Title: A Segmentation Algorithm for the Comparison of Human Limb Trajectories

Authors: Marc Pomplun and Maja J. Matarić

Running head: Segmentation Algorithm

Corresponding author:

Marc Pomplun

Department of Computer Science

York University

4700 Keele Street

Toronto, Ontario

Canada M3J 1P3

e-mail: marc@cs.yorku.ca

phone: (416) 736-2100 ext. 40629 or (416) 926-0921

A Segmentation Algorithm for the Comparison of Human Limb Trajectories

Marc Pomplun[1] and Maja J. Matarić[2]

[1]Department of Computer Science, York University

4700 Keele Street, Toronto, Ontario, Canada M3J 1P3

e-mail: marc@cs.yorku.ca

[2]Department of Computer Science, University of Southern California

Los Angeles, CA 90089-0781, USA

e-mail: mataric@cs.usc.edu

Abstract

In this article, we present an analytical approach for movement evaluation and comparison. We develop a joint-space-based segmentation and comparison algorithm that calculates pairwise similarity of arm movement trajectories. To validate the algorithm, we collected sixty sample arm movement sequences from a subject imitating video demonstrations. We present the performance of our algorithm on that data set, and discuss the contribution of each of its components.

A Segmentation Algorithm for the Comparison of Human Limb Trajectories

Comparing and evaluating movement is a known difficult problem in psychophysics, as well as in robotics. Our particular area of interest is in motor skill learning by imitation, both in humans and humanoid robots (Byrne & Russon, 1998; Matarić, 2001), where the need for analytical movement comparison is particularly acute. Our earlier work (Matarić & Pomplun, 1997, 1998) in the area of psychophysics of imitation involved studying the attentional behavior of people watching movements for subsequent imitation. In more recent experiments, we aimed to answer more precise questions about the quality of the subjects' imitation under various conditions, by recording the demonstrator's and the subjects' movements for subsequent analysis and comparison. This presented an important evaluation challenge. What is the most effective means of evaluating arm movements across subjects and tasks?

In the present paper, we try to answer this question by developing suitable evaluation metrics that allow us to make quantitative statements about human movement, resulting from imitation or otherwise. We are not aware of any other work aimed directly at analytically evaluating imitated behavior, and the work presented here is at its very beginning.

Research into copying written 2D characters, drawings, and trajectories has been done (see Cooke, Kitts, Matarić & Sekuler, 1997; Kitts, Cooke, Matarić & Sekuler, 1997). In our previous work on evaluation of motor behavior, but not imitation behavior (Matarić, Zordan & Mason, 1998), we explored end-effector jerk as a quantitative measure of motor behavior, but found it to be overly sensitive to specific parameters of each implementation. Other measures, such as minimum torque change (Uno, Kawato & Suzuki, 1989) and minimum energy (Nelson, 1983), have also been explored. None of these alone provides a sufficient description of the quality of a movement, and in particular of a motor imitation, where the underlying intentions of the demonstrator, however complex, may be relevant. Depending on the nature of the task, imitation can be as simple as achieving the observed task (so-called "task level" imitation, see Tomasello, Kruger & Rather, 1993), something that is quite easily evaluated, to being as complex as conveying the underlying message of a dance or a gesture. The latter has so far only been studied qualitatively.

Studies evaluating task-irrelevant movements (see Goncalves, Di Bernardo, Ursella & Perona, 1995; Goncalves, Di Bernardo & Perona, 1998, 1999) compared human movements to the output of computer vision systems tracking and modeling these movements by using a simple mean square error (MSE) metric. They demonstrated that in many cases even pairs of movements with large MSE were still perceived as identical by human observers. This motivates the use of more sophisticated metrics to quantitatively assess the similarity of limb trajectories. Evaluation of imitation is an open problem, and the evaluation metrics developed in the present paper are meant as a general tool for both human and robot imitation as well as motor control evaluation.

Method

Our experiment was conducted at the National Institutes of Health Resource for the Study of Neural Models of Behavior, at the University of Rochester.

Subjects. The experimental trials reported in this article were performed by one of the authors (M. M.), who is right-handed.

Stimuli. The stimuli consisted of 20 video clips, each showing an individual sequence of human arm movements. These movements had been performed by one of the authors (M. M.) and video recorded. The video clips presented her arm moving in front of a black background; each clip was about 3 to 5 seconds long. All sequences started with a straightened arm and a clenched fist and typically involved movements of the shoulder, elbow, wrist, and fingers.

Apparatus. Stimuli were presented on a 22-inch monitor at a viewing distance of 100 cm. Arm movements were tracked with the FastTrak motion tracking system. For this purpose, four position sensors were fastened to the subject's right arm, one in each of the following four positions (see Figure 1):

1. at the center of the upper arm,
2. at the center of the lower arm,
3. immediately above the wrist, and
4. on the upper phalanx of the middle finger.

———- insert Figure 1 here ———-

Each sensor was about 2×1×1 cm in size. We used elastic bands to hold the sensors and their cables in place. The cables were taken to the subject's right shoulder and then connected to the FastTrak processing computer. The spatial position of each sensor was measured every 34 ms with an accuracy of about 2 mm in each dimension.

Procedure. Prior to each recording trial, one of the 20 stimuli was shown to the subject. To initiate the imitation process after each stimulus presentation, the experimenter gave the verbal signal "ready". The subject then straightened out her arm horizontally towards the monitor screen and clenched the fist. Next, the subject received the instruction "imitate" and started to imitate the previously viewed sequence of movements as precisely as possible. When finished imitating, the subject said "done" and put the arm back into the resting position. Between the "imitate" and "done" signals, the subject's arm movements were recorded by the FastTrak system. In the trials reported here, M.M. imitated each of the 20 stimuli three times, resulting in a total of 60 reference trajectories being recorded.

Data Analysis

We developed a metric for the distance (or difference) between arm movement trajectories that can be used to evaluate the imitation of movements. Instead of comparing the spatial Cartesian coordinates of the four sensors, we chose to apply evaluation in joint-space, because joint angles are independent from the length of the subject's limbs, thus compensating for individual physical differences. We derived the following five angles from the Cartesian coordinates of the four position sensors:

1. horizontal shoulder angle
2. vertical shoulder angle
3. shoulder rotation (along the axis going from the shoulder to the elbow)
4. elbow
5. wrist (including finger angles)

Figure 2a shows these five angles as functions of time, as they were recorded during the imitation of a sample stimulus. Most conspicuously, the movement sequence involves a strong rotation of the shoulder joint by up to 160 degrees, and the repetition of a movement pattern that is about 1.8 s long.

Starting from a simple mean square error (MSE) metric between corresponding joint angle samples, we developed a more sophisticated and adequate metric by successively adding new components that reduced noise-sensitivity. As a result, such a distance metric cannot be sensitive to all of the many diverse aspects of imitation performance. It is, however, a straightforward, quantitative approach meeting the demands of future studies aimed at evaluating imitated behavior.

In the present data, the sampling rate was a constant 34ms. Consequently, the following explanation of the metrics will refer to sample indices instead of time-codes. This in no way restricts the generality of the described methods; trajectories with differing or shifting sampling rates can still be analyzed after conducting a time scaling operation we describe below.

To establish notation, let us assume that we want to compare two arm movement sequences (trajectories) α and β containing $(T_\alpha + 1)$ and $(T_\beta + 1)$ samples respectively for each of the J joint angles. We will refer to the samples as $\alpha_t^{(j)},\ j = 1, \dots, J,\ t = 0, \dots, T_\alpha$ and $\beta_t^{(j)},\ j = 1, \dots, J,\ t = 0, \dots, T_\beta$. Our aim is to define a distance metric $d(\alpha, \beta)$ yielding lower values for more similar trajectories α and β.

We now describe four comparison methods of increasing effectiveness: raw trajectory comparison, time scaling, separate segmentation, and combined segmentation.

Raw Trajectory Comparison. The mean square distance between corresponding angles in α and β trajectories, across joints and samples, is a simple and straightforward difference metric that ignores possible additional samples in the longer trajectory:

$$d(\alpha, \beta) = \sum_{t=0}^{\min(T_\alpha, T_\beta)} \sum_{j=1}^{J} (\alpha_t^{(j)} - \beta_t^{(j)})^2 \tag{1}$$

In order to compensate for differences between α and β trajectories in the starting position of the arm, before calculating (1), time-invariant offsets are added to all joint angles in both trajectories so that all angles start at 0.

The most evident shortcoming of this metric is that a slight difference in velocity can cumulatively grow into a large distance between the two trajectories being compared. To eliminate this undesirable feature, a time scaling procedure can be performed prior to the calculation of mean square distances.

Time Scaling. In order to derive a metric that is invariant to the absolute speed of movement and accounts for all data in both α and β trajectories, the shorter trajectory, say α, is expanded to contain as many samples as its counterpart. To do this, a continuous function $\tilde{\alpha}^{(j)}(\tilde{t})$ interpolating α is needed:

$$\tilde{\alpha}^{(j)}(\tilde{t}) = \alpha_t^{(j)} \quad \forall\ j = 1, \ldots, J,\ \ t = 0, \ldots, T_\alpha,\ \ \tilde{t} = t \tag{2}$$

For most kinds of trajectories, spline interpolation is an adequate method. Based on the interpolation, the samples in α can be rearranged to match the number of their counterparts in β:

$$\alpha_{t,new}^{(j)} = \tilde{\alpha}^{(j)}\left(\frac{t \cdot T_\alpha}{T_\beta}\right) \quad \forall\ j = 1, \ldots, J,\ \ t = 0, \ldots, T_\beta \tag{3}$$

$$T_{\alpha,new} = T_\beta \tag{4}$$

After time scaling, the metric $d(\alpha, \beta)$ can still be calculated as the mean square distance between α and β, according to Equation (1) above.

Separate Segmentation. Since in the present context trajectories are sequences of movements, the metric should benefit from identifying the elementary trajectory segments and comparing them individually. A useful indicator of transitions between successive segments in a trajectory α is the velocity $\dot{\alpha}_t$ calculated across all joint angles (see Figure 2b):

$$\dot{\alpha}_t = \sqrt{\frac{1}{J} \sum_{j=1}^{J} (\alpha_t^{(j)} - \alpha_{t-1}^{(j)})^2}, \quad t = 1, \ldots, T_\alpha \tag{5}$$

When examining velocity histograms within any of the recorded trajectories, we always find the highest and widest peak at a very low velocity. This peak reflects the almost motionless phases in the trajectory that occur between successive segments and also at the beginning and the end of each trajectory. The basic idea of our segmentation algorithm is to determine contiguous phases in the trajectory with velocities above a certain threshold $\theta_{\dot{\alpha}}$. Such contiguous phases exceeding a specific minimum duration are considered to be the segments of the trajectory.

We determined empirically that an appropriate value for the velocity threshold can be estimated from the position $\dot{\alpha}_{max}$ of the peak in the velocity histogram:

$$\theta_{\dot{\alpha}} = 2.5 \cdot \dot{\alpha}_{max} \tag{6}$$

For the present trajectories, the minimum duration for segments was set to 200 ms. The segmentation of the sample trajectory with these parameters is illustrated in Figure 2c.

———- insert Figure 2 here ———-

In many cases, the segmentation of the α and β trajectories to be compared leads to different numbers of segments. To calculate the distance $d(\alpha, \beta)$ in such cases, the segmentation of the trajectory with more segments, say α, is adjusted, i.e., the number of segments in α is reduced to equal those in β through the use of an iterative algorithm. In each step, the algorithm combines two successive segments in α into one segment. The segments to be combined are those whose combination leads to the smallest sum of temporal deviations between corresponding segment boundaries in α and β. This procedure is repeated until the same number of segments in α and β is reached.

Within corresponding segments, for example the first segments in α and β, time scaling is carried out as shown in Equations (2) to (4). Subsequently, the metric $d(\alpha, \beta)$ is calculated as the mean square distance across all samples and segments, which corresponds to Equation (1). We call this the "separate segmentation" algorithm, because the segmentations for the trajectories α and β to be compared are carried out separately before their distance $d(\alpha, \beta)$ is determined.

Combined Segmentation. Although the threshold estimation used in the "separate segmentation" algorithm above works well for most trajectories, there are cases in which two rather similar trajectories are segmented clearly differently, which leads to an inappropriately large distance between them. To avoid such distortions in the metric, a "combined segmentation" algorithm is employed, which finds the two thresholds that lead to the best match between the two trajectories being compared.

Instead of estimating an appropriate threshold $\theta_{\dot{\alpha}}$ for trajectory α, the "combined segmentation" algorithm determines a maximum threshold $\theta_{\dot{\alpha},max}$ limiting the range of thresholds to be examined. Empirical investigations found the following equation to yield best results:

$$\theta_{\dot{\alpha},max} = 3.5 \cdot \dot{\alpha}_{max} \tag{7}$$

After determining $\theta_{\dot{\alpha},max}$ and, analogously, $\theta_{\dot{\beta},max}$ for the trajectories α and β, the algorithm systematically performs segmentations with different combinations of velocity thresholds $\theta_{\dot{\alpha}}$, ranging from 0 to $\theta_{\dot{\alpha},max}$, and $\theta_{\dot{\beta}}$, ranging from 0 to $\theta_{\dot{\beta},max}$. For the trajectories in our data set, it was sufficient to use 20 levels for each of the two variables, resulting in 400 combinations to be tested. For each combination of segmentations, the mean square distance between α and β was calculated using the "separate segmentation" algorithm. $d(\alpha, \beta)$ was determined as the minimum distance found within these 400 combinations, i.e., the distance for the best-matching segmentations of α and β.

Results and Discussion

In order to evaluate the performance of the distance metric and the contribution of each of its four components, we compared the intra-stimulus distances to the inter-stimulus distances in our set of 60 reference movements. Intra-stimulus distances are distances between pairs of trajectories for the same stimulus, i.e., imitations of the same movement. Since we had three imitations for each stimulus, there were three trajectories and thus three distances between trajectory pairs (1-2, 1-3, 2-3), leading to a total of 60 intra-stimulus distances. The term inter-stimulus distances refers to distances between pairs of trajectories for different stimuli, i.e., imitations of different movements. Inter-stimulus distances were measured between each stimulus S and those 57 trajectories corresponding to stimuli other than S, resulting in a total of 1710 inter-stimulus distances. We would expect a more accurate and reliable metric to produce a clearer distinction between intra-stimulus distances (low values) and inter-stimulus distances (high values). In other words, a high signal-to-noise ratio, i.e., a large difference between the mean intra-stimulus and inter-stimulus values combined with low standard deviations, would indicate a desirable metric.

Figure 3a demonstrates the performance of the metric by showing intra- and inter-stimulus distances for each of the four variants of the metric (raw samples, time scaling, time scaling and separate segmentation, time scaling and combined segmentation). For each variant, these data were normalized in such a way that the mean intra-stimulus distance was always 1. As can clearly be seen, adding a component always led to an improvement in the performance of the metric, with the last component, combined segmentation, having the most pronounced effect.

The four stages of the distance metric were also evaluated in a classification test. For the first trajectory α of the first stimulus, for example, we calculated its mean distance from the second and third trajectory for each stimulus, resulting in 20 distance values. Trajectory α was then assigned to the stimulus whose second and third trajectories had the shortest mean distance from α. In our example, if the trajectory was assigned to the first stimulus, a correct classification was made. This classification test was performed on all 60 stimuli. Second trajectories were compared to first and third trajectories, and third trajectories were compared to first and second trajectories.

———- insert Figure 3 here ———-

Figure 3b presents the number of incorrect classifications produced by each of the four variants of the metric. As can be seen, each additional component reduces the number of incorrect classifications, with the most sophisticated variant, time scaling and combined segmentation, performing perfect classification on the given set of trajectories. Taken together, the above two evaluations of the metric reveal both the usefulness of all its components and the remarkably good performance of the final combined metric.

The evaluation approach we describe is general and can be applied to any joint-space representation of movements. We are currently using it to compare human performance to humanoid avatar performance in the context of implementing and evaluating an imitation model for complex humanoid robots. As discussed earlier, this is just one means of quantitatively evaluating and comparing joint-space trajectories. It does not address the more complex issue of task-level imitation. However, it does provide a reliable and robust quantitative metric for movement evaluation, imitation and motor control.

Acknowledgments

The research described here was supported by the NSF Career Grant IRI-9624237 to M. Matarić. The data used to evaluate the algorithms were gathered by the authors as part of a joint interdisciplinary project conducted at the National Institutes of Health Resource for the Study of Neural Models of Behavior, at the University of Rochester. The authors are grateful to Prof. Dana Ballard and Prof. Mary Hayhoe for hosting their experiments and to Prof. Jeff Pelz for dedicated assistance during the experiments as well as afterwards.

References

Byrne, R. W. & Russon, A. E. (1998). Learning by imitation: a hierarchical approach, The Behavioural and Brain Sciences, 21, 667–721.

Cooke, S., Kitts, B., Matarić, M.J. & Sekuler, B. (1997). Delayed and real-time imitation of complex visual gestures, in Proceedings, International Conference on Vision, Recognition, Action: Neural Models of Mind and Machine, Boston University.

Goncalves, L., Di Bernardo, E., Ursella, E. & Perona, P. (1995). Monocular tracking of the human arm in 3D, in Proceedings, ICCV 95.

Goncalves, L., Di Bernardo, E. & Perona, P. (1998). Reach out and touch space (motion learning), in Proceedings FG98, International Conf. on Face and Gesture Recognition.

Goncalves, L., Di Bernardo, E. & Perona, P. (1999). Monocular perception of biological motion in Johansson Displays, in Proceedings of ICCV 99, Corfu, Greece.

Kitts, B., Cooke, S., Matarić, M.J. & Sekuler, R. (1997). Improved pattern recognition by combining invariance methods, in Proceedings, International Conference on Vision, Recognition, Action: Neural Models of Mind and Machine, Boston University.

Matarić, M.J. & Pomplun, M. (1997). What do people look at when watching human movement?, Technical Report CS-97-194, Brandeis University.

Matarić, M.J. & Pomplun, M. (1998). Fixation behavior in observation and imitation of human movement, Cognitive Brain Research, 7, 191–202.

Matarić, M.J., Zordan, V. B. & Mason, Z. (1998). Movement control methods for complex, dynamically simulated agents: Adonis dances the macarena, in Autonomous Agents, ACM Press, Minneapolis, St. Paul, MI, 317–324.

Matarić, M.J. (2001). Sensory-motor primitives as a basis for imitation: Linking perception to action and biology to robotics, in C. Nehaniv & K. Dautenhahn, eds, Imitation in Animals and Artifacts, The MIT Press.

Nelson, W.L. (1983). Physical principles for economies of skilled movements, Biological Cybernetics, 46, 135–147.

Tomasello, M., Kruger, A. C., & Rather, H. H. (1993). Cultural learning, The Journal of Behavioral and Brain Sciences, 16, 495–552.

Uno, Y., Kawato, M. & Suzuki, R. (1989). Formation and control of optimal trajectory in human arm movement-minimum torque-change model, Biological Cybernetics, 61, 89–101.

Figure Captions

Figure 1: Arrangement of the four FastTrak sensors on the subject's right arm

Figure 2: Segmentation of a recorded sample movement sequence. (a) The five joint angles as functions of time during the movement sequence, (b) the overall velocity as a function of time, and (c) threshold-based segmentation of the sequence.

Figure 3: Evaluation of the four variants of the metric. (a) Mean intra- and inter-stimulus distances and their standard deviations, and (b) incorrect classifications (out of 60)

Note: Panels 2a, 2b, and 2c should be arranged in a column.

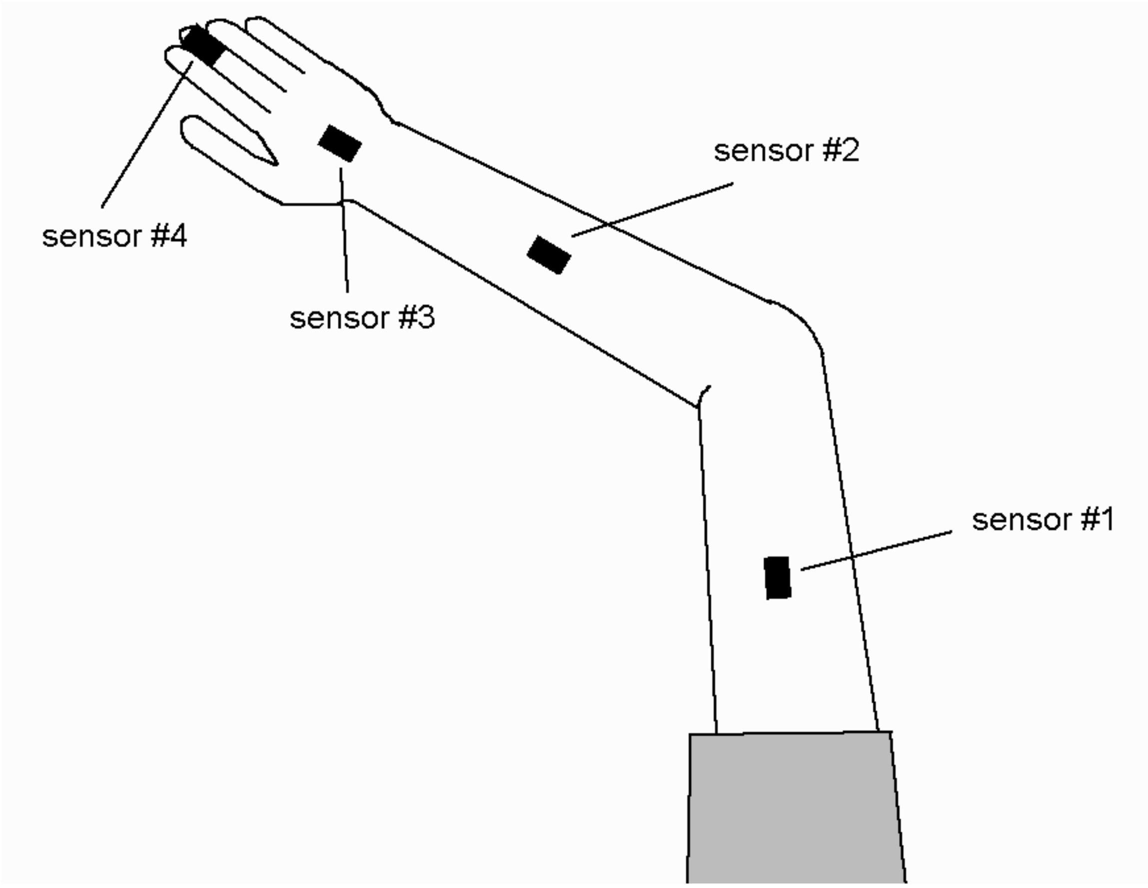

Figure 1

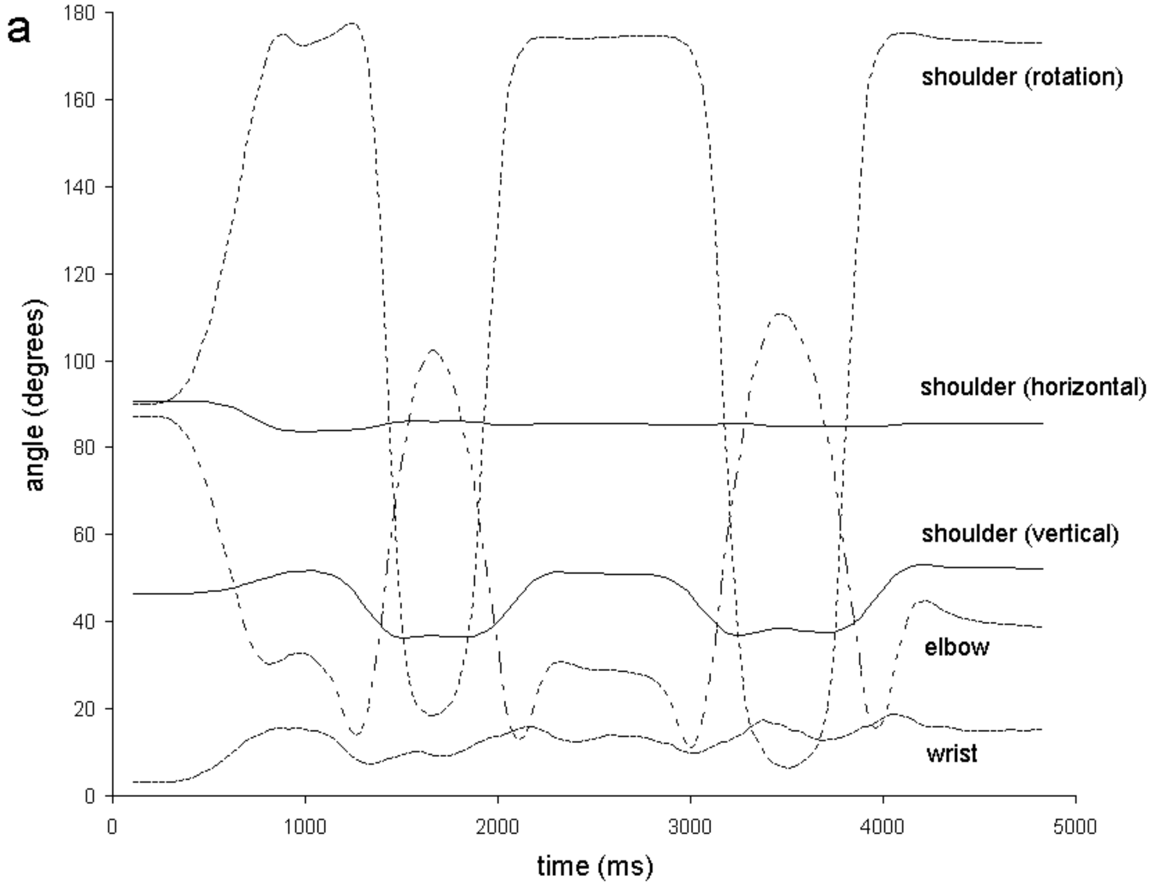
a
180
160
140
120
100
80
60
40
20
0
angle (degrees)
shoulder (rotation)
shoulder (horizontal)
shoulder (vertical)
elbow
wrist
0
1000
2000
3000
4000
5000
time (ms)

Figure 2a

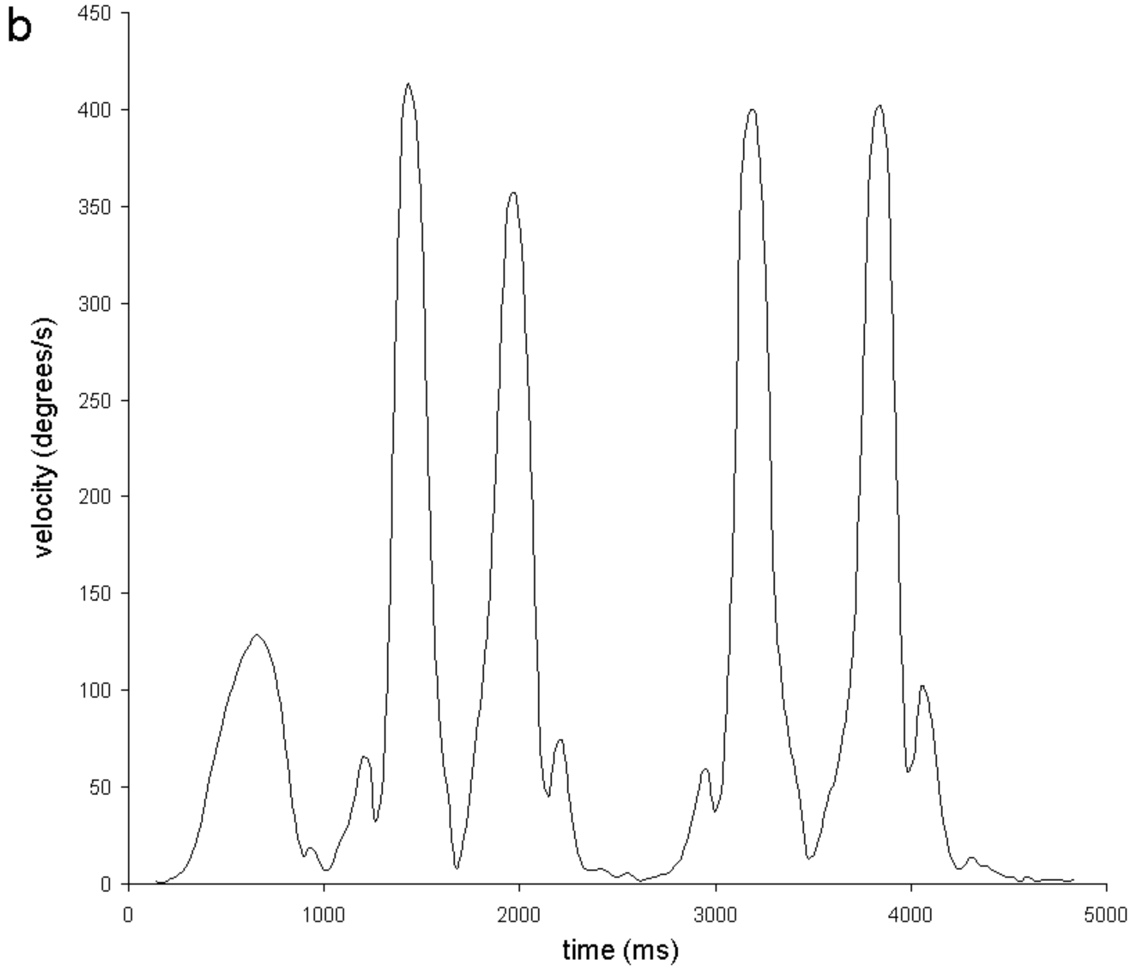
b
velocity (degrees/s)
450
400
350
300
250
200
150
100
50
0
0
1000
2000
3000
4000
5000
time (ms)

Figure 2b

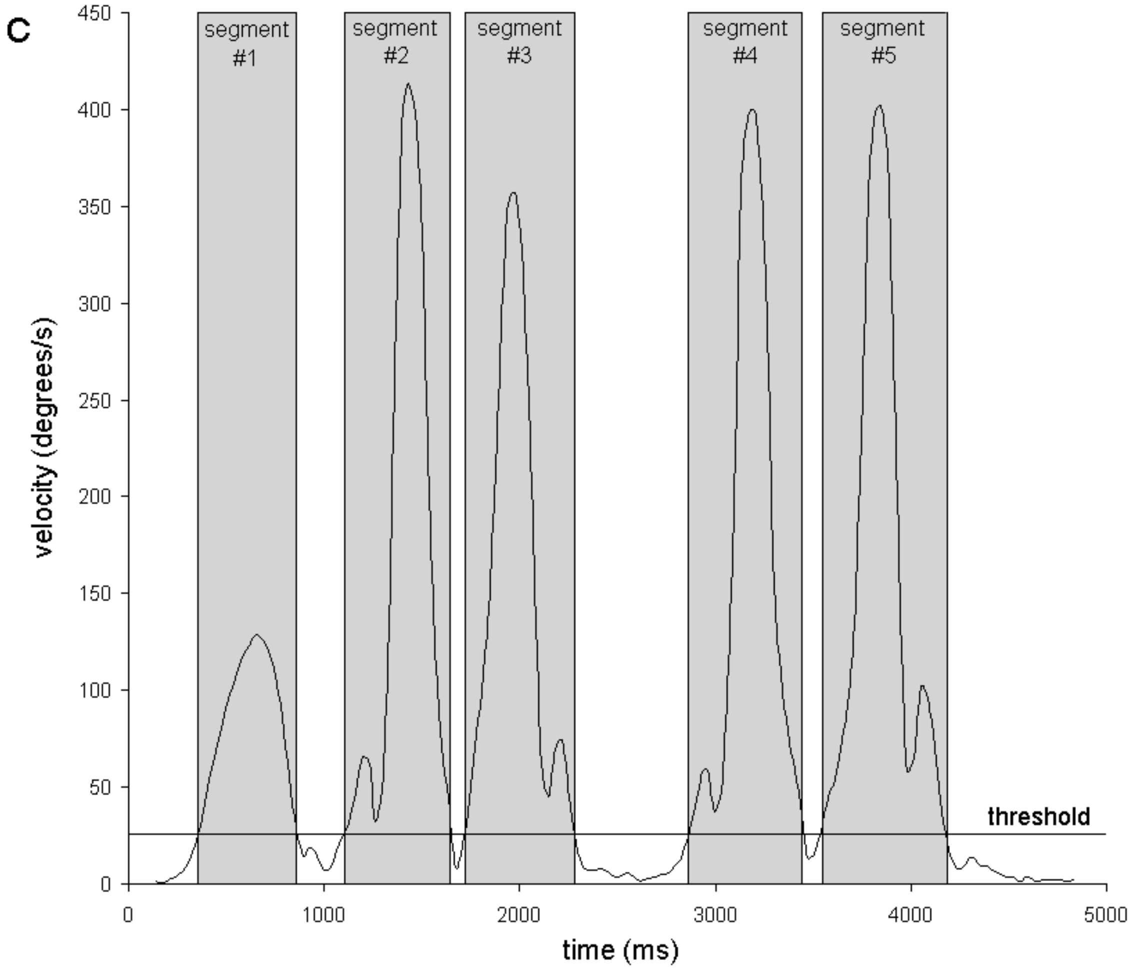
C
segment #1
segment #2
segment #3
segment #4
segment #5
threshold
velocity (degrees/s)
time (ms)
0
50
100
150
200
250
300
350
400
450
0
1000
2000
3000
4000
5000

Figure 2c

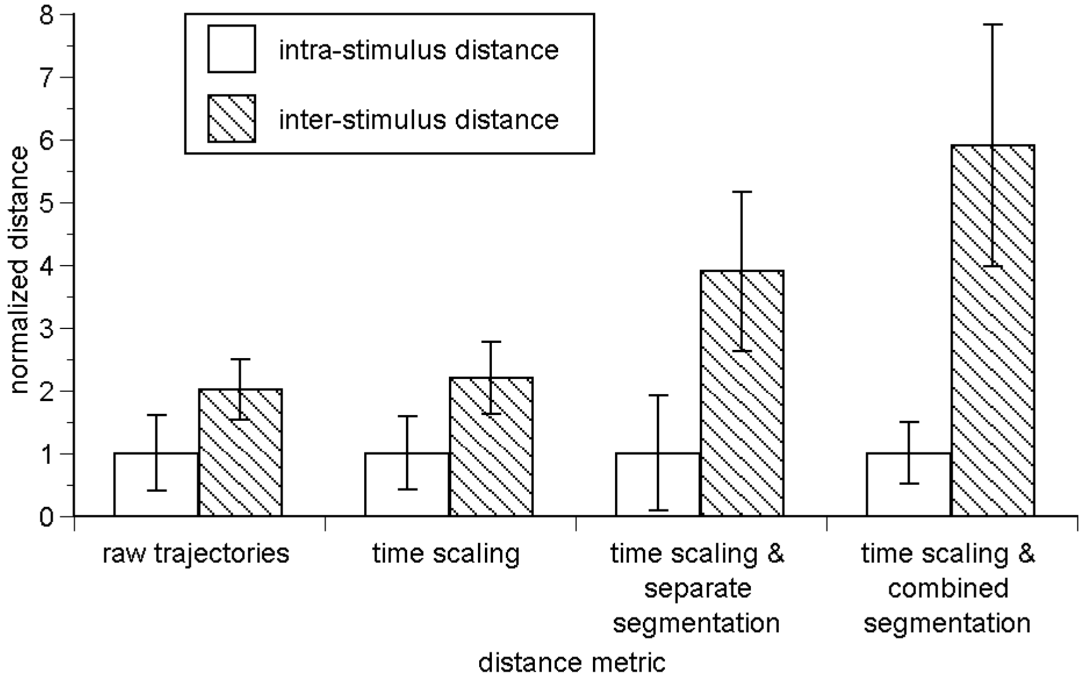
intra-stimulus distance
inter-stimulus distance
normalized distance
0
1
2
3
4
5
6
7
8
raw trajectories
time scaling
time scaling & separate segmentation
time scaling & combined segmentation
distance metric

Figure 3a

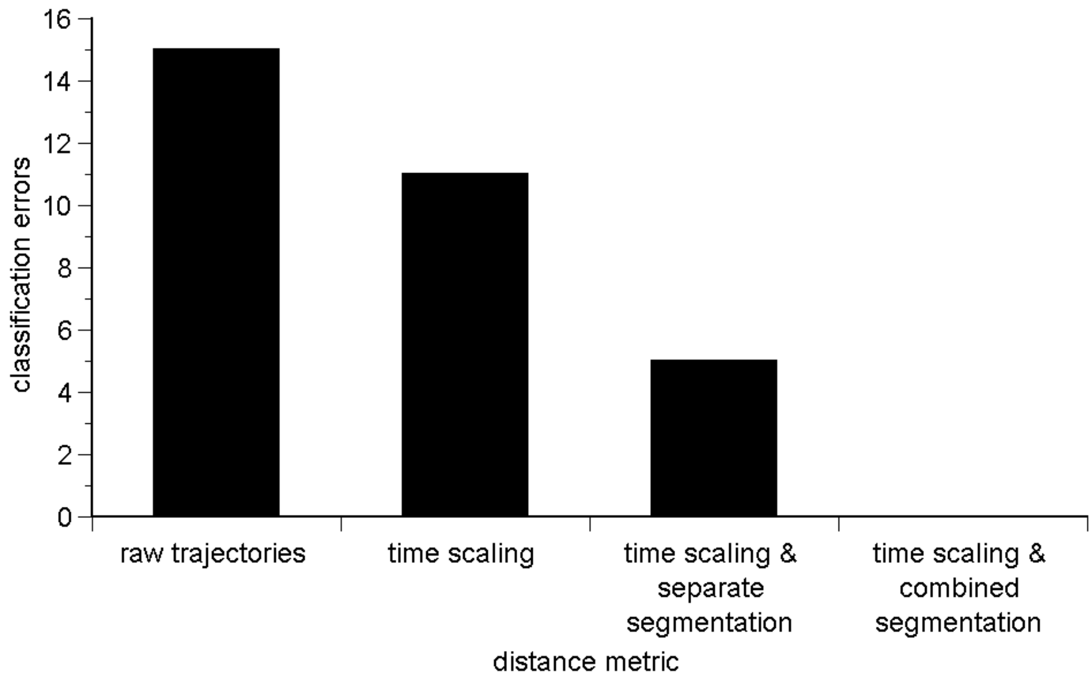
16
14
12
10
8
6
4
2
0
classification errors
raw trajectories
time scaling
time scaling & separate segmentation
time scaling & combined segmentation
distance metric

Figure 3b