

## Video Article

# Flying Insect Detection and Classification with Inexpensive Sensors

Yanping Chen<sup>1</sup>, Adena Why<sup>2</sup>, Gustavo Batista<sup>3</sup>, Agenor Mafra-Neto<sup>4</sup>, Eamonn Keogh<sup>1</sup><sup>1</sup>Department of Computer Science and Engineering, University of California, Riverside<sup>2</sup>Department of Entomology, University of California, Riverside<sup>3</sup>Institute of Mathematics and Computer Sciences, University of São Paulo - USP<sup>4</sup>ISCA TechnologiesCorrespondence to: Yanping Chen at [ychen053@ucr.edu](mailto:ychen053@ucr.edu)URL: <http://www.jove.com/video/52111>DOI: [doi:10.3791/52111](https://doi.org/10.3791/52111)

Keywords: Bioengineering, Issue 92, flying insect detection, automatic insect classification, pseudo-acoustic optical sensors, Bayesian classification framework, flight sound, circadian rhythm

Date Published: 10/15/2014

Citation: Chen, Y., Why, A., Batista, G., Mafra-Neto, A., Keogh, E. Flying Insect Detection and Classification with Inexpensive Sensors. *J. Vis. Exp.* (92), e52111, doi:10.3791/52111 (2014).

## Abstract

An inexpensive, noninvasive system that could accurately classify flying insects would have important implications for entomological research, and allow for the development of many useful applications in vector and pest control for both medical and agricultural entomology. Given this, the last sixty years have seen many research efforts devoted to this task. To date, however, none of this research has had a lasting impact. In this work, we show that pseudo-acoustic optical sensors can produce superior data; that additional features, both intrinsic and extrinsic to the insect's flight behavior, can be exploited to improve insect classification; that a Bayesian classification approach allows to efficiently learn classification models that are very robust to over-fitting, and a general classification framework allows to easily incorporate arbitrary number of features. We demonstrate the findings with large-scale experiments that dwarf all previous works combined, as measured by the number of insects and the number of species considered.

## Video Link

The video component of this article can be found at <http://www.jove.com/video/52111/>

## Introduction

The idea of automatically classifying insects using the incidental sound of their flight dates back to the earliest days of computers and commercially available audio recording equipment<sup>1</sup>. However, little progress has been made on this problem in the intervening decades. The lack of progress in this pursuit can be attributed to several related factors.

First, the lack of effective sensors has made data collection difficult. Most efforts to collect data have used acoustic microphones<sup>2-5</sup>. Such devices are extremely sensitive to wind noise and ambient noise in the environment, resulting in very sparse and low-quality data.

Second, compounding these data quality issues is the fact that many researchers have attempted to learn very complicated classification models, especially neural networks<sup>6-8</sup>. Attempting to learn complicated classification models, with a mere tens of examples, is a recipe for over-fitting.

Third, the difficulty of obtaining data has meant that many researchers have attempted to build classification models with very limited data, as few as 300 instances<sup>9</sup> or less. However, it is known that for building accurate classification models, more data is better<sup>10-13</sup>.

This work addresses all three issues. Optical (rather than *acoustic*) sensors can be used to record the "sound" of insect flight from meters away, with complete invariance to wind noise and ambient sounds. These sensors have allowed the recording of millions of labeled training instances, far more data than all previous efforts combined, and thus help avoid the over-fitting that has plagued previous research efforts. A principled method is shown below that allows the incorporation of additional information into the classification model. This additional information can be as quotidian and as easy-to-obtain as the time-of-day, yet still produce significant gains in accuracy of the model. Finally, it is demonstrated that the enormous amounts of data we collected allow us to take advantage of "*The unreasonable effectiveness of data*"<sup>10</sup> to produce simple, accurate and robust classifiers.

In summary, flying insect classification has moved beyond the dubious claims created in the research lab and is now ready for real-world deployment. The sensors and software presented in this work will provide researchers worldwide robust tools to accelerate their research.

## Protocol

### 1. Insect Colony and Rearing

1. Mosquito Colony and Rearing
  1. Rear *Culex tarsalis*, *Culex quinquefasciatus*, *Culex stigmatosoma*, and *Aedes aegypti* adults from lab colonies, which originated from wild caught individuals.
  2. Rear mosquito larvae in enamel pans under standard laboratory conditions (27 °C, 16:8 hr light:dark [LD] cycle with 1 hr dusk/dawn periods), and feed them *ad libitum* on a mixture of ground rodent chow and Brewer's yeast (3:1, v:v).
  3. Collect mosquito pupae into 200 ml cups, and place them into experimental chambers. Alternatively, aspirate the adult mosquitoes into experimental chambers within 1 week of emergence. Make sure each experimental chamber contains 20 to 40 individuals of the same species/sex.
  4. Feed adult mosquitoes *ad libitum* on a 10% sucrose and water mixture. Replace food weekly.
  5. Moisten cotton towels twice a week and place them on top of the experimental chambers to maintain humidity within the cage. In addition, place a 200 ml cup of tap water in the chamber at all times to help maintain the overall humidity level.
  6. Maintain the experimental chambers on a 16:8 hr light:dark [LD] cycle, 20.5-22 °C and 30-50% RH for the duration of the experiment.
2. House Fly and Fruit Fly Colony and Rearing
  1. Rear *Musca domestica* from a lab colony, derived from wild caught individuals. Catch wild *Drosophila simulans* individuals and rear them in the experimental chambers.
  2. Rear *Musca domestica* larvae in plastic tubs under standard laboratory conditions (12:12 hr light:dark [LD] cycle, 26 °C, 40% RH) in a mixture of water, bran meal, alfalfa, yeast, and powdered milk. Rear *Drosophila simulans* larvae in a rearing chamber and feed them *ad libitum* on a mixture of rotting fruit.
  3. Aspirate adult *Musca domestica* into experimental chambers within 1 week of emergence. Adult *Drosophila simulans* can be directly reared in the experimental chambers. Prior to data collection, make sure each experimental chamber contains no more than 10-15 individual *Musca domestica* or 20-30 individual *Drosophila simulans*.
  4. Feed adult *Musca domestica ad libitum* on a mixture of sugar and low-fat dried milk, with free access to water. Feed adult *Drosophila simulans ad libitum* on a mixture of rotting fruit. Replace food weekly.
  5. Maintain experimental chambers on a 16:8 hr light:dark [LD] cycle, 20.5-22 °C and 30-50% RH for the duration of the experiment.

### 2. Record Flying Sounds in Experimental Chambers

1. Experimental Chamber Setup
 

Note: An "experimental chamber" denotes the cage designed in our lab, in which the data was recorded. The sensor is fairly inexpensive. When built in bulk, a set up could be manufactured for less than \$10.

  1. Construct an experimental chamber, either of the larger size: 67 cm L x 22 cm W x 24.75 cm H, or the smaller size: 30 cm L x 20 cm W x 20 cm H. The experimental chamber consists of a phototransistor array and a laser line pointing at the phototransistor array. NOTE: Additionally, the chamber consists of Kritter Keepers that are modified to include the sensor apparatus as well as a sleeve attached to a piece of PVC piping to allow access to the insects.
  2. Connect the phototransistor array to an electronic board. The output of the electronic board feeds into a digital sound recorder and is recorded as audio data in the MP3 format. See the logic design of the sensor in **Figure 1.II** and a physical version of the chamber in **Figure 1.I**.
  3. Modify the lids of the experimental chambers with a piece of mesh cloth affixed to the inside in order to prevent escape of the insects. NOTE: When an insect flies across the laser beam, its wings partially occlude the light, causing small light fluctuations. The light fluctuations are captured by the phototransistor array as changes in current, and the signal is filtered and amplified by the custom designed electronic board.
2. Set Up the System to Record Sounds Produced by Flying Insects
  1. Connect the experimental chamber to a power supply. Turn on the power.
  2. On the experimental chamber, find the laser lights and photoarray. Align the laser lights to the photoarray. To achieve the proper alignment, adjust the photoarray using the magnets on the outside of the experimental chamber which correspond to the magnets attached to the photoarray on the interior of the chamber until the laser light is centered on all the individual photodiodes.
  3. Perform two sanity checks to make sure the system is properly set up.
 

Note: The first step is to make sure that the system is powered, all wires are properly connected and the laser is pointing at the photo array. The second step is to conduct further checks on the laser alignment to ensure it can capture the sound of the insects' wingbeats.

    1. Plug headphones (rather than the recorder) into the audio jack. Plunge hand in and out of the experimental chamber, near the photoarray, to break the plane of the laser light. Make sure the laser light is on (it will be a red beam of light) and that you break the plane of light a few times with your hand. Listen for changes in noise level as your hand goes in and out of the light beam. If you detect an audible difference, the sensor is able to capture sounds produced by the movement of big objects. If successful, move on to the next sanity check, otherwise, check if the headphone is properly connected and whether the laser is pointing at the photoarray. Adjust the photoarray accordingly until the sound of the hand moving in and out of the experimental chamber can be heard.
    2. Attach a thin piece of electrical wiring to an automatic toothbrush. Turn on the toothbrush, and plunge the wiring in and out of the experimental chamber close to the phototarray. Make sure the laser light is hitting the piece of wiring as it moves. If you detect an audible change in frequency, when the piece of wiring breaks the plane of laser light, the system is then ready to capture the

sounds produced by the movement of tiny objects, *i.e.*, insect sounds. If you do not detect an audible difference, go back to step 2.2.2 to re-align the laser lights and the photoarray.

4. After the system is properly set up, close the lid and add the insects.
3. Data Collection: Record Sounds Produced by Flying Insects
  1. Turn on the recorder and make a voice annotation that includes the following information: name of the species in the experimental chamber, age of the insects, date and time, current ambient RT, and relative humidity. Pause the recording.
  2. Connect the recorder to the system, via the audio cable, and resume the recording. Leave the recorder to record for 3 days, then stop the recording.
  3. Download the data from the recorder into a new folder on a PC. Empty the recorder by deleting the data.
  4. Repeat the above recording process, until the remaining insects have died off and there are no more than 5 insects left alive in the cage.

### 3. Sensor Data Processing and Detection of Sounds Produced by Flying Insects

1. Use Software to Detect Sounds Produced by Flying Insects.
 

Note: The software (detection algorithm) is much faster than real-time. It takes less than 3 hr to process a recording session, *i.e.*, three days of data, on a standard machine with Intel(R) Core™ CPU at 2.00 GHz and 8 GB RAM.

  1. For each folder containing data from a recording session, run the detection software to detect insect sounds. To run the software, open MATLAB, and type "circadian\_wbf (*dataDir*)" in the command window, where *dataDir* is the directory of the recording data. Then press "Enter" to start.
 

NOTE: Download the detection software *circadian\_wbf* from reference #16.
  2. Wait until the algorithm terminates, then check the detection results. The algorithm outputs all the detected insect sounds in a new folder named "*dataDir\_extf*", where *dataDir* is the same as in the previous step. Each sound file is a 1 sec long audio clip originally extracted from the raw recording, with a digital filter applied to remove noise. The occurrence time of each detected sound is saved in a file named "*dataDir\_time.mat*". Observe the example of a detected insect sound in **Figure 2**.
2. Detection Algorithm
  1. Use a 0.1 sec long sliding window to slide through the recording. The sliding window starts from the beginning of the recording. For each window, follow the steps below.
    1. Compute the fundamental frequency of the current window.
    2. If the fundamental frequency is within the range of 100 Hz to 1,200 Hz, then do the following:
      1. Extract the 1-sec long audio clip centering at the current window from the recording; apply a digital filter to remove the noise in the clip and save the filtered audio into the folder "*dataDir\_extf*".
      2. Save the occurrence time of the current window into the file "*dataDir\_time*".
      3. Move the sliding window to the point that immediately follows the extracted audio clip.
    3. If the fundamental frequency is NOT within the range of 100 Hz to 1,200 Hz, simply move the sliding window 0.01 sec forward.
  2. Repeat the process until the sliding window reaches the end of the recording.

### 4. Insect Classification

1. Bayesian Classification Using Just the Flying Sound
 

Note: Bayesian classifier is a probabilistic classifier that classifies an object to its most probable class.

  1. Sound Feature Computation
    1. For each insect sound, compute the frequency spectrum of the sound using the Discrete Fourier Transform (DFT). Truncate the frequency spectrum to include only those data points corresponding to the frequency range: 100 Hz to 2,000 Hz. The truncated frequency spectrum is then used in the classification as the "representative" of the insect sound.
 

NOTE: The DFT is an algorithm that transforms signals in time domain to the frequency domain. It is a built-in function in most programming libraries, and can be called in the program with just one line of code.
  2. Train a Bayesian classifier
    1. Use the kNN density estimation approach<sup>14</sup> to learn the posterior probability distribution using the sound feature. With the kNN approach, the training phase is to build a training dataset.
      1. Randomly sample a number of insect sounds from the data collected for each species of insects.
      2. Follow the steps in Section 4.1.1 and compute the truncated frequency spectrum for each sampled sound. The truncated spectrums together with the samples' class labels (insect species name) composed the training dataset.
  3. Use the Bayesian classifier to classify an unknown insect
    1. Compute the truncated frequency spectrum of the unknown insect sound.
    2. Compute the Euclidean distance between the truncated spectrum of the unknown object and all the truncated spectrums in the training dataset.
    3. Find the top *k* (*k* = 8 in this paper) nearest neighbors of the unknown object in the training dataset. Compute the posterior probability of the unknown insect sound belonging to a class as the fraction of the top *k* nearest neighbors which are labeled as class .

4. Classify the unknown object to the class that has the highest posterior probability.
2. Add a Feature to the Classifier: Insect Circadian Rhythm
    1. Learn the class-conditioned distributions of the occurrence time of insect sound, that is, the circadian rhythm for each species of insects.
      1. Obtain the occurrence time of each sound from the detection results (c.f. Section 3.2).
      2. For each species, build a histogram of the insect sound occurrence time.
      3. Normalize the histogram so that the area of the histogram is one. The normalized histogram is the circadian rhythm of the given species. It tells the probability of observing an insect, of that species, in flight within a certain time period.
    2. Classify an unknown "insect sound" by combining the "insect sound" and the circadian rhythm
      1. Given the specific point in time in which the unknown insect sound occurred, obtain the probability of observing an insect of class based on the circadian rhythm of class.  
NOTE: The circadian rhythm is a probability distribution. It is an array specifying the probability of detecting a sound, produced by a specific species of insects, at a specific time of day. So once a time is given, one can simply check the array to get the probability.
      2. Follow the steps in section 4.1.2 to compute the posterior probability that the unknown sound belongs to class using the sound features. Multiply the posterior probability to the results from the previous step to get the new posterior probability.
      3. Classify the "unknown sound" to the class that has the highest new posterior probability.
  3. Add One More Feature to the Classifier: Insect Geographic Distribution
    1. Learn the geographic distribution of the species of interest, either from data published in historical records, relevant literature, or simply gather the first-hand knowledge from field technicians/biologists. For demonstration purposes, use a simulation of the graphic distribution, as shown in **Figure 7**.
    2. Classify an "unknown insect sound" using "flying sound" and the two additional features.
      1. Given the geographic location where the insect sound was intercepted, compute the probability of observing an insect from class at that specific location using the graphic distribution of species.
        1. Follow steps in section 4.2.2 and compute the posterior probability that the "unknown sound" belongs to class using the sound features and the circadian rhythms. Multiply the outcome of this step to the results from the previous step, in order to get the new posterior probability.
      2. Classify the "unknown sound" to the class that has the highest new posterior probability.
  4. A General Framework for Adding Features
    1. Consider the Bayesian classifier that uses just sound features as the primary classifier. Follow the steps below to add new features to the classifier.
      1. In the training phase, learn the class-conditioned density functions of the new feature.
      2. In the classification phase, given the new feature of the "unknown sound", compute the probability of observing the feature in class using the density functions learned in the previous step. Multiply the new probability to the previous posterior probability corresponding to the "unknown sound" belonging to class which were computed based on just the odd features, to obtain the new posterior probability. Classify the unknown object to the class that has highest new posterior probability.

## Representative Results

Two experiments are presented here. For both experiments, the data used were randomly sampled from a dataset that contains over 100,000 objects.

The first experiment shows the ability of the proposed classifier to accurately classify different species/sexes of insects. As the classification accuracy depends on the insects to be classified, a single absolute value for classification accuracy will not give the reader a good intuition about the performance of the system. Instead, rather than reporting the classifier's accuracy on a fixed set of insects, the classifier was applied to datasets with an incrementally increasing number of species, and therefore increasing classification difficulty.

The dataset began with just 2 species of insects; then at each step, one more species (or a single sex of a sexually dimorphic species) was added and the classifier was used to classify the increased number of species (the new dataset). A total of ten classes of insects (different sexes from the same species counting as different classes) were considered, with 5,000 exemplars in each class.


The classifier used both *insect-sound* (frequency spectrum) and *time-of-intercept* for classification. **Table 1** shows the classification accuracy measured at each step and the relevant class added at that step.

According to **Table 1**, the classifier achieves more than 96% accuracy when classifying no more than 5 species of insects, significantly higher than the default rate of 20% accuracy. Even when the number of classes considered increases to 10, the classification accuracy is never lower than 79%, again significantly higher than the default rate of 10%. Note that the ten classes are not easy to separate, even by human inspection. Among the ten species, eight of them are mosquitoes, with six of them being from the same genus.

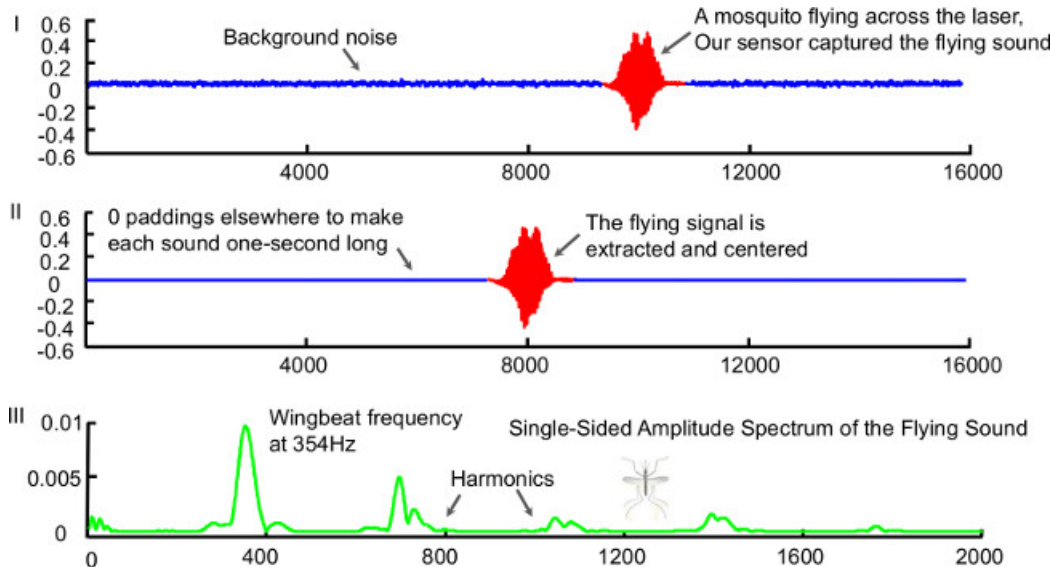
The second experiment is to show how accurately the system can sex flying insects, specifically, to distinguish male *Ae. aegypti* mosquitoes from females. For the first part of the experiment, assume that the "misclassification cost" of misclassifying males as females is the same as the cost

of misclassifying females as males. With this assumption, the classification results are shown in **Table 2.I**. The classification accuracy to sex *Ae. aegypti* is about 99.4%.

For the second part of experiment, assume the cost is not asymmetric, *i.e.*, misclassification of females as males is much more costly than the reverse. With this assumption, the decision threshold of the classifier was changed to reduce the number of high-cost misclassifications. With the threshold properly adjusted, the classification results in **Table 2.II** were achieved. Of 2,000 insects in the experiment, twenty-two males, and zero females were misclassified.



**Figure 1.** (I) One of the experimental cages used to gather the data. (II) A logical version of the sensor setup with the components annotated. [Please click here to view a larger version of this figure.](#)



**Figure 2.** (I) An example of a 1 sec audio clip containing a “flying insect sound” generated by the sensor. The sound was produced by a female *Cx. stigmatosoma*. The insect sound is highlighted in red/bold. (II) The “insect sound”, which has been cleaned and saved into a 1-sec long audio clip by centering the insect signal and padding with 0s elsewhere. (III) The frequency spectrum of the insect sound, obtained using Discrete Fourier Transform. [Please click here to view a larger version of this figure.](#)

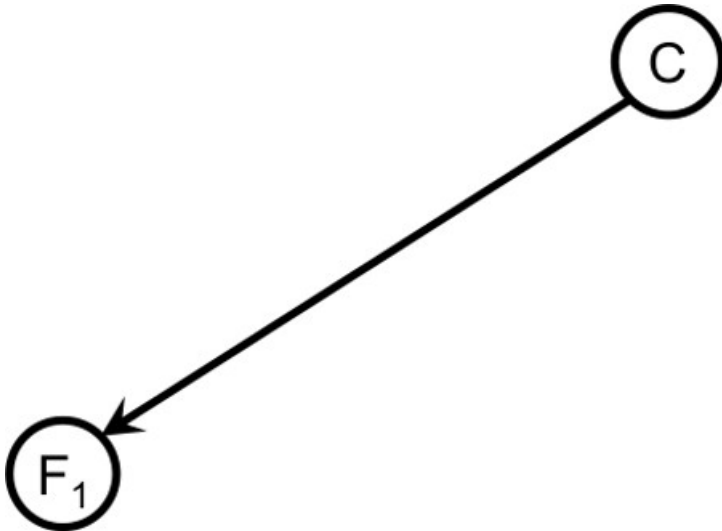


Figure 3. A Bayesian network that uses a single feature for classification.




Figure 4. A Bayesian network that uses two independent features for classification




Figure 5. The circadian rhythms of *Cx. stigmatosoma* (female), *Cx. tarsalis* (male), and *Ae. aegypti* (female), learned based on observations generated by the sensor that were collected over 1 month duration. [Please click here to view a larger version of this figure.](#)




Figure 6. A Bayesian network that uses three independent features for classification.

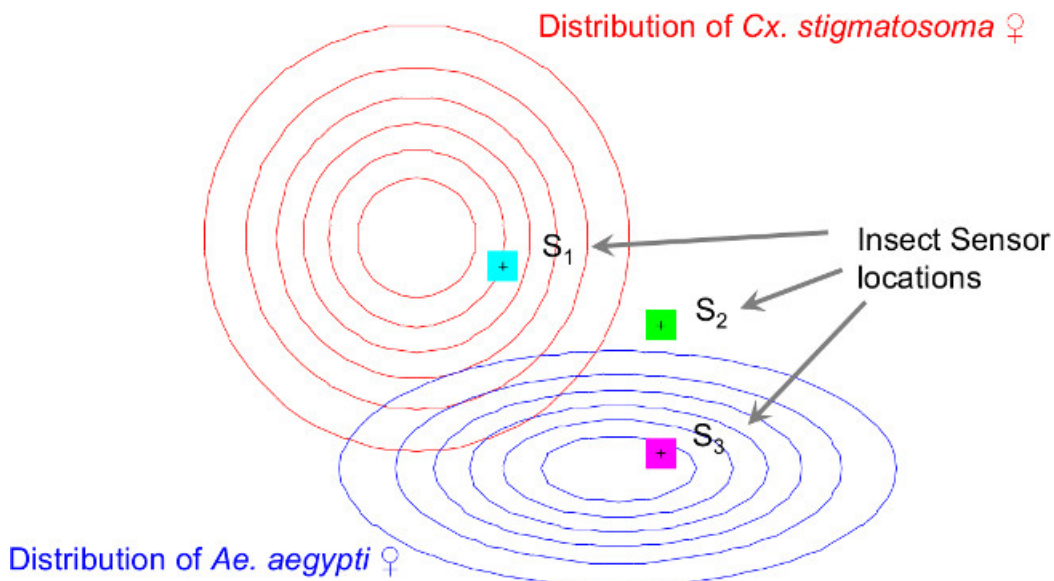


Figure 7. The assumptions of the geographic distributions of each insect species and sensor locations in the simulation to demonstrate the effectiveness of using location-of-intercept feature in classification. [Please click here to view a larger version of this figure.](#)




Figure 8. The general Bayesian network that uses  $n$  features for classification, where  $n$  is a positive integer. [Please click here to view a larger version of this figure.](#)

Step	Species Added	Classification Accuracy	Step	Species Added	Classification Accuracy
1	<i>Ae. aegypti</i> ♂	N/A	6	<i>Cx. quinquefasciatus</i> ♂	92.69%
2	<i>Musca domestica</i>	98.99%	7	<i>Cx. stigmatosoma</i> ♀	89.66%
3	<i>Ae. aegypti</i> ♀	98.27%	8	<i>Cx. tarsalis</i> ♂	83.54%
4	<i>Cx. stigmatosoma</i> ♂	97.31%	9	<i>Cx. quinquefasciatus</i> ♀	81.04%
5	<i>Cx. tarsalis</i> ♀	96.10%	10	<i>Drosophila simulans</i>	79.44%

Table 1. Classification accuracy with increasing number of classes.

		Predicted class				Predicted class	
I (Symmetric cost)		female	male	II (Asymmetric cost)		female	male
Actual class	female	993	7	Actual class	female	1,000	0
	male	5	995		male	22	978

Table 2. (I) The confusion matrix for sex discrimination of *Ae. aegypti* mosquitoes with the decision threshold for females is set at 0.5 (i.e., same cost assumption). (II) The confusion matrix of sex the same mosquitoes, with the decision threshold for females set at 0.1.

## Discussion

The sensor/classification framework described here allows the inexpensive and scalable classification of flying insects. The accuracies achievable by the system are good enough to allow the development of commercial products and could be a useful tool in entomological research.

The ability to use inexpensive, noninvasive sensors to accurately and automatically classify flying insects would have significant implications for entomological research. For example, by deploying the system in the field to count and classify insect vectors, the system can provide real-time counts of the target species, producing real-time information that can be used to plan intervention/suppression programs to combat malaria. Moreover, the system can automatically separate insects by sex, and thus it can be used to free entomologists working on the Sterile Insect Technique<sup>15</sup> from the tedious and time-consuming task of manually sexing the insects.

In using this system, the most critical step is to properly set up the sensor for data collection. If the laser and the photo array are not properly aligned, the data will be very noisy. After the insects are placed in the cage, the photo array should always be fine-tuned using the magnets on the outside of the cage. Note that flashing lights, camera flashes and vibrations near the cages will introduce noise to the data. Therefore, to obtain clean data, place the cage in a dark room, and wherever necessary, place dry towels under the cages in order to diminish the level of vibration.

The classifier presented in this work used just two additional features. However, there may be dozens of additional features that could help improve the classification performance. As the potential features are domain and application specific, the users could choose features based



on their specific need or applications. The general framework of the classifier allows users to easily add features to the classifier to improve the classification performance.

To encourage the adoption and extension of our ideas, we are making all code, data, and sensor schematics freely available at the UCR Computational Entomology Page<sup>16-17</sup>. Moreover, within the limits of our budget, we will continue our practice of giving a complete system (as shown in **Figure 1**) to any research entomologist who requests one.

## Disclosures

The authors declare that they have no competing financial interests.

## Acknowledgements

We would like to thank the Vodafone Americas Foundation, the Bill and Melinda Gates Foundation, and the São Paulo Research Foundation (FAPESP) for funding this research. We would also like to thank the many faculty members from the Department of Entomology at University of California, Riverside, for their advice on this project.

## References

1. Kahn, M. C., Celestin, W., Offenhauser, W. Recording of sounds produced by certain disease-carrying mosquitoes. *Science*. **101** (2622), 335-336 (1945).
2. Reed, S. C., Williams, C. M., Chadwick, L. E. Frequency of wing-beat as a character for separating species races and geographic varieties of *Drosophila*. *Genetics*. **27** (3), 349 (1942).
3. Belton, P., Costello, R. A. Flight sounds of the females of some mosquitoes of Western Canada. *Entomologia experimentalis et applicata*. **26** (1), 105-114 (1979).
4. Mankin, R. W., Machan, R., Jones, R. Field testing of a prototype acoustic device for detection of Mediterranean fruit flies flying into a trap. In *Proc. 7th Int. Symp. Fruit Flies of Economic Importance*. pp. 10-15 (2006).
5. Raman, D. R., Gerhardt, R. R., Wilkerson, J. B. Detecting insect flight sounds in the field: Implications for acoustical counting of mosquitoes. *Transactions of the ASABE*. **50** (4), 1481(2007).
6. Moore, A., Miller, J. R., Tabashnik, B. E., Gage, S. H. Automated identification of flying insects by analysis of wingbeat frequencies. *Journal of economic entomology*. **79** (6), 1703-1706. (1986).
7. Moore, A., Miller, R. H. Automated identification of optically sensed aphid (Homoptera: Aphidae) wingbeat waveforms. *Annals of the Entomological Society of America*. **95** (1), 1-8 (2002).
8. Li, Z., Zhou, Z., Shen, Z., & Yao, Q. Automated identification of mosquito (diptera: Culicidae) wingbeat waveform by artificial neural network. In *Artificial Intelligence Applications and Innovations*. pp. 483-489 (2005).
9. Moore, A. Artificial neural network trained to identify mosquitoes in flight. *Journal of insect behavior*. **4** (3), 391-396 (1991).
10. Halevy, A., Norvig, P., Pereira, F. The unreasonable effectiveness of data. *Intelligent Systems, IEEE*. **24** (2), 8-12 (2009).
11. Banko, M., Brill, E. Mitigating the paucity-of-data problem: exploring the effect of training corpus size on classifier performance for natural language processing. In *Proceedings of the first international conference on Human language technology research*. Association for Computational Linguistics. pp. 1-5 (2001).
12. Shotton, J. *et al.* Real-time human pose recognition in parts from single depth images. *Communications of the ACM*., **56** (1), 116-124 (2013).
13. Ilyas, P. Classifying insects on the fly. *Ecological Informatics*. (2013).
14. Mack, Y. P., Rosenblatt, M. Multivariate k-nearest neighbor density estimates *Journal of Multivariate Analysis*., **9** (1), 1-15 (1979).
15. Benedict, M., Robinson, A. The first releases of transgenic mosquitoes: an argument for the sterile insect technique. *TRENDS in Parasitology*. **19** (8), 349-355 (2003).
16. Chen, Y. *Supporting Materials*. <https://sites.google.com/site/insectclassification/> (2013).
17. Chen, Y., Why, A., Batista, G., Mafra-Neto, A., Keogh, E. Flying Insect Classification with Inexpensive Sensors. *Journal of Insect Behavior*.**27**(5), 657-677 (2014).