

An Artificial Neural Network Approach to Predict Rotor-Airframe Acoustic Waveforms

Arthur D. Wiedemann and Christopher Fuller
*Vibration and Acoustic Laboratories, Virginia Tech
 Blacksburg, Virginia, 24061, US*

Kyle A. Pascioni
NASA Langley Research Center, Hampton, VA, 23681, US

A surrogate artificial neural network/machine learning model was developed to predict the acoustic interaction for a fixed-pitch rotor in proximity to a downstream cylindrical airframe typical of small Unmanned Aerial System (sUAS) platforms. The model was trained to predict the acoustic waveform under representative hover conditions as a function of rotational speed, airframe proximity, and observer angle. Training data were acquired in an anechoic chamber on both isolated rotors and rotor-airframe configurations. Acoustic amplitude and phase of the revolution-averaged interaction were predicted, which required up to 25 harmonics to capture the impulse event caused by the blade’s approach and departure from the airframe. Prediction performance showed, on average, that the models could estimate the acoustic amplitude and phase over the relevant harmonics for unseen conditions with 86% and 75% accuracy, respectively, enabling a time-domain reconstruction of the waveform for the range of geometric and flow parameters tested.

I. Nomenclature

A	harmonic amplitude, [Pa]	R	radius of the rotor, [m]
\tilde{A}	amplitude prediction, [Pa]	W_d	weight distribution
C_T	thrust coefficient	X_r	central difference
d	diameter of the cylinder, [m]	y, y_0	revolution-average and ideal waveform, [Pa]
f	frequency, [Hz]	Δ	rotor-airframe proximity, [m]
k	harmonic number	θ	azimuth location of the observer, [°]
M_{tip}	tip Mach number	φ	elevation location of the observer, [°]
M	number of reserved runs	ψ	phase, [radian]
N	number of discrete points in waveform	$\tilde{\psi}$	phase prediction, [radian]
P_A	amplitude prediction performance	Ω	rotor rotation rate [rev/min]
P_ψ	phase prediction performance	<i>Subscripts</i>	
p'	acoustic pressure, [Pa]	i, j	index
<i>Acronyms</i>		MA	Master Aircrew
ANN	Artificial Neural Network	OASPL-A	A-weighted Overall Sound Pressure Level
BPF	Blade Passing Frequency	RAI	Rotor-Airframe Interaction
B&K	Brüel and Kjær	sUAS	small Unmanned Aerial System
CFD	Computational Fluid Dynamics	WMSE	Weighted Mean Square Error
DoF	Degree of Freedom	UIUC	University of Illinois at Urbana-Champaign

II. Introduction

For small Unmanned Aerial Systems (sUASs), the Rotor-Airframe Interaction (RAI) phenomenon contributes to and, in some cases, dominates the acoustic emission. For example, Zawodny and Boyd [1] conducted a study in which cylindrical and conical rods were placed underneath the rotor disk to measure the RAI effect in an isolated environment. The results showed increased tonal content at harmonics of the Blade Passing Frequency (BPF), and the interaction was highly directional. In addition, a Computational Fluid Dynamics (CFD) analysis showed that an approaching and departing blade produced a strong potential field on the surface of the rod, resulting in the rod emitting a strong impulse noise. It was noted that the tonal noise dropped significantly when the rotor-airframe separation increased, while the tonal content increased with larger cylindrical diameters. Wang et al. [2] investigated the RAI with Particle Image Velocimetry in an anechoic chamber. They compared experimental results

to a CFD simulation to understand the underlying behavior of the aerodynamic and acoustic interaction. They found that the airframe had minimal impact on the rotor’s aerodynamic load as the velocity fields were similar in magnitude with and without the airframe in place. Results also compared the acoustic output of a rod positioned above and below the rotor disk. After comparing the two rod positions, they found that a rod positioned above the rotor produced greater harmonic content because the airframe distorts the rotor inflow, which produces large periodic turbulent structures. In addition, smaller turbulence structures were prominent in the rod wake, creating additional broadband noise. Quinte et al. [3] compared rotors with various twist, taper, and pitch distributions operating near a cylindrical rod in hover to assess the aerodynamic and acoustic performance. They found that a rotor geometry with a higher twist and taper was beneficial for the performance metrics due to increased power loading while reducing the noise. Conversely, an increase in the pitch of the rotor from 4° to 12° resulted in a larger magnitude of pressure on the airframe. This resulted in a spectral excitation in the mid-frequency range and increased the A-weighted Overall Sound Pressure Level (OASPL-A) by 7 dBA for above-plane observers. In-plane observers saw a mild increase in OASPL-A, less than 1 dBA for the increased pitch.

The RAI is also observed in outdoor environments as Whelchel et al. [4,5] conducted experiments in an anechoic chamber and an open environment. Their results reinforced findings in Zawodny and Boyd’s work, and they demonstrated that the acoustic pressure decreases in magnitude as the observer moves away from the surface normals of the cylindrical bar. Gojon et al. [6] performed an RAI experiment to validate a CFD simulation. Their results showed that harmonic levels at $2\times\text{BPF}$ to $15\times\text{BPF}$ increased by up to 30 dB, and the RAI directivity is similar to a dipole. Recently, Gallo et al. [7] created an analytical model to predict the aerodynamic sound generated by a rod placed downstream of a rotor and applied a potential flow model to estimate the lift force using Goldstein’s method [8] to predict the far-field acoustic pressure for the BPF and the four subsequent harmonics. Compared to experimental results, the analytical model underpredicted the harmonic content, but this approach captured sound pressure level trends in rotational speed, airframe diameter, and airframe proximity.

Previous efforts, such as those mentioned above, have shown several cases where the RAI has been successfully modeled with CFD simulations. This approach is useful for obtaining insights into the RAI; however, this is computationally expensive and challenging to incorporate into the early stages of sUAS design. One potential method of modeling the RAI is to use a surrogate model. After training, these models can make accurate predictions for highly complex systems. This work uses Artificial Neural Networks (ANNs) to model and predict the RAI while avoiding the extensive computational effort required for CFD. Several ANNs were trained to model a two-bladed rotor in hover with a cylindrical airframe downstream. The data were collected in an anechoic chamber while varying rotational speed and rotor-airframe proximity from run to run. The acoustic pressure time history was then segmented with a tachometer signal to estimate the tonal acoustic waveform for all microphones. Afterward, a Fourier transformation was applied to the waveform to extract the harmonic content. Then, the ANNs were trained to predict the amplitudes and phases of the waveform as a function of rotational speed, rotor-airframe proximity, and observer location at each harmonic. The experimental setup section explains the process of collecting data and outlines the ANN architecture. The results explore trends in the experiment and the ANN’s ability to model and predict new waveforms.

III. Experimental Setup

A. Facility Layout

An experiment was conducted in the Structural Acoustic Loads and Transmission (SALT) anechoic chamber at the NASA Langley Research Center [9]. A rotor stand was positioned in the center of the room, with the rotor hub centered between the floor and ceiling wedges at the height of 3.2 m above the floor, as shown in Fig. 1a. The rotor stand supported a Scorpion SII-4020-420KV motor and a JR-3 6-DoF load cell and the two were separated by a neoprene damper to reduce vibrations transmitted to the JR-3 load cell. In addition, reflective tape was applied to the motor, allowing the optical tachometer to estimate the period for each revolution. This motor was chosen because it could quickly reach and maintain the desired rotational speed. For example, at the lowest rotational speed of $\Omega = 4000$ RPM ($M_{\text{tip}} = 0.172$), the motor had a standard deviation of 4 RPM, while the highest rotational speed of $\Omega = 9200$ RPM ($M_{\text{tip}} = 0.404$) had a standard deviation of 16 RPM. A speed controller was used to maintain the steady supply of power for the motor.

Figure 1 shows the placement of the microphones at three azimuthal locations around the test stand, marked with green, blue, and orange borders for 90° , 0° , and -90° , respectively. Each azimuthal location, θ , had a mount fixing microphones in an arc with a radial distance of 2.24 m from the rotor hub for three elevation angles; a total of nine

microphones were used in this study. Microphones M7, M8, and M9 were stationed in-plane ($\varphi = 0^\circ$) of the rotor hub and used $\frac{1}{4}$ " Type 4954-B free-field Brüel and Kjær (B&K) microphones. The out-of-plane microphones, M1 – M6, used the $\frac{1}{4}$ " free-field GRAS 46BE. Table 1 lists the type of microphone used and its location in azimuth and elevation with respect to the rotor hub.

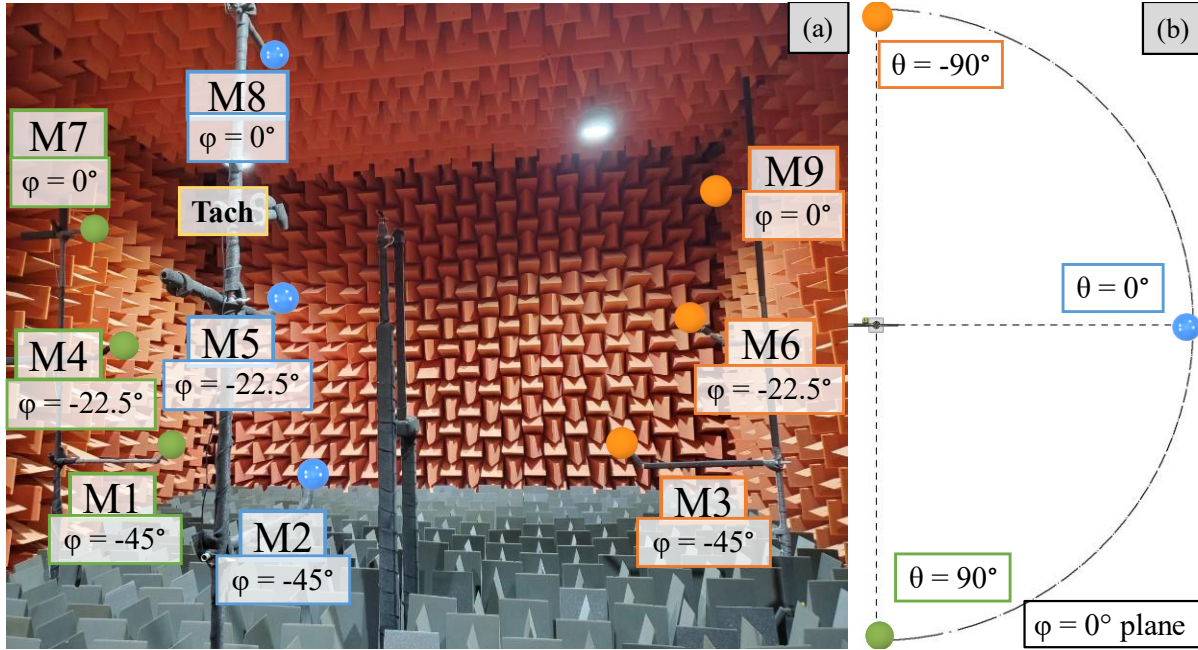


Figure 1: Experimental setup in the SALT facility with (a) nine microphones viewed from a low angle [9] and (b) a top view drawing.

Table 1: Microphone used and layout.

Microphone	M1	M2	M3	M4	M5	M6	M7	M8	M9
Mic type	GRAS	GRAS	GRAS	GRAS	GRAS	GRAS	B&K	B&K	B&K
Azimuth, θ	90°	0°	-90°	90°	0°	-90°	90°	0°	-90°
Elevation, φ	-45°	-45°	-45°	-22.5°	-22.5°	-22.5°	0°	0°	0°

The two-blade Master Airscrew (MA) 11x7 rotor has a radius of $R = 140$ mm and was positioned in proximity to a cylindrical carbon fiber rod with a constant cross-section and a diameter of $d = 25$ mm. The bar was mounted as a cantilever to a sleeve-bearing carriage, allowing the rod to change its proximity to the rotor between runs. Figure 2 shows a close-up of the stand with the cylindrical rod downstream of the rotor. The proximity between the bar and the rotor is Δ and is measured at the $0.75R$ blade span location.




Figure 2: Rotor and airframe assembly.

B. Data Acquisition

The JR-3 load cell measured thrust responses for each run. Thrust was averaged over a run, and the tare was subtracted from the measurement to return a response for the static thrust produced. Torque measurements were acquired with the load cell, but the outputs were unreliable and had to be discarded. Taring was performed between runs to monitor for drift throughout the experiment. In addition, brief intermissions were taken between runs to prevent recirculating air from contaminating consecutive runs. This was paramount as recirculated airflow increases the tonal noise once ingested by the rotor and can occur in enclosed spaces after a small-time duration [10]. Individual runs were kept to a maximum of 10 seconds to minimize recirculation effects. Spectrograms were used to monitor for recirculation, and the proceeding data could be discarded, if necessary. Data were acquired using 24-bit National Instruments PXI-4472B modules installed within a PXI-1045 chassis, synchronizing the load cell, tachometer, and microphone output with a sampling rate of 48 kS/s.

C. Test Conditions

The experiment is divided into two categories: an isolated rotor and a rotor with the airframe present. ANNs were trained to model the Fourier coefficient for each category. Table 2 provides a general overview of the number of runs conducted for each category and the range of operating conditions. The isolated rotor category only had one independent variable (tip Mach number) and could be incrementally increased as the experiment progressed, while the RAI had two independent variables under consideration, M_{tip} and Δ/R . An additional objective of this paper is to determine an effective means of sampling the domain of interest with the minimum number of runs. Previous efforts suggest that a sphere-packing scheme is an effective method of sampling the two-parameter domain. The JMP software package [11] generated a distribution of 35 sample points throughout the domain for $\Delta = 0.1R - 0.5R$ and $M_{tip} = 0.172 - 0.404$, and this initially served as the training dataset for the ANNs with additional samples, known as the unseen dataset, collected to evaluate the predictive performance of the ANNs. However, for the two-blade RAI test, the only runs that were useful in the analysis had the rotor-airframe proximity at $\Delta/R < 0.37$, resulting in 26 training and 25 unseen runs for building and evaluating the ANNs over the domain of interest represented in Table 2. On this reduced dataset, the ANNs had difficulty predicting the noise on the 26 training runs, so the training dataset was increased to 35 runs reducing the unseen dataset to 16. This augmented distribution was used to train and evaluate the RAI ANNs, and this distribution can be seen in Fig. 3. The training runs are marked with blue squares, and the unseen runs are marked with red circles. A more thorough explanation as to why runs above $\Delta/R > 0.37$ were removed from the analysis is explained in Section IV.B.

Table 2: Testing conditions used to train the two ANNs.

Category	M_{tip}	Δ/R	Number of Runs
Isolated	0.172 – 0.404	--	30
RAI	0.172 – 0.404	0.1 – 0.37	51

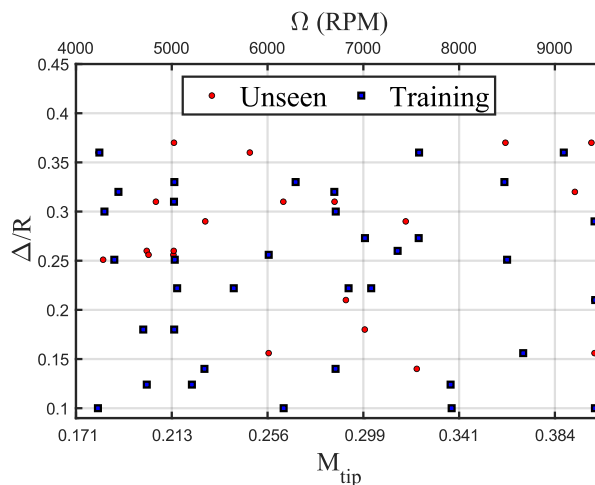


Figure 3: Distribution of sampled points over the domain of interest for the two-blade RAI.

IV. Methods

A. Postprocessing

The measured data collected in the anechoic chamber resulted in an acoustic pressure time history for each of the nine microphones. Figure 4 shows an example of the process of converting the measured signal into the revolution-averaged waveform. In order to collect the amplitude and phase data at the harmonics, the measured signal is run through a series of zero-phase filters to capture tonal content without distorting the phase. The MATLAB `filtfilt` function is applied at each harmonic with a second-order Butterworth filter using a bandpass window of ± 10 Hz up to the 30th harmonic. Afterward, the output from each filter is summed to construct a modified signal with the same time duration as the measured signal. Next, the filtered signal is segmented with pulse intervals measured by the tachometer, producing a series of single-revolution waveforms over the 10-second recording.

Depending on the tip Mach number, this would produce 700 to 1600 waveforms per run. The waveforms are overlaid and averaged to produce a steady signal and are referred to as the revolution-averaged waveform throughout the paper. The revolution-averaged waveform was also estimated with the pulse intervals applied to the measured signal, without any filters. Whether the measured data were filtered or not, the revolution-averaged waveform was consistent between processes. The time delay between the tachometer signal and the microphone measurements is ignored because this work focused on the waveform response to input parameters. With the revolution-averaged waveform constructed for a full period, the signal is extended in the time domain to increase the frequency resolution. Finally, a Fast Fourier Transformation (FFT) is applied to the waveform to extract the amplitude and phase at each harmonic and store them in a database. Phase wrapping occurs between $-\pi$ and π , and the ANN is unaware of the wrap. So, the phase dataset needs to be unwrapped before training.

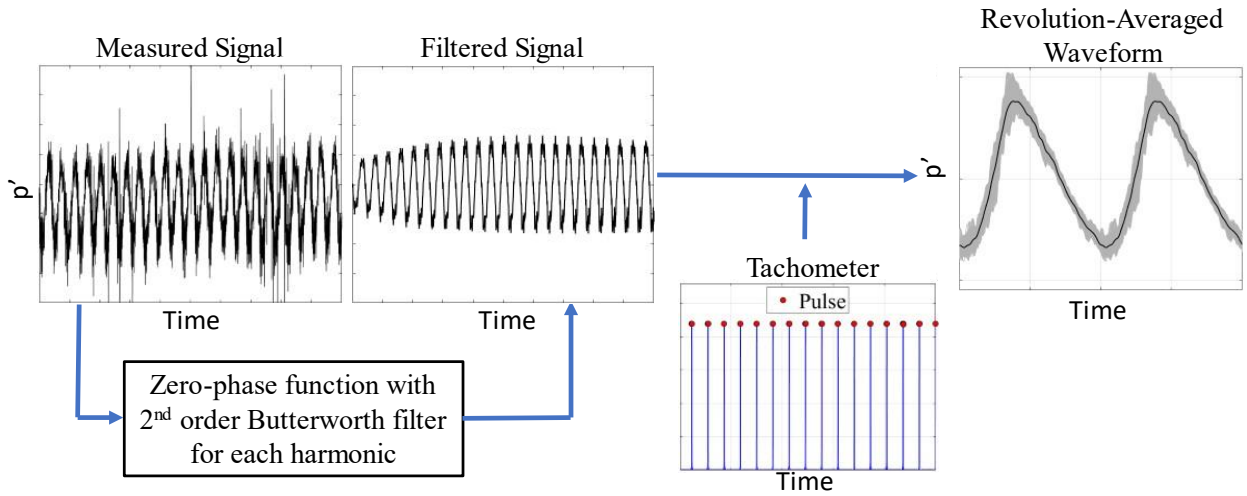


Figure 4: Postprocessing method used to estimate the revolution-averaged waveform.

B. ANN Construction

This experiment aims to build an ANN that can model and predict the acoustic waveform of the RAI for new rotational speeds and rotor-airframe proximities. A secondary goal of the ANN is to separate the steady tonal noise emitted by an isolated rotor from the interaction. Generally speaking, the acoustic waveform emitted from a rotor operating in hover is dominated by the tonal noise at the BPF and its harmonics. Fourier's Theorem states that a periodic function can be represented with a summation of sine and cosine terms. The ANN is trained to model the Fourier coefficients: amplitude, A , and phase, ψ , for each harmonic, k .

Previous work has considered modeling the waveform in the time domain [12]. This approach generates a waveform that closely matches training data. However, ANNs predicting time domain signals will introduce numerical artifacts at higher frequencies and this would contaminate the prediction because the RAI also influences the higher harmonics. Modeling the revolution-averaged waveform in the frequency domain prevents these artifacts from

occurring. It also reduces the training time and number of runs required to train a predictive model. With the Fourier coefficients known, the inverse Fourier transformation can be used to construct the time domain signal and avoid artifacts. A diagram of the frequency domain approach is presented in Fig. 5. For harmonic k , two trained ANNs predict the A and ψ as a function of M_{tip} , Δ/R , θ , and φ . Each ANN architecture consists of two hidden layers using the tansig function, which is a nonlinear activation function that is fed into the linear output layer. Using the ANNs prediction at each harmonic, the waveform can be constructed using the inverse Fourier transformation. The MATLAB Neural Network toolbox is used to train all ANNs.




Figure 5. The ANN prediction architecture and waveform reconstruction.

V. Experimental Results

A. Aerodynamics

Figure 6 shows the coefficient of thrust trends for the MA 11x7 rotor in the hover configuration while incrementing the rotational speed. For the thrust measurements taken in the anechoic chamber, there is a negligible difference between the isolated and RAI runs. Due to a drift in the loadcell, no clear relationship between airframe proximity and thrust coefficient can be established; however, previous studies support the notion that the downstream cylinder has a marginal effect on thrust [1–3]. The coefficient of thrust is positively correlated with the rotational speed because of the Reynolds number effects. The effects are especially pronounced on the rotor performance below 100,000 [13]. For this experiment, the Reynolds number varied between 57,000 – 127,000, with chord length and airspeed properties taken at the 75% span location. Finally, a comparison is made to the University of Illinois at Urbana-Champaign (UIUC) propeller database for the same rotor operating in hover in the UIUC subsonic wind tunnel [14]. The C_T levels deviated by approximately 0.01 between facilities.




Figure 6: Measured hover thrust values for the MA 11x7 rotor collected for the isolated and RAI runs compared to the UIUC database.

B. Rotor-Airframe Proximity

The effects of the rotor-airframe proximity can be seen in Figs. 7a and 7b, where the revolution-averaged waveform for eight runs, operating at a tip Mach number of 0.215 (5000 RPM), have been overlaid for incremental variations in rotor-airframe proximity measured at microphone M1. Figure 7a shows the isolated case in comparison to the three largest rotor-airframe proximities evaluated at an equivalent rotational speed. For the isolated run, the revolution-averaged waveform is dominated by the first harmonic and does not feature any impulsive perturbations. An airframe positioned at a proximity of 0.415R does not deviate from the isolated run. Once the airframe is reduced to a proximity of 0.370R, a small perturbation can be seen at $\pi/4$ and $5\pi/4$, which corresponds to the blade passing over the airframe. A further reduction in proximity to 0.310R increases the strength of the perturbation. Figure 7b overlays the other four runs, which evaluated the effect of the rotor-airframe at proximities from 0.256R to 0.124R. These four revolution-averaged waveforms show that the perturbation grows in strength as the proximity is decreased and, as a result, creates an impulse event that can approach and exceed the steady loading amplitude of the isolated rotor.




Figure 7: Revolution-averaged waveform at $M_{tip} = 0.215$ (5000 RPM) and observer location M1 ($\theta = 90^\circ$, $\phi = -45^\circ$) for (a) large and (b) small rotor-airframe proximities.

Figure 7 shows how the rotor-airframe proximity alters the revolution-averaged waveform and suggests the airframe does not contribute to the noise until it reaches a proximity of 0.370R; however, the individual revolutions show that the RAI also has an unsteady impulse contributing to the noise that is not captured with the ensemble averaging of the signal. In Fig. 8a, the series of isolated runs are overlaid against the revolution-averaged waveform and are well-represented. Figure 8b plots the response measured at M1 (0.370R). An impulse is generated for each blade at each revolution; however, the impulse is not represented in the revolution-averaged waveform. Due to the disparity, there was a concern that the signal had been phase shifted or the motor was unsteady. To address the phase shift concerns, additional methods such as time synchronous averaging, without any filters, and phase averaging were applied to the measured signal. All three methods showed excellent agreement among their revolution-averaged waveforms, and each method retained the disparity between the individual waveforms and the revolution-averaged waveform. Over the 10-second duration of the run, the motor maintains a steady rotational speed of 5023 RPM with a standard deviation of 5.3 RPM. Finally, a spectrogram was used to monitor for recirculation. There were no signs of tonal amplitude variation during this run. Furthermore, the disparity between the individual and revolution-averaged waveforms remained when the averaging was reduced to the first three seconds of the run. This suggests that the impulses generated do not synchronize between revolutions, and the RAI generates unsteady impulses at larger rotor-airframe proximities.

Next, the airframe is repositioned to 0.222R, as shown in Fig. 8c, and the unsteady impulse grows in strength. In addition, the individual impulse events begin to synchronize, and a steady impulse emerges in the revolution-averaged waveform. Once the airframe is brought to a proximity of 0.124R, as shown in Fig. 8d, a strong steady-

state impulse emerges from the waveform. There is no significant disparity between the individual and revolution-averaged waveform.

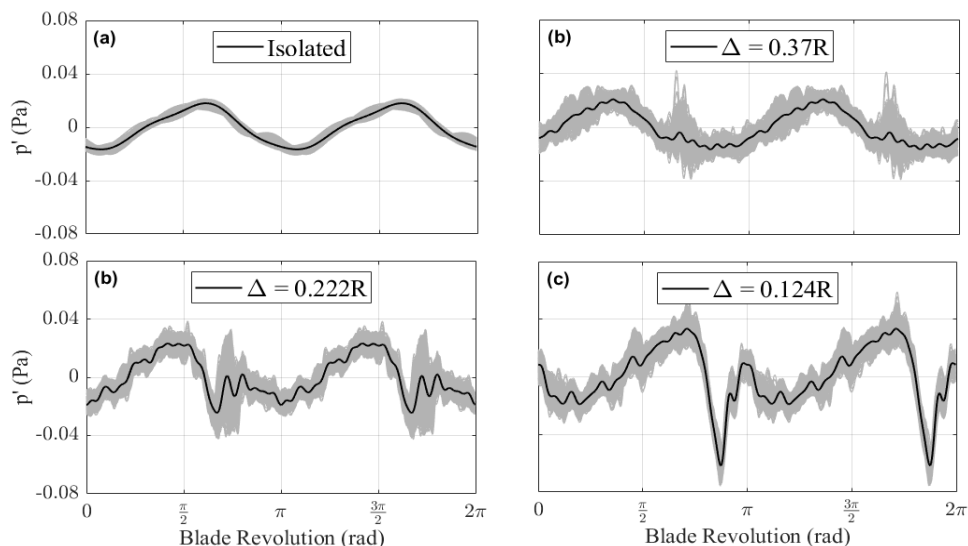


Figure 8: Revolution-averaged waveform compared to the individual waveforms for the (a) isolated run and (b-d) three rotor-airframe proximity runs for M1 ($\theta = 90^\circ$, $\phi = -45^\circ$) at $M_{tip} = 0.215$ (5000 RPM).

C. Emission Angle

The impulse was found to be directional. Thus, a separate section has been devoted to the differences in waveform characteristics at various azimuthal and elevation angles. To exemplify this directivity, Fig. 9 shows a reconstruction of the rotor stands and six microphone locations, comparing an isolated run to an RAI run. The isolated run, operating at a M_{tip} of 0.215 (5000 RPM), is plotted in black for all six locations to act as a reference for the RAI run, which had an airframe positioned at a proximity of 0.124R. The revolution-averaged waveform is plotted in two different colors for this run to denote the azimuth location of the microphones. Microphones M1, M4, and M7 are aligned at an azimuth of $\theta = 90^\circ$ and experienced a negative pressure impulse when the blade, rotating CCW as viewed from above, passed over the airframe. In comparison to the isolated run, the impulse event also creates a slightly positive overshoot when the blade is approaching and departing from the airframe. For microphones M3, M6, and M9, aligned at an azimuth of $\theta = -90^\circ$, a positive pressure impulse event is observed when the blade rotates CCW. Regarding elevation, microphones in-plane of the rotor, $\phi = 0^\circ$, experience a weaker impulse event than the observers below the plane of rotation. In addition, the microphones below the plane of rotation see the individual waveforms synchronize at larger proximities. These behaviors have strong implications for sUAS because a typical drone operates with clockwise and counterclockwise rotors. A compact sUAS would emit a series of positive and negative impulses in quick succession to an observer.

Figure 10 extends the directivity analysis to the microphones aligned at $\theta = 0^\circ$. With microphones M2, M5, and M8 coplanar to the airframe, there are no airframe surface normals directly visible for these three observers. As a result, the revolution-averaged waveforms are nearly identical between runs. This is most apparent for the in-plane microphone, M8, where the waveforms overlap for the entire revolution. At lower elevations, an impulse can be seen in the waveform, but it is significantly weaker than the impulses measured normal to the airframe. The coplanar microphones also have the rotor stand visibly blocking the airframe; however, the stand is likely not reflecting or absorbing the emission for the low- to mid-range frequencies. Instead, the small impulse measured may be emitted from the secondary source of the RAI, the rotor itself [1]. Zawodny and Boyd found similar results for their coplanar microphones. In addition, they could increase the RAI noise for coplanar observers with a conical airframe. The azimuthal variation between the three microphone stands suggests the RAI has a directivity pattern that resembles a dipole, with sound most efficiently radiating normal to the airframe. The directivity pattern has also been observed and discussed in computational results [6].

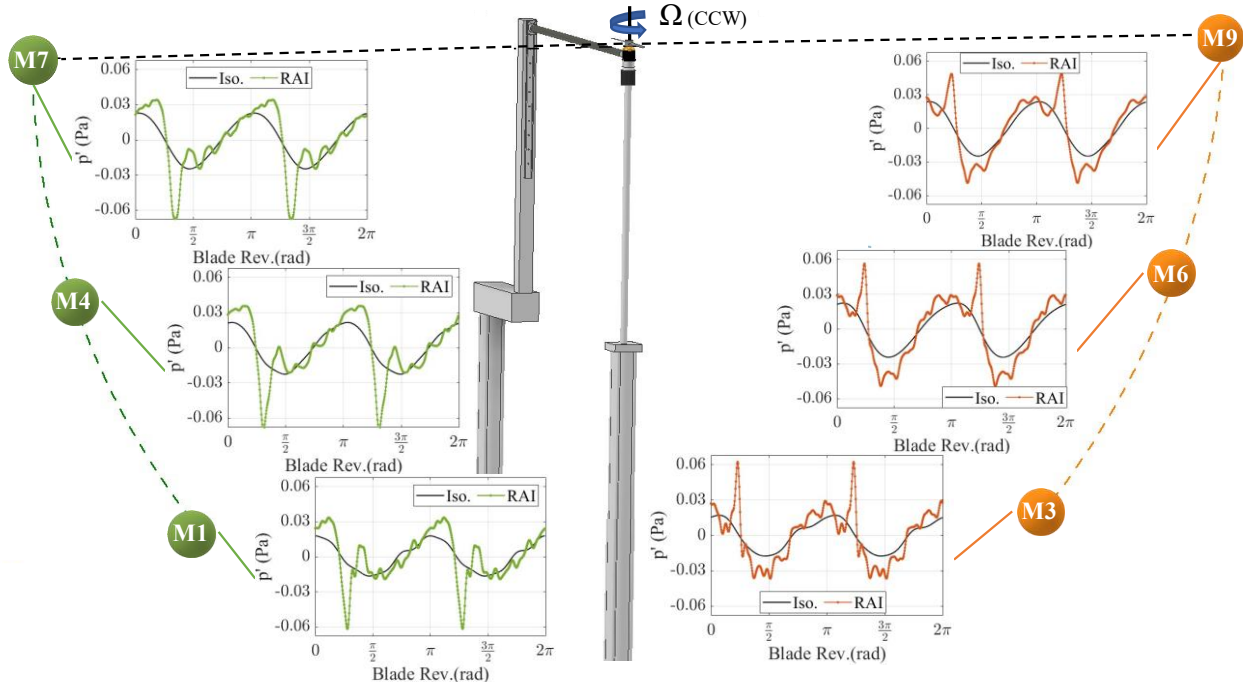


Figure 9: Revolution-averaged waveforms emitted for an RAI run positioned at $\Delta = 0.124R$ compared to an equivalent isolated run at $M_{tip} = 0.215$ (5000 RPM) for six microphones. M7, M4, and M1 are stationed at $\theta = 90^\circ$ while M9, M6, and M3 are stationed at $\theta = -90^\circ$.

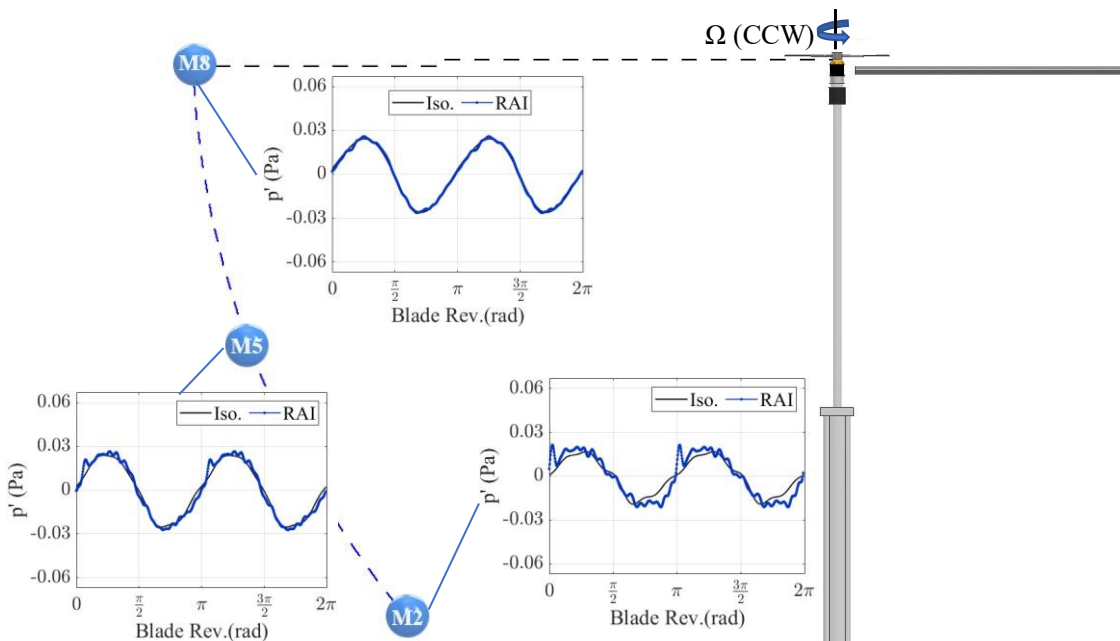


Figure 10: Revolution-averaged waveforms emitted for an RAI run positioned at $\Delta = 0.124R$ compared to an equivalent isolated run at $M_{tip} = 0.215$ (5000 RPM) for three microphones. Microphones M8, M5, and M2 are aligned at $\theta = 0^\circ$.

D. Tip Mach Number

To demonstrate the influence of tip Mach number on the RAI, three proximity runs operating at a tip Mach number of 0.215 (5000 RPM), and 0.404 (9400 RPM) are overlaid in Figs. 11a and 11b, respectively. Figure 11a

shows that the in-plane microphone M7 experiences a negative impulse event that grows in amplitude as the airframe proximity decreases. A weak impulse is generated at a proximity of 0.256R, and this revolution-averaged waveform is dominated by the loading noise emitted by the rotor. As the proximity of the airframe is shifted to 0.124R, the impulse becomes a significant secondary noise source to the rotor noise. Figure 11b shows a similar set of proximities evaluated at the higher tip Mach number. At a proximity of 0.290R, the waveform has minimal impulse noise in the steady signal; however, at a proximity of 0.100R, the RAI impulse grows in amplitude and again the RAI becomes a significant secondary noise source to the rotor noise. The revolution-averaged impulse does not necessarily form due to the rotational speed of the rotor, but the rotational speed scales the impulse.




Figure 11: RAI impulse at M7 ($\theta = 90^\circ$, $\varphi = 0^\circ$) for three airframe proximities at (a) $M_{tip} = 0.251$ (5000 RPM) and (b) $M_{tip} = 0.404$ (9400 RPM). Note the y-axis scaling between figures.

Figures 11a and 11b demonstrate the tip Mach number effect for multiple proximity settings. While insightful, it does not convey the complete interaction between rotational speed and the magnitude of the impulse. Figure 12 demonstrates the coupling effect of the RAI between airframe proximity and tip Mach number for the third, fifth, and ninth harmonic measured at microphone M3. A linear mesh has been placed at each harmonic to provide a visual aid for the amplitude trends. Across all three harmonics, the airframe proximity and the tip Mach number have similar contributions to the RAI. A combination of high tip Mach number and low proximity creates the largest amplitude. While the exact amplitudes differ from microphone to microphone, the trends shown for these three harmonics are representative of the other microphones normal to the airframe ($\theta = -90^\circ$ and $\theta = 90^\circ$). Due to the directivity of the impulse event, an observer below the plane of rotation ($\varphi < 0^\circ$) can perceive a greater number of harmonic excitations than an observer in-plane ($\varphi = 0^\circ$). It should be noted that for the isolated rotor, BPF amplitudes above the third harmonic had a negligible contribution to the revolution-averaged waveform across the range of M_{tip} measured.

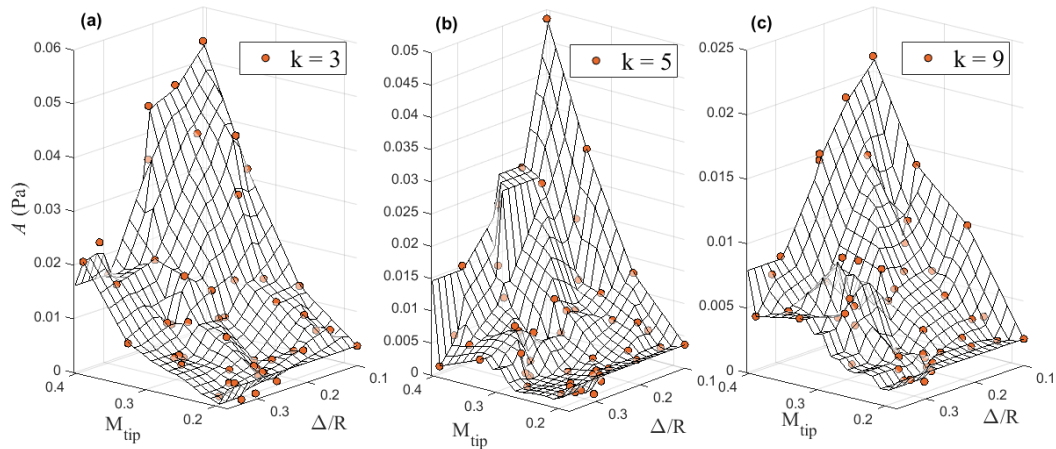


Figure 12: Amplitude trends as a function of tip Mach number and rotor-airframe proximity for the (a) third, (b) fifth, and (c) ninth harmonics for microphone M3 ($\theta = -90^\circ$, $\varphi = -45^\circ$).

VI. ANN Prediction

A. Harmonic Selection and Error Criteria

For these ANNs, the focus is on correlating the RAI impulse with rotor-airframe proximity, tip Mach number, and emission angle. Analysis in the frequency domain showed that the impulse increases the acoustic content at multiple integers of the BPF [1,2,4–6]. To model the RAI noise, an appropriate number of harmonics (amplitude and phase) must be stored in a database so the revolution-averaged waveform can be reconstructed with the inverse Fourier transformation. Runs in which the rotor-airframe proximity was small, tip Mach number was large, and the observer was normal to the airframe generate the strongest impulse and, as a result, require the most harmonics. For these operating conditions, as many as 25 harmonics are required to reconstruct the waveform, as shown in Fig. 13a, with the amplitudes of each harmonic used displayed in Fig. 13b.

The inverse Fourier transform is applied and compared to the revolution-averaged waveform at several harmonic summations. As expected, the first harmonic generates a sine wave with a similar waveform to the isolated runs shown in previous results. After summing the first ten harmonics, a pulse emerges in the signal; however, the impulse event is underestimated. A summation of 20 harmonics brings the impulse closer to the revolution-averaged waveform. Still, the signal is not well approximated until a total of 25 harmonics are used in the summation. Due to the directive nature of the impulse event, for the same run, microphone M8 only required the first two harmonics. Recording and training on 25 harmonics for all microphones across all runs would introduce noise into the training database because the higher harmonics do not always exceed the noise floor. For the purposes of this study, an algorithm has been developed to determine the number of harmonics necessary to closely approximate the impulse within the waveform. The steps used to implement the methodology are explained in the remainder of this section.

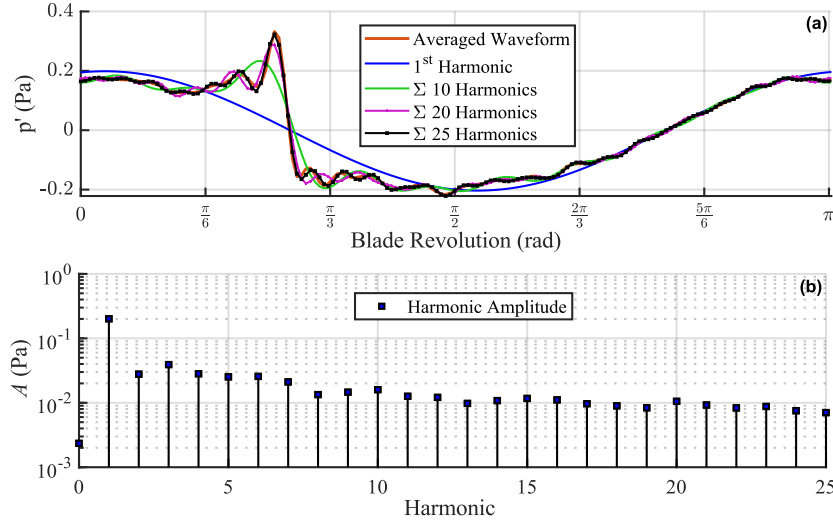


Figure 13: (a) Waveforms constructed with varying harmonic summations are compared to the revolution-averaged waveform measured at M3 ($\theta = -90^\circ$, $\varphi = -45^\circ$) for the run at $M_{tip} = 0.373$ and $\Delta = 0.156R$. (b) The amplitudes used in the summation. Note, this run had a relatively high number of harmonics.

The first step is to extract the amplitude and phase from all harmonics up to the Nyquist frequency from the revolution-averaged waveform, y . Second, an ideal waveform, y_0 , is constructed with an increasing number of Fourier coefficients until the error between the two waveforms is sufficiently small. For this analysis, the error metric used to evaluate the similarities between the two waveforms, y and y_0 , is calculated with the weighted mean square error (WMSE),

$$\text{WMSE} = \frac{\sum_{i=1}^N (y(i) - y_0(i))^2 W_d(i)}{\sum_{i=1}^N W_d(i)}. \quad (1)$$

Where N denotes the number of samples within the waveform and the weighted distribution, W_d , is used to emphasize that the error at the impulse event between y and y_0 should be minimized rather than the error for the entire waveform. The steps taken to locate and construct the weight distribution are shown in Figs. 14a and 14b. In Fig. 14a, a five-point central difference scheme, shown in equation 2, has been applied to the y signal,

$$X_r = \frac{-y(i+2) + 8y(i+1) - 8y(i-1) + y(i-2)}{12}, \quad (2)$$

enabling the steepest gradient to be easily located. A marker has been placed at the impulse locations in Fig. 14a. W_d can be any weight function that concentrates the largest weight values around the impulse location. The Gaussian distribution was chosen for this analysis. Figure 14b shows how the weight distribution is constructed with two Gaussian impulses centered at the two impulse locations found with X_r . Now, the WMSE metric can be used to compare y_0 to y , as illustrated in Fig. 14c. Additional harmonics are superimposed onto y_0 until the WMSE has an error less than $6.25e-4 \text{ Pa}^2$. This error value was chosen after visually inspecting several runs across all parameters of interest. The small error criterion is due to the absolute values of the signal and the weight distribution driving the error toward zero.

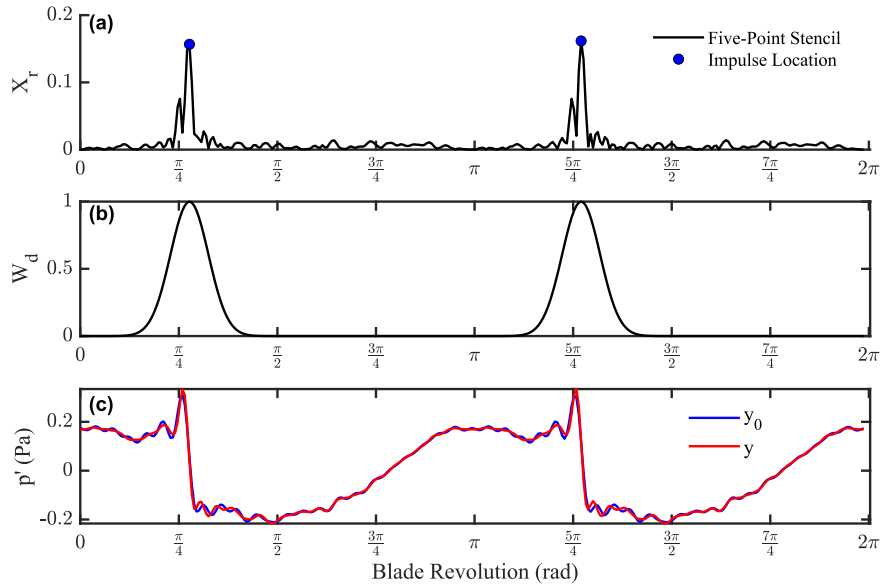


Figure 14: Determining the number of harmonics required to reconstruct the impulse by (a) finding the impulse location, (b) constructing the weight distribution, and (c) overlaying the y_0 onto y that satisfies the error criterion.

B. ANN Prediction of Isolated Rotor Waveforms

For this experiment, thirty runs were conducted at tip Mach numbers ranging from 0.173 to 0.404 (4200 – 9400 RPM). This is equivalent to increments of 200 RPM between each run and provides a large dataset to study the relationship between the number of training runs and the prediction performance. Four datasets composed of 7, 12, 18, and 23 runs were chosen to train an ANN to predict the acoustic waveform for the nine observers. Only the first two harmonics were required to reconstruct the waveform emitted by an isolated rotor. Each dataset was given to an ANN to train on and to ensure equivalent comparisons; the eight remaining runs were reserved for calculating the prediction performance on unseen amplitudes, P_A , and phases, P_ψ , for all nine microphones. Equation 3 is used to calculate the amplitude prediction by evaluating the ANN output, \hat{A}_j , against the equivalent A_j unseen measurement,

while Eq. (4) evaluates the phase offset of the ANN output, $\tilde{\psi}_j$, against the unseen phase measurements, ψ_j . M is the number of reserved runs, and the performance is measured and averaged across all microphone locations.

$$P_A = 100(1 - \frac{1}{M} \sum_{j=1}^M | \frac{A_j - \tilde{A}_j}{A_j} |) \quad (3)$$

$$P_\psi = 100(1 - \frac{1}{M} \sum_{j=1}^M | \angle e^{i(\psi - \tilde{\psi})} |) \quad (4)$$

Two separate equations are used to evaluate the performance; otherwise, using equation 3 to calculate the phase performance would overestimate the error where the phase wrap occurs. For the four datasets, the prediction performance for the first and second amplitude, P_{A1} and P_{A2} , and the prediction performance for the first and second phase, $P_{\psi1}$ and $P_{\psi2}$, which are listed in Table 3. The smallest dataset trained an ANN to predict the main BPF amplitude and phase with 93.61% and 99.27% accuracy, respectively. The second harmonic had a lower prediction performance, with an accuracy of 80.21% and 82.79% for the amplitude and phase, respectively. ANNs that were trained on the 12, 18, and 23 runs had a higher prediction performance; however, this increase is modest.

Table 3: Prediction results for the isolated ANNs trained on four different datasets.

Isolated Study	7 Runs	12 Runs	18 Runs	23 Runs
$P_{\psi1}$	99.27%	99.29%	99.31%	99.32%
$P_{\psi2}$	82.79%	86.13%	89.13%	90.53%
P_{A1}	93.61%	95.42%	96.17%	96.63%
P_{A2}	80.21%	84.85%	86.69%	88.75%

To provide context for Table 3, the ANNs trained on seven runs are evaluated on their ability to predict the unseen conditions in Fig. 15. The amplitude trends for the first and second harmonic are shown in Fig. 15a alongside the seven training runs. Note that the second harmonic does not cover the entire range. Using the WMSE method, the isolated rotor only required the main BPF for low M_{tip} . This prevented the ANN from incorporating the second harmonic until it was deemed necessary to reconstruct the waveform. Figure 15b shows the phase trends for the two harmonics. With the first harmonic acting as a constant, the fit is trivial, and the second harmonic has a good fit for the prediction because the phase has a gradual slope. Once the Fourier coefficients are estimated, the acoustic waveform can be constructed and compared to the unseen revolution-averaged waveforms. Figure 15c shows that the waveforms are nearly identical for a training dataset composed of seven runs. Larger datasets refine the prediction but are ultimately not necessary to build a useful model.




Figure 15: The ANNs, trained on seven runs, predict the (a) amplitudes and (b) phases for unseen conditions for the first and second harmonic. The revolution-averaged waveforms for the unseen runs are compared in (c) to the ANN prediction for the isolated rotor noise measured at $M7$ ($\theta = 90^\circ$, $\varphi = 0^\circ$).

C. ANN Prediction for the RAI Waveforms

The next step is to model the RAI, which is more involved because the number of harmonics required to model any given RAI case could be as little as two or as many as 25. A total of 51 runs were distributed between 35 training and 16 reserved runs for establishing the predictive performance of each ANN. The prediction performance for the amplitudes was calculated with Equation 3, and the phases were calculated with Equation 4. The results are plotted in Fig. 16. Most of the ANNs can make accurate predictions for new amplitudes above 80%, and this is sufficient for modeling the waveform for all nine microphones. The ANN phase trends do not perform as well as the amplitudes in part due to the directivity of the rotor-airframe interaction, which results in a unique phase trend for each observer. Regardless, the ANNs can predict the amplitudes and phases to reconstruct the impulse within the waveform and provide a useful prediction for unseen conditions within the domain of interest.




Figure 16: Prediction performance for the amplitude and phase ANNs for unseen conditions from the set seen in Fig. 3.

The ANNs can be further analyzed by investigating their prediction performance on the nine microphones separately. Table 4 shows the performance at each of the nine microphones, and the amplitude and phase performance are averaged over the harmonics relevant to each microphone. This means microphone M8 is averaged over the first two harmonics while M3 is averaged over all 25. This makes it difficult to compare performance between microphones, but the table shows that the amplitude prediction performance was greater than 80% for 8 of the 9 microphones. In contrast, only 4 of the 9 P_{ψ} exceed the 80% mark. M6 has a P_{ψ} less than 60%. The rest of the analysis will focus on evaluating the capabilities and limitations of the ANNs.

Table 4: Prediction performance for the RAI ANNs across all nine microphones.

	M1	M2	M3	M4	M5	M6	M7	M8	M9	Ave.
P_A	88.2%	87.6%	88.7%	86.5%	71.7%	87.0%	88.0%	93.9%	82.8%	86%
P_{ψ}	70.8%	73.0%	70.9%	83.3%	64.3%	59.7%	82.3%	92.8%	81.4%	75%

Figure 17 provides an example in which the ANN is given a run when an impulse is expected to occur. For this run, all 25 amplitude and phase ANNs are called, and estimate the impulse for their respective harmonic. Figure 17a shows the amplitudes for the revolution-averaged waveform measured at microphone M3 alongside the predicted amplitudes made with the ANN. These results closely overlap and show that the amplitude is well modeled. Figure 17b compares the revolution-averaged phase measured at M3 to the predicted phase estimated with the ANN. The inverse Fourier transform was applied to the measured run and the ANN prediction and plotted in Fig. 17c. These signals closely overlap, and the impulse is captured in the waveform.

While most ANNs estimate a phase that closely aligns with the measured value, several ANNs show small phase lags and phase leads in their estimation. For example, the second harmonic in Fig. 17b predicts a phase of 4.57 rad while the measured phase is 3.7 rad. This results in a phase lag of 0.87 rad or 50° . This is concerning because part of the interest in modeling the waveform is to provide a tool for auralization. Studies have found that small phase differences can result in artifacts while listening to reconstructed signals [15]. Additional work is required to quantitatively understand the acceptable error for phase offset in the higher harmonics.

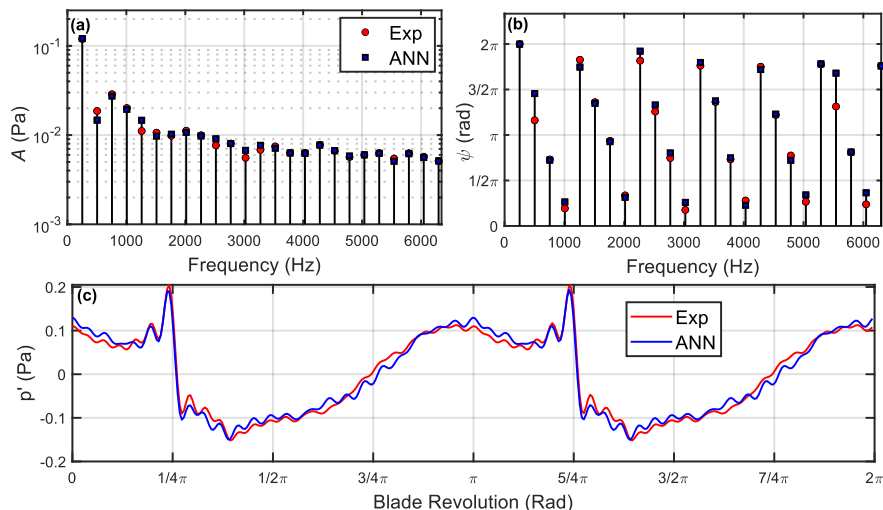


Figure 17: Comparing the experimental measurements and ANN prediction at microphone M3 ($\theta = -90^\circ$, $\varphi = -45^\circ$) operating at $M_{tip} = 0.325$ with a proximity of $\Delta = 0.14R$ for (a) amplitude, (b) phase, and (c) the resultant waveform.

Figure 18 evaluates the prediction made for microphone M1 for the same tip Mach number and blade proximity that was made for M3 to determine if the ANN can predict a negative impulse. Figures 18a and 18b show that only fifteen harmonics were required for this case. Impulses for both cases have a similar amplitude but excite the spectrum in different ways. The Fourier coefficients are well modeled for this observer and show that the ANN can predict positive and negative impulses. While the ANNs have some limitations in capturing the exact phase, the predictions are still useful in estimating the impulse for multiple observers.

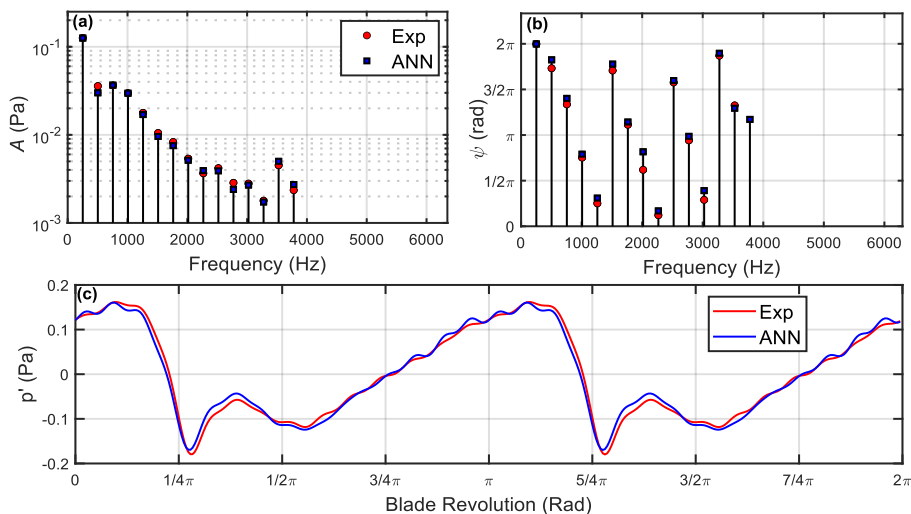


Figure 18: Comparing the experimental measurements and ANN prediction at microphone M1 ($\theta = 90^\circ$, $\varphi = -45^\circ$) operating at $M_{tip} = 0.325$ with a proximity of $\Delta = 0.14R$ for (a) amplitude, (b) phase, and (c) the resultant waveform.

Figures 17 and 18 demonstrate that the system of ANNs can predict the impulse for two observers for one run. While insightful, an overview across the training domain would provide a better representation of the predictions made by the ANN. Figure 19 provides the ANN trends for the amplitude at M3 for the third, fifth, and ninth harmonics. For the three amplitude trends, the ANNs have constructed a good fit across the domain space and reliably estimated the amplitude for unseen conditions. Figure 19 also provides the ANN fit for the phase at these three harmonics and captures the overall trends, which has a more complex response to the independent variables. It should be noted that the higher harmonics are more sensitive to adjustments in tip Mach number and proximities.

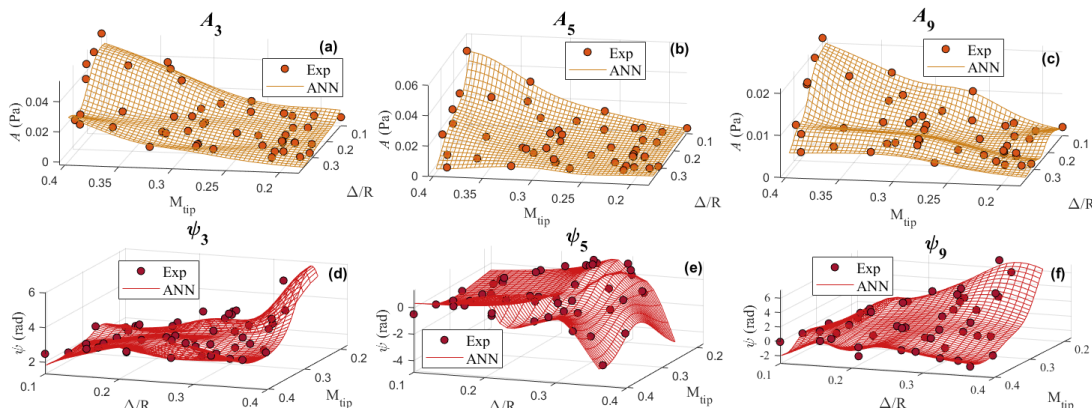


Figure 19: Amplitude trends for microphone M3 ($\theta = -90^\circ$, $\varphi = -45^\circ$) are provided for the (a) third, (b) fifth, and (c) ninth harmonics along with the phase at the (d) third, (e) fifth, and (d) ninth harmonics.

The ANN prediction trends made for microphone M3 are repeated for microphone M9. The third harmonic amplitude trends for M9 are shown in Fig. 20a and are nearly identical to M3 (seen in Fig. 19a). The significance of elevation angle becomes apparent at the higher harmonics, as seen in Figs. 20b and 20c, where only the proximities between 0.25R and 0.1R were necessary to reconstruct the waveform. Phase predictions for the third, fifth, and ninth harmonics are given in Figs. 20d, 20e, and 20f, respectively. Figure 20d shows that the third harmonic has a different scale and trend when compared to Fig. 19d, although both lie on the same azimuth with respect to the rotor stand. Contrast this with the third harmonic amplitudes found in Figs. 20a and 19a, which have similar trends and minor differences in amplitudes.

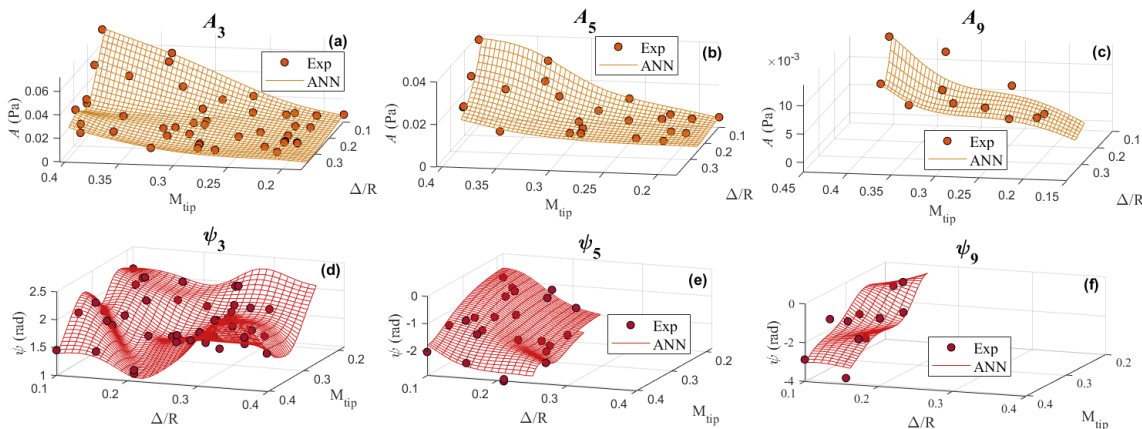


Figure 20: Amplitude trends for microphone M9 ($\theta = -90^\circ$, $\varphi = 0^\circ$) are provided for the (a) third, (b) fifth, and (c) ninth harmonics along with the phase at the (d) third, (e) fifth, and (d) ninth harmonics.

By analyzing the RAI in the frequency domain, the impulse noise in the revolution-averaged waveform can be modeled across the airframe-proximity and tip Mach number domain of interest for the nine microphone locations. As shown in Table 5, the average prediction performance for the amplitude and phase across all microphones was 86% and 75%, respectively. This performance sufficiently approximates the revolution-averaged waveform for the

unseen condition at different emission angles. While the ANNs incorporate elevation and azimuth angle, the spatial domain had insufficient observers to model directivity. Due to the experimental setup, this analysis is only valid for the MA 11x7 rotor operating in hover upstream of a 25 mm carbon fiber rod. Despite the limitations, this system of ANNs provides a low-fidelity means of estimating the acoustic emission of the impulse noise.

D. Isolating the Impulse

Two systems of ANNs have been trained to model a two-bladed rotor in hover. The first system is comprised of four ANNs that model the first two Fourier coefficients emitted by an isolated rotor, and the second system is comprised of a total of fifty-one ANNs that model the first 25 Fourier coefficients for the two-bladed rotor operating upstream of a cylindrical airframe. Since the steady loading noise of the rotor has a negligible contribution at higher harmonics, the RAI contribution at these higher harmonics can be separated from the tonal noise emitted by the rotor. This process has been carried out in Figs. 21a and 21b for microphone M6. The isolated system predicts the Fourier coefficients for the first two tones at $M_{tip} = 0.30$ (7000 RPM), and the second system estimates the Fourier coefficients for the appropriate number of harmonics at six airframe proximities between 0.3 and 0.1. Once the predictions are made, the inverse Fourier transform is applied to each case, and the resulting waveforms are overlaid in Fig. 21a. Next, the tonal noise emitted by the isolated rotor was subtracted from the RAI waveforms. The separated impulses are plotted in Fig. 21b with their corresponding proximity. As the airframe proximity decreases, the observer will experience the steepest drop in acoustic pressure when the blade is parallel with the airframe, and the departing blade brings a sharp rise in acoustic pressure [1].

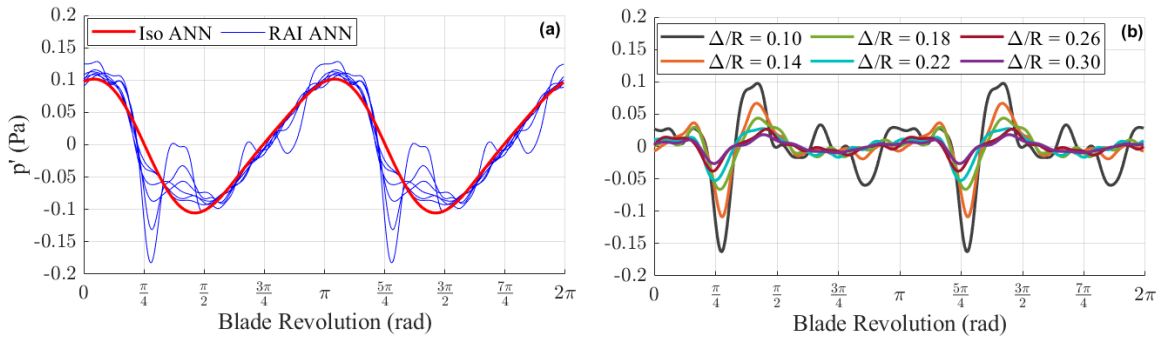


Figure 21: (a) Isolated and RAI ANNs predict the waveform for microphone M6 ($\theta = -90^\circ$, $\varphi = -22.5^\circ$) at $M_{tip} = 0.30$ for six airframe proximities. (b) Deconstructive interference is used to separate the impulse noise from the total.

The process was carried out for all microphones normal to the airframe in Fig. 22. Results for microphones M9, M6, and M3 are seen in Figs. 22a, 22b, and 22c, respectively. The microphones below the plane of rotation experience a greater impulse than the in-plane microphone. Similar observations can be made for the microphones on the opposite side of the rotor test stand as results for microphones M7, M4, and M1 are shown in Figs. 22d, 22e, and 22f, respectively. The ANNs predict the impulse with similar amplitudes and opposite phases. The system of ANNs can predict and reconstruct the impulse within the waveform and acts as a low-fidelity tool that estimates the interaction noise over the domain of interest.

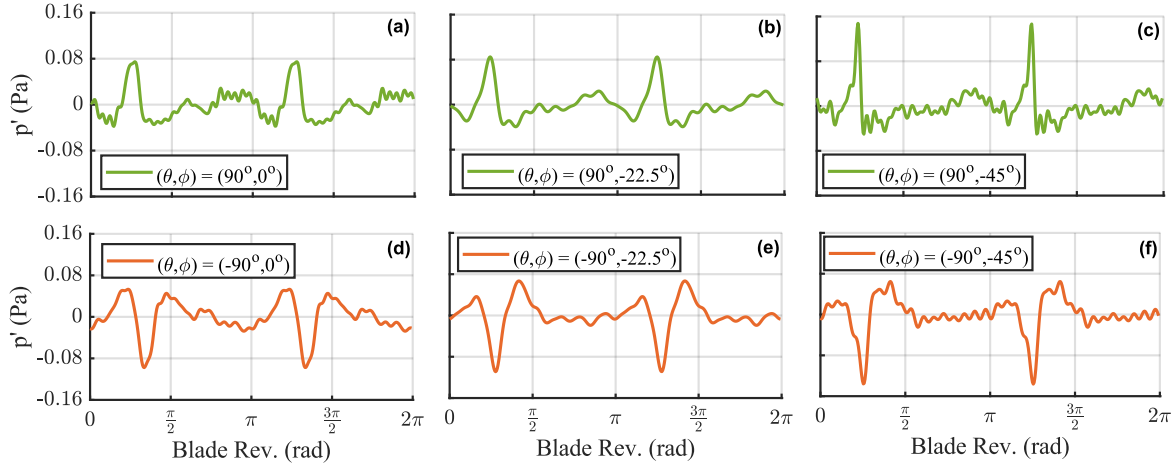


Figure 22: A prediction is made with the system of ANNs to estimate the impulse event at $\Delta = 0.14R$ and $M_{tip} = 0.30$ for observers located at (a) M7, (b) M4, (c) M1, (d) M9, (e) M6, and (f) M3.

VII. Conclusion

The rotor-airframe interaction is a category of installation noise that results in harmonic excitation. An acoustic impulse is produced as the rotor blade sweeps over the airframe component. The experiment shows that the strength of the impulse is dependent on the rotor-airframe proximity, tip Mach number, and emission angle. For configurations where the rotor-airframe proximity was small and tip Mach number was large, the observer normal to the airframe experienced the greatest increase in harmonic content relative to an isolated rotor. Results showed that the revolution-averaged impulse had a negligible contribution to the waveform when stationed at a proximity of $\Delta = 0.37R$ and above; however, an unsteady contribution from the interaction was observed at the largest proximities measured. In addition, tip Mach number positively correlated with harmonic amplitude, and the strength of the perceived impulse noise depended on azimuth and elevation angle. A system of artificial neural networks was trained to predict the rotor-airframe interaction noise and tonal noise by modeling up to the first 25 amplitudes and phases at the blade passing frequencies. The system of ANNs was able to predict the amplitudes and phases with an accuracy of 86% and 75%, respectively, across all microphones. Qualitative assessment of the reconstructed time domain waveforms displays good agreement with the measured revolution-averaged signals and provides confidence in the artificial neural network approach.

References

- [1] Zawodny, N. S., and Boyd, D. D., "Investigation of Rotor–Airframe Interaction Noise Associated with Small-Scale Rotary-Wing Unmanned Aircraft Systems," *Journal of the American Helicopter Society*, Vol. 65, 2020, pp. 1–17. doi.org/10.4050/JAHS.65.012007
- [2] Wang, Z., Henricks, Q., Zhuang, M., Pandey, A., Sutkowy, M., Harter, B., McCrink, M., and Gregory, J., "Impact of Rotor–Airframe Orientation on the Aerodynamic and Aeroacoustic Characteristics of Small Unmanned Aerial Systems," *Drones*, Vol. 3, No. 3, 2019, p 56. doi.org/10.3390/drones3030056
- [3] Henricks, Q., Wang, Z., and Zhuang, M. "Small-scale rotor design variables and their effects on aerodynamic and aeroacoustic performance of a hovering rotor," *Journal of Fluids Engineering*, Vol. 142, No. 8, 2020. doi.org/10.1115/1.4046872
- [4] Whelchel, J., Alexander, W. N., and Intaratep, N., "Propeller Noise in Confined Anechoic and Open Environments," *AIAA SciTech 2020 Forum*, 2020. doi.org/10.2514/6.2020-1252
- [5] Whelchel J. and Alexander, N. W., "sUAS Rotor-Airframe Interaction," *AIAA Aviation 2021 Forum*, 2021. doi.org/10.2514/6.2021-2212
- [6] Gojon, R., Doué, N., Parisot-Dupuis, H., Mellot, B., and Jardin, T., "Aeroacoustic radiation of a low Reynolds number two-bladed rotor in interaction with a cylindrical beam," *28th AIAA/CEAS Aeroacoustics 2022 Conference*, 2022. doi.org/10.2514/6.2022-2972
- [7] Gallo, E., Zarri, A., Bampanis, G., Fasinella, M. C., and Schram, C.F., "Simplified models for propeller potential interaction noise," *28th AIAA/CEAS Aeroacoustics 2022 Conference*, 2022. doi.org/10.2514/6.2022-2974
- [8] Goldstein, M.E., "High frequency sound emission from moving point multipole sources embedded in arbitrary transversely sheared mean flows," *Journal of Sound and Vibration*, Vol. 80, No. 4, 1982, pp. 499-522. [doi.org/10.1016/0022-460X\(82\)90495-3](https://doi.org/10.1016/0022-460X(82)90495-3)

- [9] Grosveld, F., "Calibration of the structural acoustics loads and transmission facility at NASA Langley Research Center," INTER-NOISE and NOISE-CON Congress and Conference Proceedings, Institute of Noise Control Engineering, 1999, pp. 1541–1546.
- [10] Stephenson, J.H., Weitsman, D., and Zawodny, N.S., "Effects of flow recirculation on unmanned aircraft system (UAS) acoustic measurements in closed anechoic chambers," *The Journal of the Acoustical Society of America*, Vol. 145, No. 3, 2019, pp. 1153-1155. doi.org/10.1121/1.5092213
- [11] JMP®, Version Pro 16. SAS Institute Inc., Cary, NC, 1989–2021.
- [12] Wiedemann, A.D., Fuller, C., and Pascioni, K., "Constructing a physics-guided machine learning neural network to predict tonal noise emitted by a propeller," INTER-NOISE and NOISE-CON Congress and Conference Proceedings, Vol. 264, 2022, pp. 151–162. doi.org/10.3397/NC-2022-709
- [13] Deters, R.W., Ananda, G.K. and Selig, M.S., "Reynolds Number Effects on the Performance of Small-Scale Propellers," 32nd AIAA Applied Aerodynamics Conference, 2014, doi.org/10.2514/6.2014-2151
- [14] Brandt, J., Deters, R., Ananda, G., Dantsker, O. and Selig M., Downloaded January 2023, *UIUC Propeller Database, Vols 1-4*, University of Illinois at Urbana-Champaign, Department of Aerospace Engineering, retrieved from <https://mselig.ae.illinois.edu/props/propDB.html>
- [15] Krishnamurthy, S., Tuttle, B.C. and Rizzi, S.A., "A Synthesis Plug-in for Steady and Unsteady Loading and Thickness Noise Auralization," *AIAA Aviation 2020 Forum*, 2020. doi.org/10.2514/6.2020-2597