

Model Learning for Control of a Paralyzed Human Arm with Functional Electrical Stimulation

Derek N. Wolf¹, *Student Member, IEEE*, Zinnia A. Hall², and Eric M. Scheerer¹, *Member, IEEE*

Abstract—Functional electrical stimulation (FES) is a promising technique for restoring reaching ability to individuals with tetraplegia. To this point, the complexities of goal-directed reaching motions and the shoulder-arm complex have prevented the realization of this potential in full-arm 3D reaching tasks. We trained a Gaussian process regression model to form the basis of a feedforward-feedback control structure capable of achieving reaching motions with a paralyzed upper limb. Over a series of 95 reaches of at least 10 cm in length, the controller achieved an average accuracy (measured by the Euclidean distance of the wrist to the final target position) of 3.8 cm and an average error along the path of 3.5 cm. This controller is the first demonstration of an accurate, complete-arm, FES-driven 3D reaching controller to be implemented with an individual with tetraplegia.

I. INTRODUCTION

For the approximately 170,000 individuals with some level of tetraplegia due to spinal cord injuries [1], the restoration of hand and arm function is their greatest priority to improving their quality of life [2]. Functional electrical stimulation (FES) is a promising technique for helping these individuals complete the reaching motions necessary for daily living.

Many approaches have been attempted for achieving arm function with FES by first reducing the complexity of the arm-control problem. For repetitive tasks such as grasping, the complexity of the system was reduced by using preprogrammed, repeated stimulation patterns [3]. The extension of the repeated stimulation pattern method to full-arm reaching [4] cannot achieve all daily reaching tasks because the ever changing, goal-directed nature of reaching motions would require an infeasible amount of predetermined stimulation patterns. Any everyday reaching controller must be able to automatically select the stimulation commands necessary to achieve any novel, feasible reach.

Another common approach to reducing the complexity of controlling reaching motions is to control each joint independently. FES controllers have demonstrated success in controlling individual joints such as elbow extension [5]. Extending this success to controlling multiple joints (separately), the MUNDUS program [6] used a lockable exoskeleton to lock all uncontrolled joints while a single joint was driven with FES. The current state-of-the-art FES-reaching system, the

BrainGate2 clinical trials, controls each joint independently but simultaneously [7]. Based on user intent (read by an intracortical brain computer interface), the controller selected a position along a predefined stimulation pattern for each of the joints. The controller had difficulties with multi-joint motions because the independent joint control could not account for joint interactions. The independent joint control method also does not allow for using the kinematic redundancy of the arm to complete tasks in different ways. In order to successfully control reaching motions, it is necessary to treat the arm as a complete system and not as independent joints.

Various methods of controlling the complete arm have been attempted. In computer simulations, optimized proportional-derivative control [8], combined feedforward-feedback control [9], reinforcement learning [10], and threshold control [11] have all successfully controlled reaching motions. Practically implementing these methods is difficult due to the real-world arm dynamics differing from the simulation.

Model learning, which uses data-driven machine learning models rather than parameterized physics-based models to predict the behavior of physical systems, has been used extensively to control robots (see [12] for a review) and is especially suitable for using FES to control the human arm. We intend to control multiple joints with multiple muscles – a problem that grows significantly in complexity as more joints and muscles are added. Using a physics-based model for FES control as in [13] may be effective for single joint systems with one or two muscles. However, as system complexity increases, the number of parameters needed to accurately model the arm increases. Further, guaranteeing parameter identifiability (e.g. of joint inertias) is extremely difficult given the limitations on range of motion and acceptable movement speeds for people with spinal cord injuries. Although model learning (with artificial neural networks) has been used for FES control of planar reaching in healthy persons [14], it has not been demonstrated for 3D motions in people with spinal cord injuries.

In our own prior research, we have used model learning approaches to complete steps towards full-arm reaching. We have used semiparametric Gaussian process regression (GPR) to predict joint torques produced by muscles [15]. We built on this success by using nonparametric GPR models of the arm to form the basis of a feedforward-feedback controller to hold static wrist positions [16]. With this controller, we attempted quasi-static reaches with no intermediate points (used the model of the final position as the model for the entire reach) with some success, but there was significant oscillation and large overshoot in the reaching error [17]. In simulation, we showed that adding external damping and quasi-static intermediate points improved the controller performance [18].

This work was supported by NIH NINDS grant N01-NS-5-2365, Veteran's Affairs grant B2359-C, NSF grant 1751821, NSF grant 1659559, and the Cleveland State University Graduate Student Research Award.

¹D. N. Wolf and E. M. Scheerer are with the Department of Mechanical Engineering and the Center for Human-Machine Systems, Cleveland State University and the Cleveland Functional Electrical Stimulation Center, Cleveland, OH USA d.n.wolf@vikes.csuohio.edu, e.scheerer@csuohio.edu.

²Z. A. Hall is with the University of Connecticut, Storrs, CT USA zinnia.hall@uconn.edu.

The purpose of the current study is to build upon our previous work and develop a control structure capable of achieving full-arm, 3D reaching motions driven by FES. This is an important step towards the use of FES in the home to restore the full-arm reaching motions critical to completing many activities of daily living. We present a method of developing a subject-specific model of an individual with tetraplegia’s arm and its response to electrical stimulation. We use this model as the basis of a combined feedforward-feedback controller capable of automatically determining the stimulation commands necessary to achieve desired reaching motions within the subject’s workspace.

We evaluated the performance of the controller for completing reaching tasks. In particular we quantified the accuracy of the controller for moving the wrist to a desired final position and determined if there was a difference in accuracy based on target location. These results will guide the future developments of FES-driven reaching controllers.

II. METHODS

In this study, we used a model learning based control strategy to complete reaching motions with an individual with high tetraplegia and an implanted FES neuroprosthesis. During the experiment, we 1) developed a Gaussian process regression model for the force the muscles produce as a function of the wrist position, and 2) used the model as the basis of an FES controller to move the wrist along desired paths.

The experiment took place during a four-hour time block. Experimental set up and identifying the model of the arm required approximately 1.5 hours. The participant took a half-hour break for lunch. The remaining time of the session was used to attempt randomly selected reaching motions. The participant was allowed breaks whenever requested.

A. Experimental Setup

We completed the experiments with a single human participant who has high tetraplegia and lacks voluntary control of her right arm. The participant’s abdomen is implanted with a stimulator-telemeter [19, 20, 21] that can deliver current to activate nine independent muscle groups: triceps, deltoids, latissimus dorsi, serratus anterior, biceps/brachialis, supra/infraspinatus, rhomboids, lower pectoralis, and upper pectoralis. Muscle stimulation is delivered via bi-phasic, charge balanced pulses delivered at 13 Hz. The amplitude of the pulses is constant for each muscle group. The activation of each muscle group is controlled by varying the pulse-width (referred to as the stimulation input) from 0-250 μ s. The participant’s wheelchair is equipped with a passive arm support that assists against the force of gravity to create a comfortable and achievable workspace. More details can be found in [22] (Subject 1) and [16]. Protocols used for this research were approved by the institutional review boards at Cleveland State University (IRB NO. 30213-SCH-HS) and MetroHealth Medical Center (IRB NO. 04-00014).

We gathered training data for using a HapticMaster (Moog FCS) robot with three degrees of freedom. The robot records the 3D forces and positions of its end-effector. An Optotrak

Certus Motion Capture System (Northern Digital, Inc.) captured data used to measure the position of the wrist for data gathering for modeling and feedback during reaching.

Our previous research (and that of others) has demonstrated that significant oscillations occur with feedback FES control due to the delays in the FES system (low frequency of stimulation and electrical-muscular activation delays). In simulation, we were able to improve the controller performance by adding physical damping to the arm support [18]. Due to this finding, we used the robot to create a damped environment (20 N-s/m in each direction) during the reaching experiments.

The control and data collection occurred at 52 Hz, but stimulation inputs were updated at 13 Hz.

B. Model Learning

Our model learning procedure was previously presented in [16]. We present a detailed summary here for completeness. The basis of our controller (shown in Fig. 1) was a model consisting of three parts: 1. arm statics (predicts the forces necessary to hold a desired wrist position), 2. muscle force production (the mapping from wrist position to the maximum forces produced at the wrist by each muscle), and 3. recruitment curves (the mapping from muscle group electrical stimulation to muscle activation). Our controller inverts each part of the model to determine the muscle group stimulation commands necessary to achieve a desired wrist position.

To gather the model training data, a robot held the participant’s wrist at a series of static positions within her comfortably reachable workspace. The connection of the participant’s wrist to the robot was via a ball-in-socket joint that does not transmit torque. The robot was equipped with a three-dimensional force sensor at its end-effector, and the force needed to hold the wrist stationary, $\mathbf{f}_r \in \mathbb{R}^3$, was recorded.

To determine the arm statics, the robot held the arm in a position with zero muscle stimulation, and, thus, all muscle activations, $\alpha \in \mathbb{R}^9$, were zero. Therefore,

$$\mathbf{f}_{r\text{static}} = \mathbf{p}(\mathbf{x}) \quad (1)$$

where $\mathbf{p}(\mathbf{x}) \in \mathbb{R}^3$ are the forces necessary to hold the wrist in the static position, $\mathbf{x} \in \mathbb{R}^3$.

To determine the force production of the j^{th} muscle group, the muscle group was stimulated at its maximum stimulation command so that $\alpha(j) = 1$. The forces required for the robot to hold the wrist stationary, $\mathbf{f}_{r\text{stim}j}$, are then defined by the difference of the forces with zero stimulation (i.e. the required static forces) and the forces produced by the muscle group,

$$\mathbf{f}_{r\text{stim}} = \mathbf{p}(\mathbf{x}) - \mathbf{M}(\mathbf{x})\alpha, \quad (2)$$

where $\mathbf{M}(\mathbf{x}) \in \mathbb{R}^{3 \times 9}$ is the linear mapping of muscle activation to forces at the wrist and $\mathbf{p}(\mathbf{x})$ are the forces when stimulating no muscles. Each column of $\mathbf{M}(\mathbf{x})$ represents the forces produced in each Cartesian direction by 100% activation of the corresponding muscle group. The j^{th} column of $\mathbf{M}(\mathbf{x})$ is determined by subtracting $\mathbf{f}_{r\text{stim}j}$, the recorded force during stimulation of muscle group j , from the previously identified static forces, $\mathbf{f}_{r\text{static}}$,

$$\mathbf{M}(\mathbf{x})_j = \mathbf{f}_{r\text{static}} - \mathbf{f}_{r\text{stim}j}. \quad (3)$$

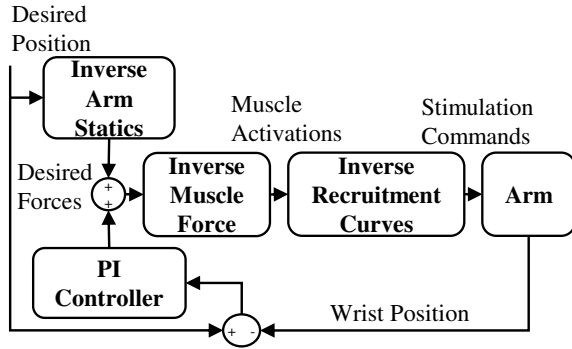


Fig. 1. Controller block diagram. The controller determines the muscle stimulation commands necessary to achieve a desired wrist position.

At a previous experimental session, we gathered training data by measuring $\mathbf{f}_{r_{static}}$ and $\mathbf{f}_{r_{stim}}$ for all muscle groups at 27 wrist positions, \mathbf{x} , within the participant’s workspace. We repeated the set of measurements three times with a random order of wrist positions and muscle group activations. The data was used to train a set of GPR models with the input being the wrist position and the output being the force in one direction measured by the robot. A model was trained for the force produced in each degree of freedom for each of the nine muscle groups as well as for the static arm (zero muscle activation) resulting in 30 total models. The squared exponential covariance function with automatic relevance detection was used, and the optimal hyperparameters were identified [23]. Using the GPR models, we can calculate $\mathbf{f}_{r_{static}}$ and $\mathbf{f}_{r_{stim}}$ and thus determine $\mathbf{p}(\mathbf{x})$ and $\mathbf{M}(\mathbf{x})$, via (3), at any position in the subject’s workspace. In the controller (Fig. 1), the GPR models form the basis of the “Inverse Arm Statics” and “Inverse Muscle Force” blocks.

To account for changes in the arm and muscles from the time of modeling, we gathered new training data from a single set of 27 wrist positions at the start of the reaching experiments session. We reused the previously found hyperparameters to train a new model with less data and thus less time. This method allows for identifying an accurate model the day of the experiment which is critical for accurate control.

The recruitment curves, the mapping from stimulation input to muscle group activation, for each muscle group were identified using the deconvolved ramp method [24].

C. Controller

The controller automatically determines the muscle stimulation commands necessary to move to a desired wrist position. The controller (Fig. 1) uses the model presented in section II-B along with feedback to determine the forces and corresponding muscle group activations necessary to achieve a desired wrist position. The stimulation inputs are then determined and applied to the arm.

The input to the controller is the desired wrist position, $\mathbf{x}_* \in \mathbb{R}^3$. The controller calculates the open-loop forces, $\mathbf{p}(\mathbf{x}_*)$, necessary to hold a desired position by using the GPR models of the inverse arm statics. Feedback is added using a positional proportional-integral controller which adjusts the

open-loop forces with corrective forces in each Cartesian direction to get the desired forces, \mathbf{f}_{des} . Next, the controller uses the GPR models of muscle force production and (3) to identify the elements of the mapping from muscle group activations to wrist forces, $\mathbf{M}(\mathbf{x}_*)$. After determining the desired forces and the muscle-force mapping, $\mathbf{M}(\mathbf{x}_*)$, we calculate the muscle activations, α which will produce the desired forces.

Determining the desired muscle activations during real-time feedback control requires overcoming two main problems at this point: 1) the arm is a redundant system in that there are more muscle groups than degrees of freedom (i.e. $\mathbf{M}(\mathbf{x}_*)$ is not square), and 2) the use of feedback means we have no control over the forces that the controller calls for and thus feedback overcompensation, the calling for forces above the greatest possible force, and muscle activation saturation can occur. Solving the system redundancy can traditionally be completed using a constrained optimization routine such as a quadratic programming routine. However, in the case of feedback overcompensation (which can happen with little feedback compensation in an individual with tetraplegia due to muscle weakness from atrophy), constrained optimizations are unable to find a feasible solution (one where the muscle forces are between zero and one) because one does not exist.

Our solution to these problems is to use the quasi-Newton method to find the α that minimizes the penalty function,

$$\begin{aligned} & \|\alpha\|_2^2 + c_1 \|\mathbf{M}(\mathbf{x}_*)\alpha - \mathbf{f}_{des}\|_2^2 + c_2 \mathbf{K} + c_3 \mathbf{T} \\ & K = \sum k_i \text{ where } k_i = \begin{cases} \alpha_i^2 & \text{if } \alpha_i < 0 \\ (\alpha_i - 1)^2 & \text{if } \alpha_i > 1 \\ 0 & \text{if } 0 \leq \alpha_i \leq 1 \end{cases}, \quad (4) \\ & T = \mathbf{f}_{des} \times \mathbf{M}(\mathbf{x}_*)\alpha \end{aligned}$$

where $\|\alpha\|_2^2$ minimizes the muscle activations, $\|\mathbf{M}(\mathbf{x}_*)\alpha - \mathbf{f}_{des}\|_2^2$ penalizes activations that do not produce the desired force, K penalizes activations which do not belong to $\alpha_i \in [0, 1]$, and T penalizes activations that produce forces in an incorrect direction. The penalty weights were chosen to be $c_1 = 100$, $c_2 = 10,000$ and $c_3 = 1,000$ because they produced feasible muscle activations with the forces in the right direction during offline testing.

When feedback overcompensation occurs, our controller aims to produce the most possible force in the direction of the desired force. The T term of the objective function works to make this happen by penalizing forces not in the desired direction. When the forces become significantly larger than the maximum possible forces in a given wrist position, the penalty function solution breaks down and can lead to solutions which do not make sense (activations greater than one or less than zero) or to an inability to find a solution.

To improve these solutions, we developed a method to restrict the forces to within a rectangular prism of the maximum forces that can be produced in each direction. The maximum force that can be produced in each of the Cartesian directions (positive and negative directions) is determined and recorded offline. In the force space, a rectangular prism is drawn with faces at each of the maximum forces. If the desired force is greater than the maximum force that can be produced in any of the Cartesian directions, the intersection of the force vector

and the rectangular prism is found, and this point becomes the new desired force. This scales the desired force back to a position closer to the feasible force space while maintaining the desired direction. While this point is not guaranteed to be feasible, this scaling resulted in more reasonable activations as a solution to (4).

Once the activations are found, the inverse recruitment curves block calculates the stimulation inputs.

D. Reaching Experiments and Data Analysis

To evaluate our controller’s ability to control reaching motions, we quantified the accuracy of the controller over a series of reaches throughout the participant’s workspace. Each reaching trial lasted for five seconds and consisted of a one-second hold at the starting position, a two-second ramp from the starting position to the target position, and a two-second hold at the target position. The straight-line ramp between positions was selected because previous quasi-static experiments showed that planning a defined path of closely spaced points between the target and goal would improve the performance (speed, smoothness, and accuracy of movement) [17, 18], and a ramp is the limit of lowest spacing and time between each quasi-static point (we use the term quasi-static because the controller is based on a static model but is used to create reaching movement). An example of a desired trajectory can be seen in Fig. 2.

Prior to completing the reaching experiments, the PI controller was tuned. The controller was tuned using a series of 3 random reaching motions. The gains were manually tuned to improve the final accuracy without increasing oscillations. The values of the proportional and integral gains were selected to be 10 N/mm and 0.3 N/mm-s respectively.

During the reaching trials, the subject’s wrist was connected to a robot which moved the subject’s wrist to the starting position for each trial and created a damped environment during the reach. At the start of each trial, to limit the effects of the transient muscle dynamics and guarantee the controller starts at the correct point, the wrist was held stationary for the first 0.5 seconds. For the rest of the reach, the wrist was allowed to move as driven by the muscle stimulation.

To select the target reaching motions, we created a grid of wrist positions with 1 cm spacing within the convex hull of the 27 positions visited during the gathering of the model training data. Start and target positions were randomly selected from this wrist position grid to create reaches of at least 10 cm in length. The average reach length was 13 cm. Once the start and target positions were selected, the complete desired reaching path was determined. For each wrist position along the path, \mathbf{x}_* , the open-loop muscle forces, $\mathbf{p}(\mathbf{x}_*)$, and the muscle force production matrix, $\mathbf{M}(\mathbf{x}_*)$ were determined offline before the trials. At each time step, the controller used $\mathbf{p}(\mathbf{x}_*)$ and $\mathbf{M}(\mathbf{x}_*)$ for the current desired position.

For each reach, the average wrist position over the final 0.5 seconds was recorded. The final accuracy of the reach was determined by the Euclidean distance between this average final wrist position and the desired target position. We also measured the accuracy over a complete reaching motion which

we refer to as the path accuracy. The path accuracy for a single trial is defined as the average Euclidean distance from the wrist position to the desired target position over all time steps. As part of the analysis, a few factors were compared to determine their effect on the controller performance. These factors included the position of the target and whether the selected path had a feasible target position. Feasible target positions are defined as a position where the model can select muscle activations capable of achieving the predicted open-loop forces. In an attempt to complete more possible reaches, non-feasible and feasible target positions were selected. We completed as many unique reaches as possible in the allotted time (95 total reaches). A 2-sample t-test was used to determine if these factors had an effect on the controller.

III. RESULTS

Over 95 trials, our controller achieved reaching motions with an average final accuracy of 3.8 cm (standard deviation of 2.2 cm) and an average path accuracy of 3.5 cm (standard deviation of 1.5 cm). Figures 2(a) and (b) show a representative reach with a final accuracy of 2.1 cm and a path accuracy of 1.9 cm. As seen, the wrist position is able to track the desired reaching path and finish near the desired position. Figure 2(c) shows the muscle group activations that the controller calculated during the trial. The activation of the deltoids muscle group (bold) shows a clear example of the feedback controller correctly adjusting the muscle activation. The deltoids muscle group produces a force at the wrist to the subject’s right (positive y direction). As seen, when the wrist position is to the left of the desired position (lesser y position), the activation of the deltoids increases to drive the wrist to the right. The deltoids activation decreases when the wrist moves to the right of the desired position. This example and the overall accuracy results demonstrate that the controller successfully calculated the muscle stimulation commands required to achieve the desired reaching motions.

Figure 2 also shows that though there were some oscillations, the amount was limited and generally low frequency which could be tolerated by the subject and would still be useful in a functional task. This experimentally demonstrates the efficacy of using a proportional-integral controller to produce reaching motions within a damped environment.

There was a significant difference ($p < 0.001$) in the final accuracy of reaching motions with feasible final target positions ($\mu = 2.1$ cm, $N=22$) and the accuracy of reaches with infeasible final target positions ($\mu = 4.3$ cm, $N=73$). There was also a significant difference ($p < 0.01$) in the final accuracy of reaching motions to the extreme right of the subject’s workspace, defined by the target being greater than 5 cm to the right of the center of the subject’s thorax ($\mu = 4.9$ cm, $N=25$) and the rest of the workspace ($\mu = 3.4$ cm, $N=70$). This difference in accuracy based on position can be seen in Fig. 3.

IV. DISCUSSION

We have used model learning to develop a controller capable of achieving arbitrary reaching motions with an FES-controlled paralyzed arm. Our controller accurately moves the

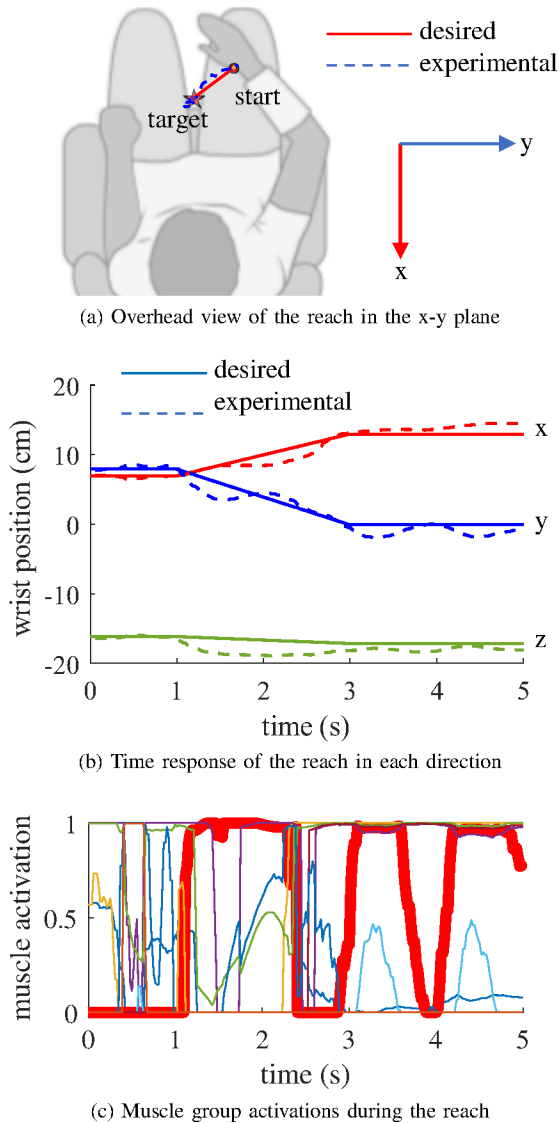


Fig. 2. This figure shows the details of an example 10 cm reach with final accuracy of 2.1 cm and path accuracy of 1.9 cm. (a) shows the overhead view in the x-y plane. (b) shows the time history of the reach in each Cartesian direction (dashed lines). As seen, the wrist was able to track the desired wrist position (solid lines). (c) shows the muscle activations which are automatically selected by the controller to track the desired path. As an example, the deltoids muscle group (highlighted) is activated to move the arm to the right (positive y) when necessary and turned off when the wrist needed to move left.

wrist along a desired path while accounting for the issues of a redundant system and feedback overcompensation/muscle activation saturation due to muscle weakness in individuals with spinal cord injuries. To our knowledge, this is the first demonstration of autonomously-selected electrical muscle stimulation for integrated shoulder and elbow control to produce 3D reaching movements in an individual with tetraplegia.

The controller accurately completed reaches throughout the participant's workspace. The final accuracy of 3.8 cm is sufficient to complete many daily reaching tasks such as grabbing a cup to drink. For points that our model predicted were feasible, the accuracy of 2.1 cm is, to our knowledge, the best reported 3D reaching accuracy achieved by FES. The improvement in accuracy for feasible points over the accuracy of infeasible



Fig. 3. This figure shows the target position and relative accuracy (represented by the size and color of each point) for all completed reaches. As seen, for targets to the right of the subject's workspace the accuracy is on average, worse than for the other target positions.

points is promising moving forward as the predicted feasible points can be used to select reaching paths that will have better overall accuracy (i.e. only traveling through/to predicted feasible points). The good overall accuracy for all points is also important as it is difficult to only choose paths that cross only feasible positions because of the limited workspace for individuals with tetraplegia.

The achieved path accuracy of 3.5 cm shows that the controller is able to accurately track a desired wrist path. This is important because when completing reaching tasks it is necessary to be able to reach a desired final hand position via differing paths. For example, when reaching out to pick up a fork off a table, the person may need to avoid bumping a cup of water with the hand. With the achieved path accuracy, our controller has the potential to achieve desired target positions while traveling along different paths.

Our accuracy was similar to the accuracy found in [14] of approximately 2 cm for planar reaching motions with healthy subjects as well as to our previous work with holding static wrist positions with an accuracy of 2.9 cm. This is encouraging because maintaining a similar accuracy while expanding to 3D reaching motions, controlling more muscles, and working with an individual with tetraplegia is critical if our controller is to restore everyday reaching motions outside of a laboratory environment.

Reaches to the extreme right hand side of the subject's workspace were less accurate than those towards the middle of the workspace. It was observed during trials to the right positions that the wrist would seem to reach a "sticking point" when trying to move out to the right. Reaching these targets often involved moving near the boundary of the workspace and sometimes a straight line to the target (as would be called for by the feedback controller) could pass through an unreachable space. Therefore, more advanced trajectory selection that guides the wrist through only reachable points may improve the performance to these extreme targets.

One major issue with any model-based FES controller is the changing muscle dynamics due to fatigue and atrophy. It is difficult to ensure that the model remains accurate over time. The performance of our controller demonstrates that our new, faster modeling procedure of developing new training data while maintaining the hyperparameters from a previously trained model is a way to update the model to maintain accu-

racy. The complete modelling procedure takes approximately three hours to gather the data and train the models. The day-of-experiment update only requires 35 minutes. This increase in model learning efficiency will allow the model to be updated more frequently and the controller to maintain its performance over time. Additionally, the performance seen during this experiment validated our previously simulated result that a damped environment can improve the performance of an FES controller [18]. It is relatively simple to create a damped environment by adding physical damping to the arm support that individuals with tetraplegia often, due to muscular atrophy, require to assist against the force of gravity and create a functional workspace.

The performance of this controller is a positive step to using an FES-controlled arm to restore everyday reaching tasks to individuals with high tetraplegia. To complete the goal of completing all possible reaching tasks, the accuracy must be improved throughout the workspace. This could be done through better path planning or through robotic assistance. Robotic exoskeletons have been shown to work cooperatively with FES to improve the accuracy of control for walking [25] and elbow flexion movements [26]. The robotic exoskeleton could be also used to replace the subject's arm support and produce the necessary damped environment. Additionally, with an accurate low-level FES controller, a brain control interface (or other input device), such as that used in the BrainGate2 study [7], could be used to determine the desired reaching target. Our controller could then automatically complete the desired reach. The subject's intent during the BrainGate2 study was able to be decoded, but controlling each joint independently made the reaching motions difficult. Our controller could replace this low level independent joint control with a complete arm controller. Combined with these possible solutions, an accurate FES-reaching controller is a critical step to restoring the reaching ability to individuals with tetraplegia in the home.

REFERENCES

- [1] NSCISC, "Spinal Cord Injury Facts and Figures at a Glance," 2018.
- [2] K. D. Anderson, "Targeting Recovery: Priorities of the Spinal Cord-Injured Population," *Journal of Neurotrauma*, vol. 21, no. 10, pp. 1371–1383, oct 2004.
- [3] 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.
- [4] 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.
- [5] W. D. Memberg, P. E. Crago, and M. W. Keith, "Restoration of elbow extension via functional electrical stimulation in individuals with tetraplegia," *Journal of rehabilitation research and development*, vol. 40, no. 6, pp. 477–86, 2003.
- [6] A. Pedrocchi, S. Ferrante, E. Ambrosini, M. Gandolla, C. Casellato, T. Schauer, C. Klauer, J. Pascual, C. Vidaurre, M. Gföhler, W. Reichenfeller, J. Karner, S. Micera, A. Crema, F. Molteni, M. Rossini, G. Palumbo, E. Guanziroli, A. Jedlitschka, M. Hack, M. Bulgheroni, E. D'Amico, P. Schenk, S. Zwicker, A. Duschau-Wicke, J. Miseikis, L. Graber, and G. Ferrigno, "MUNDUS project: MULTImodal Neuroprosthesis for daily Upper limb Support," *Journal of NeuroEngineering and Rehabilitation*, vol. 10, no. 1, p. 66, 2013.
- [7] A. B. Ajiboye, F. R. Willett, D. R. Young, W. D. Memberg, B. A. Murphy, J. P. Miller, B. L. Walter, J. A. Sweet, H. A. Hoyen, M. W. Keith, P. H. Peckham, J. D. Simeral, J. P. Donoghue, L. R. Hochberg, and R. F. Kirsch, "Restoration of reaching and grasping movements through brain-controlled muscle stimulation in a person with tetraplegia: a proof-of-concept demonstration," *The Lancet*, vol. 389, no. 10081, pp. 1821–1830, may 2017.
- [8] 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.
- [9] D. Blana, R. F. R. Kirsch, and E. K. E. Chadwick, "Combined feedforward and feedback control of a redundant, nonlinear, dynamic musculoskeletal system," *Medical and Biological Engineering and Computing*, vol. 47, no. 5 SPEC. ISS., pp. 533–542, 2009.
- [10] K. M. Jagodnik, P. S. Thomas, A. J. Van Den Bogert, M. S. Branicky, and R. F. Kirsch, "Training an Actor-Critic Reinforcement Learning Controller for Arm Movement Using Human-Generated Rewards," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 25, no. 10, pp. 1892–1905, oct 2017.
- [11] L. Lan, K. Y. Zhu, and C. Y. Wen, "Design of a Threshold FES Control System for Arm Movement," *Journal of Mechanics in Medicine and Biology*, vol. 09, no. 04, pp. 449–479, dec 2009.
- [12] D. Nguyen-Tuong and J. Peters, "Model learning for robot control: a survey," *Cognitive processing*, vol. 12, no. 4, pp. 319–340, 2011.
- [13] C. T. Freeman, A.-M. Hughes, J. H. Burrigge, P. H. Chappell, P. L. Lewin, and E. Rogers, "A Model of the Upper Extremity Using FES for Stroke Rehabilitation," *Journal of Biomechanical Engineering*, vol. 131, no. 3, p. 031011, 2009.
- [14] R. S. Razavian, B. Ghannadi, N. Mehrabi, M. Charlet, and J. McPhee, "Feedback Control of Functional Electrical Stimulation for 2D Arm Reaching Movements," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. PP, no. c, p. 1, 2018.
- [15] E. M. Scheerer, 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.
- [16] D. N. Wolf and E. M. Scheerer, "Holding Static Arm Configurations With Functional Electrical Stimulation: A Case Study," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 26, no. 10, pp. 2044–2052, oct 2018.
- [17] —, "Simple quasi-static control of functional electrical stimulation-driven reaching motions," in *2019 9th International IEEE/EMBS Conference on Neural Engineering (NER)*. IEEE, 2019, pp. 211–214.
- [18] —, "Developing a quasi-static controller for a paralyzed human arm: A simulation study," in *2019 IEEE 16th International Conference on Rehabilitation Robotics (ICORR)*. IEEE, 2019, pp. 1153–1158.
- [19] B. Smith, P. H. Peckham, M. W. Keith, and D. D. Roscoe, "Stimulator for Versatile Control of Paralyzed Muscle," *IEEE Transactions on Biomedical Engineering*, vol. BME-34, no. 7, pp. 499–508, 1987.
- [20] B. Smith, Zhengnian Tang, M. Johnson, S. Pourmehdi, M. Gazdik, J. Buckett, and P. Peckham, "An externally powered, multichannel, implantable stimulator-telemeter for control of paralyzed muscle," *IEEE Transactions on Biomedical Engineering*, vol. 45, no. 4, pp. 463–475, apr 1998.
- [21] R. L. Hart, N. Bhadra, F. W. Montague, K. L. Kilgore, and P. H. Peckham, "Design and testing of an advanced implantable neuroprosthesis with myoelectric control," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 19, no. 1, pp. 45–53, 2011.
- [22] K. H. Polasek, H. A. Hoyen, M. W. Keith, R. F. Kirsch, and D. J. Tyler, "Stimulation stability and selectivity of chronically implanted multicontact nerve cuff electrodes in the human upper extremity," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 17, no. 5, pp. 428–437, oct 2009.
- [23] C. Rasmussen and C. Williams, *Gaussian processes for machine learning*, 2006.
- [24] W. K. Durfee and K. E. MacClean, "Methods for Estimating Isometric Recruitment Curves of Electrically Stimulated Muscle," *IEEE Transactions on Biomedical Engineering*, vol. 36, no. 7, pp. 654–667, 1989.
- [25] K. H. Ha, S. A. Murray, and M. Goldfarb, "An Approach for the Cooperative Control of FES with a Powered Exoskeleton during Level Walking for Persons with Paraplegia," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 24, no. 4, pp. 455–466, 2016.
- [26] D. Wolf, N. Dunkelberger, C. G. McDonald, K. Rudy, C. Beck, M. K. O'Malley, and E. Scheerer, "Combining functional electrical stimulation and a powered exoskeleton to control elbow flexion," in *2017 International Symposium on Wearable Robotics and Rehabilitation (WeRob)*. IEEE, nov 2017, pp. 1–2.