Title
IMPROVING ACOUSTIC INSECT DETECTION USING SIGNAL PROCESSING METHODS

Author
Raju, Saachi

Publication Date
2026-06-15

IMPROVING ACOUSTIC INSECT DETECTION USING SIGNAL PROCESSING METHODS

By

Saachi Raju

A capstone project submitted for Graduation with University Honors

May 07, 2026

University Honors

University of California, Riverside

APPROVED

Dr. Eamonn Keogh

Department of Computer Science

Dr. Begoña Echeverria, Howard H Hays, Jr. Chair

University Honors

ABSTRACT

Insect monitoring is important for agriculture, biodiversity research, and public health, but many traditional monitoring methods can be expensive, labor-intensive, destructive, or difficult to apply over large areas. Acoustic insect monitoring offers a promising alternative because insects can be detected by analyzing the unique sounds produced by their wingbeats. However, real-world audio recordings collected by sensors often include background noise, environmental interference, and other non-insect sounds, making it difficult to accurately identify true insect activity. The goal of this project was to optimize acoustic insect sensor performance by developing a MATLAB-based workflow that can process insect audio recordings and identify sound segments most likely to contain insect wingbeat signals. The workflow incorporated frequency-domain signal processing techniques including Fast Fourier Transform (FFT) analysis, dominant-frequency extraction, spectral filtering, area-under-the-curve (AUC) scoring, and sliding-window subsequence detection to isolate insect-like acoustic patterns from noisy recordings. Rather than relying only on raw recordings, this project focused on improving the preprocessing and signal-analysis steps that occur before insect classification. By making these steps more reliable, the sensor can better separate useful insect data from unwanted noise and produce more accurate wingbeat frequency information. This work provides a foundation for developing acoustic insect monitoring systems that are more accurate, scalable, and practical for real-world use. With further testing and development, this type of sensor-based system could support improved pest management, ecological monitoring, and public health efforts related to disease-carrying insects.

ACKNOWLEDGEMENTS

I would like to sincerely thank my faculty mentor, Dr. Eamonn Keogh, for his consistent support, guidance, and encouragement throughout this research project. His mentorship introduced me to how computer science research is conducted in practice and provided me with the opportunity to explore interdisciplinary research connecting the computer science and entomology departments at University of California, Riverside. I am very thankful for the opportunity to participate in this research experience, which greatly expanded both my technical skills and my understanding of academic research.

TABLE OF CONTENTS

BACKGROUND

Insects play major roles in ecological and public health, yet monitoring them at large spatial and temporal scales remains difficult. Large spatial scales refer to the monitoring of insects across many different places such as across farms, wetlands, neighborhoods, and other agricultural regions[3]. Large temporal scales refer to the monitoring of insects over long periods of time[3]. Such broad-scale monitoring is necessary because insect populations vary significantly across locations and can shift over time due to seasonal climate changes and human activity. Rapid changes in insect populations and diversity over time can have immense effects agriculturally and in the daily lives of humans.

One of the key acoustic features in flying insects is the wingbeat frequency (WBF). This feature can often be used to classify insects, down to the level of sex and species. For example, male mosquitoes typically have a WBF of 600 Hz or greater, but the larger females typically have a WBF of 450 Hz or less. Extracting the true wingbeat frequency can be challenging in real recordings as insect signals may be weak, intermittent, or masked by environmental noise. Low-frequency background noise, channel inconsistencies in stereo recordings, and non-insect sounds can all lead to false frequency detection if the entire recording is analyzed without additional filtering and quality control. These challenges make the preprocessing step integral in acoustic insect monitoring.

Traditional insect sampling approaches can be labor-intensive, destructive, and difficult to standardize across agricultural environments. Some of these approaches include glue traps and camera-based systems which have several limitations. Glue traps require intense manual labor and expertise in insect identification for accurate classification. Camera-based systems are often expensive, consume significant amounts of energy, and generate large amounts of image data that can be a challenge to store and process as they require complex image analysis algorithms for automated species recognition.

Bioacoustics, the study of sound production and reception in animals, including insects, provides a useful framework for wildlife monitoring. Historically, much of the research in bioacoustics has focused on vertebrates such as birds, frogs, bats and other mammals. In recent years, however, insect sound detection research has become increasingly popular due to the profound importance of insects to the overall ecosystem. Insect monitoring is necessary for a wide variety of reasons such as disease control, tracking rates of population change, and pest management[3]. Recent studies in acoustic insect monitoring agree that sound-based methods are a promising non-destructive way to study biodiversity, monitor pests, and detect disease-carrying insects[6]. The development of effective insect monitoring systems is crucial to help farmers with a more effective approach to integrated pest management compared to traditional limited insect sampling approaches.

A recent literature review of acoustic systems for insect monitoring emphasizes that successful monitoring depends not only on the final classification algorithm, but also on the full signal-processing pipeline, including recording quality, preprocessing, feature extraction, and evaluation procedures[1]. In particular, the review highlights the need for more standardized and consistent methods for handling and analyzing acoustic data from the moment it is recorded to the final stage of insect detection[1]. The review highlights that monitoring performance can vary significantly depending on factors such as microphone technology, environmental noise, and signal-processing choices.

Another key foundational study, by Chen et al., highlights the importance of developing automatic insect classification systems that are inexpensive and easy to use, particularly as an alternative to traditional mechanical trapping methods currently used for insect monitoring[2]. This study noted that traditional microphone-based systems were highly sensitive to distance, wind interference, and surrounding environmental noise, which often reduced the quality and consistency of recorded insect signals. To address these issues, they developed an inexpensive pseudo-acoustic optical sensor that measures fluctuations caused by wing motion when an insect flies across a laser beam[2]. Their system analyzed signals using short sliding windows, allowing recordings to be divided into smaller time segments so that transient wingbeat activity could be isolated and evaluated continuously over time. Rather than relying solely on wingbeat frequency, their approach also incorporated additional signal characteristics such as harmonic structure and temporal variation, along with behavioral information including circadian activity patterns. Using Bayesian classification methods, they showed that reliable insect recognition could be achieved with relatively simple models when informative features, localized time-window analysis, and suitable sensing hardware were used.

In this study, we develop a MATLAB-based workflow for identifying the dominant wingbeat frequency in insect recordings, with initial development focused on mosquito audio data. Building on the broader need for standardized acoustic preprocessing, our approach combines fast Fourier transform (FFT)-based frequency estimation with batch processing, frequency thresholding, stereo-channel consistency checks, and a sliding-window procedure designed to isolate the cleanest insect-like portion of each recording. By focusing on the local signal peak that is most likely to reflect only insect signal, rather than the global peak which may be contaminated by spurious sound, this method aims to improve the reliability of frequency extraction from recordings that contain substantial background noise or other junk signals.

By isolating the portion of the recording most likely to contain authentic insect activity, the workflow seeks to improve frequency extraction even when recordings contain substantial junk data or environmental noise. The ultimate goal is to improve the overall accuracy of insect detection by developing a tool that can accurately distinguish true insect regions in audio samples from non-insect sounds such as music, background noise, human activity, and other environmental interference. This allows insect signals to be identified more effectively in challenging acoustic conditions where conventional full-recording FFT methods may fail. In addition to supporting more accurate wingbeat estimation, the proposed framework provides a scalable preprocessing strategy for large collections of recordings and may serve as a useful front-end step for future automated insect classification and monitoring systems.

METHODS

Project Setup

For this project, all analyses were conducted using MATLAB, where custom scripts were written to process and analyze large sets of insect wingbeat audio recordings of various species. The tools developed in this study were designed to be applicable to acoustic recordings of many insect species; however, the initial development and testing of the methods focused on recordings of the Gambie male mosquito. This dataset consisted of .wav audio files containing recordings of mosquito flight sounds captured by microphone sensors. Each file represents a short time series of the acoustic waveform produced by a species of insect.

Main Frequency Detection

The primary tool developed was the MATLAB function computeFrequencies, which was designed to detect the dominant frequency peak within the 100-1000 Hz range, a frequency band typically relevant to mosquito wingbeat activity. Each sound file was loaded into MATLAB and returned both the audio waveform and sampling rate. The waveform represents a time series of amplitude values captured by the microphone. This function served as the primary frequency extraction tool and provided the main frequency outputs used by additional functions built throughout the project. The general approach involved converting these audio signals into the frequency domain, identifying dominant frequency peaks corresponding to insect wingbeats, and implementing additional signal-processing techniques to improve the reliability of frequency detection across different insect recordings.

The first step in the execution of this function involved loading each .wav file into MATLAB and extracting the audio waveform and sampling rate. Since audio signals are stored as sequences of amplitude values measured over time, the waveform represents variations captured by the microphone during the recording. Once the waveform was loaded, it was converted from the time domain to the frequency domain using the Fast Fourier Transform (FFT) in MATLAB[4]. This transformation allowed the signal to be analyzed in terms of its frequency components, which is necessary for identifying the dominant frequency produced by insect wingbeats. The FFT generates a frequency spectrum that shows the magnitude of each frequency present in the signal. By examining this spectrum, the frequency associated with the highest magnitude peak within the selected 100-1000 Hz range was identified as the dominant frequency. This peak represents the strongest repeating component in the audio recording and is interpreted as the most likely wingbeat frequency of the insect.

Batch Processing

After implementing the algorithm to output the main frequency, an automated batch-processing script was developed so that the analysis could be applied to entire folders of recordings rather than processing files individually. This script iterated through each .wav file in a specified directory, applied the frequency detection function, and stored the resulting dominant frequency values. In addition to outputting the main frequency for each recording, the script could optionally generate frequency magnitude plots for individual files. These plots display the frequency spectrum and highlight the largest peak detected in the signal. This visualization allows for manual inspection of the results and helps verify that the algorithm was correctly identifying insect wingbeat frequencies rather than noise or unrelated signals.

Frequency Detection Validation Using Known Audio Tones

To verify that the algorithm was accurately detecting frequencies, a validation test was conducted using audio signals with known frequencies. Test audio files were generated using an online piano tone generator at several known frequencies, including 100 Hz, 200 Hz, 400 Hz, 600 Hz, and 800 Hz. These recordings were processed using the same MATLAB function used to analyze the insect recordings. The results showed that the detected frequencies closely matched the expected values. For example, the algorithm detected approximately 99.81 Hz for the 100 Hz tone from the online piano generator and approximately 799.85 Hz for the 800 Hz tone. These results confirmed that the frequency detection function was functioning correctly and could reliably identify dominant frequencies within a variety of audio signals, even non-insect ones.

Histogram Analysis of Frequency Distributions

After confirming the correctness of the frequency extraction process, the dominant frequencies from the insect recordings were analyzed collectively to understand the distribution of wingbeat frequencies present in the dataset. A histogram was generated using the frequency values obtained from all processed audio files. In this histogram, the horizontal axis represented frequency in hertz, while the vertical axis represented the number of recordings containing frequencies within each range. In addition to visualizing the distribution, summary statistics such as the mean and standard deviation of the detected frequencies were calculated. This analysis provided a general overview of the frequency characteristics of the recordings and helped identify common frequency ranges associated with insect wingbeats.

Fine-Tuning for Detecting True Insect Peak

During early testing, some recordings revealed strong low-frequency peaks that did not correspond to insect wingbeats. These peaks were likely caused by environmental noise sources such as wind, human movement, or other natural sounds. To address this issue, a frequency filtering step was introduced to remove frequencies below a specified threshold. For example, for the male mosquito, a minimum frequency cutoff of 100 Hz was applied by suppressing all spectral magnitudes below this value before searching for the dominant frequency peak. By eliminating very low frequencies from consideration, the algorithm focused on the frequency ranges that were most likely to contain insect wingbeat signals. This modification improved the accuracy of peak detection and reduced the likelihood of false detections caused by any background noise or junk data.

Stereo Consistency Checking

An additional quality-control tool called StereoCheck was developed to identify potentially unreliable recordings before further analysis was performed. Many audio recordings in the dataset were stored as stereo signals, meaning that the sound was recorded simultaneously through separate left and right audio channels. In a valid insect recording, both channels should capture approximately the same dominant wingbeat frequency because they originate from the same sound source. Large differences between the channels may indicate recording corruption, microphone imbalance, environmental interference, or strong background noise.

The StereoCheck tool automatically processed folders of wav audio files and examined each stereo recording independently. For every file, the left and right audio channels were separated and analyzed individually using the previously developed computeFrequencies function. This function computed the dominant frequency peak for each channel within a constrained insect frequency band. The tool then compared the dominant frequencies detected in the left and right channels and evaluated whether both frequencies fell within an acceptable mosquito wingbeat range.

Recordings were considered valid only if both channels produced frequencies within the expected range of approximately 100–1000 Hz and if the difference between the two detected frequencies remained below a predefined threshold. Files that failed these checks were automatically flagged for review. Several failure conditions were considered, including frequencies below the expected insect range, frequencies above the expected range, invalid frequency outputs, or large differences between the left and right channels.

The tool also generated diagnostic plots for flagged recordings. These plots displayed the left and right channel waveforms alongside their corresponding frequency waveforms, allowing visual inspection of inconsistencies between channels. This process made it possible to identify recordings dominated by noise, corrupted data, or abnormal low-frequency movement that could negatively affect later stages of analysis.

The stereo consistency check served as an important preprocessing and validation step within the overall workflow. By removing recordings with inconsistent or unrealistic frequency behavior, the quality of the dataset was improved before applying the main area-under-the-curve and sliding-window insect detection methods discussed below. This helped ensure that later analyses were performed primarily on recordings likely to contain meaningful insect wingbeat information rather than artifacts introduced by noise or recording errors.

Area-Under-the-Curve Insect Signal Detection

To support the ultimate goal of this project, a MATLAB tool called areaUnderCurveTool was developed to help identify the truest insect sequence within recordings that may contain noise or non-insect data. The motivation for this approach was to affirm that insect wingbeat recordings generally output small area-under-the curve values because they produce a peak in the waveform where the energy is concentrated at a specific region. Contrastingly, music or other non-insect sounds generally output larger area-under-the-curve (AUC) values because they contain a broader spread of energy across many frequencies.

The area-under-the-curve function computed the distribution of spectral energy within the selected frequency band of 100-1000 Hz for mosquito recordings. After loading the sound file onto MATLAB, the signal was transformed into the frequency domain using FFT. This produced the same frequency and magnitude spectrum that was used in the computeFrequencies tool. Then, the area-under-the-curve function analyzed the full spectrum within the 100-1000 band, and computed the area-under-the-curve for the signal by summing all magnitude values across that range.

The output for this function demonstrated how concentrated or spread out the energy for a given signal was. Insect wingbeat sounds were expected to produce generally small area-under-the-curve values because most of their energy is concentrated around one dominant frequency. Other non-insect data, such as music, contains broadband frequencies, which causes energy to be distributed across a wider range of the spectrum. This, in turn, produces larger area-under-the-curve values for music or non-insect signals. This tool is important because it provides a means to distinguish insect-like segments injected into noisy recordings from other non-insect background noise. This makes it possible to isolate the most insect-like sequences of data for the most accurate insect detection.

Sliding Window Signal Processing

The main goal of this project was to develop a tool capable of accurately identifying the most representative insect signal within an audio recording that may contain significant background noise or irrelevant acoustic content. In many real-world recordings, insect wingbeat signals are not uniformly present across the entire duration. Instead, they tend to appear intermittently in short, localized segments, while the remainder of the recording may include environmental noise, silence, human movement, or even unrelated sounds such as music. Performing analysis over the entire signal can therefore lead to inaccurate frequency detection. To address this limitation, a sliding-window-based approach was implemented to isolate and analyze only the most insect-like portion of the signal.

The main function used for the implementation of the sliding window tool was FindBestInsectSound, which takes in the full audio waveform, the sampling frequency, a minimum and maximum frequency constraint, and the number of samples included in each window. For the initial testing, a longer audio sample of a Female Culex was used with a sampling frequency of 44,100 Hz. The recording contained 431,472 data points, corresponding to approximately 9.78 seconds of audio. The variables minF and maxF define the frequency range used during analysis. In this specific test, frequencies below 40 Hz and above 1000 Hz were set to zero, meaning that only the 40–1000 Hz range was analyzed. These values were selected for the Female Culex sample used in this specific task, but they can be modified for other datasets depending on the expected wingbeat frequency range of the insect being studied.

The variable SubsequenceLength defines the number of samples included in each sliding window. In this example, a subsequence length of 5000 samples was used, corresponding to a sliding window length of approximately 0.11 seconds of audio. Finally, the plot variable determines whether the function generates visual output showing the full signal, the scoring profile, and the selected best insect subsequence.

The method works by moving the fixed-length window across the full recording and calculating an AUC-based score for each possible segment. For each window, the signal is converted from the time domain to the frequency domain using the Fast Fourier Transform. Frequencies outside the selected range are removed, and the remaining frequency magnitudes are used to calculate an area-under-the-curve score. This score measures how much spectral energy is spread across the selected frequency band. Clean insect wingbeat signals are expected to have energy concentrated around a narrow dominant frequency, while noise or non-insect audio tends to have energy distributed more broadly across many frequencies. Therefore, lower AUC scores indicate segments that are more likely to contain a clean insect wingbeat signal.

After scoring each window, the function identifies the location with the best insect signal. The tool then extracts the subsequence around this location and displays it separately so that it can be visually inspected. The plotted output also shows the full audio signal, the scoring profile used to determine the best segment, and the selected subsequence. This makes the method interpretable because the user can see not only the final selected segment, but also how the algorithm chose it.

Overall, the sliding-window method improves frequency detection by allowing the program to focus only on the portion of the recording that most closely resembles a true insect wingbeat. Instead of allowing noise from the full recording to influence the final frequency estimate, this approach isolates the most insect-like subsequence first and then performs further frequency analysis on that cleaner segment. This makes the workflow more reliable for longer and noisier recordings, especially when insect activity is only present during parts of the audio file.

RESULTS

Fig 1: Dominant Peak Frequency Detection for Insect and Non-Insect Audio Data in the 100–1000 Hz Range

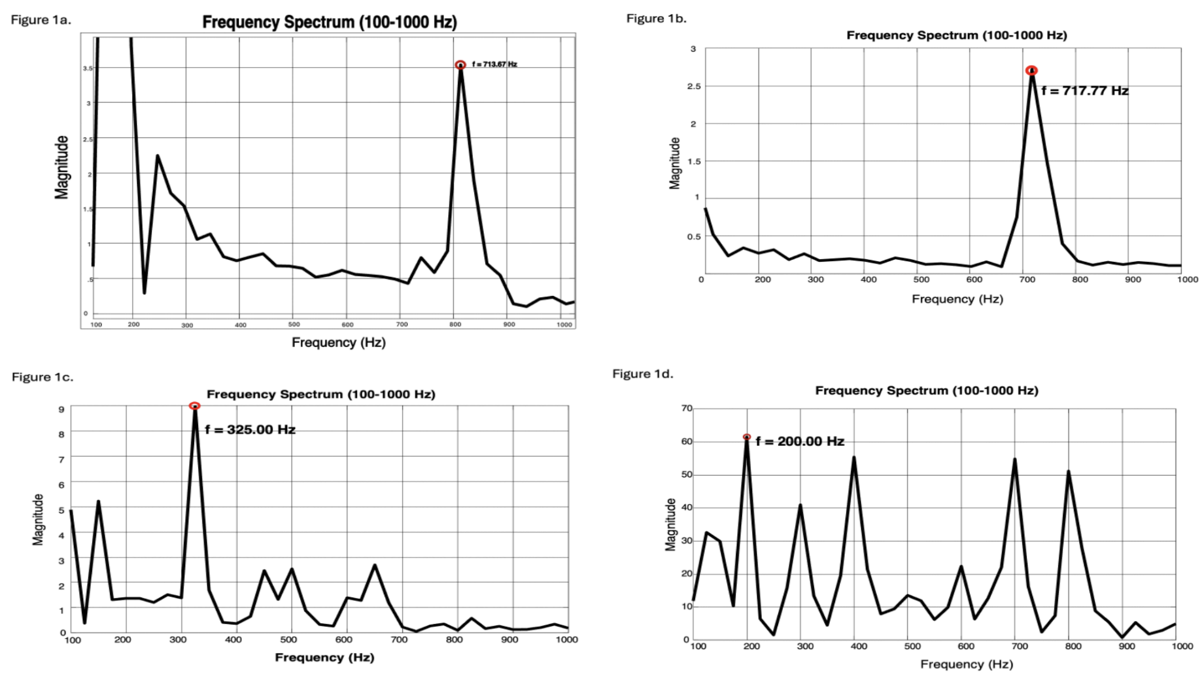

(a) *Gambie Male 0 000.wav* showing a detected dominant peak at 714.02 Hz. (b) *Gambie Male 0 001.wav* showing a detected dominant peak at 717.59 Hz. (c) Random pop music sample showing a detected dominant peak at 325.07 Hz. (d) Rock music sample showing a detected dominant peak at 199.98 Hz. Red markers indicate the dominant frequency peak identified in each spectrum.

Figure 1 demonstrates clear differences between insect and non-insect audio samples when analyzed using the FFT-based dominant frequency detection method. The two mosquito recordings, *Gambie Male 0 000.wav* (Fig. 1a) and *Gambie Male 0 001.wav* (Fig. 1b), produced dominant peaks at 714.02 Hz and 717.59 Hz, respectively. The close agreement between these values indicates strong consistency in wingbeat-frequency estimation across separate recordings of the same insect type, suggesting that the method can reliably detect characteristic mosquito flight frequencies. In contrast, the non-insect samples yielded substantially lower dominant frequencies, with the random pop music sample (Fig. 1c) peaking at 325.00 Hz and the rock music sample (Fig. 1d) peaking at 200.00 Hz. These lower values likely reflect the stronger presence of bass, rhythm, and low-frequency instrumental components that are common in music recordings, whereas mosquito wingbeats generate rapid periodic oscillations that produce higher-frequency acoustic signatures. The spectral plots also reveal qualitative differences in signal structure. The mosquito recordings display a sharper and more isolated dominant peak, indicating a stable repetitive source signal consistent with wingbeat motion. By comparison, the music samples contain multiple competing peaks spread across the spectrum, producing broader and a less distinct dominant peak due to overlapping vocals, instruments, and changing rhythmic content. These results suggest that insect recordings possess more consistent frequency signatures than many non-insect sounds, while also highlighting the importance of preprocessing methods that can distinguish true insect peaks from complex background audio.

Figure 2. Identification and correction of low-frequency contamination in mosquito wingbeat estimation (histogram)

Figure 2a.

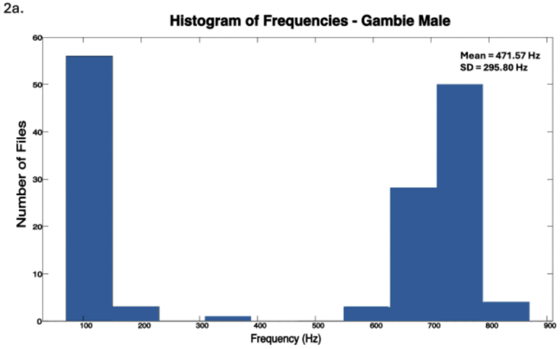

Figure 2b.

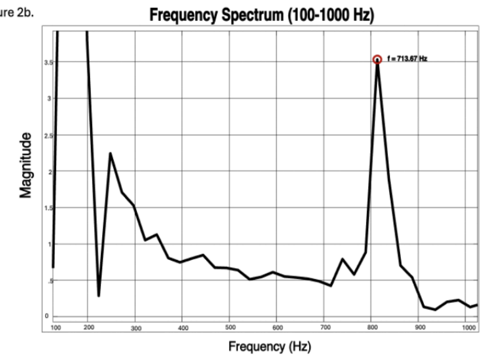

Figure 2c.

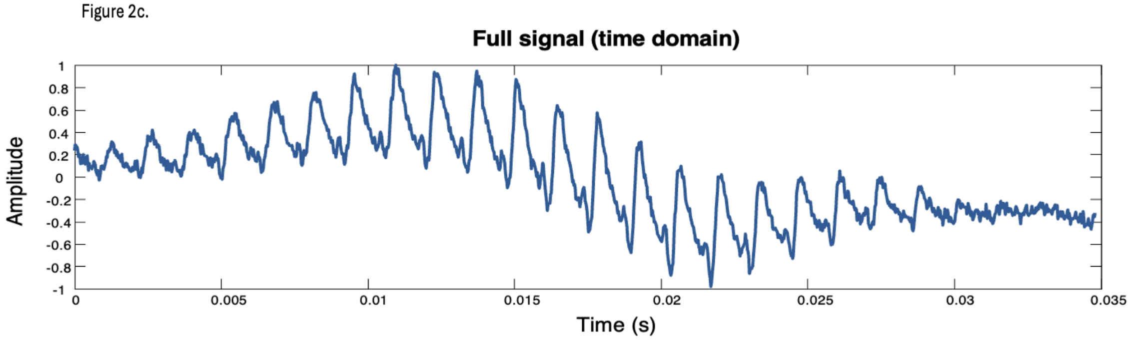

Figure 2d.

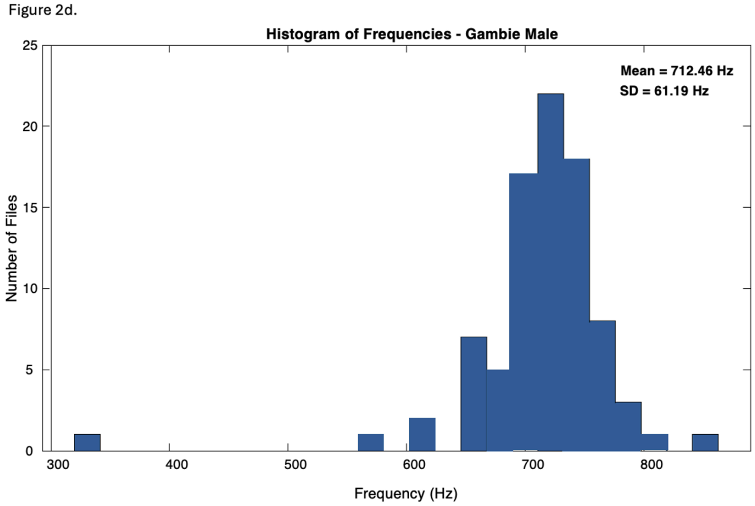

(a) Initial histogram of dominant frequencies for *Gambie Male* recordings showing a bimodal distribution, with an unexpected peak near 100 Hz.

(b) Frequency spectrum of a representative recording (*Gambie Male 0 000*), illustrating both the true insect peak near ~700 Hz and a spurious low-frequency peak near ~100 Hz.

(c) Corresponding time-domain signal, where the insect wingbeat appears as rapid, periodic oscillations, while slower, large-scale fluctuations contribute to low-frequency artifacts.

(d) Revised histogram after applying a domain-based frequency constraint (removing frequencies below 200 Hz), resulting in a unimodal, approximately Gaussian distribution centered near ~700 Hz, consistent with expected mosquito wingbeat frequencies.

Figure 2 illustrates the identification and correction of low-frequency contamination of non-insect data in the dominant frequency analysis of mosquito recordings. The initial histogram (Fig. 2a) exhibits a bimodal distribution, with one cluster centered around approximately 700 Hz and an unexpected secondary peak near 100 Hz. While the higher-frequency cluster is consistent with known mosquito wingbeat frequencies, the presence of the low-frequency peak suggests the inclusion of non-insect signal components in the analysis. To investigate the origin of this artifact, the frequency spectrum of a representative recording, Gambie Male 0 000.wav (Fig. 2b), was examined. This spectrum reveals both a strong dominant peak near ~700 Hz and a secondary peak near ~100 Hz, confirming that the low-frequency values observed in the histogram arise from features present within individual recordings.

` Further insight is provided by the corresponding time-domain signal (Fig. 2c), which highlights two distinct types of behavior within the waveform. The central portion of the signal, approximately between 0.005 s and 0.025 s, contains rapid, periodic oscillations that are characteristic of insect wingbeat motion and are responsible for the higher-frequency peak near 700 Hz. In contrast, slower, large-scale fluctuations present near the beginning and end of the signal contribute to low-frequency components around 100 Hz. These slow variations are not representative of insect activity but instead correspond to background noise, environmental disturbances, or recording artifacts.

Based on established biological constraints indicating that mosquito wingbeat frequencies typically exceed 200 Hz, a domain-specific filtering rule was introduced to exclude all frequencies below this threshold. After applying this constraint, the revised histogram (Fig. 2d) shows a unimodal, approximately Gaussian distribution centered near 700 Hz, with the low-frequency peak effectively removed. This transformation demonstrates that the original bimodal distribution was not indicative of multiple insect behaviors, but rather a mixture of valid insect signals and irrelevant low-frequency noise. By incorporating domain knowledge into the preprocessing stage, the analysis more accurately reflects the true frequency characteristics of mosquito wingbeats and improves the reliability of the overall detection method.

Figure 3. Sliding-window AUC analysis of insect and non-insect audio segments

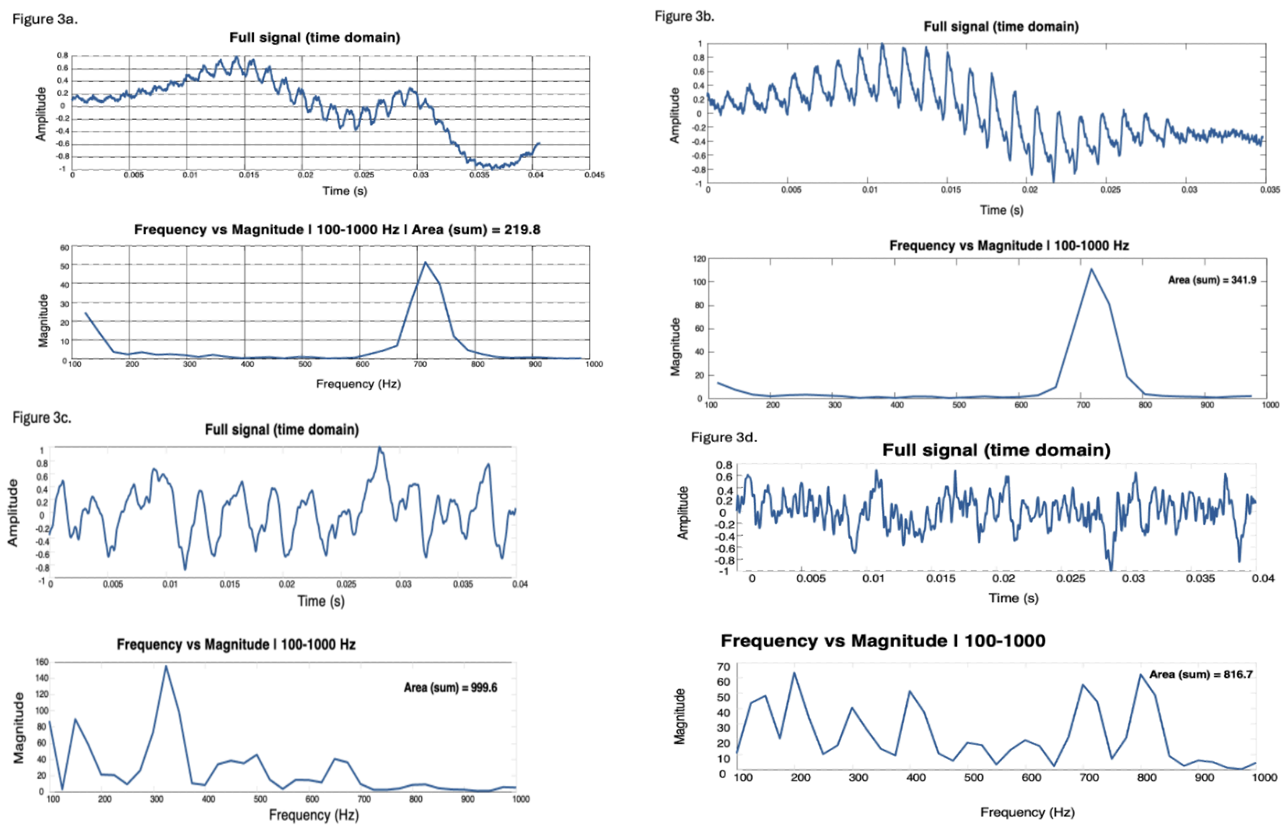

Figure 3. Time-domain waveforms and corresponding frequency-domain magnitude spectra for four recorded signals with reported area under curve values. Panels 3a–3d show the full signal amplitude over time (top) and the frequency magnitude distribution from 100–1000 Hz (bottom), with the summed spectral area reported for each signal. Figure 3 illustrates how the area-under-curve (AUC) metric distinguishes between insect recordings and non-insect audio. The Gambie Male insect recordings (Figures 3a and 3b) produce significantly lower AUC values, approximately 219.8 and 341.9, respectively. In contrast, the pop music segment (Figure 3c) and the rock music segment (Figure 3d) yield much higher AUC values, at 999.6 and 816.7, respectively.

Figure 3 demonstrates the purpose of the area-under-the-curve tool for distinguishing insect-like audio from non-insect audio. Rather than only identifying the largest frequency peak, this tool measures how spectral energy is distributed across the biologically relevant 100–1000 Hz range. This is important because insect wingbeats typically produce a narrow, concentrated frequency peak, while non-insect sounds such as music often spread energy across many frequencies.

The areaUnderCurveTool first loads the audio file, or accepts a signal segment directly when used inside the sliding-window tool. Only frequencies within the selected 100–1000 Hz band are retained, and the AUC score is calculated by summing the magnitudes across that band.

The results in Figure 3 show clear differences between insect and non-insect signals. The mosquito recordings in Figures 3a and 3b produced lower AUC values of 219.8 and 341.9, respectively. Both these figures have narrow peaks at 700 Hz with little energy spread across the rest of the band, which is expected with a stable mosquito wingbeat signal. The time-domain waveforms for these two figures also appear more periodic, reflecting regular wing motion.

In contrast, the pop and rock music samples in Figures 3c and 3d produced higher AUC values of 999.6 and 816.7, respectively. These signals contain multiple peaks and broader spectral energy likely due to overlapping instruments, vocals, and rhythm. This broader distribution increases the total area under the curve and reflects the more complex, irregular structure displayed in the time-domain waveforms and typically seen for audio data of music.

Overall, these results show that the AUC tool provides a useful measure for differentiating insect-like audio segments from non-insect audio. Lower AUC values indicate more concentrated, wingbeat-like spectral content, while higher values indicate broader, non-insect audio. Ultimately, the AUC score can identify which portions of a recording are most spectrally similar to insect wingbeats, which provides the basis for the main method, the sliding window.

Figure 4. Sliding Window Identification of the Best Insect-Like Subsequence

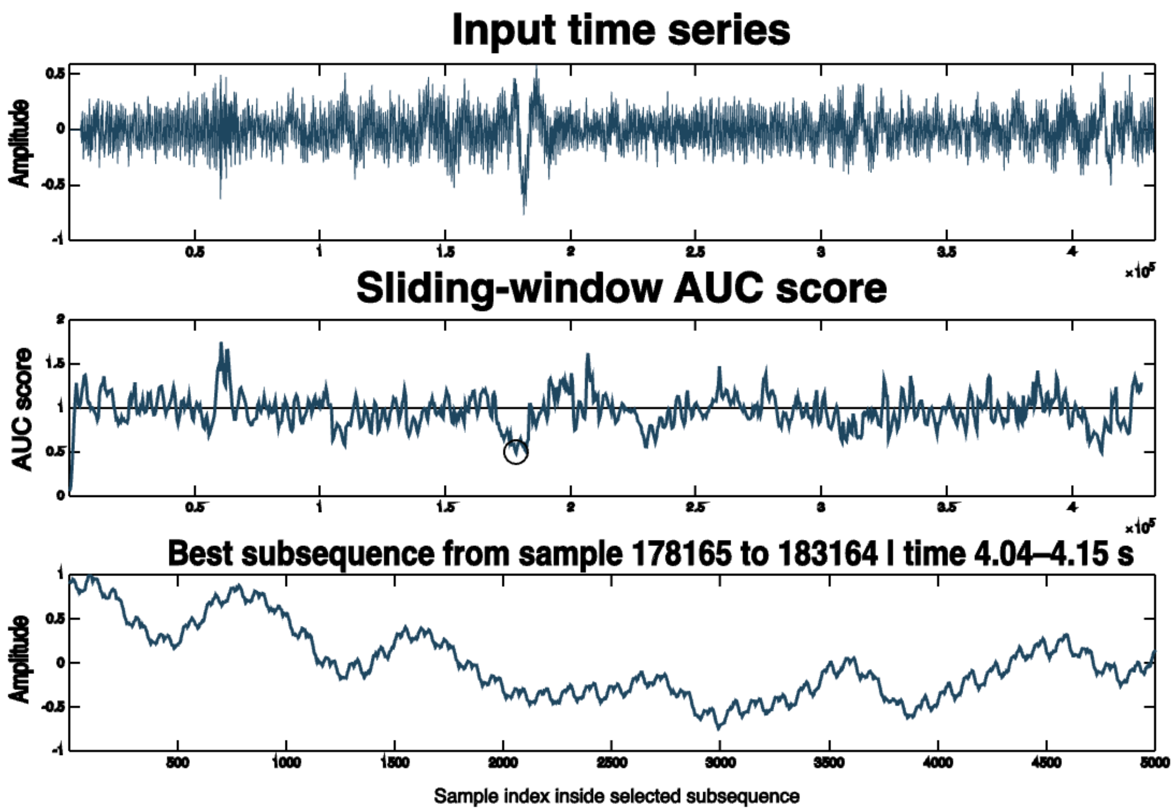

Figure 4. This figure shows the sliding-window AUC analysis performed on a longer audio recording. The top panel displays the full input time-series signal in the time domain, with amplitude plotted across the full sample range. The middle panel shows the AUC score calculated for each sliding window as it moves across the recording. The circled point marks the lowest AUC score, indicating the selected best subsequence. The bottom panel shows this selected subsequence in greater detail, from sample 178165 to 183164, corresponding to approximately 4.04–4.15 seconds. In the bottom plot, the x-axis represents the sample index inside the selected subsequence.

Figure 4 demonstrates how the sliding-window AUC tool identifies the best candidate insect-like segment within a longer recording. The full time-series signal contains many amplitude fluctuations, making it difficult to determine the best subsequence by visual inspection alone. To solve this, the tool scans the recording using a fixed-size window and calculates an AUC score for each window.

The middle panel provides the main decision metric. Each AUC score represents how much spectral energy is spread across the selected frequency range for that window. Lower AUC values indicate a more concentrated frequency distribution, which is expected for insect wingbeat audio. The circled minimum represents the window with the lowest AUC score, meaning this portion of the recording was identified as the most insect-like segment.

The bottom panel displays the selected subsequence from approximately 4.04–4.15 seconds. This zoomed-in view allows the chosen segment to be inspected separately from the full recording. Overall, Figure 4 shows that the sliding-window AUC method can automatically locate the portion of a longer signal that best matches the expected spectral structure of insect audio.

DISCUSSION

The major final experiment for this project will be a validation test of the full sliding-window workflow. While the current results show that the method can distinguish insect-like audio from non-insect audio, this experiment is intended to test the main goal of the paper which is whether the algorithm can automatically locate an insect signal hidden inside a longer, noisy recording. This represents the overall purpose of the project because a real acoustic sensor must be able to identify short insect signals within complex environmental audio, rather than only analyze clean insect recordings.

For this experiment, ten distinct 10-second non-insect audio clips will be collected, such as music, speech, machine noise, or environmental sound. A different mosquito recording will be inserted into each clip at a known location, such as the 5-second mark. The algorithm will then analyze each file without being given the insertion point and report the location it identifies as the best insect signal.

The guidelines for success will be defined before running the experiment. If the mosquito signal is inserted at 5 seconds, for example, any detected location within a reasonable range around that point, such as between 4.9–5.5 seconds, will be counted as a successful detection. Any incorrect detections outside this range will be counted as failures. The final performance will be reported as an accuracy score, such as 8 out of 10 successful detections. This experiment is important because it turns the project from a set of demonstrations into a measurable test of algorithm performance.

CONCLUSION

In this study, a MATLAB-based signal-processing workflow was developed to improve the accuracy of insect wingbeat detection in long audio recordings that may contain both insect and non-insect sounds. While FFT-based frequency extraction was useful for identifying dominant wingbeat frequencies, the ultimate goal of this project was not only to detect a frequency peak, but to identify the best and cleanest insect sequence within a larger recording. This is especially important because real-world audio data collected by sensors may include human activity, environmental noise, wind, music, or other non-insect sounds that can interfere with accurate insect detection.

The area-under-the-curve and sliding-window tools were developed to support this goal. The AUC tool helped measure how concentrated or spread out the spectral energy was within a recording segment. Since insect wingbeats tend to produce a more concentrated frequency peak, while non-insect sounds often contain broader and more scattered energy, AUC provided a way to distinguish insect-like signals from noisy or irrelevant data. The sliding-window method then applied this idea across the recording by scanning through short segments of audio and identifying the portion most likely to contain a true insect signal. By selecting the best insect-like window before performing frequency analysis, the workflow reduced the influence of junk data and improved the likelihood that the detected peak represented an actual insect wingbeat.

Overall, this project demonstrates that preprocessing is essential for making acoustic insect sensors as accurate as possible. Instead of relying on full-recording analysis, which can be distorted by noise, the proposed workflow focuses on isolating the highest-quality insect sequence from complex audio data. Although further testing across more insect species, environments, and longer recordings is needed, the AUC and sliding-window approach provides an important foundation for more reliable automated insect monitoring. With continued development, this method could improve sensor accuracy for pest management and have major entomological and public health applications.

REFERENCES

1. Fengler, Bárbara, et al. “A review of acoustic systems for insect monitoring: From recording technologies to species recognition performance.” *Ecological Informatics*, vol. 95, May 2026, p. 103733, https://doi.org/10.1016/j.ecoinf.2026.103733
2. Chen, Yu, et al. “Flying Insect Classification with Inexpensive Sensors.” *Journal of Insect Behavior*, vol. 27, no. 5, 2014, pp. 657–677, https://doi.org/10.1007/s10905-014-9454-4
3. Wagner, David L., et al. “Insect Decline in the Anthropocene: Death by a Thousand Cuts.” *Proceedings of the National Academy of Sciences*, vol. 118, no. 2, 2021, p. e2023989118, https://doi.org/10.1073/pnas.2023989118
4. MathWorks. “fft: Fast Fourier Transform.” *MATLAB Documentation*, The MathWorks, Inc., accessed 6 May 2026, https://www.mathworks.com/help/matlab/ref/fft.html
5. Erbe, Christine, and Michael L. Dent. “What Is Animal Bioacoustics?” *Acoustics Today*, vol. 13, no. 2, 2017, pp. 65–71, https://cmst.curtin.edu.au/wp-content/uploads/sites/4/2017/06/Erbe-C-AcousticsToday-2017-What-is-Animal-Bioacoustics.pdf
6. Lövei, Gábor L., et al. “The Need for a (Non-Destructive) Method Revolution in Entomology.” *Biological Conservation*, vol. 282, 2023, p. 110075, https://doi.org/10.1016/j.biocon.2023.110075