2024 33rd IEEE International Conference on
Robot and Human Interactive Communication (RO-MAN)
August 26-30, 2024. Pasadena, California, USA.

Non-expert Caregivers to Improve the Identification of a Physiologically Actuated Robot

Cameron J. LaMack[1] and Eric M. Schearer[1]

Abstract— **Functional electrical stimulation is a promising technique for restoring arm function to those with paralysis from a high spinal cord injury. While simple controllers are easy to implement, model-based controllers are likely better equipped to leverage the arm's kinematic and dynamic complexity, particularly for the high variations associated with functional arm movement. One modelling technique for a model-based controller is Gaussian Process Regression. Previous simulation work has shown promise leveraging whole-arm error data to identify the arm's various subsystems, but used perfect simulated data. We asked caregivers to correct a robotic arm's movement as simulated muscles generated torque. The simulated muscles were controlled as if they were electrically stimulated human arm muscles. This study demonstrates non-expert caregivers' ability to collect this error data via whole-arm corrections, and provides insight into their ability to improve arm subsystem models made with Gaussian Process Regression. Despite significant error in caregivers' ability to provide force corrections to hold the robot in a static configuration, these corrections were leveraged to significantly improve muscle models; the muscles that improved the most were the ones primarily used to move the physiologically actuated robot.**

I. INTRODUCTION

In the United States alone, there are an estimated 302,000 people living with a spinal cord injury, with an additional 18,000 occurring each year [1]. Of these people, an estimated 59.6% have a spinal cord injury resulting in tetraplegia, or some loss of upper limb function [1]. Regaining arm and hand function, along with the sense of independence it brings, is of the highest priority for people with tetraplegia [2]. The ability to perform tasks of daily living is highly correlated with a high quality of life as well as social community participation [3].

Functional electrical stimulation (FES) is a promising technology with the capability to restore arm function to people post-stroke [4] or after a spinal cord injury. FES is a technique in which paralyzed muscles receive targeted electrical current in a manner that provides functional movement [5]. The electrical stimulation can be applied with an implanted system [6] or with surface electrodes [7]. FES has shown success with restoring hand functions in people with lower cervical spinal cord injuries (C5-C6) [8], but for those with high cervical spinal cord injuries, restoration of the whole arm with FES is required.

This work was supported by NSF Grant 1751821.

[1]The authors are with the Center for Human-Machine Systems at Cleveland State University and the Department of Physical Medicine and Rehabilitation at the MetroHealth System, Cleveland, OH 44115-2214 USA. c.lamack@vikes.csuohio.edu; e.schearer@csuohio.edu

Controlling whole-arm movement with electrical stimulation is a challenging task. The human arm is a coupled and highly redundant system. Each joint has at least 2 muscles acting to control it, though typically more. Many of these muscles are also biarticular, which means that they actuate multiple joints. The biceps, for example, are engaged in both elbow and shoulder motion [9]. Lastly, for people with spinal cord injuries, spinal reflexes and perturbations due to spastic muscle contractions add another layer of complexity to the response of muscles to FES [10].

A variety of control strategies have been implemented in attempts to control whole-arm movement. A controller which uses fixed, pre-defined muscle stimulation patterns has been performed [6]. While fixed stimulation patterns accomplish movement and serve as a proof of concept, this method does not account for the variation in the reality of functional tasks. It also does not account for the variation of muscle capability common in people with spinal cord injuries [11]. Using fixed patterns does not account for any variation in the system or surroundings, so this is not a good long term solution for daily use.

Controllers which are capable of following any arbitrary trajectory, such as a proportional-derivative feedback controller [12] or using reinforcement learning [13], have been successful in simulations, but have challenges associated with modelling real human arms for control. One reason for this limited success is the constantly changing dynamics of human arms. As time progresses, the gains of a PD controller or numerous parameters of neural network grow non-representative of the true dynamics of the paralyzed arm.

Researchers in [14] demonstrated the feasibility of a model-based controller using a technique called Gaussian Process Regression (GPR) [15]. The controller used models of the arm's subsystems, the muscles and passive dynamics, to predict their effects on the whole arm system. The researchers demonstrated the ability to control the arm such that it accomplished a variety of arbitrary movements.

Researchers in [16] leveraged GPR to begin to overcome the problem of the arm's changing dynamics. They used large subject-specific data sets to identify parameters used in the regression. However, a smaller data set was used on the day of experiments as it was representative of the paralyzed person's arm on during the experimental reaches.

While the "day of" data collection technique addresses the problem of the arm's changing dynamics, it exacerbates an additional challenge of electrical stimulation for paralyzed arms: translation to home use from a research lab. In order for a paralyzed person to use their electrically stimulated arm, a Gaussian Processes Regression system identification expert would have to visit their home daily to develop the "day of" model. This is not a good long-term solution, as such expertise is limited.

A better long-term solution is to have the paralyzed arm identification data recollected by a non-expert caregiver, someone already at home with the paralyzed person. In our previous work [17], we have shown that whole-system error can be used to either improve subsystem models made with Gaussian Process Regression, or generate them from scratch. This modelling technique is a fantastic option for non-expert caregivers to generate "day of" arm models. The caregiver merely needs to generate the force/torque error of the whole arm while it is being moved through functional tasks, and that data can be used to produce *all* the models used to control the electrically stimulated arm.

The ability of caregivers to provide accurate whole-system error remains an unknown factor. In our simulation study to demonstrate the feasibility of our modelling technique [17], we used theoretically perfect force/torque error measurements. In a real implementation with a paralyzed person's arm, the force/torque error of the arm is unknown to a caregiver. A simple and intuitive method for the caregiver to measure the force/torque error in the paralyzed arm during functional movements is to apply a correction to the arm as it moves. As the arm deviates from a desired trajectory, the corrective force applied by a caregiver is a measurement of the whole-system error. However, the ability of a caregiver to apply accurate corrections is unknown. Additionally, the effects of human generated whole-system error measurements on model accuracy are unknown.

The first goal of this study is to measure the accuracy of human caregiver corrections to a robotic arm. The second goal of this study is to determine if corrections from non-expert human caregivers can improve the system identification models of the robotic arm. The robotic arm will be actuated with simulations of electrically stimulated human arm muscles rather than using a real human arm so there is a known ground truth to measure model accuracy and improvement.

II. METHODS

To accomplish our goals, we asked several non-expert human participants to hold a robotic arm still as it generated torques in a way that simulated electrically stimulated muscles. The muscle activations were randomly determined. The participants, who we will refer to as "caregivers," applied force to the "wrist," or robotic end-effector, which was instrumented with a force/torque sensor. We used the corrective force applied to the arm by the caregivers to update the arm subsystem models. Since the corrections were representative of all subsystems of the arm, such as the inverse statics and muscle contributions, they were able to be used as additional training data to improve the model accuracy. We then compared the models which included caregiver correction data to those which did not in order to determine the effects of including the whole-arm error data. Additionally, since the equations governing the robotic arm's motion were known, we were able to calculate the theoretically perfect amount of force for a correction. We compared the experimental corrections to these calculated perfect corrections to measure caregiver accuracy. Note that throughout this paper, "muscle equation" refers to the ground truth equations used to simulate muscle torque (1), and "muscle model" refers to predictions made with Gaussian Process Regression.

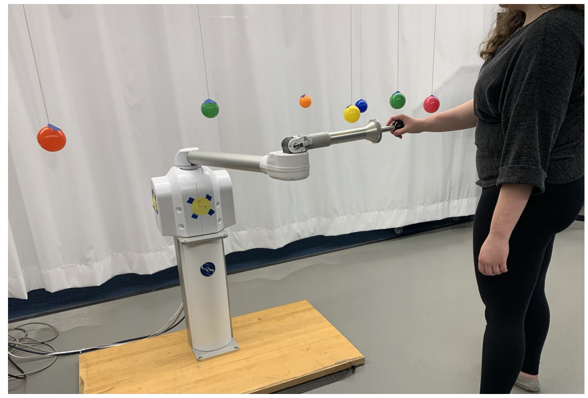

Fig. 1. An experimental caregiver applying a force to the wrist's force/torque sensor to hold the arm static. The hanging plastic balls act as frame of reference for the end-effector target position.

A. Caregiver Information

A total of 8 able-bodied caregivers were recruited to act as experimental caregivers. Their ages ranged from 20 to 28. Informed written consent was obtained for each caregiver according to the protocols approved by the Institutional Review Board at Cleveland State University (IRB-FY2016-331).

B. Experiment and Set-Up

For each trial, the arm was moved to one of 7 possible positions, each repeated 3 times. The caregiver was then verbally instructed to hold the arm static at that position, as in Fig 1. Plastic balls were hung above the goal position to provide a visual target and frame of reference to the caregivers. Then, simulated electrical stimulation was applied to the muscles actuating the arm. The electrical stimulation was constant throughout the 2 second trial, but was different between trials. The electrical stimulation was randomly determined. For each random stimulation pattern, the six individual muscles all had a different activation level drawn from a uniform distribution ranging from 0 to 1 (0% to 100%). We only used randomly generated muscle activations as our previous work [17] showed that using random muscle activations led to better muscle model improvement. This occurred because data was collected that was representative of each subsystems contributions in various states and with various control signals. Additionally, using randomized muscle activations supported our first goal of measuring caregiver correction accuracy by encouraging larger force productions in the muscles.

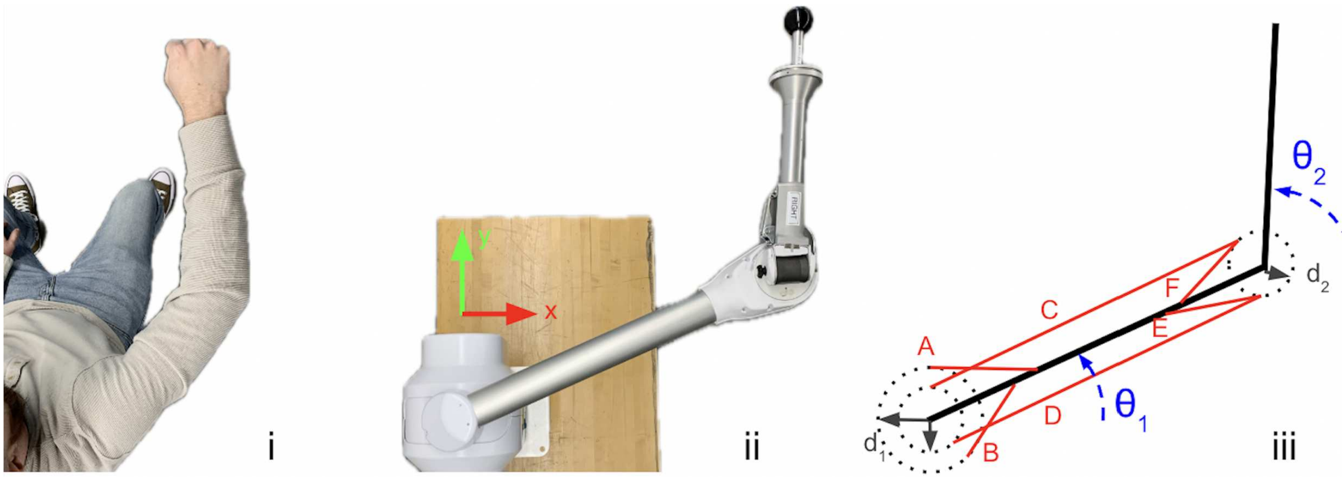

Fig. 2. The robotic arm emulated a planar human arm actuated with 6 muscles. A physical arm would have looked like (i), however we used a robot (ii). The robot was controlled based on a 2 link arm with 6 muscles (iii). Each muscle had constant moment arms. The muscles were A) the anterior deltoid, B) the posterior deltoid, C) the biceps, D) the triceps (long head), E) the triceps (short head), and F) the brachilialis.

There were a total of 21 trials with 21 different random muscle activation patterns, assigned randomly to the various arm configurations. The same 21 random muscle activation patterns were used for all caregivers, but at different arm configurations. As the muscle equations which governed the arm torque were position dependent, a muscle with a high activation might have a large force contribution in one position, but a small force contribution in another position.

We used a robotic arm controlled in such a way that it mimicked an electrically stimulated human arm so that we could calculate the ground truth of the muscle dynamics. This would not be possible with a person's arm. It was important to be able to calculate the ground truth of the muscle dynamics so we could measure model improvement, as well as calculate the force/torque error of the arm to determine the accuracy of each caregiver's correction.

1) Hardware: Caregivers interacted with a Proficio (Barrett Technology, Newton, MA) robotic arm. It has 3 degrees of freedom and an integrated 6 axis force/torque sensor.

2) Software: Custom C++ code was written so that the robotic arm emulated the planar arm shown in Figs. 2 and 3. The arm was "actuated" by a total of 6 Hill-Type [18] muscles. Each muscle's force contribution was calculated according to the equations below, and then multiplied by a constant moment arm. The sum of the muscle torques about each robot joint were then passed to the robot's motors.

$$F = \alpha F_{max} f(L_m) g\left(\dot{L_m}\right) + K_P (L_m - L_{slack})^2 \tag{1}$$

$$f(L_m) = \exp\left(-\left(\frac{L_m - L_{opt}}{0.65 L_{opt}}\right)^2\right) \tag{2}$$

$$g\left(\dot{L_m}\right) = \begin{cases} \frac{V_{max}+\dot{L_m}}{V_{max}-4\dot{L_m}} & \text{if } \dot{L_m} \leq 0 \\ \\ \frac{1.5\dot{L_m}+c}{\dot{L_m}+c} & \text{if } \dot{L_m} > 0 \end{cases} \tag{3}$$

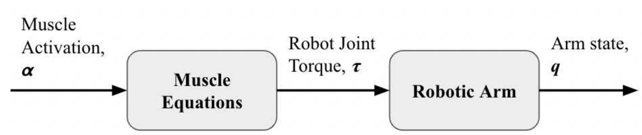

Fig. 3. The open loop controller was given randomized muscle activations, then used those muscle activations in (1) to determine the robotic joint torques.

The force F in each muscle was multiplied by constant moment arms to calculate the contribution of each muscle to the two joint torques. The muscle lengths L_m were dependant on joint angles such that $L_m = a_0 - d_1\theta_1 - d_2\theta_2$. The parameters of each muscle were taken from [12], and are presented in Table I. In addition to the active component scaled by the muscle activation $\alpha \in [0,1]$ control signal, each muscle had a quadratic spring with a stiffness of $K_P = 7500\ N/m^2$. This was much stiffer than a realistic human arm, but people with spinal cord injuries often have a passive, elastically actuated arm support [14]. This stiffness also ensured that some amount of force would be required to hold the planar arm static in any configuration.

C. Gaussian Process Regression Models

We used two types of models to control the simulated muscles with muscle activations. We used an inverse statics model to predict the amount of force required to hold the arm static in any desired configuration. We also used six muscle capability models to predict the muscle force contribution from muscle activation. In terms of (1), the single inverse statics model attempted to predict the sum of $K_P(L_m - L_{slack})^2$ from all muscles, and each muscle capability model attempted to predict $\alpha F_{max} f(L_m) g\left(\dot{L_m}\right)$ when the muscle was fully activated ($\alpha = 1$). This is the component of the muscle that is controllable with the muscle activation control signal.

To collect data to generate initial model predictions (mod-

TABLE I

THE NAMES AND PARAMETERS OF THE MUSCLES USED TO CALCULATE MUSCLE FORCES

Muscle Name	F_{max} (N)	L_{slack} (m)	L_{opt} (m)	V_{max} (m/s)	c	a_0 (m)	d_1 (m)	d_2 (m)
anterior deltoid	800	0.0538	0.1280	1.280	0.1280	0.1840	0.05	0
posterior deltoid	800	0.0538	0.1280	1.280	0.1280	0.1055	0.05	0
biceps	1000	0.2298	0.1422	1.422	0.1422	0.4283	0.03	0.03
triceps (long head)	1000	0.1905	0.0877	0.877	0.0877	0.1916	0.03	0.03
triceps (short head)	700	0.1905	0.0877	0.877	0.0877	0.2387	0	0.03
brachialis	700	0.0175	0.1028	1.028	0.1028	0.1681	0	0.03

els made before caregiver corrections), we followed the procedure presented in [19]. We assumed the robot to be a human arm with completely unknown dynamics. In such a case, we only had the ability to control muscle activation, as well as measure force/torque at the wrist, joint angles, and joint velocities.

Following the inverse statics modeling procedure in [19], we moved the arm to different states $\boldsymbol{q} = [\theta_1, \theta_2, \dot{\theta}_1, \dot{\theta}_2]$. For each configuration, there was no muscle activation applied, so the only forces acting on the joints were from the nonlinear springs in (1). We measured the force required to hold the arm static as a researcher held the force/torque sensor. We were only concerned with force acting in the plane $\boldsymbol{f} = [f_x, f_y]$. We averaged the state and force over each 0.5 seconds to make the data points. This joint state and endpoint force data were used as input-output pairs for the inverse statics model. Note that including joint velocity technically makes this an inverse dynamics model, but we use "inverse statics" to emphasize that we make predictions about the static behavior of the arm. Additionally, though the researcher attempted to hold the arm static in each configuration, there was no guarantee that the arm was perfectly static.

After collecting data for the inverse statics model, we collected data for the 6 muscle capability models. As with the entire modelling procedure, we assumed we did not have access to the muscles' underlying functions. We moved the arm to the same configurations $\boldsymbol{q}$ of the inverse statics model data. One by one, we fully activated the muscles ($\alpha = 1$) while a researcher applied force to the force/torque sensor at the arm's wrist to hold it static. Our goal was to model the portion of the muscle which scales with muscle activation in (1). To do this, we subtracted the inverse static data from the measured force, as detailed in [19]. These six sets of joint states $\boldsymbol{q}$ and end-effector forces $\boldsymbol{f}$ made up our data for the six muscle capability models.

These models served as our initial models. Their error was compared to the error of models which used the same data *and* caregiver correction data. Muscle model predictions were made with muscle-specific data and caregiver correction data using the equation

$$\begin{aligned}\boldsymbol{\mu}_i^* = & \begin{bmatrix} k(\boldsymbol{q}_i, \boldsymbol{q}^*) \\ -\alpha_i k(\boldsymbol{q}_i, \boldsymbol{q}^*) \end{bmatrix}^T \\ & \times \left(\begin{bmatrix} k(\boldsymbol{q}_i, \boldsymbol{q}_i) & -\alpha_i k(\boldsymbol{q}_i, \boldsymbol{q}_c) \\ -\alpha_i k(\boldsymbol{q}_c, \boldsymbol{q}_i) & \alpha_i^2 k(\boldsymbol{q}_c, \boldsymbol{q}_c) \end{bmatrix} + \sigma_n^2 I \right)^{-1} \begin{bmatrix} \boldsymbol{f}_i \\ \boldsymbol{f}_c \end{bmatrix}\end{aligned} \tag{4}$$

which was derived in [17]. This equation gave the predicted maximum capability $\boldsymbol{\mu}_i^*$ of the i^{th} muscle at a desired state $\boldsymbol{q}^*$. The covariance k of the functions were compared at both the researcher-collected initial input $\boldsymbol{q}_i$ and the joint state when caregivers applied corrections $\boldsymbol{q}_c$. The covariance at the corrective input was multiplied by the activation the i^{th} muscle received during the correction. The covariance determined the similarity of the functions at the desired state to previous measured states. Along with the estimated signal noise σ_n^2, this equation the determined how similar the prediction is to the researcher-collected initial force data point $\boldsymbol{f}_i$ or the force applied by the caregiver $\boldsymbol{f}_c$ to correct the robot's arm.

D. Experimental Analysis

To accomplish the goal of measuring the accuracy of human caregiver corrections, we calculated the error of the corrections by first calculating the force that should have been applied to the robot to perfectly ensure it was held static. This was done by solving for the residual force when the muscle activations were used in the governing muscle equations (1). After finding this "theoretically perfect correction," we found the norm error in the corrections applied by the caregivers. We plotted the deviation from a perfect correction versus the force required for the same perfect correction. We then fit a trend line to the scatter plot to determine if there was a correlation between the force required to hold the arm static and the error in the correction from the caregiver.

After calculating the accuracy of the correction data, we used the data to generate new model predictions. We generated predictions over the arm's range of motion using models with no corrective data, as well as models with corrective data mixed with subsystem specific data. We varied the amount of corrective data, using either 50, 100, 200, 300, 400, 500, or 672 points. For each corrective data set size, we randomly selected from the total 672 points. There were a total of 672 points as there were 8 caregivers x 21 trials x 4 data points per trial (2 second trials averaged every 0.5 seconds). We repeated the selection and prediction calculation 4096 times for each data set size, assuming different permutations of points would lead to different predictions. The total number of unique combinations for some data set sizes is on the order of 10^{199}, so we used 4096 repeated trials with different, randomly selected data sets.

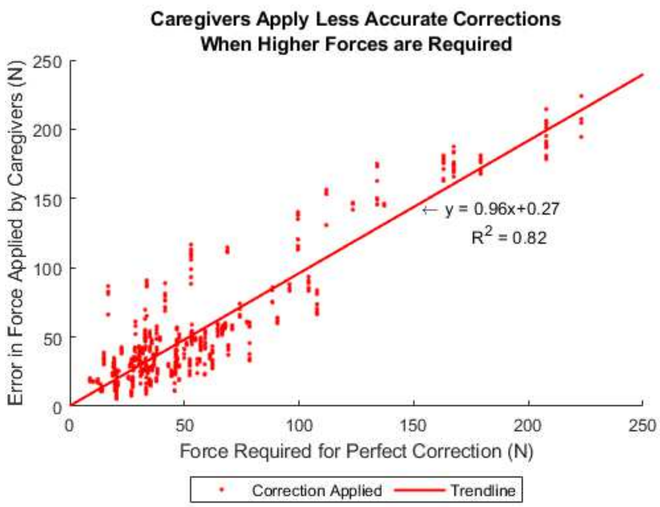

Fig. 4. This scatter plot shows the data collected from the caregivers' corrections. The horizontal axis shows the magnitude of the force required for a perfect correction. The vertical axis shows the magnitude of the error; data sitting exactly on the horizontal axis indicates a perfect correction.

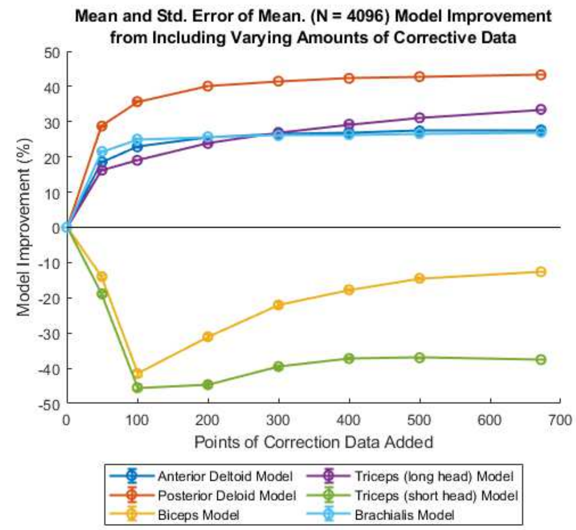

Fig. 5. This figure shows the muscle model improvement using different amounts of data. There were 4096 repeated trials for each amount of data, so the standard errors of the means are small enough to be obscured by the data markers.

We compared both predictions made without and with corrective data to the ground truth given by the governing muscle equations (1), and calculated how much the models improved with the addition of corrective data. Model improvement is the error reduction from including corrective data normalized by the error before including corrective data.

III. RESULTS

We asked several caregivers to hold a robotic arm static while its motors generated torque. The torque was calculated according to 6 muscle functions; as we knew the ground truth of the arm's dynamics, we could calculate how much force the caregivers should have been applying to the robot to hold it perfectly static. Comparing the force that the caregivers applied to the arm to the calculated "theoretically perfect correction" allowed us to calculate the error of the caregivers' applied corrections. We plotted the error magnitude versus the force magnitude of the perfect correction to show accuracy in Fig. 4. Recall each trial was 2 seconds and each data point is made of the average of 0.5 seconds, so there are 4 data points per trial. There are a total of 672 data points.

Fig. 4 shows that when a higher force is required for a correction, there is a higher error in the applied force. The trend line for the corrections has a slope close to 1 with $R^2 = 0.82$, illustrating that when a higher force is required for correction, there is a higher error in the force applied for the correction.

In addition to studying the ability or caregivers to accurately correct static arm holds, we used the data to improve model predictions of the arm's muscles. We compiled the prediction improvements, or the error reductions normalized by the errors before including the caregiver correction data, shown in Fig. 5. To elaborate, 100% improvement would indicate that a model was *exactly* perfect after including correction data.

As increasing amount of corrective data were used, four of the 6 muscle models improved; the posterior deltoid by 43% ±9.9e-13%, triceps (long head) by 33% ±4.1e-13%, the anterior deltoid by 28% ±3.8e-13%, and the brachialis by 27% ±4.0e-13%. Two muscle models, the biceps and the short head of the triceps, have an accuracy decrease with small amount of correction data, but the accuracy decrease lessens with increasing amounts of correction data. The standard errors of the means are very small, indicating that there is little deviation across the 4096 permutations of corrective data sets.

IV. DISCUSSION

The goals of this study were to measure the accuracy of human caregiver corrections to a robotic arm and determine the effects of human-generated corrections in making model predictions. The robotic arm was actuated in a way that simulated electrically stimulated human arm muscles. The results indicate that non-expert caregivers have the ability to make arm corrections that can improve the system identification accuracy, but also make the accuracy worse, as was the case with the biceps and triceps (short head) models.

When the caregivers were applying corrections to the arm, higher errors occurred when higher forces were required. This is likely because the muscles immediately turned on and generated a strong force, causing a jerking motion. The human caregivers then had to react to the jerking motion to correct the arm. To contrast, when a small force was required to correct the static hold of the arm, caregivers could much more easily apply the force immediately, without having to recover from a strong jerk. As each 2 second trial was used to make 4 data points, the data appear in vertical groups of 4. Each of these 4 points have the same force required for a perfect correction, but different error as the caregivers changed the force they applied over time.

An unaccounted for source of error for these corrections is potential bias introduced by the plastic balls, shown in Fig. 1, used to mark goal end-effector positions. Slight inaccuracies in their position, as well as the challenge of determining the downward vertical projection of the center of the sphere likely contributed to some of the caregiver's error in applying accurate correction data.

When using the corrective data to generate new predictions, four of the 6 muscle models improved. While not immediately obvious given the governing muscle equations (1), these are the muscles with the highest endpoint force in the robot's range of motion. In this range of motion, these four muscles are able to produce endpoint forces on the order of 10^2 Newtons. To contrast, the biceps and short head of the triceps are weak as a result of their large a_0 relative to L_{opt} from Table I. In the robot's range of motion, the force-length curve (2) is small; the biceps and triceps (short head) maximum endpoint forces are ~ 4 N and ~ 7.5 N, respectively. These values were from [12], and were not changed in consideration for the robotic anatomy.

The results of this study indicate that corrective data works better for stronger muscles. The corrective data is generated when all muscles are being used, and can be used to produce *all* model predictions from a single source. This is easier for the stronger muscles, as the signal to noise ratio is much larger; [17] details how whole-system data is learned from output differences, and the differences in endpoint force from the weak muscles is minor. To learn the small ground truth value of the weaker muscles requires more data, as the force contributions are minor. This trend is most evident in the biceps model in Fig. 5. With small amounts of this noisy data, the biceps model becomes 40% worse. As more data is added, the model begins to learn the small force contribution of the biceps. We expect that if more data were available, the biceps and triceps (short head) models would recover and eventually have positive improvement.

This study gives important insight how corrective data should practically be used. While we expect all models to improve with enough data, it is not a realistic solution to improve the system identification of muscles. The arm's dynamics and muscle production capabilities change with time. A practical implementation of the work presented here should be used with an evolving Gaussian Process algorithm [20], or a modelling technique where an optimal subset of data is selected based on the most informative data. This would reject the noisy data for the weaker muscles while still leading to improvement for the muscles that primarily move the electrically stimulated arm.

Accurate measurements of arm states were critical in this experiment. If we had not used joint velocity as an input to the models, there likely would not have been any model improvement with the addition of caregiver correction data. Joint velocity as a model input allowed the movement from caregiver's inaccurate corrections provide information on the dynamic behavior of the muscles, which led to improvement. To translate this modelling technique to home use, there should be a way of measuring human joint positions and velocities to account for caregiver correction error.

References

[1] NSCISC, "Traumatic spinal cord injury facts and figures at a glance," 2023.

[2] K. D. Anderson, "Targeting recovery: priorities of the spinal cord-injured population," *J Neurotrauma*, vol. 21, no. 10, pp. 1371–1383, Oct. 2004.

[3] M. Dijkers, "Quality of life after spinal cord injury: a meta analysis of the effects of disablement components," *Spinal Cord*, vol. 35, no. 12, pp. 829–840, Dec. 1997.

[4] M. Kutlu, C. Freeman, A.-M. Hughes, and M. Spraggs, "A Home-based FES System for Upper-limb Stroke Rehabilitation with Iterative Learning Control," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 12 089–12 094, Jul. 2017.

[5] P. H. Peckham and J. S. Knutson, "Functional Electrical Stimulation for Neuromuscular Applications," *Annual Review of Biomedical Engineering*, vol. 7, no. 1, pp. 327–360, 2005, _eprint: https://doi.org/10.1146/annurev.bioeng.6.040803.140103.

[6] W. D. Memberg, K. H. Polasek, R. L. Hart, A. M. Bryden, K. L. Kilgore, G. A. Nemunaitis, H. A. Hoyen, M. W. Keith, and R. F. Kirsch, "Implanted Neuroprosthesis for Restoring Arm and Hand Function in People With High Level Tetraplegia," *Archives of Physical Medicine and Rehabilitation*, vol. 95, no. 6, pp. 1201–1211.e1, Jun. 2014.

[7] S. Trier, J. Buckett, A. Campean, M. Miller, F. Montague, T. Vrabec, and J. Weisgarber, "A Modular External Control Unit for Functional Electrical Stimulation," Jan. 2001.

[8] K. L. Kilgore, H. A. Hoyen, A. M. Bryden, R. L. Hart, M. W. Keith, and P. H. Peckham, "An Implanted Upper-Extremity Neuroprosthesis Using Myoelectric Control," *J Hand Surg Am*, vol. 33, no. 4, pp. 539–550, Apr. 2008.

[9] D. Landin, J. Myers, M. Thompson, R. Castle, and J. Porter, "The role of the biceps brachii in shoulder elevation," *J Electromyogr Kinesiol*, vol. 18, no. 2, pp. 270–275, Apr. 2008.

[10] C. L. Lynch and M. R. Popovic, "Functional Electrical Stimulation," *IEEE Control Systems Magazine*, vol. 28, no. 2, pp. 40–50, Apr. 2008.

[11] V. K. Mushahwar, P. L. Jacobs, R. A. Normann, R. J. Triolo, and N. Kleitman, "New functional electrical stimulation approaches to standing and walking," *J Neural Eng*, vol. 4, no. 3, pp. S181–197, Sep. 2007.

[12] K. M. Jagodnik and A. J. Van Den Bogert, "Optimization and evaluation of a proportional derivative controller for planar arm movement," *Journal of Biomechanics*, vol. 43, no. 6, pp. 1086–1091, Apr. 2010.

[13] D. C. Crowder, J. Abreu, and R. F. Kirsch, "Improving the Learning Rate, Accuracy, and Workspace of Reinforcement Learning Controllers for a Musculoskeletal Model of the Human Arm," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 30, pp. 30–39, 2022.

[14] D. N. Wolf and E. M. Schearer, "Holding Static Arm Configurations With Functional Electrical Stimulation: A Case Study," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 26, no. 10, pp. 2044–2052, Oct. 2018.

[15] C. E. Rasmussen and C. K. I. Williams, *Gaussian processes for machine learning*, 3rd ed., ser. Adaptive computation and machine learning. Cambridge, Mass.: MIT Press, 2008.

[16] D. N. Wolf and E. M. Schearer, "Trajectory Optimization and Model Predictive Control for Functional Electrical Stimulation-Controlled Reaching," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, Apr. 2022.

[17] C. J. LaMack and E. M. Schearer, "Global system errors to simultaneously improve the identification of subsystems with mixed data gaussian process regression," *Machine Learning: Science and Technology*, vol. 5, no. 2, May 2024.

[18] F. E. Zajac, "Muscle and tendon: properties, models, scaling, and application to biomechanics and motor control," *Crit Rev Biomed Eng*, vol. 17, no. 4, pp. 359–411, 1989.

[19] E. M. Schearer, Y.-W. Liao, E. J. Perreault, M. C. Tresch, W. D. Memberg, R. F. Kirsch, and K. M. Lynch, "Semiparametric Identification of Human Arm Dynamics for Flexible Control of a Functional Electrical Stimulation Neuroprosthesis," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 24, no. 12, pp. 1405–1415, Dec. 2016.

[20] D. Petelin and J. Kocijan, "Control system with evolving gaussian process models," in *2011 IEEE Workshop on Evolving and Adaptive Intelligent Systems (EAIS)*, 2011, pp. 178–184.