

Integrating machine learning for energy management: Applications in power
system, building, and agriculture

by

Xuebo Liu

B.S., Kansas State University, 2020

AN ABSTRACT OF A DISSERTATION

submitted in partial fulfillment of the
requirements for the degree

DOCTOR OF PHILOSOPHY

Mike Wiegers Department of Electrical and Computer Engineering
Carl R. Ice College of Engineering

KANSAS STATE UNIVERSITY
Manhattan, Kansas

2025

Abstract

Energy management sits at the nexus of global sustainability objectives, striving not only to curtail greenhouse gas emissions but also to balance resource conservation and economic viability. Against this backdrop, machine learning (ML) has become an increasingly vital tool, offering flexible, data-driven methods for forecasting, optimization, and real-time decision-making in complex energy environments. By leveraging high-resolution data—from solar and wind generation to occupant behavior and crop conditions—these approaches uncover patterns and adapt strategies in real-time to enhance both operational efficiency and system resilience. This dissertation explores the application of tailored ML solutions across three critical areas. 1. For island power systems with high renewable penetration, it demonstrates how data-driven frequency nadir constraints can bolster grid stability without imposing excessive costs. 2. For occupant-centric building controls, it illustrates how multi-agent reinforcement learning can considerably reduce energy use while preserving occupant comfort. 3. Within agricultural operations, it highlights how advanced reinforcement learning frameworks can align irrigation, fertilization, and renewable resource utilization to increase yields while limiting environmental impact. Taken together, these studies exemplify the transformative role ML can play in shaping a more sustainable and cost-effective energy future.

The first section of this dissertation focuses on island or weakly interconnected power grids, which face unique challenges due to high penetrations of inverter-based renewable energy resources. As intermittent sources like photovoltaics reduce overall system inertia, operators must ensure stable frequency responses following severe disturbances such as generation outages. To this end, a data-driven unit commitment (UC) model is proposed, incorporating frequency nadir constraints derived from comprehensive dynamic simulations and year-long generation data. By capturing the relationship between scheduled generation

and post-disturbance frequency levels, these constraints effectively limit frequency deviations to acceptable ranges, thereby improving system reliability. Extensive simulations reveal that this approach yields significantly more robust frequency nadirs than a simpler minimum inertia constraint, with only a marginal increase in generation costs. The result is a more sustainable, renewable-powered grid that maintains stability and reliability without compromising economic feasibility.

Next, the dissertation addresses energy efficiency in buildings, where HVAC (Heating, Ventilation, and Air Conditioning) systems account for a considerable share of total electricity usage. Traditional control methodologies often rely on static occupant behavior assumptions, which fail to capture short-term fluctuations in clothing insulation, metabolic rates, and occupancy patterns. To overcome these limitations, a multi-agent deep reinforcement learning (MADRL) framework is introduced for multi-zone HVAC control. In this setup, each zone is managed by an intelligent agent that learns to dynamically adjust heating and cooling setpoints based on real-time occupant behavior and energy cost signals. Simulation studies demonstrate that the proposed occupant-centric approach reduces electricity expenses significantly compared to rule-based methods and by slightly saving relative to single-agent deep reinforcement learning (DRL), all while preserving occupant comfort. Such improvements underscore the role of advanced learning algorithms in achieving energy efficiency and occupant satisfaction simultaneously, thereby contributing to the broader sustainability agenda.

The final section explores the Food-Energy-Water (FEW) nexus, where agriculture intersects with renewable energy utilization. A novel framework is developed to integrate solar power, green ammonia production, and deep reinforcement learning-based optimization. By producing ammonia on-site using renewable energy, farms can leverage a valuable agricultural input while also storing excess energy for later use. The intelligent control system, driven by DRL, coordinates the timing of ammonia production, energy storage, and irrigation schedules to maximize farm revenue and optimize resource use. Simulation results confirm that this integrated approach can significantly enhance both economic and

environmental outcomes, paving the way for greener, more resilient agricultural practices.

Collectively, these three ML-driven frameworks showcase the versatility and impact of data-based decision-making in advancing sustainability and energy efficiency. Whether in stabilizing low-inertia power grids, managing building energy consumption with occupant-centric insights, or aligning agricultural operations with renewable resources, machine learning stands out as a transformative technology. By rigorously addressing each domain's distinct challenges and synthesizing overarching lessons learned, this dissertation provides a blueprint for policymakers, engineers, and researchers aiming to foster sustainable energy systems. Ultimately, the proposed methodologies demonstrate that targeted applications of ML can lead to robust, cost-effective, and resource-conscious energy management solutions, propelling global efforts toward a more efficient and low-carbon future.

Integrating Machine Learning for Energy Management: Applications in
Power System, Building, and Agriculture

by

Xuebo Liu

B.S., Kansas State University, 2020

A DISSERTATION

submitted in partial fulfillment of the
requirements for the degree

DOCTOR OF PHILOSOPHY

Mike Wieggers Department of Electrical and Computer Engineering
Carl R. Ice College of Engineering

KANSAS STATE UNIVERSITY
Manhattan, Kansas

2025

Approved by:

Major Professor
Dr. Hongyu Wu

Copyright

© Xuebo Liu 2025.

Abstract

Energy management sits at the nexus of global sustainability objectives, striving not only to curtail greenhouse gas emissions but also to balance resource conservation and economic viability. Against this backdrop, machine learning (ML) has become an increasingly vital tool, offering flexible, data-driven methods for forecasting, optimization, and real-time decision-making in complex energy environments. By leveraging high-resolution data—from solar and wind generation to occupant behavior and crop conditions—these approaches uncover patterns and adapt strategies in real-time to enhance both operational efficiency and system resilience. This dissertation explores the application of tailored ML solutions across three critical areas. 1. For island power systems with high renewable penetration, it demonstrates how data-driven frequency nadir constraints can bolster grid stability without imposing excessive costs. 2. For occupant-centric building controls, it illustrates how multi-agent reinforcement learning can considerably reduce energy use while preserving occupant comfort. 3. Within agricultural operations, it highlights how advanced reinforcement learning frameworks can align irrigation, fertilization, and renewable resource utilization to increase yields while limiting environmental impact. Taken together, these studies exemplify the transformative role ML can play in shaping a more sustainable and cost-effective energy future.

The first section of this dissertation focuses on island or weakly interconnected power grids, which face unique challenges due to high penetrations of inverter-based renewable energy resources. As intermittent sources like photovoltaics reduce overall system inertia, operators must ensure stable frequency responses following severe disturbances such as generation outages. To this end, a data-driven unit commitment (UC) model is proposed, incorporating frequency nadir constraints derived from comprehensive dynamic simulations and year-long generation data. By capturing the relationship between scheduled generation

and post-disturbance frequency levels, these constraints effectively limit frequency deviations to acceptable ranges, thereby improving system reliability. Extensive simulations reveal that this approach yields significantly more robust frequency nadirs than a simpler minimum inertia constraint, with only a marginal increase in generation costs. The result is a more sustainable, renewable-powered grid that maintains stability and reliability without compromising economic feasibility.

Next, the dissertation addresses energy efficiency in buildings, where HVAC (Heating, Ventilation, and Air Conditioning) systems account for a considerable share of total electricity usage. Traditional control methodologies often rely on static occupant behavior assumptions, which fail to capture short-term fluctuations in clothing insulation, metabolic rates, and occupancy patterns. To overcome these limitations, a multi-agent deep reinforcement learning (MADRL) framework is introduced for multi-zone HVAC control. In this setup, each zone is managed by an intelligent agent that learns to dynamically adjust heating and cooling setpoints based on real-time occupant behavior and energy cost signals. Simulation studies demonstrate that the proposed occupant-centric approach reduces electricity expenses significantly compared to rule-based methods and by slightly saving relative to single-agent deep reinforcement learning (DRL), all while preserving occupant comfort. Such improvements underscore the role of advanced learning algorithms in achieving energy efficiency and occupant satisfaction simultaneously, thereby contributing to the broader sustainability agenda.

The final section explores the Food-Energy-Water (FEW) nexus, where agriculture intersects with renewable energy utilization. A novel framework is developed to integrate solar power, green ammonia production, and deep reinforcement learning-based optimization. By producing ammonia on-site using renewable energy, farms can leverage a valuable agricultural input while also storing excess energy for later use. The intelligent control system, driven by DRL, coordinates the timing of ammonia production, energy storage, and irrigation schedules to maximize farm revenue and optimize resource use. Simulation results confirm that this integrated approach can significantly enhance both economic and

environmental outcomes, paving the way for greener, more resilient agricultural practices.

Collectively, these three ML-driven frameworks showcase the versatility and impact of data-based decision-making in advancing sustainability and energy efficiency. Whether in stabilizing low-inertia power grids, managing building energy consumption with occupant-centric insights, or aligning agricultural operations with renewable resources, machine learning stands out as a transformative technology. By rigorously addressing each domain's distinct challenges and synthesizing overarching lessons learned, this dissertation provides a blueprint for policymakers, engineers, and researchers aiming to foster sustainable energy systems. Ultimately, the proposed methodologies demonstrate that targeted applications of ML can lead to robust, cost-effective, and resource-conscious energy management solutions, propelling global efforts toward a more efficient and low-carbon future.

Table of Contents

List of Figures	xiii
List of Tables	xvi
List of Nomenclature	xviii
Acknowledgements	xxii
Dedication	xxiv
Preface	xxv
1 Introduction	1
1.1 Background	1
1.2 Introduction to Machine Learning	2
1.2.1 Supervised Learning	2
1.2.2 Reinforcement Learning	3
1.2.3 Unsupervised Learning	4
1.2.4 Selecting Machine Learning Methods	4
1.3 Machine Learning in Power System of Frequency Stability Assessment and Unit Commitment	5
1.4 Machine Learning in Buildings Energy Management	7
1.5 Machine Learning in Agriculture Energy Management	11
1.6 Research Motivations	14
1.7 Research Questions	15

1.8	Research Contributions	15
1.9	Organization of This Dissertation	17
2	Machine Learning for Frequency-Constrained Unit Commitment in Power Systems	19
2.1	Machine Learning Prediction in Frequency Nadir	21
2.1.1	Dataset and Methodology	24
2.1.2	Machine Learning Methods	24
2.1.3	Preliminary Results	28
2.1.4	Preliminary Summary	31
2.2	Unit Commitment Model with Frequency Nadir Constraint	31
2.2.1	Data Driven-based Frequency Nadir Estimation	32
2.2.2	Proposed FNC-UC Framework	35
2.3	Experiment and Case Studies	41
2.4	Summary	54
3	Machine Learning in Occupant-centric Building Energy Management	56
3.1	Optimization and occupant behavior validation in buildings	58
3.1.1	ADP-based Optimization Problem	61
3.1.2	ADP Algorithm and MPC Framework	67
3.1.3	Preliminary Results	70
3.1.4	Preliminary Summary	76
3.2	Advanced Reinforcement Learning in Multi-zone Buildings	78
3.2.1	Problem Formulation and Methodology	79
3.3	Experiment and Case Studies	92
3.3.1	Discussion	103
3.4	Summary	104
4	Machine Learning in Food-Energy-Water Nexus Energy Management	106
4.1	Planning Energy Systems for Farm-Scale Ammonia Production	108

4.1.1	Modeling Farm-scale Ammonia Production	110
4.1.2	Problem Formulation of Planning in Farm Energy	112
4.1.3	Preliminary Results	116
4.1.4	Preliminary Summary	119
4.2	Deep Reinforcement Learning based Farm Energy Management Considering Ammonia Generation	119
4.2.1	Modeling of the single farm within the FEW Nexus	120
4.2.2	Approach of FEWture-DRL Framework	129
4.3	Experiment and Case Studies	136
4.3.1	Simulation Setup	136
4.3.2	Simulation Results	138
4.3.3	Discussion	141
4.4	Summary	143
5	Conclusion and Future Work	145
5.1	Conclusion	145
5.2	Future Work	147
	Bibliography	149
A	Acronyms	181
B	Reuse permissions from publishers	185

List of Figures

1.1	Structure Diagram of this Dissertation	17
2.1	Illustration of Renewable Energy of Frequency Nadir Impacts	22
2.2	Simplified ANN model	26
2.3	Flowchart of generating UGD and TIH datasets	27
2.4	Frequency nadir versus total inertia and headroom	27
2.5	Error distribution of UGD (top) and TIH (bottom) testing	30
2.6	Illustration of the proposed frequency nadir constraint unit commitment approach	35
2.7	Flowchart of the FNC-UC model.	40
2.8	Generation capacity mix in the island system.	42
2.9	Load curve in the simulated cases.	43
2.10	Annual generation cost of the three cases.	44
2.11	The generation cost duration of the three cases.	45
2.12	The five-day generation dispatch examples of the three cases.	46
2.13	Frequency nadir for each hour of the year for the N-1 event.	48
2.14	Frequency nadir distribution of the three cases.	49
2.15	Hourly counts of the frequency nadir lower than 59 Hz in the three cases. . .	50
2.16	The system scheduled one-month total inertia in the three cases.	51
2.17	The inertia duration curves in the three cases.	51
3.1	The schematic diagram of proposed energy management.	61

3.2	Illustration of a state transition diagram considering occupant's clothing states and actions in winter (when $T = 3$).	67
3.3	Flowchart of the ADP based MPC framework.	69
3.4	The time-of-use electricity tariff.	71
3.5	Simulation results of a female winter case and a male summer case, both without the proposed HVAC energy management scheme.	72
3.6	Simulation results of a winter day for female with ADP.	73
3.7	Simulation results of a summer day for male with ADP.	75
3.8	Illustration of proposed MADRL in multi-zone HVAC control. Three modules are contained in this approach, 1) a MADRL algorithm: Multi-agent Deep Deterministic Policy Gradient; 2) A testbed: EnergyPlus simulator; 3) Stochastic occupant behavior model includes clothing, occupant presence, and metabolic rate.	80
3.9	One-day average occupant presence schedule (top) and activity level schedule (bottom) for each zone	83
3.10	One-year Tampa weather profile	94
3.11	Average episode reward during training of DDPG and MADDPG	95
3.12	Three-day behavior simulation result for all zones (base, shoulder, and peak prices are represented by white, light grey, and grey)	97
3.13	Results of Zone 1 over a three-day period, with the Rule-based method (top), DDPG (middle), and MADDPG (bottom). The red line represents heating setpoints, the blue line indicates cooling setpoints, the black line represents the zone's average temperature, and the green dashed line represents the PMV-PPD value, while the light blue line depicts the outside temperature.	98

3.14	Results of Zone 2, 3, 4, and 5 over a three-day period. The red line represents heating setpoints, the blue line indicates cooling setpoints, the black line represents the zone’s average temperature, and the green dashed line represents the PMV-PPD value, while the light blue line depicts the outside temperature.	99
3.15	Three days simulation of electricity cost and thermal comfort result (PMV-PPD) for the whole building.	100
3.16	One-month control patterns of HVAC for all zones	101
3.17	One-month HVAC decisions histogram comparison of whole building	102
4.1	Illustration of farm hybrid renewable energy system.	109
4.2	A food-energy-water abstract model in the Stella Software ¹ . No. PV panels and No. Turbines determine the PV capacity and wind capacity.	111
4.3	Flowchart of the proposed framework.	111
4.4	Illustration of farms with 1 TPD ammonia plant with its budget allocation for HRES and impact on total energy costs.	118
4.5	(a) Illustration of FEW nexus in an individual farm. (b) DQN-enabled decision-making architecture integrated into farm management.	121
4.6	Wind Power Output Generation Vs Wind Speed	123
4.7	Weather data illustration for one year of daily data	137
4.8	Comparison of (a) rewards and (b) crop yield between DRL algorithms.	139
4.9	Action comparison between algorithms	140
4.10	Action frequency comparison between algorithms	142

List of Tables

1.1	Comparison of related works in the literature (in publication time order) . .	8
1.2	Comparison of related works in the literature (in publication time order) . .	9
1.3	Literature review of DRL in HVAC control	10
2.1	Hyper-parameters of ANN	28
2.2	Evaluation metrics of UGD training	29
2.3	Evaluation metrics of UGD testing	29
2.4	Evaluation metrics of TIH training	30
2.5	Evaluation metrics of TIH testing	30
2.6	PARAMETER VALUES OF THE LINEAR REGRESSION FREQUENCY NADIR CON- STRAINT FOR THE ISLAND SYSTEM STUDIED	40
2.7	PARAMETERS OF THE PV-BESS PLANTS	42
2.8	ONE-YEAR GENERATION COSTS IN THE THREE CASES	45
2.9	PROBABILITIES OF FREQUENCY NADIR LOWER THAN 59 HZ IN THE THREE CASES	49
2.10	SYSTEM INERTIA METRICS FOR THE THREE CASES (MW*s)	52
2.11	NADIR NEAR 60 HZ IN THE MINIMUM INERTIA CONSTRAINT CASE	52
2.12	NADIR CONSTRAINT CASE HAS A LOWER NADIR THAN THE BASELINE CASE	53
3.1	Comparison of Related Works in the Literature	60
3.2	Clo values of typical business casual clothing ensembles	66
3.3	HVAC parameters and clothing conditions	71
3.4	A daily comparison between non-optimized and optimized cases.	77

3.5	Activity level and metabolic rate	84
3.6	HVAC PARAMETERS	93
3.7	MADDPG network information	94
3.8	Training and execution time comparison	96
3.9	One-year comparison between different methods	102
4.1	Renewable Energy Parameters	116
4.2	Comparison of HRES sizing between scenarios with 1 TPD ammonia plant ("Cap." is Capacity; "Total Cost" is defined in Eq. (4.2))	118
4.3	Observation Space (State Variables)	130
4.4	Action Space and Reward Function	131
4.5	Parameter settings for WOFOST and SNOMIN model simulation, including soil hydraulic and organic matter properties for four soil layers. The data is sourced from example simulations in the PCSE documentation ^{2;3}	138
4.6	Cost parameters for farm environment simulation with associated mathemat- ical symbols. NH3 refers to ammonia.	138
4.7	Summary of different farming strategies per unit area of crop land.	141

Nomenclature

Constants

H_i Inertia of unit i (s)

$\overline{P}_{p,t}$ Forecasted PV power output (MPPT) of unit p at time t (MW)

$G_{i,t}^{max}$ Maximum generation output of unit i at time t (MW)

$G_{i,t}^{min}$ Minimum generation output of unit i at time t (MW)

$T_{i,MinUp}$ Minimum uptime of unit i (h)

$T_{i,MinDn}$ Minimum downtime of unit i (h)

GSF_{l-i} Generation shift factor from bus i to line l

LP Load-shedding penalty price (\$/MWh)

H Total system inertia (MW·s)

RP/PFP Penalty price for regulation and primary frequency reserve shortage (\$/MW)

PFR_t^r Primary frequency reserve requirement at time t (MW)

SU_i Startup cost of unit i (\$)

$Reg_{u,t}^r$ Regulation-up reserve requirement at time t (MW)

$Reg_{d,t}^r$ Regulation-down reserve requirement at time t (MW)

SD_i	Shutdown cost of unit i (\$)
R_i^U	Ramp-up limit for unit i (MW/h)
R_i^D	Ramp-down limit for unit i (MW/h)
R_i^{SU}	Ramp-up limit for unit i when starting up (MW/h)
R_i^{SD}	Ramp-down limit for unit i when shutting down (MW/h)
$Limit_l$	Transmission limit for line l (MW)
$\overline{D_{b,t}}$	Forecasted power demand at load bus b at time t (MW)

Indices

b	Index for load buses
i	Index for generation units
p	Index for photovoltaic (PV) generation units
t	Index for time intervals
T	Time span or total number of time intervals
l	Index for transmission lines

Additional Parameters and Variables

Δt	Resolution of time interval
rh_t	Relative humidity at time t (%)
$\alpha_t^{\mathcal{D}}$	Discomfort weight at time t

α_t^c	Cost weight at time t
β_t	Discomfort penalty at time t
c_t	Electricity price at time t (\$/kWh)
$\mathcal{D}$	Function of discomfort
$\mathcal{C}$	Function of electricity cost
$\mathcal{F}_s$	Sobol sampling function
$(\cdot)'$	Superscript representing sampled
T_{kd}	K-D tree
s_t^i	Pre-decision state of i at time t ($^{\circ}\text{C}$)
$\dot{s}_t$	Post-decision state at time t ($^{\circ}\text{C}$)
s'_t	Sampled state at time t ($^{\circ}\text{C}$)
$\mathbf{S}$	State space
x_t^i	Decision of i at time t
$\mathbf{X}$	Decision space
$\tilde{u}_t$	Uncertain outside temperature at time t ($^{\circ}\text{C}$)

Variables

$cp_{i,t}$	Production cost of unit i at time t (\$)
$\Delta Reg_{d,t}$	System regulation-down shortage at time t (MW)

$Reg_{u,i,t}$ Regulation-up capacity provided by unit i at time t (MW)

$Reg_{d,i,t}$ Regulation-down capacity provided by unit i at time t (MW)

$Reg_{u,p,t}$ Regulation-up capacity provided by PV plant p at time t (MW)

$Reg_{d,p,t}$ Regulation-down capacity provided by PV plant p at time t (MW)

$PFR_{i,t}$ Primary frequency reserve capacity of unit i at time t (MW)

$v_{i,t}$ Commitment status of unit i at time t (binary)

$u_{i,t}$ Startup status of unit i at time t (binary)

$w_{i,t}$ Shutdown status of unit i at time t (binary)

$G_{i,t}$ Generation output of unit i at time t (MW)

$\overline{G_{i,t}}$ Maximum available generation output of unit i at time t (MW)

$P_{p,t}$ PV power output of unit p at time t (MW)

$D_{b,t}$ Scheduled demand for bus b at time t (MW)

$\Delta D_{b,t}$ Load shedding at bus b at time t (MW)

ΔPFR_t System-wide primary frequency reserve shortage at time t (MW)

ΔD_t Total system load shedding at time t (MW)

$\Delta Reg_{u,t}$ System regulation-up shortage at time t (MW)

Acknowledgments

I would like to express my deepest gratitude to Dr. Hongyu Wu. Your unwavering guidance, support, and encouragement have been instrumental throughout my entire Ph.D. journey. It has been a true honor to work under your mentorship. From 2019, when I began as an undergraduate research assistant, to now completing my Ph.D., I am sincerely thankful for your inspiration, patience, and belief in my potential. Your mentorship has shaped not only my academic development but also my professional growth, and I am forever grateful.

I also extend my sincere appreciation to my committee members from Kansas State University (KSU)—Dr. Anil Pahwa and Dr. Caterina Scoglio from the Mike Wiegers Department of Electrical and Computer Engineering, and Dr. Vincent Amanor-Boadu from the Department of Agricultural Economics—for their time, insightful feedback, and continuous support. Their teaching, guidance, and encouragement throughout my coursework and dissertation process have greatly enriched my academic experience and played a vital role in the successful completion of my PhD program.

I am deeply grateful to my mentors at the National Renewable Energy Laboratory (NREL), especially Dr. Jin Tan, for her exceptional guidance during my year-round internship. Special thanks to Dr. Yingying Wu from the Department of Interior Design and Fashion Studies at KSU, for her collaboration and insightful support in the areas of home energy management and occupant behavior modeling. I would also like to thank Dr. Mary C. Hill from the University of Kansas for her mentorship and support throughout the FEW-ture project. Her leadership and interdisciplinary approach were instrumental in shaping the broader impact of this research.

Finally, I would like to acknowledge all the members of the Kansas Smart Energy Systems Group—Dr. Bo Liu, Dr. Hang Zhang, Dr. Lawryn Kiboma, Yichen Liu, Timur Aminov, Dr. Nazif Faquiry, Li Wang, and Dr. Ananth Palani—for their collaboration,

support, and the many memorable moments we shared. I also extend my sincere thanks to my colleagues in the FEWture project team and my colleagues at the NREL for their invaluable support throughout this journey.

This work was supported by the following organizations and grants:

1. U.S. National Science Foundation (NSF)
Grant Nos. 1856084 and 2316355
2. U.S. Department of Energy (DOE), Office of Energy Efficiency and Renewable Energy,
Solar Energy Technologies Office
Award Nos. 34224 and 37772
3. National Renewable Energy Laboratory (NREL)
Operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy
under Contract No. DE-AC36-08GO28308

Dedication

This dissertation is dedicated to my beloved wife, Chunjin Zheng, whose unwavering love, encouragement, and sacrifices have been the bedrock of my journey. To my son, Felix Z. Liu, your laughter and boundless energy inspire me every day and remind me of the importance of perseverance and joy.

I am deeply grateful to my parents, Shengli Liu and Hongyan Zhu, for their endless support, wisdom, and belief in my potential. To my siblings, my little sister - Xueduo Liu, and my little brother - Xuelin Liu, thank you for your encouragement and the love that has enriched my life.

Preface

Below is a list of my peer-reviewed journal and conference publications resulting from this research effort:

Refereed Journal Publications

- [J1] **X. Liu**, L. Kiboma, V. Amanor-Boadu, A. Modarresi, M. Hill, and H. Wu, “Deep Reinforcement Learning-based Farm Energy Management Considering Small-scale Ammonia Generation: Approach Within Food Energy-Water Nexus,” In preparation, Journal of Cleaner Production.
- [J2] **X. Liu**, X. Fang, H. Yuan, J. Tan, and H. Wu, “Frequency stability constrained unit commitment: an island system study,” IEEE Open Access Journal of Power and Energy, vol. 11, pp. 141–153, 2024.
- [J3] Y. Wu, **X. Liu**, K. Morris, S. Lu, H. Wu, “An exploratory study of body measurements prediction using machine learning and 3D body scans,” Clothing and Textiles Research Journal, 2024.
- [J4] **X. Liu**, Y. Wu, and H. Wu, “Enhancing HVAC energy management through multi-zone occupant-centric approach: A multi-agent deep reinforcement learning solution,” Energy and Buildings, vol. 303, 2024.
- [J5] T. Ali, H. Wu, A. Palani, A. Sinha, and **X. Liu**, “A Multi-Cut L-Shaped Decomposition Algorithm for Stochastic SCUC with Renewable Generation,” SSRN, 2022.
- [J6] **X. Liu**, Y. Wu, and H. Wu, “Machine Learning Enabled 3D Body Measurement Estimation,” Applied Sciences, vol. 12, no. 14, 2022.
- [J7] H. Zhang, B. Liu, **X. Liu**, A. Pahwa, and H. Wu, “Voltage Stability Constrained MTD Against Net Load Redistribution Attacks,” IEEE Transactions on Smart Grid, vol. 13,

no. 5, 2022.

- [J8] H. Wu, A. Pratt, P. Munankarmi, M. Lunacek, S. Balamurugan, **X. Liu**, and P. Spit-sen, “Impact of MPC-enabled Home Energy Management on Distribution Systems,” *Advances in Applied Energy*, vol. 6, 2022.
- [J9] **X. Liu**, Y. Wu, and H. Wu, “PV-EV Integrated Home Energy Management with Residential Occupant Behaviors,” *Sustainability*, vol. 13, no. 24, 2021.
- [J10] **X. Liu**, Y. Wu, H. Zhang, and H. Wu, “Hourly Occupant Clothing Decisions in Residential HVAC Energy Management,” *Journal of Building Engineering*, vol. 40, 2021.
- [J11] **X. Liu**, H. Wu, L. Wang, and M. Faqiry, “A Stochastic Home Energy Management System via Approximate Dynamic Programming,” *IET Energy Systems Integration*, vol. 2, no. 4, 2020.

Refereed Conference Publications

- [C1] **X. Liu**, *et al.*, “Budget-Constrained Sizing of Renewable and Energy Storage Systems for Farm-Scale Ammonia Production Within the Food-Energy-Water Nexus,” in *Proc. IEEE Green Technologies Conf. (GreenTech)*, Springdale, AR, USA, 2024, pp. 218–222.
- [C2] Y. Liu, B. Liu, **X. Liu**, and H. Wu, “False Data Injection Attacks Based on Least Squares Generative Adversarial Networks with Reconstruction Loss,” in *Proc. IEEE Power & Energy Society General Meeting (PESGM)*, 2024 (accepted).
- [C3] B. Liu, Y. Liu, **X. Liu**, and H. Wu, “Deep-Learning-State-Estimation-Aided Detection Framework Against False Data Injection Attacks,” in *Proc. IEEE Kansas Power and Energy Conf. (KPEC)*, Manhattan, KS, USA, 2024, pp. 1–5.
- [C4] **X. Liu**, Y. Wu, B. Liu, and H. Wu, “A Multi-Agent Deep Deterministic Policy Gradient Method for Multi-Zone HVAC Control,” in *Proc. IEEE PESGM*, Orlando, FL, USA, 2023, pp. 1–5.

- [C5] B. Liu, Q. Yang, H. Zhang, **X. Liu**, and H. Wu, “Random-Based Hidden Moving Target Defense Against Alert False Data Injection Attackers,” in *Proc. IEEE PESGM*, Orlando, FL, USA, 2023, pp. 1–5.
- [C6] B. Liu, H. Wu, Q. Yang, **X. Liu**, and Y. Liu, “Data-Driven FDI Attacks: A Stealthy Approach to Subvert SVM Detectors in Power Systems,” in *Proc. IEEE KPEC*, Manhattan, KS, USA, 2023, pp. 1–6.
- [C7] **X. Liu**, Y. Wu, H. Zhang, B. Liu, L. Edmonds, and H. Wu, “Home Energy Management with Clothing Integrated Thermal Comfort and EV SOC Concern,” in *Proc. IEEE PESGM*, Denver, CO, USA, 2022, pp. 1–5.
- [C8] **X. Liu**, *et al.*, “A Comparison of Machine Learning Methods for Frequency Nadir Estimation in Power Systems,” in *Proc. IEEE KPEC*, Manhattan, KS, USA, 2022, pp. 1–5.
- [C9] H. Zhang, N. Fulk, B. Liu, L. Edmonds, **X. Liu**, and H. Wu, “Load Margin Constrained Moving Target Defense Against False Data Injection Attacks,” in *Proc. IEEE GreenTech*, Houston, TX, USA, 2022, pp. 51–56.
- [C10] Y. Wu, **X. Liu**, K. Morris, S. Lu, and H. Wu, “An Exploratory Study of Body Measurements Prediction Using Machine Learning and 3D Body Scans,” in *Proc. Int’l Textile and Apparel Association (ITAA) Annual Conf.*, vol. 79, no. 1, Denver, CO, USA, 2022.
- [C11] L. Edmonds, **X. Liu**, and H. Wu, “A Demand-Responsive Green Ammonia Plant and its Impact on the Electricity Distribution System,” in *Proc. IEEE PESGM*, Denver, CO, USA, 2022, pp. 1–5.
- [C12] O. Al-Ani, S. Das, H. Wu, D. Martinez-Figueroa, **X. Liu**, and R. H. Cheppally, “A Deep Architecture Using Multiple TabNets for Environmental Prediction and Control in Smart Homes,” in *Proc. IEEE Symposium Series on Computational Intelligence (SSCI)*, Singapore, 2022, pp. 998–1003.

- [C13] **X. Liu**, M. Ross, H. Bindra, and H. Wu, “Optimal Sizing of Battery Energy Storage Systems for Small Modular Reactor Based Microgrids,” in *Proc. IEEE KPEC*, Manhattan, KS, USA, 2021, pp. 1–6.

Chapter 1

Introduction

This chapter introduces the concepts studied in this dissertation. It also emphasizes the importance and necessity of this research as well as provides detailed research questions and contributions of this dissertation.

1.1 Background

Energy sustainability has become a central priority for governments, industries, and research communities worldwide, driven by mounting concerns about climate change, resource scarcity, and the need for economic resilience⁴. A pivotal aspect of achieving sustainability lies in improving energy efficiency—minimizing energy waste while meeting societal and economic demands. From residential buildings to large-scale power grids, energy efficiency initiatives can mitigate environmental impacts, reduce costs, and support long-term decarbonization goals.

Data and sensors (i.e., Internet-of-Things) provide a valuable opportunity to apply Machine Learning (ML)—to uncover complex patterns and optimize energy use across diverse domains⁵. By leveraging ML, system operators, and policymakers can make more informed decisions, manage renewable integration effectively, and streamline resource usage in the

real-time.

1.2 Introduction to Machine Learning

ML is a subfield of artificial intelligence that enables computer systems to learn and improve from data and experience, without being explicitly programmed^{6;7}. ML algorithms extract patterns and correlations from historical datasets, facilitating accurate predictions and effective decision-making strategies across diverse fields, including healthcare, finance, and particularly energy management systems⁸. Generally, ML algorithms are categorized into three distinct types: supervised learning, unsupervised learning, and reinforcement learning. Due to the scope of this dissertation, special emphasis is placed on supervised and reinforcement learning, given their extensive applications in energy management systems.

1.2.1 Supervised Learning

Supervised learning involves training algorithms on labeled datasets, where each input is paired with a specific target or outcome. The core objective of supervised learning is to accurately model the relationship between input variables (features) and output variables (predictions or classifications). Within supervised learning, two major tasks are frequently addressed:

- **Regression:** Predicts continuous values, such as energy consumption, renewable generation (e.g., solar and wind power), or electricity market prices^{9;10}.
- **Classification:** Categorizes data into discrete classes, suitable for tasks such as fault detection, equipment diagnostics, or determining operational states in power systems^{11;12}.

Widely used supervised learning algorithms include linear regression, support vector machines (SVM), decision tree-based ensemble methods (random forests, gradient boosting), and artificial neural networks (ANNs). Specifically, linear regression offers interpretability

and computational efficiency for simpler predictive tasks⁹. Ensemble methods, such as XGBoost, provide higher accuracy in load forecasting and fault detection due to their ability to handle complex nonlinear interactions¹³. Moreover, artificial neural networks and deep learning models are capable of modeling intricate relationships within data, essential for predicting intermittent renewable energy sources¹⁴. Supervised learning is highly effective when substantial labeled historical data is available, although performance heavily depends on data quality, preprocessing, and validation strategies.

1.2.2 Reinforcement Learning

Reinforcement learning (RL) addresses sequential decision-making problems by enabling computational agents to interact iteratively with dynamic and uncertain environments. RL algorithms continuously update their policies based on the feedback (rewards or penalties) received, aiming to maximize cumulative long-term rewards. This method is particularly valuable in applications requiring real-time adaptive control, such as energy system operations, robotics, and autonomous systems¹⁵.

Key reinforcement learning methods include:

- **Value-based Methods (e.g., Q-learning, Deep Q-Network (DQN))**: Estimate optimal actions by evaluating state-action values, suitable for discrete decision-making tasks¹⁶.
- **Policy Gradient and Actor-Critic Methods (e.g., Advantage Actor-Critic (A2C), Deep Deterministic Policy Gradient (DDPG))**: Directly optimize policies for continuous action spaces, essential for precise control applications like dynamic power output adjustments or temperature setpoint management^{15;17}.
- **Multi-Agent Reinforcement Learning (MARL)**: Deals with environments involving multiple interacting agents, useful for coordinated control of distributed energy resources, multi-zone climate management, and grid stability enhancement¹⁸.

Although reinforcement learning exhibits powerful adaptive capabilities, practical appli-

cations may encounter challenges related to computational complexity, data efficiency, and optimal exploration strategies.

1.2.3 Unsupervised Learning

Unsupervised learning analyzes unlabeled datasets to uncover inherent data structures or patterns without predefined targets. Common techniques, such as K-means clustering and principal component analysis (PCA), are typically employed in exploratory data analysis, anomaly detection, and dimensionality reduction tasks¹⁹. While unsupervised learning can enhance data understanding and preprocessing, it generally complements supervised or reinforcement learning rather than independently addressing predictive or sequential decision-making tasks in energy management applications.

1.2.4 Selecting Machine Learning Methods

When applying ML to energy management, practitioners should carefully select appropriate methods based on the problem context, data availability, and system dynamics²⁰:

- **Supervised Learning** is preferred when extensive historical data with clear labels is available, ideal for predictive tasks and classification scenarios.
- **Reinforcement Learning** is beneficial in environments that require sequential, real-time adaptive decision-making under uncertainty.
- **Unsupervised Learning** serves primarily in preliminary data analysis, exploratory tasks, or supportive roles in broader ML pipelines.

Understanding these distinctions ensures the effective selection and application of ML algorithms to address complex real-world problems. This dissertation examines chapters that explore the transformative role of ML across three critical domains of energy management: power systems, building energy management, and agricultural energy management. In power systems, the rapid integration of inverter-based renewable resources challenges frequency stability due to reduced inertia and heightened variability; Chapter 2 addresses this by proposing

a supervised learning-driven frequency nadir estimation integrated into a computationally efficient Mixed-Integer Linear Programming (MILP)-based Unit Commitment (UC) framework, validated through extensive real-world simulations. Subsequently, Chapter 3 shifts focus to building energy management, acknowledging buildings’ significant contribution to global energy consumption, and tackles the complexity arising from dynamic occupant behaviors—such as occupancy, clothing adjustments, and metabolic variations—through a novel occupant-centric multi-agent Deep Reinforcement Learning (MADRL) approach for multi-zone HVAC control, significantly improving energy efficiency and occupant comfort. Finally, recognizing agriculture’s critical role in human sustenance and its substantial environmental impact, Chapter 4 presents an innovative framework integrating renewable energy, green ammonia production, and reinforcement learning within the Food-Energy-Water (FEW) Nexus. This approach optimizes resource use and enhances farm-level sustainability, effectively demonstrating ML’s profound potential to address intricate sustainability challenges across diverse energy-related systems.

1.3 Machine Learning in Power System of Frequency Stability Assessment and Unit Commitment

Inverter-based resources (IBRs) deployed for energy decarbonization in power systems pose challenges to stabilizing system frequency due to the absence of conventional rotational inertia support^{21;22}. Variability and uncertainty in IBR power outputs exacerbate power imbalances, leading to larger frequency excursions and increased resource requirements for frequency regulation^{23;24}. Replacement of conventional generators with IBR reduces system inertia and the reliability of frequency regulation. Consequently, power systems with high renewable penetration face a lower frequency nadir, risking under-frequency load shedding (UFLS)^{25;26} after generation tripping. This urgent issue is particularly critical for island systems without interconnection. Some system operators in the United Kingdom adopt a minimum inertia constraint approach²⁷, used in power systems with significant penetration of

IBR, which has the drawback of increasing cost. Maintaining conventional generators online solely for inertia support incurs additional expenses for maintenance, fuel, and operations. As the IBR share increases, reliance on costly conventional generators for the inertia response grows, posing financial challenges and questioning the long-term feasibility of this approach.

To maintain frequency stability in the system operation stage, researchers integrate the frequency nadir constraint into a traditional generation scheduling problem, such as unit commitment (UC) and economic dispatch (ED). A comprehensive literature review, presented in Table 3.1, reveals that the frequency nadir constraints can be derived from two approaches: the analytical model method^{21;28–41} and the data-driven method^{42–49}. However, incorporating the frequency nadir constraint into existing Mixed Integer Linear Programming (MILP)-based UC problem through either of these methods, particularly the data-driven approach, can increase the Complexity of Adaption (CoA), leading to a larger number of decision variables and parameters during linearization or master-sub programming, and hence can increase computational complexity. Categorically, we conclude three challenges to incorporating frequency nadir constraints to MILP-based UC problems in the current situation:

First, the non-linearity of the nadir constraints makes it challenging to incorporate them into a MILP-based UC model. Many studies use Piece Wise Linear (PWL)^{21;28;29;31–37;39} and Benders’ Decomposition (BD)^{38;42–45} methods to integrate the constraint. Piecewise linearization allows for the representation of nonlinear functions with a series of linear segments, offering more flexibility in modeling complex relationships. However, using PWL requires careful selection of breakpoints and can lead to a high approximate error, while using BD could result in infeasible solutions in subproblems. These methods can increase CoA in the MILP-based UC problem by introducing more decision variables and parameters during linearization. A frequency nadir constraint has a small CoA in MILP when it is closer to a linear equation, while a nonlinear frequency nadir constraint leads to a higher CoA in MILP. Moreover, those constraints with higher CoA result in higher computational complexity during linearization.

Second, accurately estimating the frequency nadir constraint is essential to maintain

adequate inertia and headroom for fast frequency response, as well as to ensure that the system frequency nadir remains within acceptable limits⁵⁰. However, accurately estimating this constraint is challenging due to various factors such as inertia, synchronous reactance, time constants, droop settings of conventional generators and IBRs, and the number of online units. Analytical models may not capture all the nonlinear relationships involved, further complicating the estimation process. Fortunately, data-driven methods^{42–49} can offer accurate estimates by treating the frequency nadir and the input data as black boxes, considering the entire system and its output. These data-driven approaches provide a viable alternative to analytical models, leveraging the advantages of capturing complex nonlinear relationships. Therefore, this work embraces the benefits of data-driven models to estimate the frequency nadir instead of relying solely on analytical models.

Third, validation and testing of the frequency nadir constraint is crucial to ensure its accuracy, reliability, and effectiveness in the MILP solution. Long-term simulations with realistic time series data and real-world systems are necessary for comprehensive validation under various system operation scenarios. Most studies use a simulation of one day (24 hours)^{21;28;29;31;34–39;42–45}. The volatile operation conditions of high-IBR penetration power systems, including the impact of large-capacity PV-battery systems and their operating hours, cannot be fully satisfied with his short-term validation. Therefore, a "long-term" simulation with attentive analysis is crucial to confirm the frequency nadir constraint's effectiveness while considering dynamic performance.

1.4 Machine Learning in Buildings Energy Management

Buildings in the United States account for approximately 36% of total energy consumption, with HVAC systems being a significant contributor, particularly during hot summers⁵². Smart homes using IoT technologies and human-centered intelligent scheduling for HVAC control is a promising solution for energy-efficient and comfortable buildings⁵³. Understand-

Table 1.1: Comparison of related works in the literature (in publication time order)

Ref.	Frequency constraint	Real sys. w/ High IBR	Approaches	Validation Period	CoA
28 /2013	Analytical-model	✓	Piece Wise Linear	24 hours	Mid
29 /2014	Analytical-model		Piece Wise Linear	6 hours	Mid
30 /2016	Analytical-model	✓	Ordinary Differential Equation	6 hours	Mid
31 /2016	Analytical-model		Piece Wise Linear	24 hours	Mid
32 /2016	Analytical-model	✓	Piece Wise Linear	146 hours	Mid
42 /2016	Data-driven		CSC-based Benders' Decomposition	24 hours	High
43 /2019	Data-driven		CSC-based Benders' Decomposition	24 hours	High
44 /2019	Data-driven		Duality-Free Benders' Decomposition	24 hours	High
33 /2019	Analytical-model	✓	Piece Wise Linear	8760 hours	Mid
21 /2019	Analytical-model		Piece Wise Linear	24 hours	Mid
34 /2019	Analytical-model	✓	Piece Wise Linear	48 hours	Mid
35 /2020	Analytical-model		Piece Wise Linear	24 hours	Mid
36;37 /2021	Analytical-model		Piece Wise Linear	24 hours	Mid
51 /2021	Analytical-model	✓	Adaptive Piece Wise Linear	24 hours	Mid
38 /2021	Analytical-model	✓	Bilinear Benders' De- composition	24 hours	High
45 /2021	Data-driven		OCT-based Benders' Decomposition	24 hours	High
39 /2022	Analytical-model		Piece Wise Linear	24 hours	Mid
40 /2022	Analytical-model	✓	Tri-level Benders' De- composition	24 hours	High
41 /2022	Analytical-model		Gauss-Seidel Itera- tions	40 seconds	High
46 /2022	Data-driven		DNN-based Binary Encoder	24 hours	High

Table 1.2: Comparison of related works in the literature (in publication time order)

⁴⁷ /2022	Data-driven		Wasserstein-Metric Ambiguity Set	24 hours	High
⁴⁸ /2022	Data-driven		SVR-based Linearization	24 hours	Mid
⁴⁹ /2022	Data-driven	✓	Logistic Regression	24 hours	Mid
This work	Data-driven	✓	Linear Regression	8760 hours	Low

ing and incorporating Occupant-Centric Control (OCC) is crucial for effective building energy management^{54;55}. Wang *et al.*⁵⁶ investigated occupancy patterns in single-person offices and proposed a probabilistic model for occupancy prediction, emphasizing the complexity and time variation of occupancy intervals. Several studies have proposed algorithms and models for predicting occupancy to improve energy consumption and occupant comfort, including Reinhart⁵⁷, Page *et al.*⁵⁸, Klein *et al.*⁵⁹, and Fabi *et al.*⁶⁰. Furthermore, integrating clothing behaviors and clothing decisions into HVAC control strategies can further enhance building energy efficacy^{61–63}. There are a lot of literature studies on the simplified model for occupant behavior model in building energy management to reduce the electricity cost while maximizing the occupant thermal comfort in residential buildings^{64;65} and commercial buildings^{66–68}. However, occupancy and clothing behavior are not the only factors that have an effect on occupants’ thermal comfort model. Metabolic rate with activity schedule is another important factor that should be conducted in the thermal comfort⁶⁹ considered in indoor temperature control. However, dynamic occupant behavior containing the occupancy, clothing adjustment, and metabolic rate involved in thermal comfort in HVAC control is still an unsolved problem⁷⁰.

On the other hand, model-based optimization can be time-consuming when dealing with large solution spaces, making it unsuitable for real-time decision-making. In recent years, Deep Reinforcement Learning (DRL), a model-free approach, has gained traction among engineers and researchers for tackling building energy management problems⁷¹. DRL-based HVAC control methods have been proposed to address challenges posed by large state-action

spaces^{18;72} and complex indoor environments^{21;73}. However, addressing the multi-zone HVAC control problem with continuous action spaces remains a challenge despite the application of DRL techniques in previous studies. In real-world environments, multi-zone thermal control involves complex control agents and often necessitates a balance between competition and cooperation among these agents. Some studies have applied Multi-agent Deep Reinforcement Learning (MADRL) to multi-zone thermal control^{74;75}, while often lacking dynamic modeling. In Table 1, “Const.” signifies a consistent schedule applied throughout simulations, while “Dyn.” represents changing occupant activities influenced by factors such as time of day, weather, and personal choices. Although certain studies have integrated occupant presence^{76–78}, they do not account for metabolic rate and clothing adjustments. This gap will be addressed in Chapter 3 in occupant behavior modeling validation and it highlights the need for novel MADRL methods capable of addressing the complexity of multi-zone environments and dynamic occupant behaviors in HVAC control.

Table 1.3: Literature review of DRL in HVAC control

Reference	Whole building simulation	Multi-area thermal zone	Multi-agent approach	Occupant behavior		
				Occupancy	Metabolic	Clothing
¹⁸ /2017	✓	✓				
^{21;79} /2018-19	✓	✓		Const.		
⁸⁰ /2020	✓	✓		Const.	Const.	
⁸¹ /2021	✓	✓	✓		Const.	Const.
⁸² /2021	✓	✓		Const.		
⁸³ /2021	✓	✓				
⁷⁴ /2021	✓	✓		Const.		
⁷⁶ /2021	✓	✓	✓	Dyn.		
⁸⁴ /2022	✓	✓		Const.	Const.	Const.
⁷⁷ /2022	✓	✓	✓	Dyn.		
⁷⁵ /2022	✓	✓	✓			
⁷⁸ /2022	✓	✓	✓	Dyn.		
⁸⁵ /2022	✓			Dyn.		
⁸⁶ /2023	✓	✓	✓			
⁸⁷ /2023		✓	✓			
Chapter 3	✓	✓	✓	Dyn.	Dyn.	Dyn.

1.5 Machine Learning in Agriculture Energy Management

The pressing challenges of climate change, resource scarcity, and global population growth demand innovative approaches to sustainability within the Food-Energy-Water (FEW) Nexus. This concept emphasizes the interdependence of food, energy, and water systems, advocating for their integrated management to yield more effective outcomes than addressing them in isolation. For example, Huntington et al.⁸⁸ demonstrated the benefits of a FEW Nexus approach in rural areas, where renewable energy initiatives improved energy security. Similarly, Wang et al.⁸⁹ presented an integrated system combining solar desalination, power generation, and crop irrigation to address water and energy scarcity while supporting agricultural activities. On a community scale, Ma et al.⁹⁰ highlighted the role of spatial planning in optimizing FEW resource management, proposing a framework to incorporate FEW considerations into community planning. Proctor et al.⁹¹ examined diverse perspectives within the FEW Nexus, identifying key motivations like ecosystem health and waste management, and called for clearer definitions of research motivations to strengthen coherence in the field. However, as noted by Abdi et al.⁹², the FEW Nexus approach faces significant challenges due to increasing resource consumption, scarcity, and environmental degradation. They emphasized the need for integrated management strategies that support long-term sustainability.

Agriculture, as the sole provider of human food, relies heavily on energy-intensive processes, primarily powered by fossil fuels, which significantly contribute to greenhouse gas emissions and climate change^{93;94}. Meynard et al.⁹⁵ underscore the potential of renewable resources such as solar, wind, biomass, and small-scale hydro to mitigate environmental impacts while enhancing economic stability. Chel and Kaushik⁹⁶ further emphasize the balance required for agriculture, advocating for renewable energy systems like solar photovoltaic water pumps and greenhouse technologies to minimize non-renewable resource use and environmental harm. The advent of smart farming technologies, as discussed by Guven et al.⁹⁷,

leverages information and communication technologies to optimize agricultural processes through data-driven decisions, thus improving efficiency and conserving natural resources. Rahman et al.⁹⁸ review the application of renewable energy in agriculture, highlighting the disparities between developed and developing countries in adopting these technologies and stressing the need for incentive-based policies to support this transition. Mendigoria et al.⁹⁹ and Bardi et al.¹⁰⁰ explore the role of energy management systems and future restructuring of agriculture to utilize renewable electric power, advocating for significant transformations to achieve true sustainability. Collectively, these studies underscore the importance of integrating renewable energy into farm energy management as a critical strategy for reducing greenhouse gas emissions, promoting sustainable practices, and ensuring economic stability—thereby reinforcing the interconnected goals of the FEW Nexus.

Renewable energy is crucial for transforming agricultural practices into more sustainable and resilient systems. Edmonds et al.¹⁰¹ demonstrate the viability of utilizing renewable electricity for hydrogen production via electrolysis and nitrogen extraction through air liquefaction. These processes help reduce carbon emissions while also serving as demand-responsive loads and long-term energy storage. Expanding on this concept, Edmonds et al.^{102;103} investigate the coordinated operation of water and electricity distribution networks powered by variable renewable energy (VRE) sources. Then, Edmonds et al.¹⁰¹ demonstrate the viability of utilizing renewable electricity for hydrogen production via electrolysis and nitrogen extraction through air liquefaction, helping to reduce carbon emissions while serving as demand-responsive loads and long-term energy storage. In the context of small-scale energy systems, Liu et al.^{36;104} explore home energy management that integrates photovoltaic (PV) systems and electric vehicles using stochastic optimization techniques. In the other way, renewable energy integration in large-scale energy systems is explored in^{65;105} to enhance its efficiency and system’s flexibility. Additionally, green ammonia, produced using renewable energy sources, presents a promising solution for energy storage. This decarbonization transition, driven by green ammonia’s potential to substantially lower emissions has attracted significant attention. Müller et al.¹⁰⁶ highlight green ammonia’s role as a high-

density energy carrier with storage efficiency comparable to hydrogen. Bouaboula et al.¹⁰⁷ address the intermittency of renewable energy in green ammonia production, proposing a techno-economic model that optimizes ammonia plant operations, ensuring consistent production despite energy fluctuations. Their advanced energy management strategy (EMS) significantly improved the Haber-Bosch reactor’s load factor, cutting production costs. Supporting the feasibility of green ammonia, Sánchez et al.¹⁰⁸ evaluate ammonia-based power plants, emphasizing their efficiency and stable energy supply in decarbonized power systems. Sekhar et al.¹⁰⁹ discuss green ammonia as a hydrogen energy carrier, noting its benefits in cost, storage, and transport, and providing a roadmap for its contribution to achieving Sustainable Development Goals and energy sector decarbonization. Klyapovskiy et al.¹¹⁰ explore integrating green ammonia production into energy management systems for industrial clusters, demonstrating how coupling ammonia plants with renewable energy sources and demand-response strategies can optimize costs and maximize green hydrogen production. This integration is key to achieving sustainable energy systems and enhancing the economic viability of green ammonia.

RL¹¹¹ offers significant advancements for optimizing energy management systems across various applications. RL’s ability to adapt to dynamic environments and learn optimal policies without relying on predefined models makes it particularly suited for complex real-world scenarios. Recent studies^{77;84;112} demonstrate the potential of DRL in reducing energy costs, enhancing system flexibility, and managing complex, dynamic environments. These techniques, originally developed for small-scale energy systems, offer valuable insights for adapting DRL to the unique and fluctuating energy demands in farm operations. Recent studies have highlighted the effectiveness of DRL in improving irrigation efficiency, reducing water use, and enhancing crop yield. For instance, Ding et al.^{113;114} proposed a DRL-based irrigation control system that leverages soil moisture data and safety mechanisms to ensure optimal water usage in field tests, while Chen et al.¹¹⁵ demonstrated the utility of DRL for making irrigation decisions based on weather forecasts, achieving significant water conservation in rice cultivation. Similarly, Alibabaei et al.¹¹⁶ employed a DQN to

optimize irrigation scheduling for tomato fields, which led to increased productivity and reduced water consumption. Addressing climate change adaptation in agriculture requires a cumulative impact from individual farm-level decisions, which are often shaped by public policies like subsidies, taxes, and regulations, as well as corporate policies such as pricing strategies and supply chain alternatives. However, effective adaptation also depends on improved analysis of innovative approaches to ensure that they lead to meaningful outcomes rather than failed initiatives. Building on recent advancements in reinforcement learning, renewable energy integration, and green ammonia applications, Chapter 4 aims to bridge existing gaps in sustainable farm energy management within the Food-Energy-Water (FEW) Nexus. While prior research has highlighted the potential of renewable energy and RL techniques to enhance efficiency, reduce emissions, and optimize resources, few studies have offered a unified framework that supports both energy generation and resource management at the individual farm level, especially with considerations for green ammonia as a clean energy carrier and fertilizer.

1.6 Research Motivations

Combining ML techniques with advanced energy management across power systems, buildings, and agriculture is timely given:

- **Chapter 2 – High Renewable Penetrations in Grid:** Ensuring grid stability and reliability under low-inertia conditions without solely relying on expensive legacy solutions.
- **Chapter 3 – Dynamic Occupant Behavior in Building:** Addressing occupant-driven demand swings and comfort needs through real-time, occupant-centric controls in multi-zone buildings.
- **Chapter 4 – FEW Nexus Pressures in Agriculture:** Balancing agricultural productivity, resource conservation, and decarbonization at the farm level.

These motivations lead to a central goal: deploying data-driven, ML-based frameworks that maintain robust system performance, lower costs, and minimize environmental impact.

1.7 Research Questions

The following questions guide this dissertation:

Q 1: How can we efficiently incorporate frequency nadir constraints into a MILP-based unit commitment model for high-IBR or island power systems, ensuring stability and limiting the reliance on expensive conventional generators?

Q 2: In occupant-centric building energy management, can an advanced learning algorithm effectively capture dynamic occupant behaviors (occupancy, clothing insulation, metabolic rates) and maintain occupant thermal comfort at reduced energy costs?

Q 3: How can learning-based control be leveraged to coordinate on-site renewables, green ammonia production, and irrigation within the FEW Nexus, thereby improving farm revenue, resource efficiency, and sustainability?

1.8 Research Contributions

This dissertation addresses the above questions by making the following key contributions:

Q 1: Data-Driven Frequency Nadir-Constrained UC – A linear regression-based frequency nadir estimation model is proposed and integrated into a MILP-based UC problem. Validated over a year-long simulation on Maui Island’s high-IBR grid, our approach ensures reliable frequency response with minimal cost impacts.

These contributions are discussed in Chapter 2 and in the following articles:

X. Liu, J. Xie, X. Fang, H. Yuan, B. Wang, H. Wu, and J. Tan. "A comparison of machine learning methods for frequency nadir estimation in power systems," In

2022 IEEE Kansas Power and Energy Conference (KPEC), pages 1–5, Manhattan, KS, USA, 2022. ¹¹⁷.

X. Liu, X. Fang, N. Gao, H. Yuan, A. Hoke, H. Wu, and J. Tan. "Frequency nadir constrained unit commitment for high renewable penetration island power systems," IEEE Open Access Journal of Power and Energy, 11:141–153, 2024. ¹¹⁸

Q 2: Occupant-Centric Multi-Zone HVAC Control – A multi-agent DRL (MADRL) framework that explicitly models occupant presence, clothing, and metabolic rate is developed to coordinate heating and cooling setpoints. This occupant-centric paradigm achieves up to 51.09% cost savings over rule-based methods and 4.34% over single-agent DRL, all while preserving occupant comfort.

These contributions are discussed in Chapter 3 and in the following articles:

X. Liu, Y. Wu, H. Zhang, and H. Wu. "Hourly occupant clothing decisions in residential HVAC energy management," Journal of Building Engineering, 40:102708, 2021. ⁶²

X. Liu, Y. Wu, and H. Wu. "Enhancing HVAC energy management through multi-zone occupant-centric approach: A multi-agent deep reinforcement learning solution," Energy and Buildings, 303:113770, 2024. ¹¹⁹

Q 3: Integrated DRL for Sustainable Agriculture – Within the FEW nexus, a DRL strategy is proposed to manage renewable generation, ammonia production, and irrigation scheduling on farms. Simulation results reveal improved farm revenue, efficient resource usage, and reduced carbon footprints, underscoring ML's potential to modernize agricultural systems sustainably.

These contributions are discussed in Chapter 4 and in the following articles:

X. Liu, A. Modarresi, L. Kiboma, H. Wu, J. Symons, and M. Hill. "Budget-constrained sizing of renewable and energy storage systems for farm-scale ammonia production within the food-energy-water nexus," In 2024 IEEE Green Technologies Conference (GreenTech), pages 218–222, 2024. ¹²⁰




Figure 1.1: Structure Diagram of this Dissertation

X. Liu, L. Kiboma, V. Amanor-Boadu, A. Modarresi, M. Hill, H. Wu, "Deep Reinforcement Learning based Farm Energy Management Considering Ammonia Generation: Approach Within Food-Energy-Water Nexus," Submitted to Journal of Cleaner Production, under review.

1.9 Organization of This Dissertation

Figure 4.5 illustrates how the three domains addressed in this dissertation—power systems, buildings, and agriculture—fit into an overarching ML framework for energy management:

The remainder of the dissertation is structured as follows:

- **Chapter 1** (*Introduction*): Provides the overall context, motivation, research questions, and contributions, as presented here. In addition, surveys prior work on ML applications in energy management for power systems, buildings, and agriculture, identifying gaps addressed in this research.

- **Chapter 2** (*ML in Power Systems*): Presents the data-driven frequency nadir-constrained UC (FNC-UC) model, discussing its formulation, CoA improvements, and year-long validation in an islanded grid.
- **Chapter 3** (*ML in Occupant-Centric HVAC Control*): Introduces the occupant-centric multi-agent DRL scheme for multi-zone building energy management, describing occupant behavior modeling and results achieved in EnergyPlus simulations.
- **Chapter 4** (*ML in Agriculture and the FEW Nexus*): Proposes an integrated DRL framework for farm-level management, detailing its synergy with renewable generation, ammonia production, and irrigation scheduling.
- **Chapter 5** (*Conclusion and Future Work*): Summarizes key findings, discusses broader implications for sustainable energy management, and highlights avenues for future research.

Collectively, these chapters underscore the broad potential of ML-based solutions in resolving critical challenges of energy decarbonization and system optimization across power systems, buildings, and agricultural domains.

Chapter 2

Machine Learning for Frequency-Constrained Unit Commitment in Power Systems

Electric power systems are among the most critical infrastructures supporting modern society, underpinning essential services such as residential electricity, commercial operations, industrial activities, transportation, and agriculture. The U.S. National Academy of Engineering has recognized the electrification of society as one of the greatest engineering achievements of the 20th century¹²¹. Power systems involve a complex and integrated network of electricity generation, transmission, and distribution, which collectively facilitate efficient and reliable energy management and delivery across vast geographic areas. In recent years, the global push towards sustainability has significantly accelerated the deployment of renewable energy resources, such as wind and PV systems. According to the International Energy Agency (IEA), renewable energy generation has steadily increased, accounting for a growing proportion of total electricity production in the United States and worldwide¹²².

This rapid expansion, while beneficial for reducing carbon emissions and fostering energy independence, introduces several critical operational challenges to power systems, particu-

larly concerning system stability and reliability. Renewable energy sources are inherently intermittent and uncertain due to their dependence on weather conditions and natural variability. The resulting unpredictability complicates the accurate forecasting of generation, leading to power imbalances that must be carefully managed to maintain grid stability. One critical issue arising from the increased penetration of renewable resources is the decline in system inertia, traditionally provided by large synchronous generators in conventional power plants. High penetration of inverter-based renewable generation (e.g., wind turbines and solar PV) reduces overall inertia, potentially causing rapid and large frequency deviations in response to generation outages or sudden load changes^{24;123}. Such frequency instability can significantly increase the risk of system disturbances or even blackout events. ML techniques offer powerful solutions for managing the complexities and uncertainties associated with integrating high levels of renewable generation. In power systems, ML algorithms have been successfully employed to enhance load forecasting, renewable generation prediction, fault detection, and system reliability assessments⁹. More specifically, data-driven approaches, such as supervised learning and reinforcement learning, have demonstrated remarkable effectiveness in predicting critical system states and optimizing control actions, thus ensuring grid stability under variable renewable generation conditions^{13;16}.

To solve the challenges of incorporating the frequency nadir constraint into MILP of the unit commitment problem with efficiency and easy applicability by industry considered as shown in Table 1.1, this section proposes a novel data-driven frequency nadir-constrained unit commitment (FNC-UC) model, which imposes constraints on frequency nadir following the largest generation contingency. The main contributions can be summarized below:

1. A data-driven, linear regression-based frequency nadir estimation model is proposed that accounts for the impacts of the system inertia, the headroom for the PFR, and the largest generator's output capacity in a generation trip contingency. This model provides a highly accurate estimate of the frequency nadir in a realistic island system with a generation-tripping contingency.
2. The proposed frequency nadir model is embedded in a MILP-based unit commitment

formulation such that the system frequency nadir after the most significant generator outage is endogenously considered in the generation dispatch problem.

This chapter firstly focuses on leveraging ML techniques, particularly supervised learning, to address the frequency stability challenges posed by high renewable energy penetration. Then, this chapter discusses the concept of inertia and frequency stability in detail, highlighting the critical role inertia plays in mitigating frequency deviations and maintaining system reliability. Further sections delve into specific applications of ML to enhance the prediction and operational management of power systems under high renewable energy scenarios.

2.1 Machine Learning Prediction in Frequency Nadir

Energy management in power systems faces critical challenges due to the rapid expansion of renewable energy sources (RES) as a response to climate change, environmental concerns, and fossil fuel depletion. Although renewable integration offers significant environmental benefits, the variability and uncertainty such as solar photovoltaic (PV) and wind energy significantly impact traditional grid stability and reliability¹²⁴. Recent literature extensively explores integrating uncertain renewable energy into day-ahead power system scheduling, particularly focusing on Unit Commitment (UC) and economic dispatch problems^{102;125–128}.

For instance, Yin et al.¹²⁵ and Fang et al.¹²⁶ developed mixed-integer linear programming-based optimization frameworks for power systems with high PV penetration, significantly reducing market operational costs while incorporating PV and battery energy storage systems (BESS) for ancillary services. Similarly, Faqiry et al.¹²⁷ and Lawryn et al.¹⁰² optimized distributed renewable energy resource schedules and distribution locational marginal prices within day-ahead market frameworks. Additionally, Palani et al.¹²⁸ introduced stochastic methods addressing network-constrained unit commitment under high renewable penetration scenarios.

Despite these advancements, a major stability challenge emerges as renewable generation replaces conventional synchronous generation, consequently reducing system inertia and




Figure 2.1: Illustration of Renewable Energy of Frequency Nadir Impacts

threatening frequency stability¹²⁹. Figure 2.1 illustrates the frequency dynamics following a major generator trip, emphasizing how inadequate inertia can lead to severe frequency drops (frequency nadirs), potentially causing cascading outages and grid blackouts. Notably, most existing research on day-ahead market optimization neglects explicitly modeling frequency nadir constraints, posing substantial operational risks—particularly in isolated island systems characterized by high renewable penetration.

To address this critical gap, accurate and computationally efficient methods to estimate frequency nadir proactively are essential for power system operators. Due to increasing availability of historical and real-time operational data, machine learning (ML) techniques have emerged as practical solutions capable of high-accuracy frequency nadir prediction. Several ML techniques, including artificial neural networks (ANN), support vector regression (SVR), gradient boosting (GB), linear regression (LR), and extreme gradient boosting (XGBoost), have been successfully applied in various power system prediction applications, such as transient stability assessment^{50;130}, wind power forecasting¹³¹, temperature prediction^{36;62}, and frequency nadir estimation^{132;133}.

For instance, Su et al.⁵⁰ demonstrated the effectiveness of random forest regression to predict frequency nadir under various contingencies using extensive simulation data, whereas Liu et al.¹³³ applied polynomial fitting to estimate primary frequency response performance with promising accuracy. Chamorro et al.¹³² employed an ANN-based nonlinear autoregressive model to predict frequency nadir timing under various prediction horizons, further

highlighting ML’s flexibility and potential.

However, a comprehensive comparative evaluation of multiple ML techniques for frequency nadir prediction—particularly regarding computational efficiency, interpretability, and practical applicability within operational constraints—has not been thoroughly conducted yet. Moreover, literature still lacks clear guidelines on selecting appropriate ML methods based on practical criteria beyond mere prediction accuracy.

Therefore, this chapter explicitly addresses the following research objectives:

- Perform a rigorous comparative analysis of multiple ML methods (LR, GB, SVR, ANN, and XGBoost) for predicting frequency nadir, clearly evaluating accuracy, computational efficiency, and interpretability.
- Demonstrate the practical implementation of the most suitable ML method in a realistic day-ahead unit commitment (UC) framework by developing and integrating explicit frequency nadir constraints.

Through extensive dynamic simulations using the Multi-Timescale Integrated Dynamic and Scheduling (MIDAS) toolbox on a detailed Western Electricity Coordinating Council (WECC) 240-bus system, this chapter clearly identifies the most suitable ML method for operational deployment. Notably, this study reveals that although gradient boosting methods (GB, XGBoost) yield the highest prediction accuracy, **linear regression (LR)** offers substantial practical advantages in interpretability, computational simplicity, and ease of embedding as linear constraints in real-world day-ahead UC optimization frameworks. Consequently, LR is selected and further implemented as an operationally viable frequency nadir constraint within the subsequent subsections.

The remainder of this chapter is structured as follows: subsection 2.2 introduces the datasets and methodologies, including data preprocessing, dataset generation, and detailed descriptions of each ML technique compared. subsection 2.3 presents comparative numerical results and systematically evaluates ML method performance. subsection 2.4 transitions into practical application, detailing how the selected ML approach (LR) informs the creation of

explicit frequency nadir constraints within a unit commitment optimization model, with subsequent subsections providing comprehensive case studies and result analyses.

2.1.1 Dataset and Methodology

In this subsection, we review the five machine learning methods used and the background of the proposed work. In this study, we explore two ways to approach the dataset necessary for the machine learning methods: 1) the unit generation dataset (UGD) and 2) the system total inertia and headroom dataset (TIH). Both datasets are generated by running the PSS/E¹³⁴ on the Western Electricity Coordinating Council (WECC) 240-bus system^{135;136}. More details about the dataset generation are discussed in the next subsection. The UGD is a dataset with hourly generation of all units in the power system. The TIH dataset consists of the hourly total inertia and total headroom of the system, can be calculated as follows:

$$\begin{aligned} Inertia_t &= \sum_{i=0}^N v_{i,t} \cdot H_i \cdot Mbase_i, \forall i \\ PFR_t^{headroom} &= \sum_{i=0}^N (\overline{G_{i,t}} - G_{i,t}), \forall i \end{aligned} \quad (2.1)$$

where $Inertia_t$ and $PFR_t^{headroom}$ are the systems total inertia (MW*s) and total headroom (MW) at time t ; H_i is the inertia of unit i ; $Mbase_i$ is the rating in MVA of unit i ; $\overline{G_{i,t}}$ and $G_{i,t}$ are the maximum generation and the power generation of unit i at time t , respectively. Note that the headroom of the TIH dataset is considered in these three types of generators. The inertia of a unit exists only when the unit is on.

2.1.2 Machine Learning Methods

Linear regression

$$y_i = \beta x_i^T + \varepsilon_i \quad (2.2)$$

where x_i is the input variable, and y_i is the output variable. The slope of the line is β , and ε_i is the intercept (the value of y when $x = 0$)¹³⁷. Linear regression can be used to construct linear constraint of frequency nadir which is suitable for the day-ahead UC model.

Gradient boosting

$$F_0(x) = \arg \min_{\gamma} \sum_{i=1}^n L(y_i, \gamma) \quad (2.3)$$

$$F_m(x) = F_{m-1}(x) + \arg \min_{h_m \in H} \sum_{i=1}^n L(y_i, F_{m-1}(x_i) + h_m(x_i)) \quad (2.4)$$

where F_i , h_i are the objective learning function and the generic function at stage m , and L is the loss function. Parameters γ , x , and y are the initialization, input variable, and output variable, respectively¹³⁸.

Support Vector Regression

$$\min \frac{1}{2} \|w\|^2 + C \sum_{i=1}^N (\xi_i + \xi_i^*) \quad (2.5)$$

$$y_i - wx_i - b \leq \epsilon + \xi_i \quad (2.6)$$

$$-y_i + wx_i + b \leq \epsilon + \xi_i^* \quad (2.7)$$

$$\xi_i, \xi_i^* \geq 0 \quad (2.8)$$

where w refers to the hyperplane; ϵ is the boundary line; ξ_i and ξ_i^* are the upper-bound error and lower-bound error, respectively; x_i is the input variable; and y_i is the output variable. The objective of the SVR is to minimize the error in Equation (2.5) in the hyperplane, which maximizes the margin as constraints (2.6)–(2.8)¹³⁹.

Artificial Neural Network

$$W = W - \alpha \nabla J(W) \quad (2.9)$$




Figure 2.2: Simplified ANN model

$$\nabla J(W) = \left(\frac{\partial J}{\partial w_1}, \frac{\partial J}{\partial w_2}, \dots, \frac{\partial J}{\partial w_N} \right) \quad (2.10)$$

where W is the updated parameters set of the regression function, and J is the loss function, which is calculated by mean square error (MSE). Fig. 2.2 shows the graphic architecture of a fully connected neural network with two input neurals, four hidden neurals, and one output neural. The ANN is built based on the gradient descent with activation functions¹⁴⁰.

XGBoost XGBoost stands for Extreme Gradient Boosting, which is built based on gradient boosting but has more accurate approximations by enabling the regularization and second-order gradients¹⁴¹. The objective function of the XGBoost is as follows:

$$L^t = \sum_{i=1}^n l(\hat{y}_i, y_i) + \sum_k \Omega(f_k) \quad (2.11)$$

$$\Omega(f) = \gamma T + \frac{1}{2} \lambda \|w\|^2 \quad (2.12)$$

where l is the loss function that measures the difference between the prediction, $\hat{y}_i$, and the read value, y_i . The second term, Ω , penalizes the complexity of the model (i.e., the regression tree functions), which helps to smooth the final learnt weights to avoid overfitting¹⁴¹. Parameters γ , T , and w are the coefficient, the number of leaves, and the leaf weights for the corresponding regression tree, respectively.




Figure 2.3: Flowchart of generating UGD and TIH datasets




Figure 2.4: Frequency nadir versus total inertia and headroom

Table 2.1: Hyper-parameters of ANN

	activation	L2 regularization	# hidden units	optimizer
UGD	logistic	0.001	50	Adam
TIH	tanh	5e-05	50	SGD

2.1.3 Preliminary Results

The datasets of the proposed study were collected from a 1-year dynamic simulation in the WECC 240-bus system with high RES penetration. The dynamic simulation is conducted in PSS/E. The procedure of obtaining the UGD and TIH datasets is shown in Fig. 2.3. The hourly generation dispatch is obtained from MIDAS which is an open-source power-system-simulation tool for simulating the electricity market that considers multi-timescale operations with an interface for dynamic simulations. It can address the challenge of operating the grid with extremely high renewable penetrations by bridging the modeling the gaps of different timescales between economics, reliability, and stability of grid operation¹⁴².

Once the generation dispatch is given from the MIDAS scheduling, the PSS/E-based dynamic simulation is processed following the trip of the largest generator. Then, the frequency nadir points are recorded. In this procedure, the frequency nadir in the TIH dataset can be visualized as a 3D plot in Fig. 2.4. Figure 2.4 depicts three dimensions—inertia, headroom, and frequency nadir—and the z-axis is the frequency nadir point corresponding to the total inertia and the headroom calculated based on the generator dispatch point according to Equation (2.1). The outliers in Fig. 2.4 are caused by the various tripped largest generators because the different generator has different effects on frequency. In addition, it is seen that more inertia and headroom can contribute higher frequency nadir.

We randomly split the dataset into 80% and 20% for the training and testing, respectively. The hyper-parameters and architecture information of the ANN are listed in Table 2.1. Those parameters were finely tuned for an optimal MSE. In addition to the MSE, other metrics—including mean absolute error (MAE), root mean square error (RMSE), and maximum error (ME)—were used on both the UGD and the TIH datasets for training and testing evaluations in the remaining part of this subsection.

Table 2.2: Evaluation metrics of UGD training

Methods	MAE (Hz)	MSE (Hz)	RMSE (Hz)	ME (Hz)
LR	0.01014	0.00027	0.01646	0.16097
GB	0.00231	0.00001	0.00295	0.02487
SVR	0.00461	0.00250	0.04852	0.10047
ANN	0.02741	0.00161	0.04068	0.27947
XGBoost	0.00366	0.00003	0.00557	0.11243

Table 2.3: Evaluation metrics of UGD testing

Methods	MAE (Hz)	MSE (Hz)	RMSE (Hz)	ME (Hz)
LR	0.01028	0.00035	0.01813	0.17674
GB	0.00345	0.00005	0.00682	0.10395
SVR	0.04674	0.00256	0.04886	0.1351
ANN	0.02494	0.00149	0.03675	0.30963
XGBoost	0.00428	0.00006	0.00762	0.10137

The UGD training results are given in Table 2.2. As shown, the gradient boosting has the best training performance among all five machine learning methods, and the performance of the XGBoost is second. Note that a maximum error of 0.02487 from the gradient boosting indicates that the prediction error of the gradient boosting is no greater than 0.02487 Hz. The UGD testing results are shown in Table 2.3. Analogously, the gradient boosting has the best prediction performance in terms of MAE, MSE, and RMSE; however, the maximum error of the XGBoost is slightly better than that of the gradient boosting. Additionally, the SVR leads to a smaller maximum error but larger MSE and RMSE than the ANN.

The TIH training results are provided in Table 2.4. Different from the UGD dataset, the XGBoost has the best MAE, MSE, and RMSE among all five methods, whereas the gradient boosting outperforms the other methods in terms of maximum error. In addition, the evaluation metrics for both the gradient boosting and the XGBoost are quite similar to the TIH training dataset. Table 2.5 lists the TIH testing results of all the methods. Again, the XGBoost leads to the best performance; however, the ANN’s performance in this dataset is worse than that of the UGD dataset, whereas the other methods perform quite similarly between the two datasets.

Table 2.4: Evaluation metrics of TIH training

Methods	MAE (Hz)	MSE (Hz)	RMSE (Hz)	ME (Hz)
LR	0.01023	0.00028	0.0167	0.16213
GB	0.00461	0.00006	0.00753	0.07957
SVR	0.05291	0.00309	0.05557	0.10048
ANN	0.03251	0.02101	0.14496	3.87351
XGBoost	0.00403	0.00004	0.00623	0.10039

Table 2.5: Evaluation metrics of TIH testing

Methods	MAE (Hz)	MSE (Hz)	RMSE (Hz)	ME (Hz)
LR	0.01056	0.00030	0.01733	0.17674
GB	0.00499	0.00007	0.00861	0.09003
SVR	0.05383	0.00318	0.05636	0.11954
ANN	0.04853	0.08453	0.29074	4.00111
XGBoost	0.00469	0.00006	0.00781	0.07819




Figure 2.5: Error distribution of UGD (top) and TIH (bottom) testing

Figure 2.5 shows the error distribution for the (top) UGD and (bottom) TIH datasets. In general, the gradient boosting achieves the best error distribution in both datasets, whereas the XGBoost performs nearly as well. In contrast, the ANN and the SVR are not suitable for the proposed work because of the high variance distribution. In addition, both methods perform worse in the TIH testing dataset. Note that the performance of the linear regression ranks third in both datasets.

2.1.4 Preliminary Summary

In this section, we propose to use a machine-learning method, rather than dynamic simulation, to predict the frequency nadir of the largest N-1 contingency analysis based on steady-state scheduling results. Five methods—linear regression, gradient boosting, SVR, ANN, and XGBoost—are compared by testing the reduced WECC 240-bus system with high RES penetration. Results show that these methods can achieve high prediction accuracy in general. The prediction performance of the gradient boosting and the XGBoost are the best. In addition, both gradient boosting and XGBoost achieve high accuracy in the two methods’ datasets which are proposed in this paper 1) the unit generation dataset; and 2) the system total inertia and headroom dataset. In the next subsection, we will propose a machine learning-based security-constrained unit commitment model and analyze the dynamic simulation of the proposed work.

2.2 Unit Commitment Model with Frequency Nadir Constraint

In the previous section, we demonstrated that supervised machine learning methods, specifically regression-based models, can effectively and accurately predict the frequency nadir following a major generation outage. Building upon this high prediction accuracy, the next critical step is integrating these predictive models directly into power system operations to enhance reliability and stability. In this section, we present a detailed framework to achieve

this integration through a novel frequency nadir-constrained unit commitment formulation.

To effectively implement the machine learning-based frequency nadir prediction within power system operation, the following two objectives are defined:

- We propose a frequency nadir estimation model embedded within a Mixed-Integer Linear Programming (MILP)-based Unit Commitment (UC) framework. The proposed approach directly incorporates the system frequency nadir—estimated via supervised learning methods—into the generation scheduling decisions, specifically accounting for the largest potential generator outage. As illustrated in Table 3.1, the proposed FNC-UC formulation achieves a notably lower Complexity of Adaptation (CoA) compared to existing methods, significantly improving computational efficiency, scalability, and practicality for real-world industrial applications.
- To rigorously validate the proposed FNC-UC model, we utilize a real-world power system characterized by high inverter-based resource (IBR) penetration, specifically the Maui Island power grid. This system provides a realistic testbed featuring detailed load demand profiles and renewable generation scenarios over an extensive simulation period. Through comprehensive validation, we aim to demonstrate not only the technical reliability and operational effectiveness of the proposed model but also its economic benefits and the overall enhancement of system stability and performance.

In this section, we introduce the detailed mathematical formulation of the unit commitment problem incorporating the proposed frequency nadir constraint derived from supervised machine learning models. Additionally, we present the inclusion of a minimum inertia constraint, which, alongside traditional unit commitment constraints, ensures enhanced frequency stability under high renewable energy penetration scenarios.

2.2.1 Data Driven-based Frequency Nadir Estimation

In our previous subsection¹¹⁷, we evaluated five data-driven methods (LR, GB, SVR, ANN, and XGBoost) to predict frequency nadir in the power system with high renewable

penetration levels. The training and testing dataset is Total Inertia Headroom (TIH), which was generated using simulations with the Multi-Timescale Integrated Dynamic and Scheduling (MIDAS)^{143;144} toolbox. The attributes in the TIH dataset include the system-wide inertia and the PFR headroom. The results showed that the five methods had a high prediction accuracy. In particular, linear regression ranked higher than ANN and SVR, as shown in Fig. 2.5. Although LR may not have the best error performance, it can still provide reasonable accuracy after considering CoA. The hourly total inertia and total headroom can be calculated as Equation (2.1), inspired by^{145–147}, which is typically used to analyze the frequency nadir of connected machines in the entire system, the swing equation in the DC model can be simplified as

$$Inertia \cdot 2\partial\Delta f(t)\frac{\partial}{\partial t} = PFR^{headroom} - \Delta P_L \quad (2.13)$$

where ΔP_L is the power outage, $\Delta f(t)$ is the frequency deviation function of t . Therefore, it is significant to consider potential power outages as an element in the frequency nadir constraint to ensure the power system maintains the frequency above the threshold of the UFLS level. Most studies^{32;33;146;147} use the largest generation outage for this P_L .

Linear regression is one of the most common techniques in machine learning¹⁴⁰. Extensive literature^{123;148} shows that linear regression performs well in forecasting, prediction, and classification applications. It can provide stable, robust, and accurate prediction results. The goal of linear regression is to build a model whose output scalar, $y \in \mathbb{R}$, is a linear function of the input vector, $x \in \mathbb{R}^n$, by solving a regression problem. Let $\hat{y}$ be the prediction value and y be the actual value, so the equation of the linear regression should be defined as $\hat{y} = \mathbf{w}^T \mathbf{x} + b$, where $\mathbf{w}^T \in \mathbb{R}^n$ is the vector of parameters; b is an intercept term. Note that the parameters are values that control the prediction of the system. In other words, the output, $\hat{y}$, is the sum of all the weighted input vectors, $\mathbf{x}$. The more significant weight, w^i , contributes more substantially to the prediction value. The training process is listed as

follows:

$$\begin{aligned} \frac{1}{n} \nabla_w \left[\sum_{i=1}^n (y_i^{train} - \hat{y}_i^{train})^2 \right] &= 0 \\ \nabla_{\mathbf{w}} (\mathbf{X}^{train} \mathbf{w} - \mathbf{y}^{train})^T (\mathbf{X}^{train} \mathbf{w} - \mathbf{y}^{train}) &= 0 \end{aligned} \quad (2.14)$$

In a nutshell, considering an $N-1$ largest generation contingency (tripping the largest generation), the inertia of the system, the amount of headroom, and the largest power generation are the most related variables in the linear regression equation to estimate the frequency nadir of the system. The adopted formulation from the linear regression to this study is the linear regression-based nadir constraint in Equation (2.15), as shown:

$$f_{nadir} = a * Inertia + b * PFR_{headroom} + c * P_{max} + d \quad (2.15)$$

where f_{nadir} is the frequency nadir (Hz); P_{max} is the system's largest generation outage capacity (MW); and a , b , c and d are the linear coefficients obtained from linear regression. In our previous section ¹¹⁷, we analyzed the impact of inertia and PFR headroom on frequency nadir in power systems. Traditional systems exhibit a strong correlation with inertia and a weak correlation with PFR due to its slower governor model limitations, necessitating separate consideration of inertia and PFR headroom. In contrast, modern systems with high renewable penetration and integrated BESS present different dynamics, where BESS contributes to a rapid PFR. Thus, Eq. (2.15) in our study, which includes PFR headroom as a separate variable, captures the distinct influence of these factors on frequency nadir. In this way, we can obtain the frequency nadir constraint, which contains only a set of linear constraints to the MILP problem, resulting in a lower CoA than other methods, as shown in subsection III D. 2).




Figure 2.6: Illustration of the proposed frequency nadir constraint unit commitment approach

2.2.2 Proposed FNC-UC Framework

Figure 2.6 shows the flow chart to obtain the LR-based frequency nadir constraint in the flow of the red arrows. First, a conventional one-year generation scheduling simulation is performed to obtain the generation dispatch of 365 days (8,760 hours). The network-constrained unit commitment is run sequentially for each day of this year. The Maui grid system and the time series of load and renewable power output are used for the one-year simulation. Then, after obtaining the generation dispatch of each hour, a dynamic simulation is performed to simulate the system frequency response following the trip of the generator with the most significant power output at that hour. The frequency nadir of each hour is obtained from this dynamic simulation. PSSE software is used for dynamic simulation in this study¹³⁴ because system operators and researchers for dynamic simulations of power systems use it extensively. Therefore, a linear regression model is proposed to obtain the relationship of the frequency nadir to the capacity of the most significant generation power output, the system inertia, and the PFR capacity/headroom of the system. Finally, a comprehensive techno-economic analysis is proposed to evaluate the proposed constraint by the 8760-hour simulation, assessing economic, reliability, and exceptional case factors. This analysis aims to analyze the performance and outcome of the proposed data-driven frequency nadir constraint. We conducted a comprehensive assessment of machine learning methods for estimating power system frequency nadir, taking into account loading levels, contingency types, and capacities. This work involved an extensive evaluation over a one-year duration,

incorporating a wide range of load levels and capacity scenarios. Specifically, we focused on scenarios involving the contingency type of the largest generator tripping event, chosen strategically to represent a severe contingency in the power system. This deliberate emphasis on the largest generator tripping event allowed us to assess the effectiveness of the frequency nadir constraint in enhancing the system's resilience during critical incidents, establishing a robustness benchmark.

Objective Function

The objective of the FNC-UC problem includes the operating cost of traditional units—represented by their generation costs associated with their startup and shutdown costs—as well as the shortage penalties for the energy, regulation, and PFR services, as follows:

$$\min \sum_{t \in T} \sum_{i \in g} (SU_i u_{i,t} + SD_i w_{i,t} + cp_{i,t}) + LP * \Delta D_t + PFP * \Delta PFR_t + RP * (\Delta Reg_{u,t} + \Delta Reg_{d,t}) \quad (2.16)$$

$$cp_{i,t} = a_i v_{i,t} + b_i G_{i,t} + c_{i,t} G_{i,t}^2 \quad (2.17)$$

Note that the production cost, $cp_{i,t}$, of the traditional thermal unit shown in (2.16), can be approximated by a piecewise linear function from its quadratic production cost curve. In this model, we assume that the operational energy price of PV is 0. For ancillary services, the bidding prices are zero because the unit commitment model co-optimizes the energy and ancillary services, so the opportunity cost of providing ancillary services will be respected.

Constraints for the Single Unit

The constraints for traditional thermal units are similar to those in¹⁴⁹ and are presented as follows for completeness. We also consider the startup and shutdown trajectories of conventional generators in (2.20) and (2.21). In addition, (2.22)–(2.24) show the ramping

rate constraint for units, and (2.25)–(2.29) impose limitations for ancillary services.

$$u_{i,t} + w_{i,t} \leq 1 \quad (2.18)$$

$$v_{i,t} - v_{i,t-1} \leq u_{i,t} - w_{i,t} \quad (2.19)$$

$$\sum_{\tau=t-T_{i,MinUp}+1}^t u_{i,\tau} \leq v_{i,t} \quad (2.20)$$

$$\sum_{\tau=t-T_{i,MinDn}+1}^t w_{i,\tau} \leq 1 - v_{i,t} \quad (2.21)$$

$$\overline{G_{i,t}} - G_{i,t-1} \leq R_i^U v_{i,t-1} + R_i^{SU} u_{i,t} \quad (2.22)$$

$$G_{i,t-1} - G_{i,t} \leq R_i^D v_{i,t} + R_i^{SD} w_{i,t} \quad (2.23)$$

$$\overline{G_{i,t}} \leq G_{i,t}^{max} v_{i,t} \quad (2.24)$$

$$G_{i,t} + Reg_{u,i,t} \leq \overline{G_{i,t+1}} \quad (2.25)$$

$$G_{i,t} + PFR_{i,t} \leq \overline{G_{i,t+1}} \quad (2.26)$$

$$G_{i,t} - Reg_{d,i,t} \geq G_{i,t+1}^{min} v_{i,t+1} \quad (2.27)$$

$$G_{i,t} + Reg_{u,i,t} - G_{i,t-1} \leq R_i^U \quad (2.28)$$

$$G_{i,t-1} - (G_{i,t} - Reg_{d,i,t}) \leq R_i^D \quad (2.29)$$

$$v_{i,t}, u_{i,t}, w_{i,t} \in \{0, 1\} \quad (2.30)$$

System-Wide Constraints

The system constraints include the energy balance constraint for every time interval, the system regulation reserve, the PFR, and the transmission constraints, as follows:

$$\sum_{i \in \mathbf{g}} (G_{i,t} + P_{p,t}) - \sum_{b \in \mathbf{B}} D_{b,t} = 0 \quad (2.31)$$

$$D_{b,t} = \overline{D_{b,t}} - \Delta D_{b,t} \quad (2.32)$$

$$\Delta D_t = \sum_b \Delta D_{b,t} \quad (2.33)$$

$$\sum_{i \in g} PFR_{i,t} + \Delta PFR_t \geq PFR_t^r \quad (2.34)$$

$$\sum_{i \in g} Reg_{u,i,t} + \Delta Reg_{u,t} \geq Reg_{u,t}^r \quad (2.35)$$

$$\sum_{i \in g} Reg_{d,i,t} + \Delta Reg_{d,t} \geq Reg_{d,t}^r \quad (2.36)$$

$$\begin{aligned} -Limit_l &\leq \sum_{i \in Lg} GSF_{l-i}(G_{i,t} + P_{p,t}) \\ &- \sum_{b \in Lb} GSF_{l-b} D_{b,t} \leq Limit_l \end{aligned} \quad (2.37)$$

Frequency Nadir Constraints

*1) Minimal Inertia Constraint

Keeping inertia above a certain level is an intuitive method to maintain the frequency nadir for power systems with high penetration of renewable generation¹⁵⁰; therefore, some power system operators in the United Kingdom deploy a straightforward approach to maintain the stable system frequency response by imposing a minimum inertia requirement constraint in system operation²⁷. Because the inertia of conventional generation units is a static parameter, they can provide inertia response when they are online; therefore, the system minimum inertia requirement constraint can be directly formulated as follows:

$$\sum_{i=1}^N H_i \geq H_{min} \quad (2.38)$$

where H_i is the inertia of the i -th generation unit, and H_{min} is the predefined minimum inertia level that the system must maintain. Note that this study has no inertia response from PV and batteries.

*2) Frequency Nadir Constraint

From the linear regression model, (2.15), presented in the previous subsection, a linear constraint representing the frequency nadir requirement can be added to the unit commitment model. The linear constraint of the frequency nadir is shown in (2.39):

$$a * Inertia + b * PFR_{headroom} + c * P_{max} + d \geq f_{threshold} \quad (2.39)$$

Table 2.6 presents the values of the parameters a , b , c , and d for the Maui Island system. Positive values for a and b indicate that inertia and PFR headroom positively impact the frequency nadir. Conversely, the largest generation output capacity negatively affects the frequency of nadir. The parameter c plays a significant role in the linear regression-based nadir constraint. The threshold for $f_{threshold}$ in the island system is set above the first level of the UFLS scheme (assumed to be 59 Hz). Figure 2.7 illustrates the flowchart of the FNC-UC framework, with the linear regression model establishing the relationship between the frequency nadir and the significant generation capacity, system inertia, and PFR headroom (as in Eq. (2.15)). While this study focuses on an island system, it's important to note that the proposed framework is adaptable to various power system cases, including both island and traditional systems. Island systems often exhibit weaker characteristics due to limited external power support compared to traditional transmission grids. Future research could explore and evaluate the framework's applicability to different power system types, considering their unique characteristics, impacts, and efficiency. This adaptability broadens the framework's utility and contributes to a more comprehensive understanding of its potential across diverse power system scenarios. Further analysis of these parameters is presented in Section 2.3.

Table 2.6: PARAMETER VALUES OF THE LINEAR REGRESSION FREQUENCY NADIR CONSTRAINT FOR THE ISLAND SYSTEM STUDIED

Coefficients	Value
Inertia (a)	0.000847
PFR headroom total (b)	0.006293
Largest Pgen (c)	-0.04737
Intercept (d)	59.86421




Figure 2.7: Flowchart of the FNC-UC model.

2.3 Experiment and Case Studies

After reviewing the ML application in power system frequency nadir prediction and FNC-UC framework, this section mainly demonstrates the framework and simulation results. To illustrate the effectiveness of the proposed FNC-UC model, the Maui Island system is used in this study. Maui grid operators are integrating several large hybrid PV and battery power plants into their systems. The stability of the system will be simulated at high penetration levels of renewable generation, especially PV.

The Maui Island grid case with two PV-BESS plant projects is selected as our study case. It includes 23 thermal units with capacities ranging from 2.75 MW to 21.45 MW. Four wind power plants are modeled. Two existing utility-scale PV plants, with 2.87 MW each, are considered. The Maui grid model includes 61 aggregated distributed PV (DPV) units with a capacity of more than 140 MW. The DPV is not dispatchable, so the operator cannot curtail its power output. In contrast, the utility-scale PV is dispatchable, which means that its power output can be curtailed in real time; thus, the operator can curtail utility-scale PV but not DPV. The model has three large PV-BESS plants, representing the plant currently under development. The parameters of the three PV-BESS plants are summarized in Table 2.7. The PV-BESS plant model in our study represents a hybrid power plant where PV and BESS are integrated as a single unit within the power system. Specifically, we maintain the SOC at 50% by the end of each midnight, ensuring that it starts and ends each day at this level. This setup allows for flexibility in charging and discharging during the transition periods each day. Both the scheduling and dynamic simulations (in PSSE) incorporate the consistent PV-BESS model, reflecting realistic operational constraints and behaviors of the PV-BESS system. This approach ensures that SOC management is consistently represented throughout the study, providing insights into its impact on system dynamics and reliability. The generation mix is shown in Fig. 2.8. The 80-MW capacity of the utility PV includes the three PV-BESS plants. The yearly system load curve with hourly granularity is shown in Fig. 2.9. The system peak load is approximately 200 MW. The one-year simulation inherently encompasses various conditions, capturing a wide range of operational scenarios

Table 2.7: PARAMETERS OF THE PV-BESS PLANTS

	Plant 1	Plant 2
Generator type	PV+BESS	PV+BESS
Rating	60 MW/240 MWh (split into two 30-MW/120-MWh units)	15 MW/60 MWh




Figure 2.8: Generation capacity mix in the island system.

and their respective uncertainties.

The resolution of time in the FNC-UC is 1 hour, and the time span planned for the day in the FNC-UC is 24 hours. Regulation-down/-up requirements are assumed to be 3% of the system load. The PFR requirement is 4.75% of the system load. These numbers are chosen on the basis of the system's dynamic simulation to maintain reliability. A detailed model of how to decide on the ancillary service requirements¹⁴⁹ is beyond the scope of this paper.

To evaluate the performance of the proposed FNC-UC model, three different cases are simulated to demonstrate the impacts of the frequency nadir constraints on the system generation scheduling and the system stability.

1. Case 1 (Baseline Case) is without a frequency nadir constraint.
2. Case 2 (Minimum Inertia Case) is with a minimum inertia requirement constraint




Figure 2.9: Load curve in the simulated cases.

(2.38). The operator historically used 350 MVA*s as the minimum inertia requirement in this island system. This inertia level was obtained from operational experience to limit the rate of change of frequency after the generation trip.

3. Case 3 (Nadir Constraint Case) is with a linear regression-based frequency nadir constraint (2.39) but without the minimum inertia requirement constraint (2.38). The frequency nadir is chosen to be greater than 59 Hz, which is Maui grid's first UFLS frequency threshold.

In addition, the exceptional simulation results from those three cases are further analyzed in Subsubsection IV.C.

Impacts of Frequency Nadir Constraints on Generation Scheduling

The system daily generation costs versus the day of the year in the three cases are depicted in Fig. 2.10. The system generation cost with the minimum inertia constraint in




Figure 2.10: Annual generation cost of the three cases.

Case 2 is higher than in the other two cases. Fig. 2.11 shows that the addition of the linear nadir constraint trivially increases the generation cost because the production cost curves of the Baseline Case and the linear regression-based Nadir Constraint Case are very close to each other; however, the constraint of the minimum inertia requirement significantly impacts the distribution of the daily generation costs. The minimum daily generation cost for the year with the minimum inertia requirement constraint (Case 2) is much higher than in the other two cases. Note that there is no comparably low cost (less than \$250k) in the minimum inertia constraint because it has the restriction of maintaining enough inertia whether or not it is needed for that hour. Requiring a large amount of inertia leads to the introduction of more conventional generators online, resulting in a more expensive operating strategy. The higher system generation cost can be explained by investigating the detailed generation scheduling results in the three cases.

A five-day generation dispatch example for the three cases is shown in Fig. 2.12 to clarify the differences in generation scheduling. Purple and red represent the discharging and charging power of the BESS. Gray represents synchronous generation; orange and light blue represent PV and wind power, respectively. The red dashed line shows the lowest level of synchronous generation for each case. The minimum inertia requirement constraint maintains a higher inertia by bringing more synchronous generators online during the PV




Figure 2.11: The generation cost duration of the three cases.

Table 2.8: ONE-YEAR GENERATION COSTS IN THE THREE CASES

	Baseline	Minimum Inertia	In- Nadir Constr.
Total cost	\$112,997,029	\$126,814,813	\$113,864,047
% increase	—	12.23%	0.77%

hours. Additionally, the green circles highlight the difference between the Baseline Case and the Nadir Constraint Case. The generation dispatches of most hours in these cases are similar, except when the BESS is discharging and PV is starting to ramp up. The BESS usually discharges during the time without PV generation while charging during the period with PV generation, as shown in Fig. 2.12. In addition, the BESS has a better capability of providing PFR, so allocating enough headroom for the BESS and committing more conventional generators helps maintain the frequency nadir of the power system. It can also be observed that this difference in the BESS scheduling leads to different wind and solar profiles because the charging of the BESS can impact the wind and solar curtailment in this island system with a high penetration level of hybrid power plants.

The generation costs of the three cases are listed in Table 2.8. The cost of one-year generation in the baseline case is \$112,997,029. Adding the frequency nadir constraint costs 0.77% more than the baseline, while adding the minimum inertia requirement constraint costs 12.23% more than the baseline. The results in Figs. 2.10–2.12 and Table 2.8 demon-




Figure 2.12: The five-day generation dispatch examples of the three cases.

strate that the inclusion of the linear nadir constraint does not significantly increase the system generation cost. In contrast, the minimum-inertia requirement constraint significantly increases the generation operation cost. For the duration of this one-year simulation, a consistent fuel cost was maintained in the cost analysis. This methodological decision was crucial in ensuring the comparability of results. Moreover, the costs associated with the minimum inertia case were consistently higher than those in the other scenarios, a trend that held true regardless of supply curve variations while considering that fuel costs are the driver of generation costs. By maintaining constant fuel cost assumptions, our study provides a clear comparison of the economic impacts associated with those three cases.

Implication of Frequency Nadir Constraint on Stability

This subsection demonstrates the efficacy of the proposed linear regression-based nadir constraint and the minimum-inertia requirement constraint in maintaining the frequency nadir. A PSSE validation and an analysis of the simulation results are shown in this subsection. Figure 2.13 shows the frequency nadir results of three cases over one year from the PSSE dynamic simulation. Note that the dynamic simulations in PSSE of the largest generator contingency (N-1) are built for each hour's generation schedule over one year (total of 8,760 snapshots). For each hour, the generator with the most significant power level is identified and tripped in the dynamic simulation. Figure 2.13 and Figure 2.14 demonstrate that adding a frequency nadir constraint, such as the linear regression-based nadir constraint or the minimum inertia constraint, can improve the system frequency nadir. Without these frequency nadir constraints, there is a risk that the frequency nadir is lower than the predefined 59-Hz UFLS threshold if the largest generator trips. Also, note that the minimum inertia constraint cannot guarantee that the frequency nadir of all hours is higher than 59 Hz. The system will still have a risk of a low-frequency nadir even with this minimum inertia requirement. On the contrary, with the proposed linear regression-based nadir constraint, the probability that the frequency nadir is lower than 59 Hz can be significantly reduced, as shown in Table 2.9. This demonstrates that the proposed nadir constraint




Figure 2.13: Frequency nadir for each hour of the year for the N-1 event.

performs better in this system than the minimum inertia requirement constraint previously considered by the operator.

The subplot of Fig. 2.13 shows the results for April. Both frequency constraints (Nadir Constraint Case and Minimum Inertia Case) help the system maintain a higher nadir than the Baseline Case; however, there are some green points (Minimum Inertia Case) approaching 60 Hz, and some orange points (Nadir Constraint Case) are below the corresponding blue points (Baseline Case). A detailed analysis of those situations appears in the following subsection.

A histogram of the simulation results reveals differences in distribution, as shown in Fig. 2.14. In the critical range of less than 59 Hz, the Baseline Case has 159 counts, the Minimum Inertia Case has 50 points, and the Nadir Constraint Case has 18 points, respectively. In other words, the Nadir Constraint Case performs best to improve the stability of the power system and reduce the risk of load shedding. The Minimum Inertia Case results in more nadirs in the range of [59.8, 60]; this will be discussed in the next subsection.

Table 2.9 shows the numbers of total hours and violation hours (in which the frequency < 59 Hz) as well as the percentage of hours that result in violations in the three cases. The




Figure 2.14: Frequency nadir distribution of the three cases.

Table 2.9: PROBABILITIES OF FREQUENCY NADIR LOWER THAN 59 HZ IN THE THREE CASES

	Baseline	Min. Inertia	Nadir Const.
Total hours	8760	8760	8760
Violation hrs	159	50	18
Probability	1.8%	0.57%	0.20%

Baseline Case has the highest probability of a frequency nadir lower than 59 Hz among the three cases, at 1.8%. The minimum inertia case has a probability of 0. 57%, and the linear regression-based Nadir constraint case has a probability of 0. 2% of frequency violation of the nadir in the most severe dynamic simulation of N-1 contingency.

Figure 2.15 shows the number of times when the nadir is less than 59 Hz for each hour of the day. Without a frequency nadir constraint, most hours of the day have a risk of a frequency nadir lower than 59 Hz, as shown in the Baseline Case. After adding the linear regression-based nadir constraint, there are only a few hours when the frequency has a risk of falling lower than 59 Hz. The minimum inertia constraint is also helpful in reducing this risk, but is not as beneficial as the linear regression-based nadir constraint. This figure also shows that in the Maui Island system with a high penetration level of hybrid power plants (PV plus battery energy storage), the risk of a low nadir occurs mostly during non-PV hours. The main reason is that during these non-PV hours, the BESS discharges to satisfy the load in the system; thus, the headroom of the batteries reduces during these hours. Consequently,




Figure 2.15: Hourly counts of the frequency nadir lower than 59 Hz in the three cases.

the system has a smaller PFR headroom from the batteries; therefore, these hours have a relatively high risk of a lower nadir. In contrast, during the day (PV hours), the BESS is charging, which leads to ample PFR headroom from the batteries because a BESS can change from charging to discharging to provide a large capacity of PFR if needed. This is an exciting phenomenon in the island system with a high penetration of hybrid power plants. In large interconnected systems with high PV penetration levels, the risk of a low nadir usually occurs around noon, when the PV penetration level is high, and the system has a low inertia level. This demonstrates that the proposed nadir constraint will not significantly increase the system generation cost but can dramatically improve the system frequency nadir; therefore, the proposed nadir constraint is promising in the actual operation of this island system.

The one-month system scheduled inertia and the one-year inertia duration curve are shown in Fig. 2.16 and Fig. 2.17, respectively. The minimum inertia requirement frequently results in system inertia that is always higher than in the Baseline Case. In contrast, the nadir constraint does not significantly increase the system inertia. This is also shown in Table 2.10. The maximum inertia levels of the three cases are the same. The minimum inertia levels of the Baseline Case and the Nadir Constraint Case are the same. In contrast, the Minimum Inertia Case has a higher minimum inertia level, as the constraint requires. The average inertia levels of the Baseline Case and the Nadir Constraint Case are very close.




Figure 2.16: The system scheduled one-month total inertia in the three cases.




Figure 2.17: The inertia duration curves in the three cases.

Table 2.10: SYSTEM INERTIA METRICS FOR THE THREE CASES (MW*s)

	Baseline	Minimum Inertia	In- Nadir Constr.
Avg. inertia	326.4	404.6	332.3
Max. inertia	615.8	615.8	615.8
Min. inertia	91.35	351.1	91.35

Table 2.11: NADIR NEAR 60 HZ IN THE MINIMUM INERTIA CONSTRAINT CASE

ID	Max Pgen (MW)	PFR (MW)	Inertia (MW*s)	BESS HR (MW)	Est. Nadir (Hz)	Act. Nadir (Hz)
1	19.04	1.81	351.1	128.3	59.31	59.82
2	16.87	1.80	354.5	119.2	59.42	59.73
3	18.53	1.81	354.5	140.3	59.33	59.90

This is because the nadir constraint can schedule more PFR headroom instead of requiring more thermal units online for most hours, leading to a generation cost similar to the Baseline Case.

Exceptional Case Analysis

This subsection analyzes two exceptional situations mentioned in the simulation results of the previous subsections, i.e., 1) the Minimum Inertia Case’s frequency nadir close to 60 Hz and 2) a lower frequency nadir that occurred after adding the nadir constraint compared to the Baseline Case. The main reasons for two counterintuitive situations are discussed here:

Frequency Nadir of Minimum Inertia Case Is Near 60 Hz

Three examples are selected in Table 2.11 where the frequency nadir is near 60 Hz after the largest N-1 generation contingency in the Minimum Inertia Case. In all these cases, the BESS is charging, which means they consume power; therefore, once the most significant generation (much smaller than the total charging power) is tripped, the BESS can immediately stop charging (consuming power) to compensate for the tripped generation. Note that the BESS

Table 2.12: NADIR CONSTRAINT CASE HAS A LOWER NADIR THAN THE BASELINE CASE

ID	Case	Max Pgen (MW)	PFR (MW)	Inertia (MW*s)	BESS HR (MW)	Est. Nadir (Hz)	Act. Nadir (Hz)
1	Baseline	31.1	1.54	384.8	35.5	58.72	59.01
	Nadir Constr.	20	1.54	384.8	11.2	59.25	58.47
2	Baseline	31.1	1.81	453.6	25.7	58.80	59.02
	Nadir Constr.	26.2	1.96	453.6	20.2	59.01	58.92

headroom (HR) is the difference between its present power level and its maximum power output, which is greater than its rating if it is charging because the current power level is negative when charging. Observe that the highest BESS headroom results in the highest actual frequency nadir among those examples. In addition, the inertia in the system in these hours (because of the minimum inertia constraint) helps prevent a sudden drop in frequency. The BESS's high inertia and great charging power result in high-frequency nadirs in the Minimum Inertia Case.

Linear Regression-Based Nadir Constraint Case Has Lower Frequency Nadir Than Baseline

Table 2.12 shows the cases where the baseline has a higher frequency nadir than those with a frequency nadir constraint. The examples occurred on day 2, hour 19, and day 6, hour 19, respectively. These are both non-PV periods when a BESS is tripped as the largest generator. As shown in the first example, the predicted frequency nadirs from the linear regression method are 58.92 Hz and 59.25 Hz for the Baseline Case and the Nadir Constraint Case, respectively; thus, the estimated system frequency nadir is higher than the predetermined level of 59 Hz in the frequency-constrained case. However, the actual frequency nadirs simulated via PSSE are 59.01 (Baseline) and 58.47 Hz (Nadir-Constrained). In the Baseline Case, a BESS with a generation of 31.1 MW is tripped, and there is 35.5 MW of headroom in the remaining BESS; in contrast, 20 MW of generation is tripped in the frequency Nadir Constraint Case, but there is only 11.2 MW of headroom in the remaining

BESS. In these cases, the BESS is a primary provider of PFR. If the remaining BESS has a smaller PFR headroom, it leads to a lower frequency nadir. In other words, a considerable value of BESS headroom represents a better ability of the BESS to respond to the frequency drop; therefore, when the BESS is enabled to provide a fast frequency response, the risk that a BESS trip should be considered in the system operation.

Additionally, in Fig. 9, an interesting observation emerges where the minimum inertia case exhibits a lower frequency nadir than the frequency nadir constraint case. This finding suggests that while inertia is indeed a significant factor in determining power system frequency nadir, it is not the sole determinant. Our analysis reveals that increasing inertia alone is not the exclusive solution. Several factors, including reserve capacity and the influence of the largest generator, play crucial roles in shaping the frequency nadir. A direct comparison between the minimum inertia case and the frequency nadir constraint case highlights noteworthy patterns. Firstly, the minimum inertia case generally exhibits higher inertia levels compared to other scenarios. This consistent trend suggests that while inertia contributes significantly to system stability, its impact is nuanced and influenced by other system parameters. Secondly, there are instances where, despite a system with large inertia, the BESS headroom is relatively small, leading to a weaker frequency response. Additionally, limitations imposed by the frequency nadir constraint on the largest generator significantly contribute to achieving a higher frequency nadir in comparison to the minimum inertia case. Figure 12 illustrates these instances, underscoring the complexity of power system dynamics, where multiple factors, including reserve capacity, BESS headroom, and generator constraints, collectively shape the system’s frequency response rather than inertia alone.

2.4 Summary

To maintain the frequency nadir of a future grid, especially an island grid with high penetration levels of renewable generation, this paper proposes a linear regression-based fre-

quency nadir constraint in the unit commitment problem. The frequency nadir of the largest generation trip event is considered endogenously in the generation scheduling stage. Simulation results demonstrate that the frequency nadir of the system can improve significantly with this nadir constraint in a realistic island system. Therefore, the risk of a frequency nadir lower than the predetermined threshold is reduced. Furthermore, compared to the direct minimum inertia requirement constraint, the proposed linear regression-based nadir constraint trivially increases the system generation cost while achieving a better nadir response; therefore, the proposed nadir constraint can serve as a good reference for future system operation with high penetration levels of hybrid PV and battery resources.

Future work will include testing other data-driven models to improve the forecast accuracy of the proposed nadir constraint. More hybrid plant operational strategies will be investigated considering the frequency nadir risk. Another future research direction can address the complexities introduced by the virtual inertia of IBRs, specifically those with Grid-Forming Inverters in the frequency nadir constraint.

Chapter 3

Machine Learning in Occupant-centric Building Energy Management

Following our discussion on securing power system stability through supervised machine learning in Chapter 2, we now shift our focus from energy delivery to the end-users—the primary consumers of electricity. Among these consumers, building energy management stands out due to its substantial impact on global energy consumption. Although buildings are often perceived as minor individual energy consumers, collectively, they account for a significant share of total energy usage worldwide¹⁵¹. Efficient management of building energy not only holds great potential for reducing global energy consumption but also significantly contributes to enhancing overall sustainability and environmental quality, steering society toward a greener and more sustainable future.

The management of building energy presents unique challenges, primarily due to the highly dynamic and complex nature of occupant behaviors. As previously discussed, occupant behavior—including occupancy patterns, clothing insulation, and metabolic activities—can considerably influence energy consumption and occupant comfort in building environments^{54;59}. Therefore, incorporating realistic occupant behavior into energy management models is crucial to accurately capture the potential for energy savings and occupant comfort

improvement.

To systematically address these challenges, this chapter focuses on developing advanced building energy management methodologies that explicitly incorporate occupant-centric control strategies. To achieve this goal, we define two key research objectives:

- Incorporate Dynamic Occupant Behavior into HVAC Management

We first examine the critical role that occupant behaviors, specifically clothing adjustments, occupancy patterns, and activity levels, play in influencing the building’s HVAC (heating, ventilation, and air conditioning) energy consumption. A dynamic clothing behavior model is integrated into the building simulation environment to quantify the impacts of occupant-driven factors on both energy consumption and occupant comfort. The validity and effectiveness of modeling dynamic occupant behaviors are demonstrated through comprehensive simulation analyses.

- Develop and Evaluate a Multi-Agent Deep Reinforcement Learning (MADRL) Framework for HVAC Optimization

Recognizing the complexities introduced by dynamic occupant behaviors, we propose a novel multi-agent deep reinforcement learning (MADRL) framework designed to optimize HVAC setpoints across multiple thermal zones. This framework intelligently schedules cooling and heating setpoints by balancing two competing objectives: minimizing electricity costs and reducing occupant thermal discomfort. The MADRL framework is evaluated through simulations in EnergyPlus, leveraging realistic building parameters and occupant behavior datasets.

This chapter systematically addresses critical challenges in building energy management, highlighting the significant potential of occupant-centric optimization techniques to transform buildings into efficient and comfortable living and working environments.

3.1 Optimization and occupant behavior validation in buildings

Approximately 100 million single-family homes in the United States account for 36% of the electricity load, and often they determine the peak system load, especially on hot summer days when the use of residential air-conditioning (AC) is high⁵². Increased consumer adoption of home automation products such as smart thermostats for heating, ventilation and air conditioning (HVAC) is a prominent trend in residential buildings.

In addition to energy cost savings, occupant's thermal comfort is one of the most significant factors considered in the thermostat control of the HVAC. One popular index of the thermal comfort is the Predicted Mean Vote (PMV) model and Predicted Percentage of Dissatisfied (PPD), which was proposed by P.O. Fanger and his colleagues in Kansas State University and Technical University of Denmark in the 1970s¹⁵². In addition, PMV-PPD is included in ISO 7730¹⁵³ in 2005 and the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) standard 55¹⁵⁴ in 2004. Buratti et al.¹⁵⁵ developed a linear regression approach to the PMV with a wider range of clothing thermal insulation in moderate environments based on measured air temperature and relative humidity. Jia et al. applied an event-based optimization control via Markov Decision Process (MDP) with PMV while considering the effect of multi-rooms¹⁵⁶. Luo et al proposed a metaheuristic algorithm based energy management system (EMS) to optimally schedule energy resources (including residential PV) while using adaptive PMV to account for the occupant's thermal comfort⁶⁷. Clothing condition has a drastic impact on the human body thermal insulation, which in turn influences occupants' thermal comfort significantly. Therefore, the occupant's clothing condition is taken as a given parameter in building energy modeling and simulations. For example, since clothing has a drastic impact on the human body thermal insulation, it still remains unclear as to how human clothing decisions, i.e., donning and doffing, can influence optimal schedules of HVAC thermostats in terms of the occupant's electricity costs and thermal comfort.

The current state of the art in building energy modeling and simulations (e.g., EnergyPlus¹⁵⁷) is that the thermal comfort condition is calculated based on a dynamic clothing model. In this model, the clothing insulation is calculated as a function of outdoor air temperature measured at 6 o'clock and this clothing insulation value remains unchanged during the entire day. While there is no standardized guideline on how to set clothing insulation schedules, the dynamic clothing model leads to HVAC systems that are incorrectly operated and to the inaccurate assessment of occupant comfort conditions. In reality, residential occupants can frequently adjust their clothing, depending on the thermal conditions around them, on an hourly or sub-hourly basis, as opposed to a constant clothing value in an entire day. Such adjustments can improve home demand flexibility and thus enable a higher energy cost saving. The frequent clothing adjustment is even more plausible under the COVID-19 outbreak, whereby many businesses worldwide shift from traditional office-based work operations to work from home (WFH) arrangements. Therefore, the hourly or sub-hourly clothing adjustment should be fully captured and optimized in the residential home energy management to model HVAC systems realistically. There has been no work to date that addresses this issue. In addition, it is challenging to incorporate occupant's clothing adjustments in the HVAC energy management due to the complex nonlinear relationship between the clothing conditions and the occupant's thermal comfort.

Various home energy management models and solution algorithms have been proposed. Among them, approximate dynamic programming (ADP) represents a promising class of algorithms for decision-making under uncertainty^{158;159}. Past research efforts have used ADP for grid-battery management^{64;160–164}, real-time microgrid operations¹⁶⁵, and air conditioner to minimize electricity consumption while considering the the occupant thermal comfort^{104;166;167}. Nevertheless, to the best of our knowledge, there has been no work to date that uses the ADP to address the HVAC energy management issue while considering occupant's hourly clothing decisions.

This paper bridge this gap by proposing an HVAC energy management scheme that aims to minimize the electricity cost and the occupant's thermal discomfort while taking into ac-

Table 3.1: Comparison of Related Works in the Literature

References	Approaches	EMS	Comfort	Clothing factor	Study scope
¹⁵⁵	Linear regression		✓	Fixed clothing values	HVAC
¹⁵⁶	Complexity-based approach	✓	✓	Fixed clothing values	HVAC
⁶⁷	MPC	✓	✓		PV, HVAC, battery
¹⁶⁰	Temporal difference ADP	✓	✓		PV, battery
¹⁶¹	Distributed iterative ADP	✓			Battery
^{64;162;163}	Action-dependent heuristic DP	✓			PV, battery
¹⁶⁴	Iterative ADP	✓			Battery
¹⁶⁵	Spatiotemporal ADP	✓			Microgrid
^{166;167}	ADP	✓	✓		Thermal storage
¹⁰⁴	ADP-based MPC	✓	✓		HVAC, WH, EV
This work	ADP-based MPC with LSTM	✓	✓	Dynamic clothing w hourly decisions	HVAC

count the occupant’s optimal clothing decisions. Under the time-of-use (ToU) electricity tariff, the proposed scheme optimally determines the thermostat setpoints of HVAC and simultaneously provides the best recommendations on the occupant’s clothing conditions, while accounting for uncertainties in outside temperature. The proposed HVAC energy management scheme is solved by using an ADP-based algorithm. Further, a model predictive control (MPC) framework with a Long Short-Term Memory (LSTM)-based forecasting technique is developed to simulate the complex decision making considering the electricity prices, HVAC schedules, and occupants’ thermal comfort and clothing conditions. We systematically study the optimal HVAC schedules and hourly clothing decisions for both male and female occupants in the summer and winter seasons. Comparative results with and without considering the hourly clothing decisions are also revealed. The main contributions of this work are three-fold:

1) This paper proposes a novel HVAC energy management scheme that is the first of its kind to explicitly consider the occupant’s hourly dynamic clothing conditions as an additional dimension of decisions for minimizing the electricity cost and the occupant thermal discomfort under uncertainty.

2) An ADP-based algorithm is applied to include the occupant’s clothing states and actions. An MPC framework with the ADP algorithm as an optimization engine and an LSTM as an embedded outdoor temperature forecaster is proposed to simulate HVAC energy management. This framework simulates the HVAC operation under uncertainty in a more




Figure 3.1: The schematic diagram of proposed energy management.

realistic environment.

3) Proof-of-concept simulation results show that, if the occupant follows the optimal clothing decisions produced, 53.8% and 29.8% of daily electricity cost savings can be achieved respectively for a summer-male scenario and a winter-female scenario, only with negligibly compromised thermal comfort. Table 3.1 compares related work in this area to the contributions of this paper in terms of proposed approaches, energy management system, occupant comfort, clothing factor and study scope.

3.1.1 ADP-based Optimization Problem

The HVAC energy management problem is a discrete-time MDP, which contains an objective function, a group of state transition functions and physical constraints. We provide a detailed formulation in this subsection. The physical setting of the HVAC energy management scheme including the component connection, data and control flows within a residential home is illustrated in Figure 3.1. Here, the ADP-based HVAC controller is a central device that connects the sensors, meters and an HVAC thermostat. All the data from the sensors and smart meters will be sent to the cloud data center for data collection. A

weather forecaster will deliver the predicted weather data back to the ADP-based controller. Then, the controller determines optimal schedules for the setpoints of the HVAC thermostat and sends the optimally recommended clothing adjustment, if any, to the occupant through a smartphone application or a speaker in a smart home hub (e.g., Amazon Echo). The occupant decides whether to follow the clothing adjustment recommendation or simply ignore it. Then, the feedback will be sent to the ADP-based controller, in which the occupant's clothing state is updated. The above setting is available in today's smart home environment¹⁶⁸, in which various measurements (e.g., indoor and outdoor temperatures, relative humidity, and electricity prices) can be collected frequently to facilitate real-time control.

The aim of ADP is to minimize the weighted sum of the occupant's electricity consumption and thermal discomfort. This model contains discrete states, discrete actions, and stochastic factors. At each time interval t , a state consists of the room temperature s_t^H and the occupant clothing state s_t^C . A decision contains HVAC power x_t^H and occupant's clothing adjustment x_t^C , i.e., doffing or donning. The stochastic variables $\tilde{u}_t$ are the outside temperature. Additionally, the HVAC is in state s_t^H when a decision x_t^H is determined at the beginning of each time interval. Then, the decision is implemented instantly and affects the rest of the time intervals. After the decision is applied, the outdoor temperature as the uncertainty variable is realized and affects the rest of the time intervals.

The state-decision set tuple $\{s_t, x_t\} = \{(s_t^H, s_t^C), (x_t^H, x_t^C)\}$ contains HVAC- and clothing-related state and decisions for all time periods. The objective function is to minimize the expectation value of the reward:

$$\min \mathbb{E}\left\{\sum_{t=1}^T R_t(s_t, x_t, \tilde{u}_t)\right\} \quad (3.1)$$

$$R_t(s_t, x_t, \tilde{u}_t) = \alpha_t^D \cdot \mathcal{D}_t(s_t) + \alpha_t^C \cdot \mathcal{C}_t(s_t, x_t, \tilde{u}_t) \quad (3.2)$$

Equation (3.2) shows that the reward R_t consists of thermal discomfort $\mathcal{D}_t$ and the electricity

cost $\mathcal{C}_t$ of the occupant. Additional, $\alpha_t^{\mathcal{D}}$ and $\alpha_t^{\mathcal{C}}$ are respectively weight discomfort coefficient and cost coefficient. These value can be calculated based on the occupant's preferences¹⁰⁴. The occupant's discomfort $\mathcal{D}_t$ consists of PMV at time t is defined as follows:

$$\mathcal{D}_t(s_t) = \beta_t \cdot |\text{PMV}(s_t)| \quad (3.3)$$

In equation (3.3), the occupant's thermal discomfort HVAC is modeled as an absolute value of the PMV times a coefficient β . The PMV equation for the thermal comfort model is expressed as follows:

$$\text{PMV}(s_t) = a(s_t^{\mathcal{C}}) \cdot s_t^{\text{H}} + b(s_t^{\mathcal{C}}) \cdot P(s_t^{\text{H}}, rh_t) - c(s_t^{\mathcal{C}}) \quad (3.4)$$

$$P(s_t^{\text{H}}, rh_t) = rh_t \cdot 0.61121 \exp((18.678 - s_t^{\text{H}}/234.5) \cdot (s_t^{\text{H}}/(257.14 + s_t^{\text{H}})) \quad (3.5)$$

where $P(s_t^{\text{H}}, rh_t)$ is the water vapor pressure in ambient air (kPa) and rh_t is the relative indoor humidity. In order to study the occupant's dynamic clothing behaviors, we adopt a simplified PMV model from¹⁵⁵. This model is dependent only from air temperature and relative humidity because these two parameters are widely available due to their easy measurement in residential buildings. The boundary conditions (e.g., metabolic rate, air velocity) and the parameters $a(s_t^{\mathcal{C}})$, $b(s_t^{\mathcal{C}})$ and $c(s_t^{\mathcal{C}})$ in the PMV model can be found in¹⁵⁵. The contribution of electricity cost $\mathcal{C}_t$ is defined:

$$\mathcal{C}_t(s_t^{\text{H}}, x_t^{\text{H}}, \tilde{u}_t) = c_t x_t^{\text{H}} \Delta t \quad (3.6)$$

where c_t is the electricity price. In this paper, we adopt a ToU tariff such that the electricity prices are available beforehand. The cost function in (3.6) is calculated as the total electricity cost of HVAC in the scheduling horizon. The physical constraints of the HVAC are listed as follows:

The constraints of the energy management scheme are associated with the physical operational limits of the HVAC. We adopt a first-order, 1R1C (i.e., 1 resistor-1 capacitor) thermodynamic model that is widely used in the literature^{52;66;169;170}. The equivalent capacitor C represents the thermal mass of the building and the equivalent resistor R represents the resistance to heat transfer. The differential equation of 1R1C model is transferred in the Laplace domain and derived to get a set of linear Equations (3.7). The resultant linear equations consider the room temperature as a function of the room temperature at the previous time period, HVAC power consumption, and the outside temperature. The HVAC constraints are shown below:

$$s_t^H = S^M(s_{t-1}^H, x_t^H, \tilde{u}_t) = \begin{cases} \gamma^{\text{RM}} s_{t-1}^H - \gamma_c^H x_t^H + \gamma^{\text{OUT}} \tilde{u}_t, \forall t, \text{cooling} \\ \gamma^{\text{RM}} s_{t-1}^H + \gamma_h^H x_t^H + \gamma^{\text{OUT}} \tilde{u}_t, \forall t, \text{heating} \end{cases} \quad (3.7)$$

$$\underline{s}_t^H \leq s_t^H \leq \bar{s}_t^H, \forall t \quad (3.8)$$

$$\underline{x}_t^H \leq x_t^H \leq \bar{x}_t^H, \forall t \quad (3.9)$$

In (3.7), thermal coefficients are computed according to either the values of equivalent resistor and capacitor or a regression analysis^{52;170}. Equation (3.8) suggests that the states should be within the lower and upper bounds prescribed by the occupants. Equation (3.9) indicates that the decisions should be limited by the physical cooling or heating capability of the HVAC. In Table 3.3, parameters of the HVAC first-order model and details of the clothing conditions are clearly displayed.

Occupant's Clothing Conditions

Clothing insulation, measured in the unit of clo, is the thermal insulation provided by clothing as well as any layer of trapped air between skin and clothing. One unit of clo equates to $0.155 \text{ K}\times\text{m}^2/\text{W}$, which means the amount of clothing needed by a sedentary per-

son to maintain thermal comfort in an environment with 21 °C air temperature, 50% relative humidity (RH), and 0.1 m/s airspeeds. The ASHRAE Standard 55¹⁵⁴ and ASHRAE Handbook¹⁷¹ contains a list of clo values for selected garment types and formulas for estimating the insulation provided by a total clothing ensemble.

On the other hand, a simplified method of estimating clothing insulation is to multiply the weight of a clothing ensemble in lbs by 0.15. However, this method assumes there is no wind penetration in the nearby environment or body movements pumping air around. In general, the higher a clo value, the more insulating value is provided by a total clothing ensemble. Schiavon and Lee¹⁷² found that the median clothing insulation is 0.59 clo in summer and 0.69 clo in winter. Clothing adjustment behaviors, i.e. adding or reducing layers of clothing, have a direct impact on the occupants' thermal comfort, thus the optimum operative temperature changes with clo values. The effect of changing clothing insulation on the optimum operative temperature is approximately 6°C per clo for a sedentary man whose metabolic rate is approximately 1.2 met and this effect is greater with the higher metabolic rate¹⁵⁴.

Table 3.2 shows the estimated clo values of some typical business casual clothing ensembles. All the estimated clo values are from¹⁷³. The range of clo values are traditionally divided into three groups¹⁵⁵, which are Clo 1: 0.25 - 0.5, Clo 2: 0.51-1.00, and Clo 3:1.01-1.65. Studying dressing behaviors is complicated particularly with the WFH arrangement during the COVID-19 pandemic. Some people value comfort, therefore prefer to dress down when they are working at home while some others believe dressing in a routine manner, such as wearing business casual outfits while working from home, can help them maintain a sense of control, degree of normality, and productivity. Even though business casual ensembles are used to demonstrate the relationship between clo values and the occupant's donning and doffing behaviors at a residential home in this paper, this relationship exists in other types of ensembles as well.

According to Table 3.2, if a male is dressing in Clo 2 but feeling cold, he can exchange his thin sweater for a thick suit, swap the stockings for a pair of calf-length socks, trade his

Table 3.2: Clo values of typical business casual clothing ensembles

Range (clo)	Male Ensembles	Female Ensembles
Clo 1: 0.25-0.50	<u>Estimated clo = 0.42</u> <u>Outfit:</u> A short-sleeve shirt (0.19), a pair of thin straight trousers (0.15). <u>Underwear:</u> a man’s brief (0.04). <u>Footwear:</u> A pair of stockings (0.02), shoes (0.02).	<u>Estimated clo = 0.41</u> <u>Outfit:</u> A short-sleeve dress shirt (0.19), a thin skirt (0.14) <u>Underwear:</u> a bra (0.01) and a panty (0.03). <u>Footwear:</u> A pair of stockings (0.02), shoes (0.02).
Clo 2: 0.51-1.00	<u>Estimated clo = 0.82</u> <u>Outfit:</u> A long-sleeve shirt (0.25), a thin long sleeve sweater (0.25), a pair of thick straight trousers (0.24). <u>Underwear:</u> a man’s brief (0.04). <u>Footwear:</u> A pair of stockings (0.02), shoes (0.02).	<u>Estimated clo = 0.81</u> <u>Outfit:</u> A long-sleeve shirt (0.25), a thin long sleeve sweater (0.25), a thick skirt (0.23) <u>Underwear:</u> a bra (0.01) and a panty (0.03). <u>Footwear:</u> A pair of stockings (0.02), and a pair of shoes (0.02).
Clo 3: 1.01-1.65	<u>Estimated clo = 1.23</u> <u>Outfit:</u> A long-sleeve shirt (0.25), a thick single-breasted suit jacket (0.42), a pair of thick straight trousers (0.24). <u>Underwear:</u> a man’s brief (0.04), a pair of long underwear bottoms (0.15). <u>Footwear:</u> A pair of calf-length socks (0.03), a pair of boots (0.10).	<u>Estimated clo = 1.11</u> <u>Outfit:</u> A long-sleeve shirt (0.25), a thick single-breasted suit jacket (0.42), a pair of thick straight trousers (0.24). <u>Underwear:</u> a bra (0.01), a panty (0.03). <u>Footwear:</u> A pair of knee socks (thick) (0.06), and a pair of boots (0.10).

shoes for boots, and even add a pair of long underwear bottoms. Differently, if he is feeling hot when he is dressing in Clo 2, he can take off his thin sweater, exchange the long-sleeve shirt for a short-sleeve shirt, and swap the thick trousers for a pair of thin trousers. On the other hand, if a female is dressing in Clo 2 and feeling cold, she can exchange her thin sweater for a thick suit, swap her stockings for a pair of thick knee socks, and trade her shoes for boots. She may also place a blanket on her laps to warm-up her body (clo value varies based on the thickness and materials of the blanket, so its clo value is not included in Table 3.2. If she is feeling hot when she dresses in Clo 2, she can take off her thin sweater, exchange the long-sleeve shirt for a short-sleeve one, and swap the thick skirt for a thin skirt.

Based on the above analysis, we for the first time incorporates the occupant clothing conditions in the proposed model. Similar to the HVAC, the state transition of clothing

conditions is expressed as follows:

$$s_t^C = s_{t-1}^C + x_t^C \quad (3.10)$$

$$\underline{s}_t^C \leq s_t^C \leq \bar{s}_t^C, \forall t \quad (3.11)$$

$$\underline{x}_t^C \leq x_t^C \leq \bar{x}_t^C, \forall t \quad (3.12)$$




Figure 3.2: Illustration of a state transition diagram considering occupant's clothing states and actions in winter (when $T = 3$).

3.1.2 ADP Algorithm and MPC Framework

In this subsection, we first describe the ADP algorithm, including state transition, the K-D tree Nearest Neighbor (KDNN), and the Sobol Sampling Backward Induction (SSBI).

Then, we introduce the MPC framework that incorporates the ADP algorithm and an LSTM forecaster.

State Transition and ADP

In our ADP algorithm, the value function $V(s_t)$, the expected value of the state s_t while applying all feasible decisions x_t and uncertainty $\tilde{u}_t$ at time t :

$$V(s_t) = \sum_{s_{t+1}} \mathbb{P}(s_{t+1}|x_t, s_t, \tilde{u}_t) [R(s_t, x_t, \tilde{u}_t) + V(s_{t+1})] \quad (3.13)$$

where $\mathbb{P}(s_{t+1}|x_t, s_t, \tilde{u}_t)$ is the probability of s_{t+1} given x_t, s_t and $\tilde{u}_t$. An optimal action x_t^* is attained by $V_t^*(s_t)$ along with qualifying the Bellman optimality condition in (3.14)

$$V_t^*(s_t) = \min_{x_t} \left[R_t(s_t, x_t, \tilde{u}_t) + \mathbb{E}(V_{t+1}(s_{t+1})|s_t) \right] \quad (3.14)$$

Figure 3.2 shows a diagram of state transition when $T = 3$, in which a red-highlighted path represents the optimal decision sequence. At the terminal state ($t = T$), $V_t(s_t)$ is calculated according to discomfort value for all possible s_T in the state space $\mathbf{S}$ as defined in equation (3.8). Next, for all previous states ($t = 1$ and $t = 2$), we apply a backward induction to identify the optimal value $V_t^*(s_t)$ and the best decision x_t^* . While traversing backward in the state transition diagram, x_t^* and $V_t^*(s_t)$ at each time index t are collected. At the initial time period ($t = 1$), $V_1^*(s_1)$ corresponding to x_1^* is obtained.

The well-known "curse of dimensionality"¹⁵⁸ dictates that an extremely long or even infeasible run time may occur due to a vast amount of state, decision and uncertainty spaces. Here, we use KDNN and SSBI to address this challenge¹⁰⁴. SSBI is similar to the classic backward induction, but it contains a supplementary Sobol sampling function $\mathcal{F}_s$, which compacts the workload efficiently since only necessary state, action and uncertainty sets are considered. KDNN is one important approach of value function approximation. In order to achieve a fast computation, KDNN traverses the nearest point of s'_t , and looks for an

approximation value of $V(s'_t)$ instead of calculating from equation (3.13). The combination of KDNN and SSBI has shown promising performances in finding the optimal decision x_1^* . The ADP algorithm used in this paper minimizes the electricity cost and the occupant's thermal discomfort while taking into account the occupant's clothing decisions. The algorithm is modified based on Algorithm 3 in¹⁰⁴, but the main difference between them is the introduction of clothing factor into the state and decision spaces. More details about the algorithm can be found in¹⁰⁴.




Figure 3.3: Flowchart of the ADP based MPC framework.

ADP-based MPC Framework with LSTM

The methodology of the ADP-based MPC framework is illustrated in Figure 3.3, which is developed to test the effectiveness of the proposed HVAC energy management scheme in a more practical simulation that represents the state of the art in smart home applications. The MPC framework minimizes objective function (3.1) subject to the constraints (3.2)-(3.12) through a look-ahead horizon from the the present time to the future ($t = T$). As shown in Figure 3.3, the MPC framework collects three classes as inputs: 1) ADP-based

MPC configurations including parameters pertaining to HVAC and ToU prices; 2) occupant's preferences including the weight of both discomfort and cost, desired sleeping temperature, PMV and clothing conditions; and 3) an outside temperature forecast, which is updated over time. After the ADP is solved, the decisions made for the present time interval are adapted to the real-time operation simulation module. Then, the thermostat setpoint and occupant's clothing adjustment are simulated with given optimal decisions. Initial states of the HVAC, the occupant's clothing, the PMV value and the electricity cost, all for the next window, are returned. The above steps repeat in the next window to form a receding horizon control. The ADP-based MPC sums up the simulated thermal discomfort and electricity costs over the whole horizon and determines an overall objective value.

The proposed HVAC energy management model can either receive weather forecasts of hourly outside temperature from the internet, or use the embedded LSTM forecaster to predict the outside temperature. Here, we unitize an LSTM technique to predict the hourly outside temperature based on the historical data by leveraging our previous work¹³¹. The LSTM is one type of recurrent neural networks (RNNs) and is capable of learning order dependence in sequence prediction problems such as speech recognition and context recognition^{174, 175}. The superiority of the LSTM is that it solves the long-range dependence more accurately than the conventional RNN¹⁷⁶. According to the testing results, the mean absolute percentage error (MAPE) of the outside temperature forecast is about 0.02 and 0.38 for summer and winter, respectively, signifying an accurate forecast well suited in the proposed MPC framework.

3.1.3 Preliminary Results

The ADP-based energy management model is implemented in MATLAB via DYNAMO toolbox¹⁷⁷. The HVAC parameters and occupant's clothing conditions are shown in Table 3.3. The HVAC parameters in Equation (3.7) are obtained using a linear regression method on a historical dataset of a residential home in Hillsboro, Oregon^{169;170} with both a cooling case and a heating case.

Table 3.3: HVAC parameters and clothing conditions

State Range			State Revolution	
[18 30]			0.1	
Cooling case parameters				
γ^{RM}	γ^{OUT}	γ_c	COP	$\bar{x}(\text{kW})$
0.94	0.06	0.116	4.75	4
Heating case parameters				
γ^{RM}	γ^{OUT}	γ_h	COP	$\bar{x}(\text{kW})$
0.95	0.05	0.15	2.95	8
Time range		Occupant's clothing conditions		
[10pm, 6am]		Clo 1		
[6am, 10pm]		optimal hourly decision		




Figure 3.4: The time-of-use electricity tariff.

We use the ToU electricity tariff, i.e., Pacific Gas & Electric EToU-E6 as the summer rate, and the Southern California EDISON TOU-D-5-8PM as the winter rate. For the summer, this tariff contains three price levels: Base, Peak A, and Peak B prices, i.e., \$0.244/kWh, \$0.32/kWh, and \$0.436/kWh, represented by white, light grey and dark grey in Figure 3.4 (a), respectively. Peak A takes place in Hours 11-13, 20, and 21, and Peak B occurs during Hours 14-19. Base prices are utilized in the rest of the hours. Similarly, for the winter, as shown in Figure 3.4 (b), Base, Peak A, and Peak B are \$0.24/kWh, \$0.28/kWh, and \$0.42/kWh, respectively.

In order to demonstrate the benefit of using the proposed HVAC energy management scheme, simulation results in both summer and winter under a traditional fixed setpoint thermostat (i.e., without the proposed energy management scheme) are shown in Figure 3.5. We assume a medium clothing state (i.e. $s_t^C = 2$) with the desired temperature the occupant sets is 22.8 °C in winter for female and 20 °C in summer for the male. This is because 22.8




Figure 3.5: Simulation results of a female winter case and a male summer case, both without the proposed HVAC energy management scheme.

$^{\circ}\text{C}$ is the most comfortable ($\text{PMV} = 0$) temperature for the female when the cloth state s_t^C is set to Clo 2, while that is 20°C for male. Without loss of generality, we simulate the winter case for a female occupant and the summer case for a male occupant. In the winter, i.e., a heating case, the outside temperature is 22.4°C lower than the room temperature at the initial time, i.e., 6:00. For the summer, i.e., a cooling case, the outside temperature is around 2.2°C higher than the room temperature at 6:00 and becomes 11.1°C higher at 18:00. In Figure 3.5, we only show the simulation results from 6:00 to 22:00 because the other hours are considered as the occupant’s sleeping time, when the fixed setpoint of the thermostat is respectively set to 25°C and 22°C in winter and summer as a comfortable temperature for female and male occupants¹⁷⁸. It is seen the traditional fixed-setpoint thermostat keeps the room temperature around the desired temperature regardless of the electricity price and energy cost, implying the cost-saving potential for the energy management scheme.

With the proposed scheme, Figure 3.6 shows the MPC simulation results for a female occupant. Specifically, Figure 3.6 (a) shows the indoor temperature, outside temperature and the HVAC power, while Figure 3.6 (b) shows the clothing state and dressing decisions. To facilitate our comparison and focus exclusively on the hours with clothing behavior, we adopt the same fixed setpoint control during the sleeping time as in the non-optimized case. At 6:00, the user’s cloth state is Clo 1 and the indoor temperature is 25°C . From 6:00 to 12:00 (noon), there is no HVAC heating but the donning decisions are made at both 6:00 and 8:00 to eventually get a clothing state Clo 3 at 10:00. This result shows that the occupant’s




Figure 3.6: Simulation results of a winter day for female with ADP.

clothing behavior plays an additional role in maintaining a satisfactory thermal comfort even with an extremely low outdoor temperature.

Meanwhile, the indoor temperature decreases due to the dropping outside temperature until it reaches the lowest temperature during the day. The outside temperature is rolling down to the minimum while the clothing state is reaching the maximum, without any HVAC heating to save the cost. Then, the HVAC energy management scheme makes a maximum heating decision at 12:00 since it is in an off-peak period, which in turn drives the indoor temperature to rise. The turning point of the indoor temperature occurs at 12:00 while the occupant makes a doffing decision to change its cloth state from Clo 3 to Clo 2. With an increasing outside temperature, the energy management scheme makes several heating deci-

sions from 13:00-15:00. As a result, the indoor temperature increase and the clothing state also decreases from Clo 3 to Clo 1. The above results in Figure 3.6 suggest that in concert with the well-known pre-heating effect of the scheme, the ability to make optimal clothing decisions can provide an additional mechanism to better consider the occupant's thermal comfort while more effectively utilizing the thermal storage in the residential building.

At 16:00, both the outside temperature and indoor temperature come to the peak temperature of the daytime. At 17:00, in order to get a comparatively high indoor temperature (around 23.9 °C), the HVAC heats up the room and the female occupant stays with the minimum clothing state. Then, the occupant doffs at 18:00 when the indoor temperature decreases slightly. The HVAC slightly heats up the room at 19:00 and 20:00 to make the indoor temperature ends at the sleeping temperature we set at 22:00. It is worth mentioning the optimal clothing state follows the relationships between the PMV values and indoor temperatures for all three clothing states are referred from¹⁵⁵.

For a summer day, the simulation result of the HVAC energy management scheme for a male occupant is demonstrated in Figure 3.7. At 6:00, the occupant's clothing state is Clo 1 and the indoor temperature is 20.6 °C. The occupant dons (e.g., adding a pair of long underwear bottoms) and his clothing state increases from Clo 1 to Clo 2 at 6:00. There is no HVAC power during 6:00-12:00 when the outside temperature is rising. At 12:00, the HVAC energy management scheme decides to slightly pre-cool the house. With the increasing outside temperature to the maximum at 92 °F, the scheme cools the house at 15:00, 17:00, and 18:00 to maintain the maximum of indoor temperature around 23.3 °C. Accordingly, the dressing decision made is doffing from Clo 2 to Clo 1 at 14:00 (i.e., the beginning of Peak B), and keep the minimum clothing state unchanged from 14:00 to 20:00 when both the peak electricity price and outside temperature are on the peak. To save the electricity cost, the proposed scheme postpones a cooling decision at its rated power (i.e., 4kW) until 21:00 right after Peak B, to drastically brings down the room temperature. Meanwhile, The clothing state increases from Clo 1 to Clo 2 while the indoor temperature is drastically drooped to around 21.1 °C. The simulation result in Figure 3.7 shows again that the ability to make




Figure 3.7: Simulation results of a summer day for male with ADP.

optimal clothing decisions enhances the occupant's thermal sensation and provides an extra means to balance the occupant's thermal comfort and electricity cost.

In order to unequivocally show the benefit of the proposed scheme considering optimal clothing decisions, we perform a comparative study between the non-optimized and the optimized scenarios by the HVAC energy management. The comparative simulation results are listed in Table 3.4, in which two aforementioned scenarios, i.e., the summer case for a male occupant and the winter case for a female case, are compared. To facilitate the comparison, from 6:00 to 22:00, the initial and end temperatures of the non-optimal case are identical to those of the optimized case. However, the differences between the non-optimized and optimized cases are that the most comfortable temperatures under the dynamic clothing model

(i.e., $\text{Clo} = 2$ on this day) are taken as manual temperature setpoints in the non-optimized case, while the temperature setpoints in the optimized case are optimally determined by the proposed scheme. For example, in the summer-male scenario, the non-optimized thermostat setpoints are set to a most comfortable temperature, which is $19.8\text{ }^{\circ}\text{C}$ for $\text{Clo} = 2$. For the winter-female scenario, they are set to $22.8\text{ }^{\circ}\text{C}$ for $\text{Clo} = 2$.

As seen in Table 3.4, the optimized electricity cost of the summer-male scenario is 3.8 US dollars and the average PMV from 6:00 am to 10:00 pm is -0.1, whereas that of the winter-female scenario is 7.88 US dollars and the average PMV is -0.06. In either scenario, the optimized electricity cost by the HVAC energy management scheme is much superior to the non-optimized cost when the dynamic clothing model¹⁵⁷ is used. Recall that this dynamic model calculates the clothing insulation as a function of outdoor air temperature measured at 6 o'clock of the scheduling day and this clothing insulation value remains unchanged during the entire day. Therefore, this model cannot capture the intricacy of occupant clothing adjustment on an hourly basis. Through optimally determining the hourly clothing conditions, the proposed HVAC energy management scheme achieves a cost saving of 53.8% in the summer-male scenario, while a cost saving of 29.8% in the winter-female scenario. This much higher cost saving percentage in the summer-male scenario can be explained by comparing the average PMV values in Table 3.4. It is seen the average PMV value in the optimized summer-male scenario has a larger deviation from zero than that in the winter-female scenario, indicating a larger compromise in his thermal comfort. However, since the average PMV values are still close to zero, the occupant can hardly feel obvious thermal discomfort. The comparative results in Table 3.4 demonstrate that the proposed energy management scheme can significantly save the occupant's electricity cost with a negligibly compromised comfort.

3.1.4 Preliminary Summary

In this section, we propose a novel HVAC energy management scheme that is the first of its kind to take into account the occupant's hourly dynamic clothing behaviors to min-

Table 3.4: A daily comparison between non-optimized and optimized cases.

Case	Item	Non-optimized	Optimized
Summer	Cost (\$/day)	8.22	3.80
Male	Avg-PMV	-0.02	-0.10
Winter	Cost (\$/day)	11.23	7.88
Female	Avg-PMV	0.06	-0.06

imize the electricity cost and the occupant’s thermal discomfort. We adopt a PMV model that accounts for the indoor temperature, humidity, occupant gender, and a wide range of clothing thermal insulation. Then, the proposed scheme is embedded in an ADP-based MPC framework that includes the occupant’s clothing states and actions. Under a ToU tariff, the ADP optimally determines the setpoint of the HVAC thermostat and provides recommendations on occupant’s optimal clothing decisions based on typical business casual clothing ensembles for both genders, while considering outside temperature uncertainties. Such clothing recommendations can be sent to the occupant through a smartphone application or a speaker in a smart home hub. We systematically compare HVAC schedules with and without the optimal clothing decisions for both summer-male and winter-female scenarios. The proof-of-concept simulation results demonstrate the validity of the proposed HVAC energy management scheme and the effectiveness of the proposed ADP approach.

The benefits of considering the occupant’s optimal hourly clothing decisions as opposed to a constant clothing condition in an entire day are shown in this paper. In particular, simulation results show that, if the occupant follows the optimal clothing decisions produced, a 53.8% and a 29.8% of daily electricity cost savings can be achieved respectively for a summer-male scenario and a winter-female scenario, only with negligibly compromise in the occupant’s thermal comfort. Our simulation results also illuminate that the proposed HVAC energy management scheme has great capabilities of utilizing the building thermal storage in terms of pre-cooling, pre-heating, and delayed-cooling, etc. In the next section, we will focus on enhancing the proposed energy management scheme using deep machine learning techniques for a variety of behavior applications in next-generation commercial buildings.

3.2 Advanced Reinforcement Learning in Multi-zone Buildings

In this section, we aim to go beyond conventional occupant behavior models by incorporating comprehensive and dynamic occupant behaviors—including occupancy presence, metabolic rates, and clothing adjustments—into building energy management systems. We utilize advanced machine learning approaches, specifically reinforcement learning, to address the challenging multi-zone HVAC control problem, a common yet difficult-to-solve issue in practical building environments. Reinforcement learning is particularly well-suited for capturing and adapting to the stochastic and dynamic nature of occupant behaviors, enabling intelligent and adaptive setpoint control for improved energy efficiency and occupant comfort.

Specifically, the main contributions of this section are summarized as follows:

1. We propose a pioneering multi-zone HVAC energy management framework explicitly accounting for dynamic occupant behaviors, including occupancy presence, clothing insulation adjustments, and variations in occupant metabolic activity levels. The proposed framework simultaneously optimizes electricity costs and occupant thermal comfort across different zones in a building.
2. We develop a novel Multi-Agent Deep Reinforcement Learning (MADRL) algorithm to effectively manage HVAC setpoint controls, addressing critical challenges such as sequential decision-making, continuous action spaces, and stochastic, dynamically changing occupant behaviors. Our approach ensures adaptability and efficiency in realistic, complex multi-zone environments.
3. We extensively validate the practical applicability and effectiveness of the proposed MADRL framework using EnergyPlus simulations with real-world occupancy, clothing, and metabolic datasets. Simulation results at both daily and yearly timescales demonstrate the proposed model’s significant advantages, achieving electricity cost

savings of approximately 4.34% compared to single-agent DRL solutions and 51.09% relative to traditional rule-based HVAC controls, while consistently maintaining high occupant comfort levels across multiple thermal zones.

These contributions highlight the significant potential of integrating dynamic occupant behavior modeling and advanced reinforcement learning methods into building energy management, furthering the pursuit of sustainable and energy-efficient built environments.

3.2.1 Problem Formulation and Methodology

This subsection details the problem formulation and methodology adopted for developing the proposed multi-zone thermal management framework. We start by formally defining the optimization problem that addresses the complex interactions among multi-zone HVAC systems, dynamic occupant behaviors (including occupancy, clothing insulation adjustments, and metabolic rates), and energy consumption objectives. Next, we introduce a detailed mathematical representation of the optimization objectives and constraints, explicitly focusing on occupant comfort, energy efficiency, and operational limitations. Subsequently, the methodological framework leveraging MADRL is elaborated, highlighting the state, action, and reward definitions designed specifically for the multi-zone HVAC control context. Finally, the reinforcement learning training and implementation processes within the EnergyPlus environment are described, illustrating the practical applicability of our approach.

Overview of Approach

Reinforcement learning is a paradigm within machine learning wherein an autonomous agent endeavors to acquire an optimal strategy for selecting a sequence of actions within an environment to maximize its cumulative reward. The agent’s decision-making process hinges on the feedback it receives in the form of a reward value following each executed action. These decisions are contingent upon the agent’s interpretation of the environment, which is encapsulated by a state representation. This iterative process persists as the agent engages with the environment, with the aim of progressively improving its policy, ultimately




Figure 3.8: Illustration of proposed MADRL in multi-zone HVAC control. Three modules are contained in this approach, 1) a MADRL algorithm: Multi-agent Deep Deterministic Policy Gradient; 2) A testbed: EnergyPlus simulator; 3) Stochastic occupant behavior model includes clothing, occupant presence, and metabolic rate.

striving to attain predefined objectives. Figure 1 provides a visual representation of the research framework, showcasing the seamless integration of the MADRL algorithm within a complex five-zone office building environment. This environment, as shown in the middle module of Figure 1, is accurately simulated using the Building Controls Virtual Test Bed (BCVTB)¹⁷⁹ in conjunction with EnergyPlus, facilitating realistic building dynamics and HVAC power consumption modeling. Within this framework, the MADRL agent (top module of Figure 1) actively engages with the environment, effectively making decisions (heating and cooling setpoints of HVAC) pertaining to HVAC control to simultaneously optimize energy efficiency and occupant comfort. A distinctive feature highlighted in the figure is the incorporation of dynamic occupant behavior models, as shown in the bottom module of Figure 1, represented through a stochastic process. These models dynamically account for variations in factors such as clothing insulation, occupant presence, and metabolic rates over time. Furthermore, it's essential to note the bidirectional interaction, where reward values are transmitted from EnergyPlus to the MADRL module, and optimal actions are passed from MADRL to EnergyPlus during simulation, effectively capturing occupant behavior. This integration stands at the core of the research's mission to achieve efficient multi-zone HVAC control while upholding occupant thermal comfort standards within the building.

Occupant behavior Modeling

Building energy management has undergone a significant transformation with the integration of technology, sustainability considerations, and occupant comfort. Occupant-centric control (OCC) has emerged as a prominent concept, shifting the focus from a building-centric to an occupant-focused approach^{180–182}. Researchers now recognize the importance of addressing occupants' needs and optimizing building systems for their comfort and energy efficiency. In the context of ensuring thermal comfort, the Predicted Mean Vote (PMV)/Predicted Percentage of Dissatisfied (PPD) index, introduced by P.O. Fanger and colleagues, is employed to measure occupants' comfort levels based on their environmental conditions¹⁵². This index is included in ISO 7730 (2005) and the ASHRAE standard 55

(2004)^{153;154}. PMV/PPD models, which take into account the indoor environment and occupant behavior factors as input, have been integrated into OCC as a part objective of the HVAC energy management scheme.

The proposed OCC allows occupants to have control over settings such as cooling and heating points in HVAC systems, enhancing their satisfaction and productivity. However, one of the critical factors in implementing OCC is understanding and integrating occupant behavior¹⁸³. Factors such as occupant presence, activity levels, and clothing behavior are significant in optimizing thermal comfort strategies. Prior literature has often overlooked that behavior in building energy management, resulting in an incomplete understanding of occupant needs and potentially leading to inefficient control strategies. Our work contributes to addressing these challenges and developing an approach that considers and integrates these aspects into occupant-centric building energy management systems. Note that during specific time periods, such as morning opening (8 am), lunchtime (noon to 1 pm), and afternoon closing (6 pm to 7 pm), occupants are assumed to be engaged in walking activities, resulting in elevated activity levels. In the simulation, we assumed that occupants within the same zone share similar behaviors in terms of activity level and clothing insulation. This is reasonable in a multi-zone commercial building environment as the behaviors of one can easily affect others within the same zone. This assumption allowed us to group occupants within each zone and apply general patterns of behavior, which is a common approach in HVAC simulations. Initially, we established a fixed schedule and subsequently introduced a stochastic model to this schedule to simulate diverse behaviors across different zones.

Occupant Presence

Occupant presence information plays a pivotal role in modern building management systems, offering invaluable insights into space utilization and opportunities for optimizing energy consumption. As illustrated in Table 1.3, occupant presence is a prevalent factor in HVAC control as occupant behavior. In our study, we initially implement a fixed schedule for occupant presence and subsequently introduce stochastic variables to simulate occupancy patterns across the five zones, as depicted in Figure 3.9 (top). While future research avenues




Figure 3.9: One-day average occupant presence schedule (top) and activity level schedule (bottom) for each zone

may explore data-driven or Markov chain models to predict time-dependent occupant behavior within simulations^{56;184;185}, our current approach serves as an initial exploration of how dynamic occupant behavior influences HVAC control strategies.

Metabolic Rate Efficient energy management in office or residential buildings necessitates understanding occupants' metabolic rates and activity levels, which directly impact energy expenditure^{69;186}. Metabolic rate, derived from oxygen consumption, indicates an individual's energy expenditure during physical activity. Activity levels, measured in watts per square meter (W/m^2), determine the heat generated by occupants. Typical office activities range from $55 \text{ W}/\text{m}^2$ to $115 \text{ W}/\text{m}^2$.

To optimize energy usage, it is essential to consider these variations in activity levels. Table 3.5 provides an overview of different behavioral conditions based on metabolic rates (in units: W/m^2 and Met). Tailoring energy management strategies to accommodate these diverse activity levels enables the implementation of sustainable practices in office buildings.

Table 3.5: Activity level and metabolic rate

Office Activities	Activity Level (W/m^2)	Met
Reading, seated	55	1
Seated, quiet	60	1
Writing	60	1
Typing	65	1.1
Standing, relaxed	70	1.2
Filing, seated	70	1.2
Filing, standing	80	1.4
Walking (0.9 m/s)	115	2

Figure 3.9 (bottom) illustrates the average activity level in the office building. The assumption is that the schedule of this behavior is average in all thermal zones and during morning opening (8 am), lunchtime (noon to 1 pm), and afternoon closing (6 pm to 7 pm), occupants are assumed to have a high activity level. By introducing random variables to the activity level, we can introduce greater diversity in this behavior, capturing more realistic occupant activity patterns.

Clothing behaviors Recognizing the influence of clothing on the discomfort function (PPD in the reward function), it becomes vital to consider occupants’ clothing behaviors. To achieve this, the research adopts a dynamic clothing behavior model developed by Schiavon and Lee [51], as represented in Equation 1. Within this model, o^{clo} denotes the occupant’s clothing insulation, while $t_{a(out,6)}$ represents the outdoor temperature at 6 am. Schiavon and Lee’s study employed multivariable linear mixed models, with the first model accounting for outdoor air temperature and the second incorporating indoor operative temperature. These models successfully explained the total variance in clothing behavior, enhancing the realism of thermal comfort assessments in office buildings. Moreover, the inclusion of stochastic elements accommodates the inherent variability in clothing behavior, making the evaluation of thermal comfort in the office building more authentic and reflective of the diverse clothing choices and individual preferences of occupants.

$$\begin{aligned}
o^{clo} = & \begin{cases} 1, & t_{a(out,6)} < -5^\circ C \\ 0.818 - 0.0364 * t_{a(out,6)}, & -5^\circ C \leq t_{a(out,6)} < 5^\circ C \\ 10^{-0.1635-0.0066+t_{a(out,6)}}, & 5^\circ C \leq t_{a(out,6)} < 26^\circ C \\ 0.46, & t_{a(out,6)} \geq 26^\circ C \end{cases} \quad (3.15)
\end{aligned}$$

Stochasticity in Occupant behavior Model

Stochastic modeling is a mathematical approach used to analyze systems involving randomness and uncertainty which employs probability theory to describe the likelihood of various outcomes, often using simulations to estimate complex systems. In addition, stochastic modeling is essential to make probabilistic predictions and assess risks in systems influenced by chance events and variability¹⁶⁹. Previous studies have extensively explored appliance scheduling problems using various stochastic models, such as forecast errors in hot water usage¹⁰⁴, outdoor temperature¹⁷⁰ and renewable energy generation^{187;188}. On the other hand, as shown in Table 1.3, occupant presence has been one of the most popular areas in recent research. However, these studies have primarily focused on only occupant presence as the main aspect of occupant behaviors, while neglecting the dynamic aspects. In addition, varying clothing behaviors and various metabolic rates are not considered. To address this research gap, here we propose a model with occupants' behaviors influenced by the time of day and weather. In addition, we introduce stochasticity^{189;190} into the occupant behavior model to show the effect of variability and range of patterns. Specifically, the following parameters are defined:

$$\begin{aligned}
OB(t) &= \{o_t^{pres}, o_t^{clo}, o_t^{met}\} \\
X(t) &= \{X_t^{pres}, X_t^{clo}, X_t^{met}\}
\end{aligned}$$

where t is the time index; $OB(t)$ represents the estimated values of occupant behaviors without considering any stochastic effects; $X(t)$ represents Gaussian-distributed random variables at time t for occupant presence, clothing behavior and metabolic rate. Note that $OB(t)$ values are derived from the office schedules, shown in Figure 3.9 for occupant presence and

metabolic rate based on different activity levels, while the clothing behavior is determined using Equation 3.15. Note that the clothing insulation and activity level are measured in the average value through zone because EnergyPlus, as the simulation platform used in this study, has limitations in representing individual differences.

Based on those parameters, we define:

$$\widetilde{OB}_t = OB(t) + X(t) \quad (3.16)$$

where $\widetilde{OB}_t$ is the stochastic dynamic occupant behavior that provides the ability to capture the inherent time-dependent variability and uncertainty of occupant behaviors. By incorporating random variables, our model can effectively simulate the nature of occupant behavior (presence, clothing, and metabolic rate), making it the first of its kind to integrate this concept into DRL-based building energy management systems.

States, Actions and Rewards of Reinforcement Learning

Three indices are employed: $t = \{1, 2, 3, \dots, T\}$ for time slots, $m = \{1, 2, 3, \dots, M\}$ for building zones, and $j = \{H, C\}$ to distinguish between the HVAC dual modes of heating (H) and cooling (C). The Reinforcement Learning consists of sets of states and actions, denoted as S and A , respectively, where $s \in S$ and $a \in A$. Specifically, s_t^m represents the state of zone m at time slot t , and $a_t^{m,H}$ and $a_t^{m,C}$ indicate the heating and cooling setpoints for zone m at time slot t . In the same way, $\widetilde{OB}_t^m$ is the stochastic occupant behavior in zone m at time slot t . It is worth mentioning that $\widetilde{OB}_t^m$ is part of s_t^m . The reward function r_t is defined for the entire building at time slot t .

The state s_t includes two parts: external states and internal states. External states related to the building's outside environment encompass the outdoor dry bulb temperature, air relative humidity, wind speed, wind direction, diffuse solar radiation, and direct solar radiation. For each zone m , seven states are considered, comprising the zone air temperature, zone thermal comfort mean radiant temperature, zone air relative humidity, zone thermal

comfort clothing value, thermal comfort index (PMV/PPD value), occupant counts, and zone average metabolic rate. Additionally, there are two internal states associated with electricity price and HVAC power consumption. The power consumption in our simulation represents the total electricity consumed by HVAC for the entire building, with the assumption that the entire building receives a single utility bill. As a result, power consumption is treated as a single parameter, and individual HVAC power consumption for each zone is not considered. The policy of the proposed MADDPG is a mapping that takes the current multi-zone environment observation and generates a probability distribution of actions, specifically the heating and cooling setpoints. This approach empowers the agent to make informed decisions based on the current state of the building and its zones, leading to optimized HVAC control strategies that adapt to various environmental and occupant conditions. The set tuple $\{s_t, a_t\}$ is used to represent the states, including external states and internal states, and actions of all zones in the building:

$$\{s_t, a_t\} = \{(s_t^1, \dots, s_t^M), (a_t^{1,H}, a_t^{1,C}, \dots, a_t^{M,H}, a_t^{M,C})\}$$

A reward function in DRL is a numerical signal that informs an agent about the desirability of its actions in a given state, with higher values indicating favorable actions and lower values representing unfavorable ones. The agent's goal is to learn a policy that maximizes the cumulative reward it receives over time. The reward function for the whole environment is defined as follows:

$$r_t(m) = -d \cdot \sum_{m=1}^M PPD(s_t^m) - b \cdot \mathcal{C}_t(s_t, a_t) \quad (3.17)$$

In (3.17), the reward function consists of two components: 1) a discomfort function calculated using the PMV/PPD, and 2) an electricity cost function $\mathcal{C}_t$. The weighted coefficients d_t and b_t can be determined based on historical data¹⁶⁹ and represent varying preferences of occupants, including cost-saving or comfort-seeking type. The electricity cost function $\mathcal{C}_t$ is given by:

$$\mathcal{C}_t(s_t^m, a_t^m) = c_t \sum_{m=1}^M p_t^m \Delta t \quad (3.18)$$

where, c_t is the electricity price at time slot t , Δt is the simulation time interval, and p_t^m is the electricity consumption caused by the HVAC system in zone m at time slot t . Importantly, c_t and p_t^m are part of the states that can be observed from the environment.

Multi-Agent Deep Reinforcement Learning

While classical reinforcement learning algorithms like Q-learning and policy gradient have exhibited proficiency in single-agent domains, their application presents unique challenges characterized by evolving policies, non-stationary surroundings, and the imperative for agent collaboration. In response to these multifaceted demands, the machine learning community has introduced DDPG, a prominent member of the reinforcement learning family renowned for its aptitude in handling continuous action spaces¹⁷. DDPG serves as a pivotal precursor to our exploration of MADDPG, an extension tailored explicitly for multi-agent domains. MADDPG advances the DDPG paradigm by empowering agents to deliberate global states and make informed decisions predicated on the actions of fellow agents, thereby enhancing coordination and overall system performance. The ensuing discourse delves deeper into the nuanced application of DDPG and the pivotal role of MADDPG in addressing the intricacies of multi-agent scenarios¹⁹¹. MADDPG is specifically crafted for multi-agent scenarios, where multiple agents interact within the same environment, potentially requiring coordination, communication, collaboration, or competition among agents. MADDPG extends DDPG to accommodate these complex multi-agent dynamics, making it suitable for modeling a wide range of cooperative or competitive interactions among autonomous agents. MADDPG offers several positive aspects in the context of multi-agent reinforcement learning:

- 1) Cooperative Learning: MADDPG facilitates cooperative learning by allowing agents to share information and learn from each other. By considering the joint actions and obser-

vations of all agents, MADDPG promotes coordination and collaboration among the agents, leading to better overall performance.

2) Centralized Learning, Decentralized Execution: MADDPG employs a centralized training approach, where a centralized critic network is used to estimate the Q-values based on the joint actions and observations. However, during execution, each agent acts independently based on its local observations, enabling decentralized decision-making and reducing communication requirements.

3) Handling Non-Stationarity: MADDPG is designed to handle non-stationarity in multi-agent environments, where agents' policies may change during training. By incorporating a centralized critic network that considers all agents' actions and observations, MADDPG can adapt to changing dynamics and maintain stability during training.

4) Policy Exploration and Exploitation: MADDPG combines the benefits of exploration and exploitation by utilizing the DDPG algorithm. DDPG employs an exploration policy, such as adding noise to the actions, to encourage exploration and discover new strategies. At the same time, it leverages the learned policies to exploit the most promising actions and maximize performance.

The proposed MADDPG is formulated using the Bellman equation, enabling the learning of a Q-function and a multi-agent-based policy. Similar to DDPG, a deep neural network (DNN) is utilized as the Value Function Approximation (VFA) in MADDPG. This approach integrates the actor-evaluation approach and multi-agent technique, making it suitable for handling model-free, high-dimensional, and continuous action spaces in multi-zone environments. The Q-Value function, representing the value function, is given by the following expression:

$$Q^\pi(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha[r_t(s_t, a_t) + \gamma \max_{a_{t+1} \in A} Q(s_{t+1}, a_{t+1}) - Q(s_t, a_t)] \quad (3.19)$$

$\pi \in \Pi$ represents the policy, which is a set of actions with a probability distribution. α

Algorithm 1 Multi-agent Deep Deterministic Policy Gradient

Input: Environment with M zones, and occupant behavior OB

Output: M sets of actor networks

Randomly initialize evaluation network $Q(s, a|\theta^Q)$ with weights θ^Q where $a = \{a^1, a^2, \dots, a^M\}$, and M actors $\mu(s|\theta^\mu) = \{\mu^1(s|\theta^{\mu,1}), \mu^2(s|\theta^{\mu,2}), \dots, \mu^M(s|\theta^{\mu,M})\}$ with weights $\theta^{\mu,m}$, where $m \in \{1, 2, \dots, M\}$

Initialize target-evaluation network Q' with weights $\theta^{Q'} \leftarrow \theta^Q$, and M target actors $\mu' = \{\mu'^1, \mu'^2, \dots, \mu'^M\}$ with weights $\theta^{\mu',m} \leftarrow \theta^{\mu,m}$ where $m \in \{1, 2, \dots, M\}$

Initialize experience replay buffer B

for episode = 1 to E **do**

 Initialize Ornstein–Uhlenbeck process (OU) for action exploration

 Receive initial observation state s_1

for $t = 1$ to T **do**

 Select action $a_t^m = \mu(s_t^m|\theta^\mu) + OU_t$ according to the policy network and exploration noise for zone m at t

 Apply stochastic-based dynamic occupant behavior based on Equation 3.16 to the environment

 Execute action a_t and observe reward r_t and new state s_{t+1}

 Store transition (s_t, a_t, r_t, s_{t+1}) in buffer B

 Sample a random mini-batch of M transitions (s^l, a^l, r^l, s^{l+1}) from B

 Set y^l based on Equation 3.20, where $\mu'(s^{l+1}|\theta^{\mu'})$ is the sets of M actors

 Update evaluation network by minimizing the loss based on Equation 3.21, where $a_l = \{a_l^1, a_l^2, \dots, a_l^M\}$

 Update the actor policy using the sampled policy gradient for all $\mu(\cdot)$ based on Equation 3.22

 Soft update the target evaluation network and all target actor networks:

$$\theta^{Q'} \leftarrow \tau \theta^Q + (1 - \tau) \theta^{Q'}$$

$$\theta^{\mu'} \leftarrow \tau \theta^\mu + (1 - \tau) \theta^{\mu'}$$

end for

end for

denotes the learning rate, and γ is the discount factor used for future reward considerations.

With Equation (3.19) as the foundation, the VFA of DRL can be formulated as follows:

$$y_l = r_l + \gamma Q(s_{l+1}, \mu(s_{l+1}|\theta^\mu)|\theta^Q) \quad (3.20)$$

where $l \in L$ denotes the index of mini-batch L , which is sampled from the experience replay buffer B . The parameters θ^μ and θ^Q correspond to the weights of the actor neural network $\mu(\cdot)$ and the evaluation network $Q(\cdot)$, respectively. The evaluation network $Q(\cdot)$ is updated

through the minimization of the loss function:

$$Loss = \frac{1}{M} \sum_t (y_l - Q(s_l, a_l | \theta^Q))^2 \quad (3.21)$$

To facilitate the exploration of the actor-network, the Ornstein-Uhlenbeck process is utilized¹⁹². The actor networks are updated using the policy gradient technique, which applies the chain rule to compute the gradient of the expected return with respect to the actor parameters, represented by the approximated loss of the distribution J :

$$\nabla_{\theta^\mu} J \approx \frac{1}{M} \sum_l \nabla_a Q(s, a | \theta^Q)|_{s=s_l, a=\mu(s_l)} \nabla_{\theta^\mu} \mu(s | \theta^\mu)|_{s_l} \quad (3.22)$$

The MADDPG algorithm introduces specific modifications tailored for multi-zone HVAC control:

- During initialization, M actors and target actors are established, each corresponding to a zone within the building.
- Each actor contributes the action for its respective zone during policy execution, resulting in a set of zone-specific actions.
- The stochastic occupant behavior, represented by Equation 3.16, is applied to the environment (EnergyPlus) with each execution, capturing variability in occupant actions.
- The outputs of all M actors and target actors serve as inputs to the target evaluation network.
- All actor networks are updated using the sampled policy gradient, allowing the agents to learn and improve their strategies based on the environment's feedback.
- Instead of returning a single-actor network, the well-trained multiple-actor networks are returned, providing zone-specific actions for each zone in the building.

It is essential to highlight that the MADDPG algorithm employs $2M + 2$ networks (with

M being the total number of zones) compared to the four neural networks used in single-agent DDPG. This increase in the number of networks leads to longer training times for the MADDPG algorithm compared to DDPG, as will be demonstrated in subsection 3.3.1.

To learn in the multi-agent environment better, we utilized Cyclical Learning Rates (CLR)¹⁹³ to enhance neural network training in MADDPG. CLR dynamically adjusts the learning rate during training, allowing it to increase and decrease within a single run. By cycling between upper and lower bounds, the network explores a wider range of learning rates, improving performance and convergence. CLR benefits neural network training by preventing instabilities and escaping saddle points. It facilitates faster traversal across the loss landscape, leading to better solutions. Therefore, implementing CLR in MADDPG optimizes the networks' generalization and optimization capabilities, while dynamic learning rate adjustment improves parameter space exploration, resulting in superior model performance.

3.3 Experiment and Case Studies

This section evaluates the proposed MADDPG-based HVAC control in a five-zone single-floor commercial building modeled in EnergyPlus. The building includes one interior and four exterior zones, with windows on all facades and glass doors on the north and south. A variable air volume HVAC system with direct expansion cooling and gas heating serves all zones (see Table 3.6). One year of 15-minute interval weather data from Tampa, Florida (Figure 3.10), and a Time-of-Use pricing scheme from PG&E (EToU-E6) are used to reflect realistic environmental and economic conditions. We compare three strategies: (1) rule-based control using only outdoor temperature, (2) single-agent DDPG with occupant behavior, and (3) the proposed MADDPG with zone-specific control. This setup allows us to assess the benefits of multi-agent coordination and human-aware control in reducing energy costs and enhancing comfort.

Table 3.6: HVAC PARAMETERS

Feature	Parameters	Setting
General	High-Speed/Low-Speed Sensible Heat Ratio	0.75
	Nominal Capacity (W)	3500
Cooling	Rated Cooling COP (W/W)	3.0
	Internal Static Air Pressure (Pa)	450
Heating	Burner Efficiency	0.98
	Nominal Capacity (W)	3500
	Fan Total Efficiency	0.7
	Pressure Rise (Pa)	600
Fan	Maximum Flow Rate (m^3/s)	3.0
	Power Minimum Flow Fraction	0.25
	Motor Efficiency	0.9
	Motor In Air-stream Fraction	1.0

Simulation Setup

A single-floor rectangular building with five zones (containing one interior and four exterior zones) is used to simulate the practical building, which features windows on all four facades and glass doors on the south and north facades. The HVAC system incorporates a packaged variable air volume system with direct expansion cooling coils and gas heating coils, serving the five zones. The HVAC parameters are listed in Table 3.6. For the simulation, the one-year weather dataset of Tempa, Florida, USA, is employed, providing detailed measurements at 15-minute intervals. This dataset includes all needs for the external state, as shown in Figure 3.10. Time-of-Use (ToU) electricity tariff chosen is the Pacific Gas & Electric EToU-E6, which consists of three price levels: the base, shoulder, and peak prices, represented by white, light grey, and grey colors, respectively, with unit costs of \$0.244/kWh, \$0.32/kWh, and \$0.436/kWh in Figure 3.12. Note that the current validation in this study utilizes EnergyPlus as the whole-building simulator. The MADDPG algorithm is implemented in Python 3.8.10 with BCVTB serving as the interface. The computational platform used for the experiments is a PC equipped with an Intel(R) Core(TM) i7-4790 CPU and 8 GB RAM with Windows Subsystem for Linux (WSL) Version 2.

To clearly show the efficacy of the proposed approach, we compare the proposed MADDPG with 1) a single-agent DDPG counterpart (control with occupant behavior) and 2) a




Figure 3.10: One-year Tampa weather profile

Table 3.7: MADDPG network information

Actor Network	# of Neurons	Activation
Input	Input(shape=(# of state))	
Hidden 1	Dense(128, 256)	Relu
Hidden 2	Dense(256, 512)	Relu
Output	Dense(128, # of action)	Sigmoid
Critic Network	# of Neurons	Activation
State Input	Input(shape=(# of state))	
State Hidden	Dense(256, 256)	Relu
State Output	Dense(256, 256)	Relu
Action Input	Input(shape=(# of action))	
Action Output	Dense(256, 64)	Relu
Concatenate	Concatenate([State Output, Action Output])	
Hidden 1	Dense(Concatenate, 512)	Relu
Hidden 2	Dense(512, 256)	Relu
Output	Dense(1)	

rule-based control scheme (control without occupant behavior), which is discussed in [194;195](#). The rule-based control highlights its reliance on outside temperature for HVAC electricity cost reduction in a whole-year hourly time step simulation. This method optimizes electricity cost by allowing slightly wider temperature ranges while ensuring occupant comfort which is deployed in realistic buildings.




Figure 3.11: Average episode reward during training of DDPG and MADDPG

Result

This subsection presents the simulation results for three compared methods: Rule-based method (from¹⁹⁴, DDPG (from¹⁷) and MADDPG (from Algorithm 1)), all applied to the testbed of a 5-zone building with 15 minutes time interval simulation model. The neural network information, including input, hidden, and output layers for actor and critic networks, is shown in Table 3.7. Table 3.8 displays the total training time for DDPG and MADDPG, along with the single-run execution times for the rule-based method, DDPG, and MADDPG. It is important to note that the rule-based method does not utilize any deep neural network for computation, making it significantly faster in execution time compared to the learning-based approaches. The single-run execution time refers to the time taken for a one-time decision-making process. Specifically, DDPG takes 5 hours to train and 6.25 ms for one-time decision-making, while MADDPG takes 32 hours (5.4 times longer than DDPG) for training and 29.17 ms (3.66 times longer than DDPG) for one-time decision-making. Additionally, the reward function plot (Figure 3.11) displays the training rewards for the episodes. It is evident that MADDPG exhibits better reward returns after approximately 25 episodes of training compared to DDPG. It is important to note that MADDPG takes longer to train compared to DDPG due to the presence of more neural networks. However, this extended training time leads to higher reward returns after the training process.

Table 3.8: Training and execution time comparison

Item	Rule-based	DDPG	MADDPG
Training time	N/A	5 hr	32 hr
Single-run execution time	0.01 ms	6.25 ms	29.17 ms

In the following, we present the simulation results for three days, encompassing occupant behavior, decision results for all zones, and the whole building’s electricity cost with average PMV-PPD. Figure 3.12 provides a visualization of the three-day simulation of occupant behavior, illustrating fluctuations in metabolic rate, occupant presence, and clothing adjustments across all zones. The base price, shoulder price, and peak price are visually represented in white, light grey, and grey colors, respectively. It’s essential to emphasize that our simulation environment is grounded in stochastic behavior modeling rather than a fixed schedule. This deliberate choice allows us to accurately capture the significant variability in occupant thermal comfort experiences over time. Furthermore, our approach offers flexibility in adjusting stochastic modeling parameters to align with real-world scenarios, leveraging insights from historical datasets¹⁹⁶. It’s worth noting that in our modeling, we assume that clothing adjustments play a relatively minor role in the overall variations. Detailed schedules for clothing behavior, metabolic rate, and occupant presence are elaborated upon in subsection 2.2 for a more comprehensive understanding.

Figure 3.13 compares decision results, including heating and cooling setpoints of HVAC, for the three methods, along with room and outside temperatures. The weather dataset contains a wealth of information, with the outside temperature being the major factor affecting the room temperature. In the rule-based simulations, the control range (between heating and cooling setpoints) is narrow, as it solely considers the outside temperature for the control strategy. In addition, the heating and cooling set points for our rule-based method are derived from [63,64], in which high cooling and heating set points are observed in the early morning. Therefore, the rule-based method is typically less adaptive than DDPG and MADDPG which considers occupant behavior due to limited control range. As a result, the room temperature is controlled within this range without considering occupant behavior. In




Figure 3.12: Three-day behavior simulation result for all zones (base, shoulder, and peak prices are represented by white, light grey, and grey)

contrast, both DDPG and MADDPG show a wider range of temperature control as they take into account the occupant’s behavior. Specifically, DDPG cools the zone once on Day 1 around 4 pm and requires no control on Day 2, while on Day 3, it cools the zone from 9 am to 6 pm. On the other hand, MADDPG adapts to the outside environment on Day 1 and Day 2. On Day 3, it maintains climate control from 9 am to 6 pm but undergoes a two-hour shift from the grey area (high electricity price) to the white area (based price). Furthermore, there are two pre-heating phase on Day 1 around 8 am and Day 3 from about 5 am to 8 am to raise the zone temperature due to its base price period, leading to lower electricity costs while ensuring similar thermal comfort levels. The decisions made by MADDPG demonstrate a more adaptive and cost-effective approach, taking into account both the outside environment and occupant behavior, resulting in better energy efficiency and comfort management. Figure 3.14 presents a visual representation of the three-day simulation results for Zones 2 to 5, offering insights into the strategies employed by the three methods. It’s worth noting that the single-agent method demonstrates consistent actions across all zones, while MADDPG exhibits adaptive actions tailored to each zone’s unique requirements. Importantly, MADDPG effectively avoids peak ToU hours, as evidenced by its morning actions around 7 am






Figure 3.13: Results of Zone 1 over a three-day period, with the Rule-based method (top), DDPG (middle), and MADDPG (bottom). The red line represents heating setpoints, the blue line indicates cooling setpoints, the black line represents the zone’s average temperature, and the green dashed line represents the PMV-PPD value, while the light blue line depicts the outside temperature.

or 8 am, leading to cost savings. Additionally, MADDPG maintains a high level of thermal comfort, as reflected in the lower PMV-PPD values compared to the single-agent method.


Figure 3.15 illustrates the evaluation metrics for the three methods, including HVAC whole building cost and average PMV-PPD in the five zones. As expected, the rule-based method incurs the highest electricity cost due to its narrow control climate range, which limits zone temperatures and leads to increased electricity consumption. This highlights the crucial importance of considering occupant behavior in the simulation. Regarding thermal comfort, the PMV-PPD values for the first two days are similar across all methods. However, on the last day, both the rule-based and MADDPG methods control the temperature at 25




(a) Zone 2



(b) Zone 3



(c) Zone 4



(d) Zone 5

Figure 3.14: Results of Zone 2, 3, 4, and 5 over a three-day period. The red line represents heating setpoints, the blue line indicates cooling setpoints, the black line represents the zone's average temperature, and the green dashed line represents the PMV-PPD value, while the light blue line depicts the outside temperature.




Figure 3.15: Three days simulation of electricity cost and thermal comfort result (PMV-PPD) for the whole building.

$^{\circ}\text{C}$, while DDPG maintains it at 24°C . This slight difference in temperature control results in varied performance in thermal comfort, as evident in the PMV-PPD results. Notably, the pre-heating on the last day contributes to higher electricity costs in MADDPG but provides more comfortable thermal control compared to DDPG. The findings underscore the significance of incorporating occupant behavior into the simulation. By accounting for occupant preferences and adjusting temperature settings accordingly, MADDPG optimizes both electricity costs and thermal comfort, offering a more adaptive and efficient HVAC control strategy for the building. This approach provides a more personalized and responsive solution that takes into account occupant comfort while achieving energy savings.

An interesting finding is that the single-agent DDPG demonstrates a nearly identical control action for all zones, as depicted in Figure 3.16. The control pattern remains consistent over the time horizon, with a relaxation of the control range during nights and a narrowing of the range during daytime. However, a crucial limitation of this approach is that it uses similar temperature setpoints for all zones, therefore covering each other plot line in the top (rule-based) and middle (DDPG) in Figure 3.16, which does not reflect the actual real-world conditions in the simulations. Figure 3.17 displays the decision distribution over one




Figure 3.16: One-month control patterns of HVAC for all zones

month of simulation. It becomes evident that MADDPG provides more options in terms of temperature settings, while DDPG and rule-based control exhibit a more limited range of options in the simulation results. This phenomenon indicates that MADDPG employs a more detailed and zone-specific control strategy, while DDPG applies the similar control approach to all five zones. Consequently, DDPG fails to capture the distinct differences between zones within the building, resulting in different simulation outcomes. This observation highlights the advantage of using MADDPG, which allows for more personalized and adaptive control strategies for each zone, leading to enhanced performance and efficiency in the overall HVAC management within the building.

To comprehensively represent the simulation results at the long-term level, one-year worth of weather data was employed in this study from the Tampa dataset from January 1 to December 31. Table 3.9 presents a one-year comparison of the three methods, showcasing




Figure 3.17: One-month HVAC decisions histogram comparison of whole building

the average daily electricity cost (\$/day) and average PMV-PPD across all five zones during the year-long simulation. The rule-based control approach exhibits the highest electricity cost but maintains the best thermal comfort levels. However, this advantage in thermal comfort comes at the cost of overlooking occupant behavior, as the rule-based control solely relies on outside temperature, leading to non-optimal energy usage. In contrast, MADDPG stands out with exceptional performance, achieving the lowest electricity cost, which is 51.09% lower than the rule-based control and 4.36% lower than DDPG. Importantly, MADDPG not only excels in cost-efficiency but also enhances thermal comfort levels compared to DDPG. Its capacity to strike a harmonious balance between cost-effectiveness and occupant comfort positions it as a highly promising and superior choice for HVAC control in buildings.

Table 3.9: One-year comparison between different methods

Item	Rule-based	DDPG	MADDPG
Average Daily Electricity Cost (\$ / Day)	216.77	110.83	106.02
Avg. PMV-PPD (%)	10.36	11.43	10.39

MADDPG achieves this through a multi-agent control strategy, which takes into account the unique thermal transfer characteristics and occupant behaviors in different zones. The control strategy, as in DDPG, would not capture these differences and result in suboptimal multi-zone thermal control, as shown in Figure 3.17. MADDPG’s incorporation of multi-

agent techniques, training the agents to cooperate and coordinate within the environment, leads to better results in this study. The approach effectively optimizes both electricity cost and thermal comfort, demonstrating its superiority over the rule-based and single-agent DDPG methods. The consideration of occupant behavior and zone-specific thermal dynamics in MADDPG results in more adaptive and efficient HVAC control, offering a more desirable outcome for building occupants and energy management.

3.3.1 Discussion

The findings from the figures and table in this study highlight the effectiveness of the MADDPG approach in multi-zone HVAC control. MADDPG operates by making preemptive decisions to avoid peak ToC price periods while ensuring consistent thermal comfort. Its capability to customize HVAC setpoints individually for different zones, as opposed to applying uniform actions across all zones, significantly enhances its ability to reduce electricity costs compared to the rule-based approach. Furthermore, MADDPG exhibits superior thermal comfort performance when contrasted with the single-agent DDPG method. In addition, over the course of one year, MADDPG consistently outperformed the single-agent DDPG methods in terms of average daily electricity cost, while maintaining superior thermal comfort levels compared to DDPG, striking a balance between cost reduction and occupant satisfaction. This improved performance is attributed to MADDPG’s ability to adaptively design HVAC setpoints based on zone-specific characteristics and occupant behaviors, effectively utilizing multi-agent techniques to optimize both electricity cost and thermal comfort.

Regarding the issues of similar actions for all zones in DDPG, we realize that during the training phase, all agents participate in centralized learning, sharing a global value function that captures interactions and dependencies between agents and the environment. However, this collective approach does not distinguish individual agents, treating those agents as a unified entity. Consequently, the shared knowledge acquired through single-agent DDPG tends to result in similar actions among agents, particularly when they face similar environmental conditions such as a testbed of multi-zone HVAC building. However, MADDPG

is tailored for multi-agent environments, which facilitates communication and coordination among multiple agents, enabling them to learn distinct policies and make varied decisions based on observations and other agents' actions.

This study serves as a proof of concept for applying MADRL to multi-zone thermal control while considering dynamic occupant behavior. One of the critical avenues for future research is the validation of this approach in real-world environments. The validation process entails several essential steps to ensure the algorithm's practicality and effectiveness in realistic scenarios. Firstly, it involves establishing the necessary hardware and software infrastructure, gathering real-world data, including occupant behavior and weather variations, and constructing a simulation model for virtual testing using the EnergyPlus simulator. Following this, the MADDPG algorithm undergoes offline training before its implementation within a realistic building system. Continuous system operation over an extended period facilitates ongoing performance monitoring of MADDPG, including model calibration for building models and parameter updates based on user feedback. Finally, tracking electricity consumption and occupant comfort, followed by comprehensive data analysis, evaluates MADDPG's performance compared to baseline approaches in realistic buildings.

3.4 Summary

This chapter first validates the influence of dynamic occupant clothing behavior on building energy management through HVAC control optimization. Building on this, we present a multi-zone HVAC energy management strategy utilizing the Multi-Agent Deep Deterministic Policy Gradient (MADDPG) algorithm. This approach aims to minimize electricity costs while maintaining thermal comfort across individual zones. By capturing zone-specific thermal dynamics and dynamic occupant behavior, the MADDPG framework demonstrates superior performance compared to rule-based control and single-agent DDPG methods.

Simulation results reveal that MADDPG consistently achieves the best trade-off between

energy efficiency and occupant comfort, adapting intelligently to variations in external environmental conditions and internal occupant behavior. The multi-agent structure enables personalized, zone-level control, fostering greater flexibility and responsiveness in HVAC operations. Over a year-long simulation, the proposed method not only delivers the lowest electricity cost but also maintains high levels of occupant satisfaction. Furthermore, the integration of stochastic occupant modeling into the control logic highlights the potential for more realistic and adaptive building energy management systems. Future work will advance in two key directions:

Enhanced Occupant Behavior Sensing: We aim to develop and deploy a comprehensive sensing infrastructure capable of capturing detailed, real-time occupant behavior. This includes using camera-based computer vision techniques to classify occupant activities such as sitting, standing, walking, working, or being idle, and integrating sensors or image-based inference to estimate clothing insulation levels (e.g., layers, garment type, or insulation estimates). These high-resolution inputs will allow the energy management system to adaptively optimize HVAC control in response to dynamic, individualized comfort needs.

Edge AI for Real-time Building Control: We plan to leverage the NVIDIA Edge Computing platform (e.g., NVIDIA Jetson series) for on-device training and inference of machine learning models. These platforms support online learning capabilities that allow ML models—especially on-policy reinforcement learning agents—to continuously adapt using streaming data from the building environment. This approach supports low-latency, privacy-preserving, and robust control decisions, enabling truly intelligent and autonomous energy management systems in real-world deployments.

Chapter 4

Machine Learning in Food-Energy-Water Nexus Energy Management

Having discussed energy management within power systems and buildings in previous chapters, this chapter focuses on agriculture—a sector fundamental to human existence, ensuring food security and underpinning global population sustainability. Agriculture not only directly supports human life but also significantly impacts the economic stability and national security of all nations. Thus, efficient energy management in agricultural systems is critically important, particularly considering the sector’s substantial consumption of resources and energy.

However, agricultural energy management presents unique complexities and challenges due to diverse operational demands, fluctuating environmental conditions, and the intricate interdependencies within agricultural systems. The Food-Energy-Water (FEW) Nexus encapsulates these interdependencies, highlighting the critical interactions between food production, energy consumption, and water resource management. To align with global sustainability goals, it is essential to adopt integrated and efficient approaches within this

nexus, promoting the use of renewable resources and reducing environmental footprints. One promising pathway towards sustainable agricultural energy management is through the integration of green ammonia production. Green ammonia serves a dual purpose: it is essential as a sustainable fertilizer to support crop yields, and it acts as a highly efficient energy carrier, particularly in renewable energy systems. Despite its potential, effective and feasible planning and operational strategies for integrating green ammonia production within agricultural energy systems remain largely unexplored.

To address these gaps and the inherent complexity within the FEW Nexus, this chapter introduces a novel approach leveraging Deep Reinforcement Learning (DRL). The proposed DRL framework intelligently manages resources, optimizes renewable energy usage, and strategically integrates green ammonia production to enhance farm sustainability, efficiency, and profitability.

In this context, the key contributions of this chapter are summarized as follows:

- We propose a comprehensive planning methodology for the integrated design and optimal sizing of renewable energy systems combined with green ammonia production within the FEW Nexus. This addresses a significant knowledge gap and provides foundational insights for future sustainable agricultural energy system planning.
- We develop an innovative DRL-based farm energy management framework that simultaneously optimizes renewable energy deployment, green ammonia production, and resource allocation strategies. The framework ensures economic efficiency, environmental sustainability, and resilience in local agricultural systems, offering practical solutions to complex FEW Nexus challenges.

4.1 Planning Energy Systems for Farm-Scale Ammonia Production

In the transition toward sustainable agriculture, farms have emerged as eco-friendly pioneers, harnessing clean hybrid wind and solar systems to improve farm performance. A concern in this paradigm is the effective sizing of renewable energy systems to ensure optimal energy use within budget considerations. This research focuses on optimizing renewable energy sizing in small-scale ammonia production to meet specific farm demands and enhance local resilience, emphasizing the interplay between environmental and economic factors. These findings promise increased energy efficiency and sustainability in this innovative agricultural sector.

Depletion of water resources, climate change impacts, and food market challenges pose significant threats to local agricultural economies and global food supplies¹⁹⁷. The adoption of locally harnessed renewable energy within the agricultural sector emerges as a promising solution to fortify local economies and global food production and reduce carbon emissions linked to global warming, all while adhering to stringent budget considerations. This underscores the urgency of exploring renewable energy’s potential in meeting agricultural needs, with a specific focus on small-scale ammonia production, water supply, and water treatment¹⁰¹. Fossil fuel-based large-scale ammonia production has seen wild price swings (\$400 to \$1600 in 2022), and increasing global demand is expected¹⁹⁸. Local green ammonia provides a stable cost, if sometimes more expensive, alternative to large-scale fossil fuel-based ammonia. As a way to decarbonize agriculture within strict capital budget limitations, installing renewable energy systems on farms serves as a potentially promising solution within the food-energy-water nexus.

Farm Hybrid Renewable Energy Systems (HRESs) (Fig. 4.1) could form a vital part of the larger power distribution network and could be pivotal for achieving sustainable local farms. The system considered in this work integrates solar, wind, and battery facilities within the farm infrastructure. Its seamless connection to the power grid enables bidirec-




Figure 4.1: Illustration of farm hybrid renewable energy system.

tional electricity exchange, significantly enhancing the farm’s overall energy resilience. While renewable energy integration has made notable strides in residential and industrial contexts in recent years, the specific needs of farm energy systems have been somewhat overlooked in the existing literature. Moreover, the existing body of work explores the optimization of renewable energy across diverse contexts. Reference¹⁹⁹ underscores the benefits of appliance scheduling and demand response strategies, emphasizing their potential to enhance cost-effectiveness and reliability in local farms with dairy systems. Reference²⁰⁰ explores cost-effective HRES tailored to regional agricultural needs through the hybrid optimization of multiple renewables. Reference²⁰¹ explores cost-effective renewable energy solutions, with a particular emphasis on demand management in farm drip irrigation systems with pumps. It is worth noting, however, that these studies predominantly center around electricity considerations, overlooking the crucial aspect of ammonia production, which holds significance in agricultural practices within local farms. In addition, considering that financial resources are often constrained for many agricultural operations, it becomes imperative to incorporate this constraint into the optimization and analysis processes.

The contribution of this subsection lies in the effective, economical sizing of HRESs for small-scale ammonia production, aligning renewable energy with farm-specific demands. By considering farm energy usage, it shows energy utilization, emphasizing the interplay between ammonia production, water resources, and energy demand. The study’s insights enhance HRES planning while adhering to strict HRES capital budget limitations and ultimately

improving energy efficiency and reducing carbon emissions within the agriculture sector. The simulations reveal significant cost savings in farm electricity consumption, with renewable energy integration resulting in a remarkable reduction in electricity costs compared to farms solely reliant on power grid resources over a 15-year simulation period.

4.1.1 Modeling Farm-scale Ammonia Production

We propose a model developed in the Stella Architect software²⁰² that uses graph theory and considers energy, small-scale ammonia, farm, and water systems illustrated in Fig. 4.2, which shows the interrelationship among the components¹. The figure is the abstraction of a multilayer network mapped to a single-layer graph. Each module in the system is shown with a node, and edges identify the interaction among modules. Furthermore, there are different types of interactions in the system. Each type of interaction is illustrated with a different color. For example, interactions between the farm energy system (microgrid) and other systems are depicted by color-coded links. The hubs of electricity, water, ammonia, and farms are represented by orange, blue, yellow, and green colors, respectively. The black and pink links symbolize energy transfers in/out with electricity hubs. In contrast, the water module exclusively interacts with the ammonia and farm modules, illustrated by the distinct blue links denoting the flow of water. This color scheme provides a clear visualization of interdependencies and allows for the assessment of how a failure or disruption in one module may impact others in the system²⁰³.

In this model, electricity is generated with a combination of solar panels and wind turbines and consumed in other modules. The largest consumer in the model is the ammonia module. The water module uses electricity to pump water from wells and surface water, while the farm module mostly uses electricity for irrigation. The ammonia module²⁰⁴ is dependent on the farm energy and water modules to produce ammonia, while the water module depends on the farming demand and ammonia demand. The amount of electricity and water consumed in the ammonia module is fixed and depends on the capacity of the ammonia production, while the electricity usage is periodic in the farm module and increases




Figure 4.2: A food-energy-water abstract model in the Stella Software¹. No. PV panels and No. Turbines determine the PV capacity and wind capacity.




Figure 4.3: Flowchart of the proposed framework.

during the irrigation seasons. Therefore, the electricity consumption is also periodic in the water module. We use Stella for our simulation environment, which has the capability of simulating discrete and continuous models.

4.1.2 Problem Formulation of Planning in Farm Energy

In this subsection, we delve into formulating the problems and the mathematical approaches employed to address each model. The objective function is

$$\min_{PV^{Cap}, Wind^{Cap}, Battery^{Cap}} \left(\sum_{t=0}^T \frac{O_t^{total}}{(1+r)^{yr}} \right) \quad (4.1)$$

where $PV^{Cap}, Wind^{Cap}, Battery^{Cap}$ are installed capacity of solar energy, wind energy, and battery storage; the total electricity cost, O_t^{total} , for the sizing renewable energy problem in the farm.

$$O_t^{total} = CC_{t=0} + OC_t + PC_t + GC_t - REC_{yr} \quad (4.2)$$

In the given equations, the time index t operates on an hourly scale, while the parameter yr counts in yearly intervals, with yr incrementing by 1 for every 8760 hours over time horizon T . Equation (4.1) incorporates the discount rate r . Equation (4.2) features several components, including CC (capital cost), OC (operation and maintenance cost), PC (penalty cost for system imbalances), and GC (cost associated with electricity exchange between the farm and grid). Additionally, the equation includes REC , representing renewable energy certificates, which are considered an annual reward for renewable energy. As such, REC is subtracted from the total electricity cost. The capital cost calculation is as follows:

$$CC_{t=0} = CC^{PV} * PV^{Cap} + CC^{Wind} * Wind^{Cap} + CC^{BESS} * BESS^{Cap} \quad (4.3)$$

$$CC_{t=0} \leq Budget \quad (4.4)$$

$$PV^{Cap} \times l^{PV} + Wind^{Cap} \times l^W \leq Land \quad (4.5)$$

where $CC_{t=0}$ represents the initial capital cost incurred at the time of establishing the HRES on the farm. This encompasses expenses related to setting up the PV solar, wind, and battery energy storage system (BESS) denoted as CC^{PV} , CC^{Wind} , and CC^{BESS} in \$/kW,

respectively. Furthermore, we consider the installed capacities (in kW for PV and wind, kWh for battery) of these systems: PV^{Cap} , $Wind^{Cap}$, and $BESS^{Cap}$. To ensure financial feasibility and adhere to HRES budget considerations, Equation (4.4) is introduced. This constraint reflects the investment of staying within HRES budget expectations, as some farm stakeholders typically possess limited financial resources for renewable energy. Equation (4.5) imposes restrictions on the land space allocated for the HRES installation within the farm. Here, l^{PV} and l^W denote the coefficients reflecting space requirements for the hybrid energy system (in m^2/kW). The rest of the components of total electricity cost are as follows:

$$\begin{aligned} OC_t = & Om_{yr}^{PV} * PV^{Cap} + Om_{yr}^{Wind} * Wind^{Cap} \\ & + Om_{yr}^{BESS} * BESS^{Cap} \end{aligned} \quad (4.6)$$

where OC_t is the operation and maintenance cost for the hybrid energy system at time t . Parameters Om_{yr}^{PV} and Om_{yr}^{Wind} are operation and maintenance unit costs (\$/kW/yr) each year for PV solar, wind, and Om_{yr}^{BESS} is for battery systems (\$/kWh/yr), respectively.

$$REC_{yr} = \left(\sum_i^{i=\{PV, Wind\}} \sum_{t=yr_{begin}}^{yr_{end}} P_t^i \right) \times u_{rec} \quad (4.7)$$

In Equation (4.7), we show the incorporation of renewable energy certificates (RECs) into our framework. RECs are valuable, tradable commodities that symbolize both the environmental and renewable attributes of the MWh of electricity produced from a renewable energy source. Specifically, Equation (4.7) illustrates a situation where renewable energy generation exceeds the local farm's energy demand, resulting in surplus electricity being sold back to the grid. An assumption is that the farm's owner sells annually. Within this equation, several variables come into play: P_t^i signifies the renewable energy generation (kW) from source i at time t , and u_{rec} stands for the unit price associated with RECs, contributing

to the extra revenue from renewable energy within the framework.

$$GC_t = \begin{cases} P_t^{G2F} \times u_t^{buy}, & \text{if } P_t^{grid} > 0 \\ P_t^{F2G} \times u_t^{sell}, & \text{if } P_t^{grid} < 0 \end{cases} \quad (4.8)$$

Equation (4.8) delineates the grid cost, GC_t , implications associated with the exchange of electricity between the power grid and the farm. In this context, P_t^{G2F} refers to the power (kW) transferred from the grid to the farm at time t , signifying a power purchase action. The variable u_t^{buy} represents the unit price (\$/kW) of purchased electricity during this exchange. Conversely, P_t^{F2G} represents the power (kW) transferred from the farm to the grid at time t , reflecting a power selling action, while u_t^{sell} denotes the unit revenue (\$/kW) obtained from selling excess electricity back to the grid. We conservatively assume the price of power sold back to the grid is 10% of the purchase price from the grid. This equation captures the dynamics of power exchange and its associated costs within the framework.

Equation (4.9) shows the penalty cost, PC_t , caused by the unbalanced system, i.e., oversupply or power shortage, where c_{pc} is the unit price (\$/kW) for the penalty. Here, D_t^{total} is the total electricity demand at time t .

$$PC_t = |D_t^{total} - \sum_{i=\{PV, Wind, BESS, Grid\}} P_t^i| \times c_{pc} \quad (4.9)$$

Next, charging and discharging constraints regulate the farm battery behavior in Equations (4.11) and (4.10), where $\overline{P}_t^{dis}$ and $\overline{P}_t^{ch}$ represent the maximum discharge and charge rates, respectively. These constraints ensure that the battery's charge and discharge rates do not exceed their limits. Parameters $\overline{Disch}$ and $\overline{Char}$ are the maximum discharging/charging rates in the battery but are also limited by the headroom of state of charge (SOC) or the remaining energy in the battery.

$$\overline{P}_t^{dis} = \min(\overline{Disch}, SOC_t) \quad (4.10)$$

$$\overline{P}_t^{ch} = \min(\overline{Char}, BESS^{Cap} - SOC_t) \quad (4.11)$$

$$B_t^{flag} = \begin{cases} discharge, & \text{if } D_t^{total} \geq (P_t^{PV} + P_t^{Wind}) \\ charge, & \text{if } D_t^{total} < (P_t^{PV} + P_t^{Wind}) \end{cases} \quad (4.12)$$

Here, B_t^{flag} determines whether the battery should be in a charge or discharge state, depending on the relationship between the total energy demand and the combined energy generation from PV and wind sources.

$$P_t^{BESS} = \begin{cases} \min(\overline{P}_t^{dis}, |D_t^{total} - P_t^{Wind} - P_t^{PV}|), & \text{if } B_t^{flag} = discharge \\ -\min(\overline{P}_t^{ch}, |P_t^{Wind} + P_t^{PV} - D_t^{total}|), & \text{if } B_t^{flag} = charge \end{cases} \quad (4.13)$$

Then, Equation (4.13) indicates the battery's power in charging or discharging. Variable P_t^{BESS} represents the power aggregation of the battery. It is determined based on whether the battery should be in a charge (with a “-” sign) or discharge (with a “+” sign) state and ensures the battery's power output aligns with the energy demand.

$$\hat{P}_t^{Grid} = D_t^{total} - D_t^{slack} - P_t^{PV} - P_t^{Wind} - P_t^{BESS} \quad (4.14)$$

The grid power characteristic is defined in Equation (4.14), where $\hat{P}_t^{Grid}$ represents the expected power exchanged with the grid, and D^{slack} is the slack variable in the energy balance equation. It is calculated as the difference between the total energy demand and the combined energy generation from PV, wind, and the battery.

$$P_t^{grid} = \begin{cases} \min(\overline{P}_{G2F}, |\hat{P}_t^{Grid}|), & \text{if } \hat{P}_t^{grid} > 0 \\ -\min(\overline{P}_{F2G}, |\hat{P}_t^{Grid}|), & \text{if } \hat{P}_t^{grid} < 0 \end{cases} \quad (4.15)$$

Equation (4.15) serves as an upper bound for power exchange during grid interaction. Specifically, their values are determined within limits set by $\overline{P}_{G2F}$ and $\overline{P}_{F2G}$ to ensure bound-

Table 4.1: Renewable Energy Parameters

Technology	PV	Wind	BESS
Capital cost	\$1220/kW	\$1015/kW	\$500/kWh)
O&M cost (per year)	\$35/kW	\$45/kW	\$35/kWh
Land requirement	28 m ² /kW	240 m ² /kW	-

aries in the exchange between the two entities, from grid to farm and from farm to grid, respectively.

Together, Equations (4.1)-(4.15) establish the objective function, constraints, and variables integral to the optimization problem for sizing renewable energy and battery systems, taking into account cost considerations, balance requirements, and grid interactions. Figure 4.3 shows the flowchart of the proposed budget-considering approach in the farm-based HRES sizing problem. Weather data, technology cost, budget, and price information are inputs to the framework. To address this optimization challenge, the study employs the differential evolution algorithm as the chosen optimization tool²⁰⁵.

4.1.3 Preliminary Results

Simulation Setup

The simulation setup parameters, as shown in Table 4.1, encompass capital cost and decision range of PV, wind, and battery capacity. These parameters are sourced from the commercial technology NREL Annual Technology Baseline (ATB)²⁰⁶. The battery parameters are derived from reference²⁰⁷.

The electricity pricing model utilized in this analysis is a Time-of-Use (ToU) structure. We assume the selling price to the grid is set at 10% of the purchase price. The ToU pricing is from Pacific Gas & Electric (PG&E) rate EToU-E6 includes three tiers: a Base level priced at \$0.244/kWh, Peak A at \$0.32/kWh, and Peak B at \$0.436/kWh. Peak A rates are applicable during hours 11-13, 20, and 21, while Peak B rates are in effect from hours 14-19, and Base rates are at all other times. Additionally, the study assumes a REC price of \$30/MWh. Furthermore, it is based on the assumption that five acres of land space are

available for the installation of renewable energy systems on the farm. The assumption of land space per power is listed in Table 4.1. The parameters for ammonia and water are sourced from reference²⁰⁸, in which the ammonia plant requires 410 kW, and the water system requires 25 kW when running.

Simulation Results

The simulation results uniquely address the challenge of sizing renewable energy for small-scale ammonia production (We assume 1 ton per day (TPD) in this case study.) To assess the effectiveness of our budget-constrained approach, we have designed several scenarios:

- Baseline: this case serves as a reference scenario, simulating a farm relying solely on electricity from the grid. It showcases the avoided energy cost, environment, and sustainability benefits without renewable energy integration.
- Ten scenarios that span HRES capital budgets from \$100k to 1 million with HRES optimization: these scenarios evaluate this method’s effectiveness for optimizing HRESs in farm energy systems, emphasizing the reduction of grid electricity dependence. Note that budget allocation is the portion of the overall budget that is set aside for HRES capital investments.

Table 4.2 provides a cost comparison across various HRES budget scenarios for implementing an HRES for small-scale ammonia production. In the Baseline (no HRES budget allocation) scenario, where the farm relies solely on grid electricity, the total cost over 15 years amounts to approximately \$2.48M (million). Subsequent HRES capital budget scenarios ranging from \$100k to \$1M demonstrate the effectiveness of the proposed HRES optimization. As the HRES capital budget increases, cost savings become evident. For instance, with a \$100k HRES budget, the total cost drops to about \$2.13M, which reduces the total cost 13% compared to the Baseline scenario, achieved through the installation of 50 kW of PV, 40 kW of wind, and no battery. As HRES budgets grow, the farm invests in larger PV and wind systems, along with battery storage, leading to substantial cost reductions. Moreover, with

Table 4.2: Comparison of HRES sizing between scenarios with 1 TPD ammonia plant (“Cap.” is Capacity; “Total Cost” is defined in Eq. (4.2))

Allocated Budget for HRES	Total Cost (\$)	PV Cap. (kW)	Wind Cap. (kW)	Battery Cap. (kWh)
Baseline (\$ 0)	2.48M	0	0	0
\$ 100k	2.13M	50	40	0
\$ 200k	1.86M	120	60	10
\$ 300k	1.64M	230	50	10
\$ 400k	1.45M	340	40	10
\$ 500k	1.24M	450	30	10
\$ 600k	1.13M	550	20	30
\$ 700k	1.04M	630	10	40
\$ 800k	975k	720	0	90
\$ 900k	975k	720	0	90
\$ 1M	975k	720	0	90




Figure 4.4: Illustration of farms with 1 TPD ammonia plant with its budget allocation for HRES and impact on total energy costs.

HRES budgets ranging from \$800k to \$1M, the total electricity cost remains constant at approximately \$975k for 15 years. These scenarios allocate the HRES budget to install 720 kW of PV capacity and a 90 kWh battery. However, scenarios with \$900k and \$1M HRES budgets did not invest their entire allocation.

Figure 4.4 illustrates the dynamic changes caused by various HRES budget allocations in the components of the total cost. As the HRES capital budget increases, there is a notice-

able decrease in electricity purchased from the grid, denoted as “Grid cost.” It’s important to emphasize that the cost associated with electricity purchases plays a dominant role in the total cost throughout the 15-year simulation for small-scale ammonia production farms. In addition, beyond a \$680k HRES budget, the capital cost remains fixed, signifying maximal land space utilization (five acres). Therefore, further investment doesn’t reduce the total energy cost. These scenarios underscore the significant cost-effectiveness and sustainability advantages achieved by optimizing HRESs for small-scale ammonia production while adhering to HRES budget considerations.

4.1.4 Preliminary Summary

This section offers a unique view of optimizing renewable energy system sizing for small-scale ammonia production. It prioritizes alignment with farm-specific needs and strict HRES budget considerations. Considering farm energy usage with Stella architecture models enhances both the dynamics and connections between various modules in the farm. The simulations highlight substantial cost savings, emphasizing the advantages of renewable energy integration in farming, with maximum cost reduction achieved within budget constraints. Results showed that \$100k investment in an HRES can reduce approximately 13% of electricity costs in 15 years.

4.2 Deep Reinforcement Learning based Farm Energy Management Considering Ammonia Generation

After solving the planning problem in this complex system, the next step is designing an effective operation strategy in this system, which is proposed in this section.

Addressing long-term weather trends adaptation in agriculture requires a cumulative impact from individual farm-level decisions, which are often shaped by public policies like subsidies, taxes, and regulations, as well as corporate policies such as pricing strategies and


supply chain alternatives. However, effective adaptation also depends on improved analysis of innovative approaches to ensure meaningful outcomes rather than failed initiatives. Building on recent advancements in reinforcement learning, renewable energy integration, and local ammonia production, this study aims to bridge existing gaps in sustainable farm energy management within the Food-Energy-Water (FEW) Nexus.

While prior research has demonstrated the potential of renewable energy and reinforcement learning (RL) techniques to enhance efficiency and reduce emissions, few studies have introduced a unified framework that simultaneously integrates energy generation, resource management, and sustainability at the individual farm level—particularly incorporating local ammonia production as a clean energy carrier. To address this gap, we propose a novel framework for farm energy management that leverages renewable energy sources, ammonia production, and Deep Reinforcement Learning (DRL) within the Food-Energy-Water (FEW) Nexus. This framework optimizes resource allocation while enhancing the sustainability and economic viability of agricultural systems. We develop and implement DRL-based algorithms to assess their effectiveness in improving farm operational efficiency, optimizing resource utilization (irrigation and nitrogen application), and maximizing economic returns. Our results demonstrate that DRL-driven approaches outperform traditional decision-making strategies by enhancing energy efficiency and reducing costs, thereby supporting more sustainable agricultural practices.

4.2.1 Modeling of the single farm within the FEW Nexus


Efficient resource management in agriculture is critical for ensuring sustainability within the Food-Energy-Water (FEW) Nexus. Renewable energy sources alone cannot provide a continuous and stable supply due to their intermittent nature, requiring an optimized energy storage and management framework. Additionally, ammonia, traditionally produced using fossil fuels, can serve as an energy carrier if integrated into farm energy systems. The proposed models in this section aim to address these challenges by integrating renewable energy, ammonia production, irrigation management, and crop modeling into a cohesive

Environment



(a)

Agent



(b)

Figure 4.5: (a) Illustration of FEW nexus in an individual farm. (b) DQN-enabled decision-making architecture integrated into farm management.

decision-making framework. This approach enhances operational efficiency while ensuring economic and environmental sustainability. Farming systems operate at the intersection of resource constraints, economic pressures, solar needs, and environmental imperatives, making their modeling within the FEW Nexus essential. Efficient resource allocation, renewable energy integration, and sustainable management practices are crucial for addressing the challenges of long-term weather trends, resource scarcity, and increasing food demands. To effectively model the FEW nexus at the individual farm level, we propose a FEWture-DRL framework that captures the interdependencies among energy generation, and consumption. This framework is divided into two major components: the energy model and the farm model as shown in Figure 4.5 (a).

In this system, ammonia production is facilitated through a community-owned ammonia plant. The plant is shared within the community; however, farmers must supply their own energy to generate ammonia, as the plant only provides the production infrastructure. Figure 4.5 (a) illustrates the interaction between various modules, including renewable energy sources (solar and wind), energy storage (battery), irrigation systems, and the community ammonia plant. The farm is allowed to draw power from both local renewable sources and the external grid while using the energy to manage crop production and ammonia production, ensuring efficient use of resources within the FEW nexus. The specifics of energy, water, ammonia and agriculture modeling are described next:

Energy Model

The energy model is composed of renewable energy generation, storage, grid interaction, and balance and curtailment components, which are defined in the following four subsections. This integrated framework ensures the farm’s energy demands are met.

Renewable Energy Generation Model

The renewable energy generation model quantifies the energy contributions from solar and wind resources, which are critical for powering the farm’s operations. The energy

generated from these sources can be represented by the following equations:

$$E_{\text{solar}} = Cap_{\text{solar}} \times \eta_{\text{pv}} \times \frac{I_{\text{solar}}}{I_{\text{stc}}} \times t_{\text{sun}} \quad (4.16)$$

$$E_{\text{wind}} = Cap_{\text{wind}} \times \eta_{\text{wind}}(v_{\text{wind}}) \times t_{\text{wind}} \quad (4.17)$$

where E_{solar} and E_{wind} represent the energy generated from solar and wind, respectively, in kWh. Cap_{solar} and Cap_{wind} are the installed capacities of the photovoltaic and wind systems, respectively, in kilowatts (kW). I_{solar} is the solar irradiance in watts per square meter (W/m^2), I_{stc} is the standard test condition irradiance, typically 1000 W/m^2 , and η_{pv} is the efficiency of the photovoltaic panels. The ratio $\frac{I_{\text{solar}}}{I_{\text{stc}}}$ is used to adjust the energy output according to the actual solar irradiance relative to standard conditions, and t_{sun} is the duration of effective sunlight in hours. For wind energy, v_{wind} represents the wind speed in meters per second (m/s), $\eta_{\text{wind}}(v_{\text{wind}})$ is the efficiency of the wind turbine depending on wind speed, and t_{wind} is the duration of wind availability in hours. Figure 4.6 shows the relationship between wind speed and the corresponding wind power output. The data indicates that the farm can expect significant power contributions from wind when wind speeds are within the range.




Figure 4.6: Wind Power Output Generation Vs Wind Speed

Energy Storage

The energy storage system serves as a buffer, enabling the farm to manage temporal mismatches between renewable energy generation and demand. It is modeled as a 4-hour battery system, which supports efficient energy use across various farm operations.

$$E_{\text{charge}} = \max(\min(E_{\text{wind}} + E_{\text{solar}} - E_{\text{load}}, S_{\text{max}} - S, \frac{S_{\text{max}}}{4}), 0) \quad (4.18)$$

$$E_{\text{discharge}} = \max(\min(E_{\text{load}} - E_{\text{solar}} - E_{\text{wind}}, S, \frac{S_{\text{max}}}{4}), 0) \quad (4.19)$$

$$S = S^{\text{previous}} + E_{\text{charge}} - E_{\text{discharge}} \quad (4.20)$$

where E_{charge} and $E_{\text{discharge}}$ refers to charging and discharging the energy storage system, respectively. S is the current stored energy in kWh, and S_{max} is the maximum storage capacity in kWh. The function $\max(\cdot, 0)$ ensures that no negative values occur when the demand exceeds the renewable energy generation. This equation describes how the state of charge S of the energy storage system is updated over time. The current state of charge S is calculated by taking the previous state of charge S^{previous} , adding the energy that has been charged into the system E_{charge} , and subtracting the energy that has been discharged $E_{\text{discharge}}$.

In addition, the energy storage system is modeled as a 4-hour battery, meaning it is designed to discharge its full capacity continuously over a four-hour period. This duration is chosen based on typical farm energy demand cycles, ensuring that stored renewable energy can be utilized during peak operation hours, such as irrigation and ammonia production. A 4-hour battery system provides flexibility to bridge the gap between intermittent renewable energy generation and real-time farm energy consumption.

Grid Interaction

The grid interaction component ensures energy reliability by allowing the farm to import energy when renewable sources and storage are insufficient. The equations are as follows:

$$E_{\text{grid}} = \max(E_{\text{load}} - E_{\text{solar}} - E_{\text{wind}} - E_{\text{discharge}}, \bar{E}_{\text{grid}}) \quad (4.21)$$

where E_{grid} is the energy drawn from the grid in kWh, when $\bar{E}_{\text{grid}}$ is the upper boundary of the imported energy from the grid.

Energy Balance and Curtailment

Energy balance and curtailment ensure that the supply meets the farm's demand while efficiently managing any surplus through curtailment, which is governed by:

$$E_{\text{curt}} = \max(E_{\text{solar}} + E_{\text{wind}} - E_{\text{load}} - E_{\text{charge}}, 0) \quad (4.22)$$

The energy balance is:

$$\Delta E_{\text{imb}} = |E_{\text{solar}} + E_{\text{wind}} + E_{\text{discharge}} + \bar{E}_{\text{grid}} - E_{\text{load}} - E_{\text{charge}}| \quad (4.23)$$

where ΔE_{imb} denotes the energy imbalance in kWh. E_{load} is the total energy demand of the farm, while E_{curt} in kWh represents curtailed energy when generation exceeds demand. The energy balance equation ensures that all energy generated and stored is fully utilized or curtailed, enforcing $\Delta E_{\text{imb}} = 0$. Any surplus energy beyond storage and load capacity is curtailed, while deficits are met through grid imports.

Irrigation Energy Model

Irrigation requires energy for water collection and distribution. The irrigation energy demand is modeled as:

$$E_{\text{irr}} = \eta_{\text{irr}} \times a_{\text{irr}} \quad (4.24)$$

where E_{irr} is the energy required for irrigation in kWh, a_{irr} is the amount of irrigation in centimeters (cm*hectare), and η_{irr} is the conversion factor from cm-ha to kWh for irrigation. Other aspects of water use for irrigation and the relatively tiny amount of water required for ammonia production (about 0.06% for corn) are discussed by²⁰⁹⁻²¹¹. In addition, the water model is in the following section 4.2.1.

Ammonia Generation Model

Ammonia serves as a vital nitrogen fertilizer for crop growth. In FEWture-DRL, the ammonia generation model provides a framework for assessing the energy demands associated with producing ammonia for fertilizer, a critical component of agricultural productivity^{211;212}. The energy required for ammonia production is substantial and the only aspect of ammonia production considered in FEWture-DRL. The energy demand is modeled as:

$$E_{\text{NH}_3} = \eta_{\text{NH}_3} \times a_{\text{NH}_3} \quad (4.25)$$

where E_{NH_3} is the energy required for ammonia generation in kWh, a_{NH_3} is the mass of ammonia produced in kilograms (kg), and η_{NH_3} is the specific energy consumption in kWh per kg of ammonia produced.

Agriculture and Soil Model

The agriculture model integrates crop growth simulations with environmental factors such as nitrogen and water limitations. This approach helps optimize resource use and enhance crop productivity. Specifically, this study employs the World Food Studies (WOFOST) model, integrated within the Python Crop Simulation Environment (PCSE) framework, the Soil Nitrogen module for Organic and Mineral Nitrogen (SNOMIN) as the simulation en-

gine². The SNOMIN module in WOFOST introduces an advanced approach for modeling both nitrogen and water limitations on crop growth, thereby enhancing the accuracy of crop simulations.

Nitrogen limitation in SNOMIN is represented by accounting for the dynamic interactions of soil nitrogen in both organic and mineral forms. Key processes modeled include mineralization, where organic nitrogen is converted into inorganic forms of ammonium, such as ammonium nitrate that can be absorbed by plants, and nitrogen immobilization and its assimilation by soil microorganisms. The foregoing determines the amount of available nitrogen to crops and is influenced by factors such as soil temperature, moisture, and carbon content. When soil nitrogen levels fall below a critical threshold, crop growth is constrained due to inadequate protein synthesis and chlorophyll formation. Additionally, SNOMIN incorporates the vertical movement of nitrogen across different soil layers, enabling a better understanding of how the root access to nutrients changes over time based on rainfall, irrigation practices, and root development. This dynamic representation of nitrogen behavior allows SNOMIN to capture the complex factors influencing nitrogen availability, ultimately contributing to more accurate crop growth simulations. The implications of nitrogen dynamics on plant productivity are maintained by the SNOMIN model³, relevant to FEWtures-DRL are described in the following section.

Nitrogen Balance for Fertilizer in SNOMIN

The nitrogen balance model in SNOMIN ensures a representation of nitrogen dynamics in the soil, contributing to improved estimation of crop yield under varying conditions. The nitrogen balance is expressed through the following equation, which ensures the conservation of ammonium nitrogen (NH₄) across each soil layer during a given time step:

$$\begin{aligned} \text{NH4}_{\text{balance}} = & \text{NH4}_{\text{initial}} - \text{NH4}_{\text{final}} + R_{\text{NH4AM}} + R_{\text{NH4IN}} + R_{\text{NH4MIN}} \\ & + R_{\text{NH4DEPOS}} - R_{\text{NH4NITR}} - R_{\text{NH4OUT}} - R_{\text{NH4UP}} \end{aligned} \quad (4.26)$$

where $\text{NH4}_{\text{balance}}$ represents the ammonium nitrogen mass balance for the soil layer, which should be close to zero to satisfy the principle of mass conservation. $\text{NH4}_{\text{initial}}$ is the initial amount of ammonium nitrogen (NH4-N) present in the soil layer at the beginning of the time step, and $\text{NH4}_{\text{final}}$ is the final amount at the end of the time step, both expressed in kg N m^{-2} . The term R_{NH4AM} represents the rate of ammonium nitrogen application ($\text{kg NH4-N m}^{-2} \text{ day}^{-1}$), which accounts for ammonium added to the soil from external sources such as fertilizers. R_{NH4IN} is the rate of inflow of ammonium nitrogen from adjacent soil layers, representing the nitrogen moving into the soil layer from neighboring layers, while R_{NH4MIN} denotes the net rate of mineralization, which is the conversion of organic nitrogen into ammonium through biological processes. R_{NH4DEPOS} refers to the rate of ammonium deposition from sources like precipitation, which can enhance nitrogen availability. Conversely, R_{NH4NITR} is the rate of nitrification, representing the conversion of ammonium into nitrate by soil microorganisms, thereby reducing the amount of available ammonium. R_{NH4OUT} is the rate of outflow of ammonium nitrogen to adjacent soil layers, indicating the movement of ammonium out of the current soil layer into surrounding layers. Finally, R_{NH4UP} is the rate of ammonium uptake by plant roots, which represents the amount of ammonium nitrogen absorbed by plants from the soil. The nitrogen balance equation takes into account all these processes—application, inflow, outflow, deposition, mineralization, nitrification, and root uptake—to ensure that the ammonium nitrogen in each soil layer is balanced accurately over time.

Water Balance for Irrigation in SNOMIN

For modeling water limitation², SNOMIN implements a Multi-Layer Water Balance (MLWB), an enhancement over the simpler single-layer balance used in earlier WOFOST versions. The multi-layer approach enables the model to assess water availability across various soil horizons, providing insights into the water accessible to crop roots at different growth stages. Water limitation occurs when the soil moisture content in the root zone falls below the level required to meet the crop’s evapotranspiration needs, leading to reductions

in photosynthesis and other growth processes. The MLWB in SNOMIN also helps simulate the impact of water stress on nitrogen uptake, since water availability directly affects nutrient mobility within the soil profile. By considering nitrogen and water limitations interactively, SNOMIN provides a comprehensive representation of how environmental factors impact crop productivity, highlighting the critical importance of managing both water and nutrient resources to achieve optimal growth.

$$R_{\text{in}} = \min(R_{\text{in, PRE}}, Flow_{\text{top}}) \quad (4.27)$$

$$R_{\text{in, PRE}} = R_{\text{rain}} + a_{\text{irr}} + SS \quad (4.28)$$

The infiltration rate (R_{in}) in the multi-layer water balance is influenced by both rainfall and irrigation, as well as the soil's capacity to absorb water. It is calculated as the minimum of the potential infiltration rate ($R_{\text{in, PRE}}$) and the maximum downward flow at the top boundary of the soil profile ($Flow_{\text{top}}$). The potential infiltration rate ($R_{\text{in, PRE}}$) includes contributions from rainfall (R_{rain}), irrigation (a_{irr}), and surface storage (SS)—representing the amount of water that has accumulated on the soil surface. The infiltration rate is constrained by $Flow_{\text{top}}$, which is the maximum rate at which water can move downward from the soil surface based on soil properties such as conductivity and capacity. This equation ensures that the amount of water infiltrating the soil is physically realistic, accounting for both the amount of water added through irrigation and rainfall, and the soil's ability to absorb this water.

4.2.2 Approach of FEWture-DRL Framework

The FEWture-DRL framework is composed of spaces for observations, actions, reward functions, and economics. Specifically, the agent is a reinforcement learning model responsible for making sequential decisions regarding irrigation and nitrogen application to optimize

farm operations. The state represents the set of information available to the agent at each time step in the simulation. The agent does not have access to the entire time series of states in advance. Instead, it operates in a Markov Decision Process (MDP) framework, where it only observes the current state at each time step and makes decisions accordingly. As the optimization progresses, the agent builds knowledge based on past interactions through the experience replay mechanism, but it does not have foresight into future states beyond what is inherently available in the given step of the simulation. This structure ensures that the FEWture-DRL framework integrates decision-making for sustainable farm energy management.

Table 4.3: Observation Space (State Variables)

Variable	Definition	Unit (Time Step: Daily)
DAY	Current day of the growing season	Day
DVS	Crop development stage	-
WSO	Weight of storage organ (crop yield)	kg
SM	Soil moisture content	%
RFTRA	Root fractional transpiration	-
RMINT	Root mineral nitrogen uptake	kg
RIRR	Irrigation amount applied	cm
TRA	Transpiration rate	mm/day
NAVAIL	Available nitrogen in soil	kg/ha
Ndemand	Crop nitrogen demand	kg/ha
Grid Exchange	Power exchanged with the grid	kWh
Energy Storage	Stored energy level in battery	kWh
Energy Demand	Total farm energy demand	kWh
Wind Generation	Wind power generation	kWh
Solar Generation	Solar power generation	kWh
IRRAD	Solar irradiance	W/m ²
TMIN	Minimum daily temperature	°C
TMAX	Maximum daily temperature	°C
WIND	Wind speed	m/s
RAIN	Rainfall	mm

Observation Space

The observation space in this environment provides the agent with key information about the crop, soil, water, energy, and weather dynamics that affect farm management. The state consists of 19 variables, which include agricultural parameters, energy-related variables, and

weather data as shown in Table 4.3. These variables provide the necessary environmental context for energy generation and crop growth.

Action Space

The agent is tasked with controlling two action spaces: the amount of irrigation and the amount of nitrogen applied to the crop. The irrigation amount, measured in centimeters, ranges from 0 to 0.6 cm with 0.1 cm incremental, over the planting period (we assume the irrigation decision is applied around noon in the day). The nitrogen application amount, measured in kg/ha, ranges from 0 to 3.5 kg/ha with 0.5 kg/ha incremental step as shown in Table 4.4. These actions directly influence crop growth, water usage, nitrogen uptake, and energy consumption, which are critical for optimizing farm operations. This setup allows the agent to explore different combinations of irrigation and nitrogen application rates in each timestep.

Table 4.4: Action Space and Reward Function

Category	Variable	Unit (Time Step: Daily)	Action Boundaries
Action Space	Irrigation	cm	[0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6]
	Nitrogen Application	kg/ha	[0, 0.5, 1.0, 1.5, 2.0, 2.5, 3.0, 3.5]

Reward Function and Economics

The reward function is designed to balance the objectives of maximizing crop yield and minimizing operational costs, including irrigation, nitrogen application, and energy usage. The reward is calculated as crop revenue minus the total production cost.

Revenue is determined by the crop area, yield, and crop price. Here, crop yield ΔY_{crop} is represented by the weight of the storage organ (WSO) per unit area, expressed in kg ha^{-1} (kilograms per hectare). This unit is widely used in agricultural research to quantify crop productivity. The revenue equation integrates these factors and is expressed as:

$$Revenue = \Delta Y_{\text{crop}} \times p_{\text{crop}} \quad (4.29)$$

where crop yield ΔY_{crop} is the mass of the storage organ changed at the end of each step simulation, and crop price p_{crop} is a constant reflecting the market value. Costs are composed of three main components. First, the irrigation cost is proportional to the amount of water applied during each step. Second, energy cost is incurred when energy is drawn from the grid to support irrigation or nitrogen generation, which also includes the energy required for ammonia production. Note that the ammonia plant is community-owned and farmers supply only energy, so the cost of ammonia is inherently captured within the energy cost. Finally, a penalty is applied for power imbalance, which arises when the farm's energy supply from renewable sources does not match the demand. The total cost ($\text{Cost}_{\text{total}}$) is composed of three main components: irrigation cost ($C_{\text{irrigation}}$), energy cost (C_{energy}), and imbalance cost ($C_{\text{imbalance}}$):

$$\text{Cost}_{\text{total}} = C_{\text{irrigation}} + C_{\text{energy}} + C_{\text{imbalance}} \quad (4.30)$$

The irrigation cost is calculated by multiplying the unit cost of irrigation water ($p_{\text{irrigation}}$) by the energy used for irrigation (E_{IRR}). The energy cost is determined by the product of the electricity price (p_{elec}) and the energy imported from the grid (E_{grid}).

$$C_{\text{irrigation}} = p_{\text{irrigation}} \times E_{\text{IRR}} \quad (4.31)$$

$$C_{\text{energy}} = p_{\text{elec}} \times E_{\text{grid}} \quad (4.32)$$

$$C_{\text{imbalance}} = p_{\text{penalty}} \times \Delta E_{\text{imb}} \quad (4.33)$$

The imbalance cost arises when the farm's energy demand exceeds the available supply from renewable sources and storage, and the maximum allowable grid import ($\bar{E}_{\text{grid}}$) is reached. It is computed by multiplying the penalty price (we assume the penalty price

is much higher than the regular electricity price) for imbalance (p_{penalty}) with the energy imbalance (ΔE_{imb}). Since the farm integrates renewable energy sources, energy production may not always match demand in real time. If surplus energy cannot be stored or used efficiently, it is curtailed, while deficits require grid imports. The imbalance cost incentivizes the agent to optimize energy usage, minimize reliance on the grid, and ensure efficient resource allocation. The reward function is then given by the difference between revenue and production cost:

$$\text{Reward} = \text{Revenue} - \text{Cost}_{\text{total}} \quad (4.34)$$

If the crop has been harvested, the reward reflects the overall net profit after the entire growing season. Otherwise, it represents the operational efficiency up to the current timestep. The reward structure incentivizes the agent to manage both water and nitrogen efficiently while minimizing energy dependence on external sources.

Temporal Dynamics and Decision Process

The simulation operates on a daily time step, ensuring consistent updates across all state, action, and reward variables. Each day, the agent observes the environment based only on past and present data, without access to future states or predictive insights. The agent makes daily decisions on irrigation and nitrogen application, dynamically responding to changes in soil moisture, crop development, and weather conditions.

This study assumes a one-hectare planting area dedicated to potato cultivation throughout the growing season. The model does not account for crop selection, intercropping, or crop rotation. All reported values are standardized per hectare: irrigation depth is measured in centimeters, nitrogen application in kilograms, energy use in kilowatt-hours, and economic factors in dollars. While the daily decision-making framework results in a high-dimensional action space, reinforcement learning efficiently navigates this complexity to optimize long-term farm performance. At the end of the growing season, the cumulative reward is calcu-

lated based on farm profitability, resource efficiency, and energy consumption, incentivizing sustainable and cost-effective decision-making.

DQN Algorithm for Farm Energy Management

The FEWture-DRL framework applies the Deep Q-Network (DQN) algorithm^{16;111} to optimize energy management in a farm setting, considering key factors such as irrigation energy use and ammonia production. This framework captures the interactions between renewable energy sources, energy storage, and agricultural processes to enhance sustainability and operational efficiency. This study expands on this FEW Nexus by incorporating the DQN-based decision-making process, as Figure 4.5 (b) highlights the DQN agent, which is employed to efficiently manage energy resources and farm operations by selecting optimal actions based on predefined discrete action spaces. DQN is well-suited for this task as it effectively handles environments with discrete decision variables, such as varying irrigation levels and nitrogen applications. By learning from past experiences, the DQN algorithm enhances farm sustainability by balancing energy consumption, resource allocation, and economic performance.

Algorithm 2 DQN Algorithm for FEWture-DRL

- 1: Initialize action-value function $Q_\theta(s, a)$ with random weights θ ; Initialize target action-value function $\hat{Q}_\theta(s, a)$ with weights $\theta' \leftarrow \theta$; Initialize replay buffer $\mathcal{B}$ with size N
- 2: Set exploration parameters: exploration rate ϵ , decay factor ϵ_{decay} , and minimum ϵ_{min}
- 3: **for** each episode **do**
- 4: Reset environment and observe initial state s_0
- 5: **for** each time step t within the episode **do**
- 6: With probability ϵ , select a random action a_t
- 7: Otherwise, select $a_t = \arg \max_a Q_\theta(s_t, a)$
- 8: Execute action a_t in the environment, observe next state s_{t+1} and reward r_t
- 9: Store transition (s_t, a_t, r_t, s_{t+1}) in replay buffer $\mathcal{B}$
- 10: **if** replay buffer is sufficiently populated **then**
- 11: Sample a minibatch of transitions (s_i, a_i, r_i, s_{i+1}) from $\mathcal{B}$
- 12: Compute the target for each transition:

$$y_i = r_i + \gamma \max_{a'} \hat{Q}_\theta(s_{i+1}, a')$$

- 13: Update Q-network by minimizing the loss:

$$L(\theta) = \frac{1}{|B|} \sum_i (y_i - Q_\theta(s_i, a_i))^2$$

- 14: Perform a gradient descent step to update the Q-network parameters θ
- 15: Every C steps, update the target network:

$$\theta' \leftarrow \theta$$

- 16: **end if**
 - 17: Update $\epsilon \leftarrow \max(\epsilon_{min}, \epsilon \times \epsilon_{decay})$
 - 18: **end for**
 - 19: End episode when a terminal state is reached
 - 20: **end for**
-

The Algorithm 1 (DQN) algorithm approximates the optimal action-value function $Q_{\theta}(s, a)$ using a neural network, selecting actions via ϵ -greedy exploration to balance exploitation and exploration. The experience replay buffer stabilizes training by reducing temporal correlations. DQN was chosen for its high sample efficiency and suitability for discrete action spaces. For comparison, we also applied the Advantage Actor-Critic (A2C) algorithm²¹³. Unlike DQN, A2C is an on-policy method, updating its policy with newly collected experiences but lacking experience replay, leading to lower sample efficiency and potential learning instability. This difference explains why DQN outperformed A2C in our simulations, particularly in optimizing irrigation and nitrogen application, where long-term reward optimization is crucial.

4.3 Experiment and Case Studies

In this section, we provide the key parameters influencing the simulation of the farm environment and renewable energy. In addition, comprehensive simulation results reflect the effectiveness of DRL in farm energy management.

4.3.1 Simulation Setup

This section shows the farm parameters, weather data and other parameters we deployed in this research. Figure 4.7 presents an overview of the weather data used during the farm environment simulation. This includes various climatic parameters such as temperature, solar radiation, and wind speed, which are key inputs for crop growth models and renewable energy production estimation. The weather data is representative of typical seasonal variations, which helps assess the feasibility of integrating renewable energy resources into the farm’s operations.

Table 4.5 provides detailed information on the crop start and end dates, defining the crop growth period used in the simulation. Additionally, the soil parameters for a four-layer profile are essential for modeling water and nutrient dynamics in the WOFOST model. These




Figure 4.7: Weather data illustration for one year of daily data

parameters, sourced from example simulations in the PCSE documentation³, are crucial in representing realistic soil-water-plant interactions for crop growth modeling.

To evaluate the effectiveness of the proposed DQN-based farm energy management framework, two rule-based control (RBC) strategies, with no applied optimization, are introduced as benchmarks for comparison:

- **No Irrigation, No Fertilizer (NINF):** This strategy represents an extremely resource-conservative scenario where neither irrigation nor fertilizer is applied. It relies entirely on natural environmental conditions and is used as a lower-bound benchmark to assess the benefits of more active farm management strategies.
- **Full Irrigation, Full Fertilizer (FIFF):** This strategy represents a resource-intensive endpoint, where irrigation and fertilizer are applied at their maximum levels. It serves to illustrate the upper bound of input use and potential yield.

Table 4.6 summarizes the cost parameters used in the simulation, focusing on operational expenses rather than capital costs. These parameters play a crucial role in evaluating the

Crop Start Date	Crop End Date
April 21, 2020	October 24, 2020

Parameter	Values
Soil Properties and Hydraulic Parameters	
<i>Permanent Wilting Point</i>	4.2
<i>Field Capacity</i>	2.0
<i>Surface Conductivity (cm/day)</i>	75.0
Data for Soil – Four Layers	
<i>Layer Thickness (cm)</i>	[25, 15, 20, 60]
<i>CN Ratio SOMI</i>	[17, 25, 25, 25]
<i>CRAIRC (Critical Air Content)</i>	[0.09, 0.09, 0.09, 0.09]
<i>FSOMI (Fraction of Soil Organic Matter Input)</i>	[0.05, 0.02, 0.01, 0]
<i>RHOD (Soil Bulk Density, g/cm³)</i>	[1.375, 1.576, 1.633, 1.672]
<i>Soil pH</i>	[4.8, 4.5, 4.5, 4.7]

Table 4.5: Parameter settings for WOFOST and SNOMIN model simulation, including soil hydraulic and organic matter properties for four soil layers. The data is sourced from example simulations in the PCSE documentation^{2;3}.

Parameter	Value
Required energy for ammonia production (kWh/kg)	11.5 ²¹⁴
Unit cost for irrigation ($p_{\text{irrigation}}$) (\$/cm/ha)	25.51 ²¹⁵
Required energy for irrigation (kWh/cm/ha)	59.33 ²¹⁶
Crop price (p_{crop}) (\$/kg)	1.10 ²¹⁷
Electricity price (p_{elec}) (\$/kWh)	0.15 ²¹⁸


Table 4.6: Cost parameters for farm environment simulation with associated mathematical symbols. NH3 refers to ammonia.

farm’s economic viability. By providing references for each cost component, the table ensures transparency and reliability, enhancing the credibility of the simulation outcomes.

4.3.2 Simulation Results

This section presents the outcomes of the farm simulation using the proposed FEWture-DRL framework. The performance of the DQN and A2C algorithms is evaluated against two rule-based strategies: NINF and FIFF. The analysis focuses on cumulative reward, crop yield, irrigation and nitrogen application patterns, and overall resource efficiency.

Figure 4.8 compares reward trends and crop yield across four DRL algorithms (NINF,



(a) Reward comparison for DRL algorithms, i.e., DQN and A2C and Rule-based Control (RBC), i.e., NINF and FIFF (b) Crop yield comparison for DRL algorithms and Rule-based Control (RBC)

Figure 4.8: Comparison of (a) rewards and (b) crop yield between DRL algorithms.

FIFF, A2C, and DQN). Subfigure (a) illustrates the reward progression over time, reflecting the typical agricultural growth pattern. Initially, rewards are low as plants are in their early stages, followed by rapid accumulation during peak growth, and eventually stabilizing as crops mature. DQN achieves the highest rewards, showcasing its effectiveness in optimizing resource allocation. FIFF closely follows, while A2C and NINF lag behind, indicating less efficient management strategies.

Subfigure (b) presents cumulative crop yield (WSO) across the algorithms. The yield follows a typical agricultural growth pattern: slow initial growth, rapid accumulation, and a plateau as crops mature. While DQN and FIFF maintain higher yields through efficient irrigation and nitrogen application, A2C achieves moderate yields, and NINF, lacking resource inputs, results in the lowest yield. The sharp increase in yields until around day 115 reflects optimal resource allocation, after which diminishing returns set in. These results highlight how water and fertilizer decisions impact crop growth and suggest that DQN and FIFF most effectively optimize farm productivity.

Figure 4.9 illustrates the trends in irrigation and nitrogen application across the simulation period. In the top subplot, DQN applies irrigation at an average of 0.4 cm per day, demonstrating its adaptive but resource-intensive strategy. FIFF, which follows a fixed




Figure 4.9: Action comparison between algorithms

high-resource approach, consistently applies irrigation at the upper limit of 0.6 cm per day. A2C employs a more conservative approach, with irrigation levels mostly around 0.2 cm, while NINF does not apply irrigation, adhering to an extreme farming strategy. In the lower subplot, DQN frequently applies nitrogen at the maximum level of 3.0 kg/ha/day, optimizing for yield while adjusting applications dynamically. FIFF, however, rigidly applies the highest nitrogen level (3.5 kg/ha/day) throughout the simulation without adaptation, making it the most resource-intensive strategy. A2C applies nitrogen more conservatively, averaging around 1.5 kg/ha/day, while NINF applies none. These trends highlight the key difference between rule-based strategies like FIFF, which applies maximum inputs without optimization, and DRL-based strategies like DQN, which learn to balance efficiency and productivity.

Algorithm	NINF	FIFF	A2C	DQN
Total Irrigation Depth (cm)	0	101.4	42.7	58.9
Total Nitrogen Weight (kg)	0	591.5	293.0	406.3
Total Crop Yield (kg)	2962.9	12281.3	7573.2	11642.7
Total Reward (\$)	3259.2	10947.4	7241.3	11352.9

Table 4.7: Summary of different farming strategies per unit area of crop land.

4.3.3 Discussion

The FEWture-DRL framework applies the DQN algorithm^{16?} to optimize farm energy management, considering irrigation, ammonia production, and renewable energy integration. DQN, an off-policy method, uses experience replay to learn from past experiences, enhancing decision-making over time. In contrast, the A2C algorithm²¹³, an on-policy method, updates its policy using only recent experiences, making it less sample-efficient and more sensitive to short-term fluctuations.

Despite sharing the same objective function and farm model, these differences in learning mechanisms lead to notable performance variations. DQN outperforms A2C in crop yield and economic return by better balancing resource application and energy efficiency. Its off-policy learning enables recognition of delayed effects, such as the long-term impact of nitrogen on crop yield. A2C, being on-policy, prioritizes short-term gains, resulting in more conservative irrigation and nitrogen application, reducing costs but leading to lower yields. Figure 4.10 presents the frequency distribution of irrigation and nitrogen levels applied by the algorithms. The top subplot demonstrates that DQN and FIFF overwhelmingly select the highest irrigation levels (0.4 cm and 0.6 cm, respectively), emphasizing their aggressive resource utilization strategies. A2C applies lower irrigation levels more frequently, particularly around 0.1 cm and 0.2 cm, while NINF does not apply irrigation at all. In addition, nitrogen frequency distribution confirms DQN’s nearly consistent application of the maximum nitrogen level (3.0 kg/ha). FIFF also applies high nitrogen levels, predominantly at 3.5 kg/ha. A2C shows a broader distribution with a focus on moderate nitrogen levels (1.5 kg/ha), and NINF applies no nitrogen. These distributions further reinforce the distinct resource allocation strategies employed by each algorithm.




Figure 4.10: Action frequency comparison between algorithms

Table 4.7 provides a comparative summary of different farming strategies, illustrating the trade-offs between resource use and crop yield. The NINF scenario, which excludes both irrigation and fertilizer, represents an extreme case of minimal input application rather than a typical dryland farming strategy. As a result, it yields the lowest crop production and economic return, highlighting the limitations of non-interventionist approaches under simulated conditions. In contrast, FIFF applies maximum levels of irrigation and nitrogen, achieving the highest yield but at the cost of intensive resource use, reflecting a resource-intensive endpoint for agricultural practices rather than a realistic baseline.

The adaptive strategies of DQN and A2C differ significantly despite both being optimization-based approaches. DQN consistently achieves higher crop yield and economic reward than A2C while using only moderately more irrigation and nitrogen. This advantage stems from DQN’s off-policy nature and use of experience replay, which allow it to learn more effectively from past experiences and better capture long-term dependencies in crop responses. In contrast, A2C, as an on-policy method, updates its policy using only the most recent interactions, making it less sample-efficient and more sensitive to short-term fluctuations. As a result, A2C tends to adopt more conservative resource use strategies. These findings suggest that FEWture-DRL with DQN more accurately mimics the adaptive decision-making of experienced farmers, offering a robust tool for sustainable farm management.

4.4 Summary

This chapter first introduces a planning framework for agricultural energy systems that integrates local ammonia production within the FEW Nexus. Building on this foundation, DRL algorithms—specifically, DQN and A2C—are applied to optimize farm-level energy management. The simulations focus on achieving a balance between energy consumption, irrigation, and nitrogen application, while simultaneously maximizing crop yield and minimizing operational costs. Among the algorithms, DQN consistently demonstrated superior performance in terms of total reward and crop yield. This can be attributed to its more

resource-intensive approach, utilizing higher levels of irrigation and nitrogen throughout the simulation. Compared to A2C, DQN demonstrated superior sample efficiency due to its use of experience replay, allowing it to learn more effectively from past decisions. Additionally, A2C’s on-policy nature led to higher variance in learning, making it less stable in optimizing farm resource allocation, whereas DQN’s off-policy learning facilitated better decision-making over time. In terms of irrigation and nitrogen application strategies, DQN applied the highest levels of resources, reflecting an aggressive approach to maximizing crop growth. The frequency of resource application across the algorithms further highlighted these differences. DQN frequently utilized high levels of both irrigation and nitrogen, while other algorithms employed more moderate and varied levels of resource use. This difference in strategy significantly impacted the outcomes, with DQN producing the highest crop yield but at the cost of greater resource utilization. In conclusion, DQN emerged as the most effective algorithm for maximizing farm performance in terms of crop yield and rewards. Future research could focus on refining DRL algorithms by incorporating sustainability constraints, such as water conservation policies and carbon footprint reduction. Additionally, real-world testing of these models on actual farm operations could provide deeper insights into their feasibility and long-term benefits.

Chapter 5

Conclusion and Future Work

This chapter concludes the dissertation by summarizing the key research findings, highlighting the broader impacts of this work, and suggesting several promising directions for future research. Detailed analyses and chapter-specific conclusions can be found within their respective sections.

5.1 Conclusion

This dissertation has investigated how advanced Machine Learning (ML) techniques can significantly improve energy management, efficiency, and sustainability across three interconnected domains: power systems, building energy management, and agricultural operations within the Food-Energy-Water (FEW) nexus. Each study developed novel methodologies to address critical, sector-specific challenges, demonstrating ML's flexibility and effectiveness in solving complex real-world problems.

Island Power Systems. To maintain grid stability in islanded power systems with high renewable penetration, a linear regression-based frequency nadir constraint was developed and embedded within a mixed-integer linear programming (MILP)-based Unit Commitment (UC) model. Unlike traditional minimum inertia approaches, the proposed method explicitly

accounts for frequency stability after the largest generation trip event. Validation with the Maui Island grid over an extended simulation period demonstrated that the proposed model significantly enhances frequency stability and operational reliability while only marginally increasing generation costs. This work underscores the value of integrating data-driven constraints into UC scheduling, improving renewable integration and ensuring reliable operation in low-inertia grids.

Occupant-Centric Building Energy Management. In building energy management, a multi-agent deep reinforcement learning (MADRL) framework, based on the Multi-Agent Deep Deterministic Policy Gradient (MADDPG), was proposed to optimize multi-zone HVAC systems by explicitly modeling dynamic occupant behaviors, including occupancy patterns, clothing insulation changes, and metabolic activity. EnergyPlus-based simulations demonstrated remarkable efficiency gains, achieving up to 51.09% electricity cost savings compared to traditional rule-based methods and outperforming single-agent DRL approaches by 4.34%, all without compromising occupant thermal comfort. These results confirm that integrating occupant-centric behavioral dynamics within multi-agent DRL frameworks can substantially enhance building energy performance.

Agricultural Operations within the FEW Nexus. In the agricultural domain, the study introduced a DRL framework to optimize agricultural resource management within the FEW Nexus. The framework jointly manages energy consumption, irrigation scheduling, and nitrogen application to enhance crop productivity while minimizing operational costs. Two DRL algorithms—DQN and A2C—were implemented and benchmarked against a rule-based control strategy. The results reveal that DQN, with its aggressive policy, consistently maximizes crop yield through higher resource usage. This approach significantly outperforms the rule-based baseline in terms of overall reward and adaptability to environmental variability. These findings highlight the potential of DRL to enable intelligent, adaptive decision-making in agricultural systems, supporting both economic performance and environmental sustainability.

Collectively, these three interconnected studies illustrate the transformative potential of machine learning in advancing energy system sustainability, flexibility, and efficiency. By providing novel data-driven methods tailored to power systems, occupant-centric building management, and agriculture, this dissertation makes significant contributions to energy management literature and practical applications, offering scalable solutions for the sustainable development of modern energy infrastructure.

5.2 Future Work

Several promising directions emerge for further investigation, extending the methodologies and insights developed in this dissertation:

- **Advanced Frequency Regulation in Power Systems:** Future research could incorporate grid-forming inverter models and virtual inertia technologies into the data-driven frequency nadir constraint, further enhancing grid resilience and operational fidelity in systems characterized by extremely high renewable penetrations.
- **Integration of Occupant-Centric and Bioclimatic Strategies in Buildings:** Expanding upon occupant-centric DRL frameworks, future studies could integrate bioclimatic design features and detailed human behavior analytics to create comprehensive solutions that merge active controls with passive environmental designs, significantly enhancing building energy efficiency and occupant comfort.
- **Multi-objective Reinforcement Learning in Agriculture:** Future research may extend the DRL framework to explicitly include multiple conflicting sustainability objectives, such as reducing water use, minimizing nitrogen runoff, and decreasing greenhouse gas emissions alongside economic targets. Multi-objective DRL algorithms could help farmers navigate these complex trade-offs effectively.
- **Cross-domain Integration of Energy Management:** Finally, an exciting research pathway involves cross-domain studies integrating grid operations, smart building con-

trols, and agricultural energy management into cohesive ML-based platforms. Such studies would offer insights into system-wide optimization and synergies across traditionally siloed sectors, further advancing energy sustainability on broader geographic and temporal scales.

Exploring these avenues will undoubtedly enhance the impact and practical utility of the methodologies presented here, driving forward the collective goal of sustainable, intelligent, and integrated energy management solutions.

Bibliography

- [1] Alex Modarresi and John Symons. Failures to food, energy, and water systems: Mapping and simulating components to improve resilience. In *2023 International Conference on Smart Applications, Communications and Networking (SmartNets)*, 2023. doi: 10.1109/SmartNets58706.2023.10215865.
- [2] Herman N.C. Berghuijs, João Vasco Silva, Pytrik Reidsma, and Allard J.W. de Wit. Expanding the wofost crop model to explore options for sustainable nitrogen management: A study for winter wheat in the netherlands. *European Journal of Agronomy*, 154:127099, 2024. ISSN 1161-0301. doi: <https://doi.org/10.1016/j.eja.2024.127099>. URL <https://www.sciencedirect.com/science/article/pii/S1161030124000200>.
- [3] Allard de Wit. *PCSE Documentation Release 6.0*, January 22 2025. Available online at <https://pcse.readthedocs.io/en/latest/>.
- [4] Xunzhang Pan, Tianming Shao, Xinzhu Zheng, Yanru Zhang, Xueqing Ma, and Qi Zhang. Energy and sustainable development nexus: A review. *Energy Strategy Reviews*, 47:101078, 2023. ISSN 2211-467X. doi: <https://doi.org/10.1016/j.esr.2023.101078>. URL <https://www.sciencedirect.com/science/article/pii/S2211467X23000287>.
- [5] Mou Mahmood, Prangon Chowdhury, Rahbaar Yeassin, Mahmudul Hasan, Tanvir Ahmad, and Nahid-Ur-Rahman Chowdhury. Impacts of digitalization on smart grids, renewable energy, and demand response: An updated review of current applications. *Energy Conversion and Management: X*, 24:100790, 2024. ISSN 2590-1745. doi: <https://doi.org/10.1016/j.ecmx.2024.100790>. URL <https://www>.

[sciencedirect.com/science/article/pii/S259017452400268X](https://www.sciencedirect.com/science/article/pii/S259017452400268X).

- [6] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer-Verlag New York, 2006.
- [7] Michael I. Jordan and Tom M. Mitchell. Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245):255–260, 2015.
- [8] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*. Springer, 2009.
- [9] Tao Hong and Shu Fan. Probabilistic electric load forecasting: A tutorial review. *International Journal of Forecasting*, 32(3):914–938, 2016.
- [10] Anbo Meng, Peng Wang, Guangsong Zhai, Cong Zeng, Shun Chen, Xiaoyi Yang, and Hao Yin. Electricity price forecasting with high penetration of renewable energy using attention-based lstm network trained by crisscross optimization. *Energy*, 254:124212, 2022. ISSN 0360-5442. doi: <https://doi.org/10.1016/j.energy.2022.124212>. URL <https://www.sciencedirect.com/science/article/pii/S036054422201115X>.
- [11] Yuvaraju Venkatachalam and Thangavel Subbaiyan. Intelligent fault diagnosis in power systems: A comparative analysis of machine learning-based algorithms. *Expert Systems with Applications*, 265:125945, 2025. ISSN 0957-4174. doi: <https://doi.org/10.1016/j.eswa.2024.125945>. URL <https://www.sciencedirect.com/science/article/pii/S0957417424028124>.
- [12] Yunyu Cao, Jinrui Tang, Shaohui Shi, Defu Cai, Li Zhang, and Ping Xiong. Fault diagnosis techniques for electrical distribution network based on artificial intelligence and signal processing: A review. *Processes*, 13(1), 2025. ISSN 2227-9717. doi: 10.3390/pr13010048. URL <https://www.mdpi.com/2227-9717/13/1/48>.
- [13] Tianqi Chen and Carlos Guestrin. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794, 2016.

- [14] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553): 436–444, 2015.
- [15] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2018.
- [16] Volodymyr Mnih. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- [17] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In *4th International Conference on Learning Representations (ICLR)*, San Juan, Puerto Rico, Sep. 2015. doi: 10.48550/arxiv.1509.02971.
- [18] Tianshu Wei, Yanzhi Wang, and Qi Zhu. Deep reinforcement learning for building HVAC control. In *Proceedings of the 54th Annual Design Automation Conference 2017*, pages 1–6. ACM, 2017-06-18. ISBN 978-1-4503-4927-7. doi: 10.1145/3061639.3062224. URL <https://dl.acm.org/doi/10.1145/3061639.3062224>.
- [19] Kevin P. Murphy. *Machine Learning: A Probabilistic Perspective*. MIT Press, 2012.
- [20] Ioannis Antonopoulos, Valentin Robu, Benoit Couraud, Desen Kirli, Sonam Norbu, Aristides Kiprakis, David Flynn, Sergio Elizondo-Gonzalez, and Steve Wattam. Artificial intelligence and machine learning approaches to energy demand-side response: A systematic review. *Renewable and Sustainable Energy Reviews*, 130:109899, 2020. ISSN 1364-0321. doi: <https://doi.org/10.1016/j.rser.2020.109899>. URL <https://www.sciencedirect.com/science/article/pii/S136403212030191X>.
- [21] Zhiang Zhang, Adrian Chong, Yuqi Pan, Chenlu Zhang, and Khee Poh Lam. Whole building energy model for HVAC optimal control: A practical framework based on deep reinforcement learning. *Energy and Buildings*, 199:472–490, Sep. 2019. ISSN 0378-7788. doi: 10.1016/j.enbuild.2019.07.029.
- [22] S. Dong, X. Fang, J. Tan, N. Gao, X. Cui, and A. Hoke. A unified analytical method

- to quantify three types of fast frequency response from inverter-based resources. In *22nd Wind and Solar Integration Workshop (WIW 2023)*, volume 2023, pages 669–674, 2023. doi: 10.1049/icp.2023.2801.
- [23] Xin Fang, Haoyu Yuan, and Jin Tan. Secondary frequency regulation from variable generation through uncertainty decomposition: An economic and reliability perspective. *IEEE Transactions on Sustainable Energy*, 12(4):2019–2030, 2021. doi: 10.1109/TSTE.2021.3076758.
- [24] Ningchao Gao, Shuan Dong, Xin Fang, Andy Hoke, David Wenzhong Gao, and Jin Tan. Developing frequency stability constraint for unit commitment problem considering high penetration of renewables: Preprint. 9 2023. URL <https://www.osti.gov/biblio/2001487>.
- [25] V.V. Terzija. Adaptive underfrequency load shedding based on the magnitude of the disturbance estimation. *IEEE Transactions on Power Systems*, 21(3):1260–1266, 2006. doi: 10.1109/TPWRS.2006.879315.
- [26] Ningchao Gao, David Wenzhong Gao, and Xin Fang. Manage real-time power imbalance with renewable energy: Fast generation dispatch or adaptive frequency regulation? *IEEE Transactions on Power Systems*, 38(6):5278–5289, 2023. doi: 10.1109/TPWRS.2022.3232759.
- [27] Alice Grundy. National grid eso claims world first approach to inertia, awarding £328m in contracts. URL <https://www.current-news.co.uk/news/national-grid-eso-claims-world-first-approach-to-inertia-awarding-328m-in-contract>
- [28] Gary W. Chang, Ching-Sheng Chuang, Tai-Ken Lu, and Ching-Chung Wu. Frequency-regulating reserve constrained unit commitment for an isolated power system. *IEEE Transactions on Power Systems*, 28(2):578–586, 2013. doi: 10.1109/TPWRS.2012.2208126.
- [29] Hamed Ahmadi and Hassan Ghasemi. Security-constrained unit commitment with

- linearized system frequency limit constraints. *IEEE Transactions on Power Systems*, 29(4):1536–1545, 2014. doi: 10.1109/TPWRS.2014.2297997.
- [30] Leo Emil Sokoler, Peter Vinter, Runi Baerentsen, Kristian Edlund, and John Bagterp Jorgensen. Contingency-constrained unit commitment in meshed isolated power systems. *IEEE Transactions on Power Systems*, 31:3516–3526, 9 2016. ISSN 08858950. doi: 10.1109/TPWRS.2015.2485781.
- [31] Yunfeng Wen, Wenyuan Li, Gang Huang, and Xuan Liu. Frequency dynamics constrained unit commitment with battery energy storage. *IEEE Transactions on Power Systems*, 31(6):5115–5125, 2016. doi: 10.1109/TPWRS.2016.2521882.
- [32] Fei Teng, Vincenzo Trovato, and Goran Strbac. Stochastic scheduling with inertia-dependent fast frequency response requirements. *IEEE Transactions on Power Systems*, 31(2):1557–1566, 2016. doi: 10.1109/TPWRS.2015.2434837.
- [33] Vincenzo Trovato, Agnes Bialecki, and Anes Dallagi. Unit commitment with inertia-dependent and multispeed allocation of frequency response services. *IEEE Transactions on Power Systems*, 34:1537–1548, 3 2019. ISSN 08858950. doi: 10.1109/TPWRS.2018.2870493.
- [34] Luis Badesa, Fei Teng, and Goran Strbac. Simultaneous scheduling of multiple frequency services in stochastic unit commitment. *IEEE Transactions on Power Systems*, 34:3858–3868, September 2019. ISSN 15580679. doi: 10.1109/TPWRS.2019.2905037.
- [35] Payam Rabbanifar and Nima Amjady. Frequency-constrained unit-commitment using analytical solutions for system frequency responses considering generator contingencies. *IET Generation, Transmission & Distribution*, 14(17):3548–3560, 2020. doi: <https://doi.org/10.1049/iet-gtd.2020.0097>. URL <https://ietresearch.onlinelibrary.wiley.com/doi/abs/10.1049/iet-gtd.2020.0097>.
- [36] Likai Liu, Zechun Hu, Xiaoyu Duan, and Nikhil Pathak. Data-driven distributionally robust optimization for real-time economic dispatch considering secondary frequency

- regulation cost. *IEEE Transactions on Power Systems*, 36:4172–4184, 9 2021. ISSN 15580679. doi: 10.1109/TPWRS.2021.3056390.
- [37] Salar Saberi Oskouee, Sadegh Kamali, and Turaj Amraee. Primary frequency support in unit commitment using a multi-area frequency model with flywheel energy storage. *IEEE Transactions on Power Systems*, 36:5105–5119, 11 2021. ISSN 15580679. doi: 10.1109/TPWRS.2021.3074634.
- [38] Tao Ding, Ziyu Zeng, Ming Qu, Joao P.S. Catalao, and Mohammad Shahidehpour. Two-stage chance-constrained stochastic thermal unit commitment for optimal provision of virtual inertia in wind-storage systems. *IEEE Transactions on Power Systems*, 36(4):3520–3530, 7 2021. ISSN 15580679. doi: 10.1109/TPWRS.2021.3051523.
- [39] Kexin Li, Hongye Guo, Xichen Fang, Shuangquan Liu, Fei Teng, and Qixin Chen. Market mechanism design of inertia and primary frequency response with consideration of energy market. *IEEE Transactions on Power Systems*, pages 1–13, 2022. doi: 10.1109/TPWRS.2022.3223497. Early Access.
- [40] Sufan Jiang, Xin Zhao, Guangsheng Pan, Shan Gao, Chuanshen Wu, Yu Liu, and Sicheng Wang. A novel robust frequency-constrained unit commitment model with emergency control of hvdc. *Energy Reports*, 8:15729–15739, 2022. ISSN 2352-4847. doi: <https://doi.org/10.1016/j.egyr.2022.12.069>. URL <https://www.sciencedirect.com/science/article/pii/S2352484722026695>.
- [41] Diego Ortiz-Villalba, Jacqueline Llanos, Yanira Muñoz-Jadan, Rodrigo Moreno, Claudia Rahmann, and Bikash C. Pal. Optimizing system operation with nadir considerations via simulations of detailed system dynamic responses. *Electric Power Systems Research*, 212:108533, 2022. ISSN 0378-7796. doi: <https://doi.org/10.1016/j.epsr.2022.108533>. URL <https://www.sciencedirect.com/science/article/pii/S0378779622006393>.
- [42] Chaoyue Zhao and Yongpei Guan. Data-driven stochastic unit commitment for in-

- tegrating wind generation. *IEEE Transactions on Power Systems*, 31(4):2587–2596, 2016. doi: 10.1109/TPWRS.2015.2477311.
- [43] Ziliang Jin, Kai Pan, Lei Fan, and Tao Ding. Data-driven look-ahead unit commitment considering forbidden zones and dynamic ramping rates. *IEEE Transactions on Industrial Informatics*, 15(6):3267–3276, 2019. doi: 10.1109/TII.2018.2876316.
- [44] Tao Ding, Qingrun Yang, Xiyuan Liu, Can Huang, Yongheng Yang, Min Wang, and Frede Blaabjerg. Duality-free decomposition based data-driven stochastic security-constrained unit commitment. *IEEE Transactions on Sustainable Energy*, 10(1):82–93, 2019. doi: 10.1109/TSTE.2018.2825361.
- [45] Dimitris T. Lagos and Nikos D. Hatziargyriou. Data-driven frequency dynamic unit commitment for island systems with high res penetration. *IEEE Transactions on Power Systems*, 36(5):4699–4711, 2021. doi: 10.1109/TPWRS.2021.3060891.
- [46] Yichen Zhang, Hantao Cui, Jianzhe Liu, Feng Qiu, Tianqi Hong, Rui Yao, and Fangxing Li. Encoding frequency constraints in preventive unit commitment using deep learning with region-of-interest active sampling. *IEEE Transactions on Power Systems*, 37(3):1942–1955, 5 2022. ISSN 15580679. doi: 10.1109/TPWRS.2021.3110881.
- [47] Lun Yang, Zhihao Li, Yinliang Xu, Jianguo Zhou, and Hongbin Sun. Frequency constrained scheduling under multiple uncertainties via data-driven distributionally robust chance-constrained approach. *IEEE Transactions on Sustainable Energy*, 2022.
- [48] Yukang Shen, Wenchuan Wu, Bin Wang, Yue Yang, and Yi Lin. Data-driven convexification for frequency nadir constraint of unit commitment. *Journal of Modern Power Systems and Clean Energy*, 2022.
- [49] Mohammad Rajabdorri, Enrique Lobato, and Lukas Sigrist. Robust frequency constrained uc using data driven logistic regression for island power systems. *IET Generation, Transmission & Distribution*, 16(24):5069–5083, 2022. doi: <https://doi.org/>

10.1049/gtd2.12658. URL <https://ietresearch.onlinelibrary.wiley.com/doi/abs/10.1049/gtd2.12658>.

- [50] Yu Su, Hongyu Li, Yi Cui, Shutang You, Yiwei Ma, Jingxin Wang, and Yilu Liu. An adaptive pv frequency control strategy based on real-time inertia estimation. *IEEE Transactions on Smart Grid*, 12:2355–2364, 2021. ISSN 19493061. doi: 10.1109/TSG.2020.3045626.
- [51] Hao Li, Ying Qiao, Zongxiang Lu, Baosen Zhang, and Fei Teng. Frequency-constrained stochastic planning towards a high renewable target considering frequency response support from wind power. *IEEE Transactions on Power Systems*, 36(5):4632–4644, 2021. doi: 10.1109/TPWRS.2021.3066991.
- [52] Annabelle Pratt, Dheepak Krishnamurthy, Mark Ruth, Hongyu Wu, Monte Lunacek, and Paul Vaynshenk. Transactive home energy management systems: The impact of their proliferation on the electric grid. *IEEE Electrification Magazine*, 4(4):8–14, Dec. 2016. doi: 10.1109/MELE.2016.2614188.
- [53] Xiangyu Zhang, Manisa Pipattanasomporn, Tao Chen, and Saifur Rahman. An iot-based thermal model learning framework for smart buildings. *IEEE Internet of Things Journal*, 7(1):518–527, 2020. doi: 10.1109/JIOT.2019.2951106.
- [54] June Young Park, Mohamed M. Ouf, Burak Gunay, Yuzhen Peng, William O’Brien, Mikkel Baun Kjærgaard, and Zoltan Nagy. A critical review of field implementations of occupant-centric building controls. *Building and Environment*, 165:106351, 2019. ISSN 0360-1323. doi: <https://doi.org/10.1016/j.buildenv.2019.106351>. URL <https://www.sciencedirect.com/science/article/pii/S036013231930561X>.
- [55] Zoltan Nagy, Burak Gunay, Clayton Miller, Jakob Hahn, Mohamed Ouf, Seungjae Lee, Brodie W. Hobson, Tareq Abuimara, Karol Bandurski, Máira André, Clara-Larissa Lorenz, Sarah Crosby, Bing Dong, Zixin Jiang, Yuzhen Peng, Matteo Favero, June Young Park, Kingsley Nweye, Pedram Nojedehi, Helen Stopps, Lucile Sarran, Connor Brackley, Katherine Bassett, Krissy Govertsen, Nicole Koczorek, Oliver Abele,

- Emily Casavant, Michael Kane, Zheng O'Neill, Tao Yang, Julia Day, Brent Huchuk, Runa T. Hellwig, and Marika Vellei. Ten questions concerning occupant-centric control and operations. *Building and Environment*, page 110518, 2023. ISSN 0360-1323. doi: <https://doi.org/10.1016/j.buildenv.2023.110518>. URL <https://www.sciencedirect.com/science/article/pii/S0360132323005450>.
- [56] Danni Wang, Clifford C. Federspiel, and Francis Rubinstein. Modeling occupancy in single person offices. *Energy and Buildings*, 37:121–126, 2 2005. ISSN 0378-7788. doi: 10.1016/J.ENBUILD.2004.06.015.
- [57] Christoph F. Reinhart. Lightswitch-2002: a model for manual and automated control of electric lighting and blinds. *Solar Energy*, 77:15–28, 1 2004. ISSN 0038-092X. doi: 10.1016/J.SOLENER.2004.04.003.
- [58] J. Page, D. Robinson, N. Morel, and J. L. Scartezzini. A generalised stochastic model for the simulation of occupant presence. *Energy and Buildings*, 40:83–98, 1 2008. ISSN 0378-7788. doi: 10.1016/J.ENBUILD.2007.01.018.
- [59] Laura Klein, Jun Young Kwak, Geoffrey Kavulya, Farrokh Jazizadeh, Burcin Becerik-Gerber, Pradeep Varakantham, and Milind Tambe. Coordinating occupant behavior for building energy and comfort management using multi-agent systems. *Automation in Construction*, 22:525–536, 3 2012. ISSN 09265805. doi: 10.1016/J.AUTCON.2011.11.012.
- [60] Valentina Fabi, Rune Korsholm Andersen, and Stefano Corgnati. Verification of stochastic behavioural models of occupants' interactions with windows in residential buildings. *Building and Environment*, 94:371–383, 12 2015. ISSN 0360-1323. doi: 10.1016/J.BUILDENV.2015.08.016.
- [61] X. Liu, Y. Wu, and H. Wu. PV-EV integrated home energy management considering residential occupant behaviors. *Sustainability (Switzerland)*, 13(24), 2021. ISSN 20711050. doi: 10.3390/su132413826.

- [62] Xuebo Liu, Yingying Wu, Haifeng Zhang, and Hongyu Wu. Hourly occupant clothing decisions in residential HVAC energy management. *Journal of Building Engineering*, 40:102708, 2021. ISSN 2352-7102. doi: <https://doi.org/10.1016/j.jobbe.2021.102708>.
- [63] Xuebo Liu, Yingying Wu, Hang Zhang, Bo Liu, Lawryn Edmonds, and Hongyu Wu. Home energy management with clothing integrated thermal comfort and ev soc concern. In *2022 IEEE Power & Energy Society General Meeting (PESGM)*, pages 01–05, Denver, CO, USA, July 2022. ISBN 978-1-6654-0823-3. doi: 10.1109/PESGM48719.2022.9917003.
- [64] Derong Liu, Yancai Xu, Qinglai Wei, and Xinliang Liu. Residential energy scheduling for variable weather solar energy based on adaptive dynamic programming. *IEEE/CAA Journal of Automatica Sinica*, 5(1):36–46, 2018-01-10. doi: 10.1109/JAS.2017.7510739.
- [65] Hongyu Wu, Annabelle Pratt, Prateek Munankarmi, Monte Lunacek, Sivasathya Pradha Balamurugan, Xuebo Liu, and Paul Spitsen. Impact of model predictive control-enabled home energy management on large-scale distribution systems with photovoltaics. *Advances in Applied Energy*, 6:100094, June 2022. ISSN 2666-7924. doi: 10.1016/J.ADAPEN.2022.100094.
- [66] Fengji Luo, Zhao Yang Dong, Ke Meng, Junhao Wen, Haiming Wang, and Junhua Zhao. An operational planning framework for large-scale thermostatically controlled load dispatch. *IEEE Transactions on Industrial Informatics*, 13(1):217–227, 2017. doi: 10.1109/TII.2016.2515086.
- [67] Fengji Luo, Gianluca Ranzi, Can Wan, Zhao Xu, and Zhao Yang Dong. A multistage home energy management system with residential photovoltaic penetration. *IEEE Transactions on Industrial Informatics*, 15(1):116–126, 2019. doi: 10.1109/TII.2018.2871159.
- [68] Y. Liu, D. Zhang, and H. B. Gooi. Optimization strategy based on deep reinforcement learning for home energy management. *CSEE Journal of Power and Energy Systems*,

- 6(3):572–582, 2020-09. ISSN 2096-0042. doi: 10.17775/CSEEJPES.2019.02890. CSEE Journal of Power and Energy Systems.
- [69] G. Barone, A. Buonomano, C. Forzano, G.F. Giuzio, A. Palombo, and G. Russo. A new thermal comfort model based on physiological parameters for the smart design and control of energy-efficient hvac systems. *Renewable and Sustainable Energy Reviews*, 173: 113015, 2023. ISSN 1364-0321. doi: <https://doi.org/10.1016/j.rser.2022.113015>. URL <https://www.sciencedirect.com/science/article/pii/S1364032122008966>.
- [70] Xingwang Zhao, Yonggao Yin, Zhiqiang He, and Zhipeng Deng. State-of-the-art, challenges and new perspectives of thermal comfort demand law for on-demand intelligent control of heating, ventilation, and air conditioning systems. *Energy and Buildings*, 295:113325, 2023. ISSN 0378-7788. doi: <https://doi.org/10.1016/j.enbuild.2023.113325>. URL <https://www.sciencedirect.com/science/article/pii/S0378778823005558>.
- [71] Yangyang Fu, Shichao Xu, Qi Zhu, Zheng O’Neill, and Veronica Adetola. How good are learning-based control v.s. model-based control for load shifting? investigations on a single zone building energy system. *Energy*, 273:127073, 2023. ISSN 0360-5442. doi: <https://doi.org/10.1016/j.energy.2023.127073>. URL <https://www.sciencedirect.com/science/article/pii/S036054422300467X>.
- [72] Yujian Ye, Dawei Qiu, Huiyu Wang, Yi Tang, and Goran Strbac. Real-time autonomous residential demand response management based on twin delayed deep deterministic policy gradient learning. *Energies*, 14(3):531, Jan. 2021. ISSN 1996-1073. doi: 10.3390/EN14030531.
- [73] Yan Zhang, Xuemei Bai, Franklin P. Mills, and John C.V. Pezzey. Rethinking the role of occupant behavior in building energy performance: A review. *Energy and Buildings*, 172:279–294, 8 2018. ISSN 0378-7788. doi: 10.1016/J.ENBUILD.2018.05.017.
- [74] A. Kathirgamanathan, E. Mangina, and D.P. Finn. Development of a soft actor critic

- deep reinforcement learning approach for harnessing energy flexibility in a large office building. *Energy and AI*, 5, 2021. doi: 10.1016/j.egyai.2021.100101.
- [75] R.Z. Homod, H. Togun, A. Kadhim Hussein, F. Noraldeem Al-Mousawi, Z.M. Yaseen, W. Al-Kouz, H.J. Abd, O.A. Alawi, M. Goodarzi, and O.A. Hussein. Dynamics analysis of a novel hybrid deep clustering for unsupervised learning by reinforcement of multi-agent to energy saving in intelligent buildings. *Applied Energy*, 313, 2022. doi: 10.1016/j.apenergy.2022.118863.
- [76] Liang Yu, Yi Sun, Zhanbo Xu, Chao Shen, Dong Yue, Tao Jiang, and Xiaohong Guan. Multi-agent deep reinforcement learning for hvac control in commercial buildings. *IEEE Transactions on Smart Grid*, 12(1):407–419, Jan. 2021. ISSN 19493061. doi: 10.1109/TSG.2020.3011739.
- [77] A. Naug, M. Quinones-Grueiro, and G. Biswas. Deep reinforcement learning control for non-stationary building energy management. *Energy and Buildings*, 277, 2022. doi: 10.1016/j.enbuild.2022.112584.
- [78] D. Bayer and M. Pruckner. Enhancing the performance of multi-agent reinforcement learning for controlling hvac systems. pages 187–194, 2022. ISBN 9781665406642. doi: 10.1109/SusTech53338.2022.9794179.
- [79] Zhiang Zhang and Khee Poh Lam. Practical implementation and evaluation of deep reinforcement learning control for a radiant heating system. In *Proceedings of the 5th Conference on Systems for Built Environments*, pages 148–157, Shenzhen, China, Nov. 2018. ISBN 978-1-4503-5951-1. doi: 10.1145/3276774.3276775.
- [80] Donald Azuatalam, Wee Lih Lee, Frits de Nijs, and Ariel Liebman. Reinforcement learning for whole-building hvac control and demand response. *Energy and AI*, 2, 11 2020. ISSN 26665468. doi: 10.1016/J.EGYAI.2020.100020.
- [81] Jie Li, Wei Zhang, Guanyu Gao, Yonggang Wen, Guangyu Jin, and Georgios Christopoulos. Toward intelligent multizone thermal control with multiagent deep

- reinforcement learning. *IEEE Internet of Things Journal*, 8(14):11150–11162, 2021. doi: 10.1109/JIOT.2021.3051400.
- [82] D. Coraci, S. Brandi, M.S. Piscitelli, and A. Capozzoli. Online implementation of a soft actor-critic agent to enhance indoor temperature control and energy efficiency in buildings. *Energies*, 14, 2021. doi: 10.3390/en14040997.
- [83] D. Deltetto, D. Coraci, G. Pinto, M.S. Piscitelli, and A. Capozzoli. Exploring the potentialities of deep reinforcement learning for incentive-based demand response in a cluster of small commercial buildings. *Energies*, 14, 2021. doi: 10.3390/en14102933.
- [84] Z. Li, Z. Sun, Q. Meng, Y. Wang, and Y. Li. Reinforcement learning of room temperature set-point of thermal storage air-conditioning system with demand response. *Energy and Buildings*, 259, 2022. doi: 10.1016/j.enbuild.2022.111903.
- [85] Mohammad Esrafilian-Najafabadi and Fariborz Haghighat. Towards self-learning control of hvac systems with the consideration of dynamic occupancy patterns: Application of model-free deep reinforcement learning. *Building and Environment*, 226:109747, 2022. ISSN 0360-1323. doi: <https://doi.org/10.1016/j.buildenv.2022.109747>. URL <https://www.sciencedirect.com/science/article/pii/S0360132322009775>.
- [86] Xing Lu, Yangyang Fu, and Zheng O’Neill. Benchmarking high performance hvac rule-based controls with advanced intelligent controllers: A case study in a multi-zone system in modelica. *Energy and Buildings*, 284:112854, 4 2023. ISSN 0378-7788. doi: 10.1016/J.ENBUILD.2023.112854.
- [87] S. Qiu, Z. Li, Z. Pang, Z. Li, and Y. Tao. Multi-agent optimal control for central chiller plants using reinforcement learning and game theory. *Systems*, 11, 2023. doi: 10.3390/systems11030136.
- [88] Henry P. Huntington, Jennifer I. Schmidt, Philip A. Loring, Erin Whitney, Srijan Aggarwal, Amanda G. Byrd, Subhabrata Dev, Aaron D. Dotson, Daisy Huang, Barbara Johnson, Justus Karenzi, Henry J.F. Penn, Alex Anna Salmon, Daniel J.

- Sambor, William E. Schnabel, Richard W. Wies, and Michelle Wilber. Applying the food–energy–water nexus concept at the local scale. *Nature Sustainability* 2021 4:8, 4:672–679, 5 2021. ISSN 2398-9629. doi: 10.1038/s41893-021-00719-1. URL <https://www.nature.com/articles/s41893-021-00719-1>.
- [89] Meng Wang, Yen Wei, Xin Wang, Ruoxin Li, Shiyu Zhang, Ke Wang, Rupeng Wang, Haixing Chang, Chengyu Wang, Nanqi Ren, and Shih-Hsin Ho. An integrated system with functions of solar desalination, power generation and crop irrigation. *Nature Water* 2023 1:8, 1:716–724, 8 2023. ISSN 2731-6084. doi: 10.1038/s44221-023-00118-0. URL <https://www.nature.com/articles/s44221-023-00118-0>.
- [90] Ning Ma, Yukun Zhang, Rui Zhang, Wen Zhang, and Xiaojiao Li. Comprehensive review of food-energy-water nexus at the community scale. *Journal of Cleaner Production*, 420:138311, 9 2023. ISSN 0959-6526. doi: 10.1016/J.JCLEPRO.2023.138311.
- [91] Kyle Proctor, Seyed M.H. Tabatabaie, and Ganti S. Murthy. Gateway to the perspectives of the food-energy-water nexus. *Science of The Total Environment*, 764:142852, 4 2021. ISSN 0048-9697. doi: 10.1016/J.SCITOTENV.2020.142852.
- [92] Hamdi Abdi, Maryam Shahbazitabar, and Behnam Mohammadi-Ivatloo. Food, energy and water nexus: A brief review of definitions, research, and challenges. *Inventions* 2020, Vol. 5, Page 56, 5:56, 11 2020. ISSN 2411-5134. doi: 10.3390/INVENTIONS5040056. URL <https://www.mdpi.com/2411-5134/5/4/56/html><https://www.mdpi.com/2411-5134/5/4/56>.
- [93] Jirapat Phetheet, Mary C. Hill, Robert W. Barron, Benjamin J. Gray, Hongyu Wu, Vincent Amanor-Boadu, Wade Heger, Isaya Kisekka, Bill Golden, and Matthew W. Rossi. Relating agriculture, energy, and water decisions to farm incomes and climate projections using two freeware programs, fewcalc and dssat. *Agricultural Systems*, 193: 103222, 2021. ISSN 0308-521X. doi: <https://doi.org/10.1016/j.agry.2021.103222>. URL <https://www.sciencedirect.com/science/article/pii/S0308521X2100175X>.
- [94] Jirapat Phetheet, Mary C. Hill, Robert W. Barron, Matthew W. Rossi, Vincent

- Amanor-Boadu, Hongyu Wu, and Isaya Kisekka. Consequences of climate change on food-energy-water systems in arid regions without agricultural adaptation, analyzed using fewcalc and dssat. *Resources, Conservation and Recycling*, 168:105309, 2021. ISSN 0921-3449. doi: <https://doi.org/10.1016/j.resconrec.2020.105309>. URL <https://www.sciencedirect.com/science/article/pii/S0921344920306248>.
- [95] Jean-Michel Meynard, Marc Doré, Edouard Lucas, Jacques Mikic, Muriel Valantin-Morison, and Sylvain Jeuffroy. Effect of pea–barley intercropping on grain yield and yield protein concentration of both species. *Agronomy for Sustainable Development*, 31(2):273–281, 2011. doi: 10.1051/agro/2010029. URL <https://link.springer.com/article/10.1051/agro/2010029>.
- [96] A. Chel and G. Kaushik. Renewable energy for sustainable agriculture. *Agronomy for Sustainable Development*, 31:91–118, 2011. doi: 10.1051/agro/2010029. URL <https://doi.org/10.1051/agro/2010029>.
- [97] Bihter Güven, İpek Baz, Beyza Kocaoğlu, Elif Toprak, Duygun Erol Barkana, and Bahar Soğutmaz Özdemir. *Smart Farming Technologies for Sustainable Agriculture: From Food to Energy*, pages 481–506. Springer International Publishing, Cham, 2023. ISBN 978-3-031-24942-6. doi: 10.1007/978-3-031-24942-6_22. URL https://doi.org/10.1007/978-3-031-24942-6_22.
- [98] Md Momtazur Rahman, Imran Khan, David Luke Field, Kuaanan Techato, and Kamal Alameh. Powering agriculture: Present status, future potential, and challenges of renewable energy applications. *Renewable Energy*, 188:731–749, 2022. ISSN 0960-1481. doi: <https://doi.org/10.1016/j.renene.2022.02.065>. URL <https://www.sciencedirect.com/science/article/pii/S0960148122002075>.
- [99] Christan Hail R. Mendigoria, Ronnie S. Concepcion, Elmer P. Dadios, Edwin Sybingco, and Argel A. Bandala. Energy management trends for sustainability in agriculture industry of the philippines. In *2020 IEEE 12th International Conference on Humanoid, Nanotechnology, Information Technology, Communication and Control, Environment,*

- and Management (HNICEM), pages 1–6, 2020. doi: 10.1109/HNICEM51456.2020.9400098.
- [100] Ugo Bardi, Toufic El Asmar, and Alessandro Lavacchi. Turning electricity into food: the role of renewable energy in the future of agriculture. *Journal of Cleaner Production*, 53:224–231, 2013. ISSN 0959-6526. doi: <https://doi.org/10.1016/j.jclepro.2013.04.014>. URL <https://www.sciencedirect.com/science/article/pii/S0959652613002291>.
- [101] Lawryn Edmonds, Xuebo Liu, and Hongyu Wu. A demand-responsive green ammonia plant and its impact on the electricity distribution system. In *2022 IEEE Power & Energy Society General Meeting (PESGM)*, 2022. doi: 10.1109/PESGM48719.2022.9917204.
- [102] Lawryn Edmonds, Melanie Derby, Mary Hill, and Hongyu Wu. Coordinated operation of water and electricity distribution networks with variable renewable energy and distribution locational marginal pricing. *Renewable Energy*, 177:1438–1450, 2021. ISSN 0960-1481. doi: <https://doi.org/10.1016/j.renene.2021.05.168>. URL <https://www.sciencedirect.com/science/article/pii/S0960148121008570>.
- [103] Lawryn Edmonds, Peter Pfromm, Vincent Amanor-Boadu, Mary Hill, and Hongyu Wu. Green ammonia production-enabled demand flexibility in agricultural community microgrids with distributed renewables. *Sustainable Energy, Grids and Networks*, 31: 100736, 2022. ISSN 2352-4677. doi: <https://doi.org/10.1016/j.segan.2022.100736>.
- [104] X. Liu, H. Wu, L. Wang, and M.N. Faqiry. Stochastic home energy management system via approximate dynamic programming. *IET Energy Systems Integration*, 2 (4), Oct 2020. ISSN 25168401. doi: 10.1049/iet-esi.2020.0060.
- [105] Xuebo Liu, Xin Fang, Ningchao Gao, Haoyu Yuan, Andy Hoke, Hongyu Wu, and Jin Tan. Frequency nadir constrained unit commitment for high renewable penetration island power systems. *IEEE Open Access Journal of Power and Energy*, 2024.

- [106] Moritz Müller, Marcel Pfeifer, Dorian Holtz, and Karsten Müller. Comparison of green ammonia and green hydrogen pathways in terms of energy efficiency. *Fuel*, 357: 129843, 2024. ISSN 0016-2361. doi: <https://doi.org/10.1016/j.fuel.2023.129843>. URL <https://www.sciencedirect.com/science/article/pii/S0016236123024572>.
- [107] Houssam Bouaboula, Mohammed Ouikhalfan, Ismael Saadoune, Jamal Chaouki, Abdelghafour Zaabout, and Youssef Belmabkhout. Addressing sustainable energy intermittence for green ammonia production. *Energy Reports*, 9:4507–4517, 2023. ISSN 2352-4847. doi: <https://doi.org/10.1016/j.egy.2023.03.093>. URL <https://www.sciencedirect.com/science/article/pii/S235248472300330X>.
- [108] Antonio Sánchez, Elena Castellano, Mariano Martín, and Pastora Vega. Evaluating ammonia as green fuel for power generation: A thermo-chemical perspective. *Applied Energy*, 293:116956, 2021. ISSN 0306-2619. doi: <https://doi.org/10.1016/j.apenergy.2021.116956>. URL <https://www.sciencedirect.com/science/article/pii/S0306261921004323>.
- [109] S Joseph Sekhar, Melvin S Samuel, Godwin Glivin, T G Le, and Thangavel Mathimani. Production and utilization of green ammonia for decarbonizing the energy sector with a discrete focus on sustainable development goals and environmental impact and technical hurdles. *Fuel*, 360:130626, 2024. ISSN 0016-2361. doi: <https://doi.org/10.1016/j.fuel.2023.130626>. URL <https://www.sciencedirect.com/science/article/pii/S0016236123032404>.
- [110] Sergey Klyapovskiy, Yi Zheng, Shi You, and Henrik W Bindner. Optimal operation of the hydrogen-based energy management system with p2x demand response and ammonia plant. *Applied Energy*, 304:117559, 2021. ISSN 0306-2619. doi: <https://doi.org/10.1016/j.apenergy.2021.117559>. URL <https://www.sciencedirect.com/science/article/pii/S0306261921009375>.
- [111] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg

- Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, Feb. 2015. ISSN 1476-4687. doi: 10.1038/nature14236.
- [112] D. Jang, L. Spangher, S. Nadarajah, and C. Spanos. Deep reinforcement learning with planning guardrails for building energy demand response. *Energy and AI*, 11, 2023. doi: 10.1016/j.egyai.2022.100204.
- [113] Xianzhong Ding and Wan Du. Optimizing irrigation efficiency using deep reinforcement learning in the field. *ACM Trans. Sen. Netw.*, 20(4), July 2024. ISSN 1550-4859. doi: 10.1145/3662182. URL <https://doi.org/10.1145/3662182>.
- [114] Xianzhong Ding and Wan Du. Drlic: Deep reinforcement learning for irrigation control. In *2022 21st ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 41–53, 2022. doi: 10.1109/IPSN54338.2022.00011.
- [115] Mengting Chen, Yuanlai Cui, Xiaonan Wang, Hengwang Xie, Fangping Liu, Tongyuan Luo, Shizong Zheng, and Yufeng Luo. A reinforcement learning approach to irrigation decision-making for rice using weather forecasts. *Agricultural Water Management*, 250:106838, 2021. ISSN 0378-3774. doi: <https://doi.org/10.1016/j.agwat.2021.106838>. URL <https://www.sciencedirect.com/science/article/pii/S0378377421001037>.
- [116] Khadijeh Alibabaei, Pedro D. Gaspar, Eduardo Assunção, Saeid Alirezazadeh, and Tânia M. Lima. Irrigation optimization with a deep reinforcement learning model: Case study on a site in portugal. *Agricultural Water Management*, 263:107480, 2022. ISSN 0378-3774. doi: <https://doi.org/10.1016/j.agwat.2022.107480>. URL <https://www.sciencedirect.com/science/article/pii/S0378377422000270>.
- [117] Xuebo Liu, Jian Xie, Xin Fang, Haoyu Yuan, Bin Wang, Hongyu Wu, and Jin Tan. A comparison of machine learning methods for frequency nadir estimation in power

- systems. In *2022 IEEE Kansas Power and Energy Conference (KPEC)*, pages 1–5, Manhattan, KS, USA, 2022. doi: 10.1109/KPEC54747.2022.9814724.
- [118] Xuebo Liu, Xin Fang, Ningchao Gao, Haoyu Yuan, Andy Hoke, Hongyu Wu, and Jin Tan. Frequency nadir constrained unit commitment for high renewable penetration island power systems. *IEEE Open Access Journal of Power and Energy*, 11:141–153, 2024. doi: 10.1109/OAJPE.2024.3370504.
- [119] Xuebo Liu, Yingying Wu, and Hongyu Wu. Enhancing hvac energy management through multi-zone occupant-centric approach: A multi-agent deep reinforcement learning solution. *Energy and Buildings*, 303:113770, 2024.
- [120] Xuebo Liu, Alex Modarresi, Lawryn Kiboma, Hongyu Wu, John Symons, and Mary Hill. Budget-constrained sizing of renewable and energy storage systems for farm-scale ammonia production within the food-energy-water nexus. In *2024 IEEE Green Technologies Conference (GreenTech)*, pages 218–222, 2024. doi: 10.1109/GreenTech58819.2024.10520453.
- [121] ScienceDaily. Title of the article. <https://www.sciencedaily.com/releases/2000/03/000303075501.htm>, 2000. Accessed: 2025-03-24.
- [122] International Energy Agency (IEA). Renewables 2023 – electricity. <https://www.iea.org/reports/renewables-2023/electricity>, 2023. Accessed: 2025-03-24.
- [123] Celine B. Santiago, Jing Yao Guo, and Matthew S. Sigman. Predictive and mechanistic multivariate linear regression models for reaction development. *Chemical Science*, 9: 2398–2412, February 2018. ISSN 2041-6539. doi: 10.1039/C7SC04679K.
- [124] Ana Fernandez-Guillamon, Emilio Gomez-Lazaro, Eduard Muljadi, and Ángel Molina-Garcia. Power systems with high renewable energy sources: A review of inertia and frequency control strategies over time. *Renewable and Sustainable Energy Reviews*, 115:109369, 2019. ISSN 1364-0321.
- [125] Shengfei Yin, Jianhui Wang, Yanling Lin, Xin Fang, Jin Tan, and Haoyu Yuan. Prac-

- tical operations of energy storage providing ancillary services: From day-ahead to real-time; practical operations of energy storage providing ancillary services: From day-ahead to real-time. 2020. doi: 10.1109/NAPS50074.2021.9449810.
- [126] Xin Fang, Jin Tan, Haoyu Yuan, Shengfei Yin, and Jianhui Wang. Providing ancillary services with photovoltaic generation in multi-timescale grid operation. Institute of Electrical and Electronics Engineers Inc., 4 2021. ISBN 9781728181929. doi: 10.1109/NAPS50074.2021.9449700.
- [127] M. Nazif Faqiry, Lawryn Edmonds, Hongyu Wu, and Anil Pahwa. Distribution locational marginal price-based transactive day-ahead market with variable renewable generation. *Applied Energy*, 259:114103, 2 2020. ISSN 0306-2619. doi: 10.1016/J.APENERGY.2019.114103.
- [128] Ananth M Palani, Hongyu Wu, and Medhat M Morcos. A fast penalty-based gauss-seidel method for stochastic unit commitment with uncertain load and wind generation. *IEEE Open Access Journal of Power and Energy*, 8:211–222, 2021. doi: 10.1109/OAJPE.2021.3079150.
- [129] Robert Eriksson, Niklas Modig, and Katherine Elkington. Synthetic inertia versus fast frequency response: a definition; synthetic inertia versus fast frequency response: a definition. *IET Renewable Power Generation*, 12:507–514, 2017. ISSN 1752-1416. doi: 10.1049/iet-rpg.2017.0370.
- [130] Lin Zhu, Yinfeng Zhao, Yi Cui, Shutang You, Wenpeng Yu, Shengyuan Liu, He Yin, Chang Chen, Yuru Wu, Wei Qiu, Mirka Mandich, Hongyu Li, Adedasola Ademola, Chengwen Zhang, Chujie Zeng, Xinlan Jia, Weikang Wang, Haoyu Yuan, Huaiguang Jiang, Jin Tan, and Yilu Liu. Adding power of artificial intelligence to situational awareness of large interconnections dominated by inverter-based resources. *High Voltage*, 6:924–937, 12 2021. ISSN 23977264. doi: 10.1049/HVE2.12157.
- [131] Dawei Geng, Haifeng Zhang, and Hongyu Wu. Short-term wind speed prediction based

- on principal component analysis and LSTM. *Applied Sciences*, 10(13):4416, 2020. doi: 10.3390/app10134416.
- [132] Harold R. Chamorro, Alvaro D. Orjuela-Canon, David Ganger, Mattias Persson, Francisco Gonzalez-Longatt, Vijay K. Sood, and Wilmar Martinez. Nadir frequency estimation in low-inertia power systems. *IEEE International Symposium on Industrial Electronics*, 2020-June:918–922, 6 2020. doi: 10.1109/ISIE45063.2020.9152296.
- [133] Liu Liu, Weidong Li, Yu Ba, Jiakai Shen, Cuicui Jin, and Kerui Wen. An analytical model for frequency nadir prediction following a major disturbance. *IEEE Transactions on Power Systems*, 35:2527–2536, 7 2020. ISSN 15580679. doi: 10.1109/TPWRS.2019.2963706.
- [134] PSS@E — high-performance transmission planning and analysis software, . URL <https://new.siemens.com/global/en/products/energy/energy-automation-and-smart-grid/pss-software/pss-e.html>.
- [135] Haoyu Yuan, Reetam Sen Biswas, Jin Tan, and Yingchen Zhang. Developing a reduced 240-bus wecc dynamic model for frequency response study of high renewable integration. In *2020 IEEE/PES Transmission and Distribution Conference and Exposition (T D)*, pages 1–5, 2020. doi: 10.1109/TD39804.2020.9299666.
- [136] ”test case repository for high renewable study”, . URL <https://www.nrel.gov/grid/test-case-repository.html>.
- [137] G.A.F. Seber and A.J. Lee. *Linear Regression Analysis*. Wiley Series in Probability and Statistics. Wiley, 2012.
- [138] Jerome H. Friedman. Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, 29(5):1189 – 1232, 2001. doi: 10.1214/aos/1013203451.
- [139] Mariette Awad and Rahul Khanna. *Efficient Learning Machines: Theories, Concepts, and Applications for Engineers and System Designers*, pages 67–80. Apress, Berkeley, CA, 2015. ISBN 978-1-4302-5990-9. doi: 10.1007/978-1-4302-5990-9_4.

- [140] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [141] Tianqi Chen and Carlos Guestrin. Xgboost. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Aug 2016. doi: 10.1145/2939672.2939785.
- [142] Haoyu Yuan and Jin Tan. Multi-timescale integrated dynamic and scheduling model (midas-solar), 2020.
- [143] Jin Tan, Andy Hoke, Haoyu Yuan, Bin Wang, Rick Wallace Kenyon, Xin Fang, Przemyslaw Koralewicz, Emanuel Mendiola, Yingchen Zhang, Yilu Liu, Shutang You, Mirka Mandich, Annie Zhao, Jianhui Wang, Shengfei Yin, Yanling Lin, Erik Ela, Vikas Singhvi, Parag Mitra, and Robert Entrike. Final technical report: Multi-timescale integrated dynamics and scheduling for solar (midas-solar). April 2023. doi: 10.2172/1972321. URL <https://www.osti.gov/biblio/1972321>.
- [144] NREL: Multi-timescale integrated dynamic and scheduling. URL <https://www.nrel.gov/grid/midas.html>.
- [145] Prabha Kundur. *Power System Stability and Control*. McGraw-Hill, 1994.
- [146] Héctor Chávez, Ross Baldick, and Sandip Sharma. Governor rate-constrained opf for primary frequency control adequacy. *IEEE Transactions on Power Systems*, 29(3): 1473–1480, 2014. doi: 10.1109/TPWRS.2014.2298838.
- [147] Yen-Yu Lee and Ross Baldick. A frequency-constrained stochastic economic dispatch model. *IEEE Transactions on Power Systems*, 28(3):2301–2312, 2013. doi: 10.1109/TPWRS.2012.2236108.
- [148] Smita Rath, Alakananda Tripathy, and Alok Ranjan Tripathy. Prediction of new active cases of coronavirus disease (covid-19) pandemic using multiple linear regression model. *Diabetes & metabolic syndrome*, 14(5):1467–1474, September 2020. ISSN 1878-0334. doi: 10.1016/J.DSX.2020.07.045.

- [149] E Ela, M Milligan, and B Kirby. Operating reserves and variable generation. Technical Report NREL/TP-5500-51978, National Renewable Energy Laboratory, Golden, CO (United States), August 2011.
- [150] Paul Denholm, Trieu Mai, Richard Kenyon, Benjamin Kroposki, and Mark O’Malley. Inertia and the power grid: A guide without the spin. Technical Report NREL/TP-6120-73856, National Renewable Energy Laboratory, Golden, CO (United States), May 2020.
- [151] International Energy Agency (IEA). Buildings – topics. <https://www.iea.org/energy-system/buildings>, 2023. Accessed: 2025-03-24.
- [152] P. O. Fanger. Thermal comfort. analysis and applications in environmental engineering. 1970. Copenhagen: Danish Technical Press. 1970.
- [153] UNI EN ISO 7730. Ergonomics of the thermal environment — analytical determination and interpretation of thermal comfort using calculation of the PMV and PPD indices and local thermal comfort criteria. 2005.
- [154] ASHRAE Standard 55/2004. *Thermal environmental conditions for human occupancy*. 2004.
- [155] C. Buratti, P. Ricciardi, and M. Vergoni. HVAC systems testing and check: A simplified model to predict thermal comfort conditions in moderate environments. *Applied Energy*, 104:117–127, 2013. doi: 10.1016/j.apenergy.2012.11.015.
- [156] Qing-Shan Jia, Junjie Wu, Zijian Wu, and Xiaohong Guan. Event-based HVAC control—a complexity-based approach. *IEEE Transactions on Automation Science and Engineering*, 15(4):1909–1919, 2018. doi: 10.1109/TASE.2018.2844258.
- [157] United States Department of Energy. EnergyPlus, 2021. URL <https://energyplus.net/>.
- [158] Warren Powell. *Approximate dynamic programming: Solving the curses of dimensionality*. John Wiley and Sons, second edition, 2011-09-27. ISBN 978-0-470-60445-8.

- [159] Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, 4th edition edition, 2012-06-18. ISBN 978-1-886529-08-3.
- [160] C. Keerthisinghe, G. Verbič, and A. C. Chapman. Energy management of PV-storage systems: ADP approach with temporal difference learning. In *2016 Power Systems Computation Conference (PSCC)*, pages 1–7, 2016-06. doi: 10.1109/PSCC.2016.7540924.
- [161] Q. Wei, D. Liu, G. Shi, and Y. Liu. Multibattery optimal coordination control for home energy management systems via distributed iterative adaptive dynamic programming. *IEEE Transactions on Industrial Electronics*, 62(7):4203–4214, 2015. doi: 10.1109/TIE.2014.2388198.
- [162] T. Huang and D. Liu. Residential energy system control and management using adaptive dynamic programming. In *The 2011 International Joint Conference on Neural Networks*, pages 119–124, 2011-07. doi: 10.1109/IJCNN.2011.6033209.
- [163] S. Squartini, D. Fuselli, M. Boaro, F. De Angelis, and F. Piazza. Home energy resource scheduling algorithms and their dependency on the battery model. In *2013 IEEE Computational Intelligence Applications in Smart Grid (CIASG)*, pages 122–129, 2013-04. doi: 10.1109/CIASG.2013.6611508.
- [164] Q. Wei, D. Liu, Y. Liu, and R. Song. Optimal constrained self-learning battery sequential management in microgrid via adaptive dynamic programming. *IEEE/CAA Journal of Automatica Sinica*, 4(2):168–176, 2017-04. doi: 10.1109/JAS.2016.7510262.
- [165] Jianquan Zhu, Xiemin Mo, Tao Zhu, Ye Guo, Tianyun Luo, and Mingbo Liu. Real-time stochastic operation strategy of a microgrid using approximate dynamic programming-based spatiotemporal decomposition approach. *IET Renewable Power Generation*, 13(16):3061–3070, 2019-10-15. doi: 10.1049/iet-rpg.2019.0536.
- [166] F. Borghesan, R. Vignali, L. Piroddi, M. Prandini, and M. Strelec. Approximate dynamic programming-based control of a building cooling system with thermal storage.

- In *IEEE PES ISGT Europe 2013*, pages 1–5, 2013-10. doi: 10.1109/ISGTEurope.2013.6695463.
- [167] N. Ceriani, R. Vignali, L. Piroddi, and M. Prandini. An approximate dynamic programming approach to the energy management of a building cooling system. In *2013 European Control Conference (ECC)*, pages 2026–2031, 2013-07. doi: 10.23919/ECC.2013.6669287.
- [168] Siyun Chen, Ting Liu, Feng Gao, Jianting Ji, Zhanbo Xu, Buyue Qian, Hongyu Wu, and Xiaohong Guan. Butler, not servant: A human-centric smart home energy management system. *IEEE Communications Magazine*, 55(2):27–33, 2017. doi: 10.1109/MCOM.2017.1600699CM.
- [169] Annabelle Pratt, Bibaswan Banerjee, and Tonderai Nemarundwe. Proof-of-concept home energy management system autonomously controlling space heating. In *2013 IEEE Power Energy Society General Meeting*, pages 1–5, Vancouver, BC, Canada, 2013. doi: 10.1109/PESMG.2013.6672709.
- [170] Hongyu Wu, Annabelle Pratt, and Sudipta Chakraborty. Stochastic optimal scheduling of residential appliances with renewable energy sources. In *IEEE Power Energy Society General Meeting*, pages 1–5, 2015. doi: 10.1109/PESGM.2015.7286584.
- [171] ASHRAE 2017. *Handbook Fundamentals: Inch-Pound Edition*. 2017.
- [172] Stefano Schiavon and Kwang Ho Lee. Dynamic predictive clothing insulation models based on outdoor air and indoor operative temperatures. *Building and Environment*, 59:250–260, 2013-01-01. doi: 10.1016/j.buildenv.2012.08.024.
- [173] Clo - clothing and thermal insulation, 2021. URL https://www.engineeringtoolbox.com/clo-clothing-thermal-insulation-d_732.html.
- [174] F. A. Gers, J. Schmidhuber, and F. Cummins. Learning to forget: continual prediction with LSTM. In *9th International Conference on Artificial Neural Networks: ICANN*, pages 850–855, 1999-01-01. doi: 10.1049/cp:19991218.

- [175] F. A. Gers and E. Schmidhuber. LSTM recurrent networks learn simple context-free and context-sensitive languages. *IEEE Transactions on Neural Networks*, 12(6):1333–1340, 2001-11. doi: 10.1109/72.963769.
- [176] Abdelhadi Azzouni and Guy Pujolle. A long short-term memory recurrent neural network framework for network traffic matrix prediction. *CoRR*, abs/1705.05690. (2017).
- [177] Bryan Palmintier, Dheepak Krishnamurthy, and Hongyu Wu. dynamo, 2018.
- [178] M. Y. Beshir and J. D. Ramsey. Comparison between male and female subjective estimates of thermal effects and sensations. *Applied Ergonomics*, 12(1):29–33, 1981-03-01. doi: 10.1016/0003-6870(81)90091-0.
- [179] Michael Wetter. Co-simulation of building energy and control systems with the building controls virtual test bed. *Journal of Building Performance Simulation*, 4(3):185–203, Sep. 2010. ISSN 19401493. doi: 10.1080/19401493.2010.518631.
- [180] Xuezheng Wang and Bing Dong. Physics-informed hierarchical data-driven predictive control for building hvac systems to achieve energy and health nexus. *Energy and Buildings*, 291:113088, 2023. ISSN 0378-7788. doi: <https://doi.org/10.1016/j.enbuild.2023.113088>. URL <https://www.sciencedirect.com/science/article/pii/S0378778823003183>.
- [181] Han Li, Hicham Johra, Flavia de Andrade Pereira, Tianzhen Hong, Jérôme Le Dréau, Anthony Maturo, Mingjun Wei, Yapan Liu, Ali Saberi-Derakhtenjani, Zoltan Nagy, Anna Marszal-Pomianowska, Donal Finn, Shohei Miyata, Kathryn Kaspar, Kingsley Nweye, Zheng O’Neill, Fabiano Pallonetto, and Bing Dong. Data-driven key performance indicators and datasets for building energy flexibility: A review and perspectives. *Applied Energy*, 343:121217, 2023. ISSN 0306-2619. doi: <https://doi.org/10.1016/j.apenergy.2023.121217>. URL <https://www.sciencedirect.com/science/article/pii/S0306261923005810>.
- [182] Meng Kong, Bing Dong, Rongpeng Zhang, and Zheng O’Neill. Hvac energy savings,

- thermal comfort and air quality for occupant-centric control through a side-by-side experimental study. *Applied Energy*, 306:117987, 2022. ISSN 0306-2619. doi: <https://doi.org/10.1016/j.apenergy.2021.117987>. URL <https://www.sciencedirect.com/science/article/pii/S0306261921012903>.
- [183] H. Burak Gunay, William O’Brien, and Ian Beausoleil-Morrison. A critical review of observation studies, modeling, and simulation of adaptive occupant behaviors in offices. *Building and Environment*, 70:31–47, 12 2013. ISSN 0360-1323. doi: 10.1016/J.BUILDENV.2013.07.020.
- [184] H. Burak Gunay, William O’Brien, and Ian Beausoleil-Morrison. Implementation and comparison of existing occupant behaviour models in energy-plus. <http://dx.doi.org/10.1080/19401493.2015.1102969>, 9:567–588, 11 2015. ISSN 19401507. doi: 10.1080/19401493.2015.1102969. URL <https://www.tandfonline.com/doi/abs/10.1080/19401493.2015.1102969>.
- [185] Salvatore Carlucci, Marilena De Simone, Steven K. Firth, Mikkel B. Kjærgaard, Romana Markovic, Mohammad Saiedur Rahaman, Masab Khalid Annageeb, Silvia Biandrate, Anooshmita Das, Jakub Wladyslaw Dziedzic, Gianmarco Fajilla, Matteo Favero, Martina Ferrando, Jakob Hahn, Mengjie Han, Yuzhen Peng, Flora Salim, Arno Schlüter, and Christoph van Treeck. Modeling occupant behavior in buildings. *Building and Environment*, 174:106768, 5 2020. ISSN 0360-1323. doi: 10.1016/J.BUILDENV.2020.106768.
- [186] Haneul Choi, Bonghoon Jeong, Joosang Lee, Hooseung Na, Kyungmo Kang, and Taeyeon Kim. Deep-vision-based metabolic rate and clothing insulation estimation for occupant-centric control. *Building and Environment*, 221:109345, 2022. ISSN 0360-1323. doi: <https://doi.org/10.1016/j.buildenv.2022.109345>. URL <https://www.sciencedirect.com/science/article/pii/S0360132322005789>.
- [187] M. Shafie-Khah and P. Siano. A stochastic home energy management system consid-

- ering satisfaction cost and response fatigue. *IEEE Transactions on Industrial Informatics*, 14(2):629–638, July 2018. ISSN 1941-0050. doi: 10.1109/TII.2017.2728803.
- [188] Amirreza Heidari, François Maréchal, and Dolaana Khovalyg. Reinforcement learning for proactive operation of residential energy systems by learning stochastic occupant behavior and fluctuating solar energy: Balancing comfort, hygiene and energy use. *Applied Energy*, 318:119206, 7 2022. ISSN 0306-2619. doi: 10.1016/J.APENERGY.2022.119206.
- [189] Hongyu Wu, Mohammad Shahidehpour, Zuyi Li, and Wei Tian. Chance-constrained day-ahead scheduling in stochastic power system operation. *IEEE Transactions on Power Systems*, 29(4):1583–1591, 2014. doi: 10.1109/TPWRS.2013.2296438.
- [190] Hongyu Wu, Mohammad Shahidehpour, Ahmed Alabdulwahab, and Abdullah Abusorrah. Thermal generation flexibility with ramping costs and hourly demand response in stochastic security-constrained scheduling of variable energy sources. *IEEE Transactions on Power Systems*, 30(6):2955–2964, 2015. doi: 10.1109/TPWRS.2014.2369473.
- [191] Ryan Lowe, YI WU, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/68a9750337a418a86fe06c1991a1d64c-Paper.pdf>.
- [192] Daniel T. Gillespie. Exact numerical simulation of the ornstein-uhlenbeck process and its integral. *Phys. Rev. E*, 54(2):2084–2091, Aug 1996. doi: 10.1103/PhysRevE.54.2084.
- [193] Leslie N. Smith. Cyclical learning rates for training neural networks, 2017.
- [194] Dimitrios Gyalistras, Markus Gwerder, et al. Use of weather and occupancy forecasts for optimal building climate control (opticontrol): Two years progress report. *Terres-*

trial Systems Ecology ETH Zurich, Switzerland and Building Technologies Division, Siemens Switzerland Ltd., Zug, Switzerland, 158, 2010.

- [195] Dimitrios Gyalistras, Markus Gwerder, Frauke Oldewurtel, Colin N Jones, Manfred Morari, Beat Lehmann, Katharina Wirth, and Vanessa Stauch. Analysis of energy savings potentials for integrated room automation. *Clima - RHEVA World Congress*, 2010. URL <https://infoscience.epfl.ch/record/169732>.
- [196] Da Yan, William O’Brien, Tianzhen Hong, Xiaohang Feng, H. Burak Gunay, Farhang Tahmasebi, and Ardeshir Mahdavi. Occupant behavior modeling for building performance simulation: Current state and future challenges. *Energy and Buildings*, 107: 264–278, 11 2015. ISSN 0378-7788. doi: 10.1016/J.ENBUILD.2015.08.032.
- [197] Mary Hill, Susan Stover, Vincent Amanor-Boadu, Peter Pfromm, Lawryn Edmonds, Robert Barron, Hongyu Wu, and Jirapat Phetheet. FEWtures: Building Rural Economies, Reducing Carbon, and Supplying Food for the Future. In *AGU Fall Meeting Abstracts*, volume 2021, pages SY21B–07, December 2021.
- [198] Peter S. Goodman. How a fertilizer shortage is spreading desperate hunger. URL <https://www.nytimes.com/2023/10/15/business/nigeria-fertilizer-shortage.html>. Accessed: October 25, 2023.
- [199] Omolola A. Ogbolumani and Nnamdi Nwulu. Integrated appliance scheduling and optimal sizing of an autonomous hybrid renewable energy system for agricultural food production. *Lecture Notes in Mechanical Engineering*, pages 651–660, 2020. ISSN 21954364. doi: 10.1007/978-981-15-5753-8_60/COVER.
- [200] Pawan Kumar Kushwaha, Priyanka Ray, and Chayan Bhattacharjee. Optimal sizing of a hybrid renewable energy system: A socio-techno-economic-environmental perspective. *Journal of Solar Energy Engineering, Transactions of the ASME*, 145, June 2023. ISSN 15288986.
- [201] Javier Carroquino, Rodolfo Dufo-López, and José L. Bernal-Agustín. Sizing of off-grid

- renewable energy systems for drip irrigation in mediterranean crops. *Renewable Energy*, 76:566–574, April 2015. ISSN 0960-1481. doi: 10.1016/J.RENENE.2014.11.069.
- [202] Stella® architect. [StellaArchitect,https://www.iseesystems.com/store/products/stella-architect.aspx](https://www.iseesystems.com/store/products/stella-architect.aspx).
- [203] Amir Modarresi and John Symons. Modeling Resilience for Sustainable Food, Energy, and Water Systems. In *AGU Fall Meeting Abstracts*, December 2021.
- [204] Peter H. Pfromm. Towards sustainable agriculture: Fossil-free ammonia. *Journal of Renewable and Sustainable Energy*, 9(3):034702, June 2017. ISSN 1941-7012. doi: 10.1063/1.4985090.
- [205] Kenneth V. Price. *Differential Evolution*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013. ISBN 978-3-642-30504-7. doi: 10.1007/978-3-642-30504-7_8.
- [206] National Renewable Energy Laboratory. 2023 annual technology baseline. *Golden, CO: National Renewable Energy Laboratory*, 2023. <https://atb.nrel.gov/electricity/2023/about>.
- [207] Xuebo Liu, Molly Ross, Hitesh Bindra, and Hongyu Wu. Optimal sizing of battery energy storage systems for small modular reactor based microgrids. In *2021 IEEE Kansas Power and Energy Conference (KPEC)*, 2021. doi: 10.1109/KPEC51835.2021.9446203.
- [208] Catherine Obiribea Ofori-Bah and Vincent Amanor-Boadu. Directing the wind: Techno-economic feasibility of green ammonia for farmers and community economic viability. *Frontiers in Environmental Science*, 10, 2023. ISSN 2296-665X. doi: 10.3389/fenvs.2022.1070212.
- [209] Patience Bosompemaa, Sam C. Zipper, Andrea Brookfield, and Mary C. Hill. Quantifying Water Sustainability in the Central Arkansas River Basin: Evaluating Existing Tools at the Regional Scale. In *Frontiers in Hydrology 2022*, pages 115–05, June 2022.
- [210] Patience Bosompemaa, Andrea Brookfield, Sam Zipper, and Mary C. Hill. Using

- national hydrologic models to obtain regional climate change impacts on streamflow basins with unrepresented processes. *Environmental Modelling & Software*, 183:106234, 2025. ISSN 1364-8152. doi: <https://doi.org/10.1016/j.envsoft.2024.106234>. URL <https://www.sciencedirect.com/science/article/pii/S1364815224002950>.
- [211] Peter H. Pfromm and Wrya Aframehr. Green ammonia from air, water, and renewable electricity: Energy costs using natural gas reforming, solid oxide electrolysis, liquid water electrolysis, chemical looping, or a haber–bosch loop. *Journal of Renewable and Sustainable Energy*, 14(5):054701, 09 2022. ISSN 1941-7012. doi: 10.1063/5.0101709. URL <https://doi.org/10.1063/5.0101709>.
- [212] Wrya Mohammadi Aframehr, Chaoran Huang, and Peter H. Pfromm. Chemical looping of manganese to synthesize ammonia at atmospheric pressure: Sodium as promoter. *Chemical Engineering & Technology*, 43(10):2126–2133, 2020. doi: <https://doi.org/10.1002/ceat.202000154>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/ceat.202000154>.
- [213] Volodymyr Mnih. Asynchronous methods for deep reinforcement learning. *arXiv preprint arXiv:1602.01783*, 2016.
- [214] Rutherford Titan. Liquid nitrogen generators - liquid nitrogen price usa. <https://www.rutherfordtitan.com/liquid-nitrogen-generators/liquid-nitrogen-price-usa/>. Accessed: 2024-10-07.
- [215] Derbyks. Water rates. <https://derbyks.com/264/Water-Rates>. Accessed: 2024-10-07.
- [216] ScienceDirect. Irrigation water cost calculation. <https://www.sciencedirect.com/science/article/pii/S1876610218305411>. Accessed: 2024-10-07.
- [217] Expatistan. Price of potatoes in wichita. <https://www.expatisitan.com/price/potatoes/wichita>. Accessed: 2024-10-07.

- [218] ElectricChoice. Electricity prices by state. <https://www.electricchoice.com/electricity-prices-by-state/>. Accessed: 2024-10-07.

Appendix A

Acronyms

Acronyms

A2C	Advantage Actor-Critic
ADP	approximate dynamic programming
AGC	automatic generation control
ANN	artificial neural network
ASHRAE	American Society of Heating, Refrigerating and Air-Conditioning Engineers
BD	Benders' decomposition
BESS	battery energy storage system
CN	carbon-nitrogen
CoA	complexity of adaption
COP	coefficient of performance
CRAIRC	critical air content
DNN	deep neural network

DPV	distributed photovoltaic
DQN	Deep Q-Network
DRL	deep reinforcement learning
ED	economic dispatch
EMS	energy management strategy
EV	electric vehicle
FEW	food-energy-water
FIFF	full irrigation, full fertilizer
FNC-UC	frequency nadir-constrained unit commitment
GB	gradient boosting
HR	headroom
HVAC	heating, ventilation, and air conditioning
IBR	inverter-based resource
IQR	interquartile range
KDNN	K-D tree nearest neighbor
LR	linear regression
LSTM	long short-term memory
MAPE	mean absolute percentage error
MDP	Markov decision process
MDT	minimum downtime
MIDAS	multi-timescale integrated dynamic and scheduling
MILP	mixed integer linear programming

MLWB	multi-layer water balance
MPC	model predictive control
MPPT	maximum power point tracking
MUT	minimum uptime
NH₃	ammonia
NINF	no irrigation, no fertilizer
NREL	National Renewable Energy Laboratory
PCSE	Python Crop Simulation Environment
PFR	primary frequency reserve
PMV	predicted mean vote
PPD	predicted percentage of dissatisfied
PSSE	power system simulator for engineering
PV	photovoltaic
PWL	piecewise linear
RBC	rule-based control
RHOD	soil bulk density
RL	reinforcement learning
RNN	recurrent neural network
SNOMIN	soil nitrogen module for organic and mineral nitrogen
SOC	state of charge
SOMI	soil organic matter input
SSBI	Sobol sampling backward induction

SVR	support vector regression
ToU	time-of-use
UC	unit commitment
UFLS	under-frequency load shedding
VRE	variable renewable energy
WFH	work from home
WH	water heater
WOFOST	World Food Studies
WSO	weight of storage organ

Appendix B

Reuse permissions from publishers



Sign in/Register



Enhancing HVAC energy management through multi-zone occupant-centric approach: A multi-agent deep reinforcement learning solution

Author: Xuebo Liu, Yingying Wu, Hongyu Wu
Publication: Energy and Buildings
Publisher: Elsevier
Date: 15 January 2024

© 2023 Elsevier B.V. All rights reserved.

Journal Author Rights

Please note that, as the author of this Elsevier article, you retain the right to include it in a thesis or dissertation, provided it is not published commercially. Permission is not required, but please ensure that you reference the journal as the original source. For more information on this and on your other retained rights, please visit: <https://www.elsevier.com/about/our-business/policies/copyright#Author-rights>

BACK

CLOSE WINDOW



Sign in/Register



Budget-Constrained Sizing of Renewable and Energy Storage Systems for Farm-Scale Ammonia Production Within the Food-Energy-Water Nexus

Conference Proceedings: 2024 IEEE Green Technologies Conference (GreenTech)

Author: Xuebo Liu

Publisher: IEEE

Date: 03 April 2024

Copyright © 2024, IEEE

Thesis / Dissertation Reuse

The IEEE does not require individuals working on a thesis to obtain a formal reuse license, however, you may print out this statement to be used as a permission grant:

Requirements to be followed when using any portion (e.g., figure, graph, table, or textual material) of an IEEE copyrighted paper in a thesis:

- 1) In the case of textual material (e.g., using short quotes or referring to the work within these papers) users must give full credit to the original source (author, paper, publication) followed by the IEEE copyright line © 2011 IEEE.
- 2) In the case of illustrations or tabular material, we require that the copyright line © [Year of original publication] IEEE appear prominently with each reprinted figure and/or table.
- 3) If a substantial portion of the original paper is to be used, and if you are not the senior author, also obtain the senior author's approval.

Requirements to be followed when using an entire IEEE copyrighted paper in a thesis:

- 1) The following IEEE copyright/ credit notice should be placed prominently in the references: © [year of original publication] IEEE. Reprinted, with permission, from [author names, paper title, IEEE publication title, and month/year of publication]
- 2) Only the accepted version of an IEEE copyrighted paper can be used when posting the paper or your thesis online.
- 3) In placing the thesis on the author's university website, please display the following message in a prominent place on the website: In reference to IEEE copyrighted material which is used with permission in this thesis, the IEEE does not endorse any of [university/educational entity's name goes here]'s products or services. Internal or personal use of this material is permitted. If interested in reprinting/republishing IEEE copyrighted material for advertising or promotional purposes or for creating new collective works for resale or redistribution, please go to http://www.ieee.org/publications_standards/publications/rights/rights_link.html to learn how to obtain a License from RightsLink.

If applicable, University Microfilms and/or ProQuest Library, or the Archives of Canada may supply single copies of the dissertation.

BACK

CLOSE WINDOW



Sign in/Register



Hourly occupant clothing decisions in residential HVAC energy management

Author: Xuebo Liu, Yingying Wu, Haifeng Zhang, Hongyu Wu
Publication: Journal of Building Engineering
Publisher: Elsevier
Date: August 2021

© 2021 Elsevier Ltd. All rights reserved.

Journal Author Rights

Please note that, as the author of this Elsevier article, you retain the right to include it in a thesis or dissertation, provided it is not published commercially. Permission is not required, but please ensure that you reference the journal as the original source. For more information on this and on your other retained rights, please visit: <https://www.elsevier.com/about/our-business/policies/copyright#Author-rights>

BACK

CLOSE WINDOW