Automatic insect classification with Machine Learning techniques: a comparison of similarity and feature extraction approaches

Diego F. Silva[1], Eamonn Keogh[2], Gustavo E. A. P. A. Batista[1],

[1] [1]Instituto de Ciências Matemáticas e de Computação, Universidade de São Paulo
{diegofsilva,gbatista}@icmc.usp.br
[2]University of California Riverside
eamonn@cs.ucr.edu

Abstract. *Insects are intimately related to human beings, in both positive and negative ways. For example, insect pests consume and destroy around US\$40 billion worth of food each year. 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. In the last decades, many researchers have developed an arsenal of chemical, biological, mechanical and educational methods of insect control. However, to be effectively used, such methods require knowledge of the spatio-temporal distribution of the insects. Without such knowledge, the use of these techniques becomes costly and inefficient. A sensor for capturing insect information is being developed with the aim of being used as a tool to assist in the control of disease vectors and agricultural pests. The main elements of this sensor are a laser beam and an array of phototransistors. When an insect crosses the laser beam, a variation in the light is caused by partial occlusion of light due to their movements. This variation is stored as a time series and should be used to count and classify insects that cross the sensor. In this paper, we investigate the use of different approaches for time series classification that can be applied to insect recognition by the laser sensor: similarity search and feature extraction. In an experiment that includes nine species of insects, we demonstrate that the feature extraction approach can be more accurate that the similarity search. More specifically, the Support Vector Machine algorithm with RBF kernel trained with mel-cepstral coefficients achieved the best accuracy in the insect recognition task.*

1. Introduction

Insects are intimately related to human beings, in both positive and negative ways. For example, insect pests consume and destroy around US\$40 billion worth of food each year [Pimentel 2009]. 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 [Benedict and Robinson 2003]. Furthermore, many species have been used as bioindicators of environmental quality, since their presence/absence, distribution and density, indicate the quality of the ecosystem, especially in relation to contaminants in the air, soil and water [Kevan 1999].

Another example of the relationship between insects and humans are the 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 [W.H.O. 2009]. 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 [W.H.O. 2012].

For these and other reasons, many researchers have developed an arsenal of chemical, biological, mechanical and educational methods of insect control [Walker 2002]. However, to be effectively used, such methods require knowledge of the spatio-temporal distribution of the insects. Without such knowledge, the use of these techniques becomes costly and inefficient, besides aggravating the problems mentioned above.

Currently, studying the spatio-temporal distribution of insects is a costly and time-consuming task. In general, insect counts are are obtained with traps, usually adhesive, which are collected periodically and analyzed by experts who manually identify and count the collected species of insects. Besides being an expensive approach in terms of material and human resources, there is a delay between the moment when the trap is installed and when it is analyzed. Even though this range is only a week, which may represent more than half life of an adult insect, such delay may be enough for the disease to infect a large number of people until the data be available to the experts [Patnaik et al. 2007]. Therefore, there is a need for automatic and accurate sensors which can detect, classify and count insects of different species in real time.

In this paper, we describe and analyze the data collected by a laser sensor proposed to automatically count and classify insects. We also conduct an extensive experimental evaluation of Machine Learning techniques to accurately identify insects species.

The remaining of this paper is organized as follows. Section 2 presents the related work about automatic insect identifycation. Section 3 describes the sensor used in this work, as well the data collecting procedure. Section 4 presents the results obtained by the application of classification methods. Finally, we present our conclusions and directions for future work in Section 5.

2. Related Work

The idea of performing automatic classification of insects is not a novelty. In 1945, Kahn et. al. [Kahn et al. 1945] used a microphone, a signal amplifier, a low-pass filter and a sound recorder to register and study inaudible sounds produced by disease vector mosquitoes. They collected the sounds of four species: *Anopheles quadrimaculatus*, *Aedes aegypti*, *Aedes albopictus*, and *Culex pipiens*. To perform the sound collection, an environment without external noise and under ideal conditions of temperature and humidity was necessary. Different sounds that could represent the insect behaviors were identified. Furthermore, the study shows that the pitch can be used to distinguish male and female of the same species. This is possible because the sounds produced by male mosquitoes have a higher frequency than the sounds produced by female mosquitoes.

A few years later, Kahn & Offenhauser Jr [Kahn and Jr 1949] mentioned that the fast evolution of eletronic devices for sound recording would make the study of insects behavior easy, fast and accurate by using the sounds they produce. However, we note a small evolution related to the automatic identification of insects by acoustic devices.

More recently, researchers have attempted to identify species and analyze the behavior of insects through the use of microphones. The general procedure of these studies is the use of signal processing techniques for features extraction and the application of machine learning algorithms. For example, to classify crickets and cicadas, Potamitis et. al. [Potamitis et al. 2007] used a probabilistic neural network and a gaussian mixture model from features widely used in speaker recognition applications. In a more recent study, Le-Qing [Le-Qing 2011] also used features from speaker recognition application and a probabilistic neural network to classify the different behaviors of insects such as wings vibration, locomotion and alimentation in soil, wood and in other materials.

Several other studies use the approach of audio recording to analyze insects. For example, Ganchev et. al. [Ganchev et al. 2007] used this approach to classify 313 species of crickets, grasshoppers and cicadas. They used cepstral coefficients to generate classifiers based on probabilistic neural network, gaussian mixture models and hidden markov models. Based on studies of speaker recognition, Chaves et. al. [Chaves et al. 2012] used the mel-frequency cepstral coefficients (MFCC) and hidden markov models to classify 36 species of grasshoppers.

All the previous related work has performed the audio recordings in an environment with ideal conditions. Furthermore, most of the analyzed species produce very evident sounds, such as the songs produced by crickets or cicadas. However, the use of microphones in non-ideal environments leads to some difficulties. Microphones are very sensitive to external interference, such as sounds produced by cars transiting near the location of data collection.

Taking into account these difficulties, Moore et. al. [Moore et al. 1986] proposed the use of an optical sensor based on a phototransistor. The authors used the sensor to record the variation of the light caused by passages of insects. They performed an analysis of the wing-beat frequency of two species of the genus *Aedes* from both sexes. The automatic classification considering species and sex was posteriorly presented in Moore [Moore 1991].

Some years later, Moore [Moore 1998] proposed an insect data collection system. Basically, he used the previously proposed optical sensor connected to a computer with multimedia features and tools to process the obtained signal. He placed a transparent plastic jar with flying insects above the sensor. The light source used in the system is a halogen lamp located above the transparent jar.

More recently, researchers presented a new optical sensor to automatically identify flying insects [Batista et al. 2011b]. The basic components of this sensor are a laser light, an array of phototransistors and a circuit board to filter and record the variation in the light caused by the insects that cross the light. In this paper, we present results of a larger experimental evaluation using similarity and feature extraction approached, as well as combination of classifiers using ensembles.

3. Laser Insect Sensor

The main elements of the sensor are a laser beam and an array of phototransistors. When an insect crosses the laser beam, a variation of light is caused by partial occlusion of light due to the wings movements. Such a variation is stored as a short time series. Our main goal is to build a classification system that takes such a time series as input and provides counts of insects discriminated by species.

3.1. Sensor Description

The general design of the sensor used in this work is shown in Figure 1. It consists of a low-powered planar laser source pointed to an array of phototransistors. When a flying insect crosses the laser, its wings partially occlude the light, causing small light variations that are captured by the phototransistors. An electronic circuit board filters and amplifies the signal and the output is recorded by a digital sound recorder.

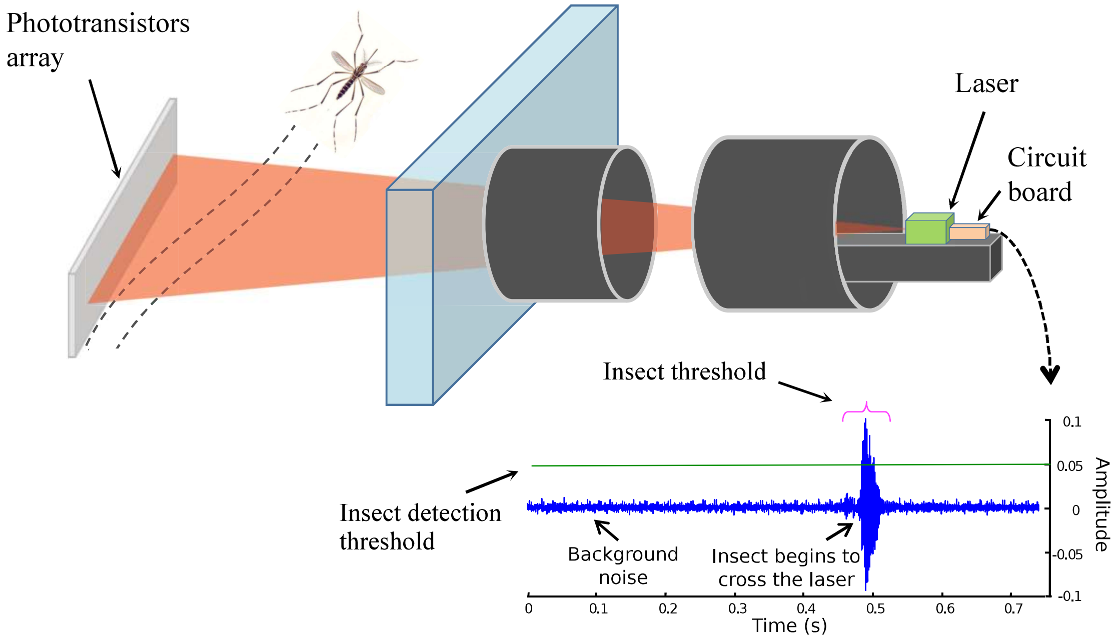

Figure 1. The logical design of the sensor. A planar laser light is directed at an array of phototransistors. When an insect flies across the laser, a light variation is registered by the phototransistors as a time series

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", result of the brief moment that an insect flew across the laser. In the next section, we provide details about the procedure used to collect and preprocess the data used in this work.

3.2. Collecting and Preprocessing Data

We use data collected in laboratory, for which ground-truth labels are available. We need to know the true class labels of each insect passage to assess the classification procedures. These data were collected in several containers ("insectaries"), each with an individual sensor attached and containing insects of a single species. Figure 2 shows some examples of these insectaries.

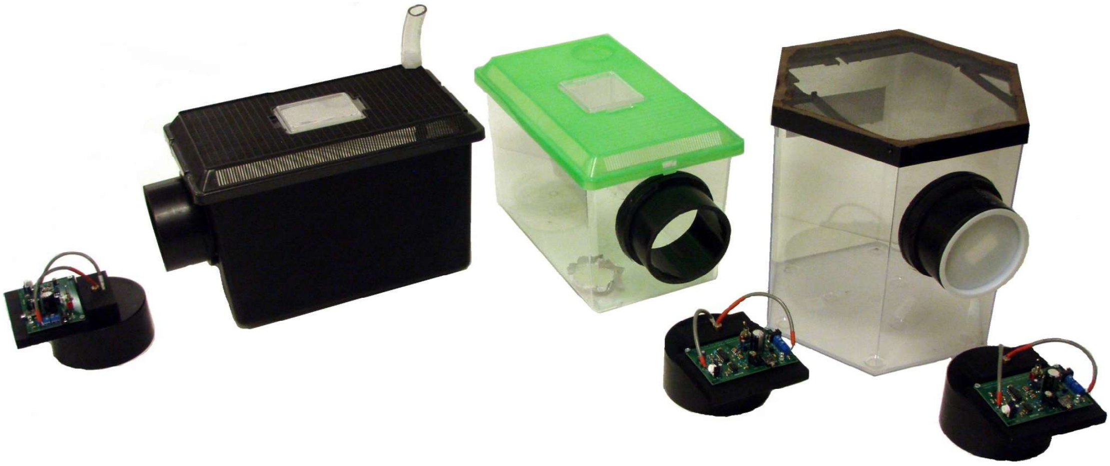

Figure 2. Examples of boxes for data collection (insectaries)

After collecting the data, we preprocessed 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 beat frequencies which range from 100Hz to 1000Hz, we used the maximum magnitude of the signal spectrum in this range as the detector confidence.

The detector uses a sliding window and calculates the magnitude of signal components within the window. Then, the maximum magnitude is taken as a confidence value for the detector. The larger the magnitude, the higher the confidence that the signal is not background noise. All signals with magnitude above a user-specified threshold are considered an event generated by an insect. The high signal to noise ratio of the data collected by the sensor allows the user to specify low values for the threshold without the risk of false positives. Figure 3 illustrates how the detector works.

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 the electronic circuit, there is some noise combined with the insect signals. So, we filtered most of the noise using a digital filter based on spectral subtraction, responsible for the removal of certain frequency ranges of signal [Boll 1979].

4. Experimental Results

In this section, we present experimental classification results using the strategies of similarity comparison and feature extraction.

4.1. Dataset description

In the experiments presented in this paper, we included four species of mosquitoes: *Aedes aegypti* (vector of filariasis, dengue, 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* also known as fruit fly, *Musca domestica* or house fly

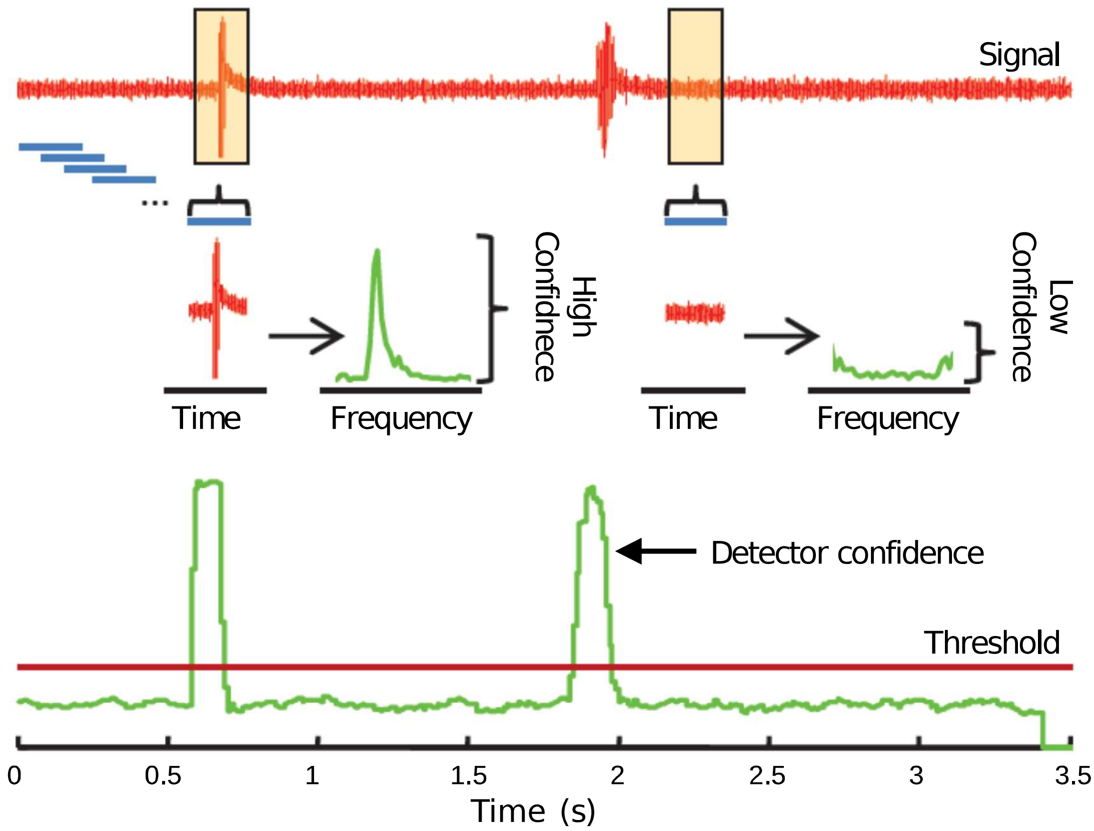

Figure 3. General design of the wing-beat detector [Batista et al. 2011a]

and *Psychodidae diptera* popularly known as moth fly; the beetle *Cotinis mutabilis* and the bee *Apis mellifera*. The number of examples of each species varies between 172 (0.95%) and $5,309$ (29.31%), for the species *Cotinis mutabilis* and *Culex tarsalis*, respectively.

The data set was divided into standard training and test partitions. This division was performed in a stratified approach, leaving 33% of the examples in the training set and the remaining in the test set.

4.2. Similarity Search

The similarity search is a simple approach for classification. In our domain, two design decisions can significantly influence the classifier performance: the distance measure and the data representation. In this section, we evaluate different distance measures applied to the spectrum and the cepstrum of the signals. The time domain was not used because the signals have different lengths and also because the alignment of important pieces of the signals (peaks and valleys) is a very sensitive issue. We start our analysis by comparing the use of different distances applied to cepstrum and spectrum. The results are presented in Table 1.

The results achieved by similarity on the spectrum were slightly superior then the ones obtained by the cepstrum. Given these results, we decided to extend the evaluation of the classification by similarity only to the spectral domain. Table 2 presents the distance measures and the classification accuracy.

Table 1. Result of classification by similarity over the spectrum and the cepstrum

Distance Measure	Accuracy (%) Spectrum	Accuracy (%) Cepstrum
Euclidean	76.14	78.66
Manhattan	80.09	67.24
Cosine	77.25	76.29
Correlation	76.60	75.34

Table 2. Result of classification by similarity in the frequency domain

Distance Measure	Accuracy (%)
Canberra	72.28
Chebyshev	71.20
Jaccard	77.26
Topsoe	81.54
Clark	75.59
Average L_1 L_∞	80.09
Squared χ^2	81.38
Additive Symmetric χ^2	81.01
DTW (band-width = 5 observations)	81.04

4.3. Feature Extraction

The feature extraction approach uses different representations of signals to identify features, which are used as input to machine learning algorithms.

In this work, we use temporal and spectral features. The interested reader can find a detailed review of these features in [Park 2004]. We use *temporal features* and *spectral features* to refer to feature vectors extracted from time and frequency domains, respectively. Table 3 lists the features that compose each of these vectors.

Table 3. List of features that compose temporal and spectral feature vectors

Domain	Feature
Temporal	Mean amplitude, Root mean square, Short-time energy, Interval, Temporal centroid, Zero-crossing rate, Complexity estimate [Batista et al. 2011c], Variance,Standard deviation, Skewness, Kurtosis, Duration
Spectral	Fundamental frequency, Inharmonicity, Tristimulus (1, 2 and 3), Flux, Spectral centroid, Energy, Spectral irregularity, Modified spectral irregularity, Variance, Standard deviation, Skewness, Kurtosis, Mean magnitude, Roll-off, Flatness

Moreover, we also use Mel-Frequency Cepstrum Coefficients (MFCC), Linear Prediction Coefficients (LPC) and Line-Spectral Frequencies (LSF). 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 issue, we also evaluated the Linear-Frequency Cepstrum (LFC) and the Log-Linear Frequency Cepstrum (LLFC).

We evaluate several machine learning techniques using these features. Most learning algorithms have parameters that can significantly influence their performance. 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.

In the case of Support Vector Machine, we use grid search [Hsu et al. 2003] to vary the parameters of the base algorithm and of the kernel. Given values of minimum, maximum and step size, we evaluate the cross-validation accuracy of each combination of parameters. This search is performed with coarse estimate, using 2-fold cross-validation. The search is then refined in regions with better results.

The learning algorithms, as well as parameter ranges, are described in Table 4.

Table 4. 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	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 / 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

Table 5 presents the results of the first experiment. For reasons of readability, we omit results obtained by Naïve Bayes and J48 classifiers, since they achieved the worst results across all feature sets. Additionally, we only show the results for SVM RBF since SVM Poly had inferior results. Finally, we also omited the results obtained by using LPC and temporal features, because they are significantly worse than the ones obtained by other feature vectors.

Table 5. Accuracy results per classifier and feature set. and optimal parameter values. The best result in each feature set is highlighted

Feature Set	Algorithm	Selected Parameter Configuration	Accuracy (%)
LFC	KNN	#c= 75. $k = 7$	81.71
	RF	#c= 80. $T = 75$	83.49
	SVM RBF	#c= 95. $c = 10.\gamma = 1$	**86.93**
	GMM	#c= 100. $G = 9$	83.17
LLFC	KNN	#c= 15. $k = 7$	74.70
	RF	#c= 20. $T = 60$	76.30
	SVM RBF	#c= 70. $c = 10^4.\gamma = 0.01$	**79.05**
	GMM	#c= 20. $G = 17$	74.03
MFCC	KNN	#c= 30. $k = 5$	83.61
	RF	#c= 35. $T = 75$	85.39
	SVM RBF	#c= 40. $c = 10.\gamma = 1$	**87.33**
	GMM	#c= 45. $G = 13$	82.42
LSF	KNN	#c= 95. $k = 5$	80.23
	RF	#c= 95. $T = 75$	84.25
	SVM RBF	#c= 100. $c = 10.\gamma = 1$	**84.97**
	GMM	#c= 75. $G = 17$	75.28
Spectral	KNN	$k = 5$	70.51
	RF	$T = 50$	**79.38**
	SVM RBF	$c = 10^5.\gamma = 0.1$	76.24
	GMM	$G = 21$	63.73

The best results were obtained with MFCC, being that LFC and LSF achieved slightly lower accuracy rates, and the spectral feature set and LLFC also slightly lower. The results obtained with temporal features and LPC were substantially lower than the other features. The best single classifier performance, 87.33%, was obtained with the SVM RBF classifier applied to MFCC, and seems to be a respectable accuracy rate given the complexity of the application. The best result obtained by similarity search was 81.87%.

We evaluated the hypothesis that the combination of different representations can provide enough diversity to improve the classification accuracy. We performed experiments with different combinations of feature sets using the same induction algorithm. To combine the results, we used the sum of classification scores.

First, we checked if different frequency scales used to extract cepstral coefficients can be complementary. So, 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).

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 noting that the combination of all feature sets improves the accuracy over the base classifiers in all cases. The best result, 88.70%, was achieved by combining the five feature sets using the sum of SVM RBF outputs.

This result may lead the reader to questions about the real contribution of each feature in classifier combinations. So far, we only used combinations of classifiers outputs, obtained by using different features. To know the real contribution of the different types of features, we built a data set with all features with the largest number of coefficients used previously. In other words, we built a dataset with 529 features: 100 LFC, 100 LLFC, 100 MFCC, 100 LSF, 100 LPC, 12 temporal features and 17 spectral features.

Due to the high dimensionality of this dataset, feature selection techniques were applied on it. Specifically, we used the Correlation-Based Feature Selection (CFS) [Hall 1999] and the Relief [Kononenko 1994] algorithms. In the case of Relief, the algorithm just creates a ranking of features according to their quality. We must then choose how many features will be used and select them according to the order established by the algorithm. To do this, we used 27, 53, 106 and 159 features (5%, 10%, 20% and 30% of total). The CFS algorithm does require have this parameter, and this algorithm automatically selected 74 features.

Interestingly, the MFCC are always selected in a large number. In all cases, CFS and variations of Relief, the feature vector with larger number of selected coefficients was always the MFCC. LFC and LSF were also taken in large numbers by the feature selection algorithms. The same happened for the spectral attributes. In contrast, the LPC and temporal features were mostly discarded.

The learning algorithms used in this phase were the KNN, SVM with RBF and Random Forest. This choice was made because these algorithms have provided the best results in previous experiments. The results are shown in Table 6.

The use of all features does not systematically improve the performance of classifiers. In one of the analyzed classifiers, this strategy achieved a lower performance than the classifier trained with only one feature vector. The same does not happen when a feature selection strategy is used. In the case of CFS, its application improved classification performance in all cases. The same happened for the algorithm Relief with certain number of selected features. In this case, 20% (106) and 30% (159) of the total.

Table 6. Classification result with all feature sets and feature selection techniques. The highlighted values are relative to those with better performance than the base classifier considering the best feature set for it

Learning Algorithm	Individual Accuracy (%)	All Features	CFS	Relief 5%	Relief 10%	Relief 20%	Relief 30%
KNN	83.61	83.51	**86.19**	83.07	82.76	**83.85**	**85.23**
RF	85.39	**86.98**	**88.37**	**85.63**	**86.16**	**86.86**	**87.54**
SVM RBF	87.33	**89.14**	**88.78**	85.88	86.96	**87.38**	**89.55**

5. Conclusion

The sensor presented in this paper is important for a range of applications. For the effective operation of the sensor, it is necessary to investigate techniques for signal classification that can be used in this application. Thus, the aim of this study was to conduct and present a comprehensive investigation on these methods. We conducted our research with two approaches for time series classification: similarity search and feature extraction. Both approaches were applied using different representations.

We demonstrated the influence of different distance measures in our data. Thirteen distance measures were evaluated with classification by similarity in frequency domain and the accuracy ranged from 71.20% to 81.54%.

With the feature extraction approach, we evaluated features from temporal, spectral and cepstral representations, as well as features based on linear predictions coefficients and its variant LSF. We observed that, in different configurations of features and classifiers, the feature extraction approach is more accurate than the classification based on similarity search. More specifically, the Support Vector Machine algorithm with RBF kernel trained with MFCC achieved accuracy of 87.33%. This result represents an improvement of nearly 7% compared to the best classifier based on similarity search.

We also evaluated different ways to combine classifiers and features. The combination of different feature vectors as input to the same learning algorithm usually improves the results. In this case, the best accuracy was 88.70%. In the case of the features being used together, creating a new data set with feature subset selection techniques, the accuracy achieved 89.55%.

References

Batista, G. E. A. P. A., Hao, Y., Keogh, E. J., and Mafra-Neto, A. (2011a). Towards automatic classification on flying insects using inexpensive sensors. In *IEEE International Conference on Machine Learning and Applications and Workshops*, volume 1, pages 364–369.

Batista, G. E. A. P. A., Keogh, E. J., Mafra-Neto, A., and Rowton, E. (2011b). SIGKDD demo: sensors and software to allow computational entomology, an emerging application of data mining. In *ACM International Conference on Knowledge Discovery and Data Mining*, pages 761–764. ACM.

Batista, G. E. A. P. A., Wang, X., and Keogh, E. J. (2011c). A complexity-invariant distance measure for time series. In *SIAM International Conference on Data Mining*, pages 699–710.

Benedict, M. Q. and Robinson, A. S. (2003). The first releases of transgenic mosquitoes: an argument for the sterile insect technique. *Trends in parasitology*, 19(8):349–355.

Boll, S. (1979). Suppression of acoustic noise in speech using spectral subtraction. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 27(2):113–120.

Chaves, V., Travieso, C., Camacho, A., and Alonso, J. (2012). Katydids acoustic classification on verification approach based on mfcc and hmm. In *IEEE International Conference on Intelligent Engineering Systems*, pages 561–566.

Ganchev, T., Potamitis, I., and Fakotakis, N. (2007). Acoustic monitoring of singing insects. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, volume 4, pages 721–724.

Hall, M. A. (1999). *Correlation-based feature selection for machine learning*. PhD thesis, The University of Waikato.

Hsu, C. W., Chang, C. C., and Lin, C. J. (2003). A practical guide to support vector classification. Technical report, Department of Computer Science, National Taiwan University.

Kahn, M. C., Celestin, W., and Jr, W. O. (1945). Recording of sounds produced by certain disease-carrying mosquitoes. *Science*, 101(2622):335–336.

Kahn, M. C. and Jr, W. O. (1949). The identification of certain west african mosquitoes by sound. *The American Journal of Tropical Medicine and Hygiene*, 29(5):827–836.

Kevan, P. (1999). Pollinators as bioindicators of the state of the environment: species, activity and diversity. *Agriculture, Ecosystems & Environment*, 74(1-3):373–393.

Kononenko, I. (1994). Estimating attributes: analysis and extensions of relief. In *European Conference on Machine Learning*, pages 171–182.

Le-Qing, Z. (2011). Insect Sound Recognition Based on MFCC and PNN. *International Conference on Multimedia and Signal Processing*, 2:42–46.

Moore, A. (1991). Artificial neural network trained to identify mosquitoes in flight. *Journal of Insect Behavior*, 4:391–396.

Moore, A. (1998). Development of a data acquisition system for long-term outdoor recording of insect flight activity using a photosensor. In *Conference on Aerobiology and Biometeorology*. American Meteorological Society.

Moore, A., Miller, J. R., Tabashnik, B. E., and Gage, S. H. (1986). Automated identification of flying insects by analysis of wingbeat frequencies. *Journal of Economic Entomology*, 79(6):1703–1706.

Park, T. H. (2004). *Towards Automatic Musical Instrument Timbre Recognition*. PhD thesis, Princeton University.

Patnaik, J. L., Juliusson, L., and Vogt, R. L. (2007). Environmental predictors of human West Nile virus infections, Colorado. *Emerging Infectious Diseases*, 13(11):1788–1790.

Pimentel, D. (2009). Environmental and economic costs of the application of pesticides primarily in the United States. In *Integrated Pest Management: Innovation-Development Process*, pages 89–111.

Potamitis, I., Ganchev, T., and Fakotakis, N. (2007). Automatic acoustic identification of crickets and cicadas. In *International Symposium on Signal Processing and Its Applications*, pages 1 –4.

Walker, K. (2002). A review of control methods for african malaria vector. Technical Report 108, Bureau for Global Health.

W.H.O. (2009). Dengue: guidelines for diagnosis, treatment, prevention and control. Technical report, World Health Organization.

W.H.O. (2012). The world malaria report. Technical report, World Health Organization.