

Acoustic Drone Package Delivery Detection

François Marcoux and François Grondin

Abstract—In recent years, the illicit use of unmanned aerial vehicles (UAVs) for deliveries in restricted area such as prisons became a significant security challenge. While numerous studies have focused on UAV detection or localization, little attention has been given to delivery events identification. This study presents the first acoustic package delivery detection algorithm using a ground-based microphone array. The proposed method estimates both the drone’s propeller speed and the delivery event using solely acoustic features. A deep neural network detects the presence of a drone and estimates the propeller’s rotation speed or blade passing frequency (BPF) from a mel spectrogram. The algorithm analyzes the BPFs to identify probable delivery moments based on sudden changes before and after a specific time. Results demonstrate a mean absolute error of the blade passing frequency estimator of 16 Hz when the drone is less than 150 meters away from the microphone array. The drone presence detection estimator has a accuracy of 97%. The delivery detection algorithm correctly identifies 96% of events with a false positive rate of 8 %. This study shows that deliveries can be identified using acoustic signals up to a range of 100 meters.

I. INTRODUCTION

Unmanned aerial vehicles (UAVs) have grown rapidly due to their affordability, maneuverability and long range capabilities. These features make them well-suited for a variety of applications, including videography, search and rescue missions and item delivery in remote areas. However, UAVs pose significant risks in no-fly zones such as airports and prisons. Authorities report a rise in hostile UAV intrusions over the past few years. In prisons, UAVs are frequently used to smuggle contraband such as cellphones, weapons, and drugs [13, 12]. Although active mitigation systems exist and can neutralize drones [3, 20], local laws often prohibit radio frequency jamming or drone interception [15]. Authorities therefore rely on multimodal detection systems to identify and track intrusions in real-time, and eventually recover packages dropped within the perimeter.

Surveys on drone detection systems [14, 2, 16] identify four sensor types used for drone surveillance: cameras, radio-frequency antennas, radar, and microphones. Cameras are relatively inexpensive but their effectiveness is highly dependent on the environment, lighting, and

field of view. Radars provide long-range object detection but are more expensive and susceptible to false positives, as small, slow-moving drones can be mistaken for birds [2]. Radio-frequency antenna eavesdrops on data sent between drones and their ground stations. Valuable data like GPS and possibly drone delivery state can be extracted. However, encrypted communication or an autonomous drone with preprogrammed flight missions make this strategy inefficient [16]. Microphone arrays are cost-effective and capable of localizing sound sources, but they have limited range and are sensitive to noise. In most scenarios, multimodal systems offer the best overall performance [14].

To the best of our knowledge, no prior research addresses package delivery detection. However, two related tasks have been explored: 1) classification of payload weight [7, 4, 8, 18] and 2) detection of payload presence [10, 11, 19, 21]. The first studies on the classification of loaded versus unloaded UAVs are presented in [10, 11]. Both use recordings from a DJI Phantom 2 and perform classification with a convolution neural network (CNN) on long audio sequences (5 to 16 seconds). Mel Frequency Cepstrum Coefficients (MFCC) [10] or Short-Time Fourier Transform (STFT) [11] can be used as input features and both achieve more than 99% accuracy. Despite their high accuracy, these methods are restricted to a specific drone model and overlook the impact of the distance between the microphone and the drone. Other studies aimed to evaluate different machine learning models and feature extraction techniques. Deep neural networks [8] and classical machine learning models [21, 7] have been tested for both payload classification and detection tasks, consistently achieving around 99% accuracy. However, data collection is limited to one or two drone models and the audio is recorded at short range (up to 15 meters). Ibrahim et al. [7] and Doster et al. [4] further demonstrate that low-cost microphones can capture drone sound up to a range of 100 m, as they achieve 99% accuracy using MFCC and a Support Vector Machine (SVM) model. They also investigate the impact of distance and wind noise level on accuracy. Accuracy decreases with distance reaching 70% at 100 m and wind noise levels above 54 dB significantly impact the accuracy of the predictions. Utebayeva et al. [19] collect data from 11 different drone models at different locations with multiple payload weights and at distances up to 100 m. Using Mel spectrograms of one second and a Recurrent Neural Network (RNN), they achieved 98% accuracy.

This work is financed through the Fonds de recherche du Québec - Nature et Technologies with the Research Support for New Academics grant. F. Marcoux and F. Grondin are with the Department of Electrical Engineering and Computer Engineering, Université de Sherbrooke, Québec, Canada

Previous work demonstrates that UAV payload detection and classification using acoustic signals is feasible. However, no studies have focused on identifying the precise moment of a delivery event. This paper introduces an acoustic drone delivery detection system. We conducted numerous test flights to collect audio data both onboard the drone and from a ground-based station. We also collect telemetry data, including GPS coordinates and motor pulse width modulated (PWM) control signals. We use pairs of audio segments and telemetry data to train a supervised convolutional recurrent neural network (CRNN) that estimates the drone activity and the blade passing frequencies (BPF) of the motors. Based on the estimated BPF, we design a detection algorithm that predicts delivery events.

The rest of this paper is organized as follow. Section II presents the drone activity and BPF estimation algorithm and the delivery detection algorithm. Section III describes the experimental setup, data collection procedure and training hyperparameters. Section IV shows the results and discussion about our algorithm performances and Section V presents the conclusion and limitations.

II. BLADE PASSING FREQUENCY ESTIMATION AND DELIVERY DETECTION

Each propeller emits a sound characterized by the blade passing frequency (BPF) and some harmonics, defined as:

$$BPF = N_p \omega, \quad (1)$$

where $N_p \in \mathbb{N}$ is the number of blades on the propeller and $\omega \in \mathbb{R}^+$ stands for the angular speed of the motor's shaft in revolutions per second. When a drone hovers, the speed of its motors depends on its weight and external perturbations such as wind. The BPFs of a drone are on average higher with a payload than without it. Figure 1 shows this difference in the BPFs obtained from the control signal sent to each motor.




Fig. 1: BPF distribution before and after a delivery. A payload shifts the distribution to the right as the effective drone weight increases. The ground truth BPF data from multiple delivery segments were used to generate this graph resulting in approximately 25 000 data points.

To detect drone delivery event, we propose a two steps approach; 1) Use the acoustic signal recorded on the ground-based microphone array to detect the drone and

extract the BPFs of two propellers using a convolutional and recurrent neural network (CRNN); 2) Use the estimated BPFs to identify possible delivery moments. The general approach is presented in Figure 2.

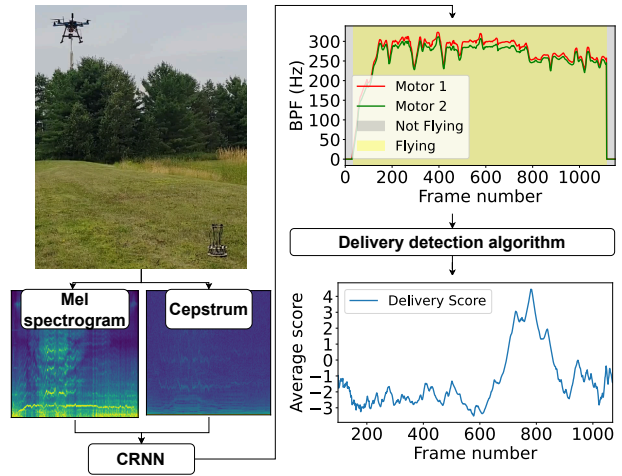


Fig. 2: Architecture overview of the proposed algorithm.

A. Blade passing frequency estimation and drone activity detection

In practice, the BPF for each motor cannot be measured directly and must be estimated from the acoustic signals. More specifically, the proposed algorithm estimates the BPFs of the two fastest motors, as these actuators are the most audible under various conditions. A binary flag also predicts whether there is a drone sound event or not. The input of the network is a 93 frames mel spectrogram generated using the parameters in Table I and a 93 frames power cepstrum which correspond to three seconds of audio. The higher quefrency bins of the power cepstrum were removed to match the dimension of the mel spectrogram. Table II shows the architecture of the proposed supervised multitask CRNN, with two branches to predict the BPF and the drone activity. The convolution blocks presented in Table III extract the common features for both tasks, which are then fed to two branches, each with three stacked bi-directional gated recurrent unit (bGRU) and a two fully connected layer. The ReLU activation function ensures BPFs are positive, and the sigmoid function restricts the activity to the interval $[0, 1]$. The BPFs are sorted from smallest to largest at each time frame to avoid ambiguity since they can overlap or cross during flights.

TABLE I: Mel spectrogram parameters

Window size	Hop size	Sampling	Nb of mel	Max freq
2048	512	16 000 Hz	128	7000 Hz

To train the multitask network, a combination of the mean squared (MSE) error loss function and the binary

TABLE II: Proposed CRNN architecture for BPF estimation and drone activity detection

Input (93 frames \times 2 \times 128 bins)	
ConvBlk(33) ConvBlk(21) ConvBlk(11) ConvBlk(3)	
Concatenate ConvBlks = (93 \times 128 (channels) \times 128 (bins)) Flatten channels and bins = (93 \times 16384)	
bGRU($p = 0.4, H = 128$) bGRU($p = 0.4, H = 128$) bGRU($p = 0.4, H = 128$)	bGRU($p = 0.4, H = 128$) bGRU($p = 0.4, H = 128$) bGRU($p = 0.4, H = 128$)
Fully Connected(256, 128) ReLU Dropout($p = 0.4$) Fully Connected(128, 2) ReLU	Fully Connected(256, 128) ReLU Dropout($p = 0.4$) Fully Connected(128, 1) Sigmoid
BPF (93 frames \times 2)	Activity (93 frames \times 1)

$p =$ dropout probability, $H =$ hidden size

TABLE III: Convolutional block definition

ConvBlk (k)
Input = (93 frames \times 2 \times 128 bins)
Conv2D($K = (3, k), S = 1, P = (1, \lfloor k/2 \rfloor), C = 16$) + BN + ReLU Conv2D($K = (3, 3), S = 1, P = (1, 1), C = 32$) + BN + ReLU
Output = (93 frames \times 32 channels \times 128 bins)

$K =$ kernel, $S =$ stride, $P =$ padding, $C =$ number of channels,

cross-entropy loss function (BCE) is used:

$$L = \alpha(\hat{\mathbf{y}}_b - \mathbf{y}_b)^2 + \beta(\mathbf{y}_d \log \hat{\mathbf{y}}_d + (1 - \mathbf{y}_d) \log(1 - \hat{\mathbf{y}}_d)), \quad (2)$$

where $\alpha, \beta \in \mathbb{R}^+$ are weights that balance both loss functions, $\mathbf{y}_b, \hat{\mathbf{y}}_b \in \mathbb{R}^+$ are the BPF label and prediction, respectively, and $\mathbf{y}_d, \hat{\mathbf{y}}_d \in [0, 1]$ are the drone activity label and prediction, respectively. In our case, $\alpha = 1.0$ and $\beta = 1.0$.

B. Delivery detection algorithm

The delivery detection algorithm estimates, for each time frame t , the likelihood of a delivery and is based on a change point detection algorithm [1]. First we define the estimated BPF as $B_t^m \in \mathbb{R}^+$ where $m \in \{1, 2\}$ is the motor index and $t \in [0, T - 1]$ is the frame index. We also define a window that spans $n \in 2\mathbb{N} + 1$ frames centered on the middle frame t as:

$$\gamma_n^m | t = \{\mathbf{b}_n^m | t, \mathbf{a}_n^m | t\}, \quad (3)$$

where

$$\mathbf{b}_n^m | t = \{B_{t-\frac{n}{2}}^m, \dots, B_{t-1}^m\}, \mathbf{a}_n^m | t = \{B_t^m, \dots, B_{t+\frac{n}{2}}^m\}. \quad (4)$$

Histograms H_a and H_b are computed for $\mathbf{a}_n^m | t$ and $\mathbf{b}_n^m | t$, respectively. The histograms are bounded between 100 Hz and 500 Hz and must contain a minimum of $n/4$ non-zero values to estimate the delivery score. This prevents outliers and sequences with high rates of negative predictions from skewing the distribution. Samples for a time frame t that fail to fulfill these criteria are discarded. The motor delivery score $s_n^m | t$ is computed using a weighted sum (with weights w_1, w_2, w_3 and w_4)

of the Chi-Squared distance D_{χ^2} , the Jensen-Shannon divergence D_{JSD} , the histogram intersection score D_{HI} and the difference in average between H_a and H_b , denoted $\Delta\mu_{a,b}$:

$$s_n^m | t = w_1 D_{\chi^2} + w_2 D_{JSD} + w_3 D_{HI} + w_4 \Delta\mu_{a,b}. \quad (5)$$

The final score d_t is obtained by averaging the delivery score of both motors for all windows:

$$d_t = \frac{1}{2N} \sum_{i=1}^N (s_{n_i}^1 | t + s_{n_i}^2 | t), \quad (6)$$

where N is the total number of windows and n_i is the windows size for each window i . The weights $w_1 = 10, w_2 = 2, w_3 = -20$ and $w_4 = -0.05$ were chosen manually after testing different values. Delivery events correspond to a time t when the score d_t exceeds a given threshold. In an online scenario, the windows size n influences the responsiveness of the algorithm. A shorter window provides a quicker response but can be affected by instant deviation while a larger window provides a slower but more precise response. We chose a total of 15 sliding windows ranging from 93 frames (3 seconds) to 465 frames (15 seconds). For every time frame, the score of each window is averaged with invalid estimation ignored. In a "real-time" scenario, the worst case lag introduced by our delivery detection algorithm is 7.5 seconds, which is half of the biggest window. Figure 3 shows an overview of the delivery detection algorithm for one window size.

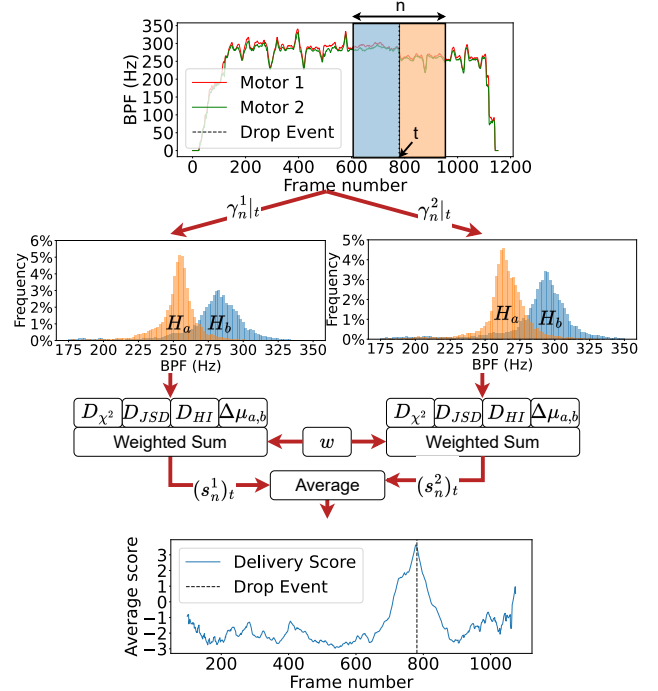


Fig. 3: Delivery detection algorithm overview. D_{χ^2} is the Chi squared distance, D_{JSD} the Jensen-Shannon divergence, D_{HI} the intersection score and $\Delta\mu_{a,b}$ the difference in average between both histograms

III. EXPERIMENTAL SETUP

A. Data collection

A modified Holybro X500 V2 drone kit¹ was used to conduct multiple flight tests. The drone is equipped with a servo motor used for package delivery and an onboard Raspberry Pi (RPi) for telemetry and audio recording. A ground-based 16-microphone array (with micro-electromechanical systems (MEMS) microphones and the 16-sound USB card [9]) captures audio data from the drone. The microphones are carefully positioned to minimize symmetry between the pairs, thereby reducing side lobes in localization algorithms. Figure 4 shows the drone and the microphone array setup used for data collection.




Fig. 4: The modified Holybro X500 V2 drone and the 16 microphone array

During each flight, the RPi records acceleration, speed, orientation, GPS coordinates, servo position, motor Pulse-Width Modulation (PWM) signals, flight status and audio signals. To ensure data synchronization, each telemetry data point is saved along with its corresponding timestamp, and the start time of the audio recording is also logged. Data synchronization is performed offline. The audio from the 16-microphone array is recorded at a sample rate of 16,000 samples/sec on a dedicated laptop. Hand claps are captured at the beginning of each flight to manually synchronize the onboard and ground-based audio recordings.

A total of 28 flights were recorded: 17 contains both onboard and array audio (143 minutes of flight time) and 11 contains only onboard audio (51 minutes of flight time). Additionally, 39 minutes of background noise recordings were collected between flights, capturing sounds such as highway traffic, voices, hand claps and chainsaw. Flight trajectories vary significantly, ranging from linear paths with small acceleration to complex maneuvers involving frequent changes in direction, altitude, and speed. The drone-to-array distance during these flights ranged from 0 to 160 m. Two additional flights were recorded with the drone hovering at various distances from the array, up to 180 meters. These flights were specifically designed to evaluate the performances of the CRNN-based BPF estimation across different distances.

¹<https://holybro.com/products/x500-v2-kits>

A total of 23 delivery events were recorded, with a payload weighting between 500g and 600g. Six of these deliveries were recorded by both the onboard microphone and the microphone array, while the remaining 17 events were recorded only by the onboard microphone. For each delivery, 30 seconds of audio before and after the event were extracted and the BPFs were estimated using the network described in Section II. To prevent overfitting, the flights that contain delivery events were excluded from the training dataset. The recorded data are available here².

The motor's PWM signal, sent by the flight controller, is compared to the actual BPF of the propeller, which was measured using an optical tachometer and reflective tape placed on the propeller blades. BPF measurements were taken for each motor across a PWM range from 1100 μ s (10%) to 2000 μ s (100%) at 100 μ s (10%) intervals, with each measurement lasting between 3 to 7 seconds. These observations were averaged across all four motors and interpolated to generate the final PWM-to-BPF calibration curve shown in Figure 5.




Fig. 5: Measured BPF for PWM signal ranging from 1100 us (10% throttle) to 2000 us (100% throttle)

One important observation made during those measurements is that battery level affects the PWM-to-BPF relationship. The curve of the first motor (100% battery) was slightly different from the last motor measurement (50% battery). Because battery state was not recorded during flight, we had to manually adjust the BPFs along every flight so they were aligned with the audible frequencies. An offset in Hz was set for the beginning and the end of the recording and the BPF were gradually adjusted throughout the recording. In the majority of the recording, the maximal offset was around 20 Hz

B. Data processing

For the BPF estimation algorithm, only flights containing the onboard and ground-based array audio were used for the training, validation and testing. Out of the

²<https://github.com/introlab/d3set>

17 available flights, 8 were used for the training, 3 for validation and 6 for testing. The test set was further split based on the recording device: one with the onboard audio and the other using the array audio. The BPFs were sorted from smallest to largest for each time frame to avoid ambiguity since BPFs from different motors can intersect. BPFs were also slightly smoothed using a Kalman filter to reduce abrupt variations that did not match the audible signal. Moreover, frame where the drone was more than 150 meters away from the microphone array were removed from the training and validation datasets.

Audio data augmentation was applied to the training dataset using the *audiomentations*³ library. The augmentation included: 1) light background noise with a SNR between 20 dB and 30 dB; 2) random gain from -5 dB to 5 dB; 3) time masking between 0.2 and 0.5 seconds; 4) frequency masking between 0 Hz and 7000 Hz. For each raw audio file, two augmented versions were generated by applying augmentations with randomized parameters every second. Each augmentation was applied with a probability of approximately 50%. AudioSet [5] was used as background noise for augmentation and as complementary noise samples in the datasets. Additionally, we manually selected and incorporated noises such as helicopter, aircraft, mosquito, white noise and sinusoids ranging from 10 Hz to 6000 Hz in the training dataset. In the validation dataset, only noise from the AudioSet was added, while the test set included noise exclusively from the Urban16K [17] dataset. Table IV shows the composition of each dataset.

TABLE IV: Datasets distribution in second

Sound type	Train	Valid	Test	Onboard	Array
Drone	4752	2648	5993	2996	2996
Drone augment	9504	0	0	0	0
Added noise	5709	1064	2997	0	0
White noise	570	0	0	0	0
Sinusoidal noise	570	0	0	0	0
Silence	570	0	0	0	0
Total	21675	3712	8990	2996	2996

IV. RESULTS AND DISCUSSION

A. BPF estimator

The performance of the BPF estimator was evaluated using two metrics: 1) the mean absolute error (MAE) in Hz; 2) a custom metric called the masked MAE (MMAE). The MMAE corresponds to the MAE when the drone activity flag exceeds a predefined threshold. For drone activity detection, metrics include accuracy, recall and precision. Table V shows the metrics for the training, validation and both test sets.

Results with the onboard test set are promising with a drone activity accuracy, recall and precision of 99%, and a BPF MAE of 6.2 Hz, which is close to the spectrogram

TABLE V: Neural network metrics for the training, validation and both test dataset

Metrics	Train	Valid	Test	Onboard	Array
MAE (Hz)	3.9	7.6	13	6.2	31
MMAE (Hz)	7.3	10.6	13	7.2	18
Accuracy (%)	99	97	97	99	91
Precision (%)	99	97	98	99	97
Recall (%)	99	98	95	99	90

frequency bin resolution of 7.8 Hz. The performances slightly degrades for the array test set, with a MAE of 31 Hz and an accuracy of 91%. This is expected as the drone acoustic signal attenuates with increasing distance.

The network performances were also evaluated across multiple drone-to-array distances and SNRs. For the distance-based evaluation, predictions from the excluded flight mentioned in Section III were divided into 9 groups based on the drone-to-microphone array distance, ranging from 0 meter to 180 meters. In each group, random background noise was added to balance the groups with 50% drone and 50% noise labels. Figure 6 shows the receiver operating characteristic (ROC) curves computed on the drone activity predictions, and the MAE and MMAE of the estimated BPF. Figure 7 shows the MAE with respect to the SNR. We used background urban noise recorded from the microphone array and white noise.


Figure 6a shows that the recall, or true positive rate for a fixed false positive rate decreases significantly with distance. In fact, the activity detection is accurate with a true positive rate of 100% at a 10% false positive rate up to 100 meters. Figure 6b shows a similar pattern with the MAE staying below 15 Hz up to 100 meters. MMAE is generally lower than MAE across groups, which confirms that the activity flag can be used at inference to obtain more accurate predictions on the blade passing frequencies, especially at higher drone-to-array distances.

Figure 7a confirms that the BPF estimator is reliable from -22 dB onward. This robustness to background noise, which is predominantly concentrated in lower frequencies, is due to the presence of high-frequency components (around 5000 Hz) known as the azimuthal mode 2 vibration generated by the rotors [6]. Figure 7b confirms this observation as white noise also masks high frequencies, which results in an higher MAE at low SNRs.

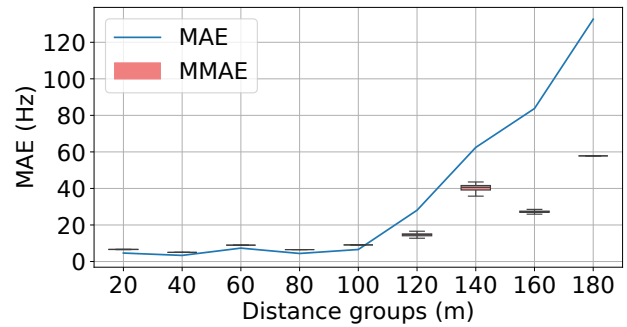
B. Package delivery detection

To evaluate the package delivery detection algorithm, each registered delivery event was extended to a 0.5 second window. This adjustment helped account for slight misalignments between the label and the event. For the 29 test cases, true positive rate and the false negative rate were computed across multiple threshold values, and the resulting ROC curve is presented in Figure 8. At a false positive rate of 8%, 96% of delivery

³<https://github.com/iver56/audiomentations>




(a) Drone activity detection ROC curves vs distances




(b) BPF MAE and MMAE vs distances

Fig. 6: Drone activity detection and BPF errors



(a) Urban noise



(b) White noise

Fig. 7: BPF MAE vs SNR




Fig. 8: Delivery detection algorithm ROC curve

frames were identified correctly, demonstrating strong performance. However, it is important to note that a relatively heavy payload (500g to 600g) was used, which caused significant changes in motor speeds at delivery. Smaller payloads would result in subtler speed variations, leading to smaller discrepancies between distributions and making detection more challenging. For now, the current event detection algorithm is agnostic of the drone position or speed, meaning that the Doppler effect, wind gust or intentional piloting maneuvers could lead to incorrect predictions.

V. CONCLUSION

This paper presents the first drone acoustic delivery detection system. Our method is based on a two steps approach; 1) Detect the drone and extract the BPFs of the two most audible propellers using a CRNN; 2)

Identify possible delivery moments in the estimated BPFs sequence using a change point detection algorithm. Experimental results showed that the CRNN accurately estimated the BPFs with a mean absolute error of 16 Hz and estimated the drone presence with an accuracy of 97% on realistic data with varying distances up to 100 m. We also conducted evaluation for multiples drone-to-array distances and SNRs. The distance evaluation highlighted that false negatives increases with distances. However, when a drone is detected, the estimated BPFs are accurate and we concluded that the proposed network is reliable up to 100 m. The SNR evaluation showed a great resilience to urban noise up to -22 dB SNR by utilizing drone's high frequency motor noise ($\approx 5000\text{Hz}$) to properly estimated the BPF even when the propeller fundamental frequencies are masked. The proposed delivery detection algorithm detected 96% of the 29 delivery frames with a false positive rate of 8%.

Limitations. One major limitation is that only one drone was used to train the BPF estimation algorithm, which could make the method sensitive to domain mismatch and limits the general applicability. Also, the delivery event detection approach relies solely on the differences between the data distribution besides an identified time frame. This difference is significant when using relatively heavy payloads compare to the drone weight ($\approx 50\%$) but would be less significant when using lighter weight resulting in more false negatives for the same true positive rate. Moreover, since the

delivery detection algorithm ignores the position of the drone, the Doppler effect, wind gusts or intentional drone maneuvers could lead to incorrect predictions. Direction of arrival estimation should be investigated in future work to make the system more robust.

REFERENCES

- [1] S. Aminikhanghahi and D.J. Cook. “A survey of methods for time series change point detection”. In: *Knowledge and Information Systems* 51.2 (2017), pp. 339–367.
- [2] J. A. Besada, I. Campaña, D. Carramiñana, L. Bergesio, and G. de Miguel. “Review and simulation of counter-uas sensors for unmanned traffic management”. In: *Sensors* 22.1 (2022), pp. 1–30.
- [3] M.O Dix, M. J. Osbourne, M. Ascolese, M. Kucharski Schwartz, C. Lindquist, M. Camello, and T. Craig. *Contraband detection technology in correctional facilities an overview of technologies for screening people, vehicles, and correctional settings*. Tech. rep. 300856. US Department of justice, 2021, p. 19.
- [4] H. Doster and B. Mullins. “UAV payload identification with acoustic emissions and cellphones”. In: *Proceedings of the International Conference on Cyber Warfare and Security*. Vol. 18. 1. 2023, pp. 523–533.
- [5] J.F. Gemmeke, D.P.W. Ellis, D. Freedman, A. Jansen, W. Lawrence, R.C. Moore, M. Plakal, and M. Ritter. “Audio set: An ontology and human-labeled dataset for audio events”. In: *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*. 2017, pp. 776–780.
- [6] B. S. Henderson, D. Huff, J. Cluts, and C. Ruggeri. *Electric motor noise from small quadcopters: Part II – source characteristics*. Tech. rep. NASA/CR-2018-219869. National Aeronautics and Space Administration, 2018.
- [7] O.A. Ibrahim, S. Sciancalepore, and R. Di Pietro. “Noise2Weight: On detecting payload weight from drones acoustic emissions”. In: *Future Generation Computer Systems* 134 (2022), pp. 319–333.
- [8] I. Ku, S. Roh, G. Kim, C. Taylor, Y. Wang, and E. T. Matson. “UAV payload detection using deep learning and data augmentation”. In: *Proceedings of the IEEE International Conference on Robotic Computing*. 2022, pp. 18–25.
- [9] P.-O. Lagacé, F. Ferland, and F. Grondin. “Ego-noise reduction of a mobile robot using noise spatial covariance matrix learning and minimum variance distortionless response”. In: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*. 2023, pp. 3533–3538.
- [10] S. Li, H. Kim, S. Lee, J. Gallagher, D. Kim, S. Park, and E. Matson. “Convolutional Neural Networks for Analyzing Unmanned Aerial Vehicles Sound”. In: *Proceedings of the International Conference on Control, Automation and Systems*. 2018.
- [11] D. Lim, H. Kim, S. Hong, S. Lee, G. Kim, A. Snair, L. Gotwals, and J.C. Gallagher. “Practically classifying unmanned aerial vehicles sound using convolutional neural networks”. In: *Proceedings of the IEEE International Conference on Robotic Computing*. 2018, pp. 242–245.
- [12] C. Norman. “A global review of prison drug smuggling routes and trends in the usage of drugs in prisons”. In: *WIREs Forensic Science* 5.2 (2023), pp. 1–33.
- [13] A. OHagan and R. Hardwick. “Behind Bars: The Truth about Drugs in Prisons”. In: *Forensic Research & Criminology International Journal* 5.3 (2017).
- [14] S. Park, H.T. Kim, S. Lee, H. Joo, and H. Kim. “Survey on anti-drone systems: components, designs, and challenges”. In: *IEEE Access* 9 (2021), pp. 42635–42659.
- [15] J. Russo, D. Woods, J.S. Shaffer, and B.A. Jackson. *Countering Threats to Correctional Institution Security: Identifying Innovation Needs to Address Current and Emerging Concerns*. Tech. rep. RAND Corporation, 2019.
- [16] J. Sadvovskis and A. Aboltins. “Modern methods for UAV detection, classification, and tracking”. In: *Proceedings of the IEEE Annual International Scientific Conference on Power and Electrical Engineering of Riga Technical University*. 2022, pp. 1–7.
- [17] J. Salamon, C. Jacoby, and J.P. Bello. “A dataset and taxonomy for urban sound research”. In: *Proceedings of the ACM International Conference on Multimedia*. 2014, pp. 1041–1044.
- [18] A. Traboulsi and M. Barbeau. “Identification of drone payload using mel-frequency cepstral coefficients and LSTM neural networks”. In: *Proceedings of the Future Technologies*. Springer International Publishing, 2021, pp. 402–412.
- [19] D. Utebayeva, L. Ilipbayeva, and E. T. Matson. “Practical study of recurrent neural networks for efficient real-time drone sound detection: A review”. In: *Drones* 7.1 (2023), p. 26.
- [20] J. Wang, Y. Liu, and H. Song. “Counter-Unmanned Aircraft System(s) (C-UAS): State of the Art, Challenges, and Future Trends”. In: *IEEE Aerospace and Electronic Systems Magazine* 36.3 (2021), pp. 4–29.
- [21] Y. Wang, F.E. Fagiani, K.E. Ho, and E.T. Matson. “A feature engineering focused system for acoustic UAV payload detection”. In: *Proceedings of the International Conference on Agents and Artificial Intelligence*. Vol. 3. 2022, pp. 470–475.