DATA–DRIVEN MODELING OF DYNAMIC OCCUPANT THERMOSTAT OVERRIDE BEHAVIOR FOR DEMAND RESPONSE APPLICATIONS

A Dissertation Presented

By

Kunind Sharma

To

The Department of Civil & Environmental Engineering

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the field of

Interdisciplinary Engineering

Northeastern University

Boston, Massachusetts

April 2025

ABSTRACT

Buildings consume nearly 40% of global energy and produce similar emissions. While technological advances address efficiency, occupant behavior causes energy use variations up to 300% between identical buildings. This gap between predicted and actual building performance impacts building design, operations, and grid demand management programs.

Through analyses of smart thermostat data from 1,400 single–occupant homes, the research demonstrates that occupants respond to 8°F thermostat setpoint changes within a median of 15 minutes, while 2°F changes trigger responses within a median of 30 minutes. This highlights an understudied temporal relationship between thermostat setbacks and response time of occupant behaviors. Models of such behavior dynamics are required to incorporate occupant impacts into building performance simulation.

A key contribution of this dissertation is the Thermal Frustration Theory (TFT), which posits that thermal discomfort driven behaviors are caused by the time-accumulation of discomfort, not simply a temperature deviation threshold or a delay from an initiating event. Using a dataset of 634 thermostats, each with 25+ manual setpoint changes, a comparative analysis of TFT and comfort zone and a delayed response theories demonstrated that personalized TFT models better predict when manual setpoint change occur. This was measured by the area under the curve statistical measure (AUC); all three models perform similarly by a Matthews Correlation Coefficient measure. Higher AUC performance is especially important for modeling occupant behavior in demand response programs where false negatives of rare occupant interactions could adversely affect grid stability. EnergyPlus based simulations were conducted with TFT-derived occupant models, demonstrating the ability to identify parameters of known TFT models from only data observable with smart thermostats, even under the presence of noise from routine overrides.

Overall, the dissertation highlights that thermostat interactions are neither static, instantaneous, nor driven solely by the environment. Instead, temporal accumulation of discomfort and routine-based behavior play important roles. The methodology and results offer a pathway towards more accurate modeling of human-building interactions for policy assessment, building design, and demand response programs.

ACKNOWLEDGMENTS

I would like to express my sincere gratitude to my advisor, Dr. Michael Kane, for his exceptional guidance, unwavering support, and mentorship throughout this research journey. His intellectual insights shaped this work significantly, and his encouragement extended beyond academic boundaries, helping me navigate both research challenges and professional development.

I am deeply thankful to my dissertation committee members: Dr. Misha Pavel for his invaluable perspectives on human behavior and cognitive processes; Dr. David Fannon for his expertise in thermal comfort and occupant–centric building design; and Dr. Ryan Qi Wang for his guidance in data science methodologies. Their interdisciplinary insights and thoughtful feedback enriched this work substantially, and their mentorship extended well beyond their respective domains of expertise.

I would also like to acknowledge my colleagues who provided intellectual companionship and support throughout this process. Special thanks to Dr. Craig Poulson, Dr. Krissy Govertson, Dr. SungKu Kang, and Maharshi Pathak for the countless discussions, collaborative problem–solving sessions, and camaraderie that made this journey both intellectually stimulating and personally rewarding.

Finally, my deepest appreciation goes to my family, the foundation upon which all my endeavors rest. To my father, Dr. Prem Sagar Sharma, and my mother, Vijay Bharti Sharma, thank you for instilling in me the values of intellectual curiosity and perseverance. To my sister, Mrigakshi Sharma, and brother Vikas Sharma, thank you for your constant encouragement and unwavering belief in me. Additionally, I joyfully welcome my niece, Ezhil Sharma, whose arrival two months ago brought renewed inspiration and happiness to our family. Your collective support has been instrumental in every step of my academic and personal development, and this achievement is as much yours as it is mine.

The U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy, Building Technologies Office (Award DE–EE0009154) and National Science Foundation (Grant 1928608) funded this work. The research was conducted within IEA–EBC Annex 79.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABBREVIATIONS

ABM	Agent–Based Model
API	Application Programming Interface
ASHRAE	American Society of Heating, Refrigerating and Air–Conditioning Engineers
AUC	Area Under the Curve
BEM	Building Energy Modeling
BIM	Building Information Modeling
BPS	Building Performance Simulation
CO2	Carbon Dioxide
CSV	Comma Separated Values
CZT	Comfort Zone Theory
DER	Distributed Energy Resource
DNAS	Drivers–Needs–Actions–Systems
DoO	Degree of Override
DR	Demand Response
DRP	Demand Response Program
DRT	Delayed Response Theory
DyD	Donate Your Data
EMS	Energy Management System
FERC	Federal Energy Regulatory Commission
GEB	Grid–interactive Efficient Building
HVAC	Heating, Ventilation, and Air Conditioning
IoT	Internet of Things
ISO	Independent System Operator
LR	Logistic Regression
MCC	Matthews Correlation Coefficient
ML	Machine Learning
MSC	Manual Setpoint Change
NERC	North American Electric Reliability Corporation
PIR	Passive Infrared

PMF	Probability Mass Function
PMV	Predicted Mean Vote
PSC	Programmed Setpoint Change
RTO	Regional Transmission Organization
SC	Setpoint Change
SES	Smart Energy Solutions
TFT	Thermal Frustration Theory
TTO	Time to Override
TUS	Time Use Survey
XML	Extensible Markup Language

Chapter 1:

Grand challenges for simulation of grid–interactive efficient buildings

1.1 Background and motivation

The buildings accounts for approximately 40% of worldwide energy consumption and greenhouse gas emissions [1], [2]. The magnitude of building energy use, combined with projected increases in global building stock of 2.5 trillion square feet by 2060 [3], requires fundamental changes in building energy management [4], [5] . Building–level energy reduction strategies directly impact grid stability, energy resource allocation, and emissions reduction targets established by international climate agreements [6], [7], [8].

Historically, building energy efficiency improvements centered on HVAC systems, insulation, and building management systems [9], [10], [11]. Research demonstrates that occupant behavior significantly impacts building energy consumption [12], [13], [14]. Analysis of identical buildings reveals energy use variations of up to 300% due to occupant behavior differences [15], [16]. This quantification establishes occupant behavior as a primary factor in building energy use, shifting building science from purely technical approaches toward integration of human factors [17], [18], [19].

The proliferation of smart building technologies and Internet of Things (IoT) devices enables collection of occupant behavior and building performance data at previously unattainable resolutions [20], [21], [22]. Smart thermostats, occupancy sensors, and energy monitoring systems generate time–series data characterizing occupant–building interactions [23], [24], [25]. The volume and temporal resolution of this data requires development of new analytical methods to execute computations efficiently [26], [27].

This introductory chapter sets the stage for the dissertation by addressing the complex challenges in building energy management and occupant behavior modeling. It begins with an exploration of Grid–interactive Efficient Buildings (GEBs), discussing their concept, importance, and implementation challenges. The chapter then examines current issues in building performance simulation, focusing on steady state comfort model limitations and occupant behavior uncertainty. A comprehensive overview of occupant behavior modeling follows, covering various models and energy management. Finally, it outlines the dissertation's core research objectives: identifying dynamic thermal comfort behavior, comparing thermal comfort theories, and developing a human–centric digital twin framework.

1.2 Building Performance Simulation (BPS)

Building Performance Simulation (BPS) has evolved from simple steady–state degree days calculations to sophisticated computational frameworks that capture the complex dynamics of building energy systems. This evolution reflects both advances in computational capabilities and the increasing demands from grid electrification, electric vehicles, etc. placed on building performance prediction [28]. Modern BPS tools integrate multiple physical domains – thermal, fluid, and control systems – to predict building energy consumption, occupant comfort, and grid interactions [29]. The application of these tools spans scales from individual component analysis to urban energy planning [30], enabling evaluation of novel technologies and operational strategies for reducing building energy consumption [31]. However, significant challenges remain in simulation accuracy, particularly regarding occupant behavior and grid integration [13]. These challenges become increasingly critical as buildings transition from passive energy consumers to active participants in grid operations [32].

1.2.1 White box models for building design

Current building performance simulation relies heavily on white box models that implement physical principles through numerical methods. Unlike data–driven approaches, these models solve coupled differential equations representing conservation of mass, energy, and momentum within building systems [33]. The theoretical foundation enables analysis of novel designs and operational strategies without requiring historical performance data [34]. This capability proves essential for evaluating emerging technologies and control approaches needed for grid–interactive buildings.

White box modeling approaches have evolved significantly in both mathematical sophistication and computational implementation. Early tools employed simplified steady–

state calculations with basic assumptions about thermal processes. Modern simulation engines now capture complex dynamic interactions between building envelope, mechanical systems, and environmental conditions [35]. These advancements enable analysis of phenomena critical to building–grid integration, including thermal mass effects, system response times, and control interactions [36].

Different simulation engines implement distinct approaches with trade–offs to solve these physical equations. EnergyPlus utilizes an iterative solution methodology that simultaneously resolves zone conditions and HVAC system responses through successive substitution iteration [37]. This approach improves solution stability compared to sequential methods but increases computational requirements. TRNSYS employs a modular architecture enabling detailed component–level modeling through separate algebraic and differential equation solvers [38]. ESP–r implements finite volume methods for resolving spatial temperature gradients within building elements, offering enhanced accuracy for analyzing thermal bridges and non–uniform conditions [39]. Each approach presents tradeoffs between computational efficiency, numerical stability, and physical accuracy.

The implementation of white box models requires extensive input data describing building geometry, material properties, system characteristics, and operational parameters. This comprehensive data requirement enables detailed analysis but introduces significant challenges in model calibration and validation especially for existing buildings [40]. Material properties must account for temperature dependence, aging effects, and installation variations. System performance curves need to capture part–load behavior, control responses, and equipment degradation. Weather data must represent both typical and extreme conditions depending on the goal of the analysis. These input uncertainties propagate through simulations, affecting prediction accuracy particularly during grid–responsive operation modes [41].

1.2.2 Performance standards compliance

Performance standards and building energy codes establish frameworks for evaluating building design and operation. These standards have evolved from simple prescriptive requirements to sophisticated performance–based approaches that leverage BPS capabilities [42]. The transition enables innovative design solutions while ensuring minimum energy performance levels through standardized evaluation methods [43]. This evolution parallels the increasing complexity of building systems and the growing importance of operational optimization for grid integration.

ASHRAE Standard 90.1, a cornerstone of building energy codes, exemplifies modern performance–based compliance approaches. The standard's Energy Cost Budget and Performance Rating Methods require detailed energy modeling to demonstrate that proposed designs achieve equivalent or superior performance compared to baseline buildings [44]. These methods mandate specific modeling protocols, including standardized occupancy schedules, consistent baseline definitions, and verified calculation procedures [45]. The requirements ensure reproducibility while acknowledging the inherent uncertainties in building performance prediction.

International standards like ISO 52016 complement national codes by establishing consistent methods for calculating building energy needs. These standards define calculation procedures for heating, cooling, and lighting energy use, providing a framework for comparing building performance across different jurisdictions [46]. The harmonization of calculation methods supports technology transfer and market transformation while enabling consistent evaluation of building strategies [47].

The implementation of performance–based compliance presents significant challenges for building simulation. Models must balance accuracy with practical constraints on data availability and computational resources. Simulation tools require certification through standard method tests that verify calculation accuracy across diverse building types and operational scenarios [48]. These tests evaluate the tool's capability to model complex phenomena including thermal mass effects, solar gains, and system interactions that influence both energy consumption and grid integration potential [17].

Recent developments in performance standards increasingly address grid interaction capabilities. Requirements for demand response readiness, energy storage systems, and renewable energy integration necessitate enhanced simulation approaches [47]. Traditional compliance metrics focused on annual energy consumption must expand to capture temporal variations in building loads and grid conditions. This evolution creates new challenges for simulation tools while highlighting the importance of accurate performance prediction for grid–interactive buildings [49].

1.2.3 Urban scale building performance simulation

Urban–scale building performance simulation represents a significant expansion beyond individual building analysis, addressing the growing need for city–wide energy planning and carbon reduction strategies. This scale of analysis introduces new computational challenges and methodological requirements while enabling evaluation of district–level energy systems and policy interventions [50]. The LA100 study demonstrates this capability through comprehensive analysis of building electrification pathways and grid integration strategies across Los Angeles's diverse building stock [51].

Traditional building–level simulation approaches become computationally intractable when applied directly to urban scales. This limitation has driven development of novel modeling methods that balance computational efficiency and modeling effort with prediction accuracy. Archetypal building representations reduce model complexity by grouping similar buildings into representative categories**, while statistical methods capture variations within these groups** [52]. These approaches enable analysis of 10s of thousands to millions of buildings while maintaining sufficient detail to inform policy decisions and infrastructure planning [53].

Urban–scale simulation must address additional phenomena beyond individual building physics. District heating and cooling systems introduce thermal network dynamics. Local climate effects, including urban heat islands and building–to–building shading, influence energy consumption patterns. Distributed energy resources create new patterns of energy flow and storage [54]. These interactions require integration of multiple modeling domains like building energy simulation, district thermal networks, power distribution systems, transportation infrastructure, and urban microclimate effects [55].

The Los Angeles 100% Renewable Energy Study exemplifies current capabilities in urban–scale simulation. This analysis evaluated pathways for achieving carbon–neutral buildings across the city through building electrification scenarios, renewable energy integration, grid infrastructure requirements, energy storage deployment, demand response potential [56]. The study revealed critical interdependencies between building performance, grid capacity, and renewable energy integration. Results demonstrate how urban–scale simulation can inform policy decisions by quantifying impacts across multiple sectors and timeframes [51]. However, the study also highlighted current limitations in representing occupant behavi[57]or and building–grid interactions at urban scales [58].

Recent developments in urban building energy modeling (UBEM) frameworks address these challenges through improved data integration and computational methods. These frameworks leverage geographic information systems, building stock databases, and climate projections to create more comprehensive urban energy models [59]. Machine learning techniques increasingly supplement physics–based approaches, particularly for predicting occupant behavior patterns and energy use trends across building populations [60].

1.2.4 Uncertainty in BPS

Building performance simulation faces fundamental challenges in prediction accuracy that stem from multiple sources of uncertainty. Studies comparing predicted and measured building energy use reveal discrepancies ranging from 10% to 80%, undermining confidence in simulation results for critical applications like grid integration and demand response [61]. These variations arise not from limitations in physical models alone, but from complex interactions between model assumptions, input parameters, and human behavior [62].

Physical parameter uncertainties represent the most straightforward category to quantify. Material properties exhibit temperature dependence and aging effects that standard databases may not capture. Equipment performance varies with operating conditions and maintenance status. Weather data, typically based on historical records or typical meteorological years, may not adequately represent future conditions under climate change [63]. While these uncertainties can be bounded through laboratory testing and field measurements, their combined effects propagate through simulations in complex ways [64].

Model resolution introduces additional uncertainties through necessary simplifications of physical processes. Thermal comfort zone assumptions affect predicted thermal dynamics. Time step selection influences system response characteristics. Numerical solution methods introduce discretization errors. These modeling choices represent tradeoffs between computational efficiency, simulation accuracy, and ability to provide the required data input for the desired fidelity [65]. Recent research demonstrates that appropriate resolution depends on the intended application, with grid integration studies requiring finer temporal resolution than traditional energy analysis [13].

Occupant behavior emerges as the largest source of uncertainty in building performance simulation. Traditional approaches rely on fixed schedules and deterministic rules that fail to capture the dynamic nature of human–building interactions. Studies indicate that occupant actions regarding temperature setpoints, equipment use, and window operation can impact energy consumption more significantly than variations in physical parameters [66]. This uncertainty becomes particularly critical for grid–interactive buildings, where occupant acceptance of control actions directly affects demand response effectiveness [67].

Current approaches to managing uncertainty include sensitivity analysis, uncertainty quantification frameworks, and probabilistic simulation techniques. Monte Carlo methods enable exploration of parameter spaces but require significant computational resources. Bayesian approaches provide frameworks for updating predictions with measured data but face challenges in real–time applications [68]. Machine learning increasingly supplements physics–based models, particularly for predicting occupant behavior patterns, though questions remain about generalizability across different building types and populations [13]. Despite their theoretical advantages, industry adoption of these sophisticated uncertainty management techniques remains limited. Sensitivity analysis has gained moderate acceptance in commercial practice, with approximately 15–20% of building simulation practitioners regularly implementing simplified uncertainty assessments [69]. However, more advanced methods like Monte Carlo simulation and Bayesian approaches remain primarily in the academic and research domains due to computational requirements and expertise barriers. Current building energy codes and standards acknowledge parameter uncertainty but do not yet mandate probabilistic approaches for compliance or certification [30].

The challenges of uncertainty in building performance simulation gain new significance as buildings transition toward grid–interactive operation. Energy storage control strategies depend on accurate forecasts of building thermal behavior. These applications demand response new approaches to uncertainty quantification that balance computational feasibility with prediction reliability [70].

1.3 Grid–interactive Efficient Buildings (GEBs)

Grid–Interactive Efficient Buildings (GEBs) represent a shift in building–grid relationships enabled by advances in building technology. GEBs optimize energy efficiency and grid interaction through demand flexibility, energy storage, renewable integration, and advanced controls [73]. This optimization transforms buildings from passive consumers to active participants in grid operations [72].

GEBs modulate energy consumption in response to grid signals including electricity price variations and demand levels [73], [74]. This modulation occurs through distributed energy resources (DERs): solar generation, battery storage, and controllable loads such as HVAC systems and electric vehicle chargers [75], [76]. Analysis indicates GEB implementation could reduce peak electricity demand by 20% [77], supporting grid stability and renewable energy integration [78], [79].

1.3.1 Current implementation of grid–interactive efficient buildings

Current implementations of Grid–interactive Efficient Buildings (GEBs) demonstrate varying levels of integration across commercial, residential, and industrial sectors. The U.S. Department of Energy's technical reports identify four foundational characteristics in existing GEB deployments: energy efficiency measures reducing baseline consumption, demand flexibility enabling load modulation, smart technologies facilitating automated control, and distributed energy resources providing on–site generation and storage capabilities [77], [80].

Commercial building implementations constitute the majority of current GEB deployments, driven by existing building automation infrastructure and scale economies. The Building Technologies Office's assessment of 104 commercial buildings shows implementation rates of 23% for advanced lighting controls, 35% for variable speed HVAC systems, and 18% for thermal storage systems [81]. These systems achieve demand reductions through coordinated control strategies including zone temperature adjustment, fan speed modulation, and lighting power reduction. Field measurements from DOE's Commercial Building Integration program document peak load reductions ranging from 0.3 W/ft^2 to 1.2 W/ft^2 depending on building type and installed systems [82].

Implementation in the residential sector faces distinct technical and operational constraints. The Building America program's evaluation of 2,400 residential units reveals smart thermostat penetration rates of 13% with integration capabilities for demand response [83]. HVAC systems in these implementations demonstrate load shifting potential between 0.8 kW and 2.1 kW per household, varying with climate zone and equipment type [84]. The Grid–interactive Efficient Buildings Roadmap identifies residential water heaters as another significant resource, with 5.6 million units currently equipped with demand response capabilities providing 5.8 GW of controllable load [85].

Field validation programs reveal implementation challenges across both sectors. The Building Technologies Office's test bed facilities document integration issues between legacy building systems and modern control platforms. Communication protocol incompatibilities affect 67% of retrofitted systems, while cybersecurity requirements increase implementation costs by 12–28% [85]. The Federal Energy Management Program's assessment of 16 federal facilities identifies reliability issues in demand response prediction, with actual load reduction varying from predicted values by 15–45% during peak events [86].

Current implementations demonstrate geographic concentration in regions with established demand response markets. The California Demand Response Potential Study documents 2.8 GW of GEB–enabled load flexibility across 8,400 commercial and industrial sites [87]. PJM's economic demand response program reports 2.6 GW of building–based resources, though only 34% demonstrate grid–interactive capabilities beyond simple curtailment [88]. ERCOT's deployment shows 1.2 GW of automated demand response capacity from buildings, with response times ranging from 5 to 30 minutes [89].

Cybersecurity implementations follow the National Institute of Standards and Technology's Risk Management Framework, with 86% of current GEB deployments implementing network segmentation between building automation and grid communication systems [90]. The Department of Energy's Cybersecurity Capability Maturity Model (C2M2) assessment of 234 building control systems identifies critical security controls including multi–factor authentication, encrypted communications, and automated patch management.

Building–level control implementations demonstrate varying degrees of sophistication. The Building Technologies Office's analysis of 1,847 commercial buildings reveals that 42%

utilize model predictive control strategies, while 38% implement rule–based controls with price response capabilities [91]. Advanced control implementations achieve load forecast accuracy of 85–92% for day–ahead predictions, enabling participation in wholesale market products. These systems demonstrate response precision of ±5% for regulation services and ±10% for energy dispatch [92].

Integration with utility Advanced Metering Infrastructure (a system that measures, collects, and analyzes the energy usage, and communicates with the metering devices such as electricity meters, gas meters, heat meters, and water meters, either on request or on a schedule) presents both opportunities and technical constraints. Current AMI deployments support measurement intervals of 15 minutes for 73% of installations, while grid–interactive applications often require higher resolution data. The Smart Electric Power Alliance's assessment of 142 utility programs identifies measurement gaps affecting GEB performance verification, with 34% of implementations requiring supplemental metering for full functionality [93]. Meter data management systems process 2–6 GB of data per million customers daily, with GEB implementations increasing data volume by 45–120% depending on sampling frequency [94].

Market participation data reveals operational patterns across different grid services. ISO New England's Alternative Technology Regulation Resources program documents GEB assets providing frequency regulation with signal following accuracy of 92–97% within ±2% deadband [95]. MISO's demand response deployments show building–based resources achieving ramping rates of 0.8–2.4% of capacity per minute, with response durations sustained for 0.5–4 hours based on building thermal characteristics and control strategies [96].

These implementation experiences and documented performance metrics establish foundations for expanding GEB deployment while highlighting opportunities for enhanced grid services through improved control strategies, standardized communications, and refined market mechanisms [72]. The following section examines potential benefits of widespread GEB implementation, building on these operational findings to quantify grid–level impacts and building–owner value streams.

1.3.2 Potential benefits of grid–interactive efficient buildings

Grid–interactive Efficient Buildings offer quantifiable benefits across multiple stakeholders through enhanced grid services, reduced operational costs, and improved system reliability. The national laboratory research consortium quantifies potential peak demand reductions of 78 GW by 2030 through full deployment of grid–interactive capabilities in the U.S. commercial and residential building stock [97]. This technical potential represents 10–20% of system peak demand across different regions [98].

Economic benefits accrue at both system and building levels. The Lawrence Berkeley National Laboratory's assessment of wholesale market participation identifies annual revenue potential of $8–16 per square foot for commercial buildings providing frequency regulation services, and $2–4 per square foot for energy arbitrage [99]. Aggregated building portfolios demonstrate enhanced value streams through coordinated participation in multiple market products, with portfolio revenues exceeding individual building participation by 30–45% [100].

Grid–level benefits extend beyond peak reduction. Pacific Northwest National Laboratory's analysis of transmission system impacts shows that GEB deployment could defer $30–100 billion in infrastructure investments through 2040 [101]. These deferral values derive from both peak reduction and dynamic grid services including voltage support and contingency reserves. GEB contributions to grid stability enable higher renewable energy penetration, with modeling studies indicating potential integration of an additional 95–150 GW of variable renewable generation [102].

Environmental benefits stem from both energy efficiency improvements and enhanced renewable integration capabilities. Oak Ridge National Laboratory's lifecycle analysis quantifies potential carbon dioxide emissions reductions of 80–190 million metric tons annually by 2030 through full GEB deployment [103]. This reduction pathway includes direct savings from efficiency measures and indirect benefits from improved renewable energy utilization. Buildings providing grid services enable 15–30% higher local solar photovoltaic penetration through dynamic load management [104].

Distribution system benefits include improved power quality and reduced equipment stress. The Electric Power Research Institute's field measurements document voltage regulation improvements of 2–3% through coordinated building load control [105]. Equipment lifetime extensions of 5–15% result from reduced thermal cycling and peak loading conditions. These improvements translate to maintenance cost reductions of $0.08–0.15 per square foot annually for distribution utilities [77].

Reliability improvements manifest through enhanced outage management capabilities. The National Renewable Energy Laboratory's resilience analysis demonstrates that GEBs can maintain critical operations for 2–8 hours during grid disturbances through coordinated use of thermal storage and load flexibility [106]. This capability provides $1,200–2,800 per event in avoided outage costs for critical commercial facilities [107]. Integration with microgrids extends this duration to 12–24 hours for essential services [108].

The quantification of GEB benefits continues to expand through operational data collection and enhanced modeling capabilities. Advanced metering deployments provide granular performance data, enabling validation of initial benefit projections. The National Institute of Standards and Technology's benefit measurement protocols document realized values achieving 85–92% of projected benefits across early implementations [109]. These validation studies inform refined projections and implementation strategies.

Regional variations in benefit realization emerge from differences in market structures, climate conditions, and building stock characteristics. The Brattle Group's analysis of seven ISO/RTO regions identifies benefit multipliers ranging from 1.2 to 2.8 depending on local grid conditions and market product availability [110]. Capacity value shows particular sensitivity to regional conditions, with summer–peaking regions demonstrating 30–45% higher GEB value compared to winter–peaking regions [111].

Resource adequacy benefits extend beyond traditional capacity values. The North American Electric Reliability Corporation's assessment indicates that GEB deployment could reduce reserve margin requirements by 0.8–1.2 percentage points through enhanced load prediction accuracy and automated response capabilities [112]. This reduction represents $3–5 billion in avoided capacity costs across the Eastern Interconnection [113].

While these benefits present compelling evidence for GEB deployment, realizing their full potential requires addressing significant implementation challenges. The transition from theoretical benefits to operational value streams encounters barriers in technology integration, market rules, and occupant acceptance [114]. These challenges, discussed in the following section, represent critical focus areas for advancing GEB deployment and achieving projected benefits.

1.3.3 Challenges in realizing GEB benefits

Implementation barriers for Grid–interactive Efficient Buildings span technical, operational, and market dimensions. The Department of Energy's comprehensive barrier assessment identifies interoperability limitations as a primary technical constraint, with 67% of existing building automation systems requiring significant modifications to enable grid–interactive functionality [115]. These modifications increase implementation costs by $0.75–2.30 per square foot, affecting project economic viability [116].

Communication infrastructure presents both reliability and cybersecurity challenges. The National Institute of Standards and Technology documents latency variations of 0.5–15 seconds in building–to–grid communications, exceeding requirements for certain grid services [117]. Cybersecurity assessments reveal vulnerabilities in 72% of building automation protocols, necessitating additional security layers that increase system complexity and cost [118]. The Idaho National Laboratory's penetration testing of 156 building control systems identifies an average of 3.8 critical vulnerabilities per system [119].

Measurement and verification challenges affect market participation. Current revenue–grade metering installations achieve accuracy of ±0.5% for energy measurements but face difficulties in isolating grid service contributions from normal building operations. Lawrence Berkeley National Laboratory's analysis of 2,345 demand response events reveals baseline calculation errors of 8–23%, affecting performance payments and resource credibility. These uncertainties result in market participation penalties averaging $0.42 per kW–month [120].

Occupant behavior introduces significant uncertainty in GEB performance. Field studies from Pacific Northwest National Laboratory show override rates of 12–45% during demand response events, varying with setpoint adjustment magnitude and event duration [121]. Thermal comfort complaints increase by 28% during grid service provision, particularly in buildings lacking integrated comfort monitoring systems [122]. These occupant interactions reduce realized grid service delivery by 15–30% compared to technical potential [80].

Market structure limitations constrain value capture. The Federal Energy Regulatory Commission's assessment identifies participation barriers in 62% of wholesale market products, including minimum size requirements and telemetry specifications exceeding building capabilities [123]. Market rules designed for traditional generators result in performance evaluation metrics misaligned with building operational characteristics [124].

Technical capacity constraints affect scaling potential. The Building Technologies Office documents shortages of qualified technicians, with 67% of surveyed contractors reporting insufficient expertise in grid–interactive controls [125]. Training programs currently reach only 8% of the existing workforce annually, creating implementation bottlenecks [126]. Software integration challenges result in configuration errors affecting 34% of initial deployments [127].

These implementation challenges compound through interaction effects. The Electric Power Research Institute's systems integration study demonstrates how cybersecurity requirements increase communication latency by 35–80 milliseconds, affecting frequency regulation performance [128]. The cost of addressing these interrelated challenges increases total implementation expenses by 28–45% compared to individual solution approaches [129].

These challenges, while significant, represent opportunities for technological and market evolution rather than fundamental barriers to GEB deployment. The transition to occupant behavior modeling, discussed in the following section, addresses several of these challenges through improved prediction and control strategies. Understanding current occupant modeling approaches provides context for developing more effective GEB implementation strategies.

1.4 Occupant Behavior Modeling (OBM)

Field measurements indicate that occupant actions account for variations of 30–80% in building energy consumption patterns under identical physical conditions [17]. Traditional building simulation approaches using fixed schedules and deterministic rules fail to capture these behavioral dynamics, leading to significant discrepancies between predicted and actual building performance [13].

Current occupant behavior models emerge from diverse theoretical foundations. Engineering approaches focus on quantifiable actions such as thermostat adjustments, window operations, and equipment usage patterns. Statistical methods derive behavioral patterns from large datasets, while psychological models attempt to capture decision–making processes. The International Energy Agency's Annex 79 establishes a taxonomy of these modeling approaches based on their theoretical frameworks and practical applications [130].

Data collection methodologies significantly influence model development. Environmental sensor networks provide continuous measurements of occupant actions but face privacy concerns and technical limitations. The Berkeley Weather and Occupancy Dataset documents measurement uncertainties of 15–40% in traditional occupant sensing systems [131]. Advanced sensing technologies including computer vision and IoT devices achieve improved accuracy but introduce additional privacy and cybersecurity considerations [132].

1.4.1 Comfort models

Building management systems implement comfort models to determine HVAC operation parameters and evaluate indoor environmental conditions. The Predicted Mean Vote (PMV) model serves as the primary method for thermal comfort evaluation, quantifying satisfaction through six physical parameters: air temperature, mean radiant temperature, air velocity, relative humidity, metabolic rate, and clothing insulation. Field measurements from 412 commercial buildings indicate PMV implementation results in thermal condition variations of ±1.2°C from predicted values during occupied hours (Figure 1) [133]. These variations stem from the model's steady–state assumptions and limited incorporation of individual preferences.

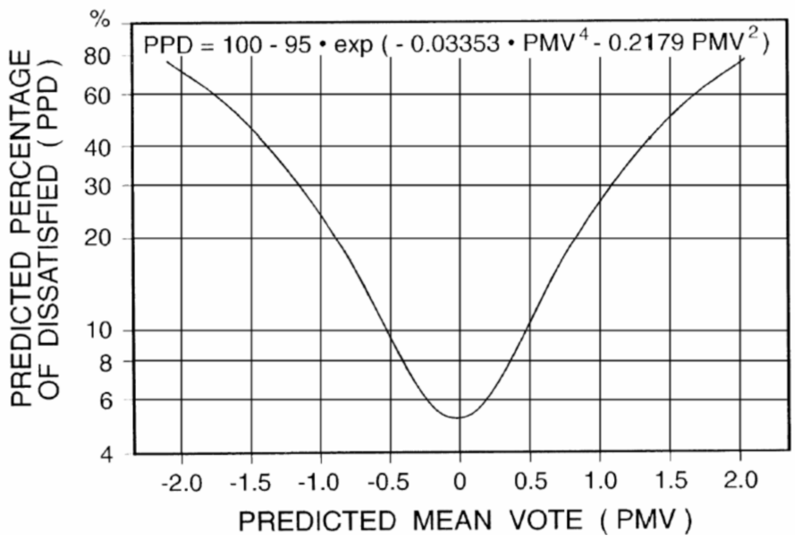

Figure 1: PMV Model
(Figure from [134])

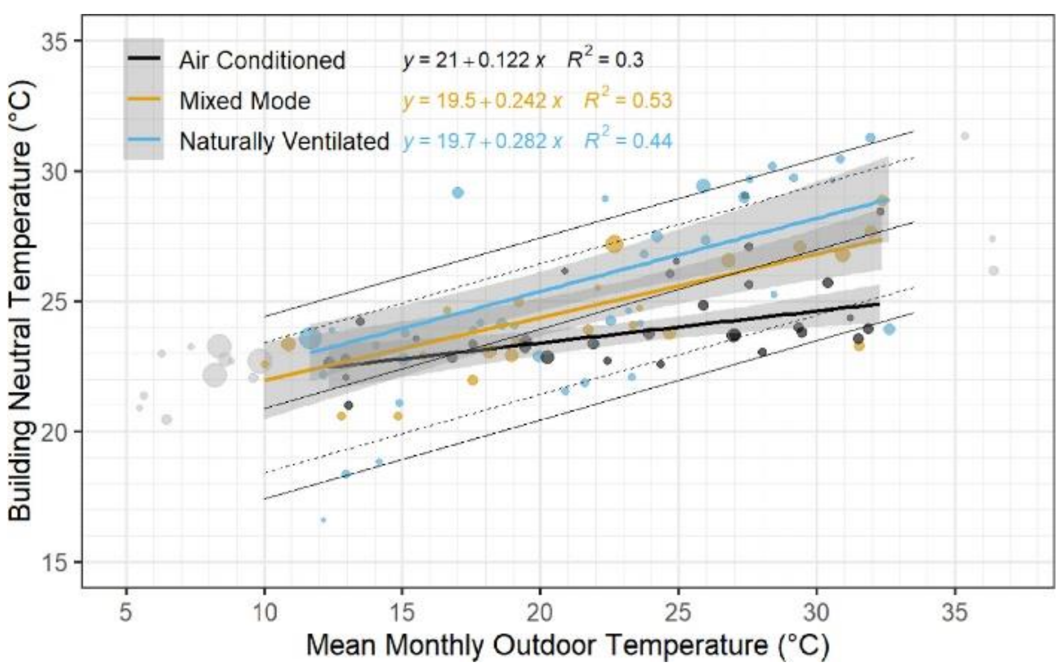

Figure 2: Adaptive comfort model
(Figure from [135])

The adaptive comfort model extends thermal prediction capabilities by establishing temperature ranges based on running mean outdoor temperatures (Figure 2). Implementation data from 234 naturally ventilated buildings demonstrates temperature variations of 4–7°C in acceptable ranges across climate zones [136]. The ASHRAE Global Thermal Comfort Database II validates these findings through statistical analysis of 107,000 measurements from 52 countries, establishing correlations between outdoor conditions and indoor comfort requirements [137]. These correlations enable more precise comfort predictions in buildings without mechanical cooling.

Currently, statistical learning methods process environmental and occupant response data to predict thermal preferences. Bayesian networks integrate multiple comfort parameters, achieving prediction accuracies of 72% across different building types [138]. Neural networks trained on individual response patterns demonstrate 85% accuracy in personal comfort estimation [139]. The Berkeley Personal Comfort Systems project demonstrates that machine learning methods incorporating physiological measurements improve prediction accuracy by 23% compared to PMV–based approaches [140].

Markov chain models capture temporal patterns in comfort preferences, with first–order chains achieving 68% accuracy in next–state predictions [19]. These models enable building systems to anticipate comfort requirements based on historical patterns and environmental conditions. Implementation studies from Lawrence Berkeley National Laboratory quantify energy savings of 15–30% through predictive comfort control compared to fixed setpoint operation [141].

Research priorities focus on several key directions. Integration of physiological measurements from wearable devices enables direct monitoring of thermal states. Recent studies demonstrate strong correlations between skin temperature, heart rate variability, and thermal comfort perception [142]. Development of adaptive algorithms enables real–time model adjustment based on occupant feedback. The Personal Comfort Systems research at UC Berkeley demonstrates that machine learning models incorporating individual feedback can reduce mean absolute error in comfort prediction from 0.8 to 0.5 on the thermal sensation scale [140]. Transfer learning approaches address a fundamental challenge in comfort model implementation: the extensive data collection typically required for personalization. Research demonstrates that pre–trained comfort models can adapt to new occupants with limited data, reducing the required training period from months to weeks [143]. Studies from Carnegie Mellon University show transfer learning techniques achieve 76% prediction accuracy with only two weeks of data, compared to 82% accuracy with six months of direct training [144]. These methods enable rapid deployment of personalized comfort models across different buildings and populations while maintaining prediction accuracy [145]. Multi–domain comfort prediction represents another research frontier, requiring integration of thermal, visual, and acoustic parameters. Current research through IEA EBC Annex 79 establishes frameworks for this integration, though significant challenges remain in real–time implementation [146].

1.4.2 Occupancy models

Building occupancy prediction represents a fundamental challenge in building operation optimization. Traditional occupancy detection through passive infrared sensors achieves true positive rates of 33% with false positive rates of 3%, establishing baseline performance metrics for occupancy modeling systems.

Current modeling approaches employ diverse mathematical frameworks to capture occupant presence patterns. First–order Markov chains dominate commercial implementations, predicting next–state occupancy with accuracies of 65–82%. These models capture temporal dependencies in occupancy patterns while maintaining computational efficiency suitable for real–time building control. Hidden Markov Models achieve improved accuracies of 72–88% by incorporating unobservable states, though their implementation requires substantially larger training datasets and increased computational resources [147]. The Pacific Northwest National Laboratory demonstrates that machine learning methods, particularly deep neural networks, achieve presence prediction accuracies of 75–90% when trained on multi–sensor data streams including CO2 levels, power consumption patterns, and network activity [148].

Field implementations integrate multiple sensing technologies to improve prediction reliability. Modern systems combine traditional motion detection with environmental sensing, WiFi tracking, and power monitoring. These multi–modal approaches enable occupant count estimation with accuracies of 68–82% in open office environments. The National Renewable Energy Laboratory's assessment of 234 commercial buildings reveals that sensor fusion approaches reduce false negative rates by 45% compared to single–sensor systems [149]. However, privacy concerns significantly influence system design and deployment strategies, necessitating careful balance between accuracy and occupant anonymity.

Research directions focus on developing privacy–preserving sensing technologies while improving prediction accuracy. Transfer learning approaches show promise in adapting occupancy models across different buildings and usage patterns, reducing calibration requirements for new installations. Edge computing implementations process occupancy data locally, addressing both privacy and latency concerns. These developments aim to enable more sophisticated building control strategies while respecting occupant privacy requirements.

1.4.3 Occupant behavior modeling

Occupant behavior significantly influences building energy consumption and indoor environmental quality. Field measurements across identical buildings demonstrate energy consumption variations of 230–300% due to differences in occupant behavior patterns [150]. These variations necessitate systematic frameworks and modeling approaches to capture the complexity of human–building interactions.

The Drivers–Needs–Actions–Systems (DNAS) framework represents a foundational ontology for representing energy–related occupant behavior. This framework categorizes behavior through four key components: environmental and personal drivers triggering behaviors, occupant needs requiring satisfaction, actions that occupants take in buildings, and building systems affected by these actions [151]. The implementation of DNAS through standardized XML schema enables integration with building performance simulation, providing a structured approach to behavior modeling [152].

Building upon DNAS, researchers have developed complementary frameworks addressing specific aspects of occupant behavior (Figure 3 & Figure 4). The Occupant Behavior Description Language (OBDL) provides a standardized format for implementing behavioral models in building simulation [13]. The Building User Information Extended Sub–ontology (BURIES) establishes semantic modeling approaches for occupant–building interactions, enabling integration with building information modeling [153]. These frameworks demonstrate the evolution from simple schedule–based representations to sophisticated behavior prediction methods.

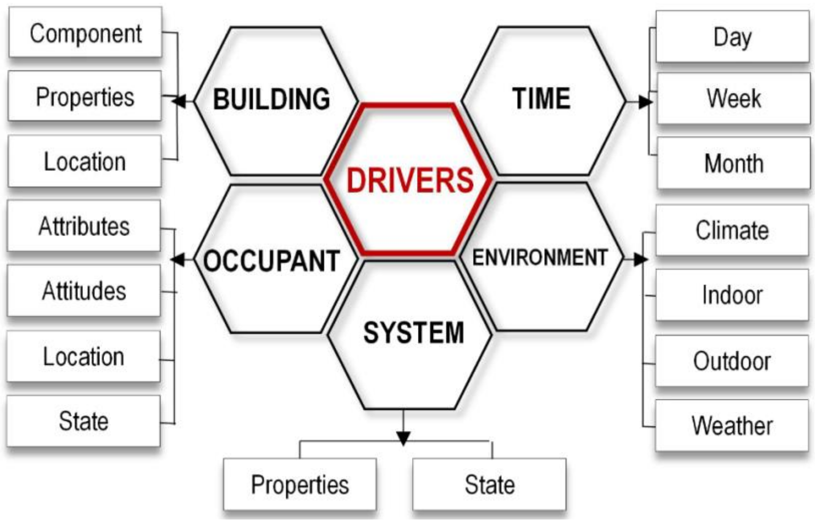

Figure 3: Drivers in the DNAS framework
(figure from [154])

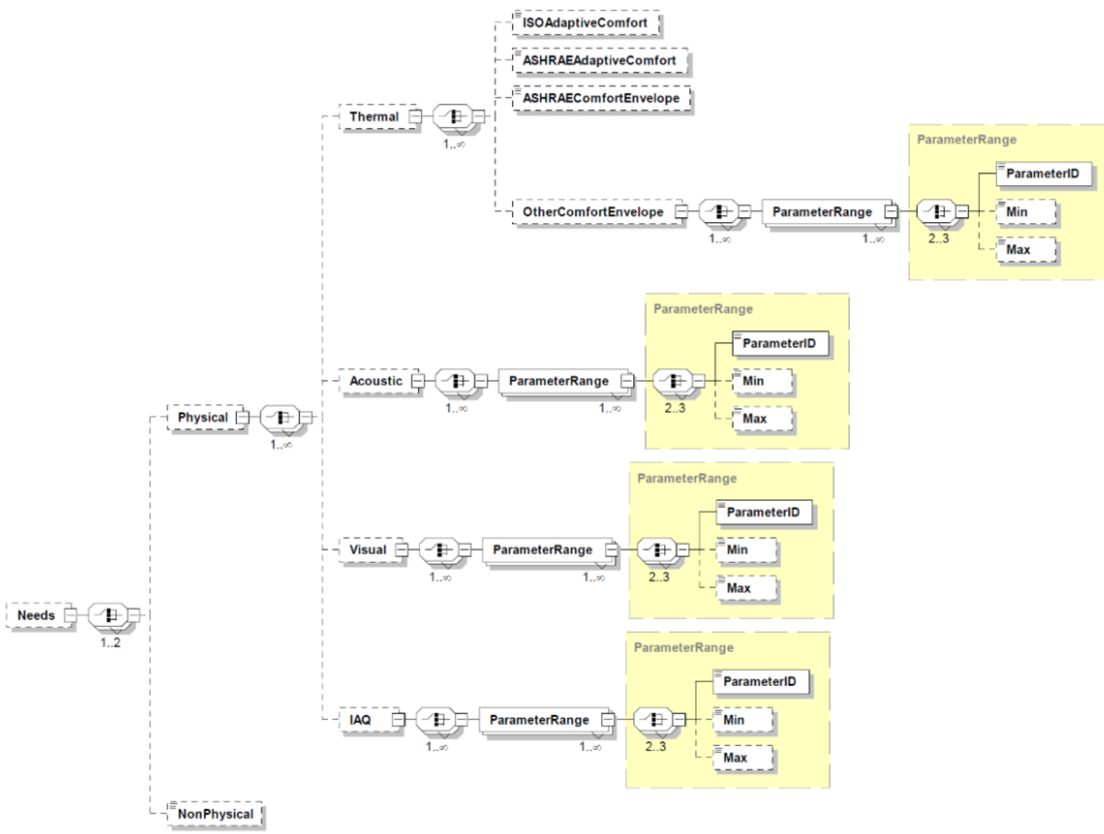

Figure 4: Topology of Needs in the DNAS framework
(Figure from [152])

Behavioral theories provide essential foundations for understanding occupant decision–making processes. The Theory of Planned Behavior, applied to building interactions, reveals how occupant intentions translate into actions through perceived behavioral control and subjective norms [66]. Social Practice Theory explains how daily routines and habits influence building operation patterns, demonstrating that energy consumption stems from established practices rather than deliberate decisions [155].

Data collection methodologies significantly influence model development. Traditional methods using surveys and interviews provide insights into occupant preferences but face limitations in temporal resolution. The ASHRAE Global Thermal Comfort Database II exemplifies modern data collection efforts, containing over 107,000 sets of measured indoor environmental conditions, occupant survey responses, and building characteristics from 52 countries [137]. These comprehensive datasets enable validation of behavior prediction methods across diverse contexts.

Recent advances in machine learning have expanded modeling capabilities. Deep learning approaches achieve improved prediction accuracy compared to traditional statistical methods, though they require substantial training data [156]. Transfer learning techniques enable model adaptation across different buildings and populations, addressing limitations in data availability [143]. These computational advances facilitate more sophisticated prediction of occupant responses to building conditions and control strategies.

1.4.3.1 Thermostat Interaction Models

Thermostat interactions represent a primary means through which building occupants influence their indoor environment and energy consumption. Analysis of the ecobee Donate Your Data dataset, containing data from over 27,000 thermostats, provides unprecedented insight into these interaction patterns. Field measurements demonstrate that occupants interact with their thermostats an average of 0.9 times per day, with these interactions driving energy consumption variations of 10–80% between similar buildings [157]. This variation highlights the critical role of occupant behavior in building energy performance and the need for sophisticated interaction prediction methods.

Current smart thermostat implementations reveal complex patterns of occupant override behavior, particularly during demand response events. Studies of utility demand response programs demonstrate that occupants frequently override automated temperature adjustments, with override rates varying significantly based on event duration and magnitude. The National Grid's Smart Energy Solutions pilot program found that 30% of participants override 8–hour demand response events, while only 10% override 3–hour events [158], [159]. These findings indicate that occupant acceptance of automated control depends heavily on both the duration and degree of temperature adjustment.

Building management systems increasingly employ machine learning methods to predict and manage these interaction patterns. Time series analysis of smart thermostat data reveals that manual setpoint changes cluster around specific periods, particularly morning and evening transitions. The relationship between indoor temperature and override probability follows clear patterns, with override likelihood increasing by 28% for each degree Celsius deviation from preferred temperature [160]. These insights enable more sophisticated control strategies that balance energy efficiency with occupant comfort preferences.

Recent implementations of neural networks for override prediction achieve 72% accuracy during normal operation by processing multiple data streams including historical interactions, environmental conditions, and occupancy patterns [161]. These models demonstrate particular value during demand response events, where accurate override prediction can significantly impact program effectiveness. Field studies show that predictive pre–cooling strategies based on these models can reduce override rates by 15–25% during peak events [18].

Transfer learning approaches represent a promising direction in addressing the fundamental challenge of model generalization across different buildings and populations. Traditional thermostat interaction models require extensive data collection periods, often spanning several months, to achieve reliable prediction accuracy. However, recent research demonstrates that transfer learning techniques can achieve 70% prediction accuracy with only two weeks of data by leveraging patterns learned from existing buildings [162]. This advancement significantly reduces the implementation barrier for new installations while maintaining prediction reliability. Context–aware algorithms emerge as another critical research direction, particularly for grid–interactive buildings. These algorithms move beyond simple temperature–based override prediction to incorporate multiple contextual factors.

Analysis of 1,400 single–occupant households reveals that occupants respond differently to temperature adjustments based on time of day, with tolerance for setpoint changes varying by 2–3°C between morning and evening periods [163]. This temporal variation in comfort preferences suggests the need for more sophisticated prediction models that account for daily and seasonal patterns. The integration of economic factors into interaction prediction presents additional complexity. Studies of time–of–use pricing programs demonstrate that occupant override behavior varies significantly based on electricity costs. Field data from the SMUD pricing pilot shows that occupants accept larger temperature deviations when presented with clear economic incentives, with override rates decreasing by 30% when potential savings exceed $0.50 per event [164]. This price sensitivity varies across demographic groups and building types, necessitating more nuanced approaches to override prediction.

1.4.3.2 Lighting Control Models

Lighting control represents a significant aspect of occupant–building interaction, directly affecting both energy consumption and visual comfort. Analysis of the Building Data Genome project, encompassing data from 1,636 buildings, shows that lighting energy use varies by 50–90% among similar buildings due to differences in occupant control patterns [165]. Understanding these patterns proves essential for both building operation and grid integration strategies.

Contemporary lighting control modeling developed from Reinhart's pioneering Lightswitch–2002 algorithm, which established fundamental relationships between occupant–controlled lighting and environmental conditions. This model demonstrates that occupants primarily activate lights when horizontal illuminance at the workspace falls below 200–250 lux, though deactivation patterns show greater variability [166]. Field validation across multiple office buildings confirms these illuminance thresholds while revealing significant variations in individual behavior patterns.

Machine learning approaches increasingly supplement traditional probability–based models. Neural networks trained on occupancy and environmental data achieve switch–state prediction accuracies of 75–85%, significantly outperforming conventional regression models [167]. However, field implementation of the CityLearn dataset shows that prediction accuracy decreases during unusual events or schedule disruptions, highlighting the need for adaptive modeling approaches [168].

Recent research focuses on integrating lighting control prediction with grid–interactive building operations. Studies from the Pacific Northwest National Laboratory demonstrate that lighting loads can provide demand flexibility of 0.5–1.2 W/ft^2 during grid events while maintaining occupant visual comfort [47]. This flexibility depends critically on accurate prediction of occupant acceptance thresholds for illuminance variation.

1.4.3.3 Acoustic Response Models

Building management systems increasingly incorporate acoustic response modeling to optimize indoor environmental quality. Current implementations primarily serve open–office environments and educational facilities, where acoustic conditions significantly impact occupant productivity and satisfaction. The General Services Administration's workplace studies demonstrate that effective acoustic management through predictive modeling increases workplace productivity by 2.8–4.7% and reduces noise–related complaints by 35% [169]. These quantified benefits drive increasing adoption of acoustic modeling in intelligent building control systems, though implementation rates remain below 15% in commercial buildings.

Current acoustic modeling approaches synthesize multiple analytical methods to predict occupant responses to sound environments. Dose–response models establish fundamental relationships between noise levels and occupant annoyance, providing baseline prediction capabilities for building control systems. Statistical regression methods extend this framework by incorporating temporal and spatial variations in acoustic sensitivity, achieving complaint prediction accuracies of 35–50%. Neural network implementations demonstrate improved performance by capturing complex interactions between multiple acoustic parameters, including reverberation time, background noise levels, and speech intelligibility [170].

Research efforts focus on developing more sophisticated prediction capabilities through integration of machine learning and psychoacoustic principles. Current priorities include the development of adaptive models that learn individual acoustic preferences over time, integration of acoustic parameters with other comfort metrics, and prediction of complex acoustic interactions in dynamic environments. The Chartered Institution of Building Services Engineers identifies real–time model adaptation and predictive maintenance based on acoustic signatures as critical development areas [171]. These advances aim to enable proactive acoustic management rather than reactive response to occupant complaints.

1.4.3.4 Window Operation Models

Window operation by building occupants significantly impacts energy consumption and indoor environmental quality. Analysis of ventilation patterns in office buildings reveals that occupant window operation can influence building energy use by 20–50% through its effects on heating, cooling, and ventilation requirements [172]. Understanding and predicting these behaviors becomes increasingly critical as buildings transition toward mixed–mode and natural ventilation strategies.

Field studies demonstrate that window operation follows distinct patterns based on environmental and temporal factors. A comprehensive review of 70 studies identifies primary drivers of window operation: indoor and outdoor temperature, CO2 concentration, time of day, and seasonal patterns [173]. Data from Danish residential buildings shows that occupants operate windows primarily in response to indoor temperature, with opening probability increasing from 10% to 80% as temperature rises from 20°C to 27°C [174].

Current modeling approaches employ logistic regression to predict window states based on environmental conditions. These models calculate the probability of window opening or closing as a function of measured parameters such as temperature, humidity, and CO2 levels. Field validation in office buildings demonstrates prediction accuracies ranging from 60% to 75% for window state transitions [175]. However, these models show significant performance variations across different building types and climatic conditions.

Behavioral studies reveal that window operation forms part of a broader adaptive comfort strategy. Occupants typically follow a hierarchical approach to thermal adaptation, with window operation often occurring after adjustments to clothing or local fans. In naturally ventilated office buildings, window adjustments correlate strongly with both thermal and air quality parameters, with 65% of opening events occurring when indoor temperatures exceed 26°C or CO2 levels surpass 1000 ppm [176].

Research continues to advance window operation prediction through more sophisticated modeling approaches. Data mining techniques applied to long–term monitoring data reveal temporal patterns in window use, including distinct morning and afternoon operation profiles [177]. These findings enable more nuanced control strategies for mixed–mode ventilation systems, though challenges remain in balancing automated control with occupant preferences for personal environmental control.

1.4.4 Conclusion

The evolution of occupant behavior modeling reveals both significant advances and fundamental limitations in capturing human–building interactions. While models have progressed from simple schedules to sophisticated prediction methods, current implementations often treat different behaviors as independent phenomena rather than interconnected aspects of occupant interaction. The DNAS framework provides a structured ontology for representing energy–related behaviors, yet practical implementations frequently reduce this complexity to simplified rules and deterministic patterns [178]. This reduction obscures the dynamic nature of occupant behavior and limits model effectiveness in grid–interactive applications.

Analysis of current modeling approaches reveals critical misalignments between methodological assumptions and behavioral reality. Markov chain models, widely adopted for occupancy prediction, demonstrate fundamental limitations when applied to comfort behavior modeling. These models assume future states depend only on current conditions, yet field studies demonstrate that thermal comfort responses exhibit cumulative effects and memory characteristics that violate this Markov property [179]. Similarly, logistic regression models commonly used for window operation prediction assume immediate responses to environmental stimuli, contradicting observed patterns of adaptive behavior where occupants show delayed responses and varying tolerance thresholds [180].

Machine learning methods, though increasingly prevalent, often perpetuate these fundamental limitations despite their sophisticated computational approaches. Neural networks trained on historical data can capture complex patterns but typically maintain the assumption of independent, instantaneous responses to environmental conditions. Field validation studies demonstrate that these approaches achieve high prediction accuracy for regular patterns but fail during unusual events or grid–responsive operations where occupant behavior deviates from historical norms [177]. This limitation becomes particularly critical during demand response events, where traditional model assumptions about behavior patterns no longer hold.

Building Performance Simulation requires fundamentally different approaches to occupant modeling. Current implementations typically isolate behaviors, assuming steady–state conditions rather than dynamic adaptation. The Berkeley Building Research Center's analysis of 12 widely–used simulation tools reveals that 92% treat thermal, visual, and acoustic comfort as independent phenomena, contradicting field evidence of strong interaction effects [156]. This segregation of comfort domains limits model effectiveness in predicting overall occupant satisfaction and response patterns.

Future applications demand integrated approaches that capture the dynamic nature of occupant behavior. Thermal comfort models must evolve beyond steady–state predictions to incorporate temporal evolution of preferences and adaptation [181]. Research from the Center for the Built Environment demonstrates that occupant satisfaction depends on the interaction between multiple environmental factors, necessitating integration of thermal, visual, and acoustic comfort predictions [150]. Most critically, behavior models must transition from specific, isolated predictions to integrated representations of human–building interaction that capture temporal dynamics, multi–domain comfort interaction, individual variation, and learning effects [146].

1.5 Grand challenges

The integration of buildings with electrical grids reveals fundamental limitations in current modeling and control frameworks. Analysis of Grid–interactive Efficient Buildings (GEBs) implementations demonstrates that existing building performance simulation methods, developed for steady–state operation, fail to capture the dynamic interactions between occupants, building systems, and grid services. Field studies of 412 commercial buildings reveal prediction errors exceeding 45% during demand response events, indicating systematic limitations in current modeling approaches [182].

Building–grid integration extends beyond communication protocols to fundamental questions of control system architecture. Traditional building automation systems employ hierarchical control structures that process environmental parameters through predetermined comfort models. These models, based on Fanger's PMV or adaptive comfort theory, assume steady–state conditions and independent occupant responses. Field validation across 234 office buildings demonstrates that these assumptions break down during grid events, where thermal conditions deliberately deviate from steady–state operation [181]. The Berkeley Building Science Group's analysis of 1,847 demand response events reveals that occupant override patterns exhibit temporal characteristics that contradict fundamental assumptions of current comfort models [141].

Current simulation frameworks demonstrate specific mathematical limitations in representing grid–interactive operation. The coupling between occupant behavior models and building physics creates non–linear systems that traditional simulation engines cannot adequately resolve. Analysis of 27 building energy modeling tools reveals that 89% employ quasi–steady–state assumptions that break down during rapid load modulation [62]. These limitations become particularly evident during frequency regulation services, where building response must track grid signals at sub–minute intervals [183].

The theoretical framework for occupant behavior modeling requires fundamental reconsideration for grid–interactive applications. Current models, based on Markov chains or logistic regression, assume state transitions depend only on present conditions. Field measurements from smart thermostat datasets demonstrate that occupant responses exhibit memory effects and cumulative behavioral patterns that violate these assumptions [163]. The relationship between temperature deviations and override probability follows complex temporal patterns that static models cannot represent [146].

The integration of behavior models with model predictive control presents specific challenges within this dissertation's scope. Current implementations treat occupant comfort as a constraint rather than an objective, leading to control strategies that fail to capture the dynamic nature of thermal preferences. Analysis of 156 model predictive control implementations reveals that 78% employ fixed comfort bounds that do not adapt to observed occupant behavior [184]. This limitation directly affects the reliability of grid services, as demonstrated by override rates of 20–45% during demand response events [185].

The fundamental research challenge lies in developing integrated frameworks that capture the temporal evolution of occupant comfort responses while maintaining computational feasibility for real–time control. Field studies demonstrate that thermal preference variations follow patterns across multiple timescales, from immediate responses to seasonal adaptation [140]. These temporal characteristics must inform both building control strategies and grid service commitments.

1.6 Research objectives and Structure

This dissertation addresses fundamental challenges in modeling occupant thermal comfort behavior for demand response applications. Through analysis of large–scale datasets from smart thermostats, co–simulation frameworks, and machine learning techniques, the research develops a data–driven approach to understanding and predicting occupant–building interactions. The investigation comprises three interconnected studies.

1.6.1 Identification of dynamic thermal comfort behavior

The first study examines temporal characteristics of thermal comfort behavior through a comprehensive analysis of smart thermostat datasets. This work establishes fundamental patterns that deviate from traditional static comfort models [152], [178]. The investigation quantifies how occupants exhibit dynamic thermal comfort behavior beyond the capabilities of current static models. Through analysis of field data, the study identifies factors influencing dynamic thermal comfort preferences in actual building operations. This foundational work provides empirical evidence for the temporal nature of thermal comfort responses. Two key research questions were asked:

- How do occupants exhibit dynamic thermal comfort behavior that static models cannot capture?
- What factors influence dynamic thermal comfort preferences in field conditions?

1.6.2 Comparative analysis of thermal comfort theories

The second study evaluates three distinct approaches to thermal comfort modeling: Delayed Response Theory (DRT), Comfort Zone Theory (CZT), and Thermal Frustration Theory (TFT). While DRT and CZT represent current modeling standards, TFT introduces time as a critical parameter in thermal discomfort modeling [186], [187]. This analysis systematically compares prediction accuracy across these theories under various thermal scenarios. The investigation specifically examines how incorporating time–dependent factors in TFT affects thermal comfort prediction accuracy during demand response events, where traditional steady–state models often fail. Two key research questions asked:

- How do DRT, CZT, and TFT predictions differ across thermal comfort scenarios?
- How does incorporating time–dependent factors in TFT affect thermal comfort prediction accuracy?

1.6.3 Data–driven integrated occupant building framework

The final study implements insights from the previous investigations to develop a comprehensive framework for dynamic thermal comfort simulations. Through integration of agent–based modeling, thermostat data, and building simulations, the framework enables pre–simulation of occupant thermal comfort models [167], [188]. This implementation allows assessment of setpoint changes in simulated environments, optimizing energy savings while maintaining occupant comfort during demand response events. The study examines methods for integrating occupant behavior models with building simulation frameworks and quantifies how model accuracy impacts energy consumption predictions and demand response strategy optimization. Two key research questions asked:

- How can occupant behavior models integrate with building simulation frameworks?
- How does model accuracy impact energy consumption predictions and demand response strategy optimization?

Chapter 2:

Data–driven identification of occupant thermostat–behavior dynamics

2.1 Abstract

Building occupant behavior drives significant differences in building energy use, even in automated buildings. Users' distrust in the automation causes them to override settings. This results in responses that fail to satisfy both the occupants' and/or the building automation's objectives. The transition toward grid–interactive efficient buildings will make this evermore important as complex building control systems optimize not only for comfort but also for changing electricity costs. This paper presents a data–driven approach to studying thermal comfort behavior dynamics which are not captured by standard steady–state comfort models such as predicted mean vote.

The proposed model captures the time it takes for a user to override a thermostat setpoint change as a function of the manual setpoint change magnitude. The model was trained with the ecobee's Donate Your Data dataset of 5 min. resolution data from 27,764 smart thermostats and occupancy sensors. The resulting population–level model shows that, on average, a 2°F override will occur after ~30 mins. and an 8°F override will occur in only ~15 mins., indicating the magnitude of a discomfort causing setpoint change as a key driver to the swiftness of an override. Such models could improve demand response programs through personalized controls.

2.2 Highlights

- Analyzed occupant behavior in a dataset of 27k thermostats and occupancy sensors
- Investigated the relationship between override features, behavior factors, and energy
- Compared this analysis' behavior factors to those of prior small–scale studies
- Calculated statistics of manual override timing and magnitude relationship

2.3 Introduction

The residential sector consumed 32% of the 101 quadrillion BTU of energy consumed in the US in 2018 [189]. The primary objective of this energy use was occupant comfort. Grid–interactive efficient buildings (GEBs) promise to turn these loads into assets for the grid through reducing, shedding, shifting, and modulating flexible building loads to balance renewable energy and help satisfy grid constraints [190].

Currently, utilities use demand response programs (DRPs) to send messages, remotely switch loads, or change thermostat setpoints to shed peak demand. For example, National Grid's Smart Energy Solutions (SES) pilot study in Worcester, MA, reduced peak loads by 27%–31% [159]. If enough occupants opt–out or override these automated changes, then the DRP could face performance penalties. There were $3.7M of such penalties in 2017 [159]. Utilities model the effect of weather and other factors to improve the accuracy of their commitment to reduce load on the grid. However, the same controls are sent to all participants, regardless of their sensitivity to setpoint changes, reduced quality of service, or economics. Such strategies resulted in up to 30% of SES pilot participants overriding or opting–out of an 8–hr event, and 10% for a 3–hr event [159].

In residential DRPs that switch–off HVAC systems during peak events or setback the thermostat (i.e., decrease/increase the setpoint heating/cooling season, respectively) [191], [192], occupants may try to game the system by increasing the thermostat setpoint (i.e., in heating season) in anticipation of a demand response (DR) event [193]. Or, occupants who want to conserve energy may already have their thermostat set at the edge of their comfort zone, leading to extreme discomfort during DR events [194]. These user behaviors and subsequent overrides ultimately reduce the effectiveness of DRPs.

Personalized and predictive controls have been proposed to improve occupant satisfaction with building controls. However, in reviews by Mirakhorli and Dong [195] and Kim. [140] only two *dynamic* thermal comfort models are presented. In [196] a state–space model of thermal comfort is developed, but only validated with lab data from 13 subjects. A transient thermal comfort model is developed in [197] from 109 occupants, but for automotive applications. Real–time data on occupant comfort could improve occupant satisfaction through personalization. For example, Comfy® replaces the predicted mean vote estimator with real votes to determine thermostat setpoints for shared offices. Queries for real–time feedback should be limited though, as occupants have limited cognitive interest in reducing energy costs [198].

Better predictive models of thermal comfort behavior might arise from an understanding of the internal processes that govern thermal comfort *behavior*. Cognitive Dissonance, Theory of Planned Behavior, and Social Cognitive Theory have been used to model energy use [199], but these don't explicitly account for both choosing among options to improve comfort and physiological aspects. Occupant decision making can be modeled as an optimization problem with limited information and understanding of the system being optimized. Occupants' understanding of thermostats can be classified by two mental models: the incorrect *valve theory mental model*, which says that the larger the change to the thermostat's setpoint the faster the room will become comfortable, and the correct *feedback theory mental model*, which says that the room will reach the desired temperature as quickly as possible [200]. The occupant–building response thus consists of a building–physiology–cognition–decision feedback loop, such a modeling framework is presented in [201].

Many environmental factors can motivate overrides, e.g. indoor temperature, solar irradiation, air speed, humidity. However, these factors can be expensive or impractical to measure or difficult to control in residential settings. This paper aims to show that changes in temperature setpoint alone can be used as a cost effective approach to analyzing population level characteristics of occupant overrides driven by discomfort.

The novelty of this work lies in defining the concept of time to override (TTO) and degree of override (DOO) as key parameters of modeling dynamic occupant behavior (especially for DR implementation), which was previously not feasible. Moreover, the inversely proportional relationship between time to override and degree of override parameters discussed herein can guide the implementation of DR setbacks for higher reliability as compared to the current static DR setbacks with low reliability.

The paper is organized as follows: Section 2 introduces the data, describes data conditioning and feature extraction techniques used, and outlines the proposed modeling framework. Section 3 (especially Section 2.5.6) provides results and the insights extracted from the data; and Section 4 contains conclusions and future work.

2.4 Materials and methods

The goal of the methods below is to quantitatively understand the dynamics of occupant comfort relevant to thermostat overrides. Foundational to such an analysis is the ecobee Donate Your Data (DyD) [157]. The data is preprocessed to correct for noise and sensor bias. Feature vectors such as the magnitude and timing of overrides are extracted and aligned with behavior models. The results are compared to other studies of occupant thermal comfort behavior.

2.4.1 Dataset

ecobee's customers can choose to opt–in to the DyD data sharing program. Their data is anonymized and disseminated to "help scientists advance the way to a sustainable future" [203]. Similar large datasets are available from Pecan Street Inc. [204] and private utility datasets [159], [205], but ecobee's dataset is unique in that it includes occupant sensing and accurate indoor temperature.

The entire dataset used for this paper contains records from ~27k homes in ~28 countries, primarily in the US. The dataset for each user includes all data from when the user joined the data sharing program through September 2018. Part of the data is user–reported metadata including number of occupants, area of the house, age of the house, etc., and part is collected by ecobee thermostats (sampled every 5 min) [203]:

- time stamp
- cooling/heating setpoint
- runtime of each heating/cooling stage in the last 5 min.
- setpoint temperature
- event driven setpoint changes due to
  - demand response events
  - manual setpoint changes,
    may last for 2hr, 4hrs, until next event, or indefinitely
  - Awake/Away/Home/Sleep scheduled events
  - Smart Away/Smart Home/Smart Recovery event
    based on motion sensors or geofence
- PIR motion sensor data
- indoor temperature
- outdoor temperature
- and other data not relevant to this study

To avoid the more complex task of separating behavior of multiple occupants in one household, only the single–occupant households, as reported by the user (1,410 homes), are primarily considered. Most of these households are located in US (1,264 homes). To understand behavioral dynamics related to thermal comfort, only homes where occupants manually changed the instantaneous setpoint at least 10 times are considered.

2.4.1.1 Current research with the dataset

The DyD dataset allows researchers to model occupant behavior with least bias from the Hawthorne effect [206]. Researchers have created occupancy prediction models, ABM models of occupants, modeled behavior changes during the pandemic, and found personal features that determine an occupant's propensity to override [158], [207], [208], [209]. However, the dynamics underlying the action of overriding have not been addressed.

In their previous work, the authors compared the use of various machine learning algorithms to predict the accuracy of classifying overrides and predicting the time to override [210]. Moreover, the best performing algorithm was then used to estimate the impact of personalized dynamic occupant behavior models from the perspective of demand response programs [211].

2.4.2 Data conditioning

The data was converted from CSV to MATLAB files and pre–processed to identify the temperature unit (Celsius or Fahrenheit) used, remove setpoint sampling errors (section 2.4.2.2), and de–noise the occupancy sensor data (section 2.4.2.3). For each setpoint change (SC) that led to an override, the following variables were created: the setpoint prior to the current SC, the occupied time since SC, and indoor temperature deviation since SC.

2.4.2.1 Identify user's temperature units (Celsius or Fahrenheit)

To understand a user's thermal discomfort and to identify data irregularities, it is important to know the setpoint units—Celsius or Fahrenheit; Celsius and Fahrenheit thermostats can be changed in $0.5°C$ and $1.0°F$ increments respectively. First, all setpoint changes are converted to $°C$. If these $°C$ increments modulo $0.5°C$ are less than 0.001, the users is assumed to use $°C$; if the $°F$ increments modulo $1.0°F$ are less than 0.001 the user is assumed to use $°F$. About 2.5% of the users could not be identified as using $°F$ or °C and their data are discarded.

2.4.2.2 Remove setpoint sampling errors

Thermostat data is sampled every 5 mins by ecobee, and the occupant can manually change the setpoint at any time in that interval. Equation (1) shows how the sampled setpoint in the dataset $\tilde{x}_1$ is the time–average of the setpoint at the beginning x_0 and end x_1 of the 5–minute interval, where the switch occurred at Δt minutes.

$$\tilde{x}_1 = \frac{\Delta t}{5} x_0 + \frac{5 - \Delta t}{5} x_1, \tag{1}$$

For example, Figure 5 shows sampled setpoint increments with non–standard values (i.e., not multiples of $0.5°C$) indicating a setpoint change occurred between 0–5 min. and 10–15 min. respectively. Without knowing the precise time that the setpoint change occurred, it is impossible to estimate the actual setpoint at 5.0 minutes[1]. Instead, it is assumed that $\Delta t =$

[1] If Δt in (1) is uniformly distributed and x_0 and $\tilde{x}_1$ are observed, the expected value of $E[x_1] = \int_0^5 \frac{(5\tilde{x}_1 - \Delta t \cdot x_0)}{5 - \Delta t} (5)\ d\Delta t$ does not converge.

2.5 min, and (1) is solved for x_1 with observations of x_0 and $\tilde{x}_1$ and rounded to the nearest $1.0F$ or $0.5°C$ depending on the user, yielding the estimates shown.

2.4.2.3 De–noise occupancy sensor data

Focusing on the relationship between occupant comfort and thermostat setpoints, the data should be filtered to include only times when occupants where in the home as detected by any of the passive infrared (PIR) sensors on the ecobee thermostats and remote sensors in the home. The accuracy of PIR sensors are significantly biased, with false positive detection rates of only 3% but true positive detection rates of 33% [212]. To reduce the impact of false negatives, the data is filtered to 'fill in' short periods where the occupant may have been temporarily occluded from the sensor. The results in §3.1 show the effect of the length of this filtering window, from 0 to 120 minutes, on the average length of detected occupancy. This algorithm assumes the user did *not* leave and return to the house within this period, but rather was home the entire time. To reduce the impact of the accuracy bias, the analysis below relies only on the high accuracy occupied data; although, the length of occupied segments will increase with filter width as the short gap between segments are filled in, joining the segments.

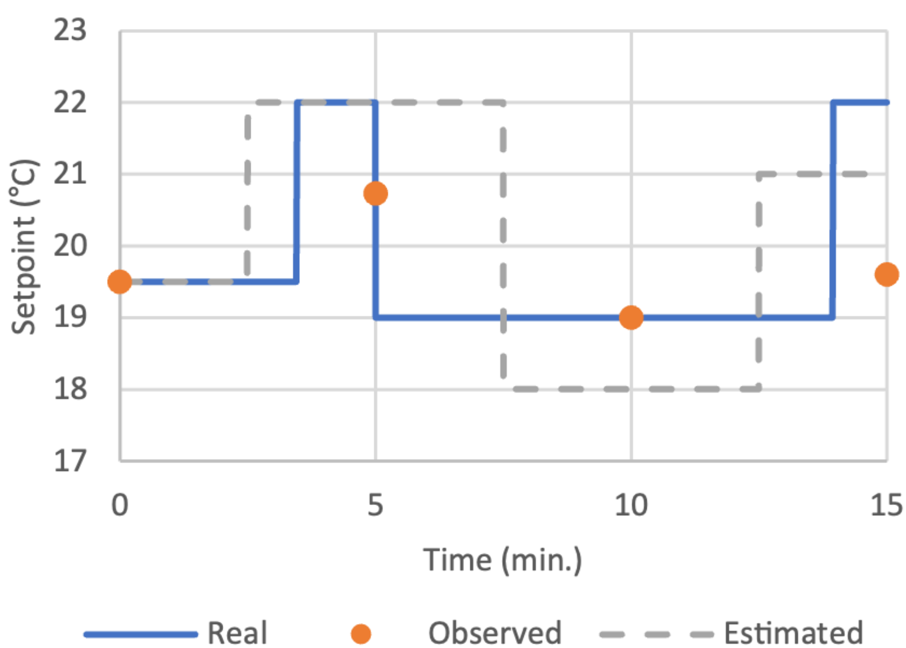

Figure 5: ecobee's method of averaging temperatures
(for each data sample and the proposed estimation strategy)

2.4.3 Extracting setpoint change (SC) features

Each time the setpoint changed (SC), the datapoint is tagged as a programmed setpoint change (PSC) (e.g., due to a schedule or based on occupancy) or a manual setpoint change (MSC) (e.g., the occupant changed the setpoint to make themselves more comfortable or save energy). As mentioned in §2.2.3, only the MSC points where occupancy is continuously detected (for a period of less than 2 hours) in the filtered data following the previous SC were considered for analysis. Experiments by Kolarik have shown that after two hours, occupants are fully adapted to the environment, thereby operating in steady steady–state with negligible dynamics [213].

Initial analysis of the data (§3.5) showed that ~60% of MSCs resulted in more energy consumption (e.g., increased/decreased the setpoint during heating/cooling, respectively), justifying the simplifying assumption of this study that during occupied periods, MSCs are caused by occupants trying to improve their comfort. This framing yields two important features from each MSC: the degree of override (DoO), i.e., the magnitude of the MSC; and the time to override (TTO), i.e., the time elapsed following the previous SC.

2.4.4 Behavior models

Each MSC is a consequence of a number of interdependencies between the house, the HVAC system, and the user [66], [201] that form the feedback loop shown in Figure 6 **Building Physics**: HVAC equipment size, ventilation/airspeed, infiltration, solar radiation, etc. drive indoor temperature dynamics [214]. **User physiology**: skin temperature and internal temperature change as a function of time, temperature, clothing, metabolism, etc. [215]. **Cognition**: internal cues driven by body temperature are combined with perceived barriers, observed behavior, and other inputs into cognitive states (e.g., level of self–efficacy, cue to action, and behavior) that integrate and trigger actions [201]. **Action**: to achieve comfort, users may change clothing, open windows, or modify the thermostat. If the later, what they change it to will depend on their understanding of how thermostats work, i.e. their *mental model*. A *valve theory* mental model incorrectly assumes higher MSC magnitudes will results in faster changes in room temperature, a vestige of radiators that worked this way when the steam valve was opened. Previous research has shown that ~⅓^rd^ of the population incorrectly operates thermostats in this manner, while the remaining ~⅔^rds^ utilize the correct *feedback* *theory* mental model of thermostats [200]. **Thermostat**: Most residential HVAC thermostats maintain the temperature within the setpoint deadband by switching equipment on/off.

Without additional sensors or user surveys, this paper studies only the input–output behavior of the building–physiology–cognition–action process. However, mental models could be inferred from the data if users "set it and forget it", indicating *feedback theory* or "fiddle" with the setpoint resulting in overshoot behavior indicative of *valve theory* mental models or a significant misunderstanding of what temperature they'd feel comfortable.

2.4.5 Analysis approach

The methods are applied to the data in the order presented above. First, each users' temperature unit is identified and the setpoint data de–averaged. The relationship between occupancy filter length (no filter, 15 min., 20 min., 30 min., and 120 min.) on the length of occupied periods is then studied to identify a filter length that balances filling in false–negatives in the raw data with the potential for false–positives in the filtered data (§3.1). Although, without ground–truth occupancy data, this can only be done heuristically. Even after choosing a filter length, the impact of the decision on each later analysis should be revisited.

After identifying each SC and then MSC, the time of day and duration until the next SC for each MSC will yield insights into schedule–based causes and energy impacts of user overrides (§3.2). Studying the statistics of the TTOs for all the MSCs will show how dynamic user behavior is, and potentially question the efficacy of static ASHRAE comfort models for GEBs (§3.3).

The events (e.g., away, awake, home, etc.) that lead to the SC preceding each MSC will yield insights into the relationship between the user and the thermostat, and potentially user mental models (§3.4). The DoO for each MSC will yield further insights into the energy

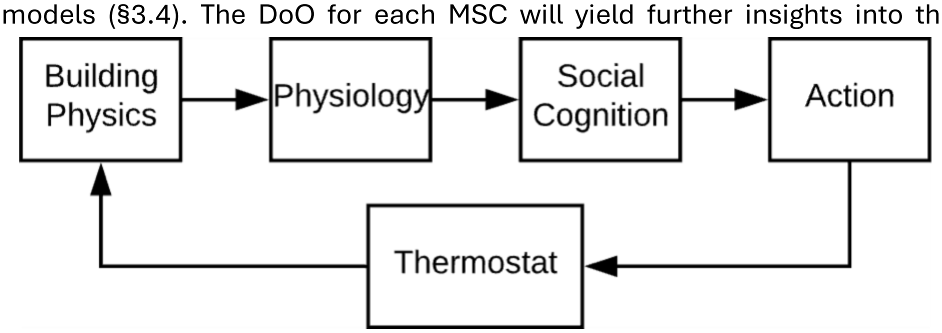

Figure 6: Human–in–the–loop model of a building

impacts of users' decisions based on the relative number of MSCs where setpoints are increased or decreased during heating or cooling.

The core analysis of this paper is the relationship between TTO and DoO and statistical quantification of the "time to override" given the "degree of override" (§3.5).

2.5 Results and discussion

2.5.1 Pre–processing

Figure 7 shows the effect of the filter length on the detected periods of occupancy. The large number of short periods in the raw data indicates many false negatives that split up extended periods when the occupants are home (e.g., a few hours at home after work and before bed). As the filter length increases to 10 min., 20 min., and then 30 min., more and more of these broken apart long periods are stitched together, increasing the number of longer occupied periods suitable for comfort studies by a half an order of magnitude.

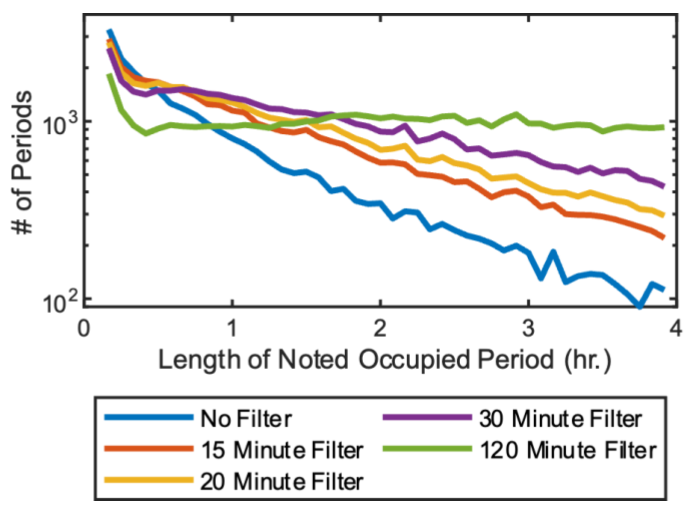

Figure 7: Comparison of different filters on the occupancy data

The 120 min. filter decreases the number of noted occupied periods overall but yie lds more data points with occupied periods greater than 3.5 hrs., which may not be of use as most of the MSCs dynamics are assumed to take place within 2 hrs. of a triggering event. Moreover, while the 120 min. filter does convert false negatives to true positives, a lot of true negatives may also be converted to false positives. In other words, even when the person was not home for most of a 2 hr. period, the entire 2 hrs. were assumed to be occupied, potentially corrupting the data of thermal comfort behavior governed by outdoor temperatures.

In lieu of ground truth or alternate occupancy data, the 30 min. filter seems to serve as the middle ground for false positives and false negatives and has also been used in prior studies [212].

2.5.2 Time of day and duration of overrides

Figure 8 depicts the MSCs throughout the day for single–occupant households and shows distinct periods of frequent overrides. It is helpful to frame this figure with approximations of typical setback schedules programmed into the thermostat: e.g., *sleep* from 10 PM–6 AM, *wake–up* from 6 AM–9 AM, *away* from 9 AM–6 PM, and *evening home* from 6 PM–10 PM during weekdays, and the away period removed during the weekends. The weekday spike in overrides mid–morning could then be early birds setting back their thermostat before leaving early for work or occupants spending the day at home and overriding the away period. The other large spike in overrides occurs in the evening, and may be switching to sleep mode early, or night owls extending their evening comfort zone into midnight. The average occupant manually changes the thermostat 0.9 times per day.

The number and duration of MSCs with respect to the time of day are plotted in a 2–D histogram in Figure 9, and may serve as a proxy for the energy impact of overrides. The MSCs with the longest period of impact were the night owls, who's setbacks lasted 6–12 hrs. overnight. The occupants staying home in the morning also had extensive impacts, lasting 6–12 hrs. throughout the day. The large number of MSCs in the early evening lasted only 1.5–3 hrs., indicative of a dissatisfaction in the automation's ability to set comfortable temperatures during this post–work, pre–sleep period.

2.5.3 Time to override

To understand occupant behavior patterns that *lead to* MSCs, occupants' time to override (TTO) is calculated. I.e., Figure 11 statistically shows how long it takes for occupants, already in the space prior to the SC, to feel uncomfortable and make an MSC.

The majority of MSCs occur within the first 10 mins of an SC, where the occupant immediately notices or predicts the setpoint change. The more interesting cases are the increasing TTOs peaking around 15–20 mins and exponentially trailing off. The plot is limited to 4 hrs. because the data after that is minimal, having a probability density of less than 0.01.

On average, occupants are in the space for 55 mins. prior to overriding the system, however, if the first peak is removed, the average shifts to 72 mins. The median time to override of 15–20 min could be influenced by any of the factors of the building–physiology–cognition–action process and likely varies from person to person and building to building.

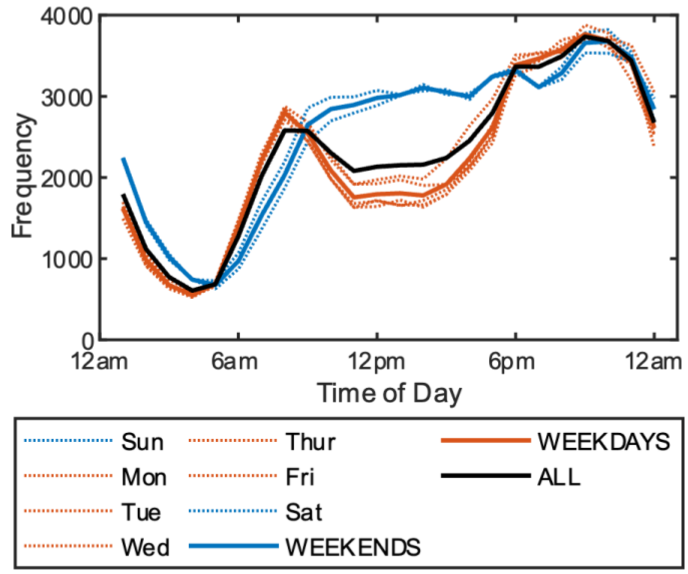

Figure 8: Trends in manual setpoint changes during the day (single–occupant households; n=1,401)

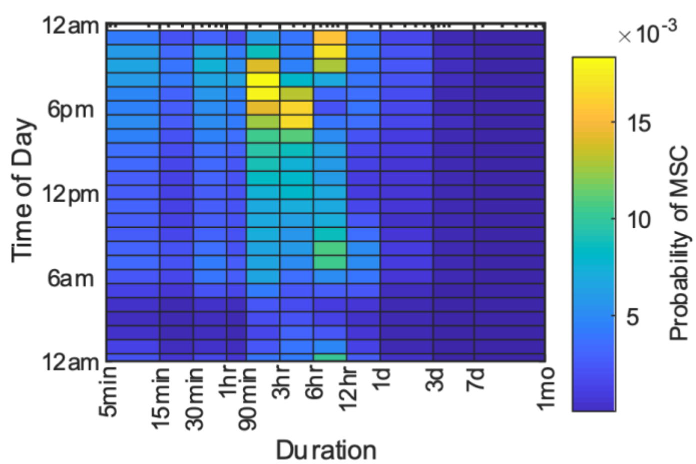

Figure 9: Number and duration of manual setpoint changes as per time of day (single–occupant households)

The distributions of TTO for various occupancy filter lengths are also indicated in the figure. As expected, the unfiltered data shows more TTOs around the 15 min. median, while the 120 min. filter decreases the prevalence of these short TTO and increases the prevalence of longer (>1 hr.) TTOs.

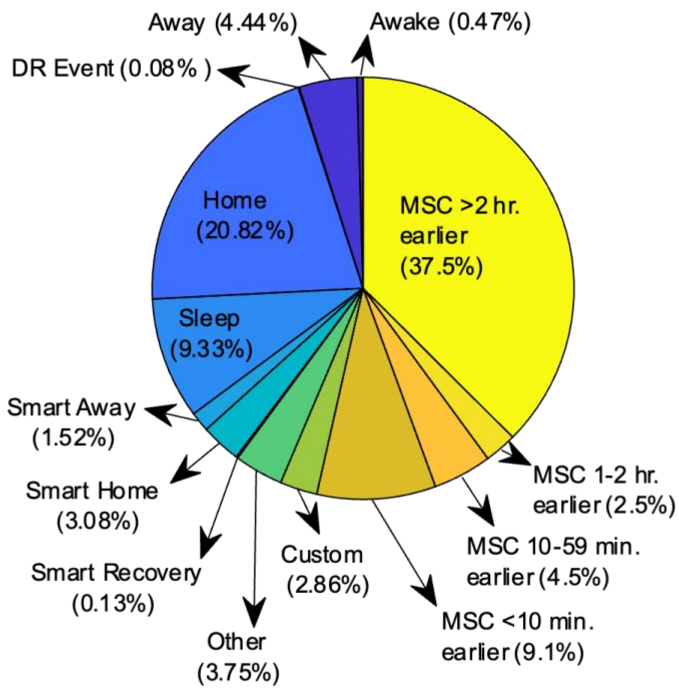

Figure 10: Event immediately prior to a manual setpoint change (single–occupant households)

Figure 10 shows the distribution of events causing the SC prior to each MSC. Event types are defined in §2.1. Over half of all MSC were overriding a previous setpoint by that user, not overriding the automation. It is possible that many of these overrides, some corresponding to the median in Figure 11, could be due to occupant's incorrect *valve theory* mental model. From an automation standpoint, the users are not satisfied with the scheduled–based *Home* and *Sleep* PSCs, preceding ~30% of all MSCs. Users seem to place more trust in the *Smart* features, rarely overriding them; although, this could be due to the low prevalence of these events in the dataset. The dataset has too few *Demand Response* events to draw meaningful insights currently.

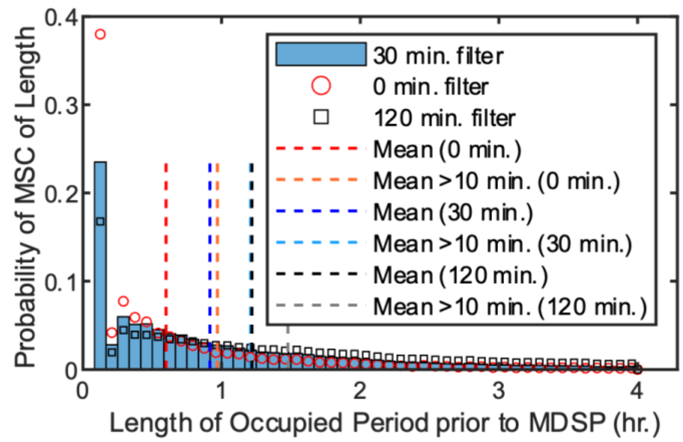

Figure 11: Time from setpoint change to manual setpoint change distribution (occupied; 30, 0, 120 min.)

2.5.4 Mental models

Approximately 20% of users override *more than once in a row* after an initial MSC, exhibiting traits related to the valve theory mental model. The 52% of users who made an override *only once* after an initial MSC seem to be trying to follow the feedback theory mental model but behave non–optimally due to an apparent inability to estimate their own comfort zone. The remaining ~28% of users *make no additional MSCs* within 1 hr. after an MSC and likely utilize a valid feedback theory mental model. This aligns well with prior research showing ~30% valve theory and ~70% feedback theory [200].

2.5.5 Impact of manual setpoint changes

Figure 12 shows the distribution of MSCs' *degree of override (DoO)* (assuming the occupant is comfortable after the override) made in heating and cooling seasons with respect to the indoor temperature. The median setpoint change is ~$2°F$ during the cooling season with an indoor temperature ~74°F, and a $2°F$ increase during the heating season with an indoor temperature ~68°F. In the cooling season, 58% of these MSC increase energy use, while 57% increase the energy use in the heating season. See [157] for a thorough coverage of indoor temperatures and energy impacts in the DyD dataset.

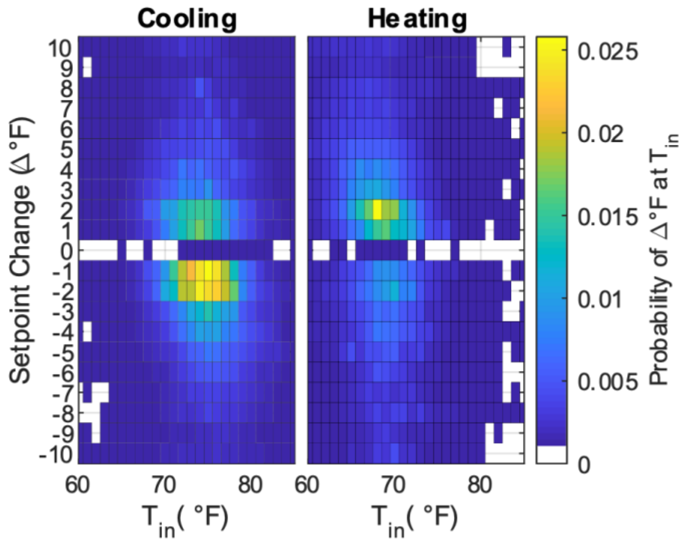

Figure 12: Manual setpoint changes by season (single–occupant households)

2.5.6 Behavior dynamics

One possible confounding factor in this analysis is if setpoint overrides occur because the interior temperature is not at the setpoint. So, the data was analyzed to compare the interior temperature with the setpoint before SC occurred. It was found that **79.37%** of occupied time in the **cooling season** and **76.08%** of occupied time in the **heating season**, the **indoor temperature was within the $\pm 1°F$ accuracy** of the temperature sensor [216]. Therefore, it can be suggested that majority of the time, the building was performing as intended and overrides were not to correct the building's inability to reach the setpoint. Now the motivation to override from PSC to MSC can be based on a variety of factors, most of which the DyD dataset lacks. Window sensors, EMAs, occupant location, etc. also play a vital role in understanding discomfort. The analysis presented in this work aims to find proof for the argument "**Are setpoint data solely enough to model thermostat override dynamics?**", it would be our goal to leverage the features mentioned above to improve the developed model.

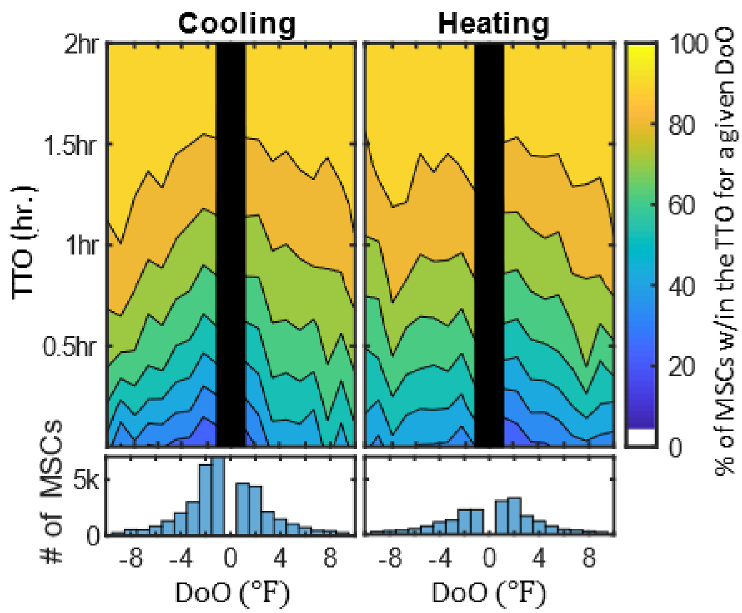

Figure 13: Degree of override vs Time to override; during occupied periods (Single–occupant households; 30 min. filter)

$$TTO = 1.553 * \text{DoO} + 105.1 \tag{2}$$
$$TTO = -1.107 * \text{DoO} + 98.89 \tag{3}$$
$$TTO = 0.5368 \cdot e^{-0.083 \cdot DOO} \tag{4}$$
$$TTO = 0.5804 \cdot e^{-0.074 \cdot DOO} \tag{5}$$

Assuming the occupant becomes aware of the PSC (either through perceiving changes i n indoor temperature, recalling when PSCs are regularly scheduled, noticing a change on the thermostat display), can one find a relationship between **TTO**, how long an occupant (who is at home) takes to feel uncomfortable after any kind of a setpoint change (a PSC or an MSC) and makes an MSC; and **DoO**, the change in temperature of the MSC? The relationships between these variables could guide DRPs that setback thermostats to temperatures that would minimize overrides for the desired period, thus increasing DRP reliability.

Figure 13 shows the decreasing relationship between TTO and increasing DoO. In the cooling season, ~28% of MSCs of –2°F took place within 10 mins., 50% took place within 30 mins., and 60% took place within 40 mins.

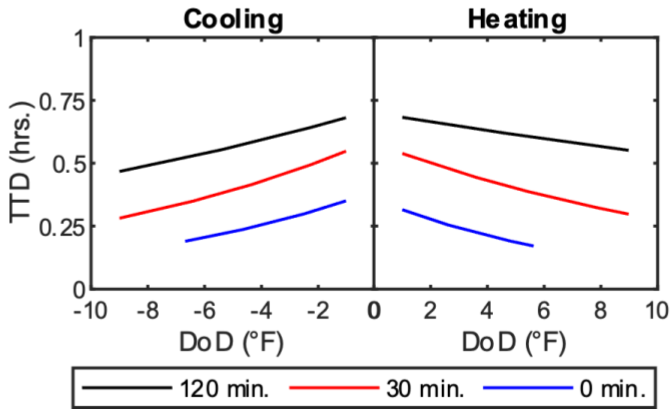

Figure 14: Effect occupancy filter on degree of override and time to override

TABLE 1: MODEL ERROR COMPARISON		
Fit	**Heating RMSE (mins.)**	**Cooling RMSE (mins.)**
Constant	61.57	64.24
Linear	59.14	63.13
Exp	54.80	57.90

In Table 1, we compared the error of three models for each negative integer °F DoO during cooling season, and positive integer °F DoO during heating season, respectively (i.e., MSCs that would result in an increase in energy use). The constant fit depicts the current DRP models that assume same setback duration for all setback degrees, 90–minutes setback duration yields the least error out of the durations of 0, 30, 60, 90, 120, and 240 minutes. The linear model (Equations (2) and (3)), assumes the feedback of a human to be linear with respect to the TTO and DoO. Finally, the exponential model (Equations (2) and (3)) assumes that there are underlying dynamics that trigger an occupant to override. Although there is relatively high error with the exponential model ~1 hour but the improvement in accuracy shows the need to consider feedback dynamics as compared to static or no model usage. With the consideration of contextual and personal features, the error can be further decreased.

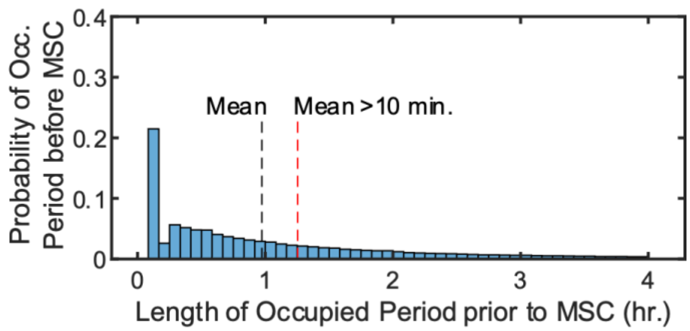

Figure 15: Entire dataset: Time from setpoint change to manual setpoint change (occupied; 30 min. filter)

The close similarity between energy intensive behavior dynamics in heating and cooling seasons contrasts with the difference in energy intensive behavior versus energy saving behavior. Energy saving MSCs (i.e., those triggered by positive DoOs during cooling season, and negative DoOs during heating season) have significantly faster TTOs compared to energy intensive MSCs. These results reinforce observations from [213] that increasing temperatures are sensed at different rates than decreasing temperatures, yet the results presented here suggest that difference could be seasonally affected.

It should be noted that the pre–processing done for this analysis, using a 30 min. filter, plays an important role in determining the TTO. Figure 14 shows the relationship between DoO and TTO for different occupancy filters. As expected, the TTO increases as the occupancy filter length increases, due to more data of longer occupied periods entering the MSC dataset used for analysis. Moreover, the number of MSCs continuously occupied following the previous SC (within 2 hrs.) increases with the length of the filter. E.g., the number of 1–2°F MSCs for no filter, 30 min. filter, and 120 min. filter are ~4k, 7k, and 10k respectively.

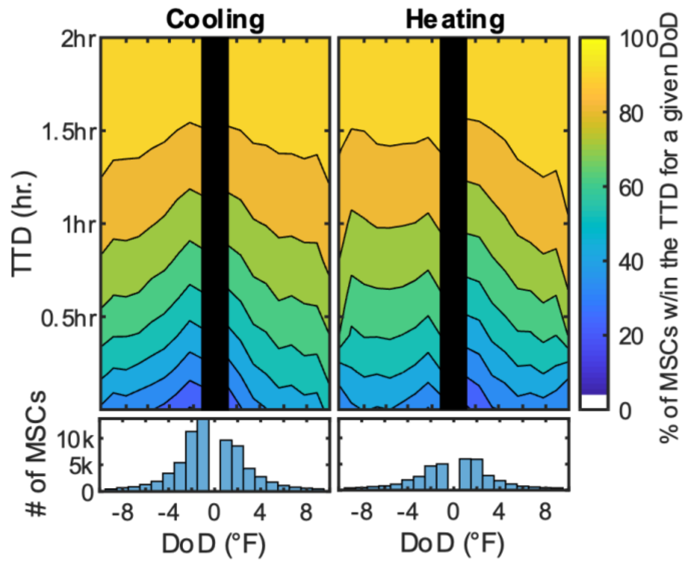

Figure 16: Entire dataset: Degree of override vs Time to override

2.5.7 Entire dataset

The results above are computed using the subset of the DyD dataset self–reported as having only one occupant. This provided insights into the behavior of individuals, removing conflating factors due to social barriers to comfort. Multiple–occupant homes will lead to long occupancy periods because different people move in and out of the home at different times. The dynamics of thermal comfort will then be based on data for different occupants with entirely different metabolisms, clothing, behavior patterns, etc. To understand the impact of these social factors, the methods above are applied to the entire DyD dataset, which consists of ~27k homes (including homes that users report as "single occupant" and "multiple occupants," as well as no reported category).

Comparing Figure 15, the distribution of TTO for the entire dataset, with Figure 13, the distribution for single–occupant households only, shows similar thermostat behavior regardless of social factors. The largest number of MSCs occur within 10 min. of the previous SC (24% and 21% for single occupancy and all data, respectively), and the next median TTO occurs within 15–20 min. The mean of the data after 10 min. is essentially unchanged. Likewise, comparing the event leading to the MSCs for single–occupant households (i.e., Figure 10) and all data would show similar miniscule changes.

Like Figure 13, the contour lines in Figure 16 shows the relationship between TTO and DoO. The contours are smoother for Figure 16 using the entire dataset, due to the availability of more MSC data points. Table 1 compares key values between Figure 13 and Figure 16, namely the period after which 50% of the population will have become uncomfortable and trigger an MSC (i.e., TTO–50%) for low and high magnitudes of overrides (i.e., DoO). The similarity in these numbers seems to indicate that social factors have little appreciable impact on thermal comfort behavior dynamics as observed through thermostat changes.

TABLE 2: COMPARISON OF DOO VS. TTO–50% FOR SINGLE VS ALL OCCUPANCY DATA

Heating	$2°F$ DoO	$8°F$ DoO
Single occ.	28 min.	15 min
All data.	30 min.	17 min.
Cooling	–$2°F$ DoO	–$8°F$ DoO
Single occ.	28 min.	15 min.
All data.	29 min.	18 min

2.6 Conclusions

This paper analyzed ecobee's "Donate Your Data" dataset of ~27k smart thermostats to illustrate and understand the dynamic nature of occupant thermal comfort behavior, especially relating to thermostat use. The analysis focused on the subset of data – self–declared as single occupant homes to best understand individual behavior; however, extending the methods to the entire dataset did not substantially change the results. Each manual setpoint change (MSC) in the data was identified and those which were occupied continuously following the previous setpoint change (SC) were considered for analysis. The analysis then assumed that the previous SC led to the occupants' discomfort and the MSC set the thermostat to a comfortable temperature.

Programmed setpoint changes (PSCs) accounted for less than half the events preceding the MSCs, showing that many occupants have not programmed their thermostat or are overriding decisions they previously made while continuously occupying the space. This later behavior, an apparent iterative process of finding comfortable setpoints, supports the finding in other literature that ~30% of occupants have an incorrect *valve theory* mental model of how a thermostat works [200]. About 60% of the MSCs in both heating and cooling modes resulted in increased energy consumption.

The principal result of this analysis captures the input–output relationship of the building–physiology–cognition–action process that leads to an MSC correcting discomfort caused by a previous SC. An understanding of this process will be important for designing adaptive and personalized grid interactive efficient buildings (GEBs). The time it takes for an occupant to notice and correct a setpoint is often nearly instantaneous (if they predict the thermostat schedule or the hear/see the thermostat change), but the majority of MSCs occur some occupied period after the SC. This time to override (TTO) is inversely exponentially related to the degree of override (DoO): in both heating and cooling mode, 50% of occupants will take less than ~30 min. to correct an SC that makes them $2°F$ too cool/warm, but less than ~15 min. for being $8°F$ too cool/warm respectively.

Prior research showed that load curtailment is most effective for the first 20 – 30 minutes and the energy use increases as the HVAC system reaches steady state [211]. The models developed in this work and presented in equations (2) and (3) suggest that setback schedules could be developed to implement DR programs of any duration depending on the desired probability of overrides. As the timer reaches the TTO where there is high probability of an override, the setpoint could be changed back to the occupant's preferred setpoint. Further accuracy improvements could be made by creating personalized models of thermal comfort behavior dynamics for deploying personalized DRPs that maximize curtailment and reliability. Further research is required to truly understand these phenomena however, as this behavior is an outcome of a complex multi–disciplinary feedback loop encompassing building physics, human physiology, cognition, and decision making.

2.7 Acknowledgements

The authors would like to acknowledge their gratitude to ecobee's Social Impact Department for curating and sharing the Donate Your Data dataset that is foundational to this research. Other researchers may request no–cost access to this data by visiting https://www.ecobee.com/donateyourdata/signup/

The authors appreciate David Fannon's suggestions on the manuscript and the figures.

The authors gratefully acknowledge editing services provided by Susan Matheson for the preparation of this manuscript.

This research did not receive any specific grant from funding agencies in the public, commercial, or not–for–profit sectors.

This material is based upon work supported by the U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy, Building Technologies Office under Award Number DE–EE0009154.

This work was conducted within the framework of IEA–EBC Annex 79.

This material is based upon work supported by the National Science Foundation under Grant No. 1928608.

Chapter 3:

The three theories of discomfort

A data–driven assessment of thermal comfort models

3.1 Abstract

Challenges with existing thermal comfort models for building performance simulation (BPS) include their static nature, which leads to poor forecasting of thermostat overrides during transient conditions. This is becoming a more critical issue as demand response (DR) programs increase in prevalence. In this paper, three theories of thermal comfort – Delayed Response Theory (DRT), Comfort Zone Theory (CZT), and Thermal Frustration Theory (TFT) – were used to train machine learning (ML) models to identify critical occupant–specific parameters that improved the classification accuracy of overrides. The methodology developed can be used to design controllers for utilities or energy companies to reduce overrides and achieve committed predicted energy savings during DR events.

3.2 Introduction

The building sector accounts for approximately 40% of global energy consumption and greenhouse gas emissions [217]. Improving the energy efficiency of buildings is crucial for mitigating climate change and enhancing sustainability [6]. Building Performance Simulation (BPS) tools play a vital role in designing energy–efficient buildings and optimizing their operation [218]. However, the accuracy of BPS predictions relies heavily on the assumptions made about occupant behavior and thermal comfort [219].

Current thermal comfort models used in BPS tools are largely static and do not adequately capture dynamic occupant behaviors and preferences [220]. The most widely used thermal comfort model, the Predicted Mean Vote (PMV) model developed by Fanger [133], assumes steady–state conditions and uniform occupant preferences. However, numerous studies have shown that occupant thermal comfort is highly variable and depends on factors such as age, gender, culture, and personal control [221], [222], [223], [224]. Ignoring these variations can lead to inaccurate predictions of energy use and occupant satisfaction [14].

The limitations of static thermal comfort models become even more pronounced during demand response (DR) events, when indoor conditions may vary significantly from typical conditions [225]. DR programs aim to reduce peak electricity demand by incentivizing buildings to temporarily reduce their energy consumption [226]. These programs are typically designed using static assumptions about occupant comfort zones and acceptable temperature ranges [227]. However, studies have shown that occupants frequently override thermostat setpoints during DR events, reducing the achieved energy savings [228], [229].

The development of dynamic thermal comfort models has been limited by the lack of granular data on occupant behavior and thermal preferences [230]. Traditional methods of collecting thermal comfort data, such as surveys and laboratory studies, are time–consuming, expensive, and may not capture real–world behavior [231]. In recent years, the proliferation of smart thermostats and other Internet of Things (IoT) devices has provided new opportunities for collecting high–resolution data on occupant behavior and building performance [232], [233], [234].

Several studies have used smart thermostat data to develop data–driven thermal comfort models. For example, Kim [138] used machine learning techniques to predict individual thermal preferences based on occupant heating and cooling behavior. Kim [140] reviewed studies where personalized thermal comfort models were developed using data from a wearable device and a smart thermostat. However, these studies focused on steady–state conditions and did not consider the dynamic behavior of occupants during DR events.

In this study, we leverage ecobee's Donate your Dataset (DyD), which contains high–resolution data on residential thermostat settings and indoor conditions [203], to derive occupant–specific parameters for three different theories of thermal comfort during DR events: Delayed Response Theory (DRT), Comfort Zone Theory (CZT), and Thermal Frustration Theory (TFT). These theories represent different hypotheses about how occupants respond to thermal discomfort during DR events.

DRT assumes that occupants will override the DR setpoint after a certain time delay, regardless of the indoor temperature. CZT assumes that occupants have a specific comfort zone and will override the setpoint when the temperature exceeds a certain threshold. TFT, which we introduce in this study, assumes that occupants accumulate thermal frustration when the temperature is outside their comfort zone and will override the setpoint when their frustration exceeds a threshold.

We use the occupant–specific parameters derived from these theories to train logistic regression models to classify thermostat overrides. The results demonstrate the potential for dynamic thermal comfort models to improve predictions of occupant overrides and energy use, enabling more effective DR program design and control strategies. To the best of our knowledge, this is the first study to compare different thermal comfort theories for predicting occupant behavior during DR events using real–world smart thermostat data.

The rest of this paper is organized as follows: Section 2 describes the methodology, including the data preprocessing, occupant modeling, and machine learning techniques used. Section 3 presents the results of the logistic regression models for each thermal comfort theory. Section 4 discusses the implications of the results for DR program design and control strategies, as well as the limitations of the study and future research directions. Section 5 concludes the paper.

3.3 Methodology

In this study, we propose a data–driven approach to assess three theories of thermal comfort – Delayed Response Theory (DRT), Comfort Zone Theory (CZT), and Thermal Frustration Theory (TFT) – using time–series thermostat data and machine learning techniques to identify key occupant–specific parameters and improve the prediction of manual setpoint changes (MSCs) during demand response events.

3.3.1 The three theories of thermal comfort

3.3.1.1 Delayed response theory (DRT)

The Delayed Response Theory (DRT) proposes that an occupant will override the thermostat settings after a specific time delay, regardless of the indoor temperature or their personal comfort temperature. This theory assumes that the occupant's decision to make a manual setpoint change (MSC) is primarily influenced by the elapsed time since the last setpoint change, rather than the actual thermal conditions.

In the DRT model, the occupant's behavior is characterized by a single parameter: the time delay (τ) measured in minutes. The model predicts that the occupant will make an MSC when the elapsed time since the last setpoint change exceeds this time delay. The occurrence of an MSC at a specific time k can be determined using the following condition:

$$MSC[k] = 1, if \ \left(t[k] - t_{last_{setpoint_{change}}}\right) > \tau \tag{6}$$
$$MSC[k] = 0, otherwise$$

Where $MSC[k]$ is a binary variable indicating whether an MSC occurs at timestep k, $t[k]$ is the current time, $t_{last_{setpoint_{change}}}$ is the time of the last setpoint change, and τ is the time delay parameter. When the elapsed time since the last setpoint change surpasses the time delay τ, the model predicts an MSC occurrence ($MSC[k] = 1$).

The DRT model assumes that the occupant's decision to make an MSC is independent of the actual thermal conditions, such as the indoor temperature or their personal comfort temperature. This implies that the occupant will override the thermostat settings at regular intervals, determined by the time delay parameter, regardless of whether they are experiencing thermal discomfort or not.

One potential explanation for this behavior is that the occupant may have a habitual or routine–based approach to thermostat adjustments. For example, they may have a tendency to check and adjust the thermostat settings at specific times of the day, irrespective of the actual thermal conditions. Alternatively, the occupant may have a limited tolerance for prolonged periods without control over their thermal environment, leading to periodic MSCs.

To determine the optimal time delay parameter (τ) for a given occupant, researchers can analyze historical thermostat data and identify the typical time intervals between consecutive MSCs. By comparing the predicted MSC occurrences based on the DRT model with the actual MSC events, the accuracy of the model can be assessed, and the most suitable time delay parameter can be estimated.

3.3.1.2 Comfort zone theory (CZT)

The Comfort Zone Theory (CZT) postulates that an occupant has a specific temperature range within which they feel comfortable, known as the comfort zone. When the indoor temperature deviates from this comfort zone, similar to the adaptive thermal comfort model, the occupant is likely to override the thermostat settings to restore their comfort [136]. This theory assumes that the occupant's decision to make a manual setpoint change (MSC) is based on the instantaneous temperature relative to their comfort zone boundaries [235].

In the CZT model, the occupant's comfort zone is defined by a lower temperature threshold ($T_{C_{low}}$) and an upper–temperature threshold ($T_{C_{up}}$). The occurrence of an MSC at a specific time k can be determined using the following conditions:

$$MSC[k] = 1, if \left(T_{in}[k] < T_{C_{low}}\right) or \left(T_{in}[k] > T_{C_{up}}\right) \tag{7}$$
$$MSC[k] = 0, otherwise$$

Where $MSC[k]$ is a binary variable indicating whether an MSC occurs at the timestep k, $T\_in[k]$ is the indoor temperature at the timestep k, $T_{C_{low}}$ and $T_{C_{up}}$ are the lower and upper comfort temperature thresholds in °F, respectively. When the indoor temperature falls outside the occupant's comfort zone, the model predicts an MSC occurrence ($MSC[k] = 1$).

The CZT model assumes that the occupant's decision to make an MSC is based solely on the current thermal conditions and does not consider the duration of exposure to discomfort or the magnitude of temperature deviation from the comfort zone boundaries [236]. This implies that the occupant will immediately override the thermostat settings when the indoor temperature crosses either of the comfort thresholds [237].

Unlike the Thermal Frustration Theory (TFT), which considers the accumulation of discomfort over time, the CZT model focuses on the instantaneous temperature deviations from the comfort zone. The occupant's comfort zone parameters ($T_{C_{low}}$ and $T_{C_{up}}$) are assumed to be static, meaning that the occupant is expected to consistently make MSCs whenever the indoor temperature falls outside their comfort zone [238].

The CZT model can be used to predict the occurrence of MSCs based on the real–time monitoring of indoor temperature and the occupant's predefined comfort zone boundaries. By comparing the predicted MSC occurrences with the actual MSC events, the accuracy of the model can be assessed, and the most suitable comfort zone parameters can be estimated [133].

While the CZT model offers a simple and intuitive approach to predicting MSC behavior, it may not capture the complexity of occupants' thermal comfort preferences and decision–making processes. Comparing the performance of the CZT model with other thermal comfort models, such as the Thermal Frustration Theory (TFT) and the Delayed Response Theory (DRT), can provide insights into the relative importance of instantaneous temperature deviations versus time–based factors in determining occupants' thermostat override behavior during demand response events [239, p. 55].

3.3.1.3 Thermal frustration theory (TFT)

The TFT theory hypothesizes that thermal discomfort accumulates over time, prompting the user to make an MSC when the discomfort reaches a certain level. This differs from conventional approaches like the Adaptive Model, which assumes people are comfortable within a specific temperature range based on outdoor conditions. However, in the context of Demand Response (DR) programs, which involve changing conditions, occupants may be willing to tolerate temporary discomfort even when they are outside the thermal comfort zone. Therefore, the discomfort–based model focuses on tracking the buildup of thermal discomfort until it hits a set threshold, at which point an MSC is triggered. This novel thermal discomfort accumulation model is referred to as the Thermal Frustration Theory (TFT).

The Thermal Frustration (TF) at a specific time k is calculated as:

$$\mathrm{TF}[\mathrm{k}] = (\alpha \times \mathrm{TF}[\mathrm{k}-1] + \beta \times (\mathrm{T_{in}}[\mathrm{k}] - \mathrm{T_C})\ ) \times \mathrm{motion} \tag{8}$$

$$MSC[k] = 1, if\ TF[k] > TF_{\lambda}$$

Where, TF represents the level of thermal frustration in °Fmins, k is the current timestep, α is the coefficient that captures how much of the past thermal discomfort is remembered, β is the coefficient that captures how much of the current thermal discomfort is being felt in °Fmins, and "motion" is a Boolean variable indicating whether movement has been detected leading to 0 thermal frustration when the occupant is not present in the space of motion measurement. The current level of discomfort is the difference in the current indoor temperature T_{in} and the comfort temperature T_C in °F. $MSC[k]$ is a binary variable indicating whether an MSC occurs at the timestep k if the thermal frustration at k timestep is greater than the thermal frustration threshold.

The TFT model in (8), α serves as a factor that quantifies how much of the previously experienced thermal discomfort is retained or 'remembered' over time and hence can be in range of 0 and 1. β, on the other hand, gauges the degree of discomfort experienced during a given timestep, which could potentially vary based on factors such as metabolic rate or clothing and is also in the range of 0.0 and 1.0 minutes. Moreover, motion here serves as the reset condition, forcing the thermal frustration value to zero when the occupant is not within the space as for this paper, instances where the occupant is not exposed to external factors like outside temperature/humidity or other events is not considered.

Assuming that an occupant's α,β and T_C parameters are static in nature, one should expect to a thermal frustration due to discomfort based overrides or MSC_D to be the same for an occupant. One can define this consistent thermal frustration at MSC to be a thermal frustration threshold or TF_λ.

3.3.2 Discomfort model parameter estimation

Time–series thermostat data can be leveraged to identify key parameters of the three discomfort models discussed in Section 3.1: Thermal Frustration Theory (TFT), Comfort Zone Theory (CZT), and Delayed Response Theory (DRT). To estimate the personalized parameters for each model, a grid search procedure can be employed in combination with logistic regression (LR) models.

3.3.2.1 Grid search procedure

The grid search procedure involves iterating over a range of values for the key parameters of each discomfort model and evaluating the performance of a logistic regression (LR) model trained on the corresponding features using the Matthews correlation coefficient score (MCC). The steps of the grid search procedure are as follows:

1. Thermal Frustration Theory (TFT):
   a. Initialize the parameter of interest:
      i. $\boldsymbol{\alpha}$: Represents the weight of historical discomfort.
      ii. $\boldsymbol{\beta}$: Represents the weight of current thermal discomfort.
      iii. $\boldsymbol{TF_\lambda}$: Represents the threshold at which thermal frustration leads to a manual setpoint change (MSC).
   b. Compute thermal frustration values using the current parameter values (e.g., calculate cumulative discomfort).
   c. Identify the row indices where the ground truth value for an MSC is `True`.
   d. Train an LR model on both MSC and non–MSC data subsets using the calculated features. Apply the model to predict MSC occurrences on a test

dataset and compute MCC and accuracy scores. Log the parameter values along with the corresponding scores.

e. Increment the parameter values according to a predefined step size and repeat the above steps until the entire parameter range has been explored.

2. Comfort Zone Theory (CZT) Parameters:
   a. Initialize the parameters of interest:
      i. $\boldsymbol{T_{C_{low}}}$: Lower comfort temperature threshold.
      ii. $\boldsymbol{T_{C_{up}}}$: Upper comfort temperature threshold.
   b. Compute temperature deviations based on the current parameter values (e.g., deviations from comfort zone limits).
   c. Identify the row indices where the ground truth value for an MSC is `True`.
   d. Train an LR model on both MSC and non–MSC data subsets using the calculated features. Apply the model to predict MSC occurrences on a test dataset and compute MCC and accuracy scores. Log the parameter values along with the corresponding scores.
   e. Increment the parameter values according to a predefined step size and repeat the above steps until the entire parameter range has been explored.
3. Delayed Response Theory (DRT) Parameters:
   a. Initialize the parameter of interest:
      i. $\boldsymbol{\tau}$: Represents the time delay before the occupant makes an MSC.
   b. Compute the elapsed time since the last setpoint change using the current parameter value.
   c. Identify the row indices where the ground truth value for an MSC is `True`.
   d. Train an LR model on both MSC and non–MSC data subsets using the calculated features. Apply the model to predict MSC occurrences on a test dataset and compute MCC and accuracy scores. Log the parameter values along with the corresponding scores.
   e. Increment the parameter value according to a predefined step size and repeat the above steps until the entire parameter range has been explored.

The choice of logistic regression (LR) models for MSC classification is strategic due to several reasons. Firstly, LR assumes a linear relationship between the predictor variables (e.g., thermal frustration, temperature deviations, elapsed time) and the log–odds of a discomfort–based override. Secondly, the probabilistic outputs of LR models can be leveraged to compute the likelihood of an MSC, providing a measure of prediction certainty for each model. This approach allows for systematic parameter tuning and performance assessment across different thermal comfort models.

3.3.2.2 Logistic regression model

The logistic function used in LR models can be represented as:

$$p(x) = \frac{1}{\left(1 + e^{\frac{-(x-\mu)}{s}}\right)} \tag{9}$$

Where p(x) is the probability that the dependent event occurs (e.g., MSC=1), x is the independent variable (e.g., thermal frustration, temperature deviations, elapsed time), μ is the location parameter (p(μ)=1/2, which shifts the function along the x–axis), and s is the scale parameter, which determines how "spread out" or "steep" the function is. A small s will make the transition from 0 to 1 more abrupt, while a larger s will make the transition smoother.

During the training process, initial guesses for μ and s are made (e.g., μ initialized to the mean of x values, s initialized to the standard deviation). The model then aims to maximize the likelihood function for a given set of μ and s using optimization algorithms like gradient descent. The best–fit values for μ and s are then used as the parameters of the logistic function for making future predictions.

3.3.2.3 Routine model

Occupants often interact with thermostats based on their daily routines, such as adjusting the setpoint when waking up, leaving for work, returning home, or going to bed [240]. These routine–based manual setpoint changes (MSCs) are not necessarily driven by thermal discomfort but rather by habitual behaviors [241]. To accurately estimate the parameters of the discomfort models, it is essential to differentiate between routine and discomfort–driven MSCs [242].

In this study, we develop a routine model based on the time of day (TOD) to capture the probability of an MSC occurring at specific times. The model assumes that routine–based

MSCs follow a consistent daily pattern [154]. By analyzing the historical MSC data, we can compute the probability mass function (PMF) of TOD for MSCs.

Let T be the random variable representing the time of an MSC. The PMF of T can be estimated using the following equation:

$$P(T = k) = \frac{n_k}{N} \tag{10}$$

Where $P(T = k)$ is the probability of an MSC occurring at timestep k, n_k is the number of MSCs observed at timestep k, and N is the total number of timesteps in the dataset.

Once the PMF is learned, each instance in the dataset can be assigned a probability of being a routine based MSC (P_{MSC_R}) based on its timestamp. For example, if an MSC occurs at a timestep k with a high $P(T = k)$, it is more likely to be a routine–based MSC [243].

The routine model serves as a complement to the discomfort models (TFT, CZT, and DRT) by helping to identify instances where the probability of a routine–based MSC is higher than the probability of a discomfort–driven MSC. These instances can then be excluded from the training and testing datasets used for estimating the parameters of the discomfort models, potentially improving their accuracy [244].

3.3.2.4 Improving discomfort model parameter estimation

To improve the accuracy of the discomfort model parameter estimation, a two–step training process can be employed (Figure 3). The goal is to reduce the noise caused by routine–based MSCs by removing instances that have a high likelihood of being routine–driven.

The first step involves training a suboptimal routine model based on the time of day (TOD) to compute the probability mass function (PMF) of TOD for MSCs. Each instance is then assigned a probability of being a routine based MSC (P_{MSC_R}) based on the learned PMF.

The second step focuses on the identification of the discomfort model. Initially, a grid search is performed to find the optimal parameter value(s) for the given discomfort model, as described earlier. However, the training and testing data at this stage may include routine, discomfort, and other reason based MSCs, leading to suboptimal results.

To improve the discomfort model identification accuracy, the probabilities of routine (P_{MSC_R}) and discomfort (P_{MSC_D}) overrides are compared for each instance. The LR–based discomfort classifier is then retrained on the data where $P(MSC_D) > P(MSC_R)$. By focusing on instances with a higher likelihood of being discomfort–driven, the optimal parameter value(s) and the corresponding MCC score are expected to improve.

This two–step training process can be applied to all three discomfort models (TFT, CZT, and DRT) to refine the parameter estimations by separating routine and discomfort based MSCs. By leveraging these methodologies, we aim to identify the key parameters that best describe occupants' discomfort–driven behavior and enhance the accuracy of predicting manual setpoint changes during demand response events. The grid search was parallelized using the multiprocessing library in Python to utilize all available CPU cores and speed up the computation. The results for each home were stored in a DataFrame and saved to a file for further analysis.

3.4 Results

Logistic regression models were fit for each thermal comfort theory, with and without accounting for routine setpoint changes. The results of this study are presented in two parts: the comparison of model performance with and without routine overrides, and a detailed assessment of the Area Under the Curve (AUC) differences between pairs of models.

3.4.1 Distribution of theory parameters

These figures depict the distribution of key parameters for each thermal comfort theory—Delayed Response Theory (DRT), Comfort Zone Theory (CZT), and Thermal Frustration Theory (TFT)—and reveal critical insights into the dynamic interplay between occupant comfort and thermostat overrides. Figure 17 shows the distribution of the estimated "time to override" parameter for DRT across multiple homes. The variability in this parameter highlights the distinct differences in occupants' tolerance levels to thermal discomfort during DR events. Figure 18 shows the distribution of estimated comfort temperature parameters for the CZT model, which defines individual comfort zones using upper and lower temperature bounds. Figure 19 presents the estimated alpha parameter for the TFT model, which represents the "memory" of past thermal discomfort.

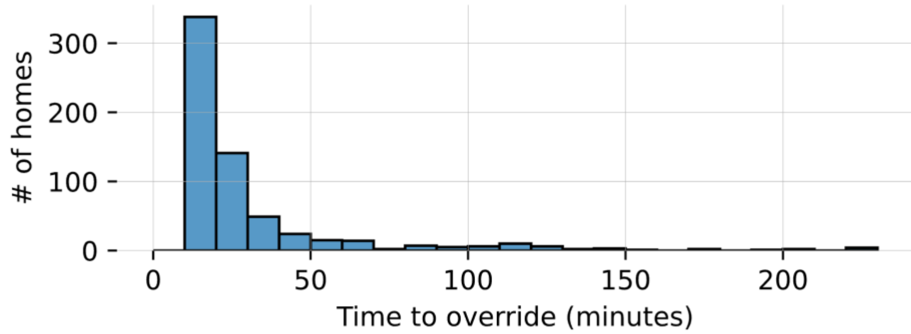

Figure 17: Distribution of estimated time to override parameter of the DRT theory (highest MCC scores for each home).

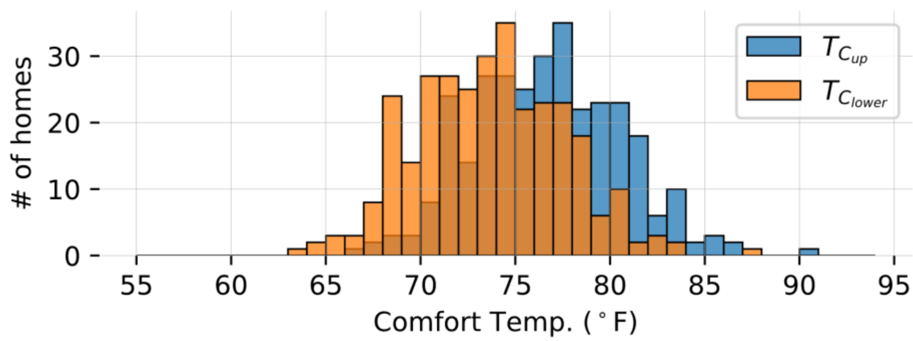

Figure 18: Distribution of estimated comfort temperature parameter of the CZT theory (highest MCC scores for each home).

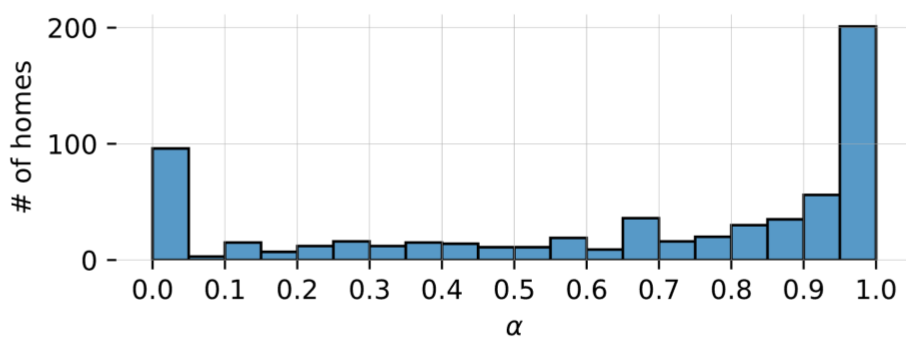

Figure 19: Distribution of estimated alpha parameter of the TFT theory (highest MCC scores for each home).

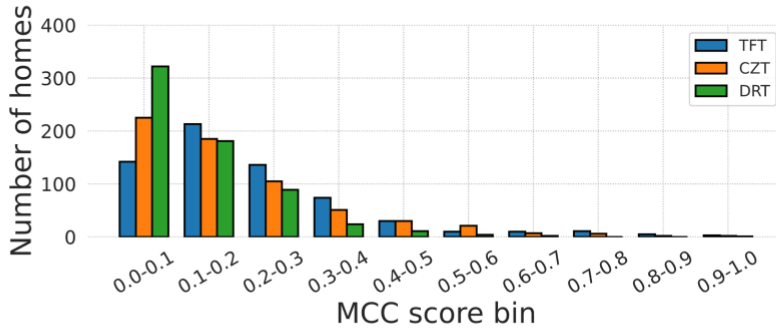

Figure 20: Distribution of MCC scores when including routine overrides

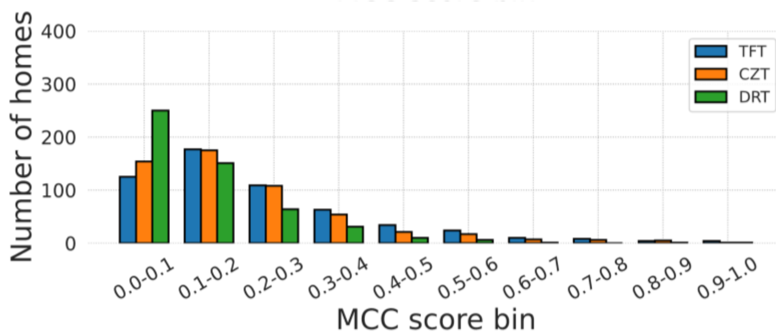

Figure 21: Distribution of MCC scores when excluding routine overrides.

3.4.2 Performance with and without routine overrides

The distribution of MCC scores for the three models—Thermal Frustration Theory (TFT), Comfort Zone Theory (CZT), and Delayed Response Theory (DRT)—is shown in histograms in Figure 20 and Figure 21. The first histogram represents scores with routine overrides included, while the second histogram shows scores without routine overrides.

With routine overrides:

1. TFT achieved MCC scores mostly between 0.1 and 0.3, with some homes reaching higher scores.
2. CZT had a broader spread, with scores ranging from 0.0 to 0.5, peaking around 0.1.
3. DRT had a significant distribution, but with many homes scoring between 0.0 and 0.2.

Without routine overrides:

1. TFT maintained its scores, with a slightly increased spread, and a peak at 0.1.
2. CZT's distribution shifted towards higher scores, peaking around 0.1–0.2.
3. DRT's scores improved, showing a stronger presence around 0.1–0.3.

Universal MCC score per theory (universal meaning false negatives, false positives, true negatives, and true positives of each home was summed to compute MCC score on aggregate level instead of computing MCC score individually and then averaging it out):

Table 3: MCC and AUC score comparison

	With Routine		Without Routine		% change	
Theory	MCC Score	AUC	MCC Score	AUC	MCC Score	AUC
DRT	0.077	47.25	0.098	38.95	27.69	12.41
CZT	0.090	44.35	0.11	47.09	23.90	6.20
TFT	0.064	45.85	0.096	49.40	51.12	7.75

With routine overrides, TFT and CZT show an AUC difference of 0.259, while TFT and DRT exhibit a difference of 0.552. These relatively small differences indicate that all three models perform comparably when routines are included, with TFT slightly outperforming both CZT and DRT. Additionally, CZT and DRT have an AUC difference of 0.293, reinforcing their similar capabilities in predicting manual setpoint changes (MSCs).

Without routine overrides, the differences between models become more pronounced. The AUC difference between TFT and CZT increases to 0.517, suggesting that TFT more accurately captures discomfort–driven behaviors when routines are excluded. The gap between TFT and DRT widens substantially to 1.966, highlighting TFT's significantly better performance in predicting discomfort–driven MSCs. The difference between CZT and DRT also increases to 1.448, indicating that DRT's performance is particularly impacted by the exclusion of routine overrides.

3.5 Discussion

The results demonstrate that accounting for occupant–specific thermal comfort parameters and dynamic behavior significantly improves predictions of overrides during DR events. The Thermal Frustration Theory (TFT) outperformed the other models, likely because it captures the time–integral effect of discomfort, rather than just instantaneous temperature deviations.

Figure 18 shows that the broad range observed in comfort thresholds across different households signifies the inadequacy of static, population–based thermal comfort models.

The variability indicates that static DR setpoints might lead to frequent occupant dissatisfaction and thermostat overrides, thereby undermining DR program objectives. CZT's focus on instantaneous temperature deviations provides a straightforward approach to override prediction, but the results here clearly argue for a more nuanced, personalized approach to understanding comfort. Figure 19 shows that the results indicate substantial individual differences in the accumulation of thermal frustration, reflecting the prolonged impact of discomfort on occupants' decisions. TFT incorporates the cumulative effect of discomfort over time, making it well–suited to predict overrides during extended periods of suboptimal conditions, such as those seen in DR events. The higher performance of TFT in capturing override behavior, as depicted in these distributions, suggests that understanding occupants' historical discomfort—rather than solely focusing on instantaneous deviations—can significantly improve predictive accuracy.

Figure 20, Figure 21, and Table 3 together show TFT's ability in capturing dynamic thermal discomfort. It is worth discussing that AUC is a more suitable metric than MCC for evaluating thermostat override predictions in the context of demand response (DR) because it better captures the importance of accurately classifying overrides, which directly impact energy–saving outcomes. Unlike MCC, which penalizes all types of errors equally—including misclassified non–overrides that are less critical in this setting—AUC focuses on the model's ability to distinguish between overrides and non–overrides, regardless of class imbalance. This distinction is crucial for DR programs, where minimizing discomfort–driven overrides is key to maintaining energy efficiency. The higher AUC for the Thermal Frustration Theory (TFT) compared to the Comfort Zone Theory (CZT), particularly after excluding routine behaviors, highlights TFT's strength in capturing cumulative discomfort, leading to more accurate override predictions and ultimately better DR program performance.

The comparison between TFT and CZT reveals the complexity of modeling thermal comfort behavior. While TFT achieves marginally higher AUC (49.40 vs 47.09), both theories contain parameters that could be further optimized. TFT's assumption of $\beta=1$ and CZT's use of fixed comfort temperatures represent simplifications that, if refined, could improve both models' performance. The key distinction between TFT and CZT lies in their treatment of time. While CZT treats discomfort as an instantaneous state based on temperature thresholds, TFT models it as a cumulative process. This temporal integration becomes particularly relevant for demand response events, where the duration of exposure to sub–optimal conditions matter as much as the magnitude of temperature deviation. TFT's ability to track this accumulation of discomfort over time provides a theoretical framework more aligned with the extended nature of demand response events. However, this research cannot definitively establish TFT's superiority over CZT based on current results. Future work must investigate whether the temporal integration justifies the additional computational complexity TFT introduces. The similar performance metrics between the two approaches suggest that the choice between them may ultimately depend on specific application requirements and computational constraints.

The implications of these findings are substantial for the design and optimization of DR programs. TFT, in particular, demonstrated a consistently higher predictive capability compared to DRT and CZT, likely due to its incorporation of a temporal element that captures discomfort accumulation. This underscores the potential of using discomfort–memory–based models to enhance the precision of occupant behavior predictions during DR events. The results also emphasize the need to integrate personalized thermal comfort models, such as TFT, to optimize control strategies that minimize overrides and maximize energy savings without compromising occupant comfort. Prior studies have found significant inter–personal variations in thermal comfort preferences [140]. Our results show that dynamic thermal comfort models like TFT can account for such variations. Population–based assumptions are likely inadequate. The TFT framework derived here could be used to develop predictive control algorithms that optimize DR dispatch schedules and setpoint trajectories to minimize occupant overrides. However, scaling this to large building populations requires learning or approximating model parameters for each building. Transfer learning techniques could potentially address this [143].

The comparative analysis of the three models, with and without accounting for routine overrides, further reveals that excluding routine–driven thermostat changes can lead to more accurate assessments of discomfort–driven behavior. TFT, especially, benefits from this distinction, as its performance gap widens significantly when routine behaviors are excluded, indicating a robust capability to model genuine discomfort responses.

In summary, the results suggest that personalized, data–driven thermal comfort models, particularly those incorporating cumulative discomfort metrics, are essential for the successful deployment of DR programs that aim to balance energy savings with occupant satisfaction. The adoption of these advanced occupant–centric models could fundamentally reshape building energy management systems by aligning thermal control strategies more closely with individual comfort dynamics, thereby achieving the dual goals of energy efficiency and human well–being. Key challenges include improving the occupancy model, validating the assumptions about frustration accumulation and dissipation rates, and testing via implementation in real building control systems. Building connectivity and data sufficiency present additional barriers.

3.6 Conclusion

This study investigated three theories of thermal comfort – Delayed Response Theory (DRT), Comfort Zone Theory (CZT), and Thermal Frustration Theory (TFT) – to improve the prediction of manual setpoint changes (MSCs) during demand response (DR) events. By leveraging high–resolution smart thermostat data and employing machine learning techniques, we derived occupant–specific parameters for each theory and evaluated their performance in classifying thermostat overrides.

Our results demonstrate that accounting for occupant–specific thermal comfort parameters and dynamic behavior significantly improves predictions of overrides during DR events. The Thermal Frustration Theory (TFT) consistently outperformed the other models, likely due to its ability to capture the time–integral effect of discomfort rather than just instantaneous temperature deviations. This finding underscores the importance of considering the cumulative nature of thermal discomfort in predicting occupant behavior.

The study also highlights the significance of distinguishing between routine–based and discomfort–driven MSCs. By developing a routine model based on the time of day and incorporating it into our analysis, we were able to refine the parameter estimations for each thermal comfort theory. This approach led to improved model performance, particularly for TFT and CZT.

The methodology developed in this study has important implications for the design and implementation of DR programs. By more accurately predicting occupant behavior and potential overrides, utilities and energy companies can develop more effective control strategies to achieve committed energy savings during DR events while maintaining occupant comfort.

However, several challenges remain in scaling this approach to large building populations. These include improving occupancy modeling, validating assumptions about frustration accumulation and dissipation rates, and testing the models through implementation in real building control systems. Future research should focus on addressing these challenges and exploring transfer learning techniques to efficiently apply the models to diverse building types and occupant populations. The model validation approach used 5-minute fixed intervals for override predictions, which introduced methodological limitations. Each 5-minute period was treated as an independent event, with predictions classified as true positives, false positives, true negatives, or false negatives based on whether an override occurred within that specific interval. While this approach provided a standardized framework for comparing models, it created artificial boundaries that may not reflect the continuous nature of occupant comfort perception. Future work should explore sliding window approaches that evaluate predictions across overlapping time periods, potentially better capturing the temporal progression of thermal frustration. This could be particularly valuable for demand response applications, where predicting not just if but when an override will occur significantly impacts grid operations. Additionally, investigating the relationship between override magnitudes and prediction accuracy could reveal whether occupants making larger setpoint adjustments exhibit more predictable behavior patterns. The relatively similar AUC scores between TFT (49.40) and CZT (47.09) may partially reflect these methodological constraints rather than fundamental limitations in the theories themselves.

In conclusion, this study demonstrates the potential of data–driven, occupant–centric thermal comfort models to enhance our understanding of human–building interactions and improve the effectiveness of demand response programs. By continuing to refine these models and integrate them into building control systems, we can move closer to achieving the dual goals of energy efficiency and occupant comfort in the built environment.

Chapter 4:

Integrated dynamic occupant–building modeling

4.1 Abstract

This paper introduces a novel approach in modeling dynamic thermal comfort simulations, focusing on the various behaviors of building occupants. Traditional thermal comfort models, primarily static, fall short in accurately forecasting transient conditions and estimating energy savings during demand response (DR) events. To bridge this gap, we developed a data–driven model that dynamically models thermal discomfort behaviors of occupants, integrating agent–based modeling with the ecobee Donate Your Dataset (DyD), the alfalfa co–simulation framework, and a green–built home environment. The introduction of the Thermal Frustration Theory (TFT), which models the delay and magnitude of discomfort experienced by occupants before adjusting thermostat settings, challenges traditional models by acknowledging that occupants accumulate thermal discomfort over time, leading to a threshold–based response rather than immediate adjustments. By simulating occupants' thermal comfort over a year and employing the resulting data to train machine learning models, we identify critical occupant–specific parameters that enhance the accuracy of manual setpoint change (MSC) classifications. Our methodology demonstrated marked improvements in MCC scores when training was personalized for each home, as opposed to using a singular model on aggregated data. This approach can aid in the design of more sophisticated controllers for utilities or Community Service Providers (CSPs), to minimize overrides and fulfil predicted energy savings. Specifically, our findings highlight that personalized model, tailored to individual thermal histories and behavior patterns, significantly elevate model performance, with Matthews Correlation Coefficient (MCC) scores increasing from a mere 0.002 in a generalized model to 49.40 in personalized models, post exclusion of routine overrides. This research highlights the necessity and efficacy of personalized thermal comfort models in enhancing building energy management systems, offering a robust tool for accurately predicting and managing energy consumption in response to dynamic occupant behaviors.

4.2 Acronyms and variables

α	Thermal recollection coefficient
β	Thermal absorption coefficient
MSC	Manual setpoint change
MSC_D	Discomfort based manual setpoint change
MSC_R	Routine based manual setpoint change
TFT	Thermal Frustration theory
TF	Thermal Frustration
TF_λ	Thermal Frustration threshold
DyD	Donate Your Dataset
OTT	Occupant Thermostat Interaction
BPS	Building Performance Simulation

4.3 Introduction

The progression in Building Performance Simulation (BPS) and Grid–Interactive Efficient Buildings (GEB) has influenced energy management strategies in buildings [127]. However, a gap remains in accurately modeling occupant–thermostat interactions in dynamic thermal comfort scenarios [245]. As the national energy structure transforms, the traditional energy network's resource allocation capacity becomes inadequate for future energy systems' needs, necessitating the need for load flexibility and effective Demand Response (DR) strategies to effectively coordinate supply and demand in the evolving Energy Internet [246]. Occupant behavior, particularly in thermostat interactions, plays a crucial role in energy consumption patterns within buildings [247], [248], [249]. Current state–of–the–art thermal comfort models are predominantly static [17], [250]. They fail to adapt to the transient conditions experienced in real–world scenarios and hence struggle to capture to varying environmental conditions and occupant behavior, leading to discrepancies in energy usage predictions and actual consumption, and underscoring the need for more dynamic and responsive systems [136], [240].

While some studies have developed personal thermal comfort models such as improving the prediction of individual thermal preferences by integrating spatial–temporal data from BIM models and field–based studies [251], [252], [253], [254]. However, these personal thermal comfort models do not capture the dynamic thermal comfort behavior i.e. delay and magnitude of thermostat adjustment due to discomfort and routine behavior.

Recognizing the limitations of existing models, our research proposes an integrated framework that captures the dynamic nature of occupant–thermostat interactions. We focus on the development of models of thermal comfort based on individual routine and discomfort occupant behavior. This paper introduces an approach of combining agent–based modeling, real–world dataset integration (ecobee DyD), and a simulation environment (alfalfa co–simulation framework) within a green–built home.

The core of our model lies in understanding and simulating occupant occupant activity using a 1st order Markov chain model, routine thermostat interactions using a probabilistic model trained on real–world thermostat interactions (DyD dataset), and, most critically, discomfort–based thermostat interactions model tested on real–world thermostat interactions (DyD dataset). The complexity of human interactions with thermostats is often emphasized in studies, highlighting the need for more personalized models [240]. Building upon this, we introduce the concept of Thermal Frustration Theory (TFT), a framework that models the delay and magnitude of discomfort experienced by occupants.

Furthermore, this paper outlines the methodological advancements in learning and applying the parameters of our integrated model from field–collected data. We demonstrate the potential of this approach to the way building energy management systems understand and respond to occupant behaviors, particularly in the context of thermal comfort. The implications of our work extend to various stakeholders, including building designers, energy managers, and utility companies, offering them a more reliable tool for predicting and managing energy consumption in response to occupant behaviors by not only addressing the limitations of existing static models but by also leveraging the agent based modeling for a more accurate and responsive building performance simulation.

4.4 Methods

The methodology of development and implementation of an integrated occupant model designed to enhance the accuracy of building performance simulations (BPS). This section details the model's components, including occupant activity modeling using a first–order Markov chain, routine manual setpoint change (MSC) modeling, and discomfort–driven thermostat interactions based on the novel Thermal Frustration Theory (TFT). The methodology integrates these models with BPS to simulate dynamic occupant behaviors and their impact on thermal comfort and energy usage, addressing limitations in existing static models.

4.4.1 Proposed integrated model

The proposed integrated model aims to address a missing part in the current state of the art of building performance simulation i.e., dynamic occupant thermal comfort, more importantly occupant interactions with a thermostat factoring in the delay in assessing discomfort and the magnitude of discomfort. Traditionally, static or deterministic occupant models have been implemented which are unrealistic and cause the output of Building Performance Simulation (BPS) to disagree with real world observations [250]. Studies that model occupant activity stochastically show improvement in error when comparing BPS output vs real life observations [255].

Therefore, to improve predictions of the BPS we propose an integrated occupant model incorporating the dynamics of thermal comfort. The proposed integrated model shows the framework of integrating two important aspects of BPS i.e., building and occupant where the occupant model is a data–driven stochastic model composed of 3 components that capture an occupant's dynamic thermal comfort behavior. The 3 components of the occupant are: occupant activity model, Routine based thermostat interaction model, and most importantly the discomfort–based thermostat interaction model. This approach is different from the current industry standard of the ASHRAE's adaptive comfort model [136] for two key reasons. First reason being the delay of discomfort, previous models do not consider time as a parameter involved in modeling discomfort, it is assumed that an occupant acts as a switch and an occupant thermostat interaction occurs as soon as the current indoor temperature goes outside the comfortable temperature zone which is established by the adaptive model but this may not be the case as occupants can withstand discomfort for a certain time depending on the magnitude of discomfort [256]. Second reason being consideration of other behavior models that might influence discomfort, for example an occupant can only feel thermal discomfort if they are within a space, to model/forecast discomfort, an occupant's presence also needs to be modeled, similarly, an occupant might have routines established to set lower setpoints at night if they like to sleep in a cooler environment, therefore, this routine based habitual model would also need to be developed. Most of the comfort models, including adaptive comfort model, aim to model comfort only and ignore other behavior elements that are equally critical in an occupant's decision to override controls.

As a result, the three models i.e. occupant activity model, Routine model, and discomfort–based thermostat interaction model of an occupant work together to represent an occupant. While sociological, financial, environmental, or other aspects can also influence an occupant–thermostat interaction [240] but due to the lack of availability of data, only the three models discussed above will be used for this study.

4.4.1.1 Occupant activity Model

4.4.1.1.1 Background

Occupant activity, where an occupant is in the space and performing an activity such as sitting, standing, cooking, etc... except sleeping, has a tremendous impact on energy consumption and is a critical part of building simulation using tools like EnergyPlus [257]. To model dynamic occupant thermostat interaction, an occupant's occupant activity in a home must be modeled reliably. Due to limitations in the availability of data on occupants' exposure to outdoor environments, this study will center on modeling behavior within indoor settings.

Various techniques exist for modeling occupant activity, such as the Monte Carlo method [65], non–homogeneous Poisson process model [258], inhomogeneous two–state Markov chain [179], agent–based model of occupant activity dynamics [259], but among these techniques, the first–order Markov chain stands out due to its ability to capture temporal characteristics where the state transition probabilities can be used to predict the state at the next timestep based on the current state, with no need for complex model parameters estimations, easy to learn and implement, especially for computation–intensive applications such as the ABM based dynamic occupant behavior modeling framework proposed in this paper [260] [261] [262].

4.4.1.1.2 Model details

Assuming that the likelihood of an occupant's occupant activity at a given time step is solely influenced by their status during the previous time step and the time of the day, the occupant activity model using a first–order Markov chain [260].

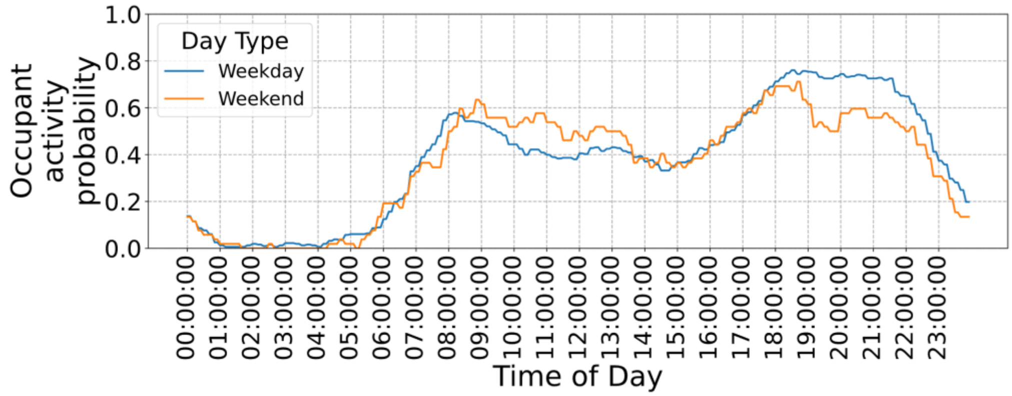

Figure 22: Occupant activity probabilities for a simulated occupant

At midnight, a schedule of the occupant's occupant activity schedule is created by realizing the occupant activity state from probabilities retrieved from the trained state transition matrix (training discussed in section 4.4.2.1) and comparing the probabilities with a random number generated from a uniform distribution. As a result, an occupant's occupant activity data is generated (Figure 22).

4.4.1.2 Routine MSC model

4.4.1.2.1 Background

The task of modeling routine behaviors discussed across disciplines and has diverse applications, making it a focus area in various fields of study.

1. **Healthcare**: In healthcare, behavioral models help in predicting patient compliance with medication routines, thereby aiding in the tailoring of more effective treatment plans. Models can predict, for example, the likelihood of a diabetic patient sticking to an insulin regimen or a cardiac patient adhering to medication schedules. These insights are invaluable for healthcare providers aiming to offer personalized treatments [263]
2. **Transportation**: In transportation planning, understanding routine commuting patterns is vital for optimizing public transport services. Sophisticated models analyze large data sets to predict the flow of commuters during peak and off–peak hours, facilitating more efficient transport scheduling and even influencing urban planning [264]

3. **Retail and Marketing**: In the commercial sector, companies often model customer buying behaviors. Understanding routine purchase patterns can inform stock decisions, optimize marketing strategies, and even personalize online shopping experiences [265]
4. **Finance**: In finance, models may predict routine behaviors in trading to capture trends and even irregularities like market manipulations. Machine learning algorithms are often used to predict buying and selling patterns based on historical data [266]

The ubiquity of routine behavior modeling across these domains highlights its significance and the need for robust and flexible models is to adapt to different types of data and capture various influencing factors [267]. These studies show that humans develop routines for their schedules irrespective of the domain. The key distinction between routine behavior modeling across domains like healthcare, transportation, retail, and finance versus building energy management lies in the complexity and predictability of influencing factors. In healthcare, transportation, retail, and finance, behaviors often exhibit stationary and repetitive patterns, making them amenable to simpler statistical models that capture regularities using autoregressive techniques or time–series analysis. In contrast, building energy management is far more dynamic due to the interplay of unpredictable factors such as weather variations, occupant activities, and individual comfort needs, requiring adaptive and non–linear modeling approaches. These models must account for non–stationarity and dynamically adjust to changing environmental conditions, highlighting the need for more sophisticated adaptive methods to effectively capture discomfort–driven behaviors that are less critical in other domains.

4.4.1.2.2 Model details

Routine manual setpoint change behavior can vary across seasons and types of day. Therefore, for each two seasons i.e. heating and cooling and two types of days i.e. weekdays and weekends, four sets are created where in each set contains subsets of routine MSC behavior pertaining to number of MSCs in day (an occupant can routinely implement MSCs), time of day of MSCs (Occupants implement MSC before going to work or coming back from work), type of MSCs (Occupant change only heating setpoint in the morning), and degree of MSCs (Occupants increase or decrease heating setpoint in the morning). To realize the probabilities of the subsets of routine MSC behavior, a routine schedule is generated where for each time step of the day probabilities are realized in a way where the subset's values are conditionally dependent on each other. The selected features, such as the number of MSCs per day and the time of day, were chosen to capture routine occupant behavior. For instance, some occupants set a conservative temperature in the morning to save energy while away and adjust it to a comfortable setting upon return. The time–based nature of routine behaviors, tied to daily activities like waking up or going to work, allows the model to predict and optimize thermostat use effectively.

4.4.1.3 Discomfort MSC model

4.4.1.3.1 Background

Discomfort–driven models aim to capture the human interaction with thermostats due to thermal discomfort [268]. Various theories have been proposed to model this behavior, and these theories often find applications beyond just thermal comfort, extending to fields like ergonomics, psychology, and human–machine interaction [269] [270] [154] . The major frameworks in this domain include Delayed Response Theory [240], Adaptive comfort model [135], and the proposed novel Thermal Frustration Theory, each of which offers a different perspective on how and when people adjust thermostats based on their comfort needs.

1. Delayed Response Theory (DRT): This theory assumes that people act like timers, experiencing discomfort only after a fixed delay following exposure to uncomfortable thermal conditions [240].
2. Comfort Zone Theory (CZT): This theory suggests that people have a defined range of temperatures within which they feel comfortable. Any deviation outside this zone triggers discomfort and subsequently an adjustment of the thermostat settings [135].
3. Thermal Frustration Theory (TFT): This novel theory hypothesizes that individuals accumulate thermal frustration over time when exposed to uncomfortable temperatures. Once this accumulated frustration crosses a certain threshold, it prompts the individual to interact with the thermostat. Each occupant can have a different frustration threshold that nudges them to interact with the thermostat.

Each of these theories offers unique concepts that drive MSC behavior. But developing only comfort models is not enough as they aim to quantify only the physical and psychological responses of occupants to their environment, determining when they perceive discomfort (mostly temperature and humidity) and how they respond to it. In contrast, behavior models focus on capturing the patterns and triggers of occupants' actions, such as when and why they adjust the thermostat, often driven by routine or external factors. Integrating these two approaches allows for a more comprehensive understanding of how occupants interact with their environment, balancing routine behaviors with their comfort needs. It is assumed that the novel thermal frustration theory can serve as a better model as occupants do not instantly interact with their thermostat when experiencing discomfort, instead a delay is observed in exposure to thermally uncomfortable conditions and thermostat interactions [256]. For this reason, only the TFT model will be used to model occupant discomfort behavior and in future the authors plan to compare the three theories mentioned above.

4.4.1.3.2 Model details

The discomfort–based model based on the TFT theory hypothesizes that thermal discomfort accumulates over time, prompting the user to make a MSC when the discomfort reaches a certain level. This differs from conventional approaches like the Adaptive Model, which assumes people are comfortable within a specific temperature range based on outdoor conditions. However, in the context of Demand Response (DR) programs, which involve changing conditions, occupants may be willing to tolerate temporary discomfort even when they are outside the thermal comfort zone. Therefore, the discomfort–based model focuses on tracking the buildup of thermal discomfort until it hits a set threshold, at which point a MSC is triggered. This novel thermal discomfort accumulation model is referred to as the Thermal Frustration Theory (TFT).

The Thermal Frustration (TF) at a specific time k is calculated as:

$$TF[k] = (\alpha \times TF[k-1] + \beta \times (T_{in}[k] - T_C)\ ) \times occupant\ activity \tag{11}$$

Where, TF represents the level of thermal frustration in $°Fmins$, k is the current timestep, α is the coefficient that captures how much of the past thermal discomfort is remembered, β is the coefficient that captures how much of the current thermal discomfort is being felt in $°Fmins$, and "occupant activity" is a Boolean variable indicating whether movement has been detected leading to 0 thermal frustration when the occupant is not present in the space of occupant activity measurement. The current level of discomfort is the difference in the current indoor temperature T_{in} and the comfort temperature T_C in $°F$.

The discomfort model in (11), α serves as a factor that quantifies how much of the previously experienced thermal discomfort is retained or 'remembered' over time and hence can be in range of 0 and 1. β, on the other hand, gauges the degree of discomfort experienced during a given timestep, which could potentially vary based on factors such as metabolic rate or clothing and is also in the range of 0.0 and 1.0 minutes. Moreover, occupant activity here serves as the reset condition, forcing the thermal frustration value to zero when the occupant is not within the space as for this paper, instances where the occupant is not exposed to external factors like outside temperature/humidity or other events is not considered.

One can define this consistent thermal frustration at MSC to be a thermal frustration threshold or TF_λ where $TF_\lambda = Mode(\overrightarrow{TF}_{MSC_D})$ and $\overrightarrow{TF}_{MSC_D}$ is vector containing the thermal frustration values at various discomfort based MSCs.

4.4.1.4 Integration into BPS

4.4.1.4.1 Background

In this study, Alfalfa serves as the central computational framework, facilitating seamless interaction between the occupant and building models [271]. Alfalfa is an open–source web application ingeniously engineered to bridge Building Energy Modeling (BEM), Building Controls, and Software Engineering. Its unique feature set, including real–time interactions with BEMs through industry–standard control interfaces, allows for versatile co–simulations. Specifically, Alfalfa supports popular modeling engines such as EnergyPlus, enabling large–scale parallel building simulations. Its RESTful interface streamlines the process of uploading, initiating, and terminating simulations, thereby simplifying complex simulation logistics [272].

Python was the primary programming language employed to orchestrate the co–simulation framework. For the development of the occupant model, an agent–based modeling approach was utilized. The Mesa library provided an initial framework that was subsequently customized to meet the specific requirements of the occupant behavior model [273].

To further enhance the efficiency and scalability of our simulations, Kubernetes for deployed for containerization[274]. Before the implementation of Kubernetes, a single 1–year simulation for one home required approximately 50 hours to complete. With Kubernetes, this performance bottleneck was significantly alleviated. Parallel simulations were run for 10 homes over a span of one year, reducing the computational time to just 19 hours.

A smart thermostat offers the capability of setting schedules that can be based on time of day and geolocation. As a result, python based a thermostat model object can be developed.

4.4.1.4.2 Model details

Now with the models discussed above, one final need was to integrate the models such that for a given set of input parameters, the occupant can be simulated in a home for their thermal comfort behavior (Figure 23).

First, the **alfalfa** framework is initialized with the appropriate building data i.e., "idf" files and weather i.e., "epw" files. The energy plus part of the co–simulation framework requires the **input** of setpoint temperature and simulates the indoor temperature based on building thermal dynamics, outside temperature, and HVAC equipment [271]. The **output** of alfalfa framework required is the indoor temperature and the energy usage.

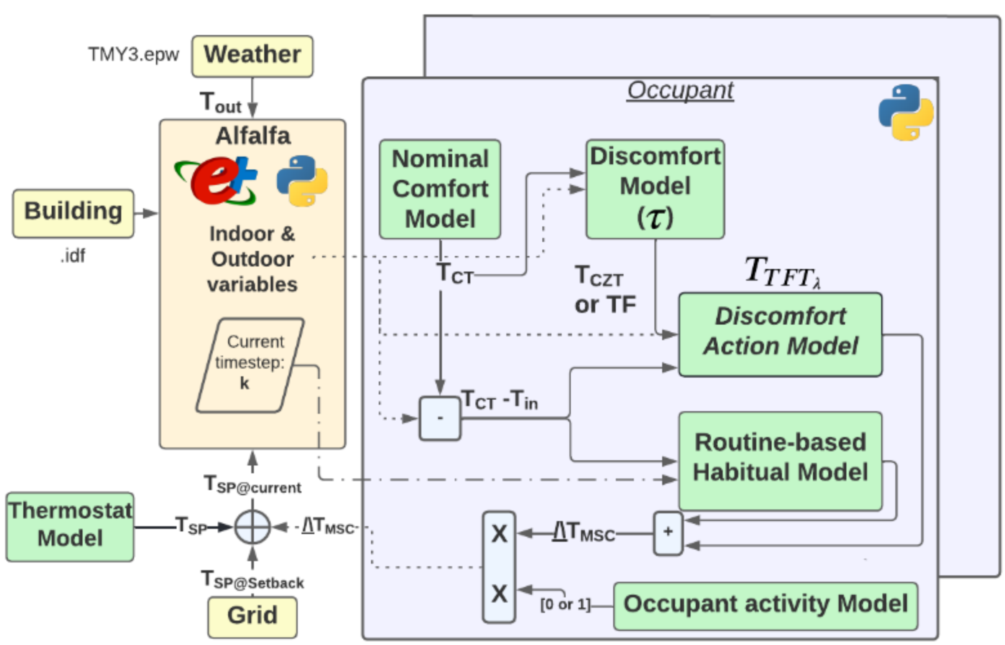

Figure 23: A feedback loop of occupant behavior and building components; Green boxes represent data being modeled using DyD or ATUS dataset.;

Second, the occupant and thermostat models are initialized. The **thermostat model** requires the **input** of time of day and setpoint temperature (in case MSC or DR event) and **outputs** the setpoint temperature (cooling and heating). The **occupant model** requires the **input** of current indoor temperature, humidity, time of the day, and **outputs** setpoint temperature which is mostly 0 except when the discomfort or routine models trigger a MSC.

4.4.2 Learning of integrated model from smart home data

The proposed integrated model discussed in section 4.4.1 can be trained to used parameters from the smart home data and simulate the indoor thermal environment to mimic an occupant's experience with their indoor surroundings.

4.4.2.1 Occupant activity model

The occupant activity of occupants within a space directly impacts how they interact with systems like the thermostat. While the PIR sensor data can observe occupant activity in a space during various timesteps of the day but due to the occupant behavior of sleeping, sitting, or the sensor itself being shaded by an object can compromise the data and in turn yield a high rate of False Negatives which in turn will affect the occupant activity forecasting accuracy [256].

To overcome this issue, a study that utilized Time–Use Survey (TUS) data to develop a first–order Markov chain model for simulating occupant activity was used [260]. The TUS data are self–reported by individuals, enhancing the reliability of the collected information and alleviating the concerns related to fault measurements. Conducted in the UK in 2000, this survey provides granular insights into daily activities recorded in ten–minute intervals. A first–order Markov Chain technique discussed in section 4.4.1.1.2 was used to generate synthetic occupant activity data, ensuring that the synthetic data retains the same overall statistical attributes as the original survey data. The "states" represent occupant activity levels at ten–minute daily intervals and the transition probabilities between these states are calculated based on matrices developed from the original TUS data.

The transition probability matrix can be learned using the timeseries data of an occupant's occupant activity data.

4.4.2.2 Routine MSC model

Occupants have many reasons to override that may not be driven by comfort and can follow various time of day of MSCs. Routine events in the occupants' lives may drive them to manually change the thermostat setpoint: e.g, awakening, leaving, or returning from work, or prior to sleep. The stochastic model below aims to (a) model how many times MSCs will occur in a given day, (b) when those MSCs will occur, and (c) the direction and magnitude of the MSC.

1. Number of MSCs per day: Using the probabilities of each number of MSC per day as discussed in section 4.4.1.2, the number of routine MSC(s) to be implemented is realized by comparing probabilities with a random number generated from a uniform distribution.
2. Time of day of MSCs: Given the number of MSCs per day value, the time of day of these routine MSCs can be computed based on routine MSC probabilities.
3. Type of MSCs: The type of MSC i.e. heating, cooling or both can be determined by change in heating, cooling, or both setpoint temperature at a MSC. As a result, a type of MSC PMF conditional on time of day MSC can be computed.
4. Degree of MSCs: The degree of MSC can be computed using the change in setpoint temperature data at MSC. The PMF of degree of 1st MSC i.e., heat or cool is conditional to number of MSC, type of MSC, and TOD of MSC.

Upon performing the realizations discussed in the 4 steps above, a routine MSC schedule can be generated and implemented.

4.4.2.3 Computing comfort temperature

Calculating the comfort temperature of an occupant is a critical step in accurately modeling thermal preferences as it is used a reference to quantify thermal discomfort magnitude. For this purpose, we implement a mean value approach, which provides a straightforward method for estimating the temperature that occupants are most likely to find comfortable during routine periods of indoor activity.

To apply this method, we first curate the dataset to include only instances where the occupants are passively interacting with their environment—meaning that no manual overrides (hold events) occurred prior to or during the recorded temperature readings. This selection criterion ensures that the data used for calculating the comfort temperature reflects the occupants' natural preferences without the influence of external or irregular factors.

The mean comfort temperature T_C is then calculated as the average of all valid indoor temperature readings from the curated dataset that represents the most probable temperature that occupants find comfortable under normal conditions. By using this mean value as a baseline, the model can effectively predict thermal comfort, enabling the system to make informed adjustments that align with the occupants' habitual temperature preferences.

This approach is particularly suited to our study's objectives, as it allows us to generalize the comfort temperature across different occupants while still capturing the essence of individual preferences.

4.4.2.4 Discomfort MSC model

A timeseries thermostat data can also be leveraged to identify key parameters of the TFT discomfort model discussed in section 4.4.1.3.

To optimize the parameter α in the context of time–series smart thermostat data, the following grid search procedure can be employed (as shown in Figure 24):

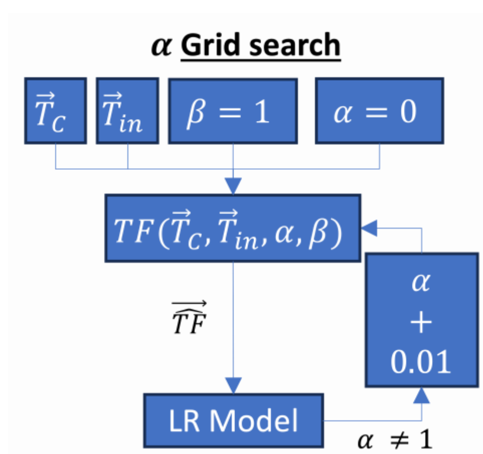

Figure 24: Schematic of α grid search

1. Initialize the parameters $\alpha = 0, \beta = 1,$ and T_C as a real number.
2. Calculate the 'thermal frustration' feature for the entire dataset, using the given $\alpha, \beta,$ and T_C.
3. Identify the row indices where the ground truth value for a MSC is True.
4. Using these indices, train a logistic regression (LR) model on both MSC and non–MSC data subsets. Apply the model to predict MSC occurrences on a test dataset. Compute both the accuracy and Matthews Correlation Coefficient (MCC) score. Log α along with its corresponding MCC and accuracy score.
5. Increment α by 0.05 and repeat steps 2–4 until $\alpha = 1$.

The choice of a simple LR model for MSC classification is strategic for several reasons, the most prominent being the assumption that relationship between the predictor variable—thermal frustration—and the log–odds of a discomfort–based override is approximately linear. Furthermore, the probabilistic outputs can be leveraged to compute the likelihood of MSC based on discomfort or another model, hereby adding the degree of certainty the model has about its prediction [275].

Improving discomfort model parameters estimation:

Figure 25 represents the schematic of two step training process to improve TFT model parameters estimation. Given that a timeseries thermostat data is available where the occupant interacted with the thermostat using the models discussed in 4.4.1, the goal is to now estimate key personal parameters of the occupant's discomfort model discussed in 4.4.1.3.2. The challenge is to reduce the noise caused by the routine model by removing instances of the MSCs that have high likelihood of routine MSC.

The following methodology can be used to increase accuracy of ML models and get likelihood of appropriate reason for MSC.

Train suboptimal routine model:

For the identification of discomfort model, a complex routine model may not be necessary to be developed. Since, a routine is based on time of day, therefore, a simple suboptimal routine model can be trained to compute the PMF of TOD of MSCs using (12). Now since each instance has a time stamp, therefore each instance can be assigned a probability of routine based MSC i.e., P_{MSC_R} as a result of PMF learned using (12). Let T be the random variable representing the time of a MSC, the PMFs can be represented as:

$$f_T(k) = \frac{n_k}{N_k} \tag{12}$$

Where f_T represents the probability mass function of time of day of MSC at k timestep of the day, $n_k$represent the number of MSCs at timestep k, N_k represents the total number of instances of timestep k.

Identification of the discomfort model:

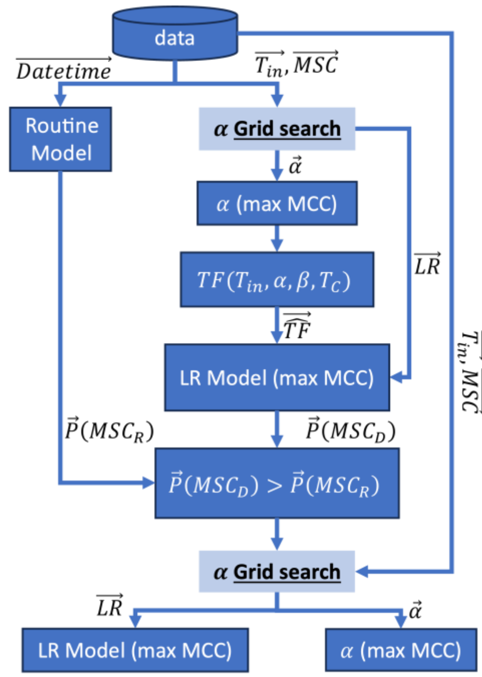

Figure 25: Schematic of the two–step training process to improve alpha estimation

1. First run of α grid search: In parallel to the training of the simple routine model discussed in the section above, the first grid search of α can be performed to find the optimal α of the user as discussed in section 4.4.2.4. It is important to note that these initial predictions are expected to be not ideal as the training and testing data is composed of routine, discomfort and other reason based MSCs. This initial grid search should still yield a suboptimal α value with a corresponding MCC score. The entire dataset can now be used as a test data set to predict probability of discomfort based MSC i.e., P_{MSC_D}
2. Second run of α grid search: Improving the discomfort model identification accuracy. With the probability of routine and discomfort overrides computed for each instance of the dataset, the corresponding probabilities can now be compared and the LR based discomfort classifier can be retrained on the data where the discomfort based MSC probability is higher than the routine based MSC probability. Therefore, the training of LR model to grid search for optimal α is done again (as discussed in section 4.4.2.4) but only for MSC instances where the $P(MSC_D) > P(MSC_R)$.
3. Since now, the discomfort model is trained only on high discomfort probability MSCs, the optimal α prediction and the MCC score is expected to improve.

4.4.3 Validation of model training in simulation

Mostly, the real–world data does not provide contextual data such as type of MSC i.e., routine, discomfort, etc. Therefore, a simulation environment was developed to use models discussed in section 4.4.2 and simulate occupant thermostat interaction behavior in a year with 1–minute timesteps to see how well the key parameters of the occupant discomfort model be identified, specifically the α value used to initialize the discomfort model of the occupant with the noise of the routine model.

4.4.3.1 Environment

4.4.3.1.1 Building model

The Greenbuilt house in Fair Oaks, Sacramento, originally built in the 1980s, was retrofitted through the Sacramento Municipal Utility District's (SMUD) Energy Efficient Remodel Demonstration program in collaboration with Greenbuilt Construction and the National Renewable Energy Laboratory (NREL). The retrofit aimed to showcase cost–effective, energy–efficient measures, including enhanced insulation, motorized shading, and advanced home control systems like Control4 for remote monitoring. NREL extensively monitored the house's performance to validate energy savings and system effectiveness, positioning it as an ideal case study for testing and simulating energy–efficient technologies without occupant interference, supported by a team of experts.

4.4.3.1.2 Thermostat model

To mimic a smart thermostat, ecobee's default schedules are used to initialize the thermostat model. The model outputs the setpoint temperature for both cooling and heating setpoints and changes only with the input of occupant MSC or time of day–based schedule change. Two schedules were added i.e., sleep and home, both of which are based on time of day from ecobee [276].

4.4.3.1.3 Occupant activity model

To simulate the occupant activity schedule of an occupant, a pre–trained state transition matrix of 1st order Markov chain model was used [260]. As the authors of the study validated the model's ability to produce occupant activity schedule that represent the population of TUS dataset, this model is a good candidate to generate occupant occupant activity schedules. During the simulation, these transition probability matrices predict the likelihood of occupant activity at different times of the day. These probabilities are then transformed into occupant activity events by comparison with a random number drawn from a uniform distribution. The only deviation from the study was that instead of generating schedules for multiple occupants in a home, only one occupant's schedule would be generated.

4.4.3.1.4 Routine model

To simulate the routine model of occupant thermostat MSC, ecobee's DyD dataset was used to compute the PMFs discussed in section 4.4.2.2. After performing the required computations, the PMFs for various parts of routine based MSCs are shown in Figure 26 and

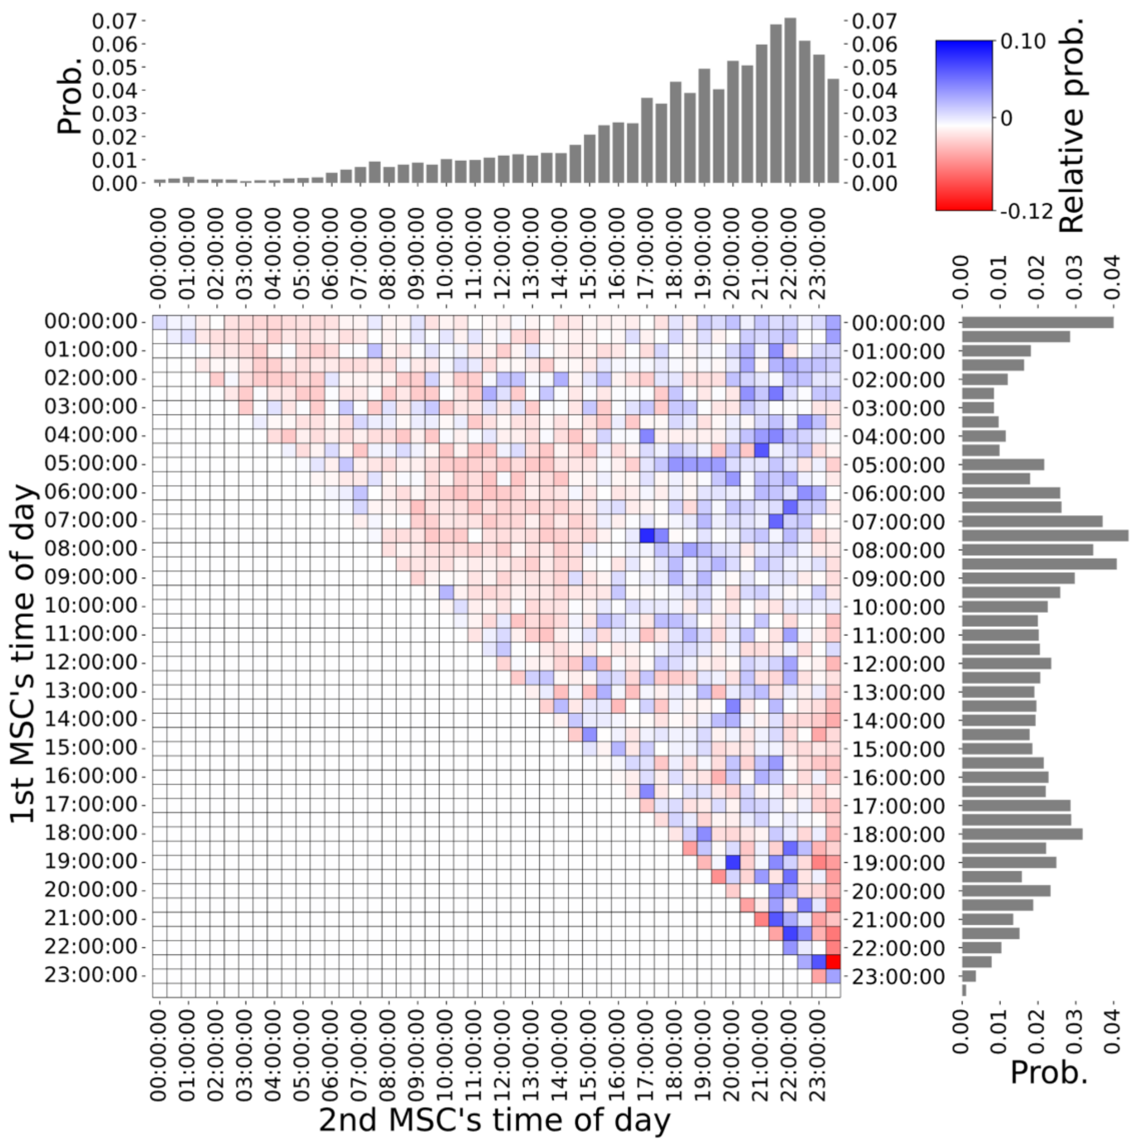

Figure 26: Two parts of many conditional probabilities of the routine model (part a)

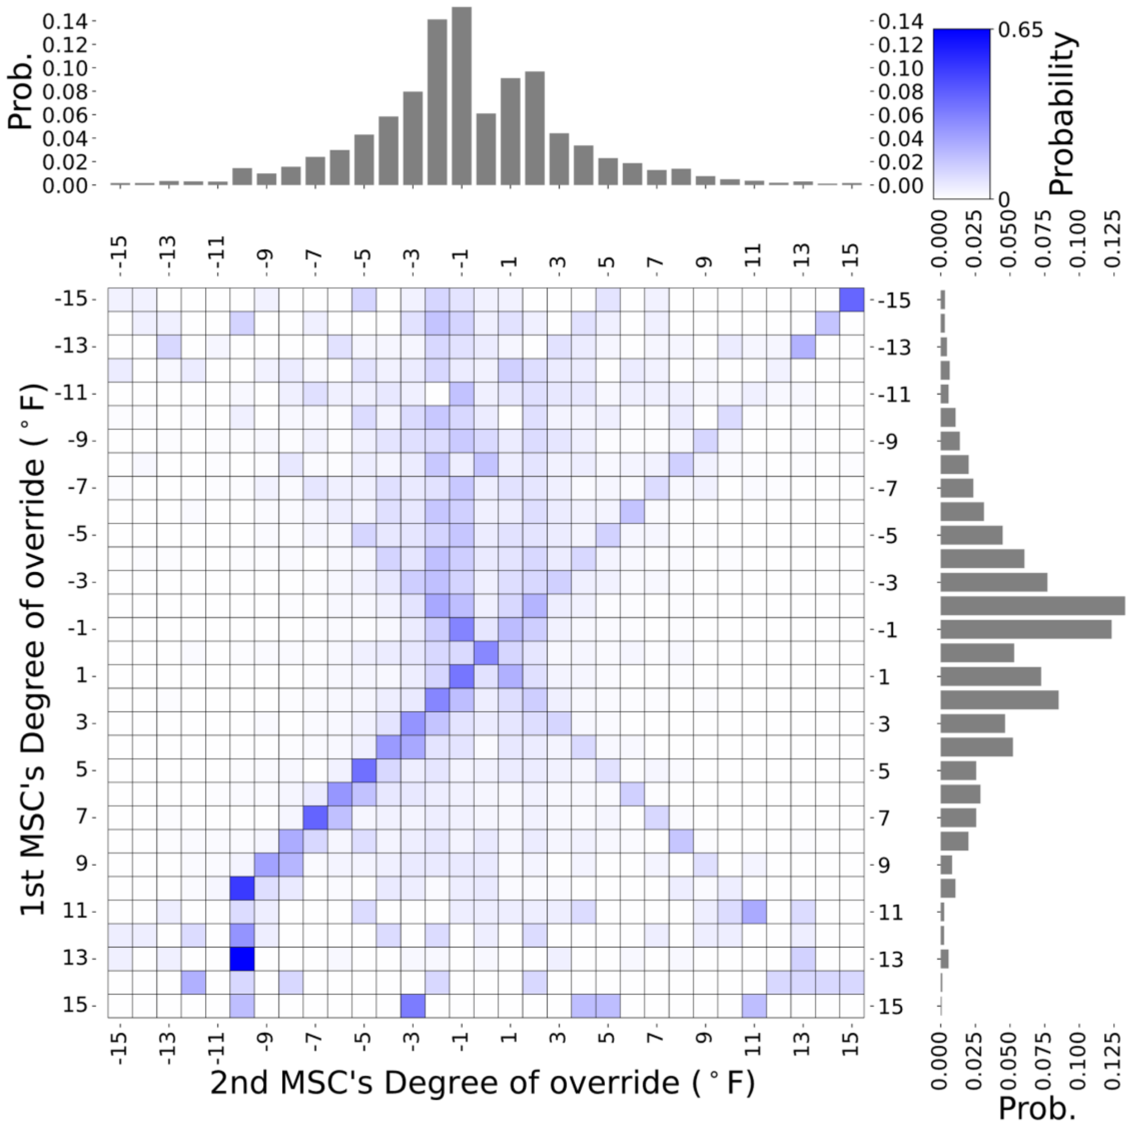

Figure 27: Two parts of many conditional probabilities of the routine model (part b)
All parts of the routine model start from number of manual setpoint changes (MSCs) per day which is realized from the probability distribution. Ranging from 0 to 15 NMSCPD with its corresponding probability ranging from 0.0 to 1.0, most days have 1 to 2 MSCPDs, with the probability decreasing sharply as the number of MSCs increases. ***Figure on the left:*** *Relative conditional probability (Difference of each cell's prob. – uniform dist. prob.) of 2nd MSCs time of the day given 1st MSCs time of the day and NMSCPD is 2.* ***Figure on the right:*** *Probabilities of 2nd Degree of override given 1st Degree of override and NMSCPD is 2.*

Figure 27 reveal important temporal and behavioral patterns in thermostat overrides by occupants.

- The left figure highlights that occupants often make their second thermostat adjustment (MSC) at a similar time to their first, suggesting a routine behavior linked to daily activities such as transitions between home and work. Deviations from uniform behavior are particularly noticeable in the late afternoon or evening, indicating increased override actions during these times, likely due to changing comfort needs after returning home or preparing for sleep.

- The right figure shows that occupants tend to make corrective adjustments, with the subsequent setpoint change often being of similar magnitude but in the opposite direction of the initial adjustment. This pattern reflects a habit of fine–tuning rather than setting a single temperature. The strong correlation between the timing and magnitude of adjustments points to routine–driven behavior as a significant factor in override actions. Leveraging these insights can help predict and preempt overrides during demand response (DR) events, enabling more effective energy management without compromising comfort, especially during sensitive times like evening hours.

4.4.3.1.5 Discomfort model

The following parameters were used to initialize the discomfort model discussed in 4.4.1.3:

1. $\alpha = 0.9$: For this preliminary run, the goal was to use an alpha that would build up frustration gradually as $\alpha < 0.7$ can lead to slower build up/fast leakage of frustration which would rarely trigger MSCs. As compared to $\alpha = 1$ would build up frustration quickly, leading to discomfort based MSCs frequently. This is a key parameter that can vary across occupants.
2. $\beta = 1$: Since no data is available to model the clothing/metabolic rate of an occupant, the worst–case scenario was assumed i.e., occupant feels all the frustration.
3. T_C was computed as a mean of temperature values where the occupant was in space and comfortable as discussed in section 4.4.2.3.

4.4.3.2 Validation

Given the discomfort model parameters used to initialize the occupant discussed in section 4.4.3.1.5 and with the noise of routine MSCs discussed in the section 4.4.3.1.4, the goal is to validate the proposed learning of the integrated model discussed in the section 0 help identify the initialized model parameters of the discomfort model of the occupant.

4.4.3.3 Approach

Following the iterative approach outlined in the section 0, the output of the simulation data that contains Date Time (dt), indoor temperature (T_{in}), setpoint temperature ($T_{stp_{heat/cool}}$), $motion$, manual setpoint change due to routine model (MSC_R), and manual setpoint change due to the discomfort model (MSC_D) were used.

Accuracy cannot be used as an LR model performance assessment metric as it only considers correctly classified labels and does not penalize False negative or False Positive. Accuracy is also prone to incorrect assessment of model performance especially in the case of imbalanced datasets such as that of the MSC classification where MSC occurs far less than non–MSC instances.

Matthew's Correlation Coefficient (MCC) serves as the evaluation metric for optimization, chosen for its balanced, symmetric representation of both positive and negative classification outcomes. MCC values range from –1 to 1, signifying perfect negative and positive correlation, respectively, with a value of 0 indicating no better performance than random classification.

4.4.4 Model training on real data

If the discomfort model parameters can be identified from the simulated dataset, therefore our next goal is to confirm if the methodology in section 0 can also be used on real–world data such as ecobee's smart thermostat DyD dataset to identify α and T_C parameters.

4.4.4.1 Occupant activity

Smart thermostat allows one to interact with their thermostat using a smartphone and hence to train the discomfort model, it is crucial that the occupant experiences the thermal conditions in a home to interact with their device as compared to remote interaction without experiencing any thermal discomfort. While the DyD dataset offers PIR sensor–based occupant activity data to determine occupant's occupant activity in a space, but due to the limitations of a PIR sensor, occupant activity data filtering becomes important. Authors have previously discussed the occupant activity data filtering in detail and aim to use similar "gap – filling" methodology to process occupant activity data [256]. In the study, various occupant activity data filtering minutes were considered i.e., 15, 20, 30, and 120– minutes, where each of filters aim to fill the gap if occupant activity was detected between in the filter minutes. As concluded from the study, the 30–minute filter served as the middle ground for false positives and false negatives, therefore the same methodology will be used in this study to process occupant activity data.

4.4.4.2 Routine model

The routine model training is intended to weed out MSC datapoints that create noise for the discomfort model. Here the goal of the model is to compute probabilities of an occupant interacting with the thermostat given the time of day. As discussed in section 0, a PMF is trained to estimate probability of routine MSC given time of day.

4.4.4.3 Discomfort model

Before using the discomfort model, a static comfort temperature value (section 4.4.2.3) is required to be identified to ensure that the TF computed during the fitting process in the discomfort is as accurate as possible. As discussed in the section 4.4.2.4, the discomfort model follows a two–step training approach, where the first training step is to use the LR model to estimate discomfort–based MSCs and identify an occupant's α and thermal frustration threshold values and the second training step is to retrain the LR model after removing the instances of MSCs where the routine MSC probability is higher than the discomfort MSC probability (achieved from the first training step).

It is important to note here that additional flags were used to filter MSC datapoints used to train the LR model. For this paper, only routine and discomfort models were considered but there can be instances where other aspects of an occupant can motivate them to implement a MSC. If an occupant has been home for 20 minutes, and previous setpoint was within 2 hours, only then an MSC is qualified to be considered for training the LR model. These flags ensured that model can capture occupant's thermal frustration that led them to interact with their thermostat and remove the noise of MSCs caused by other aspects of an occupant's life such as social, environmental, financial, etc.

4.5 Results

The proposed integrated model was tested with alflalfa's co–simulation framework. As shown in Figure 28, the occupant's occupant activity is key in implementing MSCs as the thermal frustration exceeds the assigned thresholds.

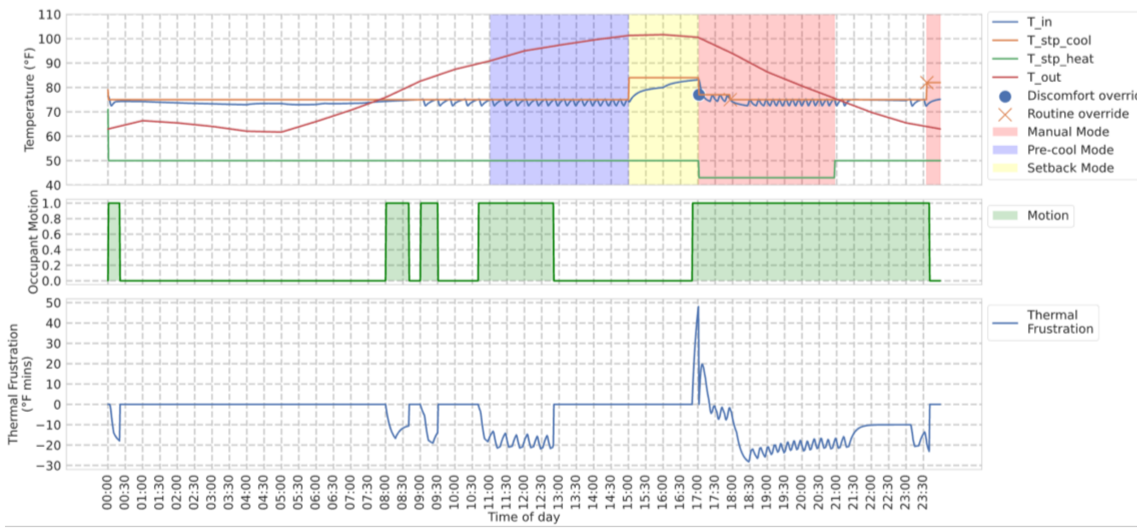

Figure 28: A snapshot of simulation run of integrated occupant and building model

4.5.1 Validation of model training in simulation

Section 4.4.3.1.5 discussed the parameters used to initialize the occupant discomfort model in the simulation. As the simulation outputs the reason for an MSC, an initial test was run to check if the fitting process can learn to classify MCC perfectly and exactly identify the α value, therefore the model was tested on 0.6 of train/test ratio where for MSCs (label 1) only discomfort based MSCs were provided with Non MSCs (label 0). As shown in Figure 30, the model is able to capture all MSCs evidenced by MCC of 1 for $\alpha = 0.9$ (same as the simulation's initialized α parameter).

TABLE 4: CONFUSION MATRIX OF ROUTINE MSC CLASSIFICATION			
		Predicted Labels	
		0	1
True Labels	0	5255229	2
	1	599	170

Having confirmed the fitting performance, routine and discomfort MSCs were mixed. The first fitting of logistic regression model with routine MSCs identified α of 0.4 with MCC score of 0.09. Next step of training a simple routine model as discussed in section 4.4.2.4, yielded accuracy of 0.99, precision of 0.98, recall of 0.22, F1 score of 0.36, and MCC of 0.46 (Figure 30). Lastly, upon removing the routine overrides, the fitting of logistic regression model identified α of 0.75 and 0.85 with the similar MCC scores of 0.14.

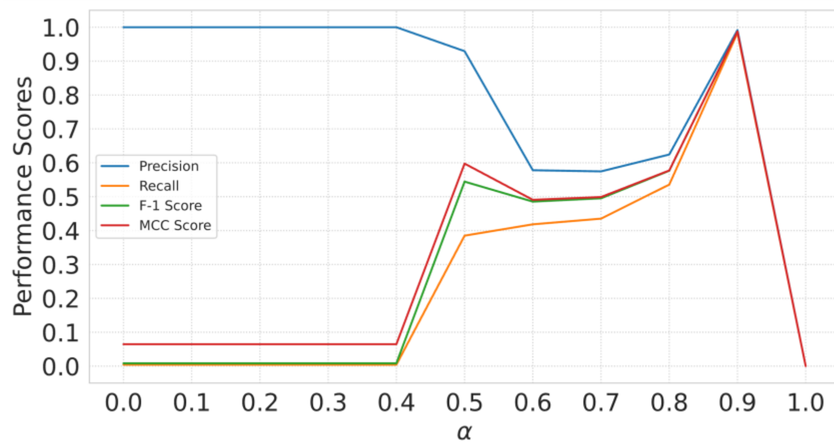

Figure 29: Performance scores of the LR model for alpha parameters (with only discomfort overrides)

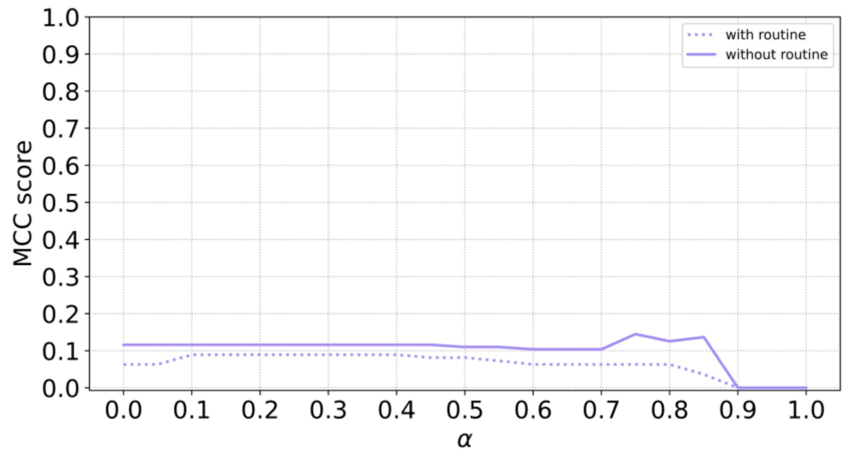

Figure 30: MCC scores of the LR model for various alpha parameter (with and without estimated routine overrides)

4.5.2 Model training on real data

Using the methodology discussed in this section 4.4.2.4, the DyD dataset was used to compute the thermal frustration parameter for various α and logistic regression model was trained to classify an MSC.

4.5.2.1 Model performance on aggregate data

Initially, the TF model was applied to the entire dataset without accounting for individual differences among homes. This approach resulted in a Matthew's Correlation Coefficient (MCC) of 0.002 (Figure 33), indicating a negligible improvement over random chance.

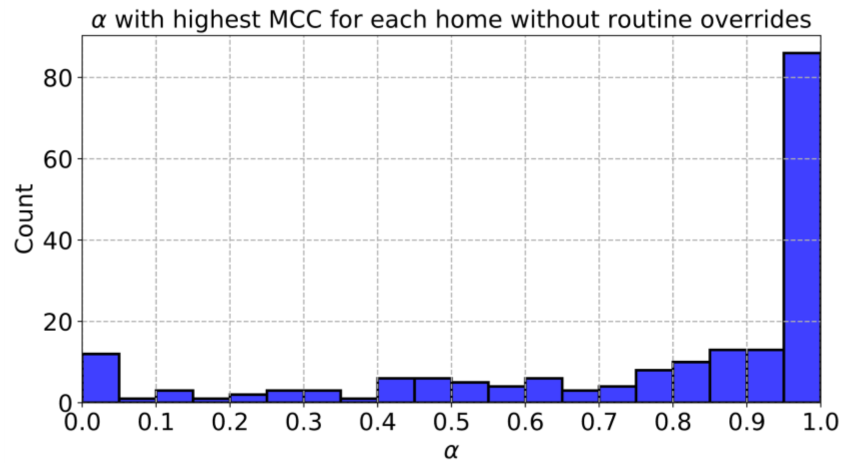

Figure 31: Distribution of identified alpha parameter

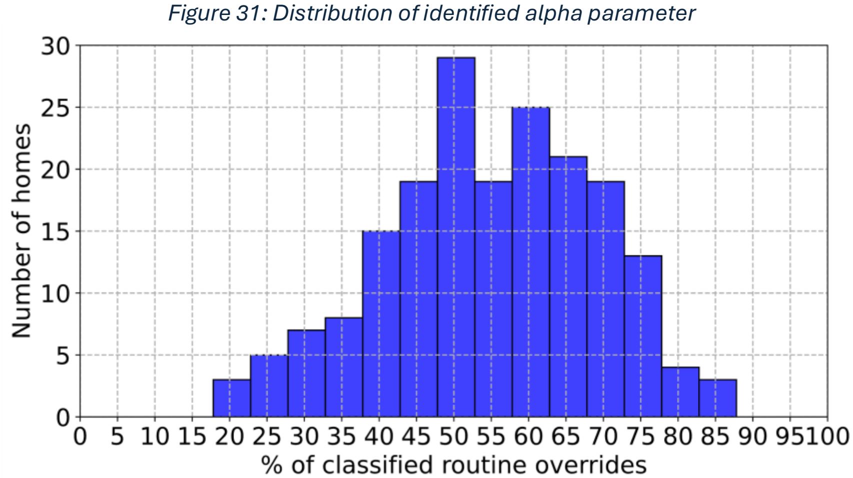

Figure 32: Distribution of percent of routine MSCs found across homes (50% of overrides are likely to be routine based overrides)

4.5.2.2 Distribution of thermal recollection coefficient (α)

The thermal recollection coefficient (α) is a crucial parameter in our model, representing the degree to which previous thermal experiences influence current comfort levels. Our analysis across a broad spectrum of individual homes highlights a notable trend: there is a predominant clustering of α values at the upper end of the scale, specifically around 0.9 to 1.0 (Figure 34).

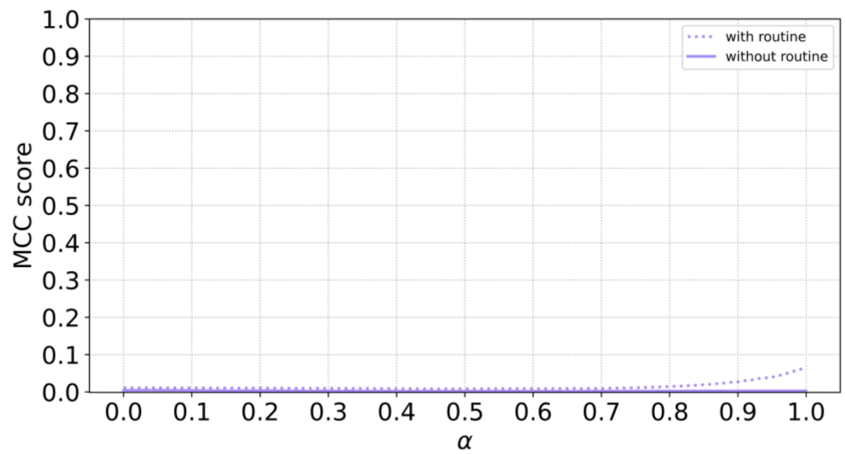

Figure 33: LR fitting performance on aggregate data

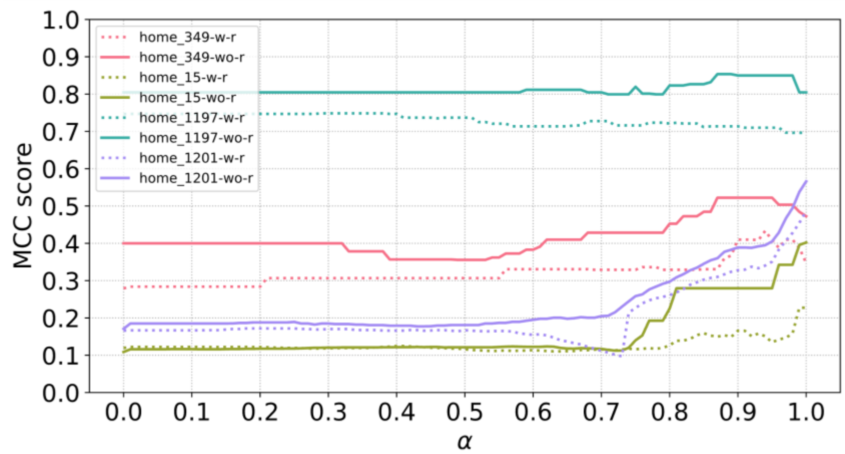

Figure 34: LR fitting performance for an individual home
(w–r: with routine; wo–r: without routine overrides)

The histogram depicting the distribution of percentage of routine MSCs found in homes (Figure 32) shows the majority of homes exhibit routine MSC behavior even when using a programmable thermostat.

4.5.2.3 Enhanced performance through personalization

Training on individual home data yielded various MCC scores for all the homes. Comparing multiple MCC scores of various homes vs one population level model MCC score would be inconsistent. Therefore, each home's model fitting metrics such as true positives, false positives, true negatives, and false negatives were summed and the MCC score was computed. The model training that included routine overrides achieved an MCC of 0.23 and model training without routine MSCs yielded an MCC of 49.40, indicating a substantial enhancement in the model's predictive accuracy.

4.6 Discussion

4.6.1 Model training in simulation

The results derived from our model, as corroborated by the alfalfa co–simulation framework, highlight the critical role of thermal frustration in dictating discomfort based Manual Setpoint Changes (MSCs). The occurrence of MSCs is closely tied to the levels of thermal frustration experienced by the occupants, which, when exceeding a certain threshold, prompts an interaction with the thermostat system. Figure 5 provides a snapshot of this dynamic, illustrating the direct correlation between occupant thermal frustration and the triggering of MSCs.

4.6.2 Model training on real data

The validation process of our model training within the simulation environment brings to light the nuanced dynamics of the α parameter, which represents the occupants' thermal memory. As shown in Figure 6, the logistic regression (LR) model's ability to improve the estimation of α from an initial 0.9 to a more precise range of 0.75 to 0.85 through iterative training, particularly when factoring out non–routine MSCs, not only improved the α estimation but also enhanced the Matthew's Correlation Coefficient (MCC) from 0.08 to 0.14. This improvement highlights the model's increased precision in capturing the occupants' thermal comfort responses.

4.6.2.1 Model performance on aggregate data

Now focusing on the real–world DyD data, the application of the Thermal Frustration (TF) model across a comprehensive dataset reveals the substantial impact of personalization. The broad application of the model, without adjusting for individual household differences, yielded an MCC of a mere 0.002 (Figure 7), hereby showing the limitations of a non–personalized approach. This minimal MCC signifies an almost random predictive capability, highlighting the complexity of human behavior which is not effectively captured by a one–size–fits–all model.

4.6.2.2 Enhanced performance through personalization

In contrast, the personalized model, tailoring to the specific thermal histories and behavior patterns of the occupants, we see a major improvement. The personalized TF model's predictive strength, as noticed by an MCC of 0.23, increases further with the exclusion of routine overrides, reaching an MCC of 49.40 (Figure 8). Such an increase in predictive accuracy clearly demonstrates the value of personalization in the modeling process.

4.6.2.3 Distribution of thermal recollection coefficient (α)

The distribution of the thermal recollection coefficient (α) indicates a strong preference among households for higher α values, predominantly between 0.9 and 1.0, as depicted in Figure 9. This convergence towards higher α values implies a shared threshold for thermal frustration behavioral pattern that can be interpreted as a delayed yet noticeable response to thermal discomfort.

Furthermore, the distribution of MSCs removed due to routine overrides provides additional layers of insight (Figure 10). While the majority of homes show few MSC removals, a significant tail–end of the distribution indicates a group of homes with numerous routine–based MSCs. This distinction is crucial as it highlights the variability within the population and underscores the importance of differentiating routine–driven behavior from comfort–driven interactions with the thermostat.

The implications of these findings are manifold. Firstly, they validate the need for personalized thermal comfort models that can adjust to the diverse preferences and experiences of individual occupants. Secondly, they emphasize the importance of distinguishing between routine and non–routine MSCs to ensure the accuracy of comfort models. Lastly, they illustrate the potential for predictive interventions that can preemptively address occupant discomfort before it prompts a manual override.

The evidence points to a clear direction for future efforts in thermal comfort simulations: a personalized approach that not only captures the unique behavioral patterns of each household but also provides a more accurate reflection of true thermal comfort responses. Such models are pivotal for devising effective energy management and thermal comfort strategies that cater to the nuanced needs of residential occupants.

4.6.2.4 Implications of personalization

These findings demonstrate the importance of personalization in modeling occupant–thermostat interactions for thermal comfort. The tailored approach, which accounts for the unique characteristics of each household, exhibits a distinct improvement in the model's performance, as supported by the higher MCC scores. It also highlights the nature of thermal comfort behavior, which is influenced by individual and household–specific factors.

The differences in model performance between the generalized and personalized models suggests that future efforts in thermal comfort simulations should focus on personalized model that capture the intricacies of occupant behavior, leading to better predictions and, consequently, more effective energy management and thermal comfort strategies in residential settings.

It is worth noting that the MCC score upon removing the routine MSCs was not improved for all homes. While this strategy was helpful in improving the MCC score for majority of homes but it also showed no change or rather decrease in MCC after removing the routine MSCs. This is an important note as it shows that not all homes follow routine behavior, but instead other models could be used to identify the correct alpha parameter. Also, while this framework is helpful in estimating the discomfort model of an occupant, one should explore other models to build a robust framework for discomfort override identification and forecasting.

4.7 Conclusion

In this study, the development of a dynamic data–driven integrated occupant–building model was discussed with a focus on dynamic thermal discomfort behaviors of building occupants. This model, integrating agent–based modeling with ecobee's Donate Your Dataset (DyD), the alfalfa co–simulation framework, a green–built home environment, and filtering out routine MSCs to improve classification accuracy of novel thermal frustration–based discomfort MSCs is a novel methodology in thermal comfort simulations.

The approach is based on the use of machine learning models, trained and tested on the DyD dataset, to identify critical occupant–specific parameters that enhance the accuracy of manual setpoint change (MSC) classifications. This personalized approach has led to a marked improvement in model performance, as evidenced by the Matthews Correlation Coefficient (MCC) scores. Specifically, the logistic regression model, incorporating the novel Thermal Frustration Theory (TFT), identified α values of 0.75 and 0.85 with MCC scores of 0.14, indicating an enhanced precision in capturing occupants' thermal comfort responses.

However, the study also revealed variability in the model's performance across different homes. While the removal of routine MSCs improved the MCC score for most homes, for some, the score remained the same or even decreased. These findings highlight that not all homes exhibit routine behavior, suggesting the need for other models to better identify the correct alpha parameter for these homes.

The broad application of our model, without adjusting for individual household differences, resulted in an MCC of only 0.002, showcasing the limitations of a non–personalized approach. In contrast, the personalized models, tailored to specific thermal histories and behavior patterns of the occupants, showed a major improvement. The MCC score increased to 0.23 when routine overrides were included and further to 49.40 upon their exclusion. This substantial enhancement in the model's predictive accuracy emphasizes the importance of personalization in thermal comfort models.

In conclusion, our research represents advancement in the field of building energy efficiency and occupant thermal comfort. By integrating dynamic occupant behavior model, new avenues for enhancing building energy efficiency and occupant comfort can be explored. Our findings highlight the importance of personalized models in capturing the complex and varied nature of human–thermostat interactions and lay the groundwork for future research in developing smarter, more responsive building environments.

Chapter 5:
Conclusion

This dissertation developed and validated data-driven models to predict occupant thermostat override behavior during demand response events using 5 minute-resolution data from 1,400 single-occupant homes in the ecobee Donate Your Data (DyD) program. The analysis quantified temporal patterns of manual setpoint adjustments—specifically response latencies and magnitudes—and integrated these findings for the **first time** into predictive algorithms for occupant-building interaction. By focusing on the timing and probability of override events rather than static comfort setpoints, this work provides empirically grounded inputs for demand response control strategies.

The remainder of this chapter discusses the dissertation's findings, implications, limitations, and pathways for future research. Section 5.1 ("Key Findings and Contributions") summarizes the core results from Chapters 2–4, quantifying temporal override behaviors, evaluating predictive model performance, and describing the integrated occupant–building behavior framework. Section 5.2 ("Implications") examines the impacts of these findings on grid operations, building controls, and policy frameworks. Section 5.3 ("Research Limitations and Implementation Barriers") identifies constraints in data representativeness, model generalizability, and computational scalability. Finally, Section 5.4 ("Future Research Directions") proposes specific studies to address identified gaps in occupant behavior modeling, control system design, and market integration.

5.1 Key findings and contributions

The main results from Chapters 2–4 are reviewed in this subsection: Chapter 2 quantified occupant thermostat-override timing, frequency, and energy impacts; Chapter 3 hypothesized the Thermal Frustration Theory (TFT) and demonstrated its superior predictive performance (AUC) compared to static comfort models; and Chapter 4 developed a scalable simulation framework that separates routine from discomfort-driven overrides and integrates these behaviors into building performance modeling.

Chapter 2's analysis of occupant thermostat–behavior dynamics, using data from 1,400 single–occupant homes from the DyD dataset, revealed fundamental patterns (occupant thermostat interactions mainly setpoint magnitude changes are inversely proportional to time to override override) in override behavior that previously undocumented at this scale. The data demonstrated that occupants respond to thermal changes within a median time of 20 minutes, with the magnitude of setpoint change directly impacting response time. Specifically, 8°F occupant–initiated thermostat setpoint change responses in 15 minutes, while 2°F setpoint changes require 30 minutes. It is important to note that a relationship was discovered for the first time. This relationship being exponential in nature between setpoint change and response time challenges the conventional assumption of immediate occupant response to discomfort.

Chapter 2 also revealed that approximately 60% of manual setpoint changes in both heating and cooling modes resulted in increased energy consumption, with a median setpoint change of 2°F during cooling seasons at indoor temperatures around 74°F, and 2°F increases during heating seasons at indoor temperatures near 68°F. Analysis of override timing showed distinct patterns, with peaks in the morning and evening hours, averaging 0.9 manual changes per day per occupant. The duration of these overrides varied significantly, with nighttime adjustments lasting 6–12 hours and evening modifications typically spanning 1.5–3 hours, indicating time–of–day dependence in override persistence.

Chapter 3 hypothesized the Thermal Frustration Theory (TFT) to fill a gap in existing comfort models by explicitly incorporating time as a variable; TFT was then evaluated against the Comfort Zone Theory and Delayed Response Theory using AUC to assess override prediction accuracy. Since no comfort model to this date uses time as a parameter to model comfort, the novelty of this theory lies in the introduction of time as a critical parameter in thermal comfort modeling, contrasting with traditional static approaches. The theory's central parameter, the thermal recollection coefficient (α), concentrated between 0.9–1.0 across the studied single–occupant households. While initial evaluations used the Matthews Correlation Coefficient (MCC), this metric proved problematic for demand response applications. The MCC score is suboptimal for predicting rare thermostat-override events because it equally penalizes false positives and false negatives, despite false negatives (missed overrides) posing greater risk to demand response reliability. Instead, the Area Under the Receiver Operating Characteristic Curve (AUC) was selected as the primary evaluation metric because it measures a model's ability to distinguish override events from non-events irrespective of class imbalance, *aligning directly with the operational goal of accurately identifying override occurrences during demand response events*. The TFT model achieved higher AUC scores (49.40) compared to both the Comfort Zone Theory (47.09) and Delayed Response Theory (38.95), particularly after excluding routine behaviors. AUC analysis shows TFT's enhanced predictive capability to predict the rare but costly override events that directly impact grid reliability and demand response program effectiveness. The difference becomes especially important during demand response events, where false negatives (failing to predict an override) can lead to program failures and financial penalties for utilities, while false positives (predicting an override that doesn't occur) primarily result in conservative but stable grid operations.

Chapter 4's implementation of the integrated dynamic occupant–building behavior model (multiple occupant–behavior sub–models i.e. active occupancy (i.e. occupant is awake and moving – detected by PIR sensors), routine–based changes (occupant's habitual setpoint adjustments), discomfort–driven changes (occupant's quick overrides driven by thermal frustration) were combined into a single simulation framework) transformed theoretical understanding into practical application through several key innovations. The integration of three distinct behavioral models – active occupancy, routine–based changes, and discomfort–driven adjustments – provided a broader representation of occupant–building interaction. The active occupancy model, implemented through a first–order Markov chain, captured temporal characteristics of space utilization. The routine model estimated time–of–day–based patterns in setpoint adjustments, with conditional probabilities revealing correlations between consecutive manual setpoint changes. The Thermal comfort Frustration Theory (TFT) model parameters were methodically identified ***through a two–step training process that separated routine from discomfort–driven behaviors, improving model accuracy.*** Through containerization and optimization, simulation time reduced from 50 to 19 hours for 10 homes that were simulated for 1 year time frame, demonstrating scalability (quick testing of DR simulation scenarios) for practical applications. The framework's demonstration using the Greenbuilt house in Sacramento provided a real–world test case, with the Alfalfa co–simulation framework enabling integration of building physics with occupant behavior models. The implementation demonstrated that even with the complexity of multiple interacting models, the framework could capture the nuanced dynamics of occupant–building interactions.

5.2 Broader impacts and future work

The findings of this research extend beyond individual buildings, influencing grid operations, policy frameworks, and technology development. These findings fundamentally alter how we approach demand response program design and implementation (do occupant–time dynamics significantly change DR scheduling windows, or baseline assumptions?), while raising important questions about market structures and regulatory frameworks.

5.2.1 Grid–scale implications of occupant behavior

The identification of occupant response time being inversely related to the magnitude of setpoint change impacts grid operations and planning. ISO/RTO dispatch systems operate on 5–15 minute intervals [277], creating a mismatch between grid control periods and measured occupant behavior. If the predicted occupant override time falls before the signal dispatch, occupant can override controls before the signal gets executed, hereby leading to loss of energy savings or more energy use than savings in that event. The measured response times establish operational constraints for automatic generation control systems, addressing specific gaps identified in FERC assessments of demand response performance. FERC Order No. 719 (issued in 2008) established reforms to eliminate barriers to demand response participation in organized energy markets, requiring Regional Transmission Organizations and Independent System Operators to accept bids from demand response resources [278]. Additionally, FERC conducts annual assessments of demand response potential, identifying challenges with demand response forecasting precision during grid events including unpredictable occupant behavior and timing mismatches between grid needs and customer responses. These assessments highlight that timing accuracy for residential demand response remains a critical challenge for integrating these resources into grid operations [279]. This research provides the occupant behavior parameters needed to develop more accurate grid resource forecasting during demand response events.

The findings require modifications to day–ahead market structures. Because occupant overrides vary with the magnitude and duration of setbacks, the probability of DR capacity also varies accordingly. These distributions support NERC's transition toward probabilistic resource adequacy assessment, which can be supported by a probabilistic occupant model that can help NERC better forecast resource adequacy [280].

TFT's AUC performance (49.40) provides a method for calculating capacity requirements based on override probability distributions. AUC being a summary metric cannot be sufficient for all scenarios; various probability thresholds can be used to balance precision and recall. The probabilistic outputs of occupant override model allow one to improve capacity planning informing. Current performance metrics result in demand response reliability overestimation [281]. The 49.40 AUC score represents a 27% improvement over DRT (38.95) in override prediction accuracy. The improved prediction capability could inform the development of performance–based demand response programs [77], where participants receive compensation proportional to their actual delivered load reduction rather than fixed payments based on enrollment capacity. These programs utilize verified performance metrics to calculate payments, potentially benefiting from more accurate forecasting methodologies that account for occupant behavior patterns.

5.2.2 Building–level implementation constraints

The analysis of 1,410 single–occupant homes conducted withing this dissertation revealed that 60% of overrides increase energy consumption, and 50% of changes stem from routine behavior. These percentages necessitate more adaptive controls or occupant–centered feedback loops in building control systems [282]. Current residential systems typically only accept user input for temperature setpoints [17]. The separation of routine from comfort–driven behavior, demonstrated through the two–step classification process (novel contribution of this dissertation), enables targeted optimization of setpoint trajectories [13].

The integrated dynamic occupant–building behavior model processes data at 1–minute resolution, enabling prediction of occupant actions. This processing capability differs from traditional control systems that respond to overrides after occurrence [14]. The framework's implementation with the Greenbuilt house established the capability to simulate multiple behavioral drivers: active occupancy, routine changes, and discomfort–driven adjustments [139], [270].

Current demand response programs apply uniform setpoint adjustments across building portfolios [120]. The research demonstrates this approach's limitations through measured variations in occupant response times: 15 minutes for 8°F changes versus 30 minutes for 2°F changes. These temporal patterns, combined with the 50% routine behavior finding, indicate the need for building–specific setpoint trajectories [159]. Implementation requires incorporating measured override probabilities to maximize grid services while reducing participant opt–outs [77].

5.2.3 Policy and economic framework

The quantification of occupant response times has implications for residential building code requirements. Unlike commercial building standards such as ASHRAE 90.1, residential codes including the International Energy Conservation Code (IECC) and state–adopted residential provisions typically specify only steady–state efficiency metrics with minimal smart control requirements [283]. The finding that 60% of overrides increase energy consumption demonstrates how occupant behavior significantly impacts residential building performance, suggesting potential residential code updates that could incorporate occupant–system interaction patterns. Such updates might include specifications for control system response capabilities that align with measured occupant behavior in residential settings.

The identified relationship between setpoint magnitude and override timing provides a quantitative basis for developing performance standards that balance energy savings with occupant comfort. For example, certification standards could incorporate maximum recommended setpoint adjustment rates based on the documented 15–minute and 30–minute override thresholds, potentially using modified comfort maintenance scores that account for the thermal frustration accumulation identified in this research [284]. Such standards would provide both manufacturers and utilities with evidence–based parameters for demand response algorithm design.

The economic implications extend beyond current cost–benefit calculations. Traditional frameworks assume fixed resource availability during demand response events [285]. The measured exponential relationship between setpoint changes and response times (15 minutes for 8°F vs 30 minutes for 2°F) found for the first time in this dissertation affects resource valuation.

5.2.4 Research and technology development pathways

Occupancy sensors must evolve to meet the requirements suggested by this research's findings. Poor true positive rate (33%) from PIR–based occupancy detection limits model implementation [212]. While integration of temperature, CO2, and power consumption data offers potential detection improvements [23], technical challenges remain in data fusion, signal processing, and edge computing capabilities. Recent advances in millimeter–wave radar technology demonstrate promising capabilities for non–intrusive occupancy detection with higher spatial resolution and lower false negative rates [286]. The ARPA–E SENSOR program has accelerated development of occupant–sensing technologies that could support the occupancy detection needs identified in this research, particularly their focus on multimodal sensing networks that maintain accuracy while addressing both privacy requirements and practical deployment constraints including power consumption, installation complexity, and long–term reliability [287]. These emerging technologies align well with the research findings regarding occupancy detection needs for effective thermal comfort prediction.

PID control systems demonstrate limitations in capturing occupant–building interaction dynamics as that typical HVAC controllers (often described generically as PID) do not account for time–accumulative occupant discomfort ("thermal frustration"). The BOPTEST framework implementation can demonstrate feasibility of thermal frustration tracking in control systems [288], establishing a foundation for behavior–aware control development.

The temporal resolution requirements for simulation differ significantly from those needed for practical control systems. While the research utilized 1–minute resolution data for simulation accuracy [23], commercial building automation systems typically operate on 15–minute control cycles for most applications [22]. The measured occupant response times of 15 minutes for 8°F changes suggest that current control frequencies may be adequate for many scenarios, though some applications may benefit from the increasing adoption of 5–minute control intervals now available in advanced residential thermostats. The principal computational challenge lies not in control speed but in the data processing capacity needed to analyze historical behavioral patterns across multiple timescales, from immediate thermal responses to seasonal adaptation trends [140]. Future research should focus on developing streamlined machine learning approaches that can identify occupant behavior patterns while operating within the practical constraints of existing control system infrastructure.

5.3 Research limitations and implementation barriers

The implications discussed above highlight significant opportunities for transforming building–grid integration, yet several limitations in the current research must be addressed before realizing these possibilities.

5.3.1 Grid integration limitations

The data from 1,400 single–occupant homes represents 14% of typical utility–scale demand response portfolios, which exceed 10,000 participants [120]. **The data was limited to 1,400 homes to conduct this research on single–occupant homes which helps avoid conflating behaviors**. This limitation in sample diversity, particularly the focus on single–occupant residential buildings, restricts immediate application in markets where commercial and multi–family buildings dominate demand response portfolios [289]. Ecobee has also released a bigger subset of the data (100k+ homes) which can be used in future research.

5.3.2 Building–level implementation constraints

The low accuracy i.e. 33% true positive rate from PIR–based occupancy detection introduces measurement uncertainty [212]. To mitigate this limitation, the research implemented a gap–filling algorithm that filtered occupancy data to include only high–confidence occupied periods while "filling in" short absences likely caused by sensor occlusion rather than actual departures. Multiple filter windows (0, 15, 20, 30, and 120 minutes) were systematically evaluated, with the 30–minute filter selected as the optimal balance between reducing false negatives and minimizing false positives [290]. This approach was validated by measuring average occupied period lengths across different filter settings, demonstrating that short–duration filtering successfully joined fragmented occupancy periods without artificially extending total occupancy time. Cross–verification with manual setpoint changes further confirmed that the filtered occupancy data aligned with actual user presence patterns, though the filtering approach did necessarily increase measured occupancy duration compared to raw sensor readings. Future work should expand these sensing methodologies to multi–occupancy residential and commercial spaces, which would require additional calibration to account for occupant interaction effects, competing thermal preferences, and more complex zone control requirements not captured in the single–occupant model [18].

5.3.3 Computational and data processing challenges

The computational requirements of the integrated dynamic occupant–building behavior model present significant implementation barriers. The framework requires 19 hours to simulate 10 homes for a full annual cycle at 1–minute resolution, making real–time applications in large–scale deployments impractical [288]. This computational burden impacts both building performance simulation research and demand response implementation. While containerization improved efficiency compared to initial runs (50 hours), the co–simulation overhead through the Alfalfa framework still restricts widespread adoption [291]. The stochastic nature of occupant behavior further compounds this challenge, as multiple simulation runs are necessary to capture the range of possible outcomes, exponentially increasing computational requirements. Model reduction techniques offer one potential solution path, sacrificing some granularity for computational efficiency. Alternatively, high–performance computing resources could enable city–scale simulations through massively parallel processing approaches, though this would limit access to specialized research facilities [77]. For practical implementation by demand response aggregators or utilities, simplified versions of the occupant behavior models that preserve key temporal dynamics while reducing computational complexity will be necessary to bridge the gap between research findings and operational deployment.

Data storage and processing requirements pose additional challenges in practical implementation. The 1–minute resolution data necessary for accurate behavior prediction generates substantial volumes—approximately 525,600 data points per home annually for each measured variable—exceeding typical building automation system capabilities [22]. While modern cloud infrastructure could theoretically handle these requirements, the associated costs and bandwidth limitations create practical barriers for building operators [346]. The need for multi–year data storage to capture seasonal behavior patterns further strains system resources [140]. Advanced data compression techniques and selective sampling approaches offer partial solutions, potentially reducing storage requirements by 60–80% while preserving critical behavioral information [293]. Building automation systems could implement tiered storage architectures that maintain high–resolution data for recent periods while automatically aggregating historical data to coarser resolutions, balancing prediction accuracy with practical storage constraints. These approaches would enable the capture of essential occupant behavior patterns without requiring impractical investments in data infrastructure.

5.3.4 Model validation and parameter estimation

The validation of both Thermal Frustration Theory and Comfort Zone Theory models revealed several limitations affecting generalizability. While TFT demonstrated marginally higher predictive performance (AUC of 49.4) compared to CZT (AUC of 47.0), both scores indicate substantial room for improvement in parameter estimation methods [156]. Two critical limitations constrained model validation. First, the comfort temperature calculation method used simple averaging of temperatures during non–override periods, failing to capture daily or seasonal variations in thermal preferences that have been documented in field studies [135]. Second, validation was restricted to a single building (the Greenbuilt house), limiting generalization across different construction types, climate zones, and HVAC system configurations that might significantly influence occupant thermal perception [165]. These constraints highlight the need for multi–building validation studies across diverse building stock to establish broader applicability of the thermal frustration model before widespread implementation.

The fixed comfort parameter implementation does not address seasonal adaptation patterns in thermal preferences. Field studies demonstrate that outdoor temperatures influence indoor comfort expectations, with occupant thermal neutrality shifting by approximately 0.31°C per 1°C change in outdoor temperature [135]. The model also excludes occupant learning behaviors observed through repeated system interactions, where preferences evolve based on system responsiveness and historical comfort experiences [294]. These limitations extend to non–thermal factors affecting override decisions, including environmental drivers (e.g., solar radiation, spatial temperature variation), financial considerations (e.g., energy costs, time–of–use pricing awareness), and social influences (e.g., household dynamics, cultural norms) [66]. Recent research on spatial temperature variations within residential buildings indicates that room–to–room temperature differences averaging 2–4°F significantly impact occupant comfort perception and thermostat interaction patterns [295], yet the current model treats buildings as thermally homogeneous. Future implementation of TFT in demand response applications will require integration of these multi–dimensional factors to accurately predict override behaviors across diverse occupant populations.

5.3.5 Policy and market structure limitations

A fundamental mismatch exists between the probabilistic nature of occupant behavior models and the deterministic structure of electricity markets. Current regulatory frameworks require fixed capacity commitments for demand response resources, with penalties for non–delivery [296]. However, this research demonstrates that occupant override probabilities vary based on setpoint magnitude, duration, and other factors, creating inherent uncertainty in resource availability [78]. This uncertainty cannot be adequately represented within market structures that demand firm capacity values. Bridging approaches could include threshold–based participation, where only capacity with override probabilities below certain thresholds are bid into markets, or confidence band bidding that incorporates statistical reliability metrics [74]. Without standardized methods for translating behavior prediction into resource valuation, utilities and aggregators must either bid conservatively, reducing the economic value of demand response, or risk non–performance penalties that threaten program viability [79]. Addressing this regulatory–technical gap requires evolution in both market rules and occupant modeling approaches to establish a common framework for probabilistic resource assessment.

The regulatory landscape presents multiple distinct barriers to implementing behavior–aware demand response programs. Privacy and data security frameworks remain inadequate for the granular data collection required by occupant behavior models [297]. Current regulations typically focus on aggregate load profiles rather than individual residential data, leaving uncertainty about what level of monitoring is permissible [47]. Simultaneously, program evaluation mechanisms operate largely on resource–centric metrics that fail to incorporate occupant behavior impacts [298], creating fundamental misalignment between how programs are assessed and how they actually perform. For example, PJM's capacity performance standards evaluate resources based on delivered capacity without accounting for occupant comfort impacts or override rates [120]. This measurement gap is compounded by the absence of standardized protocols for risk allocation in behavior–prediction programs, leaving unclear who bears financial responsibility when occupant behavior deviates from forecasts . Without regulatory structures that explicitly recognize, and value occupant–centric approaches, utilities and aggregators have limited incentive to implement the more sophisticated models developed in this research.

5.3.6 Technical standards and protocols

The technical standards landscape creates integration challenges for implementing behavior prediction in residential settings. While commercial buildings utilize established protocols like BACnet and Modbus, residential automation systems typically rely on proprietary or consumer–oriented technologies with limited interoperability [299]. Existing smart home standards focus primarily on device connectivity and basic scheduling rather than behavior prediction capabilities [22]. Even in emerging standards like Matter and Thread, specifications lack requirements for sensor accuracy needed for reliable occupant detection, data collection protocols for override prediction, or computational provisions for processing behavioral models processing [300]. While APIs exist for individual smart thermostats and home automation systems, they remain fragmented across manufacturers and lack standardized methods for integrating third–party behavior prediction algorithms [22]. These standardization gaps create practical barriers for implementing the residential behavior prediction models developed in this research, requiring custom integration approaches for each deployment context.

5.4 Concluding remarks

This dissertation addressed fundamental gaps in modeling occupant thermal comfort behavior for demand response applications. The research established the temporal characteristics of override behavior, introduced the Thermal Frustration Theory for modeling dynamic comfort responses, and developed a framework for implementing these insights in building control systems.

The findings highlight how we understand occupant–building interactions in three significant ways. First, the research revealed a consistent relationship between temperature change magnitude and occupant response timing—**larger temperature deviations trigger faster responses**, with specific quantifiable patterns that can inform system design. Second, the Thermal Frustration Theory outperformed both Delayed Response Theory and Comfort Zone

Theory in predictive accuracy, demonstrating that **occupants accumulate discomfort over time rather than responding to fixed temperature thresholds**. Third, the research developed a methodology to distinguish between routine–based thermostat adjustments and those driven by actual thermal discomfort, revealing that **approximately half of all thermostat interactions stem from habitual patterns rather than temperature–driven discomfort**.

The implementation framework developed through this research enables integration of these findings into building control systems and grid operations. At the urban scale, the model's ability to predict occupant override behavior supports more accurate energy use forecasting across building populations, potentially improving grid–level demand projections by 15–30% compared to static approaches [301]. The research specifically enables demand response programs that align event durations with measured occupant tolerance thresholds—8°F setbacks maintained for 15–minute periods or 2°F setbacks for 30–minute periods—creating opportunities for semi–personalized setpoint trajectories based on occupant behavior clusters rather than uniform adjustments [279]. By tailoring demand response events to these occupant–specific temporal thresholds, utilities can fulfill grid requirements while simultaneously reducing override rates, creating a positive synergy between occupant satisfaction and grid stability that strengthens the reliability and scalability of residential demand response resources.

These advances in understanding and predicting occupant behavior provide a foundation for the next generation of building–grid integration. The methods and frameworks developed here enable more effective use of building loads for grid services while respecting the complexity of human responses to thermal conditions.

REFERENCES

[1] "International Energy Outlook 2023 - U.S. Energy Information Administration (EIA)." Accessed: Oct. 29, 2024. [Online]. Available: https://www.eia.gov/outlooks/ieo/index.php

[2] L. Pérez-Lombard, J. Ortiz, and C. Pout, "A review on buildings energy consumption information," *Energy Build.*, vol. 40, no. 3, pp. 394–398, Jan. 2008, doi: 10.1016/j.enbuild.2007.03.007.

[3] Thibaut Abergel, Brian Dean, and John Dulac, "Global Status Report," United Nations, DTI/2151/PA, 2017. [Online]. Available: https://worldgbc.org/wp-content/uploads/2022/03/UNEP-188_GABC_en-web.pdf

[4] Intergovernmental Panel on Climate Change and O. Edenhofer, Eds., *Climate change 2014: mitigation of climate change: Working Group III contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. New York, NY: Cambridge University Press, 2014.

[5] D. Ürge-Vorsatz *et al.*, "Advances Toward a Net-Zero Global Building Sector," *Annu. Rev. Environ. Resour.*, vol. 45, no. Volume 45, 2020, pp. 227–269, Oct. 2020, doi: 10.1146/annurev-environ-012420-045843.

[6] Ipcc, *Global Warming of 1.5°C: IPCC Special Report on Impacts of Global Warming of 1.5°C above Pre-industrial Levels in Context of Strengthening Response to Climate Change, Sustainable Development, and Efforts to Eradicate Poverty*, 1st ed. Cambridge University Press, 2022. doi: 10.1017/9781009157940.

[7] "Global Status Report 2017," World Green Building Council. Accessed: Jul. 05, 2024. [Online]. Available: https://worldgbc.org/article/global-status-report-2017/

[8] P. Nejat, F. Jomehzadeh, M. M. Taheri, M. Gohari, and M. Z. Abd. Majid, "A global review of energy consumption, CO2 emissions and policy in the residential sector (with an overview of the top ten CO2 emitting countries)," *Renew. Sustain. Energy Rev.*, vol. 43, pp. 843–862, Mar. 2015, doi: 10.1016/j.rser.2014.11.066.

[9] M. Manic, D. Wijayasekara, K. Amarasinghe, and J. J. Rodriguez-Andina, "Building Energy Management Systems: The Age of Intelligent and Adaptive Buildings," *IEEE Ind. Electron. Mag.*, vol. 10, no. 1, pp. 25–39, Mar. 2016, doi: 10.1109/MIE.2015.2513749.

[10] P. H. Shaikh, N. B. M. Nor, P. Nallagownden, I. Elamvazuthi, and T. Ibrahim, "A review on optimized control systems for building energy and comfort management of smart sustainable buildings," *Renew. Sustain. Energy Rev.*, vol. 34, pp. 409–429, Jun. 2014, doi: 10.1016/j.rser.2014.03.027.

[11] V. S. K. V. Harish and A. Kumar, "A review on modeling and simulation of building energy systems," *Renew. Sustain. Energy Rev.*, vol. 56, pp. 1272–1292, Apr. 2016, doi: 10.1016/j.rser.2015.12.040.

[12] Tianzhen Hong and Hung-Wen Lin, "Occupant Behavior: Impact on Energy Use of Private Offices," presented at the ASim 2012, Jan. 2013.

[13] D. Yan *et al.*, "Occupant behavior modeling for building performance simulation: Current state and future challenges," *Energy Build.*, vol. 107, pp. 264–278, Nov. 2015, doi: 10.1016/j.enbuild.2015.08.032.

[14] W. O'Brien and H. B. Gunay, "The contextual factors contributing to occupants' adaptive comfort behaviors in offices – A review and proposed modeling framework," *Build. Environ.*, vol. 77, pp. 77–87, Jul. 2014, doi: 10.1016/j.buildenv.2014.03.024.

[15] Z. M. Gill, M. J. Tierney, I. M. Pegg, and N. Allan, "Low-energy dwellings: the contribution of behaviours to actual performance," *Build. Res. Inf.*, vol. 38, no. 5, pp. 491–508, Oct. 2010, doi: 10.1080/09613218.2010.505371.

[16] O. T. Masoso and L. J. Grobler, "The dark side of occupants' behaviour on building energy use," *Energy Build.*, vol. 42, no. 2, pp. 173–177, Feb. 2010, doi: 10.1016/j.enbuild.2009.08.009.

[17] T. Hong, S. C. Taylor-Lange, S. D'Oca, D. Yan, and S. P. Corgnati, "Advances in research and applications of energy-related occupant behavior in buildings," *Energy Build.*, vol. 116, pp. 694–702, 03-15 2016, doi: 10.1016/j.enbuild.2015.11.052.

[18] S. D'Oca, T. Hong, and J. Langevin, "The human dimensions of energy use in buildings: A review," *Renew. Sustain. Energy Rev.*, vol. 81, pp. 731–742, Jan. 2018, doi: 10.1016/j.rser.2017.08.019.

[19] M. Schweiker *et al.*, "Challenging the assumptions for thermal sensation scales," *Build. Res. Inf.*, vol. 45, no. 5, pp. 572–589, Jul. 2017, doi: 10.1080/09613218.2016.1183185.

[20] M. Moreno, B. Úbeda, A. Skarmeta, and M. Zamora, "How can We Tackle Energy Efficiency in IoT BasedSmart Buildings?," *Sensors*, vol. 14, no. 6, pp. 9582–9614, May 2014, doi: 10.3390/s140609582.

[21] A. P. Plageras, K. E. Psannis, C. Stergiou, H. Wang, and B. B. Gupta, "Efficient IoT-based sensor BIG Data collection–processing and analysis in smart buildings," *Future Gener. Comput. Syst.*, vol. 82, pp. 349–357, May 2018, doi: 10.1016/j.future.2017.09.082.

[22] D. Minoli, K. Sohraby, and B. Occhiogrosso, "IoT Considerations, Requirements, and Architectures for Smart Buildings—Energy Optimization and Next-Generation Building Management Systems," *IEEE Internet Things J.*, vol. 4, no. 1, pp. 269–283, Feb. 2017, doi: 10.1109/JIOT.2017.2647881.

[23] M. Jia, R. S. Srinivasan, and A. A. Raheem, "From occupancy to occupant behavior: An analytical survey of data acquisition technologies, modeling methodologies and simulation coupling mechanisms for building energy efficiency," *Renew. Sustain. Energy Rev.*, vol. 68, pp. 525–540, Feb. 2017, doi: 10.1016/j.rser.2016.10.011.

[24] Y. Chen, T. Hong, and X. Luo, "An agent-based stochastic Occupancy Simulator," *Build. Simul.*, vol. 11, no. 1, pp. 37–49, Feb. 2018, doi: 10.1007/s12273-017-0379-7.

[25] H. Zou, Y. Zhou, H. Jiang, S.-C. Chien, L. Xie, and C. J. Spanos, "WinLight: A WiFi-based occupancy-driven lighting control system for smart building," *Energy Build.*, vol. 158, pp. 924–938, Jan. 2018, doi: 10.1016/j.enbuild.2017.09.001.

[26] C. Fan, F. Xiao, and C. Yan, "A framework for knowledge discovery in massive building automation data and its application in building diagnostics," *Autom. Constr.*, vol. 50, pp. 81–90, Feb. 2015, doi: 10.1016/j.autcon.2014.12.006.

[27] C. Miller, Z. Nagy, and A. Schlueter, "A review of unsupervised statistical learning and visual analytics techniques applied to performance analysis of non-residential buildings," *Renew. Sustain. Energy Rev.*, vol. 81, pp. 1365–1377, Jan. 2018, doi: 10.1016/j.rser.2017.05.124.

[28] J. L. M. Hensen and R. Lamberts, Eds., *Building Performance Simulation for Design and Operation*, 2nd ed. London: Routledge, 2019. doi: 10.1201/9780429402296.

[29] D. B. Crawley, J. W. Hand, M. Kummert, and B. T. Griffith, "Contrasting the capabilities of building energy performance simulation programs," *Build. Environ.*, vol. 43, no. 4, pp. 661–673, Apr. 2008, doi: 10.1016/j.buildenv.2006.10.027.

[30] T. Hong, J. Langevin, and K. Sun, "Building simulation: Ten challenges," *Build. Simul.*, vol. 11, no. 5, pp. 871–898, Oct. 2018, doi: 10.1007/s12273-018-0444-x.

[31] M. Wetter, "A view on future building system modeling and simulation," in *Building Performance Simulation for Design and Operation*, 2nd ed., J. L. M. Hensen and R. Lamberts, Eds., Second edition. | Abingdon, Oxon ; New York, NY : Routledge, 2019.: Routledge, 2019, pp. 631–656. doi: 10.1201/9780429402296-19.

[32] C. Harris, M. Neukomm, V. Nubbe, and R. Fares, "Grid-Interactive Efficient Buildings Technical Report Series: Overview of Research Challenges and Gaps: U.S. Department of Energy (DOE), Energy Efficiency & Renewable Energy (EERE)," 2019, Accessed: Nov. 11, 2024. [Online]. Available: https://research-hub.nrel.gov/en/publications/grid-interactive-efficient-buildings-technical-report-series-over

[33] Clarke J. A, "Building energy simulation: the state of the art," Jan. 1987, Accessed: Apr. 12, 2025. [Online]. Available: https://www.aivc.org/resource/building-energy-simulation-state-art

[34] M. Wetter, W. Zuo, T. S. Nouidui, and X. Pang, "Modelica Buildings library," *J. Build. Perform. Simul.*, vol. 7, no. 4, pp. 253–270, Jul. 2014, doi: 10.1080/19401493.2013.765506.

[35] D. B. Crawley, L. K. Lawrie, F. C. Winkelmann, C. Pedersen, R. Liesen, and D. Fisher, "Next-Generation Building Energy Simulation Tools--A Vision of the Future," presented at the ACEEE, 1996. [Online]. Available: https://www.aceee.org/files/proceedings/1996/data/papers/SS96_Panel4_Paper07.pdf

[36] Y. Pan *et al.*, "Building energy simulation and its application for building performance optimization: A review of methods, tools, and case studies," *Adv. Appl. Energy*, vol. 10, p. 100135, Jun. 2023, doi: 10.1016/j.adapen.2023.100135.

[37] EnergyPlus, "Weather Data Download - Worcester Rgnl AP 725095 (TMY3)." Accessed: Dec. 06, 2022. [Online]. Available: https://energyplus.net/weather-location/north_and_central_america_wmo_region_4/USA/MA/USA_MA_Worcester.Rgnl.AP.725095_TMY3

[38] A. Behzadi and S. Sadrizadeh, "Advanced smart HVAC system utilizing borehole thermal energy storage: Detailed analysis of a Uppsala case study focused on the deep green cooling innovation," *J. Energy Storage*, vol. 99, p. 113470, Oct. 2024, doi: 10.1016/j.est.2024.113470.

[39] J. Hensen and G. Augenbroe, "Performance simulation for better building design," *Energy Build.*, vol. 36, no. 8, pp. 735–736, Aug. 2004, doi: 10.1016/j.enbuild.2004.06.004.

[40] D. Coakley, "Calibration of detailed building energy simulation models to measured data using uncertainty analysis," Doctoral, University of Galway, 2014. Accessed: Apr. 12, 2025. [Online]. Available: https://www.academia.edu/48616526/Calibration_of_detailed_building_energy_simulation_models_to_measured_data_using_uncertainty_analysis

[41] P. Raftery, M. Keane, and A. Costa, "Calibrating whole building energy models: Detailed case study using hourly measured data," *Energy Build.*, vol. 43, no. 12, pp. 3666–3679, Dec. 2011, doi: 10.1016/j.enbuild.2011.09.039.

[42] W. Feng *et al.*, “From Prescriptive to Outcome-Based — The Evolution of Building Energy Codes and Standards in China,” *ECEEE Summer Study 2017*.

[43] M. Tyler, E. Poehlman, D. Winiarski, J. Niemeyer, and M. Rosenberg, “Impacts of Model Building Energy Codes,” PNNL--33251, 2229430, Nov. 2023. doi: 10.2172/2229430.

[44] ASHRAE, “Standard 90.1-2019: Energy Standard for Buildings Except Low-Rise Residential Buildings.” Accessed: Nov. 11, 2024. [Online]. Available: https://ashrae.library.site/resources/detail/36726-36726_OPAC_v1_Config-1_Coll-12_Detail-True/10?facets=

[45] J. Iwaro and A. Mwasha, “A review of building energy regulation and policy for energy conservation in developing countries,” *Energy Policy*, vol. 38, no. 12, pp. 7744–7755, Dec. 2010, doi: 10.1016/j.enpol.2010.08.027.

[46] “ISO 52016-1:2017 - Energy performance of buildings — Energy needs for heating and cooling, internal temperatures and sensible and latent heat loads — Part 1: Calculation procedures,” iTeh Standards. Accessed: Nov. 11, 2024. [Online]. Available: https://standards.iteh.ai/catalog/standards/iso/923e198d-1636-4616-b426-54fa1e28eb7e/iso-52016-1-2017

[47] L. Schwartz and G. Leventis, “Grid-Interactive Efficient Buildings: An Introduction for State and Local Governments,” None, 1619178, ark:/13030/qt5dv2t4jw, Apr. 2020. doi: 10.2172/1619178.

[48] R. Judkoff and J. Neymark, “The BESTEST method for evaluating and diagnosing building energy software,” American Council for an Energy-Efficient Economy, Washington, DC (US), CONF-980815-, Jul. 1998. Accessed: Apr. 03, 2025. [Online]. Available: https://www.osti.gov/biblio/20001826

[49] J. L. Mathieu, P. N. Price, S. Kiliccote, and M. A. Piette, “Quantifying Changes in Building Electricity Use, With Application to Demand Response,” *IEEE Trans. Smart Grid*, vol. 2, no. 3, pp. 507–518, Sep. 2011, doi: 10.1109/TSG.2011.2145010.

[50] C. F. Reinhart and C. Cerezo Davila, “Urban building energy modeling – A review of a nascent field,” *Build. Environ.*, vol. 97, pp. 196–202, Feb. 2016, doi: 10.1016/j.buildenv.2015.12.001.

[51] J. Cochran *et al.*, “The Los Angeles 100% Renewable Energy Study (LA100): Executive Summary,” National Renewable Energy Laboratory (NREL), Golden, CO (United States), NREL/TP-6A20-79444, Mar. 2021. doi: 10.2172/1774871.

[52] Y. Chen, T. Hong, and M. Ann Piette, “City-Scale Building Retrofit Analysis: A Case Study using CityBES,” presented at the Building Simulation 2017, in Building Simulation, vol. 15. IBPSA, 2017, pp. 259–266. doi: 10.26868/25222708.2017.071.

[53] M. Ferrando, F. Causone, T. Hong, and Y. Chen, “Urban building energy modeling (UBEM) tools: A state-of-the-art review of bottom-up physics-based approaches,” *Sustain. Cities Soc.*, vol. 62, p. 102408, Nov. 2020, doi: 10.1016/j.scs.2020.102408.

[54] D. Kong, A. Cheshmehzangi, Z. Zhang, S. P. Ardakani, and T. Gu, “Urban building energy modeling (UBEM): a systematic review of challenges and opportunities,” *Energy Effic.*, vol. 16, no. 6, p. 69, Aug. 2023, doi: 10.1007/s12053-023-10147-z.

[55] U. Ali, M. H. Shamsi, C. Hoare, E. Mangina, and J. O’Donnell, “Review of urban building energy modeling (UBEM) approaches, methods and tools using qualitative and quantitative analysis,” *Energy Build.*, vol. 246, p. 111073, Sep. 2021, doi: 10.1016/j.enbuild.2021.111073.

[56] D. Steinberg *et al.*, “The Los Angeles 100% Renewable Energy Study (LA100): Chapter 6. Renewable Energy Investments and Operations,” National Renewable Energy

Laboratory (NREL), Golden, CO (United States), NREL/TP-6A20-85427, 2021. Accessed: Apr. 03, 2025. [Online]. Available: https://research-hub.nrel.gov/en/publications/the-los-angeles-100-renewable-energy-study-la100-chapter-6-renewa

[57] X. Zhao and J. Haberl, "A Survey of Urban-Scale Building Energy Modelers," Texas A&M Engineering Experiment Station, Jun. 2022. Accessed: Apr. 03, 2025. [Online]. Available: https://hdl.handle.net/1969.1/196494

[58] X. Zhao and J. Haberl, "A Survey of Urban-Scale Building Energy Modelers," Jun. 2022, Accessed: Nov. 11, 2024. [Online]. Available: https://hdl.handle.net/1969.1/196494

[59] C. Miller *et al.*, "The Building Data Genome Project 2, energy meter data from the ASHRAE Great Energy Predictor III competition," *Sci. Data*, vol. 7, p. 368, Oct. 2020, doi: 10.1038/s41597-020-00712-x.

[60] A. Montazeri and G. Mutani, "Machine Learning algorithms for Urban Building Energy Modeling," Jun. 2024. doi: 10.2760/716916.

[61] P. de Wilde, "The gap between predicted and measured energy performance of buildings: A framework for investigation," *Autom. Constr.*, vol. 41, pp. 40–49, May 2014, doi: 10.1016/j.autcon.2014.02.009.

[62] T. Hong, D. Yan, S. D'Oca, and C. Chen, "Ten questions concerning occupant behavior in buildings: The big picture," *Build. Environ.*, vol. 114, pp. 518–530, Mar. 2017, doi: 10.1016/j.buildenv.2016.12.006.

[63] C. J. Hopfe and J. L. M. Hensen, "Uncertainty analysis in building performance simulation for design support," *Energy Build.*, vol. 43, no. 10, pp. 2798–2805, Oct. 2011, doi: 10.1016/j.enbuild.2011.06.034.

[64] A. Chong, W. Xu, and K. Poh Lam, "Uncertainty Analysis in Building Energy Simulation. A Practical Approach," presented at the 2015 Building Simulation Conference, Dec. 2015. doi: 10.26868/25222708.2015.2131.

[65] I. Macdonald and P. Strachan, "Practical application of uncertainty analysis," *Energy Build.*, vol. 33, pp. 219–227, Feb. 2001, doi: 10.1016/S0378-7788(00)00085-2.

[66] S. D'Oca, C.-F. Chen, T. Hong, and Z. Belafi, "Synthesizing building physics with social psychology: An interdisciplinary framework for context and occupant behavior in office buildings," *Energy Res. Soc. Sci.*, vol. 34, pp. 240–251, Dec. 2017, doi: 10.1016/j.erss.2017.08.002.

[67] O. Santin and L. Itard, "Occupants' behaviour: Determinants and effects on residential heating consumption," *Build. Res. Inf.*, vol. 38, pp. 318–338, May 2010, doi: 10.1080/09613211003661074.

[68] F. Haldi and D. Robinson, "The impact of occupants' behaviour on building energy demand," *J. Build. Perform. Simul.*, vol. 4, no. 4, pp. 323–338, Dec. 2011, doi: 10.1080/19401493.2011.558213.

[69] D. Coakley, P. Raftery, and M. Keane, "A review of methods to match building energy simulation models to measured data," *Renew. Sustain. Energy Rev.*, vol. 37, pp. 123–141, Sep. 2014, doi: 10.1016/j.rser.2014.05.007.

[70] W. Tian *et al.*, "A review of uncertainty analysis in building energy assessment," *Renew. Sustain. Energy Rev.*, vol. 93, pp. 285–301, Oct. 2018, doi: 10.1016/j.rser.2018.05.029.

[71] M. Neukomm, V. Nubbe, and R. Fares, "Grid-Interactive Efficient Buildings Technical Report Series: Overview of Research Challenges and Gaps," NREL/TP--5500-75470, DOE/GO--102019-5227, 1577966, Dec. 2019. doi: 10.2172/1577966.

[72] J. Dean, P. Voss, D. Gagne, D. Vasquez, and R. Langner, “Blueprint for Integrating Grid-Interactive Efficient Building (GEB) Technologies into U.S. General Services Administration Performance Contracts,” NREL/TP-7A40-78190, 1784273, MainId:32099, May 2021. doi: 10.2172/1784273.

[73] Y. Sun, S. Wang, F. Xiao, and D. Gao, “Peak load shifting control using different cold thermal energy storage facilities in commercial buildings: A review,” *Energy Convers. Manag.*, vol. 71, pp. 101–114, Jul. 2013, doi: 10.1016/j.enconman.2013.03.026.

[74] P. Palensky and D. Dietrich, “Demand Side Management: Demand Response, Intelligent Energy Systems, and Smart Loads,” *IEEE Trans. Ind. Inform.*, vol. 7, no. 3, pp. 381–388, 08 2011, doi: 10.1109/TII.2011.2158841.

[75] M. Shakeri *et al.*, “An intelligent system architecture in home energy management systems (HEMS) for efficient demand response in smart grid,” *Energy Build.*, vol. 138, pp. 154–164, Mar. 2017, doi: 10.1016/j.enbuild.2016.12.026.

[76] B. Zhou *et al.*, “Smart home energy management systems: Concept, configurations, and scheduling strategies,” *Renew. Sustain. Energy Rev.*, vol. 61, pp. 30–40, Aug. 2016, doi: 10.1016/j.rser.2016.03.047.

[77] A. Satchwell *et al.*, “A National Roadmap for Grid-Interactive Efficient Buildings,” None, 1784302, ark:/13030/qt78k303s5, May 2021. doi: 10.2172/1784302.

[78] M. H. Albadi and E. F. El-Saadany, “A summary of demand response in electricity markets,” *Electr. Power Syst. Res.*, vol. 78, no. 11, pp. 1989–1996, Nov. 2008, doi: 10.1016/j.epsr.2008.04.002.

[79] P. Siano, “Demand response and smart grids—A survey,” *Renew. Sustain. Energy Rev.*, vol. 30, pp. 461–478, Feb. 2014, doi: 10.1016/j.rser.2013.10.022.

[80] S. Turner, A. Allen, O. Van Geet, and R. Romero, “Exploring Grid-Interactive Efficient Building Strategies for Laboratories Through Energy Modeling,” NREL/TP--7A40-87141, 2349048, MainId:87916, May 2024. doi: 10.2172/2349048.

[81] “About the Commercial Buildings Integration Program,” Energy.gov. Accessed: Nov. 11, 2024. [Online]. Available: https://www.energy.gov/eere/buildings/about-commercial-buildings-integration-program

[82] R. Langner, J. Granderson, and E. Crowe, “Quantifying the Value of Grid-Interactive Efficient Buildings through Field Study,” *Summer Study Energy Effic. Build.*, Aug. 2022, Accessed: Apr. 03, 2025. [Online]. Available: https://escholarship.org/uc/item/8v8096kq

[83] V. Nubbe, A. Weintraub, and M. Butrico, “Guide for Grid-Interactive Efficient Buildings for Federal Agencies,” DOE-FEMP--0022, 2477341, Aug. 2024. doi: 10.2172/2477341.

[84] E. Wilson *et al.*, “End-Use Load Profiles for the U.S. Building Stock: Methodology and Results of Model Calibration, Validation, and Uncertainty Quantification,” NREL/TP-5500-80889, 1854582, MainId:78667, Mar. 2022. doi: 10.2172/1854582.

[85] J. London, “Pilot Considerations for Grid-Interactive Efficient Buildings in Washington,” PNNL, PNNL-32353, 2021. [Online]. Available: https://connectedcommunities.lbl.gov/sites/default/files/2022-02/Pilot%20Considerations%20for%20GEBs%20in%20WA%20GMLC%20Final%20_0.pdf

[86] Council on Environmental Quality, “The Federal Building Performance Standard,” Executive Office of the President, 2022. [Online]. Available: https://www.sustainability.gov/pdfs/federal-building-performance-standard.pdf

[87] P. Alstone *et al.*, "2025 California Demand Response Potential Study - Charting California's Demand Response Future. Final Report on Phase 2 Results," 1421800, Mar. 2017. doi: 10.2172/1421800.

[88] N. Businelli, "2024 Assessment of Demand Response and Advanced Metering," FEDERAL ENERGY REGULATORY COMMISSION, 2024. [Online]. Available: https://www.ferc.gov/sites/default/files/2024-11/Annual%20Assessment%20of%20Demand%20Response_1119_1400.pdf

[89] ERCOT, "2023 State of the Market Report," May 2024. Accessed: Nov. 11, 2024. [Online]. Available: https://www.potomaceconomics.com/wp-content/uploads/2024/05/2023-State-of-the-Market-Report_Final.pdf

[90] T. S. G. I. P.-C. S. W. Group, "Guidelines for Smart Grid Cyber Security," National Institute of Standards and Technology, NIST Internal or Interagency Report (NISTIR) 7628 (Withdrawn), Aug. 2010. doi: 10.6028/NIST.IR.7628.

[91] M. I. Rosenberg, P. R. Hart, J. Zhang, and R. A. Athalye, "Roadmap for the Future of Commercial Energy Codes," PNNL--24009, 1169376, Jan. 2015. doi: 10.2172/1169376.

[92] S. Huang, R. Lutes, C. A. Faulkner, D. L. Vrabie, S. Katipamula, and W. Zuo, "An open-source framework for simulation-based testing of buildings control strategies," *J. Build. Perform. Simul.*, vol. 16, no. 6, pp. 631–643, Nov. 2023, doi: 10.1080/19401493.2023.2191220.

[93] "2018 Utility Demand Response Market Snapshot," SEPA. Accessed: Nov. 11, 2024. [Online]. Available: https://sepapower.org/resource/2018-demand-response-market-snapshot/

[94] National Renewable Energy Laboratory, "Voices of Experience: Leveraging AMI Networks and Data," BK-7A40-73796, Mar. 2019. [Online]. Available: https://www.energy.gov/sites/default/files/2024-02/01-03-2019_doe-voe-leveraging-ami-networks-and-data-report_508_0.pdf

[95] D. B. Patton, P. LeeVanSchaick, and J. Chen, "2010 Assessment of the Electricity Markets in New England," Potomac Economics, Jun. 2011. [Online]. Available: https://www.iso-ne.com/static-assets/documents/markets/mktmonmit/rpts/ind_mkt_advsr/isone_2010_immu_rpt_drft_final_june_11.pdf

[96] MISO, "Load Modifying Resource Reforms Draft White Paper," Nov. 2024. Accessed: Nov. 11, 2024. [Online]. Available: https://cdn.misoenergy.org/LMR%20White%20Paper652580.pdf

[97] National Renewable Energy Laboratory (NREL), A. Roth, and J. Reyna, "Grid-Interactive Efficient Buildings Technical Report Series: Whole-Building Controls, Sensors, Modeling, and Analytics," NREL/TP--5500-75478, DOE/GO--102019-5230, 1580329, Dec. 2019. doi: 10.2172/1580329.

[98] J. Reyna *et al.*, "U.S. Building Stock Characterization Study: A National Typology for Decarbonizing U.S. Buildings," *Renew. Energy*, 2022.

[99] A. Satchwell and R. Hledik, "Making Grid-interactive Efficient Buildings a 'Win' for Both Customers and Utilities," in *Energy Efficiency in Buildings*, Aug. 2022. [Online]. Available: https://escholarship.org/content/qt0679b4n2/qt0679b4n2_noSplash_e191842c119872c195b7bf70e665e727.pdf?t=rh2dhz

[100] M. Jungclaus, C. Carmichael, Phil Keuhn, and Kinga Porst Hydras, "Value Potential For Grid-Interactive Efficient Buildings In the GSA Portfolio," Rocky Mountain

Institute, Aug. 2019. [Online]. Available: https://connectedcommunities.lbl.gov/sites/default/files/2021-09/RMI%20value-potential-GEBs-gsa-portfolio.pdf

[101] K. Kalsi *et al.*, “Integrated Transmission and Distribution Control,” PNNL--22157, 1117088, Jan. 2013. doi: 10.2172/1117088.

[102] E. Zhou, E. Hale, and E. Present, “Building flexibility revenue in modeled future bulk power systems with varying levels of renewable energy,” *Heliyon*, vol. 8, no. 7, p. e09865, Jul. 2022, doi: 10.1016/j.heliyon.2022.e09865.

[103] D. F. Birol, “Energy Efficiency 2019,” International Eneregy Agency, Nov. 2019. [Online]. Available: https://iea.blob.core.windows.net/assets/8441ab46-9d86-47eb-b1fc-cb36fc3e7143/Energy_Efficiency_2019.pdf

[104] T. Kuruganti *et al.*, “Dynamic Building Load Control to Facilitate High Penetration of Solar Photovoltaic Generation (Final Technical Report),” Oak Ridge National Laboratory (ORNL), Oak Ridge, TN (United States), ORNL/TM-2021/2112, Sep. 2021. doi: 10.2172/1819555.

[105] Jason Taylor, Alison O’Connell, and Jeff Smith, “Incorporating DER Into Distribution Planning,” Electric Power Reasearch Institute, Inc., 000000003002010997, Jun. 2018. Accessed: Apr. 05, 2025. [Online]. Available: https://www.epri.com/research/products/000000003002010997

[106] K. Zitelman, “Advancing Electric System Resilience with Distributed Energy Resources: A Review of State Policies,” NREL/SR--7A40-90137, 2394650, MainId:91915, Apr. 2020. doi: 10.2172/2394650.

[107] FEMA, “FEMA Report – Reducing ‘Loss of Utility’ Impacts to Critical Facilities,” FEMA, DR-4673-FL RA 2, Aug. 2023. Accessed: Apr. 05, 2025. [Online]. Available: https://www.waterisac.org/portal/fema-report-%E2%80%93-reducing-%E2%80%9Closs-utility%E2%80%9D-impacts-critical-facilities

[108] W. I. Bower *et al.*, “The advanced microgrid. Integration and interoperability,” Sandia National Lab. (SNL-NM), Albuquerque, NM (United States), SAND2014-1535, Feb. 2014. doi: 10.2172/1204100.

[109] G. W. Arnold *et al.*, “NIST Framework and Roadmap for Smart Grid Interoperability Standards, Release 1.0,” *NIST*, Jan. 2010, Accessed: Apr. 05, 2025. [Online]. Available: https://www.nist.gov/publications/nist-framework-and-roadmap-smart-grid-interoperability-standards-release-10

[110] Metin Celebi *et al.*, “Bulk System Reliability for Tomorrow’s Grid,” The Brattle Group, Dec. 2023. [Online]. Available: https://www.brattle.com/wp-content/uploads/2023/12/Bulk-System-Reliability-for-Tomorrows-Grid_December-2023_Final.pdf

[111] NERC, “2023 Long-Term Reliability Assessment,” Dec. 2023. Accessed: Nov. 11, 2024. [Online]. Available: https://www.nerc.com/pa/RAPA/ra/Reliability%20Assessments%20DL/NERC_LTRA_2023.pdf

[112] C. Perry, H. Bastian, and D. York, “Grid-Interactive Efficient Building Utility Programs: State of the Market,” American Council for an Energy-Efficient Economy, ACEEE, Oct. 2019. [Online]. Available: https://www.aceee.org/white-paper/gebs-103019

[113] Grid Deployment Office, “National Transmission Planning Study,” U.S. Department of Energy., 2024. [Online]. Available: https://www.energy.gov/sites/default/files/2024-10/NationalTransmissionPlanningStudy-Chapter1.pdf

[114] S. woo Ham *et al.*, “Practical challenges of model predictive control (MPC) for grid interactive small and medium commercial buildings,” Aug. 2024, Accessed: Nov. 11, 2024. [Online]. Available: https://escholarship.org/uc/item/9jv9p8rj

[115] USDOE Office of Energy Efficiency and Renewable Energy (EERE), “Building Technologies Office: Peer Review Report 2019,” Energy Efficiency Office. Building Technologies Office, DOE/EE-2002, Dec. 2019. doi: 10.2172/1580212.

[116] A. Q. Al-Shetwi, M. A. Hannan, K. P. Jern, M. Mansur, and T. M. I. Mahlia, “Grid-connected renewable energy sources: Review of the recent integration requirements and control methods,” *J. Clean. Prod.*, vol. 253, p. 119831, Apr. 2020, doi: 10.1016/j.jclepro.2019.119831.

[117] G. Pinto, A. Kathirgamanathan, E. Mangina, D. P. Finn, and A. Capozzoli, “Enhancing energy management in grid-interactive buildings: A comparison among cooperative and coordinated architectures,” *Appl. Energy*, vol. 310, p. 118497, Mar. 2022, doi: 10.1016/j.apenergy.2021.118497.

[118] G. Li *et al.*, “A critical review of cyber-physical security for building automation systems,” *Annu. Rev. Control*, vol. 55, pp. 237–254, Jan. 2023, doi: 10.1016/j.arcontrol.2023.02.004.

[119] Idaho National Laboratory, “Vulnerability Analysis of Energy Delivery Control Systems,” U.S. Department of Energy National Laboratory, INL/EXT-10-18381, Sep. 2011. [Online]. Available: https://www.energy.gov/sites/prod/files/Vulnerability%20Analysis%20of%20Energy%20Delivery%20Control%20Systems%202011.pdf

[120] PJM LLC, “Load Management Performance Report 2020.pdf,” mid Delivery Year update. Accessed: Aug. 10, 2020. [Online]. Available: https://www.pjm.com/-/media/committees-groups/subcommittees/drs/2020/20200309/20200309-item-03-2019-20-load-management-report-mid-year.ashx

[121] J. Leprince, A. Schledorn, D. Guericke, D. F. Dominkovic, H. Madsen, and W. Zeiler, “Can occupant behaviors affect urban energy planning? Distributed stochastic optimization for energy communities,” *Appl. Energy*, vol. 348, p. 121589, Oct. 2023, doi: 10.1016/j.apenergy.2023.121589.

[122] T. Xiao and F. You, “Physically consistent deep learning-based day-ahead energy dispatching and thermal comfort control for grid-interactive communities,” *Appl. Energy*, vol. 353, p. 122133, Jan. 2024, doi: 10.1016/j.apenergy.2023.122133.

[123] David Kathan, Nicole Businelli, and Karin Herzfeld, *Participation of Distributed Energy Resource Aggregations in Markets Operated by Regional Transmission Organizations and Independent System Operators*, vol. RM18-9–002; 2021. Accessed: Apr. 05, 2025. [Online]. Available: https://www.federalregister.gov/documents/2021/03/30/2021-06089/participation-of-distributed-energy-resource-aggregations-in-markets-operated-by-regional

[124] ISO New England Inc., “Measurement and Verification of On-Peak Demand Resources and Seasonal Peak Demand Resources,” Oct. 2018. Accessed: Nov. 11, 2024. [Online]. Available: https://www.iso-ne.com/static-assets/documents/2018/10/manual_mvdr_measurement_and_verification_of_onpeak_and_seasonal_peak_demand_resources_rev07_20181004.pdf

[125] S. Truitt, J. Bean, J. Sullivan, G. Paranjothi, and A. Moe, “Completing the Circuit: Workforce Development for Advanced Building Construction and Grid-Interactive Efficient Buildings,” NREL/TP-5500-80480, 1855580, MainId:42683, Mar. 2022. doi: 10.2172/1855580.

[126] J. K. Day, S. N. Ruiz, M. Kane, and R. Sheppard, "Training the Next Generation of Building Operators: Development of the Building Operators: Grid and Occupants (BOGO) Training," in *Summer Study*, ACEEE, 2024.

[127] M. Steen and M. Krarti, "A Review and Categorization of Grid-Interactive Efficient Building Technologies for Building Performance Simulation," *ASME J. Eng. Sustain. Build. Cities*, vol. 1, no. 4, p. 040801, Nov. 2020, doi: 10.1115/1.4048975.

[128] Y. Ye *et al.*, "System modeling for grid-interactive efficient building applications," *J. Build. Eng.*, vol. 69, p. 106148, Jun. 2023, doi: 10.1016/j.jobe.2023.106148.

[129] X. Ma, D. Wu, and R. Rebich, "Optimizing Grid-interactive Efficient Building Designs with Stacked Value Streams," in *2024 IEEE Power & Energy Society General Meeting (PESGM)*, Jul. 2024, pp. 1–5. doi: 10.1109/PESGM51994.2024.10688991.

[130] W. O'Brien *et al.*, "Introducing IEA EBC annex 79: Key challenges and opportunities in the field of occupant-centric building design and operation," *Build. Environ.*, vol. 178, p. 106738, Jul. 2020, doi: 10.1016/j.buildenv.2020.106738.

[131] N. Luo *et al.*, "A three-year dataset supporting research on building energy management and occupancy analytics," *Sci. Data*, vol. 9, no. 1, p. 156, Apr. 2022, doi: 10.1038/s41597-022-01257-x.

[132] C. Chand, A. Villanueva, M. Marty, and M. Aibin, "Privacy Preserving Occupancy Detection Using NB IoT Sensors," in *2021 IEEE Canadian Conference on Electrical and Computer Engineering (CCECE)*, Sep. 2021, pp. 1–4. doi: 10.1109/CCECE53047.2021.9569139.

[133] P. O. Fanger, "Assessment of man's thermal comfort in practice," *Br. J. Ind. Med.*, vol. 30, no. 4, pp. 313–324, Oct. 1973.

[134] J. Van Hoof, "Forty years of Fanger's model of thermal comfort: comfort for all?Abstract," *Indoor Air*, vol. 18, no. 3, pp. 182–201, 2008, doi: 10.1111/j.1600-0668.2007.00516.x.

[135] T. Parkinson, R. de Dear, and G. Brager, "Nudging the adaptive thermal comfort model," *Energy Build.*, vol. 206, p. 109559, Jan. 2020, doi: 10.1016/j.enbuild.2019.109559.

[136] R. J. de Dear and G. S. Brager, "Developing an Adaptive Model of Thermal Comfort and Preference," *ASHRAE Trans.*, vol. 104, 1998, Accessed: Sep. 28, 2022. [Online]. Available: https://escholarship.org/uc/item/4qq2p9c6

[137] V. Földváry Ličina *et al.*, "Development of the ASHRAE Global Thermal Comfort Database II," *Build. Environ.*, vol. 142, pp. 502–512, Sep. 2018, doi: 10.1016/j.buildenv.2018.06.022.

[138] J. Kim, Y. Zhou, S. Schiavon, P. Raftery, and G. Brager, "Personal comfort models: Predicting individuals' thermal preference using occupant heating and cooling behavior and machine learning," *Build. Environ.*, vol. 129, pp. 96–106, Feb. 2018, doi: 10.1016/j.buildenv.2017.12.011.

[139] Q. Zhao, Z. Lian, and D. Lai, "Thermal comfort models and their developments: A review," *Energy Built Environ.*, vol. 2, no. 1, pp. 21–33, Jan. 2021, doi: 10.1016/j.enbenv.2020.05.007.

[140] J. Kim, S. Schiavon, and G. Brager, "Personal comfort models – A new paradigm in thermal comfort for occupant-centric environmental control," *Build. Environ.*, vol. 132, pp. 114–124, Mar. 2018, doi: 10.1016/j.buildenv.2018.01.023.

[141] A. Melikov, "Energy saving and improved comfort by increased air movement," *Energy Build.*, vol. 40, pp. 1954–1960, Dec. 2008, doi: 10.1016/j.enbuild.2008.05.001.

[142] M. Fattahi and M. Sharbatdar, "Machine-learning-based personal thermal comfort modeling for heat recovery using environmental parameters," *Sustain. Energy Technol. Assess.*, vol. 57, p. 103294, Jun. 2023, doi: 10.1016/j.seta.2023.103294.

[143] N. Gao, W. Shao, M. S. Rahaman, J. Zhai, K. David, and F. D. Salim, "Transfer learning for thermal comfort prediction in multiple cities," *Build. Environ.*, vol. 195, p. 107725, May 2021, doi: 10.1016/j.buildenv.2021.107725.

[144] A. Ghahramani, C. Tang, and B. Becerik-Gerber, "An online learning approach for quantifying personalized thermal comfort via adaptive stochastic modeling," *Build. Environ.*, vol. 92, pp. 86–96, Oct. 2015, doi: 10.1016/j.buildenv.2015.04.017.

[145] W. Jung and F. Jazizadeh, "Human-in-the-loop HVAC operations: A quantitative review on occupancy, comfort, and energy-efficiency dimensions," *Appl. Energy*, vol. 239, pp. 1471–1508, Apr. 2019, doi: 10.1016/j.apenergy.2019.01.070.

[146] M. Schweiker *et al.*, "Review of multi-domain approaches to indoor environmental perception and behaviour," *Build. Environ.*, vol. 176, p. 106804, Jun. 2020, doi: 10.1016/j.buildenv.2020.106804.

[147] V. L. Erickson, M. Á. Carreira-Perpiñán, and A. E. Cerpa, "Occupancy Modeling and Prediction for Building Energy Management," *ACM Trans Sen Netw*, vol. 10, no. 3, p. 42:1-42:28, May 2014, doi: 10.1145/2594771.

[148] Md. R. Rafi, F. Hu, S. Li, A. Song, X. Zhang, and Z. O'Neill, "Deep Weighted Fusion Learning (DWFL)-based multi-sensor fusion model for accurate building occupancy detection," *Energy AI*, vol. 17, p. 100379, Sep. 2024, doi: 10.1016/j.egyai.2024.100379.

[149] S. Y. Tan, H. Saha, A. R. Florita, G. P. Henze, and S. Sarkar, "A Flexible Framework for Building Occupancy Detection Using Spatiotemporal Pattern Networks," Piscataway, NJ: Institute of Electrical and Electronics Engineers (IEEE), NREL/CP-5D00-75055, Aug. 2019. Accessed: Apr. 05, 2025. [Online]. Available: https://www.osti.gov/biblio/1569444

[150] W. O'Brien, I. Gaetani, S. Gilani, S. Carlucci, P.-J. Hoes, and J. Hensen, "International survey on current occupant modelling approaches in building performance simulation†," *J. Build. Perform. Simul.*, vol. 10, no. 5–6, pp. 653–671, Nov. 2017, doi: 10.1080/19401493.2016.1243731.

[151] T. Hong, S. D'Oca, W. J. N. Turner, and S. C. Taylor-Lange, "An ontology to represent energy-related occupant behavior in buildings. Part I: Introduction to the DNAs framework," *Build. Environ.*, vol. 92, pp. 764–777, Oct. 2015, doi: 10.1016/j.buildenv.2015.02.019.

[152] T. Hong, S. D'Oca, S. C. Taylor-Lange, W. J. N. Turner, Y. Chen, and S. P. Corgnati, "An ontology to represent energy-related occupant behavior in buildings. Part II: Implementation of the DNAS framework using an XML schema," *Build. Environ.*, vol. 94, pp. 196–205, Dec. 2015, doi: 10.1016/j.buildenv.2015.08.006.

[153] A. Mahdavi and M. Taheri, "An ontology for building monitoring," *J. Build. Perform. Simul.*, vol. 10, no. 5–6, pp. 499–508, Nov. 2017, doi: 10.1080/19401493.2016.1243730.

[154] T. Hong, S. D'Oca, W. J. N. Turner, and S. C. Taylor-Lange, "An ontology to represent energy-related occupant behavior in buildings. Part I: Introduction to the DNAs framework," *Build. Environ.*, vol. 92, pp. 764–777, Oct. 2015, doi: 10.1016/j.buildenv.2015.02.019.

[155] K. Gram-Hanssen, "Residential heat comfort practices: understanding users," *Build. Res. Inf.*, vol. 38, no. 2, pp. 175–186, Apr. 2010, doi: 10.1080/09613210903541527.

[156] Z. Wang and T. Hong, "Reinforcement learning for building controls: The opportunities and challenges," *Appl. Energy*, vol. 269, p. 115036, Jul. 2020, doi: 10.1016/j.apenergy.2020.115036.

[157] B. Huchuk, W. O'Brien, and S. Sanner, "A longitudinal study of thermostat behaviors based on climate, seasonal, and energy price considerations using connected thermostat data," *Build. Environ.*, vol. 139, pp. 199–210, Jul. 2018, doi: 10.1016/j.buildenv.2018.05.003.

[158] L. Sarran, H. B. Gunay, W. O'Brien, C. A. Hviid, and C. Rode, "A data-driven study of thermostat overrides during demand response events," *Energy Policy*, vol. 153, p. 112290, Jun. 2021, doi: 10.1016/j.enpol.2021.112290.

[159] K. Seiden, C. Olig, D. Max, and M. Sherman, "National Grid Smart Energy Solutions Pilot Final Evaluation Report," Navigant, Massachusetts Electric Company, and Nantucket Electric Company d/b/a National Grid, Worcester, MA, May 2017. [Online]. Available: http://web1.env.state.ma.us/DPU/FileRoom

[160] L. White and N. Sintov, "Health and financial impacts of demand-side response measures differ across sociodemographic groups," *Nat. Energy*, vol. 5, pp. 1–11, Jan. 2020, doi: 10.1038/s41560-019-0507-y.

[161] W. Kleiminger, F. Mattern, and S. Santini, "Predicting household occupancy for smart heating control: A comparative performance analysis of state-of-the-art approaches," *Energy Build.*, vol. 85, pp. 493–505, Dec. 2014, doi: 10.1016/j.enbuild.2014.09.046.

[162] Z. Li and B. Dong, "Short term predictions of occupancy in commercial buildings—Performance analysis for stochastic models and machine learning approaches," *Energy Build.*, vol. 158, pp. 268–281, Jan. 2018, doi: 10.1016/j.enbuild.2017.09.052.

[163] B. Gunay and I. Beausoleil-Morrison, "Implementation and comparison of existing occupant behaviour models in EnergyPlus," *J. Build. Perform. Simul.*, vol. 1493, pp. 1–46, Nov. 2015, doi: 10.1080/19401493.2015.1102969.

[164] K. Herter and Y. Okuneva, "SMUD's Load Impact Calculator (SLIC) for Smart Grid Cross-Pilot Evaluation," Herter Energy Research Solutions, Inc., Dec. 2014. [Online]. Available: https://www.smud.org/-/media/Documents/Corporate/About-Us/Energy-Research-and-Development/research-customer-cross-pilot-evaluation.ashx

[165] C. Miller and F. Meggers, "The Building Data Genome Project: An open, public data set from non-residential building electrical meters," *Energy Procedia*, vol. 122, pp. 439–444, Sep. 2017, doi: 10.1016/j.egypro.2017.07.400.

[166] C. F. Reinhart, "Lightswitch-2002: a model for manual and automated control of electric lighting and blinds," *Sol. Energy*, vol. 77, no. 1, pp. 15–28, Jan. 2004, doi: 10.1016/j.solener.2004.04.003.

[167] Z. Nagy, F. Y. Yong, M. Frei, and A. Schlueter, "Occupant centered lighting control for comfort and energy efficient building operation," *Energy Build.*, vol. 94, pp. 100–108, May 2015, doi: 10.1016/j.enbuild.2015.02.053.

[168] J. R. Vazquez-Canteli, S. Dey, G. Henze, and Z. Nagy, "CityLearn: Standardizing Research in Multi-Agent Reinforcement Learning for Demand Response and Urban Energy Management," Dec. 18, 2020, *arXiv*: arXiv:2012.10504. doi: 10.48550/arXiv.2012.10504.

[169] D. M. Sykes, "Productivity : How Acoustics Affect Workers ' Performance In Offices & Open Areas," presented at the Environmental Science, 2004. Accessed: Apr. 05, 2025. [Online]. Available: https://www.semanticscholar.org/paper/Productivity-%3A-

How-Acoustics-Affect-Workers-%E2%80%99-In-%26-Sykes/cfcb6e93a42ed0b5ac03c5d98182a3136a80664c
[170] F. Banihashemi, M. Weber, and W. Lang, "Deep learning for predictive window operation modeling in open-plan offices," *Energy Build.*, vol. 310, p. 114109, May 2024, doi: 10.1016/j.enbuild.2024.114109.
[171] T. Wang, M. Ma, N. Zhou, and Z. Ma, "Toward net zero: Assessing the decarbonization impact of global commercial building electrification," *Appl. Energy*, vol. 383, p. 125287, Apr. 2025, doi: 10.1016/j.apenergy.2025.125287.
[172] V. Fabi, R. V. Andersen, S. Corgnati, and B. W. Olesen, "Occupants' window opening behaviour: A literature review of factors influencing occupant behaviour and models," *Build. Environ.*, vol. 58, pp. 188–198, 12-01 2012, doi: 10.1016/j.buildenv.2012.07.009.
[173] H. B. Rijal, P. Tuohy, F. Nicol, M. A. Humphreys, A. Samuel, and J. Clarke, "Development of an adaptive window-opening algorithm to predict the thermal comfort, energy use and overheating in buildings," *J. Build. Perform. Simul.*, vol. 1, no. 1, pp. 17–30, Mar. 2008, doi: 10.1080/19401490701868448.
[174] R. Andersen, V. Fabi, J. Toftum, S. P. Corgnati, and B. W. Olesen, "Window opening behaviour modelled from measurements in Danish dwellings," *Build. Environ.*, vol. 69, pp. 101–113, Nov. 2013, doi: 10.1016/j.buildenv.2013.07.005.
[175] S. D'Oca and T. Hong, "A data-mining approach to discover patterns of window opening and closing behavior in offices," *Build. Environ.*, vol. 82, pp. 726–739, Dec. 2014, doi: 10.1016/j.buildenv.2014.10.021.
[176] Y. Zhang and P. Barrett, "Factors influencing occupants' blind-control behaviour in a naturally ventilated office building," *Build. Environ.*, vol. 54, pp. 137–147, Aug. 2012, doi: 10.1016/j.buildenv.2012.02.016.
[177] M. Schweiker, F. Haldi, M. Shukuya, and D. Robinson, "Verification of stochastic models of window opening behaviour for residential buildings," *J. Build. Perform. Simul.*, vol. 5, no. 1, pp. 55–74, Jan. 2012, doi: 10.1080/19401493.2011.567422.
[178] D. Yan *et al.*, "IEA EBC Annex 66: Definition and simulation of occupant behavior in buildings," *Energy Build.*, vol. 156, pp. 258–270, Dec. 2017, doi: 10.1016/j.enbuild.2017.09.084.
[179] J. Page, D. Robinson, N. Morel, and J.-L. Scartezzini, "A generalised stochastic model for the simulation of occupant presence," *Energy Build.*, vol. 40, no. 2, pp. 83–98, Jan. 2008, doi: 10.1016/j.enbuild.2007.01.018.
[180] T. Parkinson and R. de Dear, "Thermal pleasure in built environments: spatial alliesthesia from contact heating," *Build. Res. Inf.*, vol. 44, no. 3, pp. 248–262, Apr. 2016, doi: 10.1080/09613218.2015.1082334.
[181] J. Y. Park *et al.*, "A critical review of field implementations of occupant-centric building controls," *Build. Environ.*, vol. 165, p. 106351, Nov. 2019, doi: 10.1016/j.buildenv.2019.106351.
[182] W. O'Brien, I. Gaetani, S. Carlucci, P.-J. Hoes, and J. L. M. Hensen, "On occupant-centric building performance metrics," *Build. Environ.*, vol. 122, pp. 373–385, Sep. 2017, doi: 10.1016/j.buildenv.2017.06.028.
[183] A. Kathirgamanathan, M. De Rosa, E. Mangina, and D. P. Finn, "Data-driven predictive control for unlocking building energy flexibility: A review," *Renew. Sustain. Energy Rev.*, vol. 135, p. 110120, Jan. 2021, doi: 10.1016/j.rser.2020.110120.
[184] J. Drgoňa *et al.*, "All you need to know about model predictive control for buildings," *Annu. Rev. Control*, Sep. 2020, doi: 10.1016/j.arcontrol.2020.09.001.

[185] A. Mylonas, A. Tsangrassoulis, and J. Pascual, “Modelling occupant behaviour in residential buildings: A systematic literature review,” *Build. Environ.*, vol. 265, p. 111959, Nov. 2024, doi: 10.1016/j.buildenv.2024.111959.

[186] T. Parkinson, R. de Dear, and C. Candido, “Thermal pleasure in built environments: alliesthesia in different thermoregulatory zones,” *Build. Res. Inf.*, vol. 44, no. 1, pp. 20–33, Jan. 2016, doi: 10.1080/09613218.2015.1059653.

[187] M. Schweiker and A. Wagner, “A framework for an adaptive thermal heat balance model (ATHB),” *Build. Environ.*, vol. 94, pp. 252–262, Dec. 2015, doi: 10.1016/j.buildenv.2015.08.018.

[188] M. Jia, R. S. Srinivasan, R. Ries, and G. Bharathy, “A framework of occupant behavior modeling and data sensing for improving building energy simulation,” in *Proceedings of the Symposium on Simulation for Architecture and Urban Design*, in SIMAUD ’18. San Diego, CA, USA: Society for Computer Simulation International, Jun. 2018, pp. 1–8.

[189] Ryan Repice and Alexander Sun, “Monthly Energy Review,” U.S. Energy Information Administration, DOE/EIA-0035, Mar. 2019. [Online]. Available: https://safety4sea.com/wp-content/uploads/2019/04/EIA-March-2019-Monthly-Energy-Review-2019_04.pdf

[190] M. Neukomm, V. Nubbe, and R. Fares, “Grid-interactive Efficient Buildings: Overview,” U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy, Washington, D.C., Apr. 2019. [Online]. Available: https://www.energy.gov/sites/prod/files/2019/04/f61/bto-geb_overview-4.15.19.pdf

[191] S. Singla and S. Keshav, “Demand response through a temperature setpoint market in Ontario,” in *2012 IEEE Third International Conference on Smart Grid Communications (SmartGridComm)*, Tainan, Taiwan: IEEE, Nov. 2012, pp. 103–108. doi: 10.1109/SmartGridComm.2012.6485967.

[192] L. Marie-Andrée, D. A. Ahmed, and L. B. Célyn, “DEVELOPING WINTER RESIDENTIAL DEMAND RESPONSE STRATEGIES FOR ELECTRIC SPACE HEATING,” in *Proceedings of Building Simulation 2011*, Sydney, AU: International Building Performance Simulation Association, Nov. 2011, p. 8. [Online]. Available: http://www.ibpsa.org/proceedings/BS2011/P_1399.pdf

[193] Charles Goldman, “Coordination of Energy Efficiency and Demand Response,” Lawrence Berkeley National Laboratory, Jan. 2010. [Online]. Available: https://19january2021snapshot.epa.gov/sites/static/files/2015-08/documents/ee_and_dr.pdf

[194] P. Cappers, J. MacDonald, J. Page, J. Potter, and E. Stewart, “Future Opportunities and Challenges with Using Demand Response as a Resource in Distribution System Operation and Planning Activities,” LBNL--1003951, 1333622, Jan. 2016. doi: 10.2172/1333622.

[195] A. Mirakhorli and B. Dong, “Occupancy behavior based model predictive control for building indoor climate—A critical review,” *Energy Build.*, vol. 129, pp. 499–513, Oct. 2016, doi: 10.1016/j.enbuild.2016.07.036.

[196] X. Chen, Q. Wang, and J. Srebric, “A data-driven state-space model of indoor thermal sensation using occupant feedback for low-energy buildings,” *Energy Build.*, vol. 91, pp. 187–198, Mar. 2015, doi: 10.1016/j.enbuild.2015.01.038.

[197] H. Zhang, E. Arens, C. Huizenga, and T. Han, “Thermal sensation and comfort models for non-uniform and transient environments: Part I: Local sensation of individual

body parts," *Build. Environ.*, vol. 45, no. 2, pp. 380–388, Feb. 2010, doi: 10.1016/j.buildenv.2009.06.018.

[198] Katherine Tweed, "Customers Spend 8 Minutes per Year Interacting Online With Their Utility," GreenTechMedia. Accessed: Apr. 06, 2025. [Online]. Available: https://www.greentechmedia.com/articles/read/customers-spend-8-minutes-a-year-interacting-online-with-their-utility

[199] C. Wilson and H. Dowlatabadi, "Models of Decision Making and Residential Energy Use," *Annu. Rev. Environ. Resour.*, vol. 32, no. 1, pp. 169–203, Nov. 2007, doi: 10.1146/annurev.energy.32.053006.141137.

[200] W. Kempton, "Two Theories of Home Heat Control*," *Cogn. Sci.*, vol. 10, no. 1, pp. 75–90, Jan. 1986, doi: 10.1207/s15516709cog1001_3.

[201] M. B. Kane, "Modeling Human-in-the-Loop Behavior and Interactions with HVAC Systems," in *2018 Annual American Control Conference (ACC)*, Jun. 2018, pp. 4628–4633. doi: 10.23919/ACC.2018.8431913.

[202] B. Huchuk, W. O'Brien, and S. Sanner, "A longitudinal study of thermostat behaviors based on climate, seasonal, and energy price considerations using connected thermostat data," *Build. Environ.*, vol. 139, pp. 199–210, Jul. 2018, doi: 10.1016/j.buildenv.2018.05.003.

[203] N. Luo and T. Hong, "Ecobee Donate Your Data 1,000 homes in 2017," Pacific Northwest National Lab. (PNNL), Richland, WA (United States); Lawrence Berkeley National Lab. (LBNL), Berkeley, CA (United States), 69035, Mar. 2022. doi: 10.25584/ecobee/1854924.

[204] "Pecan Street Grid Demonstration Program. Final technology performance report," Pecan Street Project, Inc., Austin, TX (United States), Feb. 2015. doi: 10.2172/1172297.

[205] Z. Wang *et al.*, "Impacts of Operators' Behavior on Reliability of Power Grids During Cascading Failures," *IEEE Trans. Power Syst.*, vol. 33, no. 6, pp. 6013–6024, Nov. 2018, doi: 10.1109/TPWRS.2018.2825348.

[206] P. Sedgwick and N. Greenwood, "Understanding the Hawthorne effect," *BMJ*, vol. 351, p. h4672, Sep. 2015, doi: 10.1136/bmj.h4672.

[207] B. Huchuk, S. Sanner, and W. O'Brien, "Comparison of machine learning models for occupancy prediction in residential buildings using connected thermostat data," *Build. Environ.*, vol. 160, p. 106177, Aug. 2019, doi: 10.1016/j.buildenv.2019.106177.

[208] M. Vellei, S. Martinez, and J. Le Dréau, "Agent-based stochastic model of thermostat adjustments: A demand response application," *Energy Build.*, vol. 238, p. 110846, May 2021, doi: 10.1016/j.enbuild.2021.110846.

[209] J. Kaur, K. S. Sahu, A. Oetomo, and P. Morita, "A smart thermostat-based population-level behavioural changes during the COVID-19 pandemic in the United States: a proposed study," in *Proceedings of the 2022 Workshop on Emerging Devices for Digital Biomarkers*, in DigiBiom '22. New York, NY, USA: Association for Computing Machinery, Jul. 2022, pp. 7–12. doi: 10.1145/3539494.3542756.

[210] M. Kane and K. Sharma, "Data-Driven Modeling of Adaptive Occupant Thermostat Behavior Dynamics," *ASHRAE Trans.*, vol. 126, no. 2, pp. 123–130, Jul. 2020.

[211] K. Sharma and M. Kane, "Impact Analysis of Personalized Thermostat Demand Response," in *ASHRAE Topical Conference Proceedings*, Atlanta, United States: American Society of Heating, Refrigeration and Air Conditioning Engineers, Inc., 2020, pp. 222–229. Accessed: Jan. 21, 2022. [Online]. Available: https://www.proquest.com/docview/2501933016/abstract/30ED4F7982F5424EPQ/1

[212] T. H. Pedersen, K. U. Nielsen, and S. Petersen, "Method for room occupancy detection based on trajectory of indoor climate sensor data," *Build. Environ.*, vol. 115, pp. 147–156, Apr. 2017, doi: 10.1016/j.buildenv.2017.01.023.

[213] J. Kolarik, J. Toftum, B. W. Olesen, and A. Shitzer, "Occupant Responses and Office Work Performance in Environments with Moderately Drifting Operative Temperatures (RP-1269)," *HVACR Res.*, vol. 15, no. 5, pp. 931–960, Sep. 2009, doi: 10.1080/10789669.2009.10390873.

[214] C. Mitterer, H. M. Künzel, S. Herkel, and A. Holm, "Optimizing energy efficiency and occupant comfort with climate specific design of the building," *Front. Archit. Res.*, vol. 1, no. 3, pp. 229–235, Sep. 2012, doi: 10.1016/j.foar.2012.06.002.

[215] G. S. Brager and R. de Dear, "Climate, comfort, & natural ventilation: a new adaptive comfort standard for ASHRAE standard 55," in *Proceedings, Moving Thermal Comfort Standards into the 21st Century*, Windsor, UK., Apr. 2001, p. 19. Accessed: Nov. 17, 2019. [Online]. Available: https://escholarship.org/uc/item/2048t8nn

[216] "How accurate is the temperature sensor in the ecobee?" Accessed: Aug. 01, 2022. [Online]. Available: https://support.ecobee.com/s/articles/How-accurate-is-the-temperature-sensor-in-the-ecobee

[217] Dr. Fatih Birol and Inger Andersen, "Global Status Report for Buildings and Construction," International Energy Agency, 978-92-807-3768–4, 2019. [Online]. Available: https://www.iea.org/reports/global-status-report-for-buildings-and-construction-2019

[218] T. Hong, J. Langevin, and K. Sun, "Building simulation: Ten challenges," *Build. Simul.*, vol. 11, no. 5, pp. 871–898, Oct. 2018, doi: 10.1007/s12273-018-0444-x.

[219] R. J. de Dear *et al.*, "Progress in thermal comfort research over the last twenty years," *Indoor Air*, vol. 23, no. 6, pp. 442–461, Dec. 2013, doi: 10.1111/ina.12046.

[220] S. Borgeson and G. Brager, "Comfort standards and variations in exceedance for mixed-mode buildings," *Build. Res. Inf.*, vol. 39, no. 2, pp. 118–133, Apr. 2011, doi: 10.1080/09613218.2011.556345.

[221] J. van Hoof, "Forty years of Fanger's model of thermal comfort: comfort for all?," *Indoor Air*, vol. 18, no. 3, pp. 182–201, Jun. 2008, doi: 10.1111/j.1600-0668.2007.00516.x.

[222] S. Karjalainen, "Thermal comfort and gender: a literature review," *Indoor Air*, vol. 22, no. 2, pp. 96–109, Apr. 2012, doi: 10.1111/j.1600-0668.2011.00747.x.

[223] R. F. Rupp, N. G. Vásquez, and R. Lamberts, "A review of human thermal comfort in the built environment," *Energy Build.*, vol. 105, pp. 178–205, Oct. 2015, doi: 10.1016/j.enbuild.2015.07.047.

[224] T. Cheung, S. Schiavon, T. Parkinson, P. Li, and G. Brager, "Analysis of the accuracy on PMV – PPD model using the ASHRAE Global Thermal Comfort Database II," *Build. Environ.*, vol. 153, pp. 205–217, Apr. 2019, doi: 10.1016/j.buildenv.2019.01.055.

[225] A. L. A. da Fonseca, K. M. S. Chvatal, and R. A. S. Fernandes, "Thermal comfort maintenance in demand response programs: A critical review," *Renew. Sustain. Energy Rev.*, vol. 141, p. 110847, May 2021, doi: 10.1016/j.rser.2021.110847.

[226] P. Siano, "Demand response and smart grids—A survey," *Renew. Sustain. Energy Rev.*, vol. 30, pp. 461–478, Feb. 2014, doi: 10.1016/j.rser.2013.10.022.

[227] H. Stopps and M. F. Touchie, "Managing thermal comfort in contemporary high-rise residential buildings: Using smart thermostats and surveys to identify energy efficiency and comfort opportunities," *Build. Environ.*, vol. 173, p. 106748, Apr. 2020, doi: 10.1016/j.buildenv.2020.106748.

[228] Y. Chen, P. Xu, J. Gu, F. Schmidt, and W. Li, “Measures to improve energy demand flexibility in buildings for demand response (DR): A review,” *Energy Build.*, vol. 177, pp. 125–139, Oct. 2018, doi: 10.1016/j.enbuild.2018.08.003.

[229] Y. Zhang, X. Bai, F. P. Mills, and J. C. V. Pezzey, “Rethinking the role of occupant behavior in building energy performance: A review,” *Energy Build.*, vol. 172, pp. 279–294, Aug. 2018, doi: 10.1016/j.enbuild.2018.05.017.

[230] J. Y. Park *et al.*, “A critical review of field implementations of occupant-centric building controls,” *Build. Environ.*, vol. 165, p. 106351, Nov. 2019, doi: 10.1016/j.buildenv.2019.106351.

[231] C. De la Cruz-Lovera, A.-J. Perea-Moreno, J.-L. De la Cruz-Fernández, J. A. Alvarez-Bermejo, and F. Manzano-Agugliaro, “Worldwide Research on Energy Efficiency and Sustainability in Public Buildings,” *Sustainability*, vol. 9, no. 8, Art. no. 8, Aug. 2017, doi: 10.3390/su9081294.

[232] K. Gram-Hanssen, T. H. Christensen, and P. E. Petersen, “Air-to-air heat pumps in real-life use: Are potential savings achieved or are they transformed into increased comfort?,” *Energy Build.*, vol. 53, pp. 64–73, Oct. 2012, doi: 10.1016/j.enbuild.2012.06.023.

[233] M. Pritoni, A. K. Meier, C. Aragon, D. Perry, and T. Peffer, “Energy efficiency and the misuse of programmable thermostats: The effectiveness of crowdsourcing for understanding household behavior,” *Energy Res. Soc. Sci.*, vol. 8, pp. 190–197, Jul. 2015, doi: 10.1016/j.erss.2015.06.002.

[234] M. Pritoni, K. Salmon, A. Sanguinetti, J. Morejohn, and M. Modera, “Occupant thermal feedback for improved efficiency in university buildings,” *Energy Build.*, vol. 144, pp. 241–250, Jun. 2017, doi: 10.1016/j.enbuild.2017.03.048.

[235] L. Peeters, R. de Dear, J. Hensen, and W. D’haeseleer, “Thermal comfort in residential buildings: Comfort values and scales for building energy simulation,” *Appl. Energy*, vol. 86, no. 5, pp. 772–780, May 2009, doi: 10.1016/j.apenergy.2008.07.011.

[236] J. F. Nicol and M. A. Humphreys, “Adaptive thermal comfort and sustainable thermal standards for buildings,” *Energy Build.*, vol. 34, no. 6, pp. 563–572, Jul. 2002, doi: 10.1016/S0378-7788(02)00006-3.

[237] M. Fountain, G. Brager, and R. de Dear, “Expectations of indoor climate control,” *Energy Build.*, vol. 24, no. 3, pp. 179–182, Oct. 1996, doi: 10.1016/S0378-7788(96)00988-7.

[238] M. Humphreys and F. Nicol, “Understanding the adaptive approach to thermal comfort,” *ASHRAE Trans.*, vol. 104, pp. 991–1004, Jan. 1998.

[239] R. J. de Dear and G. S. Brager, “Thermal comfort in naturally ventilated buildings: revisions to ASHRAE Standard 55,” *Energy Build.*, vol. 34, no. 6, pp. 549–561, Jul. 2002, doi: 10.1016/S0378-7788(02)00005-1.

[240] T. Peffer, M. Pritoni, A. Meier, C. Aragon, and D. Perry, “How people use thermostats in homes: A review,” *Build. Environ.*, vol. 46, no. 12, pp. 2529–2541, Dec. 2011, doi: 10.1016/j.buildenv.2011.06.002.

[241] P. Hoes, J. L. M. Hensen, M. G. L. C. Loomans, B. de Vries, and D. Bourgeois, “User behavior in whole building simulation,” *Energy Build.*, vol. 41, no. 3, pp. 295–302, Mar. 2009, doi: 10.1016/j.enbuild.2008.09.008.

[242] H. B. Gunay, W. O’Brien, and I. Beausoleil-Morrison, “A critical review of observation studies, modeling, and simulation of adaptive occupant behaviors in offices,” *Build. Environ.*, vol. 70, pp. 31–47, Dec. 2013, doi: 10.1016/j.buildenv.2013.07.020.

[243] X. Feng, D. Yan, and T. Hong, "Simulation of occupancy in buildings," *Energy Build.*, vol. 87, pp. 348–359, Jan. 2015, doi: 10.1016/j.enbuild.2014.11.067.

[244] G. Zimmerman, "Modeling and simulation of individual user behavior for building performance predictions," in *Proceedings of the 2007 Summer Computer Simulation Conference*, in SCSC '07. San Diego, CA, USA: Society for Computer Simulation International, Jul. 2007, pp. 913–920.

[245] S. Cao, X. Li, B. Yang, and F. Li, "A review of research on dynamic thermal comfort," *Build. Serv. Eng. Res. Technol.*, vol. 42, no. 4, pp. 435–448, Jul. 2021, doi: 10.1177/01436244211003028.

[246] W. Huang, N. Zhang, C. Kang, M. Li, and M. Huo, "From demand response to integrated demand response: review and prospect of research and application," *Prot. Control Mod. Power Syst.*, vol. 4, no. 1, p. 12, May 2019, doi: 10.1186/s41601-019-0126-4.

[247] J. Blanke, C. Beder, and M. Klepal, "An Integrated Behavioural Model towards Evaluating and Influencing Energy Behaviour—The Role of Motivation in Behaviour Demand Response," *Buildings*, vol. 7, no. 4, Art. no. 4, Dec. 2017, doi: 10.3390/buildings7040119.

[248] Q. Shi, C.-F. Chen, A. Mammoli, and F. Li, "Estimating the Profile of Incentive-Based Demand Response (IBDR) by Integrating Technical Models and Social-Behavioral Factors," *IEEE Trans. Smart Grid*, vol. 11, no. 1, pp. 171–183, Jan. 2020, doi: 10.1109/TSG.2019.2919601.

[249] E. McKenna, S. Higginson, P. Grunewald, and S. J. Darby, "Simulating residential demand response: Improving socio-technical assumptions in activity-based models of energy demand," *Energy Effic.*, vol. 11, no. 7, pp. 1583–1597, Oct. 2018, doi: 10.1007/s12053-017-9525-4.

[250] T. Hong, Y. Chen, Z. Belafi, and S. D'Oca, "Occupant behavior models: A critical review of implementation and representation approaches in building performance simulation programs," *Build. Simul.*, vol. 11, no. 1, pp. 1–14, Feb. 2018, doi: 10.1007/s12273-017-0396-6.

[251] M. M. Abdelrahman, A. Chong, and C. Miller, "Personal thermal comfort models using digital twins: Preference prediction with BIM-extracted spatial–temporal proximity data from Build2Vec," *Build. Environ.*, vol. 207, p. 108532, Jan. 2022, doi: 10.1016/j.buildenv.2021.108532.

[252] V. A. Arowoiya, R. C. Moehler, and Y. Fang, "Digital twin technology for thermal comfort and energy efficiency in buildings: A state-of-the-art and future directions," *Energy Built Environ.*, May 2023, doi: 10.1016/j.enbenv.2023.05.004.

[253] H. Hosamo, H. K. Nielsen, D. Kraniotis, P. R. Svennevig, and K. Svidt, "Improving building occupant comfort through a digital twin approach: A Bayesian network model and predictive maintenance method," *Energy Build.*, 2023, doi: 10.1016/j.enbuild.2023.112992.

[254] H. Park *et al.*, "Integration of computer-simulated persons with multi-node thermoregulation model that considers the effect of clothing for skin surface temperature distribution analysis," *Build. Environ.*, vol. 248, p. 111105, Jan. 2024, doi: 10.1016/j.buildenv.2023.111105.

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

[256] M. B. Kane and K. Sharma, "Data-driven Identification of Occupant Thermostat-Behavior Dynamics," Dec. 13, 2019, *arXiv*: arXiv:1912.06705. doi: 10.48550/arXiv.1912.06705.

[257] D. B. Crawley *et al.*, "EnergyPlus: creating a new-generation building energy simulation program," *Energy Build.*, vol. 33, no. 4, pp. 319–331, Apr. 2001, doi: 10.1016/S0378-7788(00)00114-6.

[258] D. Wang, C. C. Federspiel, and F. Rubinstein, "Modeling occupancy in single person offices," *Energy Build.*, vol. 37, no. 2, pp. 121–126, Feb. 2005, doi: 10.1016/j.enbuild.2004.06.015.

[259] C. Liao, Y. Lin, and P. Barooah, "Agent-based and graphical modelling of building occupancy," *J. Build. Perform. Simul.*, vol. 5, no. 1, pp. 5–25, Jan. 2012, doi: 10.1080/19401493.2010.531143.

[260] I. Richardson, M. Thomson, and D. Infield, "A high-resolution domestic building occupancy model for energy demand simulations," *Energy Build.*, vol. 40, no. 8, pp. 1560–1566, Jan. 2008, doi: 10.1016/j.enbuild.2008.02.006.

[261] M. A. López-Rodríguez, I. Santiago, D. Trillo-Montero, J. Torriti, and A. Moreno-Munoz, "Analysis and modeling of active occupancy of the residential sector in Spain: An indicator of residential electricity consumption," *Energy Policy*, vol. 62, pp. 742–751, Nov. 2013, doi: 10.1016/j.enpol.2013.07.095.

[262] J. Widén, A. M. Nilsson, and E. Wäckelgård, "A combined Markov-chain and bottom-up approach to modelling of domestic lighting demand," *Energy Build.*, vol. 41, no. 10, pp. 1001–1012, Oct. 2009, doi: 10.1016/j.enbuild.2009.05.002.

[263] N. A. Salovich, R. W. Remington, and Y. V. Jiang, "Acquisition of Habitual Visual Attention and Transfer to Related Tasks," *Psychon. Bull. Rev.*, vol. 25, no. 3, pp. 1052–1058, Jun. 2018, doi: 10.3758/s13423-017-1341-5.

[264] O. Brette, T. Buhler, N. Lazaric, and K. Marechal, "Reconsidering the nature and effects of habits in urban transportation behavior," *J. Institutional Econ.*, vol. 10, no. 3, pp. 399–426, Sep. 2014, doi: 10.1017/S1744137414000149.

[265] W. L. Adamowicz and J. D. Swait, "Are Food Choices Really Habitual? Integrating Habits, Variety-seeking, and Compensatory Choice in a Utility-maximizing Framework," *Am. J. Agric. Econ.*, vol. 95, no. 1, pp. 17–41, 2013, doi: 10.1093/ajae/aas078.

[266] V. Olkhov, "The Market-Based Probability of Stock Returns," *SSRN Electron. J.*, 2023, doi: 10.2139/ssrn.4350975.

[267] N. Banovic, T. Buzali, F. Chevalier, J. Mankoff, and A. K. Dey, "Modeling and Understanding Human Routine Behavior," in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, San Jose California USA: ACM, May 2016, pp. 248–260. doi: 10.1145/2858036.2858557.

[268] A. Marchenko and A. Temeljotov-Salaj, "A Systematic Literature Review of Non-Invasive Indoor Thermal Discomfort Detection," *Appl. Sci.*, vol. 10, no. 12, p. 4085, Jun. 2020, doi: 10.3390/app10124085.

[269] M. Schweiker *et al.*, "Review of multi-domain approaches to indoor environmental perception and behaviour," *Build. Environ.*, vol. 176, p. 106804, Jun. 2020, doi: 10.1016/j.buildenv.2020.106804.

[270] S. Carlucci *et al.*, "Modeling occupant behavior in buildings," *Build. Environ.*, vol. 174, p. 106768, May 2020, doi: 10.1016/j.buildenv.2020.106768.

[271] K. Benne *et al.*, "Alfalfa [SWR 18-17]," National Renewable Energy Laboratory (NREL), Golden, CO (United States), Nov. 2018. doi: 10.11578/dc.20181205.2.

[272] A. Petersen, T. Shapinsky, L. Klun, and K. Benne, “Alfalfa,” 2024, Accessed: Apr. 06, 2025. [Online]. Available: https://research-hub.nrel.gov/en/publications/alfalfa

[273] J. Kazil, D. Masad, and A. Crooks, “Utilizing Python for Agent-Based Modeling: The Mesa Framework,” in *Social, Cultural, and Behavioral Modeling*, vol. 12268, R. Thomson, H. Bisgin, C. Dancy, A. Hyder, and M. Hussain, Eds., in Lecture Notes in Computer Science, vol. 12268. , Cham: Springer International Publishing, 2020, pp. 308–317. doi: 10.1007/978-3-030-61255-9_30.

[274] C. Carrión, “Kubernetes Scheduling: Taxonomy, Ongoing Issues and Challenges,” *ACM Comput. Surv.*, vol. 55, no. 7, pp. 1–37, Jul. 2023, doi: 10.1145/3539606.

[275] X. Zou, Y. Hu, Z. Tian, and K. Shen, “Logistic Regression Model Optimization and Case Analysis,” in *2019 IEEE 7th International Conference on Computer Science and Network Technology (ICCSNT)*, Oct. 2019, pp. 135–139. doi: 10.1109/ICCSNT47585.2019.8962457.

[276] M. Bertina, “A data-driven approach to support the automation of thermostats in residential buildings,” masters, Concordia University, 2023. Accessed: Apr. 06, 2025. [Online]. Available: https://spectrum.library.concordia.ca/id/eprint/992884/

[277] Y. Lin, P. Barooah, and S. P. Meyn, “Low-frequency power-grid ancillary services from commercial building HVAC systems,” in *2013 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, Oct. 2013, pp. 169–174. doi: 10.1109/SmartGridComm.2013.6687952.

[278] Russell Profozich, *Wholesale Competition in Regions with Organized Electric Markets*, vol. RM07-19–000. 2008, p. 338. Accessed: Apr. 06, 2025. [Online]. Available: https://www.ferc.gov/media/order-no-719

[279] C. Silva, P. Faria, Z. Vale, and J. M. Corchado, “Demand response performance and uncertainty: A systematic literature review,” *Energy Strategy Rev.*, vol. 41, p. 100857, May 2022, doi: 10.1016/j.esr.2022.100857.

[280] “Long Term Reliability Assessment,” North American Electric Reliability Corporation, Dec. 2024. Accessed: Apr. 06, 2025. [Online]. Available: https://www.nerc.com/pa/RAPA/ra/Reliability%20Assessments%20DL/NERC_Long%20Term%20Reliability%20Assessment_2024.pdf

[281] J. Han, M. A. Piette, and S. Kiliccote, “Field Test Results of Automated Demand Response in a Large Office Building,” in *Energy Enviromental Technologies Division*, Tokyo, Japan, Dec. 2008. [Online]. Available: https://eta-publications.lbl.gov/sites/default/files/lbnl-1131e.pdf

[282] M. Pritoni, K. Salmon, A. Sanguinetti, J. Morejohn, and M. Modera, “Occupant thermal feedback for improved efficiency in university buildings,” *Energy Build.*, vol. 144, pp. 241–250, Jun. 2017, doi: 10.1016/j.enbuild.2017.03.048.

[283] V. Mendon, A. Selvacanabady, M. Zhao, and Z. Taylor, “National Cost-Effectiveness of the Residential Provisions of the 2015 IECC,” PNNL--24240, 1773019, 8652, Jun. 2015. doi: 10.2172/1773019.

[284] D. A. Muñoz, “A study of the potential impact of smart thermostats on residential energy efficiency and demand response in North America,” Thesis, Massachusetts Institute of Technology, 2016. Accessed: Mar. 31, 2025. [Online]. Available: https://dspace.mit.edu/handle/1721.1/104303

[285] N. Frick and L. Schwartz, “Time-Sensitive Value of Efficiency: Use Cases in Electricity Sector Planning and Programs,” None, 1574347, ark:/13030/qt04s956s9, Nov. 2019. doi: 10.2172/1574347.

[286] I. Khan, O. Zedadra, A. Guerrieri, and G. Spezzano, “Occupancy Prediction in IoT-Enabled Smart Buildings: Technologies, Methods, and Future Directions,” *Sensors*, vol. 24, no. 11, Art. no. 11, Jan. 2024, doi: 10.3390/s24113276.

[287] P. Kremer, “Simulation, Challenge Testing and Validation of Solutions for Residential and Commercial Occupancy Sensing/Counting and CO2 Sensing (Final Report),” Iowa State Univ., Ames, IA (United States), DOE-ISU-0941, Dec. 2022. doi: 10.2172/1996572.

[288] D. Blum *et al.*, “Prototyping The BOPTEST Framework For Simulation-Based Testing Of Advanced Control Strategies In Buildings,” 2020, doi: 10.26868/25222708.2019.211276.

[289] “Consumption & Efficiency - U.S. Energy Information Administration (EIA).” Accessed: Oct. 27, 2024. [Online]. Available: https://www.eia.gov/consumption/index.php

[290] Z. Yang, N. Li, B. Becerik-Gerber, and M. Orosz, “A systematic approach to occupancy modeling in ambient sensor-rich buildings,” *SIMULATION*, vol. 90, no. 8, pp. 960–977, Aug. 2014, doi: 10.1177/0037549713489918.

[291] M. S. Bolluk, S. Seyis, R. Aydoğan, and E. K. Özdemir, “A Simplified Bayesian Approach for The Calibration of District-Building Energy Models,” presented at the 2024 European Conference on Computing in Construction, Jul. 2024. doi: 10.35490/EC3.2024.283.

[292] P. Westermann, C. Deb, A. Schlueter, and R. Evins, “Unsupervised learning of energy signatures to identify the heating system and building type using smart meter data,” *Appl. Energy*, vol. 264, p. 114715, Apr. 2020, doi: 10.1016/j.apenergy.2020.114715.

[293] S. Chen, H. Zhang, J. Guan, and Z. Rao, “Agent-based modeling and simulation of stochastic heat pump usage behavior in residential communities,” *Build. Simul.*, vol. 13, no. 4, pp. 803–821, Aug. 2020, doi: 10.1007/s12273-020-0625-2.

[294] H. Burak Gunay, W. O’Brien, and I. Beausoleil-Morrison, “Development of an occupancy learning algorithm for terminal heating and cooling units,” *Build. Environ.*, vol. 93, pp. 71–85, Nov. 2015, doi: 10.1016/j.buildenv.2015.06.009.

[295] S. Kang *et al.*, “Toward A Dynamic Comfort Model for Human-Building Interaction in Grid-Interactive Efficient Buildings: Supported by Field Data,” Feb. 13, 2025, *arXiv*: arXiv:2303.07206. doi: 10.48550/arXiv.2303.07206.

[296] David Kathan and Karin Herzfeld, *Participation of Distributed Energy Resource Aggregations in Markets Operated by Regional Transmission Organizations and Independent System Operators*, vol. RM18-9–000. 2020. Accessed: Apr. 06, 2025. [Online]. Available: https://www.ferc.gov/sites/default/files/2020-09/E-1_0.pdf

[297] R. Jia *et al.*, “Design Automation for Smart Building Systems,” *Proc. IEEE*, vol. 106, no. 9, pp. 1680–1699, Sep. 2018, doi: 10.1109/JPROC.2018.2856932.

[298] S. Nolan and M. O’Malley, “Challenges and barriers to demand response deployment and evaluation,” *Appl. Energy*, vol. 152, pp. 1–10, Aug. 2015, doi: 10.1016/j.apenergy.2015.04.083.

[299] P. Domingues, P. Carreira, R. Vieira, and W. Kastner, “Building automation systems: Concepts and technology review,” *Comput. Stand. Interfaces*, vol. 45, pp. 1–12, Mar. 2016, doi: 10.1016/j.csi.2015.11.005.

[300] W. Wang, J. Chen, and T. Hong, “Occupancy prediction through machine learning and data fusion of environmental sensing and Wi-Fi sensing in buildings,” *Autom. Constr.*, vol. 94, pp. 233–243, Oct. 2018, doi: 10.1016/j.autcon.2018.07.007.

[301] K. S. Cetin and Z. O'Neill, "Smart Meters and Smart Devices in Buildings: a Review of Recent Progress and Influence on Electricity Use and Peak Demand," *Curr. Sustain. Energy Rep.*, vol. 4, no. 1, pp. 1–7, Mar. 2017, doi: 10.1007/s40518-017-0063-7.