Creating a Visual Interface to Display Indoor Conditions

Kayla Sprincis
Department of Computer Science
University of Virginia
Charlottesville, VA
xcx8fh@virginia.edu

Logan Brock
Department of Civil and Environmental Engineering
University of Virginia
Charlottesville, VA
ukq5gu@virginia.edu

Xiling Meng
Department of Computer Science
University of Virginia
Charlottesville, VA
wae7xs@virginia.edu

Danielle Sydow
Department of Systems and Information Engineering
University of Virginia
Charlottesville, VA
wgg4jh@virginia.edu

ABSTRACT
Modern building sensors can provide valuable information regarding occupants' health, comfort, and safety. However, the scale and complexity of collected time-series data often make the data's real-world implications inaccessible to non-technical users. This paper presents a visual interface designed to translate real-time sensor measurements for six key building metrics – temperature, carbon dioxide, volatile organic compounds, humidity, noise, and illumination – into a digitally displayed, intuitive color-coded diagram. The system combines simple time-series analytics with a large language model (LLM), which interprets sensor readings, trends, and predefined health-related thresholds to select representative hexadecimal colors. A Streamlit web application displays the model's outputs in a dynamic packed bubble chart, which adapts to emphasize the most critical metric levels at any given time. Tests for consistency demonstrate that the LLM's color assignments remain stable for identical metric inputs and vary appropriately across distinct environmental conditions. The resulting interface streamlines interpretation of indoor conditions, rendering building data more accessible to technical and non-technical audiences alike. This work showcases the capacity of AI-assisted visualization to generalize IoT data, while highlighting future opportunities to improve accessibility, interpretability, and reliability through user studies and refined prompt engineering.

CCS Concepts
Visualization systems and tools; human computer interaction (HCI); decision support systems; sensor networks; machine learning; natural language processing

Keywords
IoT Sensing; generative AI; large language models; prompt engineering; color encoding; user-centered design; accessibility

1. INTRODUCTION
Indoor building conditions are important to occupants' health. There are many different metrics to examine, but a few can be monitored in real time with sensing techniques. Sensing techniques include temperature, carbon dioxide, volatile organic compounds (VOCs), humidity, noise, and illumination. Within these metrics, some values are optimal to keep occupants healthy and comfortable. The optimal range for temperature is 20-25 degrees Celsius, for carbon dioxide it is below 1000 parts per million (ppm), for VOCs it is below 100 parts per billion (ppb), for humidity it is between 30-50%, for noise it is below 55 decibels (dB), and for illumination it is between 100-500 lux [3, 4, 5, 6, 10]. Conditions for these metrics are summarized in Table 1. By acquiring real-time data from sensor devices for these metrics, the health conditions of indoor spaces can be monitored.

Table 1. Metric Values in Relation to Indoor Health

Metric	Values				
Temperature (°C)		Low: < 20	Good: 20-25	High: > 25	
CO2 (ppm)	Good: < 700	Acceptable: 700-999	Warning: 1000-5000	Dangerous: > 5000	
VOCs (ppb)		Good: < 30	Acceptable: 30-100	High: > 100	
Humidity (%)		Low: < 30	Acceptable: 30-50	High: > 50	
Noise (dB)	Quiet: < 30	Ambient: 30-44	Audible: 45-55	Loud: 56-70	Dangerous: > 70
Illumination (lux)	Dark: < 100	Lit: 100-299	Well Lit: 300-500	Bright: > 500	

The enormous amount of data generated by sensors in Internet of Things (IoT) applications is a growing challenge. Traditional linguistic summarization methods have been explored to aggregate sensor patterns. A linguistic summary typically includes a quantity in agreement (Q), a summarizer (S), and a truth measure (T), providing user-friendly descriptions, such as "most temperatures are high (0.8)" [24]. Previously, data summarization relied on statistical approaches that often struggled to capture contextual relationships or connotative meanings in complex datasets [23].

Artificial intelligence (AI) driven summarization tools are emerging as vital technologies for converting unstructured or complex data into clear representations while preserving the original context [2]. An AI summarizer uses large language models (LLMs) to analyze and condense information into meaningful, concise summaries. These AI-powered systems leverage natural language processing (NLP) and deep learning to generate contextually relevant summaries [23]. Unsupervised and semi-supervised learning methods are crucial when ground-truth labels are unavailable, enabling models to adapt over time through reinforcement learning strategies. AI-enabled IoT systems (AIoT) can continually refine their knowledge bases through iterative learning, thereby enhancing resilience and understanding as new data is collected [13, 14]. Recent surveys reveal that transformer-based LLMs have expanded the scope of automated summarization. Today's AI summarizers can integrate textual, numerical, and multimodal data to produce summaries across a range of fields beyond scientific research [9].

A common strategy for presenting summaries and communicating data findings in an approachable format is to create a visual [7]. Visualizations enable viewers to make quicker judgments and draw more memorable conclusions from data than do text-based presentations [7]. It has been widely shown that people can process visuals faster than text, with speeds ranging from 6 to 6,000 times faster [18]. An effective visual is designed to require minimal working memory from the viewer, guide attention, and adhere to commonly established conventions within the format and subject [7]. For example, an image of a sun commonly conveys brightness, and it would be against established conventions to use that graphic to represent sound level. General best practices for visualizations include limiting animations to prevent overcomplication, reducing clutter, and highlighting significant points [7].

The use of color in visuals can affect its interpretation, and it is a factor that can be understood regardless of reading and language abilities [11]. Color can be used to both make a visual easier to understand and provide information. Using different contrast levels can draw the user's attention to certain parts of the visual that are more important to the overall message [11]. Brighter, warmer colors are more attention-grabbing and can guide the viewer's gaze [11]. Colors can also help convey information, as people associate them with different emotions, and they can influence viewers to take a positive or negative takeaway from the visual [17]. The basic six emotions of anger, disgust, fear, happiness, sadness, and surprise were found in previous work to be associated with the colors red, green, black, yellow, blue, and brightness, respectively [17]. This innate color association allows a visual to be created with the intention of influencing how the viewer feels about the subject matter.

When evaluating how users interact with data summaries and visuals, it is essential to account for differences in technological literacy across groups. Due to differences in age and physical development, both younger and older populations may perceive technology-related information differently from the general public. On the one hand, children around the age of six often begin engaging with the internet, even before they have fully developed their reading and writing skills. Research suggests that boys tend to access technology-related information more effectively when it is embedded in video game formats or interactive systems. In contrast, girls often find it more engaging and accessible when presented through storytelling approaches [12]. On the other hand, older adults, who may have less familiarity with digital technologies and experience age-related declines in sensory processing (such as vision and hearing), can benefit from supportive tools. Incorporating "virtual assistants" such as Amazon Echo or Alexa can provide a more intuitive and accessible way for them to interact with technology [8].

Recent advances in generative AI have led to models capable of producing images that surpass human-created visuals in realism, clarity, and emotional appeal. Previous studies have demonstrated that AI-generated marketing content was often perceived as more communicative than professionally designed imagery and as more real than its authentic counterpart [8]. The burgeoning visual capabilities of generative AI, combined with its ease of use via text prompts, suggest that AI may be ideal for translating data streams from building sensors into intuitive and expressive visual figures.

This project proposes a system that summarizes building data into an easy-to-understand visual. It will show occupants the current building conditions through a grid-like structure of labeled color bubbles, where the color and size depend on the data type and value. It will utilize AI to generate the appropriate color to represent the data. The layout will allow all occupants without significant visual impairment to quickly understand building conditions. Due to time and resource constraints, this study will focus solely on the development and design of the visual interface, rather than on its implementation in a physical space. This study will provide proof of concept for a visual interface to display indoor conditions. While user studies are not within the scope of this project, the research will provide a framework for future studies.

2. PROBLEM STATEMENT

Buildings equipped with sensor technology can provide their occupants with real-time data on indoor environmental conditions. However, due to the scale of the data and the technical knowledge required to understand the database, it can be difficult to quickly analyze and use this data to derive actionable insights into buildings. When occupants can easily understand the building's current condition, they can take action to improve health, comfort, or even energy efficiency.

3. MOTIVATION

Building management systems (BMS) receive a continuous stream of data from sensors, including measurements of air quality, temperature, occupancy, and energy use. This type of data is typically presented in graphical form, enabling the analyst to identify trends and anomalies easily. While a graphical representation may be suitable for an experienced user, the complex data would be challenging for an inexperienced viewer, such as a child, to understand. The project aims to develop a method for presenting data that enables a wide range of users to quickly and intuitively understand the conditions of the surrounding environment.

4. METHODOLOGY

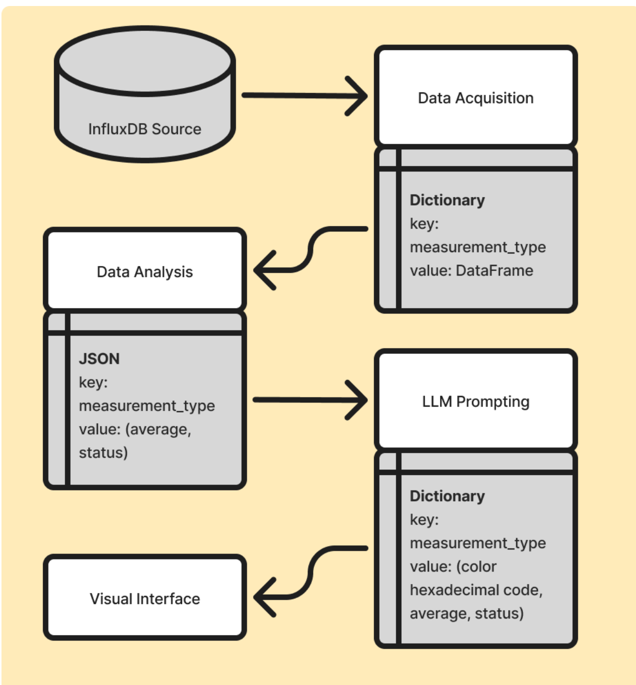

Figure 1. Process Flow Diagram

As illustrated in Figure 1, creating a visual interface for indoor conditions can be broken down into four components: data acquisition, analysis, LLM prompting, and visualization. The state of the data can be modularly broken down into its constituent parts, as discussed below and illustrated in Figure 1.

4.1 Data Acquisition

The data for this project is hosted in the Link Lab Cloud (LLC), a timeseries database built on InfluxDB [20]. The data is collected from various sensors located within the University of Virginia's Rice Hall and Olsson Hall. This project examines six metrics of indoor health: temperature (ºC), carbon dioxide (ppm), VOCs (ppb), humidity (%), average sound pressure level (dB), and illumination (lux).

A Python function called *pull_data* leverages Influx Query Language (InfluxQL) and the pandas library to acquire and format data from the LLC. The function's default parameter values are a list of the six aforementioned metrics, the start and end times for a one-hour window before the function call, and data grouping by both one-minute intervals and sensor ID. For this project, only the default arguments were used. However, the function can alternatively be called with user-specified parameters to accomplish the following:

- Select any set of metrics
- Provide specific start and end date-times
- Change the grouping criteria based on available metadata [20]
- Filter by specific location (e.g., 211 Olsson)

The function's final output is a dictionary of pandas dataframes, with metric types as keys and their respective dataframes as values.

4.2 Data Analysis

The dataframes from the data acquisition process were parsed to exclude invalid (NaN) values. The average of all sensor readings over the one-hour time window and the slope of the line of best fit of those averages were calculated for each metric. A positive, negative, or zero slope indicated an increasing, decreasing, or stable trend, respectively, and was translated into a status label of 1, -1, or 0. This information was aggregated into JavaScript Object Notation (JSON), with metric types as keys, and the one-hour average and status label as values.

4.3 LLM Prompting

The project utilized OpenAI's GPT-4.1 nano LLM through an application programming interface (API). Following the data analysis, the JSON output and the ranges corresponding to low, moderate, or high levels, as shown in Table 1, were passed to the API. The LLM was then prompted to formulate a qualitative indicator of the room's condition in the form of a representative hexadecimal color code (e.g., #76A074) for each metric over the one-hour time window.

As the system evolved, the exact language used in the LLM prompts shifted in response to the color code outputs from GPT-4.1 nano. In their research paper on an AI prompting protocol for human-AI knowledge co-construction, Robertson et al. provided a framework for iterative prompt engineering and a validation process to mitigate bias in AI outputs [15]. To initially avoid confirmation and feedback loop biases, the first version of the prompts was formulated to give the LLM freedom to choose a representative hexadecimal color code without any explicit guidance. These preliminary outputs were used to evaluate the LLM's communicativeness in its color choices. The prompts were then refined to provide the LLM with more direction for assigning colors to low, moderate, and high values for each metric. The final iteration of the prompts provided to the LLM is exemplified below.

Temperature: "Provide a single hexadecimal color code (e.g., #FF5733) representing a room temperature of {temperature}°C. The temperature is considered {temperature status} and described as {temperature condition}. Warmer temperatures (above 25 °C) should correspond to warmer colors (red/orange), and cooler temperatures (below 20 °C) to cooler colors (blue). Return only the hex code."

4.4 Visual Interface

The front-end visual interface was designed using Streamlit, an open-source framework that emphasizes building fast web applications for data [16]. The data is passed to the visual interface as a dictionary, with metric types as keys and lists containing the hexadecimal color code from LLM prompting and the JSON output from data analysis as values.

Three iterations of the visual interface were considered in pursuit of an intuitive user experience, informed by user feedback. Each iteration provided varying levels of control over the size of the metric-indication objects, starting with constant sizes and progressing to LLM-generated weighted sizes.

For all iterations, the view was adapted based on the output from data analysis. The background color for each metric-indication object was set to the hexadecimal color code specified in the LLM prompting. The visual symbol above the metric label was dynamically updated based on the one-hour average and its relationship to the values in Table 1. Below each metric label, a symbol indicates the current trend of the data discussed in section 4.2. An up arrow, down arrow, or neutral line represents an increasing, decreasing, or stable trend, respectively. Each interface was designed to direct the user's attention to the relevant information while being accessible and easy to understand. An algorithm was used to update the metric label text color, ensuring significant contrast between the background and text colors [21]. The algorithm calculates the background color's luminance (brightness) on a scale from 0 (black) to 1 (white). If the background's luminance is less than 0.5, a lighter text color is used; otherwise, a darker color is used. This algorithm provides an appropriate compatible color [1] regardless of the LLM-generated hexadecimal color code.

The first iteration employs constant-sized boxes for each metric. It used custom Hypertext Markup Language (HTML) and Streamlit components to build the visual framework, as shown in Table 5 under Version 1: Constant Size.

The second iteration used a Plotly packed-bubble chart [19] and is shown in Table 5 under Iteration 2: Human-Decided Size. This structure allowed the size of the metric-indication objects to be dynamically updated based on the values in Table 1. The location of the metric-indication object was adjusted based on each metric's data, such that the two metrics with values furthest from their acceptable ranges were placed in the center of the screen.

Use of the Plotly packed-bubble chart continued in the third and final iteration, as shown in Table 5 under Iteration 3: LLM-Generated Size. Rather than providing user-specified acceptable metric ranges, the size of the metric-indication objects was determined through additional LLM prompting. For each metric, GPT-4.1 nano was prompted to provide a weight value to determine the size of the metric-indication object. The prompt for temperature can be seen below:

"Provide a value from 0-1 that represents how a room with a temperature of {temperature}°C would impact a person's health. The value should be closer to 1 if it is dangerous and closer to 0 if it is comfortable/safe. The temperature is considered {temperature status} and described as {temperature condition}."

4.5 Test Design

Multiple runs were conducted on the same dataset to test the system's consistency. To test the system's robustness to data variations, runs over several time windows with varying conditions were conducted.

The dataset used to test the system's consistency when given the same input was collected from the time window of 1 PM to 2 PM on October 28, 2025, when the outdoor average was 12 °C. As the dataset was consistent across all ten runs and iterations, any variations in the metric-indication background color were due to the LLM-generated hexadecimal color code. The average values calculated during data analysis are represented in Table 2.

Table 2. Indoor Conditions for Consistency Test

Temperature (°C)	CO2 (ppm)	VOCs (ppb)	Humidity (%)	Noise (dB)	Illumination (lux)
21	462	123	37	57	1423

Six time windows with varying conditions were identified to test the system's robustness to data variations, as seen in Table 3. These timeframes were selected to center around the consistency test timeframe from Table 2 as a baseline. Under the presumption that variations in outdoor conditions could impact indoor conditions, the justification for each time window's selection, with the outdoor temperature, can be seen in the bulleted list as follows:

1. High temperatures on a weekday (31 °C)
2. Below freezing temperatures on a weekday (-5 °C)
3. Morning time window from October 28, 2025 (8 °C)
4. Late-night time window from October 28, 2025 (10 °C)
5. Rainy day (6 °C)
6. Weekend (14 °C)

The outdoor conditions for the consistency and variation test runs were based on the average values for Charlottesville, Virginia, for the specified one-hour time window [22].

The three iterations of the visual interface were directly compared across four time windows, where time window 0 corresponds to the consistency test in Table 2, and time windows 1, 2, and 3 correspond to the first three rows of Table 3. The most effective iteration was selected as the final view.

Table 3. Indoor Conditions for Varied Time Windows

Time Window	Temperature (°C)	CO2 (ppm)	VOCs (ppb)	Humidity (%)	Noise (dB)	Illumination (lux)
1	21	458	87	62	53	280
2	20	423	131	15	53	319
3	21	481	165	38	54	30
4	21	622	320	37	57	1435
5	20	425	64	30	52	193
6	21	462	123	37	56	1423

Table 4. Indoor Conditions for Iteration Comparison

Time Window	Temperature (°C)	CO2 (ppm)	VOCs (ppb)	Humidity (%)	Noise (dB)	Illumination (lux)
0	21	462	123	37	57	1423
1	21	458	87	62	53	280
2	20	423	131	15	53	319
3	21	481	165	38	54	30

5. RESULTS

The resulting metric-indication background colors for the consistency test are shown in Table 5 below. Each color block matches the LLM-selected color for that metric at each iteration.

Table 5. Comparison of Color Results for Consistency Test

Trial	Temperature (°C)	CO2 (ppm)	VOCs (ppb)	Humidity (%)	Noise (dB)	Illumination (lux)
1						
2						
3						
4						
5						
6						
7						
8						
9						
10						

For the robustness test, the metric-indication background colors are compared in Table 6 below. The first and second runs for each time window are denoted xA and xB, respectively.

Table 6. Comparison of Color Results for Varied Time Windows

Time Window	Temperature (°C)	CO2 (ppm)	VOCs (ppb)	Humidity (%)	Noise (dB)	Illumination (lux)
1A						
1B						
2A						
2B						
3A						
3B						
4A						
4B						
5A						
5B						
6A						
6B						

The comparison between the three versions of the visual interface is shown in Table 7.

6. DISCUSSION

As shown in Table 5, the consistency test results indicate that GPT-4.1 nano was generally capable of generating hexadecimal color codes of similar hue across repeated trials with constant inputs. However, when shades and tints are taken into account, visually distinguishable variability exists across all metrics except VOCs. Furthermore, in the results of the robustness test, there were clear discrepancies in the colors assigned during the repeated trials for some time windows (Table 6). For example, humidity values were held constant between both trials of time window 3, yet the resulting colors for trials A and B were green and orange, respectively. A similar discrepancy was observed in time window 4. Though not as noticeable, the color assignments for noise and illumination for time window 6 also raised concerns regarding the system's consistency.

Robustness testing across varied time windows demonstrated that under differing external conditions, such as date and time, the LLM generates significantly different hexadecimal color codes. Time window 3 included data from the early morning of October 28, 2025. One might infer that classrooms would be unoccupied at this time, meaning lights would be off and lux values would be lower.

Table 7. Comparison of Interface Iterations

Time Window	Iteration 1: Constant Size	Iteration 2: Human-Decided Size	Iteration 3: LLM-Generated Size
0		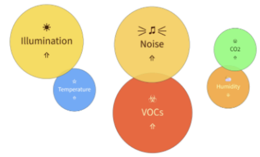	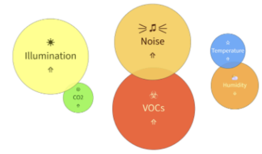
1			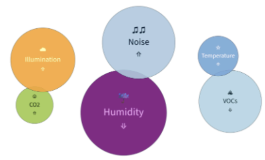
2	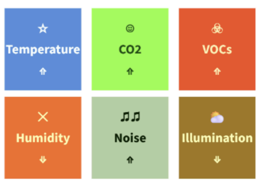		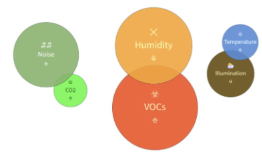
3			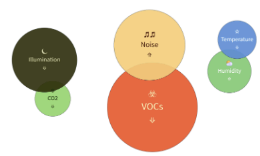

Upon close examination of Table 3, the noise levels for all six time windows ranged from 52 dB to 57 dB. As given in Table 1, noise levels between 45 and 55 dB are considered audible, while noise levels between 56 and 70 dB are considered loud. Looking again at the results from the robustness test, shown in Table 6, it should be noted that the color selections representing the noise level varied from green to blue to orange and red, despite the small differences in the input noise levels. This observed volatility in color selection could be explained by the proximity of input noise levels to the tipping point between the audible and loud categories.

Following the results from Tables 5 and 6, it is not entirely clear what the average metric value is for the color observed without having the accompanying dataset associated with the time window. While the relationship between color and temperature may be easier to intuit, it is more difficult to draw a clear, universally understood relationship between color and unobservable indoor metrics, such as CO2. A similar struggle was encountered with illumination: the original color scale requested in the prompt to the LLM ranged from black to white, but it was difficult for users to distinguish between shades of gray or intuit the color's real-world meaning with respect to indoor health and comfort. To make matters worse, a gray box displayed on a screen might appear differently depending on its surroundings. Including shades of yellow allows the user to more easily assess the room's current lighting levels. Colors alone, as shown in the first column of Table 7, Iteration 1: Constant Size, were not sufficient indicators of indoor health or the condition of the indoor space. The constant layout and size of the metric-indication objects failed to emphasize metrics at suboptimal or dangerous levels and did not convey any intuitive message to the user. Therefore, another element had to be introduced to the system to compensate for the lack of universality in color assignments by the LLM – varying bubble sizes within a packed-bubble chart.

The second iteration's dynamic augmentation of size and location for each metric-indication object effectively provided additional context to complement the background colors of those objects. The last iteration utilized LLM-generated weighted values to augment the size of the metric-indication objects. Interestingly, adding the LLM-generated weighted values did not significantly affect the layout of the metric-indication objects. This iteration was also the most resource-intensive, as additional prompting to OpenAI's GPT-4.1 nano required more API tokens and additional computing power for generating further information. It was determined that the benefit of utilizing LLM-generated weighted values did not outweigh the cost of the resource intensity.

Therefore, the final version of the visual interface shown in Figure 2 is the second iteration.

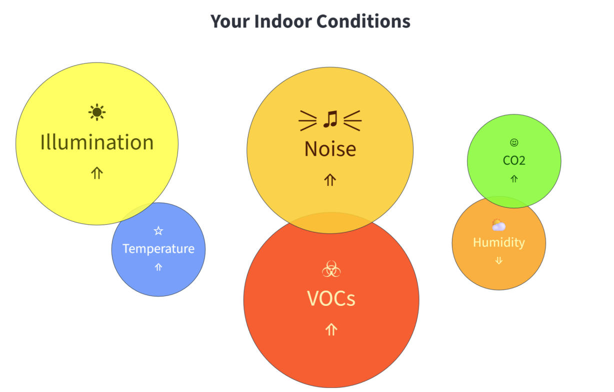

Figure 2. User View of Interface

There are some limitations to the system. The algorithm for creating the packed bubble chart often results in some of the visuals overlapping. The background color determination, as discussed previously in this section, is not entirely accurate and demonstrates variation when values are on the cusp of acceptable or unacceptable values. The data analysis portion, which assigns status labels, does not account for the human body's inability to perceive minuscule changes in the observed metrics over the one-hour time window. Future work could explore providing thresholds for assigning status labels based on human perception, rather than strict cases of greater than, less than, or equal to zero. Also, many sensors were unable to provide readings at the requested 1-minute interval. Therefore, non-numeric values (NaN) were dropped from the average calculation and subsequently the slope of the line of best fit. This lack of data over significant periods could lead to an inaccurate representation of the metric conditions over the one-hour time window.

Future development of this system could enhance its accessibility, interpretability, and consistency. In its current state, the system lacks features to support users with vision impairments or color vision deficiencies. Adding a color-blind setting to the system would begin to address these concerns, though it would require further testing of the AI's color selections.

An important next step in this research would be to conduct structured user evaluations, which may showcase the system's effectiveness or reveal additional shortcomings. Receiving feedback on the system from real participants with varying technical literacy, visual abilities, and age would provide the data to support this system's claim of simplistic interpretability. Such a study could also offer insight into the impacts of our system's lack of consistency; at present, the color determination does not always follow a predictable pattern, especially when a sensor metric's value shifts from an acceptable level to a dangerous one, or vice versa. Drastic changes in color might confuse the user, further reducing the system's interpretability. Additionally, a user study could be used to determine if the interface should be simplified to only show pressing metrics, or if the current display of all information is more helpful to occupants.

Another aspect to address in a user study would be determining the most effective way to present the visual interface. Different strategies could include a website accessed on personal devices or a publicly displayed screen in the indoor space it represents.

6. CONCLUSION

To improve indoor health and awareness, it is important that occupants can easily understand current conditions. This work illustrates the feasibility of using visuals in conjunction with LLMs to translate raw indoor environmental data on temperature, carbon dioxide, volatile organic compounds, humidity, noise, and illumination into interpretable, color-based representations. By leveraging AI to interpret time-series measurements, the system exemplifies a streamlined method for conveying indoor conditions to non-expert observers. The results of this study suggest that LLMs can serve as effective and reliable intermediaries between quantitative measurements and human-centered decision support, though their behavior remains sensitive to prompt design.

7. ACKNOWLEDGMENTS

This project was done as a part of CS 4501: Smart and Healthy Buildings taught at the University of Virginia during the fall 2025 semester. The professors, Arsalan Heydarian and Brad Campbell, provided guidance and feedback.

8. REFERENCES

[1] 101 Computing. 2020. Colour Luminance and Contrast Ratio. Retrieved October 23, 2025 from https://www.101computing.net/colour-luminance-and-contrast-ratio/

[2] Ajinkya Potdar. 2024. Intelligent Data Summarization Techniques for Efficient Big Data Exploration Using AI. *IJAIBDCMS* 5, 1 (March 2024). https://doi.org/10.63282/3050-9416.IJAIBDCMS-V5I1P109

[3] Berkeley Lab. 2025. Introduction to VOCs. *Indoor Air Quality Scientific Findings Resource Bank*. Retrieved from https://iaqscience.lbl.gov/introduction-vocs

[4] PR Boyce, HM Brandston, and C Cuttle. 2022. Indoor lighting standards and their role in lighting practice. *Lighting Research & Technology* 54, 7 (November 2022), 730–744. https://doi.org/10.1177/14771535221126413

[5] Sani Dimitroulopoulou, Marzenna R. Dudzińska, Lars Gunnarsen, Linda Hägerhed, Henna Maula, Raja Singh, Oluyemi Toyinbo, and Ulla Haverinen-Shaughnessy. 2023. Indoor air quality guidelines from across the world: An appraisal considering energy saving, health, productivity, and comfort. *Environment International* 178, (August 2023), 108127. https://doi.org/10.1016/j.envint.2023.108127

[6] Daniel J. Fink. 2017. What Is a Safe Noise Level for the Public? *Am J Public Health* 107, 1 (January 2017), 44–45. https://doi.org/10.2105/AJPH.2016.303527

[7] Steven L. Franconeri, Lace M. Padilla, Priti Shah, Jeffrey M. Zacks, and Jessica Hullman. 2021. The Science of Visual Data Communication: What Works. *Psychol Sci Public Interest* 22, 3 (December 2021), 110–161. https://doi.org/10.1177/15291006211051956

[8] Jochen Hartmann, Yannick Exner, and Samuel Domdey. 2025. The power of generative marketing: Can generative AI create superhuman visual marketing content? *International Journal of Research in Marketing* 42, 1 (March 2025), 13–31. https://doi.org/10.1016/j.ijresmar.2024.09.002

[9] Boulanouar Khedidja, Hadjali Allel, and Lagha Mohand. 2020. Data Summarization for Sensor Data Management: Towards Computational-Intelligence-Based Approaches. *IJCDS* 9, 5 (September 2020), 825–833. https://doi.org/10.12785/ijcds/090505

[10] Dolaana Khovalyg, Ongun B. Kazanci, Hanne Halvorsen, Ida Gundlach, William P. Bahnfleth, Jørn Toftum, and Bjarne W. Olesen. 2020. Critical review of standards for indoor thermal environment and air quality. *Energy and Buildings* 213, (April 2020), 109819. https://doi.org/10.1016/j.enbuild.2020.109819

[11] Hanlu Lyu. 2025. A Study of the Role of Color in Visual Communication in the Digital Perspective. *Mediterranean Archaeology & Archaeometry* 25, 1 (January 2025), 352–357.

[12] Hannah Ramsden Marston and Julie Samuels. 2019. A Review of Age Friendly Virtual Assistive Technologies and their Effect on Daily Living for Carers and Dependent Adults. *Healthcare* 7, 1 (March 2019), 49. https://doi.org/10.3390/healthcare7010049

[13] Susan McKenney and Joke Voogt. 2010. Technology and young children: How 4–7 year olds perceive their own use of computers. *Computers in Human Behavior* 26, 4 (July 2010), 656–664. https://doi.org/10.1016/j.chb.2010.01.002

[14] Subhas Chandra Mukhopadhyay, Sumarga Kumar Sah Tyagi, Nagender Kumar Suryadevara, Vincenzo Piuri, Fabio Scotti, and Sherali Zeadally. 2021. Artificial Intelligence-Based Sensors for Next Generation IoT Applications: A Review. *IEEE Sensors J.* 21, 22 (November 2021), 24920–24932. https://doi.org/10.1109/JSEN.2021.3055618

[15] Jeandri Robertson, Caitlin Ferreira, Elsamari Botha, and Kim Oosthuizen. 2024. Game changers: A generative AI prompt protocol to enhance human-AI knowledge co-construction. *Business Horizons* 67, 5 (September 2024), 499–510. https://doi.org/10.1016/j.bushor.2024.04.008

[16] Streamlit. 2025. Streamlit Docs. Retrieved November 25, 2025 from https://docs.streamlit.io/

[17] Tina M. Sutton and Jeanette Altarriba. 2016. Color associations to emotion and emotion-laden words: A collection of norms for stimulus construction and selection. *Behav Res* 48, 2 (June 2016), 686–728. https://doi.org/10.3758/s13428-015-0598-8

[18] Michaela Tuscher and Johanna Schmidt. 2022. Processing Speed and Comprehensibility of Visualizations and Texts. 2022. . Retrieved October 12, 2025 from https://www.semanticscholar.org/paper/Processing-Speed-and-Comprehensibility-of-and-Texts-Tuscher-Schmidt/d607f32a5e95f810df381a7dedc94b00e6b59c43

[19] U-Danny. 2025. Packed-bubble chart - Plotly Python. *Plotly Community Forum*. Retrieved October 30, 2025 from https://community.plotly.com/t/packed-bubble-chart/92789

[20] University of Virginia. infrastructure.linklab.virginia.edu/linklabcloud/index.html. *Link Lab Cloud*. Retrieved October 28, 2025 from https://infrastructure.linklab.virginia.edu/linklabcloud/index.html

[21] Web Accessibility Initiative. 2025. Understanding WCAG 2.1. *WCAG 2.1 Understanding Docs*. Retrieved October 23, 2025 from https://www.w3.org/WAI/WCAG21/Understanding/

[22] World Weather. 2025. Weather in Charlottesville in October 2025 (Virginia) - detailed Weather Forecast for a month. *World-Weather.info*. Retrieved November 6, 2025 from https://world-weather.info/forecast/usa/charlottesville/october-2025/

[23] Haopeng Zhang, Philip S. Yu, and Jiawei Zhang. 2025. A Systematic Survey of Text Summarization: From Statistical Methods to Large Language Models. *ACM Comput. Surv.* (April 2025), 3731445. https://doi.org/10.1145/3731445

[24] Jing Zhang and Dacheng Tao. 2021. Empowering Things With Intelligence: A Survey of the Progress, Challenges, and Opportunities in Artificial Intelligence of Things. *IEEE Internet Things J.* 8, 10 (May 2021), 7789–7817. https://doi.org/10.1109/JIOT.2020.3039359