Trajectory Optimization and Model Predictive Control for Functional Electrical Stimulation-Controlled Reaching

Derek N. Wolf[1], *Member*, *IEEE*, and Eric M. Schearer[2], *Member*, *IEEE*

***Abstract*—Functional electrical stimulation (FES) offers promise as a technology to restore reaching motions to individuals with spinal cord injuries. To date, the level of reaching necessary for everyday use has not been achieved due to the complexity and limitations of the arm and muscles of an individual with a spinal cord injury. To improve the performance of FES-driven reaching controllers, we developed a trajectory optimization and model predictive control scheme that incorporates knowledge of the person-specific muscle capabilities and arm dynamics. Our controller achieved 3D reaching motions with an average accuracy of 8.5 cm and demonstrated an ability to reach targets throughout the workspace. With improvements to the model, this control scheme has the potential to unlock many daily reaching tasks for individuals with spinal cord injuries.**

***Index Terms*—Rehabilitation Robotics, Prosthetics and Exoskeletons, Motion and Path Planning, Motion Control, Model Learning for Control**

I. Introduction

FUNCTIONAL electrical stimulation (FES) can restore reaching motions to individuals with spinal cord injuries (SCI) by applying electrical stimulation to activate the paralyzed muscles, but the success has been limited. To control reaching motions with FES, techniques such as combined feedforward-feedback control [1] and optimized PD control [2] have shown success in controlling simulated full-arm planar reaching motions. In practice, however, the two most influential FES-driven reaching studies, MUNDUS [3], [4] and the BrainGate2 clinical trials [5], controlled each degree of freedom independently and struggled due to the coupled motions of the other joints. To achieve daily reaching motions, it is necessary to control the arm as a complete system.

There have been two main attempts at practical implementation of full-arm reaching motions, the Sharif Razavian controller [6] and our own previous work [7], [8]. Both controllers used machine learning methods to find person-specific models of an individual's arm and its response to stimulation. The controllers then used a feedback control strategy to select the muscle stimulation commands necessary to drive the arm through a selected trajectory. The Sharif Razavian controller drives the hand of a healthy subject in straight line paths through a 2D tabletop workspace using a PID control. Our own control study implemented a similar PI control structure to drive the hand of an individual with SCI through 3D reaching motions with some success, but the controller was not able to accurately reach targets throughout the entire workspace [8]. This was due to the unique muscle capabilities of individuals with SCI including limited muscles which can be actuated, rapid muscle atrophy [9], and increased muscle fatigue when electrically stimulated [10]. Due to these limitations, FES stimulation cannot produce torque in any desired direction in every arm configuration. When moving in straight-lines, the simple PD control structure often called for torques that either saturated the muscles or could not be produced in the desired torque-space direction. In these situations, the arm would get stuck in a configuration. One solution is to use a more advanced controller such as model predictive control (MPC) which has been successful in controlling knee extension with FES [11]. Unlike PID controllers, MPC controllers can avoid infeasible torque commands by incorporating knowledge of the actuation limits, system dynamics, and muscle capabilities.

The contribution of this work is a novel FES reaching control strategy that uses a person-specific model of the arm's response to electrical stimulation as the basis of a trajectory optimization and MPC control strategy. This structure improves upon previous control strategies by finding and following trajectories while directly accounting for the person-specific muscle capabilities. In this study, we demonstrated our strategy with a participant with SCI.

II. Methods

We developed a controller to drive the arm of an individual with high tetraplegia due to SCI along desired reaching trajectories (see Fig. 1). We first identified a person-specific dynamic model of the participant's arm. We then used trajectory optimization to find feasible reaching motions throughout the subject's workspace and used an MPC controller to drive the arm along the desired trajectories.

The experimental session took place over a single 3.5 hour experiment. Initial set-up (attaching motion capture markers and setting up coordinate frames and finding comfortable arm configurations) required 40 minutes. Data gathering for the day of model identification and training the new model required 30 minutes. Finding feasible trajectories with trajectory optimization required an additional 20 minutes. We spent 15 minutes tuning the MPC controller, and the remainder of the time was used for completing reaches. The participant was allowed breaks whenever requested.

Manuscript received: September 9, 2021; Revised: December 6, 2021; Accepted: January 20, 2022.

This paper was recommended for publication by Editor Pietro Valdastri upon evaluation of the Associate Editor and Reviewers' comments.

This work was supported by NIH NINDS grant N01-NS-5-2365, Veteran's Affairs grant B2359-C, and NSF grant 1751821.

[1]D. N. Wolf is with the Department of Mechanical Engineering, Vanderbilt University, Nashville, TN USA, derek.wolf@vanderbilt.edu.

[2]E. M. Schearer is with the Center for Human-Machine Systems, Cleveland State University, the Cleveland Functional Electrical Stimulation Center, and the Department of Physical Medicine and Rehabilitation, MetroHealth Medical Center, Cleveland, OH USA, e.schearer@csuohio.edu.

Digital Object Identifier (DOI): see top of this page.

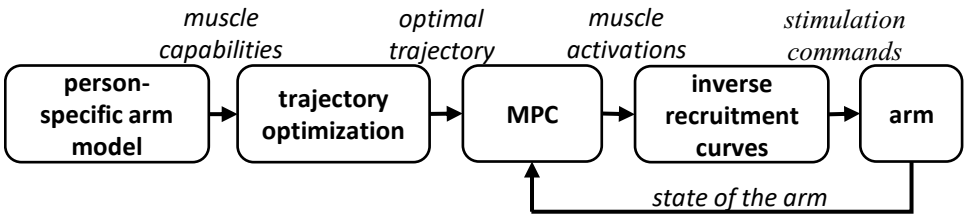

Fig. 1. Controller block diagram: We identified a person-specific model of the arm of an individual with a spinal cord injury. We used this model and a dynamic arm simulation to find an optimal trajectory between a starting and target configuration. We used a model predictive controller to determine the muscle group activations to best drive the arm along the trajectory. The corresponding stimulation commands were applied to the participant's arm.

A. Experimental setup

We worked with a single individual with high tetraplegia due to SCI. The individual sustained a hemisection of the spinal cord at the C1-C2 level and is unable to voluntarily move her right arm. 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).

The participant is implanted with the IST-12 stimulator telemeter in her abdomen [12]. The device uses intramuscular electrodes [13] and nerve cuff electrodes [14] to activate paralyzed muscles. Control signals are sent from the computer to the device via a radio frequency link. We controlled nine muscle groups with the device: 1. triceps, 2. deltoids, 3. latissimus dorsi, 4. serratus anterior, 5. biceps and brachialis, 6. supraspinatus and infraspinatus, 7. rhomboids, 8. lower pectoralis, and 9. upper pectoralis. Muscle stimulation uses a bi-phasic, charge balanced pulse delivered at 13 Hz. The muscle activation is adjusted by modifying the stimulation pulse-width. Stimulation safety limits existed.

We gathered model training data using a HapticMaster (Moog FCS) robot with three degrees of freedom. The participant's wrist was attached to the robot via a ball and socket joint. The robot recorded the 3D forces of its end-effector during model identification. During control experiments, the robot provided a supporting force which countered the arm support and the force of gravity to allow the subject's arm to move freely. The robot also created a haptic bounding box around the edge of the workspace to ensure patient safety. An Optotrak Certus Motion Capture System (Northern Digital, Inc.) was used to measure the arm configuration.

The control and data collection occurred at 52 Hz, but stimulation inputs were updated at the stimulation frequency of 13 Hz. The experiment was controlled using MATLAB xPC target on a Dell Dimension 8400 PC with a Pentium 4 3.20 GHz processor. Trajectory optimization was completed using MATLAB 2019b and IPOPT [15].

B. Person-specific dynamics model used with trajectory optimization and MPC

Our strategy to plan and control the arm to move along optimal trajectories with MPC requires a model of the arm's musculoskeletal dynamics. We developed a person-specific, muscle capability model that uses Gaussian process regression (GPR) to predict for a given arm configuration the muscle torques about each joint produced by electrical stimulation. To train the GPR models, we measured the amount of force needed for a robot to hold the arm static at arm configurations throughout the subject's workspace when the arm is passive and when each muscle group was independently fully activated. The kinematic Jacobian was used to calculate the equivalent joint torques needed to hold the arm static. The difference between the calculated torque for the passive and activated configurations determines the amount of torque produced by a given muscle group. We assumed that the muscle torques combine linearly. The limits of the participant's workspace were determined by her stated comfort. The procedure is presented in detail in [16], [17].

The dynamics model consisted of two links, a humerus and a forearm, with four degrees of freedom —shoulder plane of elevation, shoulder elevation, shoulder rotation, and elbow flexion as defined in [18]. The participant's segments were measured to be 0.315 m for the humerus and 0.253 m for the forearm (note: we define the wrist position as the position of the endpoint of the forearm). The mass, moments of inertia, and position of the center of masses for each link were estimated using the properties from [19]. The equations of motion were found using Autolev 4.3 [20].

The model made several assumptions regarding the dynamics of the real-life system. The model included passive stiffness of 1 Nm/rad and damping of 1 Nms/rad on each degree of freedom to create an equilibrium configuration which is critical for numerical stability. The stiffness also represents some portion of the dynamics produced by the elasticity in the participant's mobile arm support, though these properties were not explicitly measured. The model also did not include gravity. For the control experiments, a GPR model was used to predict the passive force needed to hold the wrist at a static position. This predicted force was then applied by the robot to support the arm against gravity and the army support and allow the arm to move freely.

To account for day-to-day changes in muscle strength, we developed a "day of" muscle model update. We gathered new training data for our GPR models at 13 wrist positions spread throughout the subject's comfortable 3D workspace while maintaining the hyperparameters of a previously trained model. This process efficiently updated our model without the need to repeat targets multiple times nor for significant computation time recomputing the hyperparameters.

C. Trajectory optimization

Using the person-specific dynamic arm model produced in II-B, we developed a trajectory optimization scheme to

find feasible trajectories which accounted for the participant's muscle capabilities and the arm dynamics.

The starting configuration for all trajectories was the participant's natural resting position which was the first configuration of the "day of" model identification. The remaining 12 training configurations were used as targets. The targets were selected to be spaced throughout the subject's comfortable workspace and spanned 12 cm in the x direction, 7 cm in y, and 15 cm in z.

We used the trajectory optimization technique described for optimizing human gait in [21]. In this method, we used the direct collocation method which transforms the optimal control problem of calculating the optimal muscle activations to achieve the desired motion to a constrained nonlinear program. IPOPT [15] was used to solve this nonlinear program. For each target configuration, we attempted to find a feasible two-second trajectory between an initial configuration, $\boldsymbol{q}_0$, and a target configuration, $\boldsymbol{q}_{\text{targ}}$, by formulating and solving the following nonlinear optimization problem,

$$\begin{aligned}
&\underset{\boldsymbol{\alpha},\boldsymbol{x}}{\text{minimize:}} \quad \text{mean}(\boldsymbol{\alpha}^2) + \gamma\ \text{mean}((\boldsymbol{q}-\boldsymbol{q}_{\text{targ}})^2)\\
&\text{subject to:}\\
&\quad \text{state constraints}\\
&\qquad \boldsymbol{x}_{\min} \le \boldsymbol{x}_k \le \boldsymbol{x}_{\max}, \quad \forall k \in \{1,2,\ldots,n\}\\
&\quad \text{activation constraints}\\
&\qquad \boldsymbol{\alpha}_{i,k} \in [0,1],\ \forall i \in \{1,2,\ldots,9\},\ \forall k \in \{1,2,\ldots,n\}\\
&\quad \text{dynamics constraints}\\
&\qquad f(\boldsymbol{x}_k, \dot{\boldsymbol{x}}_{k+1}, \boldsymbol{\alpha}_k) = 0,\ \forall k \in \{1,2,\ldots,n-1\}\\
&\quad \text{task constraints}\\
&\qquad \boldsymbol{x}_1 = [\boldsymbol{q}_0 \quad 0]^\top \qquad \boldsymbol{x}_n = [\boldsymbol{q}_{\text{targ}} \quad 0]^\top.
\end{aligned} \tag{1}$$

The first term of the objective function minimizes the average of the squared muscle activations, $\boldsymbol{\alpha}$, for all $n = 80$ nodes of the trajectory. The second term attempts to minimize the distance from each configuration across all n nodes of the trajectory, $\boldsymbol{q}$, to the target configuration, $\boldsymbol{q}_{\text{targ}}$, leading to more direct reaches. γ was selected to be $\gamma = 1$ rad^{-2} to achieve the overall goal of balancing the objectives of minimal activations and direct path reaches.

The optimization problem included constraints on the state (joint angles and joint velocities), muscle activations, dynamics, and task constraints. The joint angles were constrained to be between the minimum and maximum joint angles seen during the day of model identification with an additional 11.5° of rotation. The joint velocities could have a maximum magnitude of 10 rad/s. The combined state constraints are represented by $\boldsymbol{x}_{\min}$ and $\boldsymbol{x}_{\max}$. The muscle activations were required to remain between 0 and 1. The dynamics constraints ensured that the dynamics of the arm were satisfied throughout the trajectory and that the person-specific muscle capabilities are accounted for. The dynamics are approximated using the semi-implicit Euler method. The task constraints ensured that the first node began at the start configuration with zero velocity and the final node ended at the target configuration with zero velocity.

Due to time constraints, for each target, a trajectory optimization was attempted once for up to 1,500 iterations. For the 12 possible reaches, we were able to find 11 feasible reaching trajectories in 15 minutes of computation time.

Once a trajectory was found, a full five-second trajectory was created. The starting configuration was held for one second followed by the two second optimized trajectory. The final two seconds of the trajectory were to hold the target configuration to allow the controller time to correct for errors.

D. Controller

We linearized the nonlinear person-specific model developed in II-B to form the basis of an MPC controller to drive the arm of a subject with SCI along a desired trajectory (see Fig. 1). The input to the controller is the optimized desired trajectory and the current state of the arm. The MPC controller then calculates the desired muscle activations to best achieve the desired trajectory. The inverse recruitment curves block determines the stimulation commands that correspond to the desired muscle activations, and the stimulation is applied to the arm. The recruitment curves, the mapping from stimulation input to muscle activation, were identified using the deconvolved ramp method [22].

The state of the arm is defined by the arm configuration —shoulder elevation plane, shoulder elevation, shoulder rotation, elbow flexion —and the joint velocities. The state of the arm was measured with motion capture and a 5^{th} order moving average filter was used. The joint velocities were calculated using numerical differentiation.

We developed an MPC controller based on the incremental MPC formulation presented in [23] which incorporates integral control. The Autolev equations of motion used in the arm model from II-B were used to linearize the system model about the current state. The MATLAB function `c2d` was used to create a discretized state-space system. The output of the system and the reference trajectory included only the joint angles. To add integral action, the state is augmented with the current muscle activations, and the control input is defined as the change in muscle activations, $\Delta\boldsymbol{\alpha}$.

For a given time-step, k, the augmented, discretized state-space model of the system can be written as

$$\begin{bmatrix}\boldsymbol{x}_{k+1}\\ \boldsymbol{\alpha}_k\end{bmatrix} = \begin{bmatrix}A & B\\ 0 & I\end{bmatrix}\begin{bmatrix}\boldsymbol{x}_k\\ \boldsymbol{\alpha}_{k-1}\end{bmatrix} + \begin{bmatrix}B\\ I\end{bmatrix}\Delta\boldsymbol{\alpha}_k \tag{2}$$

$$\boldsymbol{y}_k = [C \quad D]\begin{bmatrix}\boldsymbol{x}_k\\ \boldsymbol{\alpha}_{k-1}\end{bmatrix} + D\Delta\boldsymbol{\alpha}_k. \tag{3}$$

The state-space matrices are assumed constant for the control calculations. The controller aims to select the input commands which minimize the objective function

$$J = \sum_{i=1}^{n_y} {e_{k+i}}^T e_{k+i} + \lambda \sum_{i=0}^{n_u-1} \Delta\boldsymbol{\alpha}_{k+i}{}^T \Delta\boldsymbol{\alpha}_{k+i}. \tag{4}$$

The first term of the equation minimizes the error, e_{k+i}, for a given time-step which is defined as the difference between the estimated joint positions, $\boldsymbol{y}_k$, predicted by (3) and the reference trajectory. The prediction horizon, n_y, determines for how many time steps forward the model predicts states and system error. The second term penalizes changes in muscle activations. The control horizon, n_u, determines the number of

time steps forward that the controller optimizes control inputs. For time steps $n_u < i < n_y$, $\Delta u = 0$. The lumped scalar weighting λ acts as a muscle group activation smoothness parameter by weighting the amount that the activation commands change during a time-step.

The parameters of the MPC controller were tuned on several trajectories with the goal of producing accurate reaches and comfortable stimulation profiles and motions. To solve the MPC optimization problem within the 52 Hz control loop using the active-set method, the prediction horizon was selected to be $n_y = 3$ and the control horizon was $n_u = 2$.

The key features for subject comfort were limiting oscillation of the arm and smooth stimulation profiles. To create a smooth stimulation pattern, we selected a scalar weighting on the change in input, $\lambda = 1$. During our initial tuning, it became clear that oscillations would be an issue with the controller. Our previous simulation research has demonstrated that significant oscillations occur with feedback FES control due to the delays in the FES system, and controller performance can be improved by adding physical damping to the arm support [24]. Due to this finding, we used the robot to create a damped environment of 70 Nm/s in all directions.

Due to the time constraints of the control loop, we linearized the system offline. The output of the state-space models were the joint angles, and the reference trajectory included only the discretized desired joint angles as a function of time-step. At each point during the reaching experiment, the controller would use the linearized system matrices from the desired reference state of the arm at the next time step. The reference signal for the controller was the desired arm configuration. Joint velocities were not included in the reference signal due to the aforementioned issues with system delays leading to derivative control instability.

E. Experiments and data analysis

For each target reach, the robot moved the subject's arm to the starting position. Because the internal arm joint angles were not controlled by the robot, the first second of each reach was to hold at the starting configuration to allow the controller to correct for initial configuration errors. For the duration of the five second reach, the arm was allowed move freely as driven by the muscle stimulations. The robot provided only a damped environment and support against the predicted passive forces at the wrist during the movement. The set of 11 reaching motions was repeated nine times based on the amount of time defined by the subject's schedule. A total of 99 reaches were completed.

To analyze the effectiveness of the controller, we calculated the accuracy as defined as the mean Euclidean distance of the wrist away from the desired position over the final second of each trial. The accuracy of the controller was recorded based on the wrist position as placing the hand at a desired location in space is the most important goal of a reaching controller. The wrist position of both the desired position and experimental measurements were calculated using forward kinematics. To determine the controller's effectiveness throughout the workspace, a two-sample t-test was completed for each grouping of targets to determine if there was improved accuracy to the left or right side of the workspace, up or down, and forward or backwards in the subject's workspace. The targets were grouped based on their position relative to the average position of all the targets.

III. Results

Over 99 reaching motions to 11 different targets, the controller achieved an average wrist position accuracy of 8.5 cm (standard deviation of 2.8 cm). Table I shows the accuracy results for all trials including a breakdown of the accuracy based on the position of the targets relative to the average target position. The controller was able to reach targets throughout the workspace, but it was more accurate to targets on the right side ($p < 0.001$), forward ($p < 0.001$), and down in the workspace ($p = 0.002$). These differences in accuracy based on target position are illustrated in Fig. 2. The image shows the average accuracy over the nine sets of reaches for each target. The size and color of each point represents the relative accuracy for each target.

A representative reaching trial with an accuracy of 8.5 cm is shown in Fig. 3. The target position is denoted by the red arrow in Fig. 2. The reach is able to move in the correct direction, but there are significant amounts of oscillation near the target position. Shoulder rotation is the joint with the largest error that the controller is unable to correct. As this joint moves away from the desired target, the wrist position also moves away from the desired target.

Fig. 3(c) shows the muscle activation commands for the triceps, biceps/brachialis, and the upper pectoralis muscle groups. These activations demonstrate the ability of the MPC controller built on our muscle capability models to select muscle activations which make sense physiologically. For the elbow flexion angle, the reach first requires elbow extension so the triceps, the main elbow extensor muscle, is activated. As the position overshoots, the biceps/brachialis turn on to stop the elbow extension. These two muscle groups work to control the elbow flexion as it oscillates around the desired position. Additionally, as the shoulder rotation moves away from the desired target, the integral action of the MPC controller is noticeable as the upper pectoralis, an internal rotator of the arm, increases in activation. While this is a simplified explanation of the reach (the biceps/brachialis and triceps also produce torques about the shoulder), this example reach demonstrates the potential of the control strategy to achieve reaching motions throughout the workspace.

IV. Discussion

We present a novel control structure that can achieve FES-driven reaching motions throughout the workspace of an individual with SCI. The use of trajectory optimization and an MPC controller to directly account for the subject's muscle capabilities resulted in improvements to the performance compared to past FES reaching controllers. First, though there were differences in accuracy based on the target position in the workspace, the controller could move the arm throughout the space. The MPC controller avoided getting

Target location	All	Left	Right	Forward	Back	Up	Down
mean accuracy (cm)	8.5	10.0	6.7	7.2	10.0	9.4	7.7
standard deviation (cm)	2.8	2.2	2.2	2.4	2.3	2.3	2.9
p-value		< 0.001		< 0.001		0.002	

TABLE I
WRIST POSITION ERROR FOR ALL TARGETS AND BROKEN DOWN BY THE TARGET POSITION RELATIVE TO THE AVERAGE POSITION

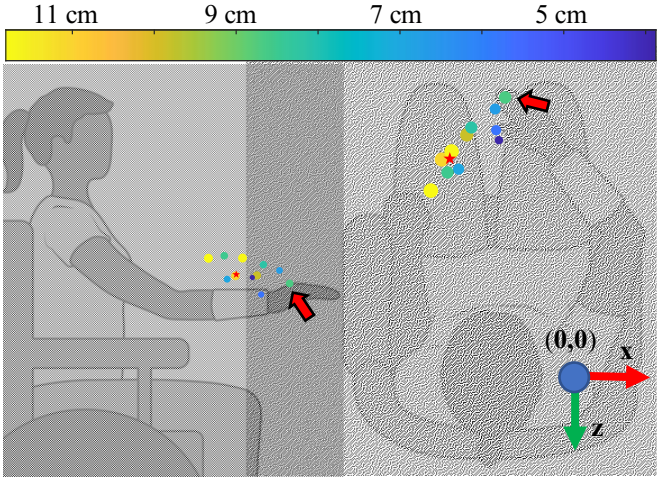

Fig. 2. This figure shows the average accuracy of each target position over nine sets of reaching trials. The starting position for every reach is denoted by the red star. The accuracy of the target is denoted by both the size and color of the circle. Relative to the average target position, the accuracy of the reaches improved for targets to the right, to the front, and down in the workspace (see Table I). The red arrow denotes the representative target position which is shown in Fig. 3. The coordinate frame and origin of the reaching Cartesian workspace is also shown.

stuck in configurations as seen with the straight-line feedback controller in [8]. Additionally, by controlling the whole arm, we are able to produce more natural motions than the reaches achieved by MUNDUS [4] once a trajectory is found. Lastly, our controller can be implemented in any individual with SCI by accounting for the person specific muscle capabilities while controllers demonstrated in healthy subjects may not immediately translate to an individual with SCI.

While the positives are significant, there were important limitations to the study. The accuracy of 8.5 cm is worse than the accuracy produced by our own previous work in [8] and other controllers [5], [4], [6]. This accuracy is not good enough to complete many daily activities including eating off a plate though some compensatory torso movements could assist these errors. One potential cause of this inaccuracy is a relatively short prediction horizon which was selected due to computational limitations of the FES system to compute activations in real-time. Improved hardware could use a longer prediction horizon to improve accuracy while still providing the benefits of accounting for muscle capabilities.

Another limitation of this study is that the spacing of the targets was limited by the subject's muscle hypertonia. Additionally, the trajectory optimization can take a significant amount of time (from ten seconds at minimum and up to five minutes if 1,500 iterations occur). For daily use, the optimization needs to occur faster to be practical. Another issue is the oscillation which makes daily tasks difficult (e.g., eating off a fork that is oscillating) and uncomfortable. To improve both the oscillation and the accuracy of the controller,

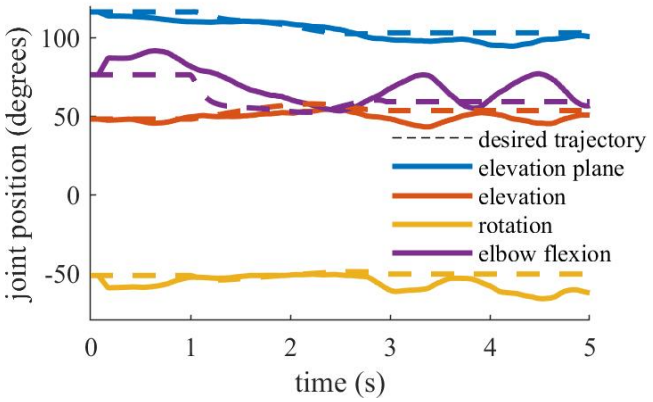

(a) configuration

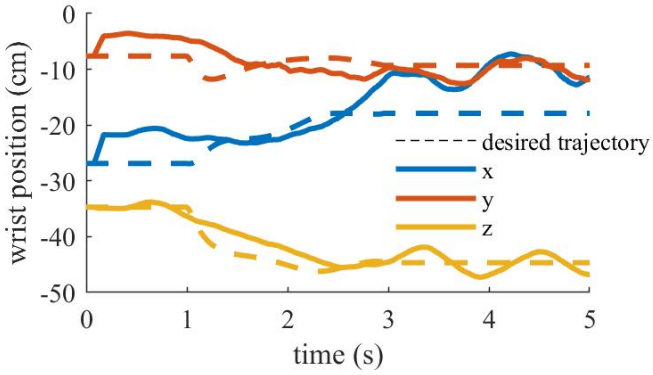

(b) wrist position

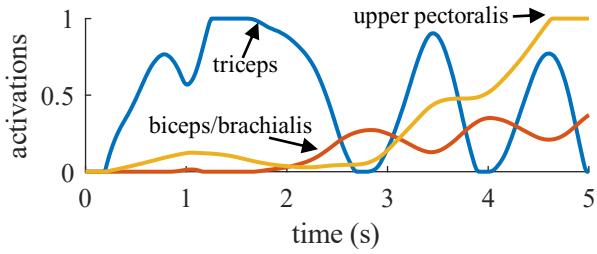

(c) muscle activations

Fig. 3. This figure shows a representative example of an FES-driven reaching motion controlled by our MPC scheme in both (a) configuration space and (b) Cartesian wrist position space. (c) shows the muscle activation commands for three muscle groups, the triceps, biceps/brachialis, and the upper pectoralis. These muscle group activations demonstrate the ability of the MPC control scheme to select muscle activations which make sense physiologically and are able to control the motion of the arm.

the MPC controller's model needs to be improved. Our current model accounts for basic multi-body dynamics of the arm with estimated mass properties and the capabilities of the muscles. A more advanced model which explicitly included muscle activation dynamics, the stiffness introduced by the arm support, and delays in control inputs because of the stimulation frequency could significantly improve the model. Our previous simulation study [24] has shown that these nonlinearities produce oscillation that derivative gain cannot account for. By improving the modeling of these sources of error, the controller could reduce oscillation.

Identifying the subject's arm dynamics is a difficult process which requires gathering a significant amount of data. One option would be to use semiparametric models of the muscle capabilities that include the arm dynamics [16]. With this model of the arm dynamics and muscle force production matrices, we could better predict the dynamics of the arm and achieve more accurate trajectories. A similar method has been used in robotic simulations using a receding horizon LQR controller and a GPR model of the dynamics [25].

This study presents a novel control strategy for FES-driven reaching motions. The use of trajectory optimization and MPC control creates a control scheme which can account for the unique muscle capabilities of individuals with SCI including muscle weakness or a complete loss of muscle function due to lower motor neuron damage. With an improved model, this control scheme has the potential to unlock many daily reaching motions for individuals with SCI.

REFERENCES

[1] 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.

[2] D. Blana, R. F. Kirsch, and E. K. Chadwick, "Combined feedforward and feedback control of a redundant, nonlinear, dynamic musculoskeletal system," *Medical & Biological Engineering & Computing*, vol. 47, no. 5, pp. 533–542, May 2009.

[3] C. Klauer, T. Schauer, W. Reichenfelser, J. Karner, S. Zwicker, M. Gandolla, E. Ambrosini, S. Ferrante, M. Hack, A. Jedlitschka, A. Duschau-Wicke, M. Gföhler, and A. Pedrocchi, "Feedback control of arm movements using Neuro-Muscular Electrical Stimulation (NMES) combined with a lockable, passive exoskeleton for gravity compensation," *Frontiers in Neuroscience*, vol. 8, no. Sep, pp. 1–16, 2014.

[4] A. Pedrocchi, S. Ferrante, E. Ambrosini, M. Gandolla, C. Casellato, T. Schauer, C. Klauer, J. Pascual, C. Vidaurre, M. Gföhler, W. Reichenfelser, 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.

[5] 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.

[6] R. Sharif Razavian, B. Ghannadi, N. Mehrabi, M. Charlet, and J. McPhee, "Feedback Control of Functional Electrical Stimulation for 2-D Arm Reaching Movements," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 26, no. 10, pp. 2033–2043, oct 2018.

[7] D. N. Wolf and E. M. Schearer, "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.

[8] D. N. Wolf, Z. A. Hall, and E. M. Schearer, "Model Learning for Control of a Paralyzed Human Arm with Functional Electrical Stimulation," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 10 148–10 154.

[9] L. R. Sheffler and J. Chae, "Neuromuscular electrical stimulation in neurorehabilitation," *Muscle & Nerve*, vol. 35, no. 5, pp. 562–590, 2007.

[10] R. J. Downey, M. J. Bellman, H. Kawai, C. M. Gregory, and W. E. Dixon, "Comparing the Induced Muscle Fatigue Between Asynchronous and Synchronous Electrical Stimulation in Able-Bodied and Spinal Cord Injured Populations," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 23, no. 6, pp. 964–972, 2015.

[11] N. Kirsch, N. Alibeji, and N. Sharma, "Nonlinear model predictive control of functional electrical stimulation," *Control Engineering Practice*, vol. 58, pp. 319–331, Jan 2017.

[12] 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.

[13] W. Memberg, P. Peckham, and M. Keith, "A surgically-implanted intramuscular electrode for an implantable neuromuscular stimulation system," *IEEE Transactions on Rehabilitation Engineering*, vol. 2, no. 2, pp. 80–91, Jun 1994.

[14] G. G. Naples, J. T. Mortimer, A. Scheiner, and J. D. Sweeney, "A Spiral Nerve Cuff Electrode for Peripheral Nerve Stimulation," *IEEE Transactions on Biomedical Engineering*, vol. 35, no. 11, pp. 905–916, 1988.

[15] A. Wächter, "An interior point algorithm for large-scale nonlinear optimization with applications in process engineering," 2002.

[16] 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.

[17] D. N. Wolf and E. M. Schearer, "Evaluating an open-loop functional electrical stimulation controller for holding the shoulder and elbow configuration of a paralyzed arm," in *2017 International Conference on Rehabilitation Robotics (ICORR)*, 2017, pp. 789–794.

[18] G. Wu, F. C. Van Der Helm, H. E. Veeger, M. Makhsous, P. Van Roy, C. Anglin, J. Nagels, A. R. Karduna, K. McQuade, X. Wang, F. W. Werner, and B. Buchholz, "ISB recommendation on definitions of joint coordinate systems of various joints for the reporting of human joint motion - Part II: Shoulder, elbow, wrist and hand," *Journal of Biomechanics*, vol. 38, no. 5, pp. 981–992, 2005.

[19] E. Chadwick, D. Blana, A. van den Bogert, and R. Kirsch, "A Real-Time, 3-D Musculoskeletal Model for Dynamic Simulation of Arm Movements," *IEEE Transactions on Biomedical Engineering*, vol. 56, no. 4, pp. 941–948, Apr 2009.

[20] D. A. Levinson and T. R. Kane, "AUTOLEV — A New Approach to Multibody Dynamics BT - Multibody Systems Handbook," W. Schiehlen, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 1990, pp. 81–102.

[21] M. Ackermann and A. J. van den Bogert, "Optimality principles for model-based prediction of human gait," *Journal of Biomechanics*, vol. 43, no. 6, pp. 1055–1060, 2010.

[22] W. K. Durfee and K. E. MaClean, "Methods for Estimating Isometric Recruitment Curves of Electrically Stimulated Muscle," *IEEE Transactions on Biomedical Engineering*, vol. 36, no. 7, pp. 654–667, 1989.

[23] H. Richter, *Advanced control of turbofan engines*. Springer Science & Business Media, 2011.

[24] D. N. Wolf and E. M. Schearer, "Developing a quasi-static controller for a paralyzed human arm: A simulation study," in *IEEE International Conference on Rehabilitation Robotics*, vol. 2019-June. IEEE, Jun 2019, pp. 1153–1158.

[25] J. Boedecker, J. T. Springenberg, J. Wülfing, and M. Riedmiller, "Approximate real-time optimal control based on sparse Gaussian process models," in *2014 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)*, 2014, pp. 1–8.