DOES A QUIETER CITY MEAN FEWER COMPLAINTS? THE SOUNDS OF NEW YORK CITY DURING COVID-19 LOCKDOWN

Mark Cartwright^{1*}, Magdalena Fuentes², Charlie Mydlarz², Fabio Miranda³, Juan Pablo Bello²

¹ New Jersey Institute of Technology, New Jersey, USA
 ² New York University, New York, USA
 ³ University of Illinois at Chicago, Illinois, USA

ABSTRACT

The COVID-19 pandemic had an unprecedented effect in human activity and city landscapes. A very notorious transformation during this period was the change in noise levels and patterns across cities. Small scale studies have show this change in noise levels across different locations in the globe. In this work, we extend these studies by using historical audio data from the SONYC sensor network deployed in New York City. We exploit machine listening models to understand not only noise levels but also patterns, by performing a sound source presence analysis. Finally, we contrast our finding from the acoustic data with noise complaints to better understand the relationship between noise and our perception of it.

Index Terms- Sensor networks, machine listening, Covid-19.

1. INTRODUCTION

The COVID-19 pandemic had a profound effect on human activity, causing an unprecedented halt on in-person activities and mobility in general, which was reflected in empty and quiet streets across the planet [1, 2]. One of the most notorious and noticeable impacts that this unrivaled event had was the decrease in global noise levels [3, 4, 5, 6]. This was even more apparent in mega cities like New York City (NYC) [7], where "the city that never sleeps" was abnormally quiet. Motivated by this unique and historic moment, many studies discuss the COVID-19 lockdown driven differences in sound pressure level (SPL) for different cities [8, 9, 10, 11, 12, 13]. Given the difficulties of long-term noise monitoring, most of these studies compare short-term noise measurements taken during the pandemic to few data points previously collected at the same locations. The dynamic and transient nature of urban noise requires longitudinal monitoring to effectively capture its true characteristics over time.

In this work we present a preliminary analysis of a unique longitudinal noise dataset from the the Sounds of New York City or SONYC project [14, 15]. This project has maintained an active network of advanced noise sensors monitoring 24/7 for over 6 years, SPL data as well as encrypted audio snippets, allowing for cutting edge noise monitoring approaches to aid in a deeper understanding of the urban noise delta over the COVID-19 lockdown period. The temporal scale of the SONYC dataset provides an unrivalled baseline of noise conditions at the sensed locations over a 4 year period.

In addition, we have the ability to investigate not only changes in noise levels, but also changes in sound sources across the city through the application of machine listening. Finally, we look into noise complaints during this period and previous years to learn about people's perception of the noise levels during this unusual time.

*This work was partially supported by NSF award 1544753.

2. SOUNDS OF NEW YORK CITY (SONYC) PROJECT SENSOR NETWORK

The Sounds Of New York City (SONYC) project deployed a network of over 55 low-cost, acoustic sensor nodes across NYC to facilitate the continuous, real-time, accurate and source-specific monitoring of urban noise [14]. The first sensor was deployed in May 2016 with the network growing in size over a three year period. This network was retired in June 2022, resulting in a vast and unique sonic dataset, enabling large-scale analysis of urban noise activities to reveal neverbefore-seen noise patterns across space and time. Cumulatively, 150 years of calibrated SPL data and 75 years of raw audio snippets was collected from the sensor network. This data has been used to identify longitudinal patterns and often overlooked occurrences of noise pollution across urban settings [16, 17]. The majority of sensors were mounted to window ledges and light poles in the Greenwich Village and NoHo areas of Manhattan, NYC, with additional clusters deployed in Midtown Manhattan, Downtown Brooklyn, the Upper East Side of Manhattan, and the Corona area of Queens, with specific locations shown in [15]. These include locations near parks, thoroughfares, universities, and transportation hubs. Deployments were always external and efforts were taken to mount the sensors away from existing sources of noise at the building edge, for example air conditioning units or vents. Care was also taken to ensure the node has an un-occluded view of the street and is as far away from main exterior walls as possible to reduce the artificial boosting of measured SPL levels when too close to large hard surfaces. A general rule of at least a one-block distance between sensor nodes was adhered to, unless there was a point of interest close-by such as a long term construction project or major roadway.

The majority of the sensors were mounted onto window ledges at heights ranging between 15–25 feet. In practice, over long deployment durations the height difference between sensor nodes would result in negligible SPL variations between the sensors ranging 15– 35 feet in height, as the most significant difference in absolute SPL measurement between different sensors occurs when a noise source is directly below the sensor for long periods of time. However, the variation in deployment characteristics such as sensor height and flat surface/pole based mounting will likely produce variable SPL measurements between deployment types. This makes the sensors more adept at measuring decibel change over time rather than directly comparing absolute SPL values across deployment types.

3. DATASET

For our analysis, we selected 11 sensors which actively collected data from early 2017 through July 24, 2020, which includes key

events in NYC's COVID response and recovery as indicated in Table 1. Nine of the sensor were located in Manhattan and two were located in Brooklyn. Each sensor collected two main types of data: SPL values, and encrypted 10 s audio snippets. In this paper, we use A-weighted SPL measured at 1 s intervals and audio snippets recorded at random intervals at an average rate of 1 snippet per minute. While we have data extending back to 2017, the target period of our analysis is the five-month period from Feb. 24, 2020 (3 weeks before the first event, E0, in Table 1) to July 24, 2020. We use historical data from 2017–2019 from the same period as a baseline to which we compare the data in our five-month target period.

 Table 1. Event IDs indicating status of New York City lockdown and reopening

1 0		
Event ID	Event	Date
EO	Schools closed	2020-03-16
E1	Bars/restaurants closed	2020-03-17
E2	New York State on pause	2020-03-22
E3	Construction halt	2020-03-28
E4	Phase 1 reopening	2020-06-08
E5	Phase 2 reopening	2020-06-22
E6	Phase 3 reopening	2020-07-06
E7	Phase 4 reopening	2020-07-19

4. THE SOUND OF A CITY IN LOCKDOWN

4.1. Measuring change in SPL

During our five-month analysis period, there were moments of what seemed to be unprecedented quiet. To confirm this, we first computed the average A-weighted SPL (LAeq) for each day for each sensor in both the analysis and the historical-baseline periods. We then averaged these values over sensors for each day and sorted the 1285 days of our dataset by increasing LAeq. The LAeq for the quietest 50 days ranged 60.0–62.4 dB (4–6 dB lower than the the mean daily LAeq of 66.5 dB in 2017–2019), and all but five of these days occurred during the pandemic period. Of the remaining five days, three were holidays (Thanksgiving 2018, Christmas 2018 and 2019), when many residents travel to their family-home states, and the other two were unusually quiet Sundays one of which was unusually cold (Feb. 17, 2019) and one which was not (Mar. 17, 2019).

To further understand the effect the city's COVID-19 response on noise levels, we performed a linear regression analysis using the daily sensor LAeq measurements as the dependent variable. To account for the seasonal characteristics of the city, we included categorical regressors coding the day of the week and month of the year. Given the non-COVID-related factors that seemed to influence the 50 quietest in our dataset, we also included a regressor indicating if the day was an official city holiday and weather-related regressors: mean daily precipitation (mm) and mean daily temperature (C). To account for the location differences of our sensors, we also included categorical regressors encoding sensor ID. Lastly, we encoded the stage of NYC's COVID-19 response as a categorical regressor, i.e. the time period between E1 and E2, the time period between E2 and E3, etc. The regression was statistically significant (Adjusted $R^2 = 0.736$, F(38, 12039) = 885.3, p < 0.001). Table 2 lists the primary regression variables of concern, along with their standardized coefficient estimates and standard error. All but one of the regressors is statistically significant (p < 0.001). By comparing the standardized coefficients, we can see how the different stages of the city's COVID-19 response affected the noise of the city, with a dampening effect increasing up to and including when construction was paused, and then decreasing as the city began reopening. The strength of these effects for the states between the time bars/restaurants closed and Phase 2 of the city's reopening were stronger than that of an official city holiday, precipitation, and temperature. It is worthwhile to note that Sundays are particularly quiet in the city (standardized coefficient of -2.064), but the quietest period was even quieter than Sundays.

Table 2. Standardized coefficients of linear regression modeling daily LAeq. Not showing seasonal (month of year, day of week) and sensor ID variables due to space. N = 12078; $R^2 = 0.736$ '***' indicates p < 0.001

indicates $p < 0.001$				
Variable	Estimate	Standard Error		
City holiday	-1.883***	(0.136)		
Precipitation (mm)	0.443***	(0.050)		
Temperature (C)	0.078***	(0.006)		
E0 < t < E1	-0.864	(0.695)		
E1 < t < E2	-2.294***	(0.318)		
E2 < t < E3	-3.499***	(0.291)		
E3 < t < E4	-4.147***	(0.094)		
E4 < t < E5	-2.624***	(0.198)		
E5 < t < E6	-1.758***	(0.193)		
E6 < t < E7	-1.443***	(0.210)		
E7 < t	-0.942***	(0.320)		

While our regression analysis investigates the factors that are predictive of the daily-averaged LAeq, it does not give us insight into the temporal patterns of the change in LAeq and how they evolved over the course of the city's response and reopening. To investigate this, we built a simple averaged model of the spatiotemporal patterns in our SPL data for 2017-2019, and we then compared our 2020 data to this model. We first created such a model for each of the K = 11sensors. To do so, we aggregated and downsampled each year of SPL data per sensor into six LAeq values for each day, i.e., fourhour intervals. We then aligned each of these yearly time series so that they start on the same day of the week. Next, we stacked and reshaped each sensor year of the 2017-2019 data into in a tensor $X_k \in \mathbb{R}^{3 \times 6 \times N}$ where 3 is the number of years, 6 is the number LAeq values per day, and N is the minimum number of days left in the year (less than 365 due to the day-of-week alignment). We then convolved this tensor with a $3 \times 1 \times 3$ averaging kernel using a dilation rate of seven in the third dimension (day). The result is a matrix $\bar{X}_k \in \mathbb{R}^{6 \times N}$ where each element in the matrix is an average of nine values from the same four-hour interval (e.g., Monday, Noon-4pm) consisting of the center point, 1 week in the past, and 1 week in the future for 2017, 2018, and 2019. This averaging should give us a robust template of the temporal noise patterns for that sensor location. We also created a resampled, aligned, and reshaped LAeq matrix for 2020, $Y_k \in \mathbb{R}^{6 \times N}$. To compare how 2020 differed from 2017–2019, we computed the difference between our 2020 noise matrix and our 2017–2019 matrix for each sensor $(Y_k - \bar{X}_k)$, and aggregated over K sensors by taking the median difference for each interval-day bin.

The resulting diel plot in Figure 1B more clearly shows the drop in LAeq during the city's COVID response and reopening, with the effect most pronounced after events E3 and E4 consisting of an approximately 4 dB drop during most times of the day. While there is a negative change for all time intervals in the day, the change is not uniform — there is a greater decrease in LAeq in the middle of the day than there is at night. Lastly, the anomalous noise increases in



Fig. 1. (A) Change in LAeq and complaint count in standard deviations (SD) by day. (B) Change in LAeq in dB by 4-hour periods and day. (C) Change in complaint count in SD by 4-hour periods and day. (D) Product of (B) in SD and (C) in SD to visualize when covarying. The white dashed lines correspond to the events in Table 1 and the grey regions indicate time periods of missing data.

late May and early June indicate that noise city's response to COVID is not the only thing that affected the noise profile of the city during our analysis period — these times correspond to protests in response to the May 25th, 2020 killing of George Floyd by a police officer.

4.2. Measuring change in class presence

The city's soundscape was disrupted not only because of the change in noise levels, but also because of the change in presence of typical sounds sources. To quantify this change, we performed a per-class analysis of sound sources typically found in NYC. To that end, we trained a multi-label multi-layer perceptron model on the SONYC-UST 2.3 dataset [18]. The model has one hidden layer with 128 units and ReLU activations, as well as AutoPool [19] pooling on the output. The audio of the dataset was recorded from the SONYC acoustic sensor network and was annotated by both volunteers on the Zooniverse citizen science platform [20] and experts from the SONYC research team. The annotations of this dataset indicate the presence of 23 classes that were chosen in consultation with the NYC Department of Environmental Protection, organized into a taxonomy of coarse and fine classes. We focus on the coarse classes: engine, machinery-impact, non-machinery impact, powered-saw, alert signal, music, human voice and dog.

Our model consists of one hidden dense layer of size 128 and a ReLU activation, followed by another dense layer of size 8 that indicates the likelihood of each coarse class. The input of the model is a pre-trained 512-dimensional OpenL3 embedding [21], which have shown to be strong embeddings for sound event detection [21]. We extract OpenL3 embeddings for each 10 s clip in our dataset, which are processed by the dense layers in parallel. Finally, we aggregate the

model's frame level mode predictions using AutoPool [19], resulting in one set of 8 likelihoods per excerpt. We use the training set of SONYC-UST to train the model, which contains clip-level annotations for the coarse classes. The presence of a given sound source within an audio excerpt is then determined by thresholding the class likelihood estimations. To do so, we select an optimal threshold for each class using the validation set of SONYC-UST.

Using the same aggregation technique described in Section 4.1 for the SPL analysis, we compute one presence matrix per-sensor per-class, summarizing the presence levels from 2017–2019 in blocks of four hours, resulting in a matrix $P_k^c \in \mathbb{R}^{6 \times N}$ for each sensor k and class c. We follow the same procedure for the class presence in 2020, obtaining another set of matrices per-sensor per-class $R_k^c \in \mathbb{R}^{6 \times N}$. To compute the change in presence, we compute the per-class per-sensor difference $(R_k^c - P_k^c)$ and aggregate over sensors by computing the median of the differences of 2020 to each one of the previous years. These diel plots are depicted in Figure 2.

Besides this intra-class presence comparison, we look at the overall change in the soundscape by looking at the presence difference when we consider all the classes together. We do this by stacking corresponding bins from all the classes from the $P_k^c \in \mathbb{R}^{6\times N}$ matrices for each sensor, resulting in a set of matrices $\hat{P}_k \in \mathbb{R}^{6\times 8\times N}$. We then compute the mean presence across bins from the matrices corresponding to the period 2017–2019, and calculate the overall change by computing the L_1 norm between the 2017–2019 median presence matrix and the 2020 matrix in the axis corresponding to the classes presence, see bottom plot of Figure 2.

Observing the diel plots in Figure 2, we notice a deep decrease of presence across all classes starting near the first lockdown event, and this decrease gets exacerbated after the pause of activities at the state level and the halt in construction (E2 and E3). The bigger change for all classes independently and combined was observed during working hours (8am-4pm), as expected given the reduction in activities. The exception being a few classes for which the decrease included the evening hours, such as human voice and alert signal. The decrease in alert signal was dominated by a decline in sirens during this period, which mirrored the cease of mobility and presence of people in the street. Not surprisingly, construction related sound sources (e.g., impact sounds, power saw) had their bigger abate after the construction halt. Lastly, an interesting pattern is shown in the dog presence, which was the only source with an increase of presence during this period. One potential explanation for this is that as people worked from home, they would walk their dogs more often during earlier afternoon hours. A second is that dog barks were more audible in the absence of significance traffic and construction noise.

5. A QUIETER CITY DOES NOT MEAN FEWER COMPLAINTS

While our sensor data gives us an indication of how the noise changed in the city, it doesn't tell us how that change affected the residents of the city. As a preliminary step in investigating this, we used noise complaints to the the city's non-emergency government services contact point, 311. We downloaded 311 service requests from Jan. 1, 2017 through July 24, 2020 through NYC OpenData [22] and filtered them to complaint types related to general noise and outdoor noise: *Noise*, *Noise* - *Park*, *Noise* - *Street/Sidewalk*, *Noise* - *Vehicle*. We then performed the same steps as described in Section 4.1 to create resample, align, and reshape, and average 2017–2019 noise complaint count data into a template to which we compare the noise complaint counts that occurred in 2020. To get a sense of the magnitude of this change, we re-scaled the change in



Fig. 2. Change in class presence. The bottom figure visualizes the L_1 norm between distributions of presence for all the classes. The white dashed lines correspond to the events in Table 1 and the grey regions indicate time periods of missing data. The color scaling is class-specific to highlight the temporal changes within a class rather than compare across classes.

complaint counts in 2020 by normalizing by the standard deviation of complaint counts in 2017–2019. Figure 1C displays the result of this process as a diel plot, and Figure 1A displays the result of this process as a line plot paired with a line plot of the change in LAeq, which was also scaled by the SD of 2017–2019 data. For a more detailed comparison of the change in LAeq to the change in noise complaint counts, we plotted the product of their SD-scaled change matrices in Figure 1D. From these plots, we can see the change in complaints was very pronounced. After construction was halted in the city (E3), we saw a reduction in complaints that followed the same daily trend as the reduction in LAeq. However, in early- to mid-may this trend shifts and noise complaints begin to increase, diverging from LAeq and increasing at a rapid rate after the Phase 1 reopening (E4), peaking on June 21, 2020 with a peak change of 7.1 SDs. By looking at Figure 1C and D, we see that the change (both decreases and increases) of complaints primarily occurred at night, again correlated to the reduction in LAeq at night and then dramatically shifting to an *increase* of complaints at night in May, diverging from LAeq. While we can't deduce the cause of this increase in complaints from our data, we speculate that the return to outdoor social activity in the summer along with the introduction of outdoor dining came as a shock to complaining residents after the very quiet lockdown. However, the data does shows that a quieter city does not necessarily mean fewer complaints and calls into question the relation of complaints to actual noise levels.

6. DISCUSSION & CONCLUSIONS

While the perceived quiet of the city during the COVID-19 lockdown was very apparent to NYC residents, we have shown in Sec 4.1 the significant influence each lockdown stage had on objective urban SPL. The pausing of urban activity had an increasing impact on these levels, with the construction halt providing the most reduction, highlighting the major contribution this source has on urban noise. With 90% of the 50 quietest days, measured by the SONYC sensor network, occurring during the lockdown period, its impact is clear.

The pronounced daytime SPL decreases observed in the ~ 2 month period of full lockdown reveal the adaptation of NYC residents to social distancing and work from home. With the May transition into warmer Spring days, the reduced evening SPL change suggests that residents were spending more time outside without the option of convening inside bars or restaurants. This reduction in SPL change continued through each reopening phase, however, measured SPL levels in our study period did not reach pre-pandemic levels after the fourth reopening phase.

With the lack of construction significantly contributing to overall reductions in urban noise levels, the observed reduction in the presence of construction related sources (such as *machinery impact, non-machinery impact and powered saw*) further confirms this. However, the largest reduction in measured class presence is seen in *human voice*. This class shows reductions in the order of 10x some other classes, suggesting that human activity and thus social distancing was followed by the majority of NYC residents, especially following the New York State on pause (E2) event. The decreasing trend of 311 noise complaints with SPL reduction give the impression of a period of resident adaptation after construction halted, followed by habituation to these quieter conditions. This appears to break down when increased human activity on the streets (seen in the increased *human voice* and *music* presence detections) in mid-to-late May produce a significant uptick in complaints for this type of noise source.

This work is not without its limitations though. The sensor locations are concentrated in a few neighborhoods. In addition, the machine listening models have a relatively small vocabulary that may be narrowing the lens of our analysis, causing us to miss significant yet unexpected sources of noise. Lastly, the models measure soundclass presence, rather than level [23], which also may hinder our analysis of the relationship of noise to complaints.

In summary, the unique dataset captured by the SONYC sensor network in combination with the application of machine listening has in this paper provided an initial set of observations into the noise conditions of a city in lockdown. Each source of data helps to shed light on a common story of significant change to urban acoustic conditions and how the city's residents responded to this.

7. REFERENCES

- "Wave of silence' spread around world during coronavirus pandemic," https://www.theguardian.com/ science/2020/jul/23/wave-of-silencespread-around-world-during-coronaviruspandemic, Accessed: 2022-10-25.
- [2] "The pandemic is turning the natural world upside down," https://www.theatlantic.com/science/ archive/2020/04/coronavirus-pandemicearth-pollution-noise/609316, Accessed: 2022-10-25.
- [3] "Coronavirus une baisse historique du bruit en ile-defrance," https://www.la-croix.com/France/ Coronavirus-baisse-historique-bruit-Ile-France-2020-03-28-1201086606, Accessed: 2022-10-25.
- [4] Dugald JM Thomson and David R Barclay, "Real-time observations of the impact of covid-19 on underwater noise," *The Journal of the Acoustical Society of America*, vol. 147, no. 5, pp. 3390–3396, 2020.
- [5] Lauren M Smith, Linyan Wang, Kuba Mazur, Michael Carchia, Glen DePalma, Reza Azimi, Samantha Mravca, and Richard L Neitzel, "Impacts of covid-19-related social distancing measures on personal environmental sound exposures," *Environmental Research Letters*, vol. 15, no. 10, pp. 104094, 2020.
- [6] Elizabeth P Derryberry, Jennifer N Phillips, Graham E Derryberry, Michael J Blum, and David Luther, "Singing in a silent spring: Birds respond to a half-century soundscape reversion during the covid-19 shutdown," *Science*, vol. 370, no. 6516, pp. 575–579, 2020.
- [7] "The coronavirus quieted city noise. listen to what's left," https://www.nytimes.com/interactive/ 2020/05/22/upshot/coronavirus-quiet-citynoise.html, Accessed: 2022-10-25.
- [8] Romain Rumpler, Siddharth Venkataraman, and Peter Göransson, "An observation of the impact of covid-19 recommendation measures monitored through urban noise levels in central stockholm, sweden," *Sustainable cities and society*, vol. 63, pp. 102469, 2020.
- [9] César Asensio, Ignacio Pavón, and Guillermo De Arcas, "Changes in noise levels in the city of madrid during covid-19 lockdown in 2020," *The Journal of the Acoustical Society* of America, vol. 148, no. 3, pp. 1748–1755, 2020.
- [10] Francesco Aletta, Tin Oberman, Andrew Mitchell, Huan Tong, and Jian Kang, "Assessing the changing urban sound environment during the covid-19 lockdown period using short-term acoustic measurements," *Noise mapping*, vol. 7, no. 1, pp. 123–134, 2020.
- [11] AA Spivak, DN Loktev, Yu S Rybnov, and VA Kharlamov, "Acoustic noise in moscow during the covid-19 quarantine period in 2020," *Doklady Earth Sciences*, vol. 496, no. 1, pp. 63–65, 2021.
- [12] Daniel Bonet-Solà, Carme Martínez-Suquía, Rosa Ma Alsina-Pagès, and Pau Bergadà, "The soundscape of the covid-19 lockdown: Barcelona noise monitoring network case study," *International Journal of Environmental Research and Public Health*, vol. 18, no. 11, pp. 5799, 2021.

- [13] Germán Said, Agustín Arias, Leonardo Carilli, and Agustín Stasi, "Urban noise measurements in the city of buenos aires during the mandatory quarantine," *The Journal of the Acoustical Society of America*, vol. 148, no. 5, pp. 3149–3152, 2020.
- [14] Juan P. Bello, Claudio Silva, Oded Nov, R. Luke Dubois, Anish Arora, Justin Salamon, Charles Mydlarz, and Harish Doraiswamy, "Sonyc: A system for monitoring, analyzing, and mitigating urban noise pollution," *Commun. ACM*, vol. 62, no. 2, pp. 68–77, Jan 2019.
- [15] Charlie Mydlarz, Mohit Sharma, Yitzchak Lockerman, Ben Steers, Claudio Silva, and Juan Pablo Bello, "The life of a new york city noise sensor network," *Sensors*, vol. 19, no. 6, pp. 1415, 2019.
- [16] F. Miranda, M. Lage, H. Doraiswamy, C. Mydlarz, J. Salamon, J. Lockerman, Y. Freire, and C. Silva, "Time lattice: A data structure for the interactive visual analysis of large time series," *Computer Graphics Forum (EuroVis '18)*, vol. 37, no. 3, pp. 13–22, 2018.
- [17] C Mydlarz, C Shamoon, and J.P Bello, "Noise monitoring and enforcement in new york city using a remote acoustic sensor network," in *In proceedings of the INTER-NOISE and NOISE-CON Congress and Conference (Hong Kong, China, August* 27th–30th). Institute of Noise Control Engineering, Reston, VA, 2017.
- [18] Mark Cartwright, Jason Cramer, Ana Elisa Mendez Mendez, Yu Wang, Ho-Hsiang Wu, Vincent Lostanlen, Magdalena Fuentes, Graham Dove, Charlie Mydlarz, Justin Salamon, Oded Nov, and Juan Pablo Bello, "SONYC Urban Sound Tagging (SONYC-UST): a multilabel dataset from an urban acoustic sensor network," https://doi.org/10.5281/ zenodo.3966543, Sept. 2020.
- [19] Brian McFee, Justin Salamon, and Juan Pablo Bello, "Adaptive pooling operators for weakly labeled sound event detection," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 11, pp. 2180–2193, 2018.
- [20] "Zooniverse," https://www.zooniverse.org/, Accessed: 2023-03-13.
- [21] Jason Cramer, Ho-Hsiang Wu, Justin Salamon, and Juan Pablo Bello, "Look, listen, and learn more: Design choices for deep audio embeddings," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing* (*ICASSP*). IEEE, 2019, pp. 3852–3856.
- [22] "311 service requests from 2010 to present," https: //data.cityofnewyork.us/Social-Services/ 311-Service-Requests-from-2010-to-Present/erm2-nwe9.
- [23] Aurora Cramer, Mark Cartwright, Fatemeh Pishdadian, and Juan Pablo Bello, "Weakly supervised source-specific sound level estimation in noisy soundscapes," in 2021 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA). IEEE, 2021, pp. 61–65.