

DOI: https://doi.org/10.48009/2_iis_2024_126

Practical wisdom and intelligent machines: toward AI-human socio-technical decision making and resource allocation

Hesham Allam, *Ball State University, hesham.allam@bsu.edu*

Colin Allen, *UCSB, colinallen@ucsb.edu*

Aaron Allred, *Ball State University, aallred@bsu.edu*

Matthew Boersma, *Ball State University, mboersma@bsu.edu*

Christopher B. Davison, *Ball State University, cbdavison@bsu.edu*

Edward J. Lazaros, *Ball State University, ejlazaros@bsu.edu*

Allen D. Truell, *Ball State University, atruell@bsu.edu*

Abstract

This project, The Machine Wisdom Project, aims to develop a framework for defining and enhancing practical wisdom in human-machine interactions. The notion of practical wisdom highlights both (1) the current limitations in Artificial Intelligence and Machine Learning Systems (AI/ML) and (2) the dangers entailed when humans must make real-time adjustments to AI/ML systems with only a partial understanding of how they work. The Machine Wisdom Project envisages strategies for mitigating both (1) and (2) that recognize the lack of self-awareness in current and foreseeable AI/ML technologies and the need to promote designs that make those limitations apparent. Feedback loops between the actions of people and artificially intelligent machines constitute socio-technical systems with the potential to alter (positively and negatively) the capacity of individuals to act ethically. The project researchers seek to deliver principles and recommendations for the design of socio-technical systems that minimize the likelihood that people adapt to the rigid incomprehension of machines in ways that restrict their ethical autonomy and responsibility, that minimize the potential for bad actors to exploit these limitations, and that improve the capacity of individuals to act ethically within the socio-technical system. This research article presents the Ethical, Multimodal, Moral Agent (EMMA) system, an HVAC socio-technical decision-making test bed for the Machine Wisdom Project.

Keywords: AI, machine learning, HVAC systems, XAI

Introduction

In recent years, the widespread integration of artificial intelligence into our daily lives has transformed human communication and interaction, extending beyond interpersonal relationships to include human-machine and machine-machine interactions. Intriguingly, the shift towards dependence on AI often goes unnoticed, making this an essential target for research. In the literature review of this research article, the authors will cover the existing research on the intricate relationship between artificial intelligence and HVAC systems, accounting for various perspectives, findings, and debates surrounding this topic. Subsequently, the authors will present the Ethical, Multimodal, Moral Agent (EMMA) system and discuss the human-AI, socio-technical decision-making process facilitated by EMMA along with the Explainable Artificial Intelligence (XAI) research aspects of EMMA.

One of the primary objectives of this article is to describe the unique aspects of EMMA that set it apart from other research focusing on explainable artificial intelligence. Unlike other studies that attempt to try to control human behavior when using AI (an example being an AI that gives you tailored ads based on your search history to influence your purchases), this study focuses on creating feedback loops between users of HVAC systems and the EMMA system to improve resource usage.

The study focuses on gathering user requests for HVAC systems and measuring how feedback provided by EMMA impacts the users' preference requests. This research hypothesizes that delivering Explainable Artificial Intelligence (XAI) feedback by EMMA within HVAC systems will significantly influence user preference requests (e.g., HVAC temperature and humidity adjustment requests). Specifically, it is expected that clear and understandable explanations offered by EMMA regarding energy usage and system adjustments will positively impact users' willingness to adjust their preferences in line with energy-saving or community-oriented approaches. Furthermore, this feedback will mitigate self-serving bias, encouraging users to consider collective benefits over personal comfort. The research question studied for this article is: Do select feedback interventions correlate to user behavior changes?

The study also analyzes the motivational implications arising from the presentation of information, examining the contrast between community-oriented and self-centered approaches. The Framing Bias Hypothesis asserts that how information about energy usage is framed significantly influences individuals' willingness to sacrifice personal comfort for the greater good. Additionally, the Self-serving Bias Hypothesis suggests that individuals tend to justify their temperature adjustment requests (situational attribution) while attributing laziness or lack of consideration for financial costs to others' requests (dispositional attribution) (Newell et al., 2022). Furthermore, The Individualized Experiences Hypothesis posits that personalized memories shape individual responses, making information resonate when presented in alignment with these unique experiences (Newell et al., 2022).

Literature Review

This review aims to provide a comprehensive understanding of the current limitations inherent in Artificial Intelligence and Machine Learning Systems (AI/ML), particularly concerning HVAC systems, shedding light on potential risks associated with real-time adjustments made by humans to AI/ML systems, even when individuals possess only a partial understanding of the underlying mechanisms. A comprehensive search across multiple databases, including IEEE Xplore, ScienceDirect, ACM Digital Library, Google Scholar, and PubMed Central, was conducted for scholarly articles, conference papers, and relevant industry reports published between 2010 and 2023. Keyword combinations such as "ai", "artificial intelligence", "machine learning", "HVAC", "smart HVAC", "heating, ventilation, and air conditioning", "autonomy", "IoT", and "internet of things" were utilized. Although there is extensive scientific exploration in the area of HVAC, including sensorization and AI models to predict and control energy consumption, there remains a gap in research focusing on AI-augmented decision-making.

HVAC and Artificial Intelligence

With some form of AI prevalent in most IoT (Internet of Things) devices, HVAC systems have followed suit, albeit in a limited capacity, focusing on energy savings and cost reduction. Modern HVAC systems utilize sensors that detect the status of the environment and measure temperature, humidity, pressure, and wind, among other things. This data is then processed and stored by an AI model, which aids in determining the best course of action for building energy management (Ahmad et al., 2016).

One of the main goals of any HVAC system is to provide occupants with comfortable, healthy, and productive indoor environments. Most households in the US use a traditional (non-AI) thermostat to control heating and cooling using positive or negative feedback to achieve a comfortable environment for the user. People in North America spend approximately 90% of their time indoors (Deng & Chen, 2020); hence, the need to provide occupants with a comfortable, healthy indoor environment remains significant. However, most occupants employ a less convenient approach to thermal comfort, which involves clothing adjustment when feeling cold or warm, drinking cold or hot liquids, and moving to other locations inside a room.

A study by (Karjalainen, 2009) concluded that thermal comfort levels were much lower in offices than in homes during summer and winter as most people in offices resort to a more adaptive approach to comfort since their environments are too hot or cold. This is because, in offices where many people work together, problems arise when people do not care about or are unaware of others' HVAC preferences and are unwilling to compromise. This can worsen the HVAC comfort levels inside the office, sometimes making it more stressful or uncomfortable. As a result, a lot of money is spent on thermal comfort. According to the US Department of Energy, buildings consume about 41% of the total energy in the United States (Dong et al., 2019), and even at that level of consumption, businesses rarely achieve the comfort level desired.

XAI / FAI

The increasing reliance on intelligent machines has created a necessity for transparent and interpretable models, particularly in data-sensitive domains where understanding and trust are paramount (Dwivedi et al., 2023). This newfound incentive has rekindled a significant interest in developing Explainable Artificial Intelligence (XAI), which provides techniques to enhance human comprehension, confidence, and clarity in AI models (Dwivedi et al., 2023). XAI is receiving intense study in healthcare diagnostics, statistical machine learning, and other highly complex systems (Dwivedi et al., 2023). The challenge of implementing XAI is due to the complex and ambiguous nature of societal problems (Bruijn et al., 2022). These problems are often called "wicked" problems because they lack precise definitions and consensus on solutions, leading to public distrust in AI systems, mainly when used for sensitive tasks concerning people's health and well-being (Bruijn et al., 2022).

Advances in statistical machine learning have focused on the need for XAI to balance predictive power with interpretability, particularly amidst growing complexity (Hassanpour et al., 2022). The literature emphasizes the need for structured explanations optimized for human decision-makers and mentions the importance of extending XAI to incorporate usability and measure the quality of explanations. Discussing how to leverage human expertise to enhance AI performance is also mentioned, as well as advocating for human-in-the-loop interactive machine learning (Hassanpour et al., 2022).

Despite the long history of explainability in AI, recent advancements in neural network architectures have introduced challenges in understanding and explaining AI decisions, particularly in highly complex systems (Hassanpour et al., 2022). While these systems exhibit remarkable performance, their lack of interpretability hinders their instructional value and ability to aid human understanding (Hassanpour et al., 2022).

An example of why more explainable AI is necessary can be seen in the recent controversy surrounding Google's Gemini AI tool. Google has aimed to address the ongoing challenges and complexities in achieving diversity and accuracy in AI-generated images (Jacobi & Sag, 2024). However, recently, the software has been accused of generating racist images and pushing diversity too far by generating politically incorrect images like black Nazis and black Vikings, further emphasizing the importance of transparency and interpretability in AI systems (Kayser-Bril, 2020). To make amends, Google acknowledges that its AI model contains inaccuracies and expresses their commitment to improvement. They discuss the need for a

more nuanced and balanced representation of diversity in their AI software (Kayser-Bril, 2020). The controversy is an excellent example of what can occur when there is a lack of explainability in an AI model.

In the context of healthcare diagnostics, some experts believe that it is more important for patients to be able to question or challenge the AI's diagnosis rather than get a detailed explanation of how the AI makes its decisions (Ploug & Søren, 2020). The concept of explanation is thus replaced with the simpler notion of 'effective contestability.' Instead of focusing on providing every detail of how the AI works, the main goal is theorized to be making sure patients have the right tools and information to challenge the diagnosis. To ensure effective contestability, Ploug and Søren (2020) identified thirteen informational requirements needed by patients if they are to be able to contest AI diagnoses in a way that is meaningful and fair. These requirements include details about how the AI system utilizes data, any biases it may have, its overall performance, and the roles of both the system and healthcare professionals in the diagnostic process.

Smart HVAC

HVAC control systems are already being transformed into smart systems. As in other application areas, it will be important to address issues stemming from discrepancies among user preferences and conflicts with system-level goals. To facilitate user understanding of AI decision-making and improve ethical outcomes AI, it will be necessary to explain to users the consequences of the interactions with the system.- Numerous earlier research projects have used different interior sensors to detect occupant occupancy. Ultrasonic sensors, CO₂ detectors, wearable devices, smartphones, radio frequency identification devices (RFIDs), passive infrared (PIR) sensors, and wireless sensor networks have been used to monitor thermal comfort or occupancy behavior.

Progress in the application of sensors is exemplified in research by Cheng & Lee (2016) where thermo-fluidic sensors and occupancy detectors were implemented to achieve smart operations of air conditioning systems. Other studies have focused more on using computational intelligence to save energy in homes (Ahmad et al., 2016). Significant advances have been made to reduce energy consumption through the application of computational intelligence to control, manage, and optimize the usage of HVAC systems. The overarching goal across the reviewed studies is to establish intelligent controllers that may significantly reduce energy consumption in buildings while increasing thermal comfort for occupants, achieving improved performance in the two critical objectives of HVAC systems. Previous studies have not systematically investigated how people interact with indoor environments and provide feedback depending on their comfort. This is a significant topic for the development of ethical, explainable interfaces to smart HVAC systems.

Limitations

While integrating artificial intelligence and smart HVAC systems holds promise for enhancing energy efficiency and thermal comfort, several limiting factors must be addressed. For example, geographical location, specialized AI modeling, generalizability, limitations in human data, and other replication challenges must be addressed before smart HVAC systems can become widely viable as energy-efficient and reliable models. Artificial intelligence has progressed in terms of functionality and precision over recent decades. However, today, most artificial intelligence systems are not designed for what smart HVAC would need to function optimally. While AI models like ChatGPT are adept at natural language processing and generating human-like responses, they may need to be more suitable for the tasks required in smart HVAC systems. ChatGPT's linguistic capabilities might not align with the complex decision-making processes and data analysis required for optimizing HVAC operations. Also, AI models heavily rely on data for training and decision-making. More accurate human input data is needed to ensure the effectiveness and fairness of

AI-driven HVAC control systems. Furthermore, user preferences and comfort levels may evolve, necessitating continuous data collection and updates to AI algorithms. It would be difficult to use current AI models to achieve optimal HVAC functionality, and creating a specialized AI model for smart HVAC systems to use would take a lot of work.

When conceptualizing the idea of global integration of smart HVAC systems, a challenge that requires attention is the significant variation of environmental conditions and regulations across the globe. Implementing AI-based solutions developed for one geographical area may not be directly applicable or effective in another. Thus, localization and customization of AI algorithms would be necessary to ensure optimal performance in diverse environments. A specific example of a problem could be if two different regions have distinct sets of regulations governing HVAC systems, including requirements for energy efficiency, ventilation rates, and indoor air quality.

The generalizability of research done at a single Midwestern university could create problems for the integration of smart HVAC systems. Research conducted at a single university may only partially capture other institutions or businesses' diverse challenges and requirements. Also, relying solely on data from one university may lead to biased conclusions and inadequate solutions. Therefore, there is a need for broader datasets and input from multiple sources to account for variability in building structures, occupancy patterns, and user preferences. Moreover, while the ideal scenario involves replicating successful AI implementations from one university to another, this process can be challenging due to differences in infrastructure, organizational culture, and management preferences. Achieving widespread adoption and scalability of AI-augmented HVAC systems requires addressing contextual variations and ensuring compatibility in a variety of settings.

AI within the context of HVAC is novel in that the Machine Wisdom Project scientists are investigating machine ethics in a real-time dynamical context (unlike judicial sentencing or loan approvals) yet within a bounded, less-than-lethal, controllable environment (unlike autonomous vehicles or the trolley problem). Moreover, this project offers a novel approach to FAI: fairness monitoring by a machine surrogate. This approach is designed to fill the gap between what is understandable to humans and the amount of data required for the complete forensic reconstruction of system decisions. While building HVAC is limiting, it does provide a bounded, controlled testbed that provides quantifiable and observable results.

Preliminary data collection for this study began in April of 2022 and culminated with over 2,000 user comfort preferences and over 20 million HVAC tuples. Tracking user intervention outcomes did not begin until March of 2023. Since then, over 1,100 complete outcome interactions have been tracked in the data. Fulfillment or requested HVAC changes gathered from the EMMA system began nearly a year later, in February of 2024; at the time of writing, this part of the dataset consists of just a few hundred tuples.

Methodology

This study adopts a comprehensive methodology that focuses on implementing and evaluating the Ethical, Multimodal, Moral Agent (EMMA) system within the Machine Wisdom Project. The EMMA system facilitates ethical decision-making in human-machine interactions, particularly in heating, ventilation, and air conditioning (HVAC) adjustments. Currently, the system is an expert system that extracts information from several knowledge bases (multi-modal sources of information) such as sensors and scheduling databases. This data is processed through a custom rules engine as user requests are received and user intervention feedback is provided to the user to effectuate more ethical HVAC consumption.

Technical Infrastructure

The EMMA system is built upon a robust client-server architecture, with the EMMA client application serving as the primary data collection tool (see Figure 1 and Figure 2). Developed using the C# programming language and Xamarin framework, the app is accessible via the Apple App Store and Google Play, ensuring widespread availability and usability. The EMMA client app collects data relevant to HVAC usage, comfort requests, and location information within buildings. This includes parameters such as temperature preferences, comfort levels, occupancy patterns, environmental conditions, and decisions made from intervention feedback.



Figure 1: Emma handheld app with cost feedback




Figure 2: Emma handheld app with exercise feedback.

Intervention feedback provided to users encompasses various dimensions, including financial implications, caloric equivalents (e.g., exercise biking or eating hamburgers), dollar costs incurred through HVAC adjustments, and considerations regarding the health and comfort of individuals in the vicinity. Moreover, the server-side AI application tracks instances of spamming, evaluates the strain imposed on HVAC systems, and collects demographic information to discern user preferences and behaviors comprehensively.

Data Transmission and Analysis

Comfort request data is transmitted securely to the EMMA server, which employs a MySQL database for storage and processing (see Figure 1). Here, the server executes sophisticated analyses to generate feedback and inform decision-making regarding HVAC adjustments. The analysis encompasses various facets, including economic considerations, health-related implications, and environmental impacts. Additionally, the strategy involves examining feedback patterns across demographic groups to identify potential biases and disparities. Statistical techniques and machine learning algorithms are leveraged to extract actionable

insights from the data. The outcome of all human-machine interactions is stored in the MySQL database. For this research article, the data was downloaded and analyzed using SPSS. An analysis of the efficacy of the feedback interventions is provided regarding the research question presented herein.




Figure 2: EMMA Server

Real-world Deployment

The EMMA AI system is not merely a theoretical construct but a fully deployed system within the Applied Technology (AT) building at Ball State University (BSU) campus. This unique arrangement involves close collaboration between EMMA system designers, campus engineers, and third-party engineers to ensure the seamless integration and operation of the system. Special attention is paid to safeguarding the campus geothermal system and other control systems against potential harm. This real-world deployment provides valuable insights into the practical implementation and efficacy of the EMMA system within the context of AI-human socio-technical decision-making interactions.

Results

Subjects were recruited via flyers placed in the AT building at the BSU campus. The flyers contained information on the AI system, IRB, the Machine Wisdom Project, and the HVAC system, as well as instructions on obtaining the EMMA system. Table 1 shows the breakdown of gender of the study participants. The Percent column shows frequencies as a percentage of all data points. The Valid Percent column shows frequencies as a percentage of all data points, including those with errors (missing fields). In this case, there were no errors, so the columns show the same percentages.

Table 1: EMMA user gender

user_Gender				
	Frequency	Percent	Valid Percent	Cumulative Percent
Female	4	6.9	6.9	6.9
Male	48	82.8	82.8	89.7
Other	1	1.7	1.7	91.4
Prefer not to say	5	8.6	8.6	100.0
Total	58	100.0	100.0	

Table 2 shows the breakdown of Emma users' community. Most EMMA users are students, followed by faculty members with offices in the AT building.

Table 2: EMMA users' reasons for being in the building.

user_Reason				
	Frequency	Percent	Valid Percent	Cumulative Percent
Faculty, no office in this building	2	3.4	3.4	3.4
Faculty, office in this building	17	29.3	29.3	32.8
Other	7	12.1	12.1	44.8
Student	32	55.2	55.2	100.0
Total	58	100.0	100.0	

Table 3 shows the frequency analysis of the EMMA user community. Within the EMMA user community, 17 people have medical conditions affected by temperature and humidity levels (indicated by a 0). Most users do not have any medication conditions affected by building comfort requests (indicated by a 1). A smaller percentage of users chose "prefer not to say" concerning any medical conditions (indicated by a 2).

Table 3: EMMA users' medical conditions

user_Condition				
	Frequency	Percent	Valid Percent	Cumulative Percent
0	17	29.3	29.3	29.3
1	32	55.2	55.2	84.5
2	9	15.5	15.5	100.0
Total	58	100.0	100.0	

As of the writing of this paper, the EMMA system has over 2,000 comfort requests from users. As indicated in Table 4 and Figure 4, most of those requests are extreme heating requests (I Feel: Extremely Cold) as a request of less than .25 on the app's slider bar. A .25 request is Somewhat Cold. A .5 request is Fine. A .75 request is Somewhat Hot, and a 1.0 request is Extremely Hot. The statistical results indicate a positive skewness, and while the Kurtosis is low (slightly positive at .066), there is a noticeable peak in the Extremely Cold (left tail) user request data.

Table 4: User comfort requests descriptive statistics

Statistics		
too_hot		
N	Valid	2003
	Missing	0
Mean		.23309
Median		.02800
Mode		.028
Skewness		1.141
Std. Error of Skewness		.055
Kurtosis		.066
Std. Error of Kurtosis		.109
Minimum		.000
Maximum		1.000




Figure 4: User comfort requests histogram.

When using the EMMA app, the user is provided random feedback interventions about their comfort request. This feedback includes monetary cost, impact on others, calories and health information, energy equivalencies, and other types of feedback. The user can cancel the request, force the request to actuate or request more feedback information. Table 5 shows the outcome decisions with 0 referring to the user canceling their request, 1 means, go ahead anyway, and 2 equals get more information. Most (close to 60%) users opted to force the comfort request.

Table 5: User outcome decisions

decision				
	Frequency	Percent	Valid Percent	Cumulative Percent
0	170	15.1	15.1	15.1
1	664	59.0	59.0	74.1
2	291	25.9	25.9	100.0
Total	1125	100.0	100.0	

After the user decides (shown in Table 5) the EMMA AI makes a fulfillment decision. Table 6 shows the frequency of EMMA request outcomes where 0 is not fulfilled, 1 is fulfilled, 2 is partially fulfilled, and spaces (missing) are not processed. The request could be unreasonable, too hard on the HVAC system, disruptive to sensitive medical conditions, or aggregated with other requests in the same area. Most user comfort requests (over 80%) still needed to be fulfilled. In table 6 In this table, the difference between the Percent and Valid Percent values is due to the 47 data points with missing fields.

Table 6: EMMA AI Fulfillment Frequencies

fulfill					
		Frequency	Percent	Valid Percent	Cumulative Percent
Valid	0	945	84.0	87.7	87.7
	1	133	11.8	12.3	100.0
	Total	1078	95.8	100.0	
Missing	System	47	4.2		
Total		1125	100.0		

Table 7 shows the fulfillment data since February 8, 2024. Like table 6 0 is not fulfilled, 1 is fulfilled, 2 is partially fulfilled, and spaces (missing) are not processed. These user comfort requests can and were

fulfilled on the BSU HVAC system. Before that date, the EMMA system could not effectuate changes to the production HVAC system, but close to half of the user comfort requests were fulfilled.

Table 7. Requests since February 2024.

		fulfill			
		Frequency	Percent	Valid Percent	Cumulative Percent
Valid	0	103	36.4	43.6	43.6
	1	133	47.0	56.4	100.0
	Total	236	83.4	100.0	
Missing	System	47	16.6		
Total		283	100.0		

Discussion

User Gender and Reason for Being in Area

The majority (90%) of the EMMA system users are males. This is consistent with the Computer Technology program's predominantly male student population. Relatedly, the faculty with offices in the building make up the second largest user population, around 30% of the EMMA users. Again, the faculty in the program are predominantly male.

These data align closely with the known demographics of the program and building. This tends to validate the demographic data within the EMMA system.

User Medical Conditions

Around 30% of EMMA users had self-disclosed medical conditions affected by indoor comfort conditions (e.g., temperature and humidity). Half of the users reported no medical conditions that would be impacted by the HVAC. Finally, 15% of the users chose not to disclose any medical conditions to the EMMA system.

User Comfort Requests

Most of the user comfort requests were for heating. As indicated in Figure 4, there is a noticeable spike in feeling cold or extremely cold. As indicated in the Limitations section, the data collection began in March of 2023, just before the end of the Spring semester. As many fewer students attend during the Summer, the data makes sense from a seasonal perspective. Additionally, the fulfillment and HVAC actuation did not begin until February 2024. Given the average temperatures in the midwest at that time, this would mostly capture heating requests as exhibited by the positive skewness of the data.

Actuation of User Comfort Requests

User actuation of a comfort request (forcing it through the system) is one of the exciting results of this study. In that area, it is interesting that most users (60%) will choose to force their comfort request (Table 5) no matter what intervention feedback is provided.

The intervention feedback the system provides is randomized as to whether medical conditions, conflicting requests, and spamming are prioritized. Table 8 show the frequency of six different types of feedback which are as follows:

- 0 = Others with medical conditions are present.
- 1 = Other users with conflicting requests/preferences are present.
- 2 = Dollar cost in Energy was provided.
- 3 = Non-branded hamburger calories offered.
- 4 = Calories, speed, and effort on an exercise bike.
- 5 = (Spamming) multiple requests from the same user in ≤ 10 minutes.

Table 8: User intervention feedback

	group_id			
	Frequency	Percent	Valid Percent	Cumulative Percent
0	108	9.6	9.6	9.6
1	170	15.1	15.1	24.7
2	229	20.4	20.4	45.1
3	238	21.2	21.2	66.2
4	241	21.4	21.4	87.6
5	139	12.4	12.4	100.0
Total	1125	100.0	100.0	

The distribution of the six types of feedback was from 10% of the time to a high of 21%. As expected, feedback 0 (users with medical conditions nearby) is the lowest provided feedback as the EMMA AI will only provide this feedback if such users are nearby. Spamming the system, operational, defined as making more than one request in 10 minutes, was the second lowest with only 139 instances (12%). There are several instances in the data where one user would send multiple (up to 8) requests back-to-back quickly. Social science research is needed to investigate the possible reasons for such behavior.

Table 9 shows all the canceled user comfort requests. There is no significant predictor of canceling a request (Table 10) based on the five types of feedback. However, the calorie content of a user's comfort request is the least likely to result in a cancellation. The calorie content is calculated by EMMA using the BTUs necessary to accomplish the user requests and convert them into non-branded hamburger caloric content. The percentage of a hamburger to eat or the number of hamburgers required to be eaten to accomplish the comfort request is provided to the user.

Table 9: Feedback causes users to cancel (Decision 0) comfort requests.

	decision			
	Frequency	Percent	Valid Percent	Cumulative Percent
0	170	100.0	100.0	100.0

	group_id			
	Frequency	Percent	Valid Percent	Cumulative Percent
0	41	24.1	24.1	24.1
1	42	24.7	24.7	48.8
2	27	15.9	15.9	64.7
3	19	11.2	11.2	75.9
4	30	17.6	17.6	93.5
5	11	6.5	6.5	100.0
Total	170	100.0	100.0	

Table 10. Canceled requests correlated to intervention feedback.

Correlations				
Spearman's rho	decision	Correlation Coefficient	decision	group_id
			.	.
			Sig. (2-tailed)	.
		N	170	170
	group_id	Correlation Coefficient	.	1.000
			Sig. (2-tailed)	.
		N	170	170

Table 11 shows all of the forced user comfort requests are shown. These are the comfort requests that the users decided to *send another preference request* after being given intervention feedback. The feedback 0 (users with medical conditions nearby) resulted in the lowest rate of users forcing their comfort requests and this is closely followed by feedback 1 (other users with conflicting comfort requests). -There is no significant predictor of forcing a request (Table 12). The hamburger calorie content was the most likely to result in users forcing their request through the system but only by a very small margin, as feedback 3, 4, and 5 are similar results.

Table 10: Feedback causes users to force (Decision 1) comfort requests.

decision				
	Frequency	Percent	Valid Percent	Cumulative Percent
1	664	100.0	100.0	100.0

group_id				
	Frequency	Percent	Valid Percent	Cumulative Percent
0	43	6.5	6.5	6.5
1	70	10.5	10.5	17.0
2	147	22.1	22.1	39.2
3	152	22.9	22.9	62.0
4	150	22.6	22.6	84.6
5	102	15.4	15.4	100.0
Total	664	100.0	100.0	

Table 12: Forced requests correlated to intervention feedback.

Correlations				
Spearman's rho	decision	Correlation Coefficient	decision	group_id
			.	.
			Sig. (2-tailed)	.
		N	664	664
	group_id	Correlation Coefficient	.	1.000
			Sig. (2-tailed)	.
		N	664	664

Conclusion

In the paper, the authors studied user HVAC comfort requests within one building of a medium-sized midwestern university and presented their findings. This study's context and enabling framework is the Ethical, Multimodal Moral Agent (EMMA) AI system. EMMA is a socio-technical decision-support AI designed to work specifically within HVAC. Users of the EMMA system can download the app (Google Play or Apple App Store), register with the system, and control HVAC settings within the building mentioned above.

The EMMA users were provided intervention feedback through dollar costs, medical condition proximity, spamming (e.g., multiple requests for the same user), caloric equivalencies, multiple (often conflicting) user requests, and exercise energy expenditure. The users were allowed to 1) Cancel their request, 2) Force their request, or 3) Request further information. As indicated in the Discussion section, most users (approximately 60%) decided to force their request through. Only about 15% of the user community opted to cancel their HVAC comfort request based on the intervention feedback provided by the EMMA AI.

Concerning the research question, no significant predictor or correlation existed between any feedback group and the accompanying user behavior. However, some intervention feedback did appear more efficacious in moderating user behavior, such as informing the user that other users with medical conditions impacted by the HVAC environment were in the immediate vicinity, resulting in less forcing of user preferences through the system as well as the highest rate of user HVAC preference cancellations. Relatedly, providing caloric equivalency in hamburger terms appears the least likely to result in comfort request cancellations and most likely to result in a user forcing their comfort request through the system.

IRB Statement

The BSU IRB determined this research to be exempt as only involving survey research and observation of public behavior with no identifiable information being retained about the subjects. The IRB assigned the following protocol number: 1838299-2.

Acknowledgement/Funding Statement

The authors of this paper would like to thank the Templeton World Charity Foundation for their generosity. This project on Practical Wisdom and Intelligent Machines ("The Machine Wisdom Project") was funded by grant number 0467 (<https://doi.org/10.54224/20467>) from the Templeton World Charity Foundation Diverse Intelligences program to Professor Colin Allen (themachinewisdomproject.org).

References

- Ahmad, M. W., Mourshed, M., Yuce, B., & Rezgui, Y. (2016). Computational intelligence techniques for HVAC systems: A Review. *Building Simulation*, 9(4), 359–398. <https://doi.org/10.1007/s12273-016-0285-4>
- Bruijn, H., Warnier, M., & Janssen, M. (2022). The perils and pitfalls of explainable AI: Strategies for explaining algorithmic decision-making. *Government Information Quarterly*, 39(2), 101666. <https://doi.org/10.1016/j.giq.2021.101666>

- Cheng, C.-C., & Lee, D. (2016). Enabling Smart Air Conditioning by Sensor Development: A Review. *Sensors*, 16(12), 2028. <https://doi.org/10.3390/s16122028>
- Deng, Z., & Chen, Q. (2020). Development and validation of a SMART HVAC control system for multi-occupant offices by using occupants' physiological signals from wristband. *Energy and Buildings*, 214, 109872. <https://doi.org/10.1016/j.enbuild.2020.109872>
- Dong, B., Prakash, V., Feng, F., & O'Neill, Z. (2019). A review of Smart Building Sensing System for Better Indoor Environment Control. *Energy and Buildings*, 199, 29–46. <https://doi.org/10.1016/j.enbuild.2019.06.025>
- Dwivedi, R., Dave, D., Naik, H., Singhal, S., Omer, R., Patel, P., Qian, B., Wen, Z., Shah, T., Morgan, G., & Ranjan, R. (2023). Explainable AI (XAI): Core ideas, techniques, and solutions. *ACM Computing Surveys*, 55(9), 1-33. <https://doi.org/10.1145/3561048>
- Guan, H., Dong, L., & Zhao, A. (2022). Ethical risk factors and mechanisms in artificial intelligence decision making. *Behavioral Sciences*, 12(9), 343. <https://doi.org/10.3390/bs12090343>
- Hassanpour, H., Mhaskar, P., & Risbeck, M. J. (2022). A hybrid machine learning approach integrating recurrent neural networks with subspace identification for modeling HVAC systems. *Canadian Journal of Chemical Engineering*, 100(12), 3620-3634. <https://doi.org/10.1002/cjce.24392>
- Jacobi, T., & Sag, M. (2024). We are the AI problem. *Emory Law Journal Online*, 1–18. <https://ssrn.com/abstract=4820165>
- Karjalainen, S. (2009). Thermal comfort and use of thermostats in Finnish homes and offices. *Building and Environment*, 44(6), 1237–1245. <https://doi.org/10.1016/j.buildenv.2008.09.002>
- Kayser-Bril, N. (2020, April 7). *Google apologizes after its Vision AI produced racist results*. AlgorithmWatch. <https://algorithmwatch.org/en/google-vision-racism/>
- Lee, D., & Lee, S.-T. (2023). Artificial intelligence enabled energy-efficient heating, ventilation and air conditioning system: Design, analysis and necessary hardware upgrades. *Applied Thermal Engineering*, 235, 121253. <https://doi.org/10.1016/j.applthermaleng.2023.121253>
- Newell, Ben R., Lagnado, D. A., & Shanks, D. R. (2022). *Straight Choices - The Psychology of Decision Making*. <https://doi.org/10.4324/9781003289890>
- Ploug, T., & Holm, S. (2020). The four dimensions of contestable AI diagnostics - A patient-centric approach to explainable AI. *Artificial Intelligence in Medicine*, 107, 101901. <https://doi.org/10.1016/j.artmed.2020.101901>