Gesture-Based Control and EMG Decomposition

Kevin R. Wheeler, *Senior Member, IEEE*, Mindy H. Chang, and Kevin H. Knuth

***Abstract*—This paper presents two probabilistic developments for the use with electromyograms (EMGs). First described is a neuroelectric interface for virtual device control based on gesture recognition. The second development is a Bayesian method for decomposing EMGs into individual motor unit action potentials (MUAPs). This Bayesian decomposition method allows for distinguishing individual muscle groups with the goal of enhancing gesture recognition. All examples presented rely upon sampling EMG data from a subject's forearm. The gesture-based recognition uses pattern recognition software that has been trained to identify gestures from among a given set of gestures. The pattern recognition software consists of hidden Markov models, which are used to recognize the gestures as they are being performed in real time from moving averages of EMGs. Two experiments were conducted to examine the feasibility of this interface technology. The first replicated a virtual joystick interface, and the second replicated a keyboard. Moving averages of EMGs do not provide an easy distinction between fine muscle groups. To better distinguish between different fine motor skill muscle groups, we present a Bayesian algorithm to separate surface EMGs into representative MUAPs. The algorithm is based on differential variable component analysis, which was originally developed for electroencephalograms. The algorithm uses a simple forward model representing a mixture of MUAPs as seen across multiple channels. The parameters of this model are iteratively optimized for each component. Results are presented on both synthetic and experimental EMG data. The synthetic case has additive white noise and is compared with known components. The experimental EMG data were obtained using a custom linear electrode array designed for this study.**

***Index Terms*—Bayesian decomposition, electromyogram (EMG), gesture recognition, hidden Markov model (HMM), motor unit action potential (MUAP).**

I. Introduction

ELECTROMYOGRAMS (EMGs) are used in the medical community to aid in the diagnosis of neuromuscular diseases, and there has been increasing interest in the use of EMGs as a means to interface with prosthetics and virtual devices [1]. For clinical applications, it is often necessary to use invasive needle electrodes to pinpoint sources of the EMGs to specific motor units. However, invasive measures are not ideal for use in the control of virtual devices. The ability to utilize surface EMG signals would enable the design of many neuroelectrically interfaced systems.

Manuscript received March 3, 2003; revised November 30, 2005. This work was supported by the National Aeronautics and Space Administration CICT/ITSR program under the Neuro-electric Machine Control project. This paper was recommended by Guest Editor S. W. Kercel.

K. R. Wheeler was with the Intelligent Systems Division, NASA Ames Research Center, Moffett Field, CA 94035 USA. He is now with the Monterey Bay Aquarium Research Institute, Moss Landing, CA 95039 USA (e-mail: kevinwheeler@ieee.org).

M. H. Chang is with the Department of Bioengineering, Stanford University, Stanford, CA 94305 USA (e-mail: mindyc@stanford.edu).

K. H. Knuth is with the Department of Physics, University at Albany (SUNY), Albany, NY 12222 USA (e-mail: kknuth@albany.edu).

Digital Object Identifier 10.1109/TSMCC.2006.875418

This paper introduces one such system in which surface EMG recordings from hand gestures are used in place of mechanical devices such as joysticks and keyboards to interface with a computer. Currently, most gesture recognition systems come in one of the following two forms.

- Gestures are recognized via an external camera, which requires sophisticated image processing and controlled lighting.
- Gestures are recognized by placing a sensing glove on the hand(s) of the participant.

We aim to achieve recognition in poor lighting conditions in extreme environments (outside of the laboratory) with minimal equipment. To date, we have accomplished this by directly connecting a person to the computer via EMG surface electrodes on the forearm. The EMG signals are sampled and digitized, and the resulting time series are passed through a pattern recognition system based upon hidden Markov models (HMMs). The recognized patterns are then transmitted as computer commands. Our first example of this was to attach four pairs of electrodes to one forearm and interpret the resulting EMG signals as joystick commands [2]. These commands were then used to fly a realistic flight simulator for a 757 transport aircraft. The acting pilot would reach out into the air, grab an imaginary joystick, and then pretend to manipulate this stick to achieve left and right banks and up and down pitches of the aircraft simulation. We also present results on pretending to type on a table (or lap) and translating the resulting sensed EMG signals into keystrokes.

The demonstration of gesture recognition through surface EMG signals leads to physiological questions of how individual sources are involved in generating these EMG signals that are distinctive for different types of movements. Voluntary limb movement occurs as a result of the brain generating a spike train that is transmitted through the nerve to a junction in the muscle known as the end-plate region. This induces an ion transfer along the length of the muscle fibers with a corresponding contraction of the muscles. The travelling waveform along the muscle fibers is known as a motor unit action potential (MUAP). This ion exchange induces a current on the surface of the skin, which can be measured as a voltage difference between two resistive electrodes. Surface EMGs measure a composite of the voltage changes produced by these individual MUAPs. Measuring from the surface of the skin presents additional complexities because the multiple MUAP sources mix as they traverse through skin, fat, muscle, and other tissues.

To separate the EMG signals into the corresponding fine motor muscle groups, we examine new ways to decompose the EMGs. Thus, the unmixed MUAPs can be used as an input to our HMMs. We present a Bayesian method to perform source separation for surface EMGs. In particular, compound motor unit action potentials (CMAPs) [3] are separated into representative MUAP waveforms. Our method is based on the differentially

1094-6977/$20.00 © 2006 IEEE

variable component analysis (dVCA) algorithm for source separation of electroencephalograms (EEG) developed by Knuth *et al.* [4], [5]. We have extensively modified this algorithm to work with surface EMGs.

In any standard Bayesian methodology, it is necessary to have a forward model and a means to optimize the parameterization of that model based upon data observations. In the case of EMGs, we have chosen to develop a model that describes the MUAPs and how they are mixed together. There has been extensive research on decomposing EMGs [6]–[9] using non-Bayesian approaches. There has also been great progress made in developing physics-based forward models for EMG signal generation as measured on the surface of the skin [10]–[14]. Unfortunately, in most of this literature, there is a gap between the methods of decomposition and model parameterization that could be bridged by following a Bayesian approach. In this paper, we detail the steps that we have taken to fill this gap with a simple mixing model.

II. Methodology

A. Gesture-Based Control

Each type of gesture set required a different methodology. The virtual joystick gesture set used four pairs of dry electrodes and four coarse grained movements. The virtual keyboard gesture set consisted of eight pairs of wet electrodes and 11 fine grained movements. The methodology that we followed consisted of the following steps:

1) gesture selection;
2) electrode application (location and number);
3) signal acquisition, filtering, and digitization;
4) feature formation;
5) pattern recognition model training and testing;
6) pattern recognition application in interactive simulation.

The process started by selecting the desired physical motions (gestures) to be used to control the virtual device. From the set of gestures, the best location for the limited number of electrode pairs (a maximum of eight in our case) was established. Then, standard signal processing practices were used to filter and digitize the signal. Transforms such as moving averages were applied to these raw digital data. The transformed data were fed into the pattern recognition software to train the models. Once the pattern recognition models were trained, they could be used for the real-time recognition task. Each of these steps will now be described in detail.

1) Gesture Selection: Our first task used coarse grained gestures to mimic manipulation of a joystick [2]. Movement of the joystick was associated with four basic gestures: up, down, left, and right. The use of the four pairs of electrodes for gesture recognition provided reasonable separation between the four gestures.

Our second task consisted of movements associated with typing on a number pad on the keys 0–9 and Enter. These movements consisted of much finer grained gestures. The first, second, and third fingers were resting over the keys, 4, 5, and 6 respectively. The first finger was used to press the keys 1, 4, and 7. The second finger was used to press the keys 2, 5, and 8. The third finger struck the keys 3, 6, and 9. The fourth finger was used for the Enter key, and the thumb was used to strike the zero key. In this case we used eight pairs of electrodes.

Fig. 1. (Top) Dry electrode sleeve for joystick-based flying. (Bottom) Wet electrodes for typing experiments.

2) Electrode Application: The placement of the electrodes depends upon the gestures that we wish to recognize and upon individual physiological differences. The joystick task was measured using four dry electrode pairs sewn into a sleeve as shown in Fig. 1. This sleeve helped to reduce variation in the placement of the electrodes. For the typing task, we chose to use eight pairs of wet electrodes due to the improved signal-to-noise characteristics of the wet electrodes over that of the dry electrodes. This was in part due to the signal amplitudes for the typing task being much smaller than that of the joystick task. The drawback of using wet electrodes is that the positions of the electrodes are difficult to replicate from one day to the next. The locations of these pairs were obtained by establishing a grid of electrodes on the forearm, and then performing the desired task; only electrodes that produced distinct signals for a gesture were used. The positions of the electrodes for the typing task were in two rings around the forearm: One near the wrist and one near the elbow also shown in Fig. 1.

Several tests were conducted to measure the effects of minor variations in placement (1–3 mm) and major displacements (1–2 cm). The minor variations had no impact but the major displacements required that the recognition models be retrained

or adapted for the individual user. Individual differences in personal physiology proved to be challenging. Differences in arm lengths and widths made it difficult to place the electrodes at the proper positions across people without considerable effort. In addition, strengths of the EMG signals varied across people and varied with the amount of training that individuals received.

3) Signal Acquisition, Filtering, and Digitizing: The EMG data were acquired by placing differential instrumentation preamplifiers near to each electrode pair with a common mode rejection ratio of 110 dB. All eight channel pairs were referenced to a common ground electrode positioned over the bone at the wrist or elbow. The signal was digitized using 16 bits at 6000 Hz, and then a 32-tap anti-alias bandpass Bessel filter was applied and down sampled to 2000 Hz unless otherwise indicated.

4) Feature Formation: The goal of the feature formation step is to separate the signals enough to allow the pattern recognition module to distinguish between gestures. Another result of working with features is to create a space smooth enough to be reliably modeled. We tried many common methods such as short time Fourier transform (STFT), wavelets, moving averages, and autoregression coefficients. At the end, moving averages, the simplest feature space, seemed to be the best. Since the EMG signals were differentially amplified, the average of the signals when presented with enough samples was approximately zero. This required that the moving average be performed on the absolute value of the signals. The windows used to form the moving averages were allowed to overlap by 75%. Note that this is purely an amplitude-based method; the frequency of the electrical activity did not seem to vary significantly from one gesture to the next.

5) Pattern Recognition: The pattern recognition method we chose to employ was an HMM. HMMs have been developed by the speech recognition community in response to their pattern recognition time-series problem [15]. The history of speech recognition reveals a process, which first attempted to recognize isolated words from a single speaker, and then isolated words from multiple speakers, followed by continuous words from a single speaker, and finally continuous words from multiple speakers. We are following a similar approach with our gesture recognition work. We have developed an isolated gesture recognition for both a single participant and multiple participants. The work in this paper will describe isolated recognition for a single typist and continuous recognition for the joystick study.

Two issues with training any model to learn from the sampled data are that the data sets are representative and the model has the appropriate number of parameters for an accurate representation. The training data sets can suffer from not having enough exemplars or being inconsistent for the sample size. In our case, we can always sample more data if we do not have enough. On an empirical basis, we have been able to use as few as 20 exemplars from each gesture to adequately model the remaining data from a single day. However, when we combine data from multiple days, it becomes readily apparent that inconsistency is a problem.

We define inconsistency as the statistics of the data varying from day to day. We could have defined this in terms of gesture-to-gesture variation but have chosen not to because this variation is more of a natural variation inherent in human behavior whereas the day-to-day inconsistencies are more an artifact of the experimental procedures.

There are many solutions to resolve this inconsistency as well as many contributions to the variations, which could be minimized. One example is electrode placement. If the electrode locations are allowed to vary from day to day, then the signal statistics will also vary. This can be reduced through the use of a fixed electrode sleeve.

Day-to-day variations related to natural behavior may not be removable, and in fact, we would benefit from modeling them. One example is the way by which people gesture may vary slightly from day to day even though their intention is to perform the gestures identically. In this case, we need to have enough data to represent the multimodal statistics and we need a way to adapt the system models over time. Our current methodology does not vary adaptively but it is our plan to include this in future work. This means that our best remedy is to recognize when day-to-day variation is too great for adequate model generalization. We can then use less data for training by using only the data similar to our current day's setup (i.e., electrode locations).

a) Training: The HMMs we used were continuous, tied mixture [16], and left-to-right models. The standard Baum–Welch training [15] was used. Models are classified as *continuous* if they use inputs that can take on a range of floating-point values. The alternative to this is to allow for only discrete values such as might be found if the inputs were transformed by quantization. *Tied mixtures* means that a fixed number of Gaussian mixtures are used throughout all of the states. Thus, any state may make use of any mixture. A *left-to-right* model means that the HMM may not go back to a previous state but may remain in a state or go to a new state.

Initialization of the models was performed using K-means clustering. The states were partitioned to equalize the amount of variance present within each state. The data sets used to train were segmented to insure that the peak of the variance was near the middle of each segment. This translated to the bulk of the energy being centered. Segments were sampled at 2000 Hz and contained 3072 samples per channel, with eight channels in total. The parameters of the HMMs that we typically varied were the number of discrete states, the number of Gaussian mixtures, the number of maximum number of iterations to train, the method used to arrive at the state partitioning (uniform versus variance based), and the method used to initialize the parameters of the mixtures (e.g., K-means clustering).

b) Recall: The real-time recall was performed using the standard Viterbi algorithm [17]. Since the system was processing streaming data, there was no knowledge as to where the peak of the variance was occurring. Because of this, the HMMs would see the data when the peak was first at the left most in the time segment, then the peak would move across from left to right, and then the final presentation was when the peak was at the right most part of the segment. Since the HMMs were trained only when the peak was centered, due to this shifting, the HMMs were required to recognize a gesture several times in a row before that gesture was selected as the one that was observed. This prevented spurious recognition when the peak was not near the center of observation.

6) Experiments: Two experiments were conducted to determine the feasibility of using bioelectric signals to substitute for, first, a joystick, and second, a keyboard.

The first experiment consisted of four pairs of dry electrodes fitted within a sleeve worn on the forearm of a participant. The participant was then asked to pretend to move a joystick left, right, up, and down. The participant performed each of these gestures 50 times. The data were separated by gesture, and segmented to have the peaks in the center of 3072 sample segments. Artifacts or incomplete gestures were removed from the data sets via manual inspection. The segmented data were then used to train four HMMs, one for each gesture. These trained models were then used to recognize gestures made on a day excluded from the training set. A confusion matrix was generated to display errors and to show which gestures were confused with one another. The system has also been used for numerous real-time demonstrations of flying a simulated 757 transport aircraft to landing [2]. A more continuous gesture recognition was implemented by decreasing the segment size.

Four methods were used to test the pattern recognition system. The first involved training the models on data from one day, and then recalling on different data obtained on the same day. We call this method *same trial acquisition and testing*. The second involved training on data from one day and recalling on data collected on a different day. We call this method *cross-trial acquisition and testing*. The third method trained on data subsampled from a large set taken across multiple days, and then recall performed on data different from the training but in the same large set. We call this method *multitrial acquisition and testing*. The final method involved training on a previously acquired single day that provided the best recognition in our real-time simulation for flying an aircraft. We call this *best trial training and real-time testing*.

The second experiment used eight pairs of wet electrodes in two rings of four each, one ring near the wrist, and the second near the elbow. The participant was asked to touch type on a printed picture of a number pad keyboard, striking the keys 0–9 and Enter. The participant was asked to type these in order, separated by a 1-s rest interval, for a total of 40 strokes on each key. These data were then segmented, and artifacts were manually removed. Data were collected on several different days. Eleven HMMs were trained, one for each gesture. These 11 models were then run in parallel during the recall.

B. EMG Decomposition

The model that we formulate for separating mixed MUAPs is dependent on how we acquire the data. Ideally, within a Bayesian framework, we would model every part of the system. We would start by modeling the sources of the potentials and observing how the shape of the potentials is changed by transmission through the tissue. This would be followed by a model of the electrodes, the amplifier, and finally the data acquisition card. The model we present relies on approximations to reduce the task of modeling all of these elements.

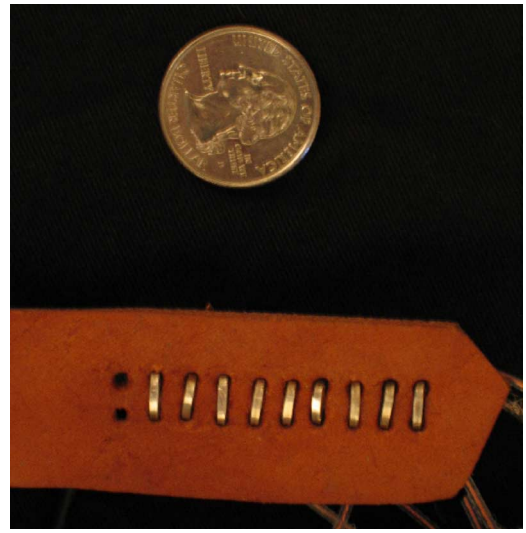

Fig. 2. Linear electrode array pictured with a U.S. quarter. (Color version available online at http://ieeexplore.ieee.org.)

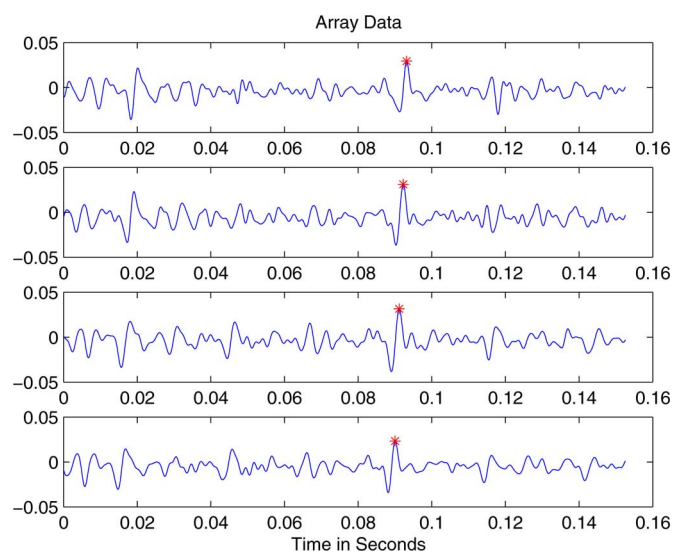

Fig. 3. EMG data from our linear electrode array. Star indicates moving MUAP over time between channels. (Color version available online at http://ieeexplore.ieee.org.)

Our model is based upon the assumption that we can observe compound MUAPs along parallel fibers of a muscle group. This assumption is facilitated by using a linear electrode array [10], [18] as shown in Fig. 2. We fabricated this electrode array with parallel silver bars spaced 5 mm apart. Fig. 3 shows the data collected by this device on four differential channels (eight total silver bars). Note that a star has been placed over one of the action potential waveforms, which is shifted between channels by an amount proportional to the conduction velocity. The muscle contraction under study was carefully controlled and can be assumed to be constant. The contraction level in this work is approximately 20% of the maximum voluntary contraction.

Our model representing the mixing process for the mth channel as a function of time can be expressed as

$$\psi_{m,t} = \sum_{n=1}^{N}\sum_{f=1}^{F} C_{mn}\alpha_{nf}s_n\left(t-(m-m_{\text{ref}})\tau_n^{\text{C}} - (f-1)\tau_n^{\text{F}} - \tau_{nf}^{\text{S}}\right) \quad (1)$$

where the subscripts index the nth component, fth firing, and mth channel; N is the total number of MUAP sources (components) being modeled; F is the number of firings; C represents the coupling between channels and sources; $s_n()$ is the source waveform; α_{nf} is the amplitude weighting; τ_n^{F} represents the time delay associated with the firing frequency of a particular source; τ_n^{C} is the delay across channels, which is proportional to the conduction velocity; and τ_{nf}^{S} is the latency for each source and firing representing the variability in firing.

There are several assumptions that underlie (1). We assume that the electrode array is positioned parallel to the muscle fibers and that the electrodes are evenly spaced. These assumptions allow us to model dominant components propagating along the muscle fibers as signals travelling from channel to channel. The time it takes to go from one channel to the next is represented by $\tau_n^{\text{C}} = d/v_c$, where $d = 5$ mm is the electrode spacing and v_{c} is the conduction velocity. We also assume that the muscle contraction is of constant force and that the sampling time is short enough that the firing rate (or time delay between firings τ_n^{F}) of any one MUAP source is effectively constant. Variation in the periodicity of the firing of a single source is modeled by τ_{nf}^{S} and is assumed small with respect to the firing rate.

The basis of model parameter estimation lies in using Bayes' theorem to maximize the *a posteriori* probability (MAP) of the model, using the likelihood of the data and the prior probability of the model parameters and other known information (symbolized by I)

$$p(\text{model}\,|\,\text{data}, I) = \frac{p(\text{data}\,|\,\text{model}, I)p(\text{model}\,|\,I)}{p(\text{data}\,|\,I)}. \quad (2)$$

Substituting the parameters of our model, this becomes

$$\begin{aligned} P &= p(C, s(t), \alpha, \tau^{\text{F}}, \tau^{\text{C}}, \tau^{\text{S}}\,|\,x(t), I) \\ &= \frac{p(x(t)|C, s(t), \alpha, \tau^{\text{F}}, \tau^{\text{C}}, \tau^{\text{S}}, I)p(C, s(t), \alpha, \tau^{\text{F}}, \tau^{\text{C}}, \tau^{\text{S}}|\,I)}{p(x(t)\,|\,I)} \end{aligned} \quad (3)$$

where the value on the left-hand side of the equation, which will be referred to as P, is the posterior probability of a model describing the data. The right side represents the product of the likelihood of data of the given model and the prior probability of the model, divided by a proportionality constant dependent on the data. A uniform distribution is assigned to the prior probabilities of each parameter, and as a result the posterior probability P becomes directly proportional to the likelihood of the data

$$P \propto p(x(t)\,|\,C, s(t), \alpha, \tau^{\text{F}}, \tau^{\text{C}}, \tau^{\text{S}}, I). \quad (4)$$

Using the principle of maximum entropy, the likelihood of the data is assigned a Gaussian distribution by introducing a new parameter σ. This parameter represents the expected squared error in prediction and is assigned a Jeffreys prior. When the likelihood is marginalized over all values of σ, the result becomes

$$P \propto (2\pi\sigma^2)^{-\frac{MT}{2}} \exp\left[-\frac{1}{2\sigma^2}Q\right] \quad (5)$$

where Q represents the square of the residuals between the data and our model, summed over all time points in all channels

$$Q = \sum_{m=1}^{M}\sum_{t=1}^{T}\left(x_m(t) - \sum_{n=1}^{N}\sum_{f=1}^{F} C_{mn}\alpha_{nf}s_n \times \left(t-(m-m_{\text{ref}})\tau_n^{\text{C}} - (f-1)\tau_n^{\text{F}} - \tau_{nf}^{\text{S}}\right)\right)^2. \quad (6)$$

To simplify calculations, we maximize P by maximizing the log of P. Using the method described by Knuth *et al.* [4], [5], the log of the posterior probability P can be written as

$$\ln P = -\frac{MT}{2}\ln Q + \text{const.} \quad (7)$$

For convenience of discussion, two expressions frequently used in the process of minimizing the difference between the data and the model are defined below. For a given component j in channel m at time t, U represents all firings of the component j deduced from the value of the actual data minus all other parameterized components

$$U(j,m,t) = x_m(t) - \sum_{\substack{n=1\\ n\neq j}}^{N}\sum_{f=1}^{F} C_{mn}\alpha_{nf}s_n \times \left(t-(m-m_{\text{ref}})\tau_n^{\text{C}} - (f-1)\tau_n^{\text{F}} - \tau_{nf}^{\text{S}}\right). \quad (8)$$

Similarly, the expression U_{F} isolates a particular firing f_0 of the jth component in channel m at time t, using the same method of deduction by also subtracting away all other firings of the jth component except for the f_0th firing

$$U_F(j, f_0, m, t) = U(j,m,t) - \sum_{\substack{f=1\\ f\neq f_o}}^{F} C_{mj}\alpha_{jf}s_j \times \left(t-(m-m_{\text{ref}})\tau_j^{\text{C}} - (f-1)\tau_j^{\text{F}} - \tau_{jf}^{\text{S}}\right). \quad (9)$$

1) Parameters: The five parameters optimized through iteration are described below.

a) Waveshape: The maximum *a posteriori* estimate of the waveshape is found by setting the partial derivative of the log probability with respect to a time point q in waveshape s_j to zero (details are given by Knuth [4]).

b) Amplitude: When taking the partial derivative of the log probability with respect to the amplitude of the f_0th firing of the jth component, the optimal estimate for the amplitude of

this particular firing becomes

$$\hat{\alpha}_{jf_0} = \frac{\sum_{m=1}^{M}\sum_{t=1}^{T} U_{\mathrm{F}} R_\alpha}{\sum_{m=1}^{M}\sum_{t=1}^{T} (R_\alpha)^2} \tag{10}$$

$$R_\alpha = C_{mj} s_j \left(t - (m - m_{\mathrm{ref}})\tau_j^{\mathrm{C}} - (f_0 - 1)\tau_j^{\mathrm{F}} - \tau_{jf_0}^{\mathrm{S}}\right). \tag{11}$$

Since the model allows for varying amplitudes between different firings of the same component, each α_{jf} term is determined irrespective of other firings by using the deduced single firing term U_{F}.

c) Firing period: To find the optimal estimate for the firing period of the jth component, one must solve

$$\hat{\tau}_j^{\mathrm{F}} = \mathrm{argmax} Y\left(\tau_j^{\mathrm{F}}\right) \tag{12}$$

$$Y\left(\tau_j^{\mathrm{F}}\right) = \sum_{m=1}^{M}\sum_{t=1}^{T} U(j,m,t) U\left(j,m,t+\tau_j^{\mathrm{F}}\right) \tag{13}$$

where the function U is defined in (8). The function $Y(\tau_j^{\mathrm{F}})$ represents the autocorrelation across each channel, summed across all channels for all firings of a given component j. The jth component is isolated by subtracting away all firings of all other components to obtain $U(j,m,t)$. Each channel is multiplied with shifted versions of itself, and assuming that the data are periodic across each channel, the latency shift that produces the maximal value will be where the second through Fth channel is closest to alignment with the first through $(F-1)$st firing of the jth component. This latency estimate is constrained to be positive and greater than 2 ms because a firing period that is significantly smaller than the time span of a single action potential is not physiologically plausible.

d) Conduction period: To find the optimal estimate for the conducting period of the jth component

$$\hat{\tau}_j^{\mathrm{C}} = \mathrm{argmax} Z(\tau_j^{\mathrm{C}}) \tag{14}$$

where

$$Z(\tau_j^{\mathrm{C}}) = \sum_{m=1}^{M}\sum_{t=1}^{T} U(j,m-1,t) U\left(j,m,t+\tau_j^{\mathrm{C}}\right) \tag{15}$$

where the function U is defined in (8). The function $Z(\tau_j^{\mathrm{C}})$ represents the cross-correlation between consecutive channels, summed across all pairs of channels for a given component j. As above, the estimate of the jth component (with all of its firings) is obtained by subtracting away all firings of all other components to obtain $U(j,m,t)$. As a convention, the first channel will be considered the reference electrode from which action potentials are first detected. As the action potentials propagate through the muscle fibers, each subsequent channel detects the action potential slightly later than the previous channel. For each pair of channels, the first channel is multiplied with shifted versions of the second channel. Assuming that the data are periodic across each channel, the latency shift between the channels that produces the maximal value will be where all firings of the jth component most nearly align between the pair of channels. The conduction period is constrained to positive values between zero and half of the firing period. Since the conduction period is significantly smaller than the firing period, we are able to apply this assumption.

e) Offset latency: To find the optimal estimate for the offset latency of f_0th firing of the jth component

$$\hat{\tau}_{jf_0}^{\mathrm{S}} = \mathrm{argmax}\ A\left(\tau_{jf_0}^{\mathrm{S}}\right) \tag{16}$$

$$A(\tau_{jf_0}^{\mathrm{S}}) = \sum_{m=1}^{M}\sum_{t=1}^{T} U_{\mathrm{F}}(j,f_0,m,t) V_{\mathrm{F}}\left(j,f_0,m,t+\tau_{jf_0}^{\mathrm{S}}\right) \tag{17}$$

where the function U_{F} is defined in (9) and $V_{\mathrm{F}}(j,f_0,m,t)$ represents the reconstruction of the f_0th firing of the jth component, using all other parameters of the jth component

$$V_{\mathrm{F}}(j,f_0,m,t) = C_{mj}\alpha_{jf_0} s_j \left(t - (m - m_{\mathrm{ref}})\tau_j^{\mathrm{C}} - (f_0 - 1)\tau_j^{\mathrm{F}} - \tau_{jf_0}^{\mathrm{S}}\right). \tag{18}$$

The function $A(\tau_{jf_0}^{\mathrm{S}})$ represents the cross-correlation between the deduced single firing of the jth component, $U_{\mathrm{F}}(j,f_0,m,t)$, based on the data after removing the firings of the other components, and the estimated single firing of the component, $V_{\mathrm{F}}(j,f_0,m,t)$, based on the jth component parameters. The deduced firing is multiplied with shifted versions of the reconstructed firing, and the latency which produces the maximal value is taken as the estimate of the offset latency.

2) Adjustments for Parameter Degeneracies

f) Latency degeneracy between τ_n^{F} and τ_{nf}^{S}: Since τ_n^{F} and τ_{nf}^{S} both represent time shifts within single channels of data, if the estimated value for τ_n^{F} is inaccurate, τ_{nf}^{S} values will increase linearly in amplitude. In other words, if the estimated τ_n^{F} value is smaller than the actual value, each successive firing will deviate from its estimate by a value larger than the previous firing with its respective estimate. For a given firing f of a given component n, the value of the net latency due to firing and offset is not affected, but the τ_n^{F} and τ_{nf}^{S} values no longer represent the firing period and offset period. To correct for this offset, every time the τ_{nf}^{S} is calculated, a linear regression on the τ_{nf}^{S} values is performed, and both the latency values are adjusted accordingly. The linear regression takes the form of $\tau_{nf}^{\mathrm{S}} = \mu_n f + \beta_n$ Using this line, the τ_n^{F} and τ_{nf}^{S} values are remapped so that τ_{nf}^{S} becomes a constant value plus or minus deviations from the regression line d_{nf} and τ_n^{F} accounts for this change

$$\begin{aligned}(f-1)\tau_n^{\mathrm{F}} + \tau_{nf}^{\mathrm{S}} &= (f-1)\tau_n^{\mathrm{F}} + (\mu_n f + \beta_n + d_{nf}) \\ &= (f-1)\tau_n^{\mathrm{F}} + (\mu_n(f-1) + \mu_n + \beta_n + d_{nf}) \\ &= (f-1)(\tau_n^{\mathrm{F}} + \mu_n) + (\mu_n + \beta_n + d_{nf}).\end{aligned}$$

The adjusted τ_n^{F} value becomes

$$\bar{\tau}_n^{\mathrm{F}} = \tau_n^{\mathrm{F}} + \mu_n \tag{19}$$

and the τ^{S}_{nf} value becomes

$$\begin{aligned}\bar{\tau}^{\mathrm{S}}_{nf} &= \mu_n + \beta_n + d_{nf} \\ &= \mu_n f + \beta_n + d_{nf} - (f-1)\mu_n \\ &= \tau^{\mathrm{S}}_{nf} - (f-1)\mu_n. \end{aligned} \tag{20}$$

g) Waveshape alignment and adjustment of τ^{S}_{nf}: A degeneracy also occurs in the time domain between the waveshape and the offset latency τ^{S}_{nf}. A time shift in the component could be characterized as either a change in waveshape or a shift in offset latency. To give the offset latency values, a relative meaning between different components, each time a stopping condition is met, as in Step 9) of the iteration process below, the peaks of all component waveshapes are aligned to match the waveshape with the earliest peak, and each τ^{S}_{nf} value is adjusted accordingly. This alignment is performed after each stopping condition to ensure that the algorithm has had a chance to estimate all parameter values before shifting all waveshapes to an earlier time. Performing the alignment during the iterative process runs the danger of shifting parts of the waveshape out of the time-frame allotted for a single waveshape into negative time, which would be invalid for this model.

2) Iterations: We optimized the parameters as follows.

Step 1) Identify the total number of firings within the dataset by human observation.

Step 2) Estimate τ^{F}_n for the component using (12) and (13).

Step 3) Estimate τ^{C}_n using (14) and (15).

Step 4) Estimate τ^{S}_{nf} using (16) and (17).

Step 5) Adjust τ^{F}_n values if τ^{S}_{nf} values show a linear rate of change using (19) and (20).

Step 6) Estimate the waveshape $s()$

Step 7) Estimate the amplitude α of each firing using (10).

Step 8) To parameterize another component, follow Steps 1)–5), using the data from which the model of the first component has been subtracted.

Step 9) Iterate through Steps 1)–5) for both components until the average change in waveshapes from the previous iteration is less than 1% or until a maximum number of iterations has been performed, and align the peaks of the component waveshapes, adjusting τ^{S}_{nf} accordingly.

Step 10) For each additional source, parameterize the new component based on the data without all other components that are already modeled, and repeat the iteration of Steps 1)–5) for all components until a stopping condition in Step 9) is reached.

Listed below are considerations used in determining the aboveiteration order.

a) Parameter initialization: Since the general waveshape of an MUAP is fairly well defined, this information is used to initiate the waveshape. The point values in $s(t)$ are determined in the manner described in Section II-B4. As mentioned earlier, the coupling matrix is set to all ones under the assumption that all detectors receive signals from all components equally well. All latency values τ^{F}_n, τ^{C}_n, and τ^{S}_{nf} are initialized as zero, and all α_{nf} values as one. Since the latency optimizations occur first in the iteration in Steps 2)–4), α_{nf} values are initialized as one. If the α_{nf} values were zero, the reconstructed component used in the τ^{S}_{nf} calculation would just be a straight line at zero. For this reason, if any value of α_{nf} becomes zero in the process of iteration, the α_{nf} value is set to one temporarily for the calculations of τ^{F}_n, τ^{C}_n, and τ^{S}_{nf} and then set back to zero.

b) Estimation of τ^{F}_j and τ^{C}_j before τ^{S}_{jf}: When considering a single component, τ^{F}_j and τ^{C}_j are parameters that characterize the component, providing information about the firing rate and conduction velocity when the distance between the electrodes is known. The optimizations of τ^{F}_j and τ^{C}_j involve correlations of the deduced component and are not directly dependent on the accuracy of the parameters of the component in question. On the other hand, τ^{S}_{jf} "picks up the slack" in the overall latency value and is restrained to be a constant with slight deviations for each firing. Effectively, this constant gives information about the relative offset between different components, and the deviation represents the time error between the model and the actual data for each firing of this particular component. τ^{F}_{jf} involves the cross-correlation of the deduced component and the reconstructed component and, therefore, is directly dependent on the accuracy of the jth component parameters.

In determining the order of this algorithm, the τ^{S}_{jf} calculation is performed after τ^{F}_j and τ^{C}_j, and so the offset calculation has the benefit of using the already parameterized firing and conduction period values when reconstructing the jth component for cross-correlation. When the latency values are remapped in Step 5), the firing period adjustment is applied to a τ^{F}_j value that represents an estimate of the firing period rather than an initial value with no significance.

c) Estimation of the waveshape before amplitude: Due to the degeneracy that could occur in the model between the scaling of the waveshape and the amplitude, the waveshape is constrained to have a peak-to-peak amplitude of one to give the α values a consistent meaning. Thus, in each set of iterations, the waveshape is estimated first and scaled peak to peak. The amplitude α is parameterized after the waveshape has been determined so that α can appropriately compensate for the waveshape scaling in each firing of the component in question.

d) Isolated optimization of new component on its first iteration: When parameterizing a new component j, all previous components have already been optimized to their stopping condition. In all parameter optimization calculations [Steps 1)–5)], the deduced component is used, in the form of either a single firing or an action potential train. For the first iteration, since the jth component has not yet been completely parameterized, using this component in calculations for other components may throw off parameter values unnecessarily. Therefore, for the first iteration of a new component j, all the parameters of j are optimized. For successive iterations, all parameters are optimized for each component in turn.

3) Synthetic Data: The waveshape of the synthetic data used for testing is based on the MUAP model developed by McGill *et al.* [19]. The source function $V'(t)$ was created by the sum of

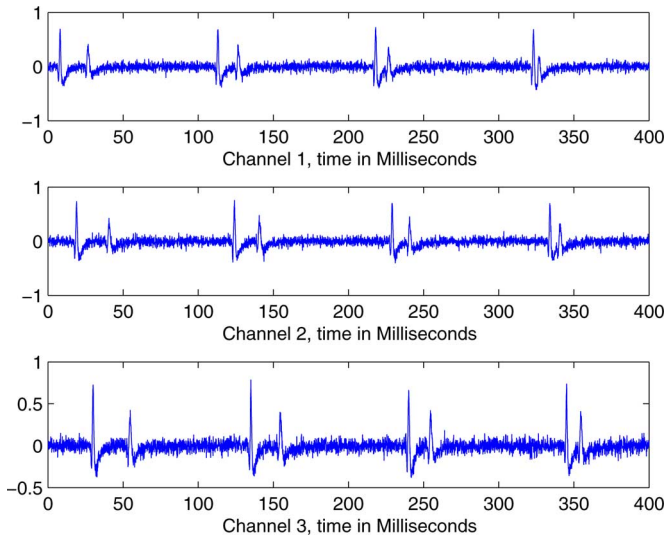

Fig. 4. Synthetically generated CMAPs with two components and white noise added. (Color version available online at http://ieeexplore.ieee.org.)

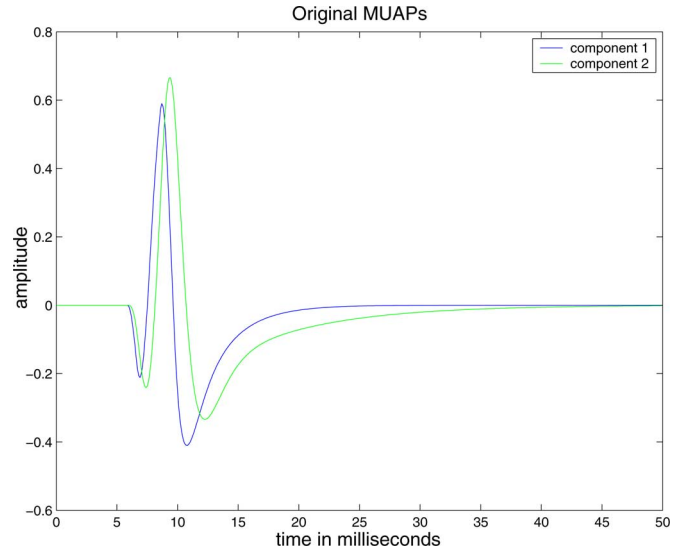

Fig. 5. Synthetic MUAPs. (Color version available online at http://ieeexplore.ieee.org.)

a scaled spike and afterpotential

$$\frac{dV(t)}{dt} = (ag'(t) + bg(t)) - \frac{b}{t_A} g(t) * (e^{-t/t_A} u(t)) \quad (21)$$

where $*$ denotes a convolution, a and b are scaling factors, t_A is a time constant of decay, and

$$g(t) = \frac{k^{n+1}}{\Gamma(n)} t^n e^{-kt} u(t) \quad (22)$$

in which n and k are adjustable constants. The standard values used were $n = 2.5$ and $k = 5.8$ [19]. This source function was used as the initial estimate of all new component waveshapes. The model described by McGill *et al.* also details spatial and temporal weighting functions to be convolved with the source function representing the waveshape distortion as it travels along the muscle fibers, as well as considerations for different lengths of muscle fibers. These factors were not implemented for this paper, but a convolution was performed using an approximate weighting function in generating the synthetic data. For simulation, varying levels of Gaussian noise were added as well. The data shown in Fig. 4 were generated with the two components shown in Fig. 5 that were mixed together with uncorrelated white noise with a signal-to-noise ratio of 3.7. This level of noise is higher than that normally observed in our experimental setting, and thus, is representative of a more difficult test.

4) Experimental Data: The subject EMG data were acquired using our electrode array positioned over the bicep. The subject was required to lift and hold a 5-lb weight and was only allowed to bend at the elbow, with the elbow supported. The data were sampled at 32 kHz using a custom-built amplifier with a gain of 1000 and an anti-aliasing filter with a 3-kHz cutoff frequency.

TABLE I
CONFUSION MATRIX FOR CROSS-TRIAL JOYSTICK DATA

Gesture	Left	Right	Up	Down	Correct
Left	15	0	26	9	30%
Right	0	50	0	0	100%
Up	0	0	50	0	100%
Down	0	0	1	49	98%

III. RESULTS

A. Joystick Gestures

1) Same Trial Acquisition and Testing: This experiment is by far the easiest to recognize because the variability associated with day-to-day differences has been eliminated. Such a variation includes conductivity levels of the skin, positioning of the dry electrode sleeve, and changes in the performance of the gestures. We noticed that we could determine when participants had used skin moisturizer before the experiment because the signal quality obtained from the dry electrodes improved. Typically, if the HMMs had an appropriate number of parameters and enough data were used, then no errors were made upon recall of the validation set.

2) Cross-Trial Acquisition and Testing: This experiment demonstrated which gesture was the hardest to separate from the others. In particular, we trained on 50 instantiations from one trial date and then validated on 50 other instantiations from a different trial date. A typical confusion matrix is shown in Table I. This indicates that day-to-day variations were significant enough to cause difficulty in separating the gesture left from other gestures. The source of the variations included electrode placement, length of the gesture, strength of gesture formation, and the form of the gesture (wrist angles). This led to the next experiment to see if the models would generalize if we trained on all of the different days together and then tested on a withheld subset.

TABLE II
CONFUSION MATRIX FOR CROSS-TRIAL SHORT JOYSTICK DATA

Gesture	Left	Right	Up	Down	Correct
Left	17	0	25	8	34%
Right	0	49	0	1	98%
Up	0	0	50	0	100%
Down	0	0	0	50	100%

3) Multitrial Acquisition and Testing: To determine the generalization capability of the HMMs, we trained on data from multiple trial dates and then recalled on points withheld from the training data for the same dates. This resulted in perfect results (100% correct for all gestures).

Of course, this does not mean that when data from yet another new day are added, the system will be able to generalize on that data. In the next experiment, we examine training on the single best day and use that for real-time testing. The real-time testing acts as new and unseen data.

4) Best-Trial Training and Real-Time Testing: The error rates determined from the previous methods were not necessarily indicative of a real-time performance. In particular, the error rates would vary across time depending upon many factors such as sleeve position (rotation), sweating, skin moisture (dry skin does not conduct well), length of time that the electrodes were worn, and fatigue (resulting in tremors). By training on only a single day's data, we were able to use the dry electrode sleeve for demonstrations on many different days. We selected the day that gave the best real-time reliability. The demonstrations consisted of flying a 757 transport aircraft in simulation to landing at San Francisco airport.

5) Continuous Recognition: In the previous experiments, a total of 3072 data samples were used to form the estimate. This introduced considerable time lag into the system (1.5 s). In an attempt to become closer to a continuous recognition process, the time segments used to train the HMMs were shortened to only contain the first part of the rise of the signal using 352 samples. In this case, the HMMs consisted of three states with nine mixtures in total. The resulting cross-trial confusion matrix is shown in Table II.

This matrix is not significantly different from the previous cross-trial longer data. The change in signal length allowed us to remove noticeable delays between the gesture action and the movement of the aircraft. We achieved a much less noticeable delays (176 ms) at the expense of a slight decrease in the robustness of the gesture recognition process. It is possible to have this response time with a small decrease in the recognition rates. Since this is intended for real-time systems, such a lag is hard to justify but it has not prevented us from successfully flying the simulated aircraft. The resulting multiday confusion matrix had no errors.

In the next set of experiments, we switched from using the dry electrode sleeve to wet electrodes. This was necessary because the EMG signals measured for the typing gestures were much smaller than those for the joystick. The wet electrodes tend to have a higher signal-to-noise ratio than the dry electrodes.

TABLE III
MULTITRIAL CONFUSION MATRIX FOR TYPING DATA

	1	2	3	4	5	6	7	8	9	%
1	46	0	0	4	0	0	1	0	0	90
2	0	48	0	0	0	0	0	3	0	94
3	0	0	49	0	0	1	1	0	0	96
4	11	0	0	38	2	0	0	0	0	75
5	1	3	0	5	36	1	3	2	0	71
6	0	1	6	0	0	42	0	0	2	82
7	0	0	0	0	0	0	51	0	0	100
8	0	0	0	0	2	1	3	44	1	86
9	0	0	0	0	0	0	0	0	51	100

B. Keyboard

In this experiment, the position of the hand above the simulated number pad was maintained in a touch-typist typing position. If the position of the hand were allowed to vary, the tasks of distinguishing between hitting the top row of keys from the bottom row of keys would greatly increase in difficulty and would require electrodes on the upper arm to sense the movement. The angle of the participant's wrist also had to be carefully maintained to avoid radically changing the sampled signals. Even with careful attention to position and maintaining electrode placement from day-to-day, the data tended to vary. There are consistent trends in the data but the variation between instantiations of the same gesture is great.

1) Multitrial Acquisition and Testing: The keyboard replication experiments had a much greater daily variation in electrode placement than with the joystick. We also had difficulty in reliably having the participants maintain a consistent hand position from trial to trial. This included ensuring that the wrist angle was similar and that the hand was consistently neither resting on the table during motion nor was in part supported by the table (i.e., bad form, but consistent bad form). Despite these uncontrollable variations, the resulting confusion matrix for multiple trials shown in Table III is remarkably good. The variability in the data caused our models to generalize gestures such that more confusion occurred. For live demonstrations, we needed to train on the same day we were giving the demonstration, and thus, used only a single day's data.

Two enhancements are planned. First, the use of wet electrodes caused unintentional misplacement. We are currently developing new dry electrode straps, which have a higher density than the sleeve and are similar in size to a wrist band. These straps will allow us to have the electrodes on a band positioned relative to each other without variability. The second enhancement is to include model correcting adaptation, which is now common in the speech recognition community. This adaptation would allow the models to be tuned to small variations, both throughout the day and whenever the current day's configuration differs from the models used to train. A calibration stage will be included so that the participant can make a gesture to issue a certain command and the computer will adapt to understand the signals as that command. Calibration will eliminate the need to require a participant to learn a fixed set of gestures. Instead, the person will be able to perform a gesture that seems natural to him or her to accomplish a given task, and the computer will simply map those signals to the correct action.

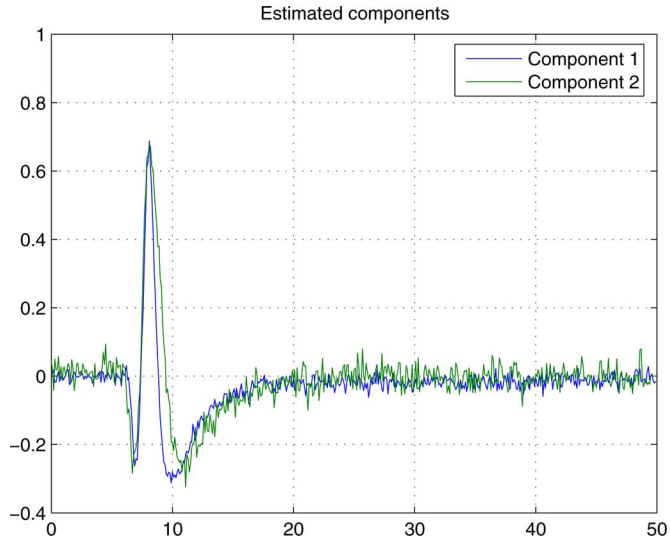

Fig. 6. Recovered components from synthetic data. (Color version available online at http://ieeexplore.ieee.org.)

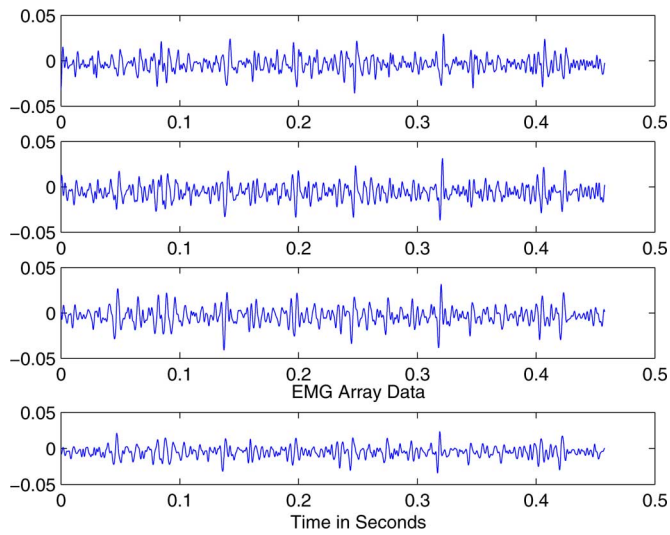

Fig. 7. Experimental EMG data. (Color version available online at http://ieeexplore.ieee.org.)

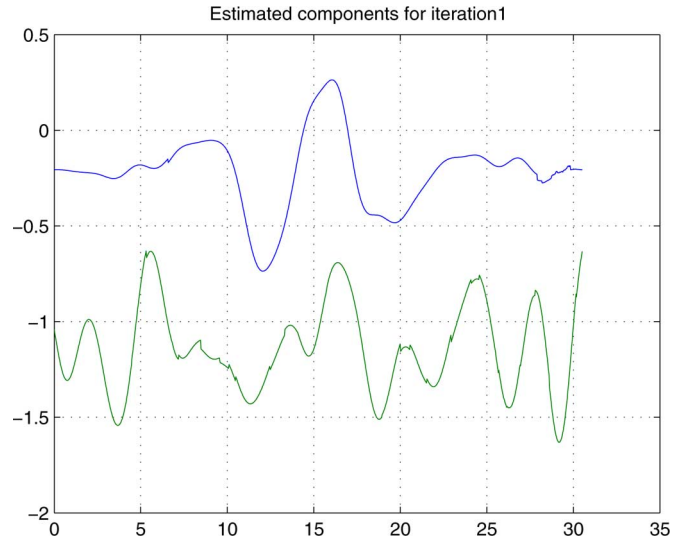

Fig. 8. Two components separated from experimental data. (Color version available online at http://ieeexplore.ieee.org.)

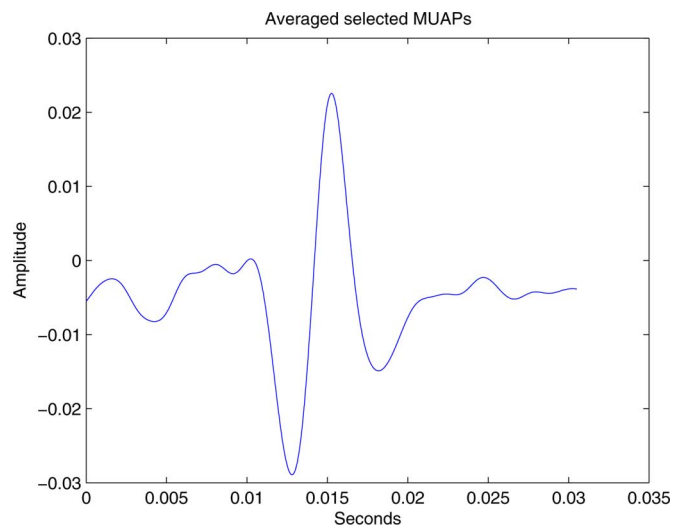

Fig. 9. Average waveform of hand-selected MUAPs. (Color version available online at http://ieeexplore.ieee.org.)

C. Muap Analysis

1) Synthetic Data: We applied the algorithm to 20 trials of synthetic data (see 4). The data consisted of three channels with only four firings across the channels. A typical decomposition is shown in Fig. 6. This resulted in a median RMS error of 0.0297.

2) Experimental Data: The experimental data, which are partially shown in Fig. 7, were decomposed into two components shown in Fig. 8. Since these were real data, we do not have the actual MUAPs to compare our decomposition, so instead we compare this with a method in which clean MUAPs were hand picked and then averaged together. This averaged MUAP waveform is depicted in Fig. 9, which illustrates that there is a qualitative agreement between the expected MUAP and the first component discovered using this algorithm. The second component contains multiple compound MUAPs because only two components were specified to be calculated.

IV. Conclusion

We have shown that it is feasible to control virtual devices via noninvasive EMG signal monitoring without explicitly modeling the EMG signals. Using HMMs for pattern recognition, we have demonstrated the ability to replicate movements associated with both joysticks and keyboards. We chose to demonstrate these input devices because they are familiar to the general computer user. Ultimately, improved interfaces will consist of more natural movements than those associated with either joysticks or keyboards. When our online adaptation software has been completed, a person will be able to make the gesture he feels is natural for a given task (assuming that we have enough electrodes to cover the muscles involved) and the computer will map from this "natural" signal space to the expected "computer command" space.

Although the dry electrode sleeve guaranteed that the electrodes would be positioned consistently relative to each other, there was no guarantee that the sleeve would be exactly at the same location on the arm. It was our hope that the sleeve would minimize the day-to-day variations, but it needs to be redesigned to assure positioning. We would also like to increase the number of electrodes to allow for spatial oversampling. The dry electrode sleeve can suffer from intermittent conductivity problems when the impedance between a dry electrode and the skin becomes temporarily elevated due to hairs lifting the sensors or variations in the moisture level of the skin. In this study, the decrease in conductivity could usually be identified and fixed by a manual readjustment of the sensors, but this inconsistency is not acceptable for daily use.

Ultimately, we envision a variety of applications for this work. The ability to naturally interface with a computer allows humans to manipulate any electrically controlled mechanical system. In addition to wearable computing applications, we are also examining interfaces to robotic arms, mobile robots for urban rescue, unmanned aircraft drones, robotic exoskeletons, and space suit interfaces. There are also side benefits to using EMG signals for control in long duration space missions. One of the side effects of living in a zero-gravity environment for extended periods is muscle atrophy. It would be possible to have astronauts train during a long flight to a distant planet by simulating the motions necessary to accomplish a given task. The EMG signals generated from these motions could be analyzed. If significant variations were detected, the astronauts could be given advance warning to change their training routines to minimize atrophy and ensure mission success.

The drawback of using moving averages of EMGs as input to the HMMs for gesture recognition was that the individual sources of EMGs could no longer be distinguished. To distinguish smaller muscle group activations and inturn improve recognition, we have presented a dVCA algorithm for EMG decomposition based on the Bayesian methodology. The original dVCA algorithm designed for EEG was substantially modified for the purpose of separating compound MUAPs measured in the surface EMGs. This modified algorithm was demonstrated on both simulated and real EMG data. The results are encouraging for both the synthetic and real cases. The most flexible part of this algorithm is that it allows the waveform $s(t)$ to vary over time. Letting the waveshape be pointwise estimated permits it to have any shape, even shapes that are not at all physiologically plausible. This flexibility was deliberate to determine whether the algorithm would discover waveshapes resembling expected MUAP shapes. Indeed, we were pleased to see that the discovered components do resemble the synthesized MUAPs. We have obtained MUAPs from surface electrodes that are remarkably similar to our expected waveform without imposing any knowledge of what we were expecting to see in terms of shape.

The EMG decomposition method shown is a simple forward mixing model, which could eventually be replaced with a much more complicated physics-based model such as the electromagnetic model of [12]. This would then allow for the automated determination of the representative tissue properties via proper model parameterization at the expense of more complex optimization.

ACKNOWLEDGMENT

The authors wish to thank the NASA IS/IDU program for support of the initial algorithm development.

REFERENCES

[1] K. R. Wheeler and C. C. Jorgensen, "Gestures as input: Neuroelectric joysticks and keyboards," *IEEE Pervasive Comput.*, vol. 2, no. 2, pp. 56–61, Apr.–Jun. 2003.

[2] C. Jorgensen, K. Wheeler, and S. Stepniewski, "Bioelectric flight control of a 757 class high fidelity aircraft simulator," presented at the World Automation Congr., Maui, HI, 2000.

[3] R. Plonsey and D. G. Fleming, *Bioelectric Phenomena*. New York: McGraw-Hill, 1969.

[4] K. H. Knuth, A. S. Shah, W. A. Truccolo, S. L. Bressler, M. Ding, and C. E. Schroeder. "Differentially variable component analysis (dvca): Identifying multiple evoked components using trial-to-trial variability," *J. Neurophysiol.*, vol. 95, pp. 3257–3276, May 2006.

[5] K. H. Knuth, W. A. Truccolo, S. L.Bressler, and M. Ding, "Separation of multiple evoked responses using differential amplitude and latency variability," in *Proc. 3rd Int. Workshop Independent Component Anal. Blind Signal Separation*, T. W. Lee, T. P. Jung, S. Makeig, and T. J. Sejnowski, Eds., 2001, pp. 463–468.

[6] K. C. McGill, "Optimal resolution of superimposed action potentials," *IEEE Trans. Biomed. Eng.*, vol. 49, no. 7, pp. 640–650, Jul. 2002.

[7] D. W. Stashuk, "Decomposition and quantitative analysis of clinical electromyographic signals," *Med. Eng. Phys.*, vol. 21, no. 6–7, pp. 389–404, Jul.–Sep. 1999.

[8] D. Zennaro, P. Wellig, V. M. Koch, G. S. Moschytz, and T. Laubli, "A software package for the decomposition of long-term multichannel EMG signals using wavelet coefficients," *IEEE Trans. Biomed. Eng.*, vol. 50, no. 1, pp. 58–69, Jan. 2003.

[9] J. Fang, G. C. Agarwal, and B. T. Shahani, "Decomposition of multiunit electromyographic signals," *IEEE Trans. Biomed. Eng.*, vol. 46, no. 6, pp. 685–697, Jun. 1999.

[10] D. Farina, E. Fortunato, and R. Merletti, "Noninvasive estimation of motor unit conduction velocity distribution using linear electrode arrays," *IEEE Trans. Biomed. Eng.*, vol. 47, no. 3, pp. 380–388, Mar. 2000.

[11] D. Farina and R. Merletti, "A novel approach for precise simulation of the EMG signal detected by surface electrodes," *IEEE Trans. Biomed. Eng.*, vol. 48, no. 6, pp. 637–646, Jun. 2001.

[12] D. Farina, L. Mesin, S. Martina, and R. Merletti, "A surface EMG generation model with multilayer cylindrical description of the volume conductor," *IEEE Trans. Biomed. Eng.*, vol. 51, no. 3, pp. 415–426, Mar. 2004.

[13] R. Merletti, L. Lo Conte, E. Avignone, and P. Guglielminotti, "Modeling of surface myoelectric signals, Part I: Model implementation," *IEEE Trans. Biomed. Eng.*, vol. 46, no. 7, pp. 810–820, Jul. 1999.

[14] R. Merletti, S. H. Roy, E. Kupa, S. Roatta, and A. Granata, "Modeling of surface myoelectric signals, Part II: Model-based signal interpretation," *IEEE Trans. Biomed. Eng.*, vol. 46, no. 7, pp. 821–829, Jul. 1999.

[15] L. R. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," *Proc. IEEE*, vol. 77, no. 2, pp. 257–286, Feb. 1989.

[16] J. R. Bellegarda and D. Nahamoo, "Tied mixture continuous parameter modeling for speech recognition," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, no. 12, pp. 2033–2045, Dec. 1990.

[17] G. D. Forney, Jr., "The Viterbi algorithm," *Proc. IEEE*, vol. 61, no. 3, pp. 268–278, Mar. 1973.

[18] T. Masuda and T. Sadoyama, "Topographical map of innervation zones within single motor units measured with a grid surface electrode," *IEEE Trans. Biomed. Eng.*, vol. 35, no. 8, pp. 623–628, Aug. 1988.

[19] K. C. McGill, Z. C. Lateva, and S. Xiao, "A model of the muscle action potential for describing the leading edge, terminal wave, and slow afterwave," *IEEE Trans. Biomed. Eng.*, vol. 48, no. 12, pp. 1357–1365, Dec. 2001.

Kevin R. Wheeler (S'86–M'90–SM'03) received the B.S. and M.S. degrees from the University of New Hampshire, Durham, in 1988 and 1991, respectively, and the Ph.D. degree from the University of Cincinnati, Cincinnati, OH, in 1996, all in electrical engineering.

After graduation, he joined the IBM Almaden Research Center to develop web-mining algorithms. In 1997, he joined the Computational Sciences Division, NASA Ames Research Center, Moffett Field, CA, where he was the Group Lead of Embedded Decision Systems. In 2006, he became the Group Leader of Instrumentation at the Monterey Bay Aquarium Research Institute, Moss Landing, CA. His research interests include applying probability theory to automating problems in the Earth sciences, man–machine interfaces, dexterous robotics, and marine instrumentation.

Mindy H. Chang received the B.S. degree in electrical engineering and computer science with a minor in biomedical engineering from Massachusetts Institute of Technology, Cambridge, in 2005. She is currently pursuing the M.S. and Ph.D. degrees in bioengineering at Stanford University, Stanford, CA.

Her research interests include signal processing and modeling of neural systems.

Kevin H. Knuth received the B.S. degree in physics from the University of Wisconsin–Oshkosh, Oshkosh, in 1988, the M.S. degree in physics from Montana State University, Bozeman, in 1990, and the Ph.D. degree in physics with a minor in mathematics from the University of Minnesota, Minneapolis, in 1995.

He held postdoctoral positions at the Louisiana State University Medical Center in 1996 and at the City University of New York and Albert Einstein College of Medicine from 1997 to 1998, where he was later an Instructor. In 1999, he was an Instructor at the Weill Medical College, Cornell University. From 1999 to 2001, he was a Research Scientist at the Center for Advanced Brain Imaging, Nathan Kline Institute. From 2001 to 2005, he was a Research Scientist at the NASA Ames Research Center, where he developed machine learning techniques and their applications in the Intelligent Systems Division. Since 2005, he has been an Assistant Professor at Physics Department and the College of Computing and Information, University at Albany, Albany, NY.