lateral impact scenarios. Ejima et al. (2012) conducted experiments to measure the response of humans under pre-crash lateral impact conditions. The contribution of muscle forces to the overall behavior of the test subjects was evaluated in detail. Ólafsdóttir et al. (2015) performed human volunteer studies to determine the muscle activation patterns in seven neck muscles in multi-direction seated perturbations. Eight volunteers were subjected to seated perturbations in eight directions, the corresponding muscle activities were measured to determine the reflex direction preference of the neck muscles. Chan et al. (2022) quantified the occupant behavior of 5[th] percentile females and 50[th] percentile males in

low-speed frontal and oblique impacts. The study found that both the male and female subjects have similar kinematic responses under braced conditions, but some of the kinematic data differed in the relaxed state.

All these volunteer studies have been useful in developing active muscle controllers in isolated head-neck models (Cappon et al., 2007a; Correia et al., 2021; Happee et al., 2017b; Panzer et al., 2011; Zheng et al., 2021) or full body models (Devane et al., 2022; Iwamoto et al., 2012; Iwamoto and Nakahira, 2015; Östh et al., 2015, 2012a). However, only one study has characterized goal-directed head movements in the axial rotation (Zangemeister et al., 1982). As one of the objectives of the RLMAC is to generate voluntary head motions in all three rotational DoFs, a limited voluntary study has been performed as a part of this dissertation to characterize fast goal-directed head motions in human volunteers. As the neck model is representative of the anatomy of a 50^{th} percentile male, the volunteer study has targeted human subjects of similar populations. The data gathered from the human subjects will be used as validation cases for the neck model coupled with RLMAC trained in chapter 4.

5.2 Methodology

Eight volunteers, all males participated in the study. The subjects did not have any prior history of injury or a neuromuscular condition that may have affected their normal day-to-day movement of the neck. The subjects were also screened for any auditory or visual issues which may affect their participation in the study. The study was approved by the University of Virginia Institutional Review Board, and the recruited subjects signed a consent form before participating in the study or any kind of data collection. The volunteers were required to undergo 10 different goal-directed voluntary motions of the head.

5.2.1 Protocol

The subjects sat in a Honda Odyssey second-row captain's chair rigidly placed on the ground and facing a flat wall at a distance of 5 feet from the wall. A centerline was drawn on the chair to provide an estimation of the mid-sagittal plane and a line parallel to the chair center was also marked at the wall using a laser level (Bosch GLL2-80). After being seated, the subjects were asked to put on headgear mounted with a laser pointer (Class 3R, wavelength – 642 nm, power – 2 mW) and a three-axis gyro sensor (Figure 5-3a). The subject's Frankfort plane was aligned with the ground using a digital protractor (PRO360, Level Development ltd, Chicago, Il). After the alignment, the laser was adjusted so that the beam was parallel to the Frankfort plane. The laser pointer on the headgear was turned on and the spot on the wall along the sagittal plane line where the laser was pointing was marked as the neutral position. The distance of the neutral spot on the wall was measured from the laser tip and this distance was used to place different target markers – 20° extension, 30° extension, 20° flexion, and 30° left axial rotation (Figure 5-1, Figure 5-3b). These values of the head target rotations were chosen from a pilot study and enabled the fast head rotations without much movement of the trunk or T1.
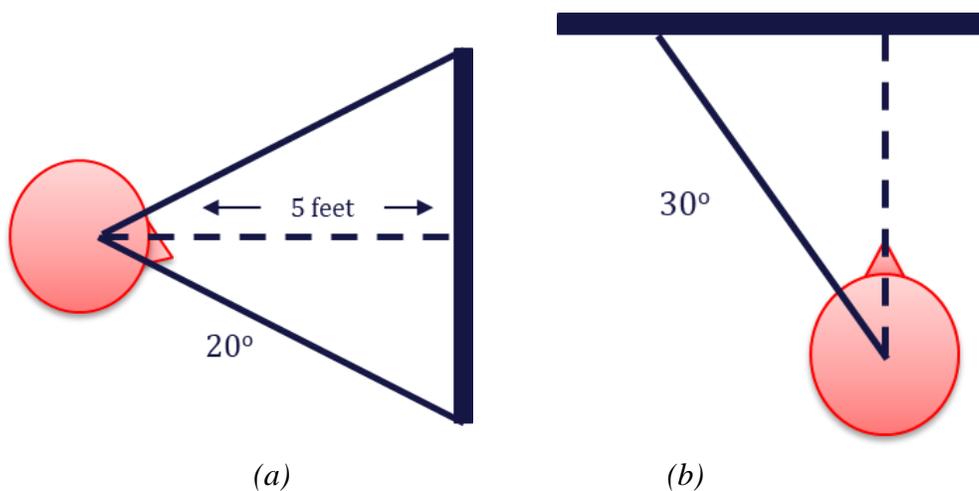
Figure 5-1: Placement of targets on the wall (a) Extension-Flexion (b) Axial rotation (c) Marker position on the wall.
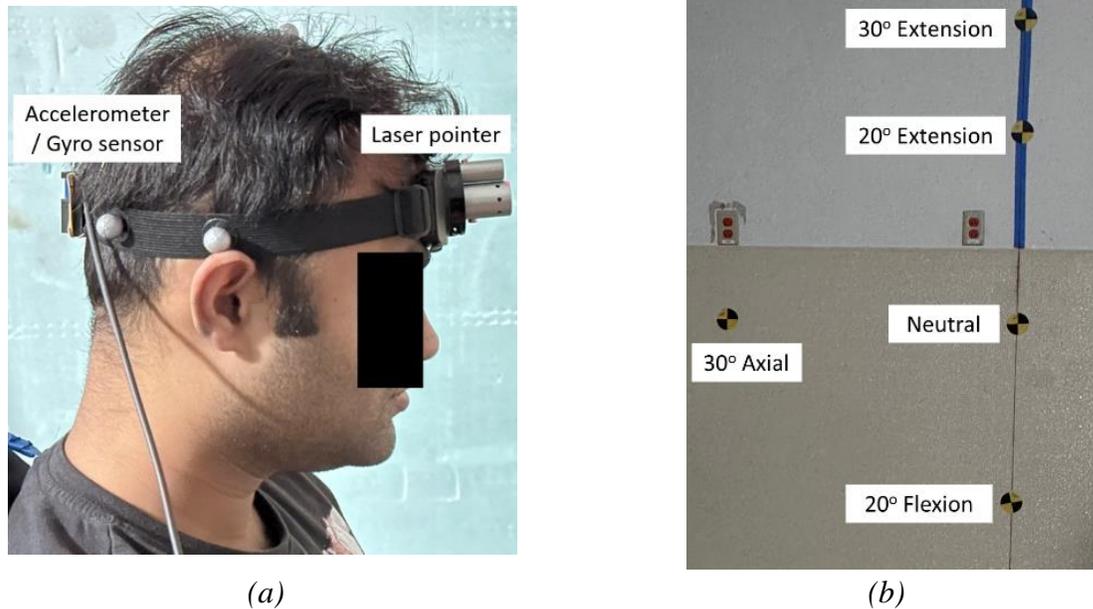
Figure 5-2: (a) Volunteer with head gear containing the laser pointer and sensors (b) The targets for goal directed head movements.

The head-mounted laser was turned on and subjects were asked to point the laser to a specific initial marker on the wall (marker 1) (Figure 5-3c). The subjects were instructed to move the head toward a different target marker (marker 2) as fast as possible on hearing an auditory signal (a beep sound). The beep sound would go on for 4 seconds and at the end of the beep sound, the subjects were asked to move their head back to marker 1 from marker 2. Subjects were allowed to practice the movement before the actual recording of data to get habituated to the movement pattern and to prevent any out-of-plane movements. Analysis of the pilot data showed that there was no observable difference between the latency and peak velocities while combining two movements (marker 1 to marker 2 and back) or while performing the movements separately. The volunteers were given a 2s warning before the triggering of the auditory signal to make them aware of the signal and prepared for the corresponding movement. The 4s gap between the to and fro motion between the two markers allowed enough time for the volunteers to stabilize and was similar to the timeframe used in Cheng et al. (2008) study. Each movement pair was performed twice and there was a small gap of approximately 5 minutes between two successive movements. The

132

subjects were informed about the initial and target markers before each trial along with the warning of the onset of the auditory signal. Table 5-1 summarizes the different head movements performed by the subjects. All the goal-directed head motions were performed one after another in a single sitting.

Table 5-1: Targets for goal-directed motions

Marker 1	Marker 2
Neutral	20° Extension
Neutral	30° Extension
Neutral	20° Flexion
20° Extension	20° Flexion
Neutral	30° Left Axial

5.2.2 Instrumentation

The head angular velocities during each motion were measured by an accelerometer, mounted on a headgear that the subjects had to wear during the study (Figure 5-3a). The headgear was adjusted according to each subject such that it fits tightly to the subjects' head with no relative motion between the subjects' head and the head gear during the head rotations. Initially, in the pilot studies, a 6 DX cube (DTX Seal Beach, CA) was used for measuring the head kinematics data. As the head angular velocity range was much lower than the maximum range of the accelerometer (18000 o/s) (Hernandez and Camarillo, 2019; Siegmund et al., 2006a; Zangemeister et al., 1982), the output from the sensors was noisy. Thus, for the tests, a three-axis gyro sensor (IES 3103-600, Braunschweig, Germany) with a max range of 600 o/s was used to measure the head rotation data.

All the sensor output signals were collected at 1000 Hz, using the DTS slice data acquisition system (DTX Seal Beach, CA). The position of the sensor was manually adjusted to place it at the

133

posterior side of the head along the sagittal plane (Figure 5-3a). The head angular velocity signals along the three anatomical planes were directly measured by the gyro sensor, which was then integrated to obtain the head angle-time data (Siegmund et al., 2001). The angle-time data were used for comparison with the output of the trained RLMAC.
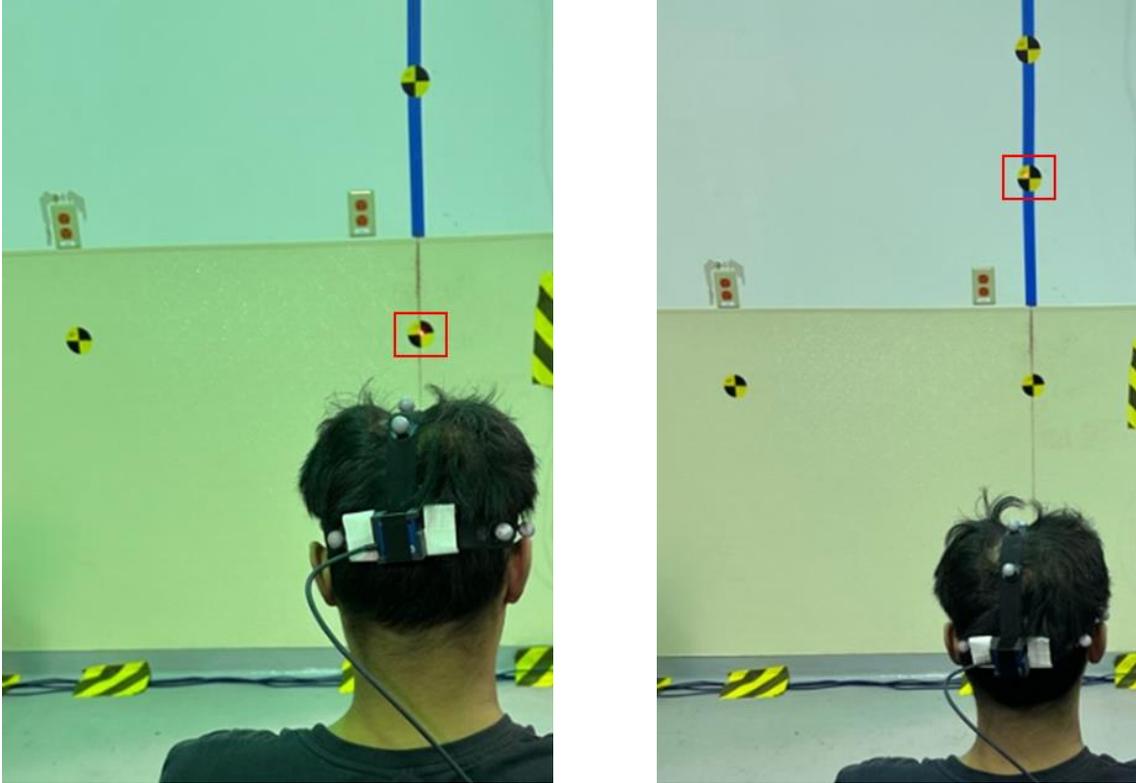
Figure 5-3: The initial and final head position for Neutral – 20° Extension. The location of the laser pointer in the wall is highlighted.

Table 5-2: Information of the study subjects

Subject	Age (years)	Height (cm)	Weight (kg)
Pilot 1	30	175.26	79
Pilot 2	43	180.34	93
Subject 1	29	180.34	89
Subject 2	23	172.72	64
Subject 3	27	177.8	86
Subject 4	28	177.8	74
Subject 5	25	190.5	100
Subject 6	27	182.88	95

5.2.3 Comparison with the simulation data

The simulation outputs from the trained RLMAC were compared with the experiment rotation-time results obtained from the volunteer study. In the simulation, the head of the model was moved to the initial position (Marker 1 from Table 5-1) before simulating the goal-directed motion. The head was stabilized at the initial position for 400 ms before prescribing the target position (Marker 2 from Table 5-1). For the simulations, a step signal was provided to the RLMAC as the target head position from the initial position, from which the angular error required for the state was calculated. The objective of the simulations was to compare the head kinematics during the goal-directed motion, and the latencies which are generally associated with a reaction to the sensory (auditory or visual) signals in the volunteer studies have not been considered in the present chapter as more information is required for model building to represent such responses.

5.3 Results

The head velocities in each trial were recorded using the 3-axis gyro for 10 seconds. After 2 seconds of the data recording initiation, an auditory beep signal was sounded, indicating the volunteer to start performing the head motions, which went on for 4 seconds. The volunteers moved their head back to the initial marker at the end of the beep sound. The velocity-time signal measured for a volunteer while performing neutral to 20° extension motion is provided in Figure 5-4. The first vertical line in the figure indicates the start of the beep sound, signaling the volunteer to move the head in extension from neutral and stabilize. The second vertical line indicates the end of the beep sound, signaling the volunteer to move the head back to the neutral position.
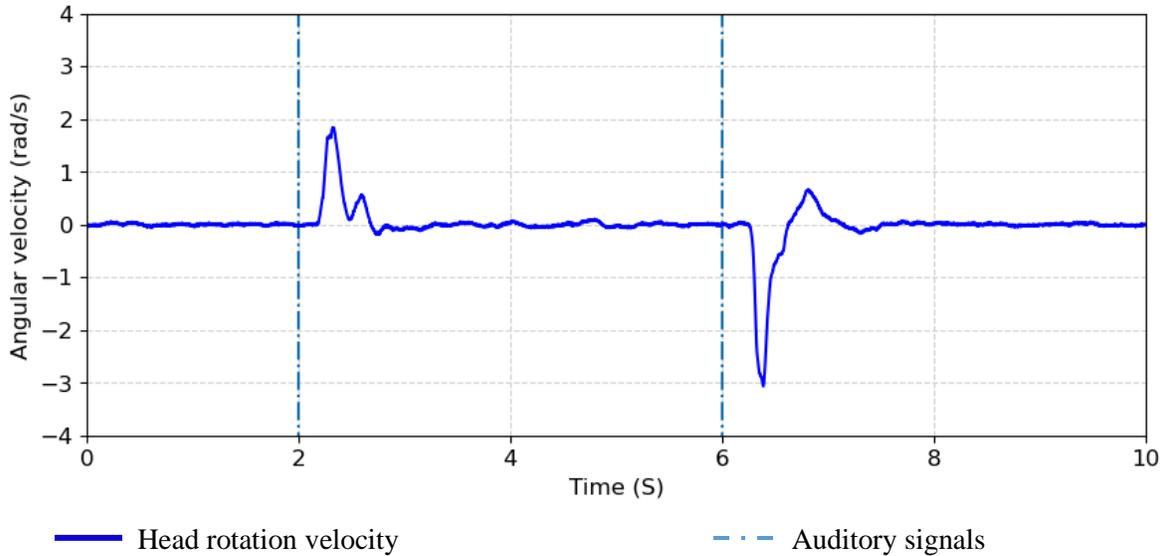
Figure 5-4: An example of the velocity-time profile of head angular velocity measured by the 3-axis gyro sensor for neutral to 20º extension and back.

A pilot study was performed initially which showed that the velocity profile while moving the head independently from extension to flexion (marker 2 to marker 1) was similar to the head kinematics while performing the combined motion (Figure 5-4). In both situations, the velocity curve had similar latencies and peaks. Figure 5-5 compares the velocity time history of head velocity while performing the 20º extension to neutral motion.
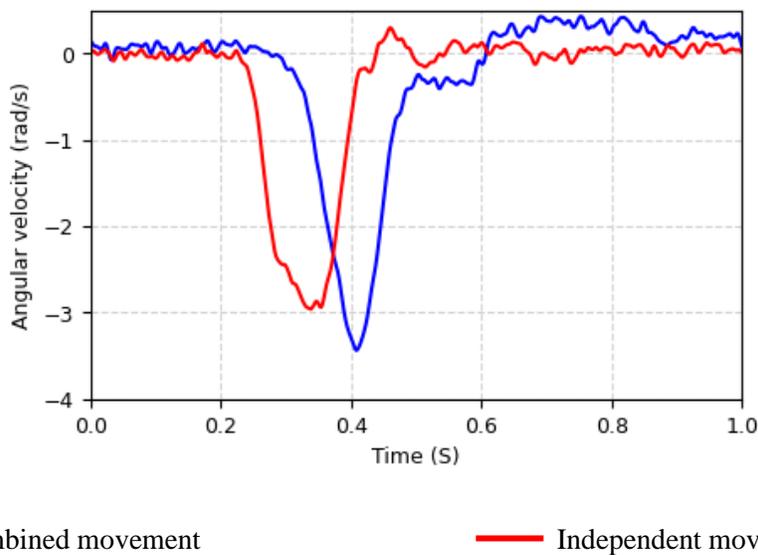
Figure 5-5: An example of head rotational velocity in 20º extension to neutral motion in combined movement and individual movement. The start time (0 s) has been adjusted for the auditory signal.

136

The pilot study was performed with the 6-axis DTX cube and the measured signals in Figure 5-5 have been filtered using a lowpass Butterworth filter at a sampling frequency of 300 Hz. The analysis of the pilot data shows that the sensor measurements obtained from the combined head kinematics situation while returning (marker 2 to marker 1) is comparable in nature to the sensor readings when a similar head motion is performed independently, and hence can be used as a validation data set for comparison with the simulation results.

(a)   (b)
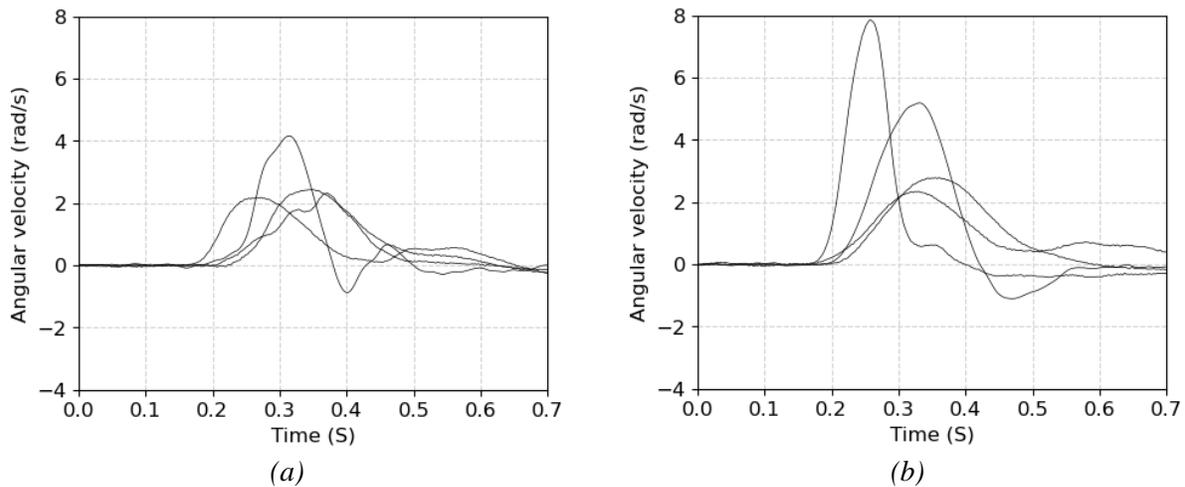
Figure 5-6: Velocity time history for 4 volunteers after the start of the buzzer sound (a) 0-20° Extension (b) 0-30° Left axial rotation.

The head velocity onset times for volunteers had latency between $150 - 250$ ms. The peak velocity and the duration of the head movements differed between the subjects (Figure 5-6). The 0 s in Figure 5-6 is the start of the buzzer sound, informing the volunteers to move their heads to the target position. The velocity profile of the volunteers was adjusted for the latency is plotted in Figure 5-7.

The plots in Figure 5-7 have been adjusted such that the head reaches 2° of angular displacement at 0s. Out of the 16 curves for each head kinematics, 4 were measured using the 6-Dx cube and hence have been filtered using a lowpass Butterworth filter. The measurements from the 3-axis

gyro have been presented without filtering. All the signals have been debiased by the mean of the

measurements 500 ms before the start of the motion.
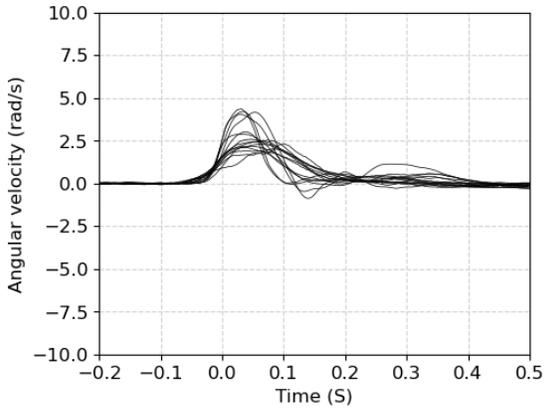
Neutral – 20° Extension
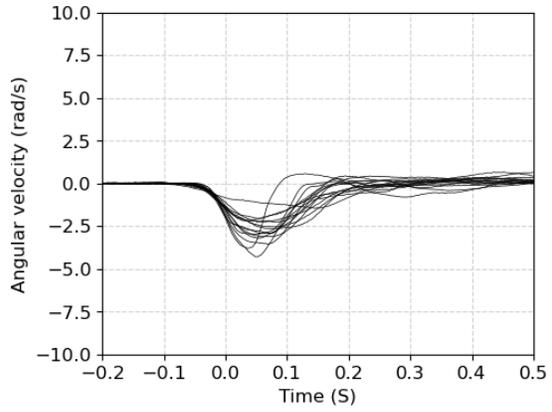
20° Extension - Neutral
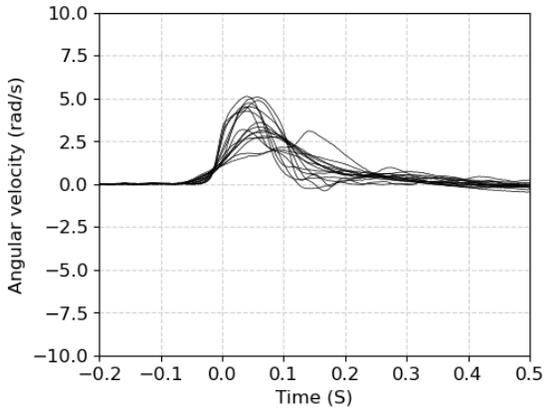
Neutral – 30° Extension
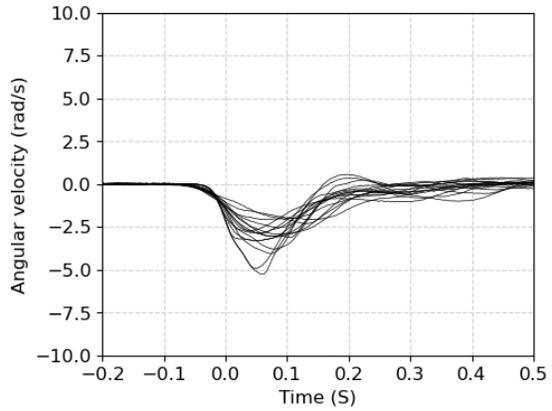
30° Extension - Neutral
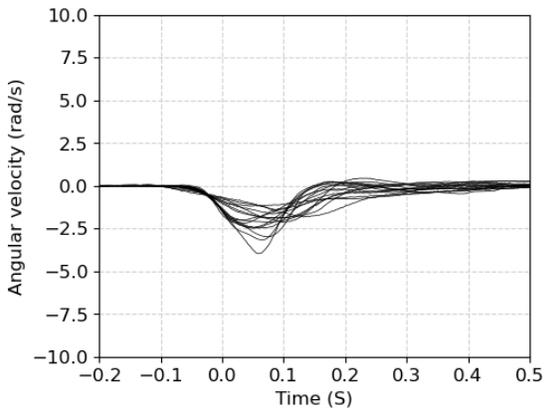
Neutral – 20° Flexion
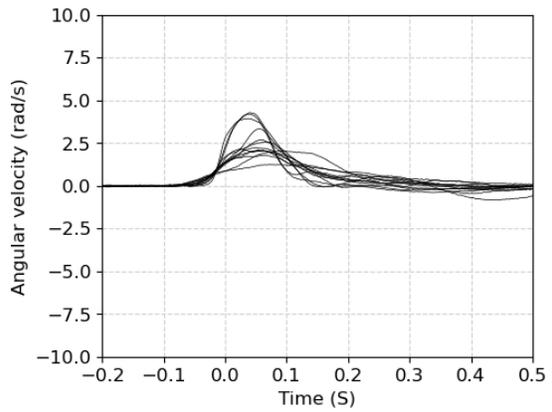
20° Flexion - Neutral
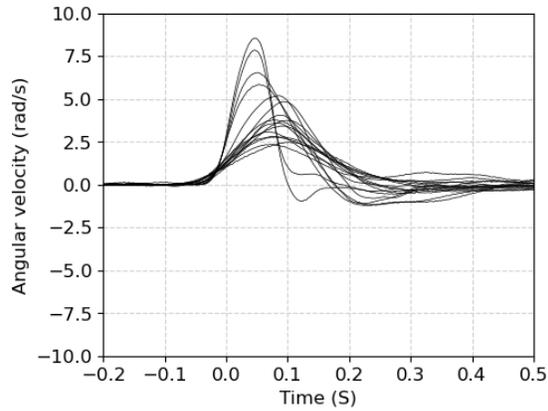
Neutral – 30º Axial rotation
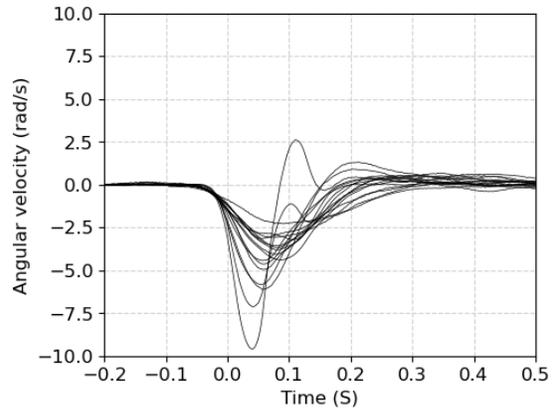
30º Axial rotation - Neutral
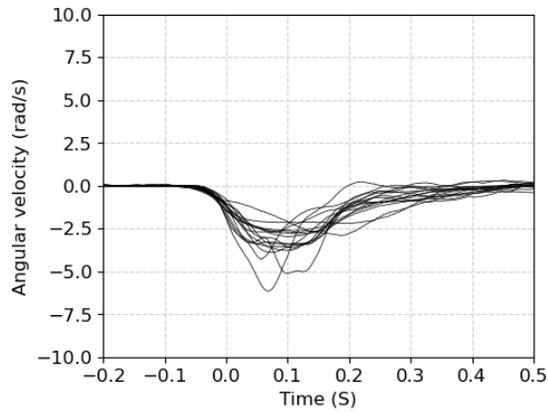
20º Extension - 20º Flexion
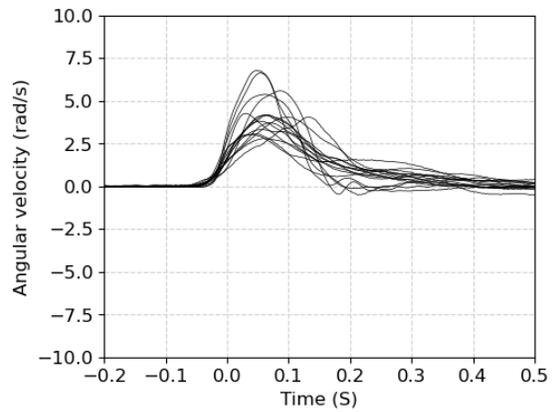
20º Flexion - 20º Extension

Figure 5-7: Velocity profiles of the goal-directed head motions for every volunteer. The velocity profiles have been adjusted for the latency for the initiation of head movement.

The axial movements - both neutral to 30º left rotation and back, had the highest average rotational velocities, while neutral to 20º flexion had the lowest angular velocity. It was commonly observed in the volunteers that the movements from the extremities to the neutral position had a higher velocity than the neutral to extremity motion, and the head also overshot while coming back to the neutral before stabilizing.

The velocity profiles in Figure 5-7 were integrated to obtain the head angle-time history corresponding to every goal-directed head motion case. The head kinematics that was performed in the sagittal plane were compared with the simulation results with the RLMAC trained in chapter

139

4. Comparisons of the simulation results of the RLMAC in the training scenarios have been presented in Figure 5-8. For the goal-directed kinematics in the sagittal plane, that were used for training in chapter 4, the RLMAC results were similar to the volunteer head kinematics. The angle-time plots in Figure 5-8 have been adjusted such that at 0 s time, the head sagittal angle is 2° for all the plots. The target angle in Figures 5-8 and 5-9 provide an approximation of initiation of the target step function at t = 0s.

Neutral – 20° Extension

Neutral – 30° Extension

Neutral – 20° Flexion
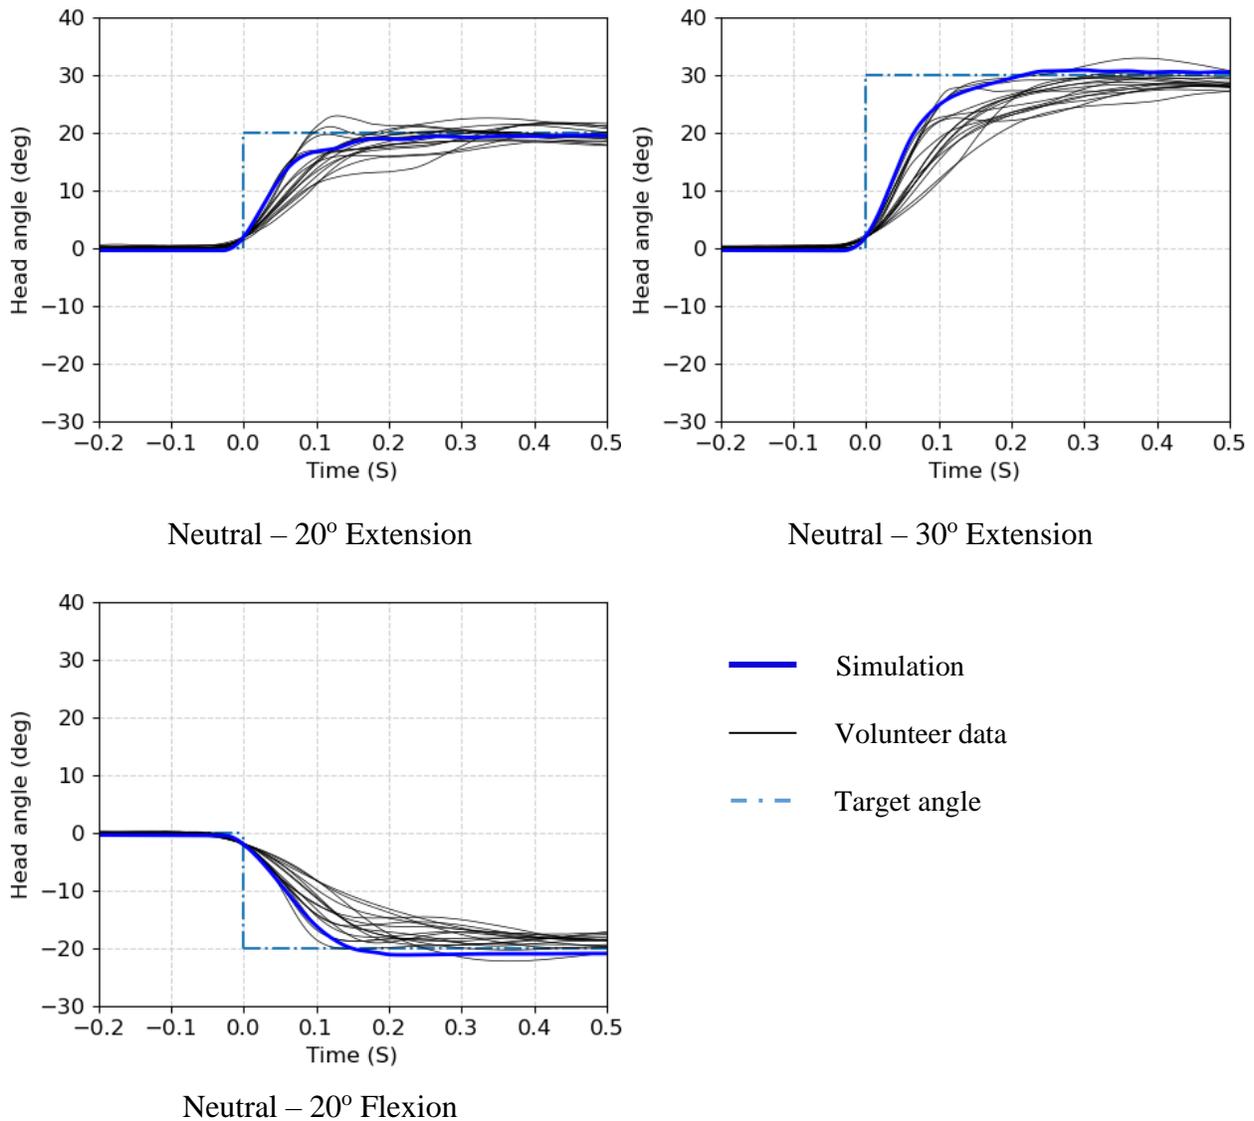
Figure 5-8: Comparison of the output of trained RLMAC with volunteer data for the training scenarios. The curves are adjusted such that the test and simulation curves are at 2° at time 0s.

140

The simulation results of the trained RLMAC were also performed with the targeted head kinematics that was not a part of the training and has been compared with the volunteer data in Figure 5-9.

20^o Extension - Neutral   30^o Extension - Neutral

20^o Extension - 20^o Flexion   20^o Flexion - 20^o Extension

20^o Flexion - Neutral
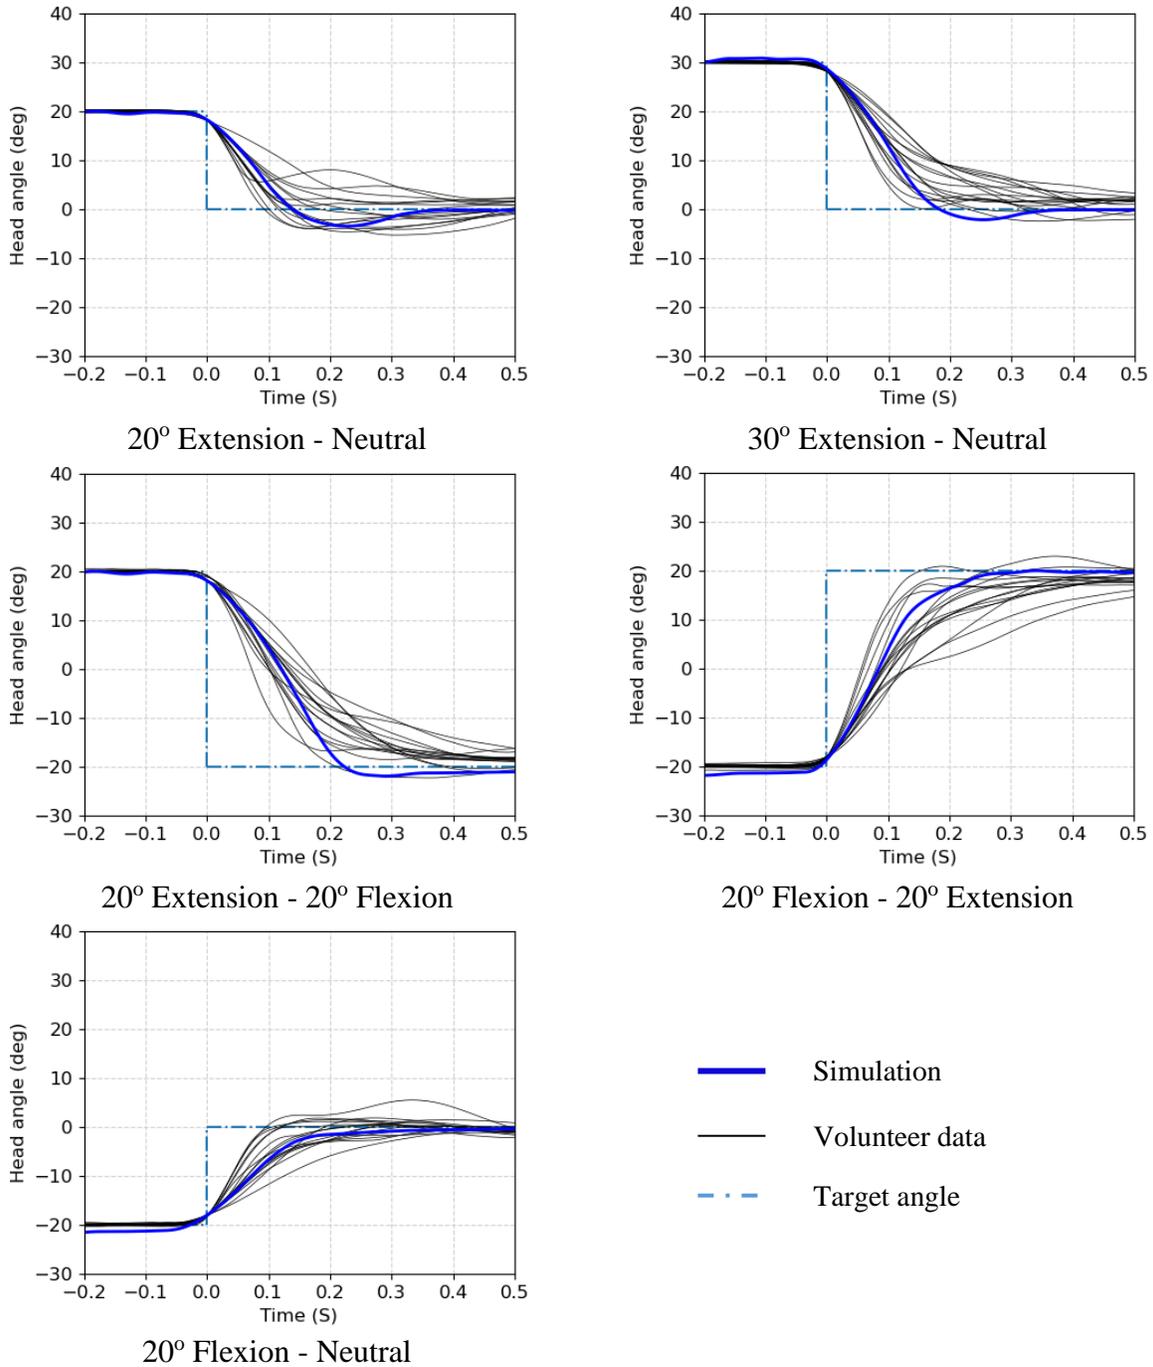
Figure 5-9: Comparison of the output of trained RLMAC with volunteer data for the scenarios not included in the training.

5.4 Discussion

In the current chapter, goal-directed head responses of male subjects with anatomy close to the 50th percentile males have been characterized. Three head movements – extension, flexion, and axial rotation have been measured with the intention of using the data as a validation case for the neck model coupled with RLMAC. Only one previous study, Zangemeister et al. (1982) has previously measured the targeted head rotations in human subjects in the axial plane, thus the data gathered in this chapter can be used as a validation dataset in other control studies as well.

Head kinematics have been measured using a 6-axis DTX cube or a 3-axis gyro sensor attached to the head of the human subject in the posterior direction. The goal-directed head motions have been measured by placing a laser pointer at the forehead of the volunteers and asking them to move their heads as fast as possible while pointing the laser between a set of markers indicating the targets. The study was approved by the UVA Institutional Review Board, and the use of laser for the study was permitted by UVA Laser Safety Officer, Environmental Health and Safety.

The head rotation targets have been chosen such that minimum motion of the T1 takes place while rotating the head. The magnitudes of the rotation targets were selected from a pilot study. The volunteers were asked to rotate their heads by reacting to an auditory signal (beep sound), and they were made aware of the sound 2 s before the onset. The auditory signal was similar to the "go" sound used in a volunteer study by Siegmund et al. (2001). The volunteers performed a total of five head movements in the sagittal and axial planes. Although the volunteers were instructed to start the head motion immediately after the auditory signal, a latency of 150-250 ms was observed before the sensors started recording any signals. The rotation velocity data measured by the sensors were used to obtain the angle-time profile for each head motion which was then used to compare with the simulations data.

On average, the rotational velocity during the axial rotations was highest. The volunteers also moved their heads towards the neutral position faster than from the neutral position to the extremities. Many volunteers found moving the head to 20° flexion difficult and slowed down considerably before finally pointing towards the flexion marker. The head angle values measured by the sensors were within 2.5° of the target head angle. Zangemeister et al. (1981) reported the head angular velocity of volunteers while performing 20° to 40° axial rotations between 3 – 5 rad/s which was within the velocity ranges measured in the study, and the velocity-time profile measured for the velocity ranges were also similar. Further analysis can be done to study the effect of physiological parameters on the head response in the volunteers.

The angle-time history data gathered from the subjects were used to validate the RLMAC which was coupled with the neck model. In chapter 4, the neck model was trained with muscle symmetry along the sagittal plane, i.e., muscles on the right side and the left side of the sagittal plane were considered symmetrical and the training was carried out such that the head in the neck model reached an extension or a flexion target position from the neutral upright position. Initially, the response of the RLMAC was compared with the volunteer response. The response of the RLMAC was comparable to the test data, as can be seen in Figure 5-8. The neck model was also tested under conditions different from that encountered in training to verify the extent to which the control model can follow the target signal. Figure 5-9 shows that the RLMAC can formulate the required muscle activations under novel control scenarios as well just from the states. The RLMAC can move the head similar to volunteers under target conditions that were not explicitly included in the RLMAC training (Figure 5-9). These results are important for general purpose use of RLMAC as it is tedious and computationally very expensive to train the RLMAC for all the possible head movement scenarios.

In the present chapter, the latencies in human response due to the time taken to process the sensory information before initiation of the muscle activations have not been considered while evaluating the response of the RLMAC. Similar latencies were also observed in previous volunteer studies for the normal responses in previous studies (Siegmund et al., 2001; Valls-Solé et al., 1999). The current study is focused on producing goal-directed head movements, and the effect of delays between the sensory input and initiation of the response was outside the scope of the current study. Incorporating the latencies in the RLMAC response will require the definition of such latencies in terms of RL state, reward, and control objectives. Siegmund et al. (2001) reported that the latencies when reacting to startle responses were shorter compared to the latencies observed in the present chapter. Recreating such human response to different sensory signals requires modifications in the MB model (activation dynamics parameter τ_{act} can be made to reduce the activation build time) as well as control architecture to capture such behavior. Future efforts can be focused on including the latencies for both 'go' and 'startle' signals on the overall control response.

5.5 Conclusions

The present chapter was focused on generating a dataset on volunteer head kinematics for validation of the response generated by the RLMAC for the neck model. The neck movements were characterized for extension, flexion, and axial rotations in eight male volunteers. The gathered data was used for validating the RLMAC in the present chapter and can serve as a source for future control studies as well. The simulation results show that RLMAC with the symmetric control architecture can synthesize a robust control response to a target signal. The control architecture will be incorporated for omnidirectional head and neck control in the neck model in the following chapter and the model responses will be evaluated with the data generated in the present chapter.

Chapter 6 – Control of head kinematics in multiple anatomical planes

In chapter 4, RLMAC was integrated with the neck model for head motion control along the sagittal plane. In the present chapter, the control framework developed in chapter 4 will be expanded for controlling the head kinematics of the neck model in all three anatomical planes. The control tasks involve stabilizing the head under gravity, as well as generating omnidirectional voluntary head kinematics – head flexion, extension, axial rotation, and lateral bending. The state and reward function were updated accordingly to include parameters required for the omnidirectional control of head rotations.

The RLMAC post training could perform the required control objectives, and the agent outputs were compared with the human subject data generated in chapter 5. The muscle activations developed by the RLMAC for different head positions were also evaluated and have been discussed. The ability of the trained RLMAC to generate head rotations different than what was used in training was also verified and presented in the current chapter. The training and simulation results demonstrate the feasibility of the use of reinforcement learning algorithms for postural stability and motor control in humans.

6.1 Introduction

The neck muscle forces control the movement and stabilize the cervical spine under various external loads, yet there is a very limited effort toward replicating the dynamic movements in computational studies. Very few studies have tried controlling head kinematics in multiple directions. Nemirovsky and Van Rooij (2010) implemented a feedback postural controller for roll, pitch, and yaw kinematics in head neck complex. The neck muscles were identified into groups based on an independent muscle pre-simulation study. Three PID controllers were used to control

the extension-flexion, axial rotation, and lateral bending of the head, and the gains for each controller were weighed to obtain the muscle synergy required for the head rotations. An open-loop co-contraction ratio was also used in the controller which could be adjusted depending on the state of the volunteers. However, the controller was only simulated for stabilizing the head in the sagittal plane, at different co-contraction ratios.

Ólafsdóttir et al. (2019) implemented a feedback control mechanism in the THUMS v3 50th percentile male model for postural control in presence of a constant 1g force field in the transverse plane. The neck muscles were divided into eight groups and the groups were activated based on the direction of the applied force and a predefined weighing function. The feedback-based muscle controller implemented in the study could limit the displacement of the head CoM under the applied forces and maintain spinal stability.

Correia et al. (2021) developed a feedback controller for regulating head movements in different impact scenarios. The muscles in the neck were grouped into four units responsible for movements in sagittal and lateral planes, and the groups were activated based on head angle and muscle length feedback. The gains of the controller were optimized considering different impact scenarios and could generate acceptable head kinematics for a range of impact severities.

Although the control models described above could generate the desired responses, the controller gains were tuned for the specific control tasks in these studies. The muscles were also assigned to different groups and the muscle synergy was predefined. However, previous volunteer studies have shown that the muscles behave differently based on the external loads and specific tasks, and thus it is not always accurate to generate identical activations for the grouped muscles (Keshner et al., 1989; Siegmund et al., 2006). The control objectives in the previous studies that considered

146

omnidirectional head kinematics were only limited to stabilizing the head, and the control mechanisms were not tested for the ability to generate voluntary rotations of the head.

In the present chapter, the RLMAC architecture developed in chapter 4 will be extended for control of head kinematics in the three anatomical planes (extension, flexion, axial rotations, and lateral bending). The RLMAC will independently activate each muscle in the neck model, and the muscle responses and the directional preferences will be compared with the data available from previous volunteer studies.

6.2 Methodology

The RLMAC developed previously (chapter 4) was updated to include the parameters essential for omnidirectional control. For the all-DoF study, both the vestibulocollic reflex (VCR) and cervicocollic reflex (CCR) were considered (Cullen, 2012; Goldberg and Cullen, 2011; Keshner, 2009). The RLMAC was trained to output the muscle activations for all 46 muscles of the neck model to rotate the head to the desired position in all the three anatomical DoFs.

6.2.1 Muscle control framework

Similar to the symmetrical control framework, a twin-delayed deep deterministic policy gradient (TD3) is used as the agent for the RLMAC (Fujimoto et al., 2018). The TD3 agent incorporates one actor network and two critic networks. The actor-network takes the state parameters as inputs and outputs the neural stimulations (u_t) of all 46 muscles. The activation dynamics relation transforms neural stimulations into muscle activations (a_t) (Zajac, 1989). The two critic networks of the TD3 agent have identical architecture and both of them map the network inputs to the Q-value. The network input consists of the state parameters and actions from the previous time step, and the lower of the two Q-values is used to update the networks every two time-step. The state is

defined by the sensory, VCR, and CCR parameters, including the head kinematics terms and neck

proprioceptor information (Table 6-1).

Table 6-1: State and reward considerations for training

	State	Reward
VCR parameters	- Skull CoM translation (S_x, S_y, S_z) - Skull CoM linear velocity (V_x, V_y, V_z) - Skull CoM angle $(\theta_x, \theta_y, \theta_z)$ - Skull CoM angle error $(\varepsilon_x, \varepsilon_y, \varepsilon_z)$ - Skull CoM rotational velocity $(\omega_y, \omega_y, \omega_y)$	$(1 - \mathrm{Sin}(2\varepsilon_y))^2 + (1 - \mathrm{Sin}(2\varepsilon_y))^2 + (1 - \mathrm{Sin}(2\varepsilon_y))^2 - \alpha\,(V_x + V_y + V_z) - \beta\,(\omega_x + \omega_y + \omega_z) - \eta\,\sum \mathrm{Joint\ energy} + \mu\,(\text{symmetric factor})$
CCR parameters	- Joint displacement - Joint velocity - Muscle activation	Equation 6-1

6.2.2 State and reward

The VCR terms in the state include head translation and rotational kinematics in the three planes

of motion. The VCR parameters consist of the head CoM translation displacements and velocities,

rotational displacements and velocities, and head position errors (Table 6-1). In the present

chapter, the RLMAC attempts to move the head to target positions in all three DoFs, therefore the

angular error of all three planes is being considered for the state.

The CCR stabilizes the spine during head movements or in the presence of any external loads. For

the all-DoF study, only the CCR scenario corresponding to the joint receptors is being considered.

Although some previous studies have considered muscle spindle lengths for CCR (Correia et al.,

2021; Happee et al., 2017a; Olafsdottir et al., 2013; Zheng et al., 2021), we found that using joint

displacements and velocities as state parameters of the RLMAC improve the spinal stability of the

neck model (chapter 4). The muscle length feedback method also requires the length change

information of all the muscle strands in the neck model, increasing the sizes of the actor and critic

network of the TD3 agent and as a result, will require more exploration and training time. Thus,

joint displacement and velocity information have been considered as the CCR parameters for the all-DoF control. The agent is also provided information on the muscle activation corresponding to the muscle tendons feedback.

The reward function was formulated such that it enables the RLMAC to minimize the position error in the three directions ($\sum(1 - Sin(2\varepsilon_i))^2 | i = X, Y, Z$). The reward function also penalizes the RLMAC for velocity (both translation and rotation) to stabilize the head at the target position. The joint energy minimization term of the reward function accounts for the CCR and stabilizes the spine during movement and considers the joint energy expenditure along all the 6 DoFs.

The symmetry term (minimize $|u_{right} - u_{left}|$) is required to equally actuate the muscles on the left and right side during sagittal motion and mirror the activations while turning left or right or during lateral bending. In some previous studies, the symmetry term was included as a loss function of the actor network (Yu et al., 2018), however, for the current study, the symmetry term is considered in the reward function, and evoked while moving the head in the sagittal plane.

6.2.3 Architecture of the actor and critic networks

Both the actor and the critic networks in the sagittal control model (chapter 4) used a feedforward network with one hidden layer. A single hidden layer actor and critic networks was found to be insufficient in the initial training runs as both the state parameters and the action space in the all-DoF control in this chapter was larger than that during the sagittal control that used parameters for only 3 DoFs. Therefore, an additional hidden layer was added between the initial and the final layers in the actor network. In the actor network, the outputs of the initial and the hidden layers were activated by the rectified linear unit (ReLU) function (Hara et al., 2015), and the output (neural stimulation) values were limited within the 0-1 range using tanh and scaling functions in

149

series. The final layer of the actor network had 46 nodes outputting the stimulations for all the

muscles in the neck model. In the critic layer, the state parameter values were activated with the

ReLU function before passing to the first hidden layer, whereas the actions were directly input to

the hidden layer without any processing or activation. Both the critic networks have the same

architecture and output the scalar Q-value individually.

(a)

(b)
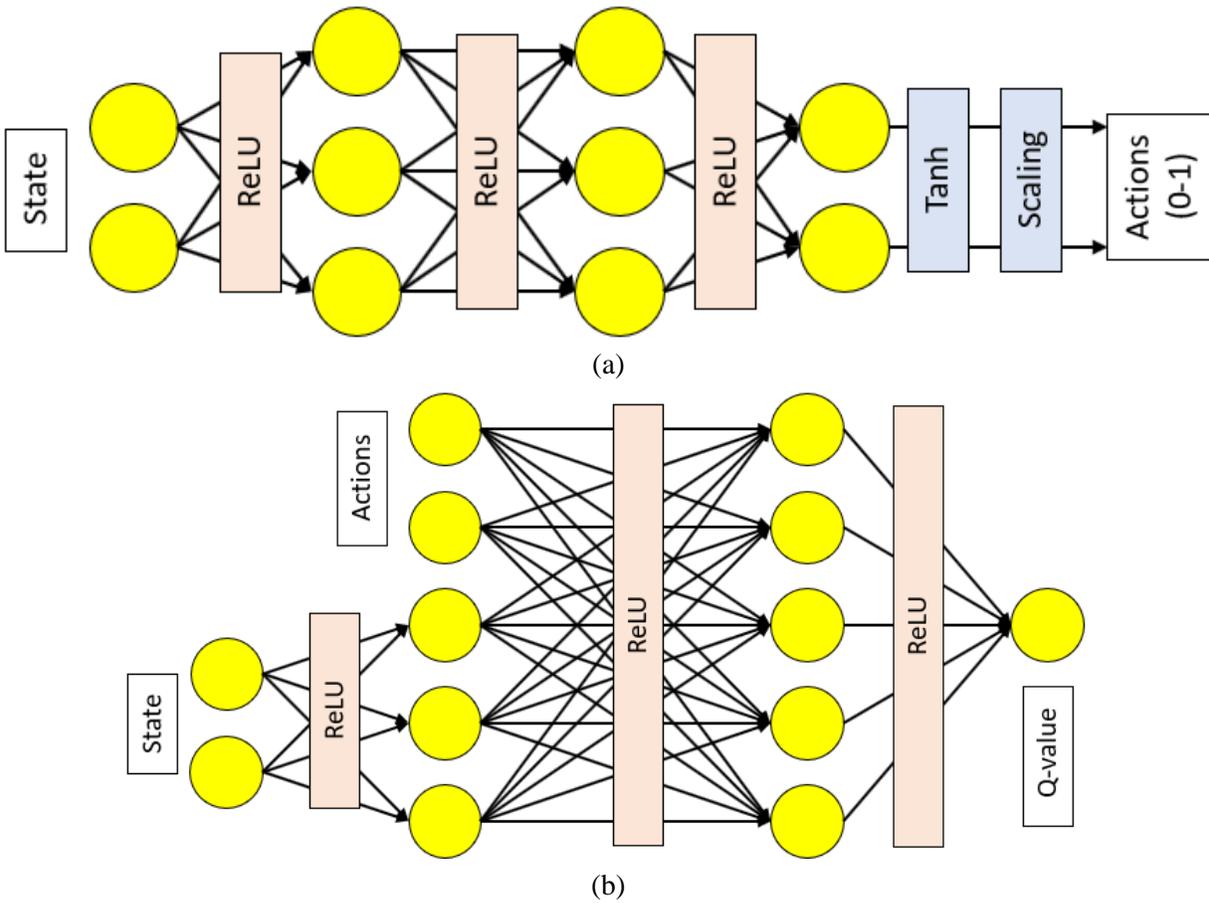
Figure 6-1: Architecture of the (a) Actor network (b) Critic network for all-DoF control. Both the critic networks in the RLMAC have identical architecture.

6.2.4 Training for omnidirectional control

The training was carried out to move the head to a desired angle in the sagittal, lateral, and

transverse planes. The head was trained to stabilize under gravity in 20% of the iteration and bend

laterally in another 20%. The RLMAC was trained for sagittal and axial rotations of the head in

30% of the iterations respectively. In one iteration, the RLMAC was tasked to rotate the head in only one plane, i.e., the training was carried out such that while the target angle in one plane was set to a non-zero value, the other two planes had 0 set as the target angle. For the sagittal plane, the target angles were set between -25 degrees (flexion) to 30 degrees (extension). For the axial rotations, the targets were set between -30 degrees to +30 degrees and between -20 degrees to +20 degrees for lateral bending. For each scenario, the target angles were sampled in multiples of 5 degrees from the target space.

6.3 Results

The neck model with RLMAC for all-DOF control was run for 15000 simulations. The RLMAC was trained to move the head to a target position from the neutral upright posture. No prior assumptions were made on the muscle synergy, the actor network of the RLMAC actuated all the muscles in the neck model individually. Figure 6-1 displays the variation of the average reward during the training.
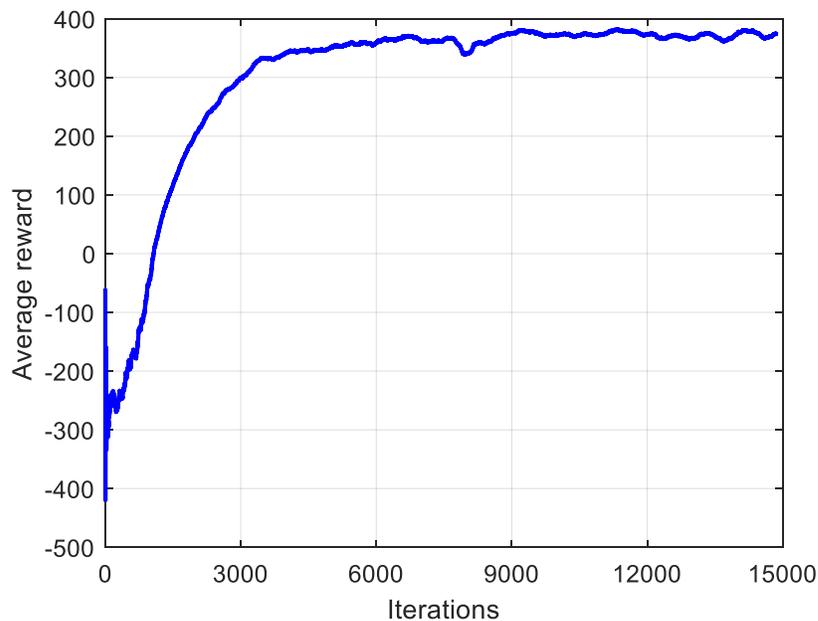
Figure 6-2: Average reward during the training for all-DoF control.

151

6.2.1 Stabilization under gravity

The trained agent (Figure 6-2) could maintain the neutral posture of the head under gravity. The head was stable with minimum compression or buckling of the spinal column. Figure 6-3 displays the orientation of the head and spine at the stable position. The head displacements at the final stabilized position were less than 2 mm from the neutral position in all directions. For the stabilization run, the angle errors (ε_x, ε_y, ε_z) were measured from 0^o in all three anatomical planes, and at the final position, the maximum error was 2.2^o in the axial rotation (Figure 6-4). The head took around 400 ms to completely stabilize.

(a)   *(b)*   *(c)*
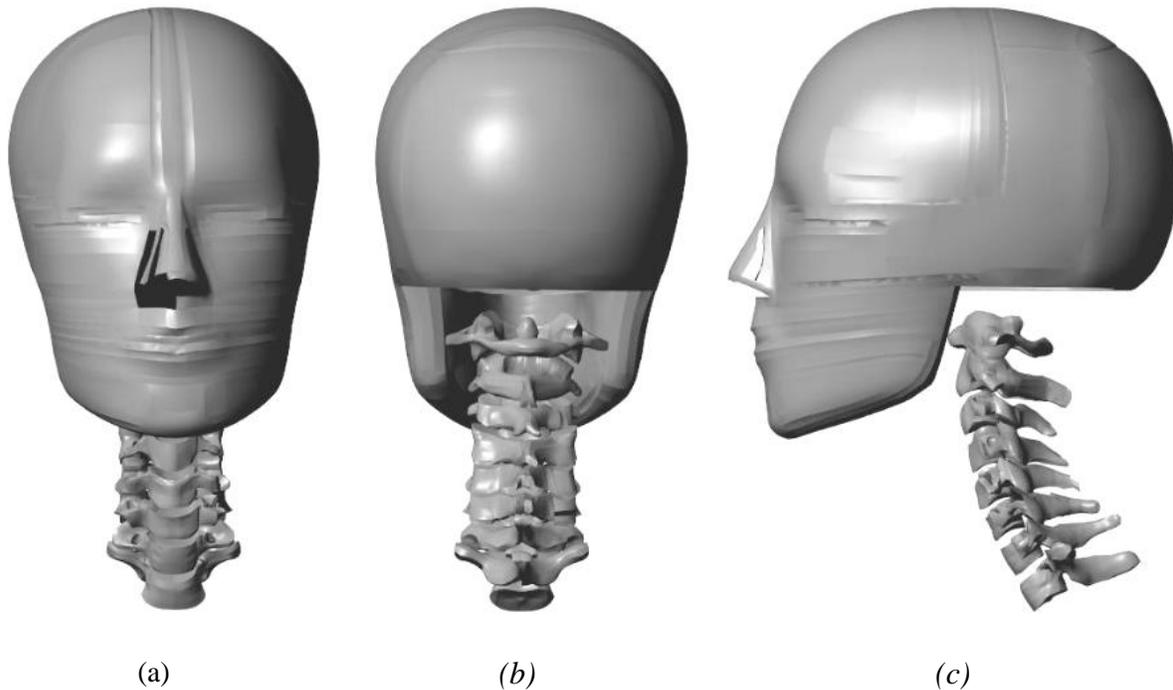
Figure 6-3: The stable posture of the head under gravity (a) Frontal (b) Rear (c) Side view.
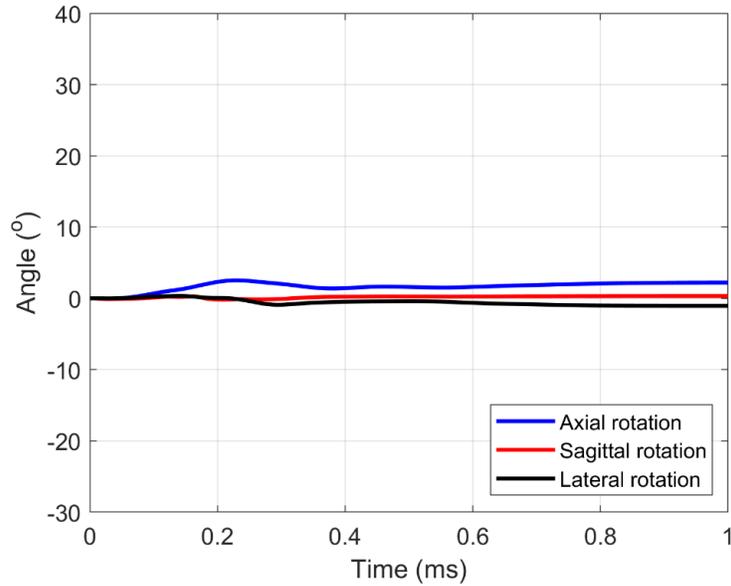
Figure 6-4: The head angle-time history during the stabilization run in the three planes.

The activation profiles of the sternohyoid and trapezius muscles during the simulation are shown in Figure 6-5. Apart from these muscles, the sternocleidomastoid, splenius cervicis, and splenius capitis were also actuated by the RLMAC. The muscles have nearly symmetric activations as the symmetry factor was prescribed in the reward.
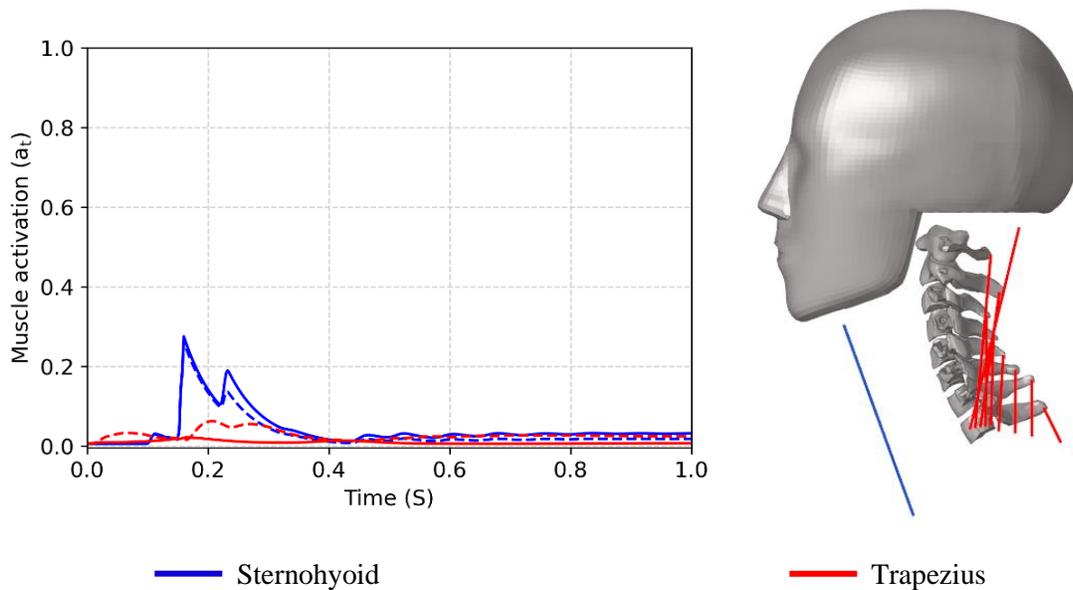
Figure 6-5: Co-contraction of the neck muscles (bold – left, dotted – right) in the stabilization simulation.

153

In the stabilization run, the rectus capitis and the multifidus were also activated, as the muscles are mostly associated with stabilizing the gaze and the spine. The rectus capitis muscle in general has high activations for most of the simulation (Figure 6-6).

Rectus capitis minor ———— Multifidus
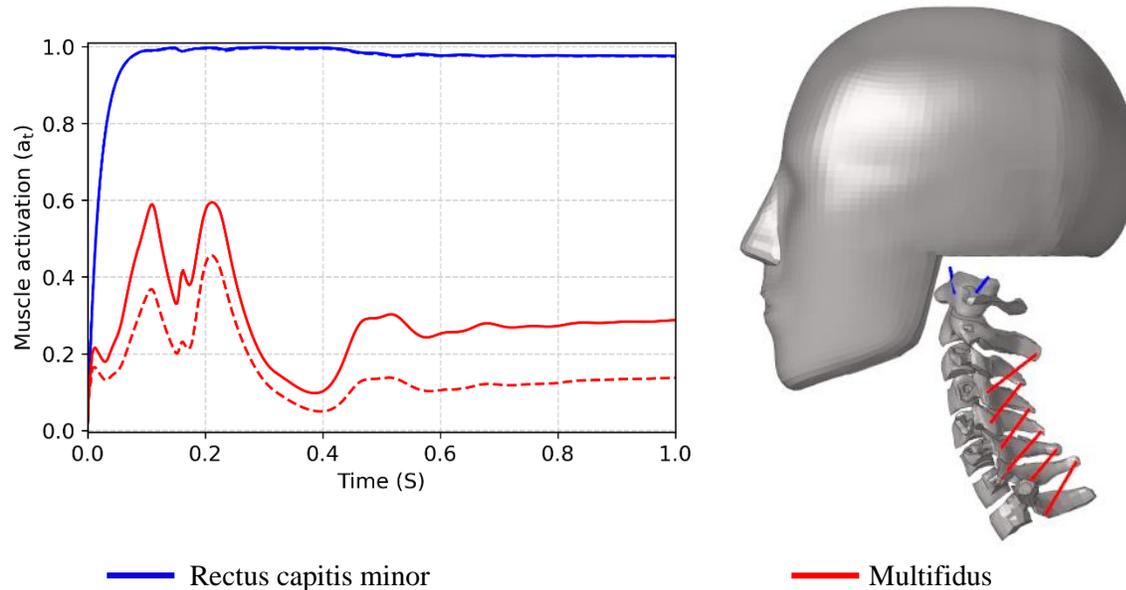
Figure 6-6: Activation in the rectus capitis minor and the multifidus during the stabilization simulation.

6.2.2 Extension of the head and neck

Initially, the head was stabilized for 400 ms before the target of 20° extension was prescribed. The trained agent could generate the extension motion of the head (Figure 6-7). The head could reach the final position without much lateral deflection (Figure 6-7a-b). Figure 6-7c shows the angle-time history of the head in the three planes when a target of 20° extension was prescribed. The head could move to the target angle in the sagittal plane while maintaining the neutral position in the other two axes.

154
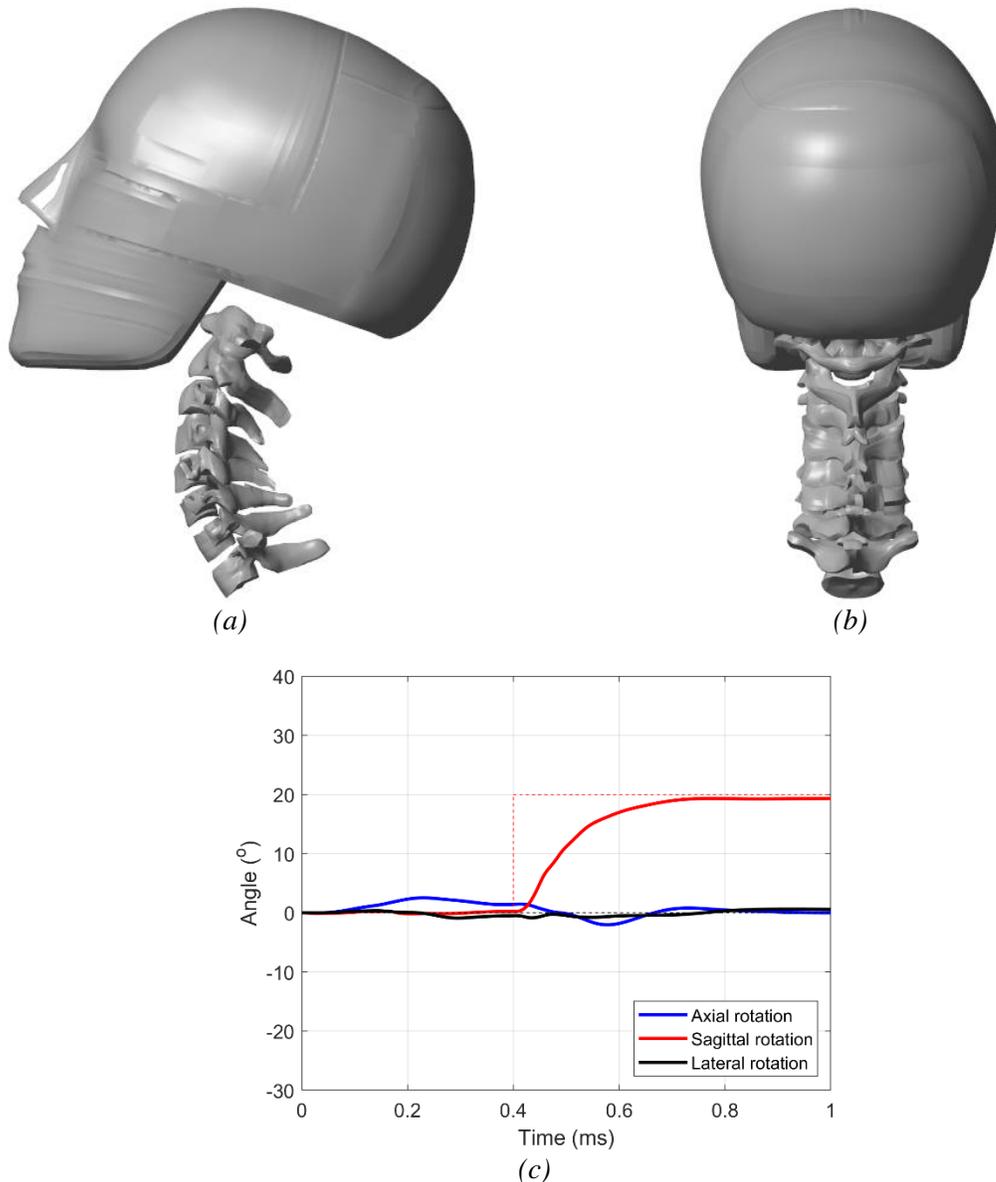
Figure 6-7: Extension motion of the head (a) Side view (b) Rear view (c) Angle-time profile of the motion (bold – head angle, dotted – target angle).

The activation histories of the sternohyoid and trapezius muscles for the extension head movement are shown in Figure 6-8. The trapezius muscles, splenius capitis, and splenius cervicis act as the agonist during the head extension whereas the omohyoid, sternohyoid, and sternocleidomastoid act as the antagonist. The activation patterns follow a triphasic pattern where activation of the antagonist and a second burst of the splenius capitis activation are required to dampen the head motion at the 20° target (Hannaford and Stark, 1983; Happee, 1992; Marsden et al., 1983). During

155

the extension, the agonists and antagonist muscles were activated on either side of the midsagittal plane, however, there were some minor differences in the magnitudes of the activation on the left and right muscles.
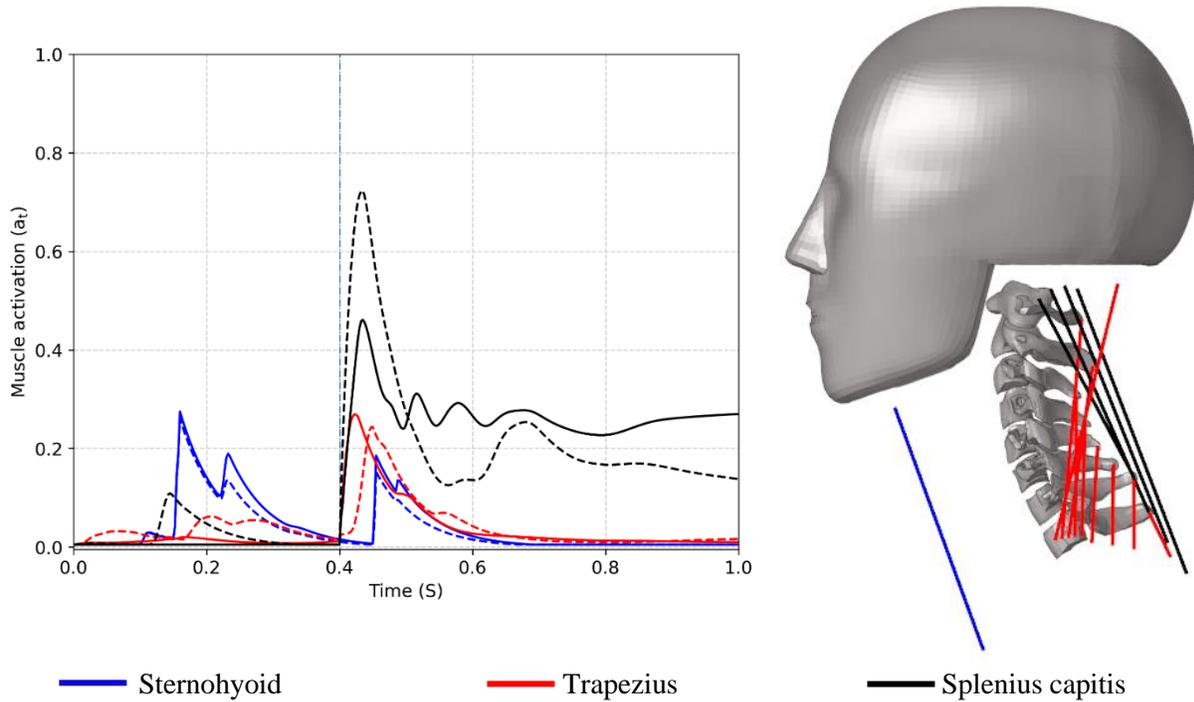
Figure 6-8: Neck muscle activations (bold – left, dotted – right) during the 20º extension. The horizontal line at 0.4 s highlights the onset of the extension target signal.

The trapezius and rectus capitis minor were the other muscles activated during the extension simulations (Figure 6-9). While the rectus capitis minor had high activation throughout the simulation, the multifidus was fully activated by the RLMAC when the extension target was prescribed at 0.4s.

156

Figure 6-9: Activation in the rectus capitis minor and the multifidus during 20° extension. The extension target was prescribed at 0.4 s.

6.2.3 Flexion of the head and neck

Similar to the extension simulation, the neck model was stabilized for 400 ms before applying the 20° flexion target as a step function. The head reached the final position while maintaining the stability of the spine. The final posture of the neck model after completing the 20° flexion motion is shown in Figure 6-10.

(a)   *(b)*

Figure 6-10: Final position of the head for the flexion motion (a) Side view (b) Rear view.

157

The angle time profile of the head CoM during the simulation is shown in Figure 6-11. The RLMAC stabilized the head at the target sagittal position while maintaining the initial angle (error $< 1.5^\circ$) in the transverse and lateral planes.
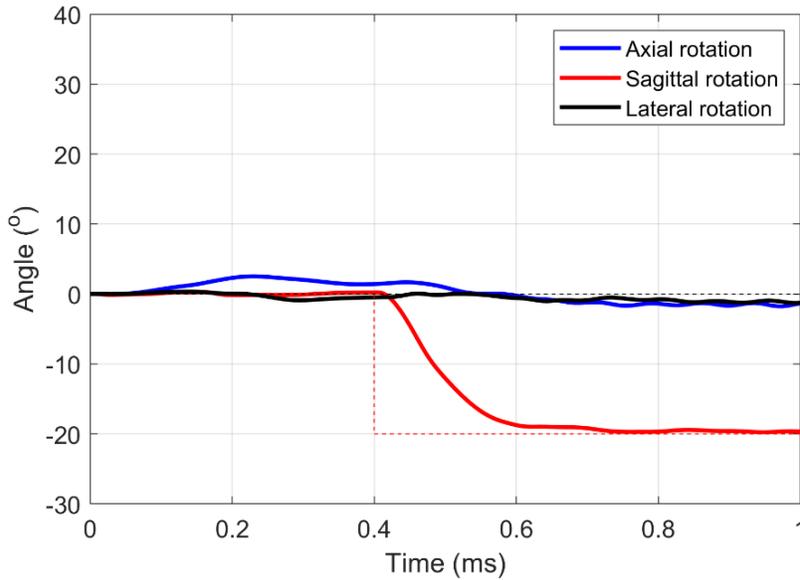
Figure 6-11: Angle-time profile of the 20° flexion (bold–head angle, dotted–target angle).
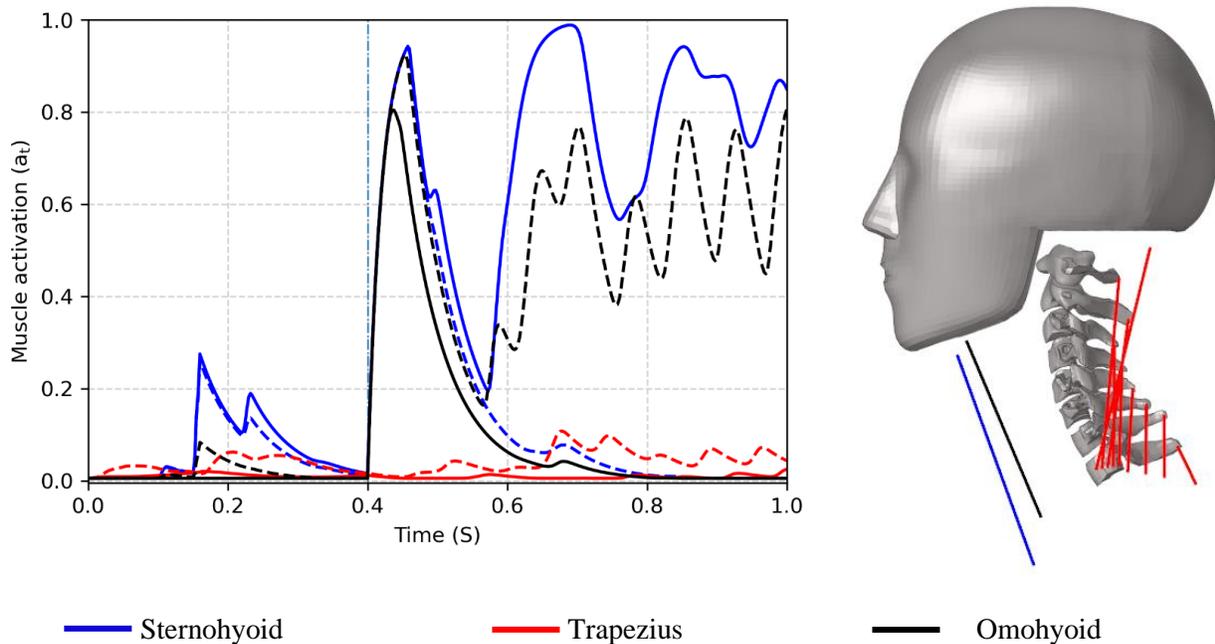
Figure 6-12: Neck muscle activations (bold – left, dotted – right) during the 20° flexion. The horizontal line at 0.4 s highlights the onset of the flexion target signal.

158

Figure 6-12 shows the activation profiles of the neck muscles during the 20° flexion. The sternohyoid, omohyoid and sternocleidomastoid are the agonists, while the trapezius, splenius capititis, and semispinalis capitis are the antagonists for flexion. Although for most muscles, the activations are symmetrical during the head movements, at the final flexed position the activations of the left sternohyoid and the right omohyoid were higher than that of the other sides.

6.2.4 Axial rotation

A target angle of 30° was assigned in the transverse plane whereas 0° target was maintained at the other two planes. Before assigning the axial target, the head was stabilized for 400 ms. The trained model could synthesize the axial rotations of the head at either side of the sagittal plane. The spine was stable while rotating the head axially on both the left and right sides, and the head was stable at the target in both rotation cases (Figure 6-12).

(a)   (b)
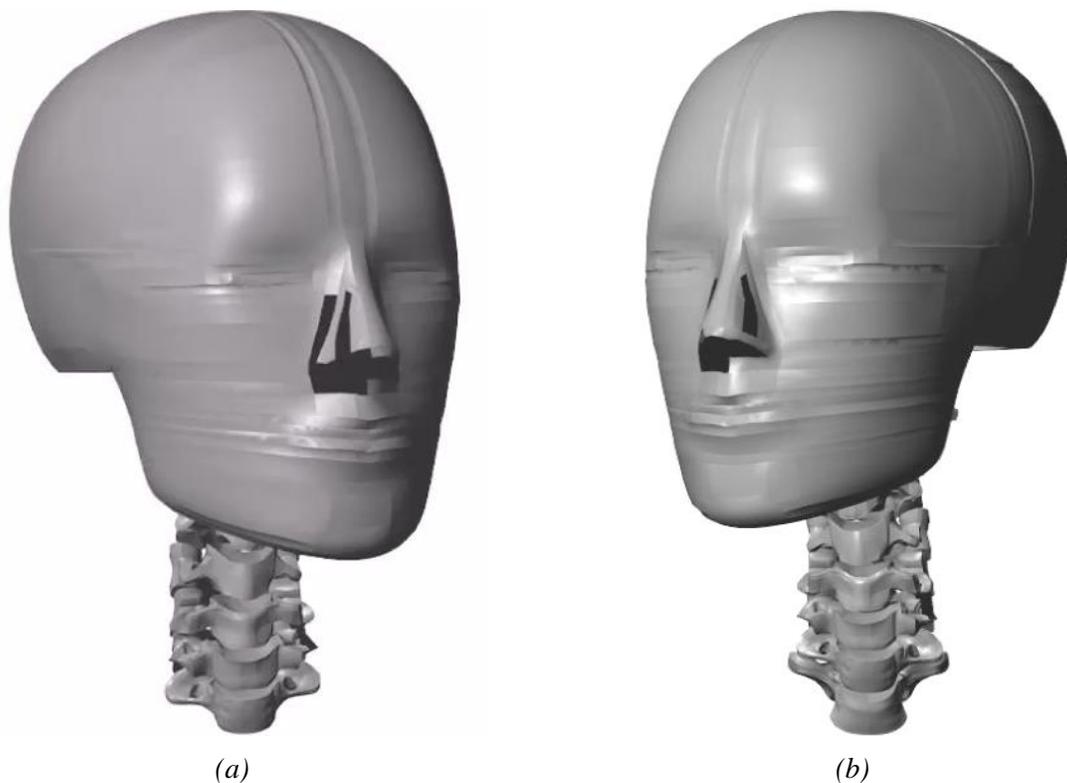
Figure 6-13: 30° axial rotation of the head in front view (a) Left rotation (b) Right rotation.

159

Figure 6-14 show the angle time profile of the head CoM during the axial rotation simulation for both the left and right sides. The head could reach the target angle in the transverse plane while maintaining an almost neutral angle in the other two planes. While turning left, the head was within 2.5^o of the target position after stabilizing, whereas the final head angle while turning right was closer to 30^o.

(a)   *(b)*
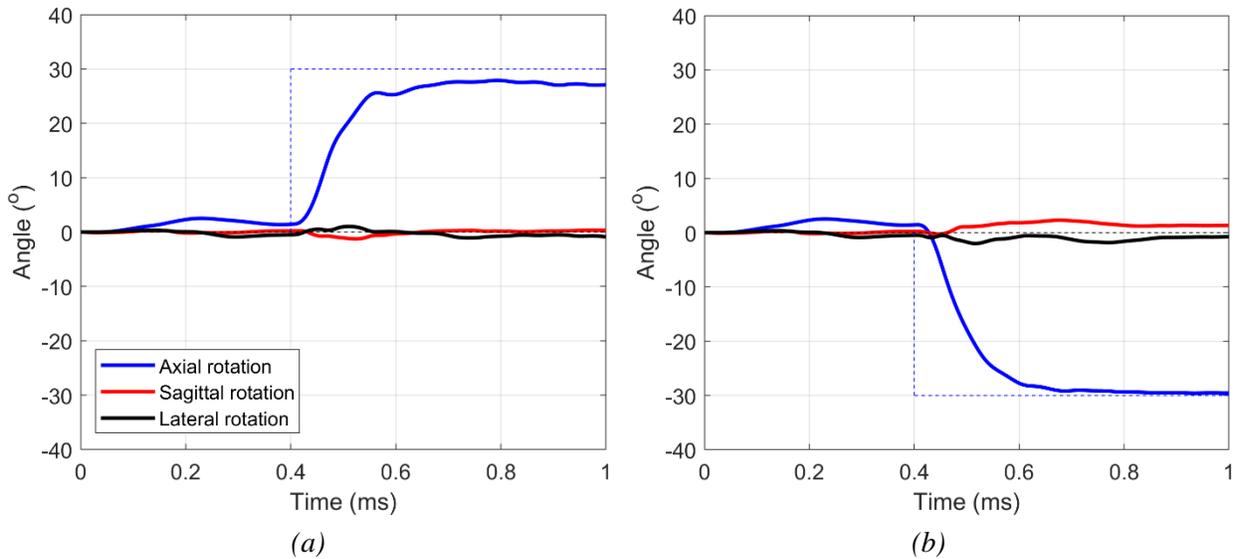
Figure 6-14: Angle-time profile of the 30^o axial rotation (a) Left turn (b) Right turn (bold–head angle, dotted–target angle).

Anterior scalene   Sterno-cleido mastoid
(a)   *(b)*
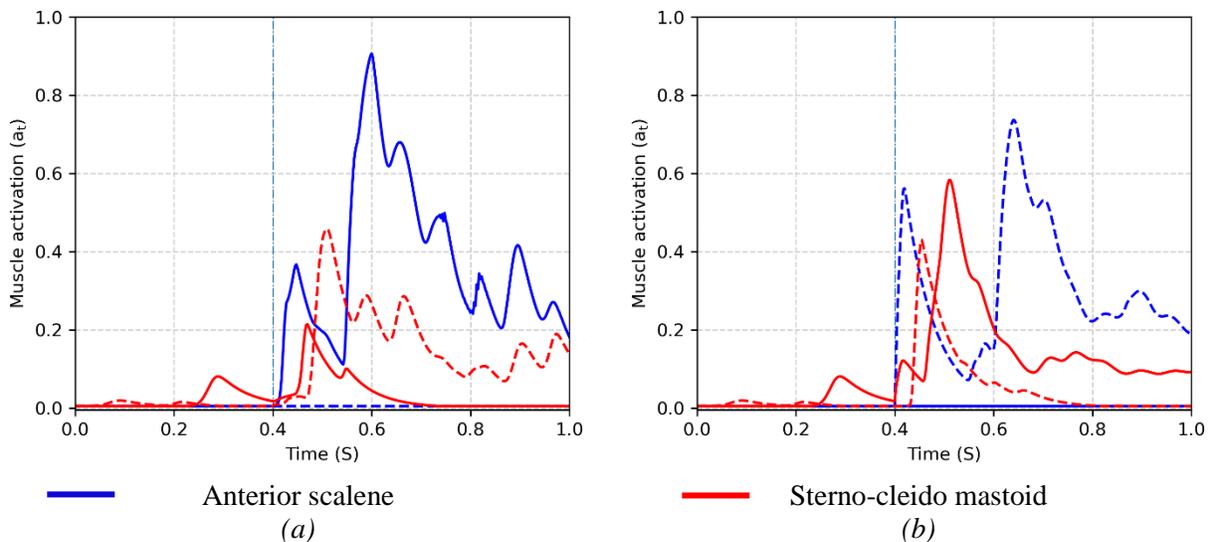
Figure 6-15: Neck muscle activations (bold – left, dotted – right) during the 30^o flexion (a) Left turn (b) Right turn. The horizontal line at 0.4 s highlights the onset of the target signal.

160

Figure 6-15 displays the activation profiles of the left and right anterior scalene and sternocleidomastoid. The muscles on the left of the sagittal planes act as agonists for the left turn and the right muscles act as antagonists, and the roles of the muscles reverse during the left turn. The agonist and the antagonist muscles follow a tri-phasic pattern. Iliocostalis cervicis, sternohyoid, mid-scalene, and levator scapula were other muscles that were activated during the simulations. The trapezius muscles acted only as antagonists, i.e., the left muscles were activated to dampen the right turn and the right muscles were activated to dampen the left turn. The right post-scalene was activated during the right turn but similar activation of the left post-scalene was not observed during the left turn.

6.2.5 Lateral bending

The ability of the trained RLMAC to produce lateral bending was also evaluated. The neck model after stabilizing, was assigned a target angle of 20^o in the lateral lane, while still maintaining the neutral target of 0^o in the other planes.

<div align="center">(a)</div> <div align="center">(b)</div>
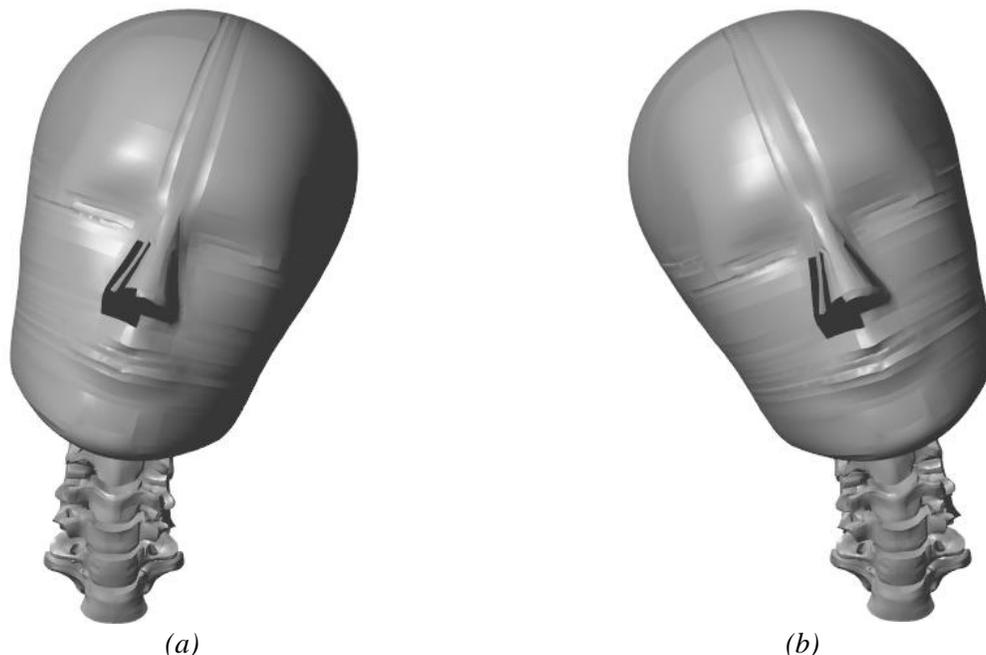
Figure 6-16: 20^o lateral bend of the head in front view (a) Left bend (b) Right bend.

The angle time history of the head CoM is shown in Figure 6-17. The head could reach the target angle on both sides in the lateral planes. The trained agent stabilized the head within 2.3° of the target of 20°. For the left turn, the head undershot the target, whereas the head overshot the target during the right turn.

(a)   *(b)*
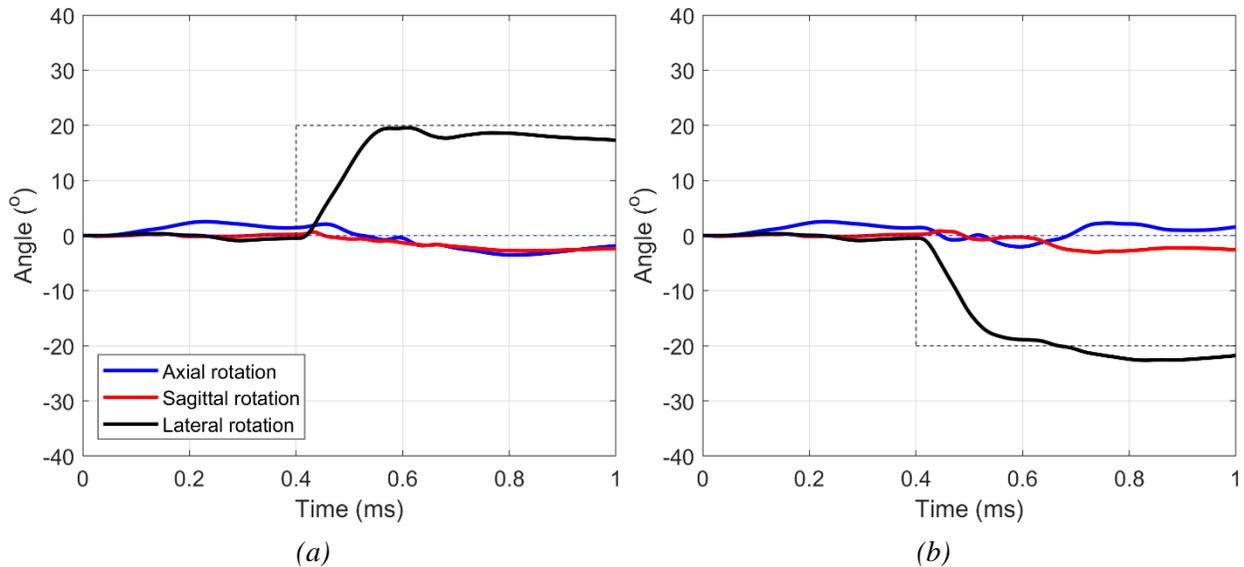
Figure 6-17: Head CoM angle time history of 20° lateral bend (a) Left bend (b) Right bend (bold–head angle, dotted–target angle).

Trapezius   Sterno-cleido mastoid   Splenius capitis

(a)   *(b)*
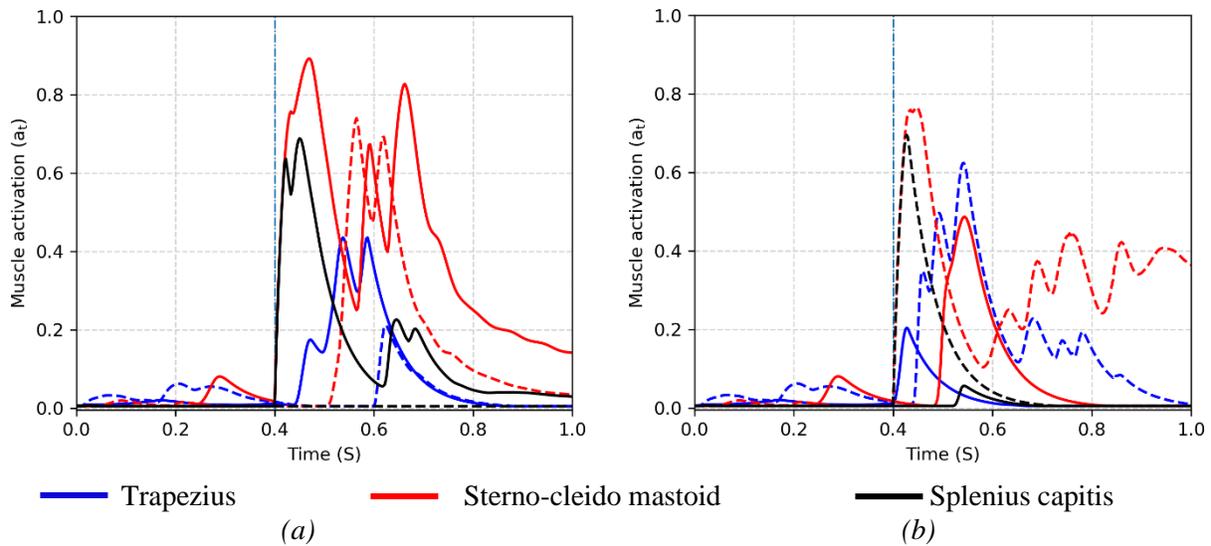
Figure 6-18: Neck muscle activations (bold – left, dotted – right) during the 30° flexion (a) Left turn (b) Right turn. The horizontal line at 0.4 s highlights the onset of the target signal.

162

Figure 6-18 shows the activation profiles of the neck muscles during the head lateral bending. The trapezius, sternocleidomastoid, and splenius capitis were the major muscles active during the lateral bending on both sides. The right mid-scalene was activated during the bending on the right side, however, the corresponding left muscle was not activated during the left bend. The sternohyoid and omohyoid muscles also had minor activations during the bending simulations.

Overall, the trained RLMAC could rotate the head to the desired position in the primary axes of motion while maintaining the head stability at the final position.

6.2.5 Comparison with test data

The head angular response from the trained RLMAC was compared to the human subject data obtained in chapter 5. The head kinematics data were obtained for extension, flexion, and left axial rotation. During the training, the RLMAC was assigned to move the head from the neutral upright posture to a target position. The neck model was initially stabilized for 400 ms before prescribing the target head angle as a step function.

In a training iteration, the target angle for the head movement was specified along individual axes, while the target angle in the other two axes was maintained at neutral position (0^o). Figure 6-19 shows the response of the RLMAC while simulating the goal-directed head movements when the prescribed target angle was similar to the training scenarios. The output angle-time curve was compared with the test data. As the human subjects were tested while rotating the head only on the left side in the transverse plane, the angle-time history obtained from the right axial simulation was flipped before comparing it with the test data. The overall kinematics of the head in the simulations with the trained RLMAC while replicating the training kinematics is similar to the tests (Figure 6-11).
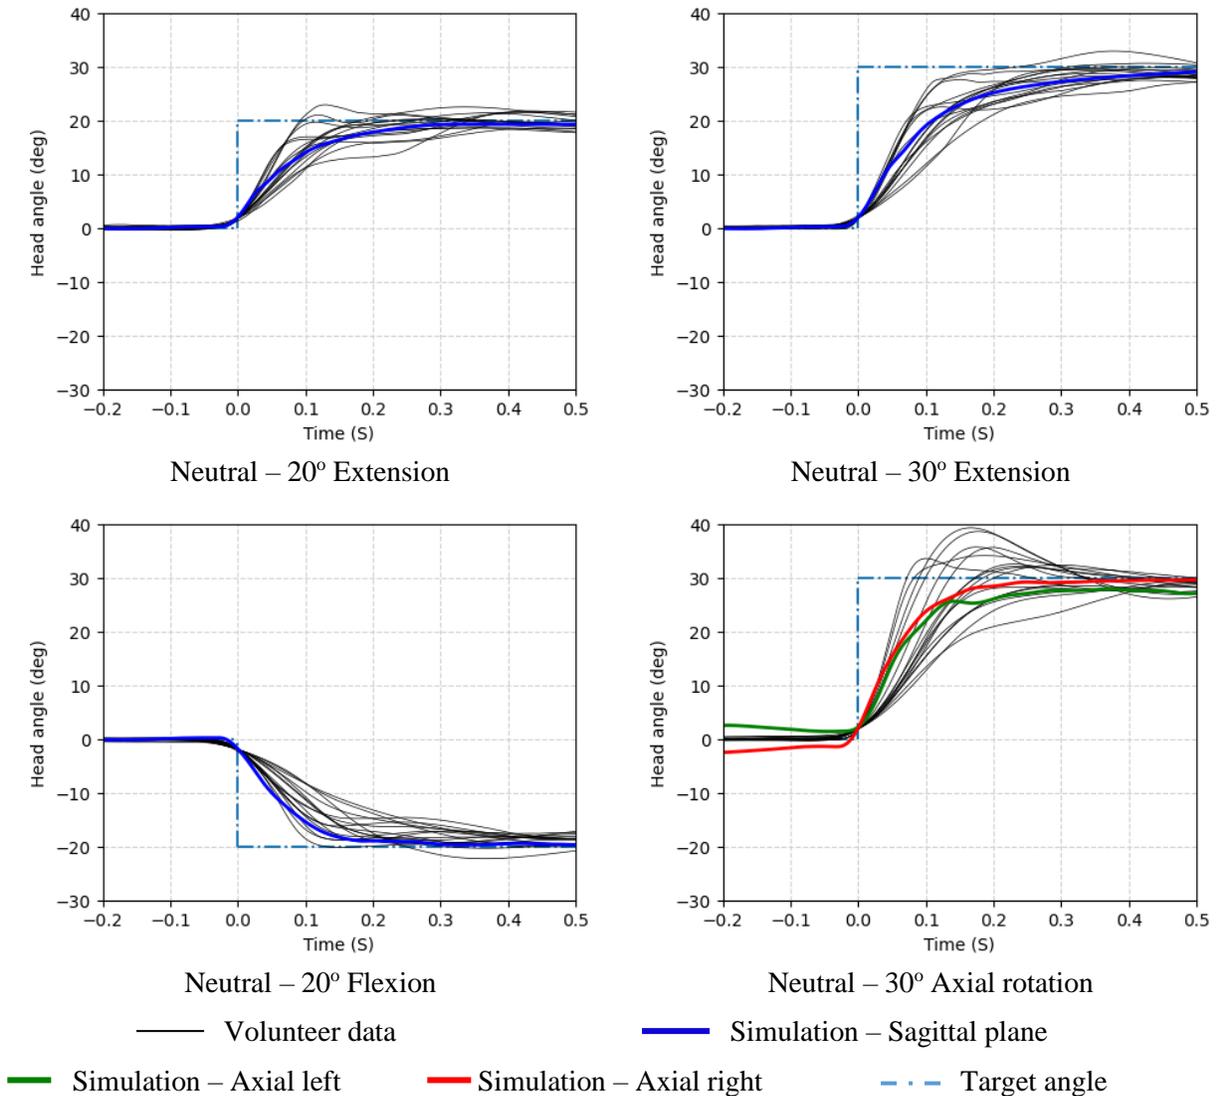
Figure 6-19: Comparison of the output of trained RLMAC with volunteer data for the training scenarios. The angle-time plots have been adjusted such that at 0 s time, the head sagittal angle is 2°.

The trained RLMAC was able to generate the target head movements even for those situations that the head did not encounter during the training. The initial and target head positions prescribed to the RLMAC along with the head CoM angle time profile are shown in Figure 6-12. From the plots in Figure 6-12, it is evident that the head in the neck model could follow the target signals even if the RLMAC was not trained for a similar scenario.
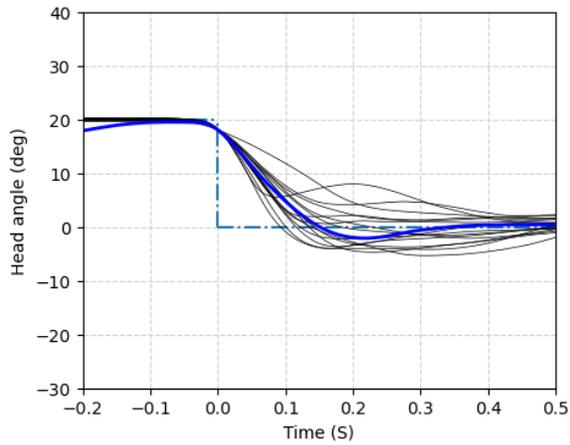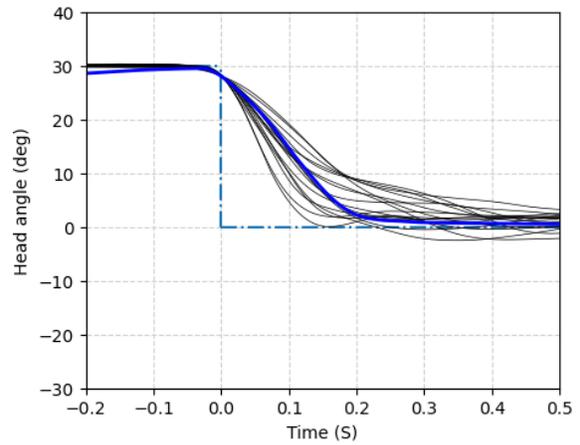
20° Extension - Neutral   30° Extension - Neutral
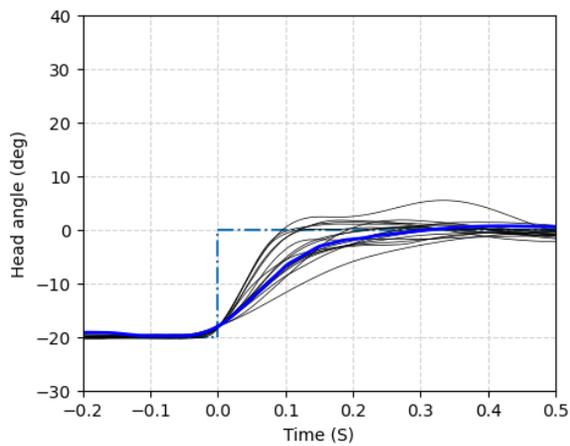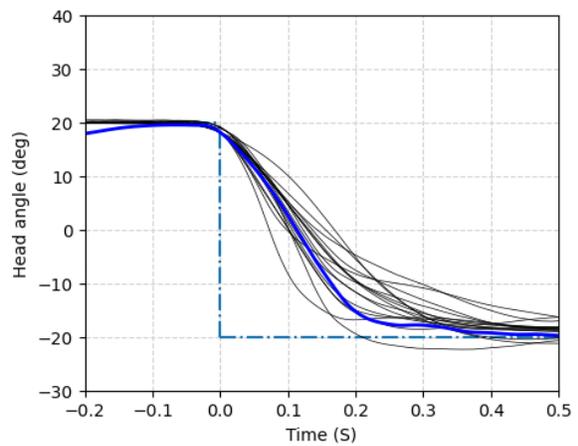
20° Flexion - Neutral   20° Extension - 20° Flexion
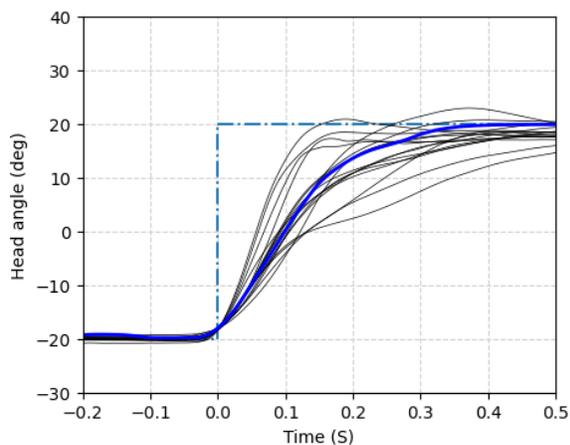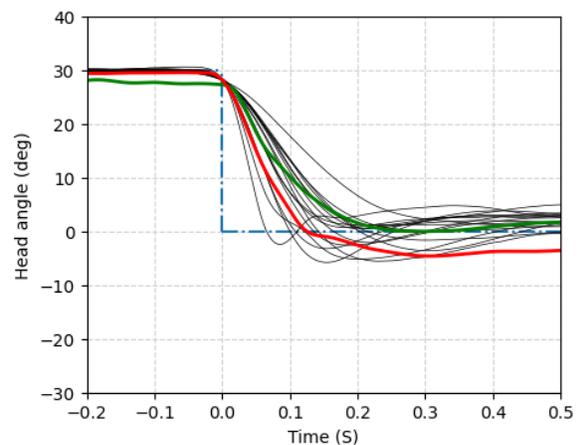
20° Flexion - 20° Extension   30° Axial rotation – Neutral

——— Volunteer data   ——— Simulation – Sagittal plane

——— Simulation – Axial left   ——— Simulation – Axial right   – · – Target angle

Figure 6-20: Comparison of the output of trained RLMAC with volunteer data for the scenarios not included in the training. The angle-time plots have been adjusted such that at 0 s time, the head sagittal angle is 2°.

The head angle time curve while rotating back to the neutral from the initial position of the right 30° axial was flipped before comparing it with the test data. The head moved to the neutral position from the right side faster than from the left side. The axial to neutral runs took a longer time to stabilize than the simulations in the sagittal plane.

6.2.5 Combined head movements along multiple axes

In the training, the head rotation was prescribed in one of the three axes. In this section, the ability of the RLMAC to respond to scenarios where the target signals are prescribed in more than one axes was evaluated. Figure 6-21 shows the response of the RLMAC when a combination of axial rotation and flexion is assigned at different times as step functions during the simulation. When the axial and flexion targets were applied simultaneously (time between 0.5s to 1.5s in Figure 6-21), the RLMAC response undershot the targets in both planes.
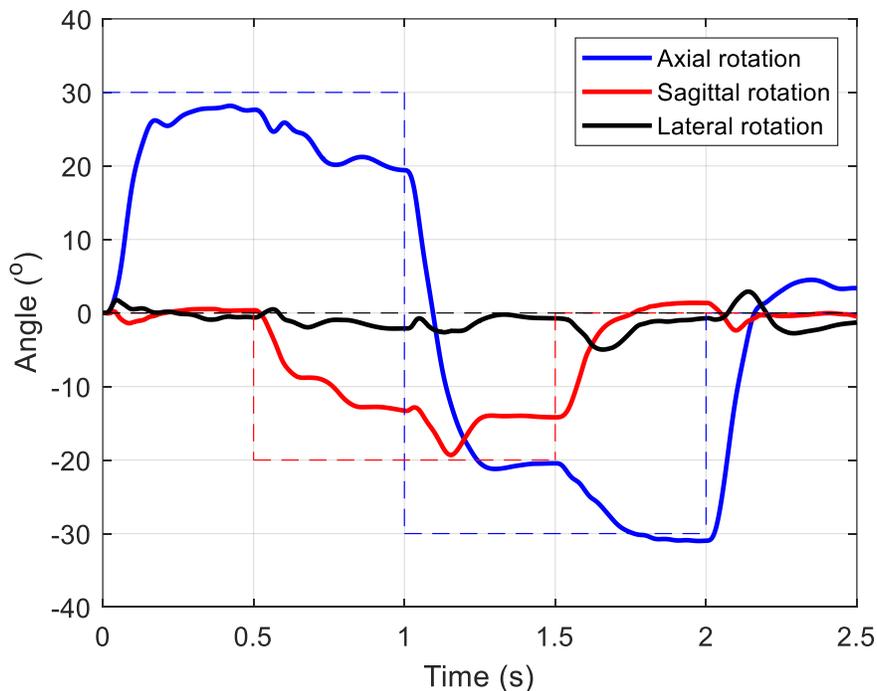
Figure 6-21: Angle time history for prescribed axial rotation and flexion (bold–head angle, dotted–target angle).

Figure 6-22 shows the response of the RLMAC in combined axial rotation and extension. While in extension, the RLMAC could maintain the axial head position nearer to the target signal. The head angle in the lateral plane was maintained near the prescribed target angle of 0^o.
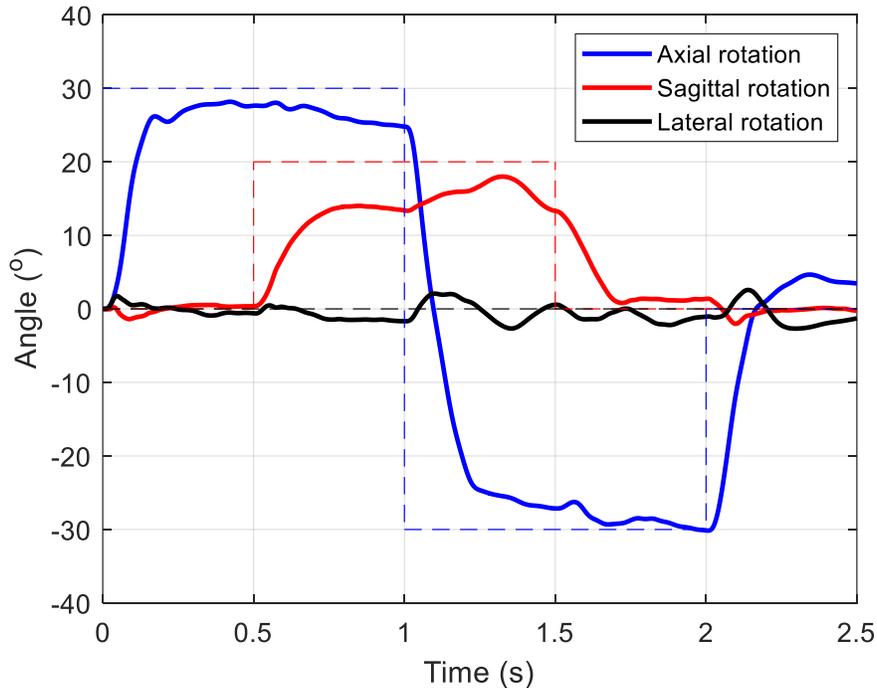
Figure 6-22: Angle time history for prescribed axial rotation and flexion (bold–head angle, dotted–target angle).

Figure 6-23 displays the head kinematics in the combined extension axial rotation case. The RLMAC could respond to the change in the target signals during the simulation. Between 0.5 to 1.5 s of the simulations, the head could rotate from left to right while being flexed (Figure 6-21) or extended (Figure 6-22). During the simulations, the neck was upright and stable even when the head was responding to the target signals in the sagittal and transverse planes. The head could return to the neutral position at the end of the simulation.
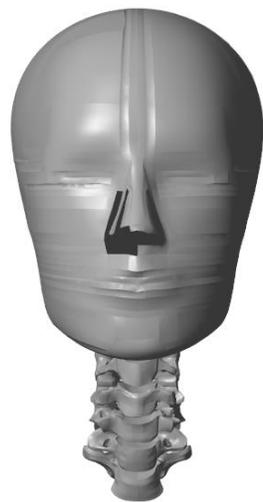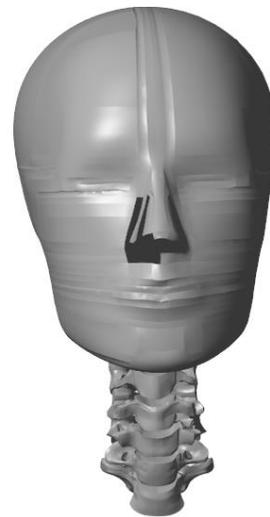
Figure 6-23: The head kinematics at different points of time during the combined extension and axial rotation simulation.

6.4 Discussion

The present chapter details the use of reinforcement learning (RL) controller for generating human head kinematics in multiple directions. The RLMAC integrated with the neck model in chapter 4 was modified to include the state parameters that could make it possible to control head kinematics in the three anatomical planes (Table 6-1). The architecture of the actor and critic network was also modified to add one additional hidden layer considering the large state parameters that the networks had to process (Figure 6-1). The reward function was also expanded to account for the change in state parameters and the overall control objective (Equation 6-1).

The RLMAC was devised to individually output the neural stimulations of the 46 muscles in the neck model. No predefined grouping or synergy of the muscles was considered in the RL architecture. A symmetry term was added to the reward function to ensure symmetry in the neck activations when the target signal prescribed in the sagittal plane during training (Equation 6-1). Without the symmetry term, the RLMAC was found to learn to output the correct kinematics response, but the activation patterns on the left and right muscles of the mid-sagittal planes were completely different.

With the systems of state and reward in Table 6-1, the RLMAC was trained to generate fast goal-directed kinematics in the three anatomical planes. Previously, very few studies have tried to develop a control model for omnidirectional head kinematics (Correia et al., 2021; Ólafsdóttir et al., 2019). However, these studies were mostly limited to maintaining head stability under external perturbations and did not focus on synthesizing targeted head kinematics. La Barbera et al. (2022) previously used an RL-based control strategy to control the neck posture in an Ostrich MB model, however, in that study the control task was mainly performed to fine-tune the stiffness of the neck joints, and the accuracy of the neck movements was not verified.

169

The RLMAC took nearly 15000 iterations to arrive at a desired control strategy. The RLMAC after training was able to maintain the stability of the neck model with minimum co-contractions of the neck muscles. The RLMAC could also produce the goal-directed motion of the neck in the three planes. The target angles used in the study varied from 20^o flexion to 30^o extension, 20^o lateral bending, and 30^o axial rotation on either side. The target angles were selected such that the angles could be reached while the T1 is constrained. Previously in a volunteer study, it was found that axial rotations are accompanied by a slight lateral bend, and similarly, lateral bends are accompanied by rotations in other planes (Schneider et al., 1975). To reduce the dependencies between rotations in the independent planes, the lower target positions were used.

The activations of the neck muscles were also analyzed for motion along each plane. While the muscle activations were largely symmetrical, there were a few inconsistencies observed in the symmetry. During the 20^o flexion movement, the activations of the left sternohyoid were balanced by the right omohyoid, although muscles on either side were responsible for initially generating the flexion of the neck. The right sternohyoid and the left omohyoid forces damped after the head reached stability, whereas the left sternohyoid and right omohyoid remained activated for maintaining head flexion. In the motion along the transverse plane, the right post-scalene was activated during the right head rotation, whereas similar activation of the left post-scalene was not observed in the left axial rotation. In the lateral bending as well, the right mid-scalene was activated in the right bend simulations but the left mid-scalene remains inactivated during the left bending. The asymmetries in the activation patterns were observed even though the reward function penalized the RLMAC for asymmetry in the muscle activations around the mid-sagittal plane (Equation 6-1). Some of the asymmetries can be reduced by increasing the coefficients of the symmetry parameter in the reward function, however, that may cause an increase in the iterations

required by the RLMAC for learning muscle synergy. The rectus capitis minor had high activations while stabilizing and while performing sagittal and lateral kinematics. The multifidus also had high activation values during extension and lateral bending (Figure 6-9). The high forces in these muscles compensate for the passive stiffness provided by some soft tissue which is not included in the neck model. The high muscle activations can also be due to the absence of any term penalizing the high activations. Also, the multifidus extends from the sacrum (T1 in the present model) to the C1, however same activation level has been considered throughout the length of the multifidus, which can also result in high activations for some neck movements.

The final head positions were within 2.3^o of the target position in all three planes, the maximum angle error occurring for the left lateral bending. The higher error in the lateral bending scenarios can be due to the agent being trained less in the lateral plane compared to the other two planes (20% of the iterations in the lateral plane compared to 30% in sagittal and transverse planes). The errors in the simulation can be reduced by further training or by decreasing the time step during the training, however, the error of 2.3^o in one of the head movement directions was considered satisfactory for the present study as the error in all the directions were less than 10% of the head movement range (50^o in sagittal plane, 60^o in transverse, and 40^o in the lateral plane).

The kinematics output from the RLMAC in the sagittal and the transverse planes were compared with the human volunteer data gathered in chapter 5. The simulation outputs correlated well with the tests for the scenarios that were encountered by the RLMAC during the training (Figure 6-19) as well as for novel target signals (Figure 6-20). During the training, the head was assigned to move from the neutral position to flexion, extension, axial rotation, or lateral bending. However, the trained RLMAC could also move the head from the extremities back to the neutral position. The RLMAC could also produce head motions by reacting to the target signals prescribed along

the sagittal and transverse planes at once (Figure 6-21 and Figure 6-22), even though the head was trained to move along only one DoF. The RLMAC could move the head towards the target signal, however, in the two cases in which the RLMAC was tested, the target could not be reached when the target signals were applied in 2 directions simultaneously. Similar observations were made by Schneider et al. (1975) in human subjects. In the study, it was found that the ranges of motion in the sagittal plane decrease while the head is rotated. These simulations show the robustness of the trained RLMAC for synthesizing biofidelic head kinematics by being able to learn the complex relationships between the actions required and the state parameters of the control model.

6.5 Conclusions

Reinforcement learning was used to simulate omnidirectional fast movement in a multibody model of the head and neck. To the best of my knowledge, this is the first study in which goal-directed rotations of the head in all three anatomical planes have been performed. The response of the RLMAC was verified with volunteer test data and the RLMAC could also generate physiologically accurate movements under novel target signals. The results from this chapter show the robustness of the RLMAC, that can be trained to adapt to the complex relationships between the passive structures, the muscle synergy, and the target positions, which can be difficult by linear feedback gains as used in previous studies.

Chapter 7 – Evaluation of the trained RLMAC to changes in anthropometry

The neck model was developed with the anatomy of a 50th-percentile male (Fice et al., 2011; Panzer et al., 2011). As the mass, stiffness, and inertia properties were assigned for rigid vertebrae, skull, and joints it was possible to scale the values to represent a different anthropometry. In this chapter, the neck model was scaled to represent the anatomy of a small (5th percentile) female and 95th percentile male. The accuracy of the scaling parameters was verified by validating the passive response of the scaled 5th percentile female model with the cadaver data.

The RLMAC which was trained to perform fast goal-directed head movements with the 50th percentile neck model was integrated with the scaled models and the ability to perform similar kinematics was verified. The extension, flexion, and axial rotation responses were obtained in the scaled models and compared with the response of the neck model. The results of this chapter are important to determine the applicability of a trained RLMAC when used with different anthropometries, which will reduce the requirement of training and validation of the RLMAC individually for each case.

7.1 Introduction

Most of the previous control studies on the head and neck body region have been performed for 50th percentile male models (Brolin et al., 2008; Cappon et al., 2007; Chancey et al., 2003; Happee et al., 2017; Ólafsdóttir et al., 2019; Zheng et al., 2021). Very few past studies have considered different anthropometries for modeling the muscle control strategy. Dibb et al (2013) performed optimization studies to maintain the stability of the head under gravity in FE models of 6 year old and 10 year old pediatric subjects. Brolin et al. (2015) implemented PID control for the spine and hip joint in a 6-year-old child MB model. The active child model was used to study the effects of

173

postural control on the responses of a child during automotive braking and steering events. Recently, there has been a push towards the development of female-specific human body models (Davis et al., 2016; John et al., 2022; Östh et al., 2017b) as it was found that females are more susceptible to injuries during motor vehicle crashes (Bose et al., 2011; Forman et al., 2019).

Devane et al. (2022) implemented feedback controllers in a 5[th] percentile female FE model. The neck PID controllers included one controller for the neck sagittal angle and 210 controllers for neck muscles. The model was simplified by assuming identical controller gains for all the neck muscles. The active 5[th] percentile female FE model improved the response in pre-crash braking and low-speed frontal tests. Putra et al. (2019) compared feedback control strategies based on head and neck kinematics (vestibulocollic control) and muscle length (cervicocollic control) in a head-neck FE model of 50[th] percentile female developed previously by Östh et al. (2017). Putra and Thomson (2022) then expanded the study to incorporate the vestibulocollic and cervicocollic parameters into a single feedback control mechanism. In the study, the neck muscles in the FE model were divided into eight groups and each group was assigned similar vestibulocollic control gains based on spatial tuning studies. The muscle controllers were separate for each muscle and sought to reduce the change in length from the initial value.

In previous studies, the controller was developed with the objective of maintaining the head stability. The controllers had the same gains throughout the simulations assuming a linear relationship between the control errors and the external environment. However, a previous study by Zheng et al. (2021) demonstrated that gains used for stabilization under gravity cannot be used for higher magnitudes of external loads.

In the present chapter, the RLMAC trained for omnidirectional control was analyzed for control of head kinematics in different anthropometries. The neck model was scaled to represent the

anatomy of a small (5th percentile) female (henceforth known as the F05 model) and a large (95th percentile) male (henceforth known as the M95 model). The RLMAC trained in chapter 6 was integrated with the scaled models to generate the goal-directed head kinematics response. The response of the scaled model with trained RLMAC was compared with the responses of the neck model obtained in chapter 6. The evaluation of the ability of RLMAC to control the head kinematics for different anthropometry is important because a robust controller eliminates the computational cost associated with training for every new anthropometry.

7.2 Methodology

The neck model consists of rigid vertebrae, the skull (represented as Hybrid III headform), intervertebral joints, and 46 neck muscles. The neck model was scaled with appropriate scaling factors to represent the 5th percentile female and 95th percentile male anatomies. The scaled geometries were integrated with the trained RLMAC to obtain the kinematics responses while reacting to target signals.

7.2.1 Scaling to different anatomies

In the neck model, the mass and inertia were assigned to each vertebra and the head, and the values were obtained from the source FE model (Fice et al., 2011; Panzer et al., 2011). The stiffness of the intervertebral joints was adapted into bilinear curves from different sources of data (Chang et al., 1992; Dibb et al., 2013; Liu et al., 1982; Shea et al., 1991; Yoganandan et al., 2007). The muscles were added as line forces between origin and insertion, with multiple routing points along the length to account for the bending of the spine during movements (Panzer et al., 2011).

As the different parameters required were prescribed during the development of the neck model, all these parameters can be scaled to arrive at a different anthropometry. For the analysis in this

175

chapter, the neck model was scaled to 5[th] percentile female and 95[th] percentile male anatomies. The scaling factors were derived from the equal stress-equal velocity scaling derived by Eppinger et al. (1984). The anisometric scaling method was used to determine the length scaling factors (λ_x, λ_y, λ_z) in the different anatomical planes (Dibb et al., 2013; Forman et al., 2006). The characteristic length (λ_z) for scaling was measured from the center of mass (CoM) of the T1 vertebra to the CoM of the head (Figure 7-1). The density of the neck model and the scaled models were considered constant; therefore, the mass scale factor (λ_m) can be calculated by Equation 7-1.

$$\lambda_m = \lambda_x \cdot \lambda_y \cdot \lambda_z \qquad\qquad \text{Equation 7-1}$$

The mass scale factor can be determined by comparing the average masses between different anthropometries found in the literature. The length scaling factors in the transverse plane were considered equal and can be obtained using Equation 7-2.

$$\lambda_x = \lambda_y = (\lambda_m / \lambda_z)^{1/2} \qquad\qquad \text{Equation 7-2}$$

All the other scaling factors could be derived from the length and mass scale factors, considering velocity scaling (λ_v) to be 1 (Eppinger et al., 1984; Forman et al., 2006).
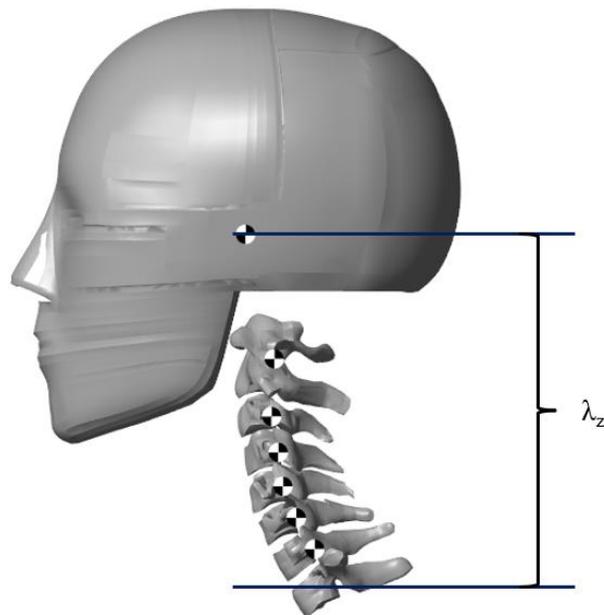
Figure 7-1: Measurement of the characteristic length scaling factor.

Table 7-1: Scaling factors for the model parameters

Scaling parameter		F05 model	M95 model
Mass	λ_m	0.667	1.31
Characteristic length	λ_z	0.787	1.132
Time	$\lambda_t = \lambda_m^{1/3}$	0.9233	1.094
Transverse length	$\lambda_x = \lambda_y = (\lambda_m/\lambda_z)^{1/2}$	0.921	1.076
Moment of Inertia	λ_I	[0.538, 0.533, 0.564]	[1.488, 1.501, 1.470]
Product moment of Inertia	λ_{PI}	[0, 0.483, 0]	[0, 1.595, 0]
Stiffness	λ_K	[0.787, 0.787, 1.077]	[1.132, 1.132, 1.022]
Damping	$\lambda_b = \lambda_K \cdot \lambda_t$	[0.726, 0.726, 0.994]	[1.045, 1.045, 0.943]
Rotational stiffness	$\lambda_{RK} = \lambda_m$	0.923	1.094
Rotational damping	$\lambda_D = \lambda_{RK} \cdot \lambda_t$	0.852	1.197
Muscle force (F_{max})	$\lambda_F = \lambda_x \cdot \lambda_y$	0.848	1.158

The mass and length factors for the F05 model were obtained by comparing the imaging data between the anatomies analyzed by Davis et al. (2016) and Gayzik et al. (2011, 2009) for the purpose of building a FE human body model. A similar mass scaling factor of 0.667 was used by Nie et al. (2019) to scale the response of the 5th percentile female foot and ankle model for validation purposes. The primary scaling factors for the M95 model were derived from the imaging study of 95th percentile males performed by Vavalle et al. (2014) and compared with the 50th percentile male data (Gayzik et al., 2009). The maximum muscle forces and the translation stiffnesses were scaled assuming constant stress (Eppinger et al., 1984). The rotational stiffness (and in turn moment) and time were assumed to be scaled equally for all directions with a scaling factor equal to the mass scaling (Forman et al., 2006; Nie et al., 2019). The inertia scaling was difficult to determine based just on the other scaling parameters for anisotropic scaling and hence

was considered to be the average value of the ratios of inertias of the vertebrae from the imaging studies (Davis et al., 2016; Gayzik et al., 2009; Vavalle et al., 2014). However, in the absence of imaging data, mass and time factors ($\lambda_m . \lambda_t^2$) can also be used to approximate the scaling value.

7.2.2 Validation of the scaled F05 model

To evaluate the accuracy of the scaling method, the F05 model was obtained by scaling with the factors mentioned in Table 7-1 and the passive response of the F05 model was evaluated in a rear impact scenario using the data gathered by Humm et al. (2021) for female head and neck specimens. The head-neck complex was isolated from the female specimens in the experimental studies. The T1 was potted and an acceleration pulse approximating a rear impact of 2.6 m/s was applied. The sagittal angles between adjacent vertebrae were measured and reported.

To replicate the testing conditions, the acceleration pulse as measured in the study was applied at the T1 of the F05 model in the horizontal direction and all the other DoF at the T1 were constrained. The angular motion between the vertebrae was outputted during the simulation and was compared with the test data to evaluate the biofidelity of the scaled F05 model and the overall scaling process.

7.2.3 Integration with the RLMAC to generate voluntary head kinematics

The F05 model and M95 model obtained by scaling were integrated with the RLMAC trained for the omnidirectional control of head kinematics (chapter 6). The RLMAC was then used to generate goal-directed kinematics in both the scaled models in the three anatomical planes. The models were stabilized under gravity for 400 ms, after which the target head position was prescribed in one of the anatomical planes. The kinematics of the head CoM was measured during the simulation, and the angle time history of the F05 model and the M95 model were analyzed and compared with the angle time data of the neck model.

7.3 Results

The mass and inertia of the rigid bodies, the stiffness and damping of the intervertebral joints, and the maximum muscle forces were scaled with scaling factors mentioned in Table 7-1 to obtain the head-neck model of two different anthropometries – a 5[th] percentile female (F05 model) and a 95[th] percentile male (M95 model).

7.3.1 Validation of the scaled F05 model

The accuracy of the scaling methodology was evaluated by comparing the passive responses of the F05 model with the test corridors in a rear impact scenario. The simulation results are compared with the test corridors in Figure 7-2.

(a)   *(b)*

(c)   *(d)*

(e)   *(f)*
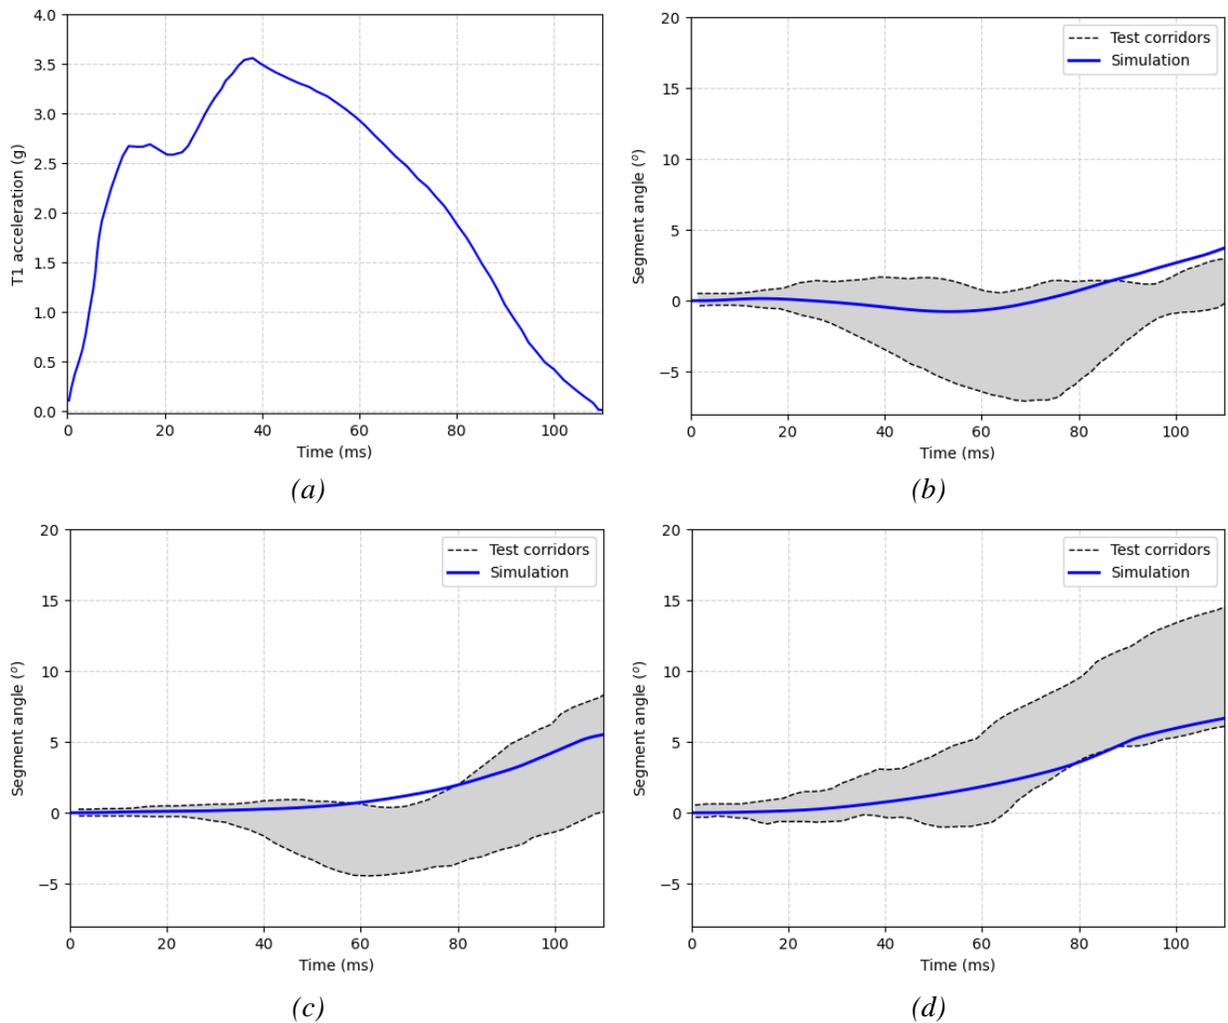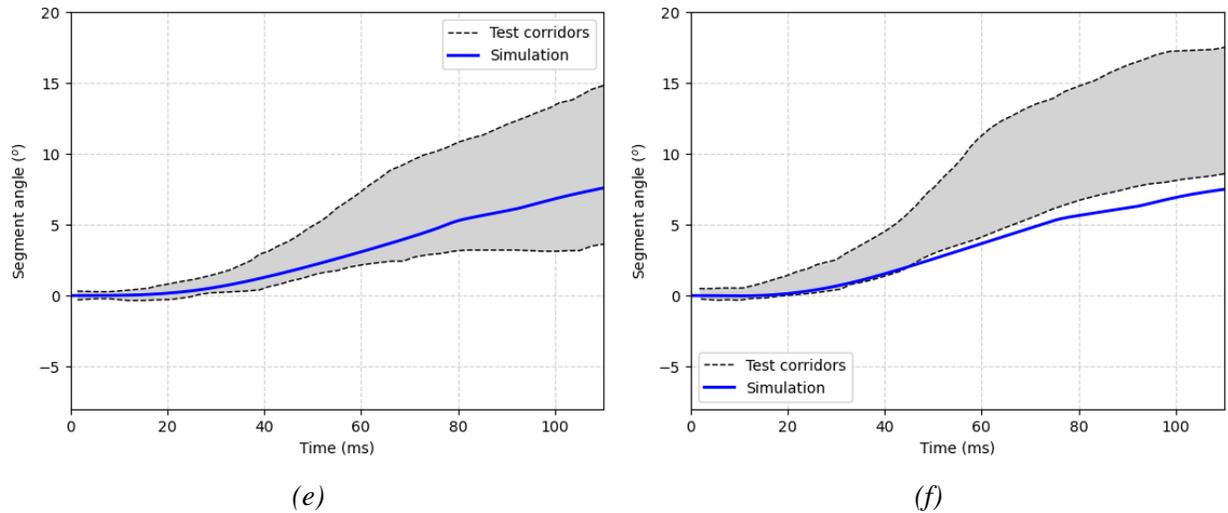
Figure 7-2: Validation of passive response of the F05 model in rear impact scenario (a) Acceleration pulse applied at T1 in the anterior direction (b) C2-C3 angle (c) C3-C4 angle (d) C4-C5 angle (e) C5-C6 angle (f) C6-C7 angle. The intervertebral angles are measured in the sagittal plane.

Figure 7-2a shows the acceleration profile imparted at T1 to replicate a 2.6 m/s rear impact scenario (Humm et al., 2021). The angles between the adjacent vertebrae in the sagittal plane were within the test corridors except for C6-C7. The C6-C7 joint was slightly stiffer compared to the test data, with the sagittal angle-time history near the lower bound of the test corridor.

7.3.2 Control of the scaled models

The scaled F05 and M95 models were integrated with the RLMAC to generate the desired head kinematics. The trained RLMAC was first used to stabilize the scaled models, followed by generating goal-directed motions in the sagittal, transverse, and lateral planes.

Figure 7-3 shows the head angle while stabilizing the F05 model and M95 model under gravity, i.e., when the target angle was set to 0 for all three planes. The RLMAC could stabilize the head in both scaled models. In the F05 model, the final angles in the three planes were less than 1^{o} (Figure 7-3a). In the M95 model, the maximum error was 2.2^{o} in the axial plane, whereas the errors in the other two planes were below 1^{o} (Figure 7-3b). In the F05 model, the head took 0.5s to completely stabilize compared to 0.4s in the M95 model.

180
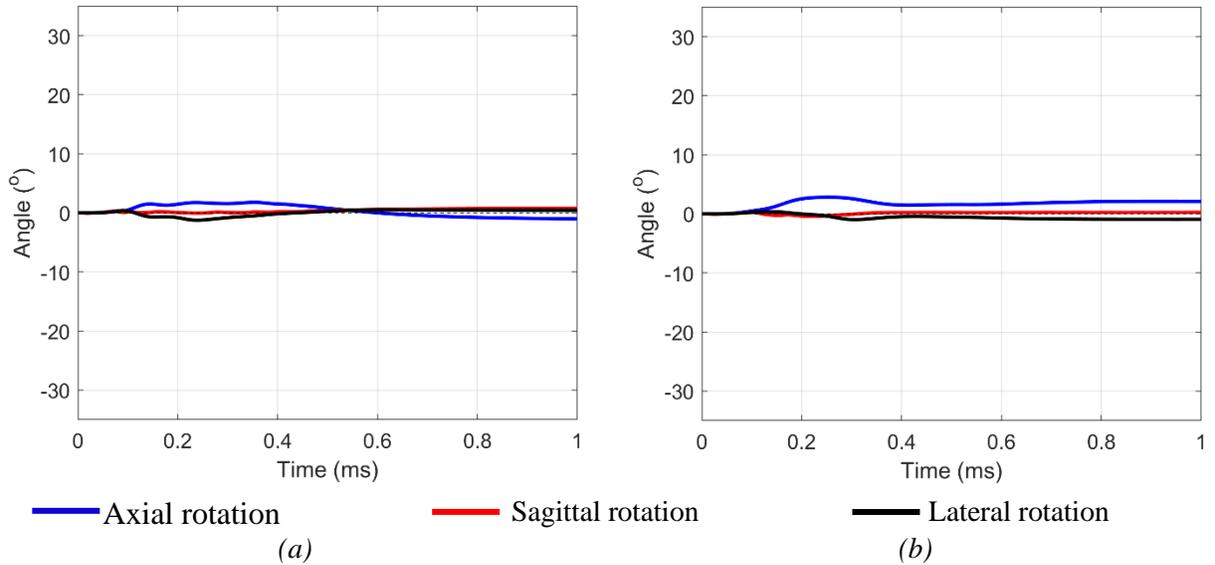
Figure 7-3: The head angle-time history during the stabilization run in the three planes (a) F05 model (b) M95 model.

Following the stabilization run, the scaled models were simulated for goal-directed head kinematics. The target angle was prescribed in the preferred plane, while targets in the other two directions were maintained at 0^o.

Figure 7-4 shows the angle time history of the F05 model in each of the anatomical planes. The F05 model could reach the target orientations in the sagittal and lateral planes with a maximum error of 2.2^o in the lateral left rotation. In the transverse planes, the head CoM reached 26^o when a target step of 30^o was applied. The F05 model undershot the target in both left and right rotations. Figure 7-5 shows the angle-time plots for M95 model under goal-directed head motions. In the M95 model, the head could follow the target signals in the three anatomical planes with maximum error less than 1.5^o. There was very little out of plane motion in the other directions as can be seen from the graphs (Figure 7-4 and Figure 7-5).
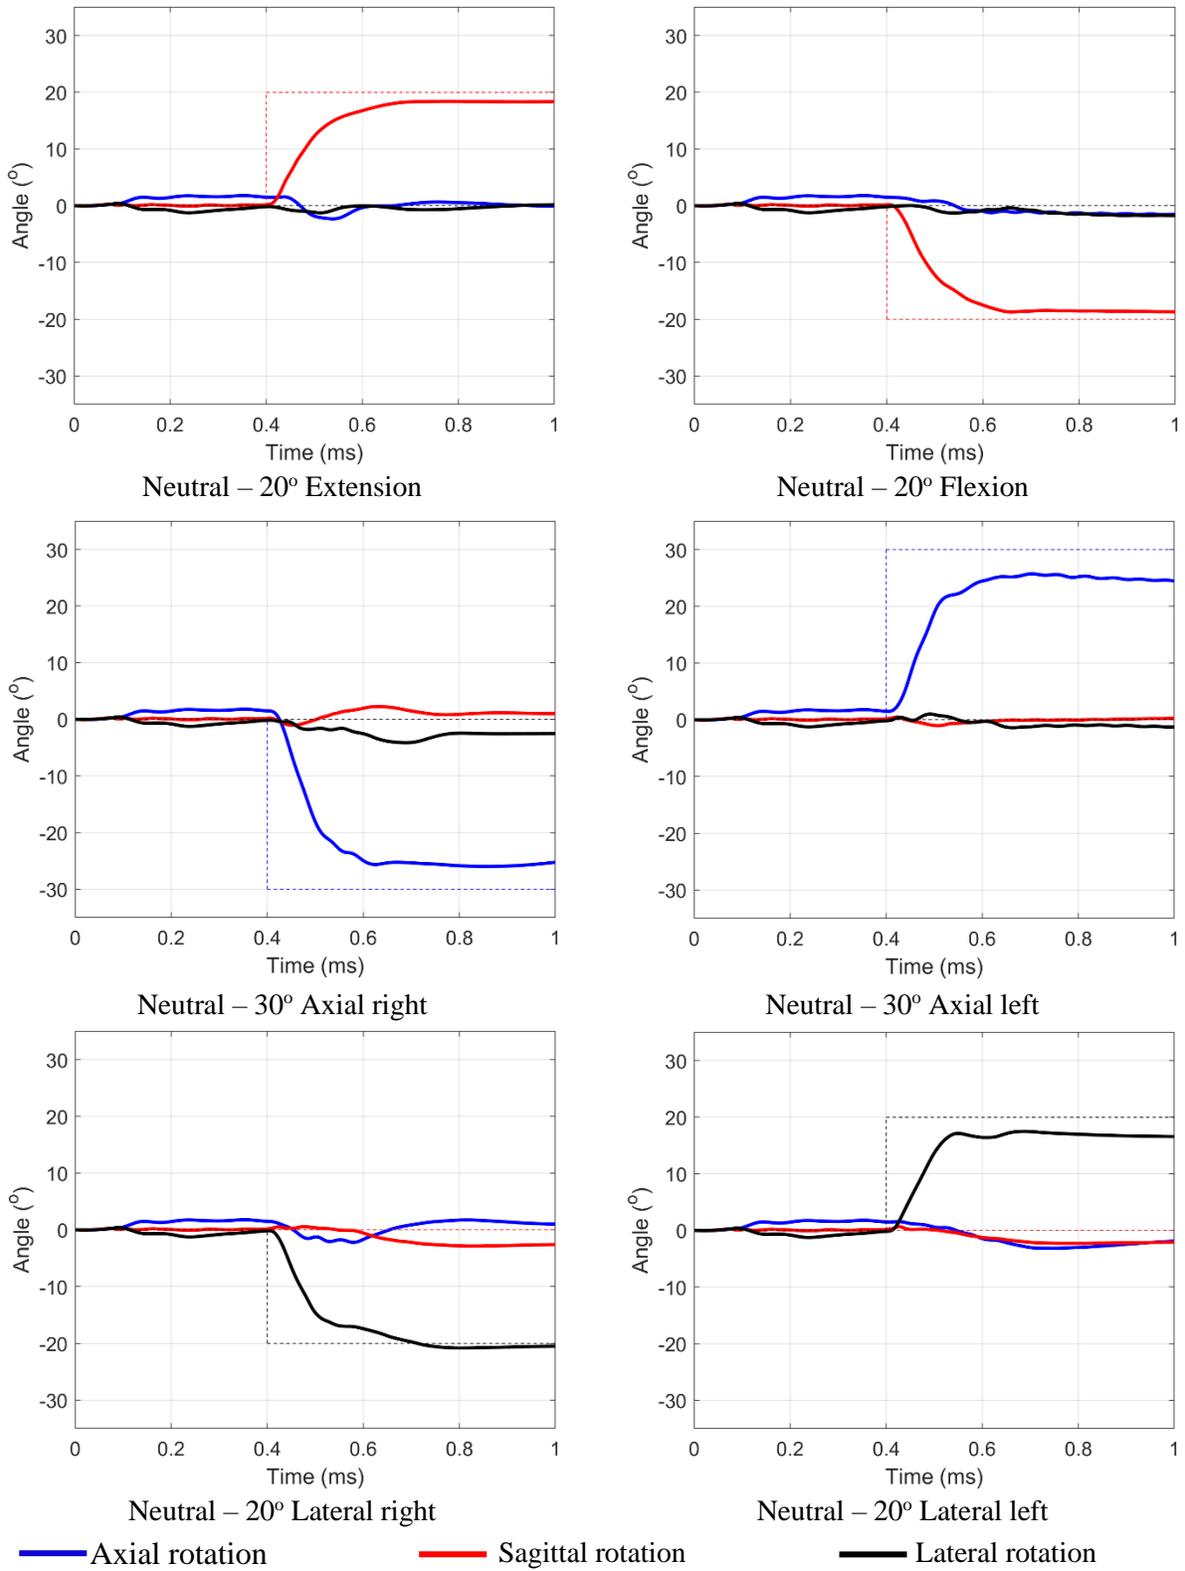
Figure 7-4: Response of the head of F05 model under goal directed head motion (bold–head angle, dotted–target angle).
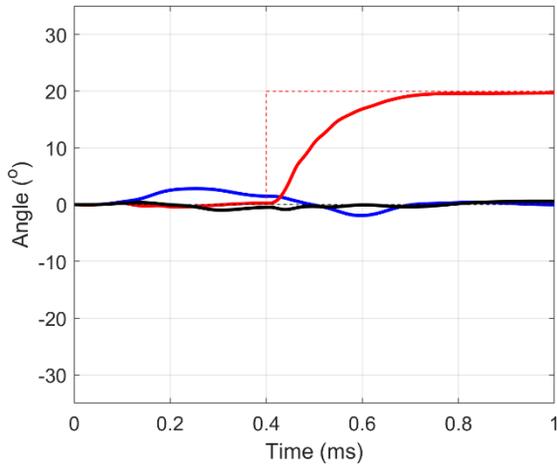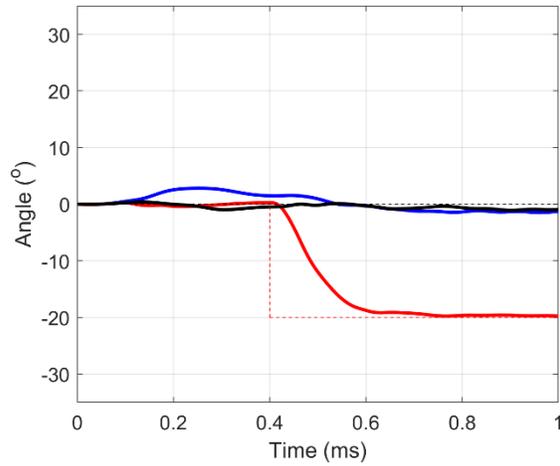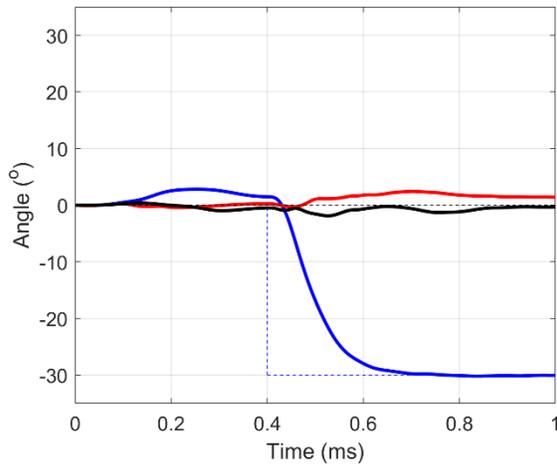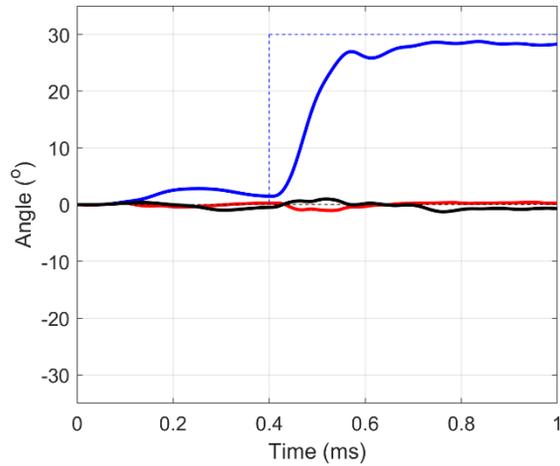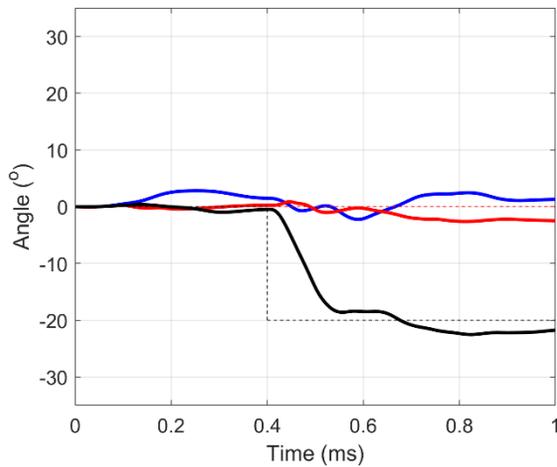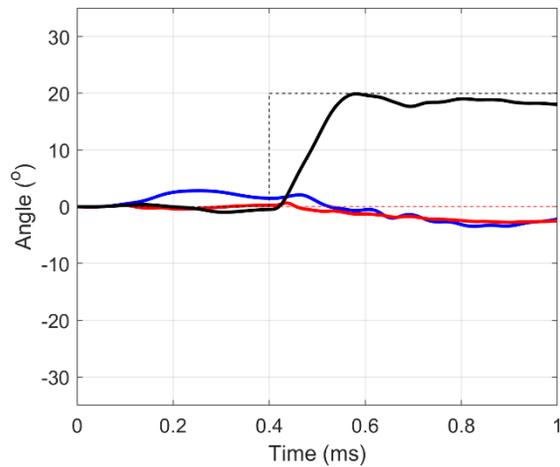
Figure 7-5: Response of the head of M95 model under goal directed head motion (bold–head angle, dotted–target angle).
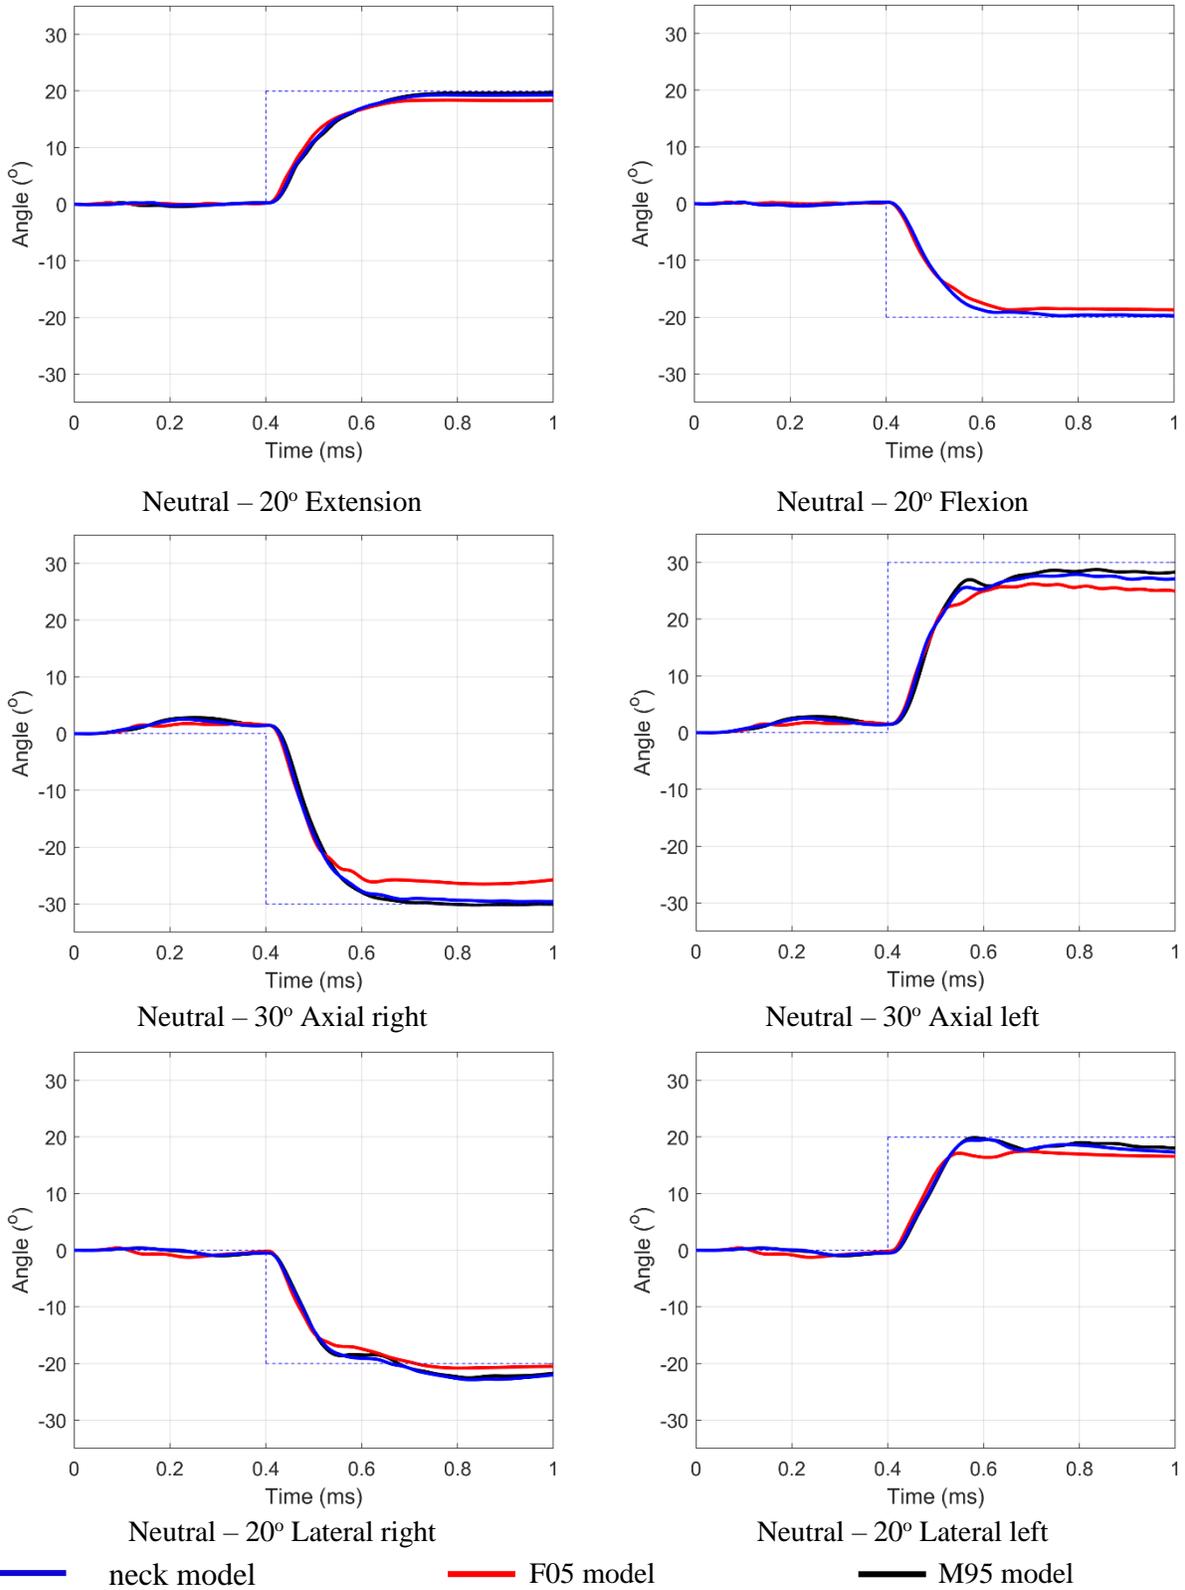
Figure 7-6: Comparison of response of the neck model and the scaled models while performing goal directed kinematics in the three planes. The blue dotted lines represent the target angles.

Figure 7-6 compares the response of the F05 and the M95 models with the neck model. The angle-time profile of the M95 model was similar to that of the neck model with very little to distinguish between the responses. The F05 model's head angle was also within 1° of the neck model in the sagittal and lateral planes. However, the angle at stabilization in the axial plane was lower in the F05 model with the difference being 2° in left turn and around 3° in right turn.

To understand the cause of the differences in the response of the F05 model from the neck model, a sensitivity study was performed evaluating the effect of each parameter on the overall response. Initially, simulation was performed by scaling the mass and the inertia of the head and the vertebra. Next, the distance between the vertebrae and the muscle lengths were also scaled to the anatomy of the 5[th] percentile female. Finally, the muscle forces and joint stiffness were scaled to arrive at the response of the F05 model.
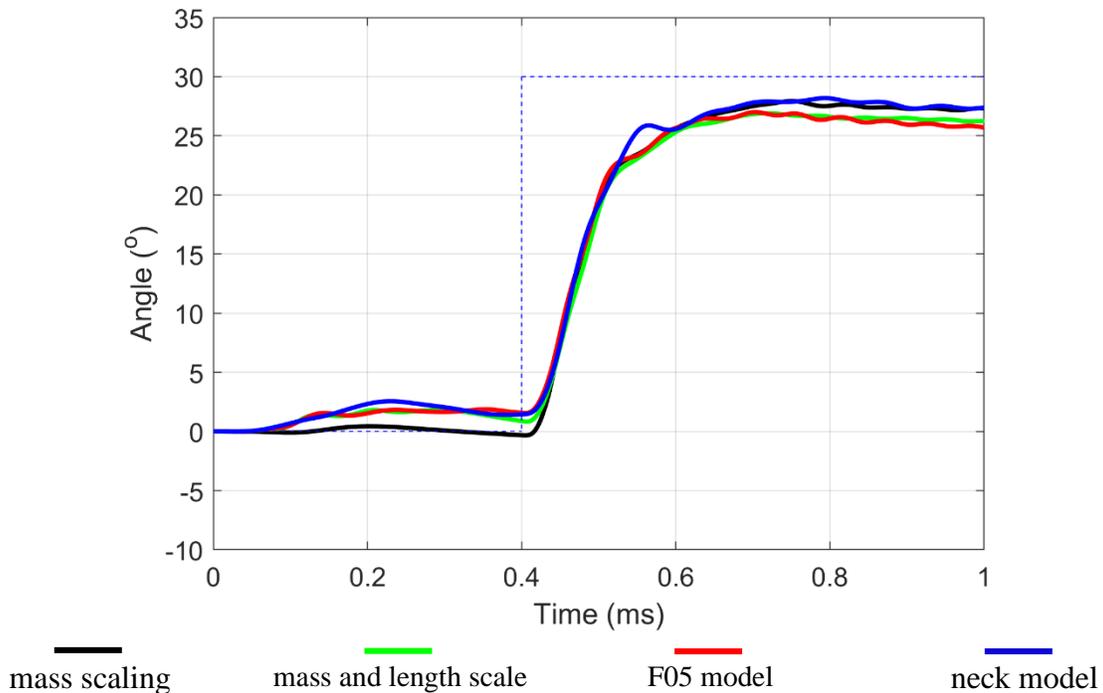
Figure 7-7: Sensitivity study of scaling parameters on the overall response of the F05 model in 30° left axial rotation. The blue dotted lines represent the target angle.

185

Figure 7-7 show the effects of various scaling parameters on the response of the F05 model. Scaling only the mass has little effect on the head rotation. In fact, scaling only the mass improves the stabilization of the neck model in the transverse plane as can be seen in the first 400 ms in Figure 7-7. Scaling the coordinates of the vertebra CoMs to the respective position representing fifth percentile female has the major effect on the model response. After the structural scaling, scaling the muscle forces and the joint stiffness again have little effect on the response of the F05 model.

7.4 Discussions

The ability of the RLMAC to adapt to changes in anthropometries have been evaluated in this chapter. The neck model was scaled to anatomies representative of a 5[th] percentile female (F05 model) and a 95[th] percentile male (M95 model). Anisotropic scaling methodology was used with the distance between the T1 CoM and head CoM as the characteristic length. The stress, density, and velocity were considered constant across the anthropometries, and the scaling factors of the other parameters were derived from the three primary parameters (Eppinger et al., 1984; Forman et al., 2006). The values of scaling factors and their calculations are provided in Table 7-1.

The accuracy of the scaling methodology was verified by evaluating the passive response of the F05 model under rear impact scenario. An acceleration pulse representative of 2.6 m/s rear impact was applied at the T1 and the pitch angle between the adjacent vertebrae were instrumented during the simulation. The simulation results show that the vertebral angles were within the test corridors. Only the C6-C7 angle was lower than the corridor, however it was within 1° of the lower bound at the conclusion of the simulation. The results show that the scaling methodologies is biofidelic for the purpose of using the scaled models for postural control.

The RLMAC trained in chapter 6 was integrated with the F05 model and the M95 models for the purpose of generating goal-directed head kinematics. The RLMAC could stabilize both the F05 model and the M95 model under gravity (Figure 7-3). Following the stability run, the targeted motions in the three planes were simulated. The scaled models were stabilized for 400 ms following which the target signal was prescribed at the preferred direction. The M95 model could synthesize the desired kinematics and the angle time response of the M95 model was comparable to the neck model in the head rotations performed (Figure 7-6). The F05 model could produce the head rotations within 1º of the neck model in the sagittal and lateral plane. However, the head angle in the F05 model was lower than the neck model in axial rotations. In the left axial rotation, the difference in head angle between the models was 1.5 whereas the error was 3º in the right axial rotation.

Sensitivity study was carried out to evaluate the effects of different scaling parameters on the response of the neck model. Scaling just the mass and inertia to that of the F05 model was found to have little effect on the head kinematics response. However, scaling the structural dimensions, i.e., the distance between the vertebra and head CoM as well as the muscle length was the major contributor to the difference in the head response observed (Figure 7-7). This can be explained as there was no reward component controlling the translation of the head, and as a result, the moment arm for the muscle forces, which may result in the minor differences observed in the F05 model. Further analysis may be required on eliminating the error, however with the current training framework the RLMAC was robust enough to keep the max error below 10% over the range of motion in the F05 model. In the M95 model, the RLMAC could reproduce the desired kinematics with angle time profile very similar to that of the neck model.

7.5 Conclusions

In this chapter, the trained RLMAC has been tested for a 5th percentile female and a 95th percentile male anatomy. The RLMAC could synthesize the desired kinematics for both the scaled models, however the F05 model fell short of the target angle in the axial rotation case. The ability of the RLMAC to reproduce the desired kinematics for change in anthropometries will drastically reduce the requirement of training. The trained RLMAC agent can be further trained for the modified anatomy without starting from scratch which reduces the computational cost associated with training a RL agent.

Chapter 8 – Exercising of the RLMAC to inertial and impact loads

In chapter 6, the RLMAC was trained to maintain the stability of the neck model under gravity and synthesize goal-directed kinematics of the head. In this chapter, the ability of the trained RLMAC to react to added inertia and impact loads will be verified. This step is important to determine the overall applicability of the trained RLMAC in different scenarios.

Inertial loads were introduced in the neck model by adding the mass and inertia of an American football helmet to the head of the model. The trained RLMAC was then coupled with the updated model to reproduce the targeted motions that were encountered during the training. The RLMAC was also incorporated to predict the motions of the head under low-speed frontal and rear impact scenarios, and lateral padded impacts to the head.

The trained RLMAC was found to be robust and was able to control the head kinematics under the mentioned load cases. Results from this chapter demonstrate the general applicability of reinforcement learning-based muscle control for static as well as dynamic loads.

8.1 Introduction

In many situations including sports or while operating a motorcycle, individuals are required to wear helmets for added safety. Helmets add inertia to the head, causing an increase in forces and moments on the neck (Harrison et al., 2015; Thuresson et al., 2005). Few previous computational studies have demonstrated the effect of muscle activation on the response of a helmeted human body model. Jin et al. (2017) performed a finite element (FE) study to evaluate the effects of muscle activations on head kinematics while wearing American football helmets. The study simulated a head-to-head lateral impact situation in passive condition in addition to two different activation scenarios: In the first scenario, all the muscles were activated to 100% around 55 ms

after the impact (late impact scenario), and in the second activation scenario, the muscles were activated at 73% at the time of impact and the activation was increased to 100% 15 ms post-impact (early impact scenario). The results showed that in the early activation simulation, the peak rotational velocity reduced by around 16% compared to the passive case. However in the late activation case, the peak rotational velocity decreased only by 8% from the passive condition. Bruneau and Cronin (2020) performed impact simulations using a helmeted 50th-percentile male FE model and a deformable impactor. Muscles were grouped as extensors and flexors and the FE model was impacted in three separate directions. A 'balanced' activation state was used for the impact simulations, in which the flexors were assigned 87% activation and a ratio of 0.15:1 was maintained between the extensor and flexor activations. The muscle activations were found to have small effect of the overall kinematics of the head with only 2% decrease of the peak rotational velocity in the 9.3 m/s lateral impact case. These studies show the importance of activation onset time as well as the muscle activation synergy on the kinematics response of a head FE model with helmet. However, the effect of adding the helmet on the voluntary motion of the head, neck, and spine has not been assessed. The relevance of such a study is related to athletes' and soldiers' inclination to wear lighter helmets that typically offer less restriction to head movement and less neck fatigue than a heavier helmet, but usually at the cost of a decrease in safety.

In dynamic events such as vehicle impact scenarios, volunteers respond differently to cadavers due to the presence of active muscle forces. Iwamoto et al. (2012) compared the results of the 8 km/hr rear impact sled test using cadavers (White et al., 2009) to a volunteer study performed under a similar condition (Ono et al., 1997) and found that the peak sagittal head angle was 14° higher in cadavers compared to the human volunteers. A robust control model should be able to effectively control the kinematics of the head while stabilizing under gravity as well as under such

dynamic scenarios where muscle activations alter the body kinematics, without the need to update the control parameters. Previous control studies using feedback controllers have used different gains for different loading cases. Happee et al. (2017) optimized the values of feedback gains based on the perturbation frequency at T1. Zheng et al. (2021) implemented feedback control based on CCR and VCR to stabilize the head under gravity and used the same architecture to predict head kinematics under impact scenarios. However, for high amplitude impact cases, the gains were modified from the stabilization condition to obtain the desired kinematics response. The methodology of varying the gains corresponding to impact loads can be useful in some cases, but not robust enough for complex impact scenarios that may involve loads of varying magnitudes that need to be accounted in the control model.

In this chapter, the ability of the trained RLMAC to adapt to changes in inertia and external impact loads has been evaluated. RLMAC trained in chapter 6 to perform omnidirectional head kinematics is subjected to added head inertia representing an American football helmet and the goal-directed head motions performed in chapter 6 are repeated. The response of RLMAC is also evaluated for dynamic environments representing automotive and lateral impacts to the head, without any modification to the controller. The RLMAC was simulated under low speed sled conditions (Fice et al., 2021b) representative of frontal and rear impact scenarios. For evaluating the control in lateral direction, the neck model was simulated under lateral padded head impacts (Reynier et al., 2020). The external boundary conditions used in this chapter will help determine the general robustness of the RLMAC in scenarios that are generally simulated with computational human body models.

8.2 Methodology

The neck model with trained RLMAC was used to evaluate the response of RLMAC to changes in inertial and dynamic loads. Mass and inertia representing an American football helmet were added to the head model to observe the effects of change in inertia on the overall response of the RLMAC. For dynamic loads, kinematics and forces obtained from various sources of volunteer data were applied to the head model to obtain the response. The responses of the RLMAC were compared with the test data to determine the applicability of the RLMAC in such scenarios.

8.2.1 Response to added inertia

Funk et al. (2018) measured the mass and inertia of 37 American football helmets. Out of the helmets measured, the Vicis zero 1 2017 model had the highest mass of 2.193 kg. For the analysis in this chapter, the mass and inertia of Vicis were added to the head of the neck model. The center of mass (CoM) of the helmet was measured with respective to the head CoM of a hybrid III dummy (Funk et al., 2018). The properties of the Vicis zero 1 helmet are provided in Table 8-1.

Table 8-1: Mass, inertia, and CoM location of the Vicis zero 1

Mass (kg)	2.193 (1.796 kg helmet shell and 0.397 kg facemask)
Inertia (kg – m^2)	$I_{xx} = 0.0245$, $I_{yy} = 0.0271$, $I_{zz} = 0.0277$ $I_{xz} = 0.0278$
CoM location (mm)	X = 13.7 anterior, Z = 20 superior

Note: The helmet CoM coordinates have been measured from the head CoM.

The helmet was assumed to be symmetric about the sagittal plane, thus the CoM of the helmet was considered to lie in the sagittal plane (i.e., Y = 0). For this reason, the product moment of inertias I_{xy} and I_{yz} were also assumed to be 0. The head model with the representation of helmet CoM location is shown in Figure 8-1.
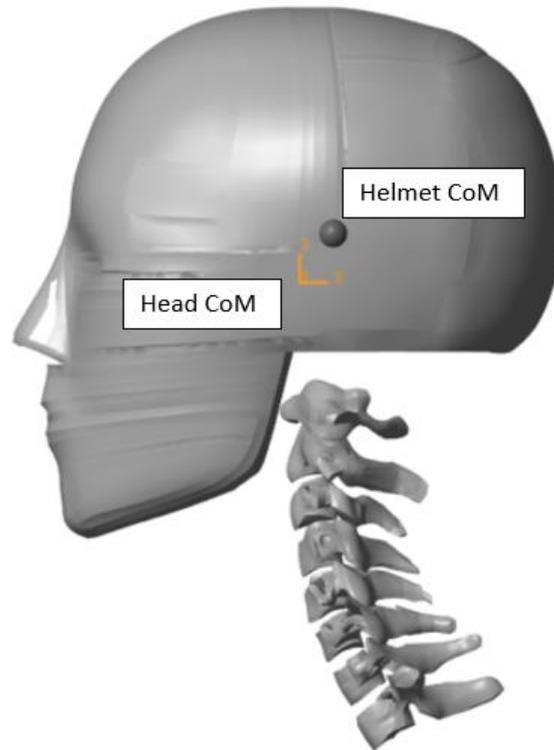
Figure 8-1: The relative position of Vicis Zero 1 center of mass with respect to head CoM.

The helmeted neck model was stabilized under gravity to evaluate the ability of the RLMAC to react to the added mass. The energy cost of stabilization (joint energy), along with the muscle activations and the compression (along Z axes or in the superior-inferior direction) of the vertebral joints were compared with the neck model in the same stabilized scenario without a helmet. The joint energy cost is calculated for all 6 DoFs (Equation 8-1).

$$Joint\ energy\ (E_j) = \sum F_j\ x\ s_j$$

<div align="right">Equation 8-1</div>

Following the head stabilization, head extension and right axial rotations were also performed as a reaction to step target signals. The head was stabilized for initial 0.4 s before applying the step target. The angle time history of the head CoM was measured and compared with that of the bareheaded neck model obtained in chapter 6. The energy cost of the head motion (Equation 8-1) was also calculated and compared for both goal-directed motions.

8.2.2 Response to automotive impact scenario

The neck model was simulated in conditions resembling low speed frontal impacts and rear impacts. Fice et al. (2021) performed volunteer sled tests to measure the head and neck kinematics in frontal and rear automotive impacts. The volunteers were seated in an automotive seat mounted over a sled and the responses were measured with their arms placed on a steering wheel under two conditions – relaxed and braced. The sled was operated with a peak acceleration of 19.9 m/s^2 with the sled pulse representative of 8 km/h frontal and rear impacts. In the experiments, the kinematics of the head CoM were measured using accelerometers. The torso kinematics were measured at the C7-T1 vertebral joints.

(a)   *(b)*

(c)

—— Test data   —— Average curve prescribed at T1
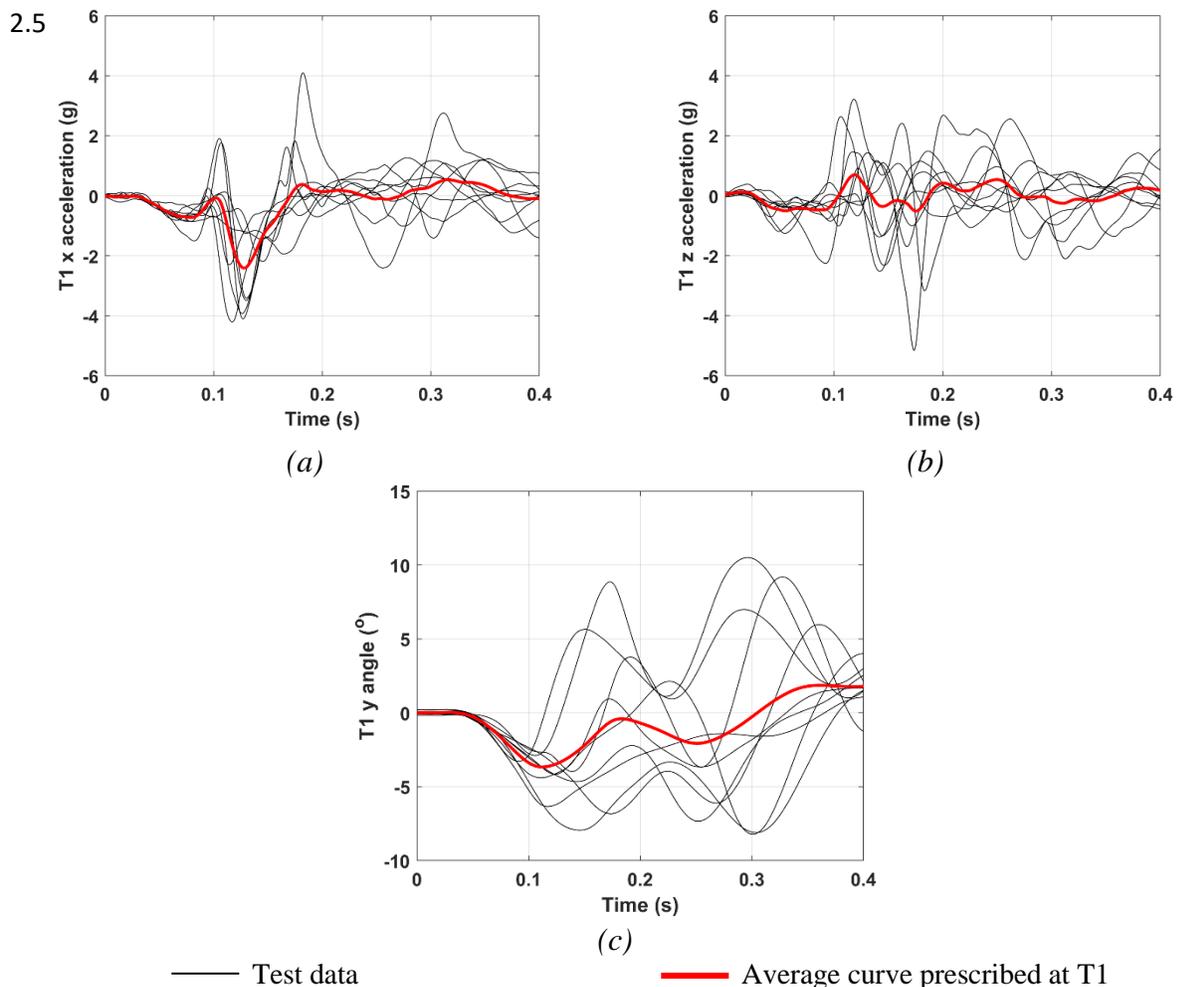
Figure 8-2: The boundary condition prescribed at T1 for relaxed frontal impact (a) T1 horizontal acceleration (b) T1 vertical acceleration (c) T1 rotation.

In the simulations, the torso kinematics resolved at the T1-C7 joint were used as the boundary conditions for the four test cases. The translation and rotations measured at the sagittal plane (X and Z translation, Y rotation) were averaged for the volunteers and prescribed at T1 of the neck model (Figure 8-2). The head CoM rotation and the head retraction, defined as the horizontal displacement of the head as measured from T1, was measured in the model and compared with the test data. Simulations were carried out for the active neck model using the RLMAC and also without the RLMAC (i.e., the passive model). The active model was stabilized for the first 400 ms before imparting the boundary conditions at T1. In the impact simulations, the target angular error was set to 0^o in all planes. A sensitivity study was also carried out to observe the effect of T1 rotation on the overall response of the RLMAC when simulated in the anterior-posterior direction.

8.2.3 Response to padded lateral impacts

To evaluate the response in the lateral direction, the neck model was simulated under padded impacts with boundary conditions adapted from human subject tests conducted by Reynier et al. (2020). In the experiments, 20 male volunteers were subjected to lateral head impacts with a steel ball weighing 3.7 kg and covered with foam at a horizontal velocity of 2m/s. The impact area on the volunteers was also padded with foam and the impacts were carried out under three muscle conditions – passive, unilateral, and co-contracted (Reynier et al., 2020). The ball was instrumented to measure the acceleration at the time of the impact. The head kinematics of the volunteers were measured using an instrumented mouthguard. To simulate the impact conditions, a horizontal force was applied to the head, 15 mm above the head CoM. The force magnitude was derived from the average impact acceleration of the steel sphere (Figure 8-3). In the volunteer studies the torso was not restrained (although a shoulder support was present to reduce the gross motion of the thorax). Thus, the T1 of the neck model was allowed to translate and rotate in the

lateral plane and stiffness of translation and rotation between the global reference and the T1 were determined from a sensitivity analysis performed on the neck model without muscle activation (passive) and compared to the volunteer response under passive muscle conditions.
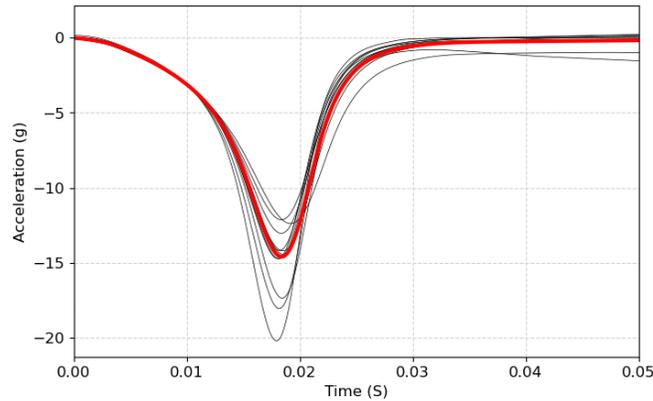
Experiment accelerations   *Mean curve used as the boundary condition*

Figure 8-3: Mean acceleration of the impactor used as the boundary condition for lateral impacts.

8.3 Results

8.3.1 Response to added inertia

RLMAC trained in chapter 6 was able to stabilize the neck model with the added helmet inertia. Figure 8-4 shows the rotation of the helmeted head under gravity. The head was stabilized around 400 ms and the maximum angular error was 2.8° in the transverse plane. In the other two planes, the angular errors were within 1°.
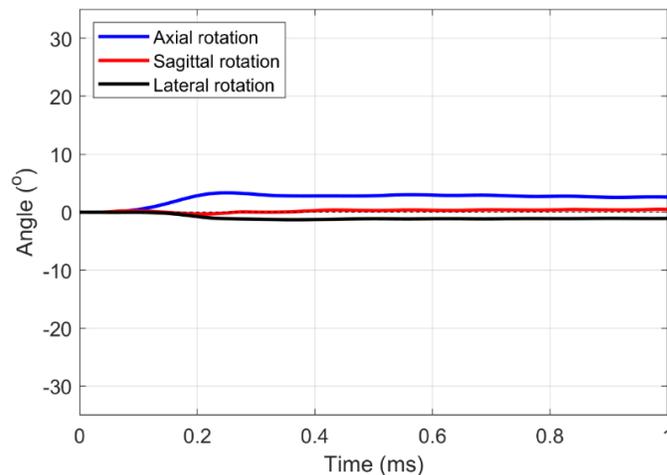
Figure 8-4: The head CoM angular rotations during the stabilization simulation.

196

The joint energy cost of the stabilization simulation was calculated during the simulation and compared with the response of the neck model without the helmet. During the head stabilization, the helmet increased the joint energy by 25% (Figure 8-5).
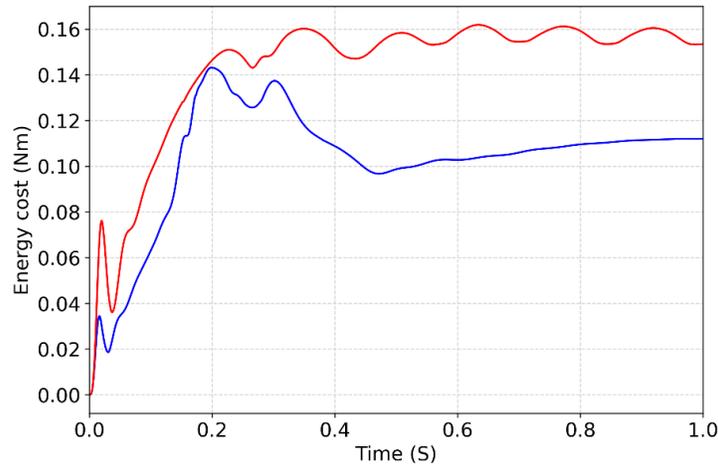
Figure 8-5: Joint energy cost during the head stabilization simulation.

The vertical compression of the spinal column at the final stabilized position was also measured and compared to the case without a helmet (Figure 8-6). The application of the helmet was found to increase the compression of the joints by around 15-20 %. While the maximum compression was seen in the C2-skull occipital joint, the joint compressions with added inertia of the helmet were still lower than 1 mm.
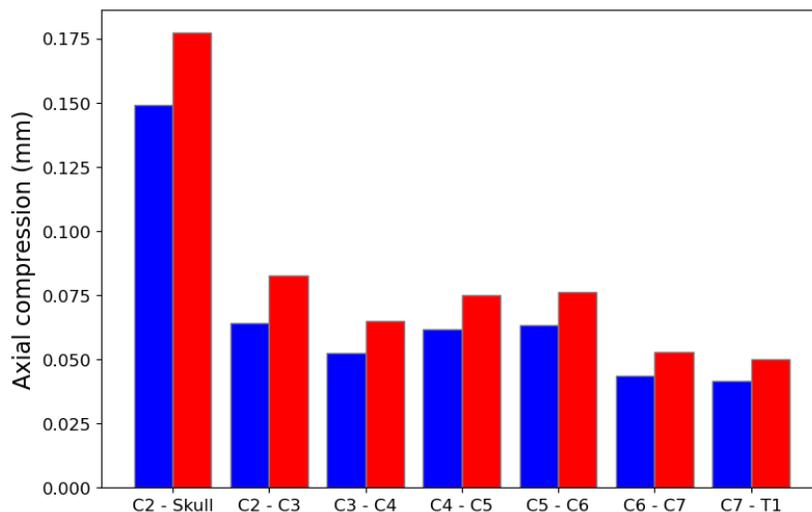
Figure 8-6: Vertical compression of the intervertebral joints during the head stabilization simulation.

The response of the RLMAC under helmeted conditions has been compared with that of the neck model without the helmet in Figure 8-7. The angle time history of the head CoM rotation has been compared for 20° extension and 30° right axial motion. The kinematics of the helmeted model was very similar to that of the bare-head model in the 20° extension motion, with the helmeted model stabilized within 1° of the target angle. In the axial rotation case, there was a small latency in the helmeted model's initial response concerning the model without a helmet. The helmeted model also overshot the response of the neck model before stabilizing at the target position around the same time.
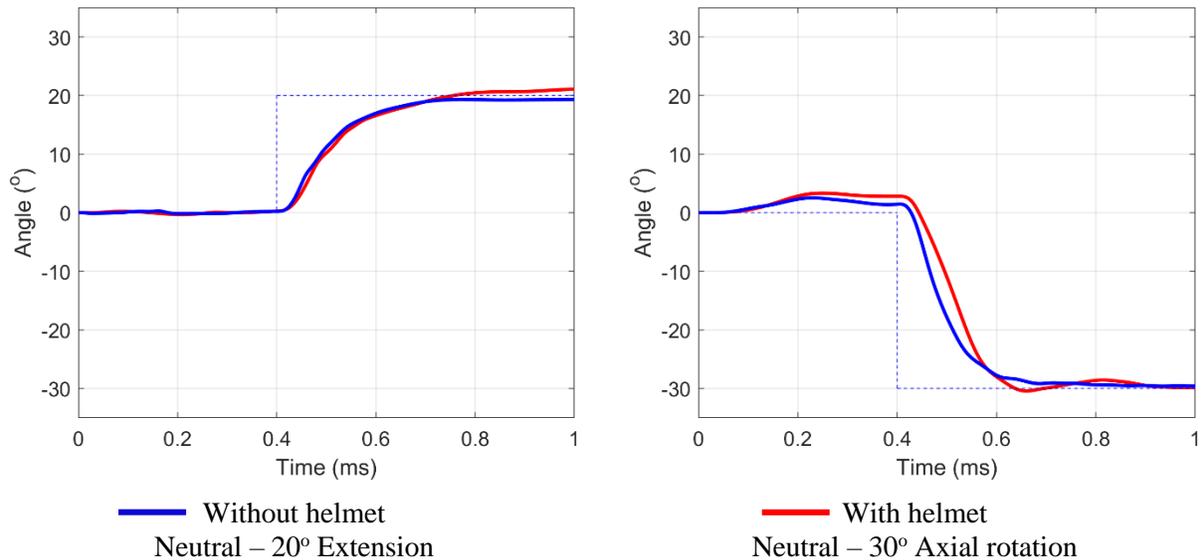
Figure 8-7: Comparison of the response of the UVa neck model while performing goal-directed head motions with and without a helmet. The dotted line shows the target angle prescribed to the RLMAC.

Figure 8-8 compares the joint energy while carrying out goal-directed kinematics. In the extension motion, the joint energy for the helmeted model at the stabilized target position was around 30% higher than the model without a helmet. The joint energy in the axial motion case was higher during the head movement in the helmeted model. Once the head stabilized, the joint energy was almost the same with less than a 10% difference at the target position.
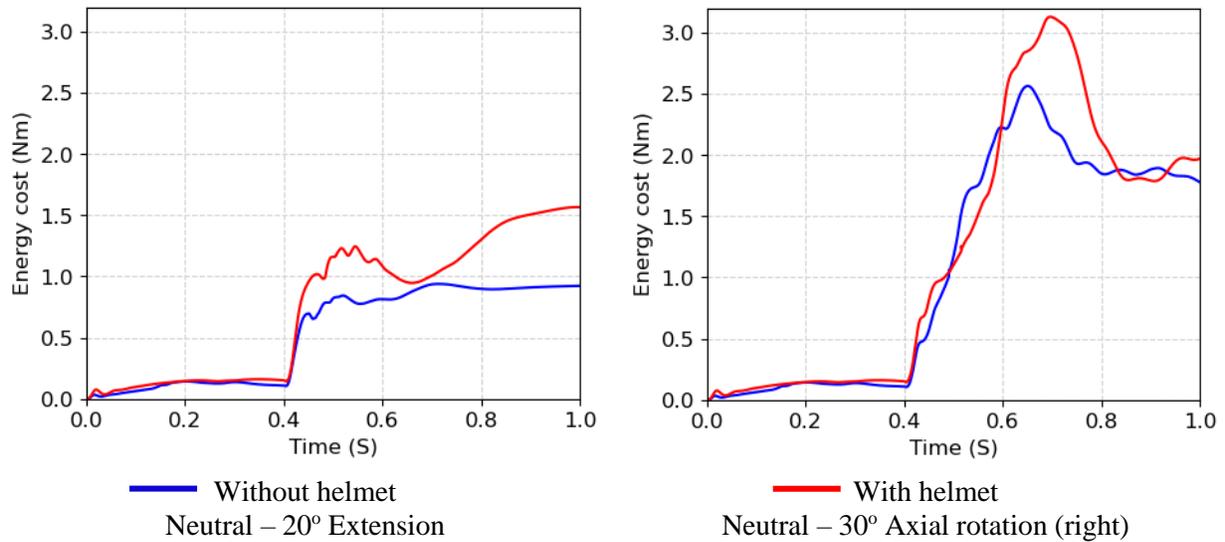
Figure 8-8: Comparison of the joint energy expenditure of the UVa neck model while performing goal-directed head motions with and without a helmet.

8.3.2 Response to automotive impact scenario

The response of the head model with RLMAC was evaluated under dynamic conditions resembling low-speed frontal and rear impacts. The response was simulated under two conditions – arm relaxed and braced against the steering wheel. The response of the RLMAC was compared with the test data as well as with the response of the passive model. As the control objective was to minimize the error between the rotation of T1 and head CoM (i.e., target head angle of 0^o when measured relative to the T1), the rotation prescribed at the T1 as boundary condition was also plotted along with the head rotations for reference. The Head CoM and T1 rotation angles have been plotted in the global reference frame.

Figure 8-9 shows the response of RLMAC and the passive neck model in frontal impact conditions. The RLMAC maintained the head within the test ranges. The passive model on the other hand was unstable in both conditions with the peak angle much higher compared to the experiments. The head CoM was also able to follow the T1 angle closely in the active case as the control target was set to minimize the angular error between the head CoM and the T1.

199
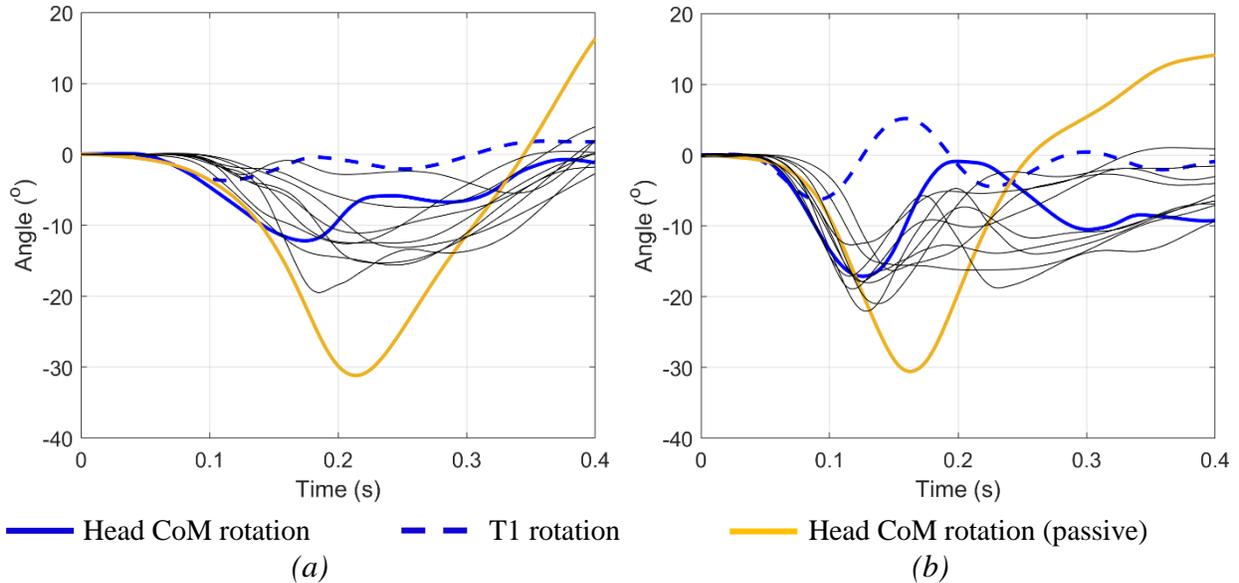
Figure 8-9: Response of the active and passive model under low-speed frontal impact condition (a) arm relaxed (b) arm braced.

Figure 8-10 displays the response of RLMAC and the passive model under the rear impact conditions. In the relaxed state, the response of the RLMAC is similar to the tests whereas the response of the passive model is outside the test bounds (Figure 8-10a). The RLMAC was initially stiff in the braced condition compared to the test data (Figure 8-10b).
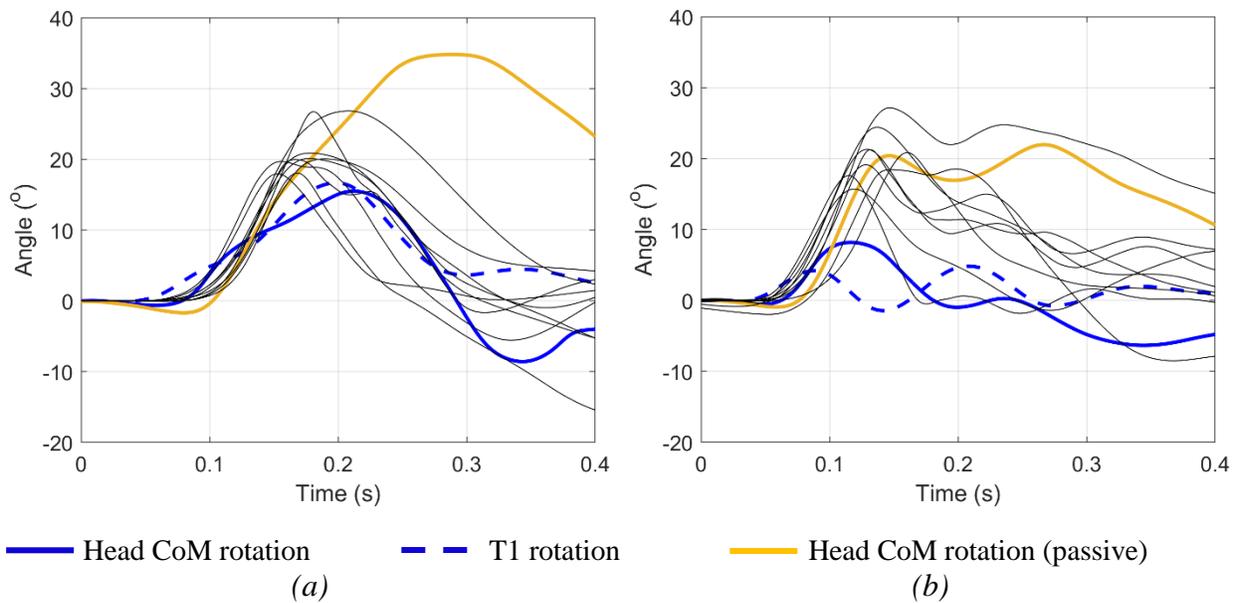
Figure 8-10: Response of the active and passive model under low-speed rear impact condition (a) arm relaxed (b) arm braced.

In the impact simulations, the head retraction was also measured, defined as the horizontal displacement between the atlantooccipital joint (AOJ) and the T1. Similar to the measurements in the tests, the AOJ was fixed at 24 mm posterior and 37 mm inferior from the head CoM. Figure 8-11 shows the comparison of the active and passive neck model with the test data. The simulation curves both active and passive condition have been resolved at the T1 reference.

(a) *(b)*

(c) *(d)*

— Active response with RLMAC   — Passive response
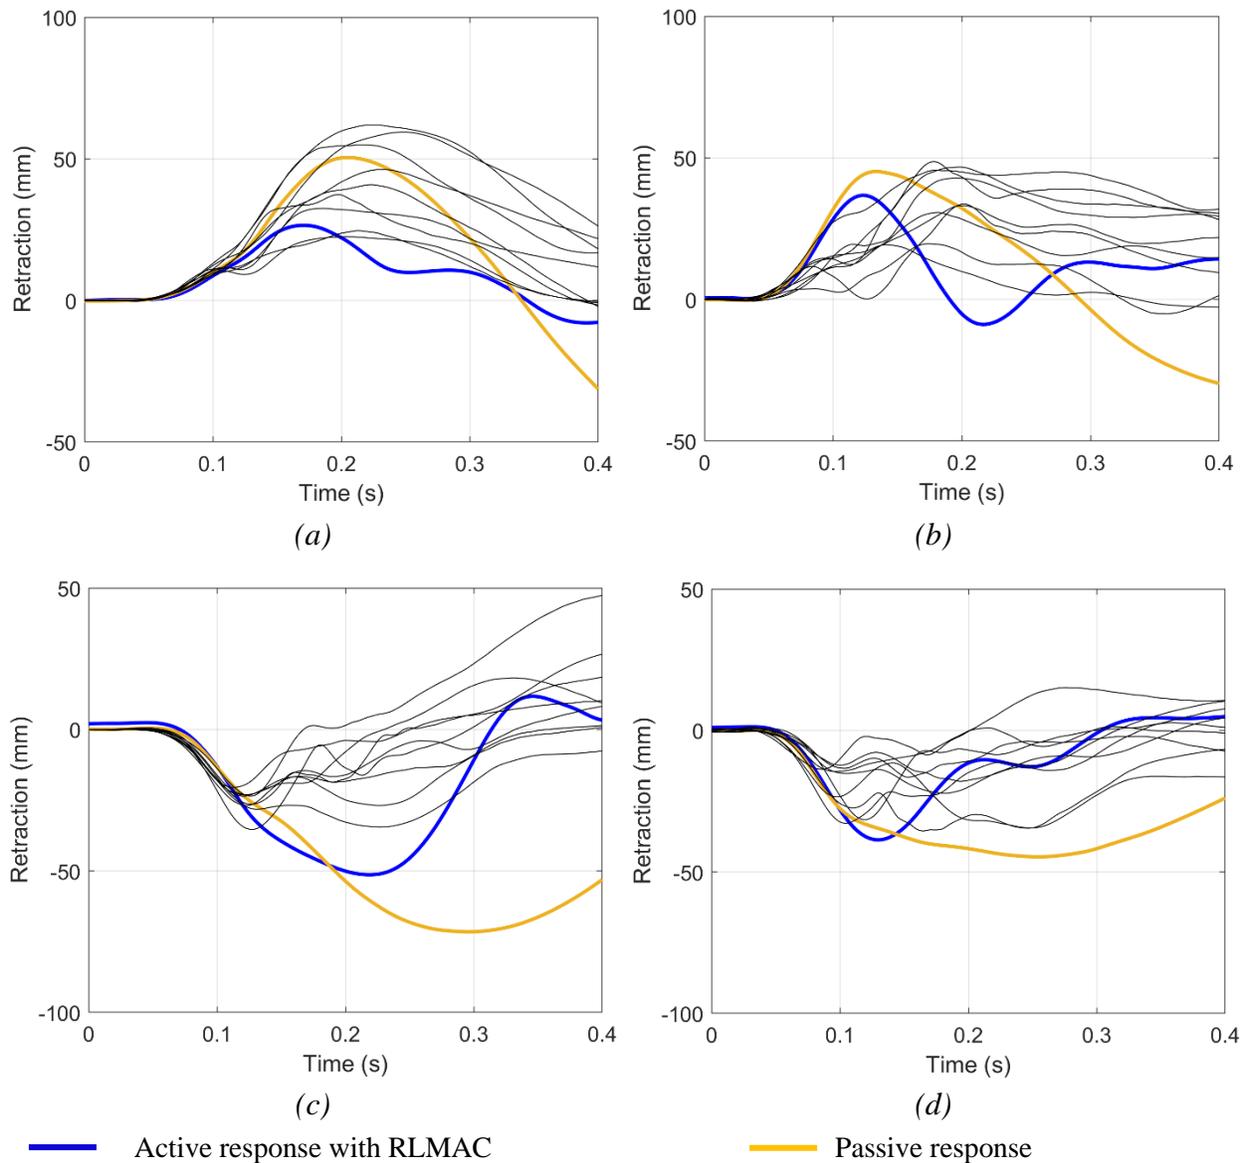
Figure 8-11: The head retraction measured for RLMAC and the passive UVa neck model (a) frontal impact arm relaxed (b) frontal impact arm braced (c) rear impact arm relaxed (d) rear impact arm braced.

The neck model with RLMAC was able to control the head motion and maintain the stability of the spine during the simulations. The active neck model was slightly stiffer in the initial part of the front impact condition with the arm relaxed, however, in all four cases the head retraction response of the active model was within the test measurements and reduced near the end of the simulation time. The passive model however was unstable in all the simulations as can be seen in the head retraction curves in Figure 8-11.

In the volunteer sled studies, the kinematics of the torso was measured between the sternal notch and the C7 spinous process (Fice et al., 2021). In the simulations, the torso kinematics was used as boundary conditions for T1. As there can be some minor difference between the kinematics measurement and the actual rotation of the T1 in the volunteers, a sensitivity study was carried out to analyze the effect of the rotation of the T1 to the head CoM rotation. The T1 rotation was scaled by a factor of 1.25 and 0.75 in the frontal and rear impact cases and the four simulations were repeated with the scaled boundary conditions (BCs). Figure 8-12 shows the T1 angle as well as the head CoM rotation with respect to the global coordinates. The control signal was set to maintain 0° error between the T1 and the head CoM. The RLMAC was able to control the head motion even when the rotations applied at T1 were scaled. The head rotations in the model followed the trend of the T1 rotations. The onset time of the peak head rotations remained almost the same in the three T1 BCs and stabilized near the end of the simulation when the T1 was also near stability. This shows that the trained RLMAC can maintain the head and neck stability even in case if the rotations at the T1 is higher than what is measured at the torso.
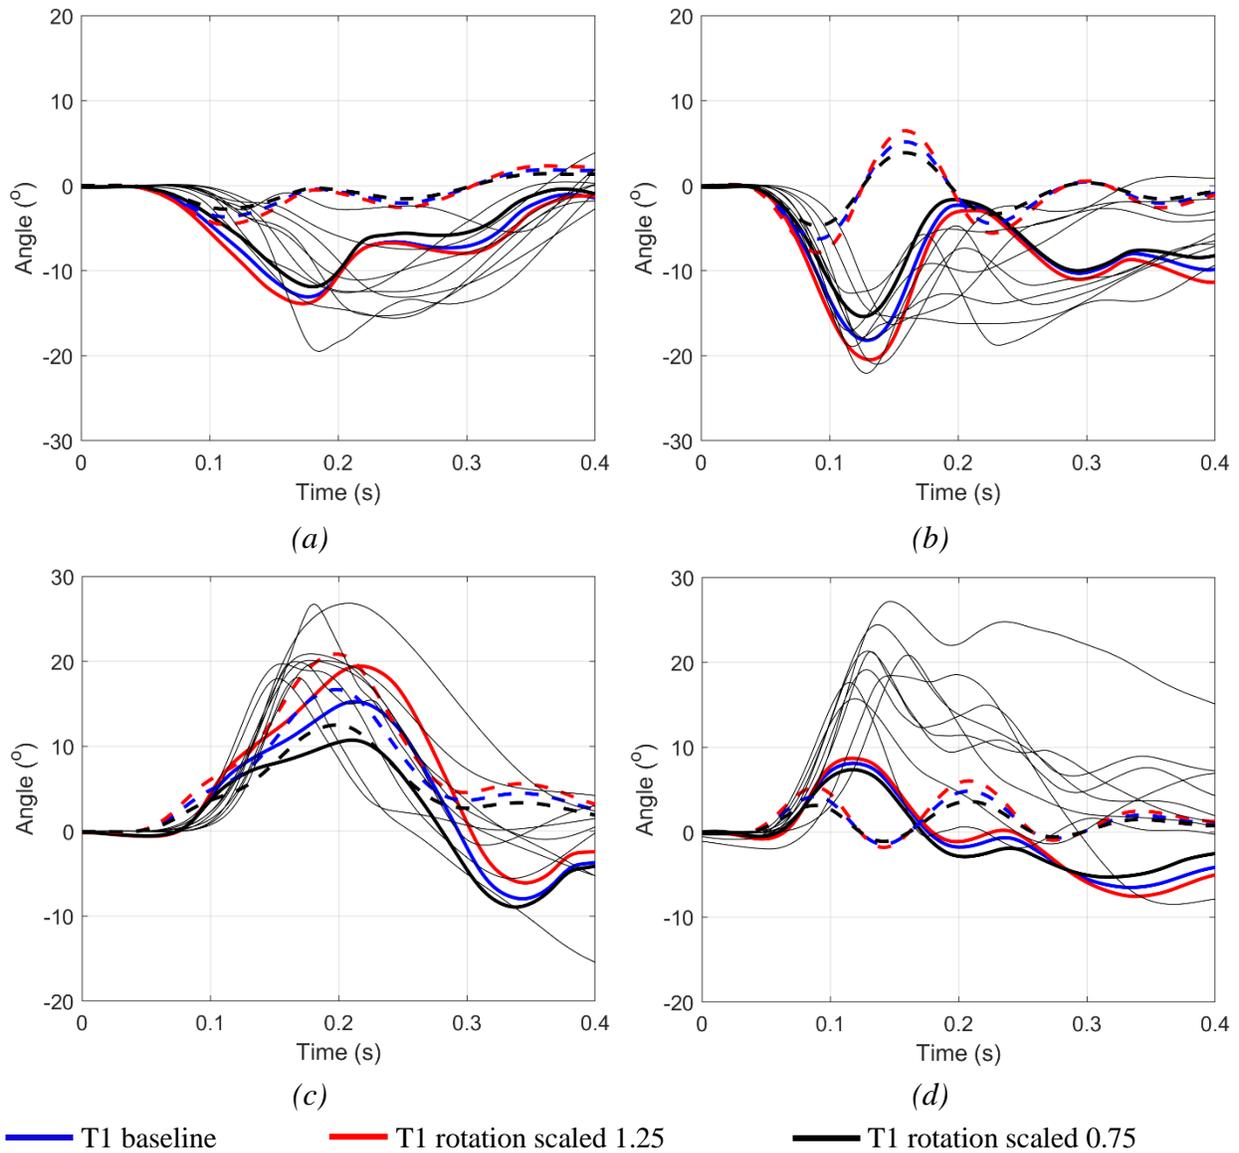
Figure 8-12: Sensitivity of head CoM rotation to the T1 rotation (a) frontal impact arm relaxed (b) frontal impact arm braced (c) rear impact arm relaxed (d) rear impact arm braced.

8.3.3 Response to padded lateral impacts

Padded lateral impacts were simulated by applying the impact forces measured in the experiments (Reynier et al., 2020) to the side of the neck model, 15 mm above the head CoM. A stiffness of 250 N/m and damping of 10 Ns/m were applied in the vertical (Z) direction to simulate the lower body weight. Damping of 10 Ns/m and 0.5 Nms/deg was applied along the lateral (Y) direction

and sagittal (X) rotation. The stiffness values were finalized after performing a sensitivity study

with the passive model to match the volunteer data under passive conditions (Reynier et al., 2020).
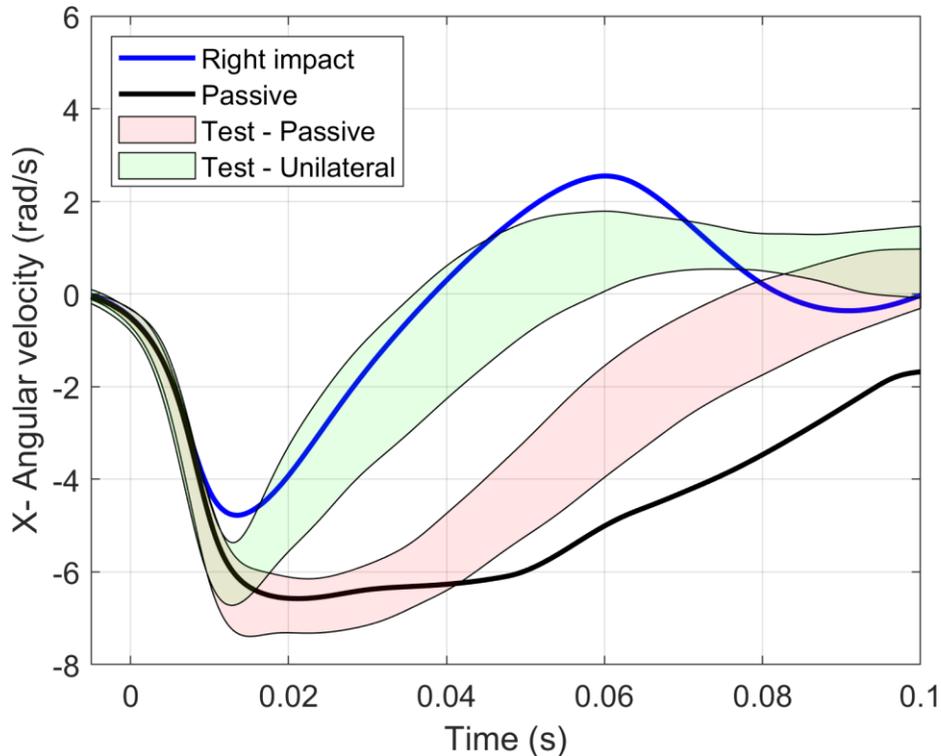
Figure 8-13: Response of the active and passive UVa head model to lateral padded impacts.

Figure 8-12 compares the angular velocity of the passive neck model and the RLMAC with the

passive and unilateral corridors from the human subject tests. The stiffness at the T1 was adjusted

for the passive simulation curve to match with the volunteer head kinematics under passive

conditions. The response of the active model was similar to the unilateral case, however, the peak

angular velocity generated by the RLMAC was lower than the test peaks.

8.4 Discussions

The ability of the RLMAC to develop responses under novel environments like added inertial loads

and impact loads was evaluated in this chapter. Inertial loads were applied by adding the mass and

inertial properties of an American football helmet (Vicis zero 1) weighing 2.193 kg. Helmets affect

the kinematics and injury risks at the head and neck region in both helmet-to-helmet impacts and impacts with stationary objects such as ground (Funk et al., 2018). Helmets also increase neck loads resulting in a greater chance of fatigue (Funk et al., 2018; Harrison et al., 2015; Thuresson et al., 2005).

In the present chapter, the inertial properties of the Vicis zero 1 helmet were added at a distance from the head CoM measured previously by Funk et al. (2018). With the added inertial loads of the helmet, the neck model with RLMAC was simulated for stability under gravity. The trained RLMAC was able to stabilize the head with the added helmet mass (Figure 8-5). The helmeted model had higher joint energy at the stabilized state however the joint energy was still lower compared to when the extension or flexion of the head was carried out. The addition of the helmet lead to 15 % higher vertical compression of the vertebral joints. Simulations were performed to rotate the head to 20^{o} extension and 30^{o} axial rotation angles. The simulations showed that during goal-directed rotations of the head, the addition of a helmet increases the joint energy. The joint energy was significantly (30%) higher at the final position in the helmeted model in the case of a 20^{o} head extension, as addition of the helmet inertia increased the moments required to stabilize the head thus putting higher load at the joints. The inclusion of added mass to the head like helmets for long term may lead to joint and muscle fatigue, causing discomfort and neck strain. Previous studies have found that adding the helmet mass shifts the combined center of mass of the head forward thus putting moment loads at the neck (Gallagher et al., 2008). Additional mass may also affect the ability of the wearer to discharge duties properly along with neck fatigue (Gaur et al., 2013). The RLMAC framework provides a tool to study the effect of added mass on the spine at different postures which can be used while developing helmets for specific purposes.

The trained RLMAC was also simulated in low-speed impact scenarios. T1 translation and rotation measured in the sagittal plane in human subject studies (Fice et al., 2021) were applied at the T1 of the neck model and the resulting translation and rotations were measured. Simulations were carried out to minimize the angular error of the head CoM with the T1 as the control objective. The RLMAC was able to control the head motion and maintain the head stability throughout the application of load, compared to the passive model which was unstable. The RLMAC was able to maintain the stability of the head in both the frontal and rear impact scenarios when the T1 kinematics corresponding to the relaxed and braced conditions were applied. The head angular response in one of the loadcaes – rear impact with arm braced was found to be stiff compared to the volunteer responses. The RLMAC was tasked to minimize any angular rotations between the head and the T1 during the simulations. It is possible that the prescribed control target differed from the intent of the volunteers in the specific case. Also, the boundary condition that was applied at T1 was averaged for the volunteers and was measured in the sternal notch. It is possible that the rotation at the T1 was slightly higher than what the accelerators measured in the volunteers. The RLMAC was also able to control and stabilize the head motion when the rotations at the T1 were scaled. The evaluation of the RLMAC to control the head movements in impact conditions is important as the RLMAC was trained while T1 was fixed and in the impact cases, the T1 is both rotated and translated. The results show that training the neck model under common everyday activities like head rotations can be extended to low speed impact conditions in which muscle activities play an important role (Iwamoto et al., 2012; White et al., 2009) and where the both the load magnitude and the boundary conditions are different than what the RLMAC was trained for.

The neck model was simulated in padded impact conditions to evaluate the response when the impact load is applied laterally. Forces representing the impact conditions were applied to the head

and appropriate stiffness was applied at the T1 to simulate the lower body inertia during the impacts. The response of the RLMAC was similar to that of the volunteers in the unilateral conditions. However, the angular velocity was lower than that observed in the tests. The stiffness at the T1 was calibrated with the passive neck model by comparing it with the response of the human subjects in passive conditions. However, even in the passive conditions, the volunteers may be having a baseline level of activation which may have stiffened the T1, thus leading to the lower peaks observed in the active model (Figure 8-13).

In some previous studies, the linear gains were varied depending on the external load cases (Happee et al., 2017; Zheng et al., 2021). However, no changes to the trained RLMAC were made in this chapter. The RLMAC trained in chapter 6 for generating goal-directed head movements was used in the present chapter to predict the head and neck responses in impact cases.

8.5 Conclusions

The ability of the trained RLMAC has been evaluated to generate responses under inertial and impact loads. The RLMAC was only trained for generating goal-directed motions, therefore the load-cases mentioned in this chapter were novel to the RLMAC controller. The RLMAC was able to control the head motions in both scenarios showing the robustness of the reinforcement learning framework. RLMAC can be further used to design and develop improved helmets and injury countermeasures by studying the effects of muscle activations in such scenarios. Reinforcement learning also gives us an opportunity to train human body models under scenarios where collecting data is convenient like goal-directed movement and weightlifting scenarios and then use the data train and predict responses where gathering data is difficult due to risk of injury. The results from this chapter show that RLMAC does not require trainings for multiple load cases which can be computationally expensive, rather, it can extrapolate the trained responses to novel loads.

Chapter 9 – Conclusions

Active muscle forces can affect the biomechanical response of humans to external forces caused by impact events and as a result, need to be considered while simulating human body models (HBMs) in certain impact conditions. From the perspective of automotive safety, active muscle forces have the most effect on occupant biomechanics in relatively low-acceleration events where the occupant has time to react. Thus, incorporating the active muscle forces in HBMs may be critical for accurate predictions of the kinematics and resulting injury risk of an occupant in scenarios such as pre-impact autonomous braking and rollover. Accuracy in predicting the occupant response in these types of scenarios is important for the proper development of injury countermeasures. Active muscle control using the closed-loop feedback mechanisms are the current state-of-the-art for HBMs used for automotive safety research, but this approach has limitations; the closed-loop controllers rely on linear feedback gains which are tuned for a narrow band of loads, and they require pre-assigned grouping of muscles as agonists and antagonists which is not always obvious to identify beforehand. Moreover, the closed-loop controllers cannot account for the adaptation of muscles to changes in external loads that result in a change in activation levels across muscles. The relationship between change in muscle activation levels to external loads is complex in nature and cannot be accurately represented with linear feedback gains.

Hence, the goal of the dissertation was to explore the use of reinforcement learning (RL) for muscle control in HBMs. To achieve this objective, a comprehensive evaluation of RL algorithms was performed for synthesizing the desired joint kinematics in HBMs. Initially, the RL muscle activation control (RLMAC) was implemented on a simplified multibody (MB) model of the elbow joint to demonstrate the feasibility of this approach. After the utility of this method was proved on the elbow MB model, the RLMAC was tested in its ability to generate biofidelic head

kinematics in the more complex head and neck body region. For this purpose, an MB model of the head and neck region was developed, and the passive responses of the model were validated. The RLMAC was then integrated with the neck model and trained for generating goal-directed head motions in the three anatomical planes. Human subject testing was performed to generate head rotation data to validate the response of the trained RLMAC. Finally, the ability of the RLMAC trained for generating targeted head rotations to produce acceptable responses was evaluated under different anthropometries, added inertial loads, and in motor vehicle environments.

9.1 Major contributions

9.1.1 Development of fast-running multibody mode of the human head and neck region

A fast-running MB model of the head and neck region was developed in MATLAB multibody toolbox for the present dissertation. The neck model had an anatomy of a 50^{th} percentile male and was based on scanning data from previous finite element (FE) studies (Fice et al., 2011; Panzer et al., 2011). The neck model was composed of eight rigid vertebrae and the head was represented visually using the Hybrid III headform; these structures were assigned inertial properties taken from the 50^{th} percentile male FE model. The stiffness between each vertebra, representing the intervertebral soft tissues (disks and ligaments), was represented using 6 degrees of freedom (DoF) non-linear joints (Chang et al., 1992; Dibb et al., 2013; Liu et al., 1982; Shea et al., 1991; Yoganandan et al., 2007). Forty-six Hill-type muscles (Bayer et al., 2017; Haeufle et al., 2014) were added into the MB model using origin and insertion data from the FE model. The passive response of the neck model was validated under a rear impact scenario (Brian D. Stemper et al., 2004). The utility of the neck model was extended to include the ability to scale the geometry and material properties to represent different anthropometries, as demonstrated in chapter 7. The

response of the neck model scaled to the anatomy of 5^{th} percentile females was validated under the rear impact dataset generated on female specimens (Humm et al., 2021).

The developed MB model can be used to simulate response under a range of loading scenarios for quick analysis of head and neck response and injury analysis. The neck model can also be integrated with other toolboxes of MATLAB for statistical analysis and optimization studies which require multiple iterations.

9.1.2 Development of RLMAC for muscle control

A reinforcement learning-based muscle control framework (RLMAC) was developed in this dissertation. A deep deterministic policy gradient (DDPG) agent (Lillicrap et al., 2019) was initially used to develop the RLMAC and integrate it with a MB model of a human arm. The RLMAC could generate the desired kinematics about the elbow joint and could generate reasonable responses to novel loading environments.

The RLMAC was then integrated with the neck model for synthesizing goal-directed motions of the neck. Due to the complexity of the head and neck region, a twin-delayed deep deterministic policy gradient (TD3) agent was used for the RLMAC as it has been shown that the TD3 agent is more efficient for outputting continuous control signals (Fujimoto et al., 2018). The RLMAC was initially trained while constraining the neck model to move only in the sagittal plane. The trained RLMAC was able to rotate the head to the prescribed target angles. The RLMAC could also generate the desired head kinematics when the prescribed target signals were different from what was used in the training.

The RLMAC architecture was expanded to include rotations in the three anatomical planes. The trained RLMAC could maintain the stability of the neck model under gravity and control the

movements of the head along the three rotational DoFs. The neck model was within 3° of the target signal and stable at the final position. The muscle activity was also instrumented during the training and found to follow a triphasic activation pattern which is generally associated with fast goal-directed motions (Hannaford and Stark, 1983; Happee, 1992; Marsden et al., 1983). To the best of my knowledge, no previous control studies have produced voluntary neck kinematics using human body models.

9.1.3 Experimental dataset for goal-directed head kinematics in human subjects

A human volunteer study was conducted on healthy male subjects to gather data on voluntary head kinematics. The data gathered was used for tuning and validation of the RLMAC control framework with the neck model. The fast goal-directed head rotations of human subjects were characterized by extension, flexion, and axial rotations. The rotational velocity–time history was measured using accelerometers placed on the head's posterior end and the velocity data were used to generate the angle-time history curve. The volunteer testing was performed to complement the modeling and simulation efforts for this dissertation and the gathered dataset can also be used as validation cases for future control studies.

9.1.4 Applicability of the RLMAC to novel loading scenarios

The range of applicability of the RLMAC was demonstrated in this dissertation. The neck model was scaled to the anatomy of a 5th-percentile female and 95th-percentile male, and the trained RLMAC was used to control the head kinematics in the scaled models. The RLMAC could stabilize the head and synthesize goal-directed head kinematics in the scaled models.

The RLMAC was simulated by attaching additional mass and inertia to the head model representing the properties of an American football helmet (Funk et al., 2018). The RLMAC was able to stabilize the head and could also generate goal-directed head movements with the added

inertial loads showing that RLMAC could adapt to external inertial loads. The effects of the added inertia on the head and spinal responses were also evaluated.

Finally, the trained RLMAC was also evaluated under dynamic conditions to maintain head stability. The neck model was simulated under automotive impact scenarios representing frontal and rear collisions by applying boundary conditions at T1 (Fice et al., 2021). Padded lateral impacts were simulated by applying forces measured in human volunteer experiments (Reynier et al., 2020). The trained RLMAC was able to control the head motions in the dynamic scenarios and maintain head stability throughout the simulation. The response of the RLMAC was compared with the passive neck model with baseline muscle activation. While the RLMAC could control the head kinematics of the neck model, the passive simulation had unstable kinematics. The simulations performed in this dissertation show the importance of incorporating active muscle behavior in HBM simulations.

9.2 Assumptions and limitations

Reinforcement learning algorithms are iterative in nature and hence require a lot of computational resources. As a result, some simplifications were made while developing the models used in this dissertation.

9.2.1 Implementation of RLMAC on the MB model of the human arm

The arm model used for preliminary analysis of RL algorithms for muscle control was simplified with the rigid representation of bones and a revolute joint. Only the extension-flexion motion was considered at the elbow neglecting any translation movement or rotation in any other DoFs. The muscles were modeled as force elements between the insertion points without considering the changing load directions due to the rotation of the elbow. The series element which accounts for

the tendon stiffness was neglected in the Hill-type muscles incorporated in the elbow. However, the passive response of the elbow was validated with available data in the literature illustrating that the simplifications in the model did not affect the overall response. The arm control analysis was a proof-of-concept study aimed at demonstrating the utility of RL algorithms for muscle control, and the assumptions of revolute joint and rigid tendons did not affect the overall objective of the analysis.

9.2.2 Development of the MB model of the human head and neck

The neck model was developed for integration with the RLMAC framework as a part of this dissertation. To some extent, the neck model was also simplified for training purposes. The non-linearities at the vertebral joints were simplified as bilinear curves for both translational and rotational degrees of freedom. No interaction or contact between the muscles was considered. The muscles with wide insertion points were split into multiple strands to capture the behavior at the attachment area to bones. However, the maximum forces (F_{max}) at each strand was considered the same. Strands belonging to the same muscle were also activated identically without considering the relative location or the length of each strand. The passive neck model, however, responds well under automotive loading scenario (Brian D. Stemper et al., 2004) and also could be integrated with the RLMAC to generate goal-directed responses which was the objective of the present study.

9.2.3 Training and muscle activation output by the RLMAC

The RLMAC was trained to generate goal-directed head kinematics in the three rotational DoFs. The training of the RLMAC took around 15000 iterations to converge, taking approximately 168 hours on 100 computer cores of the UVa high performance Rivanna cluster. Although the RLMAC took considerable computation time to converge in this scenario, the total iterations were less than what is typical in similar RL control studies (Akimov, 2020; La Barbera et al., 2022). The control

model can be further trained to improve the response of the neck to loading scenarios beyond the scope of this dissertation.

The muscle activations that were generated by the trained RLMAC were also analyzed. The muscle activations that generated the desired head and neck kinematics followed the three phase activation pattern that have been observed in various goal directed head movements (Happee, 1992; Wadman et al., 1979). However, the multifidus and the rectus capitis minor had high activations during the goal-directed head simulations. For most part of the fast head movements, the activation level of the multifidus was higher than 0.8.

(a)   *(b)*
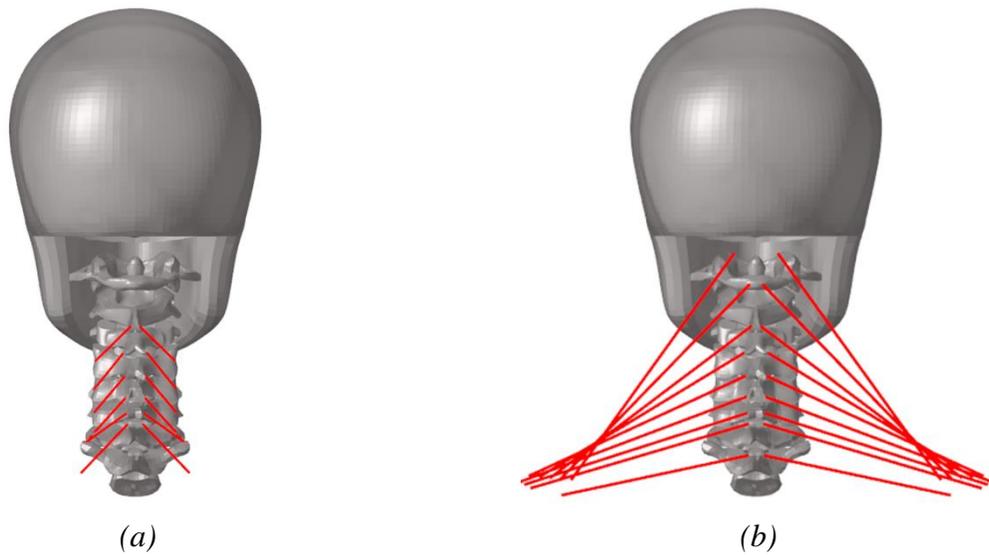
Figure 9-1: Neck muscles (a) Multifidus (b) Trapezius.

The multifidus muscle is located at the back of the spine and extends from sacrum to the skull. In the neck model, the multifidus runs parallel to the spine from T1 to the back of the head and is divided into six strands across the length of the muscle (Figure 9-1a). The maximum force (F_{max}) and the activation (a_t) of every strand is considered identical in the model. Morimoto et al. (2013) measured the muscle activations in rugby tackle position and found that the trapezius had different activation levels at the upper, middle, and lower position. The multifidus may also have different

activation levels depending upon the location of the muscle segment which was not considered in this study. The effects of activating long muscles like trapezius and multifidus (Figure 9-1b) with different a_t depending on the location along the spine can be considered in future studies. The rectus capitis minor is a small muscle at the back of the neck. The muscle may have high activations to compensate for the lack of some of the hyoid group of muscles in the model, but more investigation is needed in this area.

While training the RLMAC for rotations in three DoFs, a symmetry factor was included in the reward function which ensured muscle activations were symmetrical for head kinematics about the sagittal plane. Increasing the weight of the symmetry factor can improve the symmetric response of the RLMAC, however, changing the relative weights of the reward components will increase the iterations required for training and convergence.

9.3 Future work

A reinforcement learning-based muscle control framework was developed for this dissertation. The RLMAC was able to control the head kinematics for goal-directed head motions even for novel target signals. The trained RLMAC was also able to stabilize the head under dynamic scenarios and could be used with scaled anthropometrics. However, this is the first iteration of the RLMAC, and the response of the head and neck controlled by the RLMAC can be improved in future studies.

9.3.1 Improvement in the reward function and training

The reward function developed and used in this dissertation was successful in training the neck model for goal-directed head kinematics. The function rewards the RLMAC for reducing the angular error of the head CoM concerning a prescribed target signal. However, to reduce the

215

iterations required for convergence, the target signals were varied in multiples of 5. In future training, the target signals could be varied in finer increments to improve the training of the RLMAC. In this dissertation only the head rotations are considered, in future studies the head CoM translations could also be implemented in the reward function for incorporating the ability to move the head without changing the head CoM orientation. Also in the dissertation, it is assumed that reducing the joint energy will also result in the reduction of muscle activations, and thus muscle fatigue. However, penalizing the high muscle activities as a part of the reward in the neck muscles also needs to be evaluated in the future.

The RLMAC was developed and trained for stabilizing the head as well as performing goal-directed head movements. However, the RLMAC cannot adjust the initial activation state into relaxed or co-contracted at the stable position. Relaxed and co-contracted state of the neck muscle have showed difference in head kinematics under application of external loads (Homayounpour et al., 2021; Reynier et al., 2020). In future studies, the state of the RLMAC and the reward function can be modified to enhance the capability of the RLMAC in implementing a relaxed or co-contracted state in the neck model at the stabilized position.

In the current study, the TD3 algorithm was used for the development of the RLMAC for omnidirectional control of head kinematics. In future training, some recently developed RL algorithms like soft-actor critic (SAC) can also be implemented that have been claimed to produce better results for continuous control tasks (Haarnoja et al., 2018).

9.3.2 Implementation of the RLMAC with finite element (FE) models

While the simplified MB model of the head and neck is suitable for training and preliminary analysis, the trained RLMAC can be transferred to a more detailed FE model of the head and neck region to obtain the tissue-level responses for the prescribed control scenario. There have been some previous attempts at linking FE models with RL codes for biomechanics simulations

216

(Iwamoto et al., 2012; Joos et al., 2020; Min et al., 2018). But due to the iterative nature of the RL algorithms, the FE models needed to be simplified for training purposes. With RLMAC, the training can be performed on fast-running models like the neck model developed in the current dissertation, and the training can be transferred to a more complex and detailed FE 'twin' capable of predicting injurious responses.

9.3.3 Implementation of the RLMAC in full-body HBM

Reinforcement learning-based muscle control frameworks can be also incorporated into a full-body HBM. Multiple agents can be used to control kinematics around different joints of the HBM. The agents can be either trained simultaneously or the agents can be trained independently in isolated joints like the arm MB model and the neck model, and the trained agents can be combined into the HBM to simulate joint responses at the whole-body level (Zhang et al., 2021). The multi-agent RLMAC can act as an improved tool for the prediction of responses under scenarios like low-speed automotive impacts where the responses of the occupants are dependent on the interaction with peripheral components which is difficult to reproduce at the joint-only level. Multi-agent HBMs will also aid in the development of assistive devices for rehabilitation (Luo et al., 2021) that may require modeling the kinematics of two or more joints for accurate representation.

9.4 Summary

For accurate prediction of responses under some external loading environments, the incorporation of active muscle forces in HBMs becomes important. In this dissertation, the muscle control framework has been developed using reinforcement learning algorithms. The RLMAC developed in this dissertation was able to actuate the muscles individually without any user input of muscle coordination or synergy. The RLMAC was implemented in the MB model of the head and neck

for control of kinematics in three rotational DoFs. An experimental dataset on the head rotations in human subjects was generated to validate the response of the trained RLMAC. The trained RLMAC was able to generate biofidelic responses under novel external loads for which the RLMAC was not specifically trained. Reinforcement learning can be used to train the HBMs in scenarios that have the potential for generating human subject data, like in a weightlifting scenario (Denizdurduran et al., 2022) and the trained agent can provide a good generalization in situations where data generation is difficult (like in impact scenarios). Reinforcement learning can be implemented to study the effect of long-term exposure to inertial loads on muscle and joint fatigue as well as aid in the development of improved injury countermeasures and assistive devices for rehabilitation.

The results in this dissertation can be further improved by modifying the architecture of the RLMAC to include new state parameters and updated reward functions. It can be worthwhile to explore how the state parameters and components of the reward function affect the control response and potentially improve the RLMAC architecture. Volunteer studies can also be performed to gather data that can aid in the development of the improved control architecture.

REFERENCES

Akimov, D., 2019. Distributed soft actor-critic with multivariate reward representation and knowledge distillation. arXiv preprint arXiv:1911.13056.

Arbogast, K.B., Balasubramanian, S., Seacrist, T., Maltese, M.R., Garcia-Espana, J.F., Hopely, T., Constans, E., López-Valdés, F.J., Kent, R.W., Tanji, H. and Higuchi, K., 2009. Comparison of kinematic responses of the head and spine for children and adults in low-speed frontal sled tests. Stapp car crash journal, 53, p.329.

Bahler, A.S., Fales, J.T. and Zierler, K.L., 1967. The active state of mammalian skeletal muscle. The Journal of general physiology, 50(9), pp.2239-2253.

Bayer, A., Schmitt, S., Günther, M. and Haeufle, D.F.B., 2017. The influence of biophysical muscle properties on simulating fast human arm movements. Computer methods in biomechanics and biomedical engineering, 20(8), pp.803-821.

Beeman, S.M., Kemper, A.R. and Duma, S.M., 2016. Neck forces and moments of human volunteers and postmortem human surrogates in low-speed frontal sled tests. Traffic injury prevention, 17(sup1), pp.141-149.

Bengio, Y., Courville, A. and Vincent, P., 2013. Representation learning: A review and new perspectives. IEEE transactions on pattern analysis and machine intelligence, 35(8), pp.1798-1828.

Ben-Itzhak, S. and Karniel, A., 2008. Minimum acceleration criterion with constraints implies bang-bang control as an underlying principle for optimal trajectories of arm reaching movements. Neural Computation, 20(3), pp.779-812.

219

Bernstein, N., 1966. The co-ordination and regulation of movements. The co-ordination and regulation of movements. Pergamon Press, Oxford, New York.

Berret, B., Darlot, C., Jean, F., Pozzo, T., Papaxanthis, C. and Gauthier, J.P., 2008. The inactivation principle: mathematical solutions minimizing the absolute work and biological implications for the planning of arm movements. PLoS computational biology, 4(10), p.e1000194.

Bobbert, M., 2001. Dependence of human squat jump performance on the series elastic compliance of the triceps surae: a simulation study. Journal of Experimental Biology, 204(3), pp.533-542.

Bose, D., Segui-Gomez, ScD, M. and Crandall, J.R., 2011. Vulnerability of female drivers involved in motor vehicle crashes: an analysis of US population at risk. American journal of public health, 101(12), pp.2368-2373.

Bove, M., Fenoggio, C., Tacchino, A., Pelosin, E. and Schieppati, M., 2009. Interaction between vision and neck proprioception in the control of stance. Neuroscience, 164(4), pp.1601-1608.

Bowman, B.M., Schneider, L.W., Rohr, P.R. and Mohan, D., 1981. Simulation of head/neck impact responses for helmeted and unhelmeted motorcyclists. SAE Transactions, pp.3318-3344.

Brolin, K. and Halldin, P., 2004. Development of a finite element model of the upper cervical spine and a parameter study of ligament characteristics. Spine, 29(4), pp.376-385.

Brolin, K., Halldin, P. and Leijonhufvud, I., 2005. The effect of muscle activation on neck response. Traffic injury prevention, 6(1), pp.67-76.

Brolin, K., Hedenstierna, S., Halldin, P., Bass, C. and Alem, N., 2008. The importance of muscle tension on the outcome of impacts with a major vertical component. International journal of crashworthiness, 13(5), pp.487-498.

Brolin, K., Stockman, I., Subramanian, H., Gras, L.L. and Östh, J., 2015. Development of an Active 6-Year-Old Child Human Body Model for Simulation of Emergency Events. Proceedings-International Research Council on the Biomechanics of Injury, 9, p.11.

Bruneau, D.A. and Cronin, D.S., 2020. Head and neck response of an active human body model and finite element anthropometric test device during a linear impactor helmet test. Journal of Biomechanical Engineering, 142(2).

Budziszewski, P., van Nunen, E., Mordaka, J.K. and Kędzior, K., 2008, September. Active controlled muscles in numerical model of human arm for movement in two degrees of freedom. In IRCOBI Conference Proceedings, International Conference on the Biomechanics of impact, Bern, Switzerland (Vol. 17, p. 19).

Buhrmann, T. and Di Paolo, E.A., 2014. Spinal circuits can accommodate interaction torques during multijoint limb movements. Frontiers in computational neuroscience, 8, p.144.

Cannon, S.C. and Zahalak, G.I., 1982. The mechanical behavior of active human skeletal muscle in small oscillations. Journal of biomechanics, 15(2), pp.111-121.

Cappon, H., Mordaka, J., Van Rooij, L., Adamec, J., Praxl, N. and Muggenthaler, H., 2007. A computational human model with stabilizing spine: a step towards active safety (No. 2007-01-1171). SAE Technical Paper.

Carlsson, S. and Davidsson, J., 2011, September. Volunteer occupant kinematics during driver initiated and autonomous braking when driving in real traffic environments. In Proceedings of the IRCOBI Conference.

Chan, H., Albert, D.L., Gayzik, F.S. and Kemper, A.R., 2022. Occupant kinematics of braced 5th percentile female and 50th percentile male volunteers in low-speed frontal and frontal-oblique sled tests. In IRCOBI.

Chancey, V.C., Nightingale, R.W., Van Ee, C.A., Knaub, K.E. and Myers, B.S., 2003. Improved estimation of human neck tensile tolerance: reducing the range of reported tolerance using anthropometrically correct muscles and optimized physiologic initial conditions (No. 2003-22-0008). SAE Technical Paper.

Chancey, V.C., Ottaviano, D., Myers, B.S. and Nightingale, R.W., 2007. A kinematic and anthropometric study of the upper cervical spine and the occipital condyles. Journal of biomechanics, 40(9), pp.1953-1959.

Chang, H., Gilbertson, L.G., Goel, V.K., Winterbottom, J.M., Clark, C.R. and Patwardhan, A., 1992. Dynamic response of the occipito–atlanto–axial (C0-C1-C2) complex in right axial rotation. Journal of orthopaedic research, 10(3), pp.446-453.

Cheng, C.H., Lin, K.H. and Wang, J.L., 2008. Co-contraction of cervical muscles during sagittal and coronal neck motions at different movement speeds. European journal of applied physiology, 103, pp.647-654.

Choi, H., 2003. Quantitative assessment of co-contraction in cervical musculature. Medical engineering & physics, 25(2), pp.133-140.

Correia, M.A., McLachlin, S.D. and Cronin, D.S., 2021. Vestibulocollic and cervicocollic muscle reflexes in a finite element neck model during multidirectional impacts. Annals of Biomedical Engineering, 49(7), pp.1645-1656.

Correia, M.A., McLachlin, S.D. and Cronin, D.S., 2020. Optimization of muscle activation schemes in a finite element neck model simulating volunteer frontal impact scenarios. Journal of biomechanics, 104, p.109754.

Crowder, D.C., Abreu, J. and Kirsch, R.F., 2021. Hindsight experience replay improves reinforcement learning for control of a MIMO musculoskeletal model of the human arm. IEEE Transactions on Neural Systems and Rehabilitation Engineering, 29, pp.1016-1025.

Crowninshield, R.D. and Brand, R.A., 1981. A physiologically based criterion of muscle force prediction in locomotion. Journal of biomechanics, 14(11), pp.793-801.

Cullen, K.E., 2012. The vestibular system: multimodal integration and encoding of self-motion for motor control. Trends in neurosciences, 35(3), pp.185-196.

Davis, M.L., Koya, B., Schap, J.M., Gayzik, F.S., 2016. Development and Full Body Validation of a 5th Percentile Female Finite Element Model. Stapp Car Crash J. 60, 509–544.

Davoodi, R. and Andrews, B.J., 1998. Computer simulation of FES standing up in paraplegia: A self-adaptive fuzzy controller with reinforcement learning. IEEE Transactions on Rehabilitation Engineering, 6(2), pp.151-161.

de Bruijn, E., Van der Helm, F.C.T. and Happee, R., 2016. Analysis of isometric cervical strength with a nonlinear musculoskeletal model with 48 degrees of freedom. Multibody System Dynamics, 36, pp.339-362.

De Jager, M., Sauren, A.A.H.J., Thunnissen, J. and Wismans, J.S.H.M., 1996. A global and a detailed mathematical model for head-neck dynamics. SAE transactions, pp.1899-1911.

De Vlugt, E., Schouten, A.C. and Van Der Helm, F.C., 2006. Quantification of intrinsic and reflexive properties during multijoint arm posture. Journal of neuroscience methods, 155(2), pp.328-349.

Delp, S.L., Anderson, F.C., Arnold, A.S., Loan, P., Habib, A., John, C.T., Guendelman, E. and Thelen, D.G., 2007. OpenSim: open-source software to create and analyze dynamic simulations of movement. IEEE transactions on biomedical engineering, 54(11), pp.1940-1950.

Dempster, W.T. and Gaughran, G.R., 1967. Properties of body segments based on size and weight. American journal of anatomy, 120(1), pp.33-54.

Denizdurduran, B., Markram, H. and Gewaltig, M.O., 2022. Optimum trajectory learning in musculoskeletal systems with model predictive control and deep reinforcement learning. Biological cybernetics, pp.1-16.

Devane, K., Chan, H., Albert, D., Kemper, A. and Gayzik, F.S., 2022. Implementation and calibration of active small female and average male human body models using low-speed frontal sled tests. Traffic injury prevention, pp.1-6.

Dibb, A.T., Cox, C.A., Nightingale, R.W., Luck, J.F., Cutcliffe, H.C., Myers, B.S., Arbogast, K.B., Seacrist, T. and Bass, C.R., 2013. Importance of muscle activations for biofidelic pediatric neck response in computational models. Traffic injury prevention, 14(sup1), pp.S116-S127.

Driess, D., Zimmermann, H., Wolfen, S., Suissa, D., Haeufle, D., Hennes, D., Toussaint, M. and Schmitt, S., 2018, May. Learning to control redundant musculoskeletal systems with neural networks and SQP: exploiting muscle properties. In 2018 IEEE International Conference on robotics and automation (ICRA) (pp. 6461-6468). IEEE.

Eckner, J.T., Oh, Y.K., Joshi, M.S., Richardson, J.K. and Ashton-Miller, J.A., 2014. Effect of neck muscle strength and anticipatory cervical muscle activation on the kinematic response of the head to impulsive loads. The American journal of sports medicine, 42(3), pp.566-576.

Ejima, S., Ito, D., Satou, F., Mikami, K., Ono, K., Kaneoka, K. and Shiina, I., 2012. Effects of pre-impact swerving/steering on physical motion of the volunteer in the low-speed side-impact sled test. In Proceedings of the IRCOBI Conference (pp. 352-366).

Ejima, S., Ono, K., Holcombe, S., Kaneoka, K. and Fukushima, M., 2007. A study on occupant kinematics behaviour and muscle activities during pre-impact braking based on volunteer tests. In Proceedings of ircobi (international research council on the biomechanics of injury) conference 2007, held maastricht, the netherlands, september 2007.

Ejima, S., Zama, Y., Satou, F., Holcombe, S., Ono, K., Kaneoka, K. and Shiina, I., 2008, September. Prediction of the physical motion of the human body based on muscle activity during pre-impact braking. In Proceedings of the IRCOBI Conference (pp. 163-175).

Eppinger, R.H., Marcus, J.H. and Morgan, R.M., 1984. Development of dummy and injury index for NHTSA's thoracic side impact protection research program. SAE transactions, pp.359-387.

Fawzi, A., Balog, M., Huang, A., Hubert, T., Romera-Paredes, B., Barekatain, M., Novikov, A., R Ruiz, F.J., Schrittwieser, J., Swirszcz, G. and Silver, D., 2022. Discovering faster matrix multiplication algorithms with reinforcement learning. Nature, 610(7930), pp.47-53.

Feller, L., Kleinbach, C., Fehr, J. and Schmitt, S., 2016, September. Incorporating muscle activation dynamics into the Global Human Body Model. In Proceedings of the IRCOBI conference, Malaga, Spain.

Fice, J.B., Cronin, D.S. and Panzer, M.B., 2011. Cervical spine model to predict capsular ligament response in rear impact. Annals of biomedical engineering, 39, pp.2152-2162.

Fice, J.B., Mang, D.W., Ólafsdóttir, J.M., Brolin, K., Cripton, P.A., Blouin, J.S. and Siegmund, G.P., 2021. Neck muscle and head/neck kinematic responses while bracing against the steering wheel during front and rear impacts. Annals of biomedical engineering, 49, pp.1069-1082.

Fice, J.B., Siegmund, G.P. and Blouin, J.S., 2018. Neck muscle biomechanics and neural control. Journal of neurophysiology, 120(7), pp.361-371.

Flament, D., Hore, J. and Vilis, T., 1984. Braking of fast and accurate elbow flexions in the monkey. The Journal of Physiology, 349(1), pp.195-202.

Forbes, P.A., Dakin, C.J., Vardy, A.N., Happee, R., Siegmund, G.P., Schouten, A.C. and Blouin, J.S., 2013. Frequency response of vestibular reflexes in neck, back, and lower limb muscles. Journal of neurophysiology, 110(8), pp.1869-1881.

Forman, J., Lessley, D., Kent, R., Bostrom, O. and Pipkorn, B., 2006. Whole-body kinematic and dynamic response of restrained PMHS in frontal sled tests. Stapp car crash journal, 50, p.299.

Forman, J., Poplin, G.S., Shaw, C.G., McMurry, T.L., Schmidt, K., Ash, J. and Sunnevang, C., 2019. Automobile injury trends in the contemporary fleet: Belted occupants in frontal collisions. Traffic injury prevention, 20(6), pp.607-612.

Foust, D.R., Chaffin, D.B., Snyder, R.G. and Baum, J.K., 1973. Cervical range of motion and dynamic response and strength of cervical muscles. SAE Transactions, pp.3222-3234.

Fraga, F., van Rooij, L., Happee, R., Wismans, J., Symeonidis, I. and Peldschus, S., 2009, June. Development of a motorcycle rider model with focus on head and neck biofidelity, recurring to line element muscle models and feedback control. In International technical conference on the enhanced safety of vehicles (esv).

Fujimoto, S., Hoof, H. and Meger, D., 2018, July. Addressing function approximation error in actor-critic methods. In International conference on machine learning (pp. 1587-1596). PMLR.

Funk, J.R., Quesada, R.E., Miles, A.M. and Crandall, J.R., 2018. Inertial properties of football helmets. Journal of biomechanical engineering, 140(6), p.064501.

Gallagher, H.L., Caldwell, E. and Albery, C.B., 2008. Neck muscle fatigue resulting from prolonged wear of weighted helmets. General dynamics advanced information systems Dayton Oh.

Gaur, S.J., Joshi, W., Aravindakshan, B. and Aravind, A.S., 2013. Determination of helmet CG and evaluation of neck injury potentials using "Knox Box criteria" and neck torque limits. Indian Journal of Aerospace Medicine, 57(1), pp.37-44.

Gayzik, F.S., Hamilton, C.A., Tan, J.C., McNally, C., Duma, S.M., Klinich, K.D. and Stitzel, J.D., 2009. A multi-modality image data collection protocol for full body finite element model development. SAE Technical Paper, (2009-01), p.2261.

Geigl, B.C., Steffan, H., Dippel, C., Muser, M.H., Walz, F.H. and Svensson, M.Y., 1995. Comparison of head-neck kinematics during rear end impact between standard hybrid iii, rid neck, volunteers and pmto's (post mortal test objects). In Proceedings of the International Research Council on the Biomechanics of Injury conference (Vol. 23, pp. 261-270). International Research Council on Biomechanics of Injury.

Glorot, X., Bordes, A. and Bengio, Y., Deep sparse rectifier neural networks. InProceedings of the fourteenth international conference on artificial intelligence and statistics 2011 Jun 14 (pp. 315-323). In JMLR Workshop and Conference Proceedings.

Goldberg, J.M. and Cullen, K.E., 2011. Vestibular control of the head: possible functions of the vestibulocollic reflex. Experimental brain research, 210, pp.331-345.

Goldberg, J. and Peterson, B.W., 1986. Reflex and mechanical contributions to head stabilization in alert cats. Journal of neurophysiology, 56(3), pp.857-875.

Goldsmith, W., Deng, Y.C. and Merrill, T.H., 1984. Numerical evaluation of the three-dimensional response of a human head-neck model to dynamic loading. SAE transactions, pp.218-234.

Gribble, P.L., Ostry, D.J., Sanguineti, V. and Laboissière, R., 1998. Are complex control signals required for human arm movement? Journal of neurophysiology, 79(3), pp.1409-1424.

Grigg, P., 1994. Peripheral neural mechanisms in proprioception. Journal of Sport Rehabilitation, 3(1), pp.2-17.

Günther, M., Schmitt, S. and Wank, V., 2007. High-frequency oscillations as a consequence of neglected serial damping in Hill-type muscle models. Biological cybernetics, 97(1), pp.63-79.

Haarnoja, T., Zhou, A., Hartikainen, K., Tucker, G., Ha, S., Tan, J., Kumar, V., Zhu, H., Gupta, A., Abbeel, P. and Levine, S., 2018. Soft actor-critic algorithms and applications. arXiv preprint arXiv:1812.05905.

Haeufle, D.F.B., Günther, M., Bayer, A. and Schmitt, S., 2014. Hill-type muscle model with serial damping and eccentric force–velocity relation. Journal of biomechanics, 47(6), pp.1531-1536.

Han, J. and Moraga, C., 1995. The influence of the sigmoid function parameters on the speed of backpropagation learning. In From Natural to Artificial Neural Computation: International

Workshop on Artificial Neural Networks Malaga-Torremolinos, Spain, June 7–9, 1995 Proceedings 3 (pp. 195-201). Springer Berlin Heidelberg.

Hannaford, B. and Stark, L., 1985. Roles of the elements of the triphasic control signal. Experimental neurology, 90(3), pp.619-634.

Happee, R., 1993. Goal-directed arm movements. III: Feedback and adaptation in response to inertia perturbations. Journal of electromyography and kinesiology, 3(2), pp.112-122.

Happee, R., 1992. Time optimality in the control of human movements. Biological cybernetics, 66(4), pp.357-366.

Happee, R., de Bruijn, E., Forbes, P.A. and van der Helm, F.C., 2017. Dynamic head-neck stabilization and modulation with perturbation bandwidth investigated using a multisegment neuromuscular model. Journal of biomechanics, 58, pp.203-211.

Happee, R., de Vlugt, E., and van Vliet, B., 2015, "Nonlinear 2D Arm Dynamics in Response to Continuous and Pulse-Shaped Force Perturbations," Exp. Brain Res., 233(1), pp. 39–52.

Hara, K., Saito, D. and Shouno, H., 2015, July. Analysis of function of rectified linear unit used in deep learning. In 2015 international joint conference on neural networks (IJCNN) (pp. 1-8). IEEE.

Harrison, M.F., Coffey, B., Albert, W.J. and Fischer, S.L., 2015. Night vision goggle-induced neck pain in military helicopter aircrew: a literature review. Aviation, space, and environmental medicine, 86(1), pp.46-55.

Hayes, K.C. and Hatze, H., 1977. Passive visco-elastic properties of the structures spanning the human elbow joint. European journal of applied physiology and occupational physiology, 37, pp.265-274.

Hedenstierna, S., Halldin, P., Brolin, K. and von Holst, H., 2006. Development and evaluation of a continuum neck muscle model. Journal of Biomechanics, 39(1), p.S150.

Heess, N., TB, D., Sriram, S., Lemmon, J., Merel, J., Wayne, G., Tassa, Y., Erez, T., Wang, Z., Eslami, S.M. and Riedmiller, M., 2017. Emergence of locomotion behaviours in rich environments. arXiv preprint arXiv:1707.02286.

Hernandez, F. and Camarillo, D.B., 2019. Voluntary head rotational velocity and implications for brain injury risk metrics. Journal of neurotrauma, 36(7), pp.1125-1135.

Hill, A.V., 1938. The heat of shortening and the dynamic constants of muscle. Proceedings of the Royal Society of London. Series B-Biological Sciences, 126(843), pp.136-195.

Hochreiter, S. and Schmidhuber, J., 1997. Long short-term memory. Neural computation, 9(8), pp.1735-1780.

Hoff, B. and Arbib, M.A., 1993. Models of trajectory formation and temporal interaction of reach and grasp. Journal of motor behavior, 25(3), pp.175-192.

Homayounpour, M., Gomez, N.G., Vasavada, A.N. and Merryweather, A.S., 2021. Cervical muscle activation due to an applied force in response to different types of acoustic warnings. Annals of biomedical engineering, pp.1-13.

Howell, J.N., Chleboun, G. and Conatser, R., 1993. Muscle stiffness, strength loss, swelling and soreness following exercise-induced injury in humans. The Journal of physiology, 464(1), pp.183-196.

Humm, J., Purushothaman, Y., Umale, S., Yoganandan, N., 2021. Segmental Motion Response Corridors of Female Cervical Spines from Gx accelerative loading. In IRCOBI Conference Proceedings.

Huston, R.L., Huston, J.C. and Harlow, M.W., 1978. Comprehensive, three-dimensional head-neck model for impact and high-acceleration studies. Aviation, space, and environmental medicine, 49(1 Pt. 2), pp.205-210.

Iwamoto, M. and Kato, D., 2021. Efficient actor-critic reinforcement learning with embodiment of muscle tone for posture stabilization of the human arm. Neural Computation, 33(1), pp.129-156.

Iwamoto, M. and Nakahira, Y., 2015. Development and validation of the Total HUman Model for Safety (THUMS) version 5 containing multiple 1D muscles for estimating occupant motions with muscle activation during side impacts. Stapp car crash journal, 59, p.53.

Iwamoto, M. and Nakahira, Y., 2014. A preliminary study to investigate muscular effects for pedestrian kinematics and injuries using active thums. In Proceedings of the IRCOBI conference, IRC-14–53, Berlin, Germany (pp. 444-60).

Iwamoto, M., Nakahira, Y., Kimpara, H., Sugiyama, T. and Min, K., 2012. Development of a human body finite element model with multiple muscles and their controller for estimating

occupant motions and impact responses in frontal crash situations. Stapp car crash journal, 56, p.231.

Iwamoto, M., Nakahira, Y. and Sugiyama, T., 2011, June. Investigation of pre-impact bracing effects for injury outcome using an active human FE model with 3D geometry of muscles. In Proceedings of the 22nd international technical conference on the enhanced safety of vehicles (ESV), Washington, DC, USA.

Izawa, J., Rane, T., Donchin, O. and Shadmehr, R., 2008. Motor adaptation as a process of reoptimization. Journal of Neuroscience, 28(11), pp.2883-2891.

Jagodnik, K.M., Thomas, P.S., van den Bogert, A.J., Branicky, M.S. and Kirsch, R.F., 2016. Human-like rewards to train a reinforcement learning controller for planar arm movement. IEEE Transactions on Human-Machine Systems, 46(5), pp.723-733.

Jaśkowski, W., Lykkebø, O.R., Toklu, N.E., Trifterer, F., Buk, Z., Koutník, J. and Gomez, F., 2018. Reinforcement Learning to Run… Fast. In The NIPS'17 Competition: Building Intelligent Systems (pp. 155-167). Springer International Publishing.

Jin, X., Feng, Z., Mika, V., Li, H., Viano, D.C. and Yang, K.H., 2017. The role of neck muscle activities on the risk of mild traumatic brain injury in American football. Journal of biomechanical engineering, 139(10).

John, J., Klug, C., Kranjec, M., Svenning, E. and Iraeus, J., 2022. Hello, world! VIVA+: A human body model lineup to evaluate sex-differences in crash protection. Frontiers in bioengineering and biotechnology, 10.

Joos, E., Péan, F. and Goksel, O., 2020. Reinforcement learning of musculoskeletal control from functional simulations. In Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part III 23 (pp. 135-145). Springer International Publishing.

Kakade, S. and Langford, J., 2002, July. Approximately optimal approximate reinforcement learning. In Proceedings of the Nineteenth International Conference on Machine Learning (pp. 267-274).

Kamibayashi, L.K. and Richmond, F.J., 1998. Morphometry of human neck muscles. Spine, 23(12), pp.1314-1323.

Keshner, E.A., 2009. Vestibulocollic and cervicocollic control. Encyclopedia of neuroscience, pp.4220-4.

Keshner, E.A., Hain, T.C. and Chen, K.J., 1999. Predicting control mechanisms for human head stabilization by altering the passive mechanics. Journal of Vestibular Research, 9(6), pp.423-434.

Keshner, F.A. and Peterson, B.W., 1995. Mechanisms controlling human head stabilization. I. Head-neck dynamics during random rotations in the horizontal plane. Journal of neurophysiology, 73(6), pp.2293-2301.

Kidziński, Ł., Mohanty, S.P., Ong, C.F., Hicks, J.L., Carroll, S.F., Levine, S., Salathé, M. and Delp, S.L., 2018. Learning to run challenge: Synthesizing physiologically accurate motion using deep reinforcement learning. In The NIPS'17 Competition: Building Intelligent Systems (pp. 101-120). Springer International Publishing.

Kistemaker, D.A., Van Soest, A.K.J. and Bobbert, M.F., 2006. Is equilibrium point control feasible for fast goal-directed single-joint movements?. Journal of Neurophysiology, 95(5), pp.2898-2912.

Koelewijn, A.D., Heinrich, D. and van den Bogert, A.J., 2019. Metabolic cost calculations of gait using musculoskeletal energy models, a comparison study. PloS one, 14(9), p.e0222037.

Kolesnikov, S. and Khrulkov, V., 2020. Sample efficient ensemble learning with catalyst. rl. arXiv preprint arXiv:2003.14210.

Kumar, S., Narayan, Y. and Amell, T., 2000. Role of awareness in head-neck acceleration in low velocity rear-end impacts. Accident Analysis & Prevention, 32(2), pp.233-241.

Kuramochi, R., Kimura, T., Nakazawa, K., Akai, M., Torii, S. and Suzuki, S., 2004. Anticipatory modulation of neck muscle reflex responses induced by mechanical perturbations of the human forehead. Neuroscience letters, 366(2), pp.206-210.

La Barbera, V., Pardo, F., Tassa, Y., Daley, M., Richards, C., Kormushev, P. and Hutchinson, J., 2021. Ostrichrl: A musculoskeletal ostrich simulation to study bio-mechanical locomotion. arXiv preprint arXiv:2112.06061.

Lacquaniti, F., Bosco, G., Gravano, S., Indovina, I., La Scaleia, B., Maffei, V. and Zago, M., 2015. Gravity in the brain as a reference for space and time perception. Multisensory research, 28(5-6), pp.397-426.

Lantz, C.A., Klein, G., Chen, J., Mannion, A., Solinger, A.B. and Dvorak, J., 2003. A reassessment of normal cervical range of motion. Spine, 28(12), pp.1249-1257.

LeCun, Y., Bengio, Y. and Hinton, G., 2015. Deep learning. nature, 521(7553), pp.436-444.

LeCun, Y., Boser, B., Denker, J., Henderson, D., Howard, R., Hubbard, W. and Jackel, L., 1989. Handwritten digit recognition with a back-propagation network. Advances in neural information processing systems, 2.

Lee, S.G., Ashton-Miller, J.A. and Graziano, G.P., 1990. On muscle control strategies, agonism and antagonism. A 3-D study in human cervical muscles. In Winter Annual Meeting of the American Society of Mechanical Engineers.

Lemay, M.A. and Crago, P.E., 1996. A dynamic model for simulating movements of the elbow, forearm, and wrist. Journal of biomechanics, 29(10), pp.1319-1330.

Lieber, R.L., Jacobson, M.D., Fazeli, B.M., Abrams, R.A. and Botte, M.J., 1992. Architecture of selected muscles of the arm and forearm: anatomy and implications for tendon transfer. The Journal of hand surgery, 17(5), pp.787-798.

Lillicrap, T.P., Hunt, J.J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D. and Wierstra, D., 2015. Continuous control with deep reinforcement learning. arXiv preprint arXiv:1509.02971.

Liu, Y.K., Krieger, K.W., Njus, G., Ueno, K. and Connors, M.P., 1982. Cervical spine stiffness and geometry of the young human male. Tulane Univ New Orleans La School of Medicine.

Luo, S., Androwis, G., Adamovich, S., Nunez, E., Su, H. and Zhou, X., 2021. Robust Walking Control of a Lower Limb Rehabilitation Exoskeleton Coupled with a Musculoskeletal Model via Deep Reinforcement Learning.

Margaria, R., 1968. Positive and negative work performances and their efficiencies in human locomotion. Internationale Zeitschrift für angewandte Physiologie einschließlich Arbeitsphysiologie, 25(4), pp.339-351.

Margulies, S.S., Yuan, Q., Guccione Jr, S.J. and Weiss, M.S., 1998. Kinematic response of the neck to voluntary and involuntary flexion. Aviation, space, and environmental medicine, 69(9), pp.896-903.

Marsden, C.D., Obeso, J.A. and Rothwell, J.C., 1983. The function of the antagonist muscle during fast limb movements in man. The journal of physiology, 335(1), pp.1-13.

Martynenko, O., Kempter, F., Kleinbach, C., Schmitt, S. and Fehr, J., 2018. Integrated physiologically motivated controller for the open-source extended Hill-type muscle model in LS-DYNA. In Proceedings of the International IRCOBI Conference. IRCOBI Council, Athens, Greece (pp. 239-241).

Matsushita, T., Sato, T.B., Hirabayashi, K., Fujimura, S., Asazuma, T. and Takatori, T., 1994. X-ray study of the human neck motion due to head inertia loading. SAE transactions, pp.1623-1632.

Mayoux-Benhamou, M.A., Revel, M. and Vallee, C., 1997. Selective electromyography of dorsal neck muscles in humans. Experimental Brain Research, 113, pp.353-360.

McCrea, R.A., Gdowski, G.T., Boyle, R. and Belton, T., 1999. Firing behavior of vestibular neurons during active and passive head movements: vestibulo-spinal and other non-eye-movement related neurons. Journal of neurophysiology, 82(1), pp.416-428.

Meijer, R., Rodarius, C., Adamec, J., van Nunen, E. and Van Rooij, L., 2008. A first step in computer modelling of the active human response in a far-side impact. International journal of crashworthiness, 13(6), pp.643-652.

Micera, S., Sabatini, A.M. and Dario, P., 1999. Adaptive fuzzy control of electrically stimulated muscles for arm movements. Medical & biological engineering & computing, 37, pp.680-685.

Millard, M., Uchida, T., Seth, A. and Delp, S.L., 2013. Flexing computational muscle: modeling and simulation of musculotendon dynamics. Journal of biomechanical engineering, 135(2).

Min, K., Iwamoto, M., Kakei, S. and Kimpara, H., 2018. Muscle Synergy–Driven Robust Motion Control. Neural Computation, 30(4), pp.1104-1131.

Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D. and Riedmiller, M., 2013. Playing atari with deep reinforcement learning. arXiv preprint arXiv:1312.5602.

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G. and Petersen, S., 2015. Human-level control through deep reinforcement learning. nature, 518(7540), pp.529-533.

Moore, K.L., 2009. TRG Ph. D, AF Dalley, AMR Agur, et PWT Ph. D, Clinically Oriented Anatomy.

Morimoto, K., Sakamoto, M., Fukuhara, T. and Kato, K., 2013. Electromyographic study of neck muscle activity according to head position in rugby tackles. Journal of physical therapy science, 25(5), pp.563-566.

Mörl, F., Siebert, T., Schmitt, S., Blickhan, R. and Günther, M., 2012. Electro-mechanical delay in Hill-type muscle models. Journal of Mechanics in Medicine and Biology, 12(05), p.1250085.

Morningstar, M.W., Pettibon, B.R., Schlappi, H., Schlappi, M. and Ireland, T.V., 2005. Reflex control of the spine and posture: a review of the literature from a chiropractic perspective. Chiropractic & Osteopathy, 13, pp.1-17.

Mortensen, J., Trkov, M. and Merryweather, A., 2018. Exploring novel objective functions for simulating muscle coactivation in the neck. Journal of Biomechanics, 71, pp.127-134.

Mortensen, J.D., Vasavada, A.N. and Merryweather, A.S., 2018. The inclusion of hyoid muscles improve moment generating capacity and dynamic simulations in musculoskeletal models of the head and neck. PloS one, 13(6), p.e0199912.

Nakano, E., Imamizu, H., Osu, R., Uno, Y., Gomi, H., Yoshioka, T. and Kawato, M., 1999. Quantitative examinations of internal representations for arm trajectory planning: minimum commanded torque change model. Journal of Neurophysiology, 81(5), pp.2140-2155.

Nelson, W.L., 1983. Physical principles for economies of skilled movements. Biological cybernetics, 46, pp.135-147.

Nemirovsky, N. and Van Rooij, L., 2010, September. A new methodology for biofidelic head-neck postural control. In Proceedings of the IRCOBI Conference.

Newell, R.S., Blouin, J.S., Street, J., Cripton, P.A. and Siegmund, G.P., 2018. The neutral posture of the cervical spine is not unique in human subjects. Journal of biomechanics, 80, pp.53-62.

Nie, B., Giudice, J.S., Poulard, D., Wu, T. and Panzer, M.B., 2018. Evaluation and injury investigation of a finite element foot and ankle model for small female occupants. International journal of crashworthiness.

Nightingale, R.W., Chancey, V.C., Ottaviano, D., Luck, J.F., Tran, L., Prange, M. and Myers, B.S., 2007. Flexion and extension structural properties and strengths for male cervical spine segments. Journal of biomechanics, 40(3), pp.535-542.

Oh, K., Rymer, W.Z. and Choi, J., 2021. The speed of adaptation is dependent on the load type during target reaching by intact human subjects. Experimental Brain Research, 239(10), pp.3091-3104.

Olafsdottir, J.M., Östh, J., Davidsson, J. and Brolin, K., 2013. Passenger kinematics and muscle responses in autonomous braking events with standard and reversible pre-tensioned restraints. In Ircobi conference 2013 (No. IRC-13-70, pp. 602-617).

Ólafsdóttir, J.M., Brolin, K., Blouin, J.S. and Siegmund, G.P., 2015. Dynamic spatial tuning of cervical muscle reflexes to multidirectional seated perturbations. Spine, 40(4), pp.E211-E219.

Ólafsdóttir, J.M., Östh, J. and Brolin, K., 2019. Modelling reflex recruitment of neck muscles in a finite element human body model for simulating omnidirectional head kinematics. In IRCOBI Conf. Proc (pp. 308-323).

Ono, K., Kaneoka, K., Wittek, A. and Kajzer, J., 1997. Cervical injury mechanism based on the analysis of human cervical vertebral motion and head-neck-torso kinematics during low-speed rear impacts. SAE transactions, pp.3859-3876.

Östh, J., Brolin, K. and Bråse, D., 2015. A human body model with active muscles for simulation of pretensioned restraints in autonomous braking interventions. Traffic injury prevention, 16(3), pp.304-313.

Östh, J., Brolin, K., Carlsson, S., Wismans, J. and Davidsson, J., 2012. The occupant response to autonomous braking: a modeling approach that accounts for active musculature. Traffic injury prevention, 13(3), pp.265-277.

Östh, J., Brolin, K. and Happee, R., 2012. Active muscle response using feedback control of a finite element human arm model. Computer methods in biomechanics and biomedical engineering, 15(4), pp.347-361.

Östh, J., Mendoza-Vazquez, M., Sato, F., Svensson, M.Y., Linder, A. and Brolin, K., 2017. A female head–neck model for rear impact simulations. Journal of biomechanics, 51, pp.49-56.

Panjabi, M.M., Crisco, J.J., Vasavada, A., Oda, T., Cholewicki, J., Nibu, K. and Shin, E., 2001. Mechanical properties of the human cervical spine as shown by three-dimensional load–displacement curves. Spine, 26(24), pp.2692-2700.

Panzer, M., 2006. Numerical modelling of the human cervical spine in frontal impact (Master's thesis, University of Waterloo).

241

Panzer, M.B. and Cronin, D.S., 2009. C4–C5 segment finite element model development, validation, and load-sharing investigation. Journal of biomechanics, 42(4), pp.480-490.

Panzer, M.B., Fice, J.B. and Cronin, D.S., 2011. Cervical spine response in frontal crash. Medical engineering & physics, 33(9), pp.1147-1159.

Pedotti, A., Krishnan, V.V. and Stark, L., 1978. Optimization of muscle-force sequencing in human locomotion. Mathematical Biosciences, 38(1-2), pp.57-76.

Peng, G.C.Y., Hain, T.C. and Peterson, B.W., 1996. A dynamical model for reflex activated head movements in the horizontal plane. Biological cybernetics, 75(4), pp.309-319.

Peng, G.C., Hain, T.C. and Peterson, B.W., 1999. Predicting vestibular, proprioceptive, and biomechanical control strategies in normal and pathological head movements. IEEE transactions on biomedical engineering, 46(11), pp.1269-1280.

Peters, J., Vijayakumar, S. and Schaal, S., 2005. Natural actor-critic. In Machine Learning: ECML 2005: 16th European Conference on Machine Learning, Porto, Portugal, October 3-7, 2005. Proceedings 16 (pp. 280-291). Springer Berlin Heidelberg.

Peterson, B.W., Keshner, E.A. and Banovetz, J., 1989. Comparison of neck muscle activation patterns during head stabilization and voluntary movements. Progress in brain research, 80, pp.363-371.

Popescu, F., Hidler, J.M. and Rymer, W.Z., 2003. Elbow impedance during goal-directed movements. Experimental brain research, 152, pp.17-28.

242

Putra, I.P.A., Iraeus, J., Sato, F., Svensson, M.Y., Linder, A. and Thomson, R., 2021. Optimization of female head–neck model with active reflexive cervical muscles in low severity rear impact collisions. Annals of biomedical engineering, 49, pp.115-128.

Putra, I.P.A., Iraeus, J., Thomson, R., Svensson, M.Y., Linder, A. and Sato, F., 2019. Comparison of control strategies for the cervical muscles of an average female head-neck finite element model. Traffic injury prevention, 20(sup2), pp.S116-S122.

Putra, I.P.A. and Thomson, R., 2022. Analysis of control strategies for VIVA OpenHBM with active reflexive neck muscles. Biomechanics and Modeling in Mechanobiology, 21(6), pp.1731-1742.

Rack, P.M., 2011. Limitations of somatosensory feedback in control of posture and movement. Comprehensive Physiology, pp.229-256.

Richmond, F.J.R. and Abrahams, V.C., 1979. What are the proprioceptors of the neck?. Progress in brain research, 50, pp.245-254.

Reynier, K.A., Alshareef, A., Sanchez, E.J., Shedd, D.F., Walton, S.R., Erdman, N.K., Newman, B.T., Giudice, J.S., Higgins, M.J., Funk, J.R. and Broshek, D.K., 2020. The effect of muscle activation on head kinematics during non-injurious head impacts in human subjects. Annals of Biomedical Engineering, 48, pp.2751-2762.

Schneider, L.W., Foust, D.R., Bowman, B.M., Snyder, R.G., Chaffin, D.B., Abdelnour, T.A. and Baum, J.K., 1975. Biomechanical properties of the human neck in lateral flexion. SAE Transactions, pp.3212-3224.

Schulman, J., Wolski, F., Dhariwal, P., Radford, A. and Klimov, O., 2017. Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347.

Scovil, C.Y. and Ronsky, J.L., 2006. Sensitivity of a Hill-based muscle model to perturbations in model parameters. Journal of biomechanics, 39(11), pp.2055-2063.

Shaw, G., Parent, D., Purtsezov, S., Lessley, D., Crandall, J., Kent, R., Guillemot, H., Ridella, S.A., Takhounts, E. and Martin, P., 2009. Impact response of restrained PMHS in frontal sled tests: skeletal deformation patterns under seat belt loading. Stapp Car Crash Journal, 53, p.1.

Shea, M., Edwards, W.T., White, A.A. and Hayes, W.C., 1991. Variations of stiffness and strength along the human cervical spine. Journal of biomechanics, 24(2), pp.95-107.

Siegmund, G.P., Sanderson, D.J. and Inglis, J.T., 2004. Gradation of neck muscle responses and head/neck kinematics to acceleration and speed change in rear-end collisions (No. 2004-22-0018). SAE Technical Paper.

Siegmund, G.P., Blouin, J.S., Brault, J.R., Hedenstierna, S. and Inglis, J.T., 2007. Electromyography of superficial and deep neck muscles during isometric, voluntary, and reflex contractions.

Siegmund, G.P., Inglis, J.T. and Sanderson, D.J., 2001. Startle response of human neck muscles sculpted by readiness to perform ballistic head movements. The Journal of physiology, 535(1), pp.289-300.

Siegmund, G.P., Sanderson, D.J., Myers, B.S. and Inglis, J.T., 2003. Awareness affects the response of human subjects exposed to a single whiplash-like perturbation. Spine, 28(7), pp.671-679.

Siegmund, G.P., Sanderson, D.J., Myers, B.S. and Inglis, J.T., 2003. Rapid neck muscle adaptation alters the head kinematics of aware and unaware subjects undergoing multiple whiplash-like perturbations. Journal of biomechanics, 36(4), pp.473-482.

Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M. and Dieleman, S., 2016. Mastering the game of Go with deep neural networks and tree search. nature, 529(7587), pp.484-489.

Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D. and Riedmiller, M., 2014. Deterministic policy gradient algorithms. In International conference on machine learning (pp. 387-395). Pmlr.

Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A. and Chen, Y., 2017. Mastering the game of go without human knowledge. nature, 550(7676), pp.354-359.

Silvestros, P., Pizzolato, C., Lloyd, D.G., Preatoni, E., Gill, H.S. and Cazzola, D., 2022. Electromyography-assisted neuromusculoskeletal models can estimate physiological muscle activations and joint moments across the neck before impacts. Journal of Biomechanical Engineering, 144(3).

Smeets, J.B.J., Erkelens, C.J. and van der Gon Denier, J.J., 1990. Adjustments of fast goal-directed movements in response to an unexpected inertial load. Experimental Brain Research, 81, pp.303-312.

Song, S., Kidziński, Ł., Peng, X.B., Ong, C., Hicks, J., Levine, S., Atkeson, C.G. and Delp, S.L., 2021. Deep reinforcement learning for modeling human locomotion control in neuromechanical simulation. Journal of neuroengineering and rehabilitation, 18, pp.1-17.

Stemper, B.D., Yoganandan, N. and Pintar, F.A., 2004. Validation of a head-neck computer model for whiplash simulation. Medical and Biological Engineering and Computing, 42, pp.333-338.

Sutton, R.S. and Barto, A.G., 2018. Reinforcement learning: An introduction. MIT press.

Thomas, P.S., Branicky, M., van den Bogert, A. and Jagodnik, K., 2008. Creating a reinforcement learning controller for functional electrical stimulation of a human arm. In The... yale workshop on adaptive and learning systems (Vol. 49326, p. 1). NIH Public Access.

Thuresson, M., Äng, B., Linder, J. and Harms-Ringdahl, K., 2005. Mechanical load and EMG activity in the neck induced by different head-worn equipment and neck postures. International journal of industrial ergonomics, 35(1), pp.13-18.

Tien, C.S. and Huston, R.L., 1987. Numerical advances in gross-motion simulations of head/neck dynamics. ASME. J Biomech Eng. 109(2): 163–168.

Tu, X., Li, M., Liu, M., Si, J. and Huang, H.H., 2021, May. A data-driven reinforcement learning solution framework for optimal and adaptive personalization of a hip exoskeleton. In 2021

IEEE International Conference on Robotics and Automation (ICRA) (pp. 10610-10616). IEEE.

Umberger, B.R., Gerritsen, K.G. and Martin, P.E., 2003. A model of human muscle energy expenditure. Computer methods in biomechanics and biomedical engineering, 6(2), pp.99-111.

Uno, Y., Kawato, M. and Suzuki, R., 1989. Formation and control of optimal trajectory in human multijoint arm movement. Biological cybernetics, 61(2), pp.89-101.

Valls-Solé, J., Rothwell, J.C., Goulart, F., Cossu, G. and Munoz, E., 1999. Patterned ballistic movements triggered by a startle in healthy humans. The Journal of physiology, 516(3), pp.931-938.

Van den Kroonenberg, A., Philippens, M.M.G.M., Cappon, H., Wismans, J.S.H.M., Hell, W. and Langwieder, K., 1998. Human head-neck response during low-speed rear end impacts. SAE transactions, pp.2836-2851.

van der Horst, M.J., Thunnissen, J.G.M., Happee, R., Van Haaster, R.M.H.P. and Wismans, J.S.H.M., 1997. The influence of muscle activity on head-neck response during impact. SAE transactions, pp.4003-4023.

Van Drunen, P., Maaswinkel, E., Van Der Helm, F.C.T., Van Dieën, J.H. and Happee, R., 2013. Identifying intrinsic and reflexive contributions to low-back stabilization. Journal of biomechanics, 46(8), pp.1440-1446.

van Lopik, D.W. and Acar, M., 2004. A computational model of the human head and neck system for the analysis of whiplash motion. International Journal of Crashworthiness, 9(5), pp.465-473.

Van Mameren, H., Sanches, H., Beursgens, J. and Drukker, J., 1992. Cervical spine motion in the sagittal plane II: position of segmental averaged instantaneous centers of rotation-a cineradiographic study. Spine, 17(5), pp.467-474.

van Rooij, L., 2011. Effect of various pre-crash braking strategies on simulated human kinematic response with varying levels of driver attention. In Proceedings of the 22nd ESV Conference.

van Soest, A.J. and Bobbert, M.F., 1993. The contribution of muscle properties in the control of explosive movements. Biological cybernetics, 69(3), pp.195-204.

Vasavada, A.N., Li, S. and Delp, S.L., 1998. Influence of muscle morphometry and moment arms on the moment-generating capacity of human neck muscles. Spine, 23(4), pp.412-422.

Vasavada, A.N., Peterson, B.W. and Delp, S.L., 2002. Three-dimensional spatial tuning of neck muscle activation in humans. Experimental brain research, 147, pp.437-448.

Vavalle, N.A., Schoell, S.L., Weaver, A.A., Stitzel, J.D. and Gayzik, F.S., 2014. Application of radial basis function methods in the development of a 95th percentile male seated fea model. Stapp car crash journal, 58, p.361.

Wadman, W.J., Denier Van der Gon, J.J., Geuze, R.H. and Mol, C.R., 1979. Control of fast goal-directed arm movements. Journal of Human Movement Studies, 5(1), pp.3-17.

Walker Jr, L.B., Harris, E.H. and Pontius, U.R., 1973. Mass, volume, center of mass and mass moment of inertia of head and head and neck of the human body. Tulane Univ New Orleans La.

Walter, J.R., Günther, M., Haeufle, D.F. and Schmitt, S., 2021. A geometry-and muscle-based control architecture for synthesising biological movement. Biological cybernetics, 115, pp.7-37.

Ward, S.R. and Lieber, R.L., 2005. Density and hydration of fresh and fixed human skeletal muscle. Journal of biomechanics, 38(11), pp.2317-2320.

Werbos, P., 1974. Beyond regression: New tools for prediction and analysis in the behavioral sciences. PhD thesis, Committee on Applied Mathematics, Harvard University, Cambridge, MA.

White, N.A., Begeman, P.C., Hardy, W.N., Yang, K.H., Ono, K., Sato, F., Kamiji, K., Yasuki, T. and Bey, M.J., 2009. Investigation of upper body and cervical spine kinematics of postmortem human subjects (PMHS) during low-speed, rear-end impacts (No. 2009-01-0387). SAE Technical Paper.

Wiegner, A.W. and Watts, R.L., 1986. Elastic properties of muscles measured at the elbow in man: I. Normal controls. Journal of Neurology, Neurosurgery & Psychiatry, 49(10), pp.1171-1176.

Williams, J.L. and Belytschko, T.B., 1983. A three-dimensional model of the human cervical spine for impact simulation. J. Biomech. Eng. 105, 321–331.

Winters, J.M., 1995. An improved muscle-reflex actuator for use in large-scale neuromusculoskeletal models. Annals of biomedical engineering, 23, pp.359-374.

Wochner, I., Endler, C.A., Schmitt, S. and Martynenko, O.V., 2019. Comparison of controller strategies for active human body models with different muscle materials. In IRCOBI Conference Proceedings (pp. 133-135).

Wolpert, D.M., 1997. Computational approaches to motor control. Trends in cognitive sciences, 1(6), pp.209-216.

Yoganandan, N., Pintar, F.A., Stemper, B.D., Wolfla, C.E., Shender, B.S. and Paskoff, G., 2007. Level-dependent coronal and axial moment-rotation corridors of degeneration-free cervical spines in lateral flexion. JBJS, 89(5), pp.1066-1074.

Yu, W., Turk, G. and Liu, C.K., 2018. Learning symmetric and low-energy locomotion. ACM Transactions on Graphics (TOG), 37(4), pp.1-12.

Zajac, F.E., 1989. Muscle and tendon: properties, models, scaling, and application to biomechanics and motor control. Critical reviews in biomedical engineering, 17(4), pp.359-411.

Zangemeister, W.H., Jones, A. and Stark, L., 1981. Dynamics of head movement trajectories: main sequence relationship. Experimental neurology, 71(1), pp.76-91.

Zhang, K., Yang, Z. and Başar, T., 2021. Multi-agent reinforcement learning: A selective overview of theories and algorithms. Handbook of reinforcement learning and control, pp.321-384.

Zheng, Z., Mo, F., Liu, T. and Li, X., 2021. A novel neuromuscular head-neck model and its application on impact analysis. IEEE transactions on neural systems and rehabilitation engineering, 29, pp.1394-1402.

Zhou, B., Zeng, H., Wang, F., Li, Y. and Tian, H., 2019. Efficient and robust reinforcement learning with uncertainty-based value expansion. arXiv preprint arXiv:1912.05328.

APPENDIX

APPENDIX A: RLMAC architecture with arm MB model

Deep deterministic policy gradient (DDPG) (Lillicrap et al., 2019) was used as the reinforcement learning (RL) agent for the initial feasibility study with the arm MB model (Chapter 2). The DDPG agent is composed of two neural networks (NNs) – an actor network and a critic network.

The actor network (Figure A-1a) maps the states to actions. The input of the actor network equals the number of state parameters defined for the control environment. In the arm motion control study, the states were defined with the kinematics properties – current elbow angles, elbow velocity, elbow angular error (current and last 2 timesteps), and network action (muscle activations) from the last timestep. In the Group Activated Muscle Recruitment (GAMR) scheme, there were seven input state parameters, whereas, in the Individual Activated Muscle Recruitment (IAMR) scheme, the input size was fourteen.

(a)   *(b)*

Figure A-1: (a) Actor and (b) Critic network for the DDPG agent coupled with arm MB model for elbow kinematics. The number of nodes at each layer is provided corresponding to the layer.

The critic network (Figure A-1b) inputs state and actions, and outputs the Q-value. In the RLMAC developed for the arm MB model, the state parameters in the critic network were activated by Relu function before passing on to the hidden layer, whereas the actions were added to the hidden layer without any activation. The parameters of the DDPG agents are provided in Table A-1.

Table A-1: Training parameters for the DDPG agent

Actor network learn rate	0.0001
Critic network learn rate	0.001
Timestep (sample time)	2.5 ms
Max steps in one iteration	240
Agent discount factor	0.99
Agent standard deviation	0.09-0.1

The objective of the RLMAC was to move the elbow from an initial position to a target position within the range of motion of the arm. The reward function was devised such that the reward is maximized at the target position. The reward function penalized the RLMAC proportional to the error (target angle – current angle). The reward function awarded the agent for each timestep the error was maintained within 0.1 rad, to encourage the agent to stabilize at the target position. Muscle activations were also penalized to ensure that the muscles do not co-contract at the target position of the elbow. The coefficients of the reward function are provided in Equation A-1. n in Equation A-1 equals 2 in case of GAMR and 9 for IAMR scenario to reduce the average activation across muscles.

$$Reward = -5\,Error - \frac{2}{n}\sum a(t) + 5\,(|Error| < 0.1\,rad) \qquad \text{Equation A-1}$$

APPENDIX B: Muscle model parameters in the neck model

The coefficient of non-linearity parameter (K_{PE}) in the parallel element is given by

$$K_{PE} = F_{PEE} \frac{F_{max}}{(l_{opt}(\Delta W_{Descending} + 1 - L_{PE,0}))^{\upsilon_{PE}}}$$

Equation B-1

$L_{PE,0}$ is the normalized rest length of parallel element respect to the optimal contractile element length. F_{PEE} is the force developed in the parallel element when the length is stretched beyond ΔW (descending).

The maximum damping coefficient D_{max} is calculated as

$$D_{max} = D_{SE} \frac{F_{max} A_{rel,0}}{l_{opt} B_{rel,0}}$$

Equation B-2

Table B-1: Muscle parameters used in the neck MB model (Haeufle et al., 2014)

ΔW (descending)	0.35
ΔW (Ascending)	0.35
k_{sh} (Ascending)	3.0
k_{sh} (Descending)	1.5
$A_{rel,0}$	0.25
$B_{rel,0}$	2.25
υ_{PE}	2.5
F_{PEE}	2
D_{SE}	0.3
R_{SDE}	0.01
$\Delta U_{SEE, nl}$	0.0425
$\Delta U_{SEE, 1}$	0.017

APPENDIX C: Parameters of the neck model RLMAC

A twin-delayed deep deterministic policy gradient (TD3) (Fujimoto et al., 2018) agent was used

for RLMAC in the neck control model. The TD3 agent requires two critic networks along with the

actor network. In this dissertation, the RLMAC was initially implemented for control of head

kinematics in the sagittal plane, which was eventually extended to control of kinematics in the

three anatomical planes.

The input of the actor networks consists of the head kinematics (translation and rotation), joint

displacement of current and previous two timesteps, and muscle activations of the previous time-

step. The input parameters were normalized to reduce the magnitude of the states to a similar scale

(approximately between -2 to 2) (LeCun et al., 2002). The normalizing factors of the state

parameters are provided in Table C-1.

Table C-1: Normalizing parameters

State	Normalizing factor
Displacement (x, y, z)	0.05
Velocity (x, y, z)	1
Rotation (x, y, z)	1
Rotational velocity (x, y, z)	5
Sine of errors (x, y, z)	1
Cosine of errors (x, y, z)	1
Joint translation (x, y, z)	[0.00027, 0.0015, 0.0003]
Joint rotations (x, y, z)	[0.0792, 0.087, 0.0648]
Joint forces (x, y, z)	[16.74, 175, 500]
Joint torques (x, y, z)	[0.4, 1.14, 0.6]

The sagittal control model, the kinematics was limited to the x-z plane. The RLMAC of the control

model had head kinematics corresponding to the x and z translation and y rotation state parameters.

The total state parameters were 96 in the sagittal neck model. The actions were considered

symmetrical about the mid-sagittal plane, limiting the output of the actor network to 23. The state

in the omnidirectional neck model included parameters for the 3 DoFs, increasing the actor input

size to 193. No assumptions regarding the grouping or symmetry of the muscle actuations were made, thus 46 actions were output from the actor network.
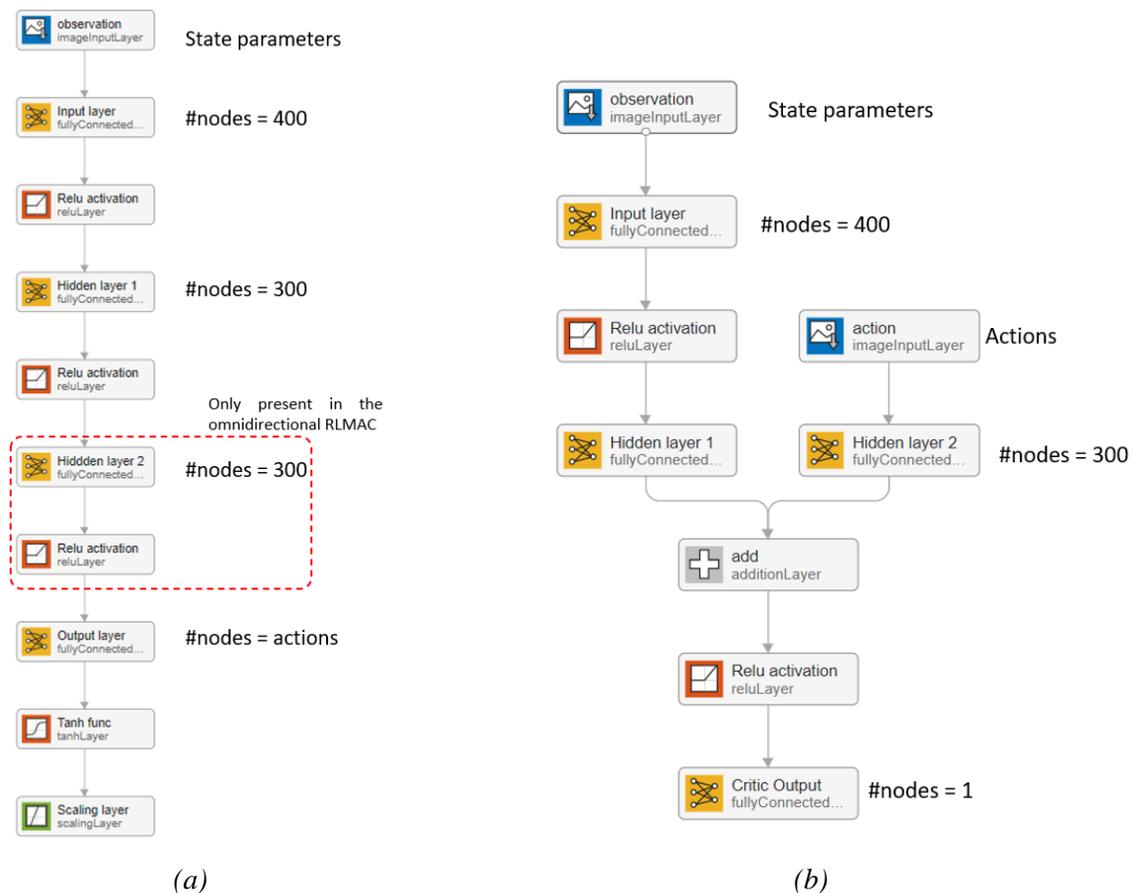
(a)   *(b)*

Figure C-1: (a) Actor and (b) Critic network for the TD3 agent coupled with neck MB model for generating head kinematics. The number of nodes at each layer is provided corresponding to the layer. An additional hidden layer was added in the actor network for the omnidirectional control model.

Figure C-1 shows the actor and critic networks of the RLMAC for the neck model. The outputs from the actor network were scaled using a combination of tanh and scaling functions to bound the muscle activations between 0 and 1. In the omnidirectional model, two hidden layers were used to process the state data before the final layer, while in the sagittal control model, only one hidden layer was present in the actor network. The critic network was the same for both the neck control models (Figure C-1b). The state parameters were processed and activated before the hidden layer whereas the actions were directly input to the hidden layer.

The objective of the RLMAC was to generate goal-directed motion of the head, i.e., move the head to a specific target orientation in the head kinematic space. The reward function was developed such that the reward is maximized if the head reaches a target orientation (Equation C-1).

$$(1 - Sin\ (2\varepsilon_y))^2 - 0.2\ (|V_x| + |V_z|) - 0.1\ |\omega_y| - \frac{0.05}{21} \textstyle\sum Joint\ energy \qquad \text{Equation C-1}$$

The translation and rotational velocity in Equation C-1 were normalized with factors in Table C-1. The joint energy was the product of normalized joint forces and displacements in the x-z plane and joint torque and rotation along the y axes. Translation and rotational velocities were penalized to damp the head at the target position. The joint energy was included in the reward to stabilize the spine under gravity.

$$\frac{1}{3} [(1 - Sin\ (2\varepsilon_x))^2 + (1 - Sin\ (2\varepsilon_y))^2 + (1 - Sin\ (2\varepsilon_z))^2] - 0.2\ (|V_x| + |V_y| + |V_z|) \qquad \text{Equation C-2}$$
$$- 0.1\ (|\omega_x| + |\omega_y| + |\omega_z|) - \frac{0.05}{42} \textstyle\sum Joint\ energy - 0.5\ symmetry$$

The reward function was expanded to include parameters in the three anatomical planes (Equation C-2). The force, displacement, and velocities used for the reward function were normalized. The symmetry factor in the reward function was defined as Equation C-3.

$$|Neural\ stimulation\ left - Neural\ stimulation\ right|\ when\ the\ training\ objective \qquad \text{Equation C-3}$$
$$was\ stability\ or\ target\ was\ prescribed\ in\ the\ sagittal\ plane$$

The symmetry factor was included in the reward function to train the RLMAC to develop a symmetric muscle activation strategy. Without the symmetry factor, the RLMAC would learn to generate the desired kinematics, however, when the RLMAV was prescribed to maintain the head stability or generate kinematics in the sagittal plane, the activation patterns of left and right side were completely different.

The training parameters of the TD3 agent for both the control model are given in Table C-2

Table C-2: Training and simulation scenarios

Actor network learn rate	0.001
Critic network learn rate	0.001
Timestep (sample time)	1.0 ms
Max steps in one iteration	500
Agent discount factor	0.99
Agent standard deviation	0.09

Both the control models, the sagittal control as well as the omnidirectional control model could generate the desired kinematics. Figure C-2 shows the weight distribution in the input layer of the trained actor network for the state parameters in the RLMAC for the sagittal and the omnidirectional control model.
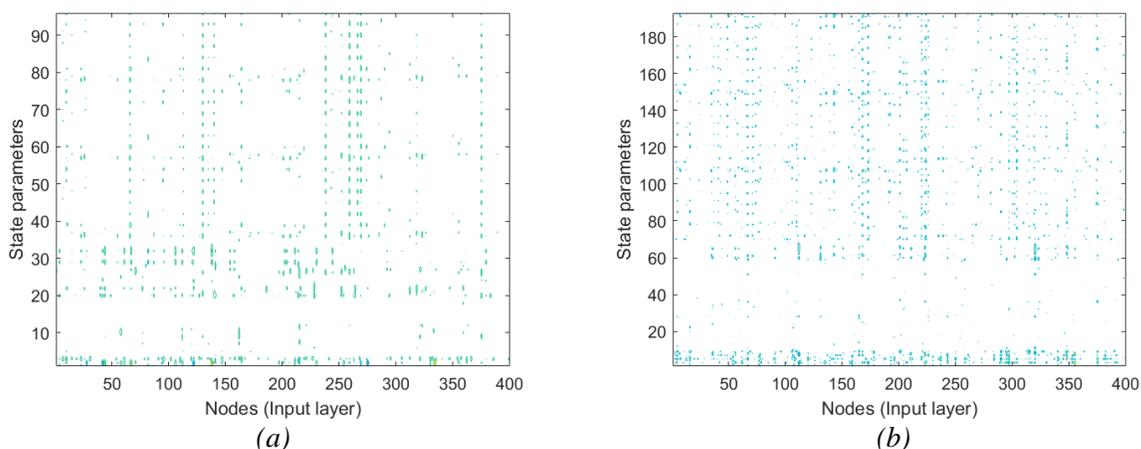

Figure C-2: Weight distribution in the first layer of the actor network (a) Sagittal control model (b) Omnidirectional control model.

The weight distribution plots show that the head rotational kinematics (translation and rotation) along with the sine and cosine of the angular errors are denser, which suggests that the head rotational kinematics are the most important factors in the RLMAC states. The rotational kinematics and error are followed by muscle actuations (action from the actor network in the previous timestep), which have low weights. It can be assumed that compared to the other state parameters, the muscle activations from previous timestep have less impact on the response. The translation kinematics and the joint displacements follow the activations, were found to be

important for the RLMAC response as can be seen by the dense representation in the plots (Figure C-2).

The trained actor networks of both the control models show a similar trend when the weights of the state parameters are plotted for each input node, even though the relative importance of the parameters in the model may differ. The muscle activation feedback to the actor network was found to have low weights. However, the reward function for the training did not have any term penalizing or awarding the muscle forces or activations. A different reward function with the actuation term may increase the contribution of the activations to the overall response, therefore further consideration is required before removing the activation terms from the state.

Another observation from the weight distribution plot is that the sagittal actor network weight matrix is sparse (Figure C-2a). Therefore, the size of the input and the hidden layer may be decreased in the actor network (and the critic network) for faster training.

APPENDIX D: Reward function components for the neck control model

The reward function for the control of the neck model in the sagittal plane was developed so that the cumulative reward in a training iteration is maximized when the head moved to the target orientation and stabilized at that position. The reward function developed for the head motion control in the sagittal pane is provided in Equation D-1.

$$(1 - Sin\ (2\varepsilon_y))^2 - 0.2\ (|V_x| + |V_z|) - 0.1\ |\omega_y| - \frac{0.05}{21} \sum Joint\ energy \qquad \text{Equation D-1}$$

The first component of the reward function represents the sensory feedback that seeks to reduce the position error of the head. The sensory term is maximized when the error is close to 0. The head angular error is penalized to minimize any head disturbances at the target angle. Adjusting the coefficient of the angular velocity is important as high influence of the angular velocity on the reward will not let any motion of the head even if the training time is increased (Crowder et al., 2021). The coefficient of 0.1 was sufficient to damp any rotations at the final position.

The translation velocities were penalized to remove any linear noise or disturbances of the head when stable. Without penalizing the linear velocities, the head will move to the target orientation, however the head CoM will have some linear disturbances which will not be reported in the head angle or angular error plots. Some previous study have considered a feedback on head acceleration corresponding to the otolith organs (Happee et al., 2017; Zheng et al., 2021), however minimizing the head acceleration did not completely remove the disturbances of the head at the target orientation. A sensitivity study was performed with different coefficients to analyze the effects of the relative importance of the velocity parameter on the response of the neck model.

Figure D-1 compares the training of the RLMAC for the different values of the velocity penalty coefficient. In the plots it can be seen that for the coefficient of 0.2, the training converges faster, and thus was used for the training in this dissertation.
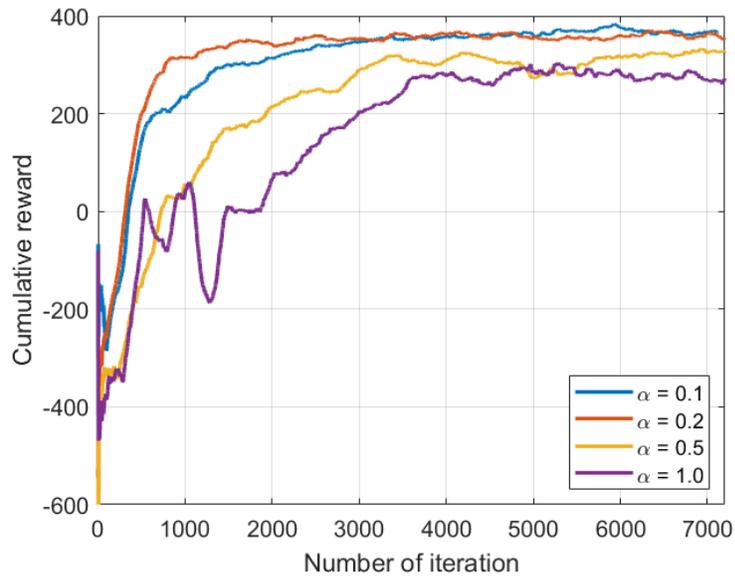
260

Figure D-1: Average reward plots for training of the sagittal neck model for different coefficients of velocity penalty in the reward function. The 'α' in the legend refers to the coefficient.

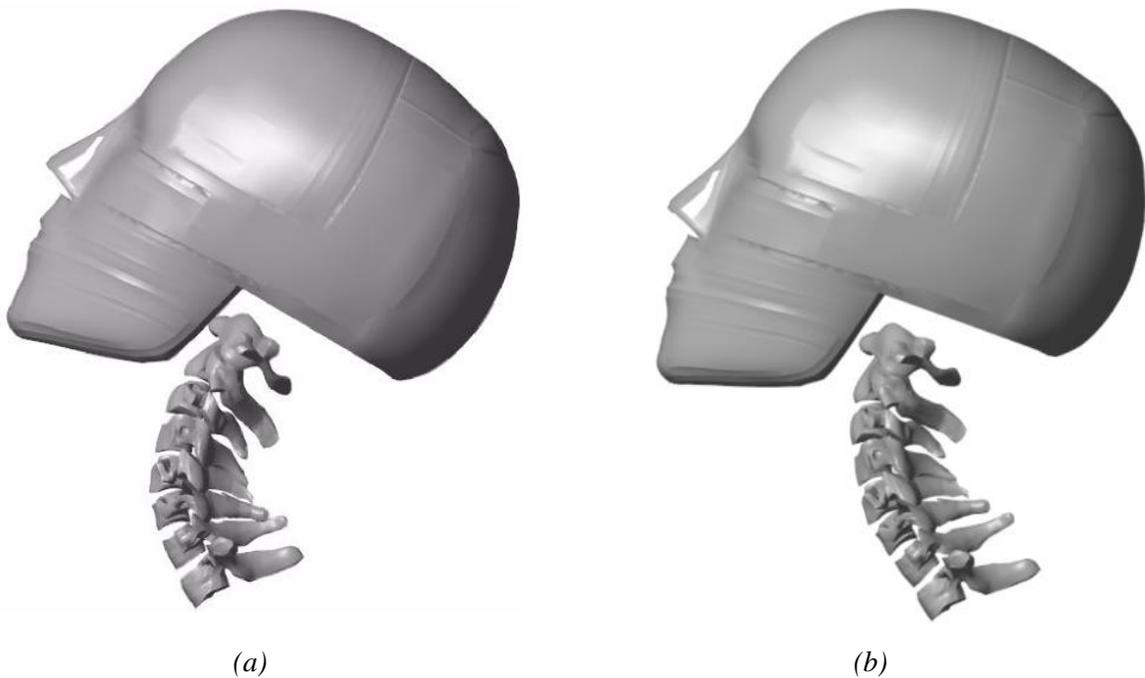<center>(a)   (b)</center>

Figure D-1: Goal directed head motion for 20^o extension with (a) Velocity coefficient = 0.2 (b) Velocity coefficient = 1.0.

261

Figure D-2 shows the comparison of the final orientation of the head for a prescribed target of 20° extension with the RLMAC trained with different coefficients of velocity penalty. The trained RLMAC with coefficient 1.0 for velocity penalty had the head CoM nearer to the neutral head position but compresses the spine more.

REFERENCES

Crowder, D.C., Abreu, J. and Kirsch, R.F., 2021. Hindsight experience replay improves reinforcement learning for control of a MIMO musculoskeletal model of the human arm. IEEE Transactions on Neural Systems and Rehabilitation Engineering, 29, pp.1016-1025.

Fujimoto, S., Hoof, H. and Meger, D., 2018, July. Addressing function approximation error in actor-critic methods. In International conference on machine learning (pp. 1587-1596). PMLR.

Happee, R., de Bruijn, E., Forbes, P.A. and van der Helm, F.C., 2017. Dynamic head-neck stabilization and modulation with perturbation bandwidth investigated using a multisegment neuromuscular model. Journal of biomechanics, 58, pp.203-211.

262

LeCun, Y., Bottou, L., Orr, G.B. and Müller, K.R., 2002. Efficient backprop. In Neural networks: Tricks of the trade (pp. 9-50). Berlin, Heidelberg: Springer Berlin Heidelberg.

Lillicrap, T.P., Hunt, J.J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D. and Wierstra, D., 2015. Continuous control with deep reinforcement learning. arXiv preprint arXiv:1509.02971.

Zheng, Z., Mo, F., Liu, T. and Li, X., 2021. A novel neuromuscular head-neck model and its application on impact analysis. IEEE transactions on neural systems and rehabilitation engineering, 29, pp.1394-1402.