

Stochastic Similarity for Validating Human Control Strategy Models

Michael C. Nechyba, *Student Member, IEEE*, and Yangsheng Xu, *Senior Member, IEEE*

Abstract—Modeling dynamic human control strategy (HCS), or human skill in response to real-time sensing is becoming an increasingly popular paradigm in many different research areas, such as intelligent vehicle systems, virtual reality, and space robotics. Such models are often learned from experimental data, and as such can be characterized despite the lack of a good physical model. Unfortunately, learned models presently offer few, if any, guarantees in terms of model fidelity to the training data. This is especially true for dynamic reaction skills, where errors can feed back on themselves to generate state and command trajectories uncharacteristic of the source process. Thus, we propose a stochastic similarity measure—based on hidden Markov model (HMM) analysis—capable of comparing and contrasting stochastic, dynamic, multidimensional trajectories. This similarity measure is the first step in validating a learned model's fidelity to its training data by comparing the model's dynamic trajectories in the feedback loop to the human's dynamic trajectories. In this paper, we first derive and demonstrate properties of the similarity measure for stochastic systems. We then apply the similarity measure to real-time human driving data by comparing different control strategies among different individuals. We show that the proposed similarity measure out performs the more traditional Bayes classifier in correctly grouping driving data from the same individual. Finally, we illustrate how the similarity measure can be used in the validation of models which are learned from experimental data, and how we can connect model validation and model learning to iteratively improve our models of human control strategy.

Index Terms—Hidden Markov models, human modeling, model validation, neural networks, similarity measure.

I. INTRODUCTION

MODELS of human skill, or human control strategy, which accurately emulate dynamic human behavior, have far reaching potential in areas ranging from robotics to virtual reality to the intelligent vehicle highway project. Significant challenges arise in the modeling of human skill, however. Defying analytic representation, little if anything is known about the structure, order, or granularity of an individual's human controller. Human control strategy is both dynamic as well as stochastic in nature. In addition, the complex mapping from sensory inputs to control action outputs inherent in human control strategy can be highly nonlinear for given tasks. Therefore, developing an accurate and useful model for this type of dynamic phenomenon is frustrated by a poor

understanding of the underlying basis for that phenomenon. Consequently, modeling by observation, rather than physical derivation, is becoming an increasingly popular paradigm for characterizing a wide range of complex processes, including human control strategy (HCS). This type of modeling is said to constitute learning, since the model is not derived from *a priori* laws of nature, but rather from observed instances of experimental data, known collectively as the training set.

The main strength of modeling by learning, is that no explicit physical model is required; this also represents its biggest weakness, however. On the one hand, we are not restricted by the limitations of current scientific knowledge, and are able to model HCS for which we have not yet developed adequate biological or psychological understanding. On the other hand, the lack of scientific justification detracts from the confidence that we can show in these learned models. This is especially true when the unmodeled process is

- 1) dynamic;
- 2) stochastic in nature, as is the case for human control strategy.

For a dynamic process, model errors can feed back on themselves to produce trajectories which are not characteristic of the source process or are even potentially unstable. For a stochastic process, a static error criterion (such as root mean square (RMS) error), based on the difference between the training data and predicted model outputs may be inadequate and inappropriate to gauge the fidelity of a learned model to the source process. Yet, most learning approaches today, including the cascade learning algorithm, utilize some static error measure as a test of convergence for the learning algorithm. While this measure is very useful during training, it offers no guarantees, theoretical or otherwise, about the dynamic behavior of the resulting learned model.

For a simple illustration of this problem, consider the following example. Suppose that we wish to learn a dynamic process represented by the following simple difference equation:

$$y(k+1) = 0.75y(k) + 0.24y(k-1) + x(k) \quad (1)$$

where $y(k)$, $x(k)$ represent the output and input of the system, respectively, at time step k . The input-output training data in Table I is provided.

Note that (1) is asymptotically stable. Now, suppose, for example, that we train simple neural networks to learn three different approximations of the system in (1)

$$\#1: y(k+1) = 0.76y(k) + 0.25y(k-1) + x(k) \quad (2)$$

$$\#2: y(k+1) = 0.76y(k) + 0.23y(k-1) + x(k) \quad (3)$$

$$\#3: y(k+1) = 0.74y(k) + 0.23y(k-1) + x(k). \quad (4)$$

Manuscript received June 6, 1996; revised January 15, 1998. This work was supported in part by a DOE doctoral fellowship. This paper was recommended for publication by Associate Editors F. Cervantes and B. Hannaford upon evaluation of the reviewers' comments.

The authors are with The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213 USA.

Publisher Item Identifier S 1042-296X(98)02920-6.

TABLE I
SAMPLE INPUT-OUTPUT TRAINING DATA

Input			Output
$y(k-1)$	$y(k)$	$x(k)$	$y(k+1)$
-0.1	0.1	0.4	0.349
0.1	0.1	0.5	0.599
-0.3	0.2	0.3	0.123
0.3	0.2	0.4	0.673
0.2	0.0	0.5	0.650
0.0	0.2	0.3	0.348




Fig. 1. The three models result in dramatically different (even unstable) trajectories.

The three models all have the same RMS error for the training set in Table I. Nevertheless, the dynamic trajectories for the three models differ significantly. For example, consider the input

$$x(k) = 0.1 \sin\left(\frac{k\pi}{100}\right). \quad (5)$$

The resulting output for the system as well as the three models is shown in Fig. 1. Model #2 approximates the system in (1) well; model #3 remains stable, but approximates the system with significantly poorer accuracy; finally, model #1 diverges into an unstable trajectory.

The difference in the three models is the distribution of the error over the training set. Thus, a static error measure, such as RMS error, does not provide sufficiently satisfactory model validation for a dynamic process. Furthermore, for stochastic systems, one cannot expect equivalent trajectories for the process and the learned model, given the same initial conditions. Thus, we require a stochastic similarity measure, with sufficient representational power and flexibility to compare multidimensional, stochastic trajectories. In this paper, we develop such a similarity measure, based on hidden Markov model (HMM) analysis, as a first step in validating models of dynamic human control strategy.

This paper is organized into three parts.

1) First, using HMM's, we derive a stochastic similarity measure capable of comparing arbitrary multidimen-

sional, stochastic trajectories. This measure makes no *a priori* assumptions about the statistical distribution of the underlying data to be compared. We demonstrate certain properties of the proposed similarity measure through both mathematical proof and simulation of known stochastic systems.

- 2) Second, we evaluate the similarity measure on human control data, by comparing driving strategies among different individuals. We show that the proposed similarity measure outperforms the more traditional Bayes classifier in correctly grouping driving data from the same individual.
- 3) Finally, we illustrate how the similarity measure can be used in the validation of models which are learned from experimental data, and how we can connect model validation and model learning to iteratively improve our models of human control strategy.

II. STOCHASTIC SIMILARITY

Similarity measures or metrics have been given considerable attention in computer vision [1]–[3], image database retrieval [4], and two-dimensional (2-D) or three-dimensional (3-D) shape analysis [5], [6]. These methods, however, generally rely on the special properties of images, and are therefore not appropriate for analyzing sequential trajectories. Other work has focussed on classifying temporal patterns using standard statistical techniques [7], wavelet analysis [8], neural networks [9], [10], and HMM's (see discussion below). Much of this work, however, analyzes only short-time trajectories or patterns, and, in many cases, generates only a binary classification, rather than a continuously valued similarity measure. Prior work has not addressed the problem of comparing long, multidimensional, stochastic trajectories, especially of human control data. Thus, we propose to evaluate stochastic similarity between two dynamic, multidimensional trajectories using HMM analysis.

A. Hidden Markov Models

Rich in mathematical structure, HMM's are trainable statistical models, with two appealing features:

- 1) no *a priori* assumptions are made about the statistical distribution of the data to be analyzed;
- 2) a high degree of sequential structure can be encoded by the HMM's.

As such, they have been applied for a variety of stochastic signal processing. In speech recognition, where HMM's have found their widest application, human auditory signals are analyzed as speech patterns [11], [12]. Transient sonar signals are classified with HMM's for ocean surveillance in [13]. Radons *et al.* [14] analyze 30-electrode neuronal spike activity in a monkey's visual cortex with HMM's. Hannaford and Lee [15] classify task structure in teleoperation based on HMM's. In [16] and [17], HMM's are used to characterize sequential images of human actions. Finally, Yang and Xu apply HMM's to open-loop action skill learning [18] and human gesture recognition [19].




Fig. 2. A 5-state HMM, with 16 observable symbols in each state.

A HMM consists of a set of n states, interconnected through probabilistic transitions; each of these states has some output probability distribution associated with it. Although algorithms exist for training HMM's with both discrete and continuous output probability distributions, and although most applications of HMM's deal with real-valued signals, discrete HMM's are preferred to continuous HMM's in practice, due to their relative computational simplicity and lesser sensitivity to initial random parameter settings [20]. In Section II-D below, we describe how we use discrete HMM's for analysis of real-valued signals by converting the data to discrete symbols through preprocessing and vector quantization. Thus, a discrete HMM is completely defined by the following triplet [12]:

$$\lambda = \{A, B, \pi\} \quad (6)$$

where A is the probabilistic $n \times n$ state transition matrix, B is the $L \times n$ output probability matrix with L discrete output symbols $\chi \in \{1, 2, \dots, L\}$, and π is the n -length initial state probability distribution vector for the HMM. Fig. 2, for example, represents a 5-state HMM, where each state emits one of 16 discrete symbols, based on some probability distribution.

For an observation sequence O of discrete symbols, we can locally maximize $P(\lambda|O)$ (i.e. the probability of the model λ given the observation sequence O) using the Baum–Welch Expectation-Maximization (EM) algorithm [12], [21]. Throughout this paper, we initialize the Baum–Welch algorithm by setting the HMM parameters to random, nonzero values, subject of course to the necessary probabilistic constraints. We can also evaluate $P(O|\lambda)$ (i.e., the probability that a given observation sequence O is generated from the model λ).

Finally, two HMM's λ_1 and λ_2 are defined to be equivalent

$$\lambda_1 \sim \lambda_2, \quad \text{iff. } P(O|\lambda_1) = P(O|\lambda_2), \quad \forall O. \quad (7)$$

Note that λ_1 and λ_2 need not be identical to be equivalent. The following two HMM's are, for example, equivalent:

$$\lambda_1 = \{[1], [0.5 \ 0.5]^T, [1]\} \\ \lambda_2 = \left\{ \begin{bmatrix} 0.5 & 0.5 \\ 0.5 & 0.5 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, [0.5 \ 0.5]^T \right\}. \quad (8)$$

B. Similarity Measure

Below, we derive a stochastic similarity measure, based on discrete-output HMM's. Let $O_i, i \in \{1, 2, \dots\}$, denote a distinct observation sequence of discrete symbols with length T_i . Also, let $\lambda_j = \{A_j, B_j, \pi_j\}, j \in \{1, 2, \dots\}$, denote




Fig. 3. Four normalized probability values make up the similarity measure.

a discrete HMM locally optimized using the Baum–Welch algorithm to maximize $P(\lambda_j|O_j)$. Similarly, let $P(O_i|\lambda_j)$ denote the probability of the observation sequence O_i given the model λ_j , and let

$$P_{ij} = \hat{P}(O_i|\lambda_j) = P(O_i|\lambda_j)^{1/T_i} \quad (9)$$

denote the probability of the observation sequence O_i given the model λ_j , normalized with respect to T_i . In practice, we calculate P_{ij} as

$$P_{ij} = 10^{\log P(O_i|\lambda_j)/T_i} \quad (10)$$

to avoid problems of numerical underflow for long observation sequences.

Using the definition in (9), Fig. 3 illustrates our overall approach to evaluating similarity between two observation sequences. Each observation sequence is first used to train a corresponding HMM; this allows us to evaluate P_{11} and P_{22} . Furthermore, we subsequently cross-evaluate each observation sequence on the other HMM (i.e., $P(O_1|\lambda_2), P(O_2|\lambda_1)$ to arrive at P_{12} and P_{21} . Given, these four normalized probability values, we now define the following similarity measure between O_1 and O_2 :¹

$$\sigma(O_1, O_2) = \sqrt{\frac{P_{21}P_{12}}{P_{11}P_{22}}}. \quad (11)$$

This measure takes the ratio of the cross probabilities over the training probabilities, and normalizes for the multiplication of the two probability values in the numerator and denominator by taking the square root.

Because λ_1 and λ_2 are necessarily trained on finite-length training sequences, rare events—either rare state transitions or rare observations—may not be recorded in either or both O_1 and O_2 , forcing corresponding parameters in the HMM's to converge to zero during training. This can cause P_{21} and P_{12} to degenerate to zero, even when O_1 and O_2 are, in fact, almost identical. For example, suppose O_1 and O_2 are identical observation sequences. Then, $\sigma(O_1, O_2)$ will evaluate to one. Replacing just one observation in O_1 with an observable not present in O_2 , however, will force P_{12} and hence $\sigma(O_1, O_2)$ to zero, even though the one rogue observation could simply be a measurement or recording error. To overcome this singularity problem and to account for

¹In [22], we proposed a different similarity measure for which properties #1, #2, and #3 do not hold. Thus, the previous similarity measure gives potentially inconsistent and misleading results for certain data. The similarity measure in (11) corrects these problems.

the possibility of rare, but unobserved events, we follow the common practice of replacing nonzero elements in the trained HMM's by some $\varepsilon > 0$ and renormalizing the model to fit probabilistic constraints [12]. Therefore, we calculate the P_{ij} not on λ_j itself, but rather $\hat{\lambda}_j$, a smoothed version of λ_j , where zero elements in the matrices $\{A_j, B_j, \pi_j\}$ are replaced by $\varepsilon = 0.0001$. This value of ε is chosen as it redistributes less than 0.1% of the probability mass in the state transition matrix A and less than 0.5% of the probability mass in the output probability matrix B .

C. Properties

For now, assume that P_{ii} is a global (rather than just a local) maximum. Then

$$1/L \leq P_{ii} \quad (12)$$

and

$$0 < P_{ij} \leq P_{ii}, \quad \varepsilon > 0. \quad (13)$$

The lower bound for P_{ii} in (12) is realized for single-state discrete HMM's, and a uniform distribution of symbols in O_i . From (11)–(13), we can establish the following properties for $\sigma(O_1, O_2)$:

$$\text{Property \#1: } \sigma(O_1, O_2) = \sigma(O_2, O_1) \quad (\text{by definition}) \quad (14)$$

$$\text{Property \#2: } 0 < \sigma(O_1, O_2) \leq 1 \quad (15)$$

$$\text{Property \#3: } \sigma(O_1, O_2) = 1 \quad \text{if} \quad (\text{a}) \quad \lambda_1 \sim \lambda_2 \\ \text{or} \quad (\text{b}) \quad O_1 = O_2. \quad (16)$$

As we have noted before, the Baum-Welch algorithm guarantees only that P_{ii} is a local maximum. In practice, this is not a significant concern, however, as the Baum-Welch algorithm converges to near-optimal solutions, when the algorithm is initialized with random model parameters [12], [20]. We can verify this near-optimal convergence property experimentally. First, for a given model λ , we generate a long observation sequence O' . Next, we train a second HMM λ' with initial random model parameters on O' . Finally, we evaluate the normalized probability values $\hat{P}(O'|\lambda)$ and $\hat{P}(O'|\lambda')$. If λ' is near-optimal or optimal for sequence O' , then

$$\hat{P}(O'|\lambda')/\hat{P}(O'|\lambda) \approx 1 \quad (17)$$

provided that O' is sufficiently long. We have used both this procedure as well as multiple training runs from different random initial parameter settings to verify the near-optimal convergence properties of the Baum-Welch algorithm for the type of data (i.e., human control data) studied in this paper.

Below, we illustrate the behavior of the similarity measure for some simple HMM's. First, for single-state HMM's, the similarity measure reduces to

$$\sigma(O_1, O_2) = \prod_{k=1}^L \left(\frac{b_{1k}}{b_{2k}} \right)^{(b_{2k} - b_{1k})/2} \quad (18)$$

which reaches a maximum when $b_{1k} = b_{2k}$, or simply, $B_1 = B_2$, and that maximum is equal to one. Fig. 4 shows




Fig. 4. Similarity measure for two binomial distributions. Lighter colors indicate higher similarity.

a contour plot for

$$B_1 = [p_1 \quad 1 - p_1]^T \quad B_2 = [p_2 \quad 1 - p_2]^T \quad (19)$$

where $0 < p_1, p_2 < 1$.

Second, we give an example of how the proposed similarity measure changes, not as a function of different symbol distributions, but rather as a function of varying HMM structure. Consider the following HMM:

$$\lambda(\alpha) = \left\{ \left[\begin{array}{cc} \frac{1+\alpha}{2} & \frac{1-\alpha}{2} \\ \frac{1-\alpha}{2} & \frac{1+\alpha}{2} \end{array} \right], \left[\begin{array}{cc} 1 & 0 \\ 0 & 1 \end{array} \right], [0.5 \quad 0.5]^T \right\} \\ 0 \leq \alpha < 1 \quad (20)$$

and corresponding observation sequences, $O(\alpha)$, stochastically generated from model $\lambda(\alpha)$. For all $\alpha \in [0, 1)$, $O(\alpha)$ will have an equivalent aggregate distribution of symbols 0 and 1—namely 1/2 and 1/2. As α increases, however, $O(\alpha)$ will become increasingly structured. For example

$$\lambda(0) = \{[1], [0.5 \quad 0.5]^T, [1]\} \\ (\text{equivalent to unbiased coin toss}) \quad (21)$$

$$\lim_{\alpha \rightarrow 1} O(\alpha) = \{\dots, 1, 1, 0, 0, \dots, 0, 0, 1, 1, \dots\}. \quad (22)$$

Fig. 5 graphs $\sigma[O(\alpha_1), O(\alpha_2)]$ as a contour plot for $0 \leq \alpha_1, \alpha_2 < 1$, where each observation sequence $O(\alpha)$ of length $T = 10000$ is generated stochastically from the corresponding HMM $\lambda(\alpha)$.² Greatest similarity is indicated for $\alpha_1 = \alpha_2$, while greatest dissimilarity occurs for $(\alpha_1 \rightarrow 1, \alpha_2 = 0)$, and $(\alpha_1 = 0, \alpha_2 \rightarrow 1)$.

In some cases, it may be more convenient to represent the similarity between two trajectories through a distance measure $d(O_1, O_2)$, rather than a similarity measure. Given

²This procedure only approximates our similarity measure definition, since $\lambda(\alpha)$ is only optimal for $O(\alpha)$ as $T \rightarrow \infty$.




Fig. 5. The similarity measure changes predictably as a function of HMM structure.

the similarity measure $\sigma(O_1, O_2)$, such a measure is easily derived. Let

$$\begin{aligned} d(O_1, O_2) &= -\log \sigma(O_1, O_2) \\ &= \frac{1}{2} [\log(P_{11}P_{22}) - \log(P_{21}P_{12})] \end{aligned} \quad (23)$$

such that

$$d(O_1, O_2) = d(O_2, O_1) \quad (24)$$

$$d(O_1, O_2) \geq 0 \quad (25)$$

$$d(O_1, O_2) = 0; \text{ if (a) } \lambda_1 \sim \lambda_2 \text{ or (b) } O_1 = O_2. \quad (26)$$

The distance measure $d(O_1, O_2)$ between two observation sequences defined in (23) is closely related to the dual notion of distance between two HMM's, as proposed in [23].

Let \bar{O}_i denote a random observation sequence of length \bar{T}_i generated by the HMM $\bar{\lambda}_i$, and let

$$\bar{P}_{ij} = P(\bar{O}_i | \bar{\lambda}_j)^{1/\bar{T}_i}. \quad (27)$$

Then, [23] defines the following distance measure between two HMM's, $\bar{\lambda}_1$ and $\bar{\lambda}_2$:

$$d(\bar{\lambda}_1, \bar{\lambda}_2) = \frac{1}{2} [\log(\bar{P}_{11}\bar{P}_{22}) - \log(\bar{P}_{21}\bar{P}_{12})]. \quad (28)$$

Unlike the observation sequences O_i , the sequences \bar{O}_i are not unique, since they are stochastically generated from $\bar{\lambda}_i$. Hence, $d(\bar{\lambda}_1, \bar{\lambda}_2)$ is uniquely determined only in the limit as $\bar{T}_i \rightarrow \infty$. Likewise for $d(O_1, O_2)$, the HMM's λ_1 and λ_2 are not unique, since P_{11} and P_{22} are in general guaranteed to be only local, not global maxima. Hence, $d(O_1, O_2)$ is uniquely determined only when P_{11} and P_{22} represent global maxima.

While in general, $d(\bar{\lambda}_1, \bar{\lambda}_2)$ and $d(O_1, O_2)$ are not equivalent, the discussion above suggests sufficient conditions for





Fig. 6. (a) First eight Walsh-ordered Walsh functions and (b) some sample human control data.

which the two notions—distance between HMM's and distance between observation sequences—do converge to equivalence. Specifically, $d(O_1, O_2) = d(\lambda_1, \lambda_2)$ in general if and only if

$$(1) \lambda_1 \sim \bar{\lambda}_1, \lambda_2 \sim \bar{\lambda}_2 \quad (29)$$

$$(2) P_{11}, P_{22} \text{ are global maxima} \quad (30)$$

$$(3) \bar{T}_i \rightarrow \infty. \quad (31)$$

D. Signal-to-Symbol Conversion

Since we use discrete-output HMM's in our similarity measure, we need to convert multidimensional, real-valued human control data to a sequence of discrete symbols. We follow two steps in this conversion:

- 1) spectral preprocessing;
- 2) vector quantization, as illustrated in Fig. 8.

The primary purpose of the spectral preprocessing is to extract meaningful feature vectors for the vector quantizer. In this work, we rely on the fast Fourier transform (FFT) and the fast Walsh transform (FWT), the $O(n \log n)$ algorithmic counterparts of the discrete Fourier transform (DFT) and the discrete Walsh transform (DWT), respectively. Instead of sinusoidal basis functions, the Walsh transform decomposes a signal based on the orthonormal Walsh functions [24]. The first eight Walsh-ordered Walsh functions are shown in Fig. 6(a). In Fig. 6(b), we show an example of human control data

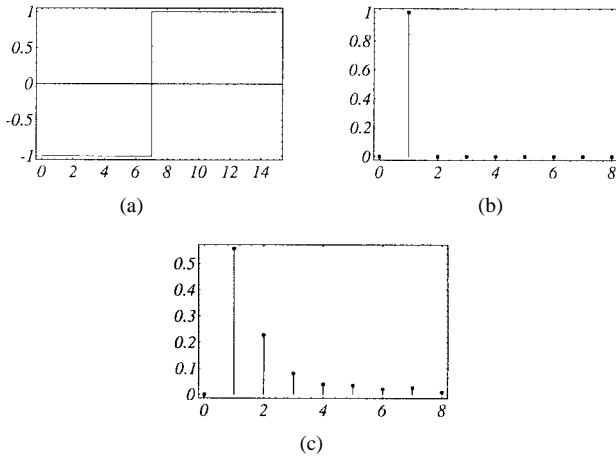


Fig. 7. (a) Sample square wave and its corresponding (b) Walsh, and (c) Fourier PSD's.

(see Section III-A below) which can be characterized better through the Walsh transform, rather than the Fourier transform, due to its sharply discontinuous, step-like profile. Consider, for example, the power spectral densities (PSD's) for the square wave in Fig. 7(a). The Walsh PSD in Fig. 7(b) is a more concise feature vector than the corresponding Fourier PSD in Fig. 7(c).

For each dimension of the human control data, we partition the data into overlapping window frames, and perform either a short-time FFT or FWT on each frame. Generally, we select the FFT for state trajectories, and the FWT for command trajectories, since these trajectories tend to have sharp discontinuities for the experimental data in this paper. In the case of the FFT, the data in each frame is filtered through a Hamming window before applying the FFT, so as to compensate for the windowing effect. The spectral coefficients are then converted to power spectral density (PSD) vectors. In preparation for the vector quantization, the PSD vectors along each dimension of the system trajectory are normalized and concatenated into one long feature vector per frame. We quantize the resulting sequence of long feature vectors using the iterative LBG VQ algorithm [25]. This vector quantizer generates codebooks of size 2^m , $m \in \{0, 1, 2, \dots\}$, and can be stopped at an appropriate level of discretization given the amount of available data and complexity of the system trajectories. Assuming that we segment the data into window frames of length k with 50% overlap, the original multidimensional, real-valued signal of length t is thus converted to a sequence of discrete symbols of length $T = \text{int}(2t/k)$.

III. COMPARING HUMAN CONTROL STRATEGIES

A. Experimental Set-Up

Fig. 9 shows the real-time graphic driving simulator, from which we collect human control data. In the interface, the human operator has full control over the steering of the car (mouse movement), the brake (left mouse button) and the accelerator (right mouse button); the middle mouse button corresponds to slowly easing off whichever pedal is currently being "pushed." The vehicle dynamics are given in (32)–(50)

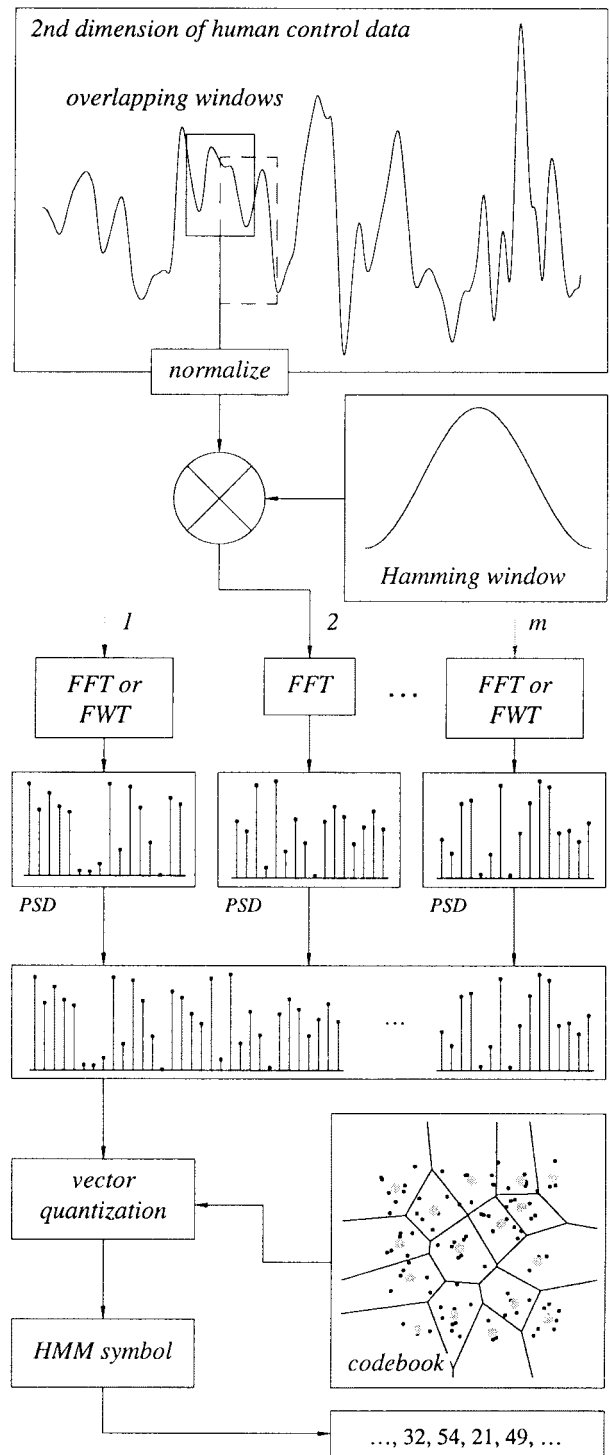


Fig. 8. Conversion of multidimensional human control data to a sequence of discrete symbols.

below (modified from [26])

$$\ddot{\theta} = (l_f P_f \delta + l_f F_{\xi f} - l_r F_{\xi r}) / I \quad (32)$$

$$\dot{\nu}_\xi = (P_f \delta + F_{\xi f} + F_{\xi r}) / m - \nu_\eta \dot{\theta} - (\text{sgn } \nu_\xi) c_D \nu_\xi^2 \quad (33)$$

$$\dot{\nu}_\eta = (P_f + P_r - F_{\xi f} \delta) / m + \nu_\xi \dot{\theta} - (\text{sgn } \nu_\eta) c_D \nu_\eta^2 \quad (34)$$

$$\begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} \nu_\xi \\ \nu_\eta \end{bmatrix} \quad (35)$$

where

$$\dot{\theta} = \text{angular velocity of the car} \quad (36)$$

$$\nu_\xi = \text{body-lateral velocity of the car}$$

$$\nu_\eta = \text{longitudinal velocity of the car} \quad (37)$$

$$F_{\xi k} = \mu F_{zk} (\tilde{\alpha}_k - (\text{sgn } \delta) \tilde{\alpha}_k^2/3 + \tilde{\alpha}_k^3/27) \cdot \sqrt{1 - P_k^2/(\mu F_{zk})^2 + P_k^2/C_k^2}, \quad k \in \{f, r\} \quad (38)$$

$$\tilde{\alpha}_k = c_k \alpha_k / (\mu F_{zk}), \quad k \in \{f, r\} \quad (39)$$

$$\alpha_f = \text{front tire slip angle} = \delta - (l_f \dot{\theta} - \nu_\xi) / \nu_\eta \quad (40)$$

$$\alpha_r = \text{rear tire slip angle} = (l_r \dot{\theta} - \nu_\xi) / \nu_\eta \quad (41)$$

$$F_{zf} = (m g l_r - (P_f + P_r) h) / (l_f + l_r) \quad (42)$$

$$F_{zr} = (m g l_f + (P_f + P_r) h) / (l_f + l_r) \quad (42)$$

$$\xi = \text{body-relative lateral axis}$$

$$\eta = \text{body-relative longitudinal axis} \quad (43)$$

$$c_f, c_r = \text{cornering stiffness of front, rear tires} \\ = 50\,000 \text{ N/rad}, 64\,000 \text{ N/rad} \quad (44)$$

$$c_D = \text{lumped coefficient of drag (air resistance)} \\ = 0.0005 \text{ m}^{-1} \quad (45)$$

$$\mu = \text{coefficient of friction} = 1$$

$$F_{jk} = \text{frictional forces}, \quad j \in \{\xi, z\}, \quad k \in \{f, r\} \quad (46)$$

$$P_r = \text{longitudinal force on rear tires} \\ = \begin{cases} 0, & P_f \geq 0 \\ k_b P_f, & P_f < 0, k_b = 0.34 \end{cases} \quad (47)$$

$$m = 1500 \text{ kg}, \quad I = 2500 \text{ kg-m}^2, \quad l_f = 1.25 \text{ m}$$

$$l_r = 1.5 \text{ m}, \quad h = 0.5 \text{ m} \quad (48)$$

and the controls are given by

$$-8000 \text{ N} \leq (P_f = \text{longitudinal force on front tires}) \\ \leq 4000 \text{ N} \quad (49)$$

$$-0.2 \text{ rad} \leq (\delta = \text{steering angle}) \leq 0.2 \text{ rad}. \quad (50)$$

Note that the separate brake and gas commands for the human are, in fact, the single P_f variable, where the sign indicates whether the brake or the gas is active [Fig. 6(b), for example, illustrates one person's P_f profile for part of one run]. The entire simulator is run at 50 Hz.

B. Similarity Results

For the first set of experiments, we ask five people:

- 1) Larry;
- 2) Curly;
- 3) Moe;
- 4) Groucho;
- 5) Harpo;

to practice driving in the simulator for a period of up to 15 min to become accustomed to the simulator's dynamics. We then record driving data for each person on three different, randomly generated 20 km roads, with only short breaks in between each run. Thus, we have a total of 15 runs.




Fig. 9. The driving simulator generates a perspective view of the road for the user, who has independent control over steering, braking, and acceleration (gas).

Each road is described by a sequence of randomly generated segments of the form $\{l, 0\}$ (straight-line segments), and $\{r, \phi\}$ (curves), connected in a manner that ensures continuous first derivatives between segments. We also place the following constraints on individual segments:

$$100 \text{ m} \leq l \leq 200 \text{ m} \\ l = \text{length of a straight-line segment} \quad (51)$$

$$100 \text{ m} \leq r \leq 200 \text{ m}, \quad r = \text{radius of curvature} \quad (52)$$

$$20^\circ \leq \phi \leq 180^\circ, \quad -180^\circ \leq \phi \leq -20^\circ \\ \phi = \text{sweep angle of curve}. \quad (53)$$

Therefore, each segment is one of the following:

- 1) straight line segment ($\phi = 0$);
- 2) left curve ($\phi > 0$);
- 3) right curve ($\phi < 0$).

No segment may be followed by a segment of the same type; a curve is followed by a straight line segment with probability 0.4, and an opposite curve segment with probability 0.6. A straight line segment is followed by a left curve or right curve with equal probability. Roads are defined to be 10 m wide (the car is 2 m wide), and the visible horizon is set to 100 m. Fig. 10(a)–(c) show the three different roads used for these experiments.

For notational convenience, let

$$R_{ij} = R(i, j) \quad i \in \{1, 2, 3, 4, 5\} \quad j \in \{1, 2, 3\} \quad (54)$$

denote the run from person (i) on road #j. Table II below reports some aggregate statistics for each of the 15 runs.

Our goal here is to see

- 1) how well the similarity measure classifies each individual's runs across different roads;
- 2) how the classification performance of the proposed similarity measure compares with a more conventional statistical technique, namely the Bayes optimal classifier.




Fig. 10. (a) Road 1, (b) road 2, and (c) road 3.

The system trajectory for the driving task is defined by the three state variables $\{\nu_\xi, \nu_\eta, \dot{\theta}\}$ and the two control variables $\{\delta, P_f\}$. For all the similarity results below (unless otherwise noted), we choose the following spectral preprocessing:

$$\{\nu_\xi, \nu_\eta, \dot{\theta}, \delta, P_f\} = \{F-16, F-16, F-16, W-16, W-16\} \quad (55)$$

where $F-k, k \geq 1$, denotes the k -point FFT, and $W-k, k \geq 1$, denotes the k -point Walsh transform, with 50% overlapping window frames. Note that the PSD vector for a k -point transform has length $(k/2 + 1)$; thus, the feature vector to be quantized is of length 45. Also, note that we choose the Walsh transform for the two control variables $\{\delta, P_f\}$, since the user's control is generally discontinuous and step-like in profile [as shown in Fig. 6(b) above]. When vector quantized, the feature vectors generated by the Walsh transform exhibit significantly reduced distortion as compared to the Fourier-transformed feature vectors for the two control variables.

From the preprocessed data, we build three vector codebooks $Q_k, k \in \{1, 2, 3\}$, each with 128 levels, and corresponding to data from $R_{ik}, \forall i \in \{1, 2, 3, 4, 5\}$. For these experiments, the number of levels in the VQ codebook is primarily constrained by the amount of available data to train the HMM's, since we want

$$\begin{aligned} (T = \text{length of each observation sequence}) \\ \gg (n = \# \text{ model parameters}) \end{aligned} \quad (56)$$

and the number of model parameters increases with the number of levels L in the VQ codebook. Now define

$$O_{ij}^k = O(i, j, k) \quad i \in \{1, 2, 3, 4, 5\}, \quad j, k \in \{1, 2, 3\} \quad (57)$$

as the observation sequence of discrete symbols vector quantized from the preprocessed feature vectors of run R_{ij} , using the codebook Q_k . We can view the observation sequences O_{ik}^k as control strategy data which we have already collected, processed, and modeled, without any information about the other data ($O_{ij}^k, j \neq k$). Thus, each sequence O_{ik}^k represents a class for individual i . On the other hand, we can view the observation sequences $O_{ij}^k, j \neq k$, as new control data, which we wish to classify, using our similarity measure, as belonging

TABLE II
AGGREGATE STATISTICS FOR HUMAN DRIVING DATA

	Run	v (mph)	$\dot{\theta}$ (rad/s)	δ (rad)	P_f (1000N)
Larry	$R(1,1)$	63.1 ± 12.2	0.00 ± 0.17	0.00 ± 0.06	1.4 ± 2.4
	$R(1,2)$	62.7 ± 9.5	0.00 ± 0.17	0.00 ± 0.06	1.3 ± 1.9
	$R(1,3)$	64.0 ± 8.6	0.00 ± 0.18	0.00 ± 0.06	1.3 ± 1.4
Curly	$R(2,1)$	70.8 ± 8.3	0.00 ± 0.20	0.00 ± 0.07	1.9 ± 3.3
	$R(2,2)$	69.1 ± 7.7	0.00 ± 0.19	0.00 ± 0.07	1.8 ± 3.3
	$R(2,3)$	71.5 ± 7.7	0.00 ± 0.20	0.00 ± 0.08	2.0 ± 3.1
Moe	$R(3,1)$	73.1 ± 9.5	0.00 ± 0.24	0.00 ± 0.09	2.2 ± 2.8
	$R(3,2)$	71.9 ± 9.0	0.00 ± 0.25	0.00 ± 0.09	2.2 ± 2.6
	$R(3,3)$	74.5 ± 9.4	0.00 ± 0.29	0.00 ± 0.11	2.6 ± 2.6
Groucho	$R(4,1)$	66.8 ± 12.4	0.00 ± 0.18	0.00 ± 0.08	1.9 ± 3.8
	$R(4,2)$	65.1 ± 13.2	0.00 ± 0.21	0.00 ± 0.09	1.9 ± 4.0
	$R(4,3)$	69.8 ± 12.3	0.00 ± 0.23	0.00 ± 0.11	2.3 ± 3.8
Harpo	$R(5,1)$	52.3 ± 12.2	0.00 ± 0.17	0.00 ± 0.05	0.9 ± 1.5
	$R(5,2)$	51.7 ± 4.2	0.00 ± 0.16	0.00 ± 0.04	0.7 ± 0.3
	$R(5,3)$	56.1 ± 5.7	0.00 ± 0.20	0.00 ± 0.06	1.0 ± 0.3

to one of the five individuals represented by the O_{ik}^k classes. For a given application, there will generally be known data from which we define our classes (represented here by the sequences O_{ik}^k), and unknown data which we require to be classified into one of the defined classes (represented here by the sequences $O_{ij}^k, j \neq k$). In this context, it would be quite burdensome to recalculate the vector codebook each time we wish to classify a new control trajectory. Having separate codebooks (for only the known data) eliminates the need for this codebook recalculation.

Tables III–V classify each of the $O_{ij}^k, j \neq k$ based on O_{ik}^k for $k = 1, k = 2$, and $k = 3$, respectively, for eight-state HMM's. Note that the maximum value in each row is highlighted. We consider $O_{ij}^k, j \neq k$, classified correctly if and only if

$$\sigma(O_{ik}^k, O_{ij}^k) > \sigma(O_{ik}^k, O_{il}^k), \quad \forall (j \neq k, l \neq i). \quad (58)$$

In other words, we expect that two runs from the same individual (but on different roads) will yield a higher similarity measure than two runs from two different individuals. From the tables, we observe that the similarity measure correctly classifies all 30 comparisons, by assigning the highest similarity between runs from the same individual, and significantly lower similarity between runs from two different individuals.

Now we compare these classification results with the Bayes optimal classifier. Define class ω_{ik} as

$$\begin{aligned} \omega_{ik} = \omega(i, k) = \{ik, \Sigma_{ik}\}, \quad i \in \{1, 2, 3, 4, 5\} \\ k \in \{1, 2, 3\} \end{aligned} \quad (59)$$

TABLE III
(a) SIMILARITY MEASURE CLASSIFICATION AND (b)
BAYES OPTIMAL CLASSIFICATION FOR ROAD #1 DATA

σ	$O(1, 1, 1)$	$O(2, 1, 1)$	$O(3, 1, 1)$	$O(4, 1, 1)$	$O(5, 1, 1)$
$O(1, 2, 1)$	0.765	0.172	0.117	0.099	0.315
$O(1, 3, 1)$	0.585	0.131	0.097	0.062	0.309
$O(2, 2, 1)$	0.187	0.857	0.358	0.287	0.012
$O(2, 3, 1)$	0.179	0.797	0.395	0.279	0.011
$O(3, 2, 1)$	0.255	0.407	0.845	0.343	0.019
$O(3, 3, 1)$	0.153	0.208	0.726	0.318	0.010
$O(4, 2, 1)$	0.073	0.189	0.287	0.644	0.006
$O(4, 3, 1)$	0.065	0.164	0.336	0.628	0.007
$O(5, 2, 1)$	0.038	0.004	0.003	0.009	0.428
$O(5, 3, 1)$	0.089	0.012	0.011	0.013	0.602

(a)

%	$\omega(1, 1)$	$\omega(2, 1)$	$\omega(3, 1)$	$\omega(4, 1)$	$\omega(5, 1)$
$R(1, 2)$	0.338	0.131	0.017	0.059	0.455
$R(1, 3)$	0.295	0.102	0.033	0.033	0.538
$R(2, 2)$	0.249	0.353	0.053	0.288	0.057
$R(2, 3)$	0.172	0.405	0.095	0.294	0.034
$R(3, 2)$	0.134	0.256	0.344	0.232	0.034
$R(3, 3)$	0.067	0.160	0.539	0.220	0.014
$R(4, 2)$	0.216	0.208	0.150	0.357	0.069
$R(4, 3)$	0.129	0.173	0.285	0.374	0.039
$R(5, 2)$	0.003	0.000	0.003	0.001	0.994
$R(5, 3)$	0.051	0.005	0.032	0.004	0.908

(b)

TABLE IV
(a) SIMILARITY MEASURE CLASSIFICATION AND (b)
BAYES OPTIMAL CLASSIFICATION FOR ROAD #2 DATA

σ	$O(1, 2, 2)$	$O(2, 2, 2)$	$O(3, 2, 2)$	$O(4, 2, 2)$	$O(5, 2, 2)$
$O(1, 1, 2)$	0.753	0.176	0.108	0.115	0.353
$O(1, 3, 2)$	0.626	0.143	0.094	0.062	0.272
$O(2, 1, 2)$	0.209	0.860	0.350	0.276	0.013
$O(2, 3, 2)$	0.180	0.786	0.384	0.270	0.010
$O(3, 1, 2)$	0.236	0.406	0.826	0.340	0.018
$O(3, 3, 2)$	0.139	0.206	0.705	0.288	0.010
$O(4, 1, 2)$	0.082	0.185	0.267	0.674	0.007
$O(4, 3, 2)$	0.076	0.171	0.331	0.637	0.007
$O(5, 1, 2)$	0.038	0.003	0.003	0.008	0.429
$O(5, 3, 2)$	0.106	0.013	0.009	0.011	0.599

(a)

%	$\omega(1, 2)$	$\omega(2, 2)$	$\omega(3, 2)$	$\omega(4, 2)$	$\omega(5, 2)$
$R(1, 1)$	0.579	0.202	0.046	0.096	0.077
$R(1, 3)$	0.635	0.090	0.041	0.025	0.209
$R(2, 1)$	0.225	0.490	0.152	0.130	0.003
$R(2, 3)$	0.232	0.500	0.150	0.117	0.002
$R(3, 1)$	0.111	0.329	0.407	0.150	0.002
$R(3, 3)$	0.061	0.203	0.539	0.196	0.001
$R(4, 1)$	0.271	0.365	0.102	0.261	0.001
$R(4, 3)$	0.115	0.267	0.294	0.322	0.000
$R(5, 1)$	0.341	0.010	0.035	0.123	0.491
$R(5, 3)$	0.256	0.006	0.044	0.053	0.641

(b)

where μ_{ik} is the mean vector for R_{ik} , and Σ_{ik} is the covariance matrix for run R_{ik} . For each road k , we have five classes, one corresponding to each individual. Each data point $\bar{x} = \{\nu_\xi, \nu_\eta, \theta, \delta, P_f\}$ in R_{ij} , $j \neq k$, is now classified into class ω_{ik} according to the Bayes decision rule [7]

$$g_{lk}(\bar{x}) > g_{ik}(\bar{x}), \quad \forall l \neq i \quad (60)$$

where

$$g_{ik}(\bar{x}) = -\frac{1}{2} \bar{x}^T \Sigma_{ik}^{-1} \bar{x} + \Sigma_{ik}^{-1} \mu_{ik} \bar{x} - \frac{1}{2} \mu_{ik}^T \Sigma_{ik}^{-1} \mu_{ik} - \frac{1}{2} \log |\Sigma_{ik}| + \log P(\omega_{ik}) \quad (61)$$

$$P(\omega_{ik}) = 1/5, \quad i \in \{1, 2, 3, 4, 5\}, \quad k \in \{1, 2, 3\}. \quad (62)$$

In Tables III–V we report the percentage of data points in R_{ij} , $j \neq k$, which are classified in class ω_{lk} for $k = 1$, $k = 2$, and $k = 3$, respectively. We consider R_{ij} to be classified correctly when a plurality of the data from R_{ij} falls into class ω_{lk} , and observe, from the tables, that the Bayes optimal

classifier misclassifies seven out 30 (23%) of the runs. The performance of the similarity measure (0% error) therefore compares quite favorably.

Next, we present results for task-based classification. We select from each run R_{ij} all the left-turn maneuvers, and all the right-turn maneuvers. We get two resulting sets of maneuvers for each person

$$\alpha_i, \beta_i, \quad i \in \{1, 2, 3, 4, 5\} \quad (63)$$

where α_i corresponds to all the left-turn maneuvers for person i , and β_i corresponds to all the right-turn maneuvers for person i . We then split each of these sets into two—one to train a VQ codebook (or calculate the Bayesian statistics), the other to determine a similarity value (or the Bayesian classification). Tables VI and VII report the results for the similarity measure (seven-state HMM's), and the Bayesian classification. Note again that the similarity measure classifies all ten sets (five left-turn, five right-turn) correctly, while, the Bayes classifier misclassifies three out of ten (30%).

TABLE V
(a) SIMILARITY MEASURE CLASSIFICATION AND (b)
BAYES OPTIMAL CLASSIFICATION FOR ROAD #3 DATA

σ	$O(1, 3, 3)$	$O(2, 3, 3)$	$O(3, 3, 3)$	$O(4, 3, 3)$	$O(5, 3, 3)$
$O(1, 1, 3)$	0.773	0.144	0.131	0.106	0.376
$O(1, 2, 3)$	0.597	0.108	0.096	0.053	0.352
$O(2, 1, 3)$	0.249	0.869	0.340	0.333	0.015
$O(2, 2, 3)$	0.190	0.803	0.394	0.285	0.009
$O(3, 1, 3)$	0.280	0.442	0.806	0.397	0.025
$O(3, 2, 3)$	0.143	0.218	0.680	0.307	0.014
$O(4, 1, 3)$	0.095	0.210	0.297	0.676	0.010
$O(4, 2, 3)$	0.065	0.163	0.317	0.679	0.009
$O(5, 1, 3)$	0.047	0.003	0.005	0.011	0.503
$O(5, 2, 3)$	0.091	0.009	0.011	0.011	0.614

(a)

%	$\omega(1, 3)$	$\omega(2, 3)$	$\omega(3, 3)$	$\omega(4, 3)$	$\omega(5, 3)$
$R(1, 1)$	0.458	0.264	0.032	0.096	0.150
$R(1, 2)$	0.513	0.172	0.014	0.057	0.245
$R(2, 1)$	0.185	0.575	0.078	0.132	0.030
$R(2, 2)$	0.208	0.588	0.037	0.129	0.037
$R(3, 1)$	0.086	0.451	0.270	0.181	0.012
$R(3, 2)$	0.108	0.424	0.255	0.196	0.018
$R(4, 1)$	0.201	0.456	0.046	0.292	0.004
$R(4, 2)$	0.114	0.440	0.083	0.361	0.001
$R(5, 1)$	0.173	0.010	0.015	0.109	0.693
$R(5, 2)$	0.059	0.000	0.001	0.000	0.940

(b)

Finally, we present classification results for data which is more difficult to classify. Moe is asked to drive over the same road on two different days, two times each day, generating four runs (#1, #2, #3, #4). Because the runs are recorded on the same road, Moe is able to improve his skill relatively quickly. As recorded in Table VIII, his average speed improves from 65.9–71.9 mi/h. from run #1–#4. We take two additional data sets, one from Larry and one from Curly, over the same road. These data sets have similar aggregate statistics compared to at least some of Moe's runs.

Now we generate a 64-level VQ codebook with data from Larry's and Moe's fourth run, and another 64-level VQ codebook with data from Curly's and Moe's fourth run. We also generate corresponding classes for the Bayes-classifier comparison. We now classify each of Moe's first three runs as either similar to Larry or Moe #4, or as either similar to Curly or Moe #4.

Table IX shows the classification results based on the similarity measure and Bayesian statistics. We observe that the similarity measure misclassifies one out of six (17%), while

TABLE VI
(a) SIMILARITY MEASURE CLASSIFICATION AND (b)
BAYES OPTIMAL CLASSIFICATION FOR LEFT TURNS

σ	$\alpha(1)$	$\alpha(2)$	$\alpha(3)$	$\alpha(4)$	$\alpha(5)$
$\alpha(1)$	0.691	0.172	0.132	0.047	0.121
$\alpha(2)$	0.072	0.832	0.263	0.200	0.007
$\alpha(3)$	0.051	0.202	0.785	0.326	0.011
$\alpha(4)$	0.026	0.161	0.280	0.720	0.013
$\alpha(5)$	0.390	0.011	0.012	0.009	0.792

(a)

%	$\alpha(1)$	$\alpha(2)$	$\alpha(3)$	$\alpha(4)$	$\alpha(5)$
$\alpha(1)$	0.478	0.223	0.037	0.050	0.212
$\alpha(2)$	0.165	0.559	0.130	0.131	0.016
$\alpha(3)$	0.042	0.275	0.527	0.143	0.012
$\alpha(4)$	0.161	0.339	0.163	0.328	0.009
$\alpha(5)$	0.122	0.015	0.012	0.034	0.817

(b)

TABLE VII
(a) SIMILARITY MEASURE CLASSIFICATION AND (b)
BAYES OPTIMAL CLASSIFICATION FOR RIGHT TURNS

σ	$\beta(1)$	$\beta(2)$	$\beta(3)$	$\beta(4)$	$\beta(5)$
$\beta(1)$	0.713	0.118	0.109	0.036	0.223
$\beta(2)$	0.119	0.801	0.276	0.199	0.008
$\beta(3)$	0.109	0.390	0.742	0.324	0.009
$\beta(4)$	0.032	0.173	0.275	0.773	0.003
$\beta(5)$	0.244	0.006	0.014	0.005	0.875

(a)

%	$\beta(1)$	$\beta(2)$	$\beta(3)$	$\beta(4)$	$\beta(5)$
$\beta(1)$	0.253	0.188	0.015	0.052	0.492
$\beta(2)$	0.147	0.616	0.075	0.108	0.054
$\beta(3)$	0.105	0.330	0.412	0.137	0.015
$\beta(4)$	0.163	0.364	0.174	0.292	0.008
$\beta(5)$	0.105	0.014	0.024	0.028	0.830

(b)

the Bayes classifier misclassifies five out of six (83%), some quite badly.

C. Discussion

Here, we discuss two issues which have not yet been addressed in the above results. First, we demonstrate why the Bayes classifier fails in some cases, where the similarity measure succeeds. Fig. 11(a) plots the distribution (over ν and P_f) of Curly's data (Table VIII), and the Gaussian approxi-

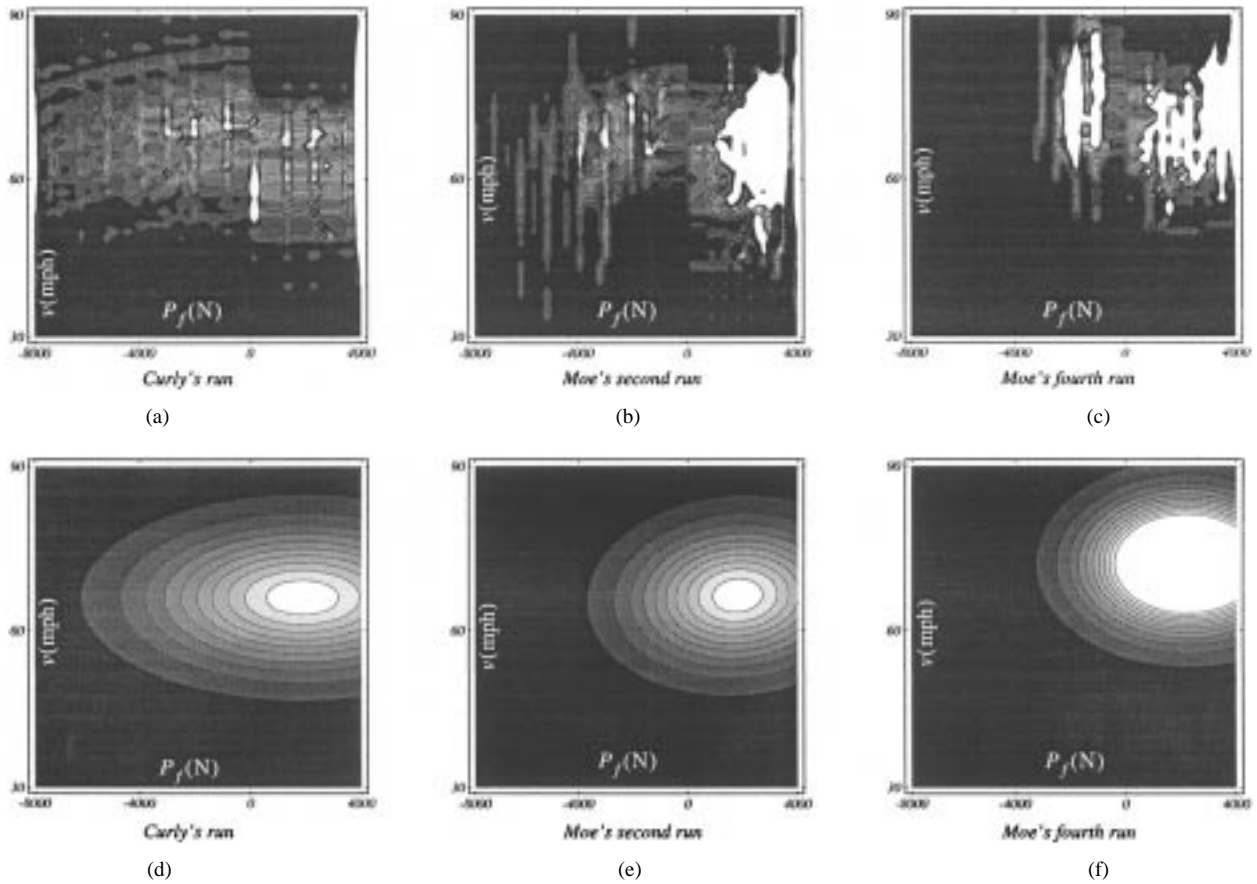


Fig. 11. Statistical distributions for (a) Curly's run, (b) Moe's second run, and (c) Moe's fourth run; and Gaussian approximations of each distribution for (d) Curly's run, (e) Moe's second run, and (f) Moe's fourth run.

TABLE VIII
AGGREGATE STATISTICS FOR ADDITIONAL HUMAN DATA

Name	v (mph)	$\hat{\theta}$ (rad/s)	δ (rad)	P_f (1000N)
Larry	61.9 ± 8.1	0.00 ± 0.18	0.00 ± 0.06	1.3 ± 1.5
Curly	65.6 ± 8.3	0.00 ± 0.20	0.00 ± 0.08	1.8 ± 3.5
Moe #1	66.1 ± 8.2	0.00 ± 0.23	0.00 ± 0.07	1.7 ± 1.8
Moe #2	65.9 ± 8.1	0.00 ± 0.21	0.00 ± 0.07	1.8 ± 2.4
Moe #3	67.4 ± 9.5	0.00 ± 0.23	0.00 ± 0.08	1.9 ± 2.6
Moe #4	71.9 ± 7.4	0.00 ± 0.24	0.00 ± 0.09	2.2 ± 2.1

mation of that distribution. Likewise, Fig. 11(b) and (c) show similar comparisons for Moe's second run and Moe's fourth run, respectively. It is clear that the Bayes classifier is doomed to fail, since the human data is distributed in a decidedly non-Gaussian manner. The similarity measure, on the other hand, succeeds because the HMM's are trained on the underlying distributions of the data sets, and make no *a priori* assumptions about each individual's distribution. We should also note that despite various attempts at improving the Bayes classifier's performance over the similarity measure—by only classifying on a subset of the vector $\{\nu_\xi, \nu_\eta, \hat{\theta}, \delta, P_f\}$ —we have yet to

TABLE IX
DIFFICULT CLASSIFICATIONS

σ	Larry	Moe #4	σ	Curly	Moe #4
Moe #1	0.572	0.528	Moe #1	0.315	0.616
Moe #2	0.435	0.540	Moe #2	0.495	0.603
Moe #3	0.258	0.728	Moe #3	0.550	0.760

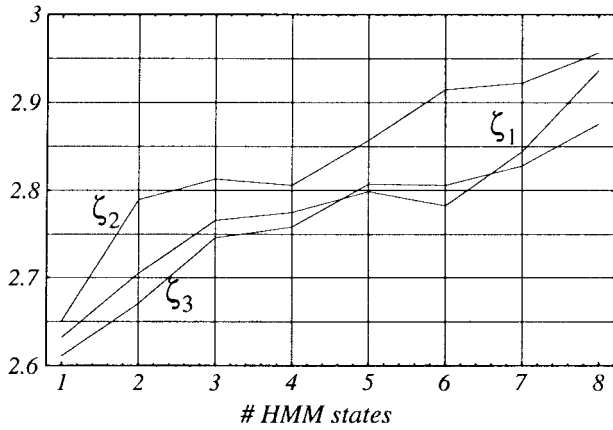
(a)

%	Larry	Moe #4	%	Curly	Moe #4
Moe #1	0.609	0.391	Moe #1	0.569	0.431
Moe #2	0.589	0.411	Moe #2	0.663	0.337
Moe #3	0.416	0.583	Moe #3	0.567	0.433

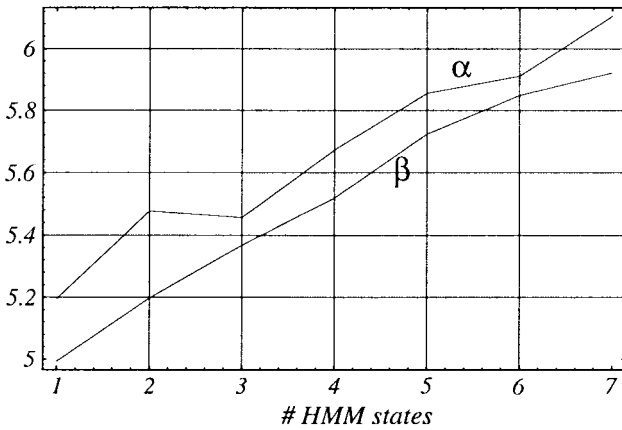
(b)

identify an example where the Bayes classifier succeeds and the similarity measure fails.

Second, we show that the increasing the number of states in the HMM's improves the discrimination capability of the similarity measure. Consider the results in Tables III–V, and



(a)



(b)

Fig. 12. The discrimination between people improves as a function of the number of states in the HMM. Note also that the administration for the task-based classification (b) is twice that for the classification across different roads (a).

define the following measure of discrimination ζ_k :

$$\zeta_k = \frac{\sum_{i,l \neq k} \sigma(O_{ik}^k, O_{il}^k)}{(n-1) \sum_{i \neq j, l \neq k} \sigma(O_{ik}^k, O_{jl}^k)}, \quad l, k \in \{1, 2, 3\}$$

$$i, j \in \{1, 2, 3, 4, 5\}, \quad n = 5. \quad (64)$$

Essentially, ζ_k forms a ratio of self-similarities over averaged cross-similarities between individuals for observation sequences vector quantized on codebook Q_k . Fig. 12(a) below plots this ratio as the number of states in the respective HMM's is varied from 1–8. Fig. 12(b) plots a similar ratio for the left-turn/right-turn classifications. From Fig. 12, we make two observations:

- 1) the discrimination of the similarity measure is affected positively by imparting structure onto the statistical model (i.e., the HMM) in the form of an increased number of HMM states;
- 2) the discrimination of the similarity measure improves significantly (two-fold in this example), when we train on specific maneuvers (i.e., left turns and right turns) rather than arbitrary roads.




Fig. 13. The cascade learning architecture adds hidden units one at a time to an initially minimal network. All connections in the diagram are feed-forward.

IV. VALIDATING HUMAN CONTROL STRATEGY MODELS

A. Learning Human Control Strategy

In learning human control strategy, such as driving, we wish to approximate the functional mapping between sensory inputs and control action outputs which guide an individual's actions. Human control strategy is dynamic, stochastic, and often highly nonlinear in nature. Little, if anything, is known *a priori* about the underlying structure, order, or granularity of an individual's internal controller. Consequently, we require a flexible, nonlinear learning architecture, capable of generating a wide spectrum of mappings from smooth to discontinuous, linear to nonlinear. Cascade neural networks [27] with variable activation functions [28]–[30] offer such flexibility.

Cascade neural networks are feed-forward neural networks. In cascade learning, the structure of the network is adjusted as part of the training process by adding hidden units one at a time to an initially minimal network. Hidden units are added in a cascading fashion, with a new hidden unit taking input from all previous hidden units (see Fig. 13). Moreover, each of these hidden units can assume a variable activation function, which is not restricted to simply the sigmoidal nonlinearity. This flexibility in functional form leads to efficiency in learning speed and good function approximation properties [30].

Below, we report cascade modeling results for human control data collected from the driving simulator described in Section III-A above. The inputs to the cascade networks for each individual include

- 1) current and previous state information (65);
- 2) previous control information (66);
- 3) description of the road (67)

$$\{\nu_\xi(k - n_s), \dots, \nu_\xi(k - 1), \nu_\xi(k), \nu_\eta(k - n_s) \dots$$

$$\nu_\eta(k - 1), \nu_\eta(k), \dot{\theta}(k - n_s), \dots, \dot{\theta}(k - 1), \dot{\theta}(k)\} \quad (65)$$

$$\{\delta(k - n_c), \dots, \delta(k - 1), \delta(k), P_f(k - n_c), \dots$$

$$P_f(k - 1), P_f(k)\} \quad (66)$$

$$\{x(1), x(2), \dots, x(n_r), y(1), y(2), y(n_r)\} \quad (67)$$

where n_s = length of state history to include as input, and n_c = length of command history to include as input. For the road description, we discretize the visible view (given a specific horizon) of the road ahead into n_r equivalently spaced, body-relative (x, y) coordinates of the road median,

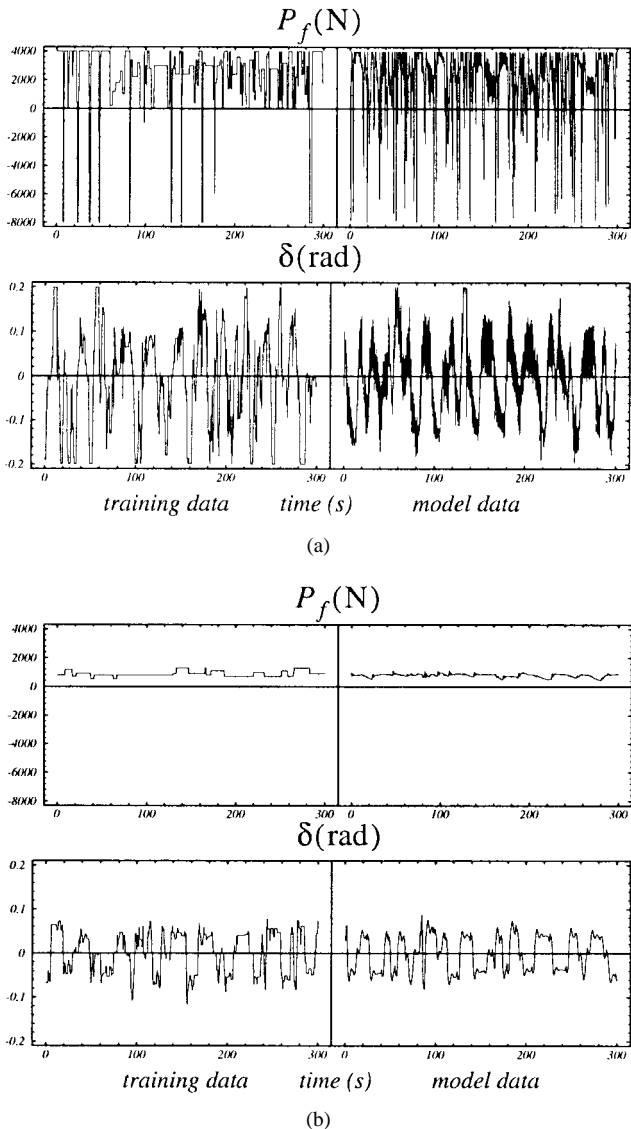


Fig. 14. (a) Oliver's driving data and (b) Stan's driving data. On the left of each figure, we show part of the source training data; on the right of each figure we show the corresponding model-generated data.

and provide that sequence of coordinates as input to the network. Thus, for fixed n_s, n_c , and n_r , the total number of inputs to the network will be

$$n_{\text{inputs}} = 3n_s + 2n_c + 2n_r. \quad (68)$$

The outputs for the cascade network are, of course, $\{\delta(k+1), P_f(k+1)\}$ (i.e., the steering and acceleration command for the next time step).

B. Similarity Results

The left sides of Fig. 14(a) and (b) show part of the driving data collected from two individuals, Oliver and Stan. Note that the driving styles for the two individuals are quite different for the same road. The right sides of Fig. 14(a) and (b) show part of the cascade model-generated command trajectories for Oliver and Stan. Since Stan's control strategy is relatively simple, his control strategy model requires 30 hidden units and only the previous two states as input (i.e., $n_s = n_c = 2$).

TABLE X
HUMAN-TO-MODEL SIMILARITY RESULTS

σ	Stan	Stan's model	Oliver	Oliver's model
Stan	1.000	0.748	0.009	0.007
Stan's model	0.748	1.000	0.006	0.004
Oliver	0.009	0.006	1.000	0.349
Oliver's model	0.007	0.004	0.349	1.000

Oliver's more complicated control strategy model, on the other hand, requires 32 hidden units and relies on the previous ten states ($n_s = n_c = 10$) to stay on the road. We determine these input histories experimentally to achieve stable road following for each HCS model. For both models, we let $n_r = 15$.

Table X below summarizes the similarity results for the data in Fig. 14. All preprocessing and vector quantization was performed in the same manner as in Section III-B.

C. Discussion and Future Work

The similarity results in Table X confirm two qualitative assessments of the data in Fig. 14(a) and (b). First, we observe that the two driving styles are objectively quite different. This fact is reflected in the low similarity measures between one individual's model and the other individual's source and model-generated data. Second, Stan's model is a better reflection of his driving style, than Oliver's model is of his, as reflected in the two respective similarity measures, 0.748 and 0.349. This is indicative that Oliver's sharply discontinuous driving strategy is more difficult to learn by a single cascade network than Stan's calmer approach. Indeed, Oliver's model generates significant oscillatory behavior, of which Oliver himself is not guilty.

Thus, the relatively low similarity measure between Oliver and Oliver's model points to a problem in the model itself, including possible improper assumptions about the controller order, granularity and control delay of the actual HCS. As a matter of fact, the values for n_s, n_c , and n_r for each model were arrived at in an essentially *ad hoc* manner. To correct this problem, we are at present working on an algorithm to improve a model's input representation based on the similarity measure. We propose combining the similarity measure with simultaneously perturbed stochastic approximation (SPSA) [31] to select the best model input representation.

In general, a given control strategy can be approximated by

$$\begin{aligned} \bar{u}[(k+1)\tau + \delta] \\ = \Gamma(\bar{x}[k\tau], \bar{x}[(k-1)\tau], \dots, \bar{x}[(k-m)\tau]) \\ \bar{u}[k\tau], \bar{u}[(k-1)\tau], \dots, \bar{u}[(k-n)\tau]) \end{aligned} \quad (69)$$

where $\Gamma(\cdot)$ is some arbitrary unknown function, $\bar{u}(k)$ is a vector of control outputs, $\bar{x}(k)$ is a vector of sensory inputs (including both state and environment variables) at time step k , τ indicates the controller resolution or granularity, and δ

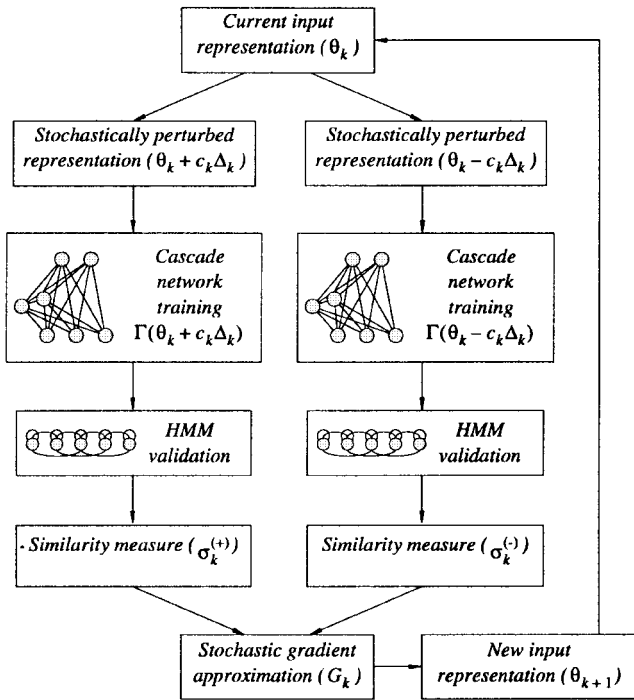


Fig. 15. Overall approach to model refinement using the similarity measure.

indicates the control delay. The order of the dynamic system is given by the constants n and m . In order to select the best input representation for the HCS model, we need to optimize the parameter vector

$$\theta = [m \quad n \quad \tau \quad \delta]^T \quad (70)$$

for some given experimental data. Let, $\Gamma(\theta)$ denote a trained HCS model with input representation θ , and let $\sigma[\Gamma(\theta)] = \sigma(\theta)$ denote the similarity measure for model $\Gamma(\theta)$. Now, the best input representation θ^* is defined in terms of the similarity measure σ , such that

$$\sigma(\theta^*) > \sigma(\theta) \quad \forall \theta \neq \theta^*. \quad (71)$$

This optimization is difficult in principle since

1) we have no explicit gradient information

$$G(\theta) = \frac{\partial}{\partial \theta} \sigma(\theta) \quad (72)$$

2) each measurement of σ is computationally expensive.

Thus, we resort to simultaneously perturbed stochastic approximation to perform the optimization. We propose to evaluate the similarity measure σ for two values of θ in order to arrive at a new estimate for θ at iteration k . Fig. 15 illustrates the overall loop, where Δ_k is a vector of mutually independent, mean-zero random variables (e.g., symmetric Bernoulli distributed), the sequence $\{\Delta_k\}$ is independent and identically distributed, and the $\{a_k\}, \{c_k\}$ are positive scalar sequences.

V. CONCLUSION

Model validation is an important problem in the area of machine learning for dynamic systems, if learned models are to be exploited to their full potential. In this paper,

we have derived a stochastic similarity measure, based on HMM analysis, by which we can compare and contrast arbitrary, multidimensional, stochastic trajectories. We have shown that this method performs significantly better than traditional statistical techniques in classifying human control strategy data. Furthermore, we have shown that the similarity measure offers a feasible means of validating a learned HCS model's performance to its training data. Finally, we have proposed an iterative algorithm, based on the similarity measure, which allows us to refine a model's input representation to improve model fidelity. Such an algorithm is of special relevance in learning human control strategy, where little is known *a priori* about the structure, order, or granularity of each individual's human controller.

ACKNOWLEDGMENT

The authors would like to thank the many people who patiently "drove" through our driving simulator. Special thanks go to O. Barkai, A. Gove, C. Lee, J. Murphy, and M. Ollis for their patience as they navigated through some 50 mi of simulated road.

REFERENCES

- [1] R. Basri and D. Weinshall, "Distance metric between 3-D models and 2-D images for recognition and classification," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 43, pp. 465-479, 1996.
- [2] M. Boninsegna and M. Rossi, "Similarity measures in computer vision," *Pattern Recognit. Lett.*, vol. 15, no. 12, pp. 1255-1260, 1994.
- [3] M. Werman and D. Weinshall, "Similarity and affine invariant distances between 2-D point sets," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 17, pp. 810-814, 1995.
- [4] R. Jain *et al.*, "Similarity measures for image databases," in *Proc. IEEE Int. Conf. Fuzzy Syst.*, vol. 3, pp. 1247-1254, 1995.
- [5] K. Y. Kupeev and H. J. Wolfson, "On shape similarity," in *Proc. 12th IAPR Int. Conf. Pattern Recog.*, 1994, vol. 1, pp. 227-231.
- [6] H. Y. Shum, M. Hebert, and K. Ikeuchi, "On 3-D shape similarity," Tech. Rep. CMU-CS-95-212, Carnegie Mellon Univ., Pittsburgh, PA, 1995.
- [7] R. O. Duda and P. E. Hart, *Pattern Classification and Scene Analysis*. New York: Wiley, 1973.
- [8] M. Sun, G. Burk, and R. J. Scabassi, "Measurement of signal similarity using the maxima of the wavelet transform," in *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Process.*, 1993, vol. 3, pp. 583-586.
- [9] J. Hertz, A. Krogh, and R. G. Palmer, *Introduction to the Theory of Neural Computation*. Redwood City, CA: Addison-Wesley, 1991.
- [10] L. G. Sotilino, M. Saerens, and H. Bersini, "Classification of temporal trajectories by continuous-time recurrent nets," *Neural Networks*, vol. 7, no. 5, pp. 767-776, 1994.
- [11] X. D. Huang, Y. Ariki, and M. A. Jack, *Hidden Markov Models for Speech Recognition*. Edinburgh, U.K.: Edinburgh Univ. Press, 1990.
- [12] L. R. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," *Proc. IEEE*, vol. 77, pp. 257-286, 1989.
- [13] A. Kundu, G. C. Chen, and C. E. Persons, "Transient sonar signal classification using hidden Markov models and neural nets," *IEEE J. Oceanic Eng.*, vol. 19, pp. 87-99, 1994.
- [14] G. Radons, J. D. Becker, B. Dulfer, and J. Kruger, "Analysis, classification and coding of multielectrode spike trains with hidden Markov models," *Biol. Cybern.*, vol. 71, no. 4, pp. 359-373, 1994.
- [15] B. Hannaford and P. Lee, "Hidden Markov model analysis of force/torque information in telemanipulation," *Int. J. Robot. Res.*, vol. 10, no. 5, pp. 528-539, 1991.
- [16] J. Yamato, S. Kurakake, A. Tomono, and K. Ishii, "Human action recognition using HMM with category separated vector quantization," *Trans. Inst. Electron., Inform. Comm. Eng. D-II*, vol. J77D-II, no. 7, pp. 1311-1318, 1994.
- [17] J. Yamato, J. Ohya, and K. Ishii, "Recognizing human action in time-sequential images using hidden Markov models," *Trans. Inst. Electron., Inform., Comm. Eng. D-II*, vol. J76D-II, no. 12, pp. 2556-2563, 1993.

- [18] J. Yang, Y. Xu, and C. S. Chen, "Hidden Markov model approach to skill learning and its application to telerobotics," *IEEE Trans. Robot. Automat.*, vol. 10, pp. 621–631, 1994.
- [19] ———, "Human action learning via hidden Markov models," *IEEE Trans. Syst., Man, Cybern A*, vol. 27, pp. 34–44, Jan. 1997.
- [20] L. R. Rabiner, B. H. Juang, S. E. Levinson, and M. M. Sondhi, "Some properties of continuous hidden Markov model representations," *AT&T Tech. J.*, vol. 64, no. 6, pp. 1211–1222, 1986.
- [21] L. E. Baum, T. Petrie, G. Soules, and N. Weiss, "A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains," *Ann. Math. Stat.*, vol. 41, no. 1, pp. 164–171, 1970.
- [22] M. C. Nechyba and Y. Xu, "On the fidelity of human skill models," in *Proc. IEEE Int. Conf. Robot. Automat.*, 1996, vol. 3, pp. 2688–2693.
- [23] B. H. Juang and L. R. Rabiner, "A probabilistic distance measure for hidden Markov models," *AT&T Tech. J.*, vol. 64, no. 2, pp. 391–408, 1985.
- [24] K. R. Rao and D. F. Elliott, *Fast Transforms: Algorithms, Analyzes and Applications*. New York: Academic, 1982.
- [25] Y. Linde, A. Buzo, and R. M. Gray, "An algorithm for vector quantizer design," *IEEE Trans. Commun.*, vol. COM-28, pp. 84–95, 1980.
- [26] H. Hatwal and E. C. Mikulcik, "Some inverse solutions to an automobile path-tracking problem with input control of steering and brakes," *Veh. Syst. Dynam.*, vol. 15, pp. 61–71, 1986.
- [27] S. E. Fahlman and C. Lebiere, "The cascade-correlation learning algorithm," Tech. Rep. CMU-CS-90-100, Carnegie Mellon Univ., Pittsburgh, PA, 1990.
- [28] M. C. Nechyba and Y. Xu, "Neural network approach to control system identification with variable activation functions," in *Proc. IEEE Int. Symp. Intell. Contr.*, 1994, vol. 1, pp. 358–363.
- [29] ———, "Human skill transfer: Neural networks as learners and teachers," in *Proc. IEEE Int. Conf. Intell. Robots Syst.*, vol. 3, pp. 314–319, 1995.
- [30] ———, "Toward human control strategy learning: Neural network approach with variable activation functions," Tech. Rep. CMU-RI-TR-95-09, Carnegie Mellon Univ., Pittsburgh, PA, 1995.
- [31] J. C. Spall, "Multivariate stochastic approximation using a simultaneous perturbation gradient approximation," *IEEE Trans. Automat. Contr.*, vol. 37, pp. 332–341, 1992.



Dr. Nechyba was an NSF Graduate Fellow as well as a Department of Energy Doctoral Fellow.

Michael C. Nechyba (S'94) received the B.S. degree in electrical engineering from the University of Florida, Gainesville, in 1992 and the Ph.D. degree from the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, in 1998.

His research interests include the learning, validation, and transfer of human control strategies, human-centered robotics, teleoperation and human interfaces, robot control, machine learning for control, computational intelligence, neural networks, and hidden Markov models.



Yangsheng Xu (M'90–SM'95) received the Ph.D. degree in 1989 from the University of Pennsylvania, Philadelphia.

He joined the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, in 1989, and has been leading research efforts in the area of space robotics and robot control. He developed several space robots and dynamically stable robot systems in the space robotics laboratory that he established in 1990. He has also contributed to advances in intelligent control in the area of human control strategy learning and the automatic transfer of human skill to robots as well as humans, based on neural networks and hidden Markov models.

Dr. Xu has been selected as a Keynote Speaker in several international conferences and as a Senior Technical Advisor for the United Nations Development Program. He is a Senior Member of the American Institute of Aeronautics and Astronautics and a Member of the New York Academy of Science.