Towards Trust-Augmented Visual Analytics for Data-Driven Energy Modeling

Akshith Reddy Kandakatla*
New Jersey Institute of Technology

Vikas Chandan†
Pacific NorthWest National Lab

Soumya Kundu‡
Pacific NorthWest National Lab

Indrasis Chakraborty§
Lawrence and Livermore National Lab

Kristin Cook¶
Pacific NorthWest National Lab

Aritra Dasgupta‖
New Jersey Institute of Technology

Abstract

The promise of data-driven predictive modeling is being increasingly realized in various science and engineering disciplines, where experts are used to the more conventional, simulation-driven modeling practices. However, trust remains a bottleneck for greater adoption of machine learning-based models for domain experts, who might not be necessarily trained in data science. In this paper, we focus on the building energy domain, where physics-based simulations are being complemented or replaced by machine learning-based methods for forecasting energy supply and demand at various spatio-temporal scales. We study the trust problem in close collaboration with energy scientists and engineers and describe how visual analytics can be leveraged for alleviating this trust bottleneck for stakeholders with varying degrees of expertise and analytical goals in this domain.

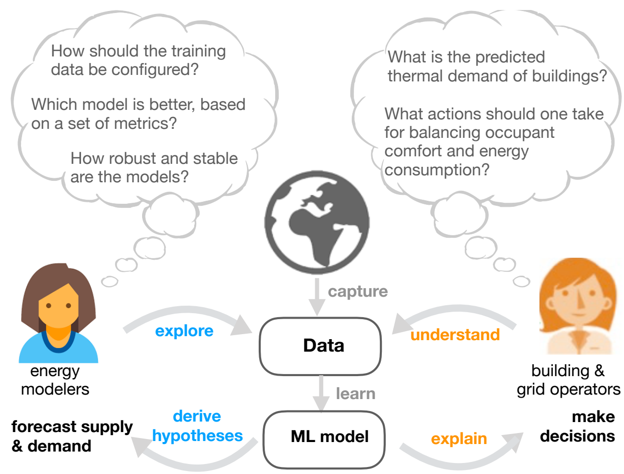

Figure 1: **Illustrating the roles of the different stakeholders for energy modeling** with respect to a data-driven predictive analysis pipeline. A trust gap exists on both sides, where energy modelers need transparent analytical methods for model accessibility and calibration, and building and grid operators need tools that can both explain model predictions and help them make better real-time decisions.

1 Introduction

The field of visual analytics was born out of the need to marry computational capabilities of automated methods with the perceptual and cognitive human faculties in the context of data analysis. As such, visual analytics is defined as the "science of analytical reasoning facilitated by interactive, visual interfaces" [12], implying that a human analyst is actively involved for analyzing the outputs of statistical or machine learning models, and for refining the models implicitly or explicitly using their domain knowledge. The degree and nature of human feedback essentially depend on the goal and background of the user. Recently, Sacha et al. [38] had reflected on the importance of users' trust in shaping data transformation and knowledge generation at various stages of the visual analytic pipeline. In this paper, we analyze how visual analytics can be applied for the purpose of fostering greater trust and reliability in the energy modeling domain for data-driven forecasting and decision-making.

Algorithmic decision-making leveraging machine learning models is becoming increasingly adopted across domains such as and finance [5], healthcare [6,27], defense [16], and crime control [2]. In the realm of energy modeling, while physics based simulations are widely used for forecasting energy demand and consumption, there is a growing recognition of the transformative effect that data-driven models can have in more accurate and proactive forecasting [1]. In Figure 1, we illustrate how different stakeholders in energy modeling can leverage predictive models, based on their goals and expertise. Energy modelers and building or grid operators might not be necessarily trained in data science methods. **Energy modelers** apply their domain knowledge in configuring training data or developing more reliable metrics for calibrating the performance of alternative models. They need to develop appropriate levels of trust in the model predictions for generating and testing alternative hypothesis. This process is usually mediated by a data scientist. On the other hand, **building and grid operators** generally inspect model predictions in a black-box fashion. They need to trust the predictions for ultimately taking real-time actions, for balancing energy demand and consumption patterns in buildings or the grid.

In visual analytics, except for a few studies [10, 14], we lack a thorough characterization of the interplay between user expertise and trust at different stages of the data transformation pipeline. In the energy modeling domain, it is critical to understand the goals and needs of the diverse stakeholders, like energy modelers and building operators, for developing and tailoring the techniques in a way they can elicit a high level of trust in both the data and the technology. To this end, in this paper, we present an initial analysis of the key trust barriers towards the adoption of data-driven predictive model through a collaboration among visual analytics researchers and energy modelers. We present a roadmap about the role of visual analytics for mitigating these trust barriers and reflect on the related interdisciplinary research directions.

*e-mail: ark68@njit.edu
†e-mail: vikas.chandan@pnnl.gov
‡e-mail: soumya.kundu@pnnl.gov
§e-mail: chakraborty3@llnl.gov
¶e-mail: kris.cook@pnnl.gov
‖e-mail: aritra.dasgupta@njit.edu

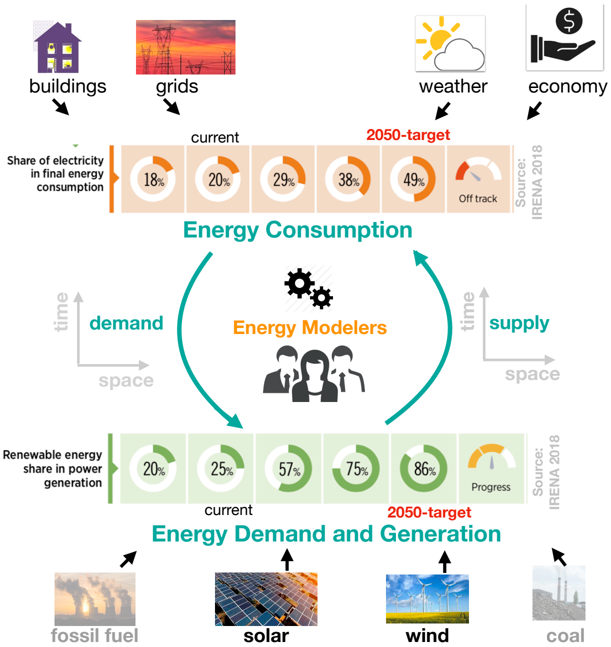

Figure 2: **Energy modeling** entails complex data analysis tasks for generating forecasts for supply and demand at different scales of space and time.

2 Characterizing Data-driven Energy Modeling

Energy generation and consumption are critical indicators of the socio-economic health of a city, state, country, or region. Commercial and residential electrical power consumption drives demand for energy in local and regional power grids, with the US Energy Information Administration report attributing 75% of all US electricity consumption in 2018 to buildings. Social behavior and built environments like smart cities also contribute to the nature of energy consumption. On the other hand, there is a growing push for clean energy generation: "scaling up electricity from renewables will be crucial for the decarbonization of the world's energy system". To achieve the United Nations (U.N.) specified sustainable development goal of achieving affordable and clean energy for all by 2030, such socio-economic dimensions of energy consumption will play a key role. Driving a sustainable, affordable, and clean energy ecosystem of the future presents evolving challenges to the knowledge workers who are tasked with analyzing and predicting ever-changing energy generation and consumption patterns; discovering novel opportunities to accommodate increasing penetration of cleaner energy resources; as well as understanding the potential and impacts of implementing advanced energy efficiency principles. In Figure 2, we illustrate the complexity of the energy modeling process and describe the role of the data and stakeholders in the context of prediction problems below.

Importance of buildings in the energy ecosystem: The building sector accounts for more than 40% of the total energy consumption in the United States [42]. Therefore, improvements in the energy efficiency of buildings have the potential for a large economic and environmental impact. Hence, efficient control of buildings that balances economic, environmental and occupant comfort goals is an important problem [8]. Also, buildings have significant potential to be utilized in grid planning and operation through demand-side management because of them being the primary drivers of electricity demand. Residential and commercial buildings consume 75% of all U.S. electricity and drive 80% of peak demand. Demand-side resources such as energy efficiency measures in buildings not only save energy and reduce costs for the owner or occupant, but also can lower electricity system costs for all customers by reducing energy and capacity needs. Moreover, energy efficiency measures can be combined with grid-interactive strategies such as demand response (DR) and distributed energy storage to further reduce and change electricity consumption to minimize consumer and electricity system costs, relieve system stress, deliver grid operational benefits through ancillary services, and integrate variable renewable energy resources.

Stakeholders: In the last couple of decades, various stakeholders such as building energy managers, utility companies, policy-makers, and researchers have identified, demonstrated, and advocated several data-driven solutions to achieve the goals of energy and cost efficiency. Some of these solutions include demand response (DR) [33], pre-cooling or pre-heating [40], optimal supervisory control [39], energy benchmarking [11] and on-site renewable generation [35]. However, almost all of these solutions require a predictive mathematical model of energy consumption as an essential element. For instance, the design of effective DR schemes for providing grid services requires a quantification of how energy demand in a building changes on a per appliance/load-type basis, and how occupant comfort is tied to the demand. Pre-cooling, pre-heating, and optimal supervisory control of HVAC systems involve optimization of set-point temperatures, which in turn requires an understanding of the relationship between set-point temperatures, energy consumption, and occupant comfort. Energy benchmarking involves identifying anomalous behavior in energy consumption patterns to identify energy wastage, and hence requires a model of nominal energy consumption, a process also known as baselining. Optimal sizing of on-site renewables requires a quantification of the building's demand profile.

A building energy modeler forecasts a building's baseline energy consumption, but also takes extraneous factors such as weather conditions, variability in power generation from renewable energy sources, or the building occupants' preferences and comfort levels, for aligning their model predictions with evolving real-world conditions. Understandably, modeling energy generation or consumption will ultimately not be purely a matter of focusing on technical dimensions but will also be driven by changes in social behavior, economic indicators, and environmental factors like emissions and weather. The glue that binds these diverse dimensions is the availability of data from many sources, such as satellite imagery, building energy consumption data, smart meters, weather sensors, open data sources, or even social media data. Such data is dynamic and reflects relevant changes in socio-environmental factors as conditions change. In an ideal scenario, energy modelers will be able to dynamically update their models based on information extracted from heterogeneous data and re-train and re-deploy them whenever and wherever necessary. The consumer of these prediction models are building or grid operators who need to make proactive or real-time decisions based on their observation and also based on the predictions made by these models.

Data characteristics: The data collected from a building management system can be classified in general [36] as **Controls:** human-provided operating inputs that at any given time are set either using a pre-programmed logic or manually by the building manager. Some examples of such variables are zone set-point temperatures, duct static pressure set-point, and mode of operation of HVAC (heating/cooling). **Exogenous inputs:** inputs such as outside weather and zone occupancy status, which drive the energy consumption in a building, but cannot be controlled. **Internal variables:** operating conditions in the building, that are a consequence of the choice of controls. This includes internal variables such as zone temperatures, airflow rates in HVAC systems, return and supply temperatures of water and air, etc. **Consumption:** measurements indicative of performance such as heating and cooling demand, active and reactive

power consumption, etc.
Prediction Problems: The modeling problem that experts typically consider involves predicting the total building HVAC thermal demand (falling under category consumption) as well as zone temperatures (falling under category internal variables) as a function of controls, and exogenous inputs. More recently, the evolution of smart buildings and cities, as well as distributed renewable generation resources, is driving a paradigmatic shift in the operation of the future energy ecosystem - requiring a tighter collaboration between the different energy infrastructures. With building energy efficiency being one of the lowest-cost energy resources, power grid utilities and operators are increasingly engaging smart buildings and communities in energy efficiency programs, with investments on the order of $6B/year [31]. Tighter coupling between the various energy infrastructures will be driving innovations in the nature of the energy modeling work, such as understanding spatio-temporal energy flexibility offered by a community of smart buildings. This might be better understood by describing a use case that leverages data driven models to minimize the cost of energy for a building. This use case involves optimizing the set-point temperature of the building to achieve this objective, since set-point temperatures have a significant impact on building energy costs. There are two specific ways in which set-point optimization helps minimize energy costs. 1. Reduction in energy consumption by utilizing the set-point flexibility provided via building code (reduce set-point in heating mode, increase in cooling mode), 2. Time-shifting of demand by utilizing building thermal capacity (pre-cool or pre-heat when energy is cheaper).

3 Analyzing the Trust Bottleneck

A persistent challenge in building energy modeling is the dependency of energy consumption on a large number of factors such as set-point temperatures, weather, occupant behavior, underlying control systems, building layout and equipment efficiencies. A classic approach to solving this problem is to retrieve insights from physics-based models that account for the thermal characteristics of the buildings and system dynamics. The EnergyPlus tool [13] is a representative for physics based modeling and has been used widely. However, the main limitation of these models as that they involve solving equations with physical parameters that are specific to the buildings at hand and therefore there is usually a great effort spent on gathering sufficient information about the physical features of the environment (building materials, heat transfer constants, boundary conditions, etc.), which can be time consuming and expensive to obtain. In this context, the availability of fine-grained spatio-temporal data by the rapid penetration of information technology (IT) in today's buildings [30] provides an opportunity to address the challenges associated with physical models. Therefore, data driven modeling paradigms such as linear regression, support vector regression, and artificial neural networks and deep learning are finding application in building energy modeling use cases [36, 42]

However, despite the presence of several publications in recent years that have set up the theoretical foundations of data driven modeling and advanced optimization based control algorithms as well as demonstrated their value on simulation test-beds and experimental systems, the state of the practice in this domain is still lagging behind. Building operators still use conventional rule based methodologies to operate them. This is despite the rapid penetration of Building Management Systems (BMS) which automate the operation of buildings. BMS providers such as GE, Siemens, Honeywell and Johnson Controls encapsulate these rules in their BMS offerings and provide an automated way of implementing them via their proprietary hardware and software platforms. These rules are often a sequence of semantic operations such as "if-then-else" constructs.

A key reason for this difference between the state of the art and state of practice in this domain is that operators of buildings as well as the providers of the BMS solutions are generally unwilling to completely trust, what appears to them a 'black-box' decision making unit that is difficult to be explained in terms of the operating framework that they are used to. For example, the output of an advanced Predictive controller can be a sequence of valve opening positions which is the solution of a dynamic optimization problem. However an operator on the ground might have been trained to operate these valves based on some situational awareness (e.g. outside temperature, indoor temperature, occupant complaints, safety needs etc.). This trust issue is further supported by findings from the survey [17], which suggest that operators and managers of commercial buildings view technology as a favorable driver to enable and support their sustainability missions. However, building operators might be resistant to deploying a technology which is opaque and whose working mechanism is not clear to them. This becomes especially important in the context of the 3-30-300 rule which specifies that the highest priority should be given to mitigating occupant complaints. The survey also provides some strategies to address the trust barrier. The most important recommendation is the need to include building operators as partners for technology penetration, by including their inputs in all phases such as design, testing and commissioning.

4 How Visual Analytics is Currently Used

In this section, we describe some of the existing visualization approaches for addressing energy modeling and model consumption tasks.
Energy consumption in buildings: Human activities are one of the primary drivers of power consumption in buildings. A comprehensive collection and analysis of power consumption is required to identify situations where residents should change their habits [9]. Commercial buildings in the US consume 19% of the electricity and the current studies have limitations in that, they are exclusively focused on households, and only very few focus on the workplaces, where energy dashboards can be used for monitoring consumption [41].
Ecological and behavior change: Several researchers have focused on ecologically responsible behavior [19, 25]. The focus areas can be categorized as follows: Developing persuasive displays: Citing that none of the existing Ambient persuasive displays seem to have any long- term persuasive effects, one research work argues that this effect could only be enabled by involving providers and potential users in the design process of an ambient persuasive display [24]. Consumption awareness: Since behaviors often get transformed to habit, researchers have studied how energy consumption gets affected by environmental consequences of human activity [23].
HVAC Systems: 80 to 90% of the yearly growth of industrial energy consumption can be attributed to Heating ventilation and Air-conditioning systems [34]. In US alone HVAC systems account 20% of global and 70% of the nation's electricity. There is a need to understand energy consumption patterns to achieve energy efficiency. To this end, there has been some emphasis for the need to combine data mining & visualization techniques in order to convey the information without cluttering the dashboard which could only be interpretable by expert user group [37].
Building energy management: It is not an easy task to determine whether the past energy consumption was really necessary or just wasted [20]. Drawbacks in the current visualizations that either display highly technical information familiar to expert users like building engineers for tuning performance or simplified displays of aggregated energy consumption over time for non-experts like the one on the electric bill and combining these creates confusion [3]. Energy performance of large building portfolios is challenging to analyze and monitor, as current analysis tools are not scalable or they present derived and aggregated data at too coarse of a level [4]. Building Information Management(BIM) domain still remains largely unexplored by the visualization community [21].

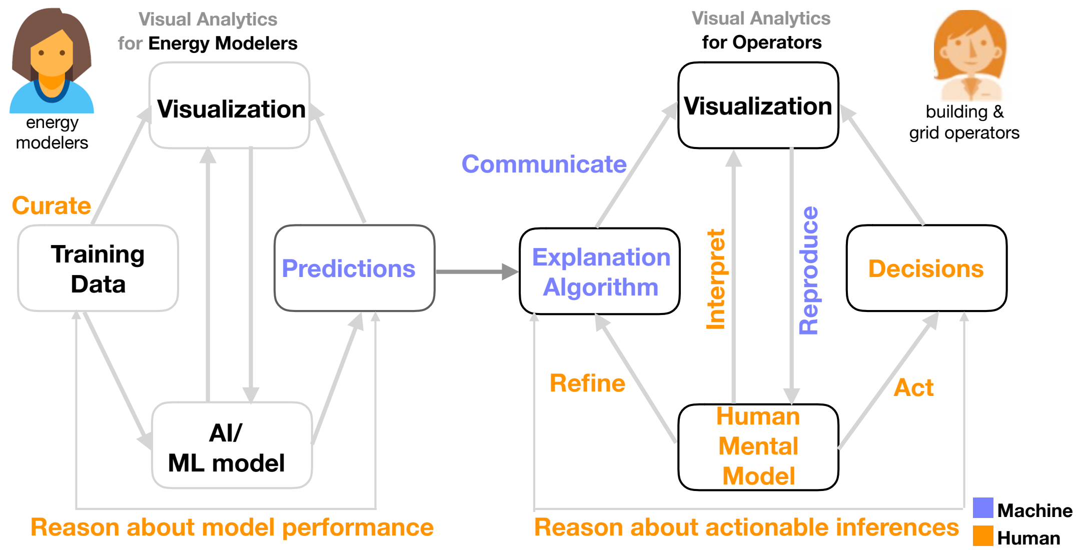

Figure 3: **Visual analytic pipelines catering to the needs of energy modelers and operators.** While energy modelers need to be involved in participatory model development stages for increasing their trust in the data and the performance indicators of the models; operators need better explanation methods and interactive interfaces for trustworthy decision-making based on the predictions. A blue color indicates machine-level tasks while an orange color indicates human-centered tasks that can can be achieved using visualization techniques.

5 Trust-Augmented Visual Analytics

In this section, we speculate on how visual analytics (Figure 3) can help energy modelers and operators increase their trust in the data, in the predictive models, and ultimately, in the inferences and decisions derived from the data-driven predictions.

Trust in the data and the prediction generation process: For energy modelers, trust in the training data is a prerequisite for trusting any of the downstream modeling steps. A visual analytic method should be able to capture this process of data generation and subsequent expert ideation and sensemaking that can inform model development processes. The challenge here is in handling heterogeneous and multi-scale data sets that can come from disparate sources, like satellite imagery, public building records, sensors, etc. This presents a key opportunity for visual analytic techniques to semantically integrate high-dimensional data [26, 29] from the different sources for faster and trustworthy calibration and prediction of building energy consumption patterns. The data here is of very high dimensionality, where each data source comprises tens or hundreds of features, many of which might not be relevant for analysis. This challenge can be addressed by developing interactive tool that ingests data from heterogeneous sources, combines statistical methods with visualization for presenting interesting correlations and clusters, and lets modelers *semantically integrate diverse pieces of information* for understanding the causal factors behind increasing and decreasing energy building energy consumption. This will further encourage participatory model development [18], where energy modelers and data scientists can engage in collaborative sensemaking using visual analytics interfaces.

Trust in the model calibration metrics: Accurately modeling an energy system is crucial for long term and short term energy prediction, along with anomaly or adversarial activity detection within the energy system. In many such cases, traditional accuracy metrics used to capture accuracy of machine learning models are insufficient for characterizing how well these models can capture expert knowledge [7]. Chakraborty et al. proposed a set of metrics relevant to short term and long term energy prediction [7], applied to machine learning for building energy consumption. Although these metrics are available in the literature, it is hard to get an unified message from those metrics regarding the performance of the energy model. This asks for inputs from experts with domain knowledge, such as energy modelers, and that is why *interactive visual comparison of various aspects of model performance* can help experts build trust in model performance and at the end, select the models that best reflect their mental model about energy demand and consumption.

Trustworthy communication in cyber-physical systems: With the recent advances in data-driven modeling in cyber-physical systems, it is being increasingly recognized that information sharing across systems and human operators can drastically improve network-wide security and trustworthiness of the decision-making processes by operators. There are challenges visual analytics techniques need to address for facilitating such *inter-operator and inter-system communication*. First, the scale and complexity of the individual systems makes it extremely challenging for operators to monitor and operate such coupled infrastructures. Second and more importantly, high-fidelity models of the component physical systems, though accurate, are often not intelligible enough for cross-domain operators. This impedes the transfer of knowledge across operators which is essential for understanding the dependencies across systems and proactively planning for contingencies. For example, let us consider a connected buildings-grid-cyber infrastructure, where buildings are used as flexible energy resources to improve grid resiliency.

Visualization techniques can ensure that the information exchange between the operators of buildings and grid is mutually intelligible and ensure secure and optimal operation of the connected infrastructure. While building models can be developed to be of arbitrarily high accuracy using data-driven approaches (e.g., machine learning), such models carry little value in grid operations. To address these problems, we posit that the expressive power of advanced, high-dimensional data visualization techniques will be effective in exposing the hidden dependencies across model outcomes from disparate domains. These visualizations can be presented to both grid and building operators, through interactive user interfaces. Using their domain knowledge, operators will incorporate their feedback about the observed patterns and correlations and help separate the signal (meaningful correlations, groupings) from the noise (e.g., spurious correlations). This will result in i) a simplified and intelligible model for operators across both domains, leading to greater trans-

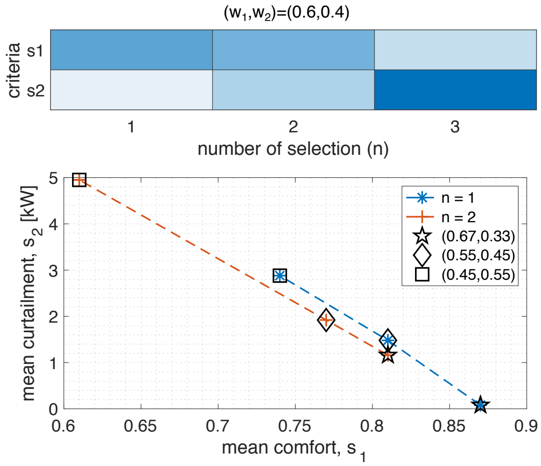

Figure 4: Impact of the decision choices of a building operator on multiple criteria (comfort and curtailment) performances (adopted from [28]. Top plot illustrates how the building operator's choice of number (n) of devices to commit for grid service affects the control performance across two different criteria - *end-user comfort* (s_1) and *amount of curtailment* (s_2) with the respective criteria weights w_1=0.6 and w_2=0.4. Bottom plot illustrates the *Pareto front* of a multi-criteria decision problem when the criteria weights (w_1,w_2) are varied, for two different choices of the number (n) of committed devices.

parency in knowledge-sharing and efficiency in operating the grid, and ii) a simplified functional relationship that represents a physical model that can then be leveraged for down the line optimization framework requirement.

Trust in inferences from black-box predictions: Outcomes from predictive models have to be ultimately communicated to building and grid operators who are traditionally used to rule-based models. The challenge here is to develop techniques that help explain model decisions and enable integration of expert feedback so that operators can generate actionable and trustworthy inferences from the model predictions. An example of an inferential reasoning scenario is the use of multi-criteria decision-making [22] to address trade-offs among energy consumption minimization, occupant comfort maximization, satisfying the grid services to the maximum possible extent, and maximizing system resiliency and reliability . Comparative visualization techniques [15] can be used where visualization will help operators take real-time decisions directly affecting the functioning of the grid. In a recent work [28], the authors investigated a multi-criteria decision problem for grid-interactive efficient buildings in which the building operator monitors the streaming data building sensors to decide which (and how many) devices to commit for a load curtailment grid service, while minimizing any adverse impact on the building occupants. A stochastic multi-criteria decision algorithm was used to assign priority ranks to each device based on the (anticipated) impact on end-user comfort and the resulting amount of load curtailment - measured via respective scores s_1 and s_2. Fig. 4 shows how the different design choices - such as the weights placed on each criteria, w_1 (for *comfort*) and w_2 (for curtailment), as well as the number (n) of devices to commit to grid service - impact the building performance measured via the different criteria metrics. The fundamental trade-off between the different control performance objectives make the real-time decision making a challenging problem for the building operators. Adopting and extending visual analytic approaches for multi-criteria decision-making [15, 32] can help synthesize explainable inferences from the streaming sensors data and enhance operator's trust in the real-time decision support systems.

6 Conclusion

In this paper, we have presented an analysis of the trust gaps that exist in the energy modeling when different stakeholders, like modelers and operators, want to leverage the data-driven predictions for accurate forecasting and decision-making. The problems and potential solutions we identified result from a collaboration among energy modelers and visual analytic researcher and can serve as a guideline towards embracing visual analytics-based technological advances in the energy sector, which is increasingly becoming the focus of national and international investments for attaining sustainable development goals.

References

[1] T. Ahmad, H. Chen, Y. Guo, and J. Wang. A comprehensive overview on the data driven and large scale based approaches for forecasting of building energy demand: A review. *Energy and Buildings*, 165:301–320, 2018.

[2] J. Angwin, J. Larson, S. Mattu, and L. Kirchner. Machine bias. *ProPublica, May*, 23, 2016.

[3] L. Bartram, J. Rodgers, and K. Muise. Chasing the negawatt: visualization for sustainable living. *IEEE Computer Graphics and Applications*, 30(3):8–14, 2010.

[4] M. Brehmer, J. Ng, K. Tate, and T. Munzner. Matches, mismatches, and methods: Multiple-view workflows for energy portfolio analysis. *IEEE transactions on visualization and computer graphics*, 22(1):449–458, 2015.

[5] C. Carter, E. Renuart, M. Saunders, and C. C. Wu. The credit card market and regulation: In need of repair. *NC Banking Inst.*, 10:23, 2006.

[6] R. Caruana, Y. Lou, J. Gehrke, P. Koch, M. Sturm, and N. Elhadad. Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1721–1730. ACM, 2015.

[7] I. Chakraborty, V. Chandan, and D. Vrabie. A sequential DNN based Baseline Energy Prediction Framework with Long term Error Mitigation. In *Proceedings of the Tenth ACM International Conference on Future Energy Systems*, pp. 508–515, 2019.

[8] V. Chandan and A. G. Alleyne. Decentralized predictive thermal control for buildings. *Journal of Process Control*, 24(6):820–835, 2014.

[9] C.-C. Chou, C.-T. Chiang, P.-Y. Wu, C.-P. Chu, and C.-Y. Lin. Spatiotemporal analysis and visualization of power consumption data integrated with building information models for energy savings. *Resources, Conservation and Recycling*, 123:219–229, 2017.

[10] J. Chuang, D. Ramage, C. Manning, and J. Heer. Interpretation and trust: Designing model-driven visualizations for text analysis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 443–452, 2012.

[11] W. Chung, Y. Hui, and Y. M. Lam. Benchmarking the energy efficiency of commercial buildings. *Applied Energy*, 83(1):1–14, 2006.

[12] K. A. Cook and J. J. Thomas. Illuminating the path: The research and development agenda for visual analytics. 2005.

[13] D. B. Crawley, L. K. Lawrie, F. C. Winkelmann, W. F. Buhl, Y. J. Huang, C. O. Pedersen, R. K. Strand, R. J. Liesen, D. E. Fisher, M. J. Witte, et al. Energyplus: creating a new-generation building energy simulation program. *Energy and buildings*, 33(4):319–331, 2001.

[14] A. Dasgupta, J.-Y. Lee, R. Wilson, R. A. Lafrance, N. Cramer, K. Cook, and S. Payne. Familiarity vs trust: A comparative study of domain scientists' trust in visual analytics and conventional analysis methods. *IEEE transactions on visualization and computer graphics*, 23(1):271–280, 2016.

[15] A. Dasgupta, H. Wang, N. O'Brien, and S. Burrows. Separating the wheat from the chaff: Comparative visual cues for transparent diagnostics of competing models. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):1043–1053, 2020.

[16] A. A. Freitas. Comprehensible classification models: a position paper. *ACM SIGKDD explorations newsletter*, 15(1):1–10, 2014.

[17] J. Goins and M. Moezzi. Case study of calstrs headquarters. 2011.

[18] A. Halfaker and R. S. Geiger. Ores: Lowering barriers with participatory machine learning in wikipedia. *arXiv preprint arXiv:1909.05189*, 2019.

[19] T. G. Holmes. Eco-visualization: combining art and technology to reduce energy consumption. In *Proceedings of the 6th ACM SIGCHI conference on Creativity & cognition*, pp. 153–162, 2007.

[20] T. Itoh, M. Kawano, S. Kutsuna, and T. Watanabe. A visualization tool for building energy management system. In *2015 19th International Conference on Information Visualisation*, pp. 15–20. IEEE, 2015.

[21] P. Ivson, A. Moreira, F. Queiroz, W. Santos, and W. Celes. A systematic review of visualization in building information modeling. *IEEE transactions on visualization and computer graphics*, 2019.

[22] S. A. A. Kazmi, S. F. Hasan, and D.-R. Shin. Multi criteria decision analysis for optimum DG placement in smart grids. In *2015 IEEE Innovative Smart Grid Technologies-Asia (ISGT ASIA)*, pp. 1–5. IEEE, 2015.

[23] J.-W. Kim, Y.-K. Kim, and T.-J. Nam. The ténéré: design for supporting energy conservation behaviors. In *CHI'09 Extended Abstracts on Human Factors in Computing Systems*, pp. 2643–2646. 2009.

[24] P. M. Kluckner, A. Weiss, P. Sundström, and M. Tscheligi. Two actors: Providers and consumers inform the design of an ambient energy saving display with persuasive strategies. In *PERSUASIVE (Adjunct Proceedings)*, 2013.

[25] C. E. Kontokosta and C. Tull. Energyviz: Web-based eco-visualization of urban energy use from building benchmarking data. In *Proceedings of the International Conference on Computing in Civil and Building Engineering*, 2016.

[26] J. Krause, A. Dasgupta, J.-D. Fekete, and E. Bertini. Seekaview: An intelligent dimensionality reduction strategy for navigating high-dimensional data spaces. In *2016 IEEE 6th Symposium on Large Data Analysis and Visualization (LDAV)*, pp. 11–19. IEEE, 2016.

[27] J. Krause, A. Dasgupta, J. Swartz, Y. Aphinyanaphongs, and E. Bertini. A workflow for visual diagnostics of binary classifiers using instance-level explanations. *IEEE Conference on Visual Analytics Science and Technology*, 2017.

[28] S. Kundu, A. Bhattacharya, V. Chandan, N. Radhakrishnan, V. Adetola, and D. Vrabie. A Stochastic Multi-Criteria Decision Making Algorithm for Dynamic Load Prioritization in Grid-Interactive Efficient Buildings. In *2020 ASME Dynamic Systems and Control Conference (DSCC)*. ASME, 2020 (accepted).

[29] S. Liu, D. Maljovec, B. Wang, P.-T. Bremer, and V. Pascucci. Visualizing high-dimensional data: Advances in the past decade. *IEEE transactions on visualization and computer graphics*, 23(3):1249–1268, 2016.

[30] P. A. Mathew, L. N. Dunn, M. D. Sohn, A. Mercado, C. Custudio, and T. Walter. Big-data for building energy performance: Lessons from assembling a very large national database of building energy use. *Applied Energy*, 140:85–93, 2015.

[31] M. Neukomm, V. Nubbe, and R. Fares. Grid-interactive efficient buildings technical report series: Overview of research challenges and gaps. Technical report, U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy, 2019.

[32] S. Pajer, M. Streit, T. Torsney-Weir, F. Spechtenhauser, T. Möller, and H. Piringer. Weightlifter: Visual weight space exploration for multi-criteria decision making. *IEEE transactions on visualization and computer graphics*, 23(1):611–620, 2016.

[33] P. Palensky and D. Dietrich. Demand side management: Demand response, intelligent energy systems, and smart loads. *IEEE transactions on industrial informatics*, 7(3):381–388, 2011.

[34] C. Ren and S.-J. Cao. Implementation and visualization of artificial intelligent ventilation control system using fast prediction models and limited monitoring data. *Sustainable Cities and Society*, 52:101860, 2020.

[35] S. Ruangpattana, D. Klabjan, J. Arinez, and S. Biller. Optimization of on-site renewable energy generation for industrial sites. In *Power Systems Conference and Exposition (PSCE), 2011 IEEE/PES*, pp. 1–8. IEEE, 2011.

[36] J. Rubio-Herrero, V. Chandan, C. Siegel, A. Vishnu, and D. Vrabie. A learning framework for control-oriented modeling of buildings. In *2017 16th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pp. 473–478. IEEE, 2017.

[37] L. Ruiz, M. Pegalajar, M. Molina-Solana, and Y.-K. Guo. A case study on understanding energy consumption through prediction and visualization (vimoen). *Journal of Building Engineering*, p. 101315, 2020.

[38] D. Sacha, H. Senaratne, B. C. Kwon, G. Ellis, and D. A. Keim. The role of uncertainty, awareness, and trust in visual analytics. *IEEE transactions on visualization and computer graphics*, 22(1):240–249, 2015.

[39] S. Wang and Z. Ma. Supervisory and optimal control of building HVAC systems: A review. *HVAC&R Research*, 14(1):3–32, 2008.

[40] P. Xu, P. Haves, M. A. Piette, and J. Braun. Peak demand reduction from pre-cooling with zone temperature reset in an office building. *Lawrence Berkeley National Laboratory*, 2004.

[41] R. Yun. Persistent workplace plug-load energy savings and awareness through energy dashboards: eco-feedback, control, and automation. In *CHI'14 Extended Abstracts on Human Factors in Computing Systems*, pp. 331–334. 2014.

[42] H.-X. Zhao and F. Magoulès. A review on the prediction of building energy consumption. *Renewable and Sustainable Energy Reviews*, 16(6):3586–3592, 2012.