

DOI:10.1145/3224204

SONYC integrates sensors, machine listening, data analytics, and citizen science to address noise pollution in New York City.

BY 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

NOISE IS UNWANTED or harmful sound from environmental sources, including traffic, construction, industrial, and social activity. Noise pollution is one of the topmost quality-of-life concerns for urban residents in the U.S., with more than 70 million people nationwide exposed to noise levels beyond the limit the U.S. Environmental Protection Agency (EPA) considers harmful.¹² Such levels have proven effects on health, including sleep disruption, hypertension, heart disease, and hearing loss.^{5,11,12} In addition, there is evidence of harmful effects on educational performance, with studies showing noise pollution causing learning and cognitive impairment in children, resulting in

decreased memory capacity, reading skills, and test scores.^{2,5}

The economic impact of noise is also significant. The World Health Organization estimates that, as of 2012, one million healthy life-years in Western Europe were being lost annually to environmental noise.¹¹ Other estimates put the external cost of noise-related health issues in the E.U. between 0.3%–0.4% of GDP¹⁴ and 0.2% of GDP in Japan.¹⁶ Studies in the U.S. and Europe also demonstrate the relationship between environmental noise and real estate markets, with housing prices falling as much as 2% per decibel (dB) of noise increase.^{21,30} Noise pollution is not merely an annoyance but an important problem with broad societal effects that apply to a significant portion of the population. It is clear that effective noise mitigation is in the public interest, with the promise of health, economic, and quality-of-life benefits.

Mitigation

Noise can be mitigated at the receiver's end by, say, wearing earplugs or along the transmission path by, say, erecting sound barriers along major roads. These strategies do not, however, reduce noise emissions but instead put the burden of mitigation on the receiver.¹² Alternatively, noise can be mitigated at the source (such as by designing aircraft with quieter engines, acoustically treating night clubs, muffling jackhammers for roadwork, and stopping unnecessary

>> key insights

- **Public exposure to noise is a growing concern in cities, leading to substantial health, educational and economic costs, but noise is ephemeral and invisible, making it difficult for city agencies to monitor it effectively.**
- **An interdisciplinary effort explores new ways to use both fixed and mobile sensors, with output annotated by citizen scientists, for training novel machine-listening models and analyzing spatiotemporal noise patterns.**
- **The resulting fine-grain and aggregate analytics layers help public agencies monitor the local environment and intervene to mitigate noise pollution.**



honking). These actions are commonly encouraged and incentivized through a regulatory framework that uses fines and other penalties to raise the cost of emitting noise.²⁰ However, enforcing noise codes in large urban areas, to the point where they effectively deter noise emissions, is far from trivial.

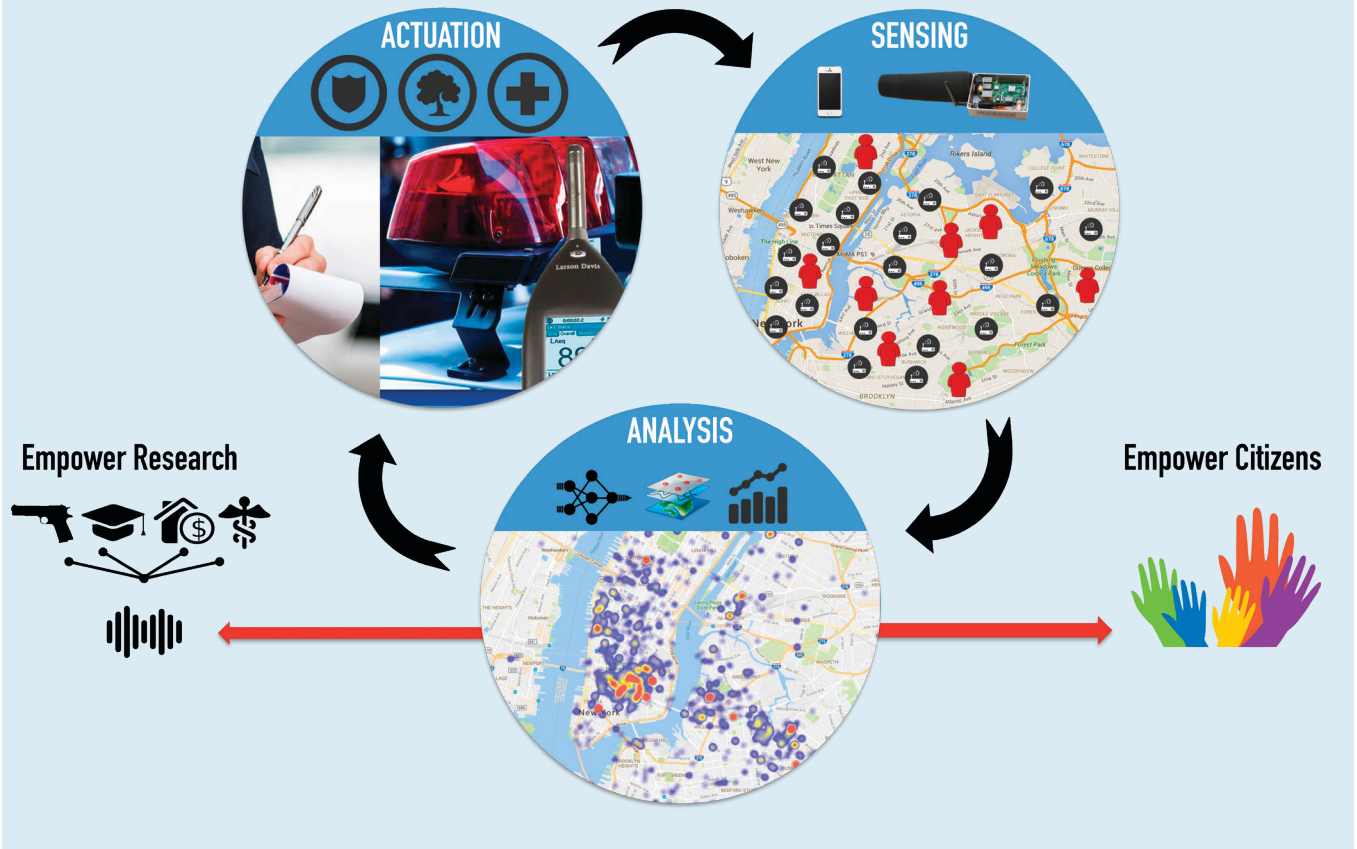
Consider New York City. Beyond the occasional physical inspection, the city government monitors noise through its 311 service for civil complaints. Since 2010, 311 has logged more than 2.7 million noise-related complaints, significantly more than

for any other type of complaint.^a This averages approximately 834 complaints a day, the most comprehensive citizen noise-reporting system in the world. However, research by New York City's Department of Health and Mental Hygiene (DOHMH) found 311 data does not accurately capture information about all noise exposure in the city.²² It identified the top sources of disruptive noise to be traffic, sirens, and construction; the effect to be similar in the boroughs of Manhat-

a <http://www1.nyc.gov/311>

tan, Brooklyn, and the Bronx; and low-income and unemployed New Yorkers among the most frequently exposed. In contrast, 311 noise-complaint data collected for the same period emphasized social noise (such as parties, car alarms, loud talking, music, and TV), with fewer complaints citing traffic or construction. Notably, residents of Manhattan, home to many affluent New Yorkers, are more than twice as likely to file 311 complaints than those in the other boroughs. This pattern clearly highlights the need to collect objective noise measurements across

Figure 1. The SONYC cyber-physical system loop, including intelligent sensing, noise analysis at city-scale, and data-driven mitigation. SONYC supports new research in the social sciences and public health while providing the data citizens need to improve their communities.



the city, along with citizen reporting, to fully characterize the phenomenon.

A closely related challenge involves how to respond to potential violations of the noise code. In New York, the subset of noise complaints pertaining to static, systemic sources (such as construction, animals, traffic, air conditioning, and ventilation units) are routed to the city’s Department of Environmental Protection (DEP), which employs approximately 50 highly qualified inspectors to measure sound levels and issue a notice of violation as needed. Unfortunately, the limited human resources and high number of complaints result in average response times of more than five days. Given the ephemeral nature of sound, a very small proportion of inspections actually result in a violation observed, let alone penalized.

To complicate matters, even when noise sources are active during inspections, isolating their individual effect is difficult. Noise is commonly measured in overall sound pressure levels (SPL) expressed in so-called A-

weighted decibels (dBA)²⁰ that aggregate all sound energy in an acoustic scene. Existing technologies are unable to isolate the effect of offending sources, especially in urban environments flooded with multiple sounds. As a result, inspectors resort to long, complicated measurement strategies that often require help from the people responsible for the violation in the first place, an additional factor contributing to the difficulty and reduced efficiency of the enforcement process.

Here, we outline the opportunities and challenges associated with SONYC, our cyber-physical systems approach to the monitoring, analysis, and mitigation of urban noise pollution. Connecting various subfields of computing, including wireless sensor networks, machine learning, collaborative and social computing, and computer graphics, it creates a potentially transformative solution to this important quality-of-life issue affecting millions of people worldwide. To illustrate this potential, we present findings from an initial study we con-

ducted in 2017 showing how SONYC can help understand and address important gaps in the process of urban noise mitigation.

SONYC

Multiple research projects have sought to create technological solutions to improve the cycle of urban noise pollution. For example, some have used mobile devices to crowdsource instantaneous SPL measurements, noise labels, and subjective responses^{3,24,28} but generally lag well behind the coverage in space-time of civic complaint systems like 311, while the reliability of their objective measurements suffers from a lack of adequate calibration. Others have deployed static-sensing solutions that are often too costly to scale up or go beyond the capabilities of standard noise meters.^{4,23,29} On the analytical side, a significant amount of work has focused on noise maps generated from sound propagation models for major urban noise sources (such as industrial activity and road, rail, and air traffic).^{13,17} However, these maps

lack temporal dynamics and make modeling assumptions that often render them too inaccurate to support mitigation or action planning.¹ Few of these initiatives involve acting on the sensed or modeled data to affect noise emissions, and even fewer have included participation from local governments.¹⁵

SONYC (Sounds of New York City), our novel solution, as outlined in Figure 1, aims to address these limitations through an integrated cyber-physical systems' approach to noise pollution.

First, it includes a low-cost, intelligent sensing platform capable of continuous, real-time, accurate, source-specific noise monitoring. It is scalable in terms of coverage and power consumption, does not suffer from the same biases as 311-style reporting, and goes well beyond SPL-based measurements of the acoustic environment. Second, SONYC adds new layers of cutting-edge data-science methods for large-scale noise analysis, including predictive noise modeling in off-network locations using spatial statistics and physical modeling, development of interactive 3D visualizations of noise activity across time and space to enable better understanding of noise patterns, and novel information-retrieval tools that exploit the topology of noise events to facilitate search and discovery. And third, it uses this sensing and analysis framework to improve mitigation in two ways—first by enabling optimized, data-driven planning and scheduling of inspections by the local government, thus making it more likely code violations will be detected and enforced; and second, by increasing the flow of information to those in a position to control emissions (such as building and construction-site managers, drivers, and neighbors) thus providing credible incentives for self-regulation. Because the system is constantly monitoring and analyzing noise pollution, it generates information that can be used to validate, and iteratively refine, any noise-mitigating strategy.

Consider a scenario in which a system integrates information from the sensor network and 311 to identify a pattern of after-hours jackhammer activity around a construction site. This information triggers targeted inspections by the DEP that results in

an inspector issuing a violation. Statistical analysis can then be used by researchers or city officials to validate whether the action is short-lived in time or whether its effect propagates to neighboring construction sites or distant ones by the same company. By systematically monitoring interventions, inspectors can understand how often penalties need to be issued before the effect becomes long term. The overarching goal is to understand how to minimize the cost of interventions while maximizing noise mitigation, a classic resource-allocation problem that motivates much research in smart-cities initiatives.

All this is made possible by formulating our solution in terms of a cyber-physical system. However, unlike most cyber-physical systems covered in the literature, the distributed and decentralized nature of the noise-pollution problem requires multiple socioeconomic incentives (such as fines and peer comparisons) to exercise indirect control over tens of thousands of subsystems contributing noise emissions. It also calls for developing and implementing a set of novel mechanisms for integrating humans in the cyber-physical system loop at scale and at multiple levels of the system's management hierarchy, including extensive use of human-computer interaction (HCI) research in, say, citizen science and data visualization, to fa-

cilitate seamless interaction between humans and cyber-infrastructure. Worth emphasizing is that this line of work is fundamentally different from current research on human-in-the-loop cyber-physical systems that often focuses on applications in which control is centralized and fully or mostly automated while usually only a single human is involved (such as in assistive robots and intelligent prosthetics). The synthesis of approaches from social computing, citizen science, and data science to advance integration, management, and control of large and variable numbers of human agents in cyber-physical systems is potentially transformative, addressing a crucial bottleneck for the widespread adoption of similar methods in all kinds of socio-technical systems, including transportation networks, power grids, smart buildings, environmental control, and smart cities.

Finally, SONYC uses New York City, the largest, densest, noisiest city in North America, as its test site. The city has long been at the forefront of discussions about noise pollution, has an exemplary noise code^b and, in 311, the most comprehensive citizen noise-reporting system. Beyond noise, the city collects vast amounts of data about everything from public

^b <http://www.nyc.gov/html/dep/html/noise/index.shtml>

Figure 2. Acoustic sensing unit deployed on a New York City street.



safety, traffic, and taxi activity to construction, making much of it publicly available.^c Our work involves close collaboration with city agencies, including DEP, DOHMH, various business improvement districts, and private initiatives (such as LinkNYC) that provide access to existing infrastructure. As a powerful sensing-and-analysis infrastructure, SONYC thus holds the potential to empower new research in environmental psychology, public health, and public policy, as well as empower citizens seeking to improve their own communities. We next describe the technology and methods underpinning the project, presenting some of our early findings and future challenges.

Acoustic Sensor Network

As mentioned earlier, SONYC's intelligent sensing platform should be scalable and capable of source identification and high-quality, round-the-clock noise monitoring. To that end we have developed an acoustic sensor¹⁸ (see Figure 2) based on the popular Raspberry Pi single-board computer outfitted with a custom microelectromechanical systems (MEMS) microphone module. We chose MEMS microphones for their low cost and consistency across units and size, which can be 10x smaller than conventional microphones. Our custom standalone microphone module includes additional circuitry, including in-house analog-to-digital converters and pre-amp stages, as well as an on-board microcontroller that enables preprocessing of the incoming audio signal to compensate for the microphone's frequency response. The digital MEMS microphone features a wide dynamic range of 32dBA–120dBA, ensuring all urban sound pressure levels are monitored effectively. We calibrated it using a precision-grade sound-level meter as reference under low-noise anechoic conditions and was empirically shown to produce sound-pressure-level data at an accuracy level compliant with the ANSI Type-2 standard²⁰ required by most local and national noise codes.

The sensor's computing core is

housed in an aluminum casing we chose to reduce RFI interference and solar heat gain. The microphone module is mounted externally via a flexible metal gooseneck attachment, making it possible to reconfigure the sensor node for deployment in varying locations, including sides of buildings, light poles, and building ledges. Apart from continuous SPL measurements, we designed the nodes to sample 10-second audio snippets at random intervals over a limited period of time, collecting data to train and benchmark our machine-listening solutions. SONYC compresses the audio using the lossless FLAC audio coding format, using 4,096-bit AES encryption and the RSA public/private key-pair encryption algorithm. Sensor nodes communicate with the server via a virtual private network, uploading audio and SPL data at one-minute intervals.

As of December 2018, the parts of each sensor cost approximately \$80 using mostly off-the-shelf components. We fully expect to reduce the unit cost significantly through custom redesign for high-volume, third-party assembly. However, even at the current price, SONYC sensors are significantly more affordable, and thus amenable to large-scale deployment, than existing noise-monitoring solutions. Moreover, this reduced cost does not come at the expense of measurement accuracy, with our sensors' performance comparable to high-quality devices that are orders of magnitude more costly while outperforming solutions in the same price range. Finally, the dedicated computing core opens the possibility for edge computing, particularly for in-situ machine listening intended to automatically and robustly identify the presence of common sound sources. This unique feature of SONYC goes well beyond the capabilities of existing noise-monitoring solutions.

Machine Listening at the Edge

Machine listening is the auditory counterpart to computer vision, combining techniques from signal processing and machine learning to develop systems able to extract meaningful information from sound. In the context of SONYC, we focus on developing computational

methods to automatically detect specific types of sound sources (such as jackhammers, idling engines, car horns, and police sirens) from environmental audio. Detection is a challenge, given the complexity and diversity of sources, auditory scenes, and background conditions routinely found in noisy urban acoustic environments.

We thus created an urban sound taxonomy, annotated datasets, and various cutting-edge methods for urban sound-source identification.^{25,26} Our research shows that feature learning, using even simple dictionary-based methods (such as spherical k-means) makes for significant improvement in performance over the traditional approach of feature engineering. Moreover, we have found that temporal-shift invariance, whether through modulation spectra or deep convolutional networks, is crucial not only for overall accuracy but also to increase robustness in low signal-to-noise-ratio (SNR) conditions, as when sources of interest are in the background of acoustic scenes. Shift invariance also results in more compact machines that can be trained with less data, thus adding greater value for edge-computing solutions. More recent results highlight the benefits of using convolutional recurrent architectures, as well as ensembles of various models via late fusion.

Deep-learning models necessitate large volumes of labeled data traditionally unavailable for environmental sound. Addressing this lack of data, we have developed an audio data augmentation framework that systematically deforms the data using well-known audio transformations (such as time stretching, pitch shifting, dynamic range compression, and addition of background noise at different SNRs), significantly increasing the amount of data available for model training. We also developed an open source tool for soundscape synthesis.²⁷ Given a collection of isolated sound events, it functions as a high-level sequencer that can generate multiple soundscapes from a single probabilistically defined "specification." We generated large datasets of perfectly annotated data in order to assess algorithmic performance as a function of, say, maximum polyphony and SNR ratio,

^c <https://nycopendata.socrata.com>

studies that would be prohibitive at this scale and precision using manually annotated data.

The combination of an augmented training set and increased capacity and representational power of deep-learning models yields state-of-the-art performance. Our current machine-listening models can perform robust multi-label classification for 10 common classes of urban sound sources in real time running on a laptop. We will soon adapt them to run under the computational constraints of the Raspberry Pi.

However, despite the advantages of data augmentation and synthesis, the lack of a significant amount of annotated data for supervised learning remains the main bottleneck in the development of machine-listening solutions that can detect more sources of noise. To address this need, we developed a framework for Web-based human audio annotation and conducted a large-scale, experimental study on how visualization aids and acoustic conditions affect the annotation process and its effectiveness.⁶ We aimed to quantify the reliability/redundancy trade-off in crowdsourced soundscape annotation, investigate how visualizations affect accuracy and efficiency, and characterize how performance varies as a function of audio characteristics. Our study followed a between-subjects factorial experimental design in which we tested 18 different experimental conditions with 540 participants we recruited through Amazon's Mechanical Turk.

We found more complex audio scenes result in lower annotator agreement and that spectrogram visualizations are superior at producing higher-quality annotations at lower cost in terms of time and human labor. Given enough time, all tested visualization aids enable annotators to identify sound events with similar recall, but the spectrogram visualization enables annotators to identify sounds more quickly. We speculate this may be because annotators are able to more easily identify visual patterns in the spectrogram, in turn enabling them to identify sound events and their boundaries more precisely and efficiently. We also found participants learn to use each interface more ef-



It is scalable in terms of coverage and power consumption, does not suffer from the same biases as 311-style reporting, and goes well beyond SPL-based measurements of the acoustic environment.



fectively over time, suggesting we can expect higher-quality annotations with only a small amount of additional training.

We found the value of additional annotators decreased after five to 10 annotators and that having 16 annotators was sufficient for capturing 90% of the gain in annotation quality. However, when resources are limited and cost is a concern, our findings suggest five annotators may be a reasonable choice for reliable annotation with respect to the trade-off between cost and quality. These findings are valuable for the design of audio-annotation interfaces and the use of crowdsourcing and citizen science strategies for audio annotation at scale.

Noise Analytics

One main SONYC promise is its future ability to analyze and understand noise pollution at city-scale in an interactive and efficient manner. As of December 2018, we had deployed 56 sensors, primarily in the city's Greenwich Village neighborhood, as well as in other locations in Manhattan, Brooklyn, and Queens. Collectively, the sensors have gathered the equivalent of 30 years of audio data and more than 60 years of sound-pressure levels and telemetry. These numbers are a clear indication of the magnitude of the challenge from a data-analytics perspective.

We are currently developing a flexible, powerful visual-analytics framework that enables visualization of noise levels in the context of the city, together with other related urban data streams. Working with urban data poses further research challenges. Although much work has focused on scaling databases for big data, existing data-management technologies do not meet the requirements needed to interactively explore massive or even reasonable-size datasets.⁸

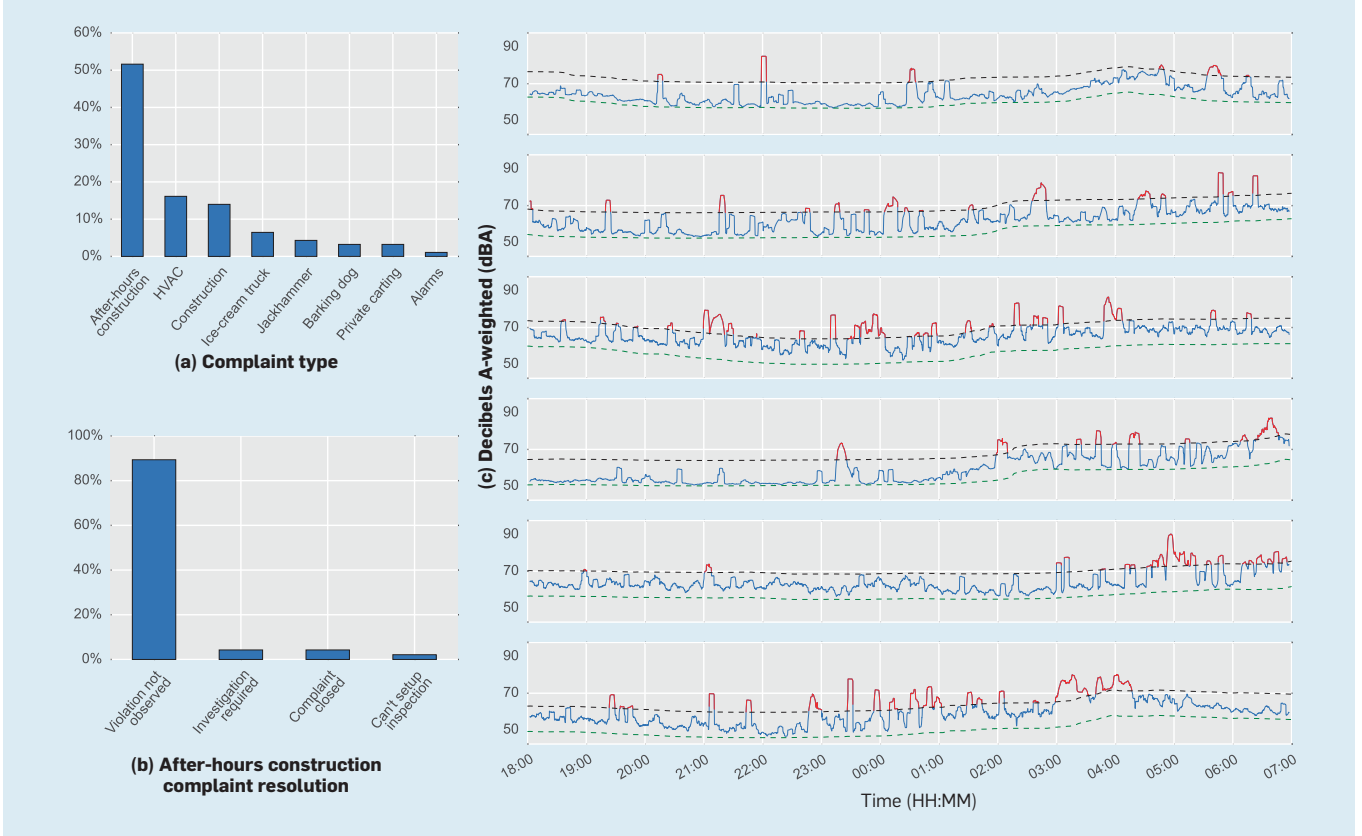
Accomplishing interactivity requires not only efficient techniques for data and query management but for scalable visualization techniques capable of rendering large amounts of information.

In addition, visualizations and interfaces must be rendered in a form that is easily understood by domain experts and non-expert users alike, in-

Figure 3. (left) Interactive 3D visualization of a New York neighborhood using Urbane. By selecting specific sensors (red pins) and buildings (purple) researchers can retrieve and visualize multiple data streams associated with these locations. (right) SPL data at various resolutions and time scales retrieved using the time lattice. Each sub-figure reflects different individual (gray) and aggregated (red) sensor data for the three sensor units highlighted in the left plot.



Figure 4. Case study involving the area around Washington Square Park: (a) Distribution of 311 outdoor noise complaints in the focus area during the study period; the bar graph shows clear predominance of after-hours construction noise. (b) Distribution of complaint resolution for after-hours construction complaints; almost all complaints result in “violation not observed” status. (c) Sensor data for the after-hours period corresponding to six complaints: continuous SPL data (blue), background level (green), event-detection threshold at 10dB above background level (black), and potential noise code violation events (red).



cluding crowdsourcing workers and volunteers, and bear meaningful relationship to the properties of the data in the physical world that, in the case of sound, implies the need for three-dimensional visualization.

We have been working on a three-dimensional, urban geographic information system (GIS) framework called Urbane⁹ (see Figure 3), an interactive tool, including a novel three-dimensional map layer, we developed from the ground up to take advantage of the GPU capabilities of modern computing systems. It allows for fast, potentially real-time computation, as well as integration and visualization of multiple data streams commonly found in major cities like New York City. In the context of SONYC, we have expanded Urbane's capabilities to include efficient management of high-resolution temporal data. We achieve this efficiency through a novel data structure we call the "time lattice" that allows for fast retrieval, visualization, and analysis of individual and aggregate sensor data at multiple time scales (such as hours, days, weeks, and months). An example of data retrieved through this capability can be seen in Figure 3, right plot. We have since used Urbane and the time lattice to support the preliminary noise analysis we cover in the next section, but their applicability goes well beyond audio.

We are currently expanding Urbane to support visual spatiotemporal queries over noise data, including computational-topology methods for pattern detection and retrieval. Similar tools have proved useful in smart-cities research projects, including prior collaborations between team members and the New York City Department of Transportation and Taxi and Limousine Commission.^{7,10}

Data-Driven Mitigation

We conducted a preliminary study in 2017 on the validity and response of noise complaints around the Washington Square Park area of Manhattan using SONYC's sensing and analytics infrastructure.¹⁹ The study combined information mined from the log of civic complaints made to the city over the study period through the 311 system,

the analysis of a subset of our own sensor data during the same period, and information gathered through interactions and site visits with inspectors from the DEP tasked with enforcing the city's noise code.

For the study we chose an area in Greenwich Village with a relatively dense deployment of 17 nodes. We established a 100-meter boundary around each node and merged them to form the focus area. From 311, we collected all non-duplicate noise complaints occurring within this area that had been routed to the DEP while neighboring sensors were active. Note this criterion discards complaints about noise from residents that are routed to the police department and tend to dominate the 311 log; see Figure 4a for a breakdown of selected complaint types.

Over an 11-month period—May 2016 to April 2017—51% of all noise complaints in the focus area were related to after-hours construction activity (6 P.M.–7 A.M.), three times the amount in the next category. Note combining all construction-related complaints adds up to 70% of this sample, highlighting how disruptive to the lives of ordinary citizens this particular category of noise can be.

Figure 4c includes SPL values (blue line) at a five-minute resolution for the after-hours period during or immediately preceding a subset of the complaints. Dotted green lines correspond to background levels, computed as the moving average of SPL measurements within a two-hour window. Dotted black lines correspond to SPL values 10dB above the background, the threshold defined by the city's noise code to indicate potential violations. Finally, we were able to identify events (in red) in which instantaneous SPL measurements were above the threshold. Our analysis resulted in detection of 324 such events we classified by noise source and determined 76% (246) were related to construction as follows: jackhammering (223), compressor engines (16), metallic banging/scraping (7), and the remainder to non-construction sources, mainly sirens and other traffic noise. Our analysis found for 94% of all after-hours construction complaints quantitative evidence in our

sensor data of a potential violation.

How does this evidence stack up against the enforcement record for the complaints? Citizen complaints submitted via 311 and routed to the DEP trigger an inspection, and public-record repositories made available by the city include information about how each complaint was resolved. Examining the records, we found that, for all complaints in this study, 78% resulted in a "No violation could be observed" status and only 2% in a violation ticket being issued. Figure 4b shows, in the specific case of after-hours construction noise, no violation could be observed in 89% of all cases, and none of the inspections resulted in a violation ticket being issued.

There are multiple possible explanations for the significant gap between the evidence collected by the sensor network and the results of the inspections. For example, we speculate it is due in part to the delay in the city's response to complaints, four to five days on average, which is too great for phenomena that are both transient and traceless. Another factor is the conspicuousness of the inspection crew that alone modifies the behavior of potentially offending sources, as we observed during our site visits with the DEP. Moreover, under some circumstances the city government grants special, after-hours construction permits under the assumption of minimal noise impact, as defined by the noise code. It is thus possible that some after-hours activity results from such permits. We are currently mining after-hours-construction-permit data to understand this relationship better.

In all cases, the SONYC sensing and analytical framework is able to address the shortcomings of current monitoring and enforcement mechanisms by providing hard data to: quantify the actual impact of after-hours construction permits on the acoustic environment, and thus nearby residents; provide historical data that can validate complaints and thus support inspection efforts on an inconspicuous and continuous basis; and develop novel, data-driven strategies for the efficient allocation of inspection crews in space and

time using the same tools from operations research that optimize routes for delivery trucks and taxis. Worth noting is that, even though our preliminary study focused on validating 311 complaints, SONYC can be used to gain insight beyond complaint data, allowing researchers and city officials to understand the extent and type of unreported noise events, identify biases in complaint behavior, and accurately measure the level of noise pollution in the local environment.


Looking Forward

The SONYC project is currently in the third of five years of its research and development agenda. Its initial focus was on developing and deploying intelligent sensing infrastructure but has progressively shifted toward analytics and mitigation in collaboration with city agencies and other stakeholders. Here are some areas we intend to address in future work:


Low-power mesh sensor network. To support deployment of sensors at significant distances from Wi-Fi or other communication infrastructure and at locations lacking ready access to electrical power, we are developing a second generation of the sensor node to be mesh-enabled and battery/solar powered. Each sensor node will serve as a router in a low-power multi-hop wireless network in the 915MHz band, using FCC-compatible cognitive radio techniques over relatively long links and energy-efficient multi-channel routing for communicating to and from infrastructure-connected base stations. The sensor design will further reduce power consumption for multi-label noise classification by leveraging heterogeneous processors for duty-cycled/event-driven hierarchical computing. Specifically, the design of the sensor node will be based on a low-power system-on-chip—the Ineda i7^d—for which we are redesigning “mote-scale” computation techniques originally developed for single microcontroller devices to support heterogeneous processor-specific operating systems via hardware virtualization.

Modeling. The combination of noise data collected by sensors and

^d <http://inedasystems.com/wearables.php>



The dedicated computing core opens the possibility for edge computing, particularly for in-situ machine listening intended to automatically and robustly identify the presence of common sound sources.



citizens will necessarily be sparse in space and time. In order to perform meaningful analyses and help inform decisions by city agencies, it is essential for the system to compensate for this sparseness. Several open datasets are available that could, directly or indirectly, provide information on the noise levels in the city; for example, locations of restaurants, night clubs, and tourist attractions indicate areas where sources of social noise are likely, while social media data streams can be used to understand the temporal dynamics of crowd behavior. Likewise, multiple data streams associated with taxi, bus, and aircraft traffic can provide indirect information on traffic-based noise levels. We plan to develop noise models that use spatiotemporal covariance to predict unseen acoustic responses through a combination of sensor and open data. We will also explore combinations of data-driven modeling, applying physical models that exploit the three-dimensional geometry of the city, sound type and localization cues from sensors and 311, and basic principles of sound propagation. We expect that through a combination of techniques from data mining, statistics, and acoustics, as well as our own expertise developing models suitable for GPU implementation using ray-casting queries in the context of computer graphics, we will be able to create accurate, dynamic, three-dimensional urban noise maps in real time.

Citizen science and civic participation. The role of humans in SONYC is not limited to annotating sound. In addition to the fixed sensors located in various parts of the city, we will be designing a SONYC mobile platform aimed at enabling ordinary citizens to record and annotate sounds in situ, view existing data contributed and analyzed by others, and contact city authorities about noise-related concerns. A mobile platform will allow them to leverage slices taken from this rich dataset to describe and support these concerns with evidence as they approach city authorities, regulators, and policymakers. Citizens will not only be more informed and engaged with their envi-

ronment, they will be better equipped to voice their concerns when interacting with city authorities.

Conclusion

SONYC is a smart-cities, next-generation application of a cyber-physical system. Its development calls for innovation in various fields of computing and engineering, including sensor networks, machine learning, human-computer interaction, citizen science, and data science. The technology will be able to support novel scholarly work on the effects of noise pollution on public health, public policy, environmental psychology, and economics. But the project is far from purely scholarly. By seeking to improve urban-noise mitigation, a critical quality-of-life issue, SONYC promises to benefit urban citizens worldwide. Our agenda calls for the system to be deployed, tested, and used in real-world urban conditions, potentially resulting in a model that can be scaled and replicated throughout the U.S. and beyond.

Acknowledgments

This work is supported in part by the National Science Foundation (Award # 1544753), NYU’s Center for Urban Science and Progress, NYU’s Tandon School of Engineering, and the Translational Data Analytics Institute at The Ohio State University. ■

References

1. Ausejo, M., Recuero, M., Asensio, C., Pavón, I., and Pagan, R. Study of uncertainty in noise mapping. In *Proceedings of 39th International Congress on Noise Control Engineering, Internoise* (Lisbon, Portugal, June 13–16). Portuguese Acoustical Society, Lisbon, 2010, 6210–6219.
2. Basner, M., Babisch, W., Davis, A., Brink, M., Clark, C., Janssen, S., and Stansfeld, S. Auditory and non-auditory effects of noise on health. *The Lancet* 383, 9925 (Apr. 2014), 1325–1332.
3. Becker, M., Caminiti, S., Fiorella, D., Francis, L., Gravino, P., Haklay, M. M., Hotho, A., Loreto, V., Mueller, J., Ricchiuti, F. et al. Awareness and learning in participatory noise sensing. *PLoS One* 8, 12 (Dec. 2013), 1–12.
4. Bell, M.C. and Galatioto, F. Novel wireless pervasive sensor network to improve the understanding of noise in street canyons. *Applied Acoustics* 74, 1 (Jan. 2013), 169–180.
5. Bronzaft, A. and Van Ryzin, G. *Neighborhood Noise and Its Consequences: Implications for Tracking Effectiveness of NYC Revised Noise Code. Special Report #14*. Survey Research Unit, School of Public Affairs, Baruch College, CUNY, New York, Apr. 2007; <http://www.noiseoff.org/document/cenyc.noise.report.14.pdf>
6. Cartwright, M., Seals, A., Salamon, J., Williams, A., Mikloska, S., McConnell, D., Law, E., Bello, J., and Nov, O. Seeing sound: Investigating the effects of visualizations and complexity on crowdsourced audio annotations. In *Proceedings of the 21st ACM Conference on Computer-Supported Cooperative*

Work and Social Computing (Jersey City, NJ, Nov. 3–7). ACM Press, New York, 2018, 29:1–29:21.

7. Doraiswamy, H., Ferreira, N., Damoulas, T., Freire, J., and Silva, C. T. Using topological analysis to support event-guided exploration in urban data. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (Dec. 2014), 2634–2643.
8. Fekete, J.-D. and Silva, C. Managing data for visual analytics: Opportunities and challenges. *IEEE Data Engineering Bulletin* 35, 3 (Sept. 2012), 27–36.
9. Ferreira, N., Lage, M., Doraiswamy, H., Vo, H., Wilson, L., Werner, H., Park, M.C., and Silva, C. Urbane: A 3D framework to support data-driven decision making in urban development. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (Chicago, IL, Oct. 25–30), 2015, 97–104.
10. Ferreira, N., POCO, J., Vo, H.T., Freire, J., and Silva, C.T. Visual exploration of big spatiotemporal urban data: A study of New York City taxi trips. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (Dec. 2013), 2149–2158.
11. Fritsch, L., Brown, L., Kim, R., Schwela, D., and Kephelopulos, S. Burden of disease from environmental noise: Quantification of healthy years life lost in Europe. World Health Organization, Bonn, Germany, 2012; <http://www.euro.who.int/en/publications/abstracts/burden-of-disease-from-environmental-noise-quantification-of-healthy-life-years-lost-in-europe>
12. Hammer, M.S., Swinburn, T.K., and Neitzel, R.L. Environmental noise pollution in the United States: Developing an effective public health response. *Environmental Health Perspectives* 122, 2 (Feb. 2014), 115–119.
13. Kaliski, K., Duncan, E., and Cowan, J. Community and regional noise mapping in the United States. *Sound and Vibration* 41, 9 (Sept. 2007), 12.
14. Maibach, M., Schreyer, C., Sutter, D., Van Essen, H., Boon, B., Smokers, R., Schroten, A., Doll, C., Pawlowska, B., and Bak, M. *Handbook on estimation of external costs in the transport sector*. CE Delft, Feb. 2008; https://ec.europa.eu/transport/sites/transport/files/themes/sustainable/doc/2008_costs_handbook.pdf
15. Manvell, D., Marcos, L.B., Stapelfeldt, H., and Sanzb, R. SADMM—Combining measurements and calculations to map noise in Madrid. In *Proceedings of the 33rd Congress and Exposition on Noise Control Engineering (Internoise)* (Prague, Czech Republic, Aug. 22–25). Institute of Noise Control Engineering, Reston, VA, 2004.
16. Mizutani, F., Suzuki, Y., and Sakai, H. Estimation of social costs of transport in Japan. *Urban Studies* 48, 16 (Apr. 2011), 3537–3559.
17. Murphy, E. and King, E. Strategic environmental noise mapping: Methodological issues concerning the implementation of the EU Environmental Noise Directive and their policy implications. *Environment International* 36, 3 (Apr. 2010), 290–298.
18. Mydlarz, C., Salamon, J., and Bello, J. The implementation of low-cost urban acoustic monitoring devices. *Applied Acoustics, Special Issue on Acoustics for Smart Cities* 117, B (Feb. 2017), 207–218.
19. Mydlarz, C., Shamoon, C., and Bello, J. Noise monitoring and enforcement in New York City using a remote acoustic sensor network. In *Proceedings of the INTER-NOISE and NOISE CON Congress and Conference* (Hong Kong, China, Aug. 27–30). Institute of Noise Control Engineering, Reston, VA, 2017.
20. National Academy of Engineering. *Technology for a Quieter America: NAEPR-06-01-A. Technical Report*. The National Academies Press, Washington, D.C., Sept. 2010; <https://www.nap.edu/catalog/12928/Technology-for-a-quieter-america>
21. Nelson, J.P. Highway noise and property values: A survey of recent evidence. *Journal of Transport Economics and Policy* 16, 2 (May 1982), 117–138.
22. New York City Department of Health and Mental Hygiene. *Ambient Noise Disruption in New York City, Data Brief 45*. New York City Department of Health and Mental Hygiene, Apr. 2014; <https://www1.nyc.gov/assets/doh/downloads/pdf/epi/databrief45.pdf>
23. Pham, C. and Cousin, P. Streaming the sound of smart cities: Experimentations on the SmartSantander test-bed. In *Proceedings of IEEE International Conference on Green Computing and Communications, IEEE Internet of Things, and IEEE Cyber, Physical and Social Computing* (Beijing, China, Aug. 20–23). IEEE, Piscataway, NJ, 2013, 611–618.
24. Ruge, L., Altakroui, B., and Schrader, A. Sound of the city: Continuous noise monitoring for a healthy city. In *Proceedings of the IEEE International Conference on*

Pervasive Computing and Communications Workshops (San Diego, CA, Mar. 18–22). IEEE, Piscataway, NJ, 670–675.

25. Salamon, J. and Bello, J. Deep convolutional neural networks and data augmentation for environmental sound classification. *IEEE Signal Processing Letters* 24, 3 (Mar. 2017), 279–283.
26. Salamon, J., Jacoby, C., and Bello, J.P. A dataset and taxonomy for urban sound research. In *Proceedings of the 22nd ACM International Conference on Multimedia* (Orlando, FL, Nov. 3–7). ACM Press, New York, 2014.
27. Salamon, J., McConnell, D., Cartwright, M., Li, P., and Bello, J. SCAPER: A library for soundscape synthesis and augmentation. In *Proceedings of the IEEE Workshop on Applications of Signal Processing to Audio and Acoustics* (Mohonk, New Paltz, NY, Oct. 15–18). IEEE, Piscataway, NJ, 2017.
28. Schweizer, I., Meurisch, C., Gedeon, J., Bärthel, R., and Mühlhäuser, M. Noisemap: Multi-tier incentive mechanisms for participative urban sensing. In *Proceedings of the Third International Workshop on Sensing Applications on Mobile Phones* (Toronto, ON, Canada, Nov. 6–9). ACM Press, New York, 2012, 9.
29. Steele, D., Krijnders, D., and Guastavino, C. The Sensor City Initiative: Cognitive sensors for soundscape transformations. In *Proceedings of GIS Ostrava 2013: Geoinformatics for City Transformation* (Ostrava, Czech Republic, Jan. 21–23). Technical University of Ostrava, 2013.
30. Theebe, M.A. Planes, trains, and automobiles: The impact of traffic noise on house prices. *The Journal of Real Estate Finance and Economics* 28, 2–3 (Mar. 2004), 209–234.

Juan Pablo Bello (jpbello@nyu.edu) is a professor of music technology and computer science and engineering at New York University, New York, USA, and director of the Center for Urban Science of Progress and of the Music and Audio Research Laboratory.

Claudio Silva (csilva@nyu.edu) is a professor of computer science and engineering and data science at New York University, New York, USA.

Oded Nov (onov@nyu.edu) is an associate professor of technology management and innovation at New York University, New York, USA.

R. Luke DuBois (dubois@nyu.edu) is co-director and an associate professor of integrated digital media at New York University, New York, USA.

Anish Arora (arora.9@osu.edu) is a professor of computer science and engineering at The Ohio State University, Columbus, OH, USA.

Justin Salamon (justin.salamon@nyu.edu) is a senior research scientist at the Music and Audio Research Laboratory and the Center for Urban Science and Progress at New York University, New York, USA.

Charles Mydlarz (cmydlarz@nyu.edu) is a senior research scientist at the Music and Audio Research Laboratory and the Center for Urban Science and Progress at New York University, New York, USA.

Harish Doraiswamy (harishd@nyu.edu) is a research assistant professor of computer science and engineering and a research scientist at the Center for Data Science at New York University, New York, USA.

Copyright held by authors.



Watch the authors discuss this work in the exclusive *Communications* video. <https://cacm.acm.org/videos/sonyc>