Online Electromyographic Control of a Robotic Prosthesis

Pradeep Shenoy*, Kai J. Miller, Beau Crawford, and Rajesh P. N. Rao

***Abstract*—This paper presents a two-part study investigating the use of forearm surface electromyographic (EMG) signals for real-time control of a robotic arm. In the first part of the study, we explore and extend current classification-based paradigms for myoelectric control to obtain high accuracy (92–98%) on an eight-class offline classification problem, with up to 16 classifications/s. This offline study suggested that a high degree of control could be achieved with very little training time (under 10 min). The second part of this paper describes the design of an online control system for a robotic arm with 4 degrees of freedom. We evaluated the performance of the EMG-based real-time control system by comparing it with a keyboard-control baseline in a three-subject study for a variety of complex tasks.**

***Index Terms*—Classification, electromyography, online, prosthetics, support vector machines.**

I. Introduction

THE surface electromyogram (EMG) provides a noninvasive method of measuring muscle activity, and has been extensively investigated as a means of controlling prosthetic devices. Amputees and partially paralyzed individuals typically have intact muscles that they can exercise varying degrees of control over. Further, there is evidence [2] that amputees who have lost their hand are able to generate signals in the forearm muscles that are very similar to those generated by healthy subjects. Thus, the ability to decode EMG signals can prove extremely useful in restoring some or all of the lost motor functionality in these individuals.

While the research community has focused on the use of sophisticated signal-processing techniques to achieve accurate decoding, clinical studies observe that widespread acceptance of prosthetic devices is difficult to achieve, and that such a prosthetic needs to be both highly accurate and intuitive to control. Thus, although offline studies [3], [4] have shown that up to six classes of gestures can be decoded from forearm electrodes with very high accuracy, the questions of online control and its expressivity and ease of use have been left open.

In this paper, we present results from a two-stage pilot study that addresses many of the issues involved in EMG control of prosthetic devices. In the first part, we extend the results obtained by other offline studies and show a classification accuracy of 92–98% on an eight-class classification problem, with a classification rate of 16 outputs/s. Our results rely on careful selection of physiologically relevant sites for recording EMG signals, and on the use of simple but powerful classification methods. These offline results provide the basis for the development of an expressive and accurate interface for online control. In the second part, we design and evaluate an online 4DOF control system for a robotic arm. Here, we address the issues of ease of use and quality of control, by choosing an intuitive gesture-to-control mapping, and by comparing control performance against a baseline obtained via keyboard-based control. We demonstrate the robustness of our method in a variety of reasonably complex online robotic control tasks involving 3-D goal-directed movements, obstacle avoidance, and pickup and accurate placement of objects.

Our results show that healthy subjects can gain significantly expressive EMG-based control of a prosthetic device, and pave the way for the design of powerful prosthetics with multiple degree-of-freedom control. We believe that our techniques can also be applied to the design of novel user interfaces based on EMG signals for human–computer interaction and activity recognition.

II. Background and Related Work

Muscle contraction is the result of the activation of a number of muscle fibers. This process of activation is mediated by the firing of neurons that recruit muscle fibers. In order to generate more force, a larger number of muscle fibers must be recruited through neuronal activity. The associated electrical activity can be measured in sum at the surface of the skin as an EMG signal. The EMG signal is thus a measure of muscle activity. Its properties have been studied extensively [5]. The amplitude of the EMG signal is correlated with the force generated by the muscle; in particular, isometric steady-state contraction of an individual muscle is proportional to the force produced by the muscle. However, this relationship is noisy, and changes significantly with the change in the shape of the muscle, fatigue in a muscle, etc. It is also very difficult to isolate activity from a single muscle using noninvasive surface measurements. In addition, the same gesture can be generated using different combinations of forces in groups of coordinated muscles. Thus, decoding the pose of the forearm or hand using EMG signals and muscle models for the individual muscles involved is a challenging task.

Manuscript received August 12, 2006; revised July 17, 2007. This work was supported in part by the National Science Foundation under Grants 0130705 and 0622252, and in part by the Packard Foundation. A preliminary version of this paper was presented as a poster at the AAAI conference [1], 2005. *Asterisk indicates corresponding author.*

*P. Shenoy is with the Department of Computer Science and Engineering, University of Washington, P.O. Box 352350, Seattle, WA 98195 USA (e-mail: pshenoy@cs.washington.edu).

K. J. Miller, B. Crawford, and R. P. N. Rao are with the Department of Computer Science and Engineering, University of Washington, Seattle, WA 98195 USA (e-mail: pshenoy@cs.washington.edu; kai@cs.washington.edu; beau@cs.washington.edu; rao@cs.washington.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TBME.2007.909536

TABLE I
ELECTRODE LOCATIONS CHOSEN ON THE FOREARM (SEE FIG. 2)

	Muscle	Function
1	Brachioradialis	forearm flexion
2	Extensor carpi ulnaris	extension and adduction of hand at the wrist
3	Pronator teres	pronation and elbow flexion
4	Extensor communis digitorum	finger extension at metacarpo-phalangeal joints, wrist extension at forearm
5	Flexor carpi radialis	hand flexion and abduction at wrist
6	Anconeus	antagonistic activity during forearm pronation
7	Pronator quadratus	initiates pronation

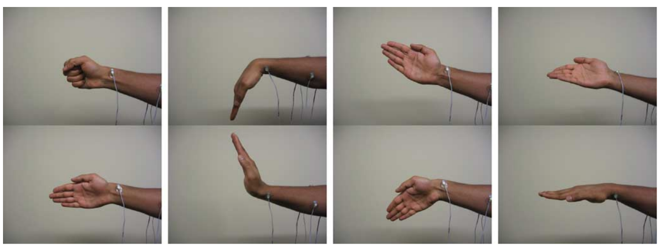

Fig. 1. Static hand gestures chosen for classification. The goal is to use gross gestures at the wrist, and decode the gesture from windows of data recorded while the gesture is maintained (the gestures in the second column show a top-down perspective). These gestures intuitively correspond to pairs of actions: grasp–release, left–right, up–down, and rotate.

For obtaining control over a prosthetic device, we need to solve a simpler problem–we only need to identify the signals produced by a small number of gestures. This means that we only need enough information from the EMG signals to distinguish between the given gestures, without explicitly identifying the source of the EMG signals, or modelling the muscles and forces involved. Several researchers have attempted to distinguish a variety of gestures using transient signals recorded at the onset of the gesture. For example, Englehart *et al.* [6] classify four discrete elbow and forearm movements and capture the transient structure in the EMG using various time–frequency representations such as wavelets. They achieve accuracy up to 93.7% with four channels of EMG data from the biceps and triceps. Reischl *et al.* [7] present multiclass classification methods for distinguishing between 5–8 classes of movement onset in amputees using two upper-arm electrodes, with errors of 4–9%. Nishikawa *et al.* [8] classified ten discrete movements of the wrist and fingers using four electrodes placed on the forearm. They propose an online learning scheme, and obtain an average accuracy of 91.5%. Sebelius *et al.* [9] also classify ten similar gestures of wrist/fingers based on EMG onset recorded at eight bipolar electrodes on the forearm. They use a virtual hand for feedback, and a data glove for monitoring movement and training their classifier. Boostani and Moradi [10] present extensive offline analysis comparing the quality of various features extracted from the EMG signal in distinguishing between the onset of a number of different movements in disabled subjects. Ju *et al.* [11] address applications in user interfaces and consumer electronics. They achieve 85% accuracy in classifying four finger movements with the aid of two electrodes placed close to the wrist. The electrode locations are suboptimal but chosen for appropriateness in the chosen applications. Carrozza *et al.* [12] compare foot action versus EMG signals for controlling grasping functions on a hand prosthesis, and find that foot movements were more effective and more easily learned than their EMG-based control scheme.

As remarked by Englehart and Hudgins [13], using transient EMG signals for control is a suboptimal choice for a variety of reasons. For example, this scheme requires initiating a gesture from a state of rest in order to produce a single command. This makes continuous control of devices cumbersome and slow. In addition, the decoding problem for transient signals is significantly harder than that of decoding steady-state signals from a statically heldhand gesture.

Thus, many recent papers [3], [4], [14] have explored continuous control, where a variety of sophisticated algorithms, such as multilayer perceptrons, hidden Markov models, and Gaussian mixtures have successfully decoded six different gestures from continuous data with an accuracy of more than 90%. As an example, Chan and Englehart [3] propose the use of hidden Markov models along with rms and autoregressive features to decode six wrist gestures using four electrodes placed around the forearm. They achieve an accuracy of 94.6% across 12 subjects.

We build on this work and extend it in several interesting directions. Similar to the successful prior work, we use simple features (rms values over windows) and continuously classify windows of data collected while the subject maintains a static hand gesture. In addition, we use knowledge of the physiology of the forearm to carefully choose electrode locations likely to have interesting information about the gestures. This fact, in combination with powerful classification techniques (support vector machines), allow us to classify eight gestures with an accuracy of 92–98% in our pilot study. Finally, we address several interesting and important issues in the design and evaluation of online controllers that have been left open by previous studies.

III. OFFLINE STUDY METHODS

A. Gestures for Robot Arm Control

Fig. 1 shows a list of the actions we chose in our study. These gestures are gross movements at the wrist and involve a number of forearm muscles. Further, they lend themselves easily to interpretation and could serve as a basis for control of a prosthesis.

B. Electrode Placement

The muscles we chose and their relevant functions are listed in Table I. Fig. 2 shows the location of the electrodes on the forearm. In contrast to the differential pair at each recording site traditionally used in the literature [5], we use an eighth electrode on the upper arm as a reference for all other electrodes, and a single electrode at each site of interest. This reference is mainly

Fig. 2. Electrode positions on the forearm chosen for our study. See also Table I.

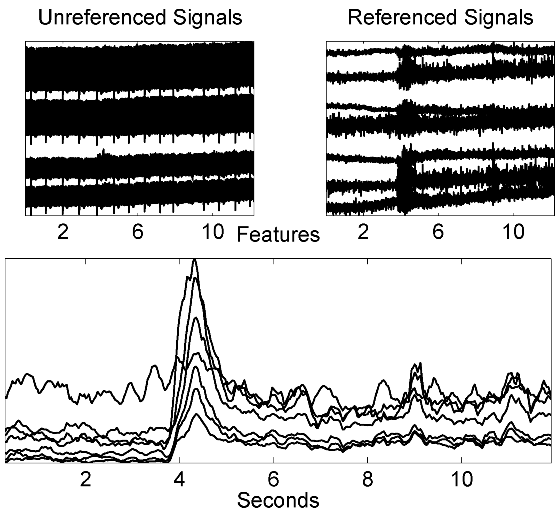

Fig. 3. Samples of EMG signals. The top left and right figure show the difference between unreferenced signals (heavily contaminated with line noise) and signals referenced from an additional electrode on the left forearm, removing the line noise. The computed features (see Section III-D) shown in the bottom frame span a rest period and onset of a gesture, and demonstrate that features during the onset of movement are substantially different from steady-state features while maintaining a gesture.

used to remove 60-Hz contamination due to line noise. Fig. 3 shows how the individual channels have line noise, but the referencing removes this noise.

The particular muscles chosen in our study are implicated in wrist-centered movements and gestures, as shown in the table. The coordinated action of these muscles spans the different movement types which we classify. Although there is redundancy amongst the actions of these muscles as well as redundancy amongst deeper muscles that contribute to the signal, we expect this to lead to robust interpretation across subjects and sessions.

The recording sites corresponding to these muscles were chosen to make the interpretation of the signal as intuitive as possible and as reproducible from subject to subject as possible. While no electrode position will isolate a single muscle, placing a given electrode on the skin directly above a given superficial muscle should ensure that the largest contribution to the signal at that electrode location is from the desired muscle. This comes with the known caveat that the muscles of deeper layers will contribute to the signal, as will adjacent superficial muscles. Since our goal is classification of discrete gestures into a discrete set of actions, and not the study of individual muscles, we rely on the classifier to extract the important features for each class from this mixture of information from each electrode.

C. Data Collection

We collected data from 3 subjects over 5 sessions each. A session consisted of the subject maintaining the 8 chosen action states shown in Fig. 1 for 10 s each. The gestures were chosen to intuitively correspond to pairs of actions: grasp-release, left-right, up-down, and rotate (see Fig. 1). The subjects cycled through the actions in order, separated by 5-s rest periods. The sessions were separated by a 20-s rest period. The subjects were instructed to relax the forearm/hand and maintain each gesture comfortably without exerting excessive force. We did not measure or restrict the force exerted by the subjects while maintaining a given hand pose.

We use five sessions in order to prevent overfitting, as a given action may be slightly different each time it is performed. Thus, in our evaluations, we use each the entire session as testing data, and average the classification results across all five splits.

D. Feature Extraction

We sample the EMG signal at 2048 Hz. Our feature extraction from this is simple: we calculate the rms of windowed steady-state EMG signals from each channel. One-hundred twenty-eight-sample windows are used, and the rms amplitude in this window is computed for each of the seven electrodes. This feature vector serves as the input to our classifier. The choice of a 128-sample window length is empirical, and results in 16 commands/s. This update rate is sufficient for developing a responsive EMG-based controller.

Other work [10] has evaluated the performance of a large number of features for distinguishing between onsets of various kinds of movements. Our scenario involved steady-state EMG signals. However, it is possible that these features may further improve our classification results. We will explore these feature sets as part of future work.

E. Classification With Linear Support Vector Machines

We use linear support vector machines (SVMs) [15] for classifying the feature vectors generated from the EMG data into the respective classes for the gestures. SVMs have proved to be a remarkably robust classification method across a wide variety of applications.

1) Binary Classification: We first consider a two-class classification problem. Essentially, the SVM attempts to find a hyperplane of maximum "thickness" or margin that separates the data points of the two classes. This hyperplane then forms the decision boundary for classifying new data points. Let $\mathbf{w}$ be the normal to the chosen hyperplane. Then, the classifier will label a data point $\mathbf{x}$ as $+1$ or -1, based on whether $\mathbf{w} \cdot \mathbf{x} + b$ is greater than 1, or less than -1. Here, $(\mathbf{w}, b)$ are chosen to maximize the margin of the decision boundary while still classifying the data points correctly.

This leads to the following learning algorithm for linear SVMs. For the classifier to correctly classify the training data points $\mathbf{x_1}, \ldots, \mathbf{x_n}$ with labels $y_1, \ldots, y_n$ drawn from ± 1, the following constraints must be satisfied [15]:

$$y_i(\mathbf{w} \cdot \mathbf{x_i} + b) \geq 1 - \xi_i \quad \forall i \quad \xi_i \geq 0 \quad \forall i.$$

This set of constraints ensures that each data point $\mathbf{x_i}$ is correctly classified, allowing for some small amount of error ξ_i since real-life data are noisy. The optimization goal for the noisy classification case is to minimize $(1/2)\mathbf{w}.\mathbf{w} + C\Sigma_i\xi_i$, where C is a user-specified cost parameter. Intuitively, the criterion is trading off the margin width with the amount of error incurred. We refer the reader to appropriate texts [15] for more technical details. This is the formulation we use, and in this formulation, the classifier has a single-free parameter C that needs to be chosen by model selection.

2) Multiclass Classification and Probabilities: The two-class formulation for the linear SVM can be extended to multiclass problems. Our system uses the following generic method for combining binary classifiers for multiclass classification [16]: for each pair of classes, a separate binary classifier is trained on data from the two classes. In order to classify a test data point, the datapoint is classified by each binary classifier, and each result is counted as a vote for the respective class. The output of the classifier is the class label with the maximum number of votes.

In our system, we use the LIBSVM [17] package which implements the SVM classification algorithm, along with support for multiclass classification. There is also support for estimating class-conditional probabilities for a given datapoint (see [18] for more details on the algorithm used). This can be useful in reducing the number of false classifications due to noisiness in the data. Specifically, the class-conditional probabilities returned can be tested against a threshold, and a "no-operation" command can be executed if the classifier is uncertain about the correct class label for the datapoint. In our online experiments, we used an empirically determined threshold to discard predicted actions that had low probabilities.

IV. Offline Results

A. Classification Accuracy

We use leave-session-out cross-validation error as a measure of performance. That is, we average the results from five runs, in each of which, an entire session of data is used as testing data for a classifier trained on the remaining four sessions.

We used the collected data to train a linear SVM classifier, and performed parameter selection using across-session cross-validation error as measure. Fig. 4 shows the SVM classifier error as a function of the cost parameter C. The graph demonstrates two aspects of the data: First, eight-class classification is performed with an accuracy of 92–98% for all three subjects. Second, the classification results are stable over a wide range of parameters for all three subjects, indicating that in an online setting, we can use a preselected value for this parameter. It is important to note, however, that careful selection is important, as the error can be significant for bad choices of C.

We also note that ten-fold cross-validation on any one session of data yielded 0–2% errors for all subjects, which is significantly less than the across-session error. Since each gesture for a session is essentially one static handpose, the data within a session is likely to be more homogeneous, and thus easier to decode. The fact that across-session error is greater implies that each session does, in fact, have different data, and using multiple sessions is essential to avoid overfitting.

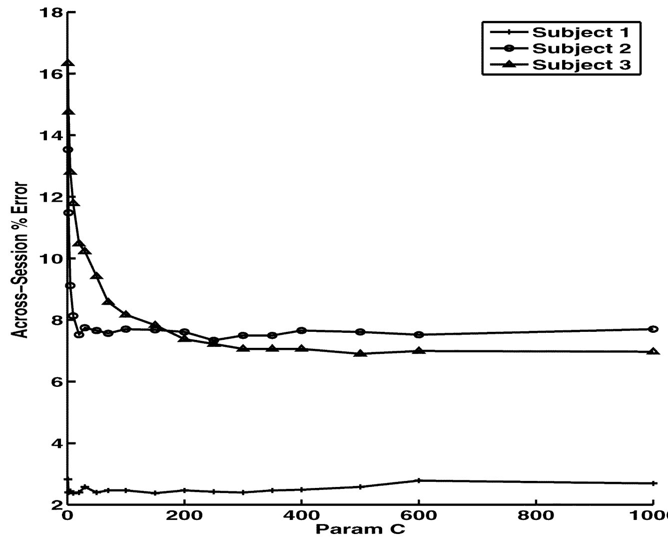

Fig. 4. Classifier error on the eight-gesture classification problem as a function of the SVM cost parameter C (see Section IV-A).

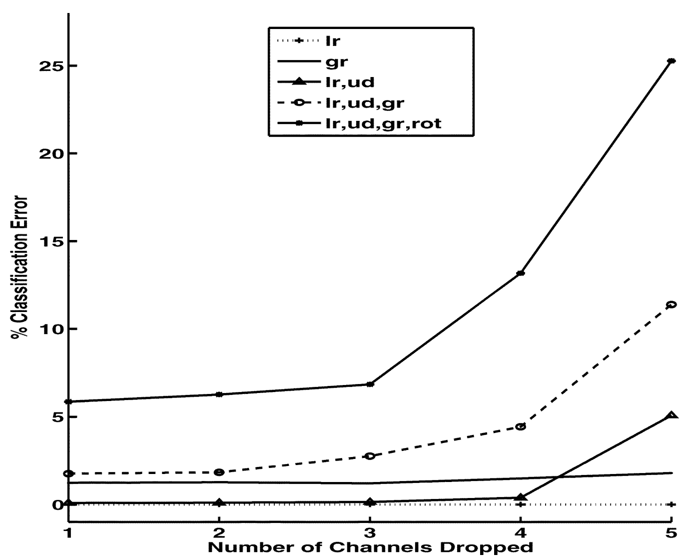

Fig. 5. Classifier error as a function of the number of channels dropped, from an initial set of seven channels. The legend describes the degrees of freedom included in the classification problem lr = $\mathrm{left\text{-}right}$, ud = up- down, gr = grasp-release, rot = rotate.

B. Evaluating Choice of Recording Locations

In previous sections, we noted that our choice of recording sites for muscle activity is motivated by the relevance of the chosen muscles to the gestures we wish to classify. We can, however, quantitatively assess aspects of this selection process. Do all of the channels of EMG data contribute to classification accuracy—is there redundancy for classification in our measurements? How many channels or electrodes would we need for highly accurate control of a given set of degrees of freedom?

Fig. 5 addresses these questions. The figure quantifies the impact of the number of electrode channels used on performance of the linear SVM classifier for various subsets of

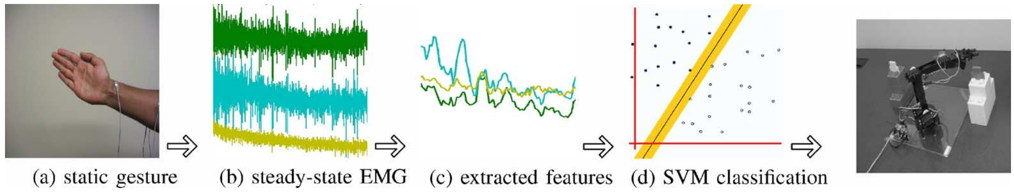

Fig. 6. Schematic for an EMG-based robotic control.

classes. For any single classification problem, the channel dropped at each step was chosen using a "greedy heuristic." That is, at each step, the channel dropped was the one that least increased the cross-validation error of the classifier trained on the remaining channels. This feature-selection procedure was carried out for the following classification problems: 1) grasp–release, 2) left–right, 3) left–right-up–down, 4) left–right–up–down–grasp–release, and 5) all eight classes. These choices represent control of an increasing number of degrees of freedom.

The figure clearly illustrates that, as expected, more degrees of freedom require more channels of information for accurate classification. For example, the two-class classification problems need only one or two channels of information, but the six-class problem requires three or more channels for a low error rate. The figure also shows that the full eight-class classification problem can be accurately solved with fewer than seven electrodes. The order in which the physical channels were dropped was different for each subject. We ascribe this to the variation in the performance of actions by different individuals, and in the spatial distribution of recording quality across individuals.

Since the *apriori* selection of the recording sites themselves was not based on a quantitative optimality criterion, it is possible that fewer electrodes, when more judiciously placed, can prove sufficient. The results of this experiment do not support selection of any particular subset of channels. More important, we do not address the question: "for x electrodes that can be placed anywhere, what is the best performance that can be obtained?" Instead, this experiment indicates that there is information redundancy in this particular choice of channel locations, and that this redundancy contributes to making our classification recipe robust across subjects.

V. Online System Design and Evaluation

Fig. 6 details our online system design: The user maintains a static hand pose that corresponds to one of a predefined set of gestures in Fig. 1. We record EMG activity from various locations on the forearm as used in the offline study. This data stream is transformed into feature vectors which are updated at 16 Hz, and classified by a linear SVM classifier. This classifier's output serves as a discrete command that moves the robotic arm by a small, fixed amount in the designated direction. Maintaining a specific hand gesture will make the arm move continuously in a chosen direction.

Fig. 7 shows the chosen mapping from gestures to degrees of freedom in the robotic arm. Care was taken to make the mapping as intuitive as possible, and the chosen gestures are appropriate metaphors for the corresponding movement of the robotic

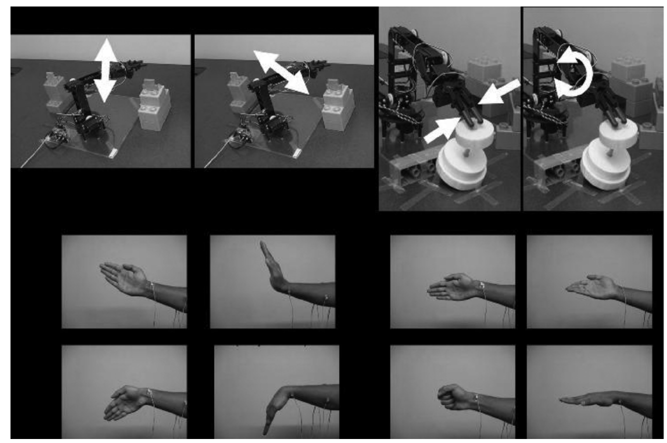

Fig. 7. Mapping between static hand gestures and degrees of control of the robotic arm. Each column shows the degree of freedom controlled, along with the two hand gestures that move it in either direction. (The gestures in the second column show a top–down perspective.).

arm. For control of prosthetic devices, one has the option of customizing the actions to better suit the device in question and the desired control, and such customization must be done on a case-by-case basis.

A. Online Experiments

1) Procedure: We had the same three subjects return for a second study and perform three real-time tasks of varying complexity with EMG-based control of the robotic arm. We retrained the classifier used for the online control system with the following prescription: The subjects were once again connected to the EMG recording device, five sessions of training data were recorded, and the SVM classifier was trained online with these five sessions and a parameter value that was recommended by the offline study. The process of collecting training data took 10 min, and the classifier was trained in less than a minute.

2) Task Selection: We chose three different tasks for our study—simple, intermediate, and complex. The metric used to quantify performance in each task was time to completion. The simple task was a gross-movement task in which the subject moved the robotic arm left and right to two specific locations in succession to knock off objects placed there. The goal was to test basic reach ability where fine control is not necessary. In the intermediate task, the robotic arm must be moved to a designated object, pick it up, and carry it over an obstacle, and drop it into a bin. In this task, accurate positioning of the arm is important, and additional degrees of freedom are needed to move the arm up and down and grasp and release the object. Fig. 8 describes the first two tasks in more detail. The third, complex,

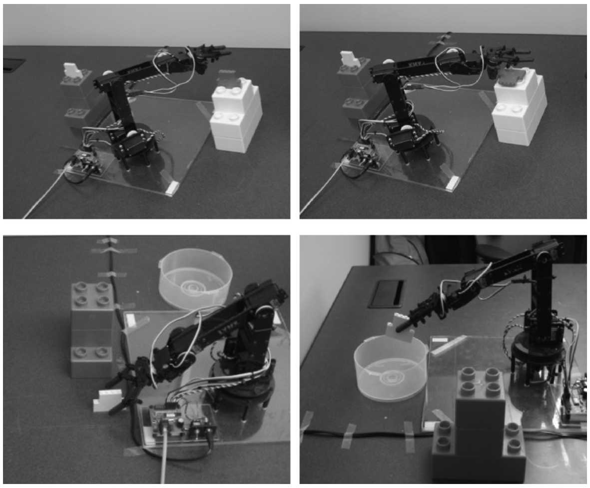

Fig. 8. Simple and intermediate online tasks. The first row shows the simple task, where the robot arm starts in the middle, and the goal is to topple the two objects placed on either side. The second row shows the intermediate task, where the goal is to pick up a designated object, carry it over an obstacle, and drop it in the bin.

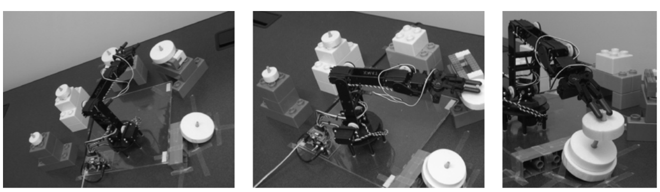

Fig. 9. Complex online task. Five pegs are placed at various fixed locations, and the goal is to stack them according to their size. The pictures show, in order, the initial layout, an intermediate step in moving a peg to the target location, and the action of stacking the peg.

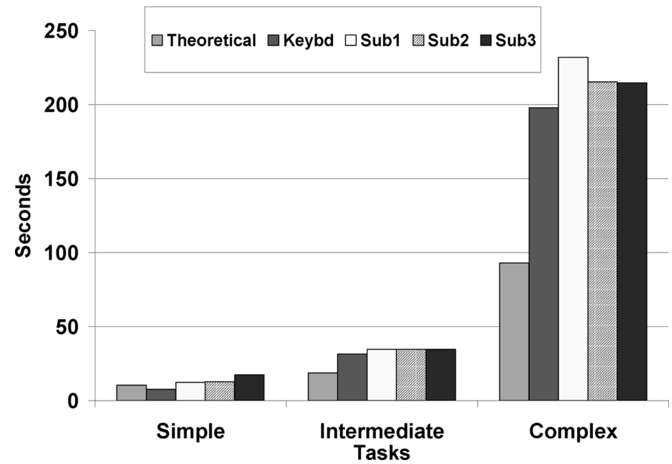

Fig. 10. Performance of three subjects using the EMG-based robotic arm control for three online tasks. The graph includes the baselines of theoretical time required, and time taken with a keyboard-based controller.

task involves picking up a number of pegs placed at chosen locations, and stacking them in order at a designated place. This requires very fine control of the robotic arm both for reaching and picking up objects, and also for placing them carefully on the stack, with an added degree of freedom to rotate the gripper, as shown in Fig. 9.

3) Measure and Baseline: The time to completion is used as the metric to assess task performance. Each subject performed each task three times, and the average time across trials was recorded. For the third task, only two repetitions were used since each task run was composed of four similar components. We use two baselines for comparison. The first of these is the theoretical time needed to perform these tasks by counting the number of commands needed for the robotic arm to perform a perfect sequence of operations, and assuming that the task is accomplished at the rate of 16 commands/s. The second was to have a fourth person perform these same sequence tasks with a keyboard-based controller for the robotic arm. This second baseline is more realistic, as it accounts for cognitive delays in planning various stages of a task as well as the time spent in making fine adjustments to the arm position for picking and placing pegs. There are differences in the skill of any given subject at task performance, as well as an important learning component where the same subject improves his or her performance at the task as he or she become accustomed to it. These issues are, however, peripheral to the scope of this paper, where the primary objective is to demonstrate that this type of EMG-to-command mapping can be robust and provide complex, intuitive control of a robotic arm for task performance in real time.

Algorithm: The control process used for the robotic arm is as shown in Fig. 6. The features and window lengths were the same as those used in the offline study. In addition, we use the probabilities returned by the classifier, along with a threshold, in order to discard commands that the controller is uncertain about. This is because the transitional periods when the user switches between different steady states may generate data that the classifier has not seen, and does not actually correspond to any of the chosen classes. Although we do not investigate this in our paper, we believe that by using a conservative threshold, the user can optimize his or her behavior to the classifier via feedback and produce more easily classifiable gestures.

B. Online Task Performance

Fig. 10 shows the performance of the three subjects and the baselines on the three online tasks. For the simple task, involving gross movements, all subjects take time close to the theoretical time required. The keyboard-based control takes less time, since the rate of keyboard control in our paradigm was faster than the EMG controller's control rate of 16 commands/s. For the intermediate task, where a moderate amount of planning and precision is required, the keyboard baseline is only slightly faster than the performance time of the three subjects.

Finally, for the complex task, it is interesting to note that the keyboard-based control also takes a comparable amount of time, thus showing that the bottleneck is not the control scheme (keyboard or EMG-classification-based control), but the task complexity and the performance of the EMG-based control regime do not add significantly to this.

C. Task Performance With Fewer Electrodes

For our final result, we present the performance of one subject on the three tasks as the number of electrodes used was dropped from seven to five, four, and three. Our offline analysis indicates that while there is redundancy in the electrode selection,

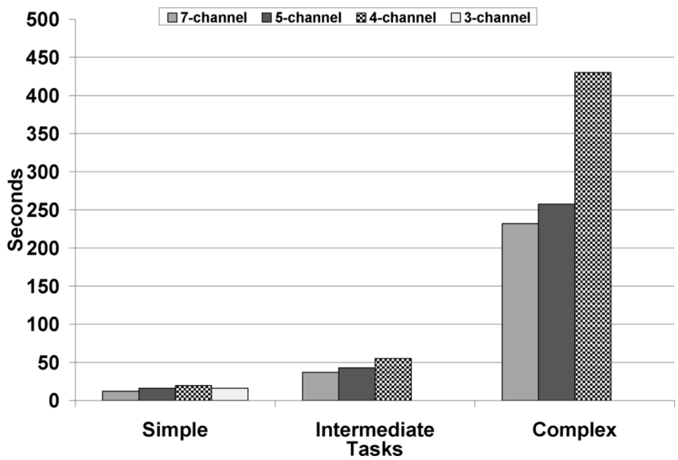

Fig. 11. Performance of subject 1 with fewer channels. Shown are the time taken by the subject on the three tasks with seven, five, four, and three electrodes. With three electrodes, the subject was unable to perform the intermediate and complex tasks.

the electrodes to drop are not consistent across subjects, and perhaps even across trials. Nevertheless, we drop two, three, and four electrodes, in order, based on the previously performed offline analysis of the subject's data. During the subject's online session, we trained four different classifiers with successively fewer channels of data. After the subject had successfully completed the tasks with the use of one classifier, we switched the classifier to the next in sequence, and the subject repeated the tasks with the new classifier in place.

Fig. 11 shows the performance of the subject on the three tasks. The results show clearly that with five, and even four channels, the subject was able to perform the tasks, although the complex task took significantly longer. With only three channels, however, the subject was no longer able to control the gripper and, thus, could not perform the intermediate and complex tasks. These data further support the robustness of our system, since the complex task could be achieved even with fewer electrodes, at the cost of efficiency.

VI. Discussion

This study established that a reliable SVM classifier-based technique could be used by individuals with intact forearm musculature to control a robotic arm in real time. While the implementation of these findings will be most useful for amputee individuals, we chose to first demonstrate its efficacy in individuals with intact forearm structures. Demonstration of the technique with amputee individuals will be useful as further proof of principle, but individual partial-limb amputee cases are each unique derivatives of the intact case. Residual muscle function will vary greatly between different amputee cases, and what is true for one amputee case will not generalize to another. Since our electrode positions were chosen anatomically to attempt to isolate individual muscles, (and, in turn, minimize the degenerate representation of a muscle across the electrode array), the reduction in electrode number simulates the reduction in musculature (shown in Fig. 5). Further, our method emphasizes learning and adaptation to the user's signals, allowing the EMG interface to be automatically tailored for an individual's musculature. This suggests that our system will be robust in the amputee setting, and ongoing studies will attempt to evaluate this hypothesis.

VII. Conclusion and Future Work

We have shown that EMG signals can be classified in real time with an extremely high degree of accuracy for controlling a robotic arm-and-gripper. We presented a careful offline analysis of an eight-class action classification problem based on EMG signals for three subjects as a function of the number of recording sites (electrodes) used for classification. Classification accuracies of more than 90% were obtained using a linear SVM-based classifier and a sparse feature representation of the EMG signal. We then demonstrated that the proposed method allows subjects to use EMG signals to efficiently solve several reasonably complex real-time motor tasks involving 3-D movement, obstacle avoidance, and pick-and-drop movements using a 4 degrees-of-freedom robotic arm.

Our ongoing work is focused on extending our results to other types of movements (e.g., discriminating finger movements). A separate effort is targeted toward replicating the results presented in this paper with actual amputees in collaboration with the rehabilitation department at our university. A parallel study [19] involves combining EEG signals from the scalp (reflecting underlying brain activity) with EMG signals for more accurate classification of motor patterns, with potential applications in brain–computer interfaces (BCIs). An interesting theoretical question that we are beginning to study is whether the EMG-based control system can be adapted online rather than only at the start of an experiment. This is a difficult problem since the subject is also presumably adapting online to generate the best muscle activation patterns possible for control and to compensate for changes in electrode conductivity with the passage of time. We intend to explore variations of our SVM-based classification technique to tackle this challenging nonstationary learning problem.

Acknowledgment

References

[1] B. Crawford, K. Miller, P. Shenoy, and R.P. N. Rao, "Real-time classification of electromyographic signals for robotic arm control," presented at the AAAI, Pittsburgh, PA, 2005.
[2] L. Eriksson, F. Sebelius, and C. Balkenius, "Neural control of a virtual prosthesis," presented at the ICANN, Skoevde, Sweden, 1998.
[3] A. D. Chan and K. Englehart, "Continuous myoelectric control for powered prostheses using hidden Markov models," *IEEE Trans. Biomed. Eng.*, vol. 52, no. 1, pp. 121–124, Jan. 2005.
[4] Y. Huang, K. Englehart, B. Hudgins, and A. Chan, "A Gaussian mixture model based classification scheme for myoelectric control of powered upper limb prostheses," *IEEE Trans. Biomed. Eng.*, vol. 52, no. 11, pp. 1801–1811, Nov. 2005.
[5] C. Deluca, *Surface Electromyography Detection and Recording*. Boston, MA: Boston Univ., Neuromuscular Res. Ctr., 1997.
[6] K. Engelhart, B. Hudgins, P. Parker, and M. Stevenson, "Classification of the myoelectric signal using time-frequency based representations," in *Med. Eng. Phys.*, 1999, vol. 21, pp. 431–438.
[7] M. Reischl, L. Groll, and R. Mikut, "Optimized classification of multiclass problems applied to EMG control of hand prostheses," presented at the ICANN, Budapest, Hungary, 2004.

[8] D. Nishikawa, W. Yu, H. Yokoi, and Y. Kakazu, "EMG prosthetic hand controller discriminating ten motions using real-time learning method," presented at the IEEE/RSJ IROS, Kyongju, Korea, 1999.
[9] F. Sebelius, M. Axelsson, N. Danielsen, J. Schouenborg, and T. Laurell, "Real-time control of a virtual hand," *Technol. Disabil.*, vol. 17, no. 3, 2005.
[10] R. Boostani and M. Moradi, "Evaluation of the forearm EMG signal features for the control of a prosthetic hand," *Physiol. Meas.*, vol. 24, no. 2, 2003.
[11] P. Ju, L. Kaelbling, and Y. Singer, "State-based classification of finger gestures from electromyographic signals," presented at the ICML, Stanford, CA, 2000.
[12] M. C. Carrozza, A. Persichetti, C. Laschi, F. Vecchi, R. Lazzarini, V. Tamburrelli, P. Vacalebri, and P. Dario, "A novel wearable interface for robotic hand prostheses," presented at the IEEE Int. Conf. Rehabilitation Robotics, Chicago, IL, 2005.
[13] K. Englehart and B. Hudgins, "A robust real-time control scheme for multifunction myoelectric control," *IEEE Trans. Biomed. Eng.*, vol. 50, no. 7, pp. 848–854, Jul. 2003.
[14] K. Englehart, B. Hudgins, and A. Chan, "Continuous multifunction myoelectric control using pattern recognition," *Technol. Disabil.*, vol. 15, no. 2, 2003.
[15] B. Scholkopf and A. Smola*, Learning with Kernels: Support Vector Machines, Regularization, Optimization and Beyond*. Cambridge, MA: MIT Press, 2002.
[16] C.-W. Hsu and C.-J. Lin, "A comparison of methods for multi-class support vector machines," *IEEE Trans. Neural Netw.*, vol. 13, no. 2, pp. 415–425, Mar. 2002.
[17] C.-C. Chang and C.-J. Lin, LIBSVM: A Library for Support Vector Machines 2001. [Online]. Available: http://www.csie.ntu.edu.tw/cjlin/libsvm.
[18] T.-F. Wu, C.-J. Lin, and R. C. Weng, "Probability estimates for multi-class classification by pairwise coupling," in *Proc. JMLR*, 2004, vol. 5, pp. 252–259.
[19] P. Shenoy and R. P. N. Rao, "Dynamic Bayesian networks for brain-computer interfaces," in *Advances NIPS*, 2005, vol. 17, pp. 1265–1272.

Pradeep Shenoy received the Master's degree in computer science from the University of Washington, Seattle, in 2004, where he is currently pursuing the Ph.D. degree in computer science.

His research interests include the application of machine learning techniques to the understanding of EEG and ECoG signals, with an emphasis on applications in brain–computer interfaces.

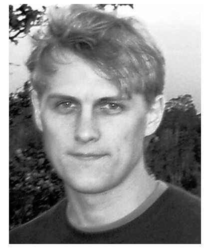

Kai J. Miller is currently pursuing the M.D./Ph.D. degree at the University of Washington, Seattle.

He is with the Departments of Physics, Computer Science, and Medicine at the University of Washington. His primary interest is the investigation of computational strategies employed by the nervous system in order to affect behavior. He has worked on the encoding of movement-related information in ECoG signals, and their application to localizing cortical representations of movement.

Beau Crawford received the B.Sc. degree in computer science from the University of Washington, Seattle.

His research interests include electromygraphic signal processing and brain–computer interfaces.

Rajesh P. N. Rao is an Associate Professor in the Computer Science and Engineering Department at the University of Washington, Seattle, where he heads the Laboratory for Neural Systems. he is the co-editor of two books *Probabilistic Models of the Brain* (2002) and *Bayesian Brain* (2007).

Prof. Rao is the recipient of a David and Lucile Packard Fellowship, an Alfred P. Sloan Fellowship, an Office of Naval Research Young Investigator Award, and a National Science Foundation Career award.