NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

MICROGRID OPTIMIZATION THROUGH INTEGRATION OF A REAL-TIME CONTROLLER

by

Drew E. Chapman

June 2023

Thesis Advisor:	Emily M. Craparo
Second Reader:	Samuel E. Buttrey

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE	*Form Approved OMB No. 0704-0188*

1. AGENCY USE ONLY *(Leave blank)*	**2. REPORT DATE** June 2023	**3. REPORT TYPE AND DATES COVERED** Master's thesis

4. TITLE AND SUBTITLE
MICROGRID OPTIMIZATION THROUGH INTEGRATION OF A REAL-TIME CONTROLLER

5. FUNDING NUMBERS

6. AUTHOR(S) Drew E. Chapman

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Naval Postgraduate School
Monterey, CA 93943-5000

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
N/A

10. SPONSORING / MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

12a. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release. Distribution is unlimited.

12b. DISTRIBUTION CODE
A

13. ABSTRACT (maximum 200 words)

Microgrid energy generation that includes non-renewable energy sources (e.g., fossil fuels) relies on long-term historical power demand data to forecast future power requirements and fuel use. Previous models seeking to optimize microgrid performance under time-varying loads introduced penalty functions to smooth the energy demand between generators and energy storage systems (ESS). However, these models require complete demand profiles prior to generating feasible solutions, making them useful for analysis but impractical for operational use. To develop a near-real-time microgrid controller, we use a sample of omniscient model outputs, a history of demand data, and time-varying demand profile variables to train several machine learning models to estimate future demand, which we then use to prescribe energy generation output based on current conditions. The models create short-term forecasts for energy generation requirements regarding future demand profiles. The models also penalize changes to the generators' energy generation loads, which reduces fuel consumption, supports energy supply resiliency, and decreases microgrid maintenance costs. We compare the results obtained using our predictive models to a previous omniscient model that requires future demand profiles; our model is shown to have an average fuel consumption rate only 0.0016 gal/hr greater than the previous omniscient model over a 75-day rolling test horizon, less than 0.09% from omniscient performance.

14. SUBJECT TERMS
optimization, machine learning, energy, energy resilience, energy infrastructure, fuel consumption, energy generation, energy storage system, ESS

15. NUMBER OF PAGES
77

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	**18. SECURITY CLASSIFICATION OF THIS PAGE**	**19. SECURITY CLASSIFICATION OF ABSTRACT**	**20. LIMITATION OF ABSTRACT**
Unclassified	Unclassified	Unclassified	UU

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. 239-18

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release. Distribution is unlimited.

MICROGRID OPTIMIZATION THROUGH INTEGRATION OF A REAL-TIME CONTROLLER

Drew E. Chapman
Major, United States Army
BS, United States Military Academy, 2011
MS, Missouri University of Science and Technology, 2016

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN OPERATIONS RESEARCH

from the

NAVAL POSTGRADUATE SCHOOL
June 2023

Approved by: Emily M. Craparo
Advisor

Samuel E. Buttrey
Second Reader

W. Matthew Carlyle
Chair, Department of Operations Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

Microgrid energy generation that includes non-renewable energy sources (e.g., fossil fuels) relies on long-term historical power demand data to forecast future power requirements and fuel use. Previous models seeking to optimize microgrid performance under time-varying loads introduced penalty functions to smooth the energy demand between generators and energy storage systems (ESS). However, these models require complete demand profiles prior to generating feasible solutions, making them useful for analysis but impractical for operational use. To develop a near-real-time microgrid controller, we use a sample of omniscient model outputs, a history of demand data, and time-varying demand profile variables to train several machine learning models to estimate future demand, which we then use to prescribe energy generation output based on current conditions. The models create short-term forecasts for energy generation requirements regarding future demand profiles. The models also penalize changes to the generators' energy generation loads, which reduces fuel consumption, supports energy supply resiliency, and decreases microgrid maintenance costs. We compare the results obtained using our predictive models to a previous omniscient model that requires future demand profiles; our model is shown to have an average fuel consumption rate only 0.0016 gal/hr greater than the previous omniscient model over a 75-day rolling test horizon, less than 0.09% from omniscient performance.

THIS PAGE INTENTIONALLY LEFT BLANK

Table of Contents

List of Figures

THIS PAGE INTENTIONALLY LEFT BLANK

List of Tables

THIS PAGE INTENTIONALLY LEFT BLANK

List of Acronyms and Abbreviations

A	amp
AC	alternating current
Ah	amp-hours
AMMPS	advanced medium mobile power sources
BESS	battery energy storage system
BOC	base operations center
C4ISR	command, control, communications, computers, intelligence, surveillance, and reconnaissance
DC	direct current
DER	distributed energy resource
DOD	Department of Defense
EABA	enhanced adaptive bat algorithm
ESS	energy storage system
FOB	forward operating base
GRU	gated recurrent unit
Hz	hertz
JP-8	jet propellant 8
kW	kilowatt
kWh	kilowatt-hours
Li-ion	lithium-ion

LRP layer-wise relevance propagation

LSTM long short-term memory

MAE mean absolute error

MILP mixed-integer linear program

MSE mean-squared error

PV photo voltaic

RNN recurrent neural network

RTE round trip efficiency

SOC state of charge

V volts

VA volt amps

VAR volt-amps reactive

Executive Summary

Microgrids are a critical capability in the Department of Defense's ability to project power, allowing forces to operate increasingly complex and connected systems while also supporting requirements for sustaining the warfighter in austere and remote locations. Microgrids are also growing in importance at permanent U.S. bases, supplementing basic energy production and providing a critical backup energy resource in emergency situations. Following her appointment to Secretary of the Army, the Honorable Christine Wormuth (2022) outlined her six priorities for the Army:

1. Put the Army on a sustainable strategic path amidst uncertainty;
2. Ensure the Army becomes more data-centric and can conduct operations in contested environments;
3. Continue efforts to be resilient in the face of climate change;
4. Build positive command climates at scale across all Army formations;
5. Reduce harmful behaviors in our Army;
6. Strategically adapt the way we recruit and retain talent into the Army.

Coupled with the Army's intent to install microgrids at all CONUS Army installations by 2035, increasing microgrid efficiency through predictive modeling has the potential to reduce Army operating costs, leverage extant but unused data, and increase energy resiliency in all domains.

Previous efforts to maximize microgrid efficiency by minimizing fuel consumption focused on developing a model that would receive a previously generated energy demand profile and generate an optimal operating policy. This previous work incorporated a battery energy storage system, and introduced a fuel consumption penalty to the minimized fuel consumption objective in order to incentivize the model to use the battery as much as possible. While the previous model succeeded in filling 100% of presented demand, increasing generator fuel efficiency, and reducing generator fluctuations, a major drawback was that the model required a future demand profile as input. This drawback made the model unsuitable for operational application.

This thesis builds on previous efforts by developing three models for predicting future demand, applying the models under different conditions, and determining which model best approaches the performance of an optimized operating profile with perfect information about future demand. The first model we create is a naïve model that predicts that the energy demand for today will be the same as the energy demand yesterday. The other two models we develop are long-short term memory (LSTM) recurrent neural networks (RNN), trained on over two-and-a-half years of historic demand data from a base operations center (BOC) at a forward operating base (FOB) in the Middle East. One of the neural network models is additionally influenced by artificially generated periodic signals associated with days, weeks, and years.

Each predictive model is used to predict a full day's worth of energy demand. The predicted demand is then given to a mixed-integer linear program (MILP) to develop an optimized generation and battery use policy over the predicted demand interval. New, rolling predictions are then generated every six, eight, or twelve hours to re-optimize the microgrid's performance according to new predicted energy demands.

Of the predictive models developed, we find that the naïve model has the least similarity to true demand due to sudden increases or decreases in the previous day's energy demand. However, because the general trends of the demand do not exhibit major changes day-over-day, the naïve model is competitive with the more sophisticated LSTM RNN models. The LSTM RNN model that was trained without artificial periodic signals generates a smoother demand curve, but consistently over-predicts future demand. Because of this, when this model is subjected to the MILP it prescribes fewer changes in generator production, but the generator runs at a higher production level and thus uses more fuel. The final model, an LSTM RNN with artificially generated periodic signals, generates the smoothest demand predictions and best fit to true demand. As such, when the predicted demand is used to develop an optimized generation policy, this model has few shifts in generator operating load and makes efficient use of the battery to minimize fuel consumption.

We evaluate and compare each model's performance over 75 consecutive days of rolling predictions. Assuming that optimal performance would come from an omniscient model that would be able to perfectly predict future demand and develop a generation policy accordingly, the omniscient model scenario yields an average fuel consumption rate of

1.7767 gal/hr. The best-performing predictive model is the LSTM RNN with artificially-generated periodic signals repredicting demand every eight hours, with an average fuel consumption rate of 1.7783 gal/hr, just 0.09% worse than optimal performance. The second and third best-performing models are the naïve predictive models at six or twelve-hour reprediction intervals, both with an average fuel consumption rate of 1.7796 gal/hr, about 0.16% worse than optimal performance. The LSTM RNN model without artificial periodic signals is the worst-performing model, regardless of reprediction interval.

The prediction and optimization models developed in this thesis demonstrate the value of leveraging extant data for the Army and within the DOD. By using historical demand records for microgrids, we have presented a method to predict future energy demands for microgrids and operate those microgrids at near optimally prescribed policies. By reducing fuel consumption, we directly decrease overall operating costs for microgrids while also reducing fuel storage and transportation requirements. We also increase the operational lifetime of the microgrid's generator by reducing the number of load shifts the generator undergoes, operating the generator in as close to a steady-state as possible.

References

Wormuth C (2022) Message from the secretary of the Army to the force. U.S. Army. Accessed February 11, 2022, https://www.army.mil/article/253814/message_from_the_secretary_of_the_army_to_the_force.

THIS PAGE INTENTIONALLY LEFT BLANK

Acknowledgments

Everything I've ever accomplished has been done so with the guidance and support of mentors, peers, friends, and family, and this thesis is no different. While I have poured countless hours into this singular work, there have been countless more milestones along the way to completion. The individuals that have shepherded me along the way deserve whatever recognition I can give and much more.

First, I must thank Dr. Emily Craparo her wisdom, patience, and presence. On one of my first meetings with Dr. Craparo, I remember asking her why—given her truly outstanding competency—she elected to spend her time teaching a rock-eater like me the fundamentals of optimization. She expressed her appreciation for the material, teaching, and tackling real problems with practical solutions, evident in every interaction thereafter. Dr. Craparo, the students you mentor are immeasurably better for the experience, and me doubly so.

I must also recognize the invaluable insight provided by Dr. Sam Buttrey. Insofar as one can be creative in statistics, Dr. Buttrey was there to offer an innovative solution and clear the path each time I ran into a statistical roadblock. His willingness to dive down rabbit holes of suppositions and hypothetical situations was essential to my learning during this thesis.

As a professional Army officer and newly minted Operations Researcher/Systems Analyst, I'd like to thank the outstanding Army leadership at NPS. Being in a functional area is a new experience for all Army students in the OR curriculum. We have been fortunate enough to have the steady guidance of COL Thoendel, LTC Schuchard, and LTC Nesbitt guiding our careers. Their efforts in introducing us to the branch, aligning our interests and strengths with our studies, and ensuring that the students make the most of our time are unlike anything I've ever experienced.

For my peers, the adage "cooperate to graduate" has never rung more true than in this program. When I've been nearly moved to tears with frustration, when life events have overtaken academic requirements, or when we want to take some time to celebrate our small victories, we have been there to support each other. During our time here, I feel like I've been more frequently in the camp of asking for help than offering it, and so I have little

doubt that I would not have found any measure of success without the support of my cohort.

Some small recognition is owed to my cat, Scout, who might as well be a co-author on this thesis for as much time as he has spent with me while writing and coding. From his favorite sleeping spot on the back of my desk, he has been a victim of my debugging efforts, an audience to my presentation practice sessions, an editing ear to my proofreading, and a very fine distraction when I don't feel like doing any work at all.

Finally, and most importantly, thank you to Danielle Chapman. For the last ten years, Dani has gracefully supported me through a tumultuous Army career that has taken us all over the world, made our house a home, and allowed me to focus on work that is exciting, interesting, and often time-consuming. And despite all this, she has managed to do amazing things everywhere we have lived, have adventures that I regularly reflect on and envy, and still be a wonderful mother to our beautiful daughter, Adeline. My days are always brighter for having you in them, I couldn't do what I do without you, and I love you.

CHAPTER 1:
Introduction

The ability to efficiently supply power to where it is needed when it is needed is a critical warfighting capability. Besides the power production systems themselves, logistic support to maintaining consistent power requires the use, transportation, and storage of fuel, repair parts, and skilled technicians to ensure a power production system functions as required. For example, the fiscal year 2022 DOD budget estimate for fuel purchases alone was over $9.2 billion, supporting long-term base and contingency operations (Assistant Secretary of Defense for Sustainment 2021). As the U.S. responds to crisis and projects forces further afield, it must also find a way to efficiently supply those forces with the power required to complete their assigned missions.

1.1 Microgrid Employment

Microgrids enjoy wide use by businesses, industrial fabrication, critical infrastructure, and personal settings. They offer a reliable and resilient energy source, especially in areas prone to power outages or with limited access to the primary power grid. Microgrids also enhance grid stability by incorporating advanced energy management systems, fulfilling energy demands during failures of the primary power grid. Because of the versatility of microgrids, they are used in areas with highly developed infrastructure and in austere or remote locations to support individual homes or compounds, ensuring an uninterrupted electricity supply for lighting, heating, and essential equipment. Microgrids are especially valuable in emergencies such as natural disasters, enabling the quick restoration of power to support aid in disaster management efforts (Farhangi and Joos 2019).

A key driver of microgrid development is application in business and industrial settings. Microgrids reduce energy dependence on the primary grid, mitigating the risk of power disruptions and associated financial losses from manufacturing interruptions. Businesses may also reduce their energy consumption costs by integrating renewable distributed energy resources (DERs) into microgrids, such as solar panels or wind turbines while taking advantage of tax incentives to reduce fossil fuel reliance and carbon emissions (Shea 2022).

The scalable nature of microgrids allows for their use in personal situations, such as off-grid vacation cabins, in recreational vehicles, on boats, or personally carried, providing electricity in isolated areas without the need for extensive infrastructure development. By harnessing renewable DERs, personal microgrids enable individuals and organizations that favor environmentally conscious power generation to consume clean energy on a small scale.

Regardless of employment methodology, several factors influence the decision to use microgrids, including the need for energy reliability, cost savings through optimized energy consumption, environmental considerations, and geographical region. Implementation of a microgrid depends on a comprehensive assessment of energy requirements, available resources, and the specific objectives and constraints of the given setting.

1.1.1 Military Microgrids and Power Generation

The U.S. military has several methods of supplying power to forward operating bases (FOBs), each of which bears the requirement to be generally self-sufficient besides logistic resupply. While operating in austere environments, these FOBs must be able to run a suite of command, control, communications, computers, intelligence, surveillance, and reconnaissance (C4ISR) capabilities and support service members for hygiene, shelter, and food preparation. Operating these small complexes requires significant power, often provided in one of two ways. First, a previously acquired base support kit comprising several containers, including tent facilities for billeting, administration, latrines, a fuel distribution system, and up to six generators and six power distribution boxes, may fill the immediate requirements of a FOB. In such a support kit, each generator can support the power requirements of several tent facilities, each acting as an "islanded" microgrid. The components of these systems, including the power generation components, are considered non-permanent facilities with limited lifetimes. More extensive power generation facilities will generally fulfill long-term, semi-permanent power requirements and be installed by Department of Defense (DOD) military units, such as the U.S. Army's Prime Power battalion, or by contracted support.

1.2 Cost Estimation

Estimating the costs of running a microgrid involves considering various factors such as generator size and output, running time, installation, transportation, and potential equipment rental costs. The generator size and output are the most critical determinants of the overall cost. The power requirements of the connected load will determine the microgrid's energy demand and, likewise, determine the microgrid's generator size—the larger the generator required, the higher the upfront cost to install a microgrid (Farhangi and Joos 2019).

The running time of the microgrid, or the duration for which it needs to operate independently, also affects the costs. Depending on the nature of the DER comprising the microgrid, longer running times may necessitate larger fuel storage capacities or higher renewable energy generation capacities, which can increase the cost. Additional recurring expenses include maintenance and operational costs, such as fuel, regular servicing, and necessary repairs and repair parts.

The installation costs of a microgrid include setting up the necessary infrastructure, such as the generator, energy storage systems (ESSs), inverters, control systems, and distribution equipment. This cost includes labor, materials, permits, and any required upgrades to the existing electrical infrastructure. Delivering microgrid components to remote or inaccessible locations may incur additional transportation costs.

Organizations or individuals may opt for renting microgrid systems instead of purchasing them outright. Rental costs vary depending on the system's size, rental duration, and any additional services provided. For example, renting a microgrid can be a more cost-effective option for short-term projects or when there is uncertainty about long-term energy requirements.

Estimating the costs of running a microgrid requires a comprehensive analysis of specific energy needs, desired level of independence, and available resources. Therefore, it is essential to consider the upfront investment costs and ongoing operational expenses to determine the most economical solution for implementing and maintaining a microgrid.

1.2.1 Short- and Long-Term Costs

While microgrids have the advantage of portability of use, their employment incurs several short and long-term costs. The type of DER being used to power the microgrid and incorporating an ESS determine a large portion of long-term costs. In the case of the DOD, short-term costs would include equipment procurement, installation, and initial testing requirements. However, fixed short-term prices can make estimating any cost-efficiency gains challenging.

Compared to short-term costs, the long-term costs of microgrid operations can grow tremendously over the lifetime of the equipment. Traditional DER, such as the fuel generators used in remote environments, require a steady fuel supply, ongoing maintenance from qualified mechanics, and a bench of replacement parts to repair broken components. Introducing non-mechanical DER, such as solar panels, or including an ESS in the microgrid architecture may offset some fuel and maintenance costs. Regardless of cost-saving measures, focusing on long-term costs offers the best venue to find efficiencies by reducing fuel usage and increasing the generator durability.

1.3 Power Generation

Microgrids' power generation methods encompass various options, including conventional generators, renewable energy sources, and energy storage systems. Conventional generators are vital in microgrid power generation and are often used to provide backup or supplementary power. Diesel generators, for instance, are commonly employed in microgrids for their reliability and relatively low upfront costs. However, they rely on fossil fuels and produce carbon emissions, which an emission control system can mitigate.

Meanwhile, renewable energy sources offer clean and sustainable power solutions. Microgrids commonly use solar photo voltaic (PV) panels, wind turbines, and hydroelectric generators to harness renewable energy. Solar PV panels convert sunlight into electricity, while wind turbines capture the wind's kinetic energy. Hydroelectric generators utilize the flow of water to generate power. These renewable sources reduce reliance on fossil fuels and contribute to environmental sustainability, but may be reliant on environmental conditions to produce energy consistently.

Energy storage systems are essential components of microgrids, enabling the storage of excess generated power and utilization of stored power. Battery storage is a widely used technology, allowing excess energy to be stored and used during periods of low or no generation. Other energy storage options include pumped hydro storage, compressed air energy storage, and flywheel systems. These energy storage technologies enhance microgrids' stability and reliability by balancing supply and demand fluctuations.

The selection and combination of power generation and energy storage methods depend on the microgrid implementation's specific requirements, available resources, and objectives.

1.4 Thesis Organization

Microgrid development and optimization remains an active field of research. As such, Chapter 2 of this thesis will review the current status of microgrid optimization efforts related to the employment of energy storage systems, energy scheduling, the application of neural networks to microgrid optimization, and the importance of resiliency for energy systems.

Chapter 3 will describe our theoretical microgrid, techniques for generating energy demand predictions, methods of optimization, and the metrics we will apply to determine model performance. This chapter will also highlight points at which we depart from the previous work discussed in Chapter 2.

Chapter 4 will examine a use case of our model, apply the prediction methodologies, and develop insights from comparing previous work and the predictive methods described herein.

Finally, Chapter 5 will summarize this study's conclusions and final thoughts and recommend avenues for future work.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 2:
Background and Literature Review

2.1 General State of Work

Since the inception of microgrids, areas of development have focused on improvement and optimization. With significant support from computational tools, the last decade has seen rapid growth in optimization methods across a wide range of actual and simulated environments. As early as 2005, Hernandez-Arambura et al. presented a method to optimize a microgrid by minimizing fuel use in a system of two traditional fuel generators, and expanded their method to include fuel generators and renewable energy sources. Their work supported the concept of having a method of communication between power sources and a method of feedback between demand nodes and power sources. Craparo et al. (2017) later created a model of a hybrid-generation microgrid to optimize power generation while incorporating short-term ensemble forecasting. In addition, they extended their model to a practical application of remote military installations and made key insights about planning horizon impacts for optimized performance.

Garcia (2017) later developed and presented two mixed-integer linear programs (MILPs) designed to provide energy to austere military locations. In the first case, Garcia minimized the cost of energy production. In contrast, the second MILP minimized fuel consumption, concluding that an optimal mix of energy sources, generator sizes, and substantial energy storage will minimize energy production costs. Dulău and Bică (2018) built a similar model for a hybrid microgrid for a civilian consumer demand profile, seeking to minimize generation cost under a steady load. In their study, Dulău and Bică arrive at a total installed capacity 160% greater than the steady load demand, accounting for changes in generation ability, such as in the case of solar or wind power, while incorporating load transport limitations. Modeling conducted by Beaton (2021) supported a distributed, hybrid microgrid with robust energy storage to maximize resiliency while proposing a formal definition for resilience in a military context. Ogunmodede et al. (2021) further examined hybrid microgrid design and energy dispatch under various scenarios and with additional constraints, though

with tendencies towards more recent developments in technological efficiencies. Their model generates a fuel and cost-optimal hybrid microgrid system to meet a site's energy needs (i.e., a small hospital or business complex). However, the demand load profile is a predetermined input.

Generally, microgrid loads have been managed simplistically, with generators automatically increasing or decreasing power production in response to changing demand requirements. In contrast, Craparo and Sprague (2019) explored the optimization of hybrid microgrids with demands whose timing can be dictated by the model. While they consider heating and cooling of structures, such flexibility in load scheduling also arises in other applications such as electric vehicle charging. The authors effectively modeled an optimal blend of generation and load scheduling and demonstrated that fuel savings can be achieved by intelligently aligning supply and demand.

Lee (2021) created an omniscient model of energy generation under time-varying loads, with the goal of reducing unnecessary fluctuations in generator output. Using a penalty term on time-varying loads, the model shows the critical role of an ESS in managing fuel consumption and maximizing efficiency. However, the model falls short of being useful for implementation because it demands a complete future demand profile. Lee's work acts as this thesis's primary point of departure as we create a near real-time controller, incorporating past and present energy demands to predict future loads and prescribing microgrid energy generation policy. In addition, generating new demand profiles unseen by the model is supported by Mallery (2021), who proposed a methodology to determine energy requirements for different mission scenarios, the subsequent loads placed on microgrid infrastructure, and the impact on energy resilience.

Most recently, microgrid optimization has come to the forefront of defense objectives in general. Following her appointment, Secretary of the Army Christine Wormuth (2022) outlined her six priorities for the Army's future:

1. Put the Army on a sustainable strategic path amidst uncertainty;
2. Ensure the Army becomes more data-centric and can conduct operations in contested environments;
3. Continue efforts to be resilient in the face of climate change;

4. Build positive command climates at scale across all Army formations;
5. Reduce harmful behaviors in our Army;
6. Strategically adapt the way we recruit and retain talent into the Army.

Developing more efficient microgrids supports the first three priorities. Real and tangible efforts continue across the Army, with plans underway to install microgrids at all installations across the continental U.S. by 2035 (Wood 2022). While traditional fuel microgrids will still have a place in these future installations—both for emergency purposes and reliability—there is an additional effort in hybridizing microgrids to increase energy resiliency and support climate change initiatives (Surash and Hughes 2022). As such, funding and research supporting microgrid optimization will continue for the foreseeable future.

2.2 Microgrids and Energy Storage Systems

Microgrids are small-scale power generation and distribution systems that may operate as an islanded power grid (independently) or supplement the capacity of other grids (Parhizi et al. 2015). A microgrid is either a single energy generation resource, such as a traditional fuel generator, or a combination of renewable energy sources (e.g., solar and wind) and ESS (batteries, pumped water reservoirs). Microgrids may supplement the electricity demands of industrial facilities, communities, and businesses, provide a resilient alternative electricity source in emergencies, or be the primary power system in remote locations. Additionally, because of their relative portability, microgrids can help reduce transmission and distribution losses in larger grids. Regardless of the final composition and use case, a microgrid is generally small-scale and comprises DERs, performance controllers, and demand load (Mohammed et al. 2019).

Microgrids are classified either through system topology or market segment. System topology classifications include alternating current (AC), direct current (DC), and Hybrid AC/DC microgrid variants and are thus hardware-based. Comparatively, market segment classification is based on the use case and includes utility, institutional, commercial-industrial, transportation, and remote-area (Ahmed et al. 2015). DOD microgrid architecture spans the system topology classifications and many market segment classifications. The DOD employs microgrids on permanent installations providing energy resiliency to critical in-

frastructure and temporary or remote military bases as the primary source of electricity (Vergun 2021).

Research in energy storage is highly active due to the need to store excess energy during low-demand times and release it during high-demand periods. The five general types of energy storage systems are thermal, electrical, chemical, mechanical, and hybrid (Faisal et al. 2018). The choice of a specific ESS depends on the microgrid's location and purpose. Electrochemical battery energy storage systems (BESSs) are popular for their portability, backup power, and ease of maintenance. While lead-acid storage was once the norm, lithium-ion (Li-ion) storage systems have become increasingly popular since the 1990s for their high energy storage efficiency (over 90%), high energy density, quick response time for charging and discharging, and low self-discharge rate (around 5%) (Faisal et al. 2018). However, Li-ion batteries are not suited for long-term energy storage and are often made abroad, so improving their performance and domestic production is a priority for the U.S. Department of Energy specifically and the U.S. in general (Federal Consortium for Advanced Batteries [FCAB] 2021).

2.3 Recurrent Neural Networks

Within the discipline of neural networks, recurrent neural networks (RNNs) are a branch of neural networks widely used in natural language processing and time series tasks that involve processing sequential data (Goodfellow et al. 2016). In recent years, there have been several advancements in the field of RNN research. Some of the latest research on RNNs includes pre-training, attention mechanisms, gated recurrent unit (GRU) with Recurrent Dropout (GRU-RD), reinforcement learning, attention mechanisms, and long short-term memory (LSTM).

Specific to the topic of this thesis, we are interested in leveraging a LSTM RNN 's forecasting capability to predict energy demand for a microgrid based on previous demand profiles. Using two real-world data sets, Wen et al. (2019) used a LSTM-based RNN to promote peak load shifting and reduce daily costs by 8.97% compared to standard microgrid controllers. They also showed that RNNs can outperform multi-layer perception networks, support vector machines, decision trees, and evolution algorithm models. Research in RNNs continues to

be active and promising, with new architectures, techniques, and applications regularly proposed and evaluated.

2.3.1 Recurrent Neural Network Forecasting

Forecasting with machine learning is a particularly active area of research. While previously referenced LSTM and GRU networks are effective in tasks such as financial forecasting (Dixon and London 2021), traffic forecasting (Baskar et al. 2022), and energy demand forecasting, some specific RNN methods show promise in a broader range of applications.

Attention-based models have been incorporated into forecasting models to allow the model to focus on specific parts of the input sequence, which improves performance in tasks such as modeling product sales (Ekambaram et al. 2020) or stock price predictions. There has been considerable effort towards improving the performance of forecasting models by pre-training them on related tasks and then fine-tuning them on a specific task, a process known as transfer learning. Like predicting the power demand for microgrids, RNN transfer learning has seen some success in supporting energy consumption estimation for electric vehicles (Hua et al. 2022).

Because of the typical deep-learning format of modern machine-learning models, some research has focused on developing explainable models which may provide insights into the underlying patterns and relationships in the data. For example, on page 18670 of their 2021 paper, Jung et al. seek to counter the "black-box problem that complicates the interpretation of the models" through explainable AI. They approach this task by comparing layer-wise relevance propagation (LRP)—a method of visualizing the amount of influence each layer has on the model—and a new method with selective LRP to provide a clearer picture and greater explanatory power (Jung et al. 2021).

2.4 Microgrid Energy Scheduling

Developing advanced control and optimization algorithms for microgrids is a dynamic area of research. Optimal microgrid energy scheduling aims to optimize the energy flow within a microgrid to minimize the cost of energy, reduce greenhouse gas emissions, and improve the

reliability of the power supply. Optimal scheduling often involves reducing the amount of fuel used to generate energy or ensuring that there is sufficient generative redundancy such that there is no supply interruption in case of failure. Optimization techniques typically use models that consider various factors, such as weather conditions, energy demand, energy prices, and the capacity and efficiency of the generative resources.

Research into algorithms that control microgrid operation by balancing power generation and demand or optimizing ESS usage is particularly active. Craparo and Sprague (2019) optimized hybrid microgrids for changing demand profiles, creating an optimal blend of generation under shifting early or late load conditions, resulting in performance comparable to consistent load conditions. Lee (2021) then built an omniscient model for energy generation under varying loads, taking a step towards a real-time controller. By adding a penalty term for time-varying loads, the model highlights the essential role of an ESS in managing fuel consumption and maximizing efficiency. However, the model's usefulness for implementation is limited because it requires a complete future demand profile before providing a power generation solution. Mallery (2021) offers a solution by proposing a method to determine energy requirements for various mission scenarios, supporting the generation of unseen demand profiles.

Beyond MILP optimization, in 2021 a group of researchers proposed a heuristic enhanced adaptive bat algorithm (EABA) for an optimal energy scheduling strategy that considers the uncertainties in renewable energy sources and the energy demand (Yang et al. 2021). The proposed strategy introduces a mechanism to share information between DERs and adaptively weight the information to generate a microgrid energy schedule, minimizing the expected cost of energy while satisfying the energy demand and energy production capacity constraints. The model considers the variability and uncertainty of solar and wind energy sources and features the possibility to evaluate inevitable microgrid equipment malfunctions.

Recent studies show that optimal energy scheduling techniques can significantly improve microgrids' performance and sustainability. Additionally, ongoing research addresses the challenges posed by the variability and uncertainty of renewable energy sources and the dynamic nature of energy demand.

2.5 Operational Resiliency in Energy Systems

Operational resiliency refers to the ability of a system, organization, or infrastructure to continue to function or recover quickly during and after a disruption or crisis. It encompasses the ability to anticipate, prepare for, and respond to disruptions in a way that minimizes damage and downtime while enabling a swift return to normal operations. In a more focused definition, Alderson et al. propose that increasing operational resilience "make[s] the system better able to absorb the impact of an event without losing the capacity to function" (2015, p. 562). They reference California's electric power infrastructure, which makes real-time decisions to increase or decrease power flow from generators by opening or closing transmission switches on the grid, balancing overall generation and demand subject to the physical constraints of grid infrastructure. They describe that models optimizing resiliency are necessarily prescriptive, and the solution to a problem indicates the best method to reconcile objectives and constraints.

Exploring reliability and resiliency from a practical perspective, Riou et al. (2021) focused on effective integration of renewable DERs by developing a better method to determine appropriate microgrid sizing. Their multi-objective optimization methodology considers the microgrids' economic, environmental, and energy reliability aspects and uses a genetic algorithm to determine the optimally reliable microgrid configuration. By accounting for the output reliability of renewable resources, the researchers were able to mitigate some requirements for optimal energy scheduling through traditional power generation methods. However, there is an identified shortfall: it is cost-prohibitive to reach 100% reliability without over-sizing the microgrid. Integrating traditional generators or ESSs mitigates this shortfall.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 3:
Methodology

3.1 Microgrid Architecture

A typical FOB microgrid can have myriad configurations. We will discuss a typical, simple configuration, its components, and its method of operation in order set the parameters for our optimization model. Figure 3.1 is a generalized smart microgrid diagram and a helpful visualization for further discussion.

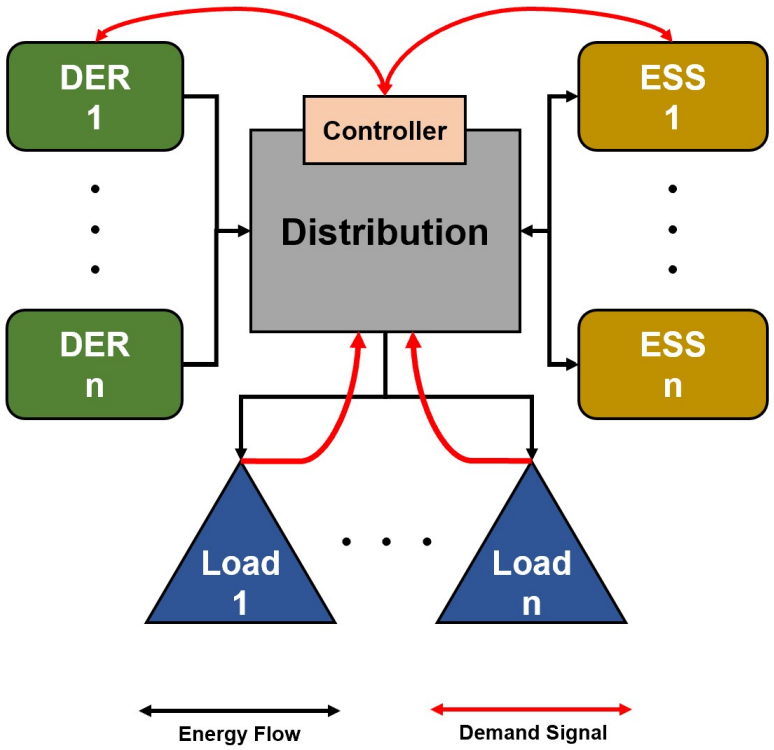

Figure 3.1. Notional smart microgrid configuration. Given some number of distributed energy resources, energy storage systems, and demand nodes, the control node predictively adjusts generation and sources to maximize efficiency. Adapted from Craparo and Sprague (2019).

An operational microgrid may have any number of DERs, ESSs, and demand loads. However, microgrids operating in austere environments are often constrained by what is easily portable, reliable, and repairable. As such, we find that microgrids will usually include one or two traditional generators, potentially one renewable DER, potentially one ESS, and any number of load demand sinks.

The advanced medium mobile power sources (AMMPS) is a widely-used U.S. Army power generator found at every echelon, from small units to major headquarters. It provides power in ranges from 5 to 60 kilowatt (kW) (United States Army Acquisition Support Center 2017). Our analysis considers a 60 kW, 60 hertz (Hz) AMMPS generator as it is one of the most common generator variants and because the three-year historic demand loads for our dataset fall within a feasible generation region for a single generator. We constrain the use of the generator to a minimum power output of 15 kW and maximum power output of 60 kW. The minimum power constraint prevents a generator condition called "wet-stacking," caused by operating a generator with a light load for long periods, resulting in the accumulation of unused fuel, carbon, and moisture in the exhaust, which may damage the generator's components. The maximum constraint is the generator's total production capacity.

We omit additional renewable DERs or multiple small generators in series, primarily due to the data indicating that a single generator could meet all load demands.

We add a single ESS to the microgrid architecture, giving it a maximum charge of 25 kilowatt-hours (kWh). We model the ESS as a Li-ion battery, and so we assume 90% round trip efficiency (RTE) and negligible charge loss due to the condition of the battery for the usage time frame (Faisal et al. 2018). As in Lee (2021), we constrain the ESS to a minimum and maximum state of charge (SOC) of 20% and 80%, respectively. While the SOC will vary between iterative policy generations throughout the day, we initialize to a 50% charge in the first period.

The number of demand load sinks that are attached to a microgrid is irrelevant to this model as long as the total energy demand for all loads at any point in time does not exceed the combined generation capacity of the DERs and discharge capacity of the ESSs. Therefore we model only a single demand load sink, which could theoretically have any number of smaller demand loads.

3.2 Recurrent Neural Network Architecture

To forecast power demand based on prior demand history, we require a tool capable of incorporating a long history of data and generating accurate predictions. To this end, the LSTM RNN is an excellent tool capable of taking valuable information from the entire corpus to generate predictions. In addition, it features the ability to selectively forget previous states if they provide little valuable input (Goodfellow et al. 2016).

The crucial component for the RNN is the self-loop, where critical observations may continue to influence model development. The weight of the self-loop is non-fixed and is based on another hidden unit, allowing the time scale to be changed dynamically. The result is that the model determines the time constants and not strictly the total input time data.

A block diagram for a shallow LSTM RNN, seen in Figure 3.2, acts as a basic model.

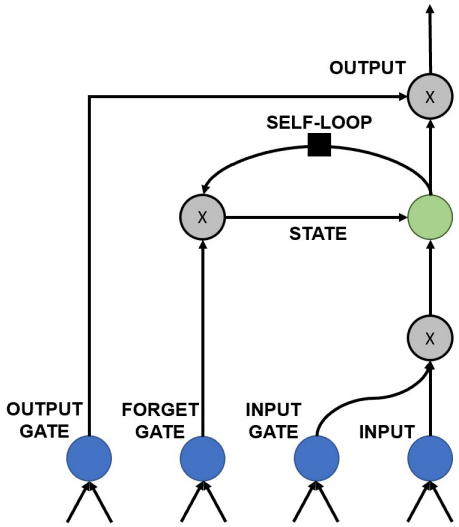

Figure 3.2. In a shallow LSTM RNN, cells are recurrently connected instead of the usual sequential hidden layers of neural networks. Nodes labeled with 'X' signify gates that are directly attached. The black square represents a single time-step. Adapted from Goodfellow et al. (2016)

.

In Figure 3.2, a value is incorporated into the state only if it can pass the sigmoidal input gate. Inputs that pass the input gate move into a linear self-loop with weights controlled by the forget gate, and the output gate ultimately controls the output of a cell. All gates

are sigmoid nonlinear. Observations are then sequentially incorporated into the network as inputs. Additional observations act as time steps to the self-loop, wherein observations must pass the forget and output gates for model development.

Our model uses a shallow LSTM RNN architecture similar to the one described above. The model uses mean-squared error (MSE) between training and validation data sets to adjust the weights assigned to inputs. We use batches of 288 training observations and periodic signals, signifying two days of demand data. After each batch of observations, the model updates its weights before moving on to the next batch and adjusting the weights to minimize MSE.

We additionally allow the model to run for a maximum of 20 epochs, which means the model can use all batches of the training data up to 20 times. Theoretically, this would allow the model to produce weights that are more effectively tuned to provide predictive accuracy, but may often result in over-training. To combat this phenomenon, we have also given the model a patience parameter, which must be some positive integer less than the maximum number of epochs. By using a patience parameter, the model will register the validation MSE for each epoch, and unless a following epoch within the range of the patience parameter has a validation MSE less than the current best epoch, the model will stop training before reaching the maximum number of epochs. The final model will then return the weights of the best-performing epoch, balancing overall model development time and predictive accuracy.

The final layer of our neural network is a single-node dense layer, which takes the outputs from the LSTM RNN layer as inputs and, in turn, outputs a single vector of predicted power demands.

After generating a series of predictions, we adjust our predictions by introducing a bias parameter shift_{RNN}, described in Equation 3.1.

$$\text{shift}_{RNN} = \left| \frac{1}{n} \sum_{1}^{n} (\text{rnnpred}_n) - \frac{1}{n} \sum_{1}^{n} (\text{naivepred}_n) \right| \tag{3.1}$$

In the equation, shift_{RNN} is the absolute difference of means between a vector of predictions

(rnnpred) and a naïve vector of predictions (naivepred), described as the vector of demand from the time increment preceding the desired predicted period. By shifting the LSTM RNN predictions per Equation 3.1, we can ensure that our predicted demands will generally overlay the actual demand of a microgrid.

3.3 Optimization Model

The optimization model we use is based mainly upon the work of Lee (2021). After describing the basic model, we depart from the single-run omniscient optimization model and introduce optimizations using recurrent prediction, thus developing a rolling forecast of energy demands.

3.3.1 Omniscient Optimization with Penalty

The MILP developed by Lee (2021) returned an energy production policy for a generator with the intent of minimizing fuel usage and penalizing changes in the generator's energy production while still meeting all energy demands placed on a single generator and ESS microgrid. Lee's objective function can be seen in Equation 3.2.

$$min \sum_{t \in T} [sl^b g_t + in^b + \sum_{h \in H} (slo_h Q_{h,t} + int_h W_{h,t})] \tag{3.2}$$

The optimization devised by Lee is composed of the generator's fuel consumption (sl^b) based on the generator's power production at a given time (g_t), the baseline fuel requirement for operating a generator (in^b), and a piecewise linear penalty term for changing generator energy production during within the policy interval ($slo_h Q_{h,t} + int_h W_{h,t}$).

Excepting those constraints and parameters that are changed to support a recurrent optimization, the architecture of the optimization model used remains unchanged from Lee's work.

3.3.2 Recurrent Predictive Optimization with Penalty

While generating a single predictive forecast for power demand is beneficial, there is a more significant application in forecasting demand recurrently at some specified triggering event (e.g., time of day, a maximum or minimum demand, a scheduled maintenance or refueling operation). Therefore, we make two updates to Lee's omniscient model to fulfill this requirement.

First, we remove the original model's constraint that the starting and ending ESS SOC must be equal to 50% of total charge capacity. The initial intent of this constraint was to represent the charge of the ESS changing between predictive intervals, and to mitigate end-of-horizon effects. Although our model exhibits end-of-horizon effects, we do not include this constraint because we will never execute this condition due to the nature of rolling-horizon predictions. However, we must ensure that the SOC is continuous between forecasting intervals such that the starting SOC for any given interval is equal to the preceding SOC when generating an energy production policy. To accomplish this, we introduce the requirement seen in Equation 3.3.

$$SOC_i = SOC_{(i-1)} \qquad \forall i \in I \tag{3.3}$$

In Equation 3.3, let i denote the forecast iteration. Our forecasts occur at regular intervals, at which point new predictions are generated (e.g., every 10 minutes, six hours, or once per day). With this condition in place, the generated microgrid policy will consider the ESS SOC at the current state of the system instead of using the starting or ending SOC for a given generation policy.

3.4 Adapting Policy to Requirements in Real Time

The predictive models described in this chapter produce a demand signal that is fed into Lee's modified optimization model to develop an operating policy for the DERs and ESSs. A shortcoming of the resulting policy is that it often fails to fulfill the total demand required at a given time t. Using the baseline predicted demands, we must therefore adjust our policy to fill actual demand as best as possible. This is most efficiently accomplished by altering our battery usage, increasing or decreasing battery output as necessary. If the battery is unable to fully respond to a demand discrepancy, then we must alter the generator plan.

3.4.1 Symbols

To describe the method for adjusting our microgrid generation policy, we introduce the following symbols.

g_t^p: prescribed DER output at time t [kW]
g_t: actual DER output at time t [kW]
d_t: actual demand in time period t [kW]
SOC_t: actual ESS state of charge at the end of time period t [kWh]
SOC_{min}: minimum ESS state of charge [kWh]
SOC_{max}: maximum ESS state of charge [kWh]
ϵ_c: ESS charging efficiency: if we send the ESS E units of energy, it receives $\epsilon_c E$
ϵ_d: ESS discharging efficiency: if the ESS sends E units of energy, we receive $\epsilon_d E$
Δt: length of each time period [hr]

3.4.2 Demand Discrepancy Cases and Our Responses

In each time period t, either $d_t > g_t^p$, $d_t < g_t^p$, or $d_t = g_t^p$. In the case where $d_t = g_t^p$, we run the DER as planned and make no changes to the ESS state (i.e., $SOC_{t+1} = SOC_t$). For cases where there is a generative and demand discrepancy, either we can adapt to the discrepancy entirely by modifying our ESS plan and operating the DER according to our original prescribed plan, or we cannot adapt to the discrepancy by modifying our ESS plan.

Case 1: Excess Demand, Battery Can Make Up the Production Shortfall

In the case that there is an excess of demand in comparison to our prescribed policy, but the ESS has sufficient charge to make up the shortfall, the following hold:

$$d_t \geq g_t^p \tag{3.4}$$

$$\epsilon_d(SOC_{t-1} - SOC_{min}) \geq \Delta t(d_t - g_t^p) \tag{3.5}$$

Because the microgrid has sufficient energy in the ESS to make up the shortfall, we should operate the generator according to plan and adjust the battery to meet the demand discrepancy:

$$g_t = g_t^p \tag{3.6}$$

$$SOC_t = SOC_{t-1} - \frac{\Delta t(d_t - g_t^p)}{\epsilon_d} \tag{3.7}$$

Case 2: Excess Demand, Battery Cannot Make Up the Shortfall

In the case that there is an excess of demand in comparison to our prescribed policy, but the ESS does not have sufficient charge to make up the shortfall, the following hold:

$$d_t \geq g_t^p \tag{3.8}$$

$$\epsilon_d(SOC_{t-1} - SOC_{min}) < \Delta t(d_t - g_t^p) \tag{3.9}$$

To meet the discrepancy, we discharge all available energy in the ESS and increase the output of the DER to fill the remaining demand:

$$g_t = d_t - \frac{\epsilon_d(SOC_{t-1} - SOC_{min})}{\Delta t} \tag{3.10}$$

$$SOC_t = SOC_{min} \tag{3.11}$$

Case 3: Excess Production, Battery Can Absorb Excess

In the case that we are over-prescribing energy production, and the battery has the ability to absorb the any extra generated energy, the following hold:

$$d_t < g_t^p \tag{3.12}$$

$$SOC_{max} - SOC_{t-1} \geq \epsilon_c \Delta t(g_t^p - d_t) \tag{3.13}$$

Then, we should operate the generator as the optimized plan prescribes and store the excess energy in the battery:

$$g_t = g_t^p \tag{3.14}$$

$$SOC_t = SOC_{t-1} + \epsilon_c \Delta t (g_t^p - d_t) \tag{3.15}$$

Case 4: Excess Production, Battery Cannot Absorb Excess

In the final case, we are over-prescribing energy production and the ESS is already at maximum capacity; therefore the following hold:

$$d_t < g_t^p \tag{3.16}$$

$$SOC_{max} - SOC_{t-1} < \epsilon_c \Delta t (g_t^p - d_t) \tag{3.17}$$

As opposed to the other three cases, here we could either leave the DER production policy alone and shed any excess production that cannot be stored in the ESS, or we could avoid shedding any excess by decreasing the production of the DER. If we adjust the DER to avoid shedding any excess:

$$g_t = d_t + \frac{SOC_{max} - SOC_{t-1}}{\epsilon_c \Delta t} \tag{3.18}$$

$$SOC_t = SOC_{max} \tag{3.19}$$

Applying the four cases as described gives us a final grid controller informed by our machine learning models, with real-time adjustments to allow dynamic load shifting.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 4:
Results and Analysis

4.1 Power Demand Case

The data for this model development comes from a U.S. FOB base in the Middle East. There are 35 records of separate microgrid demand histories relating to different use cases for this base, such as service member barracks, hygiene facilities, exchange shopping centers, and food service facilities.

Of the demand histories available, we have chosen the base operations center (BOC) microgrid to build our model. The BOC has a few features that make it attractive for building a predictive model for power demand:

1. In deployed operations, the BOC maintains 24-hour operations and should have a continuous demand signal.
2. A large BOC should send a large energy demand signal in accordance with the large staff required to maintain operations.
3. The floor area of the BOC is several thousand square feet and it is made of concrete. We expect periodic signals based on the time of day and the season to compensate for environmental conditioning.
4. The BOC microgrid has 126,037 observations, generally at 10-minute intervals. The first observation was on 09 August 2013, and the final was on 01 January 2016. There are over two years of observations within the observations without any missing values, with no need to interpolate demand.

Specific facility data is reported as a 13-column data frame, representing facility identification, commodity, the analysis component (i.e., total, average, median), timestamp, amp (A), amp-hours (Ah), kW, kWh, Power Factor, volts (V), volt amps (VA), volt-amps reactive (VAR), and Estimated. Three of these columns, namely facility identification, commodity, and the analysis component, are not valuable for model generation because they describe the corpus we are working with and are thus redundant.

We focus on the time stamp and kWh observations for the other columns. The former sets the time between observations, and the latter provides the metric of interest, the measure of energy demand. The remaining columns will necessarily report data that follows a similar pattern to kWh, so we discard them for simplicity. The first five rows of our data set appear in Table 4.1.

Table 4.1. First five observations of the BOC microgrid data

Time Stamp	kWh
2013-08-09 09:50:00	25.61
2013-08-09 10:00:00	25.35
2013-08-09 10:10:00	25.35
2013-08-09 10:20:00	25.94
2013-08-09 10:30:00	25.74

Figure 4.1 exhibits a sample of the energy demand generated by the BOC to the microgrid. The demand exhibits explicit periodic behavior, with a maximum demand of approximately 29.5 kW, and minimum demand of approximately 16 kW. The sudden demand drop near 12:00 on 05 July is unexplained, and may occur for any number of reasons not speculated on in this thesis.

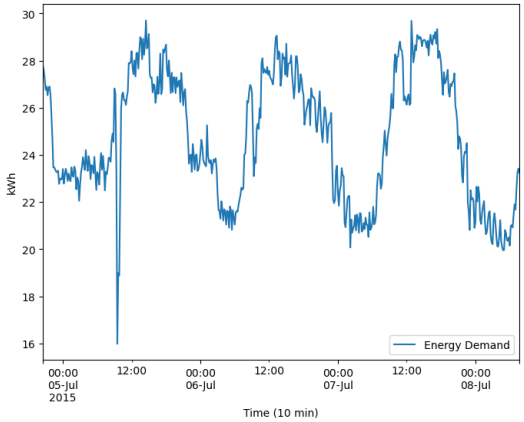

Figure 4.1. BOC demand from 05-08 July 2015.

The periodic behavior in Figure 4.1 appears to adhere to a daily cycle. This cycle makes intuitive sense based on work requirements increasing or decreasing based on the level of activity and operations a BOC would need to manage. For example, we expect power demand to be higher during the day while missions are occurring with greater frequency and to dwindle during the night when operations focus on maintaining situational awareness and operational continuity. We also expect power demand to be higher during the daylight hours in the summer due to air conditioning, and lower at night once the ambient temperature cools.

In addition to daily signals, a minor weekly signal and, potentially, long-term seasonal or yearly signals exist. The weekly signal may coincide with a standard work week, with more minor demands from Friday afternoons through Monday mornings. Despite being in a deployed, operational environment, many units will enact an extended rest cycle once weekly for up to 24 hours, often coinciding with the weekend. Yearly cycles may be especially weak but coincide with personnel inflow or outflow with 9–12 month deployments or long-term campaign changes.

With these practical realities in mind, we introduce artificial cycle waves to our data set to introduce some different periodic signals. A three-day example of the introduced signal may be seen in Figure 4.2.

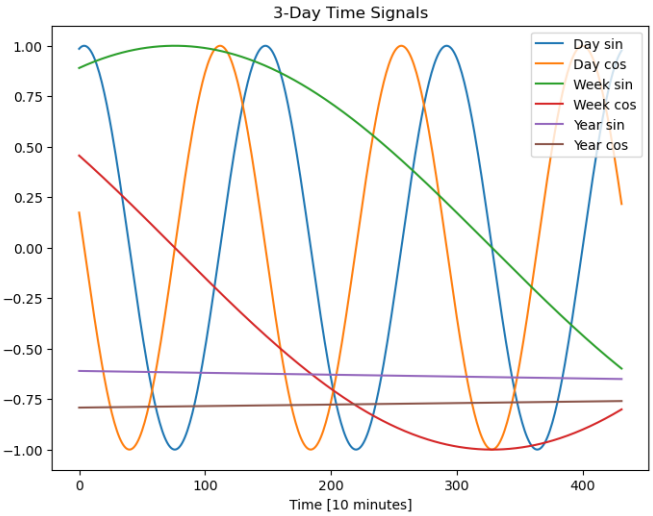

Figure 4.2. Periodic signals added to the data set, sending a stronger signal to the model that power demand may increase based on the time of day, day of the week, or seasonally.

4.2 Demand Predictions

We use three models to generate predictions on the BOC's energy demands. The first model is a naïve assumption that the demand from the previous day will be the same as the next day, called the "1-day naïve model" hereafter. The 1-day naïve model provides a reasonable approximation of the BOC's energy demands for any given day since it is largely unaffected by seasonal or long-term variations, and the prediction interval ensures that predictions will follow the same trends regarding increasing and decreasing demand cycles. A shortfall of the 1-day naïve model is that each demand observation is susceptible to sudden spikes of short-term variability and so may grossly under- or over-predict the amount of energy required for a given time t.

The second model we use is built from an LSTM RNN architecture but does not incorporate the artificial periodic signals shown in Figure 4.2. Based on this structure, it is known as the "Non-periodic LSTM RNN model." The model takes the 144 observations corresponding to the previous day's demands as input and predicts the following day's demand requirements. The periodic signals were excluded in this model so the generated demand predictions would be made naïvely, bridging an information gap between a truly naïve model and one

which could incorporate long-term trends. There is markedly less variability in the demand predictions of the non-periodic LSTM RNN model, though it does remain flexible and follow daily demand trends established by its input values.

The final model used to generate predictions incorporates the artificial periodic signals from Figure 4.2 in generating a LSTM RNN, and is thus called the "periodic LSTM RNN model." The LSTM RNN model is the only one of the three to use the periodic signals developed in 4.1, giving it additional predictive power for weekly and long-term trends. As such, the predictions generated by the periodic LSTM RNN model do not exhibit the variability seen in the naïve and non-periodic models, having a greater tendency towards smooth changes in energy demands.

Using a prediction window of one day, Figure 4.3 shows a sequence of three predictions. In each figure, the blue line represents the actual energy demand generated by the BOC. The first 144 observations, representing the previous day, are considered known data. The second 144 observations show the actual energy demand, the 1-day naïve model's prediction as a copy of the first 144 observations, and both the periodic and non-periodic LSTM RNN model predictions. Predictions are regenerated every 36 observations, representing six hours, so we can see how the demand and predictions shift accordingly.

A comparison of the three predictive models over a sample of three consecutive prediction intervals and their associated mean absolute errors (MAEs) are given in Table 4.2 and shown in Figure 4.3 on the following page.

Model	Interval 1	Interval 2	Interval 3
1-Day Naïve	0.680	0.632	0.630
Non-Periodic LSTM RNN	0.787	0.724	0.766
Periodic LSTM RNN	0.513	0.539	0.563

Table 4.2. MAEs for different models on the three sample intervals from Figure 4.3

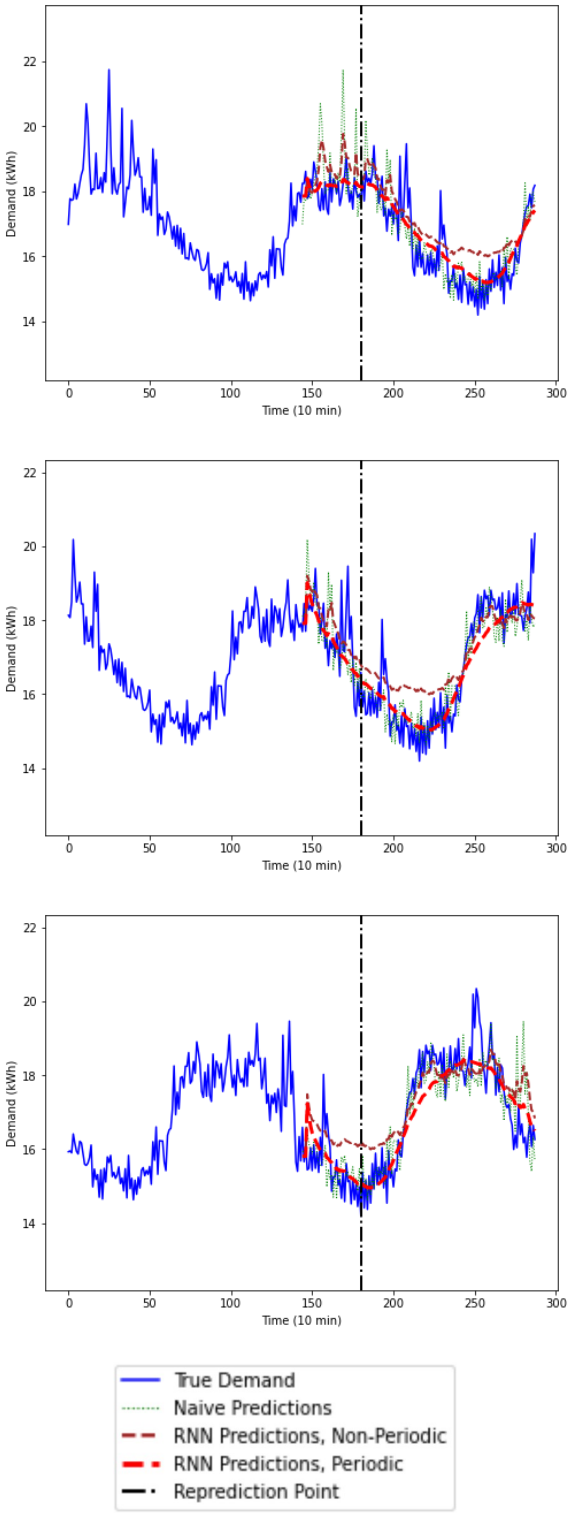

Figure 4.3. Three consecutive LSTM RNN and 1-day naïve predictions, re-generated at 6-hour intervals

4.3 Model Performance Comparison

We obtain long-term fuel consumption metrics by running the modified optimization model of Lee (2021) using each of the three LSTM RNN models as input, in a rolling horizon fashion. Each of the models examined has comparable average performance in the long term. An abbreviated summary for three different reprediction intervals–six, eight, and twelve hours–is shown in Tables 4.3, 4.4, and 4.5, with the best-performing metric emphasized in bold-italic. We selected these intervals to represent the recreation of an energy generation policy two, three, or four times per day. For these results, we assume that the DER and ESS resources exactly execute the plans prescribed by the optimization model, regardless of any discrepancies between the actual demand and the demand forecast by the corresponding LSTM RNN.

6 Hour Reprediction Interval		
Model	Average Fuel Use (gal/hr)	Average Demand Met (%)
1-Day Naïve	1.7719	83.40
Non-Periodic LSTM RNN	1.7883	***84.85***
Periodic LSTM RNN	***1.7697***	83.67

Table 4.3. Average model performance over a continuous 75-day sample, optimally solved

8 Hour Reprediction Interval		
Model	Average Fuel Use (gal/hr)	Average Demand Met (%)
1-Day naïve	1.7802	83.44
Non-Periodic LSTM RNN	1.7961	***84.90***
Periodic LSTM RNN	***1.7766***	83.69

Table 4.4. Average model performance over a continuous 75-day sample, optimally solved

12 Hour Reprediction Interval		
Model	Average Fuel Use (gal/hr)	Average Demand Met (%)
1-Day Naïve	***1.7806***	83.42
Non-Periodic LSTM RNN	1.7980	***84.86***
Periodic LSTM RNN	1.7810	83.69

Table 4.5. Average model performance over a continuous 75-day sample, optimally solved

The results presented demonstrate the competitive performance of LSTM RNN models against a truly naïve model. Regardless of the reprediction interval, the smoothing nature of the periodic LSTM RNN model meets the BOC's demand more consistently with lower fuel consumption than with the naïve predictions. Furthermore, we see that the non-periodic LSTM RNN model consistently has the best performance in meeting energy demands, but also has the highest average fuel consumption of the three models. This is because the non-periodic model's predicted demands are consistently higher than the other two models, and so the generation policy accordingly requires greater power generation output. Determining whether the increased operating cost is worth better performance in meeting average demand is dependent on operational use case.

The example predictions shown in the following sections occur in sequence at 1400 on 28 December 2015, 2000 on 28 December 2015, and 0200 on 29 December 2015. This time frame was selected because the interval has a high amount of demand variation and because it lies at the end of our rolling horizon, mitigating any assumptions about the ESS's starting SOC.

4.3.1 1-Day Naïve Prediction Performance

Figure 4.4 shows the performance-optimized generation policy provided by the 1-day naïve policy, with respect to energy generation, storage, discharge, and demand shortfall.

In the top row of Figure 4.4, we can examine when the model plans to supply, store, or discharge energy under forecast demand conditions, as well as the actual energy demand of the BOC. In the bottom row, we see the planned SOC of the microgrid's ESS. The vertical dash-dot line represents the point at which we generate a new optimal energy generation

policy. We note that the optimization model incentivizes the generator to maintain steady-state power generation with minimal shifts in policy, avoiding the fuel cost penalty prescribed in Lee's optimization model. The ESS fulfills shortfalls in the naïvely predicted demand, but overall we see that the 1-day naïve model regularly fails to meet the BOC's actual energy demands during some time periods, while also producing excess energy that is not stored in the ESS in other time periods.

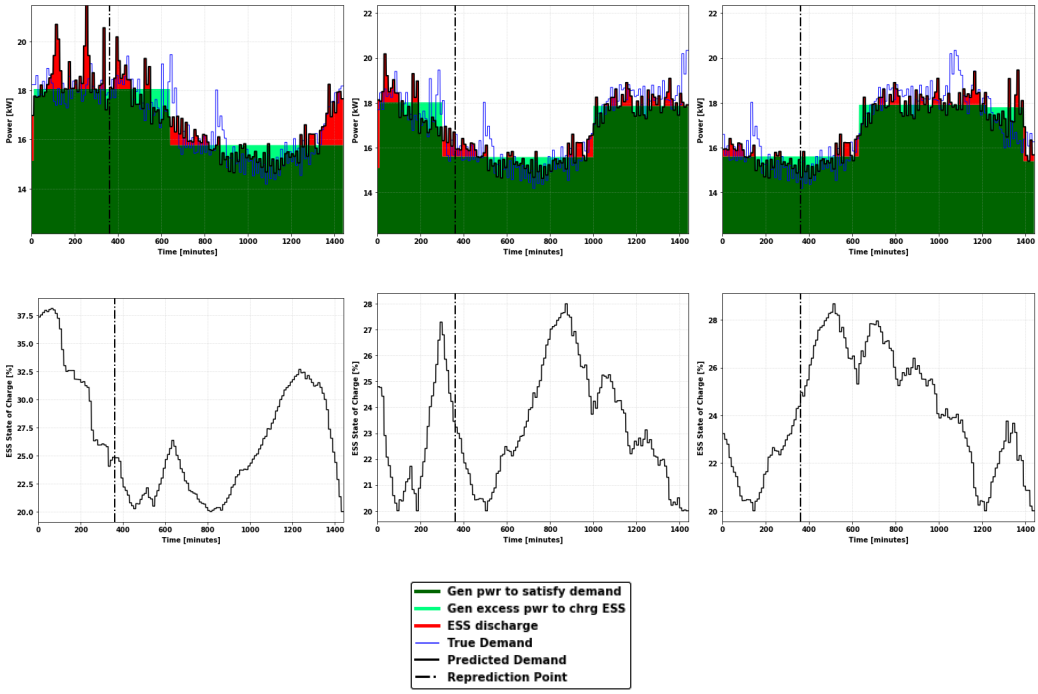

Figure 4.4. Three consecutive prediction intervals for 1-day naïve model
Top: 1-day naïve predictive performance for initial prediction, initial prediction + 6 hours, and second prediction + 6 hours
Bottom: 1-day naïve ESS SOC, given the same predictive horizons

4.3.2 Non-Periodic LSTM RNN Prediction Performance

The non-periodic LSTM RNN model's performance is shown in Figure 4.5. Compared to the 1-day naïve model, the non-periodic model better fulfills the BOC's energy demands with less variability between time steps.

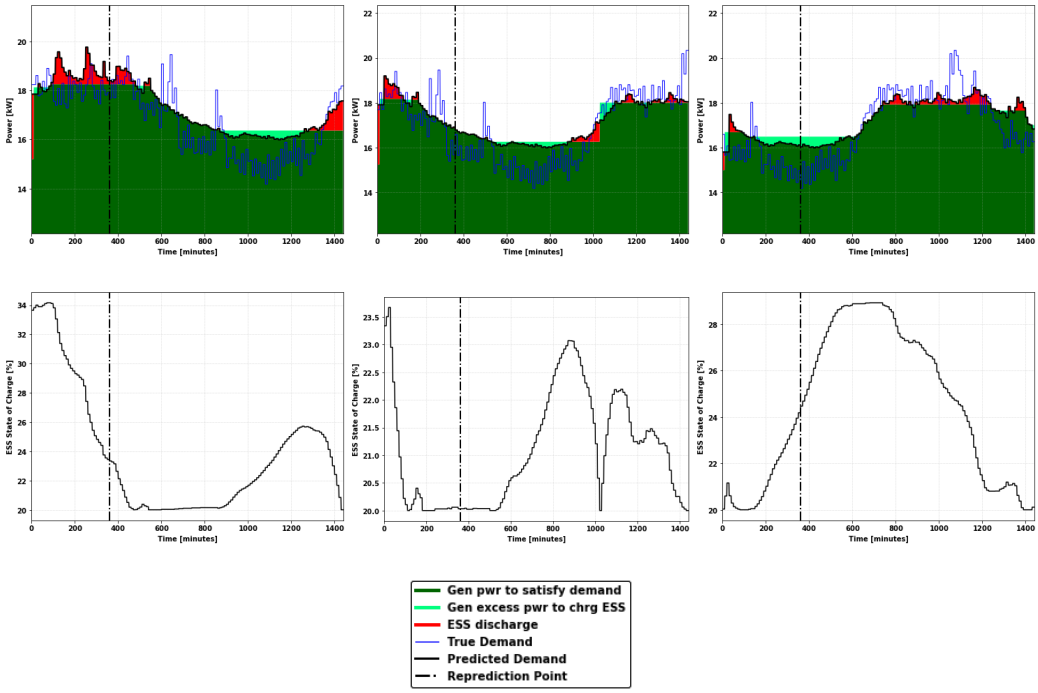

Figure 4.5. Three consecutive prediction intervals for the non-periodic LSTM RNN model
Top: Non-periodic LSTM RNN predictive performance for initial prediction, initial prediction + 6 hours, and second prediction + 6 hours
Bottom: non-periodic LSTM RNN ESS SOC, given the same predictive horizons

A shortcoming of the non-periodic LSTM RNN model's performance is its tendency to over-predict demand, leading to increased fuel consumption in an optimized power generation policy. This is especially evident during periods of low power demand, as can be seen in Figure 4.5. However, the model's predictions remain flexible to changes in demand during times of high variability, so this model may be suited to applications with dramatically shifting loads and requirements for increased reliability.

4.3.3 Periodic LSTM RNN Prediction Performance

Compared to the naïve model and non-periodic LSTM RNN model, the periodic model has a smoother range of predictions, as can be seen in Figure 4.6. It more consistently meets the BOC's energy demands than the naïve model, and has a lower fuel consumption rate than either of the previous models.

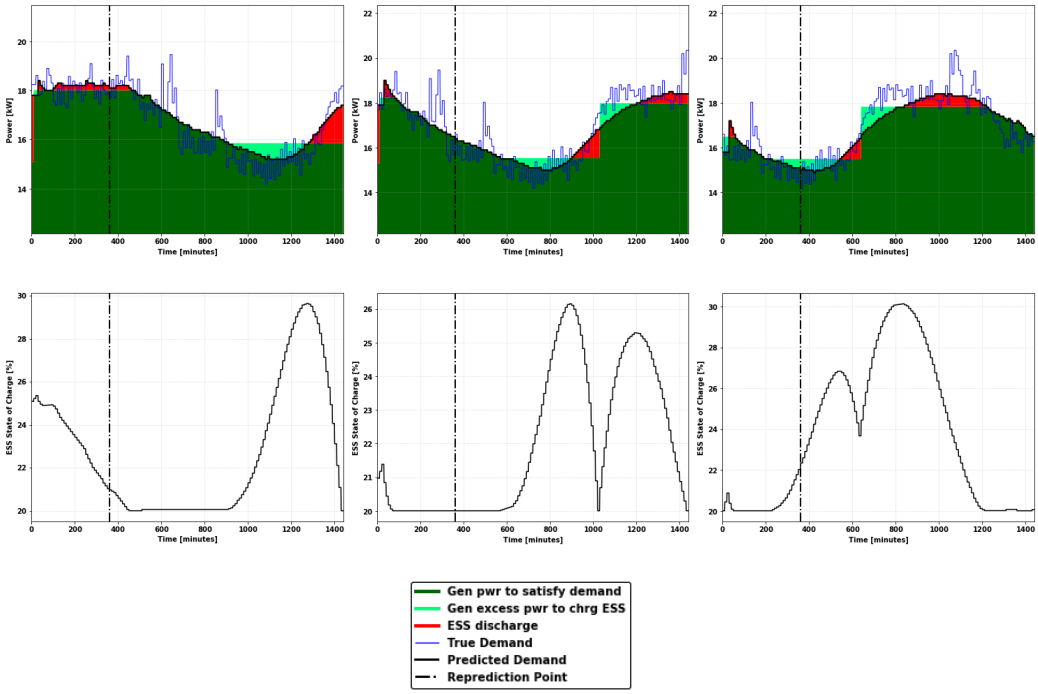

Figure 4.6. Three consecutive prediction intervals for LSTM RNN model
Top: Periodic LSTM RNN predictive performance for initial prediction, initial prediction + 6 hours, and second prediction + 6 hours
Bottom: Periodic LSTM RNN ESS SOC, given the same predictive horizons

Because of the smoother nature of demand predictions, we likewise see smoother step transitions between high and low energy generation prescriptions. Figure 4.6 also shows that the smaller demand variance between time steps results in more consolidated ESS charging and discharging events and a reduced requirement for the generator to over-produce energy to optimally meet predicted energy demand needs.

4.4 Adapting to Demands in Real Time

In Section 3.4 we described a method by which we may adjust the predicted and optimized generative policy such that a microgrid would respond to power demands in real-time. We are primarily responding to sudden increases or decreases in demand by calling on the ESS to fill the demand when able or to receive and store any energy that comes from an over-prescription of generation. While these adjustments may sacrifice some measure of optimal generation, the changes allow us to meet 100% of demand in all cases, which is seen as a hard requirement in operational settings.

4.4.1 Real-Time Naïve Prediction Performance

One of the primary shortcomings of the 1-day naïve predictive model was that sudden demand spikes or demand drop-offs resulted in either a major generative energy shortfall or waste, all while not being able to sufficiently meet the reliability requirements of an operational microgrid. The implementation of the real-time production policy remediates this shortfall. As can be seen in Figure 4.7, the primary change comes from ESS utilization, wherein we use the battery to rapidly fill unanticipated energy requirements, and depart from optimally prescribed generator operating levels only when the ESS SOC is at its minimal charge level.

In fact, we find that over a 75-day sample of real-time adjustments to the 1-day naïve model as outlined in Section 3.4, we incur only marginal increases to fuel consumption rate in exchange for meeting 100% of energy demand, regardless of which reprediction interval is used. As may be seen in Table 4.6, the additional fuel cost to meet all demand using the 1-day naïve model is roughly 0.008 gallons per hour.

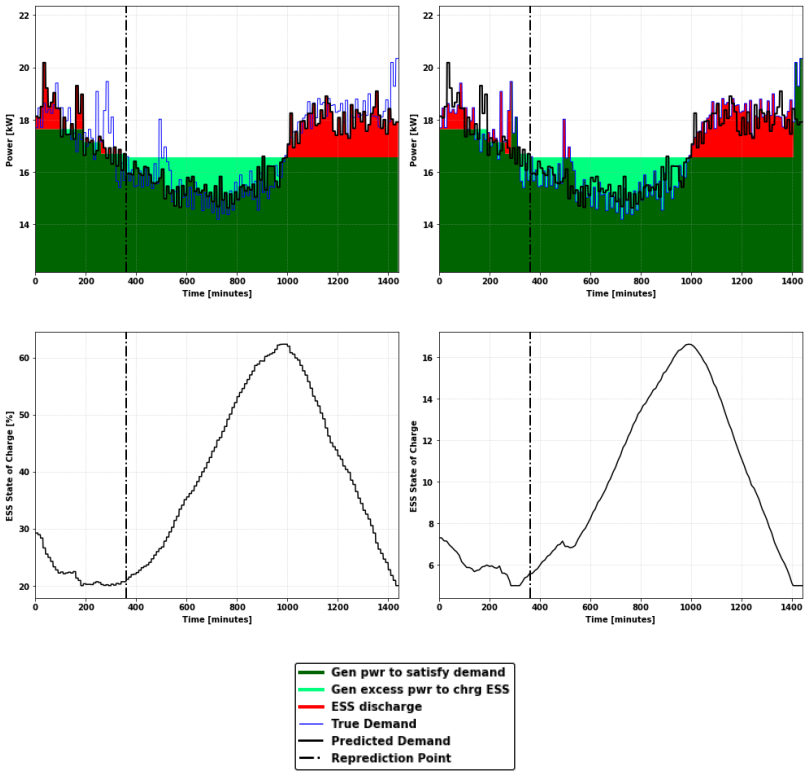

Figure 4.7. Performance comparison between the naïve predictive model without real-time adjustment (left), and with real-time load-shifting (right)

Naïve Model Performance Comparison		
Model	Non-Adjusted	Adjusted
6 Hour Reprediction Fuel Use (gal/hr)	1.7719	1.7796
Demand Met	83.40 %	100%
8 Hour Reprediction (gal/hr)	1.7707	1.7787
Demand Met	83.44 %	100%
12 Hour Reprediction (gal/hr)	1.7715	1.7796
Demand Met	83.42 %	100%

Table 4.6. Comparison of naïve model average performance over a continuous 75-day sample, with and without real-time adjustment to demand

4.4.2 Real-Time Non-Periodic LSTM RNN Prediction Performance

Compared with the 1-day naïve predictive model, both of the LSTM RNN predictive models exhibit much smoother demand predictions. Focusing on the non-periodic LSTM RNN, recall that in Section 4.3.2 we saw consistent over-prediction of demand, resulting in an optimized generative policy that regularly provided a surplus of energy which was not being stored, and was regularly shed and wasted.

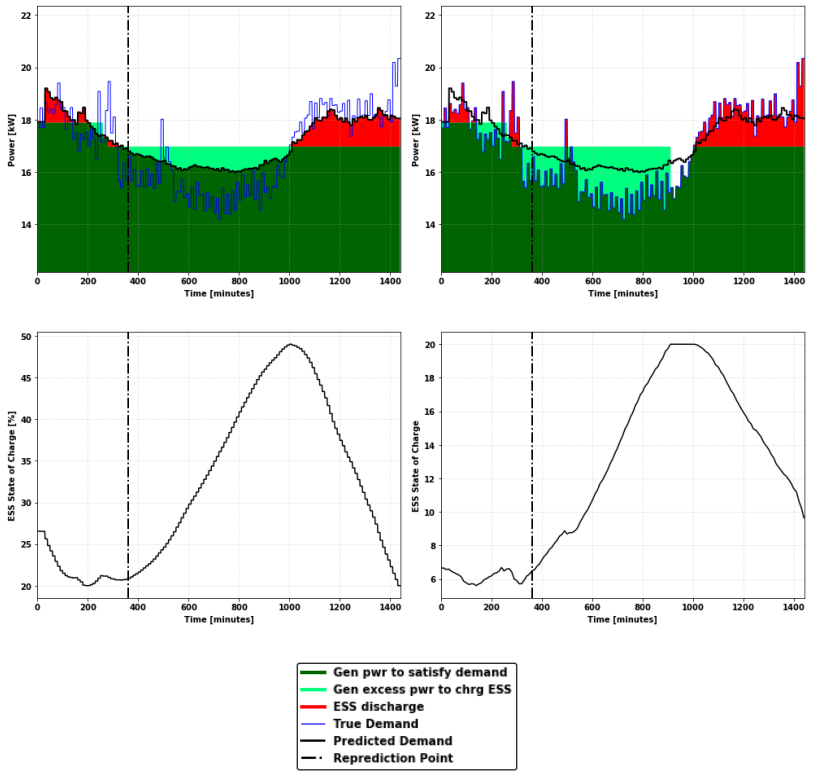

Figure 4.8. Performance comparison between the Non-Periodic LSTM RNN predictive model without real-time adjustment (left), and with real-time load-shifting (right)

Implementing the real-time adjustment cases dramatically reduces the over-generation of the non-periodic LSTM RNN model while simultaneously making better use of the ESS. As can be seen in Figure 4.8, the real-time adjustment to demand better responds to sudden increases in demand, while the prescribed levels at which the generator operates is so much greater than the actual demand levels that the ESS would reach its maximum energy storage

capacity. In our proposed policy, the ESS being unable to store any more energy has the effect of causing the generator to lower its overall output. While this rapid load shifting may be undesirable in terms of generator lifespan, the instances seen in this model are minor and come with a surprising side effect: due to the regular decreases in generator output, the real-time adjusted policy ends up being having a lower fuel consumption rate than the non-adjusted policy. The differences in fuel consumption rate compared to the percent of energy demand met by the adjusted and non-adjusted policies for 6-, 8-, and 12-hour reprediction intervals may be seen in Table 4.7.

Non-Periodic LSTM RNN Model Performance Comparison		
Model	Non-Adjusted	Adjusted
6 Hour Reprediction Fuel Use (gal/hr)	1.7883	1.7866
Demand Met	84.85 %	100%
8 Hour Reprediction (gal/hr)	1.7879	1.7859
Demand Met	84.90 %	100%
12 Hour Reprediction (gal/hr)	1.7899	1.7878
Demand Met	84.86 %	100%

Table 4.7. Comparison of non-periodic LSTM RNN model average performance over a continuous 75-day sample, with and without real-time adjustment to demand

4.4.3 Real-Time Periodic LSTM RNN Prediction Performance

Recall that of the three predictive models used, the periodic LSTM RNN had the best generalized fit to true demand excepting any sudden changes to energy demand levels. As a result, the optimized generator policy remained relatively steady with smooth demand curves that were often fulfilled by the ESS and minimal generator state change.

The real-time adjusted policy benefits from these features, as the adjusted periodic LSTM RNN model can make better use of the ESS than the non-periodic model. In the current case, the ESS rarely reaches its maximum capacity, so there are few changes to a prescribed generation policy in any given demand interval. An example of the real-time adjusted policy may be seen in Figure 4.9.

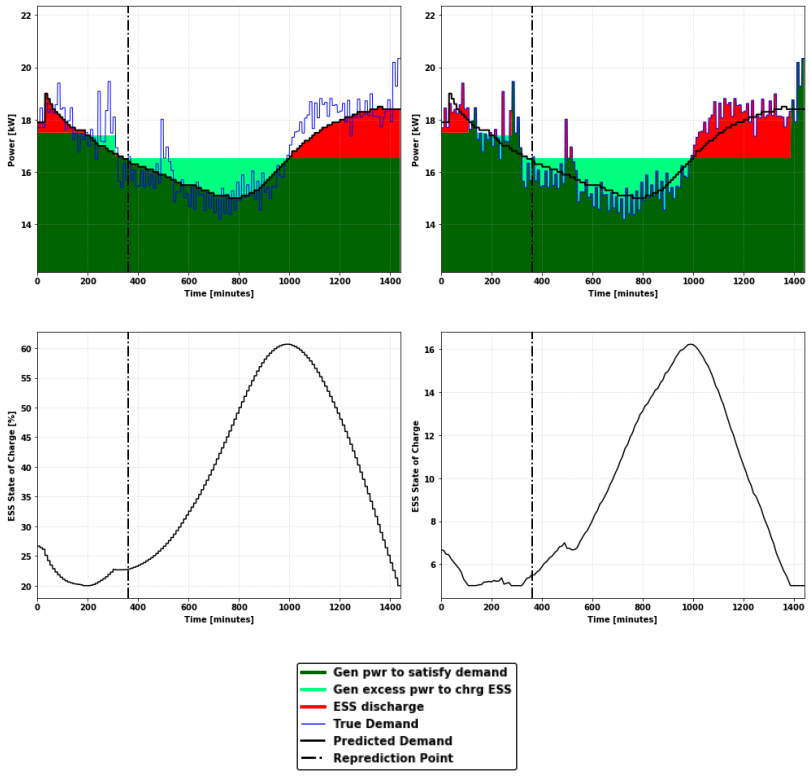

Figure 4.9. Performance comparison between the periodic LSTM RNN predictive model without real-time adjustment (left), and with real-time load-shifting (right)

Of the three models, implementing the real-time adjustment policy has the greatest effect on the periodic LSTM RNN model's fuel consumption rate, increasing fuel consumption by over 0.01 gallons per hour regardless of reprediction interval, as may be seen in Table 4.8. Of course, the small increase in fuel consumption rate that we see is still less than any rate seen in the non-periodic LSTM RNN. Additionally, we see that adjusted periodic LSTM RNN model with an 8-hour reprediction interval yields the lowest average fuel consumption rate of any model examined here.

Periodic LSTM RNN Model Performance Comparison		
Model	Non-Adjusted	Adjusted
6 Hour Reprediction Fuel Use (gal/hr)	1.7697	1.7801
Demand Met	83.67 %	100%
8 Hour Reprediction (gal/hr)	1.7672	1.7783
Demand Met	84.90 %	100%
12 Hour Reprediction (gal/hr)	1.7715	1.7820
Demand Met	84.86 %	100%

Table 4.8. Comparison of periodic LSTM RNN model average performance over a continuous 75-day sample, with and without real-time adjustment to demand

4.5 Model Performance Comparison

We have now developed three models to predict future energy demand for the BOC use case, given our notional microgrid architecture. We have also established three distinct intervals in which to regenerate predictions and generative policies for the microgrid, after which we apply a real-time adjustment to ensure that the microgrid can provide 100% of demanded energy.

Let us now establish two fuel consumption performance baselines so that we can compare our model performance to some minimum and maximum standards of performance. Consider Lee's omniscient model, which optimally solves for DER and ESS outputs for a dynamic load. Lee's model fulfills all demand presented to the model and minimizes fuel consumption, and therefore is the absolute best our system could perform if it had perfect predictive power. Conversely, a microgrid operating without an ESS is a suitable worst-performance baseline, because the DER would be required to change its energy generation output at every instance of demand change and thus be required to operate a high generation levels and incur several fuel consumption penalties every time there is a change in demand level.

Using 75 consecutive days' demand and a rolling horizon implementation, we determine

the long-term fuel consumption rate performance for each of the models. For our previously established 6-, 8-, and 12-hour reprediction intervals this results in 300, 225, and 150 average fuel consumption observations, respectively. To determine whether any differences exist between the models' performance for each reprediction interval, we now conduct a statistical analysis.

4.5.1 Statistical Analysis of Fuel Consumption

We first utilize the Friedman test (Friedman 1937), which allows us to determine whether any one model is outperforming the others in terms of fuel consumption. A major benefit of the test is that allows us to compare the models while avoiding the assumption of normality. We consider five models: our two LSTM RNN-based models and 1-day naïve model with real-time policy adjustments, Lee's omniscient model, and the no-ESS worst-case baseline. Let the null hypothesis H_0 be that the models have the same performance, and the alternative hypothesis H_a be that the models are different. Conducting the Friedman test and controlling for the reprediction intervals as a blocking factor, we arrive at the p-values in Table 4.9.

Friedman Test Results for Model Comparison	
Reprediction Interval	p-value
6-hour	< 2.2e-16
8-hour	< 2.2e-16
12-hour	< 2.2e-16

Table 4.9. Friedman Test comparing models, where H_0 is that each predictive model has identical performance, and H_a is that at least one of the predictive models is different from the others

The p-value for each reprediction interval is functionally zero, and so we strongly reject the null hypothesis in favor of the alternative.

Having determined that at least one model is performing differently than another, we next conduct a multiple comparisons test to reveal which model(s) have similar performances for each reprediction interval, and which models differ. Each reprediction interval yields the same general results for model similarity, and so an aggregate comparison is presented in Table 4.10, where statistical significance is determined at the 0.05 level.

Multiple Model Comparison After Friedman Test		
Model 1	Model 2	Statistically Similar Performance
Omniscient	1-Day Naïve	Yes
Omniscient	Periodic LSTM RNN	Yes
Omniscient	Non-Periodic LSTM RNN	No
Omniscient	No ESS	No
1-Day Naïve	Periodic LSTM RNN	Yes
1-Day Naïve	Non-Periodic LSTM RNN	No
1-Day Naïve	No ESS	No
Periodic LSTM RNN	Non-Periodic LSTM RNN	No
Periodic LSTM RNN	No ESS	No
Non-Periodic LSTM RNN	No ESS	No

Table 4.10. Multiple model comparison after Friedman test, showing which models had similar performances at a 0.05 significance level over a 75-day test interval

Interestingly, we see that, regardless of the reprediction interval, the naïve and periodic LSTM RNN predictive models with policy adjustment have a statistically similar performance to Lee's optimal, omniscient model. The naïve and periodic LSTM RNN models also exhibit similar performance to each other.

If we take a broader view and look at the performance of each model in terms of average fuel consumption, we can compare each model's fuel consumption rate for a given reprediction interval to the best case scenario. Based on the statistical similarity observed in the multiple comparisons test, we would expect both the naïve and periodic LSTM RNN models to have consumption rates approaching that of the optimal, omniscient model. The performance of all adjusted models at the 6-, 8-, and 12-hour reprediction intervals can be seen in Table 4.11, and as expected the naïve and periodic LSTM RNN models have the best performance.

Average Fuel Consumption Comparison			
Model	6-Hour Reprediction (gal/hr)	8-Hour Reprediction (gal/hr)	12-Hour Reprediction (gal/hr)
Omniscient	1.7767	1.7767	1.7767
1-Day Naïve	**1.7796 (+0.0029)**	1.7787 (+0.0020)	**1.7796 (+0.0029)**
Non-Periodic LSTM RNN	1.7866 (+0.0099)	1.7859 (+0.0092)	1.7878 (+0.0111)
Periodic LSTM RNN	1.7801 (+0.0034)	**1.7783 (+0.0016)**	1.7820 (+0.0053)

Table 4.11. Comparison between optimal performance and model performance based on fuel consumption rate (gal/hr), including difference between a model's performance and optimal performance. The best performing model for each reprediction interval is emphasized in bold.

4.5.2 Model Iterative Performance Comparison

While we have shown that the 1-day naïve and periodic LSTM RNN have a performance that is similar to an optimal, omniscient model, there are instances where the predictive models may outperform the optimal model in the short term. In fact, for each reprediction interval there are instances where all three of the predictive models have a lower fuel consumption rate than the optimal, omniscient model. Histograms of the count and magnitude of fuel efficiency differences between each model and the best-case optimal omniscient model, as well as the previously described worst-case omniscient model, may be seen in Figure 4.10.

The reason that the predictive models may be more fuel efficient is because for any iteration a model may generate a policy that over-charges the ESS, when compared to an optimal charge plan. Then, when the reprediction point occurs and the model generates a new operating policy for the DER, the ESS has the ability to assume more of the demand load than the generator. This results in a short-term reduction in fuel consumption rate, but will not be sustainable for long-term operations.

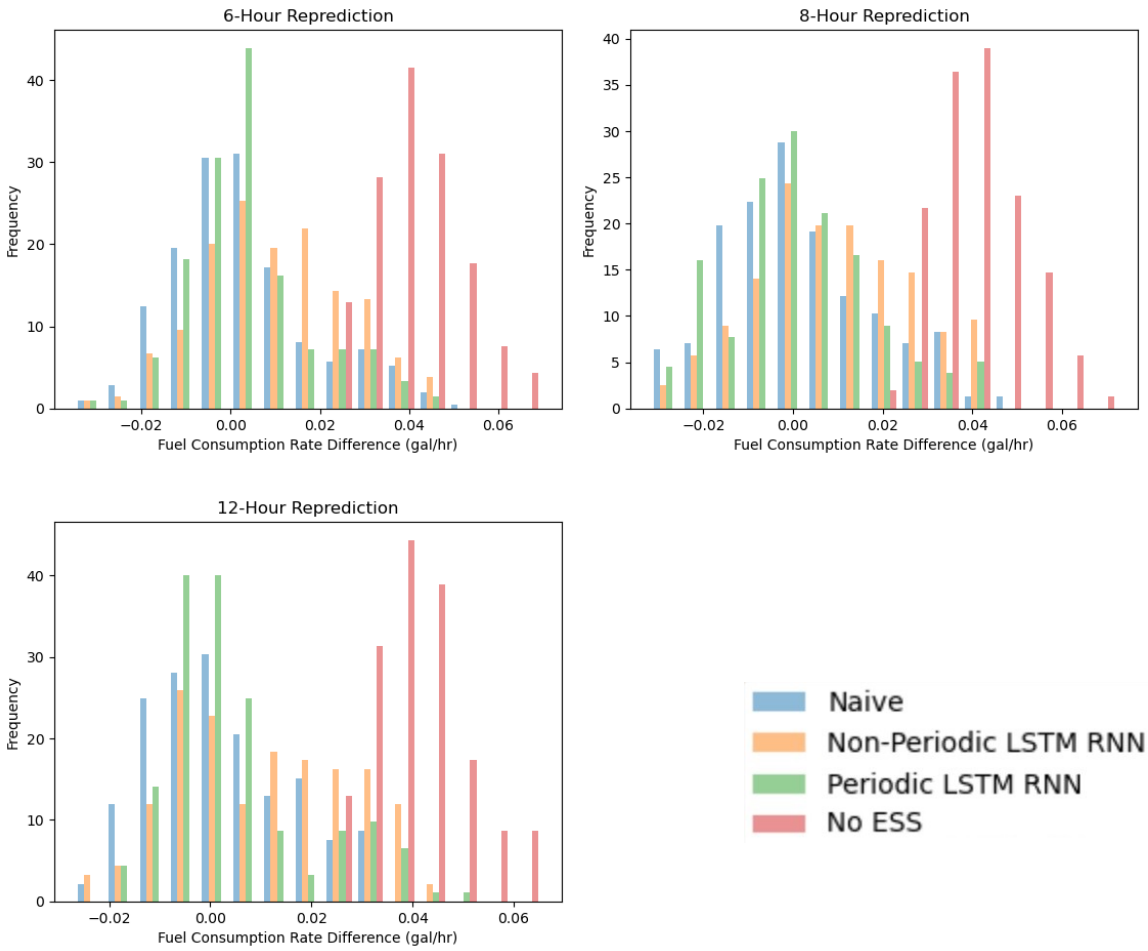

Figure 4.10. Normalized counts of fuel consumption performance difference between the optimal, omniscient model and the three predictive models, as well as the worst-case model without an ESS

4.6 Fiscal Impact

Given that we have different model performances, we are interested in how changes in fuel consumption translate to operating costs. As of March 2022, the U.S. Energy Information Administration reports that the price of jet propellant 8 (JP-8), a fuel commonly used in 60 kW AMMPS generators, is $3.285 per gallon. Assuming an operation employs 100 microgrids of the same architecture and supports similar facilities described in the use case, we can estimate the differences in fuel operating costs over the course of one year. We will

exclusively examine the case of 8-hour reprediction intervals in Table 4.12, since that is the case in which we found the best predictive performance.

Fuel Operating Costs, 8-Hour Reprediction Models			
Model	Average Fuel Use (gal/hr)	Yearly Total Cost ($)	Cost Difference ($)
Omniscient	1.7767	213,176.68	–
1-Day Naïve	1.7787	213,416.65	+239.97
Non-Periodic LSTM RNN	1.7859	214,280.54	+1,103.86
Periodic LSTM RNN	1.7783	213,368.66	+191.98

Table 4.12. Average cost difference using 8-hour reprediction interval for 100 60kW generators over one year, using the optimal, omniscient model as a baseline cost

At the referenced fuel costs and consumption rate, we incur a cost of $191.98–about 58 gallons of JP-8–in comparison to optimal performance. This is less than a 0.1% increase from the total optimal fuel cost for the year.

However, expeditionary generators are often not accompanied by a microgrid controller and ESS to increase overall microgrid fuel efficiency. Therefore, we also examine the performance of our microgrid architecture against the baseline worst-case performance model.

Fuel Operating Costs, 8 Hour Reprediction Models			
Model	Average Fuel Use (gal/hr)	Yearly Total Cost ($)	Cost Difference ($)
Baseline Worst Case	1.8173	218,048.06	–
1-Day Naïve	1.7787	213,416.65	–4,631.41
Non-Periodic LSTM RNN	1.7859	214,280.54	–3,767.52
Periodic LSTM RNN	1.7783	213,368.66	–4,679.40

Table 4.13. Average cost difference using 8-hour reprediction interval for 100 60kW generators over one year, using the optimal, omniscient model without an ESS as a baseline cost

When compared to a generator without an ESS, we see that our predictive models could save $3,700–$4,600 per year for a large, expeditionary base in raw fuel cost alone. This is equivalent to 1,100–1,400 gallons of fuel, which would introduce additional savings for fuel transportation and storage.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 5:
Conclusion

This thesis examines three methods of recurrently generating power demand predictions for a microgrid. The generated power demand predictions subsequently use a MILP to optimally prescribe a power generation policy for a microgrid incorporating an ESS, minimizing fuel usage while meeting predicted demand requirements. After producing a prescribed generation policy, we adjust the policy in real time as the true demand at a given moment is presented to the microgrid.

The three methods of generating power demand predictions are a naïve model that uses the previous day's demand as today's predicted demand, a LSTM RNN model that does not include artificial long-term periodic influence and a LSTM RNN model that provides for artificial daily, weekly and yearly signals. Each model's performance is examined using the actual microgrid power demand data generated from a BOC in the Middle East. Compared to the naïve model, the non-periodic LSTM RNN has a smoother generation policy but consistently over-predicts energy demands resulting in meeting demand requirements more frequently while also increasing overall fuel usage. The periodic LSTM RNN prediction model has the smoothest generator operating profile and is shown to provide more energy for less fuel.

By adjusting the predictive models based on presented demand, we show that using a periodic LSTM RNN model with 8-hour reprediction intervals can achieve a performance that is comparable and competitive with an omniscient, optimized model. The best performing model presented in this thesis has an average fuel consumption rate only 0.0016 gal/hr greater than a theoretical best-case scenario model.

In addition to any cost savings that may be actualized from reduced fuel usage, further savings may be found by considering fuel transportation and fuel storage reductions. Likewise, smoother generator operating profiles put less strain on traditional fuel generators, reducing maintenance requirements and increasing equipment lifetime.

5.1 Future Work

The recurrent optimization presented in this work supports leveraging RNN models to build efficiencies in microgrids. However, as the case study microgrid comprises only a traditional fuel generator and ESS, there is significant room for growth in exploring model performance while incorporating additional or different DERs. For example, microgrids powering DOD elements in remote or austere environments frequently supplement energy demands with PV or wind-based resources. The inclusion of these renewable sources may drive down fuel usage by providing power directly to the ESS, increasing the SOC at any given time, and allowing the ESS to reliably fill a more significant proportion of the demand.

In the same vein as the inclusion of renewable DERs in microgrid architecture, it may be possible to increase the accuracy of demand forecasts by developing predictive models that include forecast weather data. Due to the large amount of energy used for environmental conditioning, additional information about incoming weather patterns would provide information beyond historical daily or seasonal patterns. The increased accuracy of predictions would help prevent the over-generation of energy, allowing better sizing of a microgrid to energy demand needs, reducing fuel usage, and reducing the potential for maintenance issues.

The models developed in this thesis focused on the case study referenced in Section 4.1. Namely, the models supported a small administrative building. As such, while the naïve model can be applied directly to any facility, the LSTM RNN models are trained to a specific use case. Significant room for improvement lies in developing additional models tailored to a suite of buildings or flexible enough to provide accurate predictions regardless of the historic demand profile. Developing models for different use cases may reveal greater saving opportunities than those found in this work.

List of References

Ahmed M, Meegahapola L, Vahidnia A, Datta M (2015) Stability and control aspects of microgrid architectures—a comprehensive review. *IEEE Access* 8:144730–144766, https://doi.org/10.1109/ACCESS.2020.3014977.

Alderson D, Brown G, Carlyle W (2015) Operational models of infrastructure resilience. *Risk Analysis* 35(4):562–586, https://doi.org/10.1111/risa.12333.

Assistant Secretary of Defense for Sustainment (2021) Fiscal Year 2022 Operational energy budget certification report. Washington, D.C., https://www.acq.osd.mil/eie/Downloads/OE/FY22%20OE%20Budget%20Certification%20Report.pdf.

Baskar P, Kaluvan H, Nandkeolyar R, Sharma R (2022) Long short-term memory (LSTM) recurrent neural network (RNN) based traffic forecasting for intelligent transportation. *AIP Conference Proceedings* 2435(1), https://doi.org/10.1063/5.0083590.

Beaton D (2021) Testing whether distributed energy storage results in greater resilience of microgrids. Master's thesis, Department of Systems Engineering, Naval Postgraduate School, Monterey, CA, http://hdl.handle.net/10945/67104.

Craparo E, Karatas M, Singham D (2017) A robust optimization approach to hybrid microgrid operation using ensemble weather forecasts. *Applied Energy* 201, https://doi.org/10.1016/j.apenergy.2017.05.068.

Craparo E, Sprague J (2019) Integrated supply- and demand-side energy management for expeditionary environmental control. *Applied Energy* 233–234, https://doi.org/10.1016/j.apenergy.2018.09.220.

Dixon M, London J (2021) Financial forecasting with [ALPHA]–RNNs: A time series modeling approach. *Frontiers in Applied Mathematics and Statistics* 6, https://doi.org/10.3389/fams.2020.551138.

Dulău L, Bică D (2018) Optimization of generation cost in a microgrid. *Procedia Manufacturing* 22, https://doi.org/10.1016/j.promfg.2018.03.101.

Ekambaram V, Manglik K, Mukherjee S, Sajja S, Dwivedi S, Kaykar V (2020) Attention based multi-modal new product sales time-series forecasting. *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* 3110–3118, https://doi.org/10.1145/3394486.3403362.

Energy Information Administration (2022) U.S. kerosene-type jet fuel wholesale/resale price by refiners. Accessed April 16, 2023, https://www.eia.gov/dnav/pet/hist/LeafHandler.ashx?n=PET&s=EMA_EPJK_PWG_NUS_DPG&f=M.

Faisal M, Hannan M, Ker P, Hussain A, Mansor M, Blaabjerg F (2018) Review of energy storage system technologies in microgrid applications: Issues and challenges. *IEEE Access* 6:35143–35164, https://doi.org/10.1109/ACCESS.2018.2841407.

Farhangi H, Joos G (2019) *Microgrid Planning and Design: A Concise Guide*. https://doi.org/10.1002/9781119453550.ch10.

Federal Consortium for Advanced Batteries [FCAB] (2021) Executive summary: National blueprint for lithium batteries 2021–2030. Accessed November 10, 2022, https://www.energy.gov/eere/vehicles/federal-consortium-advanced-batteries-fcab.

Friedman M (1937) The use of ranks to avoid the assumption of normality implicit in the analysis of variance. *Journal of the American Statistical Association* 32(200):675–701, https://doi.org/10.2307/2279372.

Garcia K (2017) Optimization of microgrids and military remote base camps. Master's thesis, Naval Postgraduate School, Monterey, CA, http://hdl.handle.net/10945/56923.

Goodfellow I, Bengio Y, Courville A (2016) *Deep Learning*. https://www.deeplearningbook.org/.

Hernandez-Arambura C, Green T, Mugniot N (2005) Fuel consumption minimization of a microgrid. *IEEE Transactions on Industry Applications* 41(3):673–681, https://doi.org/10.1109/TIA.2005.847277.

Hua Y, Sevegnani M, Yi D, Birnie A, McAslan S (2022) Fine-grained RNN with transfer learning for energy consumption estimation on EVs. *IEEE Transactions on Industrial Informatics* 18(11), https://doi.org/10.1109/TII.2022.3143155.

Jung Y, Han S, Choi H (2021) Explaining CNN and RNN using selective layer-wise relevance propagation. *IEEE Access* 9:18670–18681, https://doi.org/10.1109/ACCESS.2021.3051171.

Lee J (2021) Optimizing fuel efficiency on isolated microgrid with energy storage system under varying loads. Master's thesis, Department of Operations Research, Naval Postgraduate School, Monterey, CA, http://hdl.handle.net/10945/67758.

Mallery J (2021) Defense installation energy resilience for changing operational requirements. Master's thesis, Department of Systems Engineering, Naval Postgraduate School, Monterey, CA, http://hdl.handle.net/10945/68733.

Mohammed A, Refaat S, Bayhan S, Abu-Rub H (2019) AC microgrid control and management strategies: Evaluation and review. *IEEE Power Electron* 6(2):18–31, https://doi.org/10.1109/MPEL.2019.2910292.

Ogunmodede O, Anderson K, Cutler D, Newman A (2021) Optimizing design and dispatch of a renewable energy system. *Applied Energy* 287, https://doi.org/10.1016/j.apenergy.2021.116527.

Parhizi S, Lotfi H, Khodaei A, Bahramirad S (2015) State of the art in research on microgrids: A review. *IEEE Access* 3:890–925, https://doi.org/10.1109/ACCESS.2015.2443119.

Riou M, Dupriez-Robin F, Grondin D, Le Loup C, Benne M, Tran Q (2021) Multi-objective optimization of autonomous microgrids with reliability consideration. *Energies (Basel)* 14(15):4466, https://doi.org/10.3390/en14154466.

Shea D (2022) Microgrids: State policies to bolster energy resilience. Accessed May 3, 2023, https://www.ncsl.org/energy/microgrids-state-policies-to-bolster-energy-resilience.

Surash J, Hughes R (2022) Developing microgrids to deliver energy resilience. U.S. Army. Accessed June 10, 2022, https://www.army.mil/article/255597/developing_microgrids_to_deliver_energy_resilience.

United States Army Acquisition Support Center (2017) Advanced Medium Mobile Power Source (AMMPS). U.S. Army. Accessed March 27, 2023, https://asc.army.mil/web/portfolio-item/cs-css-advanced-medium-mobile-power-source-ammps/.

Vergun D (2021) DoD demonstrates mobile microgrid technology. U.S. Department of Defense. Accessed November 9, 2022, https://www.defense.gov/News/News-Stories/Article/Article/2677877/dod-demonstrates-mobile-microgrid-technology.

Wen L, Zhou K, Yang S, Lu X (2019) Optimal load dispatch of community microgrid with deep learning based solar power and load forecasting. *Energy (Oxford)* 171:1053–1065, https://doi.org/10.1016/j.energy.2019.01.075.

Wood E (2022) Army to equip all bases with microgrids by 2035 as part of carbon-free electricity goal. Microgrid Knowledge. Accessed May 13, 2022, https://microgridknowledge.com/army-microgrid-climate.

Wormuth C (2022) Message from the secretary of the Army to the force. U.S. Army. Accessed February 11, 2022, https://www.army.mil/article/253814/message_from_the_secretary_of_the_army_to_the_force.

Yang Q, Dong N, Zhang J (2021) An enhanced adaptive bat algorithm for microgrid energy scheduling. *Energy (Oxford)* 232:121014, https://doi.org/10.1016/j.energy.2021.121014.

Initial Distribution List

1. Defense Technical Information Center
   Ft. Belvoir, Virginia

2. Dudley Knox Library
   Naval Postgraduate School
   Monterey, California