



Energy Usage Management and Optimization using IoT and Data-Driven Analytics for Food and Beverage (F&B) Establishments "Elgo"

Final Report

01.400 CAPSTONE 2

Project E02 - Elgo



Team Members:

Mohammed Fauzaan (1005404)

Kaveri Priya Putti (1005397)

Tan Xinlin (1004305)

Chow Zi Jie (1005261)

Sanat Khandekar (1005281)

Rohit Raghuram Murarishetti (1005398)

Visshal Natarajan (1005254)

Mentors:

Assistant Professor. Xiong Zehui

Mr. Ong Chee Beng

07.04.2024

Executive Summary

Website Link: [Elgo's Official Website](#)

Product Video: [Elgo's Product Video](#)

3D Model: [Elgo's 3D Model](#)

This project addresses a significant challenge in energy management for Small and Medium-sized Food and Beverage (F&B) establishments, aiming to enhance their energy efficiency and sustainability. Despite the considerable energy optimization efforts seen in larger entities, small to medium-sized F&B outlets lack tailored solutions, impacting their operational costs significantly. Given that the F&B sector accounts for a substantial portion of global energy consumption, the project aims to offer a comprehensive solution to bridge this gap.

Through extensive market research, including surveys and interviews, the team identified a critical need for energy management solutions to reduce operational costs without the hassle of significant installation downtimes. We believe that by delivering a tailored, efficient, and scalable energy management solution, the project not only addresses a pressing need within the F&B sector but also contributes towards broader sustainability goals.

Elgo's business model is designed around a zero-hassle installation approach with no upfront costs, offering F&B establishments the flexibility to subscribe and only pay based on the energy savings achieved. This approach, combined with a targeted market entry strategy and partnerships, positions the project to tap into a significant market opportunity with the potential for substantial revenue growth.

The proposed solution focuses on reducing operational expenses centered around Kitchen Appliances, Refrigeration units, HVAC systems, and Lighting systems by incorporating a comprehensive suite of algorithms aimed at cutting down energy consumption (HVAC and Lighting systems), maintenance costs (Kitchen Appliances and Refrigeration Units), and ensuring user comfort (HVAC and Lighting System). In this report, we put forward a scalable and cost-effective solution consisting of robust IoT hardware and custom algorithms hosted on AWS cloud that efficiently carries out the aforementioned objectives. The successful implementation and scalability of this solution could pave the way for a new standard in energy management for small to medium-sized F&B establishments, driving both economic and environmental benefits.

Contents

1	Problem Definition	1
1.1	Problem Statement	1
1.2	Market Opportunity	1
2	Business Development	1
2.1	Market Size	2
2.2	Competitor Analysis	2
2.3	Go To Market Strategy	2
3	Solution Overview	2
4	Kitchen Appliances and Refrigerator	3
4.1	Product Design	3
4.2	Functional Requirements	3
4.3	Solution Architecture	4
4.3.1	System Architecture	4
4.3.2	Cloud Architecture	4
4.4	Algorithms: Anomaly Detection	4
4.4.1	Point Anomalies - Isolation Forest (ISOF)	4
4.4.2	Collective Anomalies - Statistical Machine Learning (Gaussian Kernel)	5
4.4.3	Contextual Anomalies - LSTM-AE	6
4.5	Testing and Evaluation Methods	6
4.6	Iterative Prototyping	7
5	Smart Plug	8
5.1	Functional Requirements	8
5.2	Product Design	8
5.2.1	Internal Supply Design	8
5.2.2	Switching Element Design	9
5.2.3	Power Monitoring Design	9
5.2.4	MCU Integration and Firmware Design	10
5.2.5	PCB Layout	10
5.2.6	Enclosure Design	10
5.3	Iterative Prototyping	11
6	Lighting System	11
6.1	Product Design	11
6.2	Functional Requirements	11
6.3	Solutions Architecture	11
6.3.1	System Architecture	11
6.3.2	Cloud Architecture	12

6.3.3	Data Collection, Processing and Analysis	12
6.3.4	Action Execution	13
6.4	Algorithm: Optimal Action Selector	13
6.5	Hardware Input and Description	13
6.6	Testing and Evaluation methods	14
6.7	Iterative Prototyping	14
7	HVAC	14
7.1	Product Design & Functional Requirements	14
7.2	Solution Architecture	14
7.2.1	System Architecture	14
7.2.2	Cloud Architecture	15
7.2.3	Pipeline	15
7.3	Algorithm: Deep Reinforcement Learning	16
7.3.1	Algorithm Execution Flow	16
7.4	Testing and Evaluation methods	16
7.4.1	Simulation-Based Assessment	16
7.4.2	Empirical Validation in Operational Environments	17
7.4.3	Performance Metrics and Model Validation	17
7.5	Iterative Prototyping	17
7.5.1	Phase 1: Core Functionalities and Baseline Performance	17
7.5.2	Phase 2: Enhanced Control and Energy Efficiency	17
7.5.3	Phase 3: Advanced Optimization and User Experience	17
8	Cloud Architecture	18
9	User Experience - Web Application and Device Interface	18
9.1	Introduction to User Interface and Web Dashboard	18
9.2	Sequence diagrams	19
9.3	Software Architecture	19
9.4	Screenshots	19
10	Conclusion	20
	References	21
A	Figures	22
A.1	Business Development	22
A.2	Kitchen Appliances and Refrigerator	23
A.3	Smart Plug	27
A.4	Lighting Systems	35
A.5	HVAC	37
A.6	User Experience	41

B	Tables	48
B.1	Solution Overview	48
B.2	Kitchen Appliances and Refrigerator	48
B.3	Smart Plug	50
B.4	Lighting Systems	50
B.5	HVAC	50
C	Hardware Calculations	51
C.1	Internal DC Supply	51
C.1.1	<i>AC-side Protection Features</i>	51
C.1.2	<i>Rectification</i>	51
C.1.3	<i>Flyback Converter</i>	51
C.1.4	<i>Load Switch (Phase 2)</i>	55
C.1.5	<i>Buck Converter 3.3 V</i>	56
C.2	Switching Element	57
C.2.1	<i>TRIAC</i>	57
C.2.2	<i>Opto-MOSFET</i>	57
C.2.3	<i>LED Indicator Circuit</i>	57
C.3	Power Monitoring	57
C.3.1	<i>Potential Transformer (Phase 2)</i>	58
C.3.2	<i>Current Sensor (Phase 2)</i>	59
C.4	PCB Layout	60
D	Lighting Systems	61
D.1	Light Sensor PCB Schematic	62
E	Budget Allocation	63
F	Individual Contributions	64

List of Figures

1	Cloud Architecture	18
2	Competition Landscape Plane	22
3	Feature Matrix as Compared to Competitors	22
4	System Architecture Kitchen Appliances and Refrigerator	23
5	ISOF Decision Flow	24
6	Power Signature of Refrigerator over a 4-day period	24
7	Gaussian Machine Learning Decision Flow	25
8	LSTM-AE Decision Flow	25
9	Iterative Prototyping: Deliverables Achieved for Kitchen Appliances	26
10	Iterative Prototyping: Deliverables Achieved for Refrigerator	26
11	Smart Plug Functional Block Diagram	27
12	Phase 2 Schematic	27
13	Phase 3 Schematic	28
14	Phase 2 PCB Layout	29
15	Phase 2 PCB 3D Render	30
16	Phase 3 PCB Layout	31
17	Phase 3 PCB 3D Render	32
18	PCB with SMT components assembled	33
19	Smart Plug Enclosure, Assembled	33
20	Smart Plug Enclosure, Exploded-view	34
21	Smart Plug Enclosure under testing	34
22	System Architecture for Lighting Solution	35
23	Moving Average of THD, THB.	35
24	Mapping of State to Optimal Action.	36
25	Logical Control Flow.	36
26	Deep Reinforcement Learning Flowchart (HVAC)	37
27	Simulation Environment (HVAC)	37
28	Simulation Framework (HVAC)	38
29	Simulation Restaurant SketchUp 3D Model (HVAC)	38
30	Simulation Model (HVAC)	39
31	Baseline vs DRL Control for 1 month (HVAC)	39
32	Iterative Prototyping: Deliverables Achieved for HVAC	40
33	Sequence diagram for user registration and login	41
34	Sequence diagram for anomaly detection	42
35	Device registration and Anomaly Updates. This also ties in with the sequence described in Fig 34	43
36	Landing Page Screenshot	44
37	Login Page Screenshot	44
38	Dashboard Page Screenshot	45
39	Device Registration Screenshot	45

40	Anomalies Page Screenshot	46
41	HVAC Control Screenshot	46
42	Lighting Control Screenshot	47
43	Power Consumption Insights Screenshot	47

List of Tables

1	List of sequence diagrams	19
2	List of Screenshots	20
3	Device Overview	48
4	Functional Requirements for Anomaly Detection System	48
5	Refrigerator Anomaly Types	48
6	Relevant configuration parameters of ISOF and LSTM-AE.	49
7	ISOF model evaluation results on the REFIT Dataset.	49
8	ISOF model evaluation results on the FIKElectricity Dataset.	49
9	Gaussian Statistics Machine Learning model evaluation results on the REFIT Dataset	49
10	Gaussian Statistics Machine Learning model evaluation results on the FIKElectricity Dataset	49
11	LSTM-AE model performance metrics on the REFIT Dataset.	50
12	Functional Requirements for Lighting Systems	50
13	Power Consumption with and without Controller	50
14	Functional Requirements for HVAC System Optimization	50

1 Problem Definition

The problem was defined by researching the existing value chain of the electricity market and exploring various aspects of industrial & retail consumption to understand the potential of impacting these segments. Our investigation revealed a lack of sustainability and energy management solutions in the retail energy consumption sector, particularly for smaller enterprises like Food and Beverage (F&B) establishments. While major companies, like Arloid Automation, have optimized energy usage in large buildings, there's a notable gap for smaller players. Small and medium-sized F&B outlets, such as independent restaurants, small cafés, and chain restaurants, therefore lack an impactful energy optimization solution. Consequently, F&B outlets in Asia use 4 times more energy per m² than office buildings, affecting their profits by over 15% (Engie, 2022).

1.1 Problem Statement

Our capstone project targets the aforementioned market gap, focusing on retail and enterprise consumption, with the aim of providing accessible and tailored solutions to enhance energy efficiency and sustainability in the often overlooked F&B sector. This brings us to our problem statement:

“How can we build an impactful energy optimization solution for Small and Medium-sized F&B outlets, like independent restaurants, small cafés, and chain restaurants, to help them manage their electricity consumption?”

1.2 Market Opportunity

Research was conducted to validate the impact that our problem has on the F&B landscape today, which culminated in several key statistics that are presented as follows:

1. **Energy Usage:** The F&B sector contributes to 30% of global energy consumption (Egginton, 2022).
2. **Energy Cost:** An F&B establishment spends about 38% of its pre-tax profits on paying its energy bills (Miller and Othman, 1994). The establishments spend about 33 cents per meal to pay for their energy consumption, which scales up quickly when considering the number of meals served in the establishment (EasyorderApp, 2023).

These statistics are evidence of the undeniable impact of the F&B sector on the energy landscape and highlight the significant change our solution can kickstart.

2 Business Development

As part of the Entrepreneurship Capstone, a comprehensive business plan was ideated through the numerous workshops held by the VIE Office, pitching sessions, and mentor meetings to fit the purpose of our problem. Through this, we devised our business model, a zero-hassle installation, zero cost upfront, and zero lock-in period model for our customers, offering them the flexibility to subscribe to our services at their will. The cost model we adopted is to charge the clients 50% of the savings that we're able to achieve for them w.r.t. the energy consumption prior to deploying our solution (their energy consumption is audited before deployment).

2.1 Market Size

We needed to quantify the potential of our business by assessing the available market size for us to tap into. A total of 23 million F&B outlets worldwide (Statista, 2023) forms our total addressable market. Out of this population, 15 million outlets are restaurants, notes Balcan (2014), which forms our total obtainable market. If we target 5000 of these restaurants in our serviceable obtainable market over the course of three years, it would materialize into a potential annual revenue of S\$12 million for us. This figure was derived from an average medium-sized restaurant's monthly electricity bill of \$2000/month, which is obtained via the user survey. A 20% reduction in their costs yields S\$200 monthly revenue, given the 50-50 cost model. When multiplied to a total of 5000 outlets, this yields $5000 \times \$200 \times 12 = \$12m$ in annual revenue.

2.2 Competitor Analysis

As part of business development, it is critical to identify competitors in the space we're targeting. Our research highlighted a few companies in the US-EMEA and APAC markets offering energy optimization solutions to various target consumers. As shown in Figure 2, the only competitor in the ASEAN market targeting F&B outlets was TablePointer, which itself only deals with HVAC optimization. By further comparison of our competitors' offerings on a feature matrix (see Figure 3), we believe that our product strikes an optimal balance in features.

2.3 Go To Market Strategy

We plan to enter the market using a three-step approach:

1. Firstly, we want to acquire a few initial customers to test-bed our product and validate our claims while seeking feedback about what can be improved. Through our market validation journey, we successfully captured interest from a few outlet managers and owners whom we can work with for deployment.
2. Secondly, we would like to partner with channel partners whose extensive business networks and circles can help us spread the word about our product. We seek to tap into these networks to establish our pioneer customers who can help us spread the word to other restaurant owners in their circles.
3. Finally, we are looking at partnering with a governmental agency like the National Environment Agency (NEA), whose mission is to ensure a sustainable and clean Singapore. The opportunity to partner with an organization like NEA can help us push our mission of sustainability in the F&B space forward and also establish our solution as a standard for achieving a path toward sustainability for our target market.

3 Solution Overview

Our solution is aimed at optimizing energy levels and monitor appliance loads to help restaurant owners cut down on operational expenses. In order to achieve our goal, we target four main product suites: Kitchen Appliances; Refrigerator units; HVAC; and Lighting systems. Each solution suite is customized to suit the unique characteristics of the respective product.

- For Kitchen Appliances and Refrigeration units, our primary focus is on predictive maintenance, anomaly detection, and load monitoring.
- For HVAC systems, our main objective is to optimize the temperature levels for energy consumption and user comfort.

- For Lighting, we prioritize maintaining ambient lighting levels without compromising the quality of the restaurant space.

Additionally, we plan to incorporate the necessary hardware to collect the data required for achieving the aforementioned objectives. Table 3 summarizes the overview of devices and the reason we included them in our product suite.

4 Kitchen Appliances and Refrigerator

To assist restaurant owners in mitigating the substantial annual expenditure of \$4000 (Ser, 2024) allocated for the maintenance of kitchen appliances and refrigerators, we have engineered an anomaly detection system, which is designed to promptly identify anomalies within appliances which enables the clients to take preemptive action. Timely interventions are key to minimizing operational maintenance costs. This solution streamlines kitchen operations and empowers restaurant owners to adopt sustainability by reducing carbon footprint as well as energy.

4.1 Product Design

Our system is engineered with predictive maintenance at its core, enabling users to proactively schedule maintenance for their kitchen appliances and refrigerators. This strategy prevents operational anomalies from escalating into significant issues. Catering to a wide range of kitchen equipment and refrigeration units, including fryers, ovens, cookers, steamers, and more, our approach is designed for the diverse appliance ecosystem within restaurant environments.

To give restaurant owners an in-depth analysis for actionable insights, we developed three different algorithms, each targeting different anomaly types as classified by Chandola et al. (2009).

- **Point Anomalies:** are individual data instances that are anomalous w.r.t. the rest of the data points within a given time frame.
- **Collective Anomalies:** implies that a portion of the data, not necessarily individual is anomalous. Although some of these data points might seem normal, their occurrence in the wrong sequence can lead to a collective anomaly.
- **Contextual Anomalies:** are data points that are considered anomalous within a particular context. For instance, if we detect a portion of data that represents a device in the OFF state when, in fact, the device is supposed to be in an ON state - such instances can be classified as a contextual anomaly.

These algorithms underpin our system's predictive maintenance capabilities, alerting users to potential issues before they lead to downtime or costly repairs.

4.2 Functional Requirements

To meet the requirements of the end-user, the functional requirements have been summarized in table 4.

4.3 Solution Architecture

4.3.1 System Architecture

Figure 4 gives us a simplified version of the system architecture that is needed to enable the functional requirements of our system. The ingested data is pre-processed and sent to our streaming servers, which help buffer the data in our data lake S3 buckets. Once the buffer limit has been reached, the data will be streamed to our algorithms - namely Isolation Forest and LSTM-AE hosted on Sagemaker and Gaussian Machine Learning hosted on EC2. Anomalies detected are sent to the backend, which in turn is periodically polled by the vercel server. Users are instantly notified via email if an anomaly is detected. The user can then choose to take an action, if the user decides to take action, a service email is automatically scheduled with their maintenance service provider.

4.3.2 Cloud Architecture

The section highlighted in BLUE in Figure 1 illustrates the intricate infrastructure orchestrating the data flow within our pipeline. Power data, harvested through our smart plugs, is funneled via an MQTT broker into AWS IoT Core. AWS IoT Core, in turn, steers this stream into S3 buckets utilizing AWS Firehose. The collective anomaly detection algorithm resides on an EC2 instance, sifting through raw data from the ingest bucket, processing it, and then posting any anomalous records directly to the Vercel backend, thereby triggering an alert.

Point anomaly detection is seamlessly integrated into the Kinesis application, which enriches incoming data points with an appended column of anomaly scores. The contextual anomaly detection algorithm is operationalized through a SageMaker inference endpoint. This endpoint is invoked by an AWS Lambda function, which is triggered upon the Kinesis application's deposition of data into the processed data S3 bucket. Detected anomalies are subsequently conveyed to the Vercel backend table—powered by a PostgreSQL database—via a POST request. To mitigate isolated anomalies (point and contextual) from precipitating alerts, the backend employs a 30-minute cron job that reviews the anomalies. If the proportion of detected anomalies exceeds 80%, the system generates a new anomaly record that prompts an alert.

4.4 Algorithms: Anomaly Detection

4.4.1 Point Anomalies - Isolation Forest (ISOF)

Isolation Forest (ISOF) is an efficient algorithm for detecting point anomalies, as explored by [Zangrando et al. \(2022b\)](#). ISOF operates on the premise that anomalies are infrequent and structurally different from normal data points, which simplifies their identification and isolation. This quality makes ISOF particularly suited for scenarios like appliance monitoring, where anomalies need to be promptly detected and addressed.

Figure 5 represents the decision process of the Isolation Forest (ISOF) algorithm for anomaly detection:

Description of the Decision Flow:

1. *Construction of Isolation Trees:* ISOF starts by creating multiple isolation trees. For each tree, the data is recursively split by randomly selecting a feature and then a random split value between the maximum and minimum values of the selected feature.
2. *Path Length Calculation:* In these trees, instances are isolated by traversing the tree from the root node until a terminating node is reached. The path length, which is the number of splits required to isolate the

instance, becomes a measure of normality.

3. *Scoring Anomalies*: The shorter the path length to isolate an instance, the more likely it is to be an anomaly. This is because normal instances require more splits to be isolated, while anomalies are isolated sooner.
4. *Aggregation and Thresholding*: An anomaly score is calculated as the average path length across all trees in the forest. This anomaly score is then scaled such that the values range between [0,1], with 1 being an anomalous point and 0 being a normal point. A threshold is set to distinguish between normal and anomalous instances. Instances with a score greater than the threshold are marked as anomalies.

4.4.2 Collective Anomalies - Statistical Machine Learning (Gaussian Kernel)

As discussed in the previous sections, collective anomalies help detect abnormalities that take place over a period of time instead of flagging isolated spikes in power demands, which can occur during regular events like turning on the refrigerator. However, the algorithm described in this section is mainly on appliances with a cycle-based energy consumption- where there are periods of relatively high energy consumption followed by a period of low energy consumption (as shown in). In the suite of devices being targeted by Elgo, we aim to employ the usage of this algorithm mainly for refrigerators and other cooling or heating devices. This mainly includes refrigerators, water heaters, and freezers of most types. The normal power signature of the refrigerator used in our prototype has been shown in Figure 6.

To achieve this, the algorithm primarily makes use of Kernel Density Estimator (KDE) as described in [Hosseini et al. \(2020\)](#). Starting with a KDE, [Hosseini et al. \(2020\)](#) found that the best results were achieved with a Gaussian Kernel, and the algorithm described in this section will outline a decision flow that traces the power data from the input to actionable output. Figure 7 shows a high-level overview of the functioning of this algorithm.

Like the other algorithms for kitchen appliances and refrigerators, it is important to note that the algorithm involves an initial phase for normal operation data collection. The term "normal operation" may entail some ambiguity and merits further discussion. [Hosseini et al. \(2020\)](#) describes mainly two types of anomalies in the context of cooling appliances like refrigerators. Table 5 summarizes the types of anomalies in the context of refrigerators and freezers. Normal Operation refers to scenarios that fall outside of those described in 5 and have the usual energy signature.

Description of the Decision Flow:

- *Training Phase*: As per [Hosseini et al. \(2020\)](#), the recommended duration of normal operation tested across multiple devices is 113 cycles for refrigerators, which equates to approximately 3 days of normal usage. The normal operation data collection phase is conducted by running a separate training pipeline, and the data is sent to the S3 bucket. We shall call the normal operation data training data. This is set as one of the initial states.
- *Cycle Detection and Gaussian Detection Initial State*: The CycleDetection node in the Algorithm (as shown in Figure 7) trains the KNN model on the training data described previously, and its objective is to learn the subset of data describing the ON state and the OFF state. Further, the same data is used to calculate the mean and the standard deviation using the GaussianCalculator node. This is also set at the same time as the initial state.

- *Anomaly Detection:* Once the training phase is done, the initial state of the system is set, and the system starts to read live power data from the data stream based on some buffer interval defined in AWS Kinesis Firehose. The labeling of the live data stream as ON or OFF is done on the fly, based on the trained KNN model. Anomaly is detected here in two main ways.
 1. Once the current cycle is over, we will label the entire cycle as either ANOMALOUS or NORMAL based on the three-sigma rule of thumb in statistical modeling.
 2. If the current cycle is going on for too long, then it will cross the cycle length threshold and be classified as ANOMALOUS.

Anomalous cycle data is sent via an API call to the Postgres database used to serve the back-end for our UI and is discarded from our collective algorithm immediately afterward. Any live, normal cycle data is used to update the mean and the standard deviation in the GaussianCalculator node, so further calculations make use of up-to-date information on device performance.

4.4.3 Contextual Anomalies - LSTM-AE

Contextual anomalies within a dataset are critical in assessing the reliability of quasi-periodic energy-consuming appliances, such as refrigerators. These anomalies manifest when the observed data deviates from the expected pattern under specific conditions or contexts. Unlike collective anomalies that consider broad data deviations, contextual anomalies focus on discrepancies occurring within a particular context or during specific operational states of an appliance.

The Long Short-Term Memory AutoEncoder (LSTM-AE) has been identified as a particularly suitable algorithm for detecting such anomalies. By learning to reconstruct the normal operational cycles of an appliance, the LSTM-AE can effectively flag data points or sequences that are inconsistent with the learned behavior, signifying a potential contextual anomaly. The effectiveness of LSTM-AE in anomaly detection in energy consumption data series is highlighted by [Zangrando et al. \(2022a\)](#), which indicates its prowess in identifying anomalous consumption patterns that may allude to equipment malfunction or inefficiencies. Figure 8 represents the decision process of the LSTM-AE algorithm for anomaly detection:

Description of the LSTM-AE Decision Flow:

1. *Normalization and Sequence Generation:* Incoming data points, denoted as $[y_1, y_2, y_3]$, are normalized and structured into sequences for LSTM processing, ensuring consistency in the data format.
2. *LSTM Processing and Reconstruction:* The LSTM layers process the sequences to learn and reconstruct the normal data patterns, establishing a model of typical system behavior.
3. *Anomaly Detection:* The reconstruction error is calculated by comparing the reconstructed sequences against the actual data. Anomalies are detected when this error surpasses a defined threshold, indicating deviations from the learned patterns and signaling potential issues.

4.5 Testing and Evaluation Methods

Algorithm Configuration Parameters During our testing and evaluation process, specific configuration parameters (refer to Table 6) were utilized for the Isolation Forest (ISOF) and Long Short-Term Memory AutoEn-

coder (LSTM-AE) algorithms to optimize their performance for anomaly detection. Both models were trained on the REFIT refrigerator dataset to evaluate their efficiency in detecting anomalies in refrigerators. Additionally, the ISOF model was adapted and tested on the FIKElectricity dataset for kitchen appliance data.

Evaluation Metrics For anomaly detection in kitchen appliances, evaluating the precision, recall, and F1 score is vital. Precision is significant for reducing false positives, which, if high, could result in unnecessary maintenance actions. Recall ensures that the model detects as many real anomalies as possible, avoiding overlooked faults. The F1 score is especially critical, as it balances precision and recall, providing a singular metric for accuracy. In our evaluations, the ISOF model achieved the highest F1 score of 0.773 for a 3-week period on the REFIT dataset (Murray et al., 2015) (see Table 7), indicating a robust capability to detect anomalies with minimal error. The LSTM-AE model excelled with an F1 score of 0.79 over a 2-week span on the same dataset (refer to Table 11), demonstrating its effectiveness in identifying contextual anomalies. Furthermore, when applied to the FIKE dataset (Pereira et al., 2023), the ISOF model's best performance was an F1 score of 0.804, suggesting it's well-suited for varied appliance anomaly detection (Table 8 details this evaluation). Similarly, to validate the findings in Chandola et al. (2009) for collective anomalies, we constructed an evaluation matrix to see if 3 days were indeed the optimal period for training against both the REFIT Dataset as well as the FIKElectricity Dataset. The evaluation scores have been highlighted in 9 and 10. We see that the score increases consistently for the first three days of training and eventually remains stable henceforth. This suggests that a minimum of three days of training is required to obtain stable values of mean and standard deviation for collective anomalies. These high F1 scores underline the proficiency of the models in ensuring reliable monitoring and operational efficiency of kitchen appliances. We have also added graphs showcasing how point anomalies and contextual anomalies were detected in the day-to-day operation of refrigerators.

4.6 Iterative Prototyping

Iterative prototyping was a cornerstone in the development of our anomaly detection system for kitchen appliances. Across three distinct phases, we progressively enhanced our system's capabilities, continuously incorporating feedback to meet and exceed the predefined functional requirements.

During Phase 1, our focus was on establishing robust data collection mechanisms. We implemented smart plug integrations and began aggregating sensor data, laying the groundwork for real-time data analysis. At this stage, we introduced unsupervised learning algorithms for preliminary anomaly detection, fulfilling our first requirement (see Table 4).

Phase 2 saw the introduction of supervised learning algorithms, enabling the system to identify a wider range of anomalies with improved precision. This phase also marked the inception of our alert system, which provided users with timely notifications for detected anomalies, thereby meeting our third functional requirement.

In the final phase, Phase 3, we refined our inference models and integrated the system with a comprehensive dashboard for enhanced user experience. We also automated the process of booking maintenance contacts, further exceeding our functional requirements by not only alerting users of potential issues but also facilitating immediate action. The iterative approach ensured each feature was developed to its fullest potential, which can be reviewed in Figures 9 and 10. With each iteration, the system became more adept at not just identifying anomalies but also providing actionable insights to the users, marking a significant stride beyond the initial functional requirements.

5 Smart Plug

5.1 Functional Requirements

The Elgo smart plug serves as a hardware platform for wirelessly controlled switching of connected appliance loads and power monitoring. Phases 2 and 3 of the development cycle of the smart plug built upon the earlier goals outlined in the interim report. Figure 11 outlines the functional blocks required to achieve this hardware implementation.

1. **Internal Power Supply:** 230V AC mains voltage is rectified and chopped into low-voltage DC via the use of a switched-mode power supply (SMPS). It outputs a regulated 5V and 3.3V V_{cc} supply for various analog circuitries such as the microcontroller (MCU) and current and voltage sensing hardware.
2. **Voltage and Current Sensing:** AC voltage and current must be measured separately before it can be computed as the power consumption of the load. A potential transformer is used to measure the mains voltage while a hall-effect current sensor measures the current drawn through the external load in the Phase 2 design. The Phase 3 design integrates both functions directly in a single IC package.
3. **Load Switching:** A TRIAC provides switching of the AC load, *i.e.* the connected kitchen appliance, preventing unwanted quiescent power consumption during off state.
4. **MCU Integration:** An ESP32C6 module provides control action over the load switching and samples the analog sensor data through its ADC channels (Phase 2), or communicates with the power monitoring IC via a serial SPI protocol (Phase 3).

5.2 Product Design

5.2.1 Internal Supply Design

It was found that the line-frequency transformer used previously in the Phase I demo would result in unacceptable physical bulk while providing meager wattage to the internal power supply. This motivated a substantial revision of the power supply design and resulted in the selection of the flyback converter SMPS topology. The flyback converter similarly fulfills our safety requirements of galvanic isolation between its primary and secondary side, eliminating the possibility of a disastrous 'fail-short' event energizing the low-voltage DC circuitry at 230V potential. Meanwhile, it achieves the size reduction that we require since the flyback transformer element can be downsized significantly as it operates in the nominal range of 100 kHz.

The flyback converter topology is achieved by rectifying the 230V mains voltage and then chopping up the rectified voltage via a flyback controller IC, which switches the internal power MOSFET at a specific duty cycle. This duty cycle D_1 is given by:

$$\frac{V_o}{V_{in}} = \frac{N_s}{N_p} \cdot \frac{D_1}{1 - D_1}$$

Assuming a current draw not exceeding 0.5A for our 3.3V microcontroller gives us with a 1.65W power consumption just for the MCU. When other analog circuitry is included, this brings the total conservative estimate to 3W. Using this target wattage, the flyback transformer and IC was suitably sized, where a duty cycle of 17.9% achieves a 5V flyback output.

Switching the power MOSFET at such high frequency presents a serious issue of electromagnetic interference (EMI) noise which must be adequately tackled to prevent the entire circuit from becoming inoperable. The input pi filter serves as a first line of defence since it provides a low-impedance path to GND to the high-frequency noise while doubling as a differential-mode choke. In addition to the standard input filter, an RCD clamp is used to suppress dv/dt transients when uncoupled leakage inductance stores energy in the magnetic field but does not couple it to the secondary output. This arrangement is paralleled with an RC snubber to dampen the ringing effects caused by the switching noise. Additionally, a C snubber is utilized at the IC drain and source pins (of the internal MOSFET) to suppress transients without high power dissipation of a resistive element.

The 5V flyback output is ideally regulated, but its behavior cannot be predicted during transient operation. The presence of unstable 5V Vcc voltage is unacceptable as it risks damage to other analog circuitry, hence a load switch operating on active-low logic is utilised so that during transient startup, its EN pin is pulled up high and deactivated. Simultaneously, the 3.3V buck converter IC, which is immune to a variable input voltage (within rated range) will be soft-started in this transient period, eventually powering up the MCU to pull down the EN pin and switch on the 5V Vcc supply. For both Vcc supplies, the use of appropriate output filter capacitors helps reduce undesirable output ripple.

Apart from EMI noise reduction, care has to be taken to protect against over-voltage events such as transient spikes and ESD strikes. Various shunt elements such as TVS diodes and MLV varistors provide a clamping effect, especially at sensitive pins and other ESD touch-points.

5.2.2 Switching Element Design

A TRIAC is used as the primary switching element. Because our development considers only the unidirectional Type-G plug used in Singapore, only a single-pole single-throw switching element is required on the live net. The TRIAC gate is triggered by an opto-MOSFET, providing a trigger signal that enables the TRIAC to remain in conduction until the zero-crossing of every half-wave where the load current falls below the TRIAC latching current threshold. Thus for the ON state, the opto-MOSFET is triggered with a digital high signal from the ESP32 GPIO pin.

The presence of dV/dt transients typical in TRIAC circuits is mitigated by an RC snubber network to avoid spurious triggering, which is undesirable for the end user.

5.2.3 Power Monitoring Design

Voltage monitoring in the Phase 2 design is achieved using a line frequency transformer with an appropriate voltage divider. However, because the AC voltage is only scaled down proportionally, it enters a negative region every other half-cycle. For analog measurement, a negative voltage at the ESP32 analog-to-digital converter (ADC) is unacceptable because it risks damaging the ADC pin, hence a positive bias of $\frac{V_{cc}}{2}$ is applied using a rail-to-rail op-amp to provide the necessary offset so that the entire peak-to-peak waveform is within the full-scale range of the ADC.

Current monitoring is simplified with a hall-effect current sensor IC which outputs an analog signal proportional to the AC load current based on its calibration curve. This output is passed through a voltage divider for compatibility with the ESP32 ADC.

In the Phase 3 design, the adoption of the ACS37800 power monitoring IC allows us to phase out the bulky

transformer and the current sensor IC in favor of a single-IC package. Like before, the current is sensed through the internal hall-effect sensor, while the voltage is sensed through an external voltage divider and fed to the IC pins. Both are sampled by the internal 24-bit ADC, providing a significant increase in resolution compared to the 12-bit ADC of the ESP32. Finally, voltage and current parameters, stored as registers on the internal EEPROM, are sent to the MCU via an SPI communication interface.

5.2.4 MCU Integration and Firmware Design

The MCU provides various control and communication functionality for the smart plug hardware and communicates with the cloud via its integral Wi-Fi antenna. These core functions include:

- **GPIO11:** provides a digital high signal when an ON control signal is provided to it via the cloud dashboard. Hence, GPIO11 activates the opto-MOSFET input, which in turn drives the TRIAC gate to conduction so as to switch on the AC load. Simultaneously, it also controls the gate voltage of the switching MOSFET used in the LED indicator circuit.
- **GPIO10 (Phase 2-only):** the initial low state provided by the internal pull-down is pulled high 0.5s after system startup, which activates the gate of the pull-down MOSFET Q4 to turn on the EN pin of the active-low load switch. This activates the 5V Vcc after the transient response of the flyback output has stabilized.
- **GPIO0 and GPIO1 (Phase 2-only):** ADC channels used for sampling the analog output of the voltage and current sensors.
- **MOSI, MISO, CLK, and CS (Phase 3-only):** SPI communication pins used for communication with the power monitoring IC. Here, the MCU is the master device, while the ACS37800 power-monitoring IC is the peripheral.

5.2.5 PCB Layout

Careful planning of the PCB layout is required to ensure both the safe and proper functioning of the circuit.

- **Trace Widths, Clearances, and Creepages:** Strict control of these parameters include wide traces for high-current nets, namely the current paths for the AC load current; and sufficient trace clearances and creepages for high-voltage nets, *i.e.* AC and flyback primary high-voltage traces.
- **Routing Design:** Careful consideration of routing by minimizing current loops to reduce EMI radiation and routing sensitive feedback circuits away from noisy traces to avoid crosstalk.
- **GND Layers:** GND copper pours for GND1, GND2, and output GND are arranged in star connection where possible to avoid ground loops.

5.2.6 Enclosure Design

The smart plug enclosure was designed with the following factors in mind:

- **Size:** The internal components were arranged to maximize compactness, with priority given towards reducing the width to avoid inconvenience from overlapping into adjacent sockets.

- **Safety:** Sufficient gaps and insulating spacers must be included between exposed conductive surfaces within the smart plug to avoid risks of short-circuiting.
- **Load distribution:** The design must consider areas of higher load during normal use, such as compression and tension from plugging in and unplugging appliances and shear from the weight of the plug. Components must be positioned and supported such that strain is minimized.

5.3 Iterative Prototyping

The resultant schematics for the Phase 2 and Phase 3 circuits can be found in Figures 12 and 13. In Phase 2, we achieved the uprating of the components required to handle the typical load appliances found in the restaurant. The key focus of the succeeding version, Phase 3 version, was to pursue increasing miniaturization as well as mechanical design integration to use with actual plug pins and receptacles - as observed, the size footprint has been substantially reduced.

The PCB layouts for the respective Phase 2 and 3 versions are designed as shown in Figures 14 and 16. 3D renders of these PCB designs can also be found in Figures 15 and 17. The plug enclosures are shown in Figures 19 and 20.

After the components were individually hand-soldered, mechanical integration and assembly with Type G plug pins (male connectors) and receptacles (female connectors) within the 3D-printed enclosure yielded the final product as shown in Figures 18 and 21.

6 Lighting System

In the realm of culinary spaces, lighting is the orchestrator of ambience and enabler of functionality, setting the stage for memorable dining experiences and ensuring efficient work spaces in the kitchen. However, it also stands as a significant consumer of energy, typically comprising 13% of an average restaurant's total energy usage.

6.1 Product Design

The challenge lies in finding solutions that balance occupant comfort with the imperative of energy efficiency. The product will be a self-sufficient, independent system controlled by an ESP32 microcontroller. It will take in data on environmental parameters and determine the light levels required to attain the goal of occupant comfort and minimizing energy redundancy.

6.2 Functional Requirements

To meet the requirements of the end-user, the functional requirements have been summarized in table 12.

6.3 Solutions Architecture

6.3.1 System Architecture

The proposed solution adopts a straightforward system architecture wherein each light fixture incorporates an independent controller. These controllers possess the capability to autonomously and continuously analyze environmental data and make informed decisions regarding optimal actions. The solution also attempts to learn the

room occupants' average personal preference and take it into account when making decisions about the optimal lighting conditions of the room.

In Figure 22, we see the high-level overview of the system architecture. The system is led by the controller, which sends optimal actions to the switch via an MQTT REST API. The optimal action to take is determined by the state determination and optimal action selection algorithm. The algorithm, which is straightforward, will be explained in upcoming sections. Sensors serve as the interface between the physical environment and the microcontroller unit. The sensors collect data on 3 key parameters- current switch position, occupancy of the room, and current ambient light levels.

The current switch position is queried from the MQTT cloud server, where the state of the lighting (luminosity, power, occupancy) is tracked. Since the lighting is only controlled via MQTT APIs, the state can be effectively tracked during changes to light levels, as they can be persisted on a database maintained in the ESP32 file system. Occupancy is a tricky parameter to accurately track. Traditional methods of occupancy detection use ultrasonic motion detectors. However, the drawback to these methods is that motion detection does not cover the edge case when the occupant is motionless. This may cause the lighting system to falsely assume vacancy and send a signal to switch off the lights. To this end, we propose a Bluetooth-based solution. We use Bluetooth Low Energy (BLE) to establish lightweight, energy-efficient connections with occupants' smartphones. These smartphones can act as peripheral devices and emit periodic data signals to advertise themselves to a scanner. A scanner is then set up on the ESP32 microcontroller to detect these devices. A room is deduced to be occupied when a peripheral device is detected. Lastly, the ambient light levels of the room are detected using a light sensor. The initial design used a light-dependent resistor (LDR). Photo-transistors were later preferred due to their superior accuracy in detecting changes in ambient light levels. This photo-transistor is connected via the analog channel on the microcontroller.

6.3.2 Cloud Architecture

The cloud architecture is explained by the section highlighted in GREEN in Figure 1. The ESP32 maintains communication with the AWS EC2 Lighting Control Server, which interacts with and passes optimal light state messages directly to the smart plug via an MQTT broker. The user can also choose to override the current state from the frontend dashboard, which directly communicates the change to both the ESP32 and the lighting smart plug through the AWS EC2 server. The AWS EC2 server sets up a connection with the MQTT topics subscribed by the smart plug allowing it to send and receive data.

6.3.3 Data Collection, Processing and Analysis

Data collection occurs through sensors measuring three key parameters. As mentioned earlier, the switch position of the light, the occupancy status of the room and the ambient light levels within the room. A parser component is responsible for sanitizing the raw data obtained from sensors, ensuring its accuracy and reliability. Subsequently, the processed data is fed into an algorithm designed to evaluate the current state of the system and determine the most appropriate course of action. This algorithm relies on predefined values such as the threshold for darkness (THD) and the threshold for brightness (THB) to ascertain the prevailing conditions.

In Figure 23, a moving average of the THD and THB is plotted. The THD for each day is the light level at which the lights are first turned on. This indicates that any light level below the THD is too dark for the occupant. The same goes for the THB, which is the light level at which the occupant switches the lights off. This is an indication

that any light level above the THB is too bright for the occupant. We utilize the THB and THD in determining the state of the environment. Below the THD, the state is "Dark"; above the THB, the state is "Bright". Anything in between is "Comfort". The moving average is the method we use to learn the occupants' average preference for light levels over time.

6.3.4 Action Execution

Upon completion of the analysis, the controller implements the chosen optimal action at the switch, thereby influencing the lighting conditions in accordance with the determined requirements.

The code snippet 1 calls MQTT REST APIs for controlling Shelly Plug, which controls the lighting system.

This function `execute_optimal_action` is designed to execute the optimal action for controlling the lighting system using a Shelly plug. It formats the switch position and optimal action to two decimal places and constructs a URL string with these values. The URL includes parameters for the current light level (`lightLevel_current`), the Shelly plug ID (`plugID`), and the previous light level (`lightLevel_prev`). This URL is then used to perform an HTTP request to log the light level. The HTTP client configuration is set up with the URL, and the request is executed synchronously with a timeout of 10 seconds. This function can be integrated into an IoT system to control lighting levels based on optimal actions received from MQTT messages.

6.4 Algorithm: Optimal Action Selector

The algorithm is the central logical unit that provides decisions to the controller. At its core, the algorithm is a mapping of the prevailing state to the optimal action to be executed. The system's primary goal is to cut down on redundant electricity usage while ensuring the room occupant is comfortable. As such, the optimal actions keep the light level within the comfort level at all times, turning down the brightness of the lights when the ambient light is too bright and vice versa. Figure 24 summarizes the mapping between state and optimal action.

To summarize the system, the logical flow can be seen in Figure 25. The key takeaway from this diagram is that at a high level, the ESP32 runs a continuous loop, determining the state through recorded sensor readings, selects the optimal action, and sends it to the Shelly Plug to control the lighting.

6.5 Hardware Input and Description

An NPN phototransistor was used as the light-sensing element. Unlike other photosensitive elements like light-dependent resistors (LDR), P-N junction-based sensors provide faster response times to changing light intensity and higher accuracy. When used in a common-collector configuration alongside a fixed resistance at the emitter, the varying collector current can be interpreted as an analog voltage signal, which is sampled at the ADC channel of an ESP32-C6 MCU.

The deployment of the sensor package in restaurants requires a battery power supply that allows it to operate standalone. This is achieved with a nominal 6V supply of 2x CR2025/CR2032 Li-ion button cells. A buck converter is used as a switching power supply to step down the battery voltage and provide a regulated 3.3V supply to the attached ESP32-C6 MCU module. With under-voltage lock-out functionality specified at 4V, the batteries would be nearly fully discharged before the supply shuts down automatically. A basic PMOS reverse polarity protection circuit prevents catastrophic destruction of the circuit in case of user error during the insertion

of batteries. This template power supply design can be readily adapted to any other sensor configuration for our IoT suite.

6.6 Testing and Evaluation methods

In our testing phase, we left the lighting system controller for 3 days. We logged the light levels at which lights were turned on and off. The average obtained from this testing phase was used to set the THD and THB. These were then utilized in the state determiner. With the microcontroller equipped with the learnt THD and THB, the system was put to the test, and the lights were observed to turn to the correct setting in different environment states.

To evaluate the model, we operated the lights in a controlled space for a period of 3, 6, and 8 days. After this, we continued operating the lights, but this time connected to our lighting controller and smart plug for the same time periods. After evaluating the lighting system solution, we discovered that our solution has the ability to decrease power consumption by up to 50%. The results of this evaluation have been summarized in Table 13.

6.7 Iterative Prototyping

Our solution underwent iterative prototyping, refining its capabilities across phases. Initially, we established robust communication and data collection mechanisms, utilizing MQTT APIs for optimal action execution and deploying sensors for environmental data collection. Then we enhanced occupancy detection with a Bluetooth-based solution and improved light sensing accuracy with photo-transistors. Each iteration strengthened the system, ensuring it surpassed initial requirements by providing actionable insights to users. Lastly, we tied up the individual parameter components and integrated it with the optimal action selector algorithm.

7 HVAC

7.1 Product Design & Functional Requirements

The HVAC optimization system is engineered to incorporate a Reinforcement Learning (RL) model, specifically designed to predict and adapt to fluctuating environmental and occupancy conditions. The system's core objectives are to maximize energy efficiency and maintain optimal comfort levels. Table 14 enumerates the essential functional requirements that guide the system's design and operational logic. Our solution adapts the Deep RL model introduced in [Wei et al. \(2017\)](#), for our restaurant use case.

7.2 Solution Architecture

7.2.1 System Architecture

1. **Design:** The system architecture for our solution is initiated by defining the parameters and requirements of the restaurant's HVAC system (including weather conditions), where we developed and configured the spatial layout of the restaurant by 3D modeling using SketchUp. This information was used to create a virtual model within EnergyPlus ([EnergyPlus, 2017](#)), an open-source HVAC simulator.
2. **Simulation:** Using the restaurant's virtual model, we simulated various scenarios by adjusting parameters in the IDF (Input Data File), which dictate the restaurant's characteristics and system behaviors. For example, we were able to tweak the time periods, temperature, humidity, weather conditions, cooling,

and heating variables. These simulations served as a platform for initial testing, allowing us to generate comprehensive datasets that capture the complexities of real-world restaurant operations without the costs and constraints of direct experimentation.

3. **Training:** The Deep Reinforcement Learning (DRL) agent is trained using the generated simulation data, adopting a trial-and-error learning strategy to develop an understanding of the HVAC system's response to the control strategies. The agent's policy evolves over a period of time, maximizing a reward function that aligns energy consumption with desired comfort levels. We have trained the DRL agent with different time periods ranging from 1 day to 1 year.
4. **Deployment:** After training, we deployed the refined DRL model to control the actual HVAC system. The deployed model continuously receives real-time data from the environment, such as temperature and humidity to make informed decisions about adjusting the system's operation.

7.2.2 Cloud Architecture

The section highlighted in **ORANGE** in Figure 1 is our HVAC cloud pipeline, which is designed for seamless remote operation and temperature regulation. Environmental data from HVAC devices are first transmitted to the IoT Core, which serves as the central data hub. From there, a Lambda function is triggered, processing the data and sending it to Amazon SageMaker, the inference point where our trained Deep Reinforcement Learning (DRL) model resides. The model evaluates the data and infers optimal HVAC settings, which are then communicated back through the IoT Core to the devices for action. In parallel, the Lambda function also posts updates to the frontend system, allowing for user interaction and real-time system status monitoring. This architecture encapsulates the entire loop—from data collection to model inference and actuation—ensuring that our simulation-trained model effectively dictates the real-world operation of the HVAC components.

7.2.3 Pipeline

The pipeline consists of the following specific components:

1. **Data Collection:** Sensors installed throughout the restaurant gather real-time data on indoor temperature, outdoor temperature, and outdoor humidity.
2. **Agent Decision-Making:** The DRL agent processes the monitored data to select the most appropriate actions for setting the heating and cooling setpoints for temperature control to maintain comfort while optimizing energy use.
3. **Action Implementation:** Commands are sent to the HVAC system to adjust settings such as temperature levels and on/off cycles.
4. **Feedback Loop:** System responses to the DRL agent's actions are monitored to inform future decisions with rewards and penalties, closing the loop for continuous learning and improvement.
5. **Performance Evaluation:** The system's performance, pre- and post-deployment of the DRL model, are compared to assess improvements in energy consumption and occupant comfort.

This pipeline ensures that our HVAC system remains responsive to the dynamic needs of the restaurant environment, optimizing energy usage while maintaining a comfortable atmosphere for the restaurant’s customers and making it hands-free and hassle-free for the restaurant owner/manager as well.

7.3 Algorithm: Deep Reinforcement Learning

1. **Objective:** The model aims to minimize energy consumption while maintaining temperature within specific customer comfort limits in a restaurant environment.
2. **State Space:** The state space encompasses a vector of environmental factors, including:
 - Outdoor Temperature
 - Indoor Temperature
 - Outdoor Humidity
3. **Action Space:** Actions within the model are decisions made by the agent to adjust the heating and cooling setpoints for optimal temperature control.
4. **Reward Function:** The reward is calculated based on energy consumption efficiency. A penalty is applied if the indoor temperature deviates from the set comfort thresholds.
5. **Hyperparameters:** The algorithm uses several hyperparameters, which are tuned to guide the learning process. These include learning rates, optimization epochs, and exploration decay rates pertinent to the Proximal Policy Optimization (PPO) algorithm.

7.3.1 Algorithm Execution Flow

The algorithm follows a cyclical process where the current state of the environment is input into the DRL agent. The agent then predicts an action that is expected to lead to an optimal outcome. Following the action execution, the environment provides a new state and a reward based on the result of the action. The DRL agent uses this feedback to update its policy, aiming to maximize the cumulative reward over time, which corresponds to energy savings and temperature maintenance within comfort limits. A detailed reference to the algorithmic structure can be found in figure [26](#).

7.4 Testing and Evaluation methods

To rigorously assess the performance of the Deep Reinforcement Learning (DRL) model for HVAC system optimization, our evaluation protocol encompasses both simulation-based analysis and real-world empirical testing, focusing on energy efficiency and maintaining thermal comfort within specified parameters.

7.4.1 Simulation-Based Assessment

We employ the EnergyPlus simulation environment to emulate a controlled HVAC system. Our evaluations include:

1. Comparison against a **Baseline Performance** model to quantify energy efficiency improvements.
2. **Stress Testing** under peak load conditions to evaluate the resilience of the DRL-controlled system.

3. **Variability Tests** simulating energy supply fluctuations to measure the robustness of the DRL optimization strategy.

7.4.2 Empirical Validation in Operational Environments

Following simulations, we transition to real-world deployment of the DRL model. This involves:

1. Implementing **Controlled Environment Tests** to validate the model’s ability to manage real-world HVAC systems.
2. Collecting **User Comfort Feedback** to qualitatively assess the impact of the DRL system on occupant comfort levels.

Through Amazon SageMaker, we utilize custom Docker containers, S3 storage for model artifacts, and MXNet for model deployment and inference. Our RL agent’s efficacy is substantiated through comprehensive CI (Continuous Integration) tests across multiple AWS regions, ensuring the model’s robustness and transferability.

7.4.3 Performance Metrics and Model Validation

The primary metric for quantifying the DRL model’s performance is **Energy Consumption Efficiency**, derived by comparing the energy usage to a pre-established baseline and recorded during peak operational scenarios.

7.5 Iterative Prototyping

The iterative prototyping of the HVAC optimization system was executed in three successive phases as detailed further in Figure 32, methodically expanding the system’s capabilities beyond the initial functional requirements outlined in Table 14.

7.5.1 Phase 1: Core Functionalities and Baseline Performance

Phase 1 established the core functionalities, with data acquisition automated and initial deployment of DRL-based adaptive control. This phase prioritized the integration of basic energy management protocols and developed a functional user interface for system monitoring.

7.5.2 Phase 2: Enhanced Control and Energy Efficiency

Building on the initial setup, Phase 2 introduced sophisticated data processing algorithms, which led to more refined adaptive control actions and significantly better energy management. The user interface was upgraded for greater usability, focusing on actionable insights and easier system overrides.

7.5.3 Phase 3: Advanced Optimization and User Experience

The final phase iterated through the optimization process with an optimized DRL algorithm, surpassing our expected performance criteria.

Based on the evaluation data, the model undergoes iterative refinements with increasing training durations, leading to significant improvements in energy savings—**13.68%**, **17.52%**, and **19.71%** for one month, three months, and one year of training, respectively.

In each phase, the system not only met but also exceeded the stipulated functional requirements, achieving superior energy savings, user comfort, and interaction simplicity. These enhancements were not just incremental but marked by significant leaps in performance and user-centric design, validating the system’s readiness for a broad deployment.

8 Cloud Architecture

The AWS architecture plays an integral role in the data processing lifecycle of our system, acting as the critical downstream processing layer that rigorously handles and analyzes the data streamed from IoT devices. Engineered to receive, store, and process high-velocity data, this architecture leverages a suite of AWS services to perform real-time analytics and execute advanced machine learning models. The following outline elucidates the interconnected components and the sophisticated workflows that enable our system to deliver prompt and intelligent responses based on the ingested data streams. The pipeline integrates streaming data from HVAC, kitchen appliances, refrigerators, and lighting systems with the front end to allow the users to view and interact with these systems. Below is a detailed outline of the entire data pipeline:

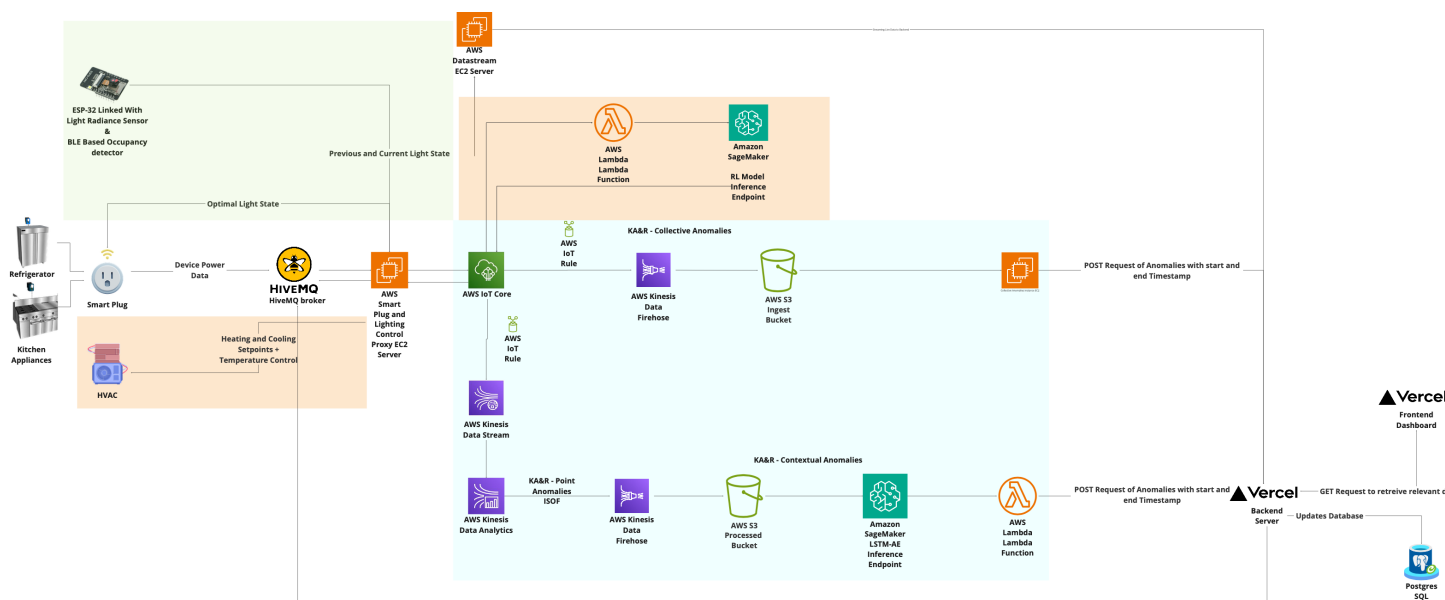


Figure 1: Cloud Architecture

9 User Experience - Web Application and Device Interface

9.1 Introduction to User Interface and Web Dashboard

- **Purpose:** Our web dashboard is where the sophisticated software architecture and algorithms come integrated with our hardware and controllers developed for controlling devices in restaurants.
- **Features:** A user’s first introduction to the problem we solve, our business model, and our pricing plan happens through the web dashboard’s landing page. After reviewing this page, they can proceed to access the features below.

The **features** are Registering and Logging in, Registering a new device, Viewing Anomaly & Power Consumption Insights, Reporting a need for maintenance from Anomaly Review, Optimizing Lighting Control, and Optimizing HVAC Control.

- **Access to Infrastructure Set Up:** To unlock insights into energy optimization, a user must inform the Elgo team about their intention to register with us, after which we can automatically link the insights from their devices to their account. This is because this step involves ML training that involves our backend interface, which is only accessible to our team.

9.2 Sequence diagrams

Refer to Table 1 for a detailed list of sequence diagrams for each scenario.

Sequence Name	Purpose	Reference
User Login and Registration	This is the entry point for all users to maintain their devices and get access to their dashboards.	Figure 33
Anomaly Detection	The sequence of actions that leads to the display of anomalies on the front end.	Figure 34
Device Registration	Once logged in, the user needs to register the devices manually in order to be able to view their energy signatures.	Figure 35

Table 1: [List of sequence diagrams](#)

9.3 Software Architecture

The architecture that enables the front end to access and present all the information to the user can be found in Figure 1. The flow highlighted in Figure 1 allows users to interact and view information about their establishments' device energy signatures.

9.4 Screenshots

Refer to Table 2 for a detailed list of screenshots for each action and feature executable in the web dashboard.

Description	Purpose	Reference
Landing Page	This is the entry point for all users to our website and get access to their dashboards.	Figure 36
User Login and Registration	A user can log in/register to access their web dashboard.	Figure 37
Dashboard Interface	Upon login, a user can visit their dashboard and see a snapshot of the energy status as of that moment.	Figure 38
Device Registration	Once logged in, the user needs to register the devices manually in order to be able to view their energy signatures.	Figure 39
Anomalies Detected	Once logged in, the user can click on the devices manually to view the anomalies detected in their usage.	Figure 40
HVAC Control	Once logged in, the user can view the optimal temperature set point and override it if they wish.	Figure 41
Lighting Control	Once logged in, the user can view the optimal lighting set point and override it if they wish.	Figure 42
Power Consumption Insights	Once logged in, users can toggle between their different registered devices to view a snapshot of their power consumption over 72 hours.	Figure 43

Table 2: [List of Screenshots](#)

10 Conclusion

Our project is aimed at developing a comprehensive solution to meet the energy optimization needs of small and medium-sized F&B establishments. As demonstrated through thorough evaluations, the project not only reduces energy costs but also paves the way for substantial cost reductions and environmental benefits.

The project successfully demonstrated a significant decrease in energy consumption, achieving close to 20% reduction in a home-grown testbed. This outcome validates the efficacy of our solution and underscores the potential for wide-scale implementation across the industry. By significantly lowering energy consumption, our project, Elgo, contributes to the broader goals of environmental sustainability and carbon footprint reduction.

References

- (2024). Benchmarking qsr refrigeration costs.
- Balcan, A. (2014). How many restaurants are in the world? <https://thewebminer.com/blog/how-many-restaurants-are-in-the-world/>.
- Chandola, V., Banerjee, A., and Kumar, V. (2009). Anomaly detection: A survey. *ACM Comput. Surv.*, 41(3).
- EasyorderApp (2023). How to reduce the energy consumption in your food and beverage business. <https://www.easyorderapp.com/en/blog/how-to-reduce-the-energy-consumption-in-your-food-and-beverage-business>.
- Egginton, T. (2022). An energy efficient future for the food and beverage sector: By brith isaksson, global food and beverage segment manager, abb. *Food & Drink International*.
- EnergyPlus (2017). Energyplus, version 00.
- Engie (2022). Startup portfolio. <https://apac.engiefactory.com/portfolio/>. Accessed: Date.
- Hosseini, S., Agbossou, K., Kelouwani, S., Cardenas, A., and Henao, N. F. (2020). A practical approach to residential appliances on-line anomaly detection: A case study of standard and smart refrigerators. *IEEE Access*, PP:1–1.
- Miller, K. C. and Othman, A. (1994). Restaurants - opportunities for energy efficiency. https://www.aceee.org/files/proceedings/1994/data/papers/SS94_Panel5_Paper17.pdf.
- Murray, D., Stankovic, L., and Stankovic, V. (2015). REFIT: Electrical Load Measurements (Cleaned).
- Pereira, L., Aguiar, V., Vasconcelos, F., et al. (2023). Fikelectricity: A electricity consumption dataset from three restaurant kitchens in portugal. *Scientific Data*, 10:779.
- Statista (2023). Number of food and beverages outlets in singapore in 2021, by establishment type. <https://www.statista.com/statistics/724783/number-of-food-outlets-in-singapore-by-establishment-type/>.
- Wei, T., Wang, Y., and Zhu, Q. (2017). Deep reinforcement learning for building hvac control. In *Proceedings of the 54th Annual Design Automation Conference 2017*, pages 1–6. ACM. Article No.: 22.
- Zangrando, N., Fraternali, P., Petri, M., Pinciroli Vago, N. O., and Herrera González, S. L. (2022a). Anomaly detection in quasi-periodic energy consumption data series: a comparison of algorithms. *Energy Informatics*, 5(Suppl 4):62.
- Zangrando, N., Herrera, S., Koukaras, P., Dimara, A., Fraternali, P., Krinidis, S., Ioannidis, D., Tjortjjs, C., Anagnostopoulos, C.-N., and Tzovaras, D. (2022b). Anomaly detection in small-scale industrial and household appliances. In *Proceedings of Artificial Intelligence Applications and Innovations. AIAI 2022 IFIP WG 12.5 International Conference*, IFIP Advances in Information and Communication Technology. Paper presented at AIAI 2022 IFIP WG 12.5 International Conference.

Appendix

A Figures

A.1 Business Development

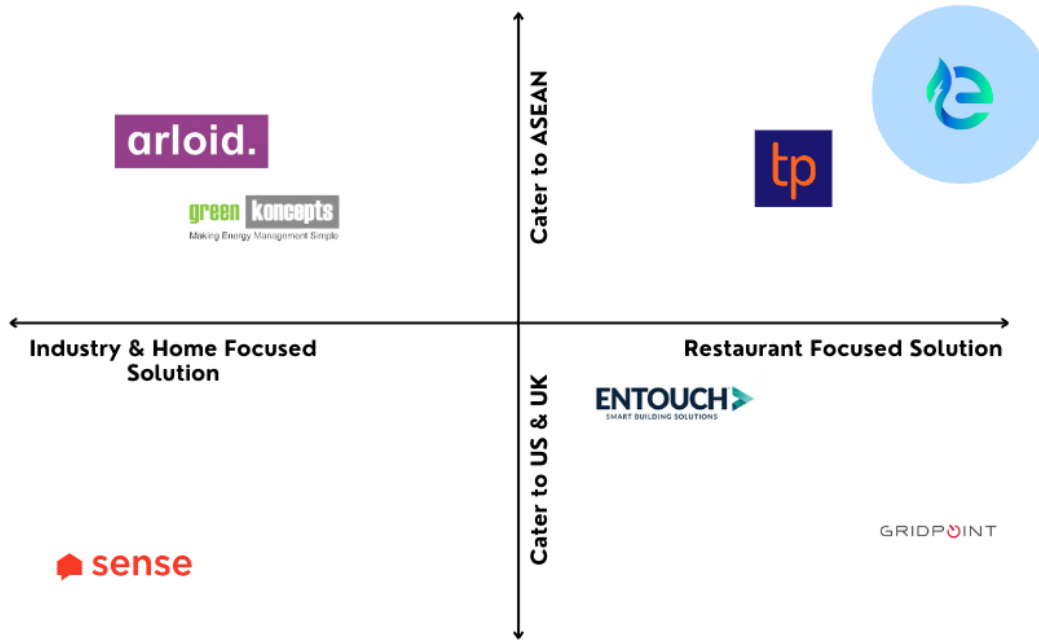


Figure 2: Competition Landscape Plane

Feature Group	Feature	Elgo	TablePointer	Sense	Arloid
Installation	Non-invasive				
	Ease of setup				
Analytics	Predictive analysis				
	Fault detection				
Intelligence	Realtime metrics				
	Remote device control				
	Recommendation engine				

Figure 3: Feature Matrix as Compared to Competitors

A.2 Kitchen Appliances and Refrigerator

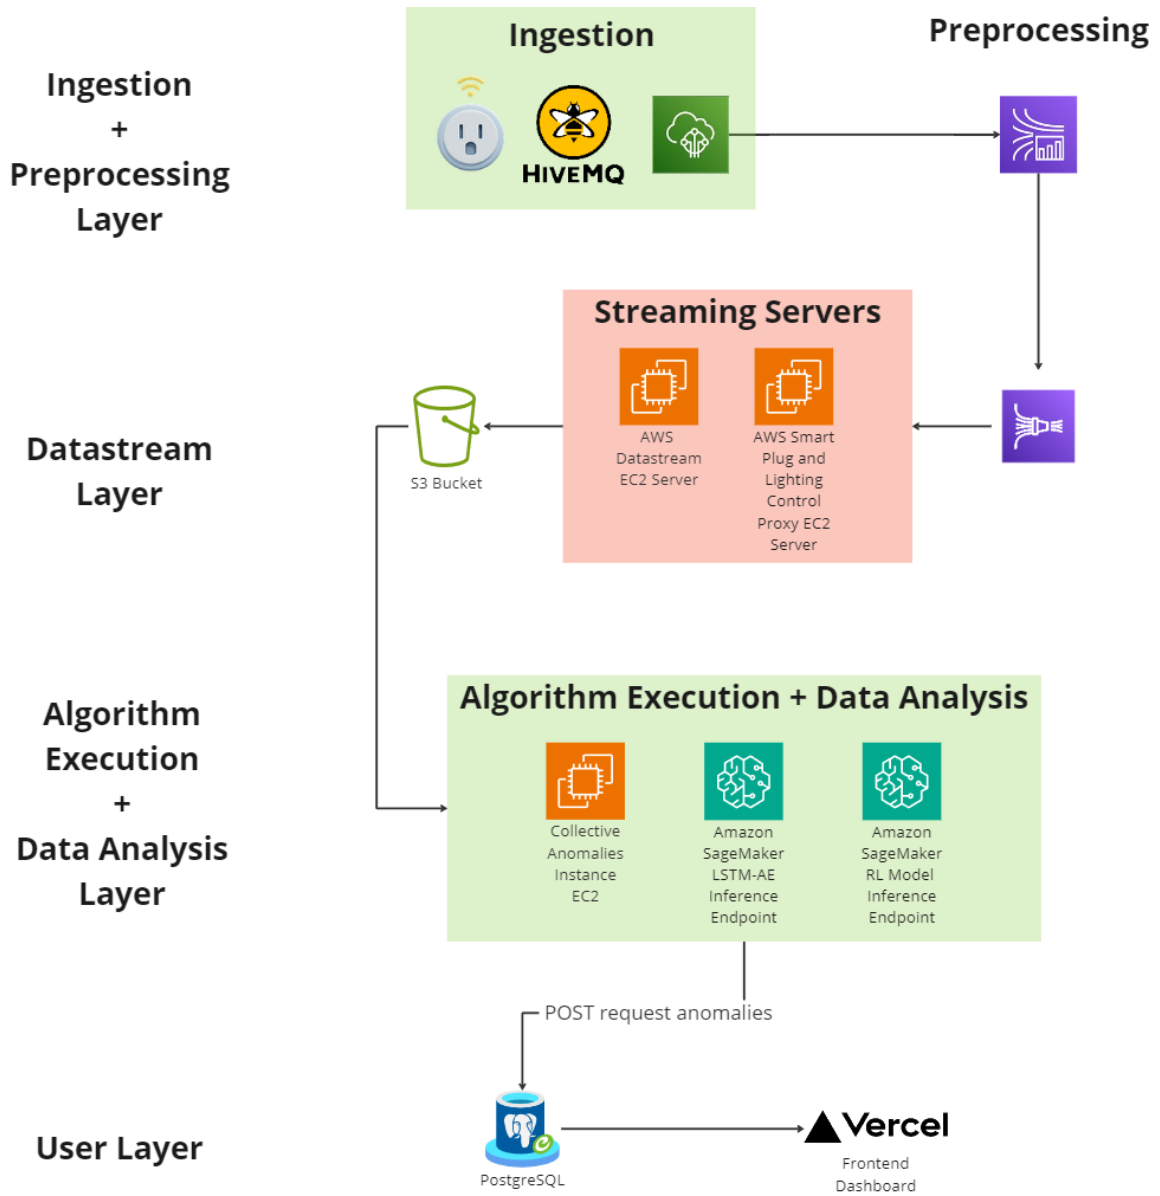


Figure 4: System Architecture Kitchen Appliances and Refrigerator

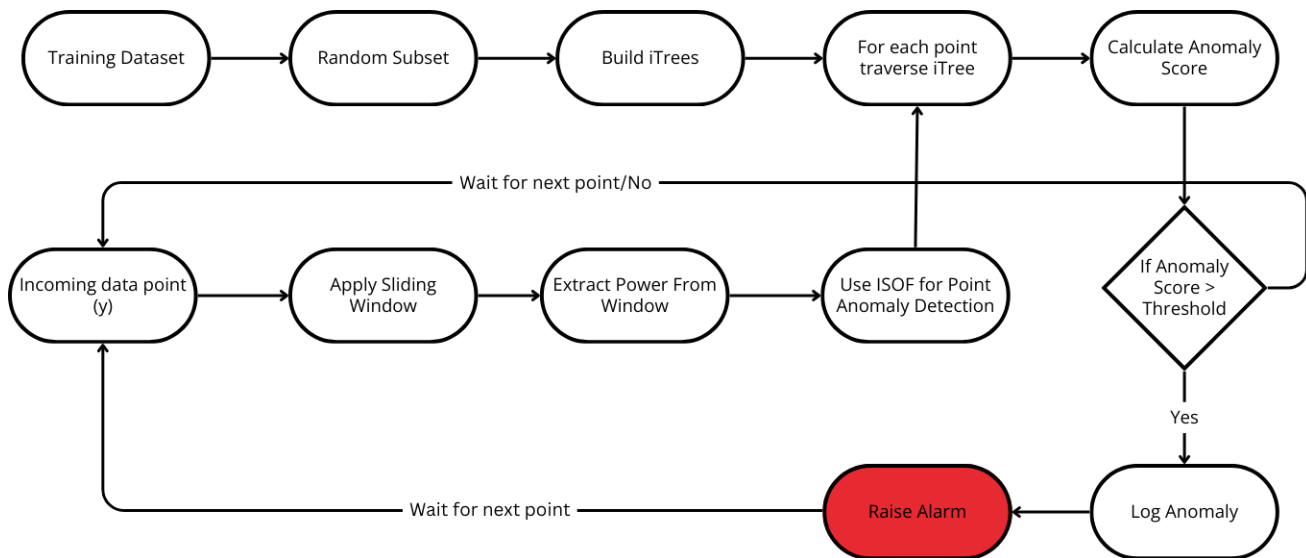


Figure 5: ISOF Decision Flow

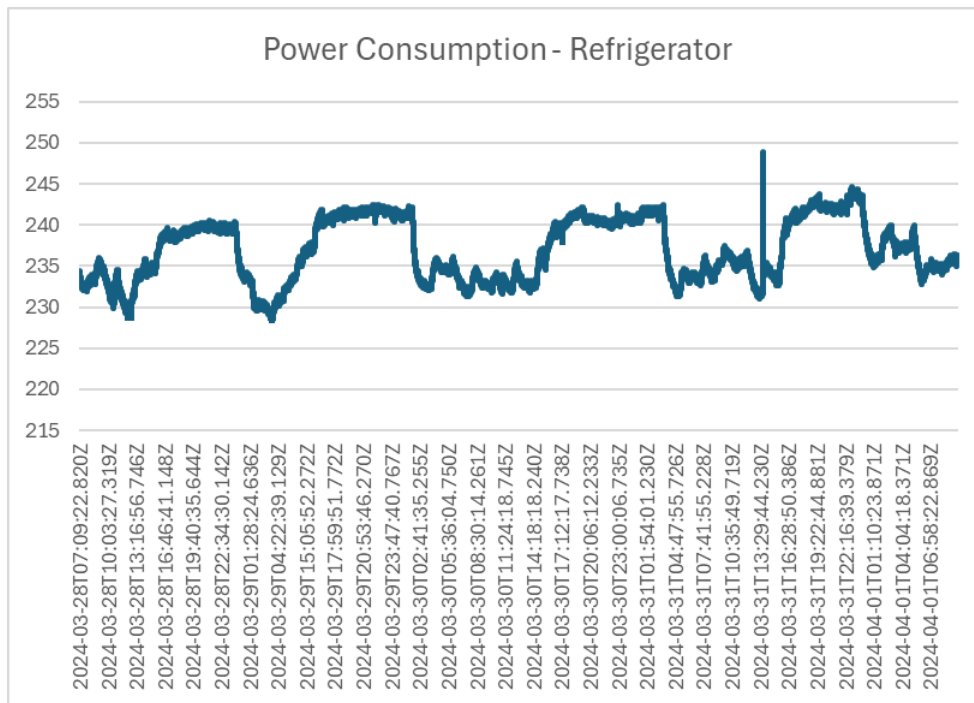


Figure 6: Power Signature of Refrigerator over a 4-day period

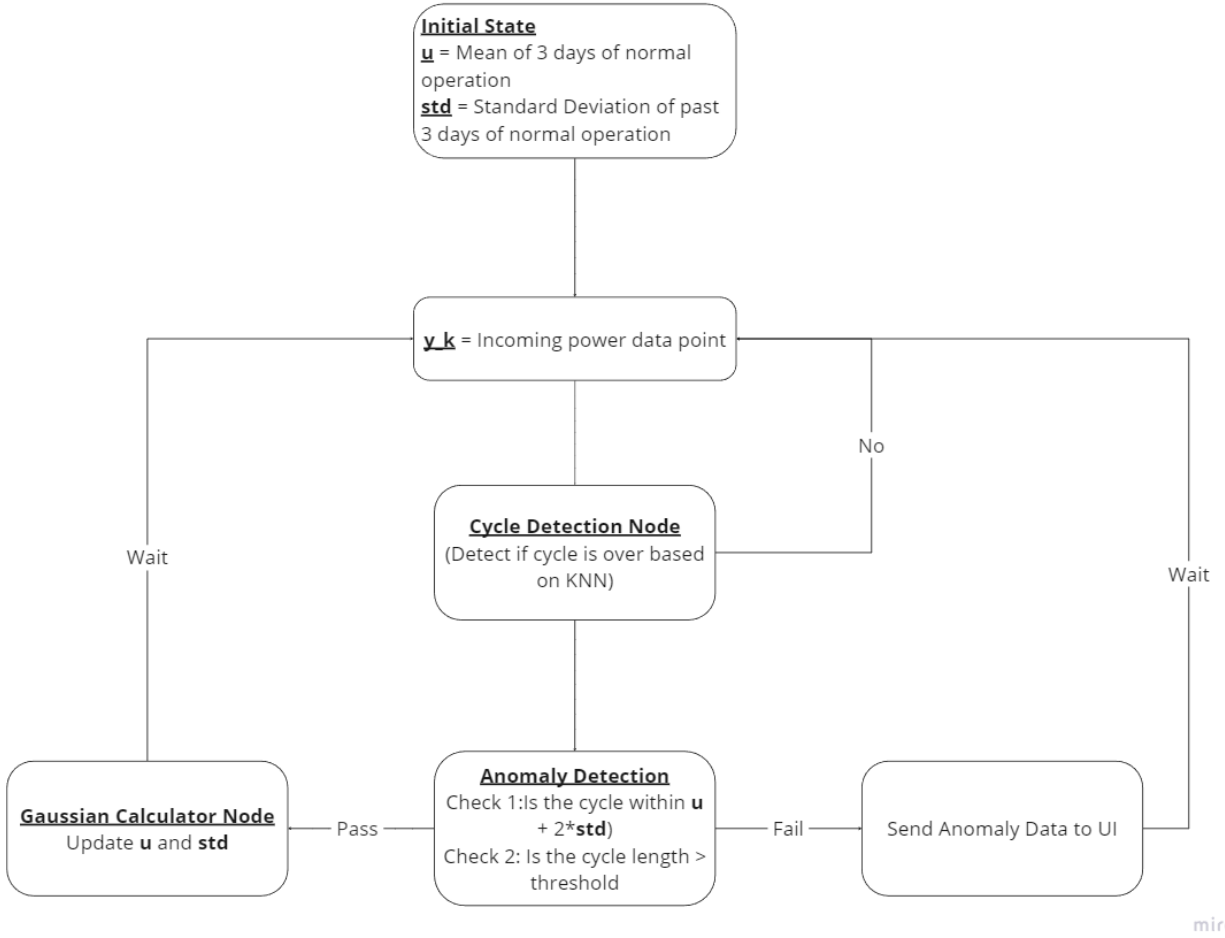


Figure 7: Gaussian Machine Learning Decision Flow

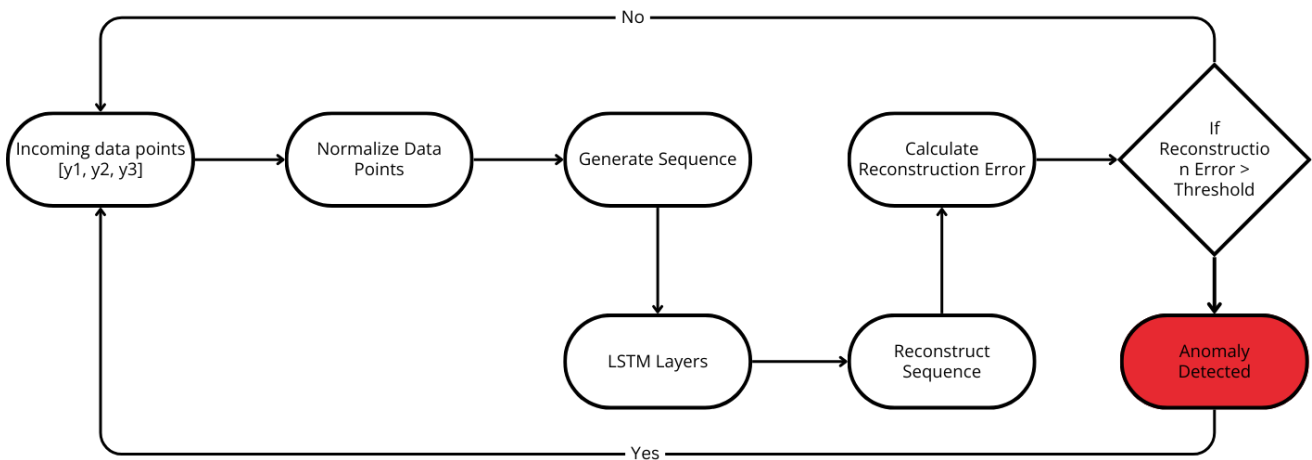


Figure 8: LSTM-AE Decision Flow
















Feature Group	Feature	Phase 1	Phase 2	Phase 3
Data Collection	Smart Plug			
	Streaming Data to IoT Hub			
Data Processing	Set up end-to-end cloud data pipeline			
	Store data			
Inference	Anomaly Detection Algorithm Development: Isolation Forest			
	Predictive Maintenance Algorithm Development: LSTM			
	Integrate with Dashboard			

Figure 9: Iterative Prototyping: Deliverables Achieved for Kitchen Appliances
















Feature Group	Feature	Phase 1	Phase 2	Phase 3
Data Collection	Smart plug			
	Sensor Data			
Data Processing	Unsupervised learning algorithms			
	Supervised learning algorithms			
Inference	Automatically book maintenance contact			
	Integrate with dashboard			

Figure 10: Iterative Prototyping: Deliverables Achieved for Refrigerator

A.3 Smart Plug

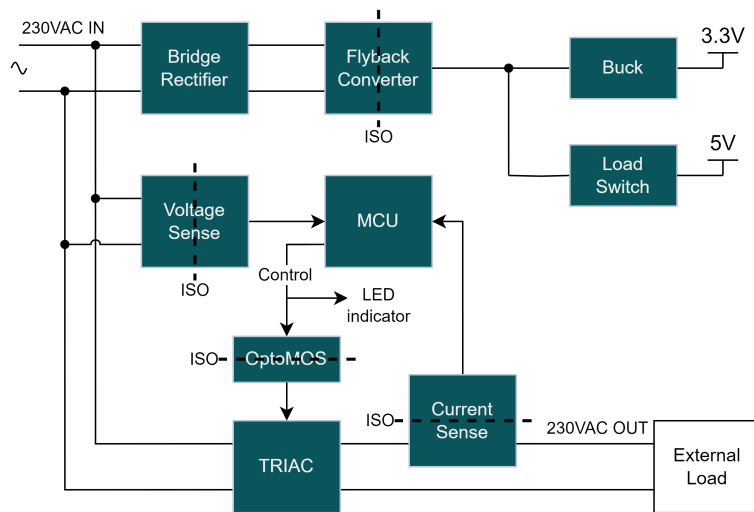


Figure 11: Smart Plug Functional Block Diagram

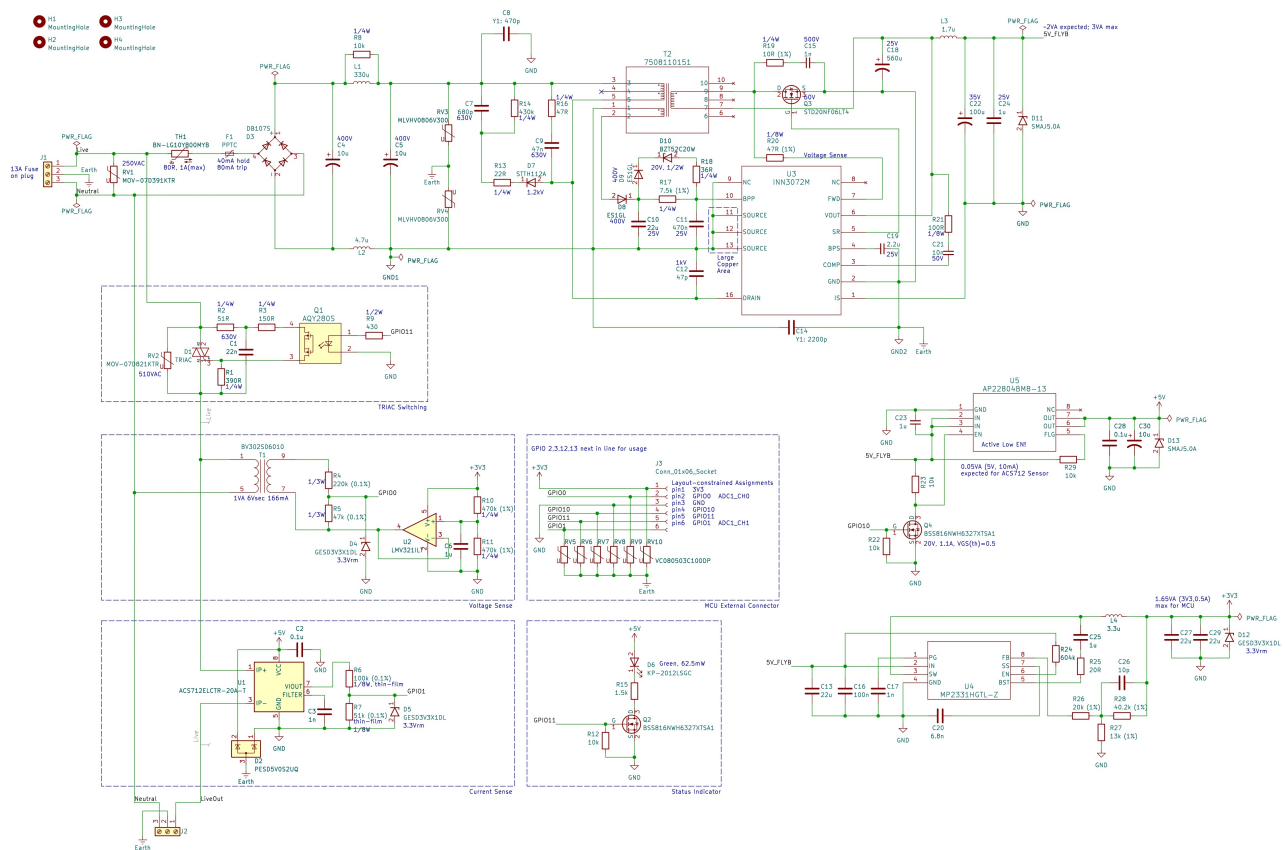
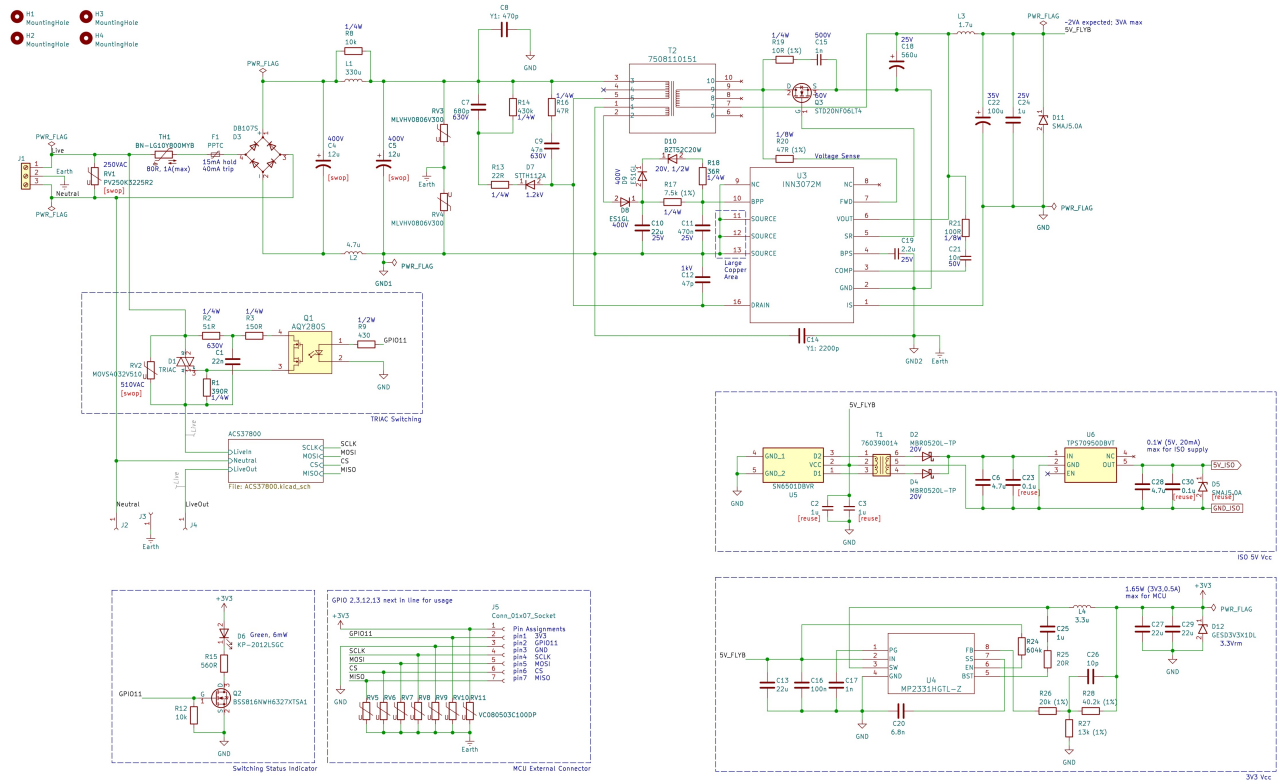
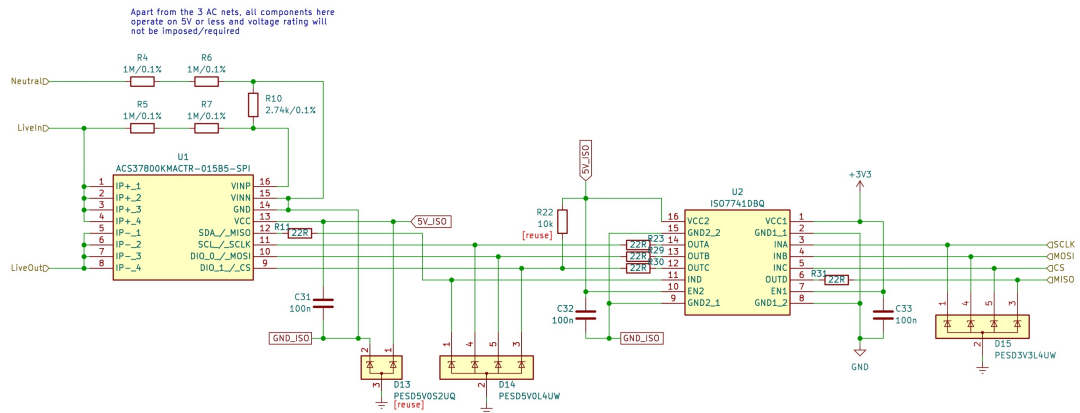


Figure 12: Phase 2 Schematic




(a) Page 1 of 2




(b) Page 2 of 2

Figure 13: Phase 3 Schematic




(a) Top View




(b) Bottom View

Figure 15: Phase 2 PCB 3D Render




(a) Top Copper Layer




(b) Bottom Copper Layer

Figure 16: Phase 3 PCB Layout




(a) Top View




(b) Bottom View


Figure 17: Phase 3 PCB 3D Render




(a) Top View



(b) Top View, rotated



(c) Bottom View



(d) Bottom View, rotated

Figure 18: PCB with SMT components assembled




Figure 19: Smart Plug Enclosure, Assembled




Figure 20: Smart Plug Enclosure, Exploded-view




Figure 21: Smart Plug Enclosure under testing

A.4 Lighting Systems

System Architecture

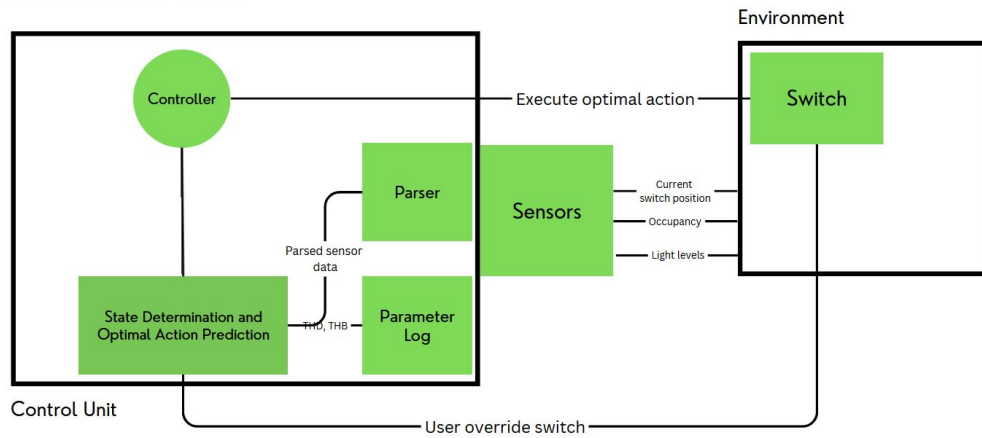


Figure 22: System Architecture for Lighting Solution




Figure 23: Moving Average of THD, THB.

Occupancy	Light switch position	Indoor light level	State	Optimal Action
Occupied	0	Dark	s1	0.5
	0	Comfort	s2	0
	0	Bright	-	Do nothing (0)
	0.5	Comfort	s3	Do nothing (0.5)
	0.5	Bright	s4	0
	1	Comfort	s5	Do nothing (1)
Unoccupied	1	Bright	s6	0.5
	0	-	s7	0
	0.5	-	s8	0
	1	-	s9	0

Figure 24: Mapping of State to Optimal Action.

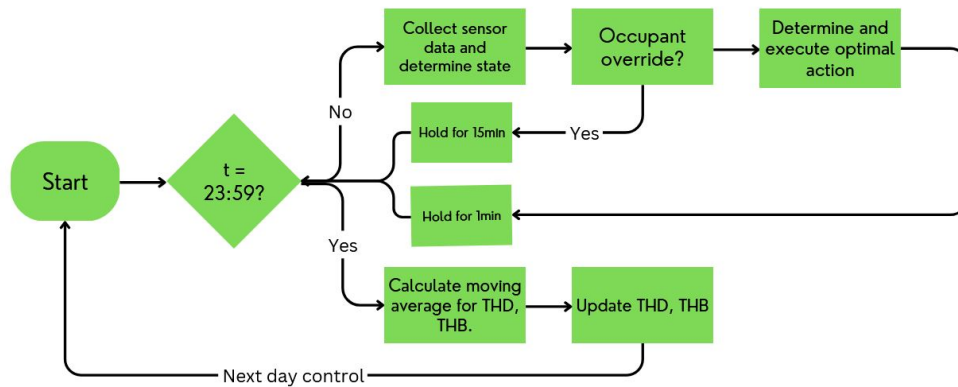


Figure 25: Logical Control Flow.

A.5 HVAC

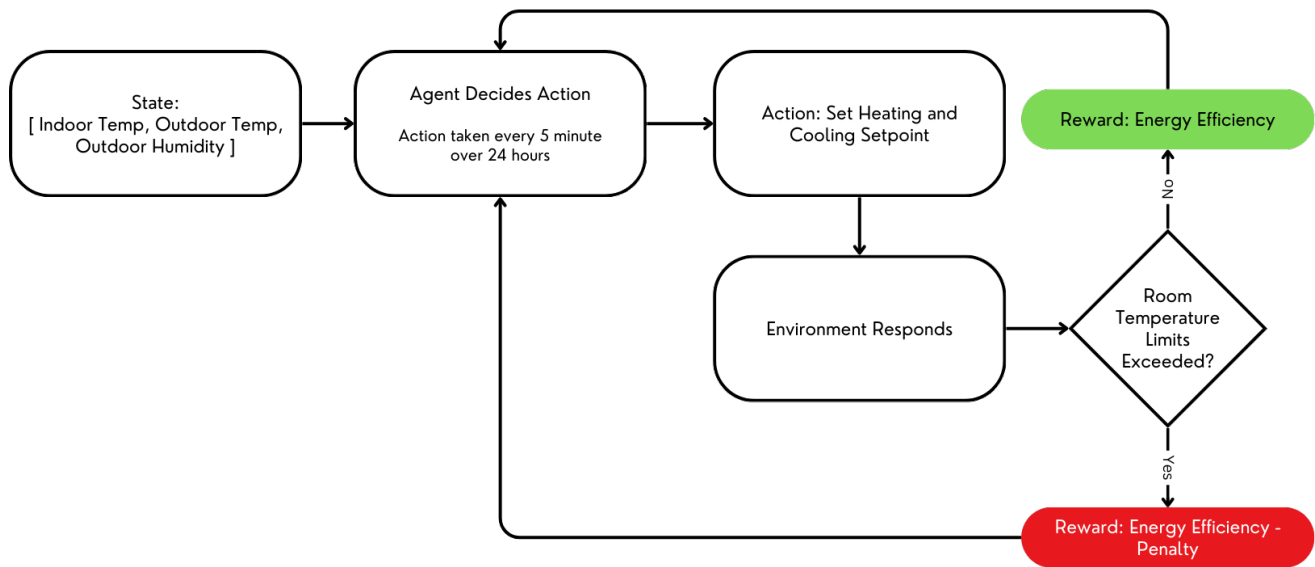


Figure 26: Deep Reinforcement Learning Flowchart (HVAC)




Figure 27: Simulation Environment (HVAC)

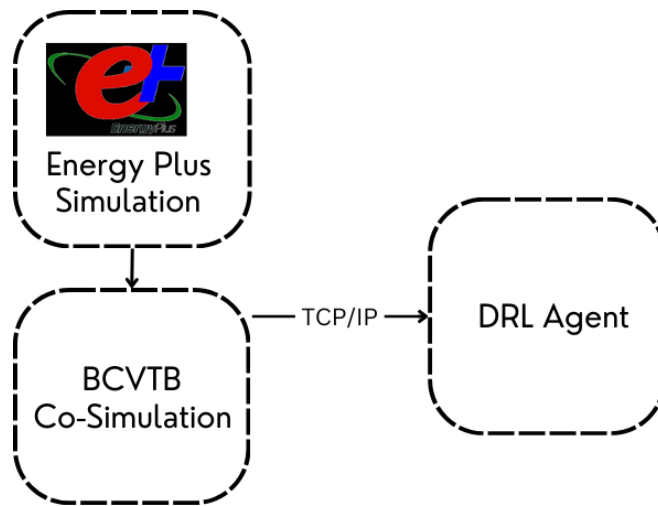


Figure 28: Simulation Framework (HVAC)



Figure 29: Simulation Restaurant SketchUp 3D Model (HVAC)




Figure 30: Simulation Model (HVAC)




Figure 31: Baseline vs DRL Control for 1 month (HVAC)
















Feature Group	Feature	Phase 1	Phase 2	Phase 3
Data Collection	Sensor Suite			
	Streaming data to IoT Hub			
Data Processing	Build Q-Learning algorithm			
	Optimizing Q-Learning			
Inference	Automate regulating temperature and heating and cooling set points			
	Integrate with dashboard			

Figure 32: Iterative Prototyping: Deliverables Achieved for HVAC

A.6 User Experience




Figure 33: Sequence diagram for user registration and login

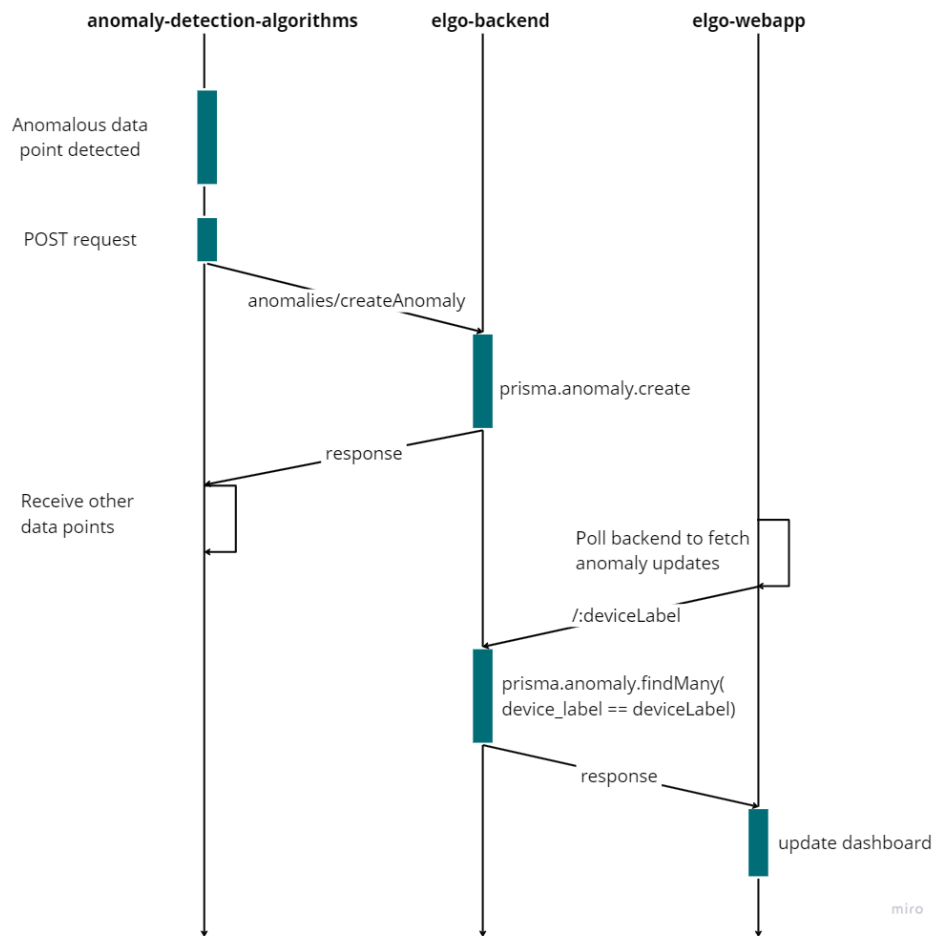


Figure 34: Sequence diagram for anomaly detection




Figure 35: Device registration and Anomaly Updates. This also ties in with the sequence described in Fig 34

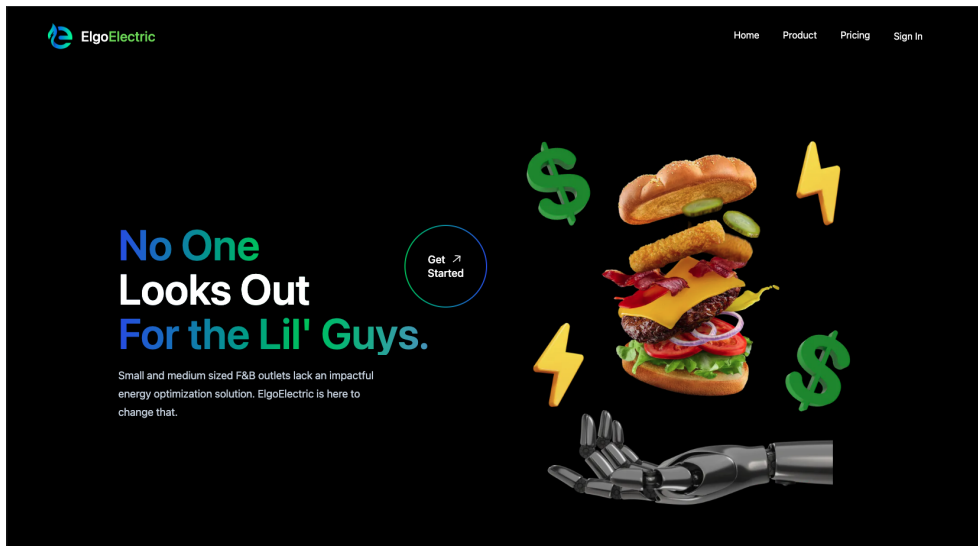


Figure 36: Landing Page Screenshot

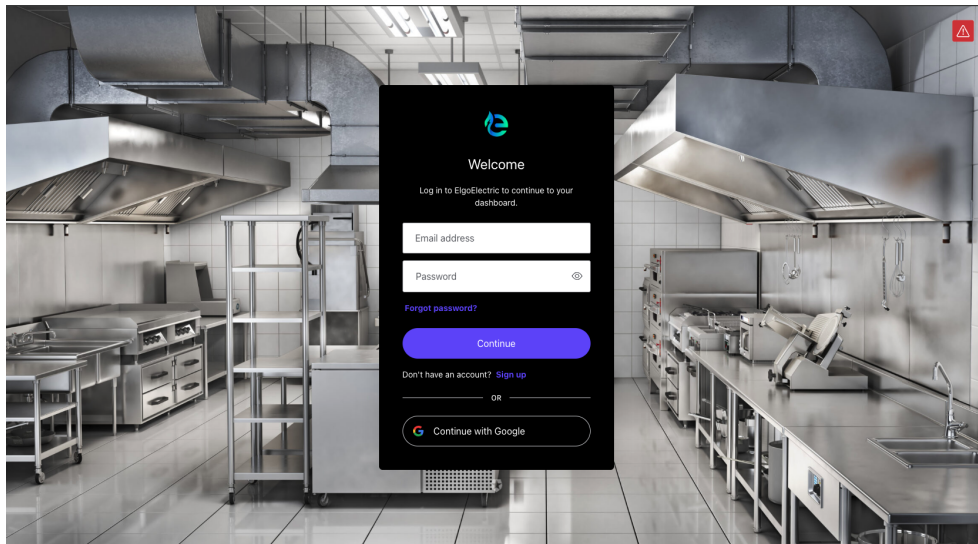


Figure 37: Login Page Screenshot

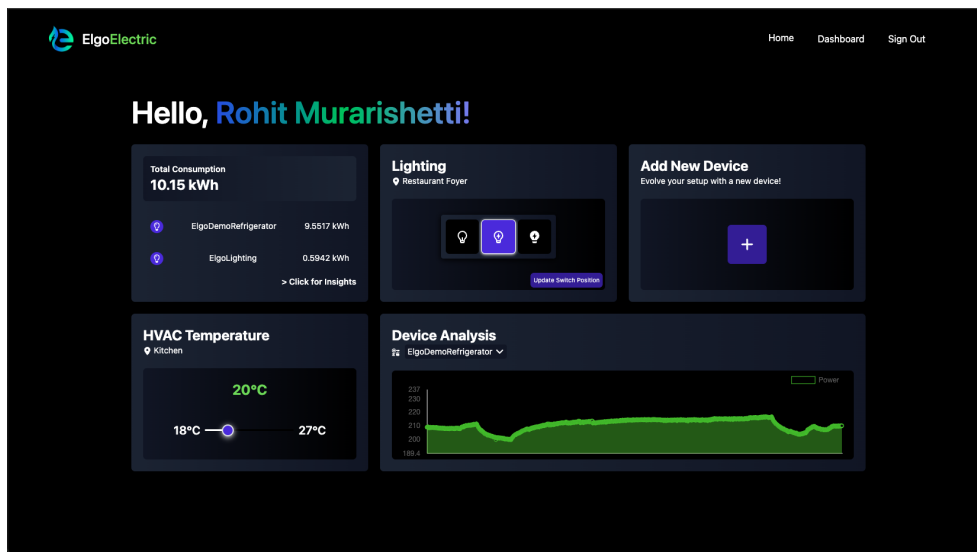


Figure 38: Dashboard Page Screenshot

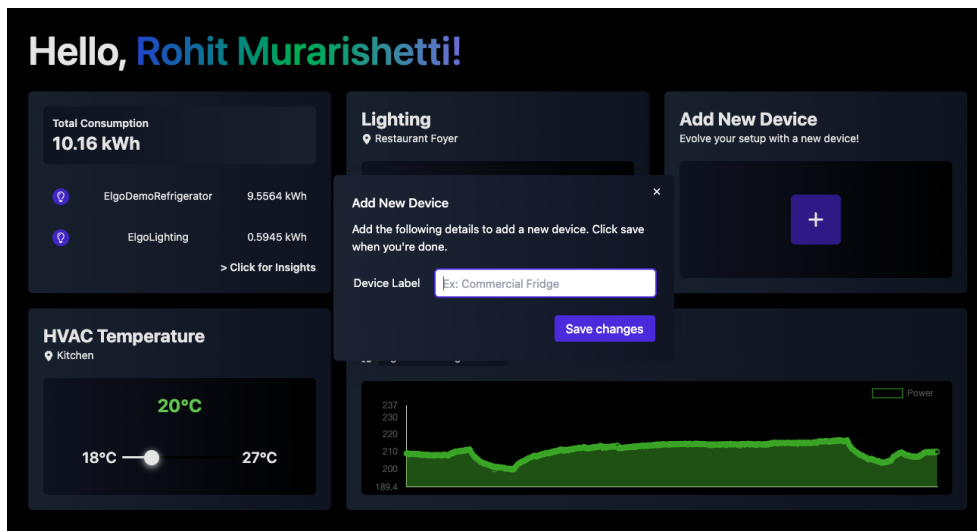


Figure 39: Device Registration Screenshot

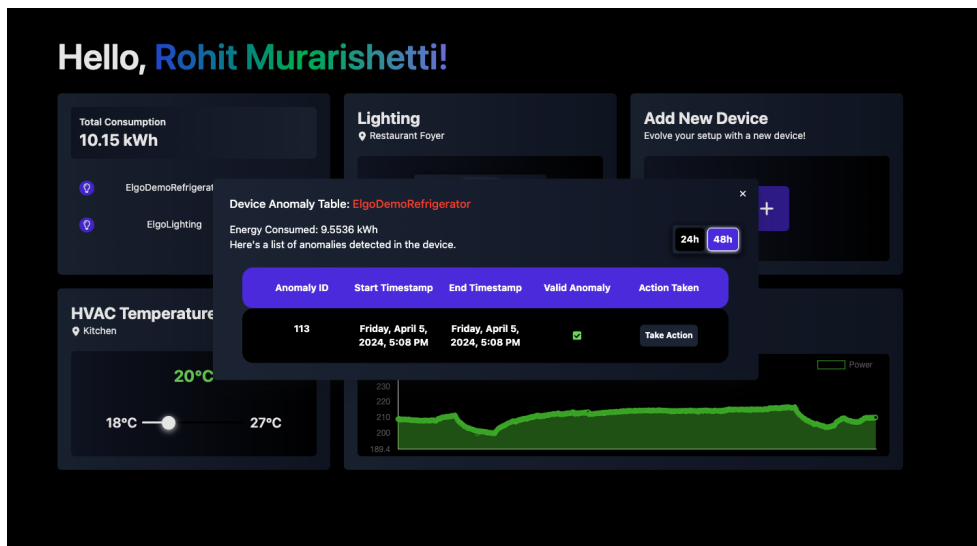


Figure 40: Anomalies Page Screenshot

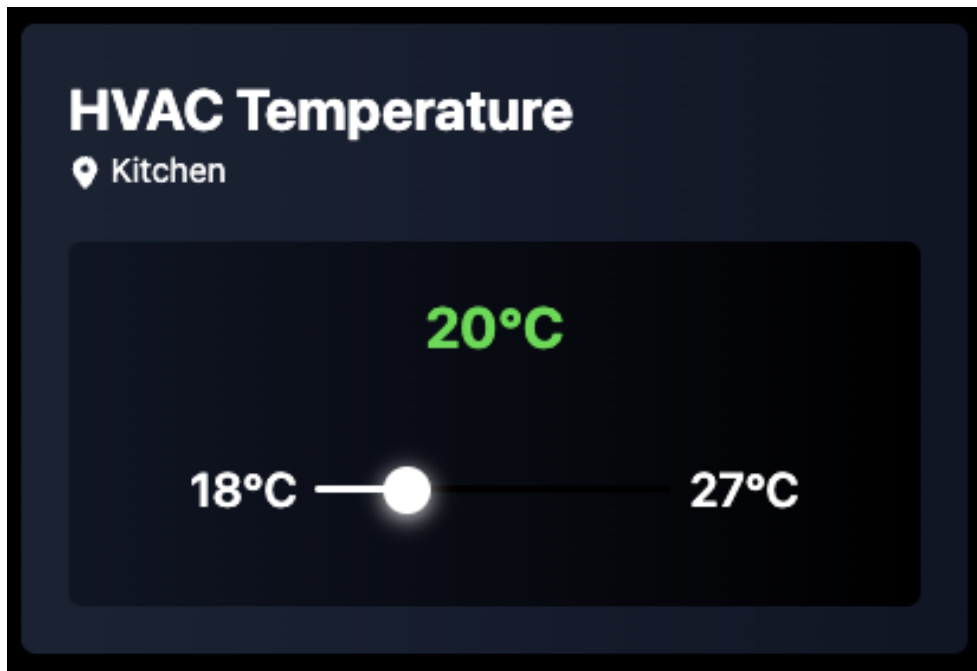


Figure 41: HVAC Control Screenshot

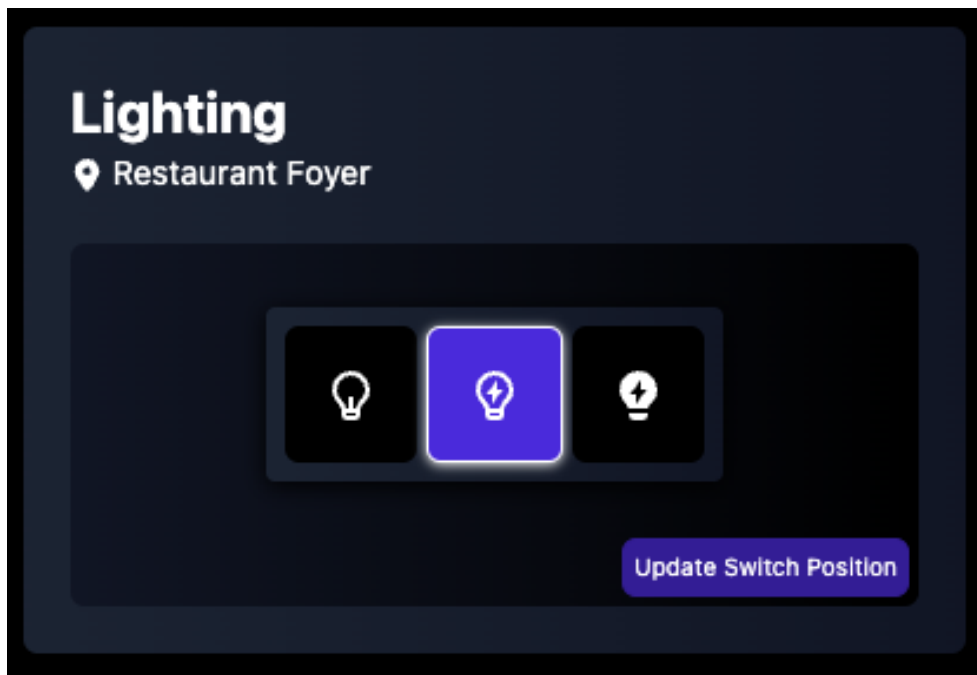


Figure 42: Lighting Control Screenshot

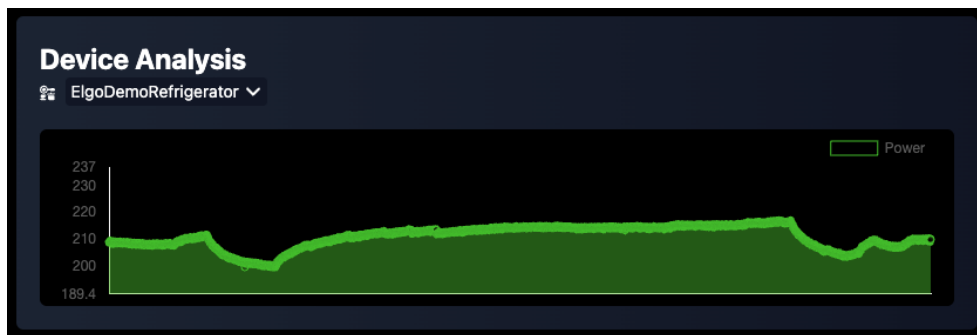


Figure 43: Power Consumption Insights Screenshot

B Tables

B.1 Solution Overview

Device Targeted	Purpose	Hardware component
Kitchen Appliance and Refrigerator	We focus on maintenance through power signatures. Appliances like refrigerator consume a large amount of power, and maintenance becomes key to cutting down costs.	Smart Plug
HVAC	HVAC systems are the biggest consumers of energy, and optimizing for indoor temperature and comfort levels will help achieve efficiency.	Smart Plug
Lighting System	Offer incremental improvements in both energy consumption optimization, as well as comfort levels.	Smart Plug + ESP 32 embedded with sensors

Table 3: [Device Overview](#)

B.2 Kitchen Appliances and Refrigerator

Requirement	Description
Data Collection	Automate real-time collection of performance and energy consumption data from kitchen appliances.
Anomaly Detection	Employ algorithms to detect point, collective and contextual anomalies .
Alerts	Provide immediate notifications to users for timely maintenance scheduling.

Table 4: [Functional Requirements for Anomaly Detection System](#)

Faulty Operation: cold air is constantly being lost.	<ol style="list-style-type: none"> 1. Door not closed properly 2. Defective Gasket
Abnormal Operation: cold air is kept inside, but has unusual power demand.	<ol style="list-style-type: none"> 1. Overly open or closed door 2. Fully loaded refrigerator

Table 5: [Refrigerator Anomaly Types](#)

Algorithm	Configuration Parameters
ISOF	Number of trees: [20, 200], Max samples: [150, 400]
Gaussian Machine Learning	Training set size: [3, 6, 9] days
LSTM-AE	Layers: 2 (128 and 64 units), Epochs: 500, Patience: 30, Batch size: 64

Table 6: Relevant configuration parameters of ISOF and LSTM-AE.

Period	F1 Score	Precision	Recall
1 Month	0.662	0.605	0.729
2 Weeks	0.772	0.651	0.950
3 Weeks	0.773	0.630	1.000

Table 7: ISOF model evaluation results on the REFIT Dataset.

Period	F1 Score	Precision	Recall
1 Month	0.737	0.584	0.998
2 Weeks	0.794	0.660	0.996
3 Weeks	0.804	0.723	0.905

Table 8: ISOF model evaluation results on the FIKElectricity Dataset.

Period	F1 Score	Precision	Recall
1 Day	0.54	0.511	0.623
2 Days	0.68	0.525	0.826
3 Days	0.81	0.533	0.689
4 Days	0.821	0.450	0.71
5 Days	0.823	0.443	0.71
6 Days	0.828	0.421	0.70

Table 9: Gaussian Statistics Machine Learning model evaluation results on the REFIT Dataset

Period	F1 Score	Precision	Recall
1 Day	0.608	0.584	0.623
2 Days	0.79	0.752	0.826
3 Days	0.80	0.81	0.689
4 Days	0.801	0.82	0.689
5 Days	0.823	0.81	0.691
6 Days	0.824	0.812	0.691

Table 10: Gaussian Statistics Machine Learning model evaluation results on the FIKElectricity Dataset

Period	F1 Score	Precision	Recall
1 Month	0.603	0.584	0.623
2 Weeks	0.79	0.757	0.826
3 Weeks	0.548	0.455	0.689

Table 11: LSTM-AE model performance metrics on the REFIT Dataset.

B.3 Smart Plug

B.4 Lighting Systems

Requirement	Description
ESP32 Integrations	Connect ESP32 to phototransistor, set up Wifi to make RESP API calls and BLE to detect peripheral devices.
Data Collection	Collect data on phototransistor analog readings at which occupant turns on and off the lights.
Action Selection	Selection of optimal action to meet goals of the system.

Table 12: Functional Requirements for Lighting Systems

Duration	Power Consumed without Controller	Power Consumed with Controller
3	1443.1W	890.2W
6	2871.6W	1523.2W
9	4118.2W	2177.6W

Table 13: Power Consumption with and without Controller

B.5 HVAC

Requirement	Description
Data Acquisition	Automate the collection of real-time data for temperature, humidity, and occupancy from sensors throughout the premises.
Adaptive Control	Utilize DRL algorithms to dynamically adjust HVAC operations based on data inputs to optimize energy use and maintain comfort.
Energy Management	Implement energy-saving strategies that respond to variable occupancy and environmental patterns, aiming for reduced energy consumption.
User Interaction	Provide an intuitive user interface for system monitoring and overrides, ensuring ease of use for non-technical personnel such as restaurant owners.

Table 14: Functional Requirements for HVAC System Optimization

C Hardware Calculations

C.1 Internal DC Supply

The internal DC supply was designed using the following calculations. The AC voltage is rectified by the bridge rectifier, and the rectified high-voltage DC is switched at the given duty cycle by the flyback controller IC on the flyback transformer primary. At the flyback secondary, synchronous rectification results in the 5V flyback output. The 5V flyback output is connected to both the 5V Vcc load switch and the 3.3V buck converter in parallel. The output of both provides the Vcc supply for other auxiliary circuits.

C.1.1 AC-side Protection Features

- Maximum internal supply wattage = 3 W, hence the maximum AC current drawn at 230 V (rms) is 13.1 mA (rms).
- Overvoltage Protection by MOV: a metal-oxide varistor (MOV) provides bidirectional clamping at specified clamping voltage, shunting current during transient voltage spikes. The working voltage of 250 V (rms) is chosen from +10% tolerance value of the 230 V mains voltage
- Inrush Current Suppression by NTC ICL Thermistor: A negative temperature coefficient thermistor exhibits high resistance during transient startup, and decreases in resistance as the device heats up from joule heating. Selection: BN-LG10Y800MYB.
 - $R_{25} = V_p / I_{inrush} = \frac{325}{13} = 25 \text{ R}$
 - This gives us leeway for uprating to a less destructive inrush current by selecting a higher (commercially available) R_{25} value, 80 R, then: $I_{inrush} = \frac{325}{80} = 4.1 \text{ A}$.
 - Under steady-state operation, $R_{I_{max}} = 2.2 \text{ R}$. Hence the maximum power dissipation under 13.1 mA max steady-state current, the power dissipation is $P = I_{ss}^2 R_{I_{max}} = (0.013)^2 (2.2) = 0.372 \text{ mW}$, which is acceptable.
- Overcurrent Protection by Poly-fuse: polymeric positive temperature coefficient fuses are chosen for 15 mA holding current and 40 mA trip current, which is minimally above the rated current at 3 W. Selection: B59890C0080A070 is rated at 150 R resistance with 25.7 mW max power dissipation.

C.1.2 Rectification

- Rectification occurs at 230 V (rms). Hence the minimal required Peak Inverse Voltage (PIV) is given by $PIV > 230\sqrt{2} - 2V_F = 230\sqrt{2} - 2(1) = 323 \text{ V}$
- The ideal PIV safety margin is given by $3V_{rms}\sqrt{2} - 2V_F = 974 \text{ V} \Rightarrow 1 \text{ kV}$ bridge rectifier rating.
- Rectified Peak Current at 230 V, $I_F = \frac{3}{230\sqrt{2}} = 9.2 \approx 10 \text{ mA}$

C.1.3 Flyback Converter

- Topology Selection:

- In the event of failure (fail short) of the offline switching IC, the DC-side will become energized at live potential. Flyback and Push-pull topologies are available to fulfill isolation requirements.
 - Complexity: Flyback topology presents reduced complexity compared to push-pull
 - Efficiency: Push-pull topology presents greater efficiency during operations, esp. for high-power applications.
 - Size: Flyback topology generally has smaller size requirements. \Rightarrow Selection: Flyback topology
- Flyback Controller IC:
 - Selection: InnoSwitch3-TN INN3072M, 750 V MOSFET, 12 W 5 V output
 - For a single output of 5 V, and an input of 325 V peak rectified from the mains, the internal regulation of the PWM is such that the following duty cycle is obtained.
 - * $\frac{V_o}{V_{in}} = \frac{N_s}{N_p} \cdot \frac{D_1}{1-D_1}$
 - * $\frac{5}{325} = \frac{1}{14.16} \cdot \frac{D_1}{1-D_1}$
 - * $D_1 = 17.9\%$
 - VOR, the flyback voltage, is given by: $VOR = V_o \cdot \frac{N_p}{N_s} = V_{in} \frac{D_1}{1-D_1} = 5(14.16/1) = 70.8 \text{ V}$
 - Using 17.9% duty cycle, auxiliary winding is given by: $\frac{V_{aux}}{325} = \frac{1}{4.47} \cdot \frac{0.179}{1-0.179} \Rightarrow V_{aux} = 15.9 \text{ V}$
 - Flyback Transformer:
 - Selection: Würth 7508110127, 82-375 VDC, 120 kHz \rightarrow 5 V 1 A Secondary
 - Turns Ratio, Pri:Sec = 14.16:1
 - Turns Ratio, Pri:Aux = 4.47:1
 - Leakage Inductance = 23 μH typical, 45 μH max
 - Inductance = 1.05 mH $\pm 10\%$
 - Input Capacitor:
 - Benchmark rating is given by $2P_{out} \text{ (W)} = C \text{ (}\mu\text{F)}$. For 3 W output, 6 μF minimum is required \Rightarrow 10 μF input capacitor, 400 VAC rating.
 - For aluminum electrolytic capacitors, lifetime rating is given by: $L = L_0 \cdot 2^{\frac{T_{max}-T_a}{10}} = 1000(2^{\frac{105-60}{10}}) = 22000 \text{ hrs} \Rightarrow 2.5 \text{ years}$
 - * L_0 is the life at rated temperature (hrs)
 - * T_{max} is the rated temperature ($^{\circ}\text{C}$)
 - * T_a is the ambient temperature expected during normal operation, *i.e.* Environmental Temp + PCB Heating Tolerance ($^{\circ}\text{C}$)
 - Input Pi-filter:
 - Cutoff frequency should adhere 1/10th of typical switching frequency (AN101E). Using 100 kHz as typical switching frequency, $f_c < 0.1f_{sw} = 0.1(100\text{k}) = 10 \text{ kHz}$

- The inductor is sized based on this desired cutoff frequency.
 - * $\frac{1}{2\pi\sqrt{LC}} < 10000 \text{ kHz}$ for -3dB attenuation
 - * $\left(\frac{1}{2\pi \times 10000}\right)^2 < L \times 10\mu$
 - * $L > 25 \mu\text{H}$
 - * With the recommended inductor value of $330 \mu\text{H}$, $f_c = \frac{1}{2\pi\sqrt{LC}} = 2.9 \text{ kHz}$ which is sufficient for our purpose.
- Primary-side RCD Clamp:
 - The RCD (resistor, capacitor, diode) clamp suppresses dV/dt surges when uncoupled leakage inductance stores energy in the magnetic field but fails to transfer it to the secondary windings. Because of energy conservation, this surge voltage then occurs across the switching MOSFET, which requires a clamp to prevent exceeding the MOSFET V_{DS} breakdown rating.
 - For 725 V switching MOSFET and 15-20% derating factor (safety margin),
 - * $V_{\text{clamp}} = 750 \times 0.8 = 580 \text{ V}$, however, we may choose a $V_{\text{clamp}} = 350 \text{ V}$, which is just above the rectified mains voltage for even more conservative sizing
 - * $V_{\text{ripple}} = 50 \text{ V}$ peak-to-peak, estimated
 - Given leakage inductance of $45 \mu\text{H}$ from the transformer specifications (otherwise the benchmark estimation is 10% of the primary winding inductance of $1.05 \text{ mH} = 105 \mu\text{H}$). The **5 V 1 A = 5 W** output power of the flyback transformer is used as an updated estimate. For the resistive element,
 - * $I_{p, \text{pri}} = I_{p, \text{sec}} \times \frac{N_s}{N_p} = 1 \times \frac{1}{14.16} = 70.7 \text{ mA}$
 - * $R_{\text{clamp}} < 2V_{\text{clamp}} \times \frac{V_{\text{clamp}} - V_{\text{OR}}}{L_{\text{leak}} I_p^2 f_{sw(\text{max})}} = 2(350) \frac{350 - 70.8}{45\mu \times 70.621\text{m} \times 132\text{k}} = 465.9 \text{ kOhm} \Rightarrow 430 \text{ kOhm}$ chosen.
 - * Power rating of R_{clamp} is: $\frac{(V_{\text{clamp}} - V_{in})^2}{R_{\text{clamp}}} = \frac{(350 - 325)^2}{430\text{k}} = 1.5 \text{ mW} \Rightarrow 1/8 \text{ W}$ rating.
 - For the capacitor,
 - * $C_{\text{clamp}} > \frac{V_{\text{clamp}}}{V_{\text{ripple}} \cdot f_{sw(\text{min})} \cdot R_{\text{clamp}}} = \frac{350}{50 \times 25\text{k} \times 430\text{k}} = 651 \text{ pF} \Rightarrow 680 \text{ pF}$
 - For the diode,
 - * PIV that withstands the maximum V_{DS} of the MOSFET (725V)
 - * Fast recovery diode required, for $f_{sw} = 100 \text{ kHz}$, $T_{sw} = 10 \mu\text{s} \Rightarrow$ a diode with nanosecond-order recovery time is required.
- RC Snubber for RCD Clamp:
 - The following method details the sizing of the RC snubber which serves as a damping network for the RCD clamp.
 1. Measure the ringing f_r
 2. Evaluate leakage reactance at f_r , given by $X_{\text{leak}} = 2\pi L_{\text{leak}} f_0$ (assumes quality factor = 1)
 3. Set $R_{\text{damp}} = X_{\text{leak}}$
 4. Set $C_{\text{damp}} = \frac{1}{2\pi f_r R}$

– Hence, assuming 100 kHz ringing due to the switching,

$$* X_{\text{leak}} = 2\pi L_{\text{leak}} f_r = 2\pi(45\mu)(100\text{k}) = 28.27 \text{ Ohm} \Rightarrow R_{\text{damp}} \text{ uses } 47 \text{ R}$$

$$* C_{\text{damp}} = \frac{1}{2\pi \times 100\text{k} \times 50} = 31.83 \text{ nF} \rightarrow 47 \text{ nF}$$

NOTE : Power dissipation increases with $P_{\text{loss}} = C_{\text{damp}} V_{\text{snubber}}^2 f_{sw}$

• C Snubber for Power MOSFET:

- C-only snubber avoids phase shift and unnecessary power dissipation associated with resistor
- 10-100 pF across the flyback IC internal MOSFET drain and source pins \Rightarrow 47 pF value chosen

• Primary Bypass OVP:

- Bypass pin (BPP) capacitor provides decoupling of the primary-side controller and is charged during startup by an internal current source.
- OVP is achieved by 7.4 mA of latching shutdown current I_{SD} entering the BPP pin. A parallel branch containing a series blocking diode, Zener diode, and resistor is used.
 - * Zener is clamped at $1.25V_{\text{aux}} = 1.25 \times 15.9 = 19.9 \text{ V} \Rightarrow 20 \text{ V}$ Zener chosen. Under normal conditions, the Zener is non-conducting, and the clamp circuit is o/c. During an OVP event, the Zener begins conducting, forming a low impedance branch (vs. the current-limiting resistor) in which the current to the BPP pin rapidly increases up to the threshold shutdown level.
 - * The blocking diode prevents reverse current from discharging the BPP capacitor during self-start.
 - * Zener Power Dissipation Rating exceeds: $20 \times 7.4\text{m} = 0.148 \text{ W}$
 - * BPP filter capacitor = 470 nF

• Secondary winding:

- In a flyback topology, the polarity dot of the primary and secondary winding lies on opposite sides (e.g., primary - top side; secondary - bottom side). This achieves the following behavior:
 - * Vin ON, current enters dot on primary \rightarrow dot on secondary is + polarity \rightarrow diode is blocking.
 - * Vin OFF, emf reverses direction \rightarrow dot on secondary is - polarity \rightarrow diode is conducting.
- The above mechanism must be strictly followed for proper operation of the flyback converter

• Output Pi-filter:

– Capacitor Selection

$$* Z_C < \frac{\Delta V}{I_{p,sec}} = \frac{0.2}{1} = 0.2 \text{ Ohm}$$

$$* jX_C = j0.2 = j \frac{1}{2\pi fC} \Rightarrow C > 8.0 \mu\text{F} \Rightarrow \text{Use default recommendation of } 560 \mu\text{F}$$

$$* \text{Rating} = 2 \times V_o = 10 \text{ V, minimal}$$

– Cutoff Frequency Estimation

$$* \frac{1}{2\pi\sqrt{LC}} = 10 \text{ kHz for } -3\text{dB attenuation}$$

$$* \left(\frac{1}{2\pi \times 10000}\right)^2 = L \times 560\mu$$

$$* L = 0.5 \mu\text{H minimum} \Rightarrow \text{Use default recommendation of } 1.7 \mu\text{H}$$

- * Actual cut-off frequency with 1.7 μ H inductor: 5.2 kHz
- EMI Filtering:
 - RC Snubber
 - * 1 nF, 10 Ohm (default value) RC snubber for SR MOSFET snubber. 100 V (C), 1/8 W (R) rating.
 - * 47 nF, 50 Ohm RC snubber for RCD damping snubber (refer above). 500 V (C), 1/8 W (R) rating.
 - * 47 pF, C snubber for Power MOSFET snubber (refer above). 1 kV (C) Rating.
 - Y1 Capacitor: decoupling between the primary and secondary winding
 - * Rating Y1 or 2x series Y2 capacitors alternatively.
 - * Class Y1 Requirements
 - Rated Operating Voltage: $V \leq 500$ VAC
 - Peak Impulse Voltage: 8 kV (exceeds flyback transformer Insulation Rating, 4 kV)
 - * Class Y2 Requirements:
 - Rated Operating Voltage: $150 \leq V \leq 300$ VAC
 - Peak Impulse Voltage: 5 kV (exceeds flyback transformer Insulation Rating, 4 kV)
 - Y2 cap does not exceed 4700 pF (4.7 nF), typical available range up to 2.2 nF
 - * Murata DE1 Series
 - 470 pF, across primary 325 V rail to secondary output GND
 - 2200 pF (2.2 nF), across primary GND1 to secondary GND2
- Grounding:
 - Primary GND1 left floating
 - Secondary GND2 tied to chassis/AC Earth plane
 - Output GND left "floating", connected to GND2 via an internal current sense resistor
 - In terms of PCB layout, GND2 should be one ground plane (In.1) and tied to the Earth plane (In.2), while GND1 and GND are separate GND planes (In.1)
- 5V Flyback Output: TVS Diode chosen with 5V reverse working voltage for unidirectional clamping.

C.1.4 Load Switch (Phase 2)

- Active Low EN logic: 0V at EN \rightarrow ON
 - 10 kOhm pull-up to 5 V (max 6.3 V tolerated) \rightarrow Load Switch OFF by default during transient operation
 - NMOS pull-down to 0 V after 3.3 V (high) gate signal provided by MCU \rightarrow Load switch ON
 - Logic-level NMOS requires $V_{GS(th)} \approx 0.5$ V for 3.3 V gate signal to turn on sufficiently.

C.1.5 Buck Converter 3.3 V

- Selection: MP2331H, 24 V input, 2 A output
- EN Pin: drive high via 604 kOhm pullup resistor to Vin
- Inductor Selection:
 - $\Delta I_L = 0.3 I_{L(\max)} = 0.3 \times 500\text{m} = 150\text{ mA}$ (30% current ripple - Lower Limit)
 - $L_1 = \frac{V_o \times (V_{in} - V_o)}{V_{in} \times \Delta I_L \times f_{sw}} = \frac{3.3 \times (5 - 3.3)}{5 \times 150\text{m} \times 1.2\text{M}} = 6.2\ \mu\text{H}$
 - $\Delta I_L = 0.3 I_{L(\max)} = 0.6 \times 500\text{m} = 300\text{ mA}$ (60% current ripple - Upper Limit)
 - $L_1 = \frac{V_o \times (V_{in} - V_o)}{V_{in} \times \Delta I_L \times f_{sw}} = \frac{3.3 \times (5 - 3.3)}{5 \times 300\text{m} \times 1.2\text{M}} = 3.1\ \mu\text{H} \Rightarrow 3.3\ \mu\text{H}$ selected
 - L_1 dc current rating exceeds $1.25 \times 0.5 = 0.625\text{ A}$ for $I_{L(\max)} = 0.5\text{ A}$
 - Saturation Limit: $I_{LSP} = I_{L(\max)} + \frac{V_o}{2f_{sw}L} \left(1 - \frac{V_o}{V_{in}}\right) = 0.5 + \frac{3.3}{2 \times 1.2\text{M} \times 6.2\mu} \left(1 - \frac{3.3}{5}\right) = 575.4\text{ mA} \Rightarrow$ fulfilled by dc current rating above.
- Input Capacitor Selection:
 - Low ESR electrolytic, tantalum, or ceramic material for 22 μF caps.
 - Use 100 nF, 1 nF ceramic cap in 0603 packages for absorbing high-frequency noise. This should be placed as close to Vin and GND pins as possible.
 - Ripple current rating should be greater than $I_{C(\text{in})} = I_{L(\max)} \sqrt{\frac{V_o}{V_{in}} \left(1 - \frac{V_o}{V_{in}}\right)} = 236.9\text{ mA}$, hence choose a capacitor with ripple current rating = $0.5 I_{L(\max)} = 250\text{ mA}$
- Output Capacitor Selection:
 - Output ripple voltage,
 - * Considering ESR = 0.1 Ohm
 - * $\Delta V_o = \frac{V_o}{f_{sw}L} \left(1 - \frac{V_o}{V_{in}}\right) \cdot \left(\text{ESR} + \frac{1}{8f_{sw}C_o}\right)$
 $= \frac{3.3}{100\text{k} \times 3.3\mu} \left(1 - \frac{3.3}{5}\right) \left(0.1 + \frac{1}{8 \times 100\text{k} \times 22\mu}\right) = 533.2\text{ mV}$
 - * Considering ESR = 10 mOhm
 - * $\Delta V_o = \frac{V_o}{f_{sw}L} \left(1 - \frac{V_o}{V_{in}}\right) \cdot \left(\text{ESR} + \frac{1}{8f_{sw}C_o}\right)$
 $= \frac{3.3}{100\text{k} \times 3.3\mu} \left(1 - \frac{3.3}{5}\right) \left(0.01 + \frac{1}{8 \times 100\text{k} \times 22\mu}\right) = 227.2\text{ mV}$
 - Ceramic/MLCC caps: capacitance dominates, with
 $\Delta V_o = \frac{V_o}{8f_{sw}^2LC_o} \left(1 - \frac{V_o}{V_{in}}\right) = \frac{3.3}{8 \times (100\text{k})^2 \times 3.3\mu \times 22\mu} \left(1 - \frac{3.3}{5}\right) = 193.2\text{ mV}$
 - Tantalum/POSCAP/Electrolytic: ESR dominates, with
 $\Delta V_o = \frac{V_o}{f_s \cdot L} \left(1 - \frac{V_o}{V_{in}}\right) \times \text{ESR} = \frac{3.3}{100\text{k} \times 3.3\mu} \left(1 - \frac{3.3}{5}\right) (0.1) = 340\text{ mV}$

C.2 Switching Element

C.2.1 TRIAC

- TRIAC Snubber: The RC dV/dt snubber and current-limiting resistor shall use the recommended values of $R_1 = 51 \text{ R}$, $R_2 = 150 \text{ R}$, $R_{GK} = 390 \text{ R}$.
- TRIAC shall be always-triggered with a high signal from the MCU. This is to prevent unwanted delays from pulsed triggering when the AC current falls below the latching current threshold.
- The opto-MOSFET switching signal, and therefore the TRIAC gate signal, is in phase with the triggering of its isolated DC-side LED.
- 510 VAC MOV required for transient overvoltage protection of the TRIAC.

C.2.2 Opto-MOSFET

- Selection: AQY-280S
 - 5 mA recommended input current
 - 1.8 mA typical LED turn-on current
 - 1.6 mA typical LED turn-off current
 - 1.14 V typical LED turn-on voltage drop
- Consider the implementation of an NMOS-switched LED indicator with a 10 kOhm pulldown resistor at the NMOS gate:
 - $R_{eff} = R_1 // R_2 = \frac{10k \cdot 430}{10k + 430} = 412.27 \text{ R}$
 - Limiting Resistor Voltage Drop, $3.3 - 1.14 = 2.16 \text{ V}$
 - $I = \frac{2.16}{412.27} = 5.3 \text{ mA} \Rightarrow$ requirement fulfilled

C.2.3 LED Indicator Circuit

- $V_F = 1.9 \text{ V}$ typical, hence the limiting resistor is given by $R = \frac{3-1.9}{2m} = 550 \text{ R}$
- Power dissipation is for the following components:
 - Limiting Resistor: $(2m)^2(550) = 2.2 \text{ mW}$
 - LED: $2m \times 1.9 = 3.8 \text{ mW}$

C.3 Power Monitoring

Power computation does not use instantaneous values of voltage or current $v(t)$ or $i(t)$, since the resulting computation $p(t) = v(t)i(t)$ is still a sinusoidally-varying quantity which does not give a meaningful result for

determining power consumption. Instead, the root-mean-square (rms) voltage V and rms current I must be calculated. When sampled and quantized at discrete instances by an ADC, we obtain the discrete-time equivalents of the instantaneous voltage and current, $v(n)$ and $i(n)$. The resulting rms equations for our use case are:

$$I = \sqrt{\frac{1}{N} \sum_{n=0}^{N-1} i(n)^2}$$

$$V = \sqrt{\frac{1}{N} \sum_{n=0}^{N-1} v(n)^2}$$

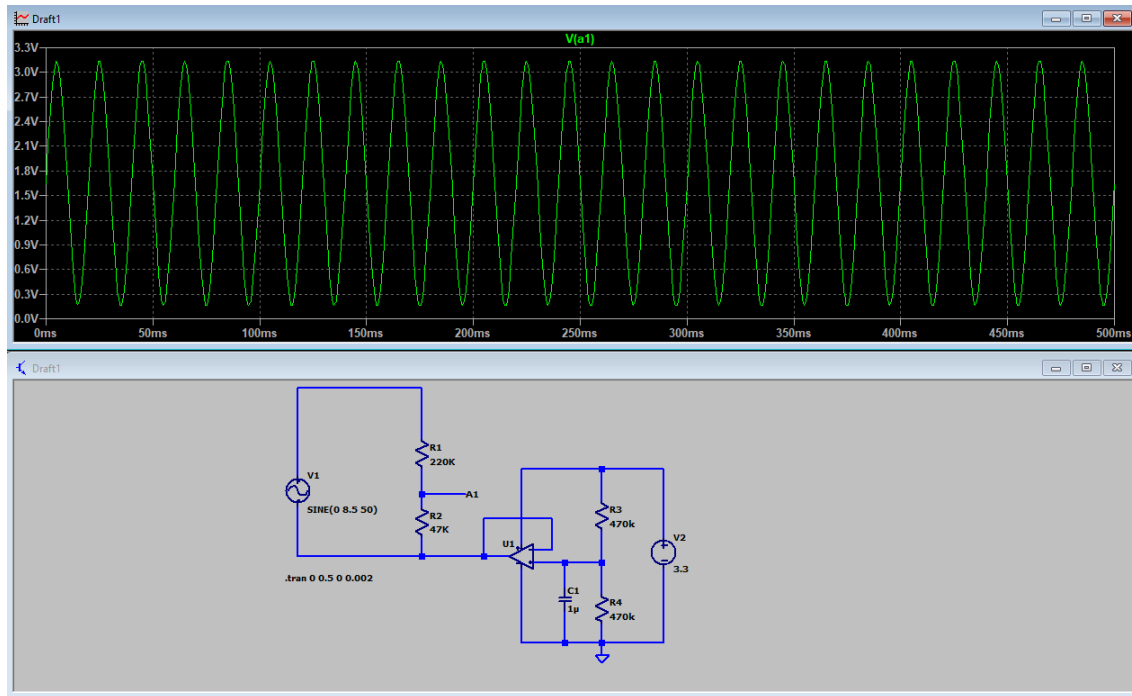
With both the rms voltage and current, the apparent power $|S|$ is computed as the product of the rms current and voltage magnitudes $|S| = VI$. In the power triangle, the complex power $S = \sqrt{P^2 + Q^2}$, where P is the real power associated with useful work done in the transfer of energy and Q is the reactive power associated with non-useful work, *i.e.* the storage and release of energy from an inductance's magnetic field or a capacitance's electric field. Hence for a voltage phase θ and current phase β , the phase difference $\phi = \theta - \beta$ determines the power factor $\text{pf} = \cos \phi$. It is ideal for the loads to operate close to unity pf, where the reactive power is zero, which is why its measurement is useful for a restaurant owner whose appliances are typically inductive (motors, coolant pumps, etc). The real power can be determined as the time average of the instantaneous power, and subsequently, the pf can be computed using the previously determined value of the apparent power.

$$P = \frac{1}{N} \sum_{n=0}^{N-1} v(n)i(n) = VI \cos \phi$$

C.3.1 Potential Transformer (Phase 2)

- For peak voltage consideration, $V_p = 6\sqrt{2} = 8.485 \approx 8.5 \text{ V} \Rightarrow V_{pp} = 17 \text{ V}$
- Max voltage input to ADC is 3.3 V, so required divider ratio achieved with $\frac{48.175\text{k}}{48.175\text{k}+200\text{k}} \times 17 = 3.3 \text{ V} \Rightarrow$ requires downsizing 48.175 kOhm to 47 kOhm to get 3.234 V_p
- Total voltage signal range is 3.234 V, with $\frac{(3.3-3.234)}{2} = 33 \text{ mV}$ buffer range at both upper and lower limits.
- To tolerate 10% variation in AC voltage, switch out 200 kOhm resistor with 220 kOhm resistor, hence:
 - $\frac{47\text{k}}{47\text{k}+220\text{k}} \times 8.5(2) \times 1.1 = 3.292 \text{ V}$
 - Normal signal range is: $\frac{47\text{k}}{47\text{k}+220\text{k}} \times 17 = 2.993 \text{ V}$
 - Revised buffer range is now $\frac{(3.3-2.993)}{2} = 153.5 \text{ mV}$
 - This tolerance adjustment is **critical** because we expect AC voltage to be hovering around 230 V (rms), unlike current, which has a wide variation with the connected AC load.
- Given a 12-bit ADC for the ESP32-C6 with a $V_{cc} = 3.3 \text{ V}$,
 - ADC Resolution = $\frac{V_{cc}}{2^{12}} = \frac{3.3}{4096} = 0.80566 \text{ mV}$
 - Sensitivity = 100 mV/A $\rightarrow 66.225 \text{ mV}$ (after divider)

- Smallest detectable current = $\frac{\text{ADC Resolution}}{\text{Sensitivity}} = \frac{0.00080566\text{V}}{0.066225\text{V/A}} = 12.2 \text{ mA}$
- Smallest load power detectable = $6 \times 12.2\text{m} = 73.2 \text{ mW}$
- Voltage Divider Current = $\frac{6}{267\text{k}} = 22.5 \mu\text{A}$
 - Power dissipation, 47kOhm resistor = $(22.5\mu)^2(47\text{k}) = 23.8 \mu\text{W}$
 - Power dissipation, 220kOhm resistor = $(22.5\mu)^2(220\text{k}) = 111.4 \mu\text{W}$
- The lower terminal of the secondary winding is biased by a 1.65V offset to make it readable at the analog pin. This bias is accomplished with a rail-to-rail op-amp, as shown below.



C.3.2 Current Sensor (Phase 2)

- Selection: ACS71240KEXBLT-030B5
- Output Voltage Range = $100 \text{ mV/A} \times 40 \text{ A} = 4 \text{ V}$ (0.5 to 4.5 V), which exceeds 3.3 V analog input range for ESP32-C6 MCU.

NOTE : there exists 2x 0.5 V buffer zones at the upper and lower ranges of the calibration curve.

- : Required divider ratio = 0.66, achieved with $\frac{100\text{k}}{100\text{k}+51\text{k}} \times 5 = 3.3113 \text{ V}$ (use 0.1% tolerance).
 - Total voltage range for a theoretical 5V original signal = 3.3113 V
 - Utilised sensor voltage range = $66.225 \text{ mV/A} \times 40\text{A} = 2.649 \text{ V}$, with 2 buffer zones each of 0.33115 V at the upper and lower limit.
 - Negative Peak Signal = 0.33115 V
 - Positive Peak Signal = $2.649 + 0.33115 = 2.98 \text{ V} \Rightarrow$ no overvoltage at ADC.

- Given a 12-bit ADC for the ESP32-C6 with a $V_{cc}=3.3\text{ V}$,
 - ADC Resolution = $\frac{V_{cc}}{2^{12}} = \frac{3.3}{4096} = 0.80566\text{ mV}$
 - Sensitivity = $100\text{ mV/A} \rightarrow 66.225\text{ mV}$ (after divider)
 - Smallest detectable current = $\frac{\text{ADC Resolution}}{\text{Sensitivity}} = \frac{0.00080566\text{ V}}{0.066225\text{ V/A}} = 12.2\text{ mA}$
 - Smallest load power detectable = $6 \times 12.2\text{ m} = 73.2\text{ mW}$

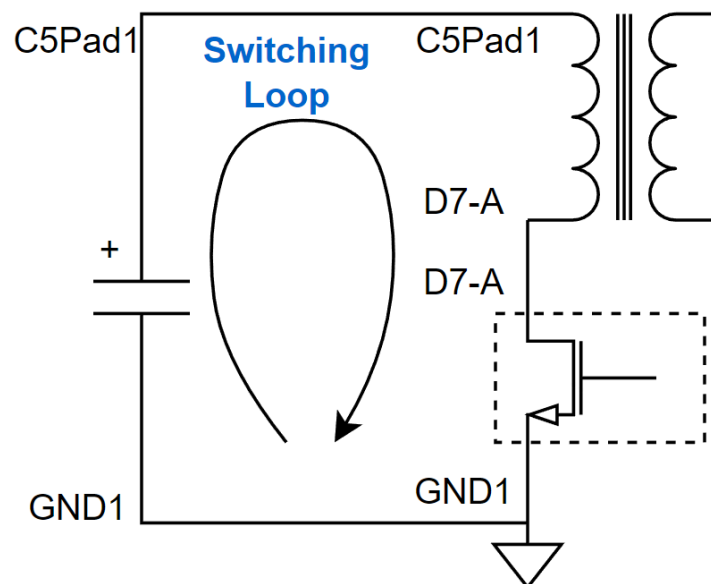
C.4 PCB Layout

The following layout guidelines were to be adhered to with regard to layout of SMPS circuitry:

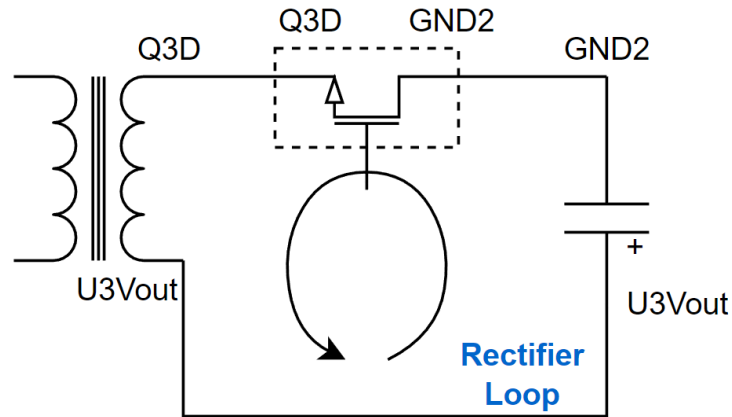
- Keep current loops tight and small. Place the high-current paths (GND, VIN, and SW) as close to the device as possible. Use short, direct, and wide traces.
 - Prioritize AC loops first (switching loop and rectifier loop)
 - Deal with input/output loops as a later priority
- High-frequency, high-current loops should be on the same side of the PCB
- Route SW away from sensitive analog areas such as FB to limit crosstalk and keep it short to reduce its ability to radiate EMI. No copper pours are allowed for the SW lines.
- Ensure all feedback, bypass, or compensation connections are short and direct; place the feedback resistors and compensation components as close to the IC as possible.

When targeting critical AC current loops for routing, the following were prioritized:

- Switching loop



- Rectifier loop



D Lighting Systems


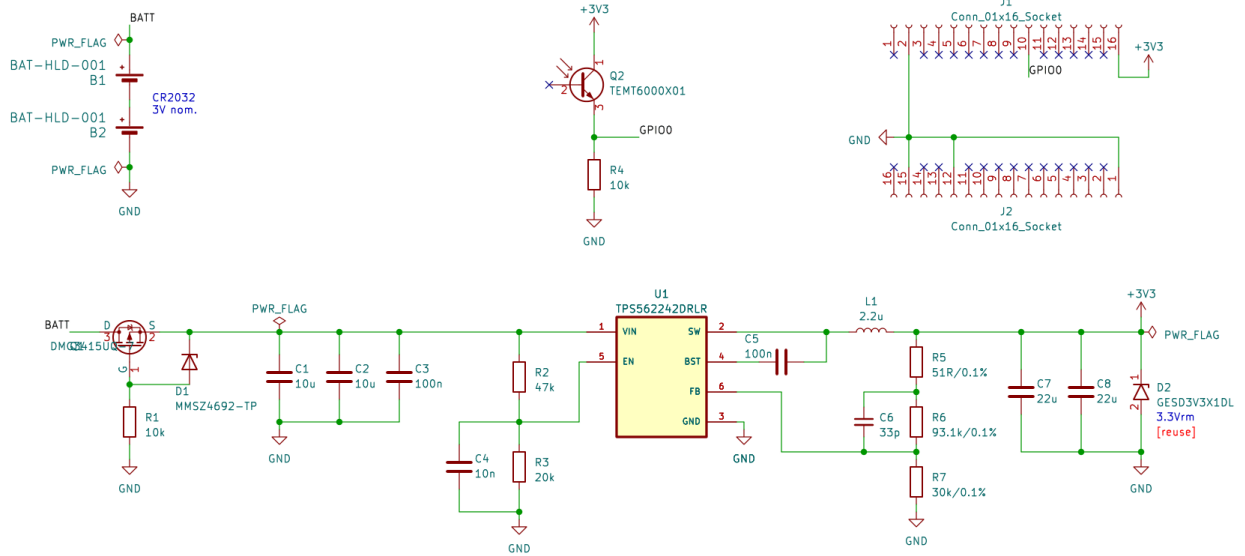
```

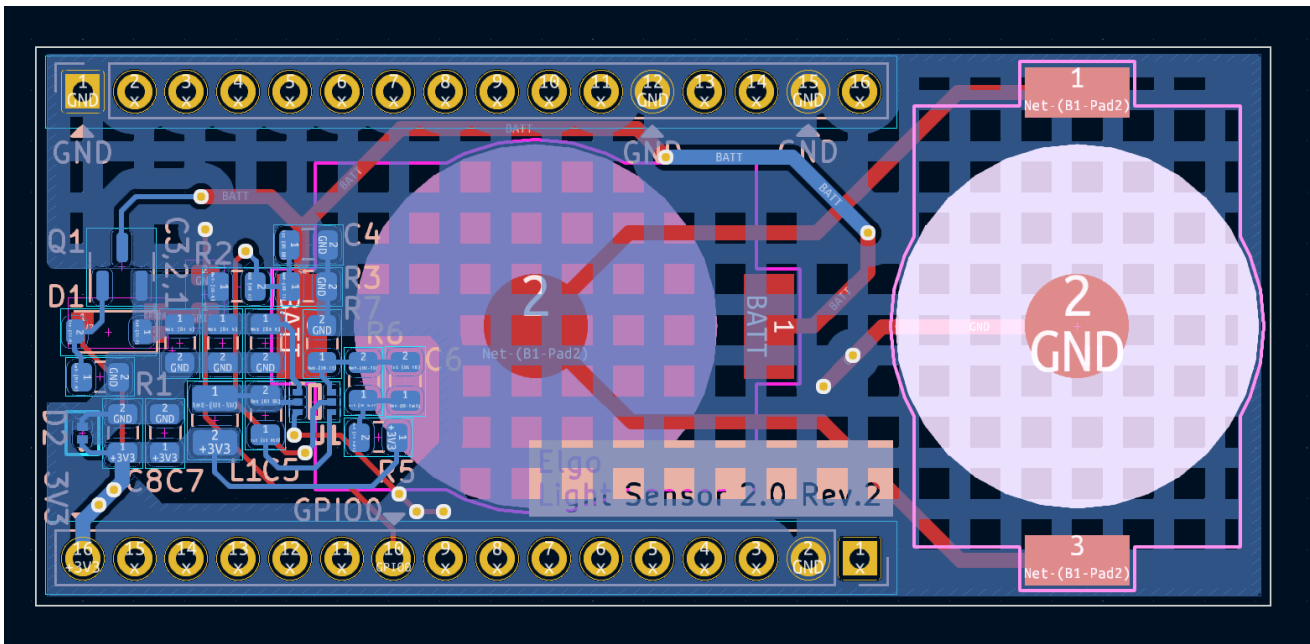
1 void execute_optimal_action(float switch_position, float optimal_action)
2 {
3     // Format switch_position and optimal_action to have two decimal places
4     char switch_position_str[20]; // Assuming maximum length of float representation
5     char optimal_action_str[20]; // Assuming maximum length of float representation
6     snprintf(switch_position_str, sizeof(switch_position_str), "%.1f", switch_position);
7     snprintf(optimal_action_str, sizeof(optimal_action_str), "%.1f", optimal_action);
8
9     // Construct the URL string with formatted float values
10    char url[200]; // Adjust size as per your requirement
11    snprintf(url, sizeof(url), "http://54.160.134.98:3000/loglightLevel?lightLevel_current
    =%s&plugID=shellyplusplugs-d4d4daec6c98&lightLevel_prev=%s", optimal_action_str,
    switch_position_str);
12
13    // Set up the HTTP client configuration with the formatted URL
14    esp_http_client_config_t config_get = {
15        .url = url,
16        .method = HTTP_METHOD_GET,
17        .is_async = false,
18        .cert_pem = NULL,
19        .event_handler = client_event_get_handler,
20        .timeout_ms = 10000};
21
22    // Initialize and perform the HTTP request
23    esp_http_client_handle_t client = esp_http_client_init(&config_get);
24    esp_http_client_perform(client);
25    esp_http_client_cleanup(client);
26 }

```

Listing 1: C code for executing optimal action

D.1 Light Sensor PCB Schematic





E Budget Allocation

Area of Expenditure	Budget Allocated (\$)	Amount Spent (\$)	Remaining (\$)
Software (e.g., UI tools, 2 AWS, Sagemaker)	3,500	442.12	3,057.88
Hardware (e.g., prototyping materials, tools & equipment)	4,500	2088.70	2411.30
Research (e.g., research tools, market validation expenses)	400	66.79	333.21
Miscellaneous (e.g., 4k challenge, biz development)	625.50	117.55	507.95
Total	9,025.50	2715.16	6310.34

Source of Funds	Amount (\$)
Capstone Fund	4,000
Stephen Riady Innovation Grant	5,000
4K Challenge Revenue	25.50
Total	9,025.50

F Individual Contributions

Names	Kaveri Priya Putti	Mohammed Fauzaan	Rohit Raghuram Murarishetti	Tan Xinlin	Sanat Khandekar	Vissal Natarajan	Zi Jie Chow
Midterm Presentation & Report - HVAC	●					●	
Midterm Presentation & Report - Lighting		●	●		●		
Midterm Presentation & Report - Kitchen Appliances	●	●				●	
Midterm Presentation & Report - Refrigerator	●	●				●	
Midterm Presentation & Report - Prototyping	●	●	●	●	●	●	●
Midterm Presentation & Report - Budgeting		●		●			

Names	Kaveri Priya Putti	Mohammed Fauzaan	Rohit Raghuram Murarishetti	Tan Xinlin	Sanat Khandekar	Visshal Natarajan	Zi Jie Chow
UI & WebApp Frontend			●				
WebApp Backend & Database		●	●			●	
WebApp Backend and Frontend Integration		●	●			●	
Cloud - Kitchen Appliances and Refrigerator	●	●				●	
Cloud - HVAC	●					●	
Cloud - Lighting		●	●		●	●	
Refrigerator Prototyping - Collective Anomaly		●					
Refrigerator Prototyping - Contextual Anomaly	●					●	

Names	Kaveri Priya Putti	Mohammed Fauzaan	Rohit Raghuram Murarishetti	Tan Xinlin	Sanat Khandekar	Visshal Natarajan	Zi Jie Chow
Refrigerator Prototyping - Point Anomaly	●					●	
HVAC Prototyping-Simulation	●					●	
HVAC Prototyping-Algorithm Development	●					●	
Lighting Prototyping			●		●		
Hardware - Lighting Prototyping			●		●		●
Hardware - Smart Plugs				●			●
CAD Modeling and 3D Printing	●			●			●
Augmented Reality Modelling	●						

Capstone Website			●				
Names	Kaveri Priya Putti	Mohammed Fauzaan	Rohit Raghuram Murarishetti	Tan Xinlin	Sanat Khandekar	Visshal Natarajan	Zi Jie Chow
Poster design	●	●	●				
Video Showcase	●			●	●		
Procurement and Financing	●	●		●		●	
Final Presentation & Report - Business			●	●			
Final Presentation & Report - HVAC	●					●	
Final Presentation & Report - Lighting		●	●		●		●
Final Presentation & Report - Kitchen Appliances	●	●		●		●	

Final Presentation & Report - UX		●	●				
Names	Kaveri Priya Putti	Mohammed Fauzaan	Rohit Raghuram Murarishetti	Tan Xinlin	Sanat Khandekar	Visshal Natarajan	Zi Jie Chow
Final Presentation & Report - Refrigerator	●	●				●	
Final Presentation & Report - Smart Plug				●			●
Final Presentation & Report - Budgeting			●	●			