

A Multi-Layer Sim-to-Real Framework for Gaze-Driven Assistive Neck Exoskeletons

Colin Rubow^{2,3,‡}, Eric Brewer^{1,‡}, Ian Bales², Haohan Zhang^{2,3}, and Daniel S. Brown^{1,3,*}

Abstract—Dropped head syndrome, caused by neck muscle weakness from neurological diseases, severely impairs an individual’s ability to support and move their head, causing pain and making everyday tasks challenging. Our long-term goal is to develop an assistive powered neck exoskeleton that restores natural movement. However, predicting a user’s intended head movement remains a key challenge. We leverage virtual reality (VR) to collect coupled eye and head movement data from healthy individuals to train models capable of predicting head movement based solely on eye gaze. We also propose a novel multi-layer controller selection framework, where head control strategies are evaluated across decreasing levels of abstraction—from simulation and VR to a physical neck exoskeleton. This pipeline effectively rejects poor-performing controllers early, identifying two novel gaze-driven models that achieve strong performance when deployed on the physical exoskeleton. Our results reveal that no single controller is universally preferred, highlighting the necessity for personalization in gaze-driven assistive control. Our work demonstrates the utility of VR-based evaluation for accelerating the development of intuitive, safe, and personalized assistive robots.

I. INTRODUCTION

Dropped Head Syndrome (DHS) is characterized by the inability of a person to move and raise their head due to neck weakness [1], [2]. DHS can result from multiple causes, such as central/peripheral neurological pathology (e.g., amyotrophic lateral sclerosis) and autoimmune conditions (e.g., polymyositis) [3]–[6]. DHS causes fatigue during head movement and, in severe cases, results in a “head-on-chest” posture [6]. This posture makes it challenging to breathe, swallow, and speak due to the bent airway. It also interferes with social interactions and walking [7], [8].

Static neck braces are currently used as a treatment for DHS. However, patients find these braces uncomfortable and ineffective [9], [10]. These braces fix the head in a static pose using support under the chin, making it difficult to open the mouth to speak or swallow. More importantly, these braces do not promote mobility necessary for daily tasks (e.g., feeding, horizontal vision, etc.) [11]. Other solutions include using a reclining wheelchair with a headrest or fixing the head to a chair using straps. These solutions, however, are not portable and similarly do not allow head-neck motions.

This work was partially supported by the National Institutes of Health (R21EB035378). Colin Rubow also received support from the National Institute of Occupational Safety and Health (5T42OH008414).

[‡]These authors contributed equally.

¹Kahlert School of Computing, University of Utah, UT 84112, USA.

²Dept. of Mechanical Engineering, University of Utah, UT 84112, USA.

³Robotics Center, University of Utah, UT 84112, USA.

*Corresponding Author (daniel.s.brown@utah.edu).

Wearable robots (e.g., powered neck exoskeletons) that can restore head-neck movement through external actuators represent a potentially transformative solution to treat DHS. Although there have been significant advances in developing the mechanical hardware of these devices [12]–[17], controlling these devices by a user remains a challenging problem. Using a hand-held device (e.g., joystick, keyboard) to control the neck exoskeleton is unintuitive and even infeasible for many patients, especially those whose DHS is a result of widespread neural degeneration (e.g., amyotrophic lateral sclerosis) and thus suffer from hand weakness [18], [19].

In contrast to hand-held devices, we propose and evaluate models that use a user’s gaze as the control input for wearable neck robots. This is biologically motivated and viable because the eyes and the head often coordinate during visual tasks through neural mechanisms such as the vestibulo-ocular reflex (VOR). In VOR, compensatory rotations of the eyes allow the eyes to remain fixated on a target while the head moves towards the target [20], [21]. Beyond VOR, it has also been found that the peak head velocity during gaze shifts is linearly proportional to the distance the head will travel [22], [23] and generally the head will travel directly towards the gaze target [20], [24]. Furthermore, it has been shown that head motions are minimized when the gaze lies in a central region in the field of view [22] for local visual exploration, which we term the deadzone. Due to these behaviors, using eye movements as control inputs for wearable neck robots to restore head movement is promising.

During early development of these controllers, the question arises as to how to evaluate them. Deploying each controller on a physical robot with a human subject is time-consuming and could result in safety hazards. We therefore employ a novel, multi-level controller selection technique that uses three layers of abstraction to eliminate poor-performing controllers early in the development process (Fig. 1). The first layer tests the ability of the models to predict head orientations during eye-head coordination tasks taken from the previous dataset. The second deploys the controllers in a safe VR environment to assist healthy human subjects in performing eye-head coordination tasks without actually applying forces to the head. The third deploys the controllers on a physical neck exoskeleton to assist human subjects to perform a real-world eye-head coordination task.

We employ two strategies for identifying gaze-driven control models. The first is to learn the parameters of an analytical, hand-crafted model based on eye-head coordination behavior described in the literature [23]. The second strategy is fully data-driven, where we train multiple neural




Fig. 1. We propose to leverage multiple levels of abstraction and emulation to assist in efficiently finding useful gaze-driven neck exoskeleton controllers. As the initial controller pool moves down through each layer, poor controllers are rejected from further study. Our goal is to arrive at a final set of controller designs suitable for restoring head-neck motions.

network models of varying architectures. These models take eye position and head orientation as inputs and predict the next desired head orientation. We train both styles of model using a large dataset of coupled eye and head movement data we collected using virtual reality (VR) [25].

One of our main contributions is a multi-layer controller evaluation framework (see Fig. 1) that enables safe, efficient rejection of poor exoskeleton controllers before physical deployment. With the goal of developing more natural eye-tracking controllers, we developed seven candidate controllers and evaluated them with our selection strategy. We were able to reject five of these controllers before physical robot deployment; and the other two controllers were deployed and approved by subjects in our user study.

II. CONTROLLER MODELS

Four categories of controllers were used in this work.

A. Quadrant (Baseline)

This controller divides the range of view into four directional regions (up, down, left, right) and a central deadzone region of 10° diameter [26]. If the gaze lies in one of the directional regions, the controller moves the head at a constant speed in that direction at $20^\circ/\text{sec}$. If in the deadzone, the controller does not move. The controller only performs motions in purely horizontal or vertical directions and thus performs diagonal trajectories in a zig-zag fashion. While this controller employs VOR [20], the speed and deadzone dimensions were not biologically chosen, and the zig-zag motions can be unnatural. However, to the best of our knowledge, this is the only eye controller developed for neck exoskeletons currently, so we use this controller as a baseline. In the following sections, we discuss our novel learned gaze-driven controllers.

B. Vector Parameterized

From literature [20]–[24], we identified three key elements of eye-head coordination: VOR, peak head velocity being proportional to head displacement, and gaze deadzone. From these elements, we developed the following control law:

$$h(x) = \frac{v(x - c)}{1 + e^{a(-(x-c)+b)}}$$




Fig. 2. Two examples of the Vector Parameterized control law for either pitch or yaw directions. The v and c values are equal for both examples. The control law transitions from a deadzone with zero velocity to a line with proportional velocity. Gray dotted lines mark inflection point locations of the plots. Blue dashed lines show what the control law transitions to and from. The Purple curve on the right has a larger but stricter deadzone.

This function takes a gaze angle x in a single dimension, i.e., pitch or yaw, and produces an angular rotation of the head in the same dimension. As x increases, h asymptotically approaches the linear function $f(x) = v(x - c)$ (two examples are shown in Fig. 2). This function induces a “soft dead zone” such that $h(x) \approx 0$ when $x \approx 0$ (i.e., the head remains still when the user is looking forward). The parameter v represents the linear asymptote slope, $1/a$ the exponential time constant and thus the “sharpness” of the inflection between the “dead zone” and asymptote, b the location in terms of x of the inflection point, and c the shift of the function from the origin.

C. Multi Layer Perceptron Model (MLP)

Our Vector Parameterized model is based on key insights from eye-head coordination; however, it follows a specific functional form and its parameters are fit to multiple subjects’ data, so the model could be sub-optimal. We thus train some simple feedforward MLP models on eye-head coordination data. We posited that a data-driven neural network would be able to induce eye tracking controllers that perform better than the analytical model we developed. As input, the model takes the 3D normalized direction of the left eye relative to the orientation of the head, that of the right eye, and that of the head itself relative to the environmental space. The MLP outputs the relative change of the head direction in the next time step in terms of angular pitch and yaw.

D. Long Short Term Memory Model (LSTM)

Due to the sequential nature of gaze and head vector data as a series over time, we also chose to evaluate the efficacy of using an LSTM [27], a popular recurrent neural network architecture, as a gaze-driven controller. An LSTM network has an inherent advantage in that it can theoretically capture inferential context over a large number of time steps. This means that the LSTM has the capability to model, for example, if a certain sequence of directional gaze vectors some number of time units in the past is currently relevant to predicting the motion of the head. In essence, the LSTM

is capable of modeling eye-head coordination in the case that this phenomenon is not totally a Markov process, i.e., that its prediction depends only on the most recent time step. It takes as input the current gaze and head directions as 3D vectors. The LSTM retains a latent hidden state vector, which enables the capturing of context over time as described above.

E. Training Details

The Parameterized Vector, MLP, and LSTM models were trained using the Adam optimizer [28]. The parameters were optimized by minimizing the mean-squared error between those models’ predicted angular head movements and the actual angular head movements in the Eye Tracking VR Dataset [25]. This dataset contains eye and head directional data for 25 participants (over 120 thousand data samples); each participant performed twelve 90-second trials (three 90-second trials for each of four distinct tasks) at 90 frames per second. We used data from 18 of the 25 participants for model training. After training these models, we evaluated them against a held-out test set consisting of data from the remaining 7 participants.

In practice, we had to downsample the input frequency for each LSTM variant by a factor of 2. This was because the ensuing VR experiment was only able to perform inference on LSTM-variant models at a rate of about 45 Hz, although the game itself runs at about 90 Hz. Thus, our LSTM models were trained using trajectories of length 4050 (90 s \times 45 Hz). Other models did not experience any unintentional slowdown in terms of latency and were trained using trajectories containing 8100 time steps.

III. AUTOREGRESSIVE EXPERIMENT

As a result of training, our models were optimized sufficiently in terms of MSE loss, as shown in Table I. Nonetheless, we conducted an additional experiment in an attempt to simulate whether these controllers would behave well autoregressively when deployed with real human subjects.

A. Methods – Autoregressive

We assume that a participant’s gaze would approximately follow the same path throughout a given trial, *regardless of* the underlying model controlling the rotation of the head. For example, in a Pursuit task, we assume that a participant would roughly look at the same sequence of points in 3D space (corresponding to the position of the object being tracked), irrespective of the actual rotation of their head.

Prior to the experiment, we calculated the sequence of **focal points** of each participant’s binocular gaze in a given trial. During each time step within our dataset, we recorded the *position* and *orientation* in 3D space of the head and the *position* and *direction* of the eyes. Given these additional data, we then defined the gaze rays (each specified by an origin and direction) for both eyes. We can then calculate the focal point as the point in space closest to the gaze rays of both eyes. This is done by first taking the cross product of the directional vectors of both eyes, $\vec{g}_f = \vec{g}_l \times \vec{g}_r$, where \vec{g}_l and \vec{g}_r are the normalized direction vectors of the left and right

eyes. This yields \vec{g}_f , the direction of the only orthogonal, hence shortest, segment between the rays. It is then possible to solve via Gaussian elimination the following:

$$\begin{aligned} \vec{p}_l + c_l * \vec{g}_l + c_f * \vec{g}_f &= \vec{p}_r + c_r * \vec{g}_r \\ c_l * \vec{g}_l - c_r * \vec{g}_r + c_f * \vec{g}_f &= \vec{p}_r - \vec{p}_l \end{aligned}$$

where \vec{p} is the origin of a particular ray, c is the magnitude of that ray, and \vec{g} is the direction. Solving the above system produces the magnitudes c for the left-eye ray, right-eye ray, and the focal segment. We defined the focal point as the midpoint of that focal segment.

For most instances in our dataset, the rays representing the gaze of the left and right eyes nearly intersected a few meters in front of the participant, as expected. However, in the case when the user was blinking, we linearly interpolate the gaze direction between the two successful recordings surrounding the blinking interval, then normalize.

Given a participant’s sequence of estimated focal points throughout a trial, the experiment was conducted as follows. First, given a single model from Section II, we maintained the location and orientation of a virtual “head” within a 3D coordinate space. We also kept track of the location of the pupils relative to the head. We then applied our assumption—that the eyes would approximately look at the calculated focal point for this time step. Instead of setting the direction of both eyes to the ground truth directions, we aimed each eye toward the focal point, subject to a small amount of Gaussian noise. Next, we fed the selected model with the novel assumed gaze. This produces both a pitch and yaw value, which we used to rotate the virtual environment. We then repeat this process through the entirety of a given 90-second trial. Essentially, we perform a simulated rollout of a given trial in which the eyes are approximately “locked on” to the ground-truth **focal point** while the head rotates at the behest of the controller.

B. Results – Autoregressive

We show in Fig. 3 sample trajectories of three different models produced by the autoregressive routine. In the first two plots—those of MLP-H16_16 and LSTM-H2—the imputed head direction roughly follows the actual head direction. The third plot, however, shows that the head direction deviated dramatically from its true course. *Qualitatively*, this indicates that this controller will exhibit undesirable behavior in the VR experiment or that of the robot.

Further, Fig. 4 compares mean squared error (MSE) between imputed and actual head directions, averaged over all trajectories, among various models. With these results, we can *quantitatively* assess the viability of different configurations for each model. Models which have consistently low MSE scores across tasks can be deemed the most viable, whereas those with high variability or with consistently high scores are unlikely to be stable in VR, an inherently autoregressive environment.

Lastly, Table I shows a comparison of model performance during training versus in the autoregressive experiment. The values in each row in the table represent the error of

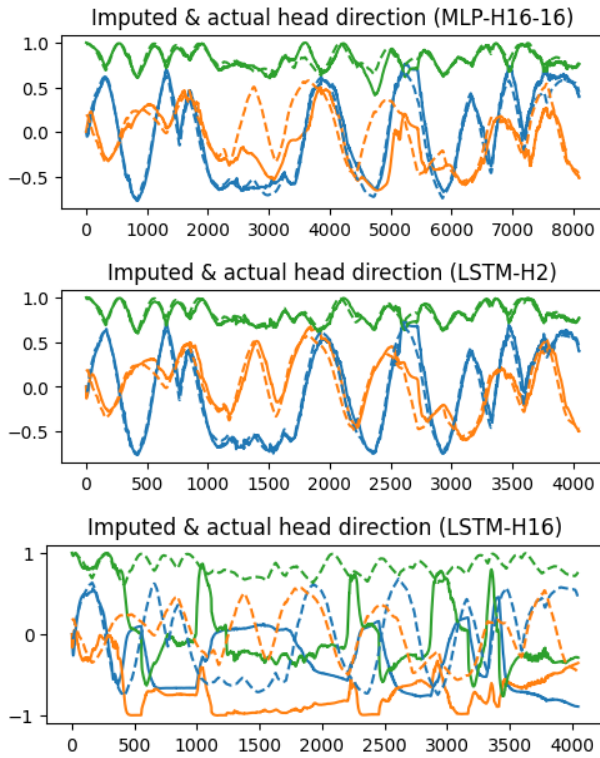


Fig. 3. Trajectories of an autoregressive rollout for one MLP and two LSTM variations. The y-axis shows the x (blue), y (orange), and z (green) values of the normalized direction of the head where $(0, 0, 1)$ represents the forward direction. The x-axis shows the time step t throughout the trajectory. Solid lines represent the direction of the simulated, rotated “virtual head” during the experiment. Dashed lines represent the ground-truth direction of the head as seen in the dataset. **(Top)** An MLP with two intermediate hidden layers each of size 16. **(Middle)** An LSTM with a context vector of size 2. **(Bottom)** An LSTM with a context vector of size 16.

the corresponding model on the held-out test set and on the ground-truth head direction during the autoregressive procedure, respectively.

TABLE I
MSE LOSS - TRAINING AND AUTOREGRESSIVE

Model	Training	Autoregressive
MLP	0.00011273	0.370922
MLP-H16_16	0.00011108	0.499707
LSTM	0.00032809	0.248416
LSTM-H2	0.00032678	0.261365
LSTM-H16	0.00037722	1.608792
LSTM-H128	0.00190842	2.125073
Quadrant (baseline)	0.01127953	1.194305
Vector (parameterized)	0.00011084	0.280166

C. Discussion – Autoregressive

Our experiment shed light on which classes of controller may perform appropriately in VR and with the exoskeleton. Notably, as the size of the MLP and the hidden state within LSTM architectures increases, the stability of the head decreases under autoregressive conditions. For this reason, we selected an LSTM and MLP model with very small hidden states to be evaluated in the VR experiment. The

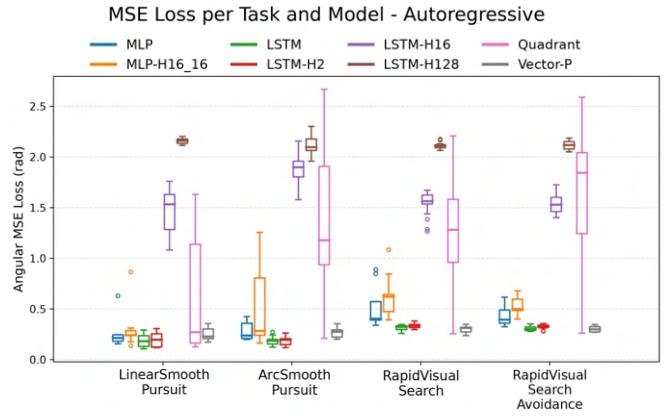


Fig. 4. Distribution of mean squared error (lower is better) between imputed head direction and ground-truth head direction for various models over trajectories from the held-out test set. Results are grouped by task.

Baseline Quadrant and Vector Parameterized controllers also moved forward to the VR experiment. Rejecting the four other controllers saves expensive time that is used instead for testing the more promising controllers.

IV. VR EXPERIMENT

Having rejected four controllers through simulation, we carried the remaining four — Baseline Quadrant, Vector Parameterized, MLP, and LSTM — forward to the next evaluation layer. While the autoregressive results suggest these controllers behave stably in simulation, simulated rollouts cannot capture a real user’s responses or preferences. We therefore deployed them in a VR environment, allowing human subjects to interact with each controller safely, before deploying them on the physical neck exoskeleton.

A. Methods – VR

30 healthy subjects (21 male, 19-50 years old) volunteered to participate in the VR portion of the experiment. A healthy cohort allows us to prevent the unsafe controllers from reaching patients with DHS. 13 subjects had previous experience with VR; 19 wore prescription glasses daily; 5 participants wore glasses while wearing the VR headset. This experiment was approved by the Institutional Review Board (IRB) of The University of Utah. Informed consent was obtained from each participant prior to the experiment.

Our VR environment employs the open source code given by the eye-head coordination dataset paper [25], developed using the Unity game engine¹. We used the VIVE Pro Eye VR headset with eye tracking technology². This allowed us to mimic neck exoskeletons by moving the VR environment according to eye positions as if the subject was moving their head, even though their head was still.

The four tasks (Fig. 5), used to both create the dataset and perform this study, represent a diverse collection of eye-head behaviors. For all tasks, targets (cubic blocks) were

¹<https://unity.com/>

²https://www.vive.com/us/support/vive-pro-eye/category_howto/about-the-headset.html

constrained within a 140° by 140° conic range. When the gaze was fixed on a target, its color switched from blue to green.

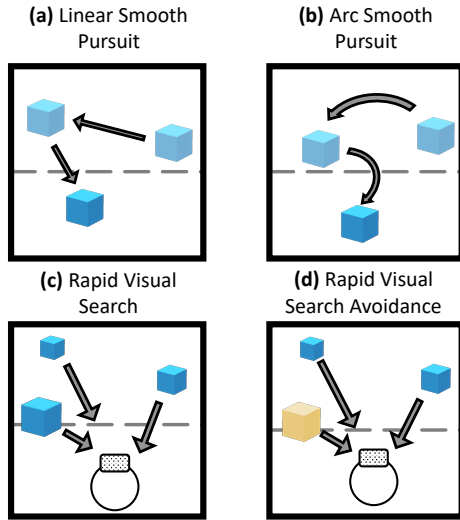


Fig. 5. Example target trajectories. (a) The target in the *Linear Smooth Pursuit* task moves in straight lines of random direction and distance. (b) The target in the *Arc Smooth Pursuit* task moves in circular lines of random direction, distance, and curvature. Both pursuit tasks feature only a single target. (c) Each of the three targets in the *Rapid Visual Search* task starts some distance from the participant and moves towards the participant. (d) In the *Rapid Visual Search Avoidance* task there are three targets to fixate on (blue) and three blocking targets to not fixate on (yellow). Each target and blocking cube starts from a plane some distance from the participant and moves towards the participant. The dashed line represents a horizon.

The first two tasks were designed to measure and test head-eye tracking behaviors for a single, slow-moving target that either follows a sequence of random linear (*Linear Smooth Pursuit*) or circular (*Arc Smooth Pursuit*) paths. In the *Linear Smooth Pursuit* task, a singular cube is instantiated in front of the user and moves in a straight line to a randomly assigned target point with a fixed speed (5 m/s). When reached, a new random target point is assigned from a uniform distribution, and the cube moves to the new point. The *Arc Smooth Pursuit* task is similar, with the change that the cube moves in circular arc patterns with a fixed angular speed (1 rad/s).

The latter two tasks, *Rapid Visual Search* and *Rapid Visual Search Avoidance*, were intended to measure and test head-eye behavior during searching for discrete targets and brief fixation. In the *Rapid Visual Search* task, three cubes are randomly instantiated within a 60° head rotation of the user at a fixed distance (10 m). The cubes move towards the participant at a fixed speed. If the user fixates their gaze on a cube for 0.3 seconds or a cube reaches the user, the cube respawns at a new random location drawn from a uniform distribution. The *Rapid Visual Search Avoidance* task is similar but also has distracting yellow cubes that turn orange on fixation. These distracting blocks do not disappear when fixated on and respawn when they reach the user.

Four controllers were selected after the simulation experiment, one from each category. The subject's goal was to perform the same VR tasks made for the eye-head

coordination dataset [25] with each controller and perform pairwise preference comparisons. For each task, the subjects first performed the task for 30 seconds without eye-tracking control to become familiar with the task. During this time, the subjects were able to explore the environment by physically moving their own head. Then, the subjects performed the task twice for two different controllers for 30 seconds each and chose their controller preference. During this time, the environment would move according to the subject's eye positions and not if the subject moved their head. The same process was repeated twice where preferred controllers moved forward to the next comparison for a total of three pairwise comparisons between four controllers. For controller comparison purposes, a point was given to the controller that was most preferred.

This comparison process was performed for each of the four tasks. Subjects were allowed to make no preference (i.e., both controllers were equally favored or disfavored) and, if no preference was chosen, one of the two controllers was randomly selected to move forward to the next comparison. If multiple controllers were equally preferred at the end of the task, the point was distributed evenly among all the preferred controllers. All pairwise orderings were randomly generated.

B. Results – VR

Since there were 30 subjects, for each of the four tasks, 30 points were distributed among the controllers. Fig. 6 illustrates this. If for each task we rank the controllers according to their accumulated points, the ranking is the same for each task: Vector Parameterized is most preferred, followed by MLP, Baseline Quadrant, and LSTM, in that order. Because the LSTM model performed so poorly in terms of preference points as well as task score, it was rejected from the potential controller pool and barred from participating in the physical neck exoskeleton experiment.

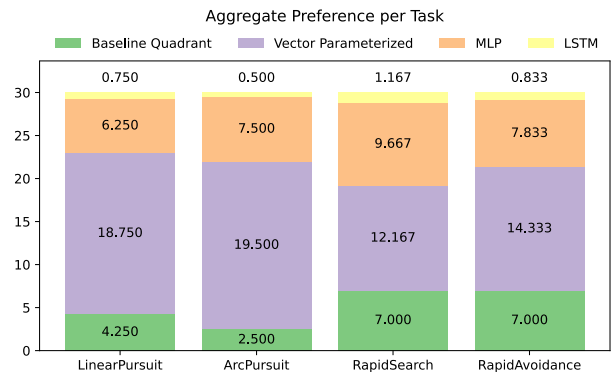


Fig. 6. Stacked bar plot showing aggregate controller preference for each VR task. Since there were 30 subjects, the total number of points distributed across each controller was 30. A bar of greater length, or higher score, correlates to a more preferred controller for that task.

C. Discussion – VR

The MLP, Baseline Quadrant, and Vector Parameterized controllers have proven effective for performing tasks in a VR environment.

The LSTM model, while promising in an autoregressive simulation, proved nonviable and unwanted for restoring head-neck motions, challenging the hypothesis that previous eye and head positions are useful for predicting future head motions. We attribute the disparity in performance of the LSTM between training and VR to the problem known as teacher forcing [29]. This occurs because, during training, predictions are discarded at each time step, whereas when used as an online controller, predictions affect the next state.

Employing VR as a tool for safely testing and rejecting bad or unsafe controllers has therefore saved more time for testing more promising controllers.

V. NECK EXOSKELETON EXPERIMENT

Even if a controller performs well in VR, this does not guarantee it will perform well in the physical robot. We therefore deployed and tested the three remaining controllers in a real-world environment with a physical neck exoskeleton.

A. Methods – Neck Exoskeleton

1) *Neck Exoskeleton Technical Details:* For the real-world robot experiment, we deployed our controllers on the Columbia Brace [30] (Fig. 7) whose design was optimized for maximizing range of motion and was thus a good fit for our healthy cohort. However, several improvements have been made to the exoskeleton.




Fig. 7. The Columbia Brace [30] neck exoskeleton. Links are manufactured from PLA 3D printed parts. Note that in this paper’s implementation, the eye tracking glasses are not attached to the robot, but are worn directly by the user on the Neon glasses frame from Pupil Labs.

For a smoother experience, velocity control was employed instead of position control with a maximum rotational speed of 1 radian per second. Software boundaries were set to prevent the robot from entering uncontrollable regions. Specifically, the robot was allowed 25° flexion, 3° extension, and 30° side-to-side rotation in both directions. Roll or bending of the head was constrained to 0° . The robot was allowed to slide along its software boundaries.

To solve the forward kinematics problem, we used a neural network model as described in [31]. However, we increased the size of the model to 4 hidden layers of 200 nodes each and trained it on 28 million kinematic data points sampled uniformly to have a resolution of 0.1° in each of the three head pose directions over the allowed robot workspace.



Fig. 8. A subject provides preference regarding eye tracking controllers applied to the Columbia Brace [30]. The eye tracking frame is placed over the subject’s glasses. The subject evaluated each controller by searching for a specific symbol and color combination on the wall. The clutter on the wall provides distraction to make the task more difficult and realistic.

To measure eye angles, we used a commercial eye tracking system (Neon, Pupil Labs, Berlin, Germany) with an accuracy of 1.8° [32]. The system does not require calibration and is resistant to factors such as slipping, lighting, and blinking, and may be worn over glasses. The gaze positional data were sampled at 200 Hz, followed by filtering using an exponential moving average with a smoothing factor of 0.1. Each of the controllers was run on a laptop CPU. Overall, the robotic system was operated at 50 Hz.

2) *Neck Exoskeleton Experiment:* All subjects from the VR evaluation experiment were invited, of whom twelve (7 male, 25-35 years old) agreed, to participate in the robotic evaluation follow-up. Signed informed consent was obtained prior to data collection, according to the IRB.

In this portion, subjects evaluated controllers during a real-world task. On a wall in front of the subject, 48 colored symbols were taped (Fig. 8). The set of symbols was $\{X, +, \backslash, /, <, >, |, -\}$ and the set of colors was $\{\text{blue, yellow, pink, orange, green, silver}\}$. The task was to seek and find as many symbols as possible in 90 seconds. None of our subjects were colorblind, 10 wore glasses regularly, and 5 wore glasses during the experiment. Eye tracker accuracy was not appreciably affected by the presence of the glasses. This task modeled a real-world visual search for specific targets among many distracting objects. The targets spanned 90° evenly across the subjects’ field of view.

Users were given 5 minutes to become comfortable controlling the neck exoskeleton. Afterward, the subject followed an experimental protocol similar to the VR experiment. Each subject performed the task with two of the controllers (MLP, Vector, or baseline) in turn and then stated a preference for either of them. The subject then performed the task with the preferred controller and the remaining controller and stated a preference for either of those. A ‘no preference’ option was allowed, and all orderings were randomized. Subjects were blinded to which controller they were using. If a subject expressed no preference between controllers in the first comparison, each controller was considered equally preferred, and the controller that moved

Quadrant	Vector		MLP	
subject	1st comparison		2nd comparison	
0	11	17	14	11
1	13	18	19	15
2	10	4	13	16
3	6	8	8	8
4	16	19	18	20
5	13	15	15	13
6	14	16	13	15
7	18	12	16	14
8	9	14	11	13
9	9	6	8	9
10	9	13	13	12
11	10	10	13	14

Fig. 9. Task score and preference rankings for neck exoskeleton experiment. The colored cells represent a trial using a particular controller; green for Baseline Quadrant, blue for Vector Parameterized model, and orange for MLP. Reading from left to right corresponds to the order each controller was given to the subjects. The score (number of found symbols) is printed on each cell which is colored with a shading that corresponds to the normalized score for that subject. Thus, for a row, a darker shaded cell corresponds to a higher score. A bold score indicates the controller was preferred in that comparison. If both scores are bold then equal preference was reported. Overall, three subjects preferred Vector Parameterized, four preferred MLP, and five preferred Baseline Quadrant.

forward was randomly selected.

B. Results – Neck Exoskeleton

The task scores and preferences for each controller for the neck exoskeleton experiment are reported in Fig. 9. By normalizing each subject’s score and averaging the score of the repeated controller, we found that the mean normalized score (standard deviation) for the Baseline Quadrant, Vector Parameterized, and MLP controllers respectively are 0.575(0.068), 0.543(0.092), and 0.596(0.078). Each controller thus performed similarly well on average. Three subjects preferred the Vector Parameterized controller, four preferred the MLP, and five preferred the Baseline Quadrant. Overall, the Baseline Quadrant controller was the most preferred in this experiment, but all controllers passed the evaluation. Five of the subjects gave preference to their highest scoring controller, while the rest preferred controllers that underperformed compared to the other controllers.

When comparing overall controller preferences in the VR experiment and the Neck Exoskeleton experiment, shown in Fig. 10, preferences were in agreement for three subjects. If we perform the same comparison for the first two VR tasks, the tracking pursuit tasks, preferences were in agreement for four subjects. The same comparison for the final two VR tasks, the search and fixate tasks, shows that preferences were in agreement for six subjects. This matches our intuition that, because the search and fixate tasks are the most similar to the real-world neck exoskeleton task, we should see that preferences over controllers are most aligned between the VR environments and the real-world exoskeleton environments in these settings. In the separate bottom row of Fig. 10, we

again see that the Vector Parameterized model was the most preferred controller in the VR experiment and the Baseline Quadrant was the most preferred controller in the Neck Exoskeleton experiment.

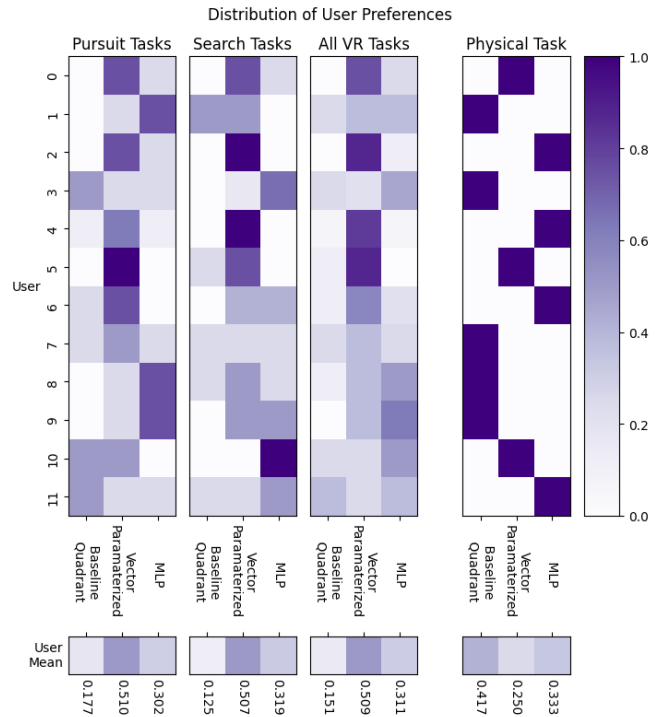


Fig. 10. User preferences per user and averaged over all users. Darker shading indicates a more preferred controller. *Linear* and *Arc Pursuit* were grouped together into Pursuit Tasks because they are similar. *Rapid Search* and *Rapid Avoidance* were also grouped together due to their similar natures. The preference over all tasks is also present in the third column and the overall controller preference during the physical neck exoskeleton experiment is present in the fourth column. Preference averaged over all subjects is presented in the bottom row.

C. Discussion – Neck Exoskeleton

Three controllers passed the VR evaluation, and all three also passed the neck exoskeleton evaluation. Employing our models in VR first has thus allowed us to find viable controllers before deployment on a physical robot.

We found that for VR tasks similar to the real-world task, half of the subjects preferred the same controller. Furthermore, we found that while the Baseline Quadrant controller was preferred the least of the three final controllers in VR, it was preferred the most with the physical neck exoskeleton. When making preference, subjects weighed task performance less than other values such as comfort or safety.

The sim-to-real gap between VR and the real world partially accounts for differences of controller preference. The physical neck exoskeleton exerts physical forces on the user’s head. The controllers in VR do not account for this dynamic. The sim-to-real gap may also account for the extra preference for the Baseline Quadrant controller on the neck exoskeleton. In the VR environment, the Baseline Quadrant controller could only move in strict horizontal and vertical

directions, which may feel jarring. The controller behaves the same algorithmically on the neck exoskeleton, but may feel less jarring due to the robot's dynamics, including damping. Indeed, one of our subjects verbally expressed that the Baseline Quadrant controller felt smooth.

Subjects also expressed that the Vector Parameterized controller moved too fast, was harder to control horizontally, and still moved slightly when fixating on a target. It was expressed that the MLP controller was harder to control vertically. The Baseline Quadrant controller may be preferred more often because it is simple and predictable.

The results of our physical experiment provide evidence that we have found two more viable controllers beyond the Baseline Quadrant that may serve the purpose of restoring head motions. Because there is a large amount of disagreement among user preferences, our results suggest that personalization is an important factor in gaze-driven controllers, motivating the need for multiple viable controllers.

VI. CONCLUSION

We demonstrated a novel multi-layer design and evaluation framework that leverages VR to safely and efficiently test for viable neck exoskeleton controllers before physical deployment, saving significant development time and resources. As a result, we developed and verified two novel gaze-driven eye-tracking controllers for restoring head-neck motion. Interestingly, our user study results reveal that no single controller is universally preferred: while some subjects favored a simple baseline, others preferred our novel data-driven models. These findings highlight the importance of personalization in assistive robotics [33] and motivate the need for future work on learning the preferences [34] and implicit constraints [35] of neck exoskeleton users.

REFERENCES

- [1] D. Lange, M. Fetell, R. Lovelace, and L. Rowland, "The floppy head syndrome," *Annals of Neurology*, vol. 20, no. 1, pp. 133–133, 1986.
- [2] J. P. Drain, S. S. Virk, N. Jain, and E. Yu, "Dropped head syndrome: a systematic review," *Clinical Spine Surgery*, 2019.
- [3] G. A. Suarez and J. J. Kelly, "The dropped head syndrome," *Neurology*, 1992.
- [4] J. Katz, G. Wolfe, D. Burns, W. Bryan, J. Fleckenstein, and R. Barohn, "Isolated neck extensor myopathy: a common cause of dropped head syndrome," *Neurology*, 1996.
- [5] J. Finsterer, M. Frank, and E. Krexner, "Steroid-responsive dropped-head-syndrome due to polymyositis," *Joint Bone Spine*, 2010.
- [6] J. D. Brodell Jr, A. Sulovari, D. N. Bernstein, P. C. Mongioli, E. Ciafaloni, P. T. Rubery, and A. Mesfin, "Dropped head syndrome: an update on etiology and surgical management," *JBJS reviews*, 2020.
- [7] M. C. Kiernan, S. Vucic, B. C. Cheah, M. R. Turner, A. Eisen, O. Hardiman, J. R. Burrell, and M. C. Zoing, "Amyotrophic lateral sclerosis," *The lancet*, vol. 377, no. 9769, pp. 942–955, 2011.
- [8] A. D. Walling, "Amyotrophic lateral sclerosis: Lou gehrig's disease," *American family physician*, vol. 59, no. 6, pp. 1489–1496, 1999.
- [9] K. M. Burke, C. Shea, V. Arulanandam, S. Sullivan, A. S. Ellrodt, C. MacAdam, K. Carney, G. Casagrande, E. Christiansen, and S. Paganoni, "Cervical collar satisfaction and functional impact in amyotrophic lateral sclerosis: A survey study," *American Journal of Physical Medicine & Rehabilitation*, pp. 10–1097, 2025.
- [10] S. Pancani, J. Rowson, W. Tindale, N. Heron, J. Langley, A. D. McCarthy, A. Quinn, H. Reed, A. Stanton, P. J. Shaw *et al.*, "Assessment of the sheffield support snood, an innovative cervical orthosis designed for people affected by neck muscle weakness," *Clinical biomechanics*, vol. 32, pp. 201–206, 2016.
- [11] C. P. Miller, J. E. Bible, K. A. Jegede, P. G. Whang, and J. N. Grauer, "Soft and rigid collars provide similar restriction in cervical range of motion during fifteen activities of daily living," *Spine*, 2010.
- [12] H. Zhang and S. K. Agrawal, "An active neck brace controlled by a joystick to assist head motion," *IEEE Robotics and Automation Letters*, vol. 3, no. 1, pp. 37–43, 2017.
- [13] S. P. Torrendell, H. Kadone, M. Hassan, Y. Chen, K. Miura, and K. Suzuki, "A neck orthosis with multi-directional variable stiffness for persons with dropped head syndrome," *IEEE Robotics and Automation Letters*, vol. 9, no. 7, pp. 6224–6231, 2024.
- [14] P. Kulkarni and S. K. Agrawal, "Design and validation of a novel robotic neck brace for cervical traction," *IEEE/ASME Transactions on Mechatronics*, vol. 29, no. 4, pp. 3092–3099, 2024.
- [15] I. Bales and H. Zhang, "Kinematic benefits of a cable-driven exosuit for head-neck mobility," *IEEE Robotics and Automation Letters*, 2024.
- [16] D. Demaree and H. Zhang, "A structurally enhanced neck exoskeleton to assist with head-neck motion," in *2023 International Symposium on Medical Robotics (ISMR)*. IEEE, 2023, pp. 1–7.
- [17] P. K. Lingampally and A. A. Selvakumar, "Kinematic and workspace analysis of a parallel rehabilitation device for head-neck injured patients," *FME Transactions*, vol. 47, no. 3, 2019.
- [18] H. Zhang, B.-C. Chang, P. Kulkarni, J. Andrews, N. A. Shneider, and S. Agrawal, "Amyotrophic lateral sclerosis patients regain head-neck control using a powered neck exoskeleton," in *2022 9th IEEE RAS/EMBS International Conference for Biomedical Robotics and Biomechatronics (BioRob)*. IEEE, 2022, pp. 01–06.
- [19] D. Demaree, J. Brignone, M. Bromberg, and H. Zhang, "Preliminary study on effects of neck exoskeleton structural design in patients with amyotrophic lateral sclerosis," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2024.
- [20] D. Tweed, B. Glenn, and T. Vilis, "Eye-head coordination during large gaze shifts," *Journal of neurophysiology*, 1995.
- [21] M. F. Land, "Eye movements and the control of actions in everyday life," *Progress in retinal and eye research*, 2006.
- [22] E. G. Freedman, "Coordination of the eyes and head during visual orienting," *Experimental brain research*, 2008.
- [23] E. G. Freedman and D. L. Sparks, "Eye-head coordination during head-unrestrained gaze shifts in rhesus monkeys," *Journal of neurophysiology*, 1997.
- [24] D. Tweed, "Three-dimensional model of the human eye-head saccadic system," *Journal of Neurophysiology*, 1997.
- [25] C. Rubow, C.-H. Tsai, E. Brewer, C. Mattson, D. S. Brown, and H. Zhang, "A dataset of paired head and eye movements during visual tasks in virtual environments," *Scientific Data*, 2024.
- [26] B.-C. Chang, H. Zhang, E. Trigili, and S. K. Agrawal, "Bio-inspired gaze-driven robotic neck brace," in *2020 8th IEEE RAS/EMBS International Conference for Biomedical Robotics and Biomechatronics (BioRob)*. IEEE, 2020, pp. 545–550.
- [27] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [28] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [29] R. J. Williams and D. Zipser, "A learning algorithm for continually running fully recurrent neural networks," *Neural Computation*, vol. 1, no. 2, pp. 270–280, 1989.
- [30] H. Zhang and S. K. Agrawal, "Kinematic design of a dynamic brace for measurement of head/neck motion," *IEEE Robotics and Automation Letters*, vol. 2, no. 3, pp. 1428–1435, 2017.
- [31] A. Prado, H. Zhang, and S. K. Agrawal, "Artificial neural networks to solve forward kinematics of a wearable parallel robot with semi-rigid links," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 14 524–14 530.
- [32] C. Baumann and K. Dierkes, "Neon accuracy test report," *Pupil Labs*, vol. 10, 2023.
- [33] J. Z.-Y. He, Z. Erickson, D. S. Brown, A. Raghunathan, and A. Dragan, "Learning representations that enable generalization in assistive tasks," in *Conference on Robot Learning*. PMLR, 2023, pp. 2105–2114.
- [34] Y. Liu, G. Datta, E. Novoseller, and D. S. Brown, "Efficient preference-based reinforcement learning using learned dynamics models," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 2921–2928.
- [35] D. Papadimitriou and D. S. Brown, "Bayesian constraint inference from user demonstrations based on margin-respecting preference models," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 15 039–15 046.