

APPLICATION OF HIDDEN MARKOV MODEL TOPOLOGY ESTIMATION TO REPETITIVE LIFTING DATA

Raymond C. Vasko, Jr.

Amro El-Jaroudi

J. Robert Boston

Department of Electrical Engineering
University of Pittsburgh, Pittsburgh, PA 15261

ABSTRACT

At ICASSP '96, we presented an algorithm that estimates the topology of a hidden Markov model (HMM) given a set of time series data. The algorithm iteratively prunes state transitions from a large general HMM topology and selects a topology based on a likelihood criterion and a heuristic evaluation of complexity. In this paper, we apply the algorithm to estimate the dynamic structure of human body motion data from a repetitive lifting task. The estimated topology for low back pain patients was different from the topology for a control subject group. The body motions of patients tend not to change over the task, but the body motions of control subjects change systematically.

1. INTRODUCTION

We present the application of our HMM topology estimator [1] to describe the dynamic structure of correlated multivariate time series data from a physical process. Specifically, we characterize the dynamic structure of body motion data obtained during a repetitive, dynamic lifting task. The task was devised to objectively evaluate the functional capacity of patients with low back pain [2]. Preliminary study of the repetitive lifting task has suggested that systematic changes in body motion can occur over the course of the task and that the changes are different for control subjects than for patients.

Patients that undergo treatment for low back pain vary widely in their responses to the treatment. A goal of functional capacity testing is to explain this variability by identifying differences among patients before treatment that predict treatment effectiveness or by identifying differences after treatment that explain why treatment was or was not effective.

Identifying differences in lift sequences between groups of subjects is a difficult problem for which traditional statistical methods (e.g. linear discriminant analysis of time-averaged data, cluster analytic methods) are of limited utility [3,4]. By applying our HMM topology estimation algorithm, we can group subjects based on the dynamic behavior of their repetitive lifting data. We applied our topology estimation algorithm to functional capacity data obtained from control subjects and from patients before and after a 25-day treatment protocol [2]. One topology was estimated for lift sequences performed by control subjects, a second topology for patients before treatment, and a third topology for patients after treatment. In the HMM representation, the states represent different types of lifts. For each population, we used the set of possible paths through the estimated topology to identify subgroups in the population. That is,

each subject was assigned to a population subgroup based on the path (i.e. state sequence) used by that subject's lift sequence data through the population model.

2. ALGORITHM APPLICATION

In this application, three features were used to describe body motion during each lift in a lift sequence. These features are defined in terms of body angles which change as a subject performs a lift. Body angles are defined by the locations of the shoulder, hip, knee, and ankle. The hip angle is defined by the shoulder-hip-knee locations and the knee angle is defined by the hip-knee-ankle locations. The first feature *mindif* is the difference between the starting hip angle and the starting knee angle, where each angle is at its minimum value as the subject begins the lift. The second feature *todif* is the difference between the midpoints of the hip and knee angles, where the midpoint is defined as the time required for the angle to reach a value half-way between the angle's maximum and minimum values. The third feature *trdif* is the difference between the risetimes of the hip and knee angles, where the risetime is defined in terms of the time required for the angle to change from the minimum to the maximum value [3].

The control data set consisted of 62 lift sequences, one per subject, with a total of 4136 lifts. The pre-treatment patient data set consisted of 101 lift sequences (i.e. subjects) totaling 2588 lifts and the post-treatment patient data set contained 73 lift sequences totaling 2417 lifts. The three-dimensional feature space spanned by the union of the three data sets was vector quantized into 32 partitions. We used 20 iterations of the K-means algorithm (with $K=32$) to train the codebook of our full-search vector quantizer (VQ). The distortion measure for our VQ was the mean-square error. Initial code vectors for training the full-search codebook were obtained by a non-uniform binary-search VQ. For each partitioning ($K=2$) in the binary-search VQ, we chose initial code vectors by computing the centroids of two clusters separated by the hyperplane orthogonal to the dimension with the largest variance. Before vector quantization, each feature was normalized by its range.

For each data set, a topology estimate was obtained using the algorithm as described in [1] with the enhancements detailed below. None of the data sets were large enough to start the pruning algorithm with one state for each observation symbol (i.e. 32 states), so we were forced to choose an initial number of states which was much smaller than the number of observation symbols. We examined the feature space spanned by each data set and observed that the data are essentially a continuum within each population. This implies that there is significant overlap in the states of the actual physical (i.e. lift sequence) process. When the states




Figure 1. Plot of Normalized $Pr(O|\lambda)$ vs. Pruning Iteration for Control Group ("elbow" indicated by arrow)




Figure 2. 3-State HMM Topology Estimate for Control Group




Figure 3. 3-State HMM Topology Estimate for both Pre-Treatment and Post-Treatment Patient Groups

of a process overlap in the feature space, decreasing the number of states in the model of the process decreases the model likelihood. For this reason, since our algorithm evaluates topologies based on likelihood, the number of states in the estimated topology will generally be the same as in the starting topology if the actual process states overlap each other. We estimated topologies for each data set starting with 2, 3, 4, and 5 states. The training of the starting topology was initialized by a K-means clustering of the data set into 2, 3, 4, or 5 partitions. As described below, when we started the algorithm with 4 and 5 states, the population subgroups that we identified by state sequence were consistent with those we obtained by starting the algorithm with 3 states. For this paper, we present the results obtained starting with 3 states so that the subgroups identified by state sequence would be large enough to allow for a meaningful analysis of subgroup differences.

As described in [1], our algorithm selects a topology estimate from the set of topologies generated over the pruning iterations. The topology chosen is the simplest topology before a substantial decrease in $Pr(O|\lambda)$. That is, the topology estimate is located at an "elbow" in the $Pr(O|\lambda)$ trajectory. The "elbow" location is determined by heuristic evaluation. Figure 1 shows $Pr(O|\lambda)$ (normalized by data length) over the pruning iterations (starting with 3 states) for the control group. The likelihood of the starting topology is shown at pruning iteration 0 in the $Pr(O|\lambda)$ plot. The "elbow" in the $Pr(O|\lambda)$ trajectory in Figure 1 was determined to be at iteration 6. The $Pr(O|\lambda)$ plots for the



Figure 4. 4-State HMM Topology Estimate for Control Group




Figure 5. 5-State HMM Topology Estimate for Control Group




Figure 6. 4-State HMM Topology Estimate for both Pre-Treatment and Post-Treatment Patient Groups




Figure 7. 5-State HMM Topology Estimate for both Pre-Treatment and Post-Treatment Patient Groups

pre-treatment and post-treatment patient groups had well-defined "elbows" as in Figure 1 with the "elbow" being reached in 4 and 3 pruning iterations, respectively.

The 3-state topology estimate for the control group is shown in Figure 2. The estimated topology for the pre-treatment patient group is shown in Figure 3. A topology identical to Figure 3 was estimated for the post-treatment patient group, though the model parameters were different. The asterisks in Figures 2 and 3 indicate non-zero initial state probabilities.

The topology estimates for all three data sets indicate a temporal structure. All control subjects started in state 1 and some remained in that state throughout the task. The other subjects had one transition to one of the two other states. For both the pre-treatment and post-treatment patient groups, the estimated topologies included one isolated state and a 2-state temporal submodel. Some patients used only the isolated state, others used only the first state of the 2-state temporal submodel, and the rest used both states of the temporal submodel.

For each data set, when we started the pruning algorithm with 4 and 5 states, our topology estimates had 4 and 5 states, respectively. For all three data sets, the estimated topologies we obtained when we started the algorithm with 3 states were submodels of the topology estimates when we started with 4 or 5 states.

Figures 4 and 5 show the 4- and 5-state topologies for the control group. With few exceptions, the subjects who used the transition between states 1 and 0 in the 3-state topology did likewise in the 4- and 5-state models. Most of the lift sequences that remained in state 1 in the 3-state model did so in the 4- and 5-state models. Also, the subjects who used states 2, 3, or 4 in the 5-state model used states 2 or 3 in the 4-state model and state 2 in the 3-state model.

Figures 6 and 7 show the 4- and 5-state topologies for the pre-treatment patient group. As with the 3-state topology estimates, the 4- and 5-state topology estimates for pre-

Table 1. State Centroids for Control Model

		State 0	State 1	State 2
<i>mindif</i> (degrees)	Mean	-31.76	-30.87	-30.64
	SD	26.03	24.28	23.94
<i>todif</i> (msec.)	Mean	-2.01	9.79	11.03
	SD	58.84	60.43	60.26
<i>trdif</i> (msec.)	Mean	-19.86	-18.78	-20.17
	SD	50.72	44.84	44.34
Lifts Using State		933	1557	1646

Table 2. Pre-Treatment Patient State Centroids

		State 0	State 1	State 2
<i>mindif</i> (degrees)	Mean	14.67	19.74	14.68
	SD	24.06	21.94	25.43
<i>todif</i> (msec.)	Mean	3.47	0.36	-1.47
	SD	66.61	66.61	75.66
<i>trdif</i> (msec.)	Mean	3.79	4.65	3.77
	SD	47.56	50.05	52.20
Lifts Using State		878	602	1108

Table 3. Post-Treatment Patient State Centroids

		State 0	State 1	State 2
<i>mindif</i> (degrees)	Mean	18.13	12.03	12.01
	SD	26.03	20.01	23.97
<i>todif</i> (msec.)	Mean	7.52	2.86	7.33
	SD	46.24	33.48	47.44
<i>trdif</i> (msec.)	Mean	-11.93	4.37	-8.05
	SD	66.09	21.05	54.70
Lifts Using State		681	478	1258

treatment and post-treatment patients were the same. In general, the subjects that used the state transition in the 3-state model also did in the 4- and 5-state models (this was true for both patient groups). By assembling population subgroups based on state sequence, we were able to determine how the isolated states in the 4- and 5-state models were merged in the 3-state model.

Our estimates of dynamic structure can be interpreted by relating each HMM state to the feature space. To do this, we compared state centroids to distinguish the types of body motion represented by the states. Tables 1, 2 and 3 contain the state centroids for the 3-state control, pre-treatment patient and post-treatment patient models.

The *mindif* feature is a static parameter not directly related to the motion of the body during the lift. A large negative value for the *mindif* feature indicates that the starting knee angle is much larger than the starting hip angle, a posture corresponding to a "back lift" (i.e. at the start of the lifting motion, the subject is bending at the waist with straight legs). A positive value for the *mindif* feature corresponds to a "squatting" posture, where the subject maintains an upright torso and bends his or her legs to start the lift. If the value of the *mindif* feature is nearly zero, the subject's body is bent at both the waist and knee. From Tables 1 through 3, the control subjects perform "back lifts" and the patients perform "squat lifts". There is little difference between the *mindif* features of the pre-treatment and post-treatment patient groups.

The *todif* and *trdif* features are dynamic parameters that depend on body motion during the lift. To aid in the interpretation of the two dynamic parameters, we use the

Table 4. Coordination Indices

Model	State 0	State 1	State 2
Control	-0.285	-0.207	-0.090
Pre-Treatment Patient	-0.173	-0.093	-0.166
Post-Treatment Patient	-0.439	-0.100	-0.283

Table 5. Control Subject Subgroups

	State Sequence		
	1 $\Rightarrow$ 0	All 1	1 $\Rightarrow$ 2
Number of Subjects	17	17	27
Avg. Number of Lifts	68	69	68
Avg. Weight Lifted (lb.)	73	62	63

Table 6. Pre-Treatment Patient Subgroups

	State Sequence		
	All 0	2 $\Rightarrow$ 1	All 2
Number of Subjects	38	25	38
Avg. Number of Lifts	23	28	27
Avg. Weight Lifted (lb.)	30	31	29

Table 7. Post-Treatment Patient Subgroups

	State Sequence		
	All 0	2 $\Rightarrow$ 1	All 2
Number of Subjects	13	24	36
Avg. Number of Lifts	28	46	32
Avg. Weight Lifted (lb.)	38	29	37

todif and *trdif* features to compute a measure of the coordination of body motion [3]. Our coordination index is the correlation coefficient between the two dynamic features, computed across all lifts in a state. Table 4 lists the coordination indices for each state of the three lift sequence models (3-state). States with large negative coordination indices represent fluid body motion, while states with small negative coordination indices represent body motion that is more deliberate or guarded. Specifically, a large negative coordination index indicates a tendency for the subject's hip and knees to complete the lift together. Conversely, with a small negative coordination index, the subject's hip and knees tend not to complete the lift together [3]. From Table 4, pre-treatment patients use a lifting motion that is less coordinated than the controls. The dynamic parameters (and coordination indices) of post-treatment patients are more like the control group than the pre-treatment patient group.

Subjects within each group were arranged into subgroups based on the path they used through the model. Tables 5, 6 and 7 list the number of subjects using each path through the control, pre-treatment patient and post-treatment patient models. From Tables 5 through 7, it is clear that the percentage of control subjects that make a state transition (72.1%) is larger than for the patients pre-treatment (24.8%) or post-treatment (32.9%).

All control subjects start their lift sequences in state 1, which represents a moderately-coordinated lifting style. Controls that show a state transition either move to a state that is more coordinated (state 0) or less coordinated (state 2). The controls that transition to state 2 lift the most weight. Post-treatment patients either transition to a state with low coordination or remain in a state with rela-

tively high coordination. The post-treatment patients that use the less-coordinated state perform more lifts but lift less weight than the other post-treatment patients. These trends suggest that the amount of work performed by subjects depends on the lift types (i.e. states) used, although the work performance measures (number of lifts and weight lifted) were highly variable, and the differences between population subgroups were not statistically significant.

3. DISCUSSION

An obvious conclusion from the topology estimates is that control subjects were more likely to show a state transition than patients. In the 3-state model, 72.1% of control subjects made a state transition as opposed to 24.8% of the pre-treatment patients and 32.9% of the post-treatment patients. Comparing the 4- and 5-state models of the control and patient groups emphasizes the observation that controls use more state transitions than patients. This result is consistent with previous analyses that used a repeated-measures analysis of variance model to compare data obtained early in the task with data obtained in the middle of the task and at the end of the task [4]. This early-middle-late analysis showed that the body motions of controls changed during the task, while the body motions of patients tended to be constant.

The treatment received by the patients doesn't change the dynamic structure of their lifting data. After treatment, patients showed a limited improvement in work performance and a moderate change in the coordination of their lifting motion, but the treatment protocol does not seem to increase the patients' tendency to adapt (relax, become more efficient, etc.) to the task. It is worth noting that the patients only received 25 days of treatment; perhaps greater changes would be observed after a longer treatment protocol. It is also possible that our results have been influenced by the patients' willingness to perform the task.

In conclusion, by applying our HMM topology estimation algorithm we have demonstrated that there are differences in dynamic structure between controls and patients in repetitive lifting data. We have shown that there is a temporal structure in repetitive lifting data, and that our estimates of dynamic structure are consistent (similar topologies are obtained when we start the algorithm with different numbers of states) and well-defined (the "elbows" in the plots of $Pr(O|\lambda)$ are distinct). Lastly, we have demonstrated how a treatment protocol can be evaluated by arranging control, pre-treatment patient and post-treatment patient populations into subgroups using our estimated topologies. Our results suggest that this approach to time-series modeling may be helpful in characterizing functional capacity testing data.

4. IMPROVEMENTS TO THE ALGORITHM

We report several changes to the algorithm as described in [1]. One modification concerns the starting topology for the algorithm. The starting topology described in [1] has a state for each observation symbol. However, if there is a limited amount of data available and the data are finely-quantized (i.e. VQ into many observation symbols), then there may not be enough data to train a topology with one state for each observation symbol. In this case, we start the algorithm with a topology that has fewer states, with multiple observation symbols per state. This starting topology is ergodic with a fully-populated state transition

matrix as described in [1]. Before the starting topology is trained, the states are initialized based on a clustering scheme.

Another change to the algorithm involves the method for removing a state from the topology. The algorithm as described in [1] performed a separate state-removal iteration to eliminate a state from the topology after its self-transition was removed by a pruning iteration. In this approach, the topology resulting from the state-removal iteration was not evaluated by the pruning iteration. Thus, the $Pr(O|\lambda)$ for the state-removed topology might be smaller than some of the candidate topologies in the pruning iteration that removed the self-transition. For this reason, we changed the algorithm to perform state removal within the pruning iteration. A pruning iteration now trains a candidate topology for each state that eliminates the state entirely.

The algorithm presented in [1] eliminated a state from the topology by pruning all transitions into and out of that state and redistributing state transition probabilities to establish direct transitions between states that were previously connected via the eliminated state. This approach results in a suboptimal initialization of the HMM parameters for training the state-eliminated topology. We have modified the algorithm to eliminate states by merging pairs of states rather than pruning single states. The set of candidate topologies evaluated by a pruning iteration now includes (in addition to topologies that remove single state transitions) topologies that merge all possible pairs of states. The primary advantage to merging states as opposed to pruning states is that the counts (e.g. number of times in state, number of times a transition is used, etc.) can be used to initialize the HMM parameters for training the state-eliminated topology.

Also, the algorithm has been modified to train each topology to the convergence of $Pr(O|\lambda)$. Instead of training by a fixed number (ten) of iterations of Baum-Welch (B-W) reestimation [1], each topology is trained until $Pr(O|\lambda_t) < 1.5Pr(O|\lambda_{t-1})$, where λ_t is the HMM after t iterations of B-W reestimation. The rate at which $Pr(O|\lambda)$ converges can vary considerably over different topologies. That is, even though topology λ_1 may have a larger $Pr(O|\lambda)$ than topology λ_2 when both have been trained to convergence, $Pr(O|\lambda_2)$ may be larger than $Pr(O|\lambda_1)$ after each has been trained by ten B-W iterations. We implemented a convergence criterion to ensure that when the algorithm evaluates candidate topologies during a pruning iteration, topology selection is not affected by the convergence rate of $Pr(O|\lambda)$. Our convergence criterion was chosen to achieve a balance between algorithm performance and computational cost.

5. REFERENCES

- [1] Vasko, R.C., El-Jaroudi, A., Boston, J.R., "An Algorithm To Determine Hidden Markov Model Topology," *IEEE Proc. ICASSP'96*, Vol. 6, pp. 3578-3582, 1996.
- [2] Rudy, T.E., Lieber, S., Turk, D.C., "Development of a Functional Capacity Protocol for Chronic Back Pain Patients", *Clin J Pain*, vol. 7, pp. 62, 1991.
- [3] Boston, J.R., Rudy, T.E., Mercer, S.R., Kubinski, J.A., "A Measure of Body Movement Coordination During Repetitive Lifting", *IEEE Trans Rehab Eng*, vol. 1, pp. 137-144, 1993.
- [4] Rudy, T.E., Boston, J.R., Lieber, S.J., Kubinski, J.A., "Body Motion Patterns During a Novel Wheel-Rotation Task", *Spine*, vol. 20, pp. 2547-2554, 1995.