Posture estimation for a canine machine interface based training system

ARTICLE *in* CONFERENCE PROCEEDINGS: ... ANNUAL INTERNATIONAL CONFERENCE OF THE IEEE ENGINEERING IN MEDICINE AND BIOLOGY SOCIETY. IEEE ENGINEERING IN MEDICINE AND BIOLOGY SOCIETY. CONFERENCE · AUGUST 2012

DOI: 10.1109/EMBC.2012.6346964 · Source: PubMed

4 AUTHORS, INCLUDING:

Rita Brugarolas
North Carolina State University

David Roberts
North Carolina State University

Barbara L Sherman
North Carolina State University

Behavior Recognition Based on Machine Learning Algorithms for a Wireless Canine Machine Interface

Rita Brugarolas[1], Robert T. Loftin[2], Pu Yang[2], David L. Roberts[2], Barbara Sherman[3] and Alper Bozkurt[1]
[1]Electrical and Computer Engineering, [2]Computer Science, [3]College of Veterinary Medicine
North Carolina State University, Raleigh, NC 27697-7911, USA

***Abstract*—Training and handling working dogs is a costly process and requires specialized skills and techniques. Less subjective and lower-cost training techniques would not only improve our partnership with these dogs but also enable us to benefit from their skills more efficiently. To facilitate this, we are developing a canine body-area-network (cBAN) to combine sensing technologies and computational modeling to provide handlers with a more accurate interpretation for dog training. As the first step of this, we used inertial measurement units (IMU) to remotely detect the behavioral activity of canines. Decision tree classifiers and Hidden Markov Models were used to detect static postures (sitting, standing, lying down, standing on two legs and eating off the ground) and dynamic activities (walking, climbing stairs and walking down a ramp) based on the heuristic features of the accelerometer and gyroscope data provided by the wireless sensing system deployed on a canine vest. Data was collected from 6 Labrador Retrievers and a Kai Ken. The analysis of IMU location and orientation helped to achieve high classification accuracies for static and dynamic activity recognition.**

Keywords—cascade learning; inertial measurement units; animal machine interfaces; body area network; canine training

I. Introduction

Day by day, more people benefit from the service and companionship of dogs. Search and rescue, guiding the blind, drug and explosive detection, and protection of people or property are some examples of the important roles dogs have in our society. Training military or service dogs costs several tens of thousands of dollars [1], requiring the extensive efforts of specialized personnel during long periods of time, and does not always achieve the desired results. Using electronic devices to provide remote commands has already become a common practice in the training and handling of dogs [2,3]. We believe that using electronic sensors and computer models of dog behaviors will provide more robust and standardized training, allowing more people to benefit from these specially-trained dogs.

In this paper, we present our latest efforts towards developing a canine body-area-network (cBAN) [4,5]. Our cBAN is a low-power wireless bioelectric system that is comprised of miniaturized wearable sensor and actuator packages, as well as algorithmic support for interpretation and feedback using those sensors and actuators. The three main components of the cBAN are a smartphone or other wireless communication device carried by the handler, a wireless sensor platform worn by the dog and a remote computational node which uses physiological and behavioral models of the dogs to process the sensor data. The current canine wearable platform includes accelerometers and gyroscopes for activity and posture estimation, with plans to add physiological sensors in the future.

The cBAN will ultimately offer a bidirectional closed feedback loop between the handler, the computational node, and the animal. The computational node will retrieve data continuously from the sensors on the dog and perform processing to provide accurate information about the canine's behavior to the handler. The handler will be equipped with a handheld wireless device to receive the analyzed data to help with a more accurate interpretation of the canine's behaviors and also to communicate wirelessly with the canine.

The capability of analyzing and quantifying daily activity levels and cycles can be used as an indicator for health assessment [6]. In our previous work, we achieved successful recognition of static dog postures (sitting, standing, lying down, standing on two legs and eating off the ground). In this study, we improve the classification algorithms to recognize dynamic activities such as walking, climbing up stairs and walking down a ramp as well. For this, we use Hidden Markov Models (HMMs) to account for the cyclic and temporal structure of these types of activities. Here, we also assess the effectiveness of using accelerometer and gyroscope data for the classification.

The contributions of this study can be summarized as follows:

- Classification of static postures (sitting, standing, lying down, standing on two legs and eating off the ground) with a decision tree classifier.
- Detection of dynamic activities (walking, climbing stairs and walking down a ramp) with various time durations by using HMMs.
- Analysis of accelerometer and gyroscope effectiveness for each classification.
- Assessment of the ssignificance of sensor location and orientation for static posture classification.
- Evaluation of the generalization of learned models between dogs.

II. Wearable Wireless Sensing System

The wireless sensing platform was an upgraded version of our earlier prototype [4]. Each node included a three-axis accelerometer (CMA3000) and a three-axis gyroscope (L3G4200D) connected to a CC2540 Texas Instruments microcontroller with an embedded high performance Bluetooth low energy transceiver. The platform was updated to support 4 sensor nodes and MATLAB was used at the base station for data collection.

III. Data Collection Protocol

For this study, we had access to five Labrador Retrievers in the school of Veterinary Medicine at North Carolina State University (NCSU). These dogs were from a cohort being trained and tested as military working dogs to detect improvised explosive devices (IEDs) [7]. All animal procedures were consistent with NIH and USDA guidelines and were approved by NCSU the Institutional Animal Care and Use Committee (IACUC).

Sensors were attached to the dogs' harnesses with Velcro in four different sensor sites: on the rump, on the chest, on the abdomen, and on the back (Fig. 1). Below, we refer to the sensor node on the rump with the x-axis of the accelerometer pointing towards the tail as 'Node1'; 'Node2' is the sensor node on the chest with the x-axis of the accelerometer pointing towards the face; 'Node3' is the sensor node located on the abdomen with the x-axis of the accelerometer pointing towards the tail; and finally 'Node4' is on the back closer to the head with the x-axis of the accelerometer pointing towards the head.

For the data collection, the dog trainer commanded the canines to perform five repetitions of each behavior (sitting, standing, lying down, eating off ground and standing on two legs) for approximately 4 s. The dogs returned to a standing position between repetitions. Their performance was video recorded for offline data processing. In this study, we also had access to a flight of stairs and ramps at the School of Veterinary Medicine training facility (Fig. 1). This enabled us to collect data on dynamic activities.

The dogs were led by a handler through the following sequence three times each: walk up the stairs, walk across a platform to the ramp, walk down the ramp, and walk back to the starting position.

The sampling rate for each of the sensors was 10 Hz and each of the behaviors was performed for a total time approximately 20 s resulting in approximately 200 instances of each behavior. Table I shows the distribution of instances available for the analysis. For Vet_dog2 and Vet_dog3 data were available for the four sensor nodes whereas in other cases only 'Node1' and 'Node2' recordings were available. The last two columns of the table (indicated as Dog1 and Dog2) correspond to recordings from additional animals for comparisons. These were two privately owned pets, a Labrador Retriever and a Kai Ken.

IV. Machile Learning Algorithms

Previously, we have shown that heuristic features from accelerometer output could be successfully used for static pose recognition [4, 5], and in this study we use such features for dynamic activity recognition. Accelerometers sensed static and dynamic acceleration. The static acceleration corresponds to the projection of gravity over the axes and the dynamic acceleration is associated with the vibration and actual motion of the sensor. During the performance of a static behavior, the accelerometer output shows the tilt angle of the sensor with respect to gravity, and because there is minimal motion, the dynamic component is very small. For dynamic activities, combinations of both components were present. Fig. 2 shows a representative sample of acceleration data from the x-axis of the accelerometer in each of the four locations.

Fig. 3 shows the classification flow we developed for this work. Maximum likelihood estimation with HMMs was used in the initial stage to identify each of the dynamic activities. HMM input was a combination of dynamic and postural recordings[8]. Data not classified as a dynamic activity was processed with a moving average filter, and fed to a decision tree cascade classifier for the recognition of the transitions and the specific static poses. In this study behavioral classification was done offline. For the trials where 'Node3' and 'Node4' were available, those recordings were not used in the HMM implementation but they were used in the posture classification.

TABLE I. Instances of Each Behavior Available for the Analysis

Behavior	**Set of instances per dog trial**								
	Vet_dog2	Vet_dog3	Vet_dog4	Vet_dog5	Vet_dog6	Dog1_Trial1	Dog1_Trial2	Dog2_Trial1	Dog2_Trial2
Sitting	542	231	366	815	204	72	476	375	400
Standing	550	449	359	1155	368	141	280	70	221
Eating off ground	244	268	51	106	188	152	148	100	123
Standing two legs	215	245	204	180	210	299	132	286	386
Lying down	0	0	0	74	235	0	0	294	292
Walk	718	432	370	0	0	0	0	0	0
Climb stairs	102	72	94	0	0	0	0	0	0
Down ramp	75	54	92	0	0	0	0	0	0
Transition	82	185	181	272	188	108	128	211	225
Total	2528	1936	1717	2602	1393	772	1164	1336	1647

Fig. 1. Diagram showing the sensor locations. The platform used to measure dynamic activities [7].

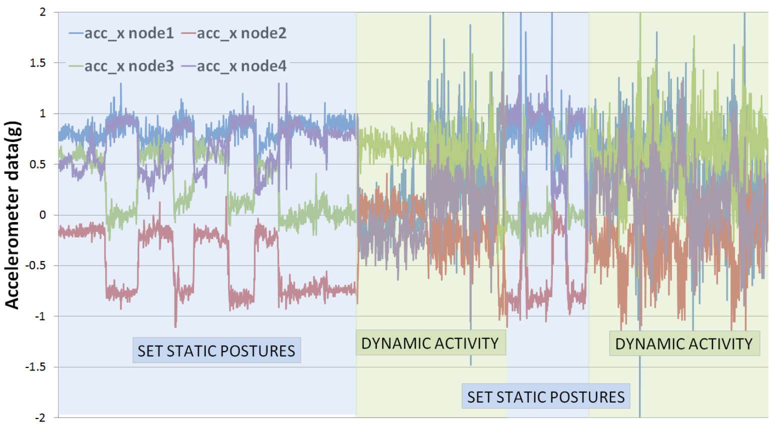

Fig. 2. Representative accelometer data of the x-axis in the four sensor location showing static and dynamic activities.

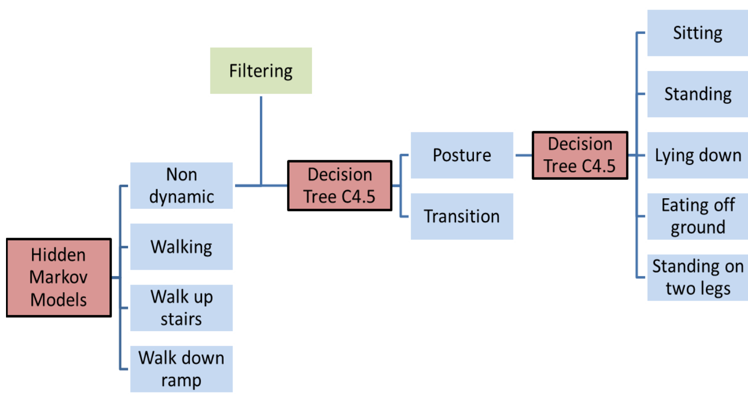

Fig. 3. Developed classification flow for dynamic and static activities.

A. *Dynamic activity recognition*

HMMs were used to represent and identify dynamic behaviors, as the classification methods used for static poses such as sitting do not explicitly account for the temporal structure of these behaviors.

Our HMMs consisted of a set of possible states that the dog could be in, along with the probabilities of transitioning from one state to the next at each time step and the probabilities of the model starting in each of the states. Each state generated an observation according to its observation distribution. In this case, the continuous sensor vectors were assumed to be drawn from a multivariate normal distribution with zero covariance, such that each feature was independent given the state.

To estimate the parameters of an HMM for a given set of sequences, the iterative Baum-Welch [9] algorithm was used to find parameters which maximized the probability of generating those sequences. As the Baum-Welch algorithm requires the number of states within the model to be known, multiple models were learned with between three and ten states, and the best of those models was selected. The HMM algorithms were implemented in the Java programming language.

An HMM was learned for each behavior, as well as a single model for all sequences in which none of the three dynamic behaviors were being performed (Fig. 3). The behavior occurring in an unlabeled sequence could then be found via *maximum likelihood* estimation, where the behavior was identified as the one for which the associated HMM has the highest probability of generating that sequence. This probability was computed using the Viterbi algorithm [8].

The signals from the sensors were split into ten step segments, using a ten step-sliding window such that the segments overlap. By also learning a model of behaviors that were not dynamic, it was possible to filter out those behaviors and pass them to the static classifiers for identification. The effectiveness of this learning approach was evaluated via five-fold cross validation, with the test sets consisting of thirty percent of the data.

TABLE II. HMM CLASSIFICATION ACCURACY

% correct	**Classification output**			
	Walking	***Up stairs***	***Down ramp***	***Non-dynamic***
Vet_dog2	100%	100%	100%	99.6%
Vet_dog3	100%	100%	92.0%	97.8%
Vet_dog4	100%	100%	100%	92.0%

TABLE III. HMM ACCURACY SENSITIVY TO ALGORITHM TRANSFER FROM A GROUP OF DOGS TO ANOTHER

% correct	**Classification output**		
	Walking	***Up stairs***	***Down ramp***
Testing against Vet_dog2	97.9%	46.3%	29.6%
Testing against Vet_dog3	66.8%	7.7%	100%
Testing against Vet_dog4	97.7%	0%	0%

TABLE IV. HMMS CLASSIFICATION ACCURARY AGAINST GYROSCOPE

% correct	**Classification output**			
	Walking	***Up stairs***	***Down ramp***	***Non-dynamic***
Vet_dog2	100%	100%	100%	97.9%
Vet_dog3	97.2%	100%	92.0%	98.9%
Vet_dog4	91.6%	100%	96%	93.3%

TABLE V. HMMS CLASSIFICATION ACCURACY AGAINST ACCELEROMETER

% correct	**Classification output**			
	Walking	***Up stairs***	***Down ramp***	***Non-dynamic***
Vet_dog2	100%	100%	100%	100%
Vet_dog3	100%	100%	100%	91.5%
Vet_dog4	99.2%	100%	98%	83%

The results in Table II show that, when learning to identify behaviors for a single dog, the HMM approach achieves high accuracy, and rarely misidentifies a sequence as being generated by another behavior. The results in Table III, however, show that this approach was ineffective when trying to apply models learned for one group of dogs to another. This is expected given the small size of the dataset.

Table IV shows the accuracy achieved by using only gyroscope data for the HMM and Table V using only accelerometer data. Accuracy results obtained with the gyroscope data were slightly better than with the accelerometer data, but neither of these achieved the accuracy of using both of the sensor data. These results support our intuition that the angular rate read by the gyroscope has an important role in recognition of cyclic activities [4].

B. Static activity recognition

For static posture classification, a two-level cascade classifier [10] was used. The C4.5 algorithm implemented in the WEKA Machine learning toolkit developed at the University of Waikato [11] and 10-fold cross-validation were used for the study. This decision tree algorithm was used in the two levels, the first to distinguish between transitions and postures, and the second to classify the specific postures (sitting, lying, standing, eating off the ground and standing on two legs) from the instances classified as postures in the first level. C4.5 builds a classifier with a tree structure from the instances in the training set. The tree leaves represent class labels and branches represent conjunctions of features that lead to those class labels. The C4.5 algorithm uses "information gain" [12] as the splitting criterion for splitting the branch. At each splitting, the decision tree algorithm chooses the feature value providing the maximum reduction in uncertainty about the class labels. So, the feature at the root of the tree is the one with the maximum information gain, and is therefore the best predictor. The feature used at the second level of the tree is the next best predictor given the value of the first [13].

Table VI shows the average accuracy of the two cascade decision tree classifiers. For Dog1 and Dog2 the average of the accuracy in the two available trials was presented. The first column shows the accuracy achieved by using only the accelerometer data as feature and the second column using only the gyroscope data. From these results, we conclude that the gyroscope data were not required for accurate pose estimation

TABLE VI. 2-LEVEL DECISION TREE CLASSIFIER (C4.5) EFFICIENCY WHEN USING ACCELEROMETER OR GYROSCOPE DATA

Dog	Sensor data used	
	Accelerometer	*Gyroscope*
Vet_dog2	99.5%	72.8%
Vet_dog3	98.9%	63.3%
Vet_dog4	94.4%	61.8%
Vet_dog5	77.8%	63.4%
Vet_dog6	97.5%	74.5%
Dog1	99.3%	62.2%
Dog2	99.4%	51.5%

TABLE VII. LOCATION AND ORIENTATION OF SENSOR ASSOCIATED TO BEST PREDICTION RESULTS FOR 2-LEVEL DECISION TREE CLASSIFIER (C4.5)

Dog	Decision Tree classifier (C4.5)	
	The best predictor	*The secondary best predictor*
Vet_dog2	node2_x	node1_x
Vet_dog3	node3_x	node1_x
Vet_dog4	node1_x	node2_x
Vet_dog5	node1_x	node2_x
Vet_dog6	node2_x	node1_x
Dog1	node1_x	node2_x
Dog2	node1_z	node1_x

Table VII shows the best and second best predictors for the C4.5 algorithm. Results show that, in general, nodes 1 and 2 in the x direction offered the best prediction results. Though there were a few exceptions, these results were basically consistent with our initial studies showing those as the optimal sensor locations [5].

V. CONCLUSIONS

In this study, we have used IMUs on 7 dogs to obtain 9 trials of data with an aim of estimating canine postures electronically. We have shown that a high level of accuracy in activity recognition can be achieved when building models for each individual dog, both for static postures and for dynamic activities like walking, climbing the stairs, and walking down a ramp. We observed that applying models learned for one group of dogs to another to be far less effective.

From the HMM algorithms used for dynamic activity recognition, we obtained better accuracy when using only gyroscope data relative to the accuracy when only using accelerometer data supporting our intuitions that gyroscope data would play an important role on cyclic activity recognition, but the combination of both sensor data provided the highest classification accuracy.

From the static pose classification, we conclude that the gyroscope data did not improve pose estimation accuracy, and we observe that nodes 1 and 2 in the x direction were, in general, the best predictors, consistent with our initial studies showing those as the optimal sensor locations [4, 5].

The implementation and computational cost of the HMM classifier were large in comparison to the decision tree classifiers for static postures, but once HMM models were learned, classifying new examples using these models were fairly computationally efficient when the Viterbi algorithm was used.

Extension of HMMs to be able to recognize different levels of intensity in activities, such as walking versus running or galloping, is an ongoing work.

ACKNOWLEDGMENT

We would like to thank Lucia Lazarowski, Beth Case, Diana Simpson at the School of Veterinary Medicine at NCSU for their help during the data collection. R.B. also thanks Fundación Caja Madrid for the financial support.

REFERENCES

[1] T. N. Y. Times, “The Dogs of War: Beloved Comrades in afghanistan,” *May 11*, 2011.

[2] “Sit means sit. Dog training.” [Online]. Available: www.sitmeanssit.com/.

[3] “Dog Schock Collars.” [Online]. Available: www.americas-pet-store.com/shock-collar.html.

[4] R. Brugarolas, D. Roberts, B. Sherman, and A. Bozkurt, “Machine Learning Based Posture Estimation for a Wireless Canine Machine Interface,” in *IEEE Topical Conference on Biomedical Wireless Technologies, Networks, and Sensing Systems (BioWireless)*, 2012, pp. 3–5.

[5] R. Brugarolas, D. L. Roberts, B. Sherman, and A. Bozkurt, “Posture Estimation for a Canine Machine Interface Based Training System,” in *Proceedings of The Engineering in Medicine and Biology Conference*, 2012.

[6] C. J. Riddoch, C. Mattocks, K. Deere, J. Saunders, J. Kirkby, K. Tilling, S. D. Leary, S. N. Blair, and A. R. Ness, “Objective measurement of levels and patterns of physical activity.,” *Archives of disease in childhood*, vol. 92, no. 11, pp. 963–9, Nov. 2007.

[7] S. Weinberger, “IDE Detector Dog Research Aims to Save Lives of U.S. Combat Troops -- Human and K-9,” *August 31*, 2012.

[8] A. Viterbi, “Error bounds for convolutional codes and an asymptotically optimum decoding algorithm,” *Information Theory, IEEE Transactions on*, vol. 13, no. 2, pp. 260–269, 1967.

[9] L. E. Baum, T. Petrie, G. Soules, and N. Weiss, “A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains,” *The annals of mathematical statistics*, pp. 164–171, 1970.

[10] P. Viola and M. Jones, “Robust Real-time Object Detection,” in *Second International Workshop on Statistical and Computational Theories of Vision*, 2001, pp. 1–25.

[11] I. H. W. Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, *The WEKA Data Mining Software: An Update*. SIGKDD Explorations, Volume 11, Issue 1, 2009.

[12] T. M. Mitchell, *Machine Learning*. The Mc-Graw-Hill Companies, Inc., 1997.

[13] J. R. Quinlan, “Induction of Decision Trees,” *Mach. Learn.*, vol. 1, no. 1, pp. 81–106, Mar. 1986.