An LLM Framework for Modeling U.S. Thermal Comfort Perceptions

Mehdi Ashayeri[1]

School of Architecture, Southern Illinois University Carbondale, Carbondale, IL

Abstract:

This study explores the use of large language models (LLMs) to analyze thermal comfort discourse on social media as a proxy for real-time thermal satisfaction in the United States. Using the OpenAI API, a longitudinal sentiment analysis was conducted on 68,396 U.S.-based tweets posted in 2019 and 2020 to examine changes before and during the COVID-19 disruption. The methodological pipeline included temporal grouping, text preprocessing, semantic interpretation, relevance classification, sentiment labeling, confidence scoring, and final filtering of thermal comfort-related content. Of the full dataset, 55,448 tweets were classified as relevant to thermal comfort, including 33,081 tweets in 2019 and 22,367 in 2020. Results show a marked decline in tweet volume in 2020, consistent with pandemic-related changes in mobility, occupancy, and daily routines. Despite this reduction, sentiment patterns remained stable: neutral tweets constituted the largest share, negative sentiment exceeded positive sentiment, and seasonal and weekday variations persisted. Lexical analysis further showed that thermal comfort discourse was dominated by direct experiential terms such as *cold*, *hot*, *humidity*, and *weather*, with greater visibility of home- and air-related terms in 2020. The findings indicate that LLM-based sentiment analysis can capture large-scale shifts in public thermal perception and provide a novel methodological framework for using digital sentiment as an indicator of social feedback on environmental comfort. This approach offers practical value for human-building interaction research and for the development of more resilient and human-centered building performance standards.

Keywords:

Thermal comfort; social media analytics; large language models; sentiment analysis; human-building interaction

1. Introduction

The evolution of the built environment has transitioned from a focus on static structural efficiency to a dynamic paradigm centered on human-building interaction (HBI). High-performance buildings are increasingly evaluated not merely by system-level mechanical output but by the complex interplay between occupant perception and psychological outcomes [1], [2]. Within this framework, building performance is recognized as an interdisciplinary challenge where environmental design must harmonize with the physiological and cognitive needs of inhabitants [3], [4]. As buildings become more 'intelligent,' their success depends on a deep understanding of occupant-centric attitudes and behavioral patterns [5]. Consequently, modern building science is shifting toward a human-centric model that prioritizes occupant satisfaction as a fundamental metric of operational success. This human-centered transition also resonates with broader performance-driven design scholarship arguing that intelligent environments should move beyond rigid control toward adaptive, multi-scalar frameworks linking buildings, cities, and socially responsive forms of decision support [6], [7], [8].

At the core of these interactions lies thermal comfort, traditionally quantified through steady-state heat balance models such as the Predicted Mean Vote (PMV) and Predicted Percentage of Dissatisfied (PPD) indices [9], [10]. However, the historical reliance on these metrics has revealed significant limitations, particularly their inability to account for the adaptive and subjective nature of human experience in real-world settings [11], [12]. Discrepancies between predicted and observed thermal sensation votes suggest that static standards often overlook individual physiological differences and environmental expectations [13]. This performance gap necessitates a move toward more flexible, data-driven approaches that can capture the longitudinal nuances of occupant comfort beyond the constraints of climate-chamber derived algorithms [14].

To address these limitations, emerging research frames building occupants as sophisticated 'biological sensors' capable of providing high-fidelity feedback on their environment [15], [3]. This communicative dimension of human-building interaction emphasizes that indoor climate control is not just a thermal problem but a social and psychological one [5], [16]. By treating subjective feedback as a primary data source, researchers can better understand the stochastic nature of occupant behavior and its impact on energy consumption [2]. The challenge remains in gathering this intensive longitudinal feedback at scale without causing occupant fatigue or intrusive monitoring [15], [4]. Leveraging the communicative nature of users offers a pathway to closing the gap between actual comfort and automated HVAC operational logic.

Digital platforms and social media have emerged as a significant proxy for public sentiment, providing an untapped repository for thermal preference data [17], [18]. In the United States, microblogging services serve as real-time channels where users express immediate reactions to their environmental conditions [19]. This user-generated content offers a unique opportunity to observe consumer sentiment and comfort preferences 'in the wild,' bypassing the limitations of traditional surveys [18]. By applying sentiment analysis to these digital footprints, researchers can extract collective attitudes toward indoor and outdoor climates across vast geographical scales. Thus, social media data provides a large-scale, cost-effective alternative to traditional longitudinal human-subject studies. Recent work has shown that social-media-derived signals can be translated into building- and city-scale energy intelligence, reinforcing the value of online discourse as a behavioral sensing layer rather than merely a communication channel [20], [21].

The computational capability to interpret this discourse has evolved from simple keyword-matching algorithms to the advanced contextual understanding of Large Language Models (LLMs) [22], [23]. Unlike earlier sentiment analysis tools, LLMs demonstrate an ability to capture context-sensitive perceptions and complex linguistic structures within thermal-related text [24]. These models can perform at levels approaching trained human coders, making them ideal for identifying subtle shifts in sentiment toward indoor environmental quality [22]. Furthermore, fine-tuned transformer architectures are now capable of discerning the sentiment phenomena inherent in diverse user-generated content, even when such content is highly informal or context-dependent [22]. This technological evolution allows for a more granular assessment of how thermal perception is communicated in the digital sphere. Related studies on indoor environmental perception and energy justice further show that transformer-based pipelines can recover occupant attitudes toward indoor air quality and inequitable building conditions from large-scale public text, which strengthens the methodological premise of the present study [25], [26].

Despite these advanced natural language processing capabilities, a distinct disconnect persists between the fields of computational linguistics and building science [22]. Specifically, there is a notable absence of comparative longitudinal studies that leverage LLM-driven sentiment analysis to track thermal perception during periods of global disruption [14]. The transition from 2019 to 2020 represents a critical pivot point where domestic and commercial thermal experiences were drastically altered by societal shifts [11], [1]. Existing literature has yet to utilize LLMs to compare public thermal sentiment across these two years, leaving a gap in our understanding of how large-scale emotional responses to indoor environments change under pressure [22], [15]. Resolving this gap is essential for developing buildings that are truly responsive to shifting human needs.

This study aims to bridge this research gap by utilizing the OpenAI API to conduct a large-scale longitudinal sentiment analysis of thermal comfort discourse in the United States. By comparing data from the pre-disruption context of 2019 with the transitional period of 2020, we seek to uncover how thermal sentiments and preferences shifted as occupants navigated new domestic and commercial environmental realities. The objective is to evaluate the efficacy of LLMs in tracking these nuanced emotional trajectories, ultimately establishing a novel methodological framework for using digital sentiment as a proxy for real-time thermal satisfaction. The findings will provide actionable insights for HBI design, demonstrating how social media can inform more resilient and human-centric building performance standards.

2- Materials and methods

2.1. Data:

This study draws on a corpus of 68,396 U.S.-based tweets related to thermal comfort. Based on the timestamp information, 40,781 tweets were posted in 2019 and 27,614 tweets in 2020, with one record excluded from yearly comparison due to an unreadable date entry. The raw dataset contains tweet text together with metadata fields such as language, posting date and time, source, and public engagement metrics including retweets, replies, likes, and quote counts. After annotation, each tweet was assigned five additional analytical fields: True Meaning, Relevance, Sentiment, Emotion, and Confidence. Of the full corpus, 55,448 tweets were classified as relevant to thermal comfort and 12,948 tweets as irrelevant. Year-specific filtering showed 33,081 relevant tweets in 2019 and 22,367 relevant tweets in 2020. The size and structure of the dataset make it suitable for comparative analysis of thermal comfort discourse before and during the COVID-19 period.

2.2. Methods:

The methodological pipeline consisted of a multi-stage process including data collection and temporal grouping, text pre-processing, semantic interpretation, relevance classification, sentiment and emotion labeling, confidence scoring, and final filtering to ensure analytical focus on thermal comfort–related discourse (Figure 1.) The proposed pipeline uses OpenAI API Key for the sentiment analysis.

At a broader methodological scale, the pipeline used here is likely to be compatible with recent large scale performance analytics frameworks that integrate heterogeneous data streams for city-scale energy inference, machine-learning-enhanced UBEM, and digital-twin-enabled management, indicating that social sensing can function as a meaningful behavioral layer within larger performance-driven modeling ecosystems [27], [28], [29].

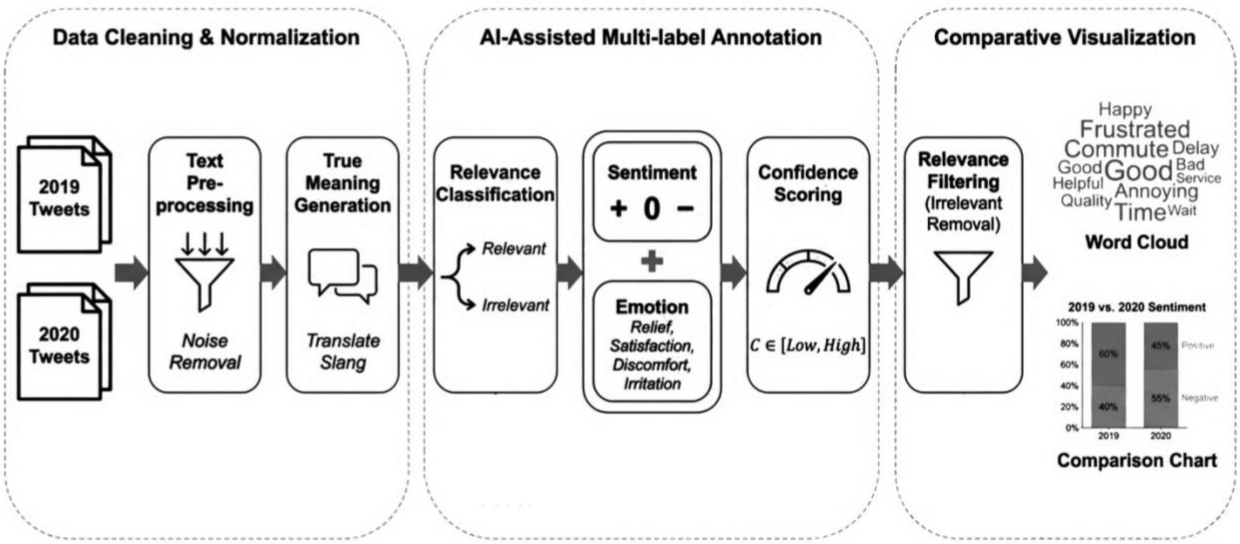

Figure 1. Flowchart of the developed methods for sentiment and emotion analysis using X data and LLMs.

The first step began with raw data collection and year-based organization. The study used a corpus of U.S.-based tweets related to "thermal comfort" and grouped them into two analytical periods, 2019 and 2020, to enable comparison between the pre-COVID and COVID-era contexts. This structure was used to examine whether public expressions of thermal comfort changed under the altered social, indoor, and urban conditions associated with the pandemic. The conceptual framing followed established thermal comfort research in building science, where thermal comfort is defined as a condition of satisfaction with the thermal environment and discomfort is treated as a measurable human response to environmental conditions [10], [12].

The second stage involved text pre-processing. Because X data are short, informal, and noisy, the raw tweets were cleaned before annotation. This included standardizing or removing URLs, user mentions, hashtag formatting, repeated punctuation, and other platform-specific artifacts while preserving the semantic content of each post. The objective was to reduce textual noise without removing the meaning necessary for interpretation. This approach is consistent with sentiment-analysis research showing that pre-processing is especially important in X-based studies because abbreviations, slang, truncated syntax, and irregular expressions can influence classification quality and downstream interpretation [30], [31].

The third stage was semantic interpretation, operationalized through the True Meaning parameter. For each tweet, a short plain-language paraphrase was generated to express the intended meaning of the post more explicitly. This was necessary because tweets often rely on compressed phrasing, slang, irony, or fragmented syntax, which can obscure the underlying message when interpreted literally. The True Meaning field therefore functioned as a semantic normalization step that translated social-media language into analytically clearer text while retaining the original intent. This interpretive normalization is aligned with prior sentiment-analysis literature, which emphasizes that short texts often require contextual reconstruction before robust downstream classification can occur [30], [31]. This semantic-normalization step is also consistent with recent NLP research on indoor-environment perception, where social-media text is translated into interpretable signals before downstream classification of environmental attitudes and health-related concerns [32], [25].

The fourth stage was relevance classification, recorded in the Relevance parameter as either Relevant or Irrelevant. A tweet was classified as relevant only when it referred to human thermal sensation, thermal comfort or discomfort, adaptive responses to feeling hot or cold, or environmental conditions affecting bodily comfort. Tweets were classified as irrelevant when the keyword appeared in unrelated technical, commercial, metaphorical, or product-based contexts that were not connected to human thermal comfort. This filtering stage was essential because the analytical interest of the study was not the word thermal in general, but thermal comfort as a human-centered built-environment concept. The coding logic was therefore grounded in standard definitions from thermal comfort research and building performance practice [10], [12].

The fifth stage involved sentiment classification, stored in the Sentiment parameter as Positive, Neutral, or Negative. Sentiment was assigned according to the overall evaluative orientation of each tweet toward the thermal experience being described. Positive tweets expressed comfort, relief, satisfaction, or pleasantness; negative tweets expressed discomfort, annoyance, overheating, cold stress, or dissatisfaction; and neutral tweets were descriptive, ambiguous, or lacked a clearly evaluative stance. This three-class polarity structure follows widely used sentiment-analysis practice for short social-media texts and is especially suitable for Twitter, where compact expressions often convey clear overall valence even when they are linguistically informal [30], [33], [34].

The sixth stage added an Emotion parameter after sentiment classification in order to capture the closest human affective state associated with the tweet's thermal-comfort meaning. This step moved beyond broad polarity by distinguishing among affective expressions such as relief, satisfaction, discomfort, frustration, irritation, or concern. In methodological terms, sentiment and emotion were treated as related but distinct layers: sentiment captured general evaluative valence, whereas emotion captured the more specific affective quality of the tweet. This distinction is supported by emotion research, including foundational work on basic emotions, as well as computational studies showing that text can be systematically associated with discrete emotion categories rather than only positive or negative polarity [35], [36].

The seventh stage generated a Confidence parameter representing the level of certainty associated with the assigned labels. Confidence was used here as a heuristic annotation-confidence score, based on the clarity of the language, the explicitness of the context, and the degree to which the tweet consistently supported the assigned meaning, relevance, sentiment, and emotion labels. Tweets with direct and unambiguous wording received higher confidence values, whereas tweets containing sarcasm, sparse context, slang-heavy compression, or semantic ambiguity received lower values. This score should not be interpreted as a formally calibrated probability of correctness. That distinction is important because model confidence and actual reliability are not automatically equivalent unless explicit calibration procedures are applied [37].

The eighth stage was relevance filtering. After all five analytical parameters had been generated, tweets classified as irrelevant were excluded from the substantive analysis. This step ensured that the comparative results reflected only discourse genuinely related to thermal comfort rather than unrelated uses of the keyword. The filtered dataset then served as the basis for year-wise interpretation. This relevance-based narrowing is particularly important in social-media research because keyword retrieval alone often captures substantial semantic noise, and meaningful built-environment analysis requires conceptual alignment between the raw text and the target phenomenon [10], [12].

The ninth and final stage was comparative analysis and visualization. Once irrelevant tweets had been removed, the remaining posts were compared between 2019 and 2020 to identify changes in sentiment, emotional tone, and dominant textual patterns related to thermal comfort. This final step was intended to reveal whether pandemic-related disruptions in occupancy, daily routines, and urban life corresponded to a shift in how thermal comfort was discussed in public discourse. To support interpretation, year-specific textual summaries, including word-cloud visualizations, were generated from the filtered relevant corpus. In this way, the final analytical stage connected AI-assisted text annotation to broader questions in building performance and urban sustainability while maintaining a human-centered understanding of thermal comfort as an experiential and socially mediated condition [30], [33], [12].

Methodologically, this final interpretive step may support emerging hybrid and health-oriented modeling frameworks that connect environmental conditions, exposure, and building performance, suggesting that social-media-derived comfort signals can eventually be combined with physics-based and exposure-aware urban analytics rather than analyzed in isolation [38], [39], [40].

3. Results and Discussion

3.1. Tweet Frequency as a Proxy for Social Behavior

Figure 2 shows that the nationwide volume of U.S. tweets about thermal comfort declined from 2019 to 2020 in both the full corpus and the subset classified as relevant. Total tweet counts decreased from 40,781 in 2019 to 27,614 in 2020, while relevant sentiment counts declined from 37,263 to 25,034. This reduction coincided with the major social and behavioral disruptions associated with the first year of the COVID-19 pandemic. The World Health Organization characterized COVID-19 as a pandemic on 11 March 2020, and in the United States mandatory stay-at-home orders were issued by 42 states and territories between 1 March and 31 May 2020, affecting 73% of U.S. counties. During the same period, mobility patterns changed substantially, with residents in states under stay-at-home orders more likely to remain close to home and to reduce daily travel. In this context, the lower tweet volume in 2020 is consistent with reduced exposure to the public, institutional, and occupational environments in which thermal comfort is commonly experienced and discussed, including offices, classrooms, retail spaces, and transit settings [41], [42], [43].

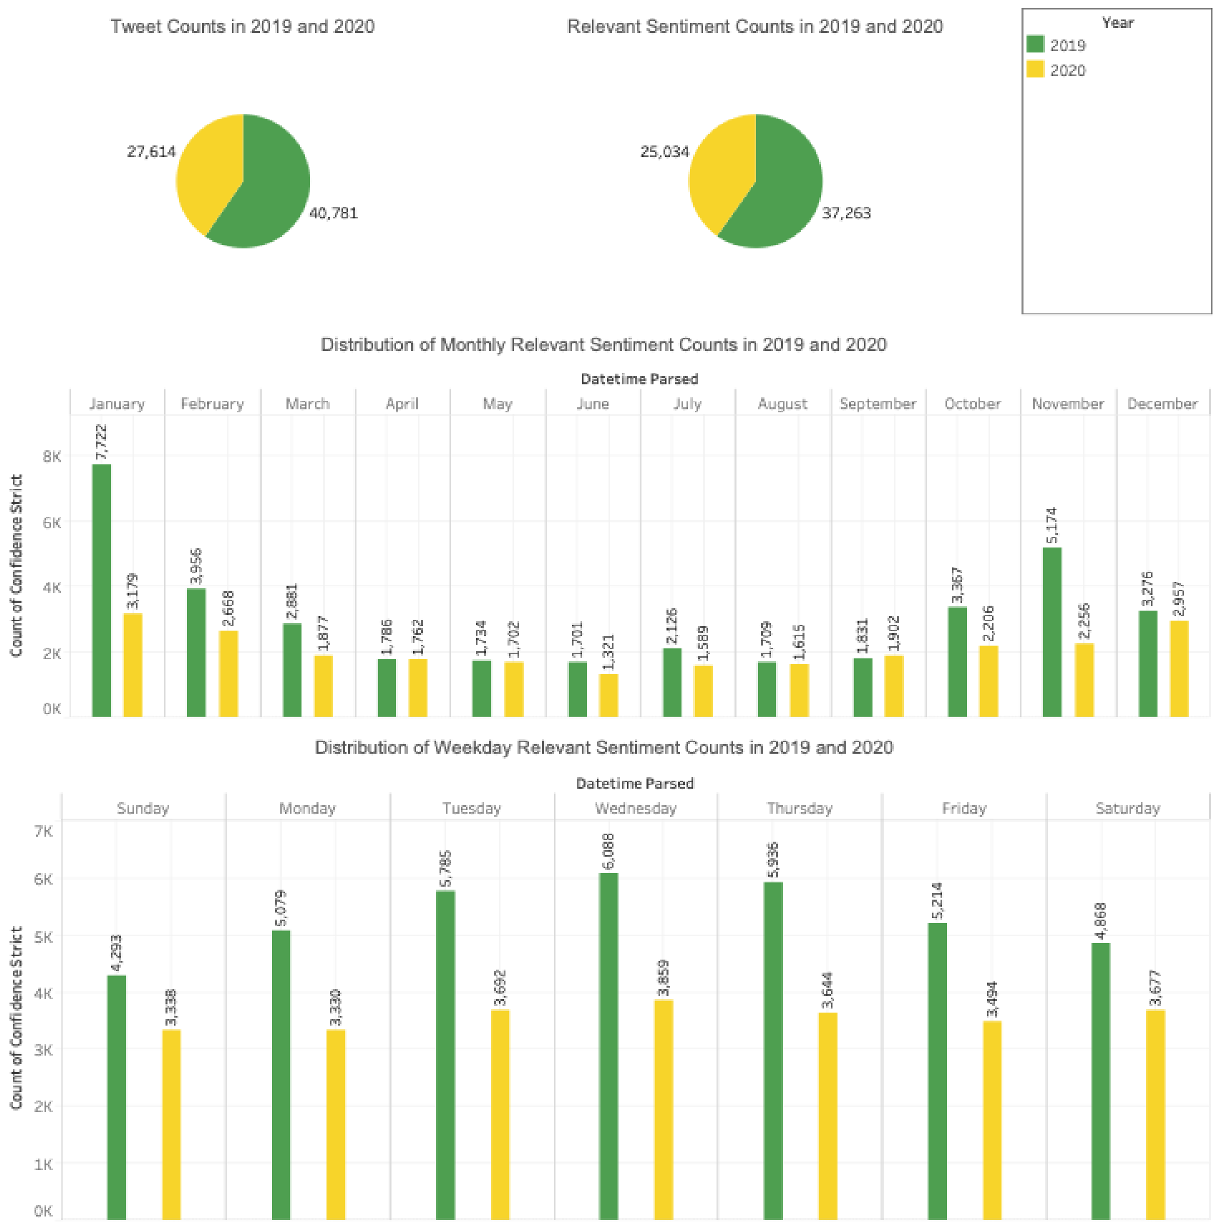

Figure 2. Annual all (top-left) and relevant (top-right), monthly (middle), and weekday (bottom) relevant tweet counts during 2019 and 2020

The monthly distribution supports this interpretation. In 2019, relevant thermal-comfort tweet counts were highest in January and rose again in October and November, whereas the 2020 pattern was lower across most months and less sharply peaked. This attenuation is consistent with the relocation of daily activity away from shared buildings and commuting spaces and toward the home. U.S. labor statistics further indicate that 35% of employed people teleworked because of the pandemic in May 2020, and 24% continued to do so in August 2020, pointing to a sustained change in where everyday life was being conducted. The weekday distribution shows a similar decline, with 2020 counts remaining below 2019 across all days of the week, although both years retain a stronger midweek pattern than weekend pattern. Taken together, Figure 2 suggests that the reduction in tweet frequency in 2020 reflected not a disappearance of concern with thermal comfort, but a contraction in the number of routine social settings in which thermal conditions were publicly noticed, compared, and discussed. The decline may also reflect intensified attention to COVID-19 on social media, which likely displaced discussion of more routine environmental experiences within nationwide online discourse [44], [45].

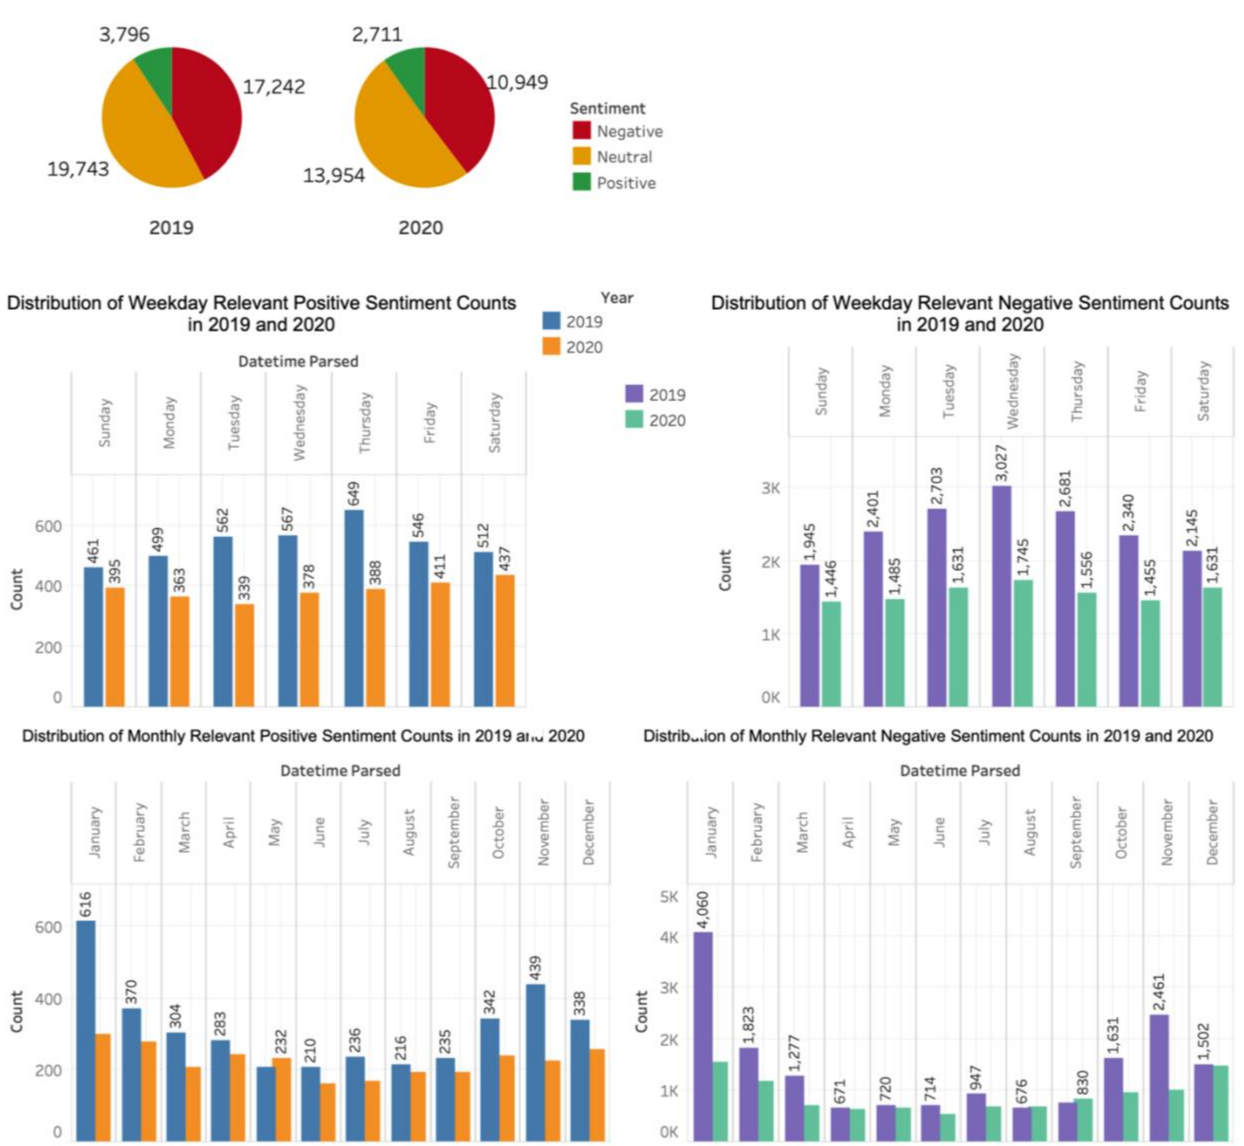

Figure 3. Distribution of sentiment in U.S. thermal-comfort tweets on X in 2019 and 2020, showing overall sentiment shares and the weekday and monthly patterns of positive and negative sentiment. Neutral sentiment was the largest category in both years, negative sentiment consistently exceeded positive sentiment, and total sentiment counts declined in 2020.

Viewed through a broader urban-performance lens, this shift toward home-centered discourse is consistent with research showing that socioeconomic conditions and microclimate variability shape both residential energy behavior and the realism of urban-scale energy modeling, especially when everyday activity patterns move across uneven environmental contexts [46], [47], [48].

Figure 3 extends the frequency analysis by showing how users evaluated thermal comfort on X during 2019 and 2020. Across both years, sentiment remained dominated by neutral and negative expression, while positive sentiment represented the smallest share. The pie charts indicate that neutral tweets were the largest category in both years, decreasing from 19,743 in 2019 to 13,954 in 2020, followed by negative tweets, which declined from 17,242 to 10,949. Positive tweets were consistently less frequent, falling from 3,796 to 2,711. This distribution indicates that thermal comfort was discussed more often as an ordinary environmental experience or a source of dissatisfaction than as a topic eliciting explicitly favorable reactions. The decline across all three sentiment categories in 2020 is consistent with the broader disruptions of daily routines during the COVID-19 period, when reduced mobility and widespread stay-at-home measures limited exposure to the shared indoor environments in which thermal discomfort is often noticed and publicly expressed [41], [42].

The weekday and monthly distributions clarify the structure of this sentiment pattern. Positive sentiment remained comparatively low and stable across the week in both years, whereas negative sentiment was substantially higher on every day, especially during the midweek period. In 2019, negative sentiment peaked on Wednesday (3,027) and remained elevated on Tuesday (2,703) and Thursday (2,681); in 2020, Wednesday remained the highest day (1,745), although all weekday counts were lower than in the previous year. The monthly distribution shows a similar contrast: negative sentiment was especially pronounced in January 2019 (4,060) and rose again in October and November (1,631 and 2,461), while the 2020 pattern remained lower and flatter. Positive sentiment showed a smaller seasonal range, with its most visible peak also occurring in January 2019 (616) and relatively modest levels thereafter. These patterns indicate that social feedback on thermal comfort was driven more by discomfort than by satisfaction, with users more likely to post when indoor conditions were perceived as problematic, disruptive, or seasonally stressful.

When read together, Figures 2 and 3 present a consistent interpretation. Figure 2 shows that the overall frequency of thermal-comfort discussion fell in 2020, while Figure 3 shows that the underlying sentiment structure remained broadly stable: neutral sentiment remained the largest category, and negative sentiment continued to exceed positive sentiment throughout the week and across much of the year. The lower 2020 counts therefore reflect a reduction in the scale of social feedback rather than a reversal in its character. Pandemic-related telework and reduced circulation through offices, schools, transit systems, and other shared settings likely narrowed the situations in which thermal comfort was publicly evaluated, while the intense focus on COVID-19 on social media further reduced the prominence of this topic in online discourse. In this sense, X functioned as a proxy for social feedback on thermal comfort, capturing both the decline in public expression during 2020 and the continued tendency for thermal discomfort to generate more explicit reaction than thermal satisfaction [44], [45].

Figure 4 provides a lexical view of how thermal comfort was discussed on X in the United States. In both years, the most frequent terms are anchored in direct bodily and environmental experience, with **"cold"** clearly dominating the corpus and with **"hot," "humidity," "outside,"** and **"weather"** also appearing prominently. This pattern is consistent with Figures 2 and 3, which showed that thermal-comfort discourse was driven by routine exposure to ambient conditions and by evaluative responses to discomfort rather than by abstract or technical discussion. The prominence of simple experiential terms suggests that users primarily framed thermal comfort through immediate sensations and everyday outdoor or indoor encounters, rather than through specialized building-science vocabulary. At the same time, the continued visibility of terms such as **"home," "air,"** and **"temperature"** indicates that users were also connecting comfort to indoor settings and environmental control.

The comparison between 2019 and 2020 suggests continuity in the basic vocabulary of discomfort, but with a modest shift in emphasis toward indoor and home-centered language in 2020. In both years, **"cold"** remained by far the most frequent term, indicating that cold-related discomfort was the most salient lexical feature of the dataset. However, the 2020 word cloud gives relatively greater visibility to terms such as **"home," "air," "temp/temperature,"** and **"humidity,"** which is consistent with the pandemic-era relocation of daily routines into domestic space. In August 2020, **24.3%** of employed people in the United States reported teleworking because of the pandemic, down from **35.4%** in May but still representing a substantial shift in where work and daily occupancy occurred. Research on home workplaces during COVID-19 likewise found that indoor environmental quality conditions in the home, including comfort-related factors, affected well-being in ways similar to office environments. Read alongside those findings, the lexical profile in Figure 4 suggests that thermal-comfort discourse in 2020 remained strongly sensation-based but became somewhat more tied to residential indoor experience than in 2019.

The increased visibility of terms related to air, humidity, and indoor conditions in 2020 is also compatible with the broader public-health context of the pandemic. The World Health Organization characterized COVID-19 as a pandemic on 11 March 2020, after which indoor air and building operation became more prominent public concerns. CDC guidance emphasized that ventilation strategies can reduce viral particle concentration indoors, and ASHRAE stated that ventilation and filtration through HVAC systems can reduce airborne concentrations of SARS-CoV-2 while also noting that unconditioned spaces can create thermal stress. In this context, Figure 4 suggests that pandemic-era thermal-comfort discourse did not simply decline in volume, as shown in Figure 2; it also acquired a somewhat more indoor-environmental vocabulary, without losing its core focus on everyday discomfort. The persistence of words such as "cold," "hot," and "humidity" therefore supports the interpretation from Figure 3 that social feedback on thermal comfort remained rooted in embodied experience, while the stronger visibility of "home" and "air" in 2020 reflects the altered environmental settings in which that experience was being lived and described. This stronger home- and air-centered vocabulary in 2020 is also consistent with pandemic-era studies linking indoor air quality, health resilience, and socially mediated perceptions of indoor environments, suggesting that thermal discourse became more tightly coupled to broader concerns about domestic environmental quality [49], [25].

3.3. Lexical Patterns in Thermal-Comfort Tweets

Figure 4 shows that the vocabulary of U.S. thermal-comfort tweets in both years was dominated by direct experiential terms, with **"cold"** clearly the most frequent word, followed by terms such as **"hot," "humidity," "outside,"** and **"weather."** This lexical profile indicates that users mainly described thermal comfort through immediate bodily sensation and everyday environmental exposure rather than through technical building-performance language. Terms such as **"home," "air,"** and **"temperature"** also appear prominently, suggesting that users connected comfort not only to outdoor weather but also to indoor conditions and environmental control within occupied spaces.

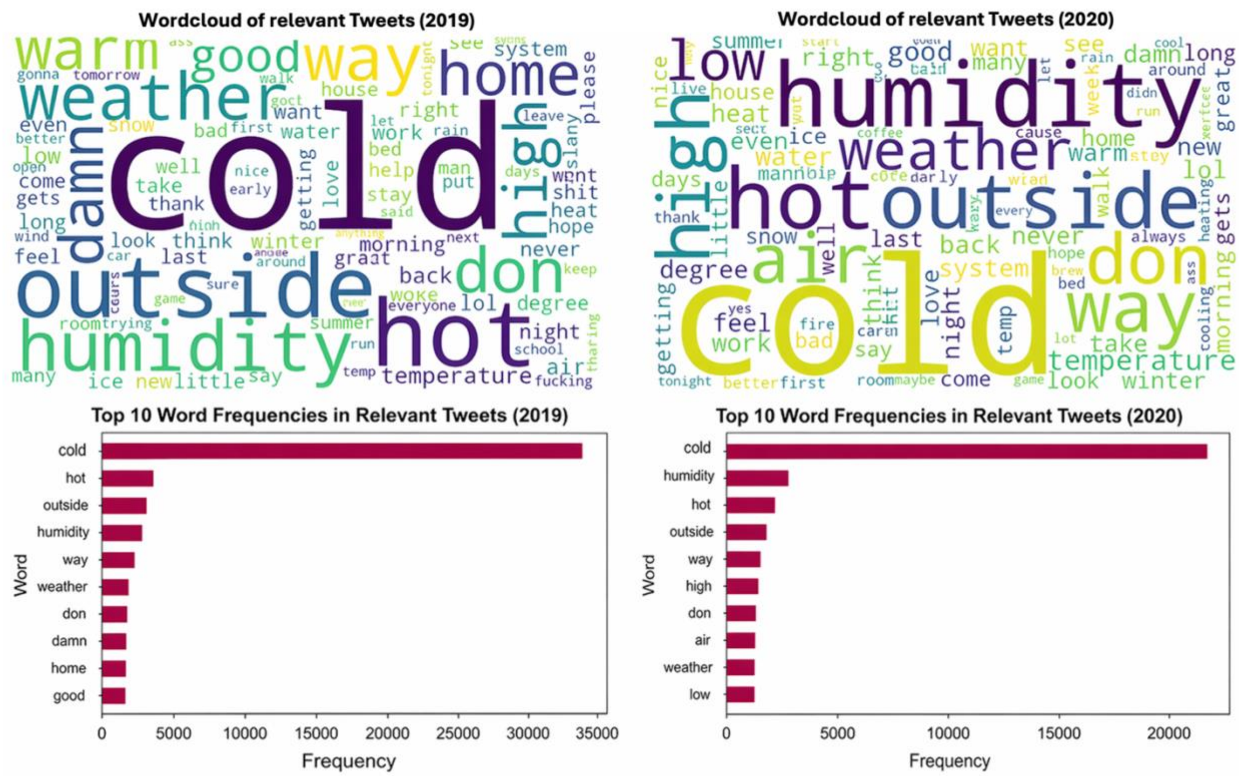

Figure 4. Word clouds and top-ten term frequencies for relevant U.S. thermal-comfort tweets in 2019 and 2020.

The comparison between 2019 and 2020 suggests continuity in the core language of discomfort, but also a greater emphasis in 2020 on indoor and home-centered experience. In both years, "cold" remained the most salient term, indicating that cold-related discomfort was a dominant feature of the dataset. At the same time, the 2020 panel gives stronger visibility to words such as "home," "air," "humidity," and "temperature," which is consistent with pandemic-related shifts in everyday occupancy. In the United States, 24% of employed people teleworked in August 2020 because of the pandemic, down from 35% in May 2020, indicating that a substantial share of daily activity had moved from workplaces into domestic settings [44]. Research on home workplaces during the COVID-19 period also found that indoor environmental quality in the home workplace was related to mental well-being, supporting the interpretation that residential indoor conditions became more salient during this period [50].

The 2020 lexical pattern is also compatible with the broader public-health context in which indoor air and ventilation became more visible in public discussion. The World Health Organization stated on 11 March 2020 that COVID-19 could be characterized as a pandemic [41]. During the pandemic, public-health guidance increasingly emphasized ventilation and indoor air management as ways to reduce exposure risk in buildings and homes, and CDC guidance specifically highlighted ventilation improvements as part of strategies to reduce the spread of airborne viruses indoors [51]. Read in that context, the greater prominence of terms such as "air" and "home" in 2020 suggests that thermal-comfort discourse did not merely decline in volume; it also reflected the changing environmental settings in which people were living and working. Figure 4 therefore complements Figures 2 and 3 by showing that, even as overall tweet counts declined, the language of thermal comfort remained rooted in embodied experience while becoming somewhat more attentive to residential indoor environments during the pandemic period.

This study should be interpreted in light of several limitations. First, X data do not represent the full U.S. population, as platform users are self-selected and may differ systematically by age, region, socioeconomic status, and posting behavior. Second, thermal-comfort tweets capture expressed perceptions rather than measured environmental conditions, which means that the analysis reflects subjective experience and willingness to post rather than direct exposure alone. Finally, the lexical and temporal patterns identified here are descriptive and associative; they reveal how thermal comfort was discussed on X, but they do not by themselves establish causal relationships between environmental conditions, social context, and online expression.

Future work should move beyond sentiment counting toward integrated social-sensing frameworks that connect online thermal discourse with urban building energy models, indoor environmental quality indicators, and adaptive control logic. Recent studies point to a promising pathway in which social-media signals inform urban-scale energy forecasting, occupant-centered ambient intelligence, and broader performance-driven design workflows [20], [21], [52]. Such integration could help translate expressed comfort needs into more responsive building operations, more equitable environmental management, and more resilient city-scale design strategies. A further future direction is to connect social-sensing outputs with design-stage optimization tools so that recurring public discomfort signals can inform context-sensitive passive and solar-responsive interventions, not only operational control [53].

4. Conclusion

This study demonstrates that X can serve as a useful proxy for examining social behavior and social feedback related to thermal comfort in the United States. Across 2019 and 2020, the results show that thermal-comfort discourse was shaped by clear temporal and evaluative patterns. Tweet frequency declined substantially in 2020, suggesting that pandemic-related disruptions reduced the volume of public expression around routine thermal experiences. At the same time, the structure of sentiment remained relatively stable: neutral tweets constituted the largest share, negative sentiment consistently exceeded positive sentiment, and lexical patterns remained centered on direct experiential terms such as cold, hot, humidity, and weather. Together, these findings indicate that thermal comfort on social media was discussed primarily as an everyday lived experience, and more often as a source of inconvenience or dissatisfaction than of positive appraisal.

The study also shows that social-media discourse on thermal comfort is responsive to broader changes in how and where daily life is organized. The reduced tweet counts in 2020, combined with the stronger visibility of home- and air-related vocabulary, suggest that pandemic conditions altered not only the scale of discussion but also the settings through which thermal comfort was experienced and described. In this sense, thermal-comfort tweets provide a complementary perspective to conventional environmental and building studies by revealing how people publicly articulate comfort, discomfort, and seasonal exposure in real time. Future work could strengthen this approach by integrating social-media data with meteorological observations, building characteristics, geographic variation, and more advanced language models in order to better understand how environmental conditions are translated into public perception and online expression. More broadly, the study supports a next generation of occupant-centered analytics in which thermal sentiment is interpreted alongside other digitally expressed indicators of indoor environmental quality, health, and energy justice to guide performance-driven design and operation [26], [52].

More broadly, these findings support a human-centered analytics agenda in which thermal discourse is interpreted together with digitally expressed indicators of indoor health and energy justice, helping performance-driven design frameworks remain socially legible as well as technically robust [32], [26].

Acknowledgment

This study utilized the Academic Twitter API (2021–2022), which provided free access for data querying from the platform. We extend our special thanks to Soroush Piri, a former graduate assistant at SIUC, for their support in data collection and querying.

References

[1] T. Hong, S. D’Oca, W. J. N., Turner, and S. C Taylor-Lange, "An ontology to represent energy-related occupant behavior in buildings", Part I: Introduction to the DNAs framework. Building and Environment, 92, 764–777, 2015. https://doi.org/10.1016/j.buildenv.2015.02.019.

[2] T. Hong, D. Yan, S., D’Oca, and C.-F Chen, "Ten questions concerning occupant behavior in buildings: The big picture", Building and Environment, 114, 518–530, 2017. https://doi.org/10.1016/j.buildenv.2016.12.006.

[3] B. Becerik-Gerber, G. Lucas, A. Aryal, M. Awada, M. Bergés, S. L. Billington, O. Boric-Lubecke, A. Ghahramani, A. Heydarian, C. Höelscher, F. Jazizadeh, A. Khan, J. Langevin, R. Liu, F. Marks, M. L. Mauriello, E. Murnane, H. Noh, M. Pritoni, Zhu et al., and R, "The field of human building interaction for convergent research and innovation for intelligent

built environments", Scientific Reports, 12, Article 22092, 2022. https://doi.org/10.1038/s41598-022-25047-y.

[4] W. O'Brien, A. Fuller, M., Schweiker, and J. K Day, "A picture is worth a thousand words: Smartphone photograph-based surveys for collecting data on office occupant adaptive opportunities", In Proceedings of the 7th International Building Physics Conference (pp. 601–606), 2018. https://doi.org/10.14305/ibpc.2018.hf-1.01.

[5] N Ma, "Human dimensions of building performance: Sensing, modeling, and predicting indoor environmental quality (Doctoral dissertation)", University of Pennsylvania, 2022. https://repository.upenn.edu/edissertations/4715.

[6] N. Abbasabadi and M. Ashayeri, Eds., Artificial Intelligence in Performance-Driven Design: Theories, Methods, and Tools. New York, NY, USA: Wiley, 2024, ISBN: 978-1-394-17206-1.

[7] N. Abbasabadi and M. Ashayeri, "Urban energy use modeling methods and tools: A review and an outlook," Building and Environment, vol. 161, Art. no. 106270, 2019, doi: 10.1016/j.buildenv.2019.106270.

[8] N. Abbasabadi and M. Ashayeri, "Towards an Adaptive Urbanism Beyond Hard Control: The Theories of Johnson and Lefebvre," in ARCC Architectural Research Addressing Societal Challenges, vol. 1. London, UK: CRC Press/Taylor & Francis Group, 2017, pp. 257-262.

[9] P. O Fanger, "Thermal comfort: Analysis and applications in environmental engineering", Danish Technical Press, 1970.

[10] ASHRAE, "ANSI/ASHRAE Standard 55-2020: Thermal environmental conditions for human occupancy", American Society of Heating, Refrigerating and Air-Conditioning Engineers, 2020.

[11] T. Cheung, S. Schiavon, T. Parkinson, P., Li, and G Brager, "Analysis of the accuracy on PMV–PPD model using the ASHRAE Global Thermal Comfort Database II", Building and Environment, 153, 205–217, 2019. https://doi.org/10.1016/j.buildenv.2019.01.055.

[12] International Organization for Standardization, "ISO 7730: Ergonomics of the thermal environment—Analytical determination and interpretation of thermal comfort", 2005. https://www.iso.org/standard/39155.html.

[13] V. A. Arowoiya, A. O. Onososen, R. C., Moehler, and Y Fang, "Influence of thermal comfort on energy consumption for building occupants: The current state of the art", Buildings, 14(5), Article 1310, 2024. https://doi.org/10.3390/buildings14051310.

[14] M. Quintana, M. Abdelrahman, M. Frei, F., Tartarini, and C Miller, "Longitudinal personal thermal comfort preference data in the wild", In Proceedings of the 4th ACM Workshop on Data: Acquisition to Analysis (pp. 556–559), 2021. https://doi.org/10.1145/3485730.3493693.

[15] P. Jayathissa, M. Quintana, M., Abdelrahman, and C Miller, "Humans-as-a-sensor for buildings—Intensive longitudinal indoor comfort models", Buildings, 10(10), Article 174, 2020. https://doi.org/10.3390/buildings10100174.

[16] T. Hong, S. C. Taylor-Lange, S. D'Oca, D., Yan, and S. P Corgnati, "Advances in research and applications of energy-related occupant behavior in buildings", Energy and Buildings, 116, 694–702, 2016. https://doi.org/10.1016/j.enbuild.2015.11.052.

[17] B Liu, "Sentiment analysis: Mining opinions, sentiments, and emotions (2nd ed.)", Cambridge University Press, 2020. https://doi.org/10.1017/9781108639286.

[18] M. A. Montaser, B. P. Ghosh, A. Barua, F. Karim, B. C. Das, R. E. R., Shawon, and M. S. R Chowdhury, "Sentiment analysis of social media data: Business insights and consumer behavior trends in the USA", Edelweiss Applied Science and Technology, 9(1), 515–535, 2025. https://doi.org/10.55214/25768484.v9i1.4164.

[19] N. C. Dang, M. N., Moreno-García, and F De la Prieta, "Sentiment analysis based on deep learning: A comparative study", Electronics, 9(3), Article 483, 2020. https://doi.org/10.3390/electronics9030483.

[20] N. Abbasabadi and M. Ashayeri, "From Tweets to Energy Trends (TwEn): An exploratory framework for machine learning-based forecasting of urban-scale energy behavior leveraging social media data," Energy and Buildings, vol. 317, Art. no. 114440, 2024, doi: 10.1016/j.enbuild.2024.114440.

[21] N. Abbasabadi and M. Ashayeri, "From Tweets to Energy Trends (TwEn2): Social sensing–informed urban building energy modeling," Frontiers in Energy Research, 2025, doi: 10.3389/fenrg.2025.1688348.

[22] W. Zhang, Y. Deng, B. Liu, S., Pan, and L Bing, "Sentiment analysis in the era of large language models: A reality check", In Findings of the Association for Computational Linguistics: NAACL 2024 (pp. 3881–3906), 2024. https://doi.org/10.18653/v1/2024.findings-naacl.246.

[23] J Elsey, "Validity of large language models for sentiment analysis: Evidence of performance comparable to human coders", PsyArXiv, 2023. https://doi.org/10.31234/osf.io/kcuwy.

[24] Y. Wang, A. Eshghi, Y., Ding, and R. D Gopal, "Echoes of authenticity: Reclaiming human sentiment in the large language model era", PNAS Nexus, 4(2), pgaf034, 2025. https://doi.org/10.1093/pnasnexus/pgaf034.

[25] M. Ashayeri, S. Piri, and N. Abbasabadi, "Exploring U.S. occupant perception toward indoor air quality via social media and NLP analysis," Journal of Environmental Science and Public Health, vol. 8, pp. 49–58, 2024, doi: 10.26502/jesph.96120205.

[26] M. Ashayeri and N. Abbasabadi, "Unraveling energy justice in NYC urban buildings through social media sentiment analysis and transformer deep learning," Energy and Buildings, vol. 306, Art. no. 113914, 2024, doi: 10.1016/j.enbuild.2024.113914.

[27] N. Abbasabadi, M. Ashayeri, R. Azari, B. Stephens, and M. Heidarinejad, "An integrated data-driven framework for urban energy use modeling (UEUM)," Applied Energy, vol. 253, Art. no. 113550, 2019, doi: 10.1016/j.apenergy.2019.113550.

[28] N. Abbasabadi and M. Ashayeri, "Machine Learning in Urban Building Energy Modeling," in Artificial Intelligence in Performance-Driven Design: Theories, Methods, and Tools, 1st ed. New York, NY, USA: Wiley, 2024.

[29] N. Abbasabadi and M. Ashayeri, "Digital Twin for Citywide Energy Modeling and Management," in Artificial Intelligence in Performance-Driven Design: Theories, Methods, and Tools, 1st ed. New York, NY, USA: Wiley, 2024.

[30] R Feldman, "Techniques and applications for sentiment analysis", Communications of the ACM, 56(4), 82–89, 2013. https://doi.org/10.1145/2436256.2436274.

[31] M. A., Palomino, and F Aider, "Evaluating the effectiveness of text pre-processing in sentiment analysis", Applied Sciences, 12(17), Article 8765, 2022. https://doi.org/10.3390/app12178765.

[32] M. Ashayeri, "Decoding Global Indoor Health Perception on Social Media via NLP and Transformer Deep Learning," in Artificial Intelligence in Performance-Driven Design: Theories, Methods, and Tools, 1st ed. New York, NY, USA: Wiley, 2024.

[33] C. J., Hutto, and E Gilbert, "VADER: A parsimonious rule-based model for sentiment analysis of social media text", Proceedings of the International AAAI Conference on Web and Social Media, 8(1), 216–225, 2014. https://doi.org/10.1609/icwsm.v8i1.14550.

[34] M. Thelwall, K., Buckley, and G Paltoglou, "Sentiment strength detection for the social web", Journal of the American Society for Information Science and Technology, 63(1), 163–173, 2012. https://doi.org/10.1002/asi.21662.

[35] P Ekman, "An argument for basic emotions", Cognition and Emotion, 6(3–4), 169–200, 1992.

[36] S. M., Mohammad, and P. D Turney, "Emotions evoked by common words and phrases: Using Mechanical Turk to create an emotion lexicon", In Proceedings of the NAACL HLT Workshop on Computational Approaches to Analysis and Generation of Emotion in Text (pp. 26–34). Association for Computational Linguistics, 2010.

[37] C. Guo, G. Pleiss, Y., Sun, and K. Q Weinberger, "On calibration of modern neural networks", In Proceedings of the 34th International Conference on Machine Learning (pp. 1321–1330). PMLR, 2017.

[38] M. Ashayeri and N. Abbasabadi, "A Hybrid Physics-based Machine Learning Approach for Integrated Energy and Exposure Modeling," in Artificial Intelligence in Performance-Driven Design: Theories, Methods, and Tools, 1st ed. New York, NY, USA: Wiley, 2024.

[39] M. Ashayeri and N. Abbasabadi, "A framework for integrated energy and exposure to ambient pollution assessment (iEnEx) toward low-carbon, healthy, and equitable cities,"

Sustainable Cities and Society, vol. 78, Art. no. 103647, 2022, doi: 10.1016/j.scs.2021.103647.

[40] M. Ashayeri, N. Abbasabadi, M. Heidarinejad, and B. Stephens, "Predicting intraurban PM2.5 concentrations using enhanced machine learning approaches and incorporating human activity patterns," Environmental Research, vol. 196, Art. no. 110423, 2021, doi: 10.1016/j.envres.2020.110423.

[41] World Health Organization, "WHO Director-General's opening remarks at the media briefing on COVID-19 – 11 March 2020", 2020.

[42] A. Moreland, C. Herlihy, M. A. Tynan, and et al, "Timing of state and territorial COVID-19 stay-at-home orders and changes in population movement—United States, March 1–May 31, 2020", MMWR Morbidity and Mortality Weekly Report, 2020.

[43] K. J Bourassa, "State-level stay-at-home orders and objectively measured movement in the United States during the COVID-19 pandemic", Psychosomatic Medicine, 83(4), 358–362, 2021.

[44] Bureau of Labor Statistics, "One-quarter of the employed teleworked in August 2020 because of COVID-19 pandemic", U.S. Bureau of Labor Statistics, 2020.

[45] J. Xue, J. Chen, R. Hu, C. Chen, C. Zheng, Y., Su, and T Zhu, "Twitter discussions and emotions about the COVID-19 pandemic: Machine learning approach", Journal of Medical Internet Research, 22(11), e20550, 2020.

[46] N. Abbasabadi and M. Ashayeri, "Socioeconomic Determinants of Public Health and Residential Building Energy use in Chicago," in 2021 International Conference: 27th World Congress of Architects (UIA 2021 RIO), ACSA, vol. II, 2021, ISBN: 978-1-944214-31-9.

[47] A. Worthy, M. Ashayeri, and N. Abbasabadi, "Leveraging Earth observational data and machine learning to enhance urban building energy modeling with microclimate effects," Sustainable Cities and Society, 2025, doi: 10.1016/j.scs.2025.106544.

[48] A. Worthy, M. Ashayeri, M. Julian, and N. Abbasabadi, "Bridging the simulation-to-reality gap: A comprehensive review of microclimate integration in urban building energy modeling (UBEM)," Energy and Buildings, 2025, doi: 10.1016/j.enbuild.2025.115392.

[49] M. Ashayeri and N. Abbasabadi, "Energy justice, indoor air quality, and community health resilience against COVID-19 pandemic," in AMPS Architecture, Media, Politics, Society Conference Proceedings Series, vol. 26, pp. 1–10, 2021.

[50] B. Boegheim, R. Appel-Meulenbroek, D., Yang, and M Loomans, "Indoor environmental quality (IEQ) in the home workplace in relation to mental well-being", 2022.

[51] Centers for Disease Control and Prevention, "How to improve ventilation in your home", n.d.

[52] N. Abbasabadi and M. Ashayeri, "Occupant-Driven Urban Building Energy Efficiency via Ambient Intelligence," in Artificial Intelligence in Performance-Driven Design: Theories, Methods, and Tools, 1st ed. New York, NY, USA: Wiley, 2024.

[53] S. Luitjohan, M. Ashayeri, and N. Abbasabadi, "An optimization framework and tool for context-sensitive solar-driven design using cellular automata (SDCA)," in IEEE Xplore, 2022, doi: 10.23919/ANNSIM55834.2022.9859496.