

Developing a Dynamic Control Algorithm to Improve Ventilation Efficiency in a University Conference Room

A Technical Report submitted to the Department of Engineering Systems and the Environment

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Avery Walters

Spring, 2022

Technical Project Team Members

Matthew Caruso

Jason Jabbour

Caleb Neale

Alden Summerville

Avery Walters

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Dr. Arsalan Heydarian, Department of Engineering Systems and the Environment

Developing a Dynamic Control Algorithm to Improve Ventilation Efficiency in a University Conference Room

Matthew Caruso	Jason Jabbour	Caleb Neale	Alden Summerville	Avery Walters
<i>Systems Engineering</i>	<i>Systems Engineering</i>	<i>Systems Engineering</i>	<i>Civil Engineering</i>	<i>Civil Engineering</i>
<i>University of Virginia</i>	<i>University of Virginia</i>	<i>University of Virginia</i>	<i>University of Virginia</i>	<i>University of Virginia</i>
Charlottesville, Virginia	Charlottesville, Virginia	Charlottesville, Virginia	Charlottesville, Virginia	Charlottesville, Virginia
mmc2nh@virginia.edu	jjj4se@virginia.edu	can4ku@virginia.edu	ads3pu@virginia.edu	acw4mhp@virginia.edu

Abstract—A robust heating, ventilation, and air conditioning (HVAC) system is needed to maintain a healthy and comfortable indoor environment. However, HVAC systems are responsible for significant energy usage in the United States, and enhancing current systems and implementing additional HVAC sensing are primary strategies for reducing energy consumption. This research developed an HVAC control algorithm (CA) that optimized ventilation operations within a conference room in the University of Virginia Link Lab. Using indoor air quality (IAQ), occupancy, weather, and HVAC operation data streams, the CA recommended a decision to ventilate or not ventilate the conference room every 15 minutes by comparing the cost of lost occupant productivity due to poor IAQ to the energy cost of ventilating the space. The ventilation decision with lower total cost was recommended. This project addressed scheduling inefficiencies of the current HVAC control system, which operates at full power throughout the day regardless of occupancy status. The CA reduced ventilation during unoccupied periods. The CA was tested over two months of historical data from October to December 2021 and recommended ventilating the conference room 15.13 percent of the time. During the same period, the standard system ventilated the conference room 49 percent of the time. Energy savings due to decreased operation were considerable and averaged 424 dollars per month, although these energy savings came at the cost of lost occupant productivity totaling 522 dollars per month. Future work on lost occupant performance will more accurately model the effects of reduced ventilation. However, annual energy savings of 5,000 dollars from a single conference room is encouraging, and scaling a similar CA to consider a set of rooms or an entire floor of a building could result in substantial energy conservation.

Index Terms—Indoor Air Quality, HVAC Ventilation, Control Algorithm, Energy Efficiency, Optimization, Simulation

I. INTRODUCTION

The SARS-CoV-2 pandemic has demonstrated the need for robust ventilation systems, yet implementing these systems often incurs a significant energy cost: 30% of commercial energy usage is due to HVAC operation. However, a 2017 report from the U.S. Department of Energy cites “Technology Enhancements for Current Systems” as one of four high priority interventions for reducing energy usage, with the top-ranked technology, “Advanced HVAC Sensors”, projected to cut current annual commercial energy use by 3.5 percent [1]. Given that Americans spend 90 percent of their time indoors,

increasing HVAC efficiency while providing high indoor air quality (IAQ) is paramount [2]. Intensive HVAC operation maintains high IAQ, but at a significant energy cost. This research investigates reducing HVAC operation through automation while maintaining high IAQ.

A. Guidelines for Indoor Air Quality

Carbon dioxide (CO₂), volatile organic compounds (VOCs), and fine particulate matter (PM_{2.5}) are the primary effluents that adversely affect productivity and health. Carbon dioxide is a byproduct of metabolic activity and is released into the air through exhalation. In enclosed spaces, CO₂ concentrations can approach levels that cause decreases in productivity [3]–[5]. The American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) recommends indoor CO₂ concentrations not exceed 1,300 parts per million (ppm), yet some offices fail to meet this guidance [6]. VOCs are emitted from solvents in paint, cosmetics, dry-erase markers, and cleaning products and are often present indoors at levels as high as ten times that of outdoor air. High levels of VOCs can cause short-term irritation to the eyes, nose, and throat, or more serious long-term effects like liver damage or cancer [7]. The World Health Organization notes that VOC levels become marginal around 200 parts per billion (ppb) and should not exceed 600ppb [8]. Particulate matter less than 2.5 microns in width is classified as PM_{2.5}, which is generated from vehicle exhaust, burning fossil fuels, cooking, and chemical reactions in the atmosphere [9]. PM_{2.5} can be filtered from building air streams using HEPA (high-efficiency particulate air) or high MERV (Minimum Efficiency Reporting Value) filters, yet buildings that lack these technologies can have elevated indoor PM_{2.5} levels, leading to negative health effects [2]. The EPA maintains a 24-hour maximum PM_{2.5} exposure standard of 12 µg/m³ [10]. Table 1 summarizes acceptable IAQ levels and their impact on productivity.

B. Energy Considerations

Although increased levels of CO₂, VOCs, and PM_{2.5} negatively impact health and performance, the operational cost of

TABLE I
IAQ GUIDELINES AND PRODUCTIVITY IMPACT

Species	Baseline	Moderate	High	Productivity Effect
CO ₂ (ppm)	600 [3]	1000 [3], [4]	2500 [4]	-21% for every 400 past 600 [3], -44-94% at 2500 [4]
VOCs (ppb)	50 [3], [8]	200 [11]	500 [3]	-13% at 100 [3]
PM _{2.5} (μg/m ³)	2 [2]	6	12 [10]	Health effects at 24hr exposure of 12 [10]

continuous ventilation is high. HVAC systems account for 30% of commercial energy consumption, and commercial buildings consume 35% of the electricity use in the United States [1], [12]. Eliminating unnecessary ventilation saves energy, aiding the environment and cutting energy costs. Maintaining high IAQ while reducing energy costs is challenging but feasible. The EPA claims “protecting indoor environmental quality in energy efficiency projects need not hamper the achievement of energy reduction goals” and the CA developed in this research is a strong step forward [13].

II. METHODS

A. Testbed

The study was conducted within a conference room in the Link Lab at the University of Virginia. The Link Lab is equipped with over 300 sensors for IAQ, room occupancy, and Bluetooth connectivity to be used in various research projects. The room is 490 ft² with an approximate volume of 4575 ft³ and can accommodate up to 20 occupants at a central table. The Trane HVAC system that serves the conference room is robust, with efficient components and a MERV-13 filtering system. The primary users of the conference room are graduate students and faculty who conduct research in the Link Lab, and during the period of study (October to December 2021), University of Virginia COVID-19 guidelines required that occupants wore masks. IAQ metrics were pulled from the room using an Awair© brand air quality monitor which provided CO₂, TVOC, PM_{2.5}, and temperature readings. Historical data from the Link Lab’s Building Automation System was pulled to determine energy usage of the HVAC during the study.

B. System Overview

The CA recommends behavior of the Variable Air Volume (VAV) box that ventilates the conference room using a series of three connected models: a statistical model that forecasts binary occupancy status (occupied/not-occupied) within the room, a mathematical model that computes future IAQ values over the next hour for each ventilation state (on/off), and calculations to compute the energy use of each ventilation state. Using findings from [3]-[5], [8], and [10]-[11], the CA then assigns a cost of lost occupant productivity given the modeled IAQ values (see section II-D). The total cost of lost occupant productivity is added to the energy cost of ventilation over the next hour to determine the total cost of a decision to ventilate or not ventilate the room. The output of the CA is a binary decision to ventilate

or not ventilate the room for the next hour based on which state has a predicted lower cost and is computed at 15-minute intervals. Fig. 1 provides an overview of the algorithm:

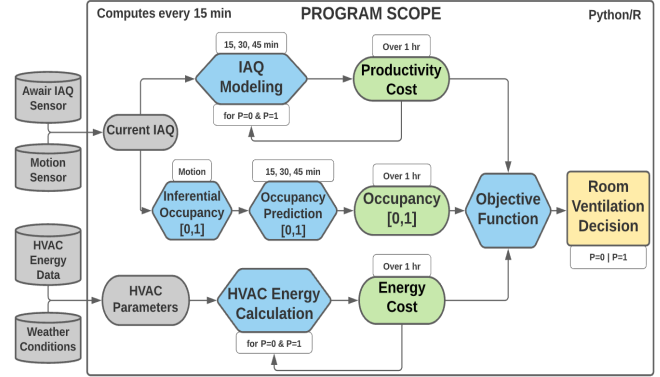


Fig. 1. System Diagram

C. Indoor Air Quality Modeling

Determining future values of IAQ given a decision to ventilate or not ventilate is a key step towards understanding the future impact of a current decision on HVAC operation. Future modeling of each IAQ metric over 15-minute timesteps was handled separately under the decision of ventilation or non-ventilation, and in the cases of the conference room being occupied and unoccupied. Ambient values of CO₂ and total VOCs (TVOCs) were set at 420 ppm and 50 ppb, respectively. The equations in Table 2 return the future value of each IAQ metric in 15 minutes (one timestep) given the current value. Equations for modeling CO₂ and TVOCs were based on information from [14], for PM_{2.5}, [15] was used. In modeling CO₂ and TVOCs, the room air change rate of 2.62 ACH (air changes per hour) was used. Temperature was modeled without strict equations.

TABLE II
IAQ MODELING EQUATIONS

State	CO ₂	TVOCs	PM _{2.5}	Temp
Ventilation, Occupied	1.648x - 272.463	1.648x - 32.436	1.05x	1.02x
Ventilation, Unoccupied	0.519x + 201.834	0.519x + 24.028	0.519x	0.95x
No Ventilation, Occupied	1.40x	1.40x	1.15x	1.10x
No Ventilation, Unoccupied	0.95x	0.95x	0.95x	x

Modeling accuracy, as displayed in Table 3, was usually within 5% of the actual IAQ readings.

D. Indoor Air Quality Cost Calculation

To determine the cost of poor IAQ, a method for converting IAQ levels to a dollar cost of human productivity was developed. Optimal productivity was valued at \$40 per hour per person, with each IAQ metric of CO₂, TVOCs, PM_{2.5}, and temperature contributing \$10 worth of value. Research

TABLE III
IAQ MODELING COMPARISON TO ACTUAL IAQ

IAQ Modeling Comparison to Actual IAQ				
Timestep	CO ₂	TVOC	PM _{2.5} ^a	Temp
15 min	1.41%	-4.46%	-0.34	-0.91%
30 min	-1.41%	-9.03%	-0.54	-2.41%
45 min	-3.44%	-9.72%	-0.66	-3.84%
Net	-1.15%	-7.74%	-0.51	-2.39%

^aNote: % not calculated for PM_{2.5} as the median value is 1ppm

from [3]-[5] defined the loss of productivity due to CO₂ concentrations, which was built on a baseline of 600 ppm, with 20% loss at 1000 ppm, 50% loss at 1500 ppm, and 100% loss at 3000 ppm. These data were trend-fitted in Microsoft Excel to develop the following loss function for CO₂:

$$y = -1.02 \times 10^{-8}x^3 + 4.27 \times 10^{-5}x^2 + 1.67 \times 10^{-3}x - 14.17 \quad (1)$$

The effect of TVOCs on productivity was defined using [3], [8], and [11], and a curve was built on a baseline of 200 ppb, 50% loss at 600 ppb, 75% loss at 1000 ppb, and 100% loss at 2000 ppb. These data were similarly trend-fitted and produced the following loss function for TVOCs:

$$y = 2.85 \times 10^{-8}x^3 - 1.29 \times 10^{-4}x^2 + 0.213x - 37.798 \quad (2)$$

PM_{2.5} influences health more than it influences productivity. However, due to significant health effects influenced by high PM_{2.5} concentrations, PM_{2.5} was included in the objective function. Using data from [10] and [16], the loss curve was built with 0% loss at 2 µg/m³, 25% loss at 6 µg/m³, 50% loss at 12 µg/m³, and 100% loss at 35 µg/m³. This curve had the equation:

$$y = -8.68 \times 10^{-2}x^2 + 6.21x + 0.213x - 11.06 \quad (3)$$

Temperature was included in the model due to its effect on occupant comfort. The loss function for temperature was based on a “goal zone” between 20 and 22.5 degrees C, with significant losses mounting below 15.5 degrees C and above 26.8 degrees C. The curve had the following equation:

$$y = -9.21 \times 10^{-3}x^4 + 7.79 \times 10^{-1}x^3 - 22.69x^2 + 264.32x - 960 \quad (4)$$

In each loss function, the current value of the IAQ metric of consideration is passed in as the independent variable. The function returns a “loss factor” at that value of the metric. Multiplying the value of productivity allotted to that metric over the next 15-minute timestep by the loss factor returns the cost of productivity due to the specific metric over the next 15 minutes. Given that optimal productivity is valued at \$40 per hour, or \$10 per 15 minutes, each IAQ metric can affect a maximum of \$2.50 of loss per 15 minutes, given that value is distributed equally across each metric. Modeled future IAQ values from section II-C are passed to the productivity cost generator to predict future costs of lost productivity. Productivity cost over each timestep is summed to determine the total cost of productivity across the next hour.

E. HVAC Energy Cost Calculation

To determine the cost of energy for operating the HVAC system without direct energy metering, the energy cost of ventilating was divided into four components, with equations for each component defining the relationship between system dynamics and energy cost in dollars. The components calculated are fan energy, heating/cooling energy, dehumidifying energy, and zone reheat energy.

The general process of conditioning air takes two main states: cooling and heating. The system enters a cooling state when outside air is warmer than the desired internal temperature. When in a cooling state, a mix of outdoor and indoor resupply air is passed to the chilled water cooler in the main AHU at the building level. The energy used in cooling is as follows [17]:

$$h_s = 1.08qdt \quad (5)$$

where h_s is the sensible heat energy used by cooling in BTU/hr (then converted to kW/hr), q is the volume (cfm) of air being treated, and dt is the temperature differential (°F) before and after cooling. The air is cooled to 55°F in order to dehumidify the air using latent heat reduction; however, because the cooling process also removes moisture from the air, an additional equation is needed to account for energy used in dehumidification [18]-[21]:

$$h_l = \rho h_{we} q dw_{kg} \quad (6)$$

where h_l is the latent heat energy (kW) used for dehumidification, ρ an assumed constant density of air (kg/m³), q is the volume (cms) of air being treated, h_{we} is the latent heat of vaporization water (kJ/kg), and dw_{kg} is the humidity ratio difference (kg/kg) before and after dehumidification. Once the air has been cooled and dehumidified, it is ducted to zones within the building which can reheat air if necessary. Zone reheating energy is calculated with (5).

When outside air is cooler than the desired internal temperature setpoint, the system enters a heating state. In this state, a mix of outdoor and resupply air is heated to approximately 55 °F at the AHU heating coil. Dehumidification is not a concern in the heating state as cold air holds less moisture than warm air. Air then moves to VAV boxes within the building for reheating if necessary. AHU heating energy is calculated with (5).

During operation, return and supply air fans are used to move air through the HVAC system. The return fan pulls “used” air from the building to be exhausted or reconditioned, while the supply air fan pushes newly conditioned air from the AHU into the building. For each fan, energy consumption is calculated as follows [22]:

$$E_{fan} = 0.746P_{hp}VFD^3 \quad (7)$$

where E_{fan} is the current fan energy consumption (kW), P_{hp} is the horsepower rating of the fan in question, and VFD is the percentage activation of the variable frequency drive controlling the fan.

For the specific case study of the conference room, consideration was only given to AHU-level energy consumption

and VAV-level reheating of the conference room. In a fully developed CA, multiple occupancy forecasts and multiple zone level energy calculations must be made to fully understand IAQ and energy costs across the entire building, yet this was out of the scope of this paper.

In order to estimate the energy impact of a decision to ventilate or not ventilate the conference room, two costs are calculated, depending on if the CA makes a decision that matches historical data.

Cost of ventilating: If the historical data used for simulation was ventilating at a timestep that the CA also determined it was optimal to ventilate, the cost of ventilation was the actual operating cost based on equations (5)-(7) using the historical data of operating parameters.

In the case the CA decided to ventilate at a time without historical operating parameters, an energy cost must be determined using estimated operating parameters. The parameters across all timesteps in which the actual system was ventilating were averaged and used as parameters in the energy calculation equations for an estimated cost of operation. In both of the above cases, actual historical weather data was used.

Cost of not ventilating: Ventilating a room incurs costs at the VAV and AHU level. In the case the CA determines it is optimal not to ventilate, the VAV supply air flow is set to zero. However, the AHU also has a reduced burden in this case, and as such, the AHU supply air flow is adjusted by subtracting the historical VAV supply air flow from the AHU supply. This accounts for costs at both the VAV and AHU level, allowing the approximation of the effects and energy savings of ventilation of just the room under consideration.

Total energy cost under ventilation and nonventilation decisions is calculated by summing equations (5)-(7) to calculate total kilowatt usage, and multiplied by an energy cost of \$0.094 per kilowatt-hour to determine the dollar cost of the decision.

F. Occupancy Inferential Model and Forecasting

The final component of the CA is occupancy. Provided IAQ is only a concern in occupied spaces, understanding occupancy patterns in the room could minimize wasted ventilation. An existing low-cost motion sensor was used to collect motion data within the conference room. That data was sampled at 15-minute intervals to collect binary motion data over the study period. Day-of-week, hour-of-week, and 15-minute period features were created and used to train a statistical model in order to approximate the weekly occupancy patterns of the space. A random forest classifier was trained and tested to be a suitable model, achieving 85% accuracy. It should be noted that the False Negative Rate—the model classifying an unoccupied period when the space is actually occupied—is 12.4%, which should be minimized to achieve optimal IAQ. The performance of this model as shown in Fig. 2 appears reasonable: the model learned the conference room was always vacant during weekends and early/late hours of the day.

G. Objective Function Optimization

The CA was backtested on two months of historical data. The data consists of IAQ and motion readings pulled from sensors




Fig. 2. Occupancy Pattern Learned for Any Given Week in Olsson 211

in the conference room, HVAC operation readings pulled from a UVA Facilities Management database, and weather data pulled from a weather station within 0.15 miles of the building. These fields were cleaned, merged, and filtered for 15 minute intervals using mean resampling and filling any NA values using forward filling. Historical data for each 15 minute timestep was used to calculate the productivity and energy costs of each ventilation state. The lower total cost (productivity + energy) of ventilating or not ventilating serves as the recommendation for each timestep.

There is one case the algorithm handles differently that occurs when the actual system is ventilating, but the algorithm decides to not ventilate: The IAQ values read from the sensors can no longer be used to model the future IAQ as those readings are influenced by the actual system ventilating. To account for this “build up” of the IAQ metrics caused by the algorithm not ventilating, the modeled IAQ values from the last timestep are pushed through and used for the future IAQ modeling. A cascading effect occurs until the productivity cost exceeds the cost of turning the system on, at which the algorithm will recommend to ventilate. The actual form of the optimization equation is as follows:

Objective Function:

$$\text{Min}\{P * C_{\text{En}} + \sum_{t \in T} [C_{\text{IAQ}}(\text{CO}_{2t}, \text{VOC}_t, \text{PM}_{2.5t}, T_t) * O_t]\} \quad (8)$$

Where P is the HVAC ventilation state $\{0, 1\}$, C_{En} is the calculated cost of HVAC energy usage over the next hour, $C_{\text{IAQ}}(\text{CO}_{2t}, \text{VOC}_t, \text{PM}_{2.5t}, T_t)$ is the cost of lost productivity due to IAQ values at timestep, and O_t is the occupancy at timestep $T = \{0, 15, 30, 45\}$

III. RESULTS

The energy and productivity costs of HVAC operation under the CA decisions versus actual operation were calculated and compared. Over the two-month period, the total energy saving using the CA was \$848.14, an average of **\$424.07/month**. This value reflects the dollars saved from decreasing the ventilation of the conference room. The algorithm recommended ventilation **15.13%** of the time, compared to the 49% scheduled operation of the actual system. As seen in Fig. 3, which details one week of the CA decision-making, this decrease in operation is mostly on the weekends when occupancy is low.




Fig. 3. CA Recommendation vs Actual System Ventilation

While the energy saving is impressive, it came at an estimated cost of \$1,043.98 productivity dollars over the two-month period, for an average loss of **\$522/month**. Therefore, the CA incurred a net cost of \$97.93/month. However, the CA achieves energy savings of \$5,089/year, and productivity losses are limited: the average productivity loss under the CA is only \$0.29/hour with a maximum of \$8.69/hr, compared to actual average loss of \$0.11/hour with a maximum of \$3.20/hour.

IV. DISCUSSION

A. Results Discussion

A primary concern of these results is the high productivity cost. Maintaining high IAQ is important for health and performance, but the CA was not able to simultaneously improve IAQ and reduce energy consumption. However, “productivity cost” is a calculated parameter with less concrete significance than energy savings. Additionally, HVAC operation is already adept at maintaining healthy indoor air: IAQ metrics rarely reach unhealthy levels, and the average hourly productivity cost is below \$1. Under the CA’s ventilation decisions, occupants lose less than 2.5% of their performance due to decreased IAQ compared to standard ventilation. Due to the marginal decrease in IAQ, energy savings are a prime object of optimization, and the energy savings of the control algorithm are justified. A comparison of IAQ values under standard HVAC operation and projected IAQ values under operation of the CA is shown below in Table 4: the standard HVAC system, as well as the CA, both maintain healthy IAQ according to benchmarks defined in Table 1. An important limitation of the CA that explains the high maximum IAQ values is the limited IAQ “build up” methodology. As explained in section II-G, the IAQ modeling equations lead to artificially high readings because there is no ceiling to the adjustment equation. In reality, IAQ would approach a steady state during unoccupied periods. However, the CA offers a valuable foundation that could be made more accurate with additional occupancy detail.

With additional time and funding, the assumptions and limitations of this research could be more fully developed. Core assumptions and significant limitations should be duly noted, and present exciting opportunities for future research.

TABLE IV
AVERAGE AND MAXIMUM IAQ PARAMETERS DURING STUDY PERIOD
UNDER ACTUAL AND CA OPERATION

Species	Actual Operation		CA Operation ^a	
	Average	Maximum	Average	Maximum
CO2 (ppm)	464.1	1053.0	577.8	21,594.0
TVOC (ppb)	180.0	4388.9 ^b	211.0	6926.6
PM2.5 ($\mu\text{g}/\text{m}^3$)	1.5	9.0	1.3	12.9
Temperature ($^{\circ}\text{C}$)	22.4	24.5	23.47	72.49

a. Limitations of IAQ buildup. b. Likely result of sensor malfunction.

B. Assumptions

The energy calculations conflate the cost of ventilating the VAV box that serves the conference room with the cost of operation of the AHU, which serves half of the entire building floor. This simplification causes calculated energy costs to be much larger than the actual cost to ventilate just the conference room of study. Assumptions were also made in computing energy savings. As stated in section II-E, the energy savings are calculated as the energy saved by only altering VAV operation. This is accomplished by subtracting the VAV supply airflow from total supply airflow at the AHU. However, the HVAC system parameters are complexly linked, and additional parameters such as VFD setpoints would be affected. These changes were not accounted for in the energy calculations. Attaining a direct energy meter reading would simplify this matter. Assumptions of occupancy must also be addressed. The “actual room occupancy” was determined using a single infrared motion sensor: a dedicated occupancy count sensor would provide a stronger occupancy determination. Finally, optimal productivity was valued at \$40/hour to reflect the general salaries of the most probable room occupants (undergraduates, graduates, faculty). Changing this value directly affects the calculated IAQ costs, and more research could provide a more accurate estimate.

C. Limitations

The main identified limitations of this project are as follows: the short time period of testing (2 months) cannot account for seasonal changes present in the system; IAQ is modeled using generalized mathematical equations rather than a model specifically trained for this use case and under the given system dynamics; due to the occupancy prediction method, this algorithm will only work during the academic school year as it was not trained on data for winter, spring, and summer breaks, and is not currently set up to learn new patterns online; productivity cost is only for a single occupant due to a binary occupancy forecasting, where a room could have n occupants and therefore should charge $\$40*n$ per hour instead of the assumed \$40 per hour; and, the algorithm can currently only be run on historical data.

V. CONCLUSION

Optimizing HVAC control is a thorny problem, but given the limited timeframe of this research, the results are promising.

Producing energy savings of almost \$5,000/year for the optimization of a single conference room is remarkable, although that saving comes at the expense of decreased productivity due to marginally worse IAQ. Primary takeaways include the confirmed difficulty of optimizing HVAC, calculating the energy cost of ventilating a single room solely using energy equations, and predicting room occupancy. Nevertheless, this project is a strong proof of concept. Along with addressing the assumptions and limitations above, other areas of improvement include: implementing a more robust set of energy cost calculations (i.e. specific cost for each room), extending the occupancy classification and prediction from a binary value to an occupancy level (low, medium, high, or specific values), backtesting on a wider timeframe, developing a real-time system and addressing security concerns, and generalizing the algorithm to any room or building.

ACKNOWLEDGMENTS

The authors thank Alan Wang and Nabeel Nasir for technical assistance, the members of the UVA Link Lab for maintaining the indoor sensor networks and Doug Livingston from the University of Virginia Facility Management office for providing data and mentorship.

REFERENCES

- [1] W. Goetzler, R. Shandross, J. Young, O. Petrichenko, D. Ringo, and S. McClive, "Energy savings potential and RDD opportunities for commercial building HVAC systems," U.S. Department of Energy, Washington D.C., USA, 2017. <https://www.energy.gov/sites/prod/files/2017/12/f46/bto-DOE-Comm-HVAC-Report-12-21-17.pdf>. [Accessed: Mar. 12, 2022].
- [2] J. G. Allen and J. D. Macomber, *Healthy Buildings: How Indoor Spaces Drive Performance and Productivity*. Boston, MA: Harvard University Press, 2020. [E-book]. <http://www.jstor.org/stable/j.ctvz0h97h.6>
- [3] J. G. Allen, P. MacNaughton, U. Satish, S. Santanam, J. Vallarino, J., and J. D. Spengler, "Associations of cognitive function scores with carbon dioxide, ventilation, and volatile organic compound exposures in office workers: A controlled exposure study of green and conventional office environments," *Environmental Health Perspectives*, vol. 124, no. 6, pp. 805-812, 2016. 10.1289/ehp.1510037. [Accessed: Oct. 16, 2021].
- [4] U. Satish, M. J. Mendell, K. Shekhar, T. Hotchi, D. Sullivan, S. Streufert, and W. J. Fisk, "Is CO₂ an indoor pollutant? Direct effects of low-to-moderate CO₂ concentrations on human decision-making performance," *Environmental Health Perspectives*, vol. 120, no. 12, pp. 1671-1677, 2012. <https://doi.org/10.1289/ehp.1104789>. [Accessed: Oct. 9, 2021].
- [5] J. G. Cedeño Laurent, P. MacNaughton, E. Jones, A. S. Young, M. Bliss, S. Flanigan, J. Vallarino, L. J. Chen, X. Cao, and J. G. Allen, "Associations between acute exposures to PM_{2.5} and carbon dioxide indoors and cognitive function in office workers: A multicountry longitudinal prospective observational study," *Environmental Research Letters*, vol. 16, no. 9, article 094047, 2021. 10.1088/1748-9326/ac1bd8. [Accessed: Nov. 1, 2021].
- [6] American Society of Heating, Refrigerating and Air-Conditioning Engineers, "Ventilation for acceptable indoor air quality," ANSI/ASHRAE Addendum d to ANSI/ASHRAE Standard 62.1-2016, 21 Feb. 2018. https://www.ashrae.org/File/%20Library/Technical/%20Resources/Standards/%20and/%20Guidelines/Standards/%20Addenda/62.1-2016/62_1_2016_d_20180302.pdf. [Accessed: Mar. 29, 2022].
- [7] Environmental Protection Agency, "Volatile organic compounds' impact on indoor air quality," *United States Environmental Protection Agency*, 2021. <https://www.epa.gov/indoor-air-quality-iaq/volatile-organic-compounds-impact-indoor-air-quality>. [Accessed: Nov. 23, 2021].
- [8] World Health Organization Occupational and Environmental Health Team, "Guidelines for air quality," *World Health Organization*, 2000. <https://apps.who.int/iris/handle/10665/66537>. [Accessed: Nov. 7, 2021].
- [9] New York State Department of Health, "Fine particles (PM_{2.5}) questions and answers," *New York State*, 2018. https://www.health.ny.gov/environmental/indoors/air/pmq_a.htm. [Accessed: Nov. 23, 2021].
- [10] United States Environmental Protection Agency, "National ambient air quality standards (NAAQS) for PM," *United States Environmental Protection Agency*, 2020. <https://www.epa.gov/pm-pollution/national-ambient-air-quality-standards-naaqs-pm>. [Accessed: Mar. 29, 2022].
- [11] D. P. Wyon, "The effects of indoor air quality on performance and productivity," *Indoor Air*, vol 14, no. 7, pp. 92-101, 2004. 10.1111/j.1600-0668.2004.00278.x. [Accessed: Oct. 8, 2021].
- [12] Office of Energy Efficiency and Renewable Energy, "About the commercial buildings integration program," *U.S. Department of Energy*, 2020. <https://www.energy.gov/eere/buildings/about-commercial-buildings-integration-program#:~:text=Commercial%20Building%20Basics&text=Commercial%20buildings%20consume%2013.6%20quads,all%20U.S.%20carbon%20dioxide%20emissions>. [Accessed: Apr. 6, 2022].
- [13] United States Environmental Protection Agency, "Energy cost and IAQ performance of ventilation systems and controls – Executive summary," EPA Office of Air and Radiation, EPA-4-2-S-01-001, 2000. https://www.epa.gov/sites/default/files/2015-01/documents/energy_executive_summary.pdf. [Accessed: Apr. 3, 2022].
- [14] J. Allen, J. Spengler, E. Jones, and J. Cedeño-Laurent, "5-step guide to checking ventilation rates in classrooms," Harvard T.H. Chan School of Public Health, 2020. <https://schools.forhealth.org/ventilation-guide/>. [Accessed: Mar. 3, 2022].
- [15] N. S. Panji, and M. P. Varnosfaderani, "Indoor airborne transmission of COVID-19: Studying the impact of building ventilation and portable air purifying systems using low-cost PM_{2.5} sensors," 2021, unpublished. [Accessed: Apr. 3, 2022].
- [16] United States Environmental Protection Agency, "EPA to reexamine health standards for harmful soot that previous administration left unchanged," *United States Environmental Protection Agency*, 2021. <https://www.epa.gov/newsreleases/epa-reexamine-health-standards-harmful-soot-previous-administration-left-unchanged>. [Accessed: Mar. 29, 2022].
- [17] Engineering ToolBox, "Cooling and heating equations," Engineering ToolBox, 2004. https://www.engineeringtoolbox.com/cooling-heating-equations-d_747.html. [Accessed: Mar. 20, 2022].
- [18] Engineering ToolBox, "Moist air – Enthalpy," Engineering ToolBox, 2004. https://www.engineeringtoolbox.com/enthalpy-moist-air-d_683.html. [Accessed: Mar. 20, 2022].
- [19] Engineering ToolBox, "Air-humidity ratio," Engineering ToolBox, 2004. https://www.engineeringtoolbox.com/humidity-ratio-air-d_686.html. [Accessed: Mar. 20, 2022].
- [20] Engineering ToolBox, "Mixing of humid air," Engineering ToolBox, 2004. https://www.engineeringtoolbox.com/mixing-humid-air-d_694.html. [Accessed: Mar. 20, 2022].
- [21] Engineering ToolBox, "Moist air – Cooling and dehumidifying," Engineering ToolBox, 2004. https://www.engineeringtoolbox.com/cooling-dehumidifying-air-d_695.html. [Accessed: Mar. 20, 2022].
- [22] Engineering ToolBox, "Fan affinity laws," Engineering ToolBox, 2003. https://www.engineeringtoolbox.com/fan-affinity-laws-d_196.html. [Accessed: Mar. 20, 2022].