BreathEasy: Exploring the Potential of Acoustic Sensing for Healthy Indoor Environments

Bhawana Chhaglani, Camellia Zakaria, Jeremy Gummeson, Prashant Shenoy
Manning College of Information and Computer Sciences, University of Massachusetts Amherst
Amherst, USA
bchhaglani@umass.edu,nurcamellia@umass.edu,jgummeso@umass.edu,shenoy@umass.edu

ABSTRACT

The exploration of our envisioned system, *BreathEasy*, offers the alternative utilization of ambient acoustic sensing techniques to promote user awareness and ensure healthy indoor environments. There has been a renewed interest in providing optimal ventilation indoors after the pandemic. While increasing ventilation at all times leads to high energy consumption, prior work uses occupancy-based or air quality-based approaches to modulate ventilation. However, risk assessment is very complex and requires information about activities performed by occupants, space distribution among occupants, and wearing of masks, on top of other indoor environmental factors such as ventilation rate, air filtration, and room dimensions. Here, we investigate the feasibility of using acoustic sensing mechanisms to produce key parameters essential for airborne transmission risk assessment of occupants in indoor spaces.

CCS CONCEPTS

• **Human-centered computing** → *Ubiquitous and mobile computing*; • **Hardware** → *Signal processing systems.*

KEYWORDS

airborne transmission, air quality, audio, ML, privacy

ACM Reference Format:
Bhawana Chhaglani, Camellia Zakaria, Jeremy Gummeson, Prashant Shenoy. 2023. BreathEasy: Exploring the Potential of Acoustic Sensing for Healthy Indoor Environments. In *1st International Workshop on Advances in Environmental Sensing Systems for Smart Cities (EnvSys '23), June 18, 2023, Helsinki, Finland*. ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3597064.3597338

1 INTRODUCTION

Breathing is essential for basic metabolic functions. The air we breathe, primarily oxygen, carbon dioxide (CO_2), and other aerosol particles, can directly affect users' physical health because of the direct interaction between these air pollutants and blood circulating through the lungs. The recent COVID-19 pandemic exemplifies the adverse effects of breathing in unhealthy air simply from performing everyday activities, especially when users are indoors. To minimize the harmful effects of the pandemic, health organizations around the world have emphasized guidelines to ensure optimal ventilation indoors [7]. However, increasing ventilation or filtration at all times can significantly increase HVAC energy consumption [16]. Prior work suggests schedule-based [25], occupancy-based [11], or CO_2-based [19] modulation of ventilation to ensure healthy air circulation while reducing energy consumption. While these techniques effectively save energy, neither occupancy nor CO_2 measures accurately indicate the risk of airborne transmission in indoor spaces. Directly estimating the amount of aerosols generated by human beings that potentially carry infectious diseases, requires tracking the movements and activities across complex human behaviors, which requires overcoming fundamental challenges that have long faced mobile computing.

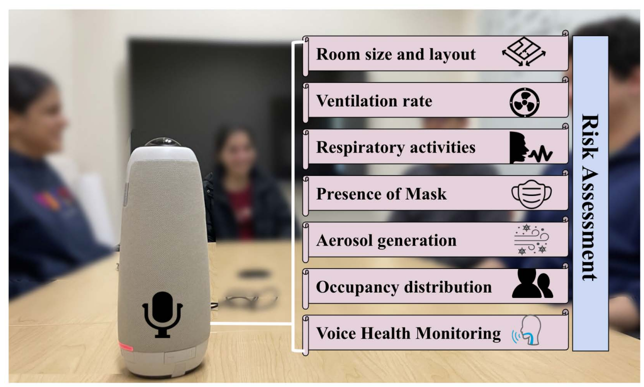

Figure 1: Conceptual visualization of an audio-based sensing system for airborne transmission risk assessment.

The ability to estimate airborne transmission risk an occupant is exposed to from being in an indoor space requires discriminating multiple parameters such as the room size and layout, activities that occupants are involved in, presence of occupants wearing a mask, the distance between occupants, and ventilation rate [3]. Manoharan *et al.* algorithmically identifies the maximum allowable occupancy and corresponding occupant placement in office space by jointly optimizing bio-safety and comfort [20]. As they attempt to optimize occupant placement, the work assumes the knowledge of many HVAC and environment parameters, including the average value of quanta emission rates due to basic human respiratory activities such as counting, whispering, speaking, and breathing. In considering the average emission rate, assessment

EnvSys '23, June 18, 2023, Helsinki, Finland

ACM ISBN 979-8-4007-0214-3/23/06…$15.00
https://doi.org/10.1145/3597064.3597338

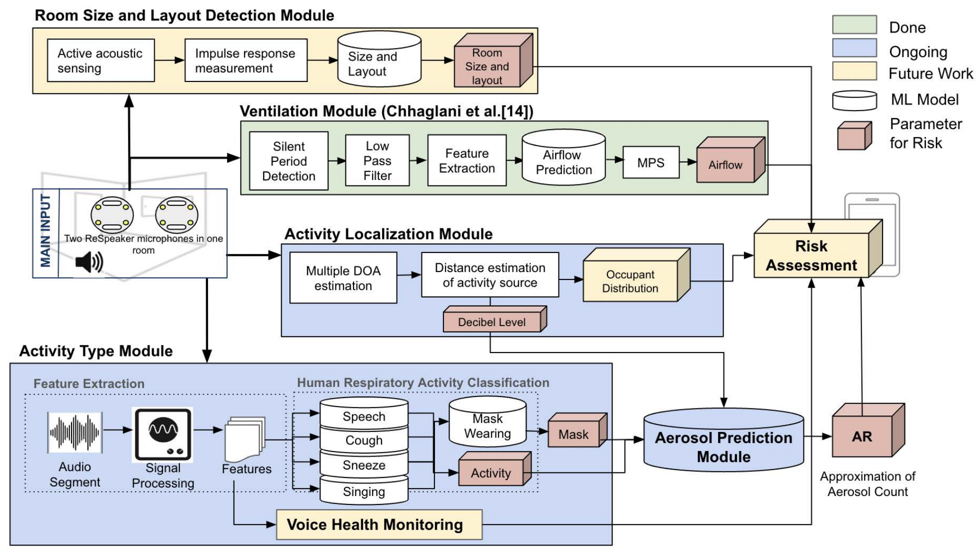

Figure 2: *BreathEasy*: Envisioned system pipeline, where the red boxes indicate key parameters the system can produce to collectively estimate the risk of airborne transmission in indoor spaces.

can be inaccurate since not all occupants in the room may pose a threat to emitting aerosols; this factor is highly dependent on the activities that individual users perform (e.g., staying quiet should not be assessed as risky, while loudly talking should). Kharvari *et al.* proposes a risk estimation tool to identify the maximum number of people occupying a building during the COVID-19 pandemic [16]. However, this assessment relies on extracting the total HVAC load manually and only considers the number of occupants as a parameter for measuring risk exposure. The capabilities presented in these works either assume the knowledge of critical HVAC parameters or determine a subset of the key parameters required for making risk assessments. Thus, there remains a research gap in continuously sensing multiple parameters (i.e., activities type and locations, mask, room size and layout, ventilation rate) automatically and continuously in tandem to accurately determine the risk of airborne transmission.

To address this gap, our work investigates the development of *BreathEasy* using acoustic sensing techniques to accurately predict decision parameters crucial in determining the risk of airborne transmission, as conceptually visualized in Figure 1. Our decision to standardize the sensing modality to an audio source is from considering the recent advancements and increased adoption of smart speakers in common indoor settings such as homes and offices. Further, acoustic sensing has been used for a myriad of applications as microphones and speakers are ubiquitous, including subject localization [8], vital signs monitoring [27], activity detection [18], subject authentication [6, 29], and enhancing user interfaces [14]. However, this technique has yet shown feasibility in ensuring healthy indoor air environments. Specifically, our system aims to achieve accurate prediction of multiple key parameters in guiding proper ventilation and mobilizing healthy indoor air.

2 ENVISIONED AUDIO-SENSING SYSTEM

Our envisioned audio-sensing system is illustrated in Figure 2. Specifically, *BreathEasy* employs two microphone array, each recording samples at a 16 kHz sampling rate and having a built-in capability of direction-of-arrival (DoA) estimation algorithms. These microphone arrays are connected using a Raspberry Pi 3B+. All components of feature extraction and model prediction are done locally on the edge.

Using audio as a sensing modality inherently presents practical challenges related to privacy, primarily if this system is designed for public space deployments. Our system adopts *privacy-first* as its key design choice, where non-reconstructible audio features are extracted on the edge for different processing capabilities. A small set of these features will be used for the four key subcomponents, namely, a) Ventilation Module, b) Activity Type Module, c) Activity Localization Module, and d) Room size estimation and voice health monitoring. To date, several subcomponents of this system have seen successful accomplishments.

2.1 Ventilation Module

Our first subcomponent is named *FlowSense*. It is built with the motivation to sense ventilation rate while empowering occupants with this critical information to assist them in making informed decisions about improving airflow for their surroundings [9]. Through preliminary experiments, we determined that airflow sound from vents resides at a low frequency, and its amplitude correlates with the airflow rate. This key observation led us to develop *FlowSense*, which is a two-step mechanism that utilizes a low-pass filter to obtain low-frequency audio signals and employs machine learning techniques to predict the state of an air vent—whether it is on or off—and the flow rate through active vents. We used Modern devices Rev.p airflow sensor [10] for collecting airflow ground truth and a variety of off-the-shelf microphone sensors built in common smartphones. This work saw vital progress in realizing a larger system because it overcame a significant challenge in preserving user privacy by using audio signals to predict ventilation rate while preserving user privacy.

Specifically, we use Silent Period detection to confine sensing. That is, the system only analyzes audio signals when silent periods are detected (i.e., does not pick up on speech or other ambient noises) to extract audio features. Further, we propose a novel post-processing technique, Minimum Persistent Sensing (MPS), to reduce interference from ambient noise, including ongoing human conversation, office machines, and traffic noises. MPS relies on continuous and constant airflow sound. At the same time, ambient noises are transient and intermittent. Accordingly, the technique extracts the lowest consistent predictions from a fixed number of predictions to report the airflow rate.

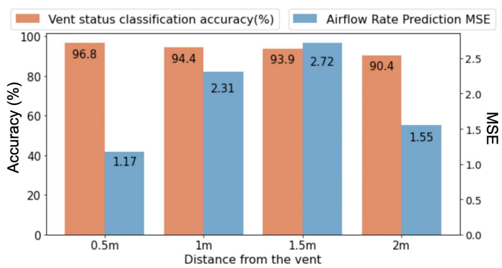

Figure 3: Airflow prediction with varying distances

As per Figure 3, a preliminary experiment on Smartphone (LG Stylo 4) microphone data found high accuracy >90% for classifying vent status (on or off) when a smartphone is within 2 meters distance. As expected, we found a decreasing trend in accuracy with increase in distance between the smartphone's microphone and vent. The MSE of airflow prediction regression model show slight inconsistencies between 1.17-2.72 MSE with increasing distance. However, this MSE is very less and can provide a reasonable estimate of airflow. The proposed algorithm, including the utility of Silent Period detection and MPS, significantly improves the system performance by 77% in ambient noise. In proposing *FlowSense*, we recommend the placement of a microphone source to be within 1.5m away from the vent to predict the rate of airflow.

2.2 Activity Type Module

The second key subcomponent in our system pipeline is the ability to predict aerosol accumulation from subjects in the room. Specifically, prior studies [3] suggest that human respiratory activities such as speaking, coughing, sneezing, and singing are the primary contributors of aerosols produced by humans in indoor environments. Further, the amount of aerosol generation depends on "what is said" and "how loud it is said" [4]. This foundational understanding led us to extend our system implementation to detect everyday human-respiratory activities, including determining whether subjects are wearing masks. The presence of a subject wearing a mask while performing a detected activity can significantly reduce the amount of aerosol generation. As suggested by prior work, aerosol deposition from humans can be reduced drastically, as much as 90% for cough and 74% for speech [2]; thus, it is a vital factor in determining the risk of airborne transmission.

In realizing this functionality, our work is currently investigating the feasibility of detecting activity types using non-reconstructible audio features while maintaining user privacy. The module begins with extracting information from audio signals at the device layer, thereby only collecting a small set of audio features. Accordingly, these features are used as input to orchestrate different activity binary classification models to detect multiple human respiratory-typed activities that might co-occur.

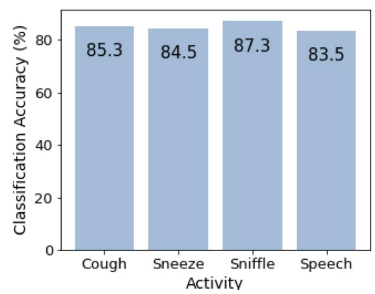

Figure 4: Activity classification Accuracy

We use training data from Google AudioSet [12] and label it using the praatIO library. We use Gradient Boosting Classifiers to train the activity detection models. As per Figure 4, our models achieved accuracies above 80% for classifying cough, speech, sneeze, and sniffle by just using the privacy-preserving features. The model outputs a list of activities (e.g., [talking, coughing, talking, sneezing]) without yet determining the presence of a mask automatically.

Mask detection is a relatively complex problem, since sounds produced by a person can vary with mask type. However, Nguyen *et al.* found that high-frequency components of audio attenuate when a mask is present [24]. This finding motivated our first steps of exploring the possibility of detecting mask presence from discerning masked audio versus non-masked audio.

2.3 Activity Localization Module

Another critical factor that impacts aerosol generation from humans is the loudness of the respiratory activity. For example, particle emission during normal human speech is positively correlated with the loudness ranging from approximately 1 to 50 particles per second (0.06 to 3 particles per cm^3) for low to high amplitudes [3]. The third component of an activity localization module is to achieve two goals; the first is determining activity loudness, and the second is to localize the source of aerosol in a confined space. To localize the activity, we use DoA at two microphone arrays to calculate the distance between the sensor and the activity source (using trigonometric equations). We use this distance estimate to compute the activity loudness at the source.

Similar to the aforementioned subcomponent, this module is a work in progress where we explore the feasibility of detecting the level/loudness of the activity by extrapolating the decibel level of the activity to predict aerosol generated from humans performing common respiratory activities such as talking, coughing and sneezing. Aerosol emission is significantly different from other groups of air pollutants, such as dust, CO_2, and VOCs, as it primarily influences airborne-transmitted virus spread.

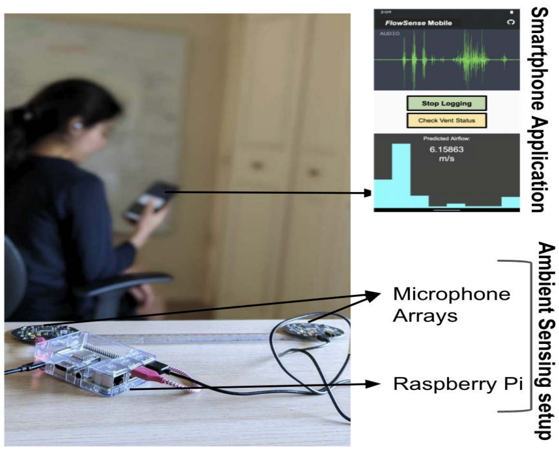

Figure 5: We developed a mobile application to alert users of airflow rate in realtime.

Conclusively, the completed and ongoing efforts of these modules, as shown in Figure 5, can produce three key parameters that will be vital for the implementation of our envisioned system; the ventilation rate and aerosol generation by detecting occupants' activity and influenced by activity loudness.

In what follows, we describe the future advancements we plan to achieve.

2.4 Expanding the Potential of Acoustic Sensing

The advancements in acoustic sensing prove the immense potential of developing a wide range of applications. We believe our envisioned system can be extended to sense even more parameters essential in determining risk levels in indoor environments.

(1) **Room Size and Layout Sensing**: Room size plays a vital role in indoor risk assessment [17]. Acoustic sensing has previously been used for acoustic indoor space mapping [26]. One way is by emitting a chirp signal and analyzing the room impulse response (RIR) to estimate the location of reflectors in the room. Our plan is to build on our current implementation (as per Figure 2) of two microphone arrays placed at a fixed distance apart. In doing so, we can use the impulse responses captured by both microphones to estimate the room size and layout.

(2) **Occupant Distribution**: The number of occupants and their distribution in a room is critical in analyzing the risk of airborne transmission. For example, if people sitting close to each other are participating in various high aerosol generating activities, it will lead to a high risk scenario. Detecting occupant count and distribution is a very hard problem as there are many potential corner cases. Prior work [13, 30] explores the differences between voices of different people by using acoustic features like mel frequency cepstral coefficients (MFCC) to estimate occupancy. However, these techniques are not accurate; for example, when people are not talking or when they are talking simultaneously it isn't possible to correctly identify unique voices. We can combine this approach with other indoor localization-based approaches to infer more accurate occupants' distribution. We can also explore information from ambient noises like detecting door opening/closing and walking noise to improve the occupancy estimation.

(3) **Voice Health Monitoring**: The proposed system can also be extended to monitor the health of occupants using voice features as a digital biomarker. There has been a rich literature on using voice features for continuous monitoring of Parkinson's disease [31], mental health [5], COVID-19 detection [21], and voice pathology detection [15], just to name a few. We plan to implement the above capabilities to enable continuous long-term monitoring of voice health, complementing wearable health sensing.

(4) **Risk Assessment**: Ultimately, solving the automatic detection and estimation of airborne transmission as an indoor risk to occupants is not a straightforward problem. It requires knowledge of a wide range of parameters. For example, Kharvari *et al.* [16] proposes a tool for COVID-19 risk estimation by considering several factors while assuming others. Broadly, it requires a prior information regarding room size/layout, activities that occupants are involved in, mask presence, the distance between occupants, and ventilation/filtration parameters [3]. All these factors are crucial in determining overall risk in indoor spaces. As discussed earlier, we can sense these parameters using ambient acoustic sensing, enabling us to estimate the risk of airborne transmission. Depending on the estimated risk, we can provide alerts to occupants, limit occupancy or exposure time, and/or guide/modulate ventilation or filtration.

(5) **Personalized Sensing**: A future direction includes miniaturizing this capability to run on smartphones or wearable devices where users can be alerted of the risk level they are exposed to in the environment in real time. It is also worth considering earables as a sensing modality since the placement of the device is very close to a user's mouth and nasal cavities, creating an opportunity to obtain more precise features that can improve the overall system accuracy. For example, we can detect breathing rate [1] using wearable devices, which is indicator of amount of aerosols inhaled by the person.

3 RESEARCH CHALLENGES

Enabling the diverse set of applications using solely audio-based sensing requires addressing significant research challenges. We discuss the key design considerations and their associated challenges:

(1) **Detection and Localization of simultaneous activities**: Detecting and localizing simultaneous activities is a challenging problem. We can use the techniques proposed by Wang *et al.* [28] that use a geometric layout of microphones

on the array to determine the unique relationship among signals from the same source along the same arriving path.

(2) **Mobility**: Mobility is another significant challenge in estimating occupants' count or distribution. In a high-mobility environment, tracking the location and number of occupants using only ambient acoustic sensing becomes difficult; thus requires the continuous exploration of more accurate localization techniques, including leveraging other sensing modalities.
(3) **Electronic noise**: Ideally, the proposed system should not detect activities like speech replayed from an electronic speaker. To eliminate the electronic speaker voice, we can use the state-of-the-art techniques proposed in Voice liveness detection [22].
(4) **Hardware heterogeneity**: Different microphone hardware can cause problems generalizing the system. Further, different mobile devices employ different microphones, and their non-linearities can lead to significant errors.
(5) **Low audio activities**: Since we are using ambient acoustic sensing, it is difficult for the sensor to detect very low amplitude activities that might be important for predicting aerosol generation, such as a user's heavy breathing. We can explore radar-based approaches [23] for detecting these activities.

4 CONCLUSION

Our work investigates the possibility of developing *BreathEasy*, an acoustic-based sensing system to provide users with critical parameters informative of airborne transmission risks. Generally speaking, a virus spread in such cases depends on the rate of pollutants generated and the rate of contaminants dissipating. Beyond existing indoor air quality systems measuring for pollutants such as dust, VOCs, CO_2, one group of pollutants that can harm occupants' health is aerosol particles generated from everyday human respiratory activities. In envisioning an audio-based system to predict airborne transmission risk, the system will have to include key subcomponents that can accurately determine activities performed by occupants, space distribution among occupants, and wearing of masks, on top of other indoor environmental factors such as ventilation rate, air filtration, and room dimensions. To date, we have successfully built the first subcomponent of predicting the ventilation rate in a room using only audio features. Our work continues to investigate measuring different parameters representative of aerosol generated by human activities and other dilution factors relevant to this assessment. Through this paper, we discuss the potential applications and research challenges of using acoustic sensing for healthy indoor spaces.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their comments for improving the paper. This research was supported in part by NSF grants 2211302, 2211888, 2213636, 2105494 and and US Army contract W911NF-17-2-0196. Any opinions, findings, conclusions, or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the funding agencies.

REFERENCES

[1] Ahmed, T., Rahman, M. M., Nemati, E., Kuang, J., and Gao, A. Mouth breathing detection using audio captured through earbuds. *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023.
[2] Asadi, S., Cappa, C. D., Barreda, S., Wexler, A. S., Bouvier, N. M., and Ristenpart, W. D. Efficacy of masks and face coverings in controlling outward aerosol particle emission from expiratory activities. *Scientific reports*, 10(1):1–13, 2020.
[3] Asadi, S., Wexler, A. S., Cappa, C. D., Barreda, S., Bouvier, N. M., and Ristenpart, W. D. Aerosol emission and superemission during human speech increase with voice loudness. *Scientific reports*, 9(1):1–10, 2019.
[4] Asadi, S., Wexler, A. S., Cappa, C. D., Barreda, S., Bouvier, N. M., and Ristenpart, W. D. Effect of voicing and articulation manner on aerosol particle emission during human speech. *PloS one*, 15(1):e0227699, 2020.
[5] Chang, K.-h., Chan, M. K., and Canny, J. Analyzethis: Unobtrusive mental health monitoring by voice. *CHI'11 Extended Abstracts on Human Factors in Computing Systems*, pages 1951–1956. 2011.
[6] Chauhan, J., Hu, Y., Seneviratne, S., Misra, A., Seneviratne, A., and Lee, Y. Breathprint: Breathing acoustics-based user authentication. *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*, pages 278–291, 2017.
[7] Chen, C.-Y., Chen, P.-H., Chen, J.-K., and Su, T.-C. Recommendations for ventilation of indoor spaces to reduce covid-19 transmission. *Journal of the Formosan Medical Association*, 120(12):2055, 2021.
[8] Chhaglani, B., Acer, U. G., Jang, S. Y., Kawsar, F., and Min, C. Cocoon: On-body microphone collaboration for spatial awareness. *Proceedings of the 24th International Workshop on Mobile Computing Systems and Applications*, pages 89–95, 2023.
[9] Chhaglani, B., Zakaria, C., Lechowicz, A., Gummeson, J., and Shenoy, P. Flowsense: Monitoring airflow in building ventilation systems using audio sensing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 6(1):1–26, 2022.
[10] Device, M. Wind Sensor Rev. P - Low Cost Anemometer. https://moderndevice.com/product/wind-sensor-rev-p. Online; accessed 23 January 2022.
[11] Erickson, V. L. and Cerpa, A. E. Occupancy based demand response hvac control strategy. *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*, pages 7–12, 2010.
[12] Gemmeke, J. F., Ellis, D. P., Freedman, D., Jansen, A., Lawrence, W., Moore, R. C., Plakal, M., and Ritter, M. Audio set: An ontology and human-labeled dataset for audio events. *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 776–780. IEEE, 2017.
[13] Ghaffarzadegan, S., Reiss, A., Ruhs, M., Duerichen, R., and Feng, Z. Occupancy detection in commercial and residential environments using audio signal. *INTERSPEECH*, pages 3802–3806, 2017.
[14] Gupta, S., Morris, D., Patel, S., and Tan, D. Soundwave: using the doppler effect to sense gestures. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1911–1914, 2012.
[15] Hossain, M. S., Muhammad, G., and Alamri, A. Smart healthcare monitoring: a voice pathology detection paradigm for smart cities. *Multimedia Systems*, 25:565–575, 2019.
[16] Kharvari, F. and O'Brien, W. C-hvac: A practical tool for assessing ventilation capacity for hvac systems during the covid-19 pandemic. *Proceedings of the 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, pages 338–339, 2020.
[17] Kurnitski, J. Ventilation rate and room size effects on infection risk of covid-19. *REHVA Eur. HVAC J*, 57(5):26–31, 2020.
[18] Laput, G., Ahuja, K., Goel, M., and Harrison, C. Ubicoustics: Plug-and-play acoustic activity recognition. *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*, pages 213–224, 2018.
[19] Li, J., Wall, J., and Platt, G. Indoor air quality control of hvac system. *Proceedings of the 2010 International Conference on Modelling, Identification and Control*, pages 756–761. IEEE, 2010.
[20] Manoharan, P., Nagarathinam, S., and Vasan, A. Keep it open, keep it safe: Maximizing space utilization in intelligent bio-safe buildings. *Proceedings of the 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, pages 100–109, 2020.
[21] Meister, J. A., Nguyen, K. A., and Luo, Z. Audio feature ranking for sound-based covid-19 patient detection. *Progress in Artificial Intelligence: 21st EPIA Conference on Artificial Intelligence, EPIA 2022, Lisbon, Portugal, August 31–September 2, 2022, Proceedings*, pages 146–158. Springer, 2022.
[22] Meng, Y., Li, J., Pillari, M., Deopujari, A., Brennan, L., Shamsie, H., Zhu, H., and Tian, Y. Your microphone array retains your identity: A robust voice liveness detection system for smart speaker. *USENIX Security*, 2022.
[23] Nandakumar, R., Gollakota, S., and Watson, N. Contactless sleep apnea detection on smartphones. *Proceedings of the 13th annual international conference on mobile systems, applications, and services*, pages 45–57, 2015.

[24] Nguyen, D. D., McCabe, P., Thomas, D., Purcell, A., Doble, M., Novakovic, D., Chacon, A., and Madill, C. Acoustic voice characteristics with and without wearing a facemask. *Scientific reports*, 11(1):5651, 2021.

[25] Nweye, K. and Nagy, Z. Hvac scheduling based on wi-fi derived occupancy. *Proceedings of the 7th ACM international conference on systems for Energy-Efficient buildings, cities, and transportation*, pages 340–341, 2020.

[26] Pradhan, S., Baig, G., Mao, W., Qiu, L., Chen, G., and Yang, B. Smartphone-based acoustic indoor space mapping. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(2):1–26, 2018.

[27] Wang, T., Zhang, D., Wang, L., Zheng, Y., Gu, T., Dorizzi, B., and Zhou, X. Contactless respiration monitoring using ultrasound signal with off-the-shelf audio devices. *IEEE Internet of Things Journal*, 6(2):2959–2973, 2018.

[28] Wang, W., Li, J., He, Y., and Liu, Y. Symphony: localizing multiple acoustic sources with a single microphone array. *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*, pages 82–94, 2020.

[29] Wang, Z., Ren, Y., Chen, Y., and Yang, J. Toothsonic: Earable authentication via acoustic toothprint. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 6(2):1–24, 2022.

[30] Williams, J., Yazdanpanah, V., and Stein, S. Privacy-preserving occupancy estimation. 2023.

[31] Zhang, H., Song, C., Wang, A., Xu, C., Li, D., and Xu, W. Pdvocal: Towards privacy-preserving parkinson's disease detection using non-speech body sounds. *The 25th annual international conference on mobile computing and networking*, pages 1–16, 2019.