Annoyance Model Assessments of Urban Air Mobility Vehicle Operations

Stephen A. Rizzi[1], Andrew W. Christian[2], and Stefan J. Letica[3]
NASA Langley Research Center, Hampton, VA 23681, USA

and

Shane V. Lympany[4]
Blue Ridge Research and Consulting, LLC, Asheville, NC 28801, USA

As urban air mobility (UAM) vehicles begin to enter service, consideration must be given to vertiport siting and flight routing to help reduce community noise impact and promote adoption by the community. Notwithstanding environmental justice concerns, a possible early strategy is to operate out of existing heliports and fly along established helicopter routes, many of which follow uninhabited waterways and/or roadways (with the thought for the latter that an already noisy ambient environment will mask the sound of the aircraft). However, no prior annoyance model has been fielded that both takes audibility into account and that may be applied to such a real-world strategy. This paper reviews a recently developed annoyance model that includes audibility as a factor and applies it to two flight operations. A simple overflight case is first undertaken to demonstrate the approach. A point-to-point operation in the New York City area is then considered to demonstrate how annoyance varies across an urban soundscape. The cases considered use modeled UAM vehicle noise propagated to a set of ground observers as the signals and either recorded or modeled ambient acoustic data as the maskers.

I. Introduction

CURRENT aircraft noise regulations in the United States, whether for aircraft type certification [1] or for community noise [2], do not take into account the other environmental noise ("background" or "ambient") that may be present in their operational environments. This is perhaps sufficient when the aircraft noise is the prominent source, for example, in the vicinity of airports, but may be insufficient when the aircraft noise blends into the existing soundscape. Such could be the case for some urban air mobility (UAM) vehicles that are intended to operate across cities transporting people and goods. In these situations, the noise of the vehicle (the signal) may be "masked" by the preexisting background sounds (the masker). Past psychoacoustic testing has shown that sounds are less annoying when they are masked. Christian [3] provides a review of contemporary knowledge of this effect.

The purpose of this paper is to assess annoyance to UAM vehicle operations over a representative community using a recently developed annoyance model in which the effect of masking is taken into account. No prior research has fielded a model of annoyance to aircraft noise that includes audibility as a factor and that is extensible to real-world conditions. A brief review of the annoyance model is first provided. It is possible to exercise the annoyance model entirely with recordings of the signal and the masker. It is not practical, however, to make recordings over a wide area to produce annoyance maps. Instead, signal data are generated by acoustic analyses consisting first of prediction of source noise hemispheres that are subsequently flown in simulation to produce noise estimates over a large grid of ground observers. Next, two types of ambient noise data are introduced. Recordings are used to demonstrate how temporal variations in the masker affect audibility and annoyance for a straight and level flyover. Modeled ambient acoustic data are subsequently used to demonstrate the effect of spatial variations in the masker over a fictitious point-to-point route from lower Manhattan, New York City, to nearby Newark Liberty International Airport. This work is aligned with a recommendation of a recent white paper [4] that "*Validated models for audibility,*

[1] Senior Researcher for Aeroacoustics, Aeroacoustics Branch, stephen.a.rizzi@nasa.gov, AIAA Fellow
[2] Research Engineer, Applied Acoustics Branch, andrew.christian@nasa.gov
[3] Research Engineer, Aeroacoustics Branch, stefan.j.letica@nasa.gov
[4] Principal Engineer, Shane.Lympany@BlueRidgeResearch.com, AIAA Member

noticeability, and annoyance to UAM aircraft noise be developed to assess their utility for assessing community noise impact," and further that "*such models be developed and validated over a wide range of operating conditions and demand scenarios, taking into account a representative range of ambient/background conditions*."

II. Annoyance Model in the Presence of Masking

When developing a model for annoyance in the presence of masking, it is helpful to consider the edge cases of the response function. When the signal is prominent over an ambient, i.e., it is clearly audible, there is not a strong effect of the ambient noise on annoyance. This finding is supported by the meta-analysis of Fields [5]. Conversely, when a signal is at a low enough level that it is at least partially masked by ambient sounds, a reduction in annoyance to the signal has been observed in both laboratory [6,7] and community [8] studies.

These observations form the basis for a model developed by Christian [9] that uses unweighted one-third octave band (OTOB) sound pressure level (L) data to estimate detectability index (d' or d-prime) and annoyance. The "discounted" annoyance model is based on a combination of those estimates. d' is a measure of the detectability of a signal that is useful because it can apply to many different scenarios, i.e., it treats data from laboratory and real-world tests equally. As a rough guide, a d' value of about 1 indicates a signal level that is so close to a masker level that it is only (confidently) heard about 50% of the time. A doubling of d' corresponds to roughly a 3 dB increase in the signal level. The discounted sound pressure level, that is, the reduced level due to the presence of the masker, is computed at each OTOB and time step according to the relation:

$$L_{i,t}\big|_{Disc} = L_{i,t} - \frac{\alpha}{\left(\frac{d'_{i,t}}{\delta}\right)^{\rho}}, \tag{1}$$

in which $L_{i,t}$ is the unweighted sound pressure level and $d'_{i,t}$ is the detectability index of the i^{th} one-third octave band at time step t, respectively. The parameters α, δ, and ρ of the discount function are determined through human subject testing [7] and govern how the discount impacts the prediction of annoyance as d' gets smaller and approaches zero. For convenience, α is set to 3 dB, a level at which the discount defines a halving of the power of the target sound (similar to the definition of the bandwidth of a filter). The value δ then defines the d' value where this 3 dB "knee" will be encountered. The parameter ρ defines the discount "rate," or the manner in which the function rolls off as d' gets smaller. δ and ρ can therefore be manipulated to produce nearly any possible roll-off of the discount function in line with the two edge cases described above. In this work, values of δ = 14 and ρ = 1 are used and represent the median values of model parameters determined on a subject-by-subject basis from a recent test using a novel test data analysis method [10]. The discount function using these model parameters is represented by the blue trace in Figure 1.

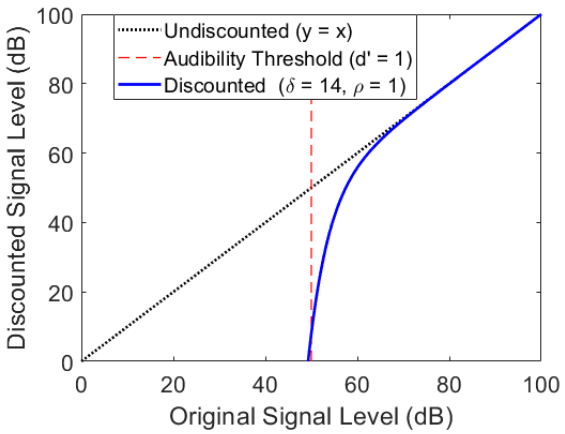

Figure 1: Relationship of original and discounted signal levels with model parameters determined through human subject testing.

The discounting function is (ideally) independent of the method used to compute d'. In this work, d' is predicted using a model that is based on OTOB data and those data are specified at 0.5 s intervals. This allows standardized time-integrated aircraft noise metrics [1] (in particular, the A-weighted sound exposure level L_{AE}) to serve as baseline predictors of annoyance in the absence of masking. The use of OTOB data is probably the simplest approach able to

yield a reasonable prediction of detectability. Many more-accurate detection algorithms need at least a narrowband spectrum as input, if not a complete time history of the signal and masker. Thus, only a simple "power spectrum model of masking" can be brought to bear on such OTOB data. The approach is meant to produce a computationally reasonable audibility prediction program. This model is effectively a signal-to-noise ratio that is augmented with several empirical factors based on human hearing characteristics. For an OTOB of index i, and time step of index t, $d'_{i,t}$ is computed as:

$$d'_{i,t} = \mathbb{D}_i \cdot \frac{s_{i,t}}{n_{i,t} + e_i} \tag{2}$$

in which $s_{i,t}$ is the unweighted OTOB sound power of the signal (in mean-squared pressure), $n_{i,t}$ is the OTOB masker power, e_i is the derived "equivalent auditory system noise" sound power that encodes the auditory threshold, and $\mathbb{D}$ is a frequency-dependent detection efficiency factor that contains corrections for OTOB and auditory filter bandwidths. The derivation and origin of the components of Eq. (2) are given in the Appendix.

In practice, the discounted annoyance model starts with an OTOB spectrogram of the signal and the masker at each observer point. In this way, both the signal and the masker may vary in time, e.g., an aircraft flyover under a steady or changing operating condition with a time-varying ambient. Next $d'_{i,t}$ is computed using Eq. (2) to generate a d' spectrogram and discounted sound pressure levels are calculated using Eq. (1). Finally, A-weighting is applied and the discounted data are summed over frequency to obtain the overall discounted A-weighted levels at each time step, L_A, and summed over time to obtain discounted L_{AE}. A dose-response relation between annoyance and L_{AE} may then be used to estimate mean annoyance to the masked and unmasked signals (see Section V.A).

Referring again to the edge cases, at high signal levels, the signal is very prominent over the masker and hence the annoyance estimated by the model is independent of the masker. As the level of the signal decreases, there is a corresponding reduction in estimated annoyance. If the signal is lowered so far that it becomes at least partially masked, the estimated annoyance to the signal will be reduced relative to the no-masker condition. Eventually, when the signal is entirely masked, the estimated annoyance to it is negligible. The discounted annoyance is expected to be a more effective predictor of annoyance in the presence of a masker than annoyance based on the original L_{AE} generated from the signal alone. It should be noted that although this work is motivated by, and directed at, annoyance to UAM aircraft noise in the presence of a masker, it is applicable in principle to any other transportation noise source.

III. Aircraft Signal Generation

The quadrotor reference vehicle developed under the NASA Revolutionary Vertical Lift Technology (RVLT) Project is chosen as the UAM vehicle in this investigation, see Figure 2. It is sized for a 1200 lb. (544 kg) payload (up to six passengers) executing a representative mission profile [11]. The quadrotor is an all-electric variant with a gross weight of 6469 lb. (2934 kg) and a maximum airspeed V_{max} of 109 knots true airspeed (KTAS) (202 km/h). Its four three-bladed rotors operate under collective pitch control at a 20 Hz blade passage frequency (BPF). Additional details on this configuration can be found in Silva et al. [12].

Figure 2: NASA quadrotor RVLT reference vehicle configuration considered in this study.

The processes for predicting flyover noise data were discussed in detail by Rizzi et al. [13] and are briefly summarized here. The periodic loading and thickness noise and broadband self-noise components were computed for 42 operational states defined by pairs of airspeed (knots) and climb angle (degrees), using the process depicted in Figure 3. The resulting source noise hemispheres were 'flown' in simulation using the NASA second generation Aircraft Noise Prediction Program (ANOPP2) Mission Analysis Tool (AMAT) [14]. The AMAT can simulate a flyover using one or more source noise hemispheres along a user-defined trajectory. At each observer location, the AMAT generates an OTOB spectrogram at 0.5 s intervals. In this work, a new capability in the AMAT found in ANOPP2 v1.5, referred to as explicit observer time, was used to synchronize the reception time across a set of ground observers. The simulations were performed to represent the signal at a dense grid of 4 ft. (1.2 m) observer locations

spaced 98.4 ft. (30 m) apart. The resulting OTOB spectral time histories, in the OTOB center frequency range from 31.5 Hz to 4 kHz, serve as the signal inputs in Eq. (2).

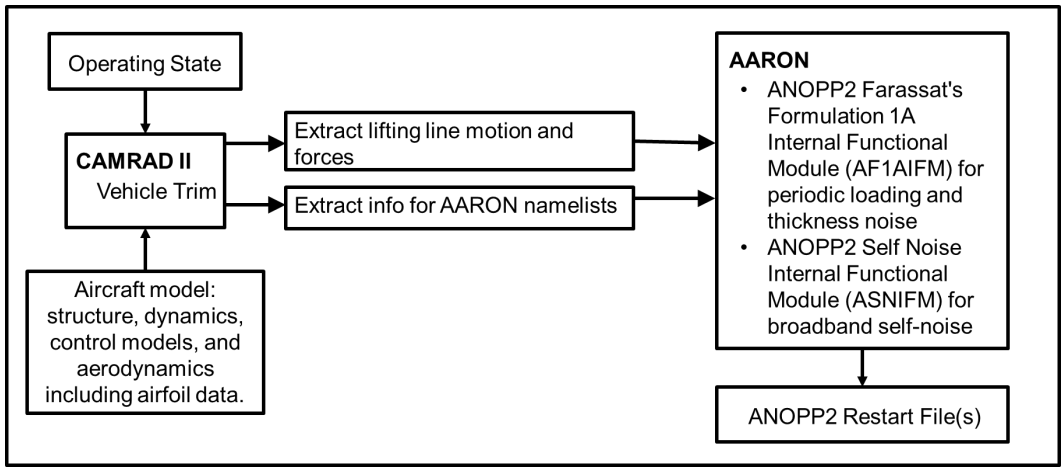

Figure 3: NASA process for generating source noise data for each operating state.

It should be noted that although the area of interest is a very large metropolitan area with many high-rise buildings and other geographical features, a flat ground plane with uniform impedance is modeled herein. More complicated modeling of atmospheric propagation in urban environments is beyond the scope of this paper. Additional simulation details may be found in Letica and Rizzi [15].

IV. Masking Noise

Both recordings and modeled ambient noise data are used in this work. Recordings are used for assessing audibility and annoyance for masking noise that varies over short time scales, i.e., seconds or minutes. Modeled acoustic data corresponding to daytime and nighttime were used for assessment over a wide area of observers that would not be practical using recordings.

A. Recorded Masker Noise

The recordings were made in New York City in September 2018. One recording was made in City Hall Park in lower Manhattan and the other was made in the Sheep Meadow in Central Park. The recording sites are shown in Figure 4. The recordings were made using a Zoom H4n field recorder with an Earthworks M23 microphone vertically oriented 4 ft. (1.2 m) above the ground. Time-averaged narrowband spectra for both locations at a resolution of 1 Hz are shown in Figure 5. The levels are higher at City Hall than at Central Park, with the difference increasing with increasing frequency.

Figure 4: Photographs of recording sites in Central Park (left) and near City Hall (right) [Source: NASA].

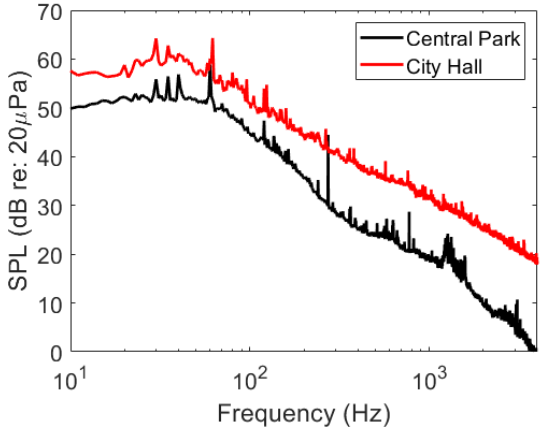

Figure 5: Time-averaged narrowband spectra of masker noise near City Hall & Central Park.

The difference in levels is also apparent in the OTOB spectrograms shown in Figure 6. The City Hall masker ranges from about 75 dB at the lower frequency limit to about 45 dB at the upper frequency limit. In contrast, the Central Park masker ranges from about 60 dB at the lower frequency limit to about 20 dB at the upper frequency limit. Segments of each recording are available for download [16].

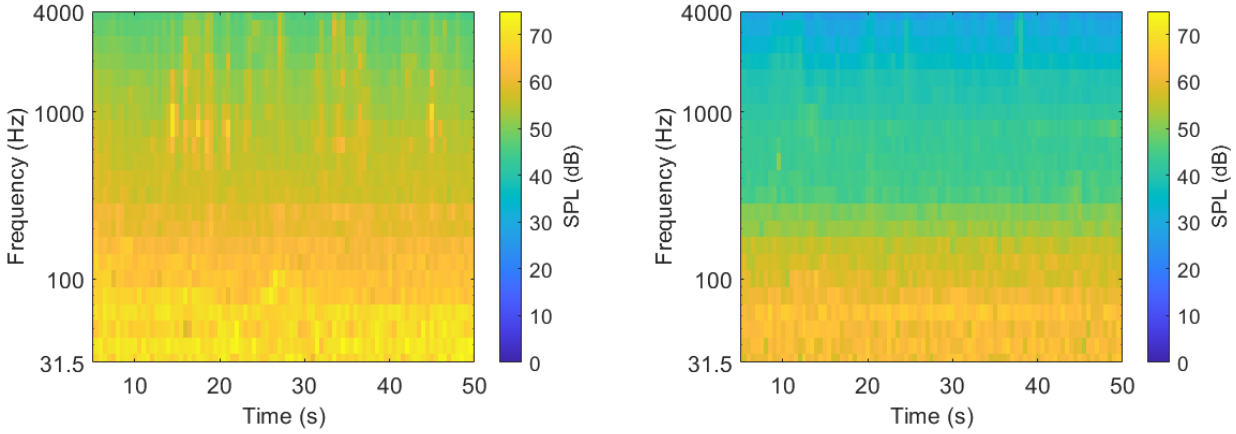

Figure 6: Spectrograms of masker sound pressure level at City Hall (left) and Central Park (right).

B. Modeled Masker Noise

The modeled ambient sound levels were generated by the Blue Ridge Research and Consulting (BRRC) AMBIENT model, a physics-informed machine learning model of ambient soundscapes [17]. The ambient soundscape, or acoustic environment, is composed of anthropogenic, biological, and geophysical sounds. The AMBIENT model generates spatially, temporally, and spectrally varying maps of the ambient sound levels produced by all anthropogenic, biological, and geophysical sources across various environments. Here, the model was applied to predict average ambient sound levels across a selected region of the New York City metropolitan area between lower Manhattan and Newark, NJ.

1. Mapping Ambient Sound Levels Using Physics-Informed Machine Learning

The BRRC AMBIENT model combines the strengths of machine learning with the strengths of physics-based models. Physics-based models are appropriate for sound sources with known locations and characteristics, such as traffic noise, whereas machine learning models are appropriate for sound sources that cannot be modeled physically, such as chirping birds. The AMBIENT model builds upon the geospatial modeling approach developed by Mennitt, et al. [18], who applied machine learning regression algorithms to model the statistical relationship between ambient sound levels and geospatial variables, and extends their approach by incorporating physics-based traffic noise predictions as a geospatial feature.

In the AMBIENT model, the ambient sound level at any location is modeled as a function of the geospatial features and the predicted traffic noise at that location,

$$\text{ambient sound level} = f(\text{geospatial features, physics-based noise}) \tag{3}$$

in which the function, f, is unknown. The unknown function f is determined by fitting an ensemble of machine learning regression models to the data at locations where the geospatial features, predicted traffic noise, and ambient sound levels are known. Once the ensemble of regression models is fitted, it is applied to predict ambient sound levels at any locations where the geospatial features and predicted traffic noise are known.

The AMBIENT model uses a training dataset of measured ambient sound levels to fit the regression models. As shown in Figure 7, the database of ambient acoustic measurements includes 962 unique sites across North America measured by the National Park Service (NPS), the Environmental Protection Agency (EPA), and a team led by BRRC. At each measurement site, the measuring organization deployed a research-grade microphone and a Class 1 sound level meter to sample ambient acoustic spectra at 1-second intervals for days, weeks, or months at a time. Together, these measurements reveal the spatially, temporally, and spectrally varying ambient sound levels in many different acoustic environments.

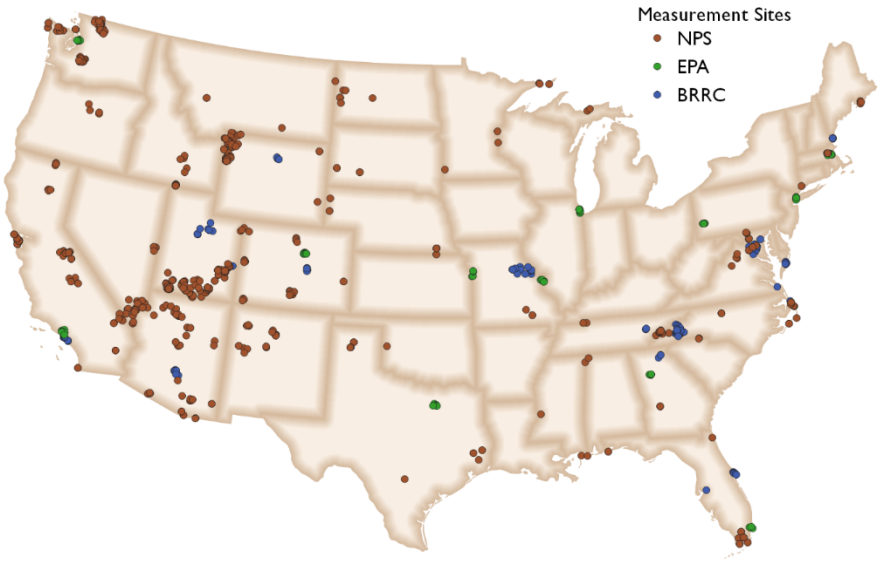

Figure 7: Acoustic measurement sites in the contiguous United States used to train the AMBIENT model.

The AMBIENT model also uses a database of geospatial features to fit the regression models at acoustic measurement sites and to apply the fitted models to predict ambient sound levels at other locations. The geospatial database includes 23 features at 30 m resolution across the contiguous United States. The resolution and spatial extents of the geospatial features determine the resolution and spatial extents of the ambient sound level predictions. At 30 m resolution, the predictions describe spatial variability at finer resolution than a typical city block. The geospatial features include population density, land cover, topography, climate, and physics-based transportation noise. Figure 8 shows the physics-based predictions of the average A-weighted L50, or median sound level, of road traffic noise, computed using the BRRC AMBIENT|TRAFFIC model, in an area between lower Manhattan and Newark, NJ.

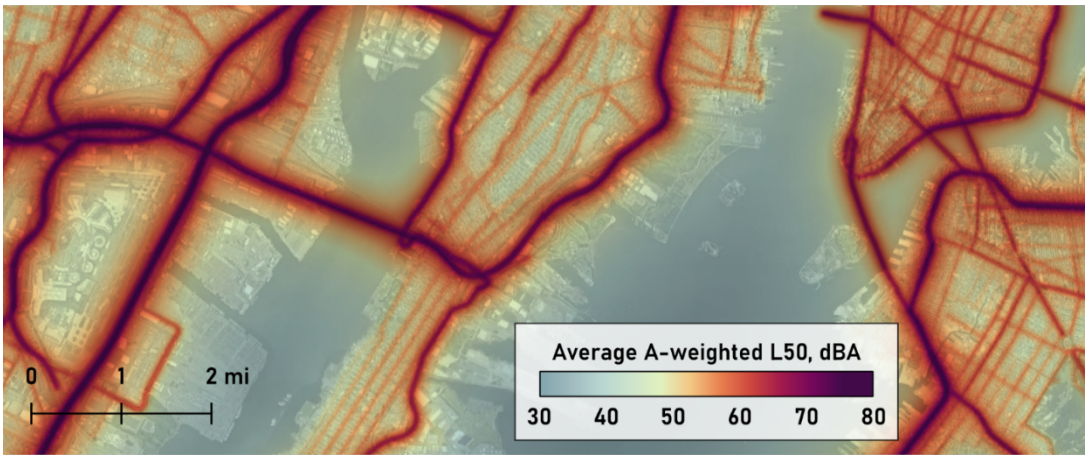

Figure 8: Physics-based traffic noise prediction in the area between lower Manhattan and Newark, NJ.

The BRRC AMBIENT|TRAFFIC model incorporates physics-based source models and simplified acoustic propagation into the machine learning models for a single transportation noise source. Aviation noise can be another important transportation noise source at locations near airports, but physics-based aviation noise is not yet included as a geospatial feature in the AMBIENT model. For this reason, the ambient noise is likely underestimated in the vicinity of airports. However, aviation noise typically does not affect the L50, or median sound level, except at locations immediately surrounding busy airports.

The AMBIENT model uses an ensemble of four machine learning regression algorithms—gradient-boosting regression, kernel ridge regression, multi-layer perceptron, and support vector regression—to predict the ambient sound level as a function of the geospatial features. The predicted ambient sound level is given by the median of the

four machine learning algorithms. Different ensemble models are trained to predict ambient sound levels at different times of day and for different OTOBs.

The AMBIENT model scales and combines the physics-based traffic noise predictions with other geospatial features to predict the overall ambient sound levels. At locations far from roads, the predicted ambient sound level is greater than the predicted traffic noise due to contributions from other anthropogenic, biological, and geophysical sound sources. At locations immediately adjacent to roads, where the ambient sound level is dominated by traffic noise, the AMBIENT model sometimes predicts ambient sound levels that are lower than the predicted traffic noise level. One cause of this apparent contradiction is that the AMBIENT model is anchored to the range of ambient acoustic measurements, and few measurement sites are immediately adjacent to busy roads. Additional ambient acoustic measurements along busy roads are required to validate the physics-based traffic noise and ambient soundscape model predictions in these locations.

2. Modeled Ambient Sound Levels in the New York City Metropolitan Area

The BRRC AMBIENT model was used to produce unweighted OTOB spectra over an extended grid of points with a 98.4 ft. (30 m) grid spacing. The data correspond to median daytime and nighttime levels. Maps of median daytime and nighttime A-weighted sound pressure levels are shown in Figure 9. These maps demonstrate significant differences in ambient sound levels at different times of day and at different locations. Figure 10 shows unweighted OTOB spectra for a high (▲) and a low (▼) ambient point of interest (POI) identified on the maps. Both the magnitudes and spectral shapes vary with location. These and data at other points serve as spatially varying maskers for the point-to-point route from JRB and EWR (see Section V.B).

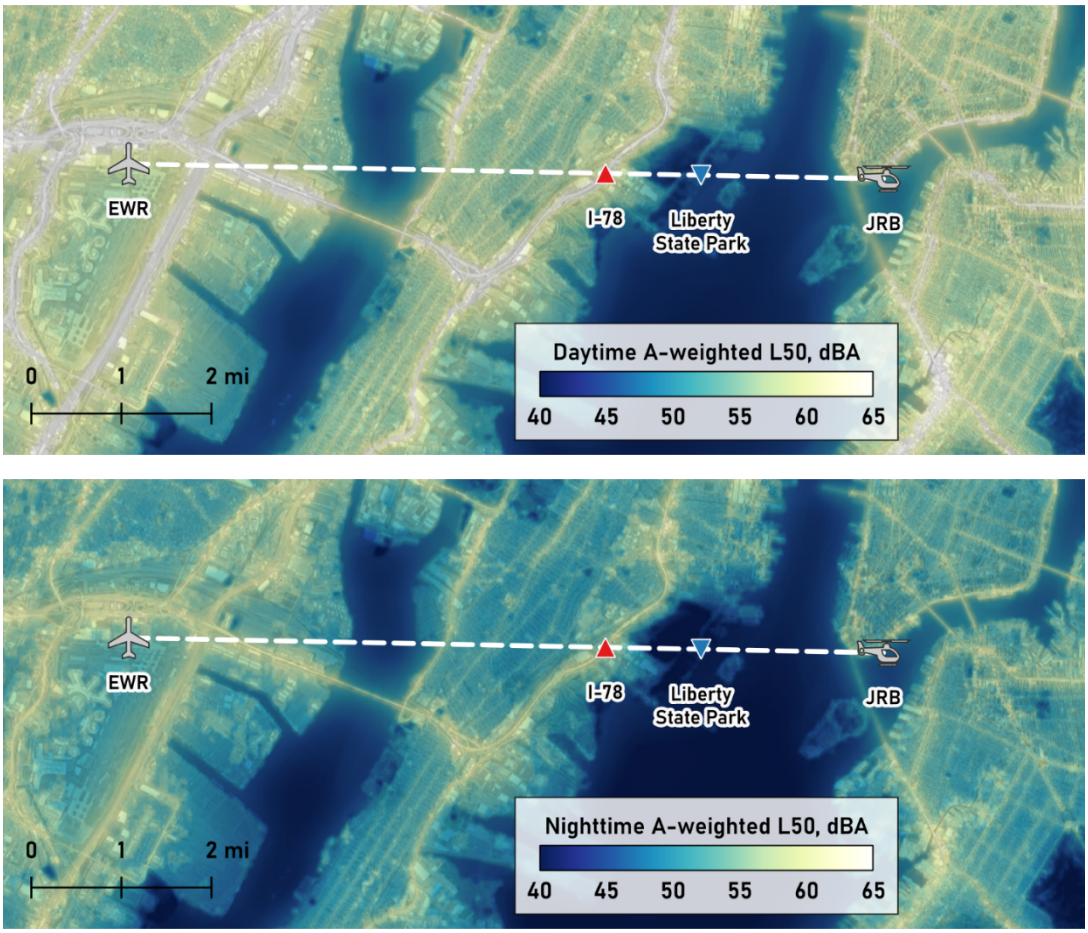

Figure 9: Maps of modeled daytime (top) and nighttime (bottom) ambient A-weighted sound pressure levels in the area between lower Manhattan and Newark, NJ.

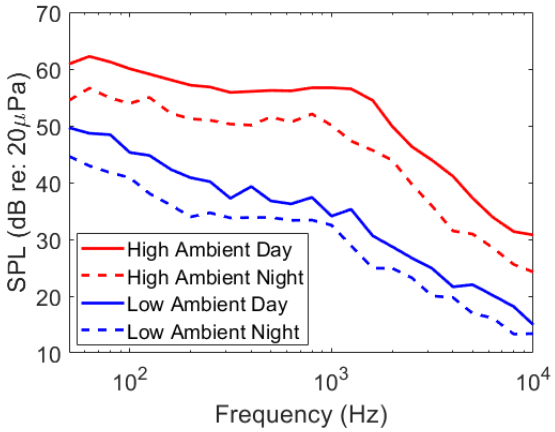

Figure 10: Modeled daytime and nighttime unweighted OTOB spectra at high ambient (along Interstate I-78) and at low ambient (Liberty State Park) POIs.

V. Results

Two sets of analyses are performed. The first case serves as a step-by-step demonstration of how the annoyance assessment is performed and how differences between two different temporally varying and spatially uniform ambient soundscapes affect annoyance. The second case is more applied; it considers a fictitious straight-line route from the Downtown Manhattan Heliport (JRB), a helicopter pier in lower Manhattan, New York City to the helipad at Newark Liberty International Airport (EWR) located about 8.4 mi. (13.5 km) away. This analysis serves to demonstrate how annoyance changes with spatially varying, temporally invariant ambient conditions.

A. Overflight Case

The process of estimating discounted annoyance at a single observer point is first demonstrated. This is followed by mappings of intermediate quantities to ultimately show how annoyance to a single-event overflight varies from one observer point to another and how annoyance varies with the masker. In the following, and for convenience, the City Hall and Central Park maskers were applied uniformly at all observer locations.

A high-fidelity simulation of the RVLT quadrotor reference vehicle was performed for a steady straight and level overflight at 90 knots (202 km/h) and 1000 ft. (304.8 m) altitude. The spectrogram of the signal at an observer point on the ground track is shown in Figure 11. Because the overflight represents a steady condition, the spectrogram for any point on the ground track is the same, while points lateral to the track differ. Asymmetry about the peak is due to nonuniform fore-aft (polar) directivity of the source. On the unweighted basis, the frequency range below about 500 Hz is dominated by tonal harmonics of the BPF and above about 500 Hz by the broadband self noise component. To give the reader a more natural sense of the signal characteristics, an auralization on the ground track was made using the NASA Auralization Framework [19] with plugins for periodic loading and thickness noise and broadband self noise [20]. The auralization is available for download [16].

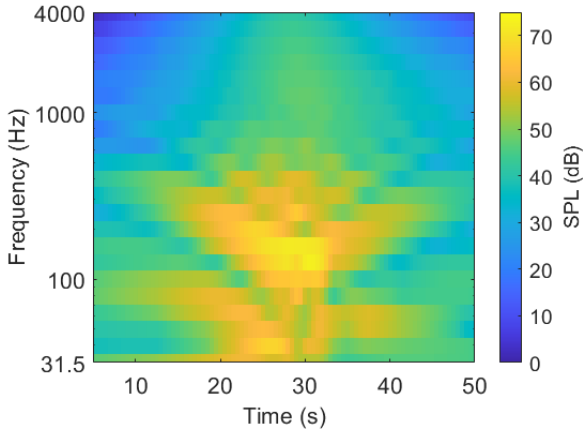

Figure 11: Spectrogram of signal sound pressure level L at an observer on the ground track for the overflight case.

The detectability index, $d'_{i,t}$ at each 0.5 s interval at each OTOB is computed using Eq. (2) to generate the d' spectrograms shown in Figure 12. Values are plotted on a logarithmic scale so that values below zero ($d' = 1$) indicate that the signal is functionally inaudible in the presence of the masker. The greater duration and frequency of detection indicated for the Central Park masker relative to the City Hall masker is expected. This is because its magnitude over the entire frequency range is lower than that of the City Hall masker and the difference increases with increasing frequency, see Figure 6. The spectrograms in Figure 12 may also be used to identify the time and OTOB with the highest probably of detection. These appear at an observer time of about 27 s and at the 125-160 Hz OTOB for the City Hall masker and at the 4 kHz OTOB for the Central Park masker. These spectrograms clearly illustrate that the detectability of the signal strongly depends on the spectral characteristics of the masker.

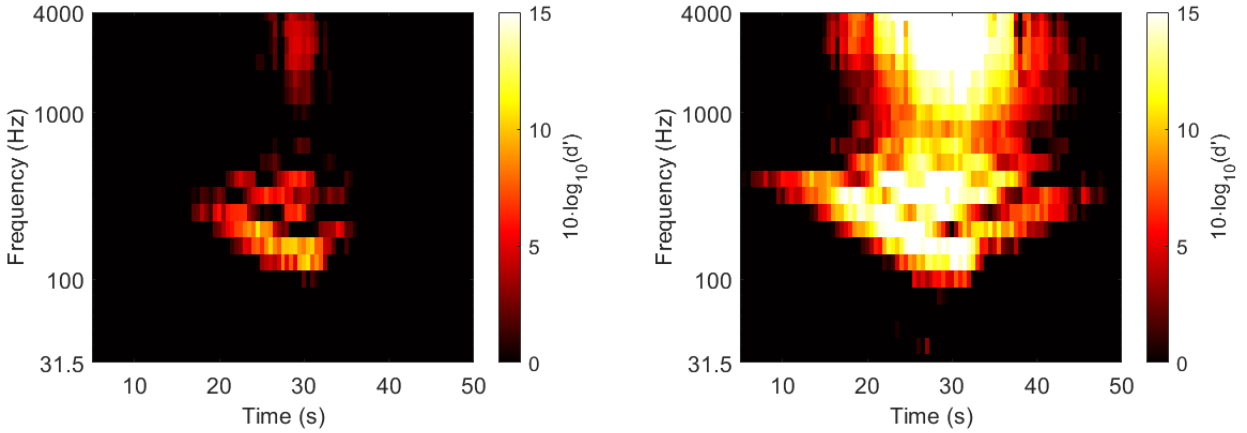

Figure 12: Detectability index spectrograms corresponding to the City Hall (left) and Central Park (right) maskers.

Discounted sound pressure level spectrograms are calculated using Eq. (1) with detectability index spectrogram data from Figure 12 and the original signal spectrogram from Figure 11. The resulting discounted spectrograms, shown in Figure 13, indicate what part of the signal is heard by the observer over the masker. In these plots, areas of 0 dB (and below) indicate the signal is fully masked, i.e., inaudible. The effect of the discount is greater with the City Hall masker than for the Central Park masker.

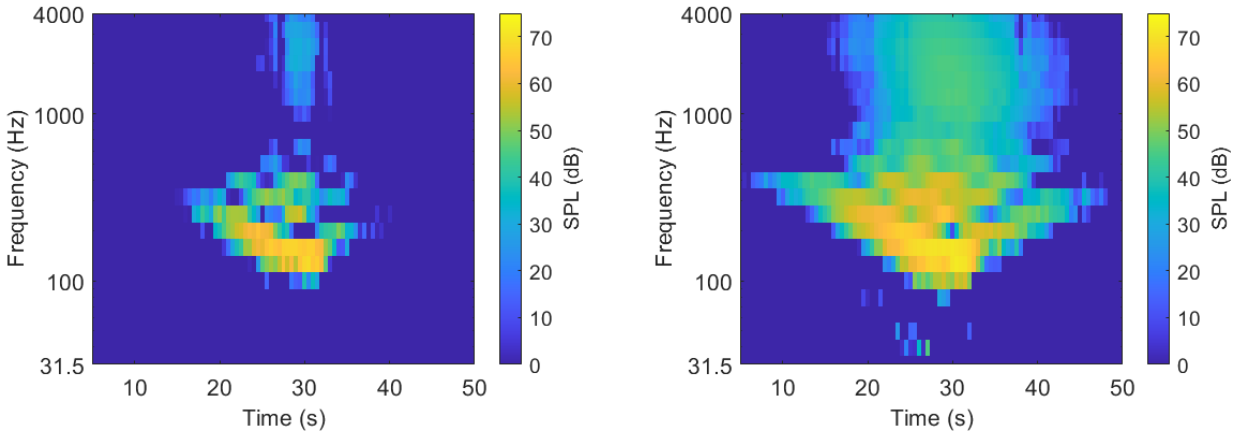

Figure 13: Discounted signal spectrogram L_{Disc} by City Hall (left) and Central Park (right) maskers.

Next, A-weighting is applied to the original and discounted spectrogram data. Both are summed over frequency to obtain the overall A-weighted level, L_A, and summed over time to obtain L_{AE}. The original and discounted L_A and L_{AE} are shown in Figure 14 for both maskers. The discounted levels, both in terms of L_{AE} and the maximum A-weighted level, L_{Amax}, are lower in magnitude for the City Hall masker relative to the Central Park masker, making the difference between the original and discounted levels greater for the City Hall masker than the Central Park masker.

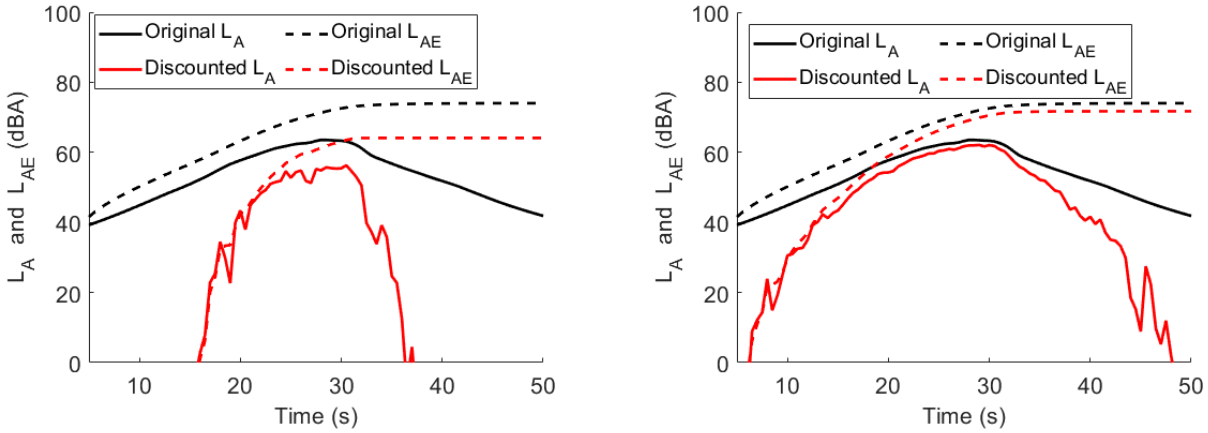

Figure 14: Original and discounted L_A and L_{AE} for City Hall (left) and Central Park (right) maskers.

Having calculated a discounted value of L_{AE} for a single flyover, a discounted day-night sound level L_{dn} for a 24 h period can be computed simply as an accumulation of multiple flyovers. A dose-response relation [21-23] specific to UAM noise could then be used to estimate the percent of the population that are highly annoyed in the presence of masking noise. However, such a relationship cannot be established until UAM aircraft are operational. Further, since flight recordings currently remain scarce or proprietary, laboratory studies of annoyance to UAM operations also remain lacking. Nevertheless, in order to demonstrate the next step relating L_{AE} (and hence discounted L_{AE}) to annoyance, laboratory test data are used.

The data are taken from a laboratory test comparing the annoyance of ground vehicles to small unmanned aerial systems (sUAS) [24]. Only the sUAS data are used here. Figure 15 shows a fit of the test subjects' mean annoyance rating with L_{AE}. The test was conducted using the 5-pt ICBEN scale [25] with ratings of "Not At All Annoyed", "Slightly Annoyed", "Moderately Annoyed", "Very Annoyed", and "Extremely Annoyed." These were assigned numerical values of 2, 4, 6, 8, and 10, respectively, on a scale of 1-11. The equation of the best fit line is given as:

$$\text{Annoyance} = 0.125L_{AE} - 3.61\,. \tag{4}$$

To the extent that this relation applies to UAM vehicles, substitution of the original and discounted L_{AE} gives estimates of original and discounted annoyance, as provided in Table 1. These data suggest that at this particular observer location, located directly under the ground track, the effect in terms of L_{AE} is significant for the City Hall masker (> 6 dBA) and not for the Central Park masker (< 1 dBA). The effect in terms of annoyance, however, is modest for both maskers, reducing the original annoyance estimate from one that is closer to being moderately annoyed (6) to one that is closer to being slightly annoyed (4) in the case of the City Hall masker.

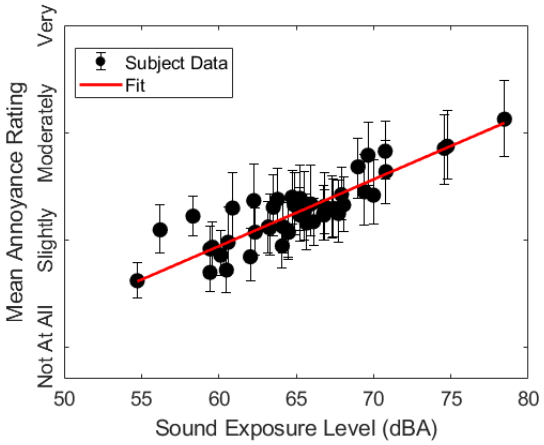

Figure 15: Relationship between annoyance and L_{AE} established from a psychoacoustic test simulating single-event flyover noise from small UAS aircraft [24].

Table 1: Comparison of estimated original and discounted L_{AE} and annoyance ratings.

	Original	City Hall (Discounted)	Central Park (Discounted)
L_{AE} (dBA)	74.04	67.82	73.15
Annoyance Rating	5.64	4.87	5.53

While the above analyses are informative about how discounted annoyance is assessed at one observer location, performing the same calculation over a grid of observers illustrates how the estimated response changes spatially. Contours of d' are shown in Figure 16 for both maskers. Along the ground track (sideline distance = 0 m), the signal is audible in the presence of either masker. Lateral to the track, the signal becomes inaudible at sideline distances greater than about 985 ft (300 m) for the City Hall masker but is audible for all or most of the domain for the Central Park masker. The variation in the index with downrange distance reflects the fact that the masking noise changes over time as the vehicle traverses the domain.

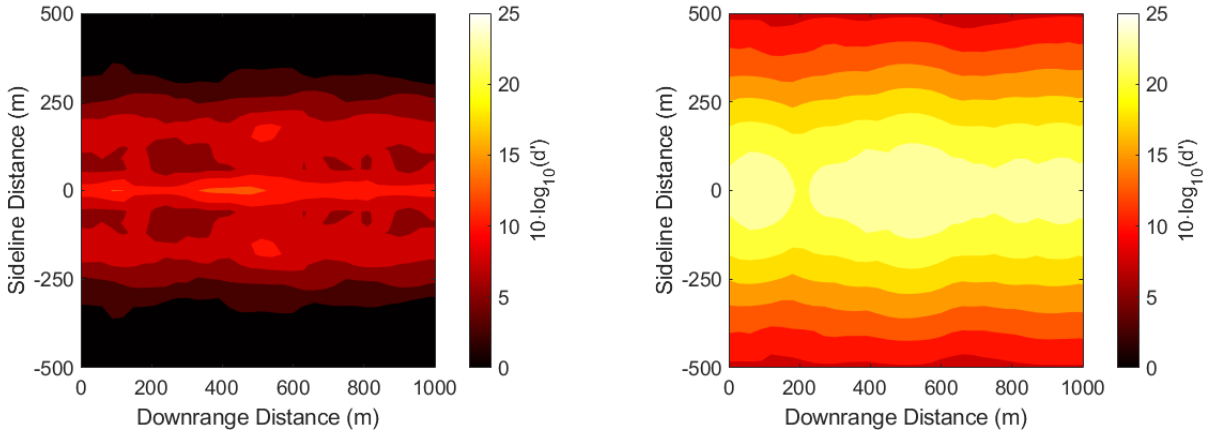

Figure 16: Detectability index contours corresponding to the City Hall (left) and Central Park (right) maskers.

Contours of the frequency of maximum d' are shown in Figure 17. As in Figure 12, detection occurs at low frequencies for the City Hall masker. The lower OTOB frequencies in the signal are dominated by higher harmonics of the BPF of the quadrotor vehicle. In contrast, detection occurs at the higher frequencies for the Central Park masker. The higher OTOB frequencies in the signal are dominated by broadband self noise. At the greatest sideline distances, atmospheric absorption reduces the high-frequency content of the signal and changes the frequency of maximum d' from high to low frequency.

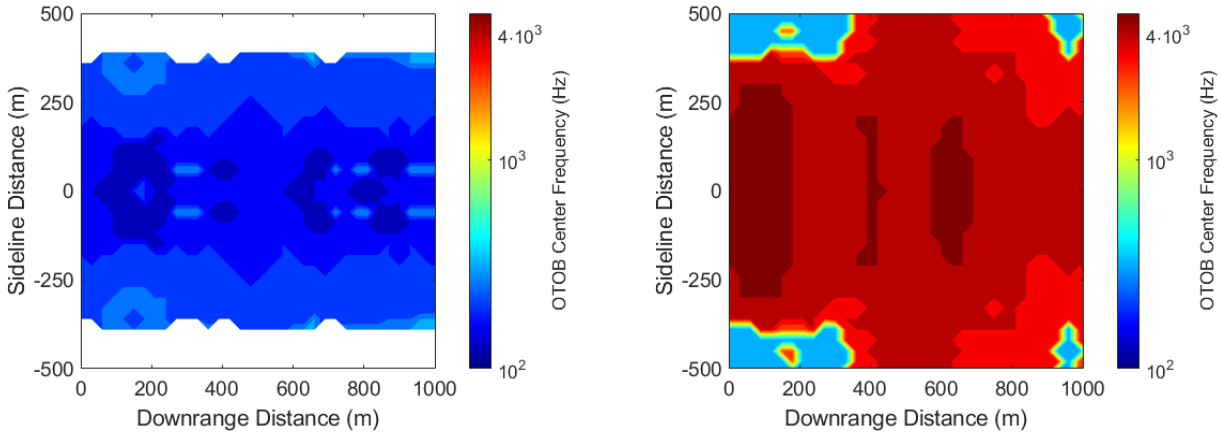

Figure 17: Frequency of maximum d' corresponding to the City Hall (left) and Central Park (right) maskers.

Original and discounted sound exposure level contours are shown in Figure 18. The original sound exposure varies only transverse to the ground track because the source noise is nearly symmetric about its azimuthal centerline [13]. Points lateral to the track differ due to increased spreading loss and atmospheric absorption, and differing ground reflection and Doppler shift. As suggested by Figure 16, there is a much greater discount for the City Hall masker than for the Central Park masker. Exposure levels of about 65 dBA are found at sideline distances of 1640 ft. (500 m) in the original and are found at sideline distances of less than 164 ft. (50 m) for the City Hall masker and at about 820 ft. (250 m) for the Central Park masker.

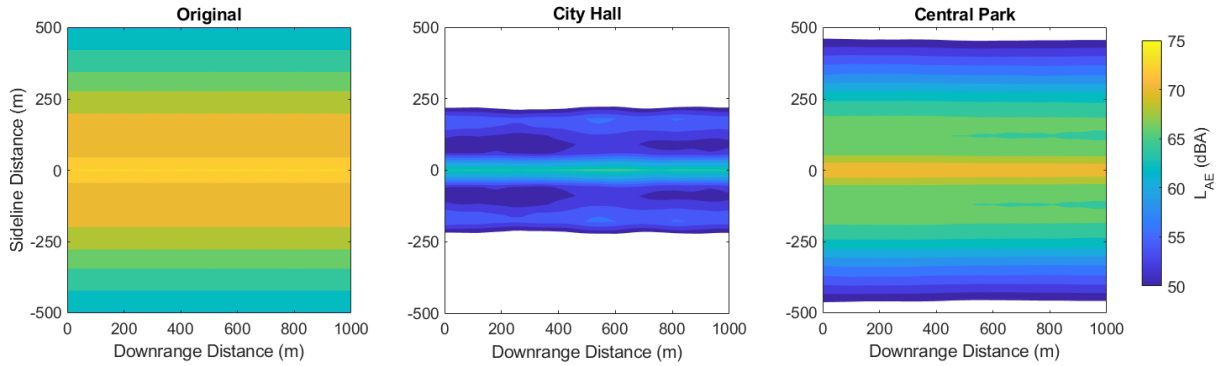

Figure 18: Original and discounted sound exposure contours for overflight case.

Estimated annoyance maps based on the original and discounted sound exposure data are shown in Figure 19. Annoyance based on the original data spans the range of "slightly annoyed" (4) at points furthest from the ground track to "moderately annoyed" (6) along the ground track. For the Central Park masker, estimated annoyance along the ground track is "moderately annoyed" but falls off to a little over "not at all annoyed" (2) at points furthest from the ground track. In contrast, the highest estimated annoyance along the ground track for the City Hall masker is "slightly annoyed" and falls off to "not at all annoyed" at a sideline distance of less than 655 ft. (200 m).

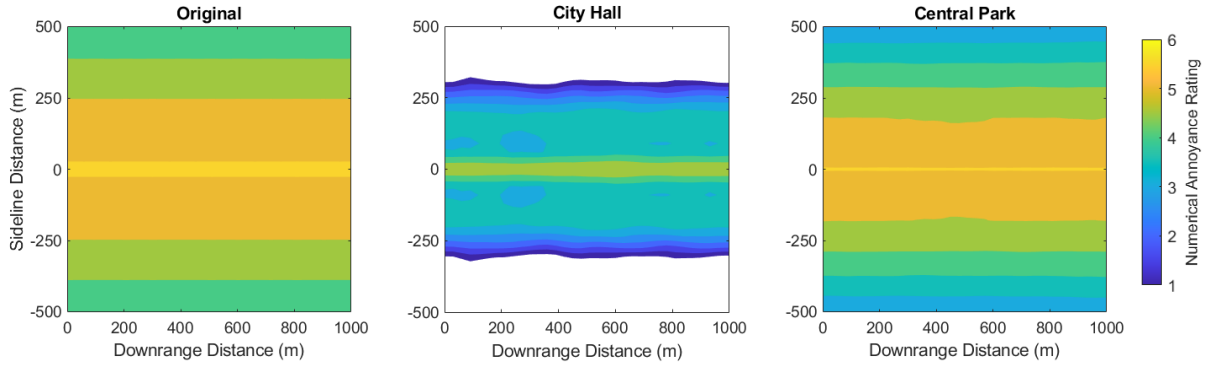

Figure 19: Original and discounted annoyance contours for overflight case.

In summary, this case illustrates how annoyance varies with short duration temporal variations of the masker, with the spectral characteristics of the signal and masker, and with distance from the source (affecting the spectral characteristics of the signal) due to atmospheric propagation.

B. Point-To-Point Case

The point-to-point case is modeled as a departure from JRB, followed by an overflight segment, and ending with an approach at EWR. The departure and approach profiles were derived from standard profiles for a Bell 206L helicopter in the AEDT Aircraft Noise and Performance database [26], as investigated by Letica and Rizzi [15]. Modifications to the standard profiles were made to ensure that the prescribed trajectories are within the flight envelopes of the quadrotor vehicle. The overflight segment was an extended version of the straight and level overflight considered in Section V.A. The east-to-west route traverses lower Manhattan, then crosses the Hudson River, the Bayonne Peninsula, and Newark Bay before landing at EWR. The sound exposure level map is shown in Figure 20. As expected, the point-to-point route shows elevated levels near the takeoff and landing areas, as observed in prior analyses [27]. In the following annoyance model analyses, an OTOB center frequency range of 50 Hz to 4 kHz was

used. This represented the intersection of the frequency range of the signal (31.5 Hz – 4 kHz) and the modeled ambient masker data (50 Hz – 10 kHz).

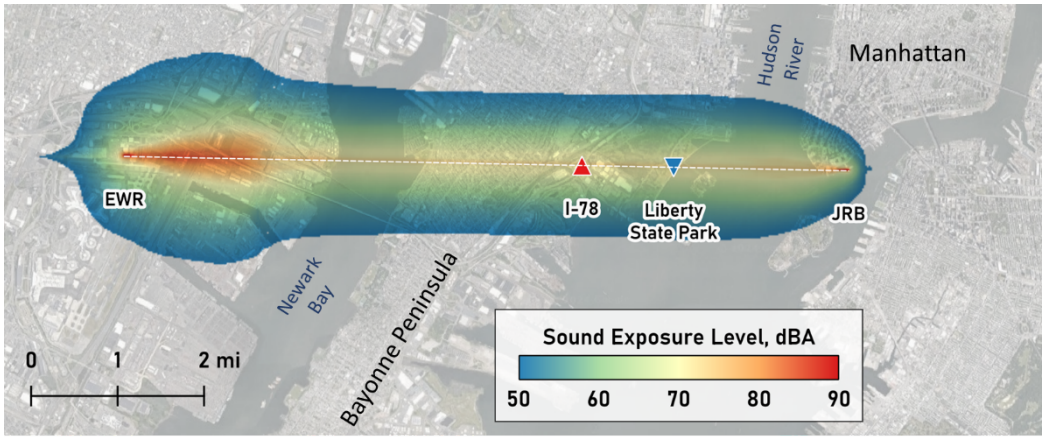

Figure 20: Sound exposure from RVLT quadrotor reference vehicle on point-to-point route from JRB to EWR.

Referring to Figure 9 and Figure 10, the original and discounted daytime levels, both in terms of L_{AE} and L_A, are shown at the high ambient POI on Interstate I-78 and at the low ambient POI at Liberty State Park (LSP) in Figure 21. The discounted levels are lower in magnitude for the I-78 masker relative to the LSP masker, making the difference between the original and discounted levels greater for the I-78 masker over the LSP masker, as expected. The L_A traces are smoother than those in Figure 14 because the I-78 and LSP maskers do not vary in time as did the City Hall and Central Park maskers.

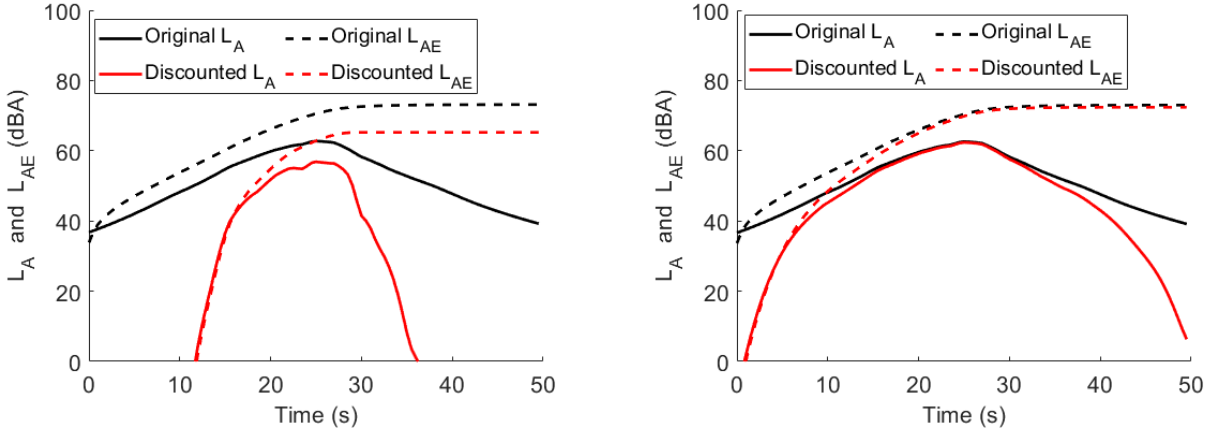

Figure 21: Original and discounted daytime L_A and L_{AE} for Interstate I-78 (left) and LSP (right) maskers.

An extended section of the overflight segment of the point-to-point route, including the I-78 (▲) and LSP (▼) POIs, is next considered. Maps of d' are shown in Figure 22. In the areas along Interstate I-78 and Route 440 (unmarked on the west side of the Bayonne Peninsula), estimates of daytime d' are significantly reduced relative to surrounding areas. This is less the case at nighttime. Estimates of d' at LSP indicate that the signals are audible at all times. Over the majority of the Bayonne Peninsula, the signal is audible over a strip of about 0.5 mi (800 m) wide, with that width being slightly greater during the nighttime.

Maps of the frequency of maximum d' for daytime and nighttime are also revealing, see Figure 23. For both daytime and nighttime, the maps indicate that detection occurs at low-frequency low-order harmonics of the BPF directly under the flight path primarily in the region over the Bayonne Peninsula. Lateral to the flight path, the frequency of maximum d' alternates between high and low frequency, perhaps due to changes in the source directivity, e.g., see Figure 6 in [13]. At somewhat larger lateral distances from the track, the frequency of maximum d' occurs in the range of 300-400 Hz due to higher-order BPF harmonics. This is also likely due to changes in the source directivity since increased atmospheric attenuation with increased lateral distance would drive detection to

lower frequencies, as seen in Figure 17 for the Central Park masker. Note that at these locations, however, the detectability is very low (near $d'=1$), see Figure 22. Over water, the frequency of maximum d' is more dominated by high frequency close to the track and moves to lower frequency at lateral locations due to an increasing amount of high-frequency atmospheric absorption. For this scenario, there is generally more high-frequency detection at nighttime than there is at daytime.

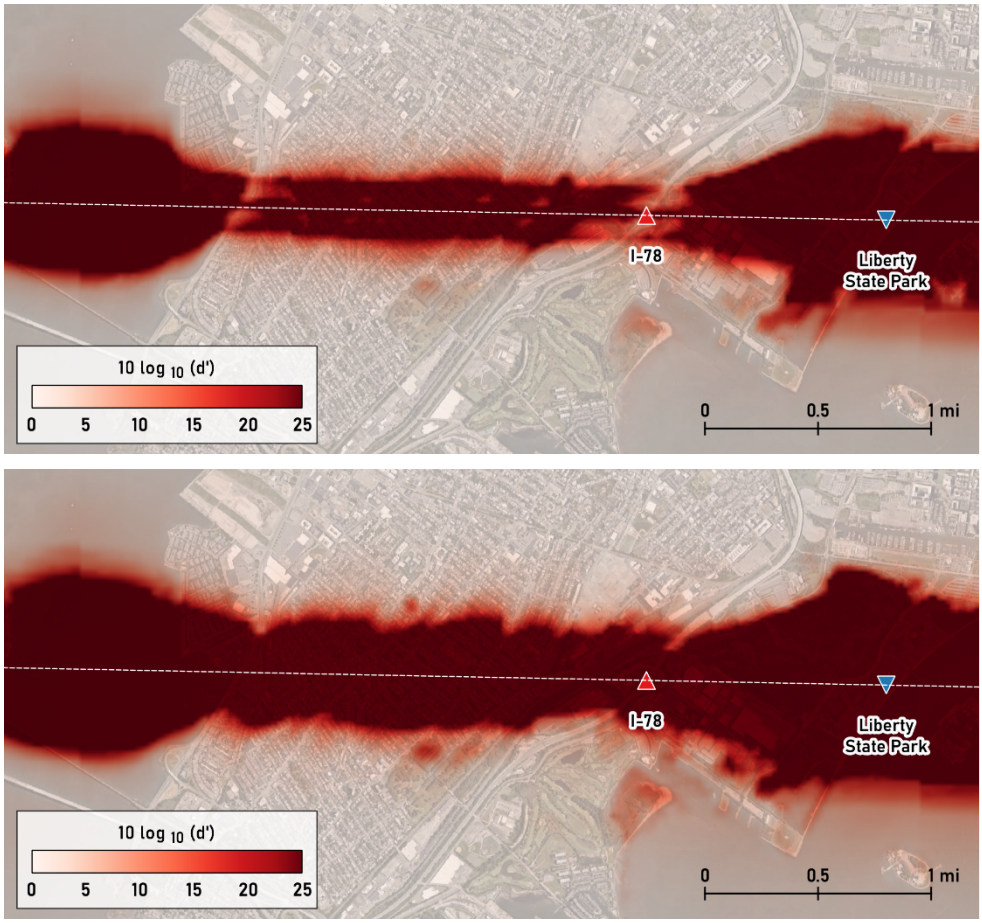

Figure 22: Maps of d' for daytime (top) and nighttime (bottom) for section of overflight segment of point-to-point route.

For brevity, maps of discounted sound exposure are not shown. Annoyance maps for the overflight segment are shown in Figure 24. Estimates of the original annoyance calculated without masking do not vary with position along the track and indicate moderate annoyance (6) levels directly below the flight track that uniformly reduces to less than not at all annoyed (2) with increasing lateral distance.

In contrast, the daytime annoyance estimates are slightly reduced below the flight track, are significantly reduced over the Bayonne Peninsula (except for the low ambient area around LSP ▼) and are less reduced over the two rivers to the east and west. In the areas along Interstate I-78 (▲ east) and Route 440, the annoyance estimates are reduced to less than slightly annoyed (4). Other geographic features are also apparent in the annoyance map. The nighttime annoyance estimates show less of a reduction due to masking than the daytime estimates. Nevertheless, reduced annoyance can be seen along Interstate I-78 and Route 440.

VI. Summary and Future Work

A method has been developed for estimating annoyance to UAM noise in the presence of masking noise. The annoyance model uses the relationship between annoyance and L_{AE} derived from a recent psychoacoustic study. An overflight case using recordings of two different maskers in NYC demonstrated how temporal variations in the masker affect the detectability index d' and the frequency corresponding to the maximum value of d'. The discounted annoyance was subsequently calculated from estimates of discounted sound exposure level using laboratory test data

relating annoyance to small unmanned aerial systems to sound exposure level. A more applied case using a dense grid of ambient noise estimates from the BRRC AMBIENT model demonstrated how differences between high and low ambient areas affect estimated annoyance.

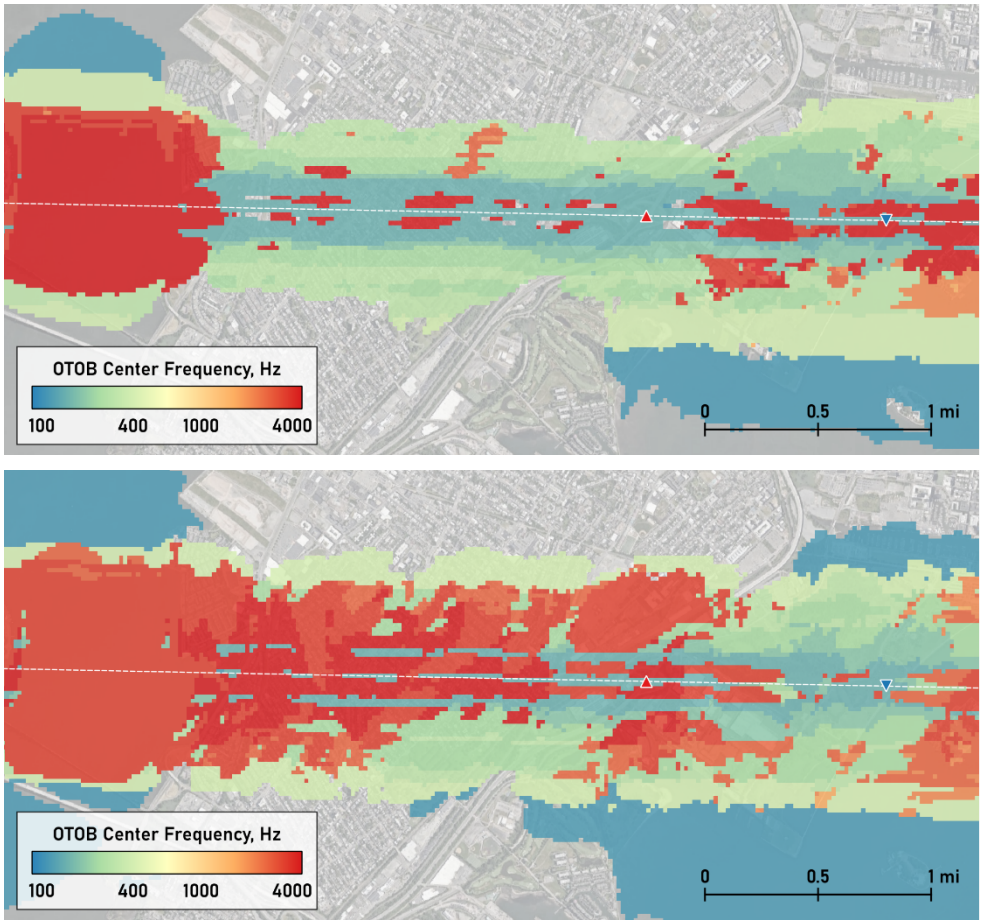

Figure 23: Maps of frequency of maximum d' for daytime (top) and nighttime (bottom) for section of overflight segment of point-to-point route.

For the cases considered, it was found that the estimated annoyance directly below the flight track was largely unaffected by the ambient noise except for very high ambient conditions, e.g., near major roadways. Lateral to the flight track however, the estimated annoyance was almost always reduced relative to the unmasked condition, and the frequency at which detection occurred increasingly tended toward low-frequency high-order harmonics of the BPF with increasing lateral distance from the ground track. The observed reductions in estimated annoyance near roadways lend some support for the proposed routing of UAM traffic over roadways. However, this study also demonstrates large differences between the lateral extent of air vehicle noise relative to road traffic noise. Additional analyses on the basis of day-night average sound level L_{dn} are still needed to determine the impact of fleet operations. More complex vehicle noise contours associated with departures from, and arrivals to, a vertiport also require further analysis. Higher sound levels in the vicinity of vertiports are likely to engender greater annoyance than found en route (see Figure 20).

While this work constitutes a plausible first step toward development of a useful tool to aid urban route planning, additional work is needed. First, the psychoacoustic data upon which the median annoyance model parameters are based are limited to a handful of subjects in a single laboratory study. More work is needed to establish an estimate applicable to the larger population. Given the diversity of UAM vehicle architectures and the sound signatures they generate, consideration of a range of aircraft must be given before making any kind of generalization about a fleet of aircraft. And to that end, a relationship between annoyance and L_{dn} for UAM vehicles is needed so that annoyance estimates may be made using more relevant data. However, that will have to wait until a number of actual flight vehicles are available and flying. Given these caveats, it should be apparent that the results provided herein are meant

to demonstrate the utility of the approach and are not intended to make a definitive statement about how UAM aircraft are likely to be perceived in any particular locality.

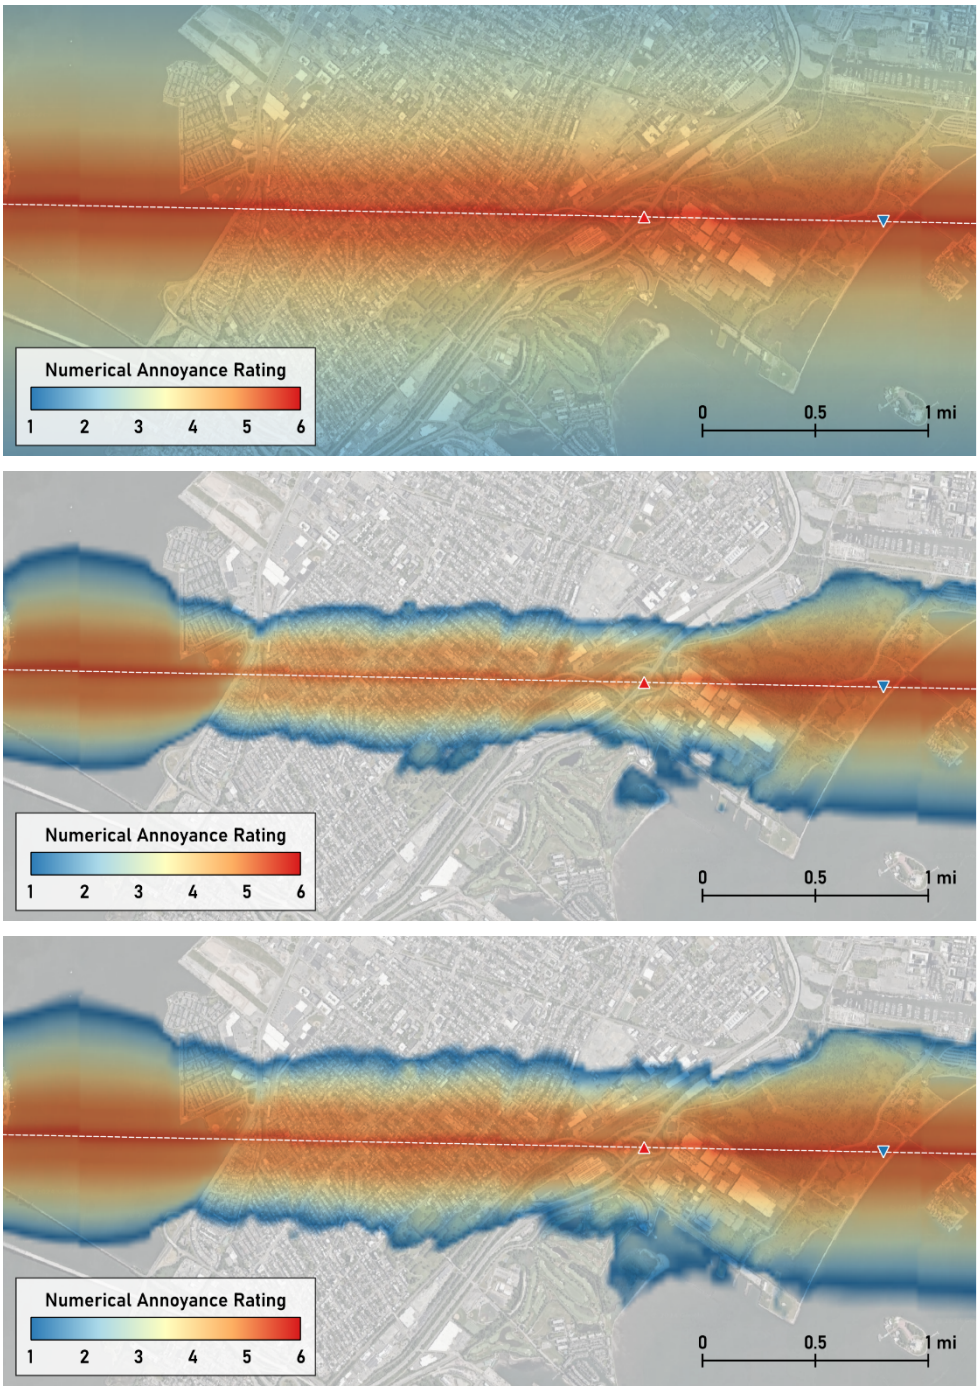

Figure 24: Original (top) and discounted annoyance maps for daytime (middle) and nighttime (bottom) for section of overflight segment of point-to-point case. [2 = “Not At All Annoyed”, 4 = “Slightly Annoyed”, 6 = “Moderately Annoyed”]

Appendix

This appendix describes the OTOB-based detection algorithm used in this research. It first introduces the concept of the power spectrum model of masking and then describes the individual components that constitute Eq. (2).

A. The Power Spectrum Model

The detection algorithm used here is generally known as a "power spectrum" model [28], as it is based on ratios of the power of the signal and masker in various frequency bands. While there are many other aspects of sound that can impact detection performance – beyond banded signal-to-noise ratios – due to the coarseness of the input data, it is not (easily) possible to incorporate such effects (see [3]).

This model arises from a conception of the peripheral auditory system as being a set of bandpass filters. Roughly, the transduction mechanism of the ear (from external pressure fluctuations to an internal neural firing code) operates as a massive bank of bandpass "auditory filters" [28]. In the ear these filters overlap significantly in their frequency selectivity and sensitivity. In the model, in order to avoid "double counting" incoming energy, the algorithm needs to use a bank of bandpass filters that sum to a flat frequency response. The OTOB data are used to fill the role of this filter bank, with the benefit of then being able to use commonly available data and digital signal processing tools. The signal-to-noise ratio coming through each filter forms the basis of the model. However, OTOB bandwidths, while following a similar trend to auditory filter widths, are never an exact match, so a correction needs to be applied. Note that any nonlinearity of the system with sensitivity (absolute level) is ignored.

The model is based on the task of detecting a single tone signal in a locally white noise masker – that is, the masker is approximated to be white for the width of each OTOB. Using this scenario allows the corrections to be formulated, though these clearly may not relate to many (more complex) real-world scenarios. These scenarios include signals having closely spaced tones and broadband components, a masker composed of common sounds each with their own characteristics, spatial effects impacting detection performance, etc. This simple approach, however, has formed the basis of past practical detection algorithms [6], and is likely a good first approximation. "Tuning" the performance of such an algorithm using data specific to a UAM-in-city scenario may be the subject of future research (though, with necessary loss of generality, unless the input data were to become richer).

There are two main effects that compete in a tone-in-noise detection scenario. The first is that of the bandwidth-time or "BT" product of the observation. This is a quantity, derived from signal detection theory [29], that indicates that if the masker is broadband (stochastic in nature), there will be more variability between observation windows of the masker for shorter duration maskers or those with more limited bandwidth. This can perhaps be thought of as the opposite of "regression towards the mean" (although that concept is prone to misconception as well). If the goal is to detect a signal in such a masker – the signal being stationary or known, such as a tone – then higher variability in the masker will lead to a reduction in detection performance on a moment-to-moment basis. The variability scales with the square roots of the auditory filter bandwidth w_i and the length of the signal Δt, that is, $d' \propto \sqrt{w_i \cdot \Delta t}$. This indicates that there is a statistical advantage to using wider auditory filters (or to being patient).

The other force in the model is that of the amount of noise coming through each auditory filter. As an auditory filter widens, more masking noise is allowed to enter the band, while the tone – a single spectral line – will maintain the same filtered power. The masker power will therefore scale linearly with the bandwidth, i.e., $n_i \approx w_i \cdot n_o$, for a white masker with specific signal power n_o defined in Pa^2/Hz. This will cause the signal to become harder to detect as the auditory filter gets wider. Thus, these forces are in competition with each other, but one is stronger. In the end, the total effect will be that the *signal level needed for detection will tend to increase as the square root of the auditory filter bandwidth for constant signal-to-noise ratios*.

It is important to note that reading the work of previous authors on this topic can be quite confusing. Depending on how the equations are written, the same bandwidth effect can seem to be described in very different ways. For instance, if an author uses the specific sound power in the band (n_o in Pa^2/Hz, as above), then a factor of $1/\sqrt{w_i}$ may appear in the equation. If they use the total power in the band (as is done here), then the increase in masker power in the band will be bookkept within n_i and a factor of $\sqrt{w_i}$ will appear. Some previous works freely go between these two conventions, making things even more confusing. This extends to the way that s and n are defined. For instance, Sneddon et al. [6] define their s and n in terms of RMS Pascals and then write the signal-to-noise ratio unambiguously as $\sqrt{(s/n)^4}$ – an equivalent expression to that used here.

In the end, the important thing to remember is that the italicized statement above is empirical, backed by decades of human response data, and should be the correct behavior of any detection prediction system for this simple tone-in-noise case (see, e.g., Figure 1 in [30]). A useful approximation to this fact comes from Fastl and Zwicker [31]. They report that a tone is detectable at the same level in a white noise masker up to 500 Hz (the bandwidth of the masker being arbitrarily wide). Above that frequency, the level needed for equal detection increases at 10 dB per decade. This heuristic is plotted in Figure 25 and is used as a benchmark for the detection algorithm.

With this background, the unknown quantities from Eq. (2) can now be discussed. There are two such elements in this equation: the detection efficiency factor $\mathbb{D}$, and the equivalent auditory system noise e_i. $\mathbb{D}$ must be determined first based upon the parameters of the simulation and input data, after which e_i may be solved for.

B. Detection Efficiency Factor

The frequency-dependent detection efficiency factor $\mathbb{D}$ has several elements that are built up from various sources in the literature.

1. There is an overall (frequency-independent) "efficiency factor" relative to the performance of a perfect detector, usually denoted in the literature as η, which defines the general capacity for human detection relative to the raw signal-to-noise ratio. Sources typically give this as being between 0.25 and 0.33 (e.g., [32]). A value of 0.3 is used here.

2. There is a factor of the square root of the time window size that arises from the BT product described above. This factor is only useful to predict changes in *relative* detection ability over a narrow range of time scales – it does not extend to very short sounds (less than 0.1 s, cf. [30]), or to infinity (giving infinite detection capacity to patient observers). In practice, this factor is often set to 1 s, sometimes simply by way of omission (e.g., [6]). In this work, the time step size of 0.5 s is used, and, in a practical sense, is just another component of η.

3. The frequency dependence stemming from the BT product appears as $\sqrt{w_{\mathrm{OTOB},i}}$ due to the formalism adopted here.

4. Human auditory filter bandwidths, while similar to OTOB bandwidths near 250 Hz, deviate significantly from OTOBs at high and low frequencies. A useful equation for the bandwidth of auditory filters comes from Moore [28] as:

$$w_{AF,i} = 24.7\left(1+\frac{4.37}{1000}\cdot f_{c,i}\right). \tag{5}$$

   A comparison of filter bandwidths is shown in Figure 25. A correction between this formula, computed at the OTOB center frequencies $f_{c,i}$, and the OTOB bandwidths is made as:

$$\sqrt{\frac{w_{\mathrm{OTOB},i}}{w_{AF,i}}}. \tag{6}$$

   This correction factor is unintuitive, as it is doing the opposite of what one would expect – dividing out $w_{\mathrm{OTOB},i}$ and replacing it with $w_{AF,i}$. However, the effect it has on the performance of the algorithm is as intended: If $w_{AF,i} < w_{\mathrm{OTOB},i}$, d' should increase based upon the above discussion. (Cf. a past similar correction made by Fidell and Horonjeff [33].)

5. The final component is one of engineering judgement. It is a high-frequency adjustment to the detection efficiency that keeps the algorithm from putting too much weight on detection occurring at very high frequencies. This occurs in practice due to recordings and specifications of background OTOB levels rolling off in the kHz range while, for instance, an auralization retains these frequencies perfectly. This dB correction is referred to as "the kick" and is defined for OTOBs with $f_{c,i} \geq 2500$ Hz:

$$\kappa_i = -6\left[\log_{10}\left(\frac{f_{c,i}}{2500}\right)\right]^2 \tag{7}$$

   The kick factor is shown on a linear scale, i.e., $10^{\kappa_i/10}$, in Figure 25. Other similar schemes have simply cut off the computation at 5 or 10 kHz (e.g., [6]). This factor allows for a smooth transition away from the

algorithm reporting that detection is occurring at 16 kHz (which is unreasonable in a practical sense), while allowing that energy to still be computed and counted in the rest of the computation. In the end this factor makes very little functional difference (i.e., on the resultant discounted L_{AE}) because UAM sounds do not have significant signal energy at such high frequencies.

Taken the above all together, the frequency-dependent detection efficiency factor $\mathbb{D}$ is given as:

$$\mathbb{D}_i = \eta \cdot \sqrt{\Delta t} \cdot \sqrt{w_{\mathrm{OTOB},i}} \cdot \sqrt{\frac{w_{\mathrm{OTOB},i}}{w_{AF,i}}} \cdot \kappa_i \,. \tag{8}$$

The total detection efficiency factor is plotted in Figure 25. The performance of the algorithm is also shown in Figure 25. This is computed for each band using a white noise masker of constant n_o and signals that are tones located at f_c. It is computed without the e_i factor - an omission that makes the algorithm independent of absolute level - so the result is plotted on an arbitrary dB scale. As can be seen, the performance closely matches that of the benchmark described above. At very high frequencies, the algorithm without the kick is trending towards better (probably unreasonably good) detection performance, whereas the performance with the kick begins to degrade above 5 kHz.

C. Equivalent Auditory System Noise (EASN)

Once the factor $\mathbb{D}$ is established, the equivalent auditory system noise e_i may be solved for. This term is derived from the hypothetical situation in which there is a single tone located at an OTOB center frequency but there is no masking noise presented to the subject ($n_{i,t}$=0). The level at which the typical subject will just be able to hear this tone is called the "auditory threshold" or "minimum audible field." These tone levels are provided by the ISO 226 standard [34] and encoded as signal magnitudes for each OTOB ($s_{MAF,i}$). In order to make Eq. (2) reproduce the equivalent behavior of the auditory threshold, an equivalent amount of noise can be added in the denominator that, in concept, originates inside the auditory system (hence the name of the term).

An assumption must be made about what the ISO data series represents in terms of detectability – that is, what was the value of d' that was sought in the experiment(s) that were used as an input to that standard. Most psychoacoustic detection experiments target d' values around 1.5 or 2.5, with higher d' values being favored (e.g., [35,36]). Thus, an assumption that the target d' =2 is used here. e_i can then be solved for via an inversion of Eq. (2):

$$e_i = \mathbb{D}_i \cdot \frac{s_{MAF,i}}{2} \tag{9}$$

This term is plotted in Figure 25, again with and without the addition of the kick factor. As can be seen, e_i is a very small amount of signal energy – on the order of 10^{-9} Pa^2 in the middle frequencies.

D. Summing Detectability Between Bands

In practical detection situations, researchers often seek an overall d' value for a complex signal. In a laboratory situation, if a psychoacoustic test were being run, only one value of d' would be able to be measured at a time for an entire signal. If the signal spans across multiple frequency bands, there are competing theories as to how to aggregate d' to form a single representative value (see, importantly, [37]).

In an earlier implementation of the discount concept, a single value of d' was computed for each time step – d' was aggregated across frequency – and the entire signal was discounted based on this single value. This research has used a simpler conception where each frequency band is discounted based on its own detectability, avoiding the need to worry about the addition of d' between bands. Determining which concept is more concordant with discounting as observed in human subjects (i.e., via psychoacoustic test) is left for future work.

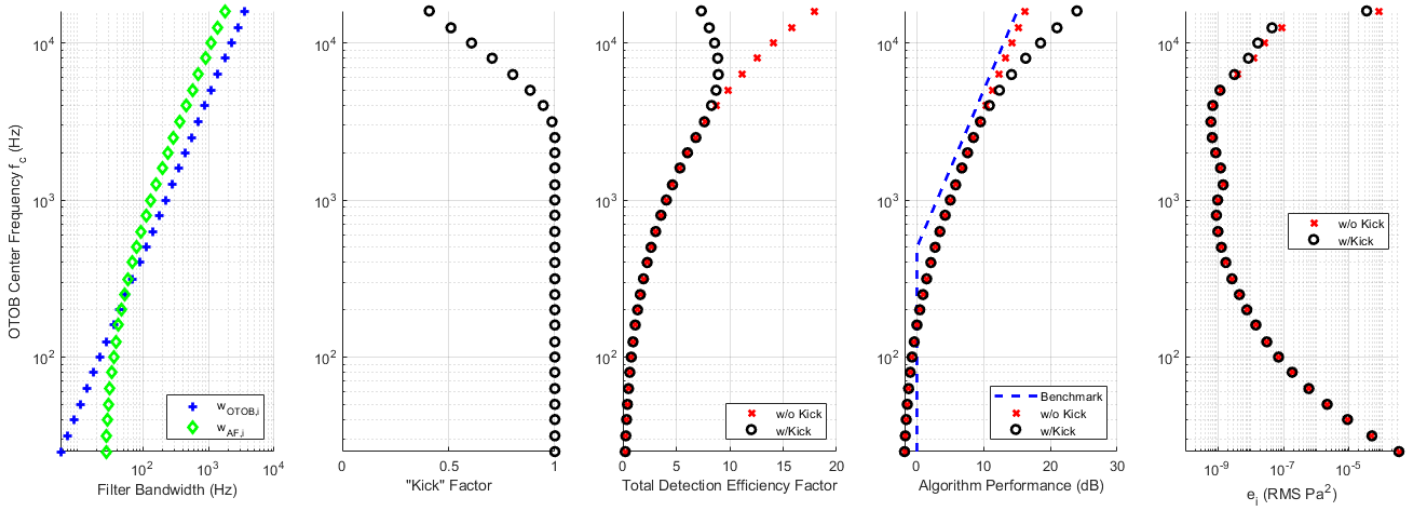

Figure 25: Dependency of power spectrum model parameters on OTOB center frequency. Shown are OTOB and auditory filter bandwidths (left) per Eq. (5), the "kick" factor (second from left) per Eq. (7), the detection efficiency factor $\mathbb{D}$ (middle) per Eq. (8) , the algorithm performance (second from right), and the EASN (right) per Eq. (9).

Acknowledgments

This work was supported by the NASA Aeronautics Research Mission Directorate, Revolutionary Vertical Lift Technology Project. Development of the BRRC AMBIENT model was performed under United States Army-funded Small Business Innovative Research program contract W911W6-18-C-0028.

References

[1] "Noise standards: aircraft type and airworthiness certification," CFR Title 14, Chapter 1 - Federal Aviation Administration, Department of Transportation, Part 36, https://www.ecfr.gov/current/title-14/chapter-I/subchapter-C/part-36, 2023.

[2] "Airport noise compatibility planning," CFR Title 14, Chapter 1 - Federal Aviation Administration, Department of Transportation, Part 150, https://www.ecfr.gov/current/title-14/chapter-I/subchapter-I/part-150, 2023.

[3] Christian, A., "The effect of background noise on human response," In *Assessment and reduction of propeller and rotor noise from unmanned aircraft systems*, NATO, STO-TR-AVT-314, pp. 24, 2023.

[4] Rizzi, S.A., Huff, D.L., Boyd Jr., D.D., Bent, P., Henderson, B.S., Pascioni, K.A., Sargent, D.C., Josephson, D.L., Marsan, M., He, H., and Snider, R., "Urban air mobility noise: Current practice, gaps, and recommendations," NASA TP-2020-5007433, 2020.

[5] Fields, J.M., "Reactions to environmental noise in an ambient noise context in residential areas," *The Journal of the Acoustical Society of America*, Vol. 104, No. 4, 1998, pp. 2245-2260, https://doi.org/10.1121/1.423770.

[6] Sneddon, M., Pearsons, K., and Fidell, S., "Laboratory study of the noticeability and annoyance of low signal-to-noise ratio sounds," *Noise Control Engineering Journal*, Vol. 51, No. 5, 2003, pp. 300-305, https://doi.org/10.3397/1.2839726.

[7] Boucher, M.A., Christian, A.W., Tracy, T.D., Krishnamurthy, S., Begault, D.R., Rizzi, S.A., and Shepherd, K.P., "A psychoacoustic test on the effect of masking on annoyance to urban air mobility vehicle noise," *186th Meeting of the Acoustical Society of America,* Ottawa, Canada, 2024.

[8] Lim, C., Kim, J., Hong, J., and Lee, S., "Effect of background noise levels on community annoyance from aircraft noise," *The Journal of the Acoustical Society of America*, Vol. 123, No. 2, 2008, pp. 766-771, https://doi.org/10.1121/1.2821985.

[9] Christian, A., "A construction for the prediction of noise-induced annoyance in the presence of auditory masking," *181st Meeting of the Acoustical Society of America,* Seattle, WA, 2021, https://doi.org/10.1121/10.0008224.

[10] Tracy, T.D., Boucher, M.A., Christian, A.W., Krishnamurthy, S., Rizzi, S.A., Begault, D.R., and Shepherd, K.P., "An annoyance model for urban air mobility vehicle noise in the presence of a masker," *Noise-Con 2024*, New Orleans, LA, 2024.

[11] Patterson, M.D., Antcliff, K.R., and Kohlman, L.W., "A proposed approach to studying urban air mobility missions including an initial exploration of mission requirements," *AHS International 74th Annual Forum and Technology Display*, Phoenix, AZ, 2018.

[12] Silva, C., Johnson, W.R., Solis, E., Patterson, M.D., and Antcliff, K.R., "VTOL urban air mobility concept vehicles for technology development," *AIAA AVIATION Forum*, AIAA-2018-3847, Atlanta, GA, 2018, https://doi.org/10.2514/6.2018-3847.

[13] Rizzi, S.A., Letica, S.J., Boyd Jr., D.D., and Lopes, L.V., "Prediction of noise-power-distance data for urban air mobility vehicles," *AIAA Journal of Aircraft*, Vol. 61, No. 1, 2024, pp. 166-182, https://doi.org/10.2514/1.C037435.

[14] Lopes, L.V. and Burley, C.L., "ANOPP2 User's Manual: Version 1.2," NASA TM-2016-219342, 2016.

[15] Letica, S.J. and Rizzi, S.A., "On the modeling of urban air mobility vehicle takeoff and landing operations in the FAA Aviation Environmental Design Tool " *Noise-Con 2024*, New Orleans, LA, 2024.

[16] "Aircraft Flyover Simulation," https://stabserv.larc.nasa.gov/flyover, NASA, 2024.

[17] Lympany, S.V., Calton, M.F., Cook, M.R., Gee, K.L., and Transtrum, M.K., "Mapping ambient sound levels using physics-informed machine learning," *183rd Meeting of the Acoustical Society of America,* Nashville, TN, 2022, https://doi.org/10.1121/10.0015498.

[18] Mennitt, D., Sherrill, K., and Fristrup, K., "A geospatial model of ambient sound pressure levels in the contiguous United States," *The Journal of the Acoustical Society of America*, Vol. 135, No. 5, 2014, pp. 2746-2764, https://doi.org/10.1121/1.4870481.

[19] Aumann, A.R., Tuttle, B.C., Chapin, W.L., and Rizzi, S.A., "The NASA Auralization Framework and plugin architecture," *InterNoise 2015*, San Francisco, CA, 2015.

[20] Krishnamurthy, S., Aumann, A.R., and Rizzi, S.A., "A synthesis plugin for auralization of rotor self noise," *2021 AVIATION Forum, 27th AIAA/CEAS Aeroacoustics Conference*, AIAA 2021-2211, Virtual Meeting, 2021, https://doi.org/10.2514/6.2021-2211.

[21] "Federal agency review of selected airport noise analysis issues, Final report: Airport noise assessment methodologies and metrics," Federal Interagency Committee on Noise (FICON), Washington, D.C., 1992.

[22] "International standard 1996-1, Acoustics - Description, measurement and assessment of environmental noise, Part 1: Basic quantities and assessment procedures, 3rd edition," International Organization for Standardization (ISO), 2016.

[23] Miller, N.P., Czech, J.J., Hellauer, K.M., Nicholas, B.L., Lohr, S., Jodts, E., Broene, P., Morganstein, D., Kali, J., Zhu, X., Cantor, D., Hudnall, J., and Melia, K., "Analysis of the Neighborhood Environmental Survey," Federal Aviation Administration, DOT/FAA/TC-21/4, 2021.

[24] Christian, A. and Cabell, R., "Initial investigation into the psychoacoustic properties of small unmanned aerial system noise," *23rd AIAA/CEAS Aeroacoustics Conference, AIAA AVIATION Forum*, AIAA-2017-4051, Denver, CO, 2017, https://doi.org/10.2514/6.2017-4051.

[25] Fields, J.M., De Jong, R.G., Gjestland, T., Flindell, I.H., Job, R.F.S., Kurra, S., Lercher, P., Vallet, M., Yano, T., Guski, R., Felscher-Suhr, U., and Schumer, R., "Standardized general-purpose noise reaction questions for community noise surveys: Research and a recommendation," *Journal of Sound and Vibration*, Vol. 242, No. 4, 2001, pp. 641-679, https://doi.org/10.1006/jsvi.2000.3384.

[26] "Aviation Environmental Design Tool (AEDT) technical manual, Version 3e," U.S. Department of Transportation, Volpe National Transportation Systems Center, Cambridge, MA, DOT-VNTSC-FAA-22-04, 2022.

[27] Rizzi, S.A. and Rafaelof, M., "On the modeling of UAM aircraft community noise in AEDT helicopter mode," *2023 AIAA AVIATION Forum*, AIAA-2023-3363, San Diego, CA, 2023, https://doi.org/10.2514/6.2023-3363.

[28] Moore, B., *An introduction to the psychology of hearing*, Sixth ed, Brill, 2013.

[29] Green, D.M. and Swets, J.A., *Signal detection theory and psychophysics*, John Wiley & Sons, New York, 1966.

[30] Dai, H. and Wright, B.A., "The lack of frequency dependence of thresholds for short tones in continuous broadband noise," *The Journal of the Acoustical Society of America*, Vol. 100, No. 1, 1996, pp. 467-472, https://doi.org/10.1121/1.415859.

[31] Fastl, H. and Zwicker, E., *Psychoacoustics: facts and models (3rd ed.)*, Springer, Berlin, 2006.

[32] Tanner, W.P., Jr. and Birdsall, T.G., "Definitions of d′ and η as psychophysical measures," *The Journal of the Acoustical Society of America*, Vol. 30, No. 10, 1958, pp. 922-928, https://doi.org/10.1121/1.1909408.

[33] Fidell, S. and Horonjeff, R., "A graphic method for predicting audibility of noise sources," U.S. Air Force Wright Aeronautical Laboratories, Wright-Patterson Air Force Base, OH, AFWAL-TR-82-3086, 1982.

[34] "Acoustics, Normal equal-loudness-level contours," International Standards Organization, ISO 226:2023, 2023.

[35] Green, D.M., "Stimulus selection in adaptive psychophysical procedures," *The Journal of the Acoustical Society of America*, Vol. 87, No. 6, 1990, pp. 2662-2674, https://doi.org/10.1121/1.399058.

[36] Green, D.M., "Erratum: ‘‘Stimulus selection in adaptive psychophysical procedures’’ [J. Acoust. Soc. Am. 87, 2662–2674 (1990)]," *The Journal of the Acoustical Society of America*, Vol. 88, No. 5, 1990, p. 2486, https://doi.org/10.1121/1.400366.

[37] Buus, S., Schorer, E., Florentine, M., and Zwicker, E., "Decision rules in detection of simple and complex tones," *The Journal of the Acoustical Society of America*, Vol. 80, No. 6, 1986, pp. 1646-1657, https://doi.org/10.1121/1.394329.