2013 12th International Conference on Machine Learning and Applications

Applying Machine Learning and Audio Analysis Techniques to Insect Recognition in Intelligent Traps

Diego F. Silva, Vinícius M. A. de Souza, Gustavo E. A. P. A. Batista
ICMC/USP – São Carlos, Brazil
{diegofsilva,vsouza,gbatista}@icmc.usp.br

Eamonn Keogh
University of California
Riverside, USA
eamonn@cs.ucr.edu

Daniel P. W. Ellis
Columbia University
New York, USA
dpwe@ee.columbia.edu

***Abstract*—Throughout the history, insects have had an intimate relationship with humanity, both positive and negative. Insects are vectors of diseases that kill millions of people every year and, at the same time, insects pollinate most of the world's food production. Consequently, there is a demand for new devices able to control the populations of harmful insects while having a minimal impact on beneficial insects. In this paper, we present an intelligent trap that uses a laser sensor to selectively classify and catch insects. We perform an extensive evaluation of different feature sets from audio analysis and machine learning algorithms to construct accurate classifiers for the insect classification task. Support Vector Machines achieved the best results with a MFCC feature set, which consists of coefficients from frequencies scaled according to the human auditory system. We evaluate our classifiers in multiclass and binary class settings, and show that a binary class classifier that recognizes the mosquito species achieved almost perfect accuracy, assuring the applicability of the proposed intelligent trap.**

***Keywords*-insect classification; sensor data; feature extraction; intelligent trap;**

I. Introduction

Throughout the whole of human history, insects have been intimately related to our wellbeing, in both positive and negative ways. For example, insect pests consume and destroy around US$40 billion worth of food each year [1]. In contrast, insects pollinate at least two-thirds of all the food consumed in the world, with bees alone responsible for pollinating one-third of this total [2]. Furthermore, many species have been used as bioindicators of environmental quality, since their presence/absence, distribution and density, define the quality of the ecosystem, especially in relation to contaminants in the air, soil and water [3].

Another example of the relationship between insects and humans is the fact that insects are vectors of diseases that kill millions of people every year and leave tens of millions sickened. It is estimated that dengue, a disease transmitted by mosquitoes of the genus *Aedes*, affects between 50 and 100 million people every year and it is considered endemic in more than 100 countries [4]. Malaria, transmitted by mosquitoes of the genus *Anopheles*, affects around 6% of the world's population and it is estimated that there are over 200 million cases per year and about 7 million lethal cases in the last decade [5].

Due to the lack of vaccines and specific and effective medicines for the treatment of several diseases transmitted by insects, a common practice is to fight the root of the problem by using insecticides to control vector insects. In agriculture, pesticides are also widely used to control pest insects. However, over 98% of sprayed insecticides and 95% of herbicides reach a destination other than their target species, including non-target species, air, water and soil, reducing biodiversity, contributing to pollinator decline and threatening endangered species [6]. Furthermore, reports of malaria, dengue and filariasis have shown that the vectors' resistance to insecticides has steadily increased since insecticides were first introduced [7], [8].

Thus, we propose a novel approach to control insect populations in agriculture and the environment that potentially can have an enormous impact in food production and public health. It consists of an intelligent trap, a simple low-cost device able to identify the insect species in real-time and trap target insect species. Our approach has minimal impact on the environment, including other insect species such as pollinators and food sources for other animals.

The intelligent trap will have a significant impact on the insect control for two major reasons:

1) It will identify and selectively trap target species, allowing other non-harmful insects to be released back into the environment. Therefore, the trap will have no relevant impact on important insects such as pollinators and on the environment in general;
2) It will provide population estimates of the target insects in the trap area. Insect population estimates are important in understanding the spatial-temporal distribution of the insects, and serve as an alarm system for outbreaks of insect pests and disease vectors, or as bioindicators of environmental quality.

There are several scientific and technological challenges involved in developing such an intelligent trap, but the most challenging is signal classification. In order to selectively capture insect species of interest, we need an accurate classifier able to correctly associate a species label to a

978-0-7695-5144-9/13 $31.00 © 2013 IEEE
DOI 10.1109/ICMLA.2013.24

IEEE computer society

very short signal fragment collected by the sensor. In this paper, we describe the signal feature extraction as well as the machine learning techniques used to obtain such classifiers.

This paper is organized as follows. Section II starts with a description of our application. Section III presents the data collection procedure as well as the preprocessing steps for the data used in the experimental evaluation. We continue with a description of the features extracted from the data in Section IV; and of the experimental setup and a discussion about the results obtained in our experiments in Section V. Finally, Section VI concludes the paper and presents directions for future work.

II. Intelligent Insect Trap

The proposed trap is an extremely simple device. It consists of two actuators and one sensor. The actuators are an inexpensive open-close door responsible for trapping or releasing the insects, and a fan with enough thrust to drive the insect through the trap until it is released or retained in a chamber. The chamber is coated with a fine net (to allow airflow) and uses a sheet of sticky paper to definitely trap captured insects. Figure 1 shows the logical design of the our prototype trap.

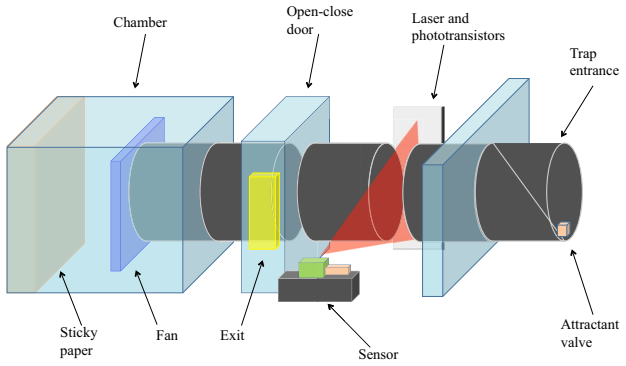


Figure 1. Logical design of the intelligent trap with a sensor that classifies the insects that enter the trap

The trap operates by spraying an attractant in the air. The attractant is used to bait the target insects to the trap entrance. For example, carbon dioxide is an attractant for many mosquito species [9]. When an insect flies in front of the trap entrance, it is gently pushed by the airflow in direction of the laser sensor. As the insect passes across the laser, it is classified in real-time by the sensor. A decision is made whether to trap or release the insect. If the insect is released, an exit door is opened and the airflow is reversed to provide a short burst. The insect is then dispensed away from the trap entrance. If the insect is trapped, it is pushed by the airflow into the chamber and retained by the sticky paper.

The sticky paper sheets can be replaced and used to estimate the density of the insect population. This procedure is similar to the one used with regular sticky traps. However, regular sticky traps capture every insect that lands or walks on its surface and, consequently, can be very expensive in terms of time and human resources, since experts manually inspect these sheets.

The laser sensor used in our intelligent insect trap is shown in Figure 2. It consists of a low-powered planar laser source pointed at an array of phototransistors. When a flying insect crosses the laser, its wings partially occlude the light, causing small light variations captured by the phototransistors. An electronic circuit board filters and amplifies the signal and the output is recorded by a digital recorder.

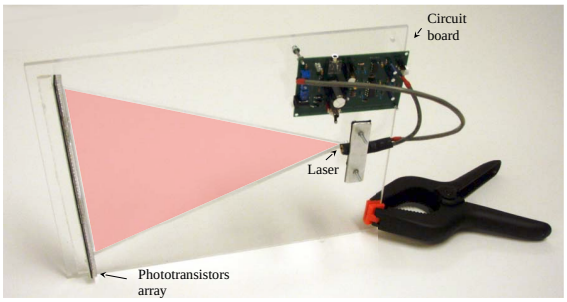


Figure 2. The laser sensor used to record insect data

The sensor signal is very similar to an audio signal captured by a microphone, even though the data are obtained optically. However, the sensor is totally deaf to any agent that does not cross the light; therefore, the sensor does not suffer any external interference such as bird sounds, cars, or airplane noise.

The data captured by the sensor are constituted, in general, of background noise with occasional "events", resulting in the brief moment that an insect flies across the laser. The procedure used to collect and preprocess the data used in this work, as well a description of the resulting data set are presented in next section.

III. Collecting and Preprocessing Data

To evaluate and compare classification systems, we employ laboratory data for which ground-truth labels are available, i.e., we need to know the true class labels of each insect passage to assess the classification procedures. These data were collected in several controlled environments ("insectaries"), each consisting of insects of a single species with an individual sensor attached.

More specifically, we used nine insectaries, each with dozens of specimens of a single species. In the experiments presented in this paper, we included four species of mosquitoes: *Aedes aegypti* (vector of filariasis, dengue fever, yellow fever, and West Nile virus), *Anopheles gambiae* (vector of malaria), *Culex quinquefasciatus* (vector of lymphatic filariasis) and *Culex tarsalis* (vector of St. Louis Encephalitis and Western Equine Encephalitis); three species

of flies: *Drosophila melanogaster, Musca domestica* and *Psychodidae diptera*; the beetles *Cotinis mutabilis* and the bees *Apis mellifera*.

After collecting the data, we pre-processed the recordings and detected the insect passages in raw data. We designed a detector responsible for identifying the events of interest and separating them from background noise. The general idea of the detector is to move a sliding window across the raw data and calculate the spectrum of the signal inside the window. As most insects have wing beats which range from 100Hz to 1000Hz, we used the maximum magnitude of the signal spectrum in this range as the detector confidence. Due to lack of space, we direct the interested reader to [10] for a more detailed description of the detector.

The detector outputs audio fragments which usually last for a few tenths of a second and have at least one insect passage. Due to the simplicity of the design of our electronic circuit, there is some noise mixed with the insect signals. We filter most of the noise using a digital filter based on spectral subtraction responsible for the removal of certain frequency ranges of signal [11]. An example of a filtered and segmented signal is shown in Figure 3.

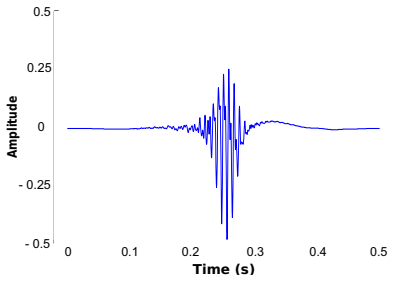


Figure 3. Filtered data segment obtained from an *Aedes aegypti* passage on the sensor laser

There was a total of 18,115 insect passages. For evaluation purposes, we divided the data into two stratified sets: 33% for training and 67% for testing. Table I presents a description of the training and testing split sizes.

Table I
DESCRIPTION OF THE DATA SET, TRAINING AND TEST SPLITS

Species	Train	Test	All	Class prior (%)
Aedes aegypti	1,585	3,171	4,756	26.25
Anopheles gambiae	470	941	1,411	7.79
Apis mellifera	170	341	511	2.82
Cotinis mutabilis	58	114	172	0.95
Culex quinquefasciatus	1,045	2,092	3,137	17.32
Culex tarsalis	1,770	3,539	5,309	29.31
Drosophila melanogaster	259	518	777	4.29
Musca domestica	448	895	1,343	7.41
Psychodidae diptera	233	466	699	3.86
Total	6,038	12,077	18,115	100

IV. FEATURE EXTRACTION

The feature extraction phase is crucial to the application's success, since machine learning algorithms are fully dependent on the predictive attributes of their input features. In particular, in signal domains such as audio, the raw data are often composed of a huge amount of very weak features. Typical machine learning algorithms will have difficulty with such weak data (or will require enormous training sets), mainly due to the curse of dimensionality [12]. Since our insect sensor generates signals very similar to audio data, we expect that feature extraction will be very important, and we investigate various feature extraction techniques borrowed from audio processing.

The most popular features in audio signal recognition are the Mel-Frequency Cepstral Coefficients (MFCCs). These are derived from a frequency spectrum calculated on the *mel* scale, an early approximation to human frequency sensitivity [13]. MFCCs are then obtained as the Discrete Cosine Transform of the logarithm of the spectral magnitudes on this frequency axis.

Linear Predictive Coding coefficients (LPC) are another feature set widely used in audio classification tasks such as spoken digit recognition [14]. Line Spectral Frequencies (LSF), a LPC-based feature set, have also proved to be effective in similar tasks [15].

Both temporal and spectral representations of the signal can provide important information for classification. Examples of temporal representations include simple signal statistics such as energy, magnitude, root mean square energy, and zero-crossing rate, as well as signal duration. Additional representations based specifically on the frequency spectrum include spectral centroid, spectral flatness measures, flux, and inharmonicity. An important higher-level property derived from the spectrum is the fundamental frequency, for signals that exhibit a dominant periodicity. For our domain, this information is associated with wing beat frequency, an attribute frequently used in entomology as a basis for differentiating species [16] [17].

In our experiments, we use MFCC, LPC, LSF, temporal, and spectral feature sets. Certain feature sets, such as MFCC, use a scale based on the human perception of sound. However, there is no a priori reason to limit our approach to the limited frequency range and resolution of human hearing. To circumvent this, we also evaluated the Linear-Frequency Cepstrum (LFC) and the Log-Linear Cepstrum (LLFC).

Lacking more space in this paper, we have provided detailed description of our feature extraction techniques on the companion website for this paper [18]. The website also includes numerical results, source code, and the data employed in our experiments.

For interested readers, we recommend [19] for a didactic description of cepstral coefficients, LPC, and LSF. For the temporal and spectral features, we recommend [20].

V. EXPERIMENTAL EVALUATION

Using the features described in the previous section, we now evaluate different machine learning techniques. Most

learning algorithms have parameters that can significantly influence their performance, and our first experiment consists of a search for the parameters that maximize classification accuracy. Since the use of test data is restricted to the final classifiers' evaluation, we used 10-fold cross-validation on the training data to search the parameter values. For each possible combination of parameter values, the accuracy of the classifier was measured in the "internal" cross-validation test sets. We use the best combination of parameter values for a given learning algorithm as the final setting, then use this combination to learn over the entire training set and evaluate the resulting classifier on the test set.

In the case of Support Vector Machine (SVM), we varied the parameters of the base algorithm and of the kernel using grid search [21]. Given values of minimum, maximum and step size, we evaluate the cross-validation accuracy of each combination of parameters. However, this search was performed with only a coarse, 2-fold cross-validation. The search was then refined in regions with better results.

The learning algorithms, as well as parameter ranges, are described in Table II.

Table III presents the results of the first experiment. For reasons of space, we omit results from by Naïve Bayes and J48 classifiers, since they gave the worst results across all feature sets. Additionally, we only show the results for SVM RBF since SVM Poly gave inferior results. The complete set of results can be found on the paper website [18].

Table III
ACCURACY RESULTS PER CLASSIFIER AND FEATURE SET, AND OPTIMAL PARAMETER VALUES. THE BEST RESULT IN EACH FEATURE SET IS IN BOLDFACE

Feature Set	Algorithm	Selected Parameter Configuration	Accuracy (%)
LFC	KNN	#Coeff. = 75, k = 7	81.71
	SVM RBF	#Coeff. = 95, C = 10, γ = 1	**86.93**
	GMM	#Coeff. = 100, #Gaussians = 9	83.17
	RF	#Coeff. = 80, #Trees = 75	83.49
LLFC	KNN	#Coeff. = 15, k = 7	74.70
	SVM RBF	#Coeff. = 70, C = 10000, γ = 0.01	**79.05**
	GMM	#Coeff. = 20, #Gaussians = 17	74.03
	RF	#Coeff. = 20, #Trees = 60	76.30
MFCC	KNN	#Coeff. = 30, k = 5	83.61
	SVM RBF	#Coeff. = 40, C = 10, γ = 1	**87.33**
	GMM	#Coeff. = 45, #Gaussians = 13	82.42
	RF	#Coeff. = 35, #Trees = 75	85.39
LPC	KNN	#Coeff. = 45, k = 21	56.18
	SVM RBF	#Coeff. = 45, C = 100000, γ = 0.1	**66.85**
	GMM	#Coeff. = 40, #Gaussians = 19	54.15
	RF	#Coeff. = 65, #Trees = 75	60.90
LSF	KNN	#Coeff. = 95, k = 5	80.23
	SVM RBF	#Coeff. = 100, C = 10, γ = 1	**84.97**
	GMM	#Coeff. = 75, #Gaussians = 17	75.28
	RF	#Coeff. = 95, #Trees = 75	84.25
Temporal	KNN	k = 11	50.91
	SVM RBF	C = 100000, γ = 0.1	**60.62**
	GMM	# Gaussians = 19	42.76
	RF	#Tress = 75	60.13
Spectral	KNN	k = 5	70.51
	SVM RBF	C = 100000, γ = 0.1	76.24
	GMM	# Gaussians = 21	63.73
	RF	#Tress = 50	**79.38**

The best results were obtained with MFCCs, with LFC and LSF achieving slightly lower accuracies, and the spectral feature set and LLFCs slightly lower again. The results obtained with the temporal set and LPC features were substantially lower that the other features. The best single-classifier performance of 87.33% was obtained with the RBF SVM classifier applied to MFCC features, and appears to be a respectable accuracy rate given the complexity of this application.

The second experiment investigates combining the output of different classifiers. Different algorithms and different features may make errors on different examples, so we can construct classifier ensembles to exploit this diversity. To combine the results, we used three different strategies. The first and simplest is voting: Each classifier votes for the predicted class and the final answer is given by the class with the highest number of votes. In case of a tie, the class with highest prior probability is chosen.

The other two strategies use sum and product functions, respectively, over the output score of each classifier. One possible advantage of these strategies in relation to voting is that they consider the fact that classifiers can assign similar score values to two different classes when an object is close to borderline regions. Therefore, the classification of borderline cases can potentially benefit from such forms of classifier combination.

Table IV presents the results of ensembles on different algorithms using the same feature set. We show two results for each feature set: the first one obtained with the combination of all classifiers on that feature set; and the second obtained by combining only the three best classifiers.

Table IV
RESULTS ACHIEVED BY THE COMBINATION OF DIFFERENT CLASSIFIERS ON THE SAME FEATURE SET. THE HIGHLIGHTED RESULTS REPRESENT AN ACCURACY GAIN OVER THE BEST BASE CLASSIFIER

Feature Set	Best Accuracy (%)	Combined Algorithms	Accuracy (%) Sum	Product	Voting
LFC	86.93	SVM RBF, KNN, GMM, RF	84.86	84.70	86.07
		SVM RBF, GMM, RF	83.58	83.94	86.29
LLFC	79.05	SVM RBF, KNN, GMM, RF	77.94	77.75	**79.12**
		SVM RBF, GMM, RF	77.48	77.88	78.68
MFCC	87.33	SVM RBF, KNN, GMM, RF	85.48	85.22	86.69
		SVM RBF, KNN, RF	85.30	85.80	86.59
LSF	84.97	SVM RBF, KNN, GMM, RF	81.72	80.49	84.64
		SVM RBF, KNN, RF	83.78	84.15	84.84
Spectral	79.38	SVM RBF, KNN, GMM, RF	73.82	72.55	77.02
		SVM RBF, GMM, RF	77.22	77.51	78.41

The results clearly show that the combination of different classifiers using the same feature set does not improve classification accuracy systematically. The accuracy rates obtained by the ensembles were higher than the best base classifier in only one (3.33%) of the analyzed cases. Even in this case, the gain was not significant.

We also evaluate the hypothesis that the combination of different representations can provide enough diversity to improve the classification accuracy. We performed experiments

Table II
LEARNING ALGORITHMS WITH THEIR RESPECTIVE PARAMETER RANGES

Algorithm	Acronym	Parameters	Parameters range (initial:step:final)
Decision Tree (J48 implementation)	J48	Pruning factor	P = 0.1:0.1:0.5
Gaussian Mixture Models	GMM	Number of components (Gaussians)	N = 3:2:21
K-Nearest Neighbors	KNN	Number of neighbors	K = 1:2:25
Naïve Bayes	NB	-	-
Random Forest	RF	Number of trees	N = 5:2:75
Support Vector Machine - Polynomial Kernel	SVM Poly	Complexity C / Polynomial Degree	C = 10^i, i = -7:1:5 / D = 1:1:3
Support Vector Machine - RBF Kernel	SVM RBF	Complexity C / γ	C = 10^i, i = -7:1:5 / γ = 10^i, i = -4:1:0

with different combinations of feature sets using the same induction algorithm.

First, we checked if different frequency scales used to extract cepstral coefficients can be complementary. Thus, we created combinations of LFC, LLFC and MFCC. We also used LSF and Spectral features in combination with MFCC, since they are the best known and most used cepstral features and achieved some of the best results in our first experiment, and LFC, which obtained competitive results in comparison to MFCC. In addition, we also evaluated the combination of all feature sets (LFC, LLFC, MFCC, LSF and Spectral). Table V shows the results.

Table V
RESULTS ACHIEVED BY THE COMBINATION OF DIFFERENT FEATURE SETS WITH THE SAME LEARNING ALGORITHM. THE HIGHLIGHTED RESULTS REPRESENT AN ACCURACY GAIN OVER THE BASE CLASSIFIER

Algorithm	Best Accuracy (%)	Combined Feature Sets	Accuracy (%) Sum	Product	Voting
SVM RBF	87.33	LFC, LLFC, MFCC	**87.46**	87.27	**87.91**
		LFC, LSF, Spectral	86.83	86.44	87.09
		MFCC, LSF, Spectral	86.85	86.35	87.14
		All feature sets	**88.70**	**88.47**	**88.44**
KNN	83.61	LFC, LLFC, MFCC	**85.48**	**85.57**	**84.57**
		LFC, LSF, Spectral	**83.94**	83.56	82.46
		MFCC, LSF, Spectral	**84.82**	**84.45**	83.05
		All feature sets	**86.15**	**86.00**	**85.18**
GMM	83.17	LFC, LLFC, MFCC	**85.50**	**86.35**	**84.72**
		LFC, LSF, Spectral	83.17	**84.16**	81.49
		MFCC, LSF, Spectral	82.86	82.68	81.18
		All feature sets	**86.20**	**86.01**	**85.50**
RF	85.39	LFC, LLFC, MFCC	**86.69**	**86.93**	84.82
		LFC, LSF, Spectral	**86.50**	**86.36**	84.76
		MFCC, LSF, Spectral	**86.99**	**86.89**	**85.44**
		All feature sets	**87.83**	**87.97**	**86.14**

The combination of different feature sets provided a significant number of accuracy improvements. In total, 31 (64.58%) of the analyzed cases showed some improvement. It is worth noticing that the combination of all feature sets improves the accuracy over the base classifiers in all cases.

So far, we have evaluated our classifiers in a multiclass setting. Although this setting provides an overall assessment of our classifiers, not all classes are equally important in most applications. Generally, an intelligent trap is used to capture a single specie of interest, such as a disease vector or an agricultural pest. All other species can be regarded as a negative class, and are set free by the trap. Therefore, many practical applications require just a binary classifier. In this context, we analyzed the performance of classifiers that consider disease-carrying insects as positive class and other species as negative class.

In analyzing the classification with such division, we cause considerable variation in the distribution of the data. The class distribution, for instance, is greatly changed. In the complete data set (training and test) the number of samples generated by insect vectors of disease is 14613, compared to just 3502 examples of the other species. Therefore, in this experiment we do not evaluate only the accuracy, but also the area under the ROC curve (AUC), more recommended for evaluating the classification in problems involving imbalanced classes [22].

In this task, we used 40 MFCC to train a SVM-RBF, since we found that this configuration achieved the best result. We also evaluated the result obtained by the combination (ensemble) of all cepstrum scales, LSF, and spectral features, using the best classifier for each and combined by sum. In the first case, we achieved an accuracy of 97.82% and a AUC of 0.966. In the second case, combining classifiers, we achieved 98% of accuracy and 0.968 of AUC. The gain in accuracy is due to the observation that the errors in the multiclass classifiers were concentrated among species with similar characteristics, such as different mosquito species. With these results, it is evident that the proposed trap is efficient in the classification and selective capture of insects.

VI. CONCLUSIONS

We have presented for the first time an intelligent insect trap that selectively captures harmful insect species. The trap uses a laser sensor to classify the species of every insect that enters the trap. We reported on a broad experimental evaluation of feature extraction techniques and machine learning algorithms for automatic insect classification using the signals recorded by our sensor.

Feature extraction techniques used in audio analysis achieved very good results in our experiments. The best results achieved in this work were obtained with MFCCs, a traditional technique used in audio processing that extracts coefficients from frequencies scaled to imitate the sensitivity of the human auditory system. We also showed that combining multiple different classifiers based on the same feature set was not effective. In contrast, a systematic accuracy

improvement was obtained by ensembles over different feature sets, obtained from different signal representations.

By posing our task as a binary classification of whether or not a particular insect should be captured, we achieved an accuracy close to 99%. Our experimental results indicate that a trap employing such classifiers has great potential for selectively capturing target insect species.

The results presented in this work are important in themselves, but they are also useful as a guide for future work. This could include the evaluation and adaptation of techniques presented to anytime and anyspace classification, as well as streaming analysis in non-stationary environments. Anytime and anyspace techniques could be useful given the memory and processing limitations imposed by the sensor's hardware. The analysis of the signal as a non-stationary data stream, looking for concept drifts, could be useful in the case of environmental variations that cause changes in some aspect of the insect features.

ACKNOWLEDGMENT

This work was funded by São Paulo Research Foundation (FAPESP). Grants #2011/04054-2, #2011/17698-5 and #2012/50714-7.

REFERENCES

[1] D. Pimentel, "Environmental and economic costs of the application of pesticides primarily in the united states," in *Integrated Pest Management: Innovation-Development Process*. Springer, 2009, pp. 89–111.

[2] M. Q. Benedict and A. S. Robinson, "The first releases of transgenic mosquitoes: an argument for the sterile insect technique," *Trends in parasitology*, vol. 19, no. 8, pp. 349–355, 2003.

[3] P. Kevan, "Pollinators as bioindicators of the state of the environment: species, activity and diversity," *Agriculture, Ecosystems & Environment*, vol. 74, no. 1-3, pp. 373–393, 1999.

[4] W.H.O., "Dengue: guidelines for diagnosis, treatment, prevention and control," World Health Organization, Tech. Rep., 2009.

[5] ——, "The world malaria report," World Health Organization, Tech. Rep., 2012.

[6] G. T. Miller, "Sustaining the earth." Cengage Learning, 2004, ch. 9, pp. 211–216.

[7] J. Hemingway and H. Ranson, "Insecticide resistance in insect vectors of human disease," *Annual review of entomology*, vol. 45, no. 1, pp. 371–391, 2000.

[8] S. Rajatileka, J. Burhani, and H. Ranson, "Mosquito age and susceptibility to insecticides," *Transactions of the Royal Society of Tropical Medicine and Hygiene*, vol. 105, no. 5, pp. 247–253, 2011.

[9] D. Kline, W. Takken, J. Wood, and D. Carlson, "Field studies on the potential of butanone, carbon dioxide, honey extract, l-octen-3-ol, l-lactic acid and phenols as attractants for mosquitoes," *Medical and veterinary entomology*, vol. 4, no. 4, pp. 383–391, 1990.

[10] G. E. A. P. A. Batista, Y. Hao, E. Keogh, and A. Mafra-Neto, "Towards automatic classification on flying insects using inexpensive sensors," in *Proceedings of the 10th International Conference on Machine Learning and Applications*, vol. 1, 2011, pp. 364–369.

[11] S. Boll, "Suppression of acoustic noise in speech using spectral subtraction," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 27, no. 2, pp. 113–120, 1979.

[12] K. Beyer, J. Goldstein, R. Ramakrishnan, and U. Shaft, "When is "nearest neighbor" meaningful?" in *Proceedings of the International Conference on Database Theory*. Springer, 1999, pp. 217–235.

[13] S. S. Stevens, J. Volkmann, and E. B. Newman, "A scale for the measurement of the psychological magnitude pitch," *Journal of the Acoustical Society of America*, vol. 8, no. 3, pp. 185–190, 1937.

[14] M. Sambur and L. Rabiner, "A statical decision approach to the recognition of connected digits," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 24, no. 6, pp. 550–558, 1976.

[15] D. F. Silva, V. M. A. Souza, G. E. A. P. A. Batista, and R. Giusti, "Spoken digit recognition in portuguese using line spectral frequencies," in *Advances in Artificial Intelligence – IBERAMIA 2012*. Springer, 2012, vol. 7637, pp. 241–250.

[16] C. Hyatt and D. Maughan, "Fourier analysis of wing beat signals: assessing the effects of genetic alterations of flight muscle structure in diptera," *Biophysical journal*, vol. 67, no. 3, pp. 1149–1154, 1994.

[17] R. Buchwald and R. Dudley, "Limits to vertical force and power production in bumblebees (hymenoptera: Bombus impatiens)," *Journal of Experimental Biology*, vol. 213, no. 3, pp. 426–432, 2010.

[18] D. F. Silva, V. M. A. Souza, G. E. A. P. A. Batista, D. P. W. Ellis, and E. Keogh. Website to this work. http://sites.labic.icmc.usp.br/dfs/ICMLA2013.

[19] J. Benesty, M. M. Sondhi, and Y. Huang, Eds., *Springer Handbook of Speech Processing*. Berlin: Springer, 2008.

[20] T. H. Park, "Towards automatic musical instrument timbre recognition," Ph.D. dissertation, Princeton University, NJ, USA, November 2004. [Online]. Available: http://www.tulane.edu/~park/publications/ParkPrincetonThesis2004.pdf

[21] C. W. Hsu, C. C. Chang, and C. J. Lin, "A practical guide to support vector classification," Department of Computer Science, National Taiwan University, Tech. Rep., 2003.

[22] G. E. A. P. A. Batista, R. C. Prati, and M. C. Monard, "A study of the behavior of several methods for balancing machine learning training data," *SIGKDD Explor. Newsl.*, vol. 6, no. 1, pp. 20–29, Jun. 2004. [Online]. Available: http://doi.acm.org/10.1145/1007730.1007735