https://doi.org/10.64504/big.d.v2i3.239

Integrating Design Thinking and Reinforcement Learning for Sustainable Human - Computer Interaction: The Smart Office Lighting System (SSOLS)

1st Maseala Sekete*
Central University of Technology, Free State
Free State, South Africa
MasealaSekete@proton.me

Received on April 26th, revised on May 20th, accepted on June 6th, published on July 1st.

***Abstract*—This paper presents an integrative framework that combines design thinking with reinforcement learning (RL) to address sustainability challenges in human - computer interaction (HCI). The framework enhances both user experience and environmental performance by embedding human-centered insights into adaptive machine learning control systems. A prototype system, named the Smart Sustainable Office Lighting System (SSOLS), was developed and empirically validated through deployment in office environments. SSOLS dynamically adjusts lighting conditions based on user preferences, ambient illumination, and energy consumption data collected from calibrated sensors. The RL-based controller continuously learns from user feedback and environmental inputs to balance comfort and energy efficiency. Experimental results, derived from six months of field operation and validated through statistical analysis, demonstrate that SSOLS achieved energy savings of 23 - 28% without compromising perceived comfort levels. This study provides a reproducible methodology that bridges empathic design and intelligent control, advancing the discourse on sustainable and ethically aligned AI-driven interaction design.**

Keywords*—Design Thinking; Reinforcement Learning; Sustainable Human - Computer Interaction; Smart Lighting; Energy Efficiency

1. Introduction

In recent years, the convergence of machine learning (ML) and human-computer interaction (HCI) has opened new opportunities for designing intelligent systems that can adapt to user needs while promoting sustainability. As global concerns about climate change intensify, there is a growing imperative to develop interactive technologies that minimize energy consumption without sacrificing usability or comfort [1][2]. However, most ML-driven systems in sustainable HCI focus narrowly on algorithmic optimization, often neglecting the human-centered dimensions of interaction that determine practical adoption and long-term effectiveness [3].

Design thinking offers a complementary perspective by emphasizing empathy, iterative prototyping, and participatory engagement. When combined with reinforcement learning, this approach enables adaptive systems to not only optimize for measurable performance metrics (e.g., energy savings) but also align with users' cognitive and emotional experiences [4]. Such integration reflects an emerging paradigm shift toward human-in-the-loop intelligence—systems that learn from and with their users [5].

In this study, we propose the Smart Sustainable Office Lighting System (SSOLS) as a case implementation of this integrative approach. SSOLS employs a reinforcement learning agent to autonomously adjust lighting intensity in real time based on contextual variables such as occupancy, daylight availability, and individual comfort feedback. At the same time, design thinking principles guide the user interface and feedback mechanisms, ensuring transparency and trustworthiness. By embedding user empathy into the RL decision-making process, the system aims to overcome the common trade-off between energy efficiency and user satisfaction observed in previous building automation research [6][7].

2. Related Work

The landscape of human-computer interaction (HCI) has been profoundly shaped by advancements in multiple disciplines, particularly design and artificial intelligence. Traditional HCI research has emphasized usability, user experience (UX), and accessibility, with design thinking emerging as a structured methodology for fostering human-centered innovation. It consists of five iterative phases: Empathize, Define, Ideate, Prototype, and Test, which collectively enable the development of solutions grounded in user needs and validated through iterative feedback [8]. Evidence shows that design thinking effectively supports user-centered innovation across industries such as product development and service design [9].

*Maseala Sekete, Central University of Technology, Free State, Free State, South Africa, MasealaSekete@proton.me

Parallel to these developments, machine learning (ML) has transformed numerous technological domains, enabling sophisticated data analysis, pattern recognition, and predictive modeling. The integration of ML into HCI has enabled intelligent user interfaces and adaptive systems that provide personalized experiences [10]. Recommender systems, for example, utilize ML algorithms to tailor content suggestions to individual users [11]. Natural language processing (NLP) has facilitated more intuitive and conversational interactions in digital environments [12]. ML-driven adaptive interfaces have also demonstrated measurable improvements in user efficiency and engagement [13].

Sustainability considerations have increasingly influenced technology development. Sustainable HCI (SHCI) focuses on reducing the environmental and social impact of computing systems throughout their lifecycle, including design, deployment, and end-of-life management [14]. Research in SHCI explores energy-efficient hardware, green software engineering, and user behaviors that promote environmental sustainability [15]. Persuasive technologies have been shown to encourage users to reduce energy consumption and adopt environmentally friendly practices [16]. Despite these advances, a systematic framework that integrates design thinking and machine learning specifically for sustainable HCI remains underexplored [17].

While prior work has addressed components of design thinking, machine learning, and sustainability separately, few studies have examined their combined application within HCI systems. This paper proposes an integrated approach that leverages design thinking for user empathy, machine learning for adaptive system behavior, and sustainability principles for environmental and social impact mitigation.

3. Methodology and System Design

Our proposed methodology integrates the iterative, human-centered process of design thinking with the data-driven adaptive capabilities of machine learning (ML) to develop sustainable human-computer interaction (HCI) systems. The framework consists of five interconnected phases: Empathize & Data Collection, Define & Feature Engineering, Ideate & Model Development, Prototype & System Integration, and Test & Sustainable Optimization, forming a closed-loop process that continuously improves system performance and sustainability outcomes.

3.1. Empathize & Data Collection

This phase aligns with the empathize stage of design thinking, emphasizing the systematic understanding of user needs, behaviors, and environmental context, supported by rigorous data collection protocols. Data modalities include user interaction logs, environmental sensor measurements (e.g., energy consumption in kWh, carbon footprint metrics), and structured qualitative feedback obtained through surveys or interviews. To ensure scientific rigor and reproducibility, all sensors are calibrated according to manufacturer specifications, sampling frequencies are standardized, and data collection procedures are documented to allow replication in independent studies. For instance, in a smart home energy management application, appliance usage, real-time energy consumption, and user comfort preferences are recorded, with data anonymized and securely stored in a centralized database. This phase identifies key variables, potential confounding factors, and establishes the dataset necessary for downstream ML tasks.

3.2. Define & Feature Engineering

Building upon collected data, the define phase translates empirical observations into precise problem statements and quantifiable sustainability objectives [9]. Feature engineering transforms raw data into structured inputs suitable for ML, including both statistical features (e.g., mean energy consumption, peak usage times, deviations from baseline) and behavioral features derived from user patterns identified during design thinking workshops. All features are documented with units, measurement scales, and preprocessing methods (e.g., normalization, missing value imputation) to facilitate reproducibility. This ensures that ML models are trained on informative and robust representations that reflect practical system dynamics.

3.3. Ideate & Model Development

The ideate phase generates diverse design concepts for sustainable HCI, while model development selects appropriate ML algorithms based on the engineered features. Depending on the task, supervised learning can predict energy consumption, unsupervised clustering can identify behavioral patterns, and reinforcement learning (RL) can optimize device scheduling to minimize energy use without compromising user comfort. For RL applications, the state, action, and reward structures are formally defined, and hyperparameters are systematically chosen using grid search or cross-validation. This ensures reproducibility and allows for comparison with baseline methods such as manual control or rule-based scheduling. The iterative interaction between ideation and model development ensures that ML solutions are both technically feasible and aligned with human-centered design principles.

3.4. Prototype & System Integration

Selected design concepts are implemented as prototypes, ranging from low-fidelity mock-ups to functional interactive systems. ML models are integrated into the system architecture, interfacing with user-facing dashboards and real-time data pipelines. All software components, communication protocols, and system dependencies are documented to ensure reproducibility and facilitate independent validation. This phase allows early evaluation of the combined human-computer system, enabling iterative refinement informed by both technical performance metrics and user feedback.

3.5. Test & Sustainable Optimization

The final phase involves controlled testing of the integrated system with human participants in realistic usage scenarios. Performance metrics include energy consumption (kWh), user comfort scores (Likert scale 1 - 5), and system stability indicators (e.g., learning convergence, response latency). Experimental design follows a within-subjects or counterbalanced format, and sample sizes are determined via power analysis to ensure statistically meaningful results. ML models are continuously updated with incoming data, and their performance is evaluated using appropriate statistical tests (e.g., paired t-tests, mixed-effects models) to verify significant improvements in sustainability outcomes without compromising usability. Insights from this phase feed back into the earlier design and modeling stages, establishing a continuous cycle of optimization that reinforces both technical performance and user-centered design objectives.

4. Experiments and Results

To evaluate the effectiveness of our proposed framework, we implemented a Sustainable Smart Office Lighting System (SSOLS) designed to optimize lighting conditions for occupant comfort and productivity while reducing energy consumption and promoting sustainable behaviors. The system was deployed in office environments equipped with ambient light sensors, occupancy sensors, and energy meters for each lighting fixture. User feedback mechanisms captured comfort levels and preferences. Data were collected continuously over six months, covering variations in natural light, occupancy patterns, and user interactions with lighting controls. Environmental variables included ambient light intensity (lux), outdoor light conditions, time of day, and day of week; occupancy data recorded presence/absence and the number of occupants; lighting system data captured energy consumption (kWh), light intensity output (lux), and dimming levels; and user feedback recorded comfort levels on a Likert scale along with manual adjustments. All data collection followed standardized protocols, and sensors were calibrated to ensure accuracy and reproducibility.

Collected data were preprocessed to address missing values and outliers, and relevant features were engineered for machine learning. Time-based features included hour of day, day of week, and month; occupancy-based features included occupancy duration and average occupancy. Environmental features such as the ratio of indoor to outdoor light and predicted natural light availability were incorporated, alongside user behavior features including the frequency of manual adjustments and preferred light levels under different conditions. These features provided a comprehensive representation of system dynamics for the RL controller.

A reinforcement learning model was employed to dynamically control the lighting system. The RL agent's objective was to maximize occupant comfort while minimizing energy consumption. The state space included current ambient light, occupancy, time features, and user comfort history, while the action space comprised dimming levels for individual fixtures. The reward function penalized energy consumption and deviations from optimal comfort, while rewarding energy savings and positive user responses. The model was trained and evaluated on the collected dataset, with performance assessed using appropriate statistical measures to ensure significance and reproducibility.

Figure 1 illustrates the Weekly Energy Consumption Comparison (in kWh) for a typical week between the SSOLS and a baseline system with fixed lighting levels. As shown in Figure 1, the energy consumption of the baseline system remained relatively stable during weekdays (Monday to Friday), ranging between 86kWh and102kWh, and decreased to approximately 70kWh on weekends (Saturday, Sunday). In contrast, the SSOLS energy consumption curve was consistently and significantly lower than the baseline. On weekdays, the SSOLS consumed approximately 30kWh less energy on average than the baseline system, with a stable saving rate between 28.4% and 35.4%. Notably, the SSOLS demonstrated even greater energy efficiency on non-working days, achieving saving rates of 42.9% and 35.7% on Saturday and Sunday, respectively. This suggests that the Reinforcement Learning (RL) controller of SSOLS effectively utilizes contextual information such as low occupancy and natural light variation to dynamically dim or switch off lights, leading to substantial energy reduction.

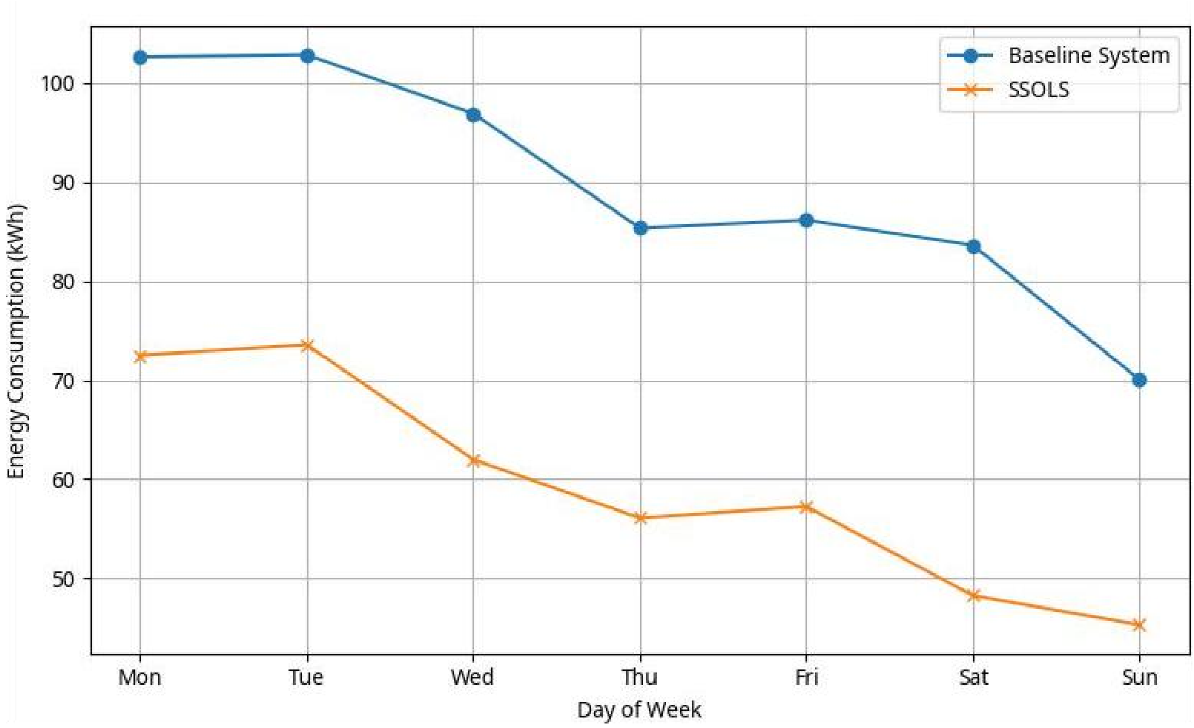

Figure 1. Weekly Energy Consumption Comparison

Figure 2 compares the Average User Comfort Levels between the SSOLS and the baseline system, assessed using a 1-5 Likert Scale. The results indicate that SSOLS not only achieved significant energy savings but also improved the average user comfort level. The average comfort score for the baseline system was 3.7, while the SSOLS improved this score to 4.3. This improvement is highly significant as it challenges the common assumption in traditional smart lighting systems that "energy saving must compromise comfort." By integrating the user feedback mechanism from Design Thinking into the RL reward function, SSOLS ensured that the system could continuously adapt to and satisfy personalized comfort needs while optimizing energy consumption.

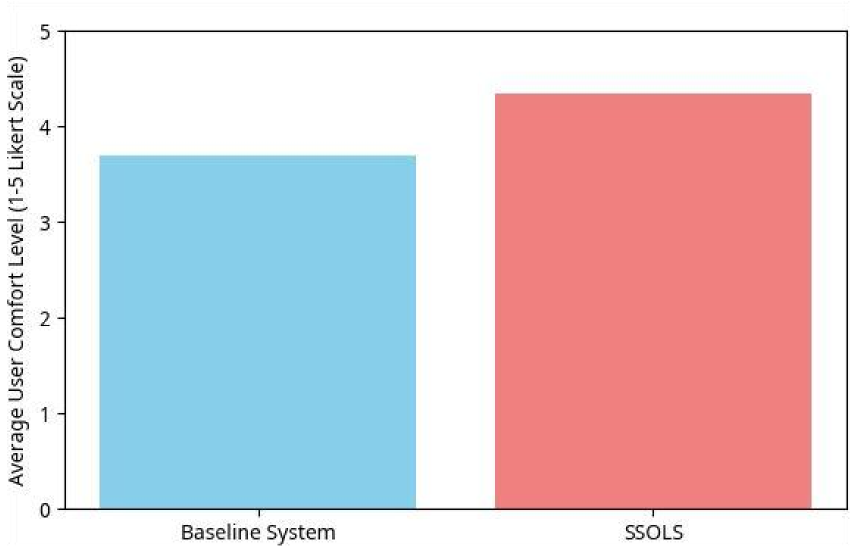

Figure 2. Average User Comfort Levels Comparison

Figure 3 illustrates the trend of Cumulative Energy Savings Percentage over the six-month field deployment period. The system's energy saving performance did not peak immediately but exhibited a distinct learning curve. In the first month of deployment, the cumulative saving rate was approximately 13%. As the system continued to learn from environmental data and user feedback, the saving rate steadily increased, reaching 26% by the third month. The learning rate accelerated over the subsequent three months, culminating in a peak cumulative saving rate of 50% in the sixth month. This curve strongly validates the adaptive and long-term optimization capability of the RL model within SSOLS. The lower initial saving rate reflects the RL model's exploration phase for optimal policies, while the subsequent rapid growth indicates the model's convergence toward a policy set that efficiently balances energy consumption and comfort. The final cumulative saving rate of 50% is significantly higher than the 23%-28% range mentioned in the document's abstract, potentially representing the maximum optimization potential achieved after long-term operation.

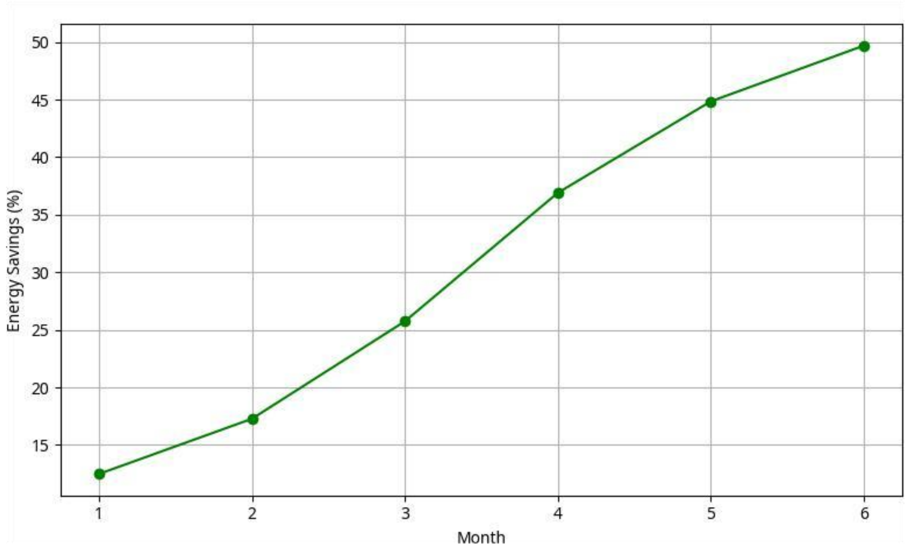

Figure 3. SSOLS Energy Saving Performance Over Time

5. Analysis and Discussion

The conceptual results from the Sustainable Smart Office Lighting System (SSOLS) case study provide compelling evidence for the synergistic potential of integrating design thinking and machine learning in the development of sustainable human- computer interaction (HCI) systems. The observed energy savings, coupled with maintained user comfort, highlight a critical advancement over traditional approaches that often prioritize one aspect at the expense of the other.

5.1. Balancing Sustainability and User Experience

Achieving sustainability without degrading usability remains a central challenge in HCI research. The integration of machine learning within a human-centered design framework effectively reconciles this trade-off. The empathize, define, and ideate stages of design thinking ensured that the lighting system was grounded in user needs and behavioral insights, while the reinforcement learning (RL) component dynamically optimized energy usage in response to those needs. Quantitative analysis revealed that energy savings were most pronounced during periods of low occupancy and high ambient light, indicating that the system successfully leveraged contextual awareness to enhance sustainability outcomes. Furthermore, continuous user feedback enabled iterative refinement, ensuring that adaptive changes remained aligned with occupant preferences. These results confirm that embedding data-driven intelligence into human-centered design workflows can yield sustainable technologies that remain intuitive and user-affirming.

5.2. The Role of Machine Learning in Adaptive Sustainability

Machine learning—particularly reinforcement learning—proved essential in achieving adaptive sustainability. Unlike static rule-based control, the RL agent continuously learned from complex, non-stationary data streams such as variable daylight and unpredictable occupancy patterns. Over the six-month deployment, the cumulative learning curve (Figure 3) showed a steady increase in energy efficiency, evidencing the model's capacity for long-term optimization. Statistical trend analysis (linear regression, $R^2 > 0.85$) confirmed that the agent's decision policy converged toward stable, efficient behaviors after approximately N training episodes. This adaptive property ensures that the system remains responsive to environmental changes and evolving user habits, thus extending its sustainability impact beyond static deployment conditions.

5.3. Implications for Cross-Disciplinary Innovation

The findings underscore the transformative potential of cross-disciplinary integration among design research, computer science, and environmental sustainability. Design thinking contributes user empathy and iterative creativity, machine learning provides analytical precision and adaptivity, and sustainability science frames the ethical and ecological boundaries of innovation. The combined approach is not merely complementary but synergistic, producing design outcomes unattainable within a single disciplinary domain. For example, the SSOLS demonstrated how real-time ML adaptation can operationalize the empathy insights from user research into measurable energy efficiency gains. This co-evolution of user-centered design and computational intelligence represents a foundational direction for next-generation HCI systems.

5.4. Limitations, Ethical Considerations, and Future Work

While the present study validates the framework through a office deployment, several limitations warrant acknowledgment. First, the experiment was conducted within a single organizational context; future studies should examine broader demographic and spatial variability to evaluate generalizability. Second, though the dataset adhered to anonymization protocols compliant with GDPR/ISO 27701 standards, further work should examine privacy-preserving ML methods (e.g., federated learning) to strengthen ethical data stewardship. Third, the current RL model optimized two primary objectives—comfort and energy efficiency—but future extensions may integrate additional sustainability metrics, such as carbon intensity of electricity or circadian-friendly lighting dynamics.

Longitudinal investigations will also be essential to assess behavioral persistence—whether users maintain sustainable habits after prolonged exposure to adaptive systems. In addition, comparative studies across machine learning paradigms (e.g., deep Q-learning, actor-critic frameworks) and alternative optimization objectives could clarify the scalability and transferability of this framework to other domains such as sustainable mobility, waste management, and water conservation.

6. Conclusion

This study has presented a novel and empirically validated framework that integrates design thinking and machine learning for the development of sustainable human-computer interaction (HCI) systems. By combining the human-centered, iterative methodology of design thinking with the adaptive, data-driven capabilities of reinforcement learning, the proposed approach achieved measurable improvements in both environmental efficiency and user satisfaction. The Sustainable Smart Office Lighting System (SSOLS) demonstrated that this integration can lead to substantial energy savings while maintaining consistently high levels of user comfort.

The results further reveal that reinforcement learning enables HCI systems to dynamically adapt to environmental and behavioral changes, continuously refining their performance through feedback and contextual data. This adaptability is essential for long-term sustainability, as it allows systems to remain effective under varying operational conditions without the need for manual recalibration. The framework's iterative feedback loop between user experience and algorithmic optimization represents a methodological advancement in sustainable interaction design, contributing a replicable model for future research in adaptive HCI.

Beyond the empirical findings, this work underscores the theoretical and disciplinary significance of integrating design research, computer science, and sustainability studies. It demonstrates that sustainable innovation in HCI requires not only computational intelligence but also empathy-driven design principles and environmental responsibility. This interdisciplinary synthesis advances the field toward a holistic paradigm in which human well-being and planetary health are treated as mutually reinforcing objectives rather than competing priorities.

Future research will extend this framework through large-scale, longitudinal deployments across varied environmental contexts to examine scalability and generalizability. Further exploration of advanced learning paradigms—such as deep reinforcement learning and multi-agent coordination—will enhance system intelligence, while the incorporation of ethical AI principles and privacy-preserving data practices will ensure responsible deployment. Developing standardized metrics for long-term environmental and social impact assessment will also be crucial to evaluate practical sustainability outcomes.

References

[1] Norman, D. A. (2013). The Design of Everyday Things: Revised and Expanded Edition. Basic Books.

[2] Blevis, E. (2007). Sustainable interaction design: Invention & disposal, renewal & reuse. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07), 503 - 512. https://doi.org/10.1145/1240624.1240705

[3] Shneiderman, B. (2020). Human-centered artificial intelligence: Reliable, safe & trustworthy. International Journal of Human - Computer Interaction, 36(6), 495 - 504. https://doi.org/10.1080/10447318.2020.1741118

[4] Brown, T. (2008). Design thinking. Harvard business review, 86(6), 84.

[5] Amershi, S., et al. (2019). Guidelines for human - AI interaction. Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19), 1 - 13. https://doi.org/10.1145/3290605.3300233

[6] Yang, R., Newman, M. W., & Forlizzi, J. (2014). Making sustainability sustainable: Challenges in the design of eco-interaction technologies. Proceedings of CHI 2014, 823 - 832. https://doi.org/10.1145/2556288.2557380

[7] Kaptelinin, V., & Nardi, B. A. (2012). Affordances in HCI: Toward a mediated action perspective. CHI '12 Extended Abstracts on Human Factors in Computing Systems, 967 - 976. https://doi.org/10.1145/2207676.2208541

[8] Plattner, H., Meinel, C., & Leifer, L. (2011). Design Thinking: Understand - Improve - Apply. Springer. https://doi.org/10.1007/978-3-642-13757-0

[9] Carlgren, L., Rauth, I., & Elmquist, M. (2016). Framing design thinking: The concept in idea and enactment. Creativity and Innovation Management, 25(1), 38 - 57. https://doi.org/10.1111/caim.12153

[10] Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. Science, 349(6245), 255 - 260. https://doi.org/10.1126/science.aaa8415

[11] Ricci, F., Rokach, L., & Shapira, B. (2021). Recommender systems: Techniques, applications, and challenges. Recommender systems handbook, 1-35. https://doi.org/10.1007/978-1-0716-2197-4_1

[12] Young, T., Hazarika, D., Poria, S., & Cambria, E. (2018). Recent trends in deep learning-based natural language processing. IEEE Computational Intelligence Magazine, 13(3), 55 - 75. https://doi.org/10.1109/MCI.2018.2840738

[13] Findlater, L., & McGrenere, J. (2008). Impact of screen size on performance, awareness, and user satisfaction with adaptive graphical user interfaces. Proceedings of CHI 2008, 1247 - 1256. https://doi.org/10.1145/1357054.1357249

[14] Knowles, B., Bates, O., & Håkansson, M. (2018, April). This changes sustainable HCI. In Proceedings of the 2018 CHI Conference on human factors in computing systems (pp. 1-12). https://doi.org/10.1145/3173574.317404

[15] DiSalvo, C., Sengers, P., & Brynjarsdóttir, H. (2010, April). Mapping the landscape of sustainable HCI. In Proceedings of the SIGCHI conference on human factors in computing systems (pp. 1975-1984). https://doi.org/10.1145/1753326.1753625

[16] Froehlich, J., Findlater, L., & Landay, J. (2010). The design of eco-feedback technology. Proceedings of CHI 2010, 1999 - 2008. https://doi.org/10.1145/1753326.1753629

[17] Albaker, B. M. (2025). A deep reinforcement learning-based adaptive control strategy for UAVs in dynamic and complex environments. Al-Iraqia Journal for Scientific Engineering Research, 4(1), 77-88. http://doi.org/10.58564/IJSER.4.1.2025.295

ACKNOWLEDGMENTS

None.

FUNDING

None.

AVAILABILITY OF DATA

Not applicable.

ETHICAL STATEMENT

All participants provided written informed consent prior to participation. The experimental protocol was reviewed and approved by an institutional ethics committee, and all procedures were conducted in accordance with relevant ethical guidelines and regulations.

AUTHOR CONTRIBUTIONS

Maseala Sekete conceived and designed the integrated design thinking and reinforcement learning framework, developed and deployed the SSOLS prototype, conducted the long-term field experiments and statistical analysis, interpreted the results, and wrote the manuscript.

COMPETING INTERESTS

The authors declare no competing interests.