Interactive Recommendation of Sustainable Energy Management Strategies through Diversified Rule Mining and Preference-Aware Embeddings

Ruben Geword
Department of Computer Science, University of Houston, Houston, TX, USA.
rubenhoward95@uh.edu

Chengying Han
Department of Computer Science, George Mason University, Fairfax, VA, USA.
chengyingh@gmu.edu

Abstract

The transformation of energy infrastructures toward sustainability demands decision-support systems that can translate vast operational data into actionable strategies while accommodating diverse stakeholder preferences and dynamic grid conditions. This paper presents a system-level investigation into an interactive recommendation framework that integrates diversified rule mining with preference-aware embeddings to generate personalized energy management strategies. We examine the architecture, trade-offs, and governance dimensions of coupling interpretable rule discovery with latent preference representations in large-scale socio-technical energy systems. The discussion foregrounds the structural tension between diversity of recommendations and personalization accuracy, the computational and infrastructural requirements for real-time interactivity, and the need for robustness under concept drift and data heterogeneity. We analyze how diversified top-k rule mining, guided by user embeddings, can surface a broad set of actionable patterns from smart meter, weather, and market data, while embedding-based user models enable the system to refine suggestions through iterative feedback. The paper further addresses fairness implications across consumer segments, the role of transparency and auditability in algorithmic governance, and the alignment of such systems with energy policy objectives such as demand response, renewable integration, and energy equity. By framing recommendation as a continuous interactive dialogue rather than a static optimization, we highlight pathways for responsible deployment of AI in critical energy infrastructures. We conclude with reflections on future research directions that bridge adaptive rule learning, human-in-the-loop design, and sustainable energy transitions.

Keywords

Sustainable energy management, interactive recommendation, diversified rule mining, preference-aware embeddings, socio-technical systems, fairness.

1. Introduction

The accelerating adoption of distributed energy resources, smart metering infrastructure, and real-time grid monitoring has transformed the landscape of energy management from a centralized, static discipline into a dynamic, data-rich environment. Sustainable energy strategies now require granular, context-aware decision-making that integrates consumption patterns, generation forecasts, market signals, and user behavior. A growing body of research has emphasized the role of renewable energy strategies in achieving long-term sustainability, while smart grid surveys underscore the necessity of bidirectional communication and intelligent control for managing this complexity [1][2]. Concomitantly, the proliferation of big data in energy systems has created opportunities to extract insights at unprecedented scale, yet the translation of raw data into actionable recommendations for diverse end-users remains a significant challenge [3].

Recommendations for energy management span a spectrum of interventions, from appliance-level scheduling to building-wide retrofits and community-scale storage dispatch. Unlike conventional recommender systems in e-commerce or media streaming, energy contexts introduce hard constraints related to physical infrastructure, grid stability, safety regulations, and long-term sustainability goals. Furthermore, user preferences are multidimensional, encompassing cost savings, comfort, carbon footprint reduction, and risk tolerance. The heterogeneity of these preferences, combined with the need for explanations that foster trust and behavioral change, calls for approaches that are simultaneously interpretable, personalized, and exploratory. This motivates the integration of two complementary paradigms: rule mining, which offers transparent, human-readable patterns, and preference-aware embeddings, which capture nuanced user inclinations in dense vector spaces suitable for interactive refinement.

Rule mining has long been valued for its ability to generate explicit associations between conditions and outcomes, making it a natural fit for domains requiring auditability. However, traditional association rule mining often yields an overwhelming number of patterns, many of which are redundant or trivial. The need for diversified rule sets—summaries that cover distinct aspects of the data without sacrificing utility—has become apparent, particularly in decision-support settings where human cognitive capacity is limited. Meanwhile, the success of embedding techniques in natural language processing and collaborative filtering has demonstrated that latent representations can encode complex preference structures and facilitate efficient similarity computations. Translating these capabilities to the energy domain requires careful system design that respects real-time constraints, data privacy, and the interpretability mandate.

This paper presents a comprehensive system-level analysis of an interactive recommendation architecture that combines diversified rule mining with preference-aware embeddings to deliver sustainable energy management strategies. We examine the foundational technologies, architectural design principles, and socio-technical dimensions that shape the viability of such a system in practice. The analysis is deliberately interdisciplinary, drawing on perspectives from large-scale systems engineering, artificial intelligence, human-computer interaction, and energy policy. By structuring the discussion around structural trade-offs, governance, robustness, and deployment, we aim to provide a holistic view that can guide both future research and real-world implementation.

2. Background and Motivation

The intellectual foundations of the proposed framework rest on three pillars: rule mining for knowledge discovery, interactive recommendation systems, and representation learning for preference modeling. Association rule mining, originating in market basket analysis, has evolved into a versatile tool for extracting frequent patterns and causal structures from transactional and temporal data [4]. Its applicability to energy data is compelling: rules such as “if outdoor temperature exceeds 30°C and time is between 14:00 and 17:00, then cooling load increases by 40% with 85% confidence” provide actionable, verifiable insights that can directly inform load shifting or demand response programs. However, raw rule sets are often massive and highly correlated, necessitating post-processing for diversity and relevance.

Interactive recommendation systems extend traditional recommenders by incorporating user feedback loops, allowing the system to adapt to changing preferences and to explore the decision space collaboratively [5]. Unlike one-shot recommendations, interactive cycles enable users to express latent preferences through critique, selection, or demonstration, which are then used to update the model. In energy management, this interactive paradigm aligns well with the iterative nature of strategy refinement: a building manager might reject a recommendation that reduces cost but increases occupant discomfort, prompting the system to propose alternatives that balance both objectives. The diversity of the recommended strategies is crucial, as presenting a set of complementary options increases the likelihood that at least one aligns with the user's unstated constraints. Research on diversity in recommender systems has demonstrated that topic diversification can improve user satisfaction and discovery, a principle that we extend to the rule space [6].

Preference-aware embeddings draw on vector space models, notably the family of techniques popularized by word2vec and extended to user and item representations [7]. By embedding energy strategies and user profiles into a shared latent space, the system can compute similarity, anticipate user reactions, and interpolate between known preferences to generate novel suggestions. This embedding layer serves as a differentiable memory of user feedback, enabling smooth updates as new information arrives. Critically, the embedding can be learned not only from explicit ratings but also from implicit signals such as dwell time, adjustment frequency, or deployment outcomes, thereby reducing the burden on users to articulate their preferences. The combination of explicit rules and implicit embeddings creates a dual representation that supports both transparency and adaptability.

The motivation for integrating these components is twofold. First, the sustainability transition demands that energy strategies be not only optimal in a narrow technical sense but also acceptable and adoptable by a wide range of stakeholders. Interpretable rule-based explanations can demystify algorithmic decisions and build trust, a recognized necessity in high-stakes domains [8]. Second, energy environments are non-stationary; consumption patterns shift with seasons, technological upgrades, and behavioral trends. A purely static model becomes obsolete quickly, whereas an interactive system that continuously mines diversified rules and updates user embeddings can maintain relevance over time. This system-level perspective necessitates a design that carefully balances computational efficiency, data freshness, and user engagement.

3. System Architecture and Design Principles

The envisioned architecture comprises four interconnected layers: a data ingestion and preprocessing layer, a diversified rule mining engine, a preference-aware embedding module, and an interactive recommendation interface. The data ingestion layer aggregates heterogeneous streams from smart meters, weather services, grid pricing signals, and building management systems. This layer must address challenges of data velocity, veracity, and format heterogeneity, employing stream processing techniques to provide a unified, timestamped event log. Privacy considerations, particularly for residential data, demand that personally identifiable information be anonymized or processed locally before transmission, aligning with emerging regulations on smart meter data [13]. The preprocessing pipeline also performs temporal alignment, missing value imputation, and feature engineering to derive semantically meaningful attributes such as peak load ratio, load flexibility index, and carbon intensity of consumed electricity.

The diversified rule mining engine operates on the preprocessed data to extract a set of high-quality, non-redundant rules. A critical design choice is the metric that governs diversity. Structural diversity might favor rules that involve different attributes or time scales, while coverage diversity seeks rules that describe distinct subsets of the data. The engine leverages a user-guided embedding to steer the search toward rules that are likely to be relevant to the current user or user segment, thereby integrating personalization at the mining stage itself. The fast diversified top-k rule discovery approach that uses user-guided embeddings has been shown to significantly improve the efficiency and relevance of rule sets in large databases, and we adopt this as a core algorithmic component [9]. The engine outputs a rule inventory that is indexed for rapid retrieval and scoring.

The preference-aware embedding module maintains a latent representation for each user, energy strategy, and contextual condition. Initial user embeddings can be bootstrapped from demographic data, past energy audits, or a brief onboarding questionnaire, and are then continuously updated via online learning from interactive feedback. Strategy embeddings are derived from the textual or structural description of the rules themselves, as well as from historical performance metrics such as actual cost savings, comfort ratings, and reliability. The embedding space is trained to minimize a composite loss that encourages similarity between a user and strategies that they have accepted or deployed, while pushing apart those rejected or ignored. An important architectural decision concerns the placement of the embedding model: a centralized server-side model offers global consistency but raises latency and privacy concerns, whereas a federated learning approach can keep raw user data on edge devices, training local embeddings that are aggregated to improve the global model without exposing sensitive information.

The interactive recommendation interface closes the loop by presenting a curated set of strategies to the user, collecting feedback, and triggering updates to both the embedding model and the rule inventory. The presentation layer must balance information richness with cognitive load, perhaps using visualizations that highlight trade-offs between cost, carbon, and comfort. Feedback can be explicit, such as liking or discarding a strategy, or implicit, such as the time spent examining a rule’s details. The interface also supports user-driven exploration, allowing the user to modify the preference sliders or query for rules containing specific conditions, thereby blending recommendation with direct manipulation.

4. Diversified Rule Mining for Actionable Energy Management Patterns

The extraction of meaningful rules from energy data streams is subject to several domain-specific constraints that distinguish it from generic pattern mining. Energy data are inherently temporal and often exhibit strong periodicity at daily, weekly, and seasonal scales. Consequently, rule antecedents and consequents must be capable of expressing temporal relationships, such as “if the average load over the past three hours is above a threshold, then the peak load in the next hour will exceed the demand limit with high probability.” Traditional frequent itemset mining is insufficient to capture such dynamics, requiring extensions to sequential pattern mining or temporal logic. The rule mining engine must therefore be equipped with a rich temporal query language and efficient algorithms for mining over sliding windows, ensuring that the rule base remains current as new data arrive.

Diversification serves multiple purposes in this context. From a cognitive perspective, presenting a user with fifty variations of the same underlying pattern is inefficient and may obscure more novel insights. From a strategic perspective, a diverse set of rules covers a wider action space, increasing the chance of surfacing non-obvious strategies, such as combining battery storage dispatch with pre-cooling schedules or exploiting time-of-use arbitrage in deregulated markets. The diversification metric must be carefully tuned to avoid fragmenting the strategy space into overly niche rules that lack statistical support. A common approach is to define a similarity function between rules based on attribute overlap, confidence, or embedding proximity, and then to select a top-k set that maximizes a diversity objective while preserving aggregate quality. By employing user-guided embeddings, the diversification process can be biased toward rules that fall within the user's preference neighborhood, thus achieving a purposeful balance between coverage and personalization [9]. This method avoids the pitfall of recommending irrelevant albeit diverse rules, which can frustrate users and reduce trust.

Actionability is another critical filter. A rule may be statistically sound but operationally infeasible if it relies on variables that the user cannot control or observe. For instance, a rule that prescribes reducing electricity import when the grid's marginal emission factor is low is actionable only if real-time emission signals are accessible to the user's energy management system. The rule mining pipeline must therefore integrate a feasibility model that prunes rules based on user capabilities, contractual constraints, and regulatory limitations. This feasibility check requires a knowledge base that encodes domain expertise about building physics, appliance characteristics, and market rules, demonstrating the deep coupling between data-driven mining and domain engineering.

The lifecycle management of mined rules is a system-level concern that extends beyond initial discovery. As operating conditions change, previously valid rules may become obsolete, a phenomenon known as concept drift [18]. The system must monitor rule performance continuously, detecting decay in confidence or support and triggering re-mining or rule retirement. Archiving historical rules along with their performance trails can also support ex-post analysis and regulatory audits, providing a transparent record of how and when particular strategies were recommended and adopted.

5. Preference-Aware Embeddings for Personalized Strategy Alignment

Preference modeling in sustainable energy management must accommodate the multi-objective and often context-dependent nature of user goals. A single utility function that aggregates cost, carbon, and comfort with fixed weights is rarely adequate, as users may shift priorities depending on external factors such as weather extremes, financial stress, or public awareness campaigns. Embedding-based representations excel at capturing such shifting landscapes by learning a manifold in which different preference modes correspond to distinct regions or trajectories. The model can infer that a user who has recently started exploring carbon-intensive strategies may be experiencing a budget constraint, and adapt the recommendations accordingly.

The training of preference embeddings from interactive feedback poses a challenging cold-start problem, both for new users and for existing users who encounter novel types of strategies. Transfer learning can mitigate this by initializing new user embeddings from aggregate segment models, such as "residential prosumer with rooftop PV" or "small commercial building with HVAC optimization," and then individualizing through online gradient updates. The embedding space can also be regularized to enforce structure, such as requiring that strategies with similar rule conditions map to nearby points, which improves generalization. Moreover, the embedding module can incorporate contextual factors as additional dimensions, such as season, occupancy schedule, or electricity price, enabling contextual bandit-style adaptation where the recommendation policy shifts smoothly with changing conditions.

A key system design question is the degree of autonomy granted to the embedding-driven recommender. In a fully automated control paradigm, the system would directly execute the top-ranked strategy, such as adjusting thermostat setpoints or charging batteries. However, in many real-world settings, human approval is required for interventions that affect comfort or incur costs, pushing the system toward a mixed-initiative interaction model. The embedding module thus serves not only to rank strategies but also to predict the probability that a user will accept a recommendation, which can guide the interface in deferring low-confidence suggestions. The interactive loop enables the system to refine its confidence estimates rapidly, turning initial hesitancy into learned precision.

It is important to note the potential for feedback loops to narrow the diversity of presented strategies over time. If the embedding model becomes overfitted to the user's past choices, it may reinforce an echo chamber, suppressing exploration of unfamiliar but potentially superior strategies. To counteract this, the system must maintain an exploration budget, injecting diversified rules that lie outside the current embedding neighborhood but have high objective sustainability metrics. Balancing exploration and exploitation in this setting is a structural requirement that touches both the rule mining diversification objective and the embedding-based scoring.

6. Infrastructure, Scalability, and Deployment Challenges

Deploying an interactive recommendation system at the scale of a regional distribution grid or a nation-wide utility requires careful attention to the computational and communication infrastructure. The diversified rule mining engine, if run centrally, must process high-velocity data streams from millions of endpoints, potentially generating trillions of candidate rules. Efficient algorithms that leverage prefix-tree structures, parallelization, and pruning heuristics are essential to keep latency within interactive bounds. The embedding training and inference pipeline also demands substantial computational resources, particularly if deep neural architectures are used for encoding rule semantics. In practice, a hybrid architecture that delegates privacy-sensitive computations to edge gateways and aggregates anonymized model updates at cloud servers can distribute the load while respecting data locality [12].

Latency constraints are particularly stringent in real-time demand response scenarios, where recommendations must be generated within seconds to respond to grid frequency deviations or price spikes. The system must therefore be designed with a clear separation between batch and online processing. The diversified rule base can be refreshed periodically, such as daily or weekly, while the embedding-based ranking and interactive feedback loop operate in real time, querying the precomputed rule inventory. Caching and indexing strategies, including approximate nearest neighbor search in the embedding space, become critical to meet latency targets. The architecture must also be fault-tolerant, with redundant components and graceful degradation modes so that a failure in the recommendation engine does not compromise the underlying energy management controls.

Data governance adds another layer of infrastructure complexity. Energy data are subject to a patchwork of regulations that vary by jurisdiction, ranging from the strict personal data protections of the GDPR in Europe to sector-specific rules for utility data access in the United States. The system must implement fine-grained access controls, data minimization principles, and auditing mechanisms to ensure compliance. The use of preference embeddings, which can inadvertently memorize sensitive attributes, must be accompanied by privacy-preserving techniques such as differential privacy during training or homomorphic encryption for secure inference. These technical measures must be integrated into the deployment blueprint from the outset rather than retrofitted, as they affect fundamental design choices about data centralization versus federation.

Testing and validation present additional deployment hurdles. The safety-critical nature of energy systems demands rigorous evaluation before recommendations are acted upon. A shadow deployment phase, where the system runs in parallel with existing decision processes without affecting operations, allows for comparison with historical baselines and identification of unstable recommendations. Continuous monitoring for fairness metrics, recommendation quality, and system health is essential post-deployment, feeding into a feedback-driven DevOps cycle that can trigger retraining or rule updates.

7. Fairness, Robustness, and Governance in Energy Decision Support

Algorithmic fairness in energy management is a multifaceted concern that intersects with energy equity, affordability, and environmental justice. Recommendations that systematically benefit certain consumer groups—for instance, those with flexible loads or smart appliances—while excluding others can exacerbate existing disparities. A rule-based recommendation system could inadvertently propose load-shifting strategies that rely on owning electric vehicles or home batteries, marginalizing low-income households. The embedding model, if trained on biased historical feedback, might similarly associate certain demographic segments with lower willingness to adopt sustainable practices, creating a self-fulfilling prophecy. Addressing fairness requires explicit design interventions, such as counterfactual fairness criteria [14] or multi-objective optimization that includes equity metrics as a dimension alongside cost and carbon.

Robustness to distribution shift is equally critical. Energy systems are subject to sudden changes from policy interventions, technological breakthroughs, and climate extremes. A rule that was optimal under last year's tariff structure may become detrimental after a rate redesign. The system must incorporate mechanisms for detecting such shifts and rapidly re-evaluating the validity of its knowledge base. Online learning of preference embeddings can adapt quickly to user behavior changes, but rule mining operates on a slower refresh cycle, creating a temporal misalignment that must be managed through sensitivity analysis and scenario testing. The governance framework should mandate that high-stakes recommendations, such as those involving significant financial commitments or safety risks, be subject to human review and conservative defaults.

Transparency and explainability are foundational to the governance of AI in critical infrastructures. The dual-representation architecture of this system provides a natural path to explanation: every recommendation can be accompanied by the specific rule or set of rules that support it, presented in natural language or graphical form. The embedding-based personalization can be explained by highlighting similar strategies that the user has previously accepted. Audit trails that log the provenance of rules, the sequence of user interactions, and the model updates provide accountability and support regulatory oversight. Establishing independent auditing bodies or algorithmic impact assessment processes, as advocated in emerging AI governance frameworks [19], can foster public trust and ensure alignment with societal values.

The governance model must also delineate the boundaries of system autonomy. In fully automated demand response, the recommender may directly control loads, raising liability and safety questions. A layered governance structure that distinguishes between advisory recommendations, semi-automatic actions requiring confirmation, and fully automatic emergency controls can balance efficiency with accountability. This structure should be co-designed with regulators, utilities, and consumer advocates to reflect the risk tolerance and ethical norms of the communities served.

8. Policy Implications and Socio-Technical Considerations

The deployment of interactive recommendation systems in the energy sector both influences and is influenced by the policy landscape. Energy efficiency mandates, carbon pricing, and renewable portfolio standards create the economic incentives that make sustainable strategies valuable, and the system's recommendations must be aligned with these policy signals. For instance, a rule that encourages charging electric vehicles during midday solar surplus supports renewable integration goals, while one that shifts load to high-carbon periods would contradict climate policies. Embedding these policy constraints directly into the rule mining feasibility filter and the embedding objective function ensures that the system promotes societally beneficial outcomes.

Behavioral and social dimensions are equally important. Research in sustainable human-computer interaction has shown that technology alone rarely induces lasting behavioral change without attention to motivation, social norms, and feedback design [17]. The interactive interface of our system can incorporate behavioral insights by framing recommendations in terms of social comparisons, environmental impact, or long-term savings, and by providing gentle nudges towards lower-carbon options. The preference embedding model can learn which framings are most effective for different users, adapting the presentation without altering the underlying strategy. However, such persuasive techniques must be deployed ethically, with transparency and respect for user autonomy, to avoid manipulative patterns that erode trust.

Energy justice considerations require that the system's benefits be accessible across diverse socioeconomic and geographic contexts. Strategies that depend on high-speed internet, smartphone ownership, or advanced metering infrastructure may exclude rural or underserved populations. A responsible deployment strategy must include alternative interaction channels, such as SMS-based feedback or printed reports, and must adapt the recommendation logic to the available infrastructure. The rule base should include strategies tailored to low-tech environments, such as passive cooling techniques or manual load shifting schedules, ensuring that the system does not inadvertently reinforce a digital divide in the clean energy transition.

The long-term sustainability of the recommendation system itself must be considered. Continuous mining and online learning incur operational energy costs from data centers and network equipment, which should be benchmarked against the energy savings achieved. A life-cycle assessment perspective encourages the design of efficient algorithms and the use of carbon-aware computing resources, such as scheduling batch mining jobs during periods of high renewable availability. This reflexivity aligns the tool with the very sustainability principles it seeks to promote.

9. Conclusion

The interactive recommendation of sustainable energy management strategies through diversified rule mining and preference-aware embeddings represents a promising confluence of interpretable machine learning, interactive systems design, and energy domain knowledge. This paper has provided a system-level analysis that highlights the architectural decisions, structural trade-offs, and socio-technical imperatives inherent in building such a system. We have argued that the coupling of explicit, diversified rule sets with latent preference embeddings creates a dual representation that balances transparency and personalization, supporting both user trust and adaptive exploration. The discussion has underscored the importance of designing for diversity not only in rule output but also in user engagement, fairness, and infrastructure resilience.

The perspective offered here extends beyond purely algorithmic considerations to embrace the governance, policy, and deployment contexts that determine real-world impact. Energy systems are deeply embedded in social, economic, and regulatory fabrics, and any AI-driven tool must be evaluated against criteria of equity, robustness, and sustainability of its own operation. Future research should advance empirical evaluations of such integrated systems in living lab settings, develop rigorous metrics for interactive diversity and fairness in energy recommendations, and explore federated architectures that preserve privacy while enabling collective learning. By bridging the gap between data-driven pattern mining and human-centered design, interactive energy recommender systems can contribute meaningfully to the global pursuit of a just and sustainable energy future.

References

1. Lund, H. (2007). Renewable energy strategies for sustainable development. Energy, 32(6), 912–919.

2. Fang, X., Misra, S., Xue, G., & Yang, D. (2012). Smart grid—The new and improved power grid: A survey. IEEE Communications Surveys & Tutorials, 14(4), 944–980.

3. Zhou, K., Fu, C., & Yang, S. (2016). Big data driven smart energy management: From big data to big insights. Renewable and Sustainable Energy Reviews, 56, 215–225.

4. Agrawal, R., Imieliński, T., & Swami, A. (1993). Mining association rules between sets of items in large databases. ACM SIGMOD Record, 22(2), 207–216.

5. Konstan, J. A., & Riedl, J. (2012). Recommender systems: from algorithms to user experience. User Modeling and User-Adapted Interaction, 22(1), 101–123.

6. Ziegler, C. N., McNee, S. M., Konstan, J. A., & Lausen, G. (2005). Improving recommendation lists through topic diversification. In Proceedings of the 14th International Conference on World Wide Web (pp. 22–32).

7. Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781.

8. Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv preprint arXiv:1702.08608.

9. Han, Z., Chen, W., Han, Y., Mao, R., & Qin, J. (2026). Fast Diversified Top-k Rule Discovery via User-Guided Embeddings. IEEE Transactions on Knowledge and Data Engineering.

10. Koren, Y., Bell, R., & Volinsky, C. (2009). Matrix factorization techniques for recommender systems. Computer, 42(8), 30–37.

11. Palensky, P., & Dietrich, D. (2011). Demand side management: Demand response, intelligent energy systems, and smart loads. IEEE Transactions on Industrial Informatics, 7(3), 381–388.

12. Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. IEEE Internet of Things Journal, 3(5), 637–646.

13. Efthymiou, C., & Kalogridis, G. (2010). Smart grid privacy via anonymization of smart metering data. In 2010 First IEEE International Conference on Smart Grid Communications (pp. 238–243).

14. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., & Zemel, R. (2012). Fairness through awareness. In Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (pp. 214–226).

15. Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. ACM Computing Surveys, 46(4), 1–37.

16. Allcott, H., & Mullainathan, S. (2010). Behavior and energy policy. Science, 327(5970), 1204–1205.

17. DiSalvo, C., Sengers, P., & Brynjarsdóttir, H. (2010). Mapping the landscape of sustainable HCI. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 1975–1984).

18. Joachims, T., Swaminathan, A., & Schnabel, T. (2017). Unbiased learning-to-rank with biased feedback. In Proceedings of the Tenth ACM International Conference on Web Search and Data Mining (pp. 781–789).

19. Winfield, A. F., & Jirotka, M. (2018). Ethical governance is essential to building trust in robotics and artificial intelligence systems. Philosophical Transactions of the Royal Society A, 376(2133), 20180085.

20. Watson, R. T., Boudreau, M. C., & Chen, A. J. (2010). Information systems and environmentally sustainable development: Energy informatics and new directions for the IS community. MIS Quarterly, 34(1), 23–38.