MULTIVARIATE TIME SERIES MODELLING OF THERMOGRAPHIC DATA FOR SUBSTATION INSPECTION

A thesis submitted to The University of Manchester
for the degree of Doctor of Philosophy
in the Faculty of Science and Engineering

2022

Alastair Alan Straker

Department of Electrical and Electronic Engineering

Contents

Word Count: 28228

List of Tables

List of Figures

Abstract

This thesis aims to investigate long-term condition monitoring of power substation assets with thermal imaging techniques. Due to limited inspections and measurement techniques which are very susceptible to noise from environmental factors and human error, the inspection process can fail to detect fault pre-cursors in equipment. This thesis aims to contribute to this problem by monitoring equipment continuously, which collecting data on the electrical load and weather conditions that influence the measurements, in order to characterise the thermal response of equipment using LSTM modelling techniques.

Two scenarios are presented in the form of two experiments, comprising of electrically loaded overhead lines and cable terminations. These are energised for multiple day periods with realistic load patterns, with wind effects emulated by an industrial fan, while thermal images, electrical load and environmental conditions are collected. The data are used as input to linear regression and LSTM recurrent neural networks.

The work contributes the use of multiple low-cost non-calibrated thermal imaging sensors for data collection, and the novel application of LSTM recurrent neural network modelling methods to produce accurate time-series models of the thermal output of points of interest on substation equipment in a laboratory environment. It also contributes a large-scale experimental rig facilitating the long-term monitoring of high voltage power equipment at high currents in laboratory environments, enabling multi-directional thermal imaging monitoring. Lastly it provides a case study into the thermal behaviour of 66 kV cable-sealing ends when energised long-term.

Recommendations for further work are outlined, including extending data collection for longer periods, conducting long-term monitoring outdoors and utilising the generated models in implementing fault-detection methods.

Declaration

No portion of the work referred to in this thesis has been submitted in support of an application for another degree or qualification of this or any other university or other institute of learning.

Copyright

Acknowledgements

I would like to thank my supervisors, Dr. Carrasco and Dr. Podd, for their support, open-mindedness and senses of humour throughout our time working together.

My gratitude extends to our partners at EPSRC, National Grid and the High Voltage Laboratory at the University of Manchester, especially Dr. Cwikowski, whose enthusiasm for the work was infectious and inspiring, and also for the generous funding and support, without which the work would not have been possible.

I would like to express heartfelt appreciation for the support offered by those I hold closest. My partner, Beth, for her constant love and patience. Danny and Conor, for the runs, battles, walks and talks. The robotics lab, for making work feel like home. And finally, my Mum, Dad and sister, for always being there.

Chapter 1

Introduction

1.1 Background

Air-insulated power substations, such as that shown in Figure 1.1 are facilities for the distribution and management of electrical power. There are over 300 substations in the UK and inspecting them requires specialist knowledge of their operation. They vary in size and function but are invariably hazardous environments. High load on substation componentry means that accelerated degradation can cause equipment failures. Other risks include lethal electric shock. Inspection engineers survey the facilities with sensors, taking measurements as they go. They then use these measurements and the context within which they are taken to make subjective judgements on the health of the asset. Often engineers rely on comparisons between the three phases to determine if a fault might be present. The specialist knowledge and access to the equipment means that there are only a small number of engineers qualified to make these judgements and due to this, and the large number of facilities, inspections are generally completed on a three-monthly basis in the UK. This limitation on when inspections may be completed means that there is no way to control the environmental or load conditions the measurements are taken in.

There are 3 main sensing modalities used in evaluating asset health during inspection: thermal, partial discharge and gas in oil measurements. Of these, the former two are passive measurements and are both dependent on the environmental conditions. Thermal inspections are the only consideration for this work though there is potential for future expansion into other modalities, using similar data analysis techniques.

Figure 1.1: A National Grid Substation in North-West England

Thermography, or thermal imaging, is the practice of generating images which display an indication of the thermal output of the items in scene, rather than a representation of the physically visible features. Generally, this is completed by measuring infrared radiation emitted from a surface, as every surface with a temperature above 0 K emits infra-red radiation [31].

Thermal imaging has a multitude of applications. It is commonly used in security, hunting, maintenance and surveying, to present a non-exhaustive list. In maintenance and condition monitoring, it has been established as a standard measurement technique, due in part to the fact that hot-spots are a common precursor to the occurrence of a fault. This is especially true within power equipment maintenance, where a hot-spot caused by a loose connection or faulty insulation can clearly indicate a problem.

Thermal images convey only the surface temperature of equipment to inspection engineers. This is affected by various conditions, including the equipment load, environmental conditions and emmissivity. Hot-spots, a common precursor to faults, are most visible under heavy electrical load. A piece of equipment under heavy load, in cold and windy conditions may not appear to be heating significantly when thermally inspected, for example, therefore increasing the chance that a false-negative occurs. Long-term monitoring of equipment in varied environmental conditions should mitigate against this effect and allow better judgements of asset health to be made.

1.2 Motivation

The primary drivers of thermal output of electrical power equipment within substations are: electrical load, wind speed and direction, ambient temperature solar radiation, and precipitation. Due to the large range in size of equipment present in a typical substation, there is also a large range in thermal capacity between, for example, a bus-bar conductor and a circuit breaker. This means that the effects of the thermally influencing factors can take many hours to be visible to a sensor, and an understanding of the conditions at a given instance in time is therefore insufficient to understand the thermal output at that instance. Inspections are typically conducted as snapshots - the engineer will capture a small time period of sensor readings for a particular piece of equipment. Aging infrastructure means that access to live or recent electrical load data for a given circuit is difficult to secure at the time of inspection, and even with portable weather station equipment, it is difficult to accurately interpret the influence of both the electrical and environmental conditions for an inspection engineer. It is also impossible to conduct inspections in similar conditions to previous inspections, limiting the ability of the engineer to compare sensor readings to those of previous site visits.

Figure 1.2 shows a circuit breaker, a substation asset with high thermal capacity, appearing hot under the influence of solar radiation. Figure 1.3 shows a set of three cable terminations, with comparatively low thermal capacity, displaying hot areas corresponding to the cable conductors entering the underside of the termination units. These figures provide a demonstration of the variety of scenarios encountered by inspection engineers.

The combination of short-term inspections, long-term thermal influences, the variation in conditions in which inspections are undertaken and the imprecise nature of thermal imaging inspections mean that the inspection process can fail to detect fault-precursors. This incurs a monetary cost, as faults are more developed when they are finally detected, requiring downtime to diagnose and repair. Furthermore it is dangerous, as undetected faults may lead to equipment failure, causing further damage to infrastructure, property and potentially harming humans.

The intention of this work is therefore to investigate the utility of a continuous, long-term monitoring approach to thermal inspections of substations. This approach will allow the long-term effects of thermally influencing factors to be captured by the monitoring system, thereby allowing the thermal output to be

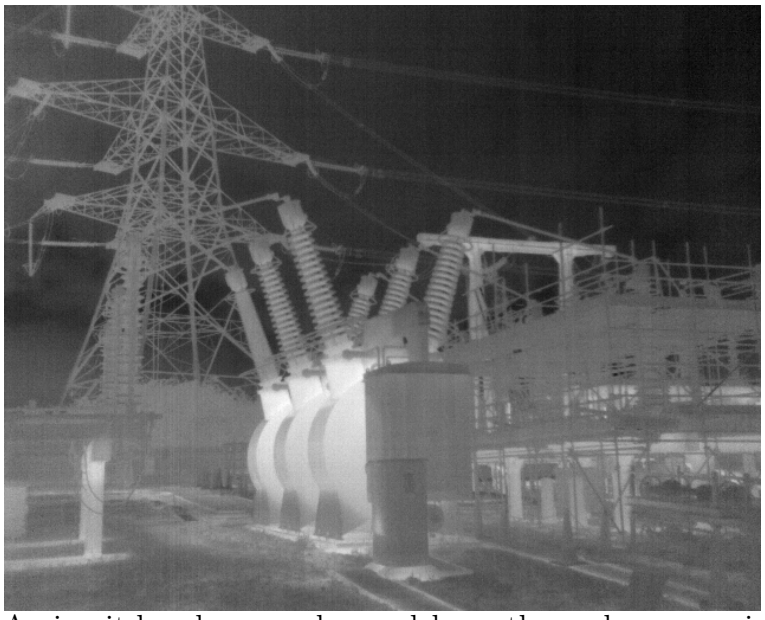

Figure 1.2: A circuit-breaker, as observed by a thermal camera, in a National Grid Substation

Figure 1.3: A set of three cable terminations, as observed by a thermal camera, in a National Grid Substation

modelled and understood. This will provide a baseline to use in the detection of fault-precursors in the case of data which do not fit the model output.

1.3 Aims

The specific aims of the project are as follows:

- gain understanding of the effects of environmental data on thermal images;
- to develop well-performing multivariate statistical models of thermal output from thermal cameras;
- to characterise directional or asymmetrical heating and environmental effects in power engineering assets.

The project objectives set in order to achieve these aims are:

- specify and implement experimental procedures in order to gather long-term thermal and environmental monitoring data;
- investigate and select statistical time-series modelling techniques;
- critically analyse laboratory data to inform modelling process;
- create models based on laboratory data.

Given the prospect of future long-term thermography monitoring, the high performance of LSTM recurrent neural network technology in time-series modelling, and the potential to frame the thermal output of substation equipment as multivariate time-series systems, it is proposed to:

- generate long-term continuous time-series thermographic data of electrically loaded power equipment;
- gather appropriate electrical and environmental conditions;
- create LSTM recurrent neural network time-series models with the data.

This would constitute a contribution to knowledge, exploring the use of modern time-series modelling techniques on long-term thermal, electrical and environmental data.

The delivery of these aims and objectives is presented in this thesis.

1.4 Contribution

The work presented in this thesis is funded by an EPSRC iCase studentship in collaboration with National Grid. It therefore comprises an academic contribution to knowledge, and industrial contributions to the sponsor.

In the academic context, the project comprises multiple contributions. Long short-term memory (LSTM) recurrent neural networks are applied to long-term, multivariate, environmental and electrical loading laboratory data, to create statistical time-series models. This method of time-series modelling on multivariate thermal image and environmental data is not found in the literature, and where other methods are used, the LSTM recurrent neural network models typically outperform them. Furthermore, the work models the directional effects of an environmental factor, specifically wind. It is additionally novel that this work utilises low-cost, non-calibrated equipment.

Previously undocumented insights are gained into the temperature profile of a pair of specific 66 kV cable-sealing ends at high current load, before and after one of the assets under-went an accelerated aging process. Differences are detected between the two thermal output of the cable-sealing ends afterwards, implying such aging processes generate measurable differences in output.

In the industrial context, the project supports National Grid in the development of new condition monitoring practices. With the advent of Industry 4.0, big-data, the Internet of Things and robotic inspection, continuous multivariate data collection in substation environments is inevitable. Specific benefits of automated continuous monitoring include:

- removal of the requirement for engineers to enter the hazardous environment for inspections;
- extra capacity for specialised inspection engineers to focus their attention on high-importance cases;
- repeatability of measurements;
- capability for data-fusion with environmental data;
- earlier warning of deviations from ‘normal’ output;
- enhanced consistency between different substations.

Determining a potentially useful application for data gathered in such a manner is valuable to the sponsor, as increasingly affordable thermal cameras facilitate extensive monitoring coverage of substations. Additionally, complex modelling solutions are increasingly accessible through open-source APIs and GPU computation. Furthermore, there are inherent difficulties with condition monitoring using thermal imaging. The methods presented in this work can result in costly false-positives and false-negatives, depending on the load and environmental context, and the work presented here assists in mitigating against that.

1.5 Scope

The scope of this project is to analyse the potential of the application of LSTM recurrent networks in modelling time-series data acquired in a laboratory setting. There is no analysis on data acquired from live assets in in-service substation environments. The project explores the performance of a single modelling method, which is compared to a traditional method in the early stages of the project. Due to their inherent high price, and prioritising features such as resolution, over absolute temperature measurement, the scope of the work is limited to non-calibrated thermal cameras. Full justification of this choice is made in Chapter 3. No attempt is made to present a comprehensive review of time-series modelling methods, though the justification for the choice of the LSTM recurrent neural network method is present in Chapter 3. Furthermore, the work does not include a comprehensive exploration of LSTM methods, instead it focuses on finding an implementation which produces satisfactory results given the context of model-based fault detection.

1.6 Thesis Outline

The thesis is arranged in five chapters excluding this introduction. First, the literature relevant to the work is reviewed and summarised, before placing this work in the context of that literature.

Chapter 3 presents the equipment and software used throughout the thesis, before detailing an exploratory experiment, consisting of two overhead line (OHL) conductors in circuit with a high-current DC power supply, an industrial fan, a single un-calibrated thermal camera, environmental weather station sensor and a

controller unit. This work was presented at Thermosense XLI 2019 [55].

Chapters 3, 4 and 5 contain the technical contribution of the thesis. Chapter 3 details an exploratory experiment, consisting of two overhead line (OHL) conductors in circuit with a high-current DC power supply, an industrial fan, a single un-calibrated thermal camera, environmental weather station sensor and a controller unit. This work was presented at Thermosense XLI 2019 [55]. Chapter 4 presents a development of the concepts presented in Chapter 3. A pair of 66 kV cable sealing-ends in a current loop, electrically loaded by a high-current transformer, are observed from all angles in the horizontal plane by two sets of four un-calibrated thermal cameras, providing the opportunity to observe any asymmetric heating and environmental effects. Chapter 5 presents a case study, exploring the utility of the non-calibrated cameras within the context of condition monitoring of a 66 kV cable sealing end. The chapter, utilising data-sets captured for Chapter 4, presents analysis of the heating profile of the cable sealing ends under load.

The final chapter of the thesis consolidates the discussion presented in the technical chapters, placing each within the wider context of the thesis. It summarises the contributions made, detailing how they constitute a development of the current literature. Finally, recommendations are made for future work that could further advance the work presented here.

Chapter 2

Literature Review

2.1 Introduction

This chapter presents the literature review conducted in the context of the work. This literature review primarily seeks to establish the state of the art in infrared thermography, or thermal imaging, condition monitoring and fault detection, specifically in the context of power equipment and outdoor, air-insulated substations.

Section 2.2 provides a review of the condition monitoring strategies in place in substations. Section 2.3 presents an overview of thermal imaging broadly, while Section 2.4 presents a review of the use of analysis techniques with regards to thermal imaging both in the context of substation condition monitoring and otherwise.

Given the review completed, an area of possible work is identified.

2.2 Substation Condition Monitoring

Electrical substations operate at high voltages and currents, so when failure occurs it can be catastrophic. Predictive and Preventative Monitoring (PPM) [9] and Condition Based Monitoring (CBM) [48] are the two most prominent preventative monitoring and maintenance methodologies. In the UK, maintenance and monitoring in National Grid facilities is condition-based [22] [14], with regular inspections providing data on equipment condition and inform maintenance strategy. Han and Song [15] present a review of conditioning monitoring strategies commonly utilised in the electrical context.

This project aims to investigate long-term monitoring as a method of condition monitoring. Due to the nature of working in high-voltage environments, clearance from live equipment is a requirement [47] Three main monitoring modalities in power transformers are identified – hot spots, gas in oil, and PD. Hot spots and PD, can be remotely monitored through thermography and antenna arrays respectively – they require no physical contact with the device. Issues with hot spot monitoring via thermography are identified as requiring internal sensors for internal equipment temperatures, or alternatively a good model of the equipment to predict the external thermal radiation. PD monitoring research has been focused on identifying the symptomatic signals from interference, and localising the source of the signals. On-line calibration of sensors has been identified as key for successful on-line monitoring, regardless of the sensor type.

A potential cause of the apparent lack of long-term or constant thermal monitoring in substation environments is simply the lack of infrastructure. Many substations are decades old, can be very large and include many assets. With thermal cameras remaining expensive, sufficient coverage could be prohibitively expensive. There is evidence of on-line monitoring in substations that does not yet include thermal imaging [30]. However, the emergence of internet-of-things, 'Industry 4.0', smart-grid and substation-specific standards on automation (IEC 61850) [38] [46] have led to work detailing plans for system architecture for such a thermography installation [61]. This provides clear motivation for exploring both low-cost thermal camera technology and long-term condition monitoring.

Robotic implementations of continuous monitoring exist, for example, Wang et al. [64] developed the SmartGuard robotic inspection system with the Shandong Electric Power Institute. The institute is a leader in the field of substation robotics, with numerous robots installed in active duty in substations throughout China. SmartGuard features autonomous navigation, autonomous battery charging and autonomous equipment recognition. SMP robotics offer the S3 Electrical Substation Robot [51]. The product features autonomous control and navigation, a Wi-Fi communications link and a data-acquisition platform consisting of a thermal camera and a visible-light camera on a 360 degree pan-tilt platform.

2.3 Thermal Imaging

Thermal imaging is a key aspect of substation inspection. Initial work in the field was focused around manual inspections by trained professionals, providing guidance on effective use of equipment and interpretation of the data collected.[11] [18].

Quantitative analysis is possible with radiometric thermography equipment but prohibitively high costs mean that most analysis is qualitative. [11]. This is partially because the sensors used in infra-red cameras (which are the most common type of thermograph) are sensitive to environmental conditions such as temperature. Generally two forms of thermal camera are available: calibrated and non-calibrated. In long-wave infra-red thermal cameras, the most common form of sensor array is a VOx micro-bolometer, these are typically used for substation inspection. These sensors respond to thermal flux incident on the sensor array, which is proportional both to the temperature of the imaged surface and the temperature of the sensor. A calibrated camera has undergone a procedure to characterise the output of the sensor at a given range of usable operating temperatures. A non-calibrated camera has output which fluctuates with sensor temperature.

An aspect of thermography into which there has been a lot of research is the effects of external factors on the readings provided by the sensors. As substations are often outdoors and exposed to weather, these effects can mask issues [11] and cause false negatives. Particularly, work has been carried out focusing on modelling the effect wind [3] and the equipment's electrical loading [4] on thermal images. Emissivity has also been found to be key in thermography. Small errors in the calculation of emissivity have been found to lead to large measurement errors [53]. Another paper has found a near-logarithmic relationship between emmissivity and temperature reading [60]. Findings have shown that an accurate emmissivity estimation for the characteristics (wavelength) of the camera in order to get accurate readings between targets with different finishes [39].

2.4 Thermal Image Processing and Modelling

2.4.1 Thermal Imaging Image Processing

Bagavathiappan et al. [1] present a review of applications of infra-red thermography in condition monitoring. This covers the technology, application domains, the relevant standards used as reference for qualitative analysis, and various analysis techniques used to extract information from the images. Consideration is given to both active and passive thermography. Given the presence of live high-voltage equipment in the context of substation monitoring, only passive thermography is relevant to this work.

Jadin and Taib present a further summary [26], Usamentiaga presents another in the context of non-destructive testing [61] Consistent with their findings, it has been discovered that the literature covers the image processing techniques used to segment thermal images into regions of interest, or background/foreground. Jadin and Taib [26] acknowledge that qualitative measurement can be sufficient for successful condition monitoring, suggesting that lower-cost non-calibrated sensors may be utilised. They also state that automated fault detection in the field of electrical power applications is still in the early stages, providing justification for this work. As thermal intensity is represented by pixel intensity, standard thresholding operations are common for this application [45] [8] [27] [25].

Duarte et al. provide a summary of these techniques in a medical context, which is also relevant here [10]. Much work has been undertaken to computerise the processing of the raw thermography data as captured by the sensor. A common first step is to 'segment' the thermal image using the otsu threshold [45]. This is an adaptive threshold, and is used to flexibly remove unwanted data from an image. [8]. Jung [27] uses an thresholding method in their work superimposing thermal images on-top of visible light images, while Jadin [25] [24] presents a brief investigation of thresholding techniques, and their own bespoke solution.

2.4.2 Fault Detection Analysis Techniques

Modern machine learning techniques can be broadly separated into two categories: image (matrix) analysis using convolutional neural networks and sequence (text, speech or time-series) analysis using recurrent neural networks. This section seeks to establish the state of the art in both with regards to applications in condition

monitoring.

As image segmentation can be framed well as a classification task, machine learning techniques and deep neural networks are commonly used to identify regions of interest. Li [34] uses a support vector machine (SVM) to recognise the equipment in-frame. Zhao [68] uses a SVM to detect insulator strings in images, albeit after a pre-processing method they have termed as Binary Feature Pooling. Attempts have been made at automating the interpretation of thermography data since 1997. Moja [42] attempted this using statistical and neural network techniques, correctly identifying the problem with simply comparing a thermal image with a reference image. There are many examples of using machine learning techniques to classify potential faults in thermal images of substation equipment [59] [58] [12] [37] [49] [19]. These methods typically consider individual frames of scenes, extracting features and performing statistical analysis to determine whether a fault is present in the scene. Limitations of this work include availability of data and consideration of environmental factors influencing the images. These methods have made significant progress in automating the process of identifying faults in good thermal images, however do not necessarily improve the ability to extract information from images in difficult conditions. Furthermore, the utility of the methods presented is limited to snapshots in time, and does not consider the potential of long-term monitoring in order to detect less obvious faults. It is clear from this survey of the literature that the 'classification' task is widely studied with regards to thermal imaging fault detection.

While it is common for individual thermal image frames to be used to extract information, there is merit in using thermal cameras to capture multiple frames in order to observe thermal changes with time. The work by Bortoni et al. [3] [4] considers sequences of thermal images, extracting values and framing them as time-series in order to perform regression analysis. This work provides a basis for exploring long-term monitoring of equipment using thermal imaging, framed as a time-series problem, however has limitations in lacking a thorough exploration of the accuracy of the models and the consideration of multi-directional environmental effects. This type of time-series analysis of thermal images is common in active thermography, where a material's response in time to a thermal impulse determines the outcome of the experiment, and is frequently used in building inspection [6] [7] [13]. This method is also commonly found in medical applications, which are reviewed by by Lahiri et al. [32]. There are well-performing

examples of recurrent neural network technology being utilised with time-series data from thermal cameras. Fang and Maldague [12] present a recurrent deep-learning based method for quantifying defect depth in carbon fiber reinforced polymer while Wang et al. [65] present the use of thermal images as time series, fed through LSTM recurrent neural networks, suggests that the memory function of the long short-term memory (LSTM) [17] type allows it to learn signal characteristics. Li et al. [34] present a LSTM based method with 99% accuracy at anomaly fault detection. These pieces of work suggest that there is merit in exploring LSTM based modelling techniques in the condition monitoring context, developing the work of Bortoni et. al. to develop accurate models with modern techniques. However, it is notable that the fields of medical analysis and building inspection concern controlled environments, where external environmental influences on the sensor equipment are minimal. Furthermore the equipment is often highly specified in order to be able to detect rapid thermal transients. This draws into question the applicability of these methods when using low-cost thermal cameras in less-controlled environments. There is no published literature concerning the application of LSTM neural network techniques to thermal imaging in the condition monitoring context, providing a clear gap in the current knowledge.

2.5 Summary and Determination of Contribution

It has been found that extensive research into the application of thermography has been completed. A large body of literature exists with regards to the application of thermography in a substation monitoring context, and there are many other applications. The nature of the research differs between application areas, with a focus on classification of faults and image segmentation in the substation context. Limitations to thermography have been explored, highlighting emissivity and the influence of environmental factors as key problems, including the aspect of human error in determining the extent of the effect of each. Long-term thermal monitoring is uncommon to find in practice, however proposals, standards and architecture are published detailing the form this may take.

Chapter 3

Modelling of an Overhead Line

3.1 Introduction

This chapter presents the details of the sensors and software selected for use throughout the work, in Sections 3.2, 3.3 and 3.4. An experimental scenario regarding the monitoring and modelling of an electrically loaded overhead-line conductor is introduced and discussed. The experiment design is described in Section 3.5, the output datasets in Section 3.6 with results and analysis presented in Section 3.7. Discussion of the results is presented in Section 3.8.

This chapter presents the results obtained using data from the first of two completed experiments over the course of the work. Raw data output is presented and described, with notable features highlighted. Example images from data points selected at random are presented in detail. A detailed description of the configuration of the modelling techniques utilised is provided.

The aim of the work is to evaluate whether good auto-regressive and LSTM models can be created based on data of this type from these sensors. The successful creation of such models would support the viability of the modelling work and provide a foundation from which more complex electrical power systems could be modelled, If model performance is satisfactory, it would suggest that such methods could be employed in an industrial setting, where model-based fault detection techniques could be used in order to detect potential faults in equipment autonomously, and potentially earlier than currently used methods achieve.

3.2 Sensors: Thermal Imaging

3.2.1 Thermal Imaging

This section will provide a functional overview of typical thermal imaging technology, focused on areas relevant tp the context of the usage of thermal imaging over the course of the work.

3.2.1.1 Microbolometer

A bolometer is a device used to measure electro-magnetic radiation incident on itself using temperature-dependent electrical resistance. A microbolometer, the primary sensor type for uncooled thermal imaging cameras, is an array of small bolometer sensor elements, sensitive in the 7-14 μm frequency range. Typically radiation from the environment is focused onto the sensor plane using infrared-transmitting optics, typically germanium. The incident radiation causes small changes in the resistance of the sensor elements it hits. Sensor elements are chosen due to their high coefficients of thermal resistance. Lower-cost long-wave infrared (LWIR) thermal cameras utilise vanadium oxide (VOx) or amorpous silicon sensor elements. Sensors made from these materials do not require active cooling for operation, resulting in lower cost. At higher price-points, actively cooled medium-wave infrared (MWIR) and short-wave infrared (SWIR) cameras are available, commonly used for gas-leak inspection and very long distance (>10 km) thermal detection through smog and/or clouds respectively. LWIR thermal cameras are most frequently used for industrial inspection and other similar applications. [21]

Measurement circuitry detects the changes in resistance of the sensor, allowing the signals to be digitised and to enter the image pipeline of the device. Typically at this stage the image undergoes a sequence of processing stages with two purposes: to correct the image (for bad pixels, lens distortion and non-uniform areas), apply gain (for enhancing visible contrast). Finally, any annotations, zoom and colourisation procedures are performed on the image and it is output to the user, via direct display, a communication interface, or both.

Images are output as a matrix of values proportional to the thermal radiation incident on the sensor elements. This quantity is referred to as 'digital counts' of thermal radiation. Digital counts are measured per pixel, allowing variations in the number of digital counts incident on the sensor from across a physical scene, to be displayed as an image of varying thermal intensity. When multiple

images are captured of a scene over time, both the spatial and temporal thermal variations can be measured.

3.2.1.2 Sensor output

The raw sensor output is proportional to thermal flux incident upon itself. Heat flux in 1-D, ϕ_q, at the position x, is defined as:

$$\phi_q(x) = -k\frac{dT(x)}{dx}, \tag{3.1}$$

where T is the absolute temperature. More specifically, the equation for radiative heat transfer between two objects (a with temperature T_a and b with temperature T_b) is defined as:

$$\phi_q = \epsilon\sigma F(T_a^4 - T_b^4) \tag{3.2}$$

where ϵ is emissivity factor, σ is the Stefan-Boltzmann constant, F is the view factor between the two surfaces and T is absolute temperature.

As flux is proportional to the difference in temperature between the two surfaces, it can be inferred that thermal cameras are not only sensitive to the temperature of that which they are measuring, they are sensitive to the sensor temperature. Provided the imaged scene maintained a stable temperature, if the temperature of the sensor increases, the flux incident on the sensor would decrease, resulting in a lower apparent scene temperature. This is an especially important consideration for thermal cameras mounted in assemblies with other electronics, especially if dissipating a large amount of power, or in otherwise thermally volatile environments. Thermal camera manufacturers correct for this effect varying amounts, further details are provided in Section 3.2.1.3.

Referring back to Equation 3.2, the other primary influencing factor is emissivity, ϵ. Emissivity is defined as the efficiency at which a surface will emit thermal radiation. An ideal emitter, a blackbody, has an emissivity $\epsilon = 1.0$, while all real materials have emissivities $0.0 < \epsilon < 1.0$, given by the ratio of radiation emitted from that surface to radiation emitted from a blackbody of equal size, as given by the Stefan-Boltzmann law. Emissivity is influenced by both the surface material and the surface finish. Given two objects of equal temperature, size, shape and material, if one is polished and another has a rough finish, the emissivity will differ markedly. For example, polished copper has an emissivity of ~ 0.04, while oxidised copper has a value of ~ 0.87 [44]. Emissivity is non-trivial to measure,

requiring controlled experiments to determine a value for a material. Generally, in the context of thermal imaging, emissivity is estimated for objects in scene using a table of reference of typical emissivity values for materials. For this reason, it can be a large source of error in the measurement of temperature using thermal imaging [40].

Objects are reflective at infra-red wavelengths in a similar manner to at visible wavelengths. This provides another important consideration for thermal imaging applications, as particularly reflective objects in a scene can appear to be hotter than they are due to being surrounded by other hot objects emitting thermal radiation onto them. This can be particularly problematic in outdoor settings where the infrared radiation from the sun being reflected from surfaces can cause significant differences in apparent temperatures.

Further factors which influence the sensor output due to an effect on infrared transmission include:

- humidity;
- precipitation;
- view distance;
- geometry of imaged scene.

Of this list, humidity, precipitation and view distance pertain to the environment through which the thermal radiation must travel between the imaged object and the sensor. Any particles present in the path between the object and the sensor contribute to both scattering and absorption of thermal radiation. As humidity and precipitation increase, the number of particles in that environment increases, resulting in an equivalent increase in absorption and scattering. The effect of these variables is proportional to the view distance. The geometry of the scene influences the apparent thermal output, firstly due to the fact that complex geometries with many heated components can cause large numbers of thermal reflections to occur. Secondly, geometry must be considered due to the effect it has on the view factor between the sensor and the imaged scene. As the quantity of radiation emitted from a surface changes markedly with angle [41], both the geometry of the scene and the position of the region of interest (ROI) within the sensors field-of-view (FOV) can impact the appearance of the scene.

In interpreting thermal images, it is important to consider the numerous thermal influences on the imaged scene, especially when the intention is to interpret the level of thermal output from an object in a scene. The following environmental variables can have a significant direct influence on the temperature of an imaged object:

- wind speed and direction;
- ambient temperature;
- precipitation;
- solar radiation;
- in-scene active heating.

These factors change the temperature of the imaged scene rather than influencing how the temperature appears to the sensor, however, it may be the case that they impact an assessment of a thermal image. Particularly important is solar radiation, as this is susceptible to directly impacting the temperature of objects while also creating adverse reflection effects.

3.2.1.3 Calibration

Thermal cameras are available from manufacturers with varying levels of calibration. The most basic, low-cost cameras are have minimal calibration completed before shipping, while more expensive products have extensive in-factory calibration completed. For example, FLIR Systems, one of the larger manufacturers of thermal cameras, generally provides, what it refers to as 'radiometric' and 'non-radiometric' cameras. Calibration is a procedure involving capturing thermal images of a device with known temperature and emissivity (a black-body source) with the sensor at a known temperature to determine the relationship between the sensor core and the output digital count value. Once this relationship is quantified, it can be applied to raw output thermal images, allowing the surface temperature of the viewed scene to be estimated from the measured incidental flux on the sensor. After radiometric calibration, thermal cameras offer a temperature accuracy of ± 2.5 K. They require user input of measurements or estimates of 'scene emissivity, atmospheric temperature and transmission, background temperature, and parameters to account for the recommended window

included on a fully integrated system' [57] in order to reach such accuracy. Each of the user-provided quantities is a potential source for further error. Calibrated cameras vary in whether every pixel is calibrated for, or they have a central area which can be used for temperature measurement in the centre of the frame, called a spot-meter.

Non-calibrated thermal cameras typically provide only a digital count value per pixel to the user as output. As detailed in Section 3.2.1.2 there are a number of factors that can influence this digital count value, making it difficult to accurately estimate a scene temperature value from it.

3.2.2 Camera Choice: FLIR Boson

3.2.2.1 Core Choice

Thermal camera core choice was made based upon the following factors:

- affordable enough to buy on the project budget;
- availability of sensor through official university suppliers.
- prior expertise with FLIR sensors in research group;
- sensor resolution;
- access to the image pipeline - providing the raw sensor data to the user before gain is applied.

The methodology employed by inspection engineers to extract a spot-temperature reading from a thermal imaging device includes making a number of assumptions. A typical inspection camera, for example a FLIR T620, requires inputs from the user regarding relative humidity, ambient temperature, emmissivity of the equipment and distance to the equipment. Further factors influence the appearance of a thermal image: angle of surface to sensor and thermal reflections, for example. The potential error in determining each of these factors, especially emmissivity, which is difficult to estimate accurately, introduces uncertainty in the final temperature output of the camera. The experimental set up allows the above factors to be assumed as constant or measurable between sequential images, and as the primary focus of the experiment is thermal response over time, this is considered satisfactory. A summary of the features and capabilities of the two types of cameras is displayed in Table 3.1.

Table 3.1: Camera Features of the consumer-level FLIR T620 and OEM FLIR Boson

Spec	T620	Boson
FOV (°)	15x11	50 (horizontal)
Resolution	640x480	320x256
Temp Range (°C)	-40 to +150	to +140
Thermal Sensitivity/ NETD (mK)	40 @30°	60
Atmospheric transmission correction	Automatic, based on inputs for distance, atmospheric temperature and relative humidity	None
Emissivity Correction	Variable from 0.01 to 1.0 or selected from materials list	None
Automatic Image Adjustment	Continuous, histogram based	User configurable
Cost (£)	17,548.00	1,325.00
Radiometric	Yes	No

3.2.3 Usage: Pre-processing, Normalisation and Image Display

Due to lack of provision of telemetry (which provides access to camera core temperature) on the USB channel of the FLIR Boson communication interface, measures were introduced to mitigate against temperature transients that the cameras may undergo during operation.

Potential temperature transients include: self-heating, experimental forced environmental conditions (wind), ambient temperature change, direct heating from the laboratory climate control and heating from the experiment conductor.

The laboratory within which the overhead-line experiments were conducted was found experimentally to have relatively stable environmental conditions (see Section 3.6, therefore the direct environmental influence on the camera itself was assumed to be negligible for the purposes of these experiments.

Figure 3.1 shows the FLIR Boson camera within it's 3D printed mount, alongside prototypes of the mount.

FLIR Boson camera cores have an automatic histogram equalisation functionality (Automatic Gain Control or AGC in FLIR terminology) [56], designed to increase visibility of thermal features in-scene. This is computed based on the

Figure 3.1: FLIR Boson thermal camera core in custom mount, alongside prototype mounts.

range of thermal intensities per frame. Examples of images with and without the equalisation process applied are presented in Figure 3.2, with significantly more features visible for human inspection in the equalised image. Corresponding histograms of the images are presented in Figure 3.3. Due to the desire to compare sequences of frames of thermal images in the experiments presented here, all images were retrieved before AGC was applied in the image pipeline.

3.3 Additional Sensors

This section provides details on the equipment employed to collect data other than thermal images over the course of the project.

3.3.1 Weather and Environmental Data

A Vaisala WXT520 [62] weather station was used to collect data about the environmental conditions in which the experiments were completed. The weather station can collect air temperature, humidity, pressure, wind speed and direction and precipitation data. For the purposes of this project, all data but precipitation were collected. Data sample rates were as default for the sensor unit.

Data are output via a USB serial link to a controller device. The controller device, when required to record the data, called a blocking routine awaiting the most recent data from the sensor device, recording it when received.

This weather station was selected as it is of a good standard and representative

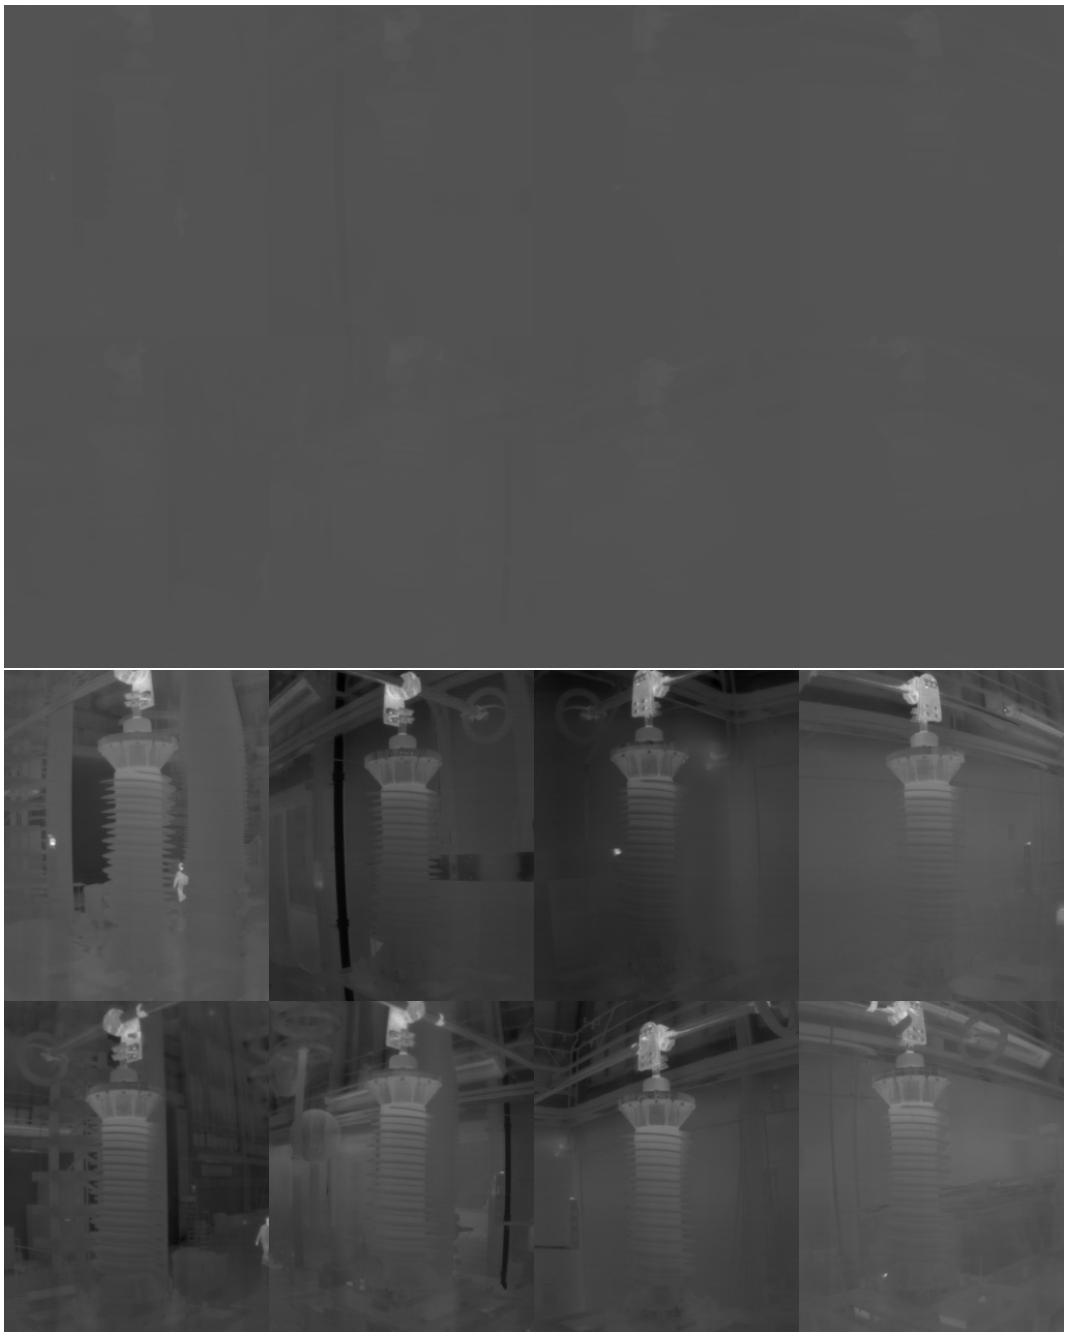

Figure 3.2: Top: a grid of eight non-normalised thermal images; Bottom: a grid of 8 normalised thermal images.

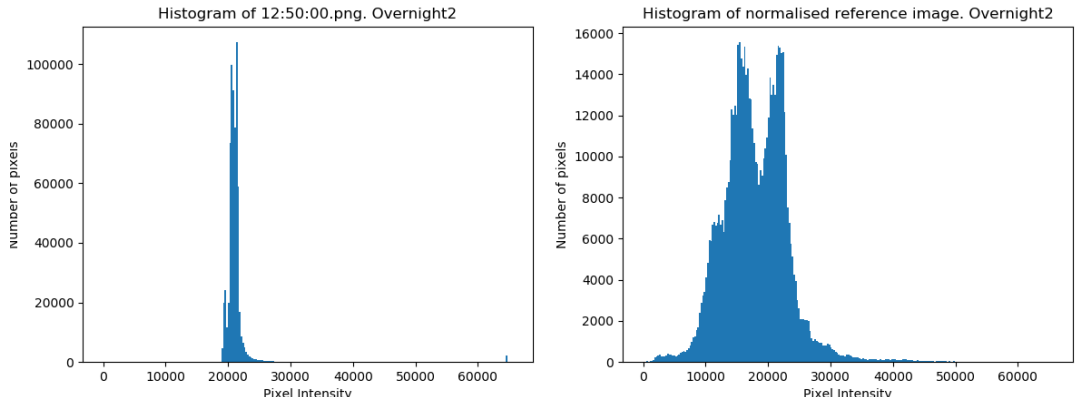

Figure 3.3: Left: histogram of a non-normalised sample thermal image; right: histogram of an independently normalised thermal image.

Table 3.2: Data output units for Vaisala WXT520 weather station

Data type	Unit
Air Temperature	°C
Humidity	%
Pressure	Bar
Wind speed	m/s
Wind direction	°
Telemetry	N/A

of what may be used in industry to collect environmental data where required.

Output data units are detailed in Table 3.2.

3.3.2 Current

Measurement of current over the course of the OHL experiment was completed by a Magna Power [23] MSA16-2700 high current power supply. The controller single-board computer (SBC), a Raspberry Pi, interfaced with the power supply using an RS-232 serial link. The power supply's communication interface provides functionality for the current to be read on demand by a device, using the

$$ser.write(”MEAS:CURR?\backslash r\backslash n”)$$

command.

The communication interface is used in order to both control and monitor the electrical load.

3.3.3 Temperature

A full complement of k-type thermo-couples were utilised. These were positioned in a range of positions, with particular care to cover the hottest areas of each experiment, using kapton tape when attached directly to a loaded conductor. The purpose of the thermo-couples was to provide a thermal safety cut-out switch, if temperatures around the experiment exceeded particular thresholds. The thermo-couples were interfaced with over USB via a PicoScope TC-08 data acquisition unit. The thermo-couple data was not utilised during the modelling process as surface contact measurements are not representative of what may be found in service in substations.

3.4 Modelling, Machine Learning and Software

Decisions on the modelling methods utilised during the course of the work were made based on the findings of the literature review (Chapter 2), on the form of the data being modelled and on the proposed future application context of the methods proposed.

Since thermal systems inherently take finite time to respond, and electrical power equipment is often large and has high thermal-mass, it can take an extended period of time for the influence of input variables to take their final effect on the output of the system. This characteristic of the physical systems determined that any modelling technique to be used should consider inputs at prior time-steps as well as the present input.

The proposed methods are defined for use in extended continuous monitoring applications, and the experiments have been designed to collect data in this manner. Continuous monitoring provides a significant benefit over a single sample, as previous values of the output variable may be used to form the model. While there may be merit in modelling precisely the impact of the input variables on the output, it is true that the extent the output can change by in the period of a time-step is finite, so must be a function of the output at the previous time step. Due to this intended application context, modelling methods that consider prior values of output as system inputs, namely auto-regressive methods, are considered.

The work of Bortoni et al [3] [4] [2] and Santos et al [52], which also considers thermal output extracted from thermal cameras as time-series, used auto-regression methods. For this reason multivariate-auto-regression (MAR), a variant of linear regression modelling, was selected as the initial modelling technique to be utilised in the work.

Elsewhere, success has been reported in recent literature utilising LSTM recurrent neural network technology in time-series modelling. This method was selected due to widespread usage, reported success and platform support in the early stages of the work. Temporal convolutional networks (TCNs) [63] have widely reported excellent multivariate time-series performance [33] [16] since the project commenced. These are not explored here but remain an area of exploration with regards to the modelling challenges of the project.

Note: a link to API developed during the course of this work is provided in Appendix A.

3.4.1 Linear Regression Modelling

A multivariate auto-regressive linear regression model equates the output of a system to some combination of coefficients and vectors of inputs. M is the model order (number of lagged values considered), N is the number of causal variables, with N+1 variables contributing to the output, due to the auto-regressive aspect. Equations 3.3 to 3.7 describe the model form in summed and matrix forms. y is the output variable, x is the input, θ is a vector of trainable parameters, and ϕ is a bias term.

$$y_t = \sum_{j=1}^{M}(y_{t-j}\theta_{0j}) + \sum_{i=1}^{N}\Big(\sum_{j=1}^{M}(x_{i(t-j)}\theta_{ij})\Big) + \phi \tag{3.3}$$

In matrix notation:

$$y_t = \phi + \boldsymbol{y}_{prev}\boldsymbol{\theta}_0 + \sum_{i=1}^{N}\boldsymbol{x}_i\boldsymbol{\theta}_i + \phi \tag{3.4}$$

Where

$$\boldsymbol{y}_{prev} = \Big[y_{t-1}, y_{t-2}, \ldots, y_{t-M}\Big], \tag{3.5}$$

$$\boldsymbol{x}_i = \left[x_{i(t-1)}, x_{i(t-2)}, \ldots, x_{i(t-M)} \right], \tag{3.6}$$

$$\boldsymbol{\theta}_i = \begin{bmatrix} \theta_{i1} \\ \theta_{i2} \\ \vdots \\ \theta_{iM} \end{bmatrix} \tag{3.7}$$

The order of the model is determined experimentally in the course of this work, by sweeping through candidate values and evaluating the effect on the output.

The model is optimised using gradient descent with a mean-squared error cost function. The learning rate and epochs are selected experimentally. The multivariate lagged values forming the inputs allow the regression to model more complex time-dependant systems than simple linear regression.

3.4.2 LSTM Techniques

Long short-term memory (LSTM) neural networks are a form of recurrent neural network (RNN) proposed in 1997 by Hochreiter [17]. LSTM networks utilise a data pipeline in their recurrent layer, which allows data to be maintained for long time periods in the dataset. This enhances their ability to 'learn' long term causal relationships between data points, without falling victim to exploding or vanishing gradients in the training process, a problem commonly faced by simpler recurrent neural networks, especially when dealing with longer series of data (more lags). A result of these characteristics is that LSTM networks are suitable for the problem of time-series prediction. Furthermore, being artificial neural networks, they are suitable for scaling up to large numbers of features.

$$f_t = \sigma(W_f x_t + U_f h_{t-1} + b_f) \tag{3.8}$$
$$i_t = \sigma(W_i x_t + U_i h_{t-1} + b_i) \tag{3.9}$$
$$\tilde{c}_t = tanh(W_c x_t + U_c h_{t-1} + b_c) \tag{3.10}$$
$$c_t = f_t \cdot c_{t-1} + i_t \cdot \tilde{c}_t \tag{3.11}$$
$$o_t = \sigma(W_o x_t + U_o h_{t-1} + b_o) \tag{3.12}$$
$$h_t = o_t \cdot tanh(c_t) \tag{3.13}$$

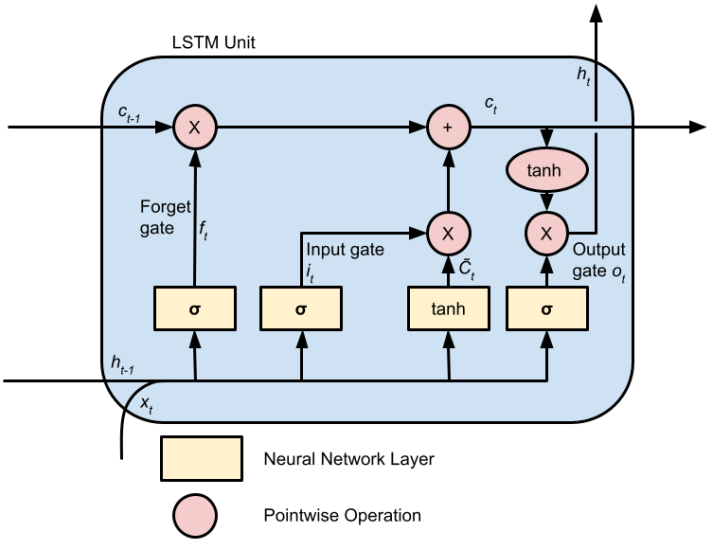

Figure 3.4: A graphical representation of an LSTM unit

The equations dictating the weights and their corresponding activation functions for an LSTM unit are presented in Equations 3.8 to 3.13, where x is input, h is hypothesis output, c_t is the cell state, f is the forget gate value, i is the input gate value and $\tilde{c}_t$ is the candidate value for the next cell state. There are two weight matrices for an LSTM cell: W and U. W contains the weights applied to input values, while U contains the weights applied to the previous prediction value from the previous cell (h_{t-1}). b is a vector of bias values. These are included in the LSTM implementation utilised for this work but are not a requirement.

The units these weights correspond to are shown in Figure 3.4. Referring to the figure, the capability of this type of recurrent neural network to retain information over many timesteps, is conferred by the horizontal data pipeline for c_t and the corresponding 'forget gate'.

In the case presented here, a single continuous output value is required for a sequence of inputs. The output of the LSTM layer of the network is fed to a single neuron, through a dropout layer, in order to minimise the risk of over-fitting. The dropout layer sets a proportion of weights to zero. This is particularly useful in ensuring the network does not train to a persistence model, as is a common issue in auto-regressive processes.

Due to the nature of the application of the generated models, L1 and L2

bias regularisation values are configured into the LSTM layer. The regularisation should assist in ensuring there is minimal bias in the model output, increasing the chances of a model residual with a mean of zero.

The model is trained on mean absolute error, with the 'Adam' [29] optimiser. Data is shuffled in training. Further configuration of the training process is detailed in Chapter 4, Section 4.5.2.1.

3.4.2.1 Input Shaping

This section describes the manner in which the input variables are manipulated specifically for input into the Keras LSTM layer. This input shaping dictates the patterns that may be learned by the recurrent neural network.

The dataset is first framed as as an input/output pairing. The output is y(t). The inputs, as in the linear regression model case, are:

$$\boldsymbol{y}_{prev} = \left[y_{t-1}, y_{t-2}, \ldots, y_{t-M} \right], \tag{3.14}$$

$$\boldsymbol{x}_i = \left[x_{i(t-1)}, x_{i(t-2)}, \ldots, x_{i(t-M)} \right], \tag{3.15}$$

where M is the number of lags. The dataset is initially shaped as [rows, columns], designated as [samples, variables]. A shuffling process is used such that each lagged time variable is a column of data. This is depicted with dummy data in Figure 3.5. In the provided example, there are three causal variables: x1, x2 and x3 and a single output variable: y. New columns are created with lagged values of the causal and output variables. In this format, each row of the dataset table forms an input/output pairing. After the shifting has been completed, the dataset is split into train and test sets, as denoted by the labels and the shading. Finally, the data is reshaped into a 3-dimensional data structure of shape: [data length, number of lags, number of variables], effectively grouping the variable types together. The Keras LSTM layer then receives sequences of length equal to the data set length.

3.4.3 Data Pre-processing

Environmental, electrical and camera casing temperature sensor data are output from the experiment in .csv files.

Index	x1	x2	x3	y(t)
1	200	10	50	1
2	196	12	45	2
3	192	14	55	3
4	188	16	51	4
5	184	18	46	5
6	180	20	56	6
7	176	22	52	7
8	172	24	47	8
9	168	26	57	9
10	164	28	53	10
11	160	30	48	11
12	156	32	58	12
13	152	34	54	13
14	148	36	49	14
15	144	38	59	15
16	140	40	55	16
17	136	42	60	17
18	132	44	50	18
19	128	46	56	19

	Index	x1(t-2)	x2(t-2)	x3(t-2)	y(t-2)	x1(t-1)	x2(t-1)	x3(t-1)	y(t-1)	y(t)
Training-set	1	-	-	-	-	-	-	-	-	1
	2	-	-	-	-	200	10	50	1	2
	3	200	10	50	1	196	12	45	2	3
	4	196	12	45	2	192	14	55	3	4
	5	192	14	55	3	188	16	51	4	5
	6	188	16	51	4	184	18	46	5	6
	7	184	18	46	5	180	20	56	6	7
	8	180	20	56	6	176	22	52	7	8
	9	176	22	52	7	172	24	47	8	9
Test-set	10	172	24	47	8	168	26	57	9	10
	11	168	26	57	9	164	28	53	10	11
	12	164	28	53	10	160	30	48	11	12
	13	160	30	48	11	156	32	58	12	13
	14	156	32	58	12	152	34	54	13	14
	15	152	34	54	13	148	36	49	14	15
	16	148	36	49	14	144	38	59	15	16
	17	144	38	59	15	140	40	55	16	17
	18	140	40	55	16	136	42	60	17	18
	19	136	42	60	17	132	44	50	18	19

Figure 3.5: Representation of the organisation of data into inputs, output, test-set and training-set. On the left, data pre-shifting. On the right, data post-shifting.

In order that all data was evenly sampled, and at the same time-base, a resampling procedure was required. First, all text was stripped from data fields. Data was up-sampled to a 1 Hz frequency using linear interpolation, then down-sampled to a sample period of 60 seconds. After all files were re-sampled, the data can be consolidated into a single .csv file. Image data was similarly stored on the host experiment controller Raspberry Pi.

Typically an image was selected for normalisation and manually inspected. This normalisation process was critical to be capable of correctly selecting ROIs/POIs in the images, as it enabled the identification of physical features of the monitored equipment in the scene as in Figure 3.6.

To ready the images for input into modelling processes, they must be re-sampled to an exact frequency, as the time-series methods used assume this to be true. Images are re-sampled using the same method as the sensor data .csv files, treating pixels as time series. For the purposes of the CSE experiments, software reads each image in the sequence in turn, extracting the mean value of a 3x3 kernel around the POI pixel, appending it to a data structure and storing in a file, ready for modelling. For the purposes of the overhead line experiments, an ROI mask is generated, providing a region of interest encapsulating the energised overhead line conductor, by using a combined threshold and connected-components operation.

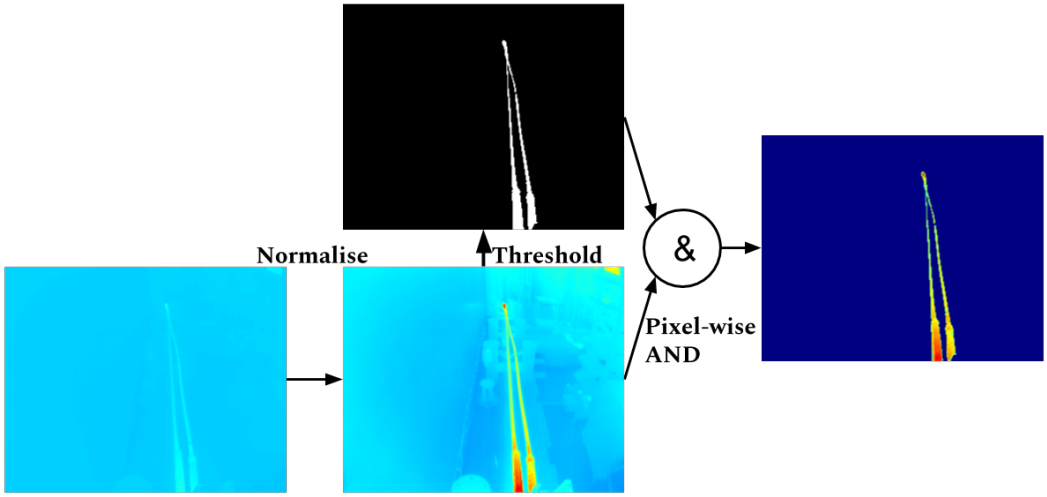

Figure 3.6: The thermal image pre-preocessing workflow for the overhead line scenario.

The sequence of images is then masked using a pixel-wise binary AND operation This masked image is unrolled into a 1D array, zeroes are discarded and the mean value of the array is taken. This average pixel value is taken as the thermal output of the system. This image pre-processing pipeline is depicted in Figure 3.6.

All data are normalised by scaling between 0 and 1, per variable, before input into modelling processes. This scaling ensures that all variables cover an equal numerical range, so there is no bias for or against any variable when assigning coefficients or weights to the variables during the process of fitting the regression or training the recurrent neural network. Particularly in the case of the recurrent neural network, this important for the training process. The non-linear nature of the activation functions used in the neural network mean that variables of different orders of magnitude can lead to large or small weights, in turn leading to increased or decreased sensitivity to certain inputs.

3.4.4 Recursive Predictions

Referring to Section 3.4.2.1, it is clear that care must be taken when inspecting the output of models generated using data of this form. While a plot of the test-set portion of a dataset, showing observed values and model output values, may appear to be a long-term time-series prediction, it is rather a series of single-step predictions, plotted in sequence.

Table 3.3: Software libraries and APIs used in the work

Library	Description
CUDA	Low level (close to the metal) framework for enabling matrix operations on NVidia GPUs. Provides potential significant speed-up
TensorFlow	Machine and deep learning library used for building, training and testing machine learning models
Keras	Interface with deep learning libraries such as TensorFlow
SciPy/SkLearn	Python Machine learning and data analysis package
Pandas	Python data analysis and manipulation package
Numpy	Python mathematical library
OpenCV	Open-source image processing library

For visualisation purposes, recursive predictions (or forecasts) are used at various points in the thesis. Figure 3.5 shows a sequence of dummy data. When computing recursive predictions, the model output is computed for time t, for example. This generates model output $\hat{y}_t$. In order to make the recursive prediction for time $t+1$, $\hat{y}_t$ is shuffled back into the row of inputs that are fed into the model, in the column corresponding to y_{t-1}. From this new row of input data, the model output can be generated for $\hat{y}_{t+1}$.

3.4.5 Implementation

This section will provide a brief overview of the software technologies used in implementing the methods described in this chapter.

3.4.5.1 Python

Python is a widely used, open-source interpreted programming language. There are numerous scientific and numerical libraries available for use and software can be written for use in console sessions (in a similar manner to MatLab), in scripts, or in full-size programs. This flexibility and utility make it ideal for data analysis in a research setting. The Python version used over the course of this project is 3.7.

Table 3.3 lists the libraries frequently used during the course of this work.

3.5 Experiment Design and Data Collection

3.5.1 Context

Based on the findings in the literature review (Chapter 2), and the resulting potential contribution, a set of requirements were compiled for an experiment that would provide the means to acquire the data on which models would be created (see Section 3.4). The requirements are as follows. The experiment must:

1. Generate time-series data
2. Comprise of genuine power-engineering equipment
3. Utilise electrical loading representative of in-service equipment
4. Comprise some form of measurable environmental variation
5. Record a sequence of thermal images
6. Record environmental data
7. Be autonomous and safe, for long-term usage

The justification for this experiment is that it provides real laboratory data from sensors akin to those used in condition monitoring practice, with which to investigate the utility of modelling the thermal output of equipment. This will serve as a simplified test-case in order to determine whether the concept of using the selected cameras and modelling methods is valid and useful. Using simple aluminium conductors as the test-object maintains a simple relationship between the input and output of the system, while still being genuine power equipment that could be found in an in-service substation environment.

3.5.2 Experiment Design

In order to collect data to facilitate the creation of models as described in Chapter 2 the overhead line experiment was devised. A current loop comprising of two Rubus type ([5]) all-aluminium-alloy (AAC) overhead line conductors terminated by custom crimps, a Magna Power high-current power supply [23] and a conducting block of aluminium was created. A FLIR Boson OEM thermal camera core (see Section 3.2.2) was selected and acquired. Software was adapted to

extract individual files from the camera, at the appropriate point in the image pipeline. A Vaisala WXT520 (Section 3.3) weather station was selected, acquired and software was written to read appropriate data from it's serial interface. A large industrial fan was sourced within the university. A high-current relay-based control box was developed to allow the fan to be autonomously controlled. A weather-proof, power-over-ethernet (POE), visible light security camera was integrated into the system, to provide potential utility for future experiments, but was ultimately unused. Safety provision was provided by setting appropriate over-current and over-voltage fault limits on the power supply, and with an over-temperature fault condition using thermo-couples, as described in Section 3.3.3. Environmental and electrical data was stored in human-readable .csv files, with a matching sequence of 16-bit resolution '.png' image files from the thermal camera.

Experiments were controlled and data was logged by a Raspberry Pi SBC. A Python program scheduled events, as defined by input files with load and wind patterns, which prompted changes in current, wind-fan state (on or off) and scheduled data capture events, during which it triggered thermal image capture, environmental data capture and recorded current and voltage values as reported by the high-current power supply. The FLIR Boson camera was positioned in order that the majority of the length of the conductor loop was within its field of view. A suitable position for this was on top of the power supply unit. The wind-fan was positioned such that it blew air at a portion of the conductors which were within the field of view of the camera, in order that its effects could be characterised.

The resistance of the conductor loop was found experimentally to be $r \approx 0.01\ \Omega$. A current range of $0\ A \leq I \leq 800\ A$ was selected to provide heating of $P \leq 6400\ W$. Further testing showed this current range to provide sufficient heating for the experiments while remaining under the safe operation temperature threshold of 90 °C, as recommended by the laboratory operators.

Based on these requirements a pair of experiments were carried out. These are presented in further detail in Chapter 3.

To perform an input and output validation check on the experiment a step-test was performed. The load and wind profiles for this test can be seen in Chapter 3, Figure 3.9. The thermal time constant was found to be $\tau \approx 20$ min during this test.

To approximate real-world conditions, National Grid national loading data

Figure 3.7: The overhead line experimental setup, including power supply, conductors, fan and weather station.

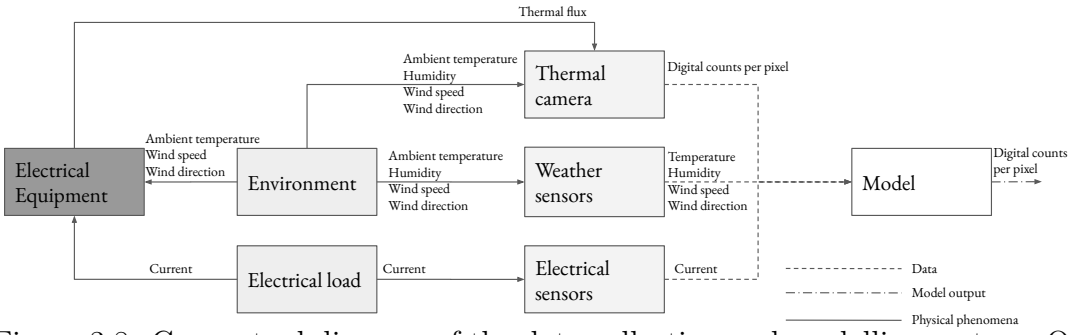

Figure 3.8: Conceptual diagram of the data collection and modelling system. On the left is the physical experiment, which is perceived by sensors in the centre, which contribute data to the model, which produces digital count pixel value outputs.

was normalised to within safe experimental levels and used to define the load pattern for the 'realistic' test. The load pattern time period was determined to be approximately four days, due to laboratory constraints. The test was run from 2018-07-19 12:56 to 2018-07-23 13:56, encompassing four varied high-load periods, corresponding to day-time in the national load figures, and four low-load periods, corresponding to night-time. The load pattern peaks were aligned with the times at which the real world data would peak, attempting to capture as much realism as possible, in the case of environment cooling in the laboratory overnight. The industrial fan was enabled in a binary pattern of a differing frequency to the loading data. Models were created based upon these datasets in order to predict the average pixel value of the region of interest based on the load and environmental conditions. The setup can be seen in Figure 3.7.

A diagram depicting the conceptual layout of the system is shown in Figure 3.8. This figure describes the relationship between the physical system, the sensors that observe the system and the model that attempts to encapsulate the effects of the system on pixel output of a given pixel of a thermal camera. As discussed in Section 3.4, the model takes data from prior time-steps as input. In the condition monitoring context, the model output of a pixel intensity (proportional to thermal output) serves as a baseline for how a piece of equipment appears during normal operation. Given a sufficiently general model, when the observed pixel output diverges from the model output, it can be taken as a signal that the thermal output is not normal and therefore warrants investigation as a potential fault-precursor.

Table 3.4: Overhead line system output and inputs. * indicates variables unused in the step-test experiment.

System output	System inputs
Digital counts	Wind speed*
	Current
	Air temperature
	Humidity

3.6 Datasets

This section presents the data collected over the course of the experiment. The raw, unprocessed results of the two tests can be seen in Figure 3.9. The data have been scaled for readability. The data variables are presented in Table 3.4.

3.7 Results and Analysis

3.7.1 System Step-response Experiment

This section presents the results from the step-test characterisation experiment. The expectation is that the simple input patterns will allow the system and it's behaviour to be understood, while also providing a test-bed for creating a multivariate model based on this form of data.

3.7.1.1 Preliminary Analysis

A number of primary findings were made from initial analysis of the output data set.

Firstly, it was found that the system output corresponding to the cable conductors temperature exhibits a first-order step response, as would be expected for a thermal system. This finding is significant to the work as it demonstrates that fundamentally the camera can provide a valid indication of the heating and cooling in a given scene over time.

The second observation made is that the wind measurement is noisy compared to the other variables in the system. The fan runs at constant speed when enabled. Two potential sources of the noisy signal have been identified. Firstly, the weather station sensor is near to a wall, potentially causing turbulence in the air-flow as it changes direction on impact with the wall. Secondly, it may be a function of the internal configuration of the weather station. This measurement noise flags

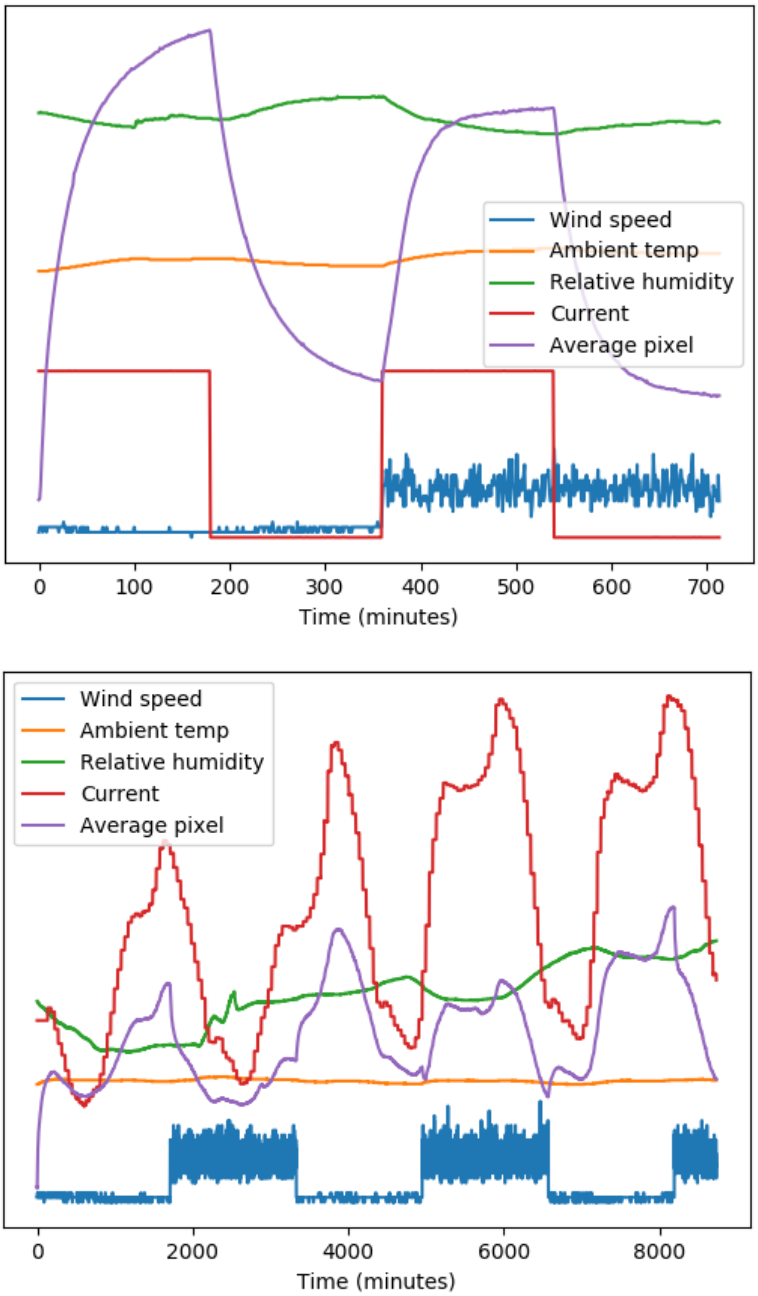

Figure 3.9: Data from; top: overhead lines ystem step-response test experiment; bottom: overhead line realistic experiment. All series are scaled in order to fit visibly on the same axis.

the variable for consideration for smoothing post-processing before input into learning algorithms in order to minimise the potential impact of the sensor noise on the model output, however this is dependent on the performance of models using the raw wind data. It is expected that a coarse indication that wind is present would be sufficient.

Significantly, the impact of the wind on the system output can be observed. When the fan is enabled in the latter half of the experiment run, the average output digital count value peak is 84% than that in the first half of the experiment. Furthermore the average output digital count value drops to a lower value in the trough. This impact is as expected for the forced convection cooling effect [50] on the OHLs and evidence of it in the output data validates the approach.

It can be observed that the ambient air temperature is relatively steady throughout the course of the test, from 25.6 to 27.8 °C, a range of only 2.2 °C From Figure 3.9, it is clear the ambient temperature changes correlate with the hot and cool periods of the output, suggesting the electrically loaded OHLs are contributing to the temperature of the laboratory. Referring to Figure 3.10, the effects of the minor fluctuations in ambient temperature can be seen in the plot for the background pixel. The fluctuation of 250 digital counts over the period of operation is considered small enough to be negligible for this test. This plot also highlights the self-heating effect [28] present in the FLIR Boson thermal camera: as the camera heats up in the initial period of operation the output for a given background pixel is impacted significantly. However referring to the ROI pixels in the remaining two plots it is clear that the self-heating effect is minimal compared to the observed changes in thermal output in the scene. The effects of heating and cooling the cameras are further explored in Chapter 4. For the purposed of the Overhead Line experiment the effects are treated as negligible due to the relatively consistent ambient temperature. Figure 3.10 shows that the influence of the wind can be clearly detected by the thermal camera. The wind influenced pixel exhibits a clear difference in the maximum number of output digital counts, along with a decreased rate of increase. With the fan enabled, the minimum value at the end of the test is lower than the midpoint, without the fan enabled. The non-wind-influenced pixel displays a closer similarity between the first and second halves of the experiment. The initial rate of increase is similar between the two halves. There is some difference in the manner in which the digital counts settle. This is likely attributed to the change in temperature in different parts of

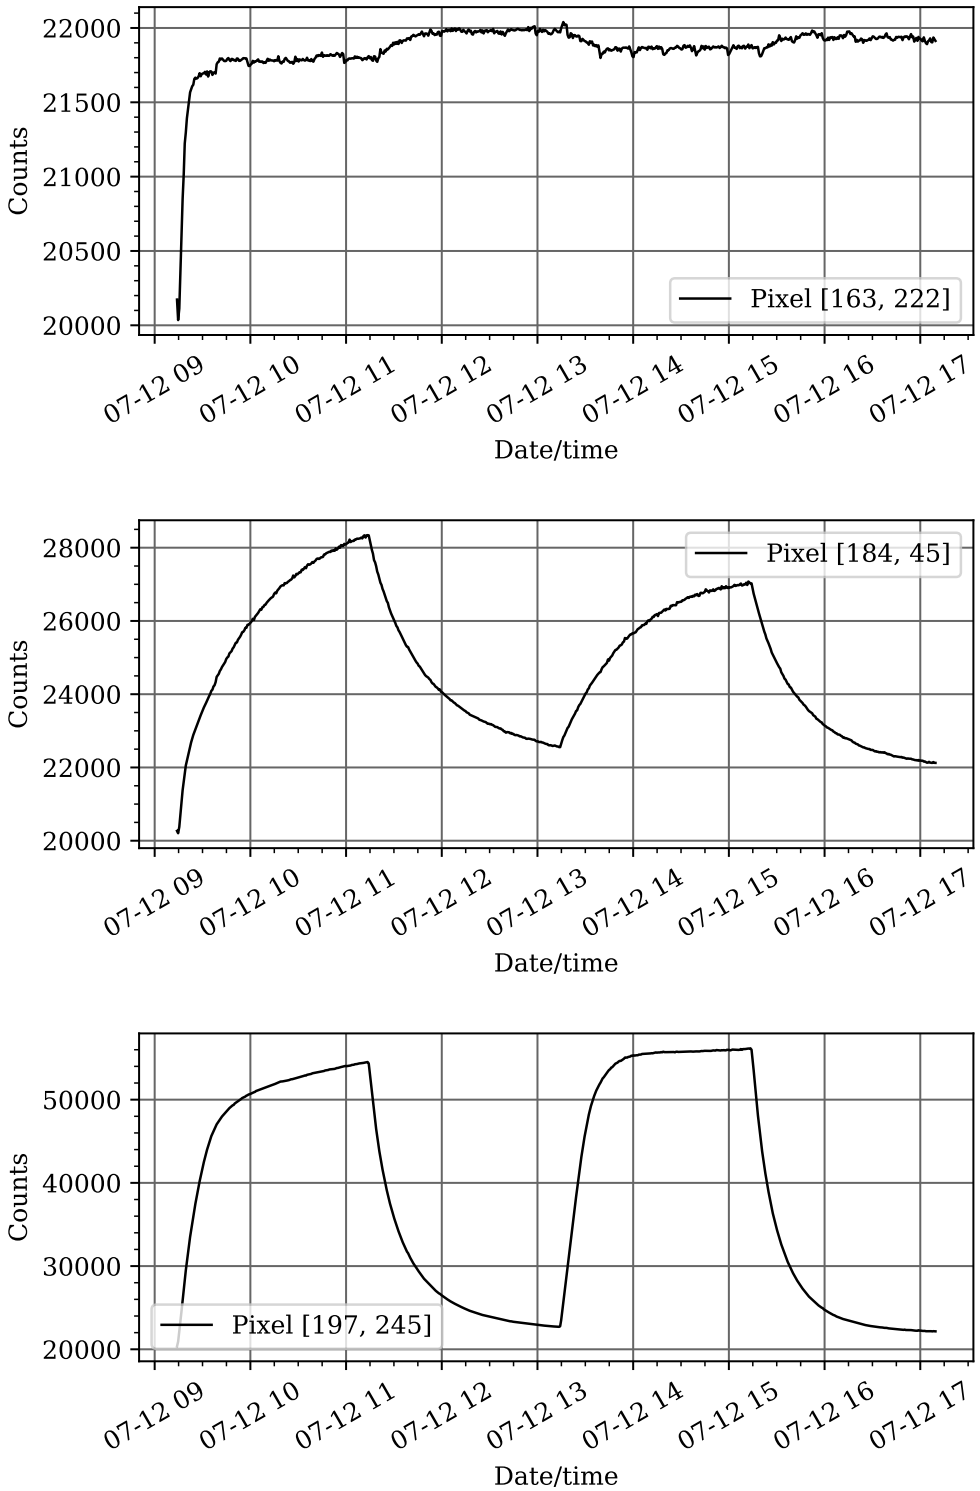

Figure 3.10: System step response of three distinct pixels; top: background pixel (outside of region of interest); middle: wind-influenced pixel, at opposite end to power supply; bottom: non-wind-influenced pixel, close to power supply. Note x axis is datetime, format 'mm-dd hh'.

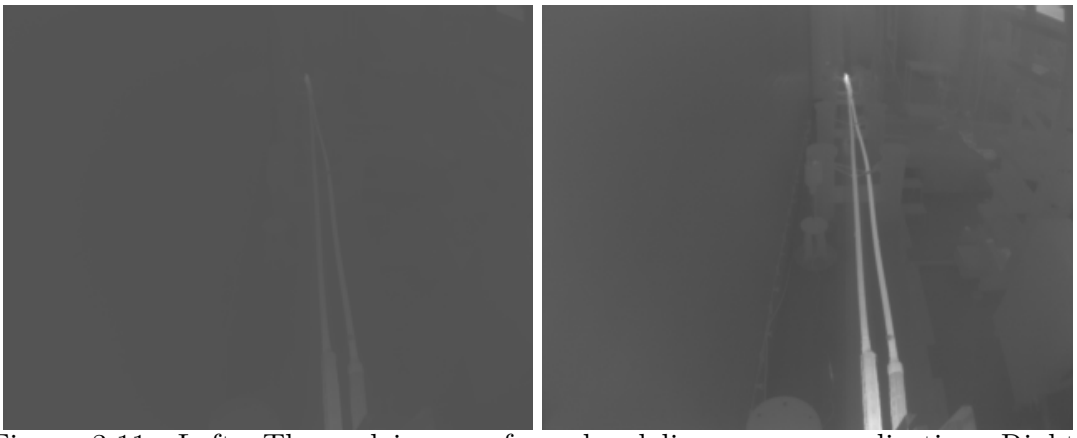

Figure 3.11: Left: Thermal image of overhead lines pre-normalisation; Right: Thermal image of overhead lines post-normalisation

the cable causing differences in effective resistance across the length of the cable.

Finally, it can be determined that the thermal time constant of the system output is approximately 20 minutes. Determining this value allows the load patterns for other experiment runs to be set at appropriate levels for detectable output changes. Inspection of the thermal step-response of the system provides some insight into how it may be modelled in terms of an electrical equivalent circuit. It is clear the overhead lines have significant thermal capacity given the time constant. Inspection of the normalised thermal images of the overhead lines in Figures 3.11 and 3.12 demonstrates the low thermal resistance of the system. This is confirmed in comparing the wind-influenced pixel and non-wind-influenced pixel in Figure 3.10: the heating effects are immediately apparent in both plots. While the thermal resistance is low, it's effect can be witnessed in that the wind-influenced pixel (nearest the power-supply) gets closer to steady-state than the non-wind-influenced pixel in the first cycle of the test, when thee fan is disabled. Though the ambient temperature remains stable during the course of the experiments, it requires consideration as an environmental factor. In terms of the equivalent electrical circuit, as the ambient temperature increases, it decreases the thermal difference between the heated conductor and the environment, therefore decreasing the thermal gradient.

As generating models is not the primary focus of the step-test experiment, the analysis performed is brief, and is included in the Section 3.7.1.2.

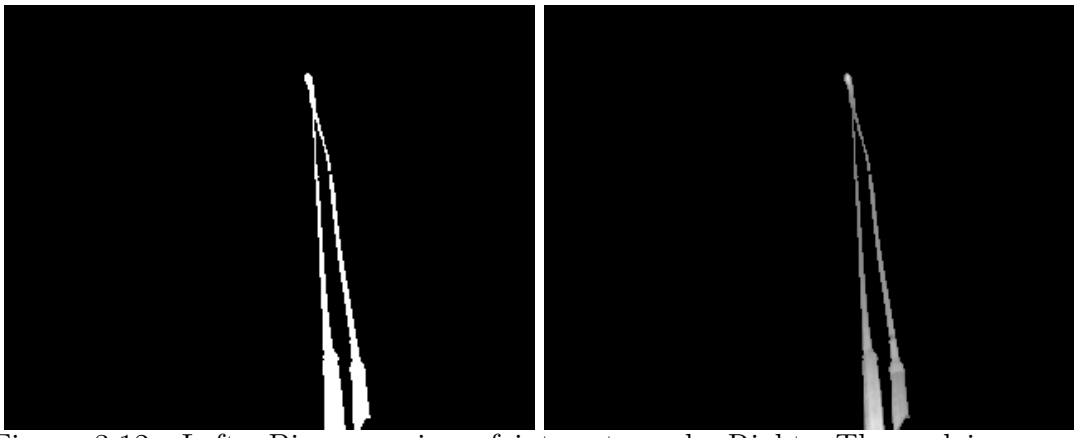

Figure 3.12: Left: Binary region of interest mask; Right: Thermal image of overhead lines with binary mask applied

3.7.1.2 Analysis

The step-test experiment, as described, was intended to provide familiarity with the experimental set-up, and to aid understanding of the system inputs and outputs. It also provided an opportunity to explore the modelling techniques on the sensor data acquired from this type of experiment. The expectation was, given that the dataset is so limited, the model performance would be poor, especially if a model was created on a training/test split.

The test was a success. Insights were gained into the output of the non-calibrated FLIR Boson thermal camera. Inspection of the camera output, both in terms of a region-of-interest and in terms of individual pixel values over time, revealed that the camera can be used to understand how the thermal output of objects in a scene changes over time. It was revealed that there is a warm-up period when the camera is initially powered-on, matching the findings in UAV surveying literature [28], informing strategy for future testing. The effect of the ambient temperature on the camera output, through changes in the cameras own temperature, was revealed by inspecting pixels corresponding to areas of the scene that are not actively heated. It was confirmed that the industrial fan effectively cools the electrically loaded overhead-lines, and that this change in thermal output is visible in the thermal camera output.

3.7.2 Realistic Load Experiment

This section presents the results of the experiment and load pattern detailed in Section 3.5. Some initial manual data analysis is performed, providing insight into

the structure of the data. Modelling results as obtained using the linear regression method (detailed in Section3.4.1) are presented for a selected set of possible hyper-parameters. Modelling results as obtained using the LSTM method are then presented with discussion of the effects of various hyper-parameters.

3.7.2.1 Preliminary Analysis

On manual inspection the realistic load pattern dataset reveals features that comprise a logical progression from the findings of the step test load pattern.

Firstly, it is confirmed that the primary driver of the output is the current input. A visual correlation between the two variables is clear.

As in the step test load pattern experiment, the ambient air temperature stays consistent for the period of the test, with a range of only 1.6 °C. This provides assurance that there is only minimal environmental influence on the thermal camera, and therefore that the output can be treated as approximately stable throughout the test period. The minimal ambient temperature variation is also a limitation of the dataset: the overhead line system does not experience much environmental thermal fluctuation, which could feasibly impact the thermal output of the loaded conductor. The effects of ambient temperature fluctuations are further explored in Chapters 4 and 5.

Manual inspection confirms that a suitable range of operating conditions are experienced over the course of the experiment. There are clear periods of high electrical load while the fan is turned both on and off, and there are periods of low electrical load while the fan is turned both on and off. As described in Section 3.5, the experiment inputs were designed for this to be the case.

It is clear that there is a slight lag between the changes in the electrical current through the system and the changes in the system output. This matches the expectation for thermal systems.

It is confirmed that the system output matches intuition for this type of thermal system. In particular, that the system output is lowest when the electrical load is minimal and the wind fan is enabled: effectively when the heating effect is minimal and the cooling effect is maximal.

Finally, it is noted that there is a visible hot-spot in the output images of the system. This provides a potential opportunity for manipulation of the data to determine whether it is possible to detect such hot spots using this data and modelling method.

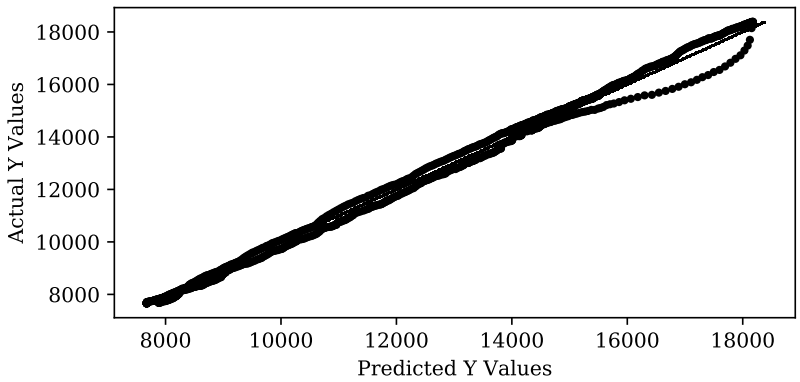

Figure 3.13: Realistic Test, MAR Linear Regression Model. Data-points of the test-set, prediction vs actual value scatter graph. The line indicating a prediction equal to the actual value is shown on the diagonal.

3.7.2.2 Linear Regression Model

This section details the multivariate auto-regression linear regression models of the realistic load pattern test of the overhead line experiment. The selected configuration for the modelling is provided in Table 3.5. This configuration was selected based on the approximate thermal time constant for the system and on an investigative parameter sweep.

Table 3.5: Selected MAR linear regression configuration for overhead line, realistic test

Parameter	Configuration
Epochs	7000
Learning rate	0.002
Lags	20

Figure 3.17 shows the model output as formed by the linear regression model. This result was obtained by including 20 lagged input values in the model. An r^2 score of 0.995 was obtained. While the score appears to suggest an accurate model with low error, it is apparent from inspection of both the model output and the residuals that the model fails to capture the influence of all inputs correctly. From Figure 3.17 in conjunction with Figure 3.13 and Figure 3.14 it is clear that the predicted output is consistently higher than the actual output until it drops significantly (at approximately 1200 minutes), from when the predicted output is lower than the actual output. Inspection of the residuals (Figures 3.13 and 3.15) of the test dataset provides some insight into the deficiencies of the

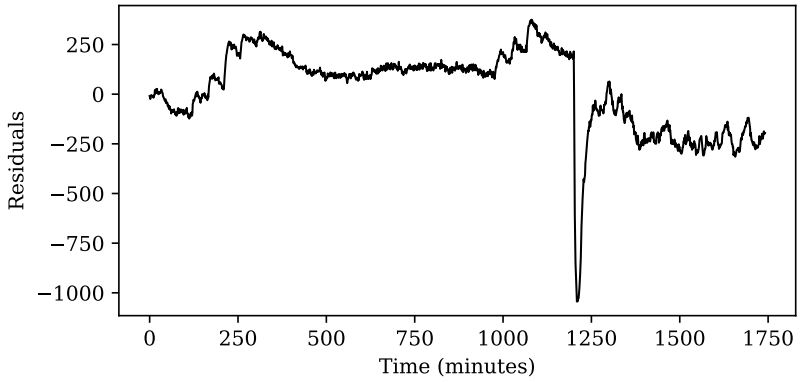

Figure 3.14: Realistic Test, MAR Linear Regression Model. Residuals: real value - predicted value for the time period forming the test-set portion of the dataset.

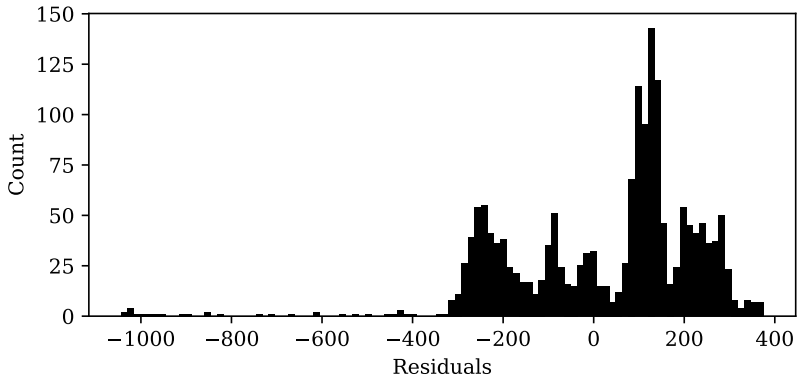

Figure 3.15: Realistic Test, MAR Linear Regression Model. Histogram plot of the residual values.

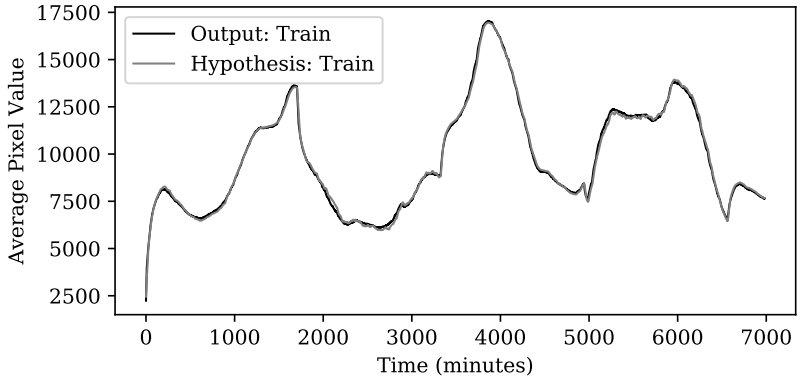

Figure 3.16: Realistic Test, MAR Linear Regression Model. Actual output and predicted (hypothesis) output for train-set portion of the dataset.

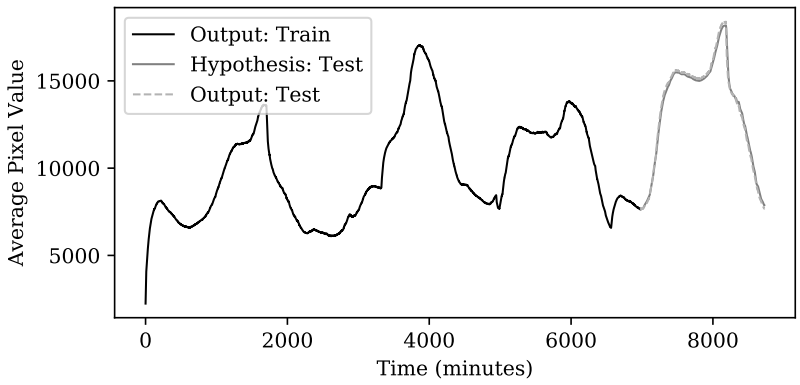

Figure 3.17: Realistic Test, MAR Linear Regression Model. Actual output for train-set, and actual and predicted (hypothesis) output for test-set portion of the dataset.

model. The residuals are widely spread and do not follow a normal distribution, exhibiting a clear positive skew, suggesting the model may be unsuitable for use in a fault-detection context. The standard deviation of the residuals, as detailed in Table 3.6 provide a further indicator of this. While there is apparent positive bias of the model during the test portion of the dataset, when the model output is plotted against the real output for the train portion of the data set (Figure 3.16) it is clear there is no consistent bias in the model output.

Table 3.6: Standard Deviation of residuals, for MAR linear regression model of the overhead line experiment, realistic test

Data	Standard Deviation of residuals
Train-set	134
Test-set	207
Combined	152

3.7.2.3 LSTM Model

The hyper-parameter sweep process used in order to determine a suitable configuration of the LSTM recurrent neural network is detailed in Section 3.4.2. The configuration selected following the hyper-parameter sweep is shown in Table 3.7.

This configuration results in loss scores that descend gradually, converging on a common point. This is shown in Figure 3.18. The convergence to a common value suggests that the model has not over-fit to the training dataset and is generalised across both portions of the dataset. Bias regularisation and dropout are used to help ensure the model does not over-fit.

Table 3.7: Selected LSTM configuration for overhead line, realistic test

Parameter	Configuration
LSTM Units	300
Lags	20
Dropout	50%
L1 Regularisation	0.01
L2 Regularisation	0.01

Inspecting the residuals of the model, as presented in Figures 3.18, 3.19 and 3.20 it is clear the model has captured the influence of the inputs on the system. The residuals are centred around 0, and have a relatively narrow distribution. The standard deviation of the distribution of the residuals for each portion of the dataset is provided in Table 3.8. The low standard deviation here means that predictions can be made using the generated model and that the predictions will have relatively narrow 95% prediction intervals. The zero-mean residuals with a relatively small spread suggest the model may be suitable for use in model-based fault detection methods.

Performing visual inspection of the model output in Figure 3.21 confirms the accuracy of the model. The plot on the left shows the model output and real output values for the test portion of the dataset. The plot portrays the output of the model as it is configured to generate it: using Very small deviations from the real output can be seen but generally the model output tracks the real output closely.

Due to the closeness of the model to the real output, and the narrow distribution of the residuals, recursive one-step predictions can be made, as shown in Figure 3.22. The process for making these predictions is detailed in Section 3.4.4. The predictions made track the true system output closely, while being based on previous predictions. This demonstrates that the model has applied appropriate weight to the influence of especially the current, this being the main driver of the output for this portion of the test. 95% prediction intervals are presented in this plot, but but these were calculated as if the forecast was a naive forecast, and as such are only useful as an indication of the spread of the residuals.

3.7.2.4 Analysis

The realistic test experiment was intended to create a long-term dataset, with an electrical loading pattern that is representative of in-service equipment and

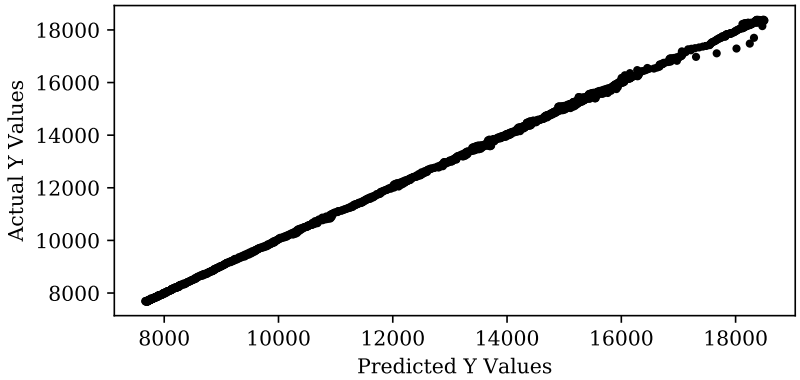

Figure 3.18: Realistic Test, LSTM Model. Data-points of the test-set, prediction vs actual value scatter graph. The line indicating a prediction equal to the actual value is shown on the diagonal.

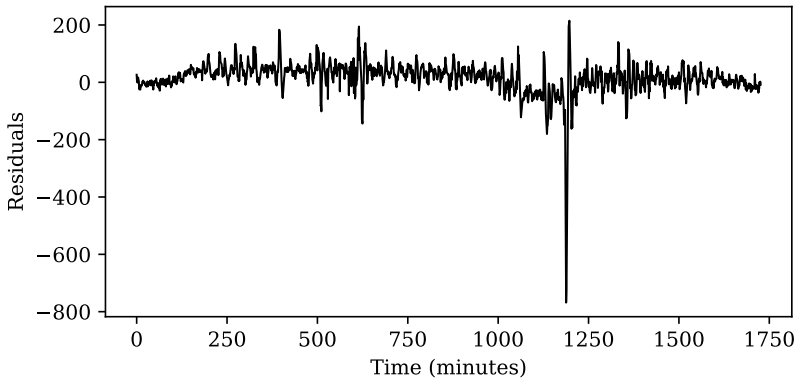

Figure 3.19: Realistic Test, LSTM Model. Real value - predicted value for the time period forming the test-set portion of the dataset.

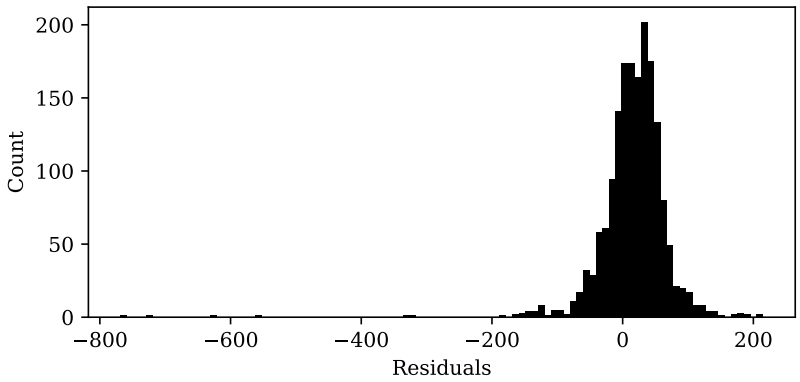

Figure 3.20: Realistic Test, LSTM Model. Histogram plot of the residual values.

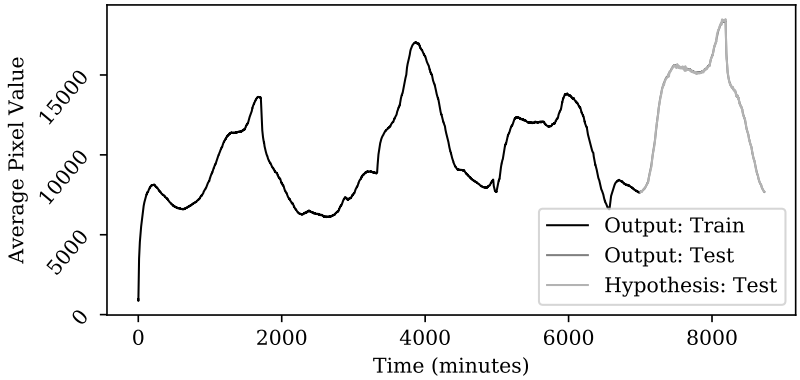

Figure 3.21: Realistic Test, LSTM Model. Actual output for train-set, and actual and predicted (hypothesis) output for test-set portion of the dataset.

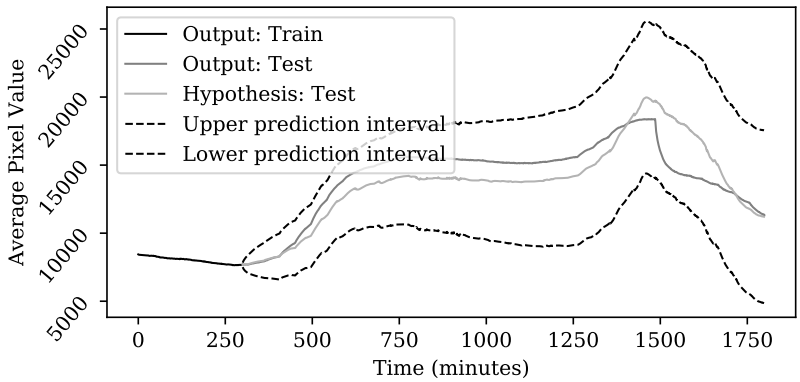

Figure 3.22: Realistic Test, LSTM Model. Rolling predictions with 95% prediction intervals.

Table 3.8: Standard Deviation of residuals, for LSTM model of the overhead line experiment, realistic test

Data	Standard Deviation of residuals
Train	28.851
Test rate	43.352
Combined	34.405

more varied environmental conditions, and to apply the methods used in the step-test experiment to model the system output. A dataset consisting of 8732 samples spanning 97 hours was created, with a load pattern based on national electrical power usage, scaled for the cable's operational current rating. A periodic wind profile was applied, at a non-matching frequency to the load pattern. Only limited variation of ambient temperature was possible in the laboratory in which the experiment was housed.

Models were created based on the data using both multivariate auto-regression and LSTM recurrent neural networks. The multivariate auto-regression model appeared to have partially captured the relationship between the system inputs and the output. Despite considerable configuration changes and adjustments to the number of lagged input values considered, the best performing model achieved is presented in Section 3.7.2.2. The training score suggests that the model has converged well, with low (normalised) error, but significant issues exist with regards to the distribution of the residuals. While the model output does trace the observed output well, the bias and wide spread of the residuals render the model unusable in the context of model-based fault detection.

The LSTM recurrent neural network results were very encouraging. As described in Section 3.7.2.3, visually the model output tracks the observed output very well, and inspection of the mean and distribution of the residuals reveal a well-performing model. Furthermore, the training score for both the training-set and the validation (test) set converge at similar rates and to similar values, suggesting a model that is not under or over-fit. Other than a single anomalous residual value, the spread of the residuals is narrow and centred close to zero. While there are limitations to the study, as detailed in Section 3.8, the result is determined to perform well enough to warrant further investigation into it's usage as a modelling tool in the context of time-series data of electrically loaded substation equipment.

3.8 Discussion

While the application of modelling techniques to data generated in the overhead line experiment has been successful, there are limitations present. Firstly, the output of the system, in terms of modelling, is defined as the average pixel value of the region of interest. This has a dampening effect on the data: the average pixel value does not rise and fall as much as, for example, the pixels with the most thermal output. It also has limited functionality as a fault-detection output, as the averaging process may mean that a potential hot spot only results in a small increase in the average pixel value. A recommendation for further work, based on this limitation, is that specific points of interest are used as outputs of the model, providing a more granular impression of the system behaviour.

The modelling process was repeated multiple times in order to find an LSTM configuration resulting in an accurate model. Accurate models were created during this process, but the specific differences of the parameters are not recorded, nor is average performance over multiple runs of the same configuration. This limitation is accepted as the work is presented as a proof of concept for more complex modelling context, however a thorough investigation would confirm the performance of the utilised method.

The experiment is inherently two-dimensional, while environmental effects are not. Due to the small size of the overhead line conductors, this does not present a significant problem with this experiment, but if moving to larger equipment, a single observation perspective may limit the ability to garner a general understanding of the impact of current loading and environmental effects. This is addressed with additional sensors configured in a novel manner in Chapter 4.

The experiment only provides limited environmental variation. While it is useful to prove that the effect of the wind can be modelled, there are a multitude of environmental factors which influence the output of thermal cameras and this experiment does not explore these. It is recommended that future experiments include more types of environmental variation, though it is accepted that this can be difficult experimentally and furthermore presents potential problems when using non-calibrated thermal cameras.

While the experiment is a proof of concept, there is value in performing additional analysis on the dataset. Due to the fact that there is an area of high impedance at the connection between the two conductors, there is an area of localised heating present in the thermal images. Generating models based on

areas of conductor exhibiting typical levels of heating, and utilising data from the hottest area, treating is as 'faulty' data, and implementing model-based fault detection using it, could be a valuable experiment in evaluating the utility of the theme of the thesis.

It is noteworthy that, while the LSTM recurrent neural network produced more accurate models than the linear regression method, the LSTM method is highly non-linear, with multiple non-linear activation functions per LSTM unit in the network. There is scope for improving the regression analysis by introducing non-linear elements, while avoiding the inherent complexity, computational cost and abstraction of a recurrent neural network.

Given that the experiment is controllable, an investigation into formal system identification methods could provide value. This is not explored here as common input stimulation used in system identification practice is not representative of what may be experienced by in-service substations, and the opportunity to stimulate assets in substations with such load patterns are very rare. Nevertheless, a comparison between these formal methods and the modelling techniques presented, in this context, would provide insight into the LSTM and multivariate auto-regression methods performance, and provide a useful benchmark. Furthermore, rigorous system identification methods would allow the influence of individual inputs to be extracted from the model, which would be valuable to the project sponsor. It is not trivial to determine how well a system identification process completed in a laboratory environment would translate to in-service equipment, however.

3.9 Summary

In this chapter, initial findings regarding the output of the selected experiment sensors are presented. It is confirmed that the output of the non-calibrated FLIR Boson thermal camera provides a high-resolution indication of the heating effects of electrical loading on genuine power equipment, by grouping subsequent images into a set and treating pixels as time-series. It is found that the movement of air as caused by a large industrial fan in a laboratory environment can be measured by a Vaisala WXT520 weather station sensor array, and that the cooling effect this wind has on the electrically loaded conductors can be detected by inspection of pixels corresponding to conductor at which the fan is pointed.

It is found that the output of the thermal system, defined as an average pixel value (in units of digital counts of thermal flux) can be modelled using multivariate linear regression and LSTM recurrent neural network methods. A long-term pseudo-realistic dataset was captured and is presented in order to facilitate this finding. The LSTM method was found to consistently outperform the multivariate auto-regression method, and provided indications that the model output could be suitable for use in a fault-detection context.

The chapter presents data, equipment and sensors designed to emulate a long-term thermal and environmental monitoring system envisioned in a substation, observing a simple, electrically-loaded conductor. It advocates for the use of such long-term monitoring by indicating that good time-series models can be created from the multi-variate data, and that these models include the impact of a simplified set of environmental factors. This contributes towards addressing some general limitations of condition monitoring using thermal cameras, reducing environmentally caused false-negatives and providing indications of potential developing faults. This chapter contributes a development of the work presented by Bortoni et. al. in replicating the time-series analysis approach to thermal condition monitoring of substation power equipment, using LSTM recurrent neural network techniques and low-cost sensors in the laboratory environment.

Chapter 4

Modelling of a Pair of Cable Sealing Ends

4.1 Introduction

This chapter presents a development of the experiment and modelling techniques detailed in Chapter 3, using large and comparatively complex cable terminations observed by multiple thermal cameras. The experiment and it's justification are detailed in Section 4.2, with information about how the methodology differs from that presented in Chapter 3. Two datasets are presented in Section 4.3, comprising a system step-response experiment, and a long-term 11 day long experiment utilising a realistic load pattern and pseudo-random wind-pattern. Section 4.4 contains preliminary analysis and exploratory LSTM-based modelling of the system step-response experiment, while Section 4.5 contains preliminary analysis and LSTM modelling of the realistic 11 day experiment. Analysis and discussion are presented in Section 4.5.3.

The aim of this chapter is to prove that the methods presented in Chapter 3 are applicable to more sophisticated equipment that may be found in a substation environment, and that the methods are capable of modelling outputs under the influence of more input variables, namely the camera temperature and the direction of any wind.

4.1.1 Context

The application context is the same as previously presented in the Chapter 3, that is, to use these machine learning methods to create a multi-variate model of the thermal camera output under a range of normal operating conditions. The intention is that these models, if they prove to perform suitably, could be used to implement model-based fault detection methods, as described in Chapter 3, Section 4.2.3. Successful models created in this way will capture the effects of the changes in the environment in the laboratory, the electrical loading present in the assets, and the effects of the direct environmental influence (the wind). The successful creation of models from laboratory data could be assumed to be a suitable proof of concept, providing a solid basis for a potential data-collection scheme to be implemented in an operational outdoor substation. An installation in a substation environment would expose the system to the true range of environmental conditions experienced by equipment, and would be the ultimate test of whether the thermal output of equipment could be modelled by the proposed methods.

Due to the performance of the LSTM recurrent neural network compared to the multi-variate auto-regression in Chapter 3, only the LSTM method is considered here.

4.1.1.1 Description of Cable Sealing End Function

This section provides a brief description of cable sealing ends, for context throughout the thesis.

Ye et al. [66] provide a review of HVDC cable terminations for reference. A cable sealing end is a form of high voltage termination, typically used when a high voltage underground cable requires connection to an asset above ground, such as a substation. Underground high voltage cables typically comprise of a conductor core, surrounded by an insulator and covered by an earthed sheath. When connecting such a cable to external assets, the sheath is required to be stripped back from the cable, leaving only the conductor. Cable sealing ends house the point at which the cable sheath and insulation is stripped back from the cable assembly. The fully housed cable enters the bottom of the cable sealing end, and only electrical conductor emerges from the top.

Due to the potential difference between the conductor core, and the earthed sheath, there is an inherent area of high electric field between the sheath and

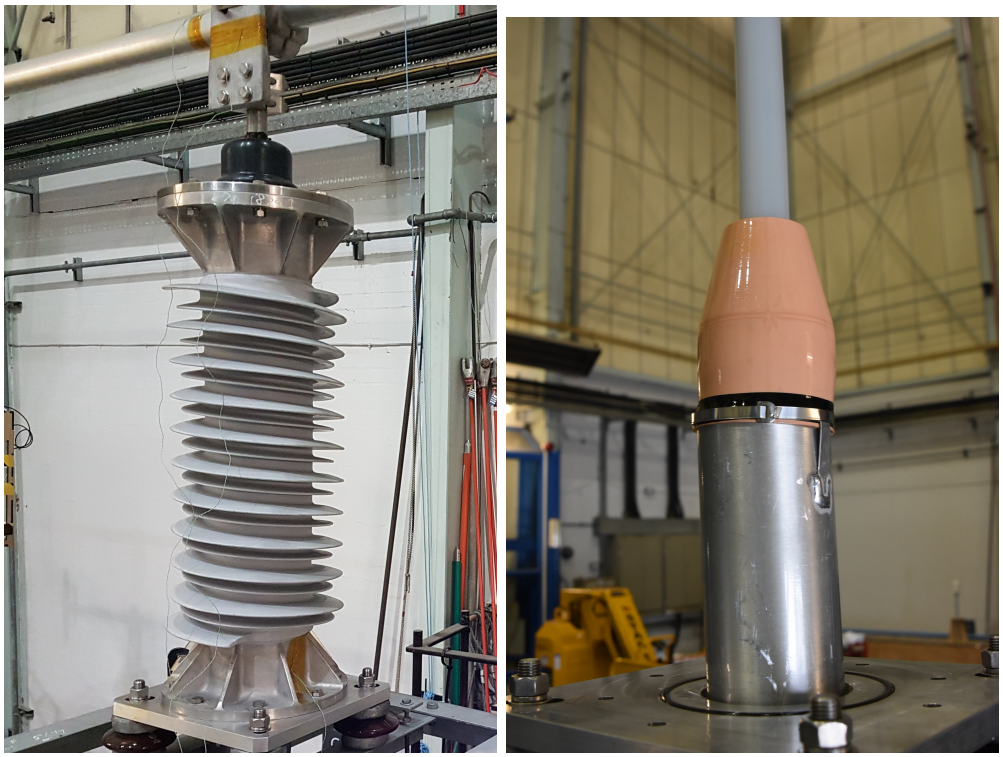

Figure 4.1: A 66 kV cable sealing end with and without outer sheath.

conductor. When the sheath is stripped from the conductor it is housed in a rubber cone, in order to minimise the sharp edges and resultant areas of very high field. This assembly is commonly housed in a cylinder filled with insulating oil. The cylinder is sealed from the environment, has polymer, porcelain or epoxy insulating sheds on its surface and is topped with a corona shield. Images of a cable sealing end with and without its outer sheath are displayed in Figure 4.1

The cable sealing ends utilised in the course of the work presented here are 66 kV cable terminations, with polymer sheds.

4.1.2 Justification

One of the aims of this chapter is to extend the methodology presented in the previous chapter to more sophisticated equipment. This should prove the utility of the method as a general tool for predictive maintenance in substations. It would prove that the modelling techniques work on both simple and complicated equipment. The optical coverage of the cable-sealing ends from 360° in the horizontal plane, provides an opportunity to investigate the capability of the adopted

Figure 4.2: A cable-sealing end in the B23 high-voltage laboratory, as seen from the approximate perspective of camera 8.

methods to model a more complete thermal profile of a device. Additionally, it provides an opportunity to stimulate a spatially varied thermal profile on the cable sealing end, by blowing air across it at different angles. This comprises a suitable increase in complexity, between that of the overhead line experiment, and data from an in-service experiment, further proving the utility of the method.

Significantly, a cable-sealing end is representative of what is commonly found in substation environments: a high-voltage conductor housed in a substantial amount of insulating material, with a connection to a different type of conductor. The insulating material is designed as an electrical insulator, but provides significant thermal insulation too. This means that any internal thermal changes will take longer to present on the surface of the equipment, than if the whole unit were an electrical conductor, as in the case of the overhead line. This may contribute towards a thermal system which is more difficult to model. The insulator in the case of the model of CSE used in this experiment (see Figure 4.2) consists of the oil contained within the CSE, the thick rubber sheds on the exterior and within the lower half of the CSE, the sheath housing around the cable. A further factor that contributes to the CSE representing a progression from the Overhead Line experiment, is the external geometry of the units. The insulating sheds are arranged in spiralling helix shapes down the length of the CSE cylinder, alternating between two diameters. The CSE is capped by metallic units, consisting of a core, bolts, and supporting fins. Both this cylindrical geometry, with extended sheds, and the simpler conductor geometry, are typical of what may be found in substations. The added complexity of the CSE increases the difficulty of selecting points of interest, introduces a greater variation in angles between the thermal

camera and the infrared emitting surface and that any effects of external heating or cooling are less trivial to quantify. The multiple materials the CSE consists of means there are varying emissivities within the thermal imaging field-of-view, potentially contributing to a less physically representative initial thermal image, a further consideration.

The CSE experiment is mounted on a large scaffold, meaning it is restricted to the largest high-voltage laboratory in the university, Laboratory B23. B23 houses the large high-voltage equipment, such as a DC test-set and an impulse generator. The size of these pieces of equipment necessitates a high ceiling, meaning the laboratory contains a large volume of air and therefore is costly to heat. For this reason, the laboratory climate control is configured to only activate when the temperature falls below a threshold, above which value the temperature may fluctuate freely, warming during the day and cooling at night. Additionally, there is a large cargo door, which is used for deliveries, contributing further to the thermal volatility of the laboratory space. This volatility, while not equivalent to real-world outdoor environmental conditions, is still added complexity when compared to the stable air temperature of the overhead line experiment.

In order to manage the additional complexity of the different environment and equipment, the FLIR Boson cameras have been supplemented with additional equipment. The approach was to quantify and dampen the environmental effects on the cameras. A temperature sensor was affixed to the outer body of the camera and interfaced with the controller Raspberry Pi of that particular camera. This camera and sensor assembly was then enclosed in a polystyrene container, with an opening for the lens, to insulate it. This allowed the temperature of the camera body to be measured and recorded, while protecting the camera units from any extreme thermal gradients present in the laboratory.

The additional complexities listed in this section constitute justification for a second experiment, as it iterates on the work completed modelling the overhead line, making it more representative of in-service substation equipment monitoring.

Complementing this work, some additional analysis regarding the heating profiles of cable-sealing end assets over time and space has been completed and is presented in Chapter 5.

4.2 Experiment Design and Data Collection

This section describes the specification, design and implementation of the experimental setup used to acquire data for the purposes of Chapters 4 and 5.

4.2.1 Context

The specification of a second experimental setup was a result of the success of modelling work completed on data acquired using the overhead line experiment. The requirements from the overhead experiment remained:

1. Generate time-series data
2. Comprise of genuine power-engineering equipment
3. Utilise electrical loading representative of in-service equipment
4. Comprise some form of measurable environmental variation
5. Record a sequence of thermal images
6. Record environmental data
7. Be autonomous and safe, for long-term usage

The following requirements were added, in order to comprise a substantial development of the findings from the overhead experiment and form a contribution to the field. The experiment must:

8. Thermally observe equipment from multiple perspectives
9. Comprise of power-engineering equipment more thermally complex than a simple conductor
10. Comprise of power-engineering equipment with more complex physical geometry
11. Include a form a directional environmental variation

A partnership between the University of Manchester and National Grid, the project sponsor, was beginning an investigation into cable-sealing ends during the period in which the second experiment was being specified. A National Grid

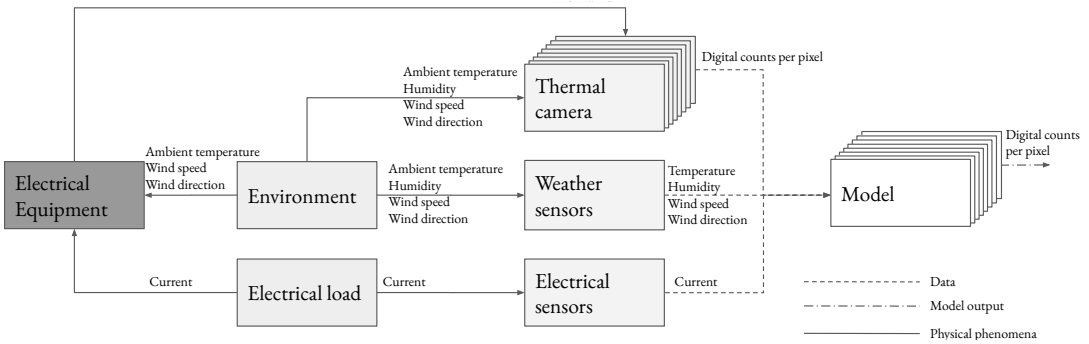

Figure 4.3: Conceptual diagram of the data collection and modelling system. On the left is the physical experiment, which is perceived by sensors in the centre, which contribute data to the model, which produces digital count pixel value outputs. Eight thermal cameras are shown, with a model for each.

Electricity Transmission (NGET) Innovation Award was applied for and granted, securing £50,000 of funding to support the development of this experiment with equipment and expertise.

A diagram depicting the conceptual layout of the system is shown in Figure 4.3. Note that the scenario is fundamentally the same as that presented in Chapter 3, other than there being multiple thermal cameras and corresponding models for the pixels of these cameras. This development in the concept provides a foundation for understanding the differences in environmental effects in space around a substation asset.

4.2.2 Experiment Design

As part of the National Grid and University of Manchester investigation, the high-voltage laboratories at the university housed a large scaffold assembly consisting of two cable sealing ends in a current loop. The two CSEs were connected at the bottom by a length of sheathed high voltage conductor, and at the top by a section of aluminium bus-bar. The current loop was capable of being electrically loaded by a 0 - 3000 A cable loading transformer. The primary coil of the loading transformer was connected to a motorised variac (variable transformer), allowing the current through the current loop to be controlled. Electronic control of the motorised variac was possible using digital logic inputs. The assembly was positioned in a corner of the laboratory, with the CSE nearest the corner denoted as CSE1, and the other CSE2.

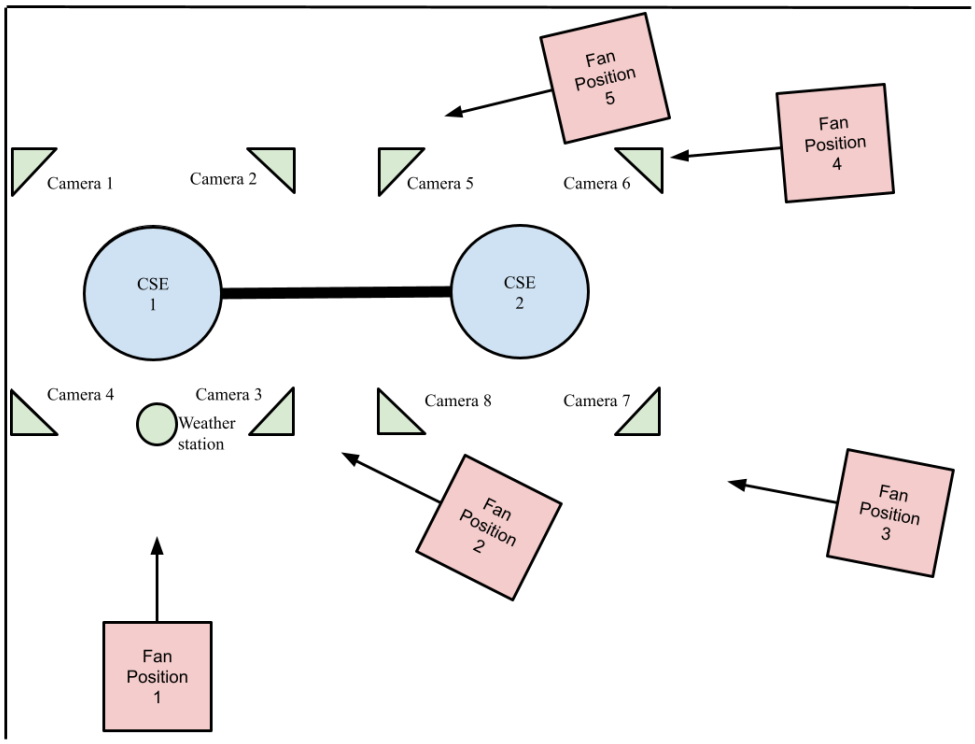

Figure 4.4: five wind fan positions, relative to the cable sealing end experiment rig

An experiment was devised which met the requirements detailed in Section 4.2.1, while utilising the National Grid cable sealing end energising assembly.

The Vaisala WXT 520 weather station, used in the overhead line experiment, was again utilised here, positioned near cable-sealing end 1. The larger laboratory where the CSE assembly was housed was found to have a more volatile environmental temperature than the small lab where the overhead line experiment was housed, serving to provide a desired aspect of environmental variation. The large wind-fan used in the overhead line experiment was re-purposed for use in this experiment. In order that the wind could be effectively measured, the fan was aimed at only one of the cable sealing ends, near the position of the weather station. For the purposes of this experiment, the fan speed was also varied. The five wind-fan positions used during the course of the experiment are shown in Figure 4.4 in a top-down diagram view of the experiment. The fan is typically mounted low, tilted slightly upwards towards the CSE body. Presence of air disturbances was physically checked for each position.

A set of eight FLIR Boson thermal camera cores with 4.3mm lenses were

Figure 4.5: The full cable sealing end experimental rig

acquired, allowing a 50° horizontal field of view. This configuration allows full vertical coverage of the 1.3 m tall CSE at a 1.4 m clearance. Wooden arm-mounts for the cameras were designed and manufactured, positioning the cameras around the cable sealing ends at 90° intervals. The length of the arms and lens choice were chosen in order that the entire length of the cable sealing end appear in the frame of each of the cameras. Each camera was specified with it's own Rasberry Pi host, in order to avoid issues synchronously capturing images from eight USB thermal cameras on single low-power SBC. The Raspberry Pis were powered and networked using POE, allowing a single cable connection to each. Software was written such that broadcast messages were sent on the Ethernet network requesting an image capture event when required. The Raspberry Pi internal clocks were synchronised to a 'parent' Raspberry Pi device housed on a base station using network-time-protocol. The parent Raspberry Pi device controlled the experiment, triggering image and environmental data capture events, and controlling the input current and wind. The Raspberry Pi controlled current by sending a target set-point to a micro-controller capable of controlling the motorised variac control relays. It received current feedback utilising a FLUKE i2000 current

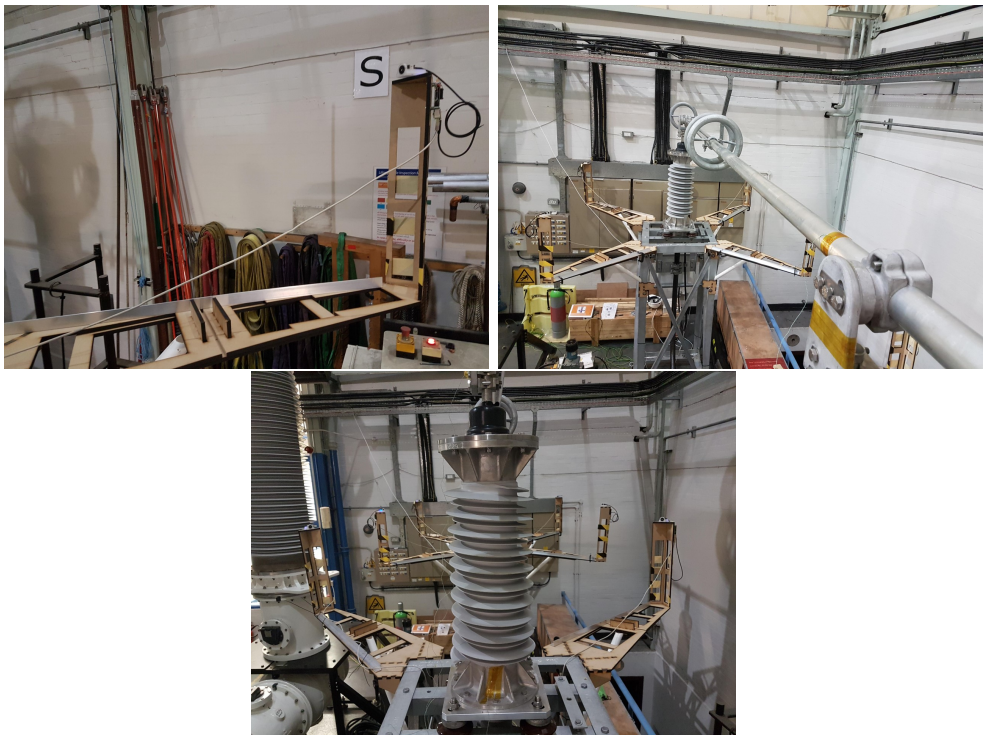

Figure 4.6: Top left: image of a camera supporting arm, with a thermal camera visible at the top; top right: image of the top conducting bus-bar spanning to the second CSE; bottom: image showing CSE and two of the four camera supporting arms.

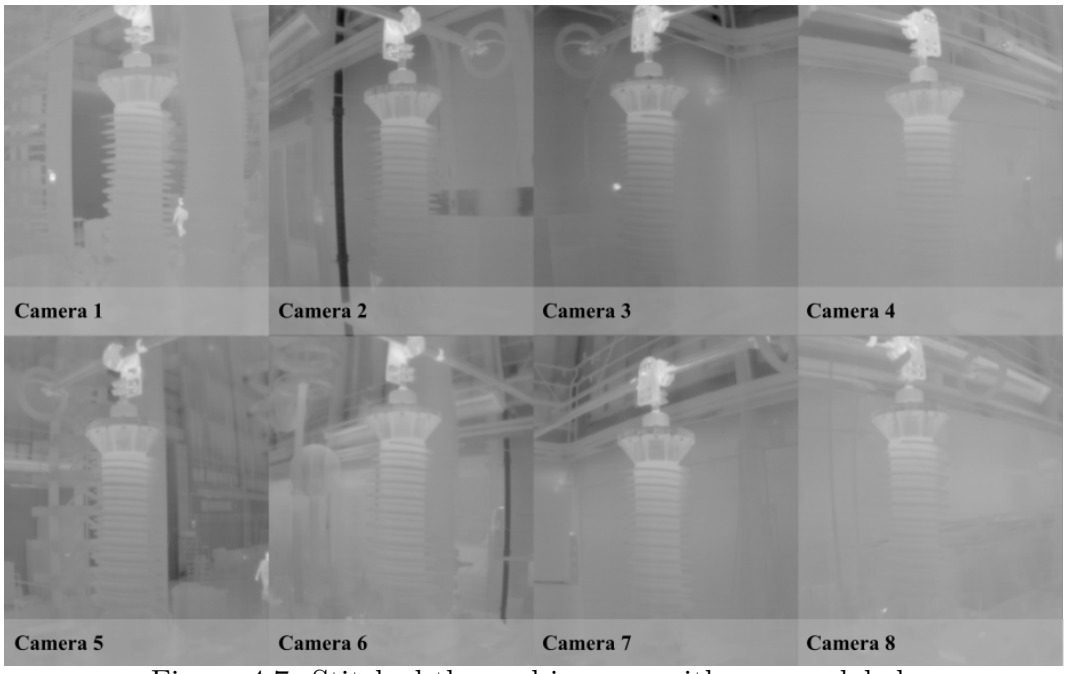

Figure 4.7: Stitched thermal images, with camera labels

clamp, which output it's values to a serial-enabled digital multimeter, which then output the values to the Raspberry Pi.

Safety measures included: software-based over-current triggers, a high-current relay between the variac and cable loading transformer, and temperature monitoring with thermocouples.

Images of the described experimental setup can be seen in Figure 4.5 and Figure 4.6.

Output images, arranged into a grid with corresponding camera labels, are depicted in Figure 4.7.

Various experiments were run on the setup, including two primary research experiments. The first of these was a 40 hour step test. During this experiment current was stepped to 300 A, and maintained at that level for 40 hours, with data samples every minute. The second primary research test was a realistic loading test. During this experiment the cable sealing ends were loaded with a realistic load pattern based on national energy usage, scaled to fit within the operating current range of the assets, for 11 days. The load pattern was simplified for the weekend periods, minimising changes of current using the motorised variac during extended periods during which the laboratory was un-occupied. During various periods of the experiment the wind-fan was enabled and disabled. Four wind positions were selected and used, and the fan speed level was selected at random.

Table 4.1: Point of interest pixel location reference. All pixels are listed in (row,column) format.

CSE	Camera	POI	Pixel	CSE	Camera	POI	Pixel
1	1	1	(021,127)	2	5	1	(368,123)
1	1	2	(058,128)	2	5	2	(405,120)
1	1	3	(107,126)	2	5	3	(449,122)
1	1	4	(167,128)	2	5	4	(514,124)
1	1	5	(238,128)	2	5	5	(571,126)
1	1	6	(265,130)	2	5	6	(609,133)
1	2	1	(041,371)	2	6	1	(349,378)
1	2	2	(078,376)	2	6	2	(384,385)
1	2	3	(123,378)	2	6	3	(433,383)
1	2	4	(197,380)	2	6	4	(491,386)
1	2	5	(254,382)	2	6	5	(564,390)
1	2	6	(285,387)	2	6	6	(591,394)
1	3	1	(055,630)	2	7	1	(380,633)
1	3	2	(092,630)	2	7	2	(415,635)
1	3	3	(145,636)	2	7	3	(460,636)
1	3	4	(211,635)	2	7	4	(533,637)
1	3	5	(282,640)	2	7	5	(591,639)
1	3	6	(300,643)	2	7	6	(623,641)
1	4	1	(042,918)	2	8	1	(367,890)
1	4	2	(083,916)	2	8	2	(402,888)
1	4	3	(129,918)	2	8	3	(452,891)
1	4	4	(195,920)	2	8	4	(518,886)
1	4	5	(253,924)	2	8	5	(576,893)
1	4	6	(290,927)	2	8	6	(608,897)

These two primary research experiment runs are further detailed in Chapter 4.

Points of interest were selected for each camera perspective, corresponding to condition monitoring strategy (see Chapter 5). The points of interest selected are listed in Table 4.1 and depicted in Figure 4.8.

4.2.2.1 Thermal Camera

In order to minimise the potential negative effects on camera output due to thermal transients [56], the impact of changes in temperature were both minimised and quantified. In order to achieve this, external temperature sensors were fitted to the outer casing of the camera body. The camera and sensor assembly was then placed inside an insulated enclosure, with an opening for the optical

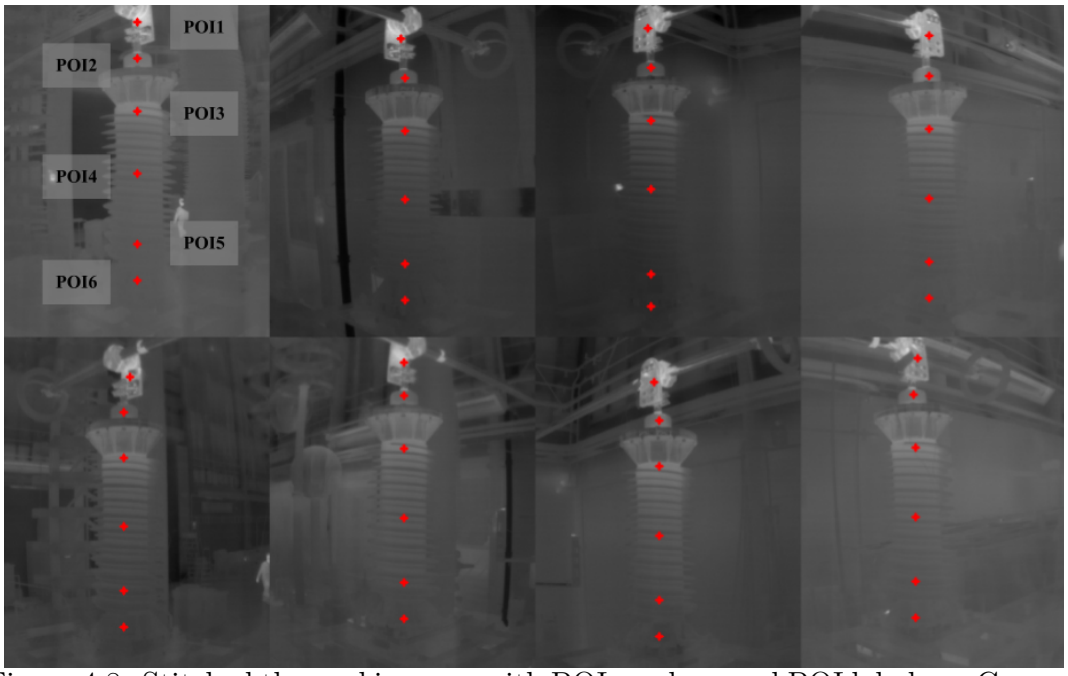

Figure 4.8: Stitched thermal images, with POI markers and POI labels on Camera 1 perspective

lens, intended to insulate the camera from thermal transients. The camera in the enclosure, with temperature sensor attached is depicted in Figure 4.9.

The temperature sensor used for this was the DS18B20. The sensor operates on a 1-wire interface, allowing easy integration into the Raspberry Pi based system (see Sections 3.5 and 4.2). The sensor can measure temperatures from -10 °C to 85 °C with ±0.5 °C accuracy (and extended temperature range at lower accuracy) with 12-bit digital resolution.

It is notable that there is a difference in the smoothness of the output of the Boson used in Chapter 3 and the Bosons used in Chapters 4 and 5. FLIR were contacted regarding this difference, and provided assurances that the difference was due to updated flat-field correction provision in the firmware of the later models used in Chapters 4 and 5.

Thermal cameras are susceptible to environmental conditions. As previously mentioned, a number of factors are acknowledged to influence perceived temperature of an observed object. The Unmanned Aircraft Systems community has studied this extensively, as the environmental factors influencing thermal images are exaggerated somewhat when the camera is subject to wind-speed from moving aircraft and inevitably high view-distance (due to an effort to maximise coverage per flight and battery charge).

Figure 4.9: Left: FLIR Boson in custom mount enclosed in insulated case; right: external temperature sensor attached to camera housing inside insulated case

Kelly et al. [28] published a comprehensive paper on considerations for using a camera similar to the FLIR Boson chosen for this work. A summary of the recommendations made relevant to this work is:

- allow 15 minutes of stabilisation time;
- enable frequent NUC;
- shelter camera from wind;
- correct for temperature drift post-flight.

The listed points all pertain to controlling and correcting for external temperature influence on the thermal camera. This is applicable to the CSE laboratory tests.

4.2.2.2 Current

A current measurement device was required for the CSE experiment, as, unlike the OHL experiment, the cable conductor was loaded by a series of transformers with no direct measurement functionality. A Fluke i2000 current measurement clamp was used to take the initial measurement. The i2000 output a voltage waveform at an amplitude of $1mV/A$. This was output directly to a digital multimeter with serial communication functionality, allowing the amplitude of the voltage waveform to be transmitted to the controller computer, where the value could be converted back to Amps.

4.2.3 Modelling and Fault Detection

One aim of the project is to investigate the viability for these methods of inspection for modelling thermal output. This means the effectiveness of the model needs to be evaluated. There are three metrics that can be used to evaluate the initial model skill: r^2 score, Root Mean-Squared Error (RMSE) and mean absolute error. These provide a method of evaluating whether the model is a good fit, and whether the model is overtrained for the training-set or is a more general model. They provide little insight into whether the dynamics of the system have been captured by the model. To evaluate the capability of the model to capture the dynamics of the system, methods from classical regression analysis are utilised. The residuals of the model are calculated. The standard deviation of the residuals is calculated. If the standard deviation of the residuals is not centred around 0, the model is effectively discarded - the model created by the network has a DC gain or bias and as such has skewed residuals. If the standard deviation of the residuals is roughly centred around zero and has approximately 'normal' distribution, the model is accepted. The creation of models to characterise the system is deemed to be useful in the context of fault detection in control theory as presented by Hwang et al [20]. The concept presented is that a healthy system should be characterised by a model with zero-mean residuals. Following this, fault detection can be achieved by 'using statistical methods to test if the residuals have significantly deviated from zero'. While implementation of this type of fault detection is out of the scope of the project, it forms a significant part of the justification for the attempts to generate models of the power equipment thermal output.

4.3 Datasets

This section presents the datasets used for the modelling discussed in this chapter, with pertinent features highlighted and discussed. The sensors and variables used for the modelling process described in this chapter are presented in detail in Chapter 3. The variables selected for the modelling process are presented in Table 4.2.

Table 4.2: CSE Experiment system output and inputs. * indicates variables unused in the step-test experiment.

System output	System inputs
Digital counts	Wind speed* Wind direction* Current Air temperature Humidity

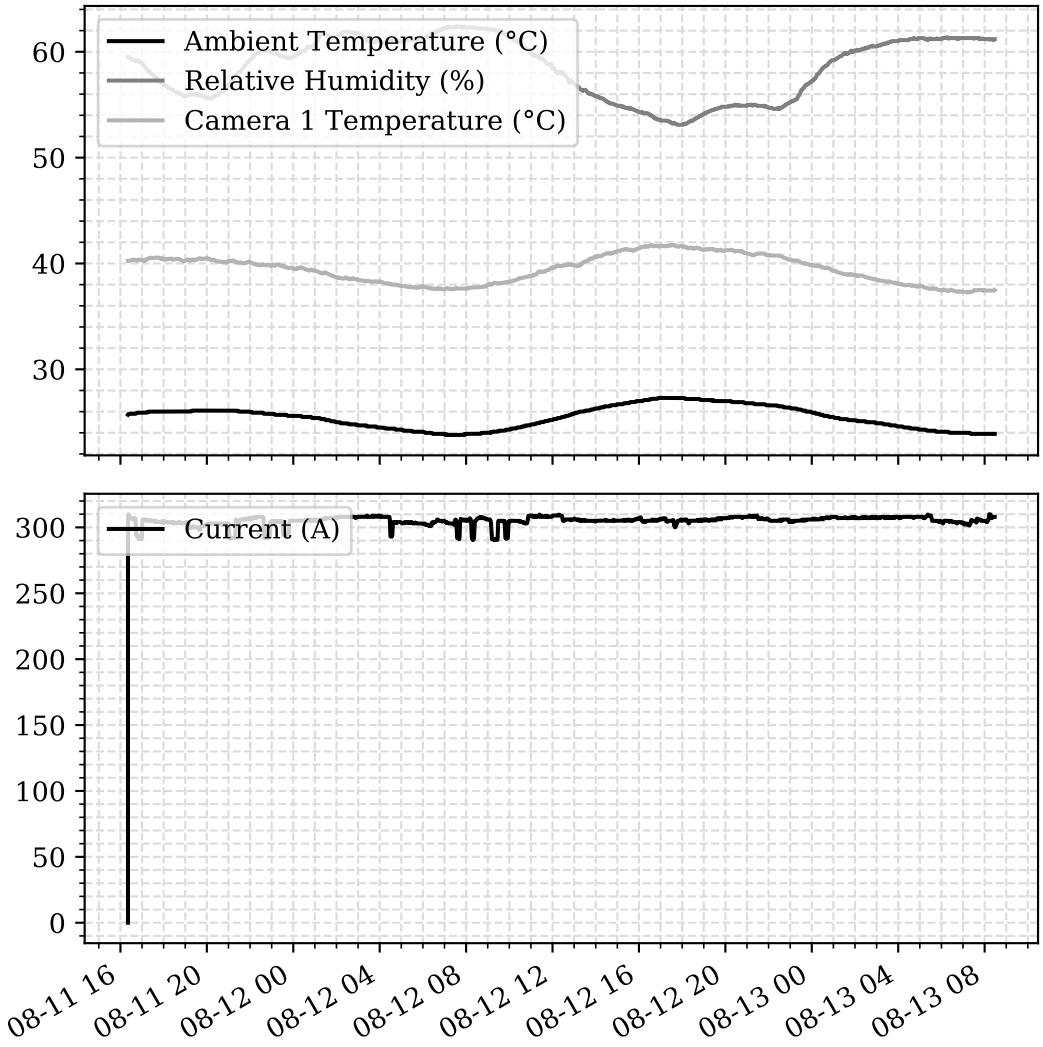

Figure 4.10: Cable sealing end characterisation step-test response experiment, input data time series. Note x-axis labels are in the datetime format 'mm-dd hh'.

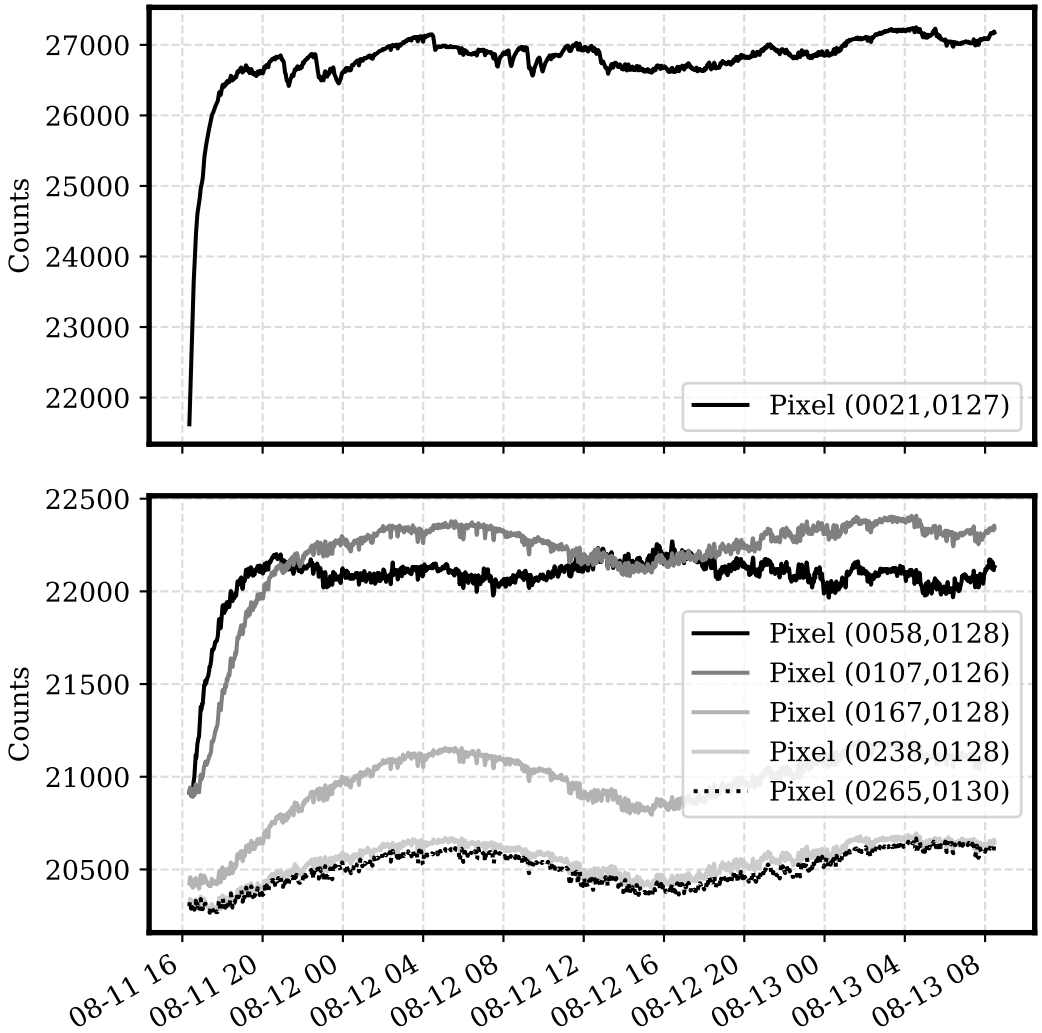

Figure 4.11: Cable sealing end characterisation step-test response experiment, output data time series. Note x-axis labels are in the datetime format 'mm-dd hh'.

4.3.1 System Step-response Test

Figures 4.10 and 4.11 show the inputs and outputs of the characterisation step-test respectively.

The test runs over a 36 hour period, allowing time for the current-driven output to settle and to observe the effects of the daily fluctuations in laboratory temperature. This fluctuation is clear for points of interest 3, 4, 5 and 6. Inspecting the system inputs in Figure 4.10, it is clear the apparent thermal output responds inversely to the camera core temperature, which is fluctuating due to the change in ambient temperature. This matches the expectation, on the principle of the camera measuring incident thermal flux. Points of interest 1 and 2 respond differently to 3-6. Point of interest 1, corresponding to the bus-bar clamp, heats a significant amount more than the rest of the cable-sealing end assembly. The greater magnitude of the heating effect means that the fluctuations in output due to momentary drops in input current are more visible than those of other points of interest. The impact of the air temperature changes is present, a peak following '08-12 04:00' matches the peak of the other points of interest, with the exception of point of interest 2. Point of interest 2 corresponds to the rubber bulb surrounding the conductor as it emerges from the top of the cable sealing end. The response of this point in unique within the set of 6 points of interest. After the initial period of current-driven heating, the apparent thermal output drops, while the others continue to rise. Potential causes of this discrepancy include: a different thermal time-constant of the insulator, the difference in emissivity due to the colour of the rubber. The data still shows a similar fluctuation, with a magnitude of approximately 250 digital counts, but in this case it peaks when the other points of interest trough. The different time-constant could cause the different two heating effects of the input current and the lab temperature to interact, causing the change in output. Wind is not included in the descriptive figures, as it was not used for this test.

4.3.2 Long Realistic Test

Figure 4.13 presents the output values of the CSE points of interest as observed from the perspective of thermal camera 1. Table 4.3 presents the minimum and maximum numerical values of digital counts for each point of interest during the experiment. Referring to Figures 4.12 and 4.13, initial exploratory analysis can be

Table 4.3: Minimum and maximum numerical values for output 'Digital Count' of each point of interest during the cable sealing end realistic experiment.

Cam.	POI	min	max	range	Cam.	POI	min	max	range
1	1	20448	25084	4636	5	1	20422	21614	1193
1	2	20463	22006	1543	5	2	20419	20898	479
1	3	20445	22120	1676	5	3	20410	21050	640
1	4	20297	21391	1094	5	4	20381	20917	536
1	5	20217	21208	991	5	5	20332	20830	498
1	6	20216	21191	975	5	6	20338	20813	475
2	1	20705	24729	4024	6	1	20918	21927	1008
2	2	20764	21847	1084	6	2	20930	21372	442
2	3	20714	22054	1341	6	3	20914	21491	577
2	4	20677	21250	573	6	4	20887	21342	455
2	5	20631	21183	552	6	5	20870	21235	365
2	6	20655	21182	527	6	6	20913	21225	313
3	1	21085	23736	2652	7	1	20653	21235	582
3	2	21127	22086	959	7	2	20607	21062	455
3	3	21042	22217	1175	7	3	20554	21068	514
3	4	20999	21467	468	7	4	20506	20853	348
3	5	20964	21393	429	7	5	20475	20724	249
3	6	21019	21421	403	7	6	20458	20801	343
4	1	20331	23257	2926	8	1	20654	21356	701
4	2	20284	21817	1534	8	2	20656	21015	359
4	3	20258	21802	1544	8	3	20551	21122	572
4	4	20200	21031	831	8	4	20513	20939	426
4	5	20135	20926	791	8	5	20497	20857	360
4	6	20127	20917	790	8	6	20499	20837	338

completed. It is clear from inspection that there are two primary output drivers. Firstly, it can be seen that current effects the thermal output. This can be most clearly seen for the pixels corresponding to points of interest high up on the cable sealing end unit, or on the bare conductor of the busbar clamp. The current influence is present on the lower pixels, but to a lesser extent. The second of the two main output drivers is the wind. The impact of the wind on the output is less intuitive than that of the current, because there are two effects of the wind. Firstly, there is a cooling effect on the cable sealing end. This is a desirable facet of the experiment, as it provides a real-world environmental effect to model with the data.

The second, more significant effect of the wind, is the effect it has on the

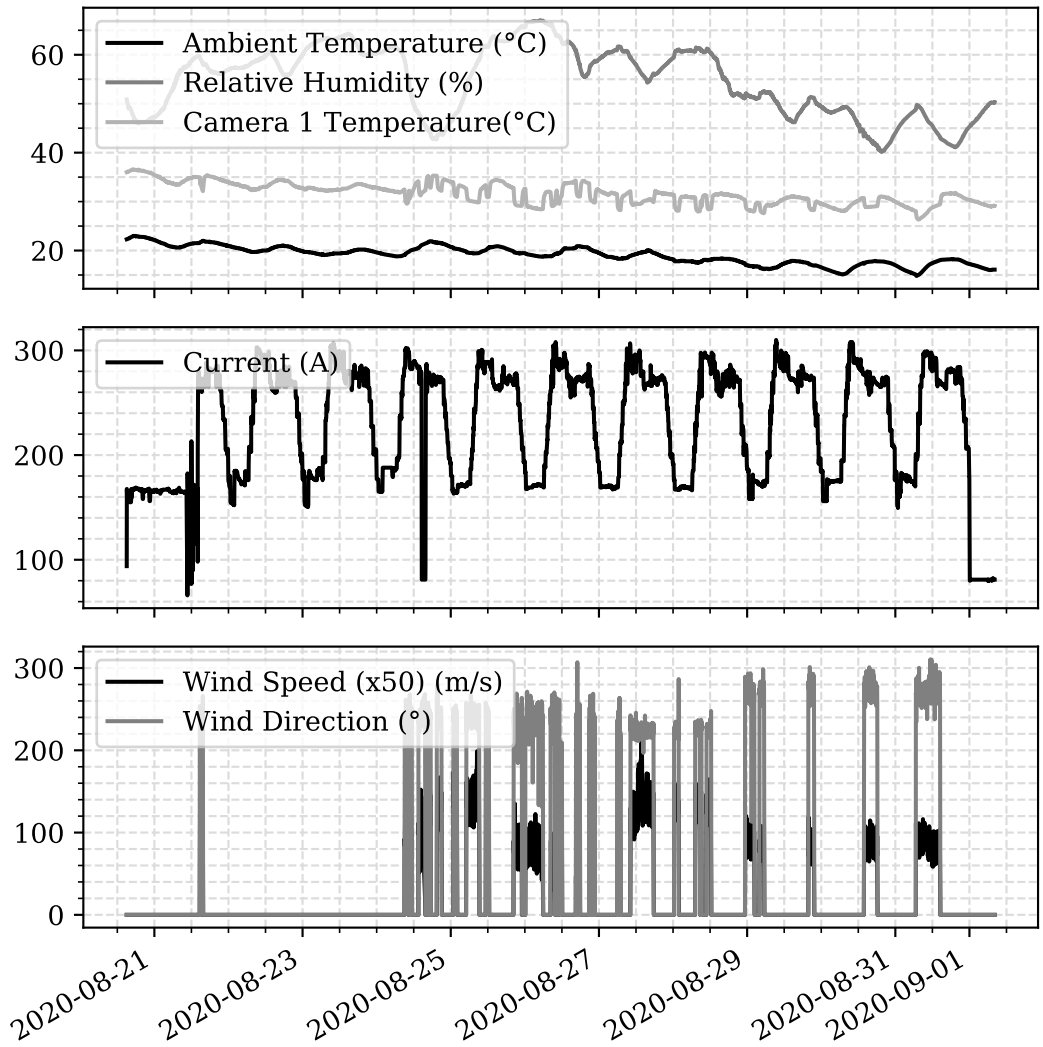

Figure 4.12: Cable sealing end realistic experiment; input data time series. Note x-axis labels are in the date format 'yyyy-mm-dd'.

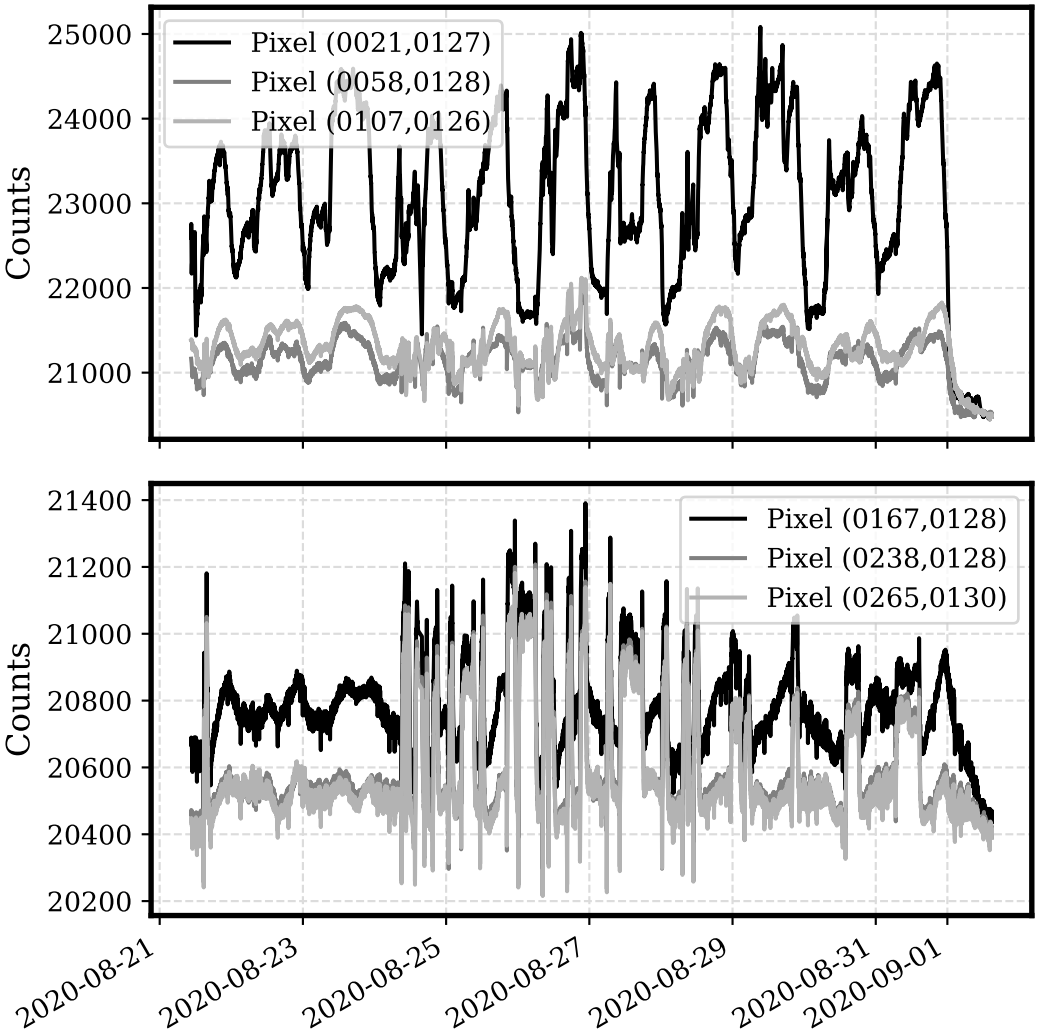

Figure 4.13: Cable sealing end realistic experiment; output data time series. For clarity, the 6 pixels are shown in two graphs. Note x-axis labels are in the date format 'yyyy-mm-dd'.

camera temperatures. 'Camera 1' is mounted to monitor 'CSE 1' and due to being near to the CSE, is in the path of the cooling effect of the wind. As described in Section 3.2, the camera cores are sensitive to their temperature. In Figure 4.13, particularly the lower axes, some periods of relatively high output fluctuation are present. With reference to Figure 4.12 it is clear that the fluctuation is a result of the wind input. The camera temperature These fluctuations in output are present on the upper plot, corresponding to the higher points of interest too, however due to the larger range in the axes, they appear less prevalent. It is clear the wind has a significant impact on the output, and while the net impact is not representative of what a calibrated camera would present in a real-world situation, there remains merit in modelling the influence as it allows the feasibility of using this type of un-calibrated camera to be evaluated. Furthermore it still provides a temporally aligned input-output relationship to model.

The final observation is that the heating effect is relatively small, especially when compared with the overhead line experiments in Chapter 3 The outputs from this camera are typical of what is seen from each of the camera perspectives, though due to the positioning of the wind fan at different points in the test, the wind impact on the output at each perspective differs slightly. The output from Cameras 5-8 is not impacted by the wind as the fan was pointed exclusively as CSE1: it was decided to focus the wind-influence study on a single cable sealing end.

4.4 System Step-response Experiment

4.4.1 Preliminary Analysis

The manual analysis for the step-test on the cable sealing end experiment is intended to confirm that the system behaves as expected under electrical load and establish an understanding of the influencing factors when considering data from multiple un-calibrated thermal cameras. Figure 4.11 is a plot of a number of different pixels from the same perspective, Figures 4.14 and 4.15 are plots of a matching points of interest viewed from different camera perspectives.

On inspection of Figure 4.11, it can be determined that there is an approximate first order step response to the input current, as expected from thermal behaviour. The time constant varies between the points of interest, with notable

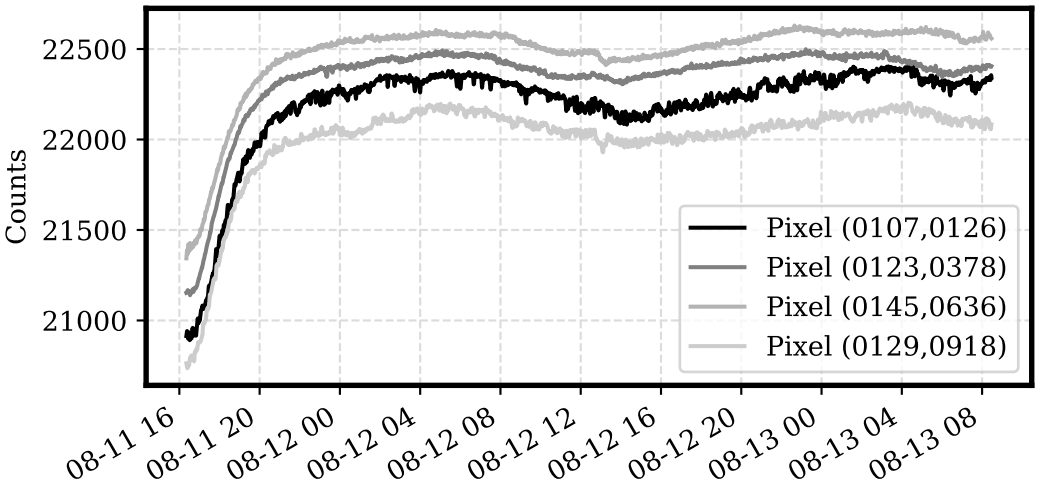

Figure 4.14: POI3 for the duration of the step-response test, from camera perspectives 1, 2, 3 and 4. Note x-axis labels are in the datetime format 'mm-dd hh'.

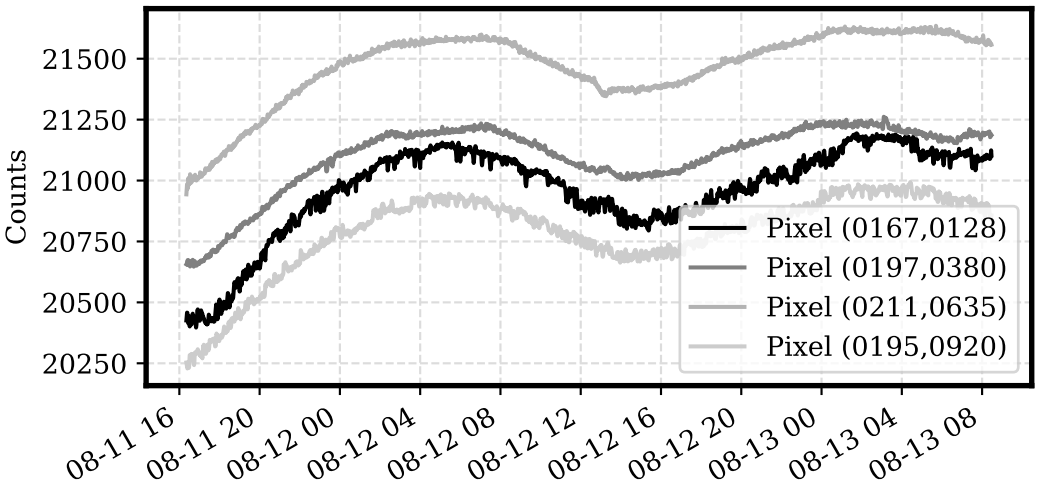

Figure 4.15: POI4 for the duration of the step-response test, from camera perspectives 1, 2, 3 and 4. Note x-axis labels are in the datetime format 'mm-dd hh'.

differences between the points of interest corresponding to loaded conductors (POI1) and those corresponding to components insulated from the conductor (POI2, 3, 4, 5, 6). Clearly, POI1 displays the largest change in output, due to the busbar:cable-sealing end being an area of high resistance relative to the rest of the conductor loop. Notably, POI1 is one of only two metallic points of interest, the other being the base, POI6. These two points of interest differ in that POI1 is the electrical conductor, while POI6 is part of the insulated cable-sealing end shell.

POIs 3,4 and 5 represent three points down the length of the rubber sheds of the cable-sealing end body. There is a temperature gradient present on the CSE, with the hottest portion of the unit corresponding to the top, getting cooler down the length. This is confirmed inspecting image data from the test, as per Figure 4.8.

POI 2 corresponds to the top bulb of the CSE, a rubber dome surrounding the conductor protruding from the top. This is notable for it's similar material to POI2 but darker colour, potentially causing a difference in emissivity and therefore apparent thermal output.

While the digital count output corresponding to the overhead line experiment was smooth, displaying little in the way of high frequency fluctuations, these are present in the cable-sealing end data output. The FLIR Boson camera module used for the overhead line experiment was purchased 18 months prior to those used in the cable sealing end experiment. As such, they were shipped with different firmware versions. The fluctuations occur at a frequency of 0.2 Hz, matching the maximum period between 'Flat Field Corrections' of the FLIR Boson.

Largely the plots match well between perspectives. There is some bias inherent in the cameras, which is expected for un-calibrated thermal cameras (this will be explored more in Chapter 5). This manifests itself as a steady state difference in the plots. There are some instances where there is a fluctuation in the output of a subset of all of the cameras. It is expected that these occur due to localised environmental conditions in the laboratory. For example, if the cargo door of the laboratory is opened for a period, cameras closest to the door may experience a thermal transient. The higher frequency fluctuations are present in every one of the camera perspectives, but the extent of the fluctuations varies between cameras.

Table 4.4: Hyper-parameter selection for cable sealing end step-response test

Hyper-parameter	Value
L1 Regularisation	0.01
L2 Regularisation	0.01
Dropout Layer Rate	20%
Shuffle	True
Batch size	1000
Train:test ratio	85:15
No. observations	2396

4.4.2 LSTM Analysis

This section presents initial attempts made to model the data from CSE experiments. The intention is to prove that the LSTM time-series modelling method is capable of modelling the pixel output in a similar manner to that of the overhead line experiments. The significance of this is that, as described in Section 4.1 the cable sealing end is a more sophisticated system in a more complex environment. Proving that the method can model the pixel output of a simple test, given these factors, proves the concept, preceding attempts to model a longer, more complicated dataset.

In order to do this, two points of interest were selected. The first is POI1, from the perspective of Camera1, as it corresponds to a part of the test object which is a simple metallic conductor. This is effectively the point of the cable-sealing end system that is most similar to the simple conductor of the overhead line experiment, while being a legitimate point of interest for inspection engineers. The second point of interest selected is POI3 from the same camera perspective. This point corresponds to the upper end of the sheds on the cable sealing end body, which has a higher thermal mass than a simple conductor and more complex physical geometry.

The results pertaining to POI1 are presented in Figures 4.16, 4.17, 4.18 and 4.19 and those pertaining to POI3 are presented in Figures 4.20, 4.21, 4.22 and Figure 4.23. The hyper-parameter configuration for these tests is consistent between the two, and is presented in Table 4.4. Plots of the input and output data are presented in Figures 4.10 and 4.11.

On inspection of the results for POI1, it can be determined that the modelling was successful. The model output (hypothesis) closely tracks the real output

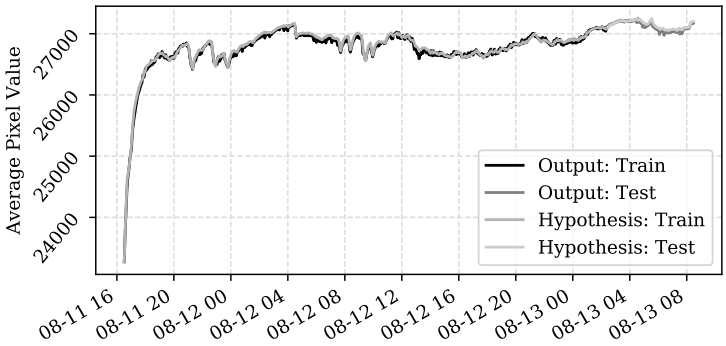

Figure 4.16: CSE, POI1, step-response test, LSTM model. Training score plotted for training time of 300 epochs, for both test-set and train-set data.

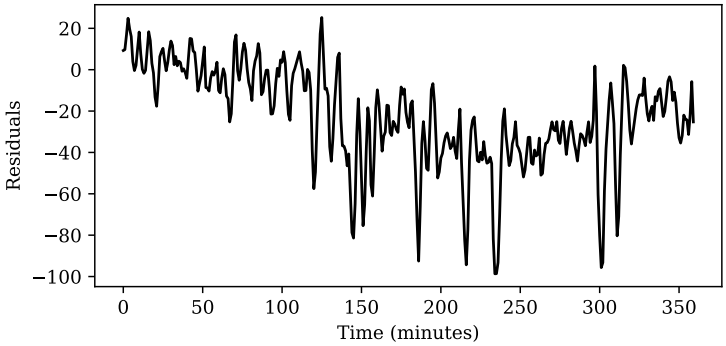

Figure 4.17: CSE, POI1, step-response test, LSTM model. Real value - predicted value for the time period forming the test-set portion of the dataset.

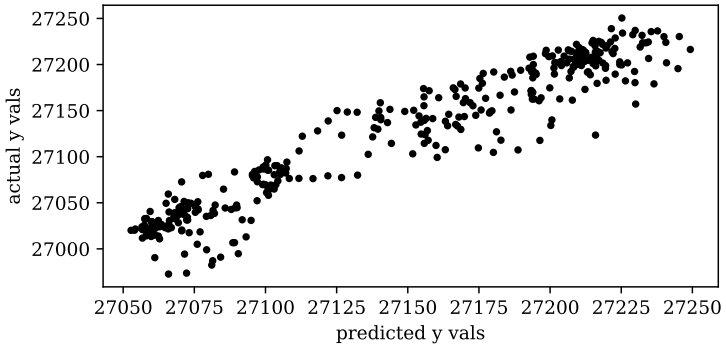

Figure 4.18: CSE, POI1, step-response test, LSTM model. Data-points of the test-set, prediction vs actual value scatter graph.

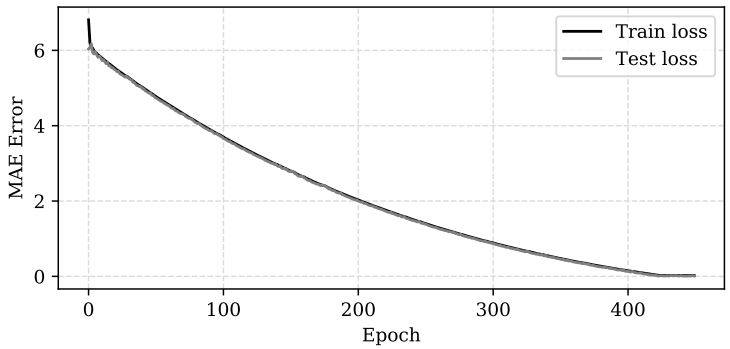

Figure 4.19: CSE, POI1, step-response test, LSTM model. Training score plotted for training time of 300 epochs, for both test-set and train-set data.

throughout the time-series. Both the high and low frequency fluctuations of the output signal are able to be modelled. The high frequency fluctuation represents the response of the output to the changes in the input current, causing the most significant visible changes in the output signal. The low frequency fluctuation shows the response of the output to the changes in air temperature in the lab environment. This can be seen clearly at the trough in the output signal at '08-12 16:00' and the peak at '08-13 04:00'. The laboratory cools down in the evening, causing a drop in the temperature of the camera and a corresponding increase in the flux incident on the camera. Figure 4.19 shows that the model is not overfit to the training dataset, as the costs fall at similar rates, settling at similar values. Inspection of the plots before and after '08-13 02:00' confirms that the model performs similarly both before and after the train/test dataset split. This is particularly good in the context of a small dataset like this one, where it may be expected that there are not enough data to allow the model to generalise well. The residuals of the model do imply that with adjustments to the method, the model could be improved, however for the purposes of proving that the method can successfully model the output of the cable-sealing end experiment, the model is deemed sufficient.

The results of the modelling of POI3 are similarly successful. The timeseries plot of the output is similar in terms of the high-frequency response. There is a large first-order step response at the beginning of the test, with some fluctuations as the current controller adjusts the input during the test. However, the low-frequency response differs from the response shown by POI1. At '08-12 16:00' there is a peak in the data, while in the POI1 test this was a trough. The cause of this discrepancy between the two tests is suggested to be the increased

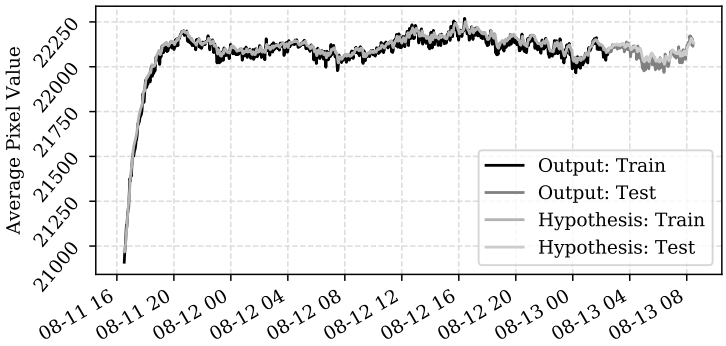

Figure 4.20: CSE, POI3, step-response test, LSTM model. Training score plotted for training time of 300 epochs, for both test-set and train-set data.

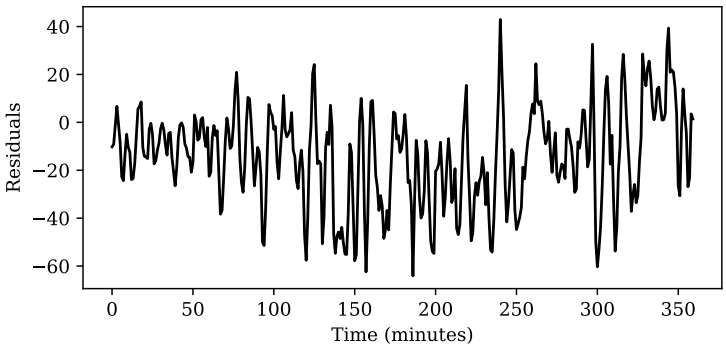

Figure 4.21: CSE, POI3, step-response test, LSTM model. Real value - predicted value for the time period forming the test-set portion of the dataset.

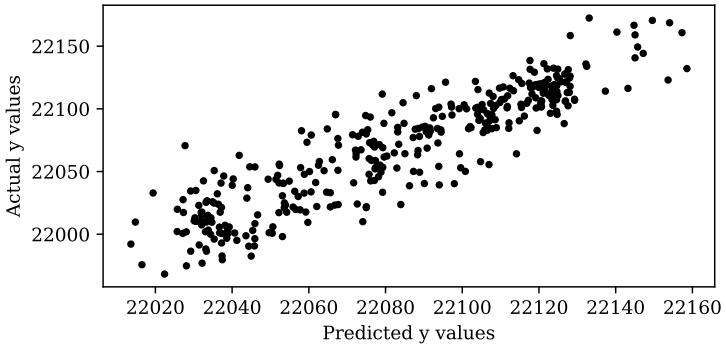

Figure 4.22: CSE, POI3, step-response test, LSTM model. Data-points of the test-set, prediction vs actual value scatter graph.

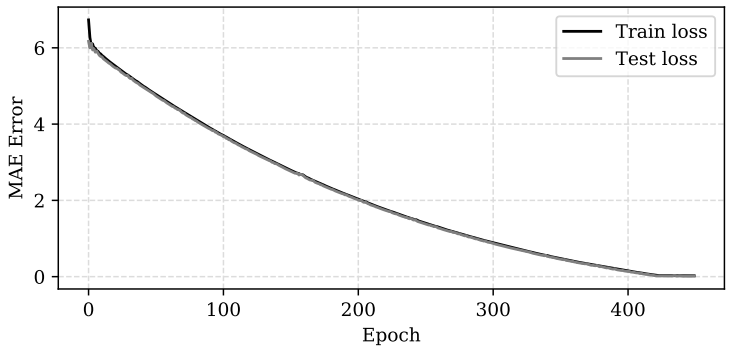

Figure 4.23: CSE, POI3, step-response test, LSTM model. Training score plotted for training time of 300 epochs, for both test-set and train-set data.

thermal mass of the cable-sealing end body increasing the time constant of the system significantly, causing a delay in the presence of the heating effects of the air temperature in the output. This observation affirms the assertion that the cable-sealing end provides a more challenging system to model. Otherwise, the LSTM modelling is successful, as in POI1, the model hypothesis output is similar to the true output. Again, inspection of the cost scores suggests the model is generalised across the entire dataset and that the residuals suggest there are possible improvements that could be made to the model, but that these improvements are out of scope for the experiment.

There are limitations to this shorter test dataset. Most prominently, that there is no wind input to the system, meaning that the system response cannot be observed. It is predicted that, since thermal imaging allows surface temperature to be measured, and cooling by convection as in the case of wind primarily effects the surface temperature, there will not be any unexpected effects detected at the output. An improvement of the test would be if there were a significant drop in the current, as in the first set of overhead line experiments. However, since the current controller of the experimental rig is only able to coarsely control the current, there are smaller scale drops in the input current present. The drops in output pixels due to these smaller drops in current are observable, especially in the more responsive case of POI1.

Note there are no attempts made to predict pixel outputs due to the fact that this will be investigated as part of the larger realistic load pattern.

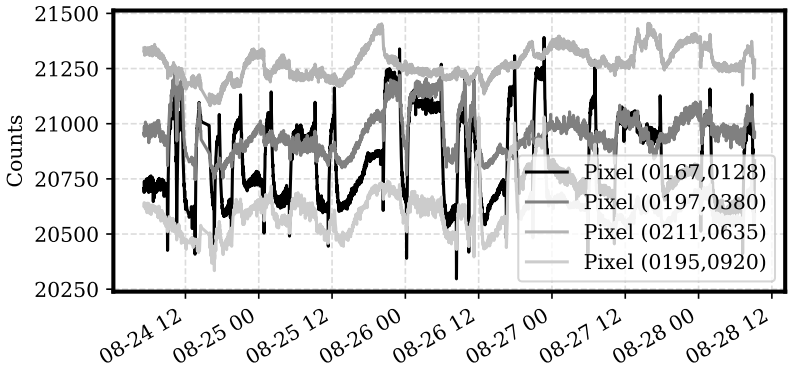

Figure 4.24: CSE, realistic data set, POI3, realistic test. Output plots from cameras 1, 2, 3 and 4.

4.5 Long Realistic Test

4.5.1 Preliminary Analysis

This section details the manual analysis performed in order to gain an understanding of the long realistic dataset. A more detailed analysis of both the camera outputs and the investigation into the thermal output of the cable-sealing end are located in Chapter 5.

Figure 4.25 shows a subset of pixel output time-series, from the four different camera perspectives around CSE1. There are two inputs that primarily drive the surface temperature of the imaged units: the current and the wind. The current directly influences the pixel output, it heats the conductor through the cable-sealing end, and this is visible on the unit surface. The wind influences the output pixel values in two distinct ways. Firstly it has a convective cooling effect on the surface of the objects in the field of view of the camera. Secondly, it has the same convective cooling effect on the cameras positioned nearby the path of the wind. The fact that these two influences occur at the same time may make it difficult to separate the individual impact of each, however it is not anticipated that this will cause additional difficulties when modelling the system.

On inspection of Figures 4.24 and 4.25, it is clear that the camera temperature has a significant impact on the output. The temperature of camera 3 is most stable for the duration of the test, and the corresponding output (pixel (211,635)) is the most stable of the four. Due to the fact that this relationship is frequently experienced through the datasets and that the camera temperature is recorded, it is anticipated that the machine learning method will succeed in modelling it. A

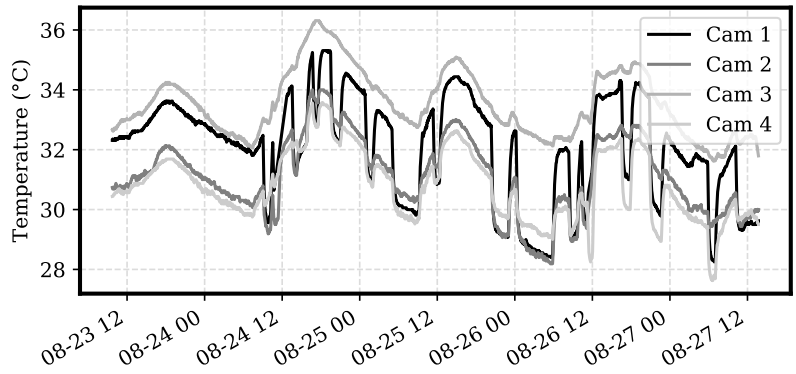

Figure 4.25: CSE, realistic data set, POI3, realistic test. Camera temperatures from cameras 1, 2, 3 and 4.

result of this analysis is to ensure that when models are created from data from a particular camera, only the temperature data corresponding to that camera is included as an input, while temperatures from other cameras are disregarded. This should assist in the training of the network, as an aspect of feature selection is removed.

4.5.2 LSTM Analysis

This chapter describes the generation of models based on the LSTM recurrent neural-network technique, using time series data from laboratory experiments. The hyper-parameter sweep used to select the neural-network configuration is detailed. The results obtained using the selected configuration are presented, with illustrative subset of examples provided.

4.5.2.1 Parameter Selection

An exploratory initial analysis of model performance was completed. Each point of interest has a unique output time-series to model. Combined with the stochastic nature of neural-network based machine learning methods, this means that the networks for each point of interest take varying numbers of epochs to train, and ultimately train to different final validation scores. Simply picking a large quantity of epochs is not viable. Besides being time consuming and inefficient, it carries the risk of resulting in over-trained models, which do not generalise well to data not available to the network in training.

Keras provides the functionality to provide a validation score at which to

stop the training of the neural network: *stop_score*. In order to determine the validation score at which to stop training in each case, a large batch of exploratory learning runs was completed, using a large number of epochs to train.

The learning curve of each run was evaluated to determine a point in training at which a model may produce satisfactory output, and is not over-trained or under-trained. Generally, this is achieved by selecting a point at which the validation metric (in this case, test-set mean-absolute-error) is at it's minimum for the learning run. Another consideration is whether the error scores for the test-set and training-set are similar (the output solution is well generalised). Due to the stochastic nature of the training process, there is no guarantee that a future training run would match the performance of the exploratory run. Equally, it may surpass the performance. Due to the large number of learning runs required to infer that a result is not an outlier, time spent per run must be considered. Therefore, in a trade-off between learning time and ultimate training performance, a constant is added to the *stop_score* determined by the described process, ensuring that if future training runs do not achieve the equivalent performance, they do not over-train. The number of epochs is capped at the point at which the exploratory run reached the goal *stop_score*, again with a small constant added, further mitigating against over-training.

For all learning runs for this experiment, a batch size of 1000 is used. Larger batch sizes such as this cause the learning process to take more epochs, with smaller changes in error per epoch. This allows the network and training process to be tuned for performance, as the smaller updates in network weights allow increased precision in selecting when to stop training. Large batches also allows the software to better utilise the 8 GB GPU memory, with less loading of data in and out of memory per epoch. Batches are shuffled between epochs, helping to ensure the network can interpret different sets of data-points per epoch, providing a more general solution. Large batches mean there are less weight updates to complete per epoch, resulting in the short execution time of a single epoch.

In the process of modelling the data generated by the over-head line experiment (Chapter 3), it was determined that the number of lags provided to the network has a significant effect on the model performance and training time. For this reason, the large-epoch exploratory modelling runs were completed for a set of candidate numbers of lags: 1, 10, 30 or 60. Qualitative inspection of a number of performance metrics including residual spread, visual model fit and minimum

Table 4.5: Hyper-parameter selection for cable sealing end realistic test

Hyper-parameter	Value
L1 Regularisation	0.01
L2 Regularisation	0.01
Dropout Layer Rate	20%
Shuffle	True
Batch size	1000
Train:test ratio	70:30
No. observations	15669

validation cost-score informed the final choice of number of lagged values for future modelling attempts. The selected value based on these metrics was 30. A further investigation into the performance of 120 lagged values, corresponding to two hours and therefore hypothetically encapsulating long-term effects, was intended. Considering the strong performance of the models with less lagged inputs, and the significant increase in training time, this was not completed.

Insights gained during the overhead line experiment were used to inform the selection of a number of the hyper-parameters. The hyper-parameters used to counter over-fitting (dropout layer)[54] was given particular attention, due to the increased risk of an over-fit model due to the auto-regressive nature of the input data. Models occasionally train to become effectively persistence models:

$$y_{(t+1)} = x_{(t)} \tag{4.1}$$

which have very good apparent performance scores, when over-fitting is not prevented. These measures each force weights in the network to small values during the training process, so selecting them well involves balancing final model performance, training time and consistency of model performance between test and training sets. The values presented in Table 4.5 were selected using qualitative evaluation of these metrics, and were used across all final learning runs.

L1 and L2 bias regularisation [43] were used in order to ensure bias was minimal in the network.

Table 4.6: Performance metrics used to evaluate models from the realistic test

Performance metric	Description
r^2	Coefficient of determination score as evaluated on the test-set
RMSE	Root mean-squared error as evaluated on the test-set
Pred. RMSE	Root mean-squared error as evaluated after 30 step rolling predictions

4.5.2.2 Model Performance

Model performance is presented in terms of metrics described in Table 4.6. These metrics are selected to provide an indication of how well the model performs primarily on the test-set portion of the dataset and on the recursive predictions. Generally, performance on the training-set portion can be assumed to be equivalent or better, though this is not included as it is the metric on which the network is trained. Model generation for each point of interest, for every camera, was repeated thirty times, in order to determine typical performance.

The coefficient of determination (r^2) is used as it gives a unit-less indicator of the model skill. This is accompanied by the root mean-squared error, in the units of the experiment output: digital counts, which is proportional to the thermal flux incident on the sensor. The number of epochs used to train the presented model is provided as an indicator of the training time for that particular model. Since all models have equal number of training points and batch sizes, the epochs approximately equal amounts of time, with variations due to other operations being performed on the PC, so the quantities of epochs are approximately proportional to the training time.

Thirty time-step recursive predictions are made for each model. The root mean-squared error of the predicted output compared to the observed output is presented, providing an insight into the limitations of the model and how this increases with time.

The mean value of the residuals, e is provided, giving an indication of any skew in the residuals of the model. The standard deviation of the residuals is the

final metric, quantifying the spread of the error between the model and reality.

Full tables of results are presented in Tables 4.7 and 4.8.

The number of epochs taken to train the network varies significantly. The minimum value is 121 and the maximum is 1676, representing over ten times the training time. There is a trend towards a longer training time for points of interest with more thermal fluctuation, with the shortest training time generally corresponding to POI6 on the base of the CSE, a point at which the camera output does not vary significantly. Table 4.3 shows POI6 fluctuating from 338 digital counts for camera 8 to 975 for camera 1. In comparison, POI1 fluctuates up to 4636 digital counts.

The model performance as measured by the r^2 score is consistently high. The lowest r^2 score is 0.8524, corresponding to POI6 of camera5. There are two other POIs that report scores below 0.9. The score tends to decrease as the point of interest moves down the length of the CSE. For CSE2, the best scores are consistently POI1. For CSE1, the scores follow the same trend, with the most thermally volatile points achieving the best performance, however POI1, 2 and 3 all score similarly for this CSE. The mean performance for CSE1 is 0.9884 and for CSE2 is 0.9618. Examples of well-performing and over-fit model outputs can be found in Figure 4.30

Root mean-squared error is highest for POI1 and lowest for POI6 for CSE1 generally. This trend is visible for CSE2, with a small number of outliers, namely camera 7, POI5, which is higher than expected and camera 8 POI1, which is lower. Camera 5 displays high RMSE for all models compared to the other cameras around CSE2, but in the context of all 8 cameras, the errors are not significantly large. However, the error for camera 5 POI6 is the largest of the entire experiment.

The root mean-squared error after a 30 time step recursive prediction is consistently higher than that for the test-set portion of the dataset. This metric varies a large amount, the largest value being 132 (to 0 d.p.) for camera 2 POI1, with the smallest being 8 (to 0 d.p.) for camera 3 POI4. There are significant differences in the performance of different models for each POI. For the worst case (camera 2 POI1) the worst performing prediction has a large error of 471 (to 0 d.p.), while the minimum is 24 (to 0 d.p.). In the best case (camera 3 POI4) the range of error is smaller: from 12 to 6. These results are presented in Figure 4.26. Examples of successful and unsuccessful recursive predictions are

Table 4.7: Results summary from Cameras 1, 2, 3 and 4 (CSE1) for the realistic test. Values are the mean of the metrics for repeated experiments for each point of interest. Predictions refers to recursive predictions, made for 30 1 minute timesteps.

Camera	POI	Epochs	r^2	RMSE	Pred. RMSE	e	std(e)
1	1	738	0.9969	14.8427	36.3187	-4.491	13.539
1	2	825	0.9986	10.7793	25.9383	-0.786	10.47
1	3	838	0.9986	10.6576	29.3457	-1.136	10.162
1	4	748	0.9898	11.5597	27.6798	-5.791	10.259
1	5	942	0.9887	11.5932	26.5290	-6.099	9.45
1	6	569	0.9880	11.7266	24.6368	-5.812	10.003
2	1	1224	0.9996	17.6487	132.5584	0.320	17.466
2	2	800	0.9983	15.4099	81.5186	0.442	12.739
2	3	845	0.9982	14.1582	65.0181	-0.486	11.512
2	4	521	0.9905	13.5777	53.4443	-2.994	11.247
2	5	475	0.9762	13.4080	47.1655	-5.400	11.139
2	6	517	0.9644	13.0213	42.7473	-2.273	10.642
3	1	1676	0.9694	13.2660	52.5335	3.467	14.027
3	2	970	0.9984	8.8238	31.9435	0.946	8.671
3	3	1002	0.9991	7.4288	25.8071	0.498	7.349
3	4	451	0.9930	5.6432	7.9028	-0.070	5.623
3	5	427	0.9802	5.2290	8.6100	-0.734	5.150
3	6	322	0.9471	5.6891	10.3444	-0.917	5.598
4	1	759	0.9995	15.5548	95.1908	-1.054	15.316
4	2	604	0.9984	10.2925	26.8027	-0.887	10.222
4	3	842	0.9991	7.6109	35.9655	0.102	7.53
4	4	597	0.9880	6.9937	14.2902	-0.934	6.841
4	5	309	0.9831	7.8320	15.3449	-1.758	7.565
4	6	121	0.9791	8.5639	18.1597	-1.457	8.380

Table 4.8: Results summary from Cameras 5, 6, 7 and 8 (CSE2) for the realistic test. Values are the mean of the metrics for repeated experiments for each point of interest. Predictions refers to recursive predictions.

Camera	POI	Epochs	r^2	RMSE	Pred. RMSE	e	std(e)
5	1	837	0.9964	16.1685	28.9700	0.861	16.067
5	2	529	0.9260	14.6723	25.0377	-0.011	14.644
5	3	550	0.9625	12.2019	24.6910	-0.770	12.082
5	4	658	0.9155	18.4741	22.7646	-1.426	18.244
5	5	689	0.9119	19.9217	23.1933	2.141	19.528
5	6	825	0.8524	21.6324	23.7487	2.150	21.195
6	1	928	0.9987	8.3697	26.1546	-0.158	8.288
6	2	493	0.9922	6.3609	21.4047	-0.231	6.336
6	3	710	0.9862	8.0636	19.0216	-1.008	7.784
6	4	204	0.9687	8.8590	19.3145	-3.506	7.624
6	5	131	0.9728	7.0830	11.7606	-1.531	6.894
6	6	392	0.9683	6.5039	14.4698	-1.592	6.189
7	1	1300	0.9984	6.1873	14.5579	0.396	6.072
7	2	687	0.9943	5.9928	10.5519	-0.527	5.959
7	3	963	0.9886	6.8740	9.3681	-0.911	6.782
7	4	710	0.9859	7.3426	9.0375	-2.925	6.558
7	5	1392	0.8965	14.1571	13.0904	-3.182	13.652
7	6	521	0.9726	5.6756	11.525	-0.037	5.649
8	1	1148	0.9970	8.1489	16.1593	-0.688	8.034
8	2	1048	0.8753	15.5751	9.7097	-0.584	15.530
8	3	491	0.9724	9.2762	9.9648	-0.898	9.165
8	4	660	0.9854	9.1916	13.7150	-1.730	8.869
8	5	572	0.9835	8.3510	11.5178	-0.923	8.249
8	6	646	0.9818	7.7689	12.9193	-0.454	7.694

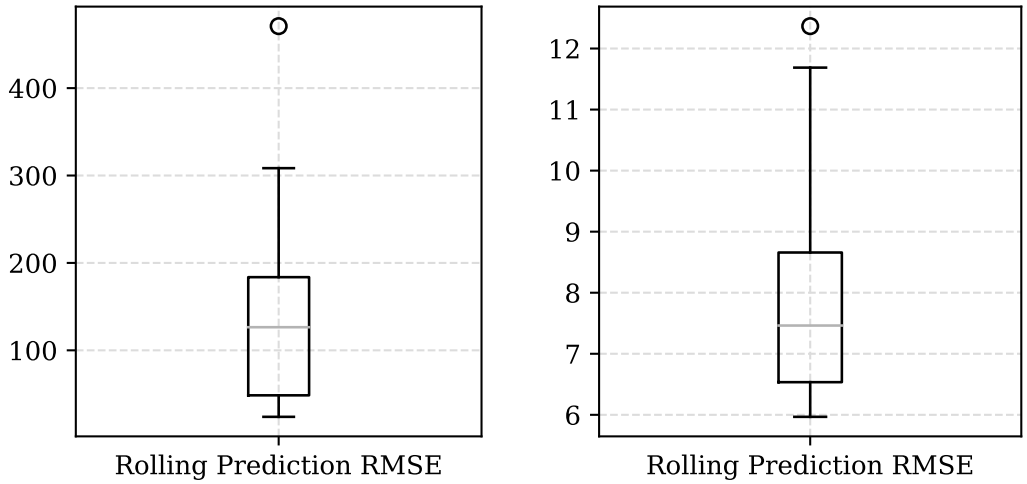

Figure 4.26: Left: RMSE for recursive predictions made for camera 2 POI 1 (worst case) for the realistic test; right: camera 3 POI 4 (best case)

depicted in Figure 4.27

The modulus of the mean value of the residuals of the models is consistently less than 7 for the entire experiment. The mean is generally closer to zero for CSE2 than for CSE1, with the largest absolute value for CSE2 being close to 3.5. Five of 24 of the the mean residual values for CSE1 exceed this, with four of them corresponding to models from camera 1. The majority of deviation from zero happens in the negative direction. There are only three values greater than 1 in the experiment, while there are eighteen values less than -1. Examples of histograms of model normalised residuals are presented in Figure 4.28, plotted with a normal distribution with zero-mean.

The distribution of the residuals is wider for the more thermally volatile points of interest, in most cases. Particular exceptions are camera 5, where the standard deviation of the residuals is both generally larger than for other camera perspectives, and also increases for the less thermally volatile points of interest. Another exception is camera 7, which exhibits large standard deviation for POI5. The standard deviation is generally low: frequently less than 10 and only greater than 20 in one case.

4.5.3 Analysis and Discussion

This section presents the findings of the work based on the chosen performance metrics as described in Section 4.5.2.2. The performance of models from data

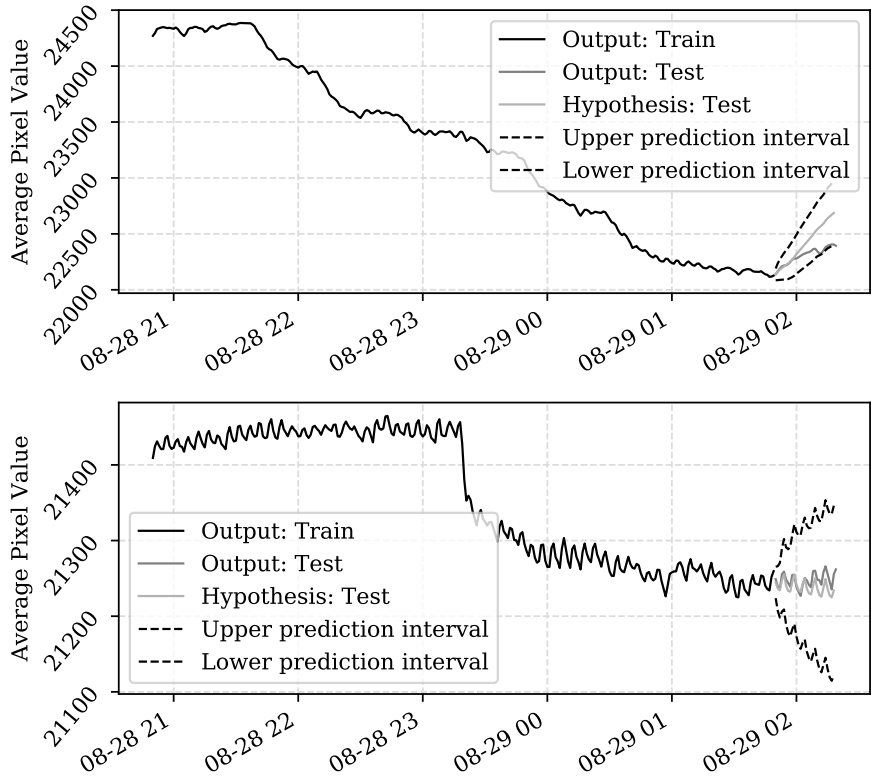

Figure 4.27: Examples of recursive predictions. Top plot is an example of a poor result (camera 2, POI 1, run 3, rolling RMSE 148), bottom plot is an example of a good result (camera 3, POI4, run 9, rolling RMSE 12). Note the prediction intervals are calculated as per a naive forecast and are therefore only for visualisation.

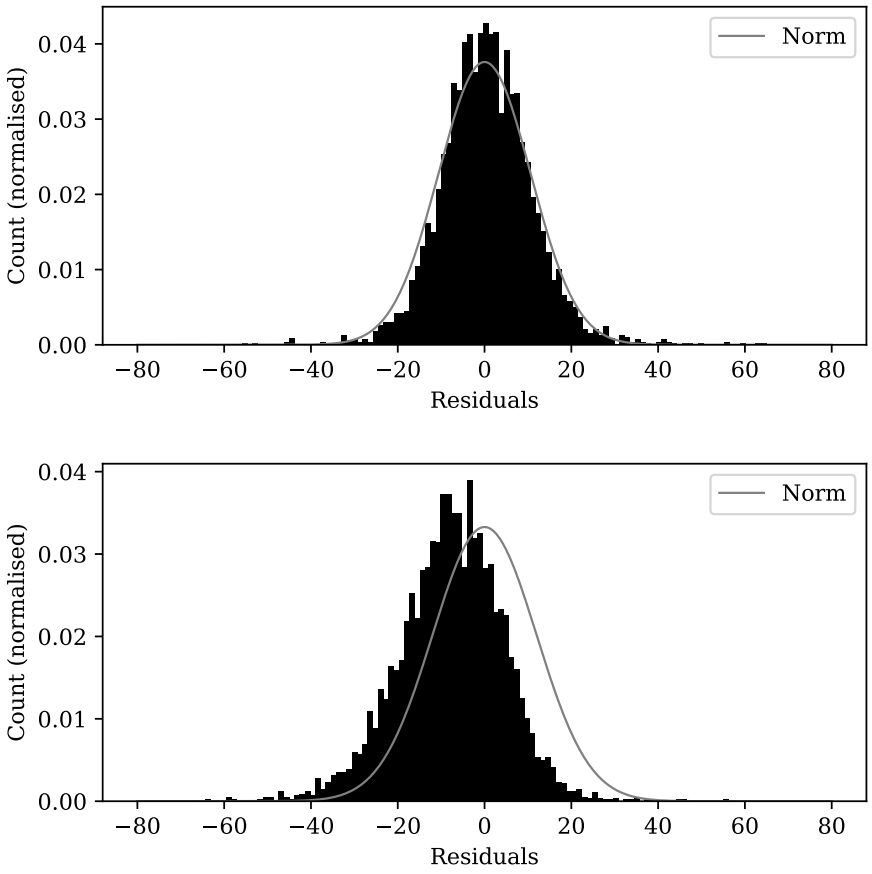

Figure 4.28: Normalised residuals with zero-mean normal distribution. Top plot is an example of a good result (camera 1, POI1), bottom plot is a poor result (camera 2, POI4)

from various perspectives is analysed and placed in the context of the wider aim of the project. The perceived limitations of the method and results are presented and discussed.

4.5.3.1 Performance

Regarding the epochs taken to train each network, it can be determined that a greater level of thermal fluctuation results in a longer training time. Since more thermally variable points of interest in the scene are effectively under greater influence from a larger number of variables, it is intuitive that it would take more time to properly converge the weights on values that will capture the large changes. The gradient of the cost-score over the course of the training indicates that there is a point after which in training, the improvements in performance are significantly diminished. This can be observed in Figure 4.29. This characteristic could be exploited if the method were applied in an industrial setting, whereby analysis into the difference in performance could be evaluated: for each of the selected metrics, firstly at the beginning of the diminished returns period of training, and at another reasonable stopping point. Since the process of selecting the number of training epochs and the stop score is qualitative and considers each model exclusively individually, caution should be used when attempting to gain insight about the method by comparing the train time for various models. It is a limitation of the method that despite the otherwise identical neural network configuration and data-types, so much tuning for training time is required. During the process of determining training time, severely over-fit models with large fluctuations above and below the true output were encountered, highlighting the need to ensure the training time is well selected. This effect could be mitigated by selecting a more sophisticated metric to configure the neural network to train on. Examples of well-fit and over-fit model outputs are provided in Figure 4.30. The over-fit model demonstrates one of the potential effects of over-training a model. In this example, the network continued to return a descending cost-score for the training-set portion of the dataset during the training process. In order to gain insight into the model skill from the r^2 score (coefficient of determination), it is considered alongside the RMSE. The r^2 score was selected as a metric of performance as it is independent of the output unit, and it considers the variance of the output. As a result, given two models with equal absolute residuals, based on two datasets with different variance, the model based on the more varied data

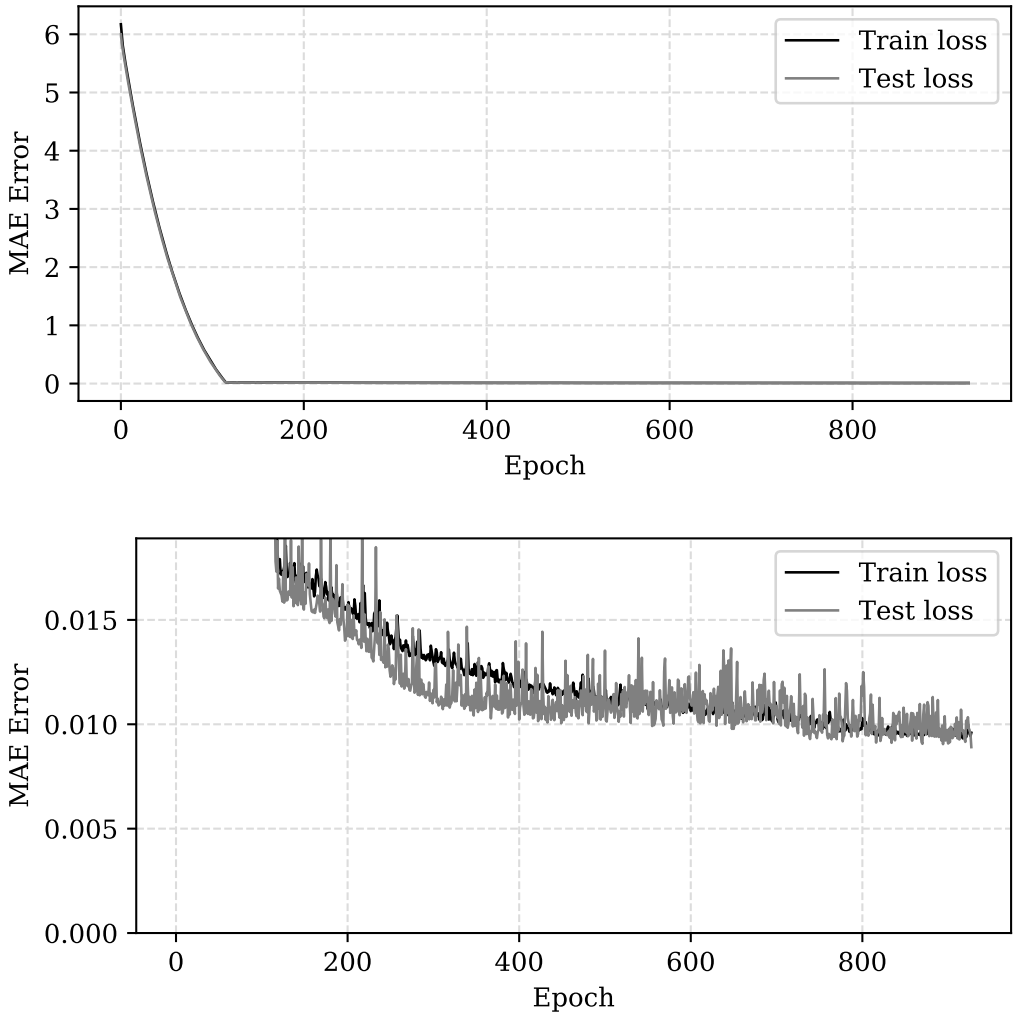

Figure 4.29: Randomly selected learning curves for the realistic experiment. Camera 2, POI3, run 16.

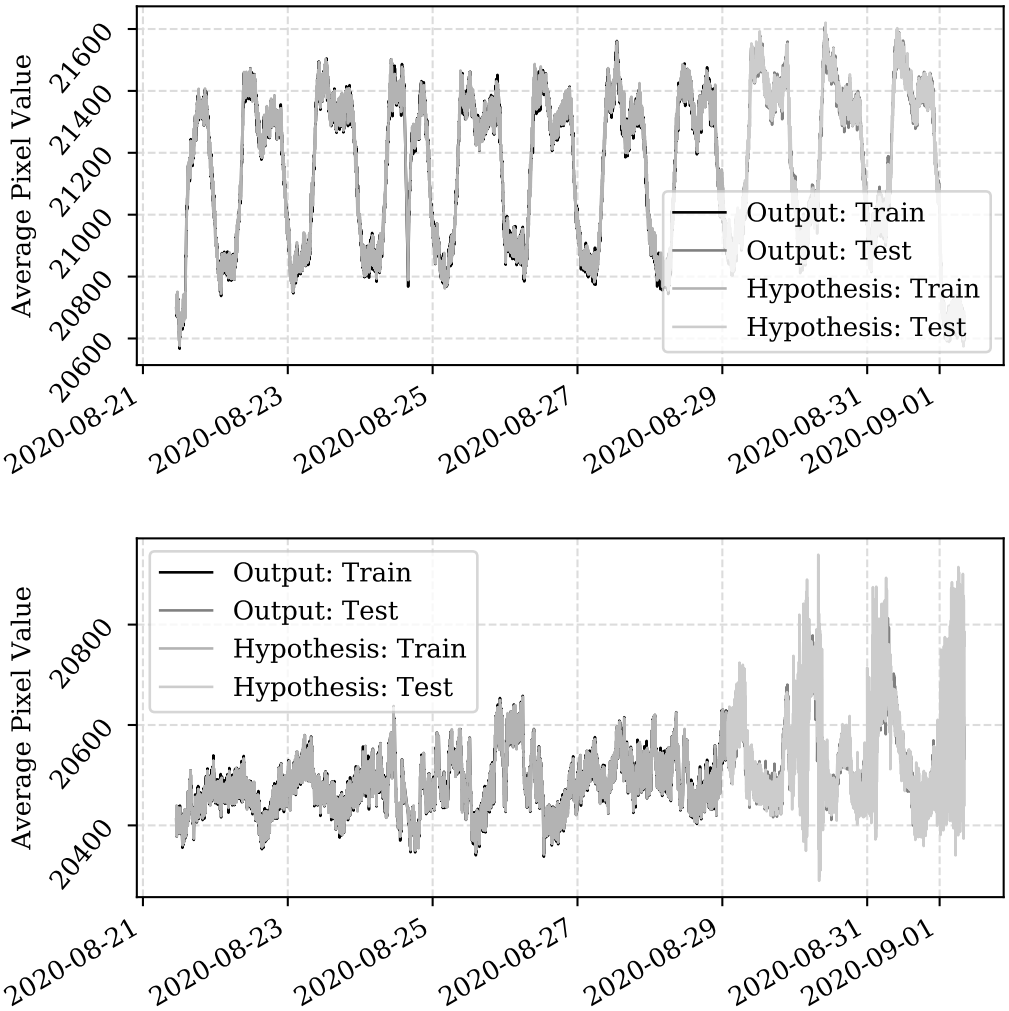

Figure 4.30: Examples of a well fit model for the realistic experiment (camera 5, POI 1) and an over-fit model (camera 5, POI 6).

will achieve the better score. The RMSE conversely is presented in the units of the output variable and is independent of the variance.

The primary finding on inspection of these two metrics, is that the performance is good and accurate models are produced. Even given the limitations of the r^2 score with regards to time series, it does confirm that the fit of the model is good when compared to the observed output. The RMSE confirms this. That the RMSE values are all less than 22 digital counts, considering that the range of some of the output time series is greater than 4000, reinforces that the results are good. Where the range is lower and the error proportionally large, this is reflected in the r^2 score. As with any system, there is noise present in the experiment. Other than any noise that could be attributed to the individual sensors, this noise is assumed to be consistent across the whole experiment. Due to the fact that different points of interest in the scene respond in different magnitudes to the primary inputs of the system, the noise comprises a more significant part of the output signal corresponding to the points of interest that heat up less. This results in the lower performance scores for the cooler POIs, as the noise is not modelled and in these cases is more prominent.

Given the quantity of scores > 0.9, the distribution of the scores of the three points of interest with r^2 scores < 0.9 are presented in Figure 4.31. These three points are camera 5 POI6, camera 7 POI 5 and camera 8 POI 2. For comparison, the distribution of a high-scoring model is included. Camera 5 POI 1 is selected for this, a result more representative of the rest of the experiment and from a camera with an anomalously low result. From inspection it can be seen that these models generally perform well but have outlier results which affect the mean. Placing this in the context of the wider accurate models, the experiment is a success. However, to determine the suitability of the models for using for fault detection, the residuals must be inspected.

There are also valid issues with the naive use of r^2 score in the context of time series models. For this reason, the r^2 score should be considered in the context of the RMSE for the same model.

Frequently the results show a skew to the mean of the residuals, a concern as many model based fault detection methods are predicated on a zero-mean. There is a trend towards a negative mean, showing that models often predict values higher than the true observed values. At the end of the test period, the current is set to zero and the cable-sealing ends cool. The current is not set to zero for a

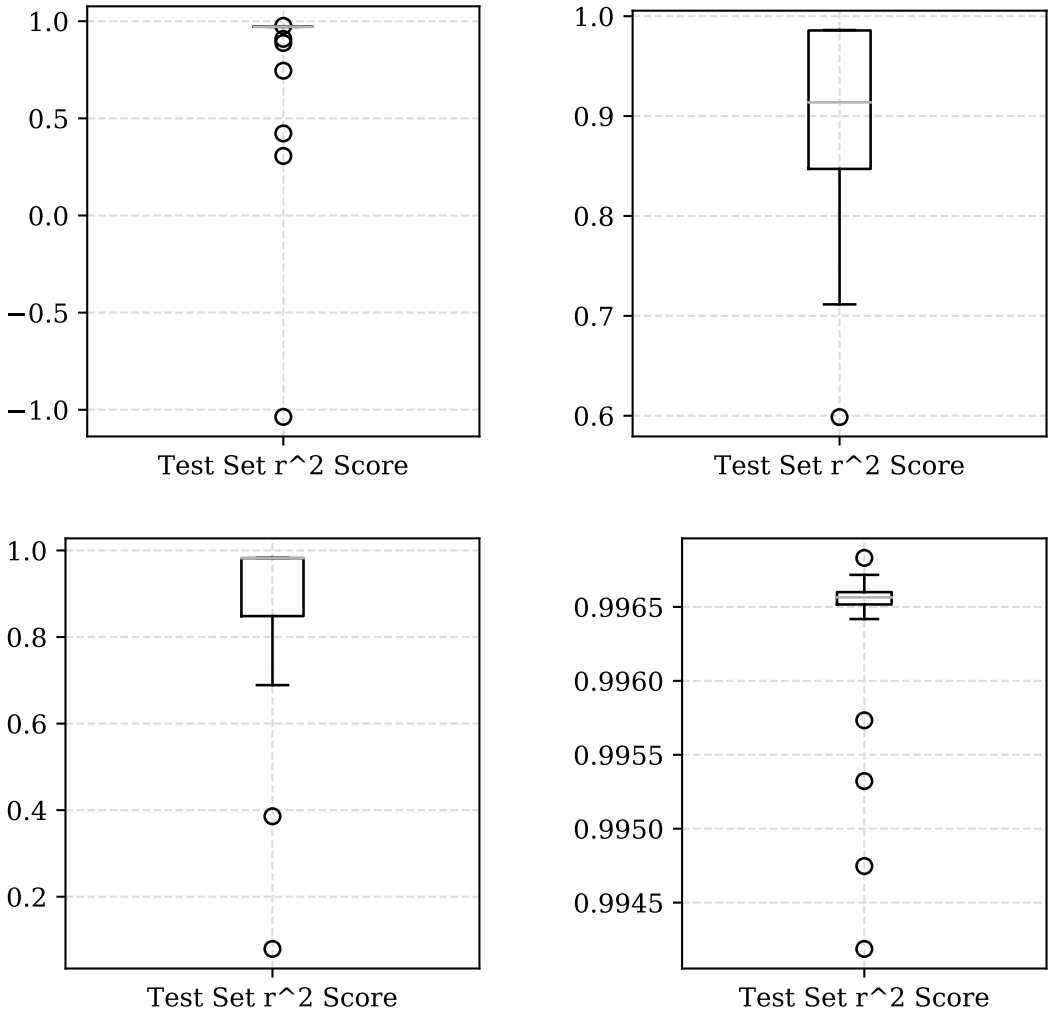

Figure 4.31: Distribution of r^2 scores for realistic test, clockwise from top left: camera 5, POI 6; camera 7, POI 5; camera 8, POI 2; camera 5, POI 1.

significant period at any point in the training data set, potentially contributing to the difficulty modelling the test portion of the data set properly. However, this should not be over-stated, as many model runs achieved good performance for the test-set, with residual mean very close to zero. This does reinforce the notion that the model skill is limited by the dataset.

Camera 1 POI 1, camera 1 POI 4, camera 1 POI 6 and camera 2 POI 5 were each manually investigated due to the large skew to the mean of their residuals. On inspection of the results, it was found that the skew was not consistent through each of the thirty modelling runs for each point of interest: particularly poorly performing models influence the mean result of the thirty runs. It is notable that camera 1 and camera 2 are significantly effected by the wind during the experiment, each of the cameras experience thermal transients which influence the camera output. These transients are fast (relative to temperature changes in-scene) and the output changes are of similar magnitude of those caused by the heating of the cable-sealing end. That the neural network converges on zero-mean residuals less often for more volatile output is a limitation of the approach of using a single configuration for the whole experiment and infers that further data may be required. Due to the fact that the distribution and skew of the residuals is key for fault detection, it is reasonable to presume that if the work were continued to implement fault detection, a more thorough model selection procedure could be utilised.

The performance of camera 3 POI 4 and camera 7 POI 6 are notable, the mean of the residuals for each are close to zero. This performance is consistent across the thirty models for each of the points of interest, suggesting that very good performance can be achieved on data such as these. It is an encouraging result that a point of interest from the perspective of camera 3 models so well, as the face of the cable-sealing end imaged by camera 3 is under direct influence of the wind. It also suggests that further tuning of the modelling process is required for performant models, and that determining what causes the neural network to perform well on some points of interest and less well on others may be difficult. The mean residuals results overall demonstrate the volatility of neural network modelling while also suggesting that accurate models can be acquired if care is taken to manage this volatility. They also reveal a limitation of the dataset: despite being a relatively long dataset, the environmental variation present is not as extensive as desired.

Considering the spread of the residuals, presented as the standard deviation, the assertion that good models can be achieved is reinforced. Where there is a narrow spread, it can be inferred that model-based fault detection would determine when a systems performance is diverging from what is expected more easily. Images from camera 5 exhibit a relatively wide spread throughout the points of interest, suggesting a systemic sensor-specific element to the result. The FFC normalisation process creates large fluctuations in the output for camera 5 and could be the source of the increased spread of the residuals. This highlights that care must be taken in viewing each result in the context of the specific sensors used, especially when considering complex, sensitive devices such as thermal cameras.

Recursive predictions are made. The methodology for these rolling forecasts is described in Chapter 3, Section 3.4.4. Briefly, the method forecasts at time t+1, then feeds that forecast back into the model for the following forecast at time t+2. This provides an indication of the utility of the model as a forecasting tool. Inspection of the model output for the test-set provides instead only a series of single step forecasts, which, while providing a good indicator of the model skill, do not portray the forecasting utility. This is because the input variables of the test-set of the data set are formed using lagged values of observed output.

When viewed in the context of the other performance metrics presented for the models, it is clear that forecasting using such models should be approached with caution. Where models appear to be very good, such as in the case of camera 2 POI 1 (see Figure 4.27), the recursive prediction for thirty time steps diverges significantly from the real output. A significant limitation of the LSTM method used here is that, due to the multiple weights within each LSTM unit and the different effects they have, it is difficult to quantify where the limitations of the model lie. There are cases in which the recursive predictions track the observed output closely, though it is difficult to determine which models produce this type of output. Two models with similar performance metrics may produce significantly different recursive predictions provided the same input. The forecast intervals are presented as if the system was making a naive forecast, not 'correct' but an indication nonetheless of the confidence as a proportion of the residuals.

4.6 Summary

This chapter presents findings based on data acquired using a unique experimental setup, consisting of two electrically loaded cable sealing ends, thermally influenced by a large industrial fan and observed by a Vaisala WXT520 weather station, and eight non-calibrated FLIR Boson thermal cameras. It briefly explores factors that contribute to the challenge of modelling output from such an experiment and presents samples of data with descriptions of the behaviour exhibited in the samples.

An 11-day experiment is presented, during which the cable sealing ends experience an electrical load pattern that is representative of what may be experienced by a similar asset in commission in a substation. The wind fan is directed at one of the cable sealing ends from a number of angles and is set to operate at a number of different intensities over the course of the experiment, constituting a substantial variation of wind conditions, designed to exploit the ability of the equipment to monitor the cable sealing ends from multiple angles. The laboratory temperature is uncontrolled for the period of the experiment and varies significantly on day/night cycles and also cools gradually over the course of the 11 days.

Fourty-eight points of interest corresponding to condition monitoring spot measurement points are selected. Multivariate LSTM recurrent neural network models are created based on the pixel outputs corresponding to these points of interest, following a procedure to determine a good general LSTM configuration. The modelling process is repeated thirty times for each of the fourty-eight points of interest, resulting in a good indication of model performance for a variety of data.

The model performance is good throughout the points of interest, with low root mean-squared error and high r^2 scores. Inspection of the residuals of the models reveals that system dynamics are well captured, as residual mean is close to zero and generally narrowly distributed. In cases where this is not the case, it is determined that tweaks to the training process and/or the model configuration could improve performance further. It is assumed that model accuracy is good enough for the implementation of model-based fault detection but acknowledged that this needs to be quantified.

The chapter represents a step-change in complexity relative to Chapter 3, further advocating for continuous thermal and environmental monitoring in substations, by successfully demonstrating that the thermal output of high thermal-mass

assets with complex exterior geometry can be modelled well, in three dimensions. This is additionally significant due to the fact that such equipment is targeted for monitoring in condition monitoring surveys because of the high internal electric field and fluctuating internal operating temperature potentially contributing to aging.

These findings suggest that a condition monitoring policy utilising long-term or constant monitoring, using either static cameras or mobile platforms, could offer value to the project sponsor in reducing employee risk and reducing cost due to equipment degradation.

Chapter 5

Case Study: Thermal Cable Sealing End Condition Monitoring

5.1 Introduction

This chapter presents a case study, exploring long-term monitoring of 66 kV CSE assets in a laboratory environment, in order to gather information on the utility of non-calibrated thermal cameras in this context. Using data captured during the course of the modelling work presented in Chapter 4, the chapter aims to determine how the CSEs respond to the heating effects of the electrical loading. It explores the heating profile of the assets around the circumference and along the length, and how each of these develop in time with continued electrical loading.

The nature of non-calibrated thermal cameras dictates that each sensor typically has a differing response to the thermal flux incident on it's sensor array, which is non-linear with respect to the camera sensor temperature. This chapter explores some of the challenges faced in extracting equivalent data from a number of these different sensors, in order to determine what differences in thermal intensity are due to asymmetric heating around the region of interest, and what differences are due to inherent sensor differences.

Finally, the chapter evaluates the thermal output of the two CSEs in relation to each other. It discusses the operational health of each of the assets, making suggestions on further investigation to determine what the cause of any differences may be, and whether or not they constitute a fault precursor.

Section 5.2 presents a overview of some literature relevant to the case study before the datasets are described in Section 5.3. Sections 5.4, 5.5, 5.6 and 5.7 present the study of the thermal response of the CSEs. Section 5.4 concerns the thermal profile of the CSEs over time, Section 5.5 concerns the thermal variation along the length of the CSEs, Section 5.6 the thermal variation around the circumference of the CSEs and 5.7 concerns the difference between the thermal profiles of the CSEs.

5.2 Literature

Work investigating the health of CSE assets exists in the literature. Zazachariades et al. [67] investigate the thermal output of a 132 kV CSE, comparing it to finite element analysis output and establishing an understanding of the temperature profile down the length of the external sheds. Based on their investigation, and other investigations into the condition assessment of high-voltage cables and cable terminations [36] [35], there is an additional academic contribution in establishing an understanding of the thermal output of the 66 kV CSEs housed on the experimental rig.

5.3 Datasets

Data from the various experiments conducted over the course of this work were utilised in order to form the investigation presented in this chapter. In order to establish a general view of the behaviour of the CSE, data from a wide range of experiments are used. An investigation into aging the stress cone inside the CSE was completed in January 2020. The two selected datasets at constant 300 A loading offer an opportunity to compare the response of the CSEs in summer and winter, and to understand if there is any change in the response as a result of the aging process. Unfortunately, due to safety concerns, the rig was not able to be energised for a period outside of normal laboratory operating hours until after the aging process was completed. Nevertheless, data gathered in August 2019 is included as a reference, as both CSEs were known to be healthy at this point. Finally, the long-term, realistically loaded experiment is included, providing insight into ranges of output in normal operating conditions. These are presented in Table 5.1.

Table 5.1: Datasets used to investigate CSE heating

Date	Electric Loading	Duration
30/08/2019	Varying	7 hrs
06/02/2020	300 A	19 hrs
11/08/2020	300 A	40 hrs
21/08/2020	Realistic	11 days

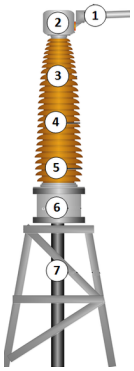

Figure 5.1: Guidance diagram on spot temperature measurement locations

Points of interest taken at various points down the length of the CSEs are utilised throughout this chapter. These points are introduced in Chapter 4, Section 4.3. The POI selection guide (Figure 5.1) and annotated thermal image (Figure 5.2) are included here for reference too.

5.4 Thermal profile of CSEs over time

5.4.1 Presentation of Data

In order to understand the temperature profile of the CSE in operation, selected images are taken from various data sets, normalised to exaggerate features, and inspected.

Three data sets are selected to inform this understanding. Firstly, a short experiment with a varying load profile, in August 2019. This experiment was conducted with two healthy CSE assets. The data points selected from this experiment were the beginning and end of the test, corresponding to the points with the least and most heating. These are presented in Figure 5.3.

The second experiment is an overnight step-test performed in February 2020. Again, the data points selected here were the beginning and end of the test. These

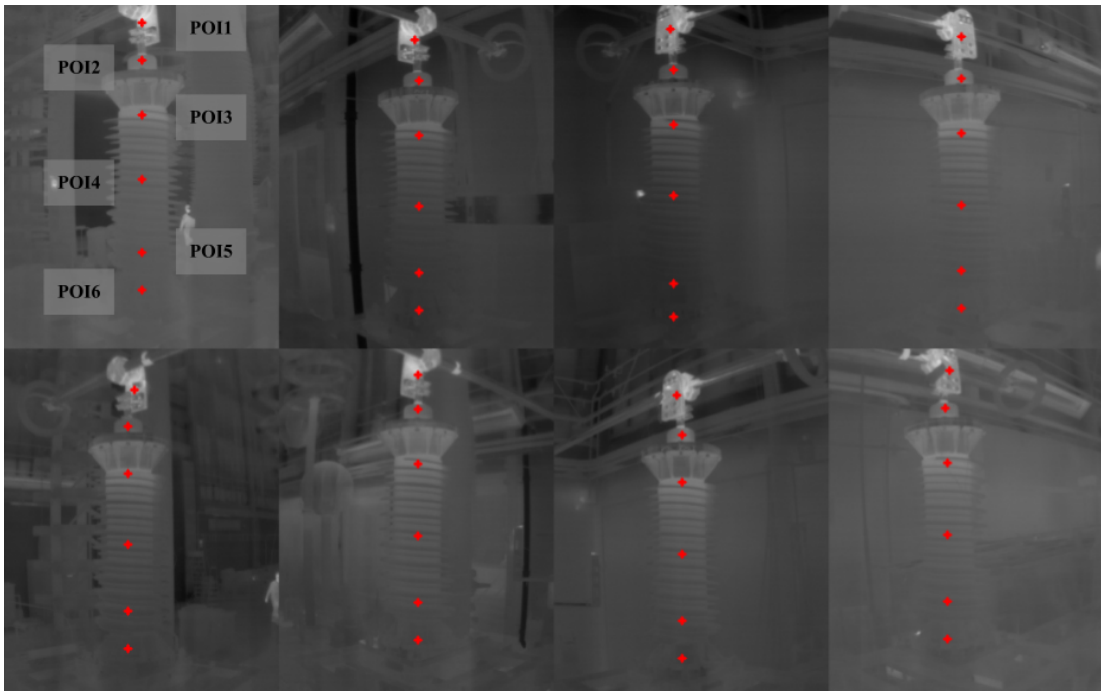

Figure 5.2: Stitched thermal images, with POI markers and POI labels on Camera 1 perspective

are presented in Figure 5.4. This test was performed after modifications were made to one of the assets and, due to being completed in winter, was performed in a lower air temperature in the laboratory.

The final test is an extended step-test performed in August 2020, after the fault has been present in CSE 1 for an extended period. In order to select the sample points to present here, a plot of a point on the CSE is selected and plotted against time. As per the previous two experiments, the start and end points are presented. In addition, a point at which the CSE is heated to approximately 50% of it's maximum (at POI3) is presented, alongside a point when the system first reaches it's steady state for the experiment current (300 A). These are presented in Figure 5.5 and 5.6.

5.4.2 Qualitative Findings

Heating is visible down the entire length of the CSE. At the end of each of the experiments there is clear visible heating at the top of the asset, and after the extended step-test, the heating is visible in the lower part of the CSE body, between the sheds. The heating profile is a pronounced gradient, hottest at the

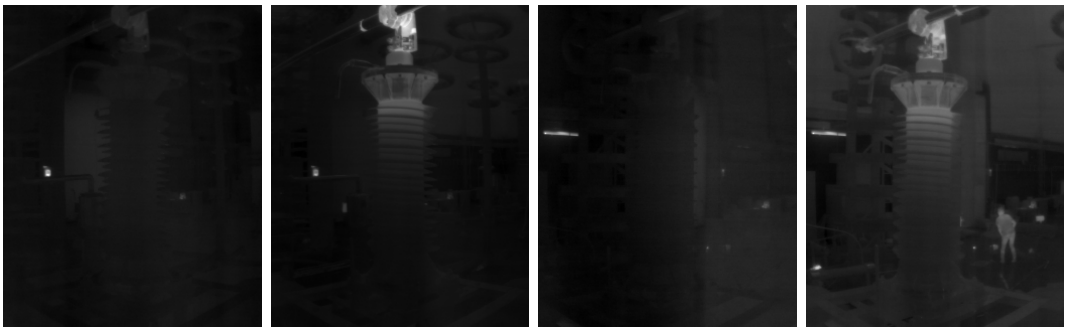

Figure 5.3: Thermal images from a six hour cable sealing end experiment. Varying input current. Left to right: camera 1, test start; camera 1, test end; camera 5, test start; camera 5 test end.

top, with decreasing temperature down the body.

The heating effects are horizontally uniform, with no visible drop in intensity near the edges of the asset. A drop in intensity may be expected here, due to the optical effects of the increasing view angle between the camera and the CSE, however this is countered by the increased surface area corresponding to each pixel after the cylinder is projected onto pixel plane.

In order to determine how the profile develops over time, Figure 5.5 and 5.6 can be inspected. These figures present data from the start of the test, after one hour of loading, after four hours of loading and after 40 hours of loading. Figure 5.7, a time-series of the digital count output at POI3 on CSE1, is referenced when selecting these points.

The output for this extended step experiment is unique within those presented in that there is some heating present at the top of the CSE at the beginning. A sanity check was performed to ensure this was not a sensor anomaly, confirming that this mild heating is present from every camera perspective in the test. It is difficult to determine the cause of this, as there is no data from before the first thermal images are taken. It is assumed that some testing had occurred prior to this test, or that there had been a period of low current.

The output of this experiment demonstrates the manner in which the temperature profile develops with time. The upper part of the asset heats first. This temperature increase spreads down the length of the CSE, while the upper part continues to increase in temperature at a lower rate with time, before settling. This suggests that changes in the thermal dynamics of the CSEs would be likely to be visible near the top first.

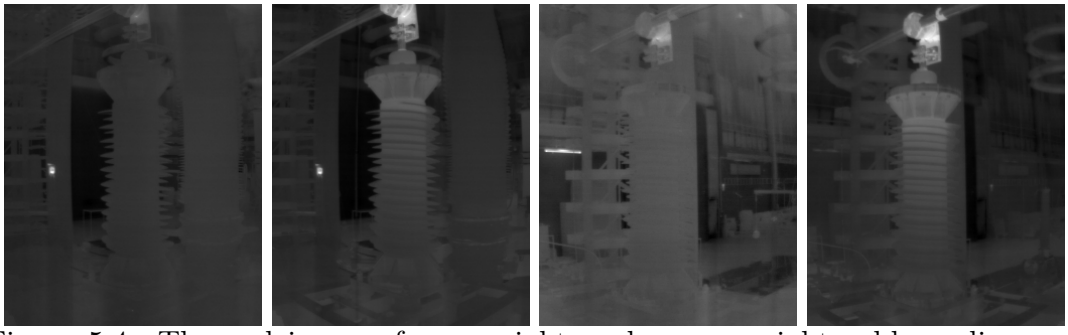

Figure 5.4: Thermal images from a eighteen hour overnight cable sealing end experiment. 300 A input current. Left to right: camera 1, test start; camera 1, test end; camera 5, test start; camera 5 test end.

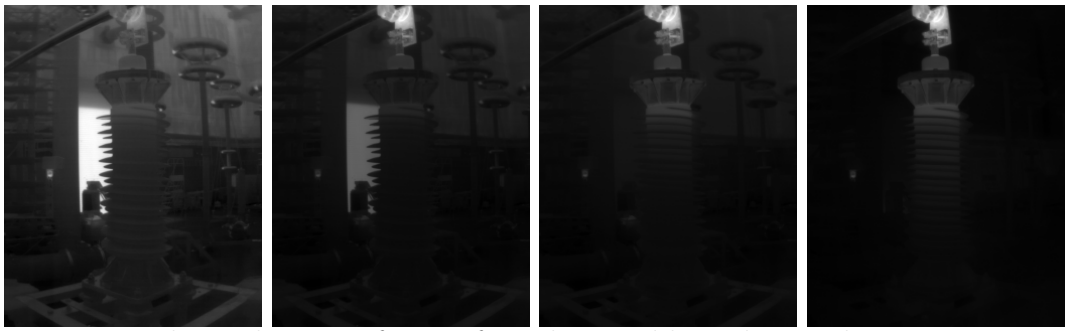

Figure 5.5: Thermal images from a forty hour cable sealing end experiment. 300 A input current. Left to right: Ccmera 1, test start; camera 1, mid-heating end; camera 1, start of steady-state; camera 1 test end.

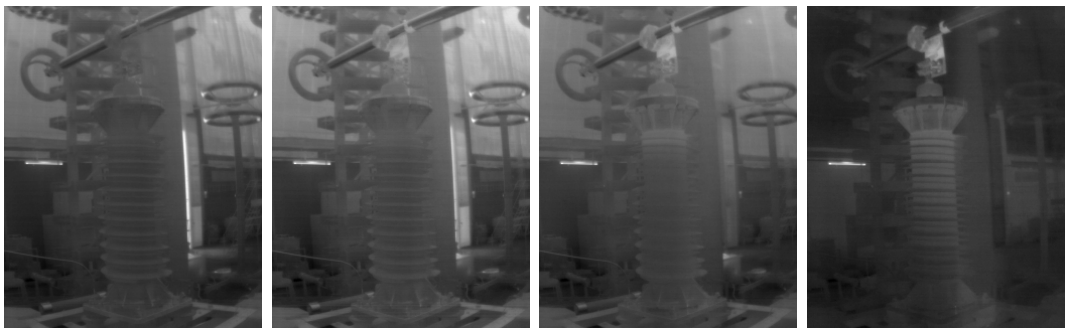

Figure 5.6: Thermal images from a forty hour cable sealing end experiment. 300 A input current. Left to right: camera 5, test start; camera 5, mid-heating end; camera 5, start of steady-state; camera 5 test end.

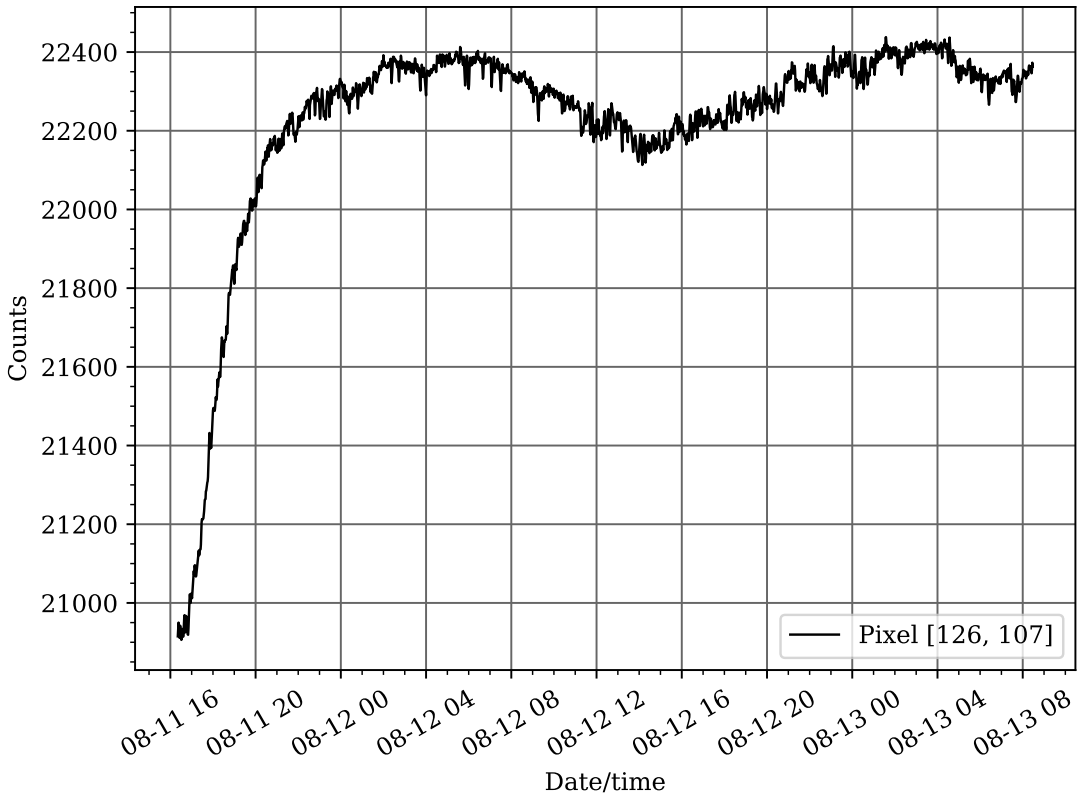

Figure 5.7: Camera 1 POI3 timeseries from a forty hour cable sealing end experiment. 300 A input current.

5.5 Thermal Variation along the length of CSEs

5.5.1 Introduction

An important metric in the condition monitoring process is the temperature difference along the length of the CSE. Spot temperatures are taken at standard positions, or POIs along the CSE as defined by the condition monitoring procedure, these can be seen in Figure 5.1. Engineers can use these spot measurements to inform their evaluation of the condition of the equipment. They may seek to understand the temperature gradient down the CSE, for example, or to establish whether the temperatures in these spots are within an expected range. For this reason, it is useful to establish the range of temperatures experienced by an electrically loaded CSE in a laboratory environment, as it can inform future condition monitoring methodologies.

Firstly, it is possible to inform the methodologies by establishing a total thermal range that may be experienced during testing. To accomplish this, the differences between the busbar clamp and the CSE base are taken, as these are the hottest and coolest points within the region of interest in the scene. Because the base and busbar clamp are each metallic with similar surface finishes, there is a potential difference in emissivity, which is difficult to determine without empirical testing. The differences in digital counts output from the camera are presented in Table 5.2.

It is possible to further inform the condition monitoring strategy with regards to CSEs by inspecting the range of thermal output for the CSE body, corresponding to the insulator sheds. The volume within the sheds is where the sheath of the cable conductor is stripped back and terminated, creating an area of high electric field between the conductor and earthed parts of the CSE. This part of the CSE is therefore most susceptible to aging and could be a potential source of additional generated heat. Due to the presence of the stress-cone and cable insulation, the bottom of a CSE typically experiences the least heating, while the upper parts heat more due to the cable being un-insulated (though surrounding by oil) and due to the absence of oil at the highest points of the internal volume. It is useful to establish an understanding of this temperature gradient under normal loading conditions in the laboratory in order that it can be understood when the typical heating profile is deviated from.

This experiment is limited by the use of non-calibrated cameras, as it is only

able to provide an indication of the temperatures that may be present instead of a direct temperature. However, there is still insight to be gained by investigating the differences in output.

5.5.2 Method and Findings

Locations corresponding to positions 1 - 6 in Figure 5.1 were chosen in each of the 8 output camera frames. These locations are marked on the reference image and displayed in Figure 5.2. Position 7 is out of the field of view of the camera for this experiment. Counting from the uppermost point, the points of interest (POI) correspond to the busbar connection, top bulb, upper middle and lower CSE body, and bottom housing. The ΔY between the POIs is kept as consistent as possible between different images, and an attempt is made to ensure the POIs are located in positions corresponding to similar surface finished between images. For example, POI 1 is on the flat surface of the metal of the joint, away from fittings like nuts and bolts. In order to ensure differences in features are minimised between images, especially in difficult to interpret areas at the bottom of the frames, a 3x3 area is taken around the selected pixel and the mean of the 9 pixels is calculated.

Two datasets are used to extract values for this section. First is the 36 hour 300 A step-test, as presented in Figure 5.5 and 5.6. This test represents the highest consistent electrical loading experiment completed. The results are presented in Table 5.2. Second is the 11 day realistic load pattern test. This test represents the longest period of continuous operation recorded. The results are presented in Table 5.3. Both of these tests are undertaken in similar laboratory environmental conditions, though the realistic load-pattern test includes the use of the wind fan.

Treating each of these results as a standalone tests gives significant range of the potential difference in thermal output from the busbar to the base of the CSE. In the step-test, CSE1 shows a range of 2000 digital counts. The difference from camera 1 is almost 50 % more than that of camera 3. CSE2 shows a similar pattern of thermal output as each camera perspective is considered around the CSE. That the cameras in positions 1 and 4 output the highest difference in intensity, and the cameras in positions 3 and 7 output the lowest difference, with cameras in positions 2 and 4, for CSE 1, and 6 and 8 for CSE 2 matching the pattern, suggests some environmental impact. The cameras in positions 1 and

Table 5.2: Difference in output digital counts between the busbar and base, and between the uppermost and lowest point on the CSE body; 36 hour step test; 300 A current.

CSE	Camera	Max. Diff. Busbar/Base	Max. Diff. Body
1	1	6593	1770
1	2	5996	1617
1	3	4216	1416
1	4	4584	1673
2	5	821	368
2	6	666	337
2	7	579	377
2	8	646	282

Table 5.3: Difference in output digital counts between the busbar and base, and between the uppermost and lowest point on the CSE body; 11 day realistic load-pattern test.

CSE	Camera	Max. Diff. Busbar/Base	Max. Diff. Body
1	1	4569	1259
1	2	3853	1120
1	3	2640	1030
1	4	2819	1228
2	5	1079	392
2	6	888	355
2	7	672	433
2	8	710	362

5 face out into the large laboratory, with walls behind them. The cameras in positions 3 and 7 face the walls, with the laboratory behind them. The walls may be causing thermal reflections, impacting the measurement and causing the surface of the CSE opposite the walls to appear hotter than the surface opposite the open space of the laboratory. Another potential contribution to this is that the busbar clamp is non-uniform and therefore difficult to select a consistent point of interest for around the various camera perspectives. This highlights a potential source of measurement error in condition monitoring surveys.

Compared with the busbar to base results, the CSE body results are consistent. The pattern around the CSE is present, with the highest differences present from the perspectives of cameras 1 and 5. The results are relatively consistent around the diameter of the CSE, with a range of results spanning under 100 digital counts for CSE2, and under 400 digital counts for CSE1. This consistency highlights the care that is required in selecting points of interest: where it is relatively easy to select a point, such as on the CSE body, consistent results can be achieved. Where it is more difficult to select a consistent point of interest, such as on the busbar clamp, which has complex geometry, it is more difficult to get a consistent result from different angles. In order to make the selection of the points of interest as easy as possible, the images are normalised to enhance visible features so the points of interest can be properly located. Without this normalisation, it would be difficult to ensure that a point of interest on the body of the CSE corresponded to the core cylinder and not instead one of the sheds. As the sheds heat less than the core of the CSE (see Figures 5.3 - 5.6), this could lead to erroneous readings.

The differences present on the shorter, higher current test, are larger than those of the longer, lower average current test. This further informs the findings of Section 5.4, suggesting that while continued heating allows the base of the CSE to heat, the highest temperatures are a response to high current, and are not experienced when extended loading occurs.

An inconsistency in the results is clear with regards to the busbar clamp to base measurement, when considering both CSEs. Here the short test results in a higher temperature difference in CSE1, while a higher temperature difference is present on CSE2 for the 11 day realistic test. The difference between results for CSE2 is small, but consistent between cameras 5, 6, 7 and 8. This inconsistency may be a result of the aforementioned difficulties in correctly selecting points of

interest, or may be a result of the aging process CSE1 has undergone, as discussed in Section 5.7.

The primary finding of this investigation is insight into the inherent sensitivity of thermal cameras to non-intuitive factors. Significantly, there is very large variance in the amount of potential difference present in a scene. Measurement noise is a large factor in the results, as can be seen in the results pertaining to the busbar clamp. Where it is relatively simple to correctly select a point of interest on the CSE (the body), the results are fairly consistent between cameras. The discrepancies between cameras in this case are, as described, likely due to environmental differences. However, where it is difficult to consistently choose the point of interest (on the busbar clamp), there is a very large variation in results. It is accepted though that without knowing categorically the health of the CSEs used at this point of the experiment, no definitive conclusions can be reached about the cause of the output inconsistencies.

5.6 Thermal variation around the circumference of CSEs

5.6.1 Method I: Utilising Overlapped Pixels

The experiment was designed such that the FOV of each of the camera perspectives would overlap. The arrangement of eight bolts and supportive fins at the top of each CSE provides a reference point for determining where the overlap occurs, as can be seen in Figure 5.8. Furthermore the double-helix arrangement of the sheds can be used as a reference point between frames, specifically the start point of the helix at the top end and bottom end. The pixels selected to correspond to these overlapping points are provided in Figure 5.9 and Table 5.4.

The benefit using this kind of overlap calculation is that it guarantees that the surface emissivity and surface temperature will be approximately equal between the two camera perspectives. If the pixel location is chosen correctly, the view angle between the surface of the CSE and each of the cameras should be approximately equal, as the ideal pixel to use as the overlap is equidistant between the two cameras. The detriment of using this kind of overlap calculation is that it is difficult to guarantee that the pixel corresponds to the same physical spot in each

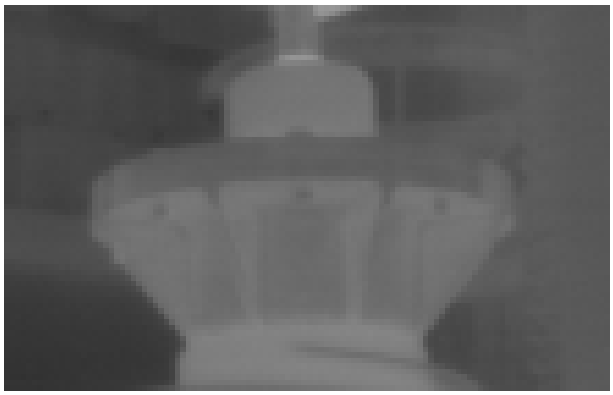

Figure 5.8: Thermal image of the top cap of a CSE, with fins as used for overlap reference

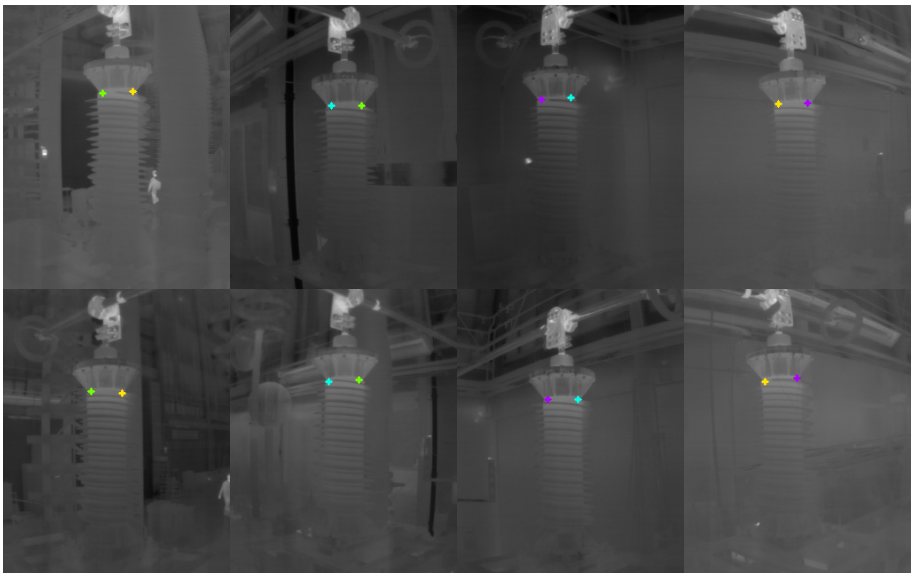

Figure 5.9: Overlap pixel spot selection overlaid on thermal images of the CSEs

Table 5.4: Overlap pixel locations for overlap analysis

CSE	Camera Pair	Colour	Cam A px(y, x)	Cam B px(y, x)
1	1, 2	Green	(99, 112)	(113, 403)
1	2, 3	Turquoise	(113, 370)	(104, 640)
1	3, 4	Purple	(106, 607)	(110, 907)
1	4, 1	Yellow	(111, 874)	(97, 146)
2	5, 6	Green	(435, 99)	(422, 401)
2	6, 7	Turquoise	(424, 367)	(444, 648)
2	7, 8	Purple	(444, 614)	(420, 895)
2	8, 5	Yellow	(424, 859)	(437, 134)

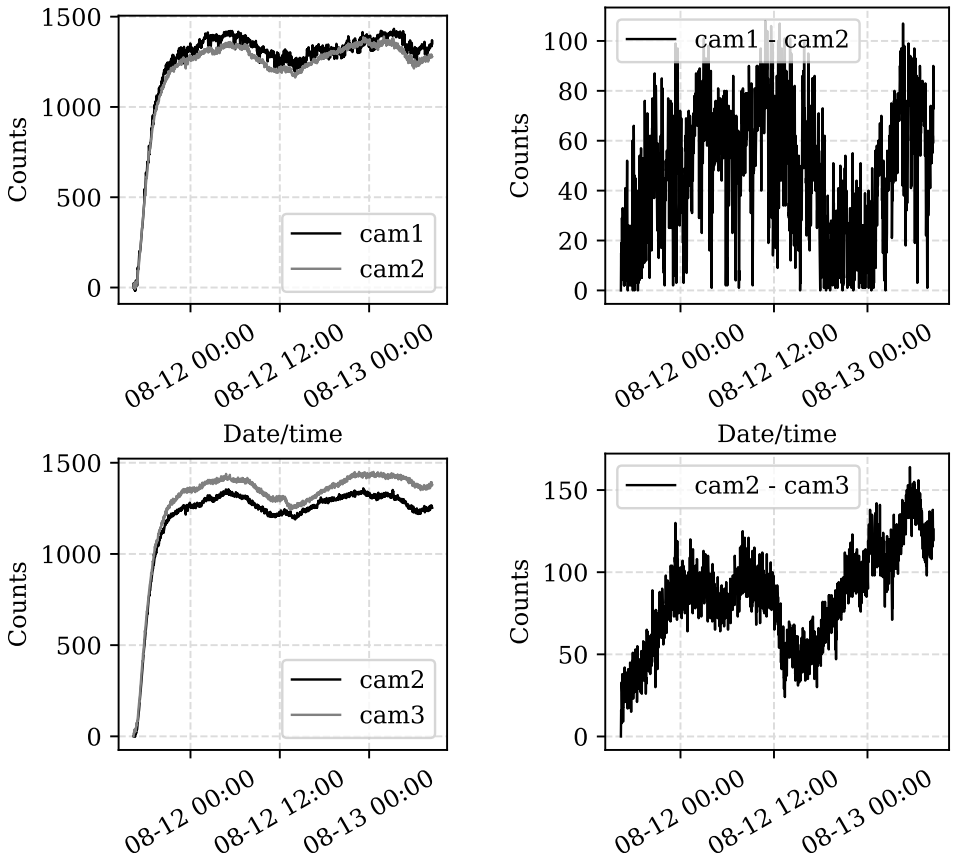

Figure 5.10: Left: overlap pixel intensities from cameras 1, 2 and 3; right: difference between camera outputs; 36 hour step-test.

of the images. If the pixel location is chosen correctly, the angle of emission of the IR radiation from the surface to each of the cameras should be approximately equal, as the ideal pixel to use as the overlap is equidistant between the two cameras. Using the background radiation as a bias calculator is not as reliable in this respect - the rig is situated in a busy laboratory, with various equipment surrounding it, and walls at different distances in every direction.

This section explores the difference in output for different cameras for pixels corresponding to the same physical location. The aim of this section is to provide insight into how the different cameras respond differently to equivalent thermal output.

This may be used in order to provide information about any asymmetric heating effects present on the assets. To summarise, points with relatively large thermal variation (near the top) are selected using the physical features visible on the CSEs. The digital count output corresponding to these points is plotted. For

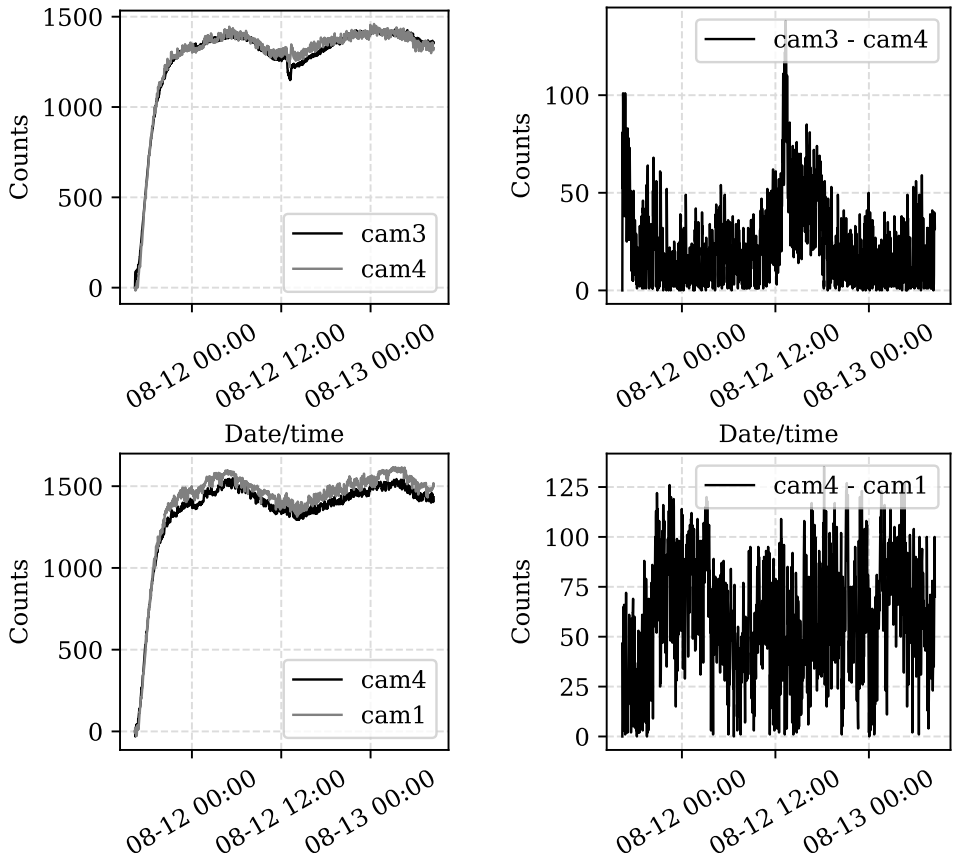

Figure 5.11: Left: overlap pixel intensities from cameras 3, 4 and 1; right: difference between camera outputs; 36 hour step-test.

the purposes of the plots, the digital count output value at time t=0 is subtracted from the entire time series, to remove bias. The difference between the output for the two camera perspectives is also plotted.

Figures 5.10 and 5.11 show a plot of the selected pixels over the period of the 36 hour step-test data set, alongside a plot of the difference between the two plots. The plots show encouraging results. The difference between the output for pairs of cameras frequently drops below 100 digital count units, in the context of the 16-bit resolution of the camera this is a close match. Broadly, the matches between the eight pairs of cameras are close, and vary in a similar manner. The plots show that, with the exception of the camera 3 and camera 4 pairing, the match is closest as the beginning of the test, diverging while the system heats. One potential cause of this is the selection of the overlap pixel, as all pixels are likely to be similar values when there is no active heating, though discrepancies are revealed when the CSE is heated. Another is that the response of each of

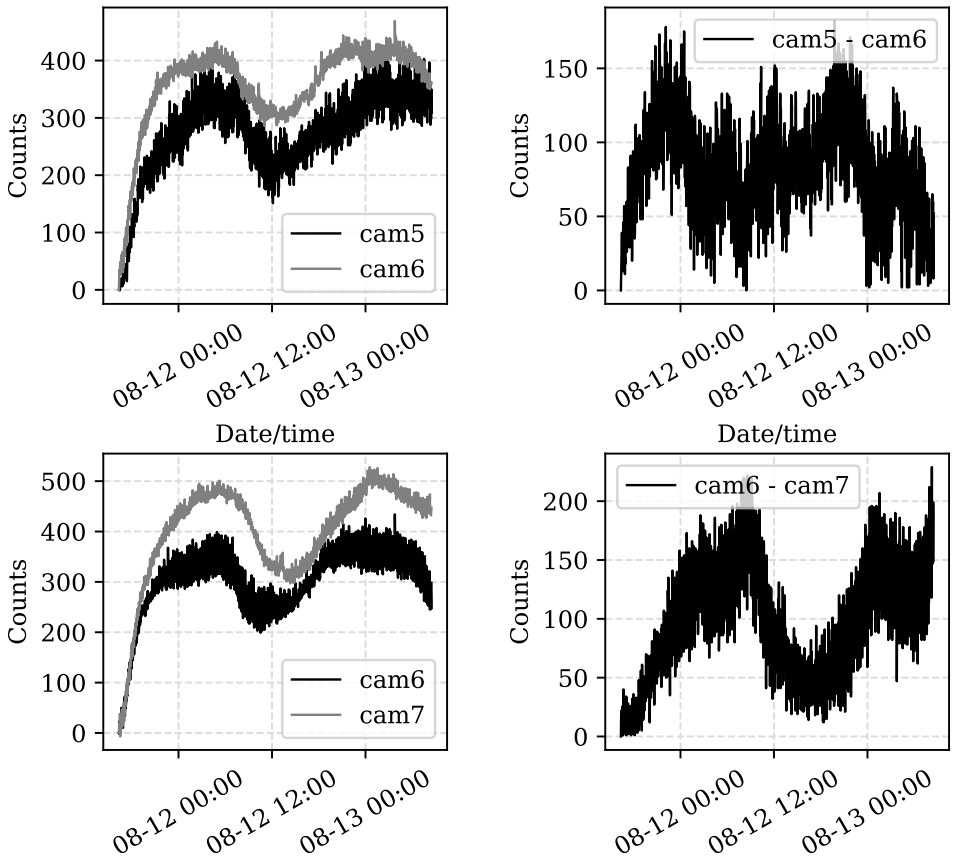

Figure 5.12: Left: overlap pixel intensities from cameras 5, 6 and 7; right: difference between camera outputs; 36 hour step-test.

the cameras is closest in the warmer periods of the tests (the beginning, mid and end-points) and responses diverge most when the cameras are coldest during the night.

In proportion to the total heating effect, the discrepancies between cameras is more significant for the cameras around CSE2. This is shown in Figures 5.12 and 5.13. CSE2 heats the least of the two, and generally the discrepancies are larger, with the exception of the pairing of cameras 5 and 8, as shown in Figure 5.13. Given that CSE2 heats less, it is marginally more difficult to correctly select the overlap pixels on the images, a potential source of the larger discrepancy.

It is known that each non-calibrated camera will have differing output: the sensors have slightly different responses to flux at different temperatures and the calibration procedure corrects for this. This simple investigation into removing the initial bias suggests that coarse calibration using overlapping pixels may be possible, in order to calibrate the camera output responses to match each other.

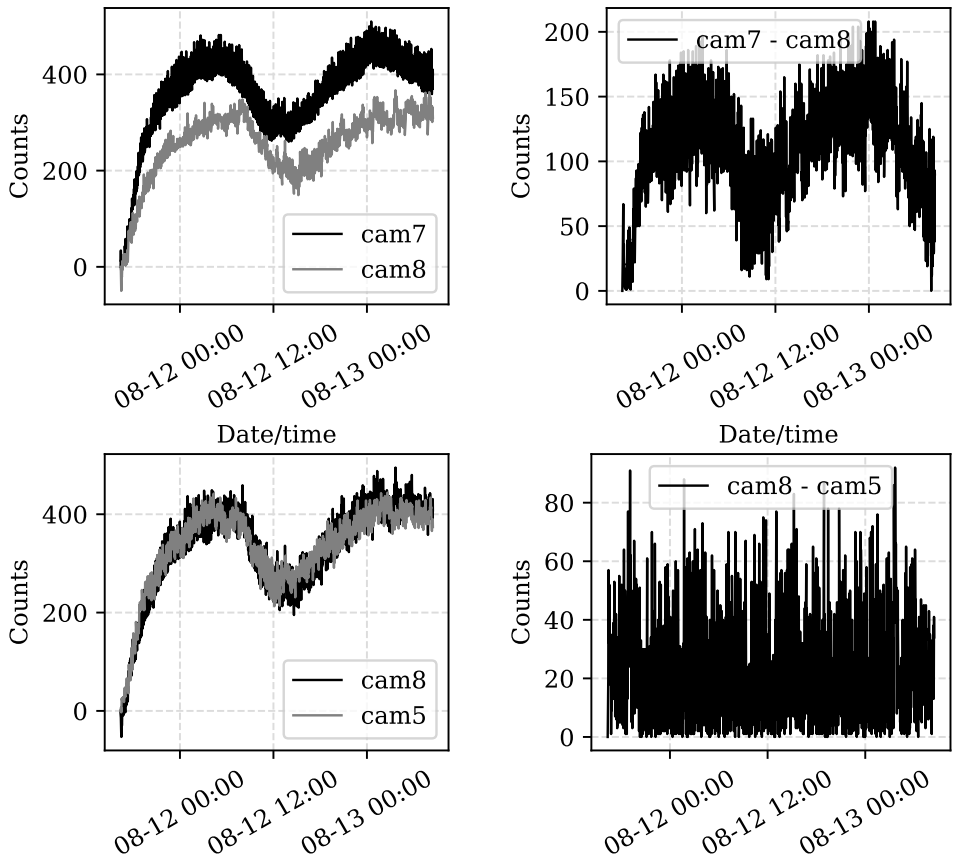

Figure 5.13: Left: overlap pixel intensities from cameras 7, 8 and 5; right: difference between camera outputs; 36 hour step-test.

If this were completed in future work, it would provide an opportunity to create a seamless 3D thermal image of an asset, facilitating ease of inspection.

5.6.2 Method II: Utilising Horizontal Intensity Lines

Lines of pixels around the top of the CSE, from each camera perspective, at the start and end of the 36 hour step-test, were selected and plotted. These are presented in Figures 5.14 and 5.15. These plots confirm the findings of the overlap tests: that the output profile does vary between the perspectives. If the output effects were identical, the lines would be equidistant for each set of cameras at each sample point presented, however the distances between the lines is not consistent. Based on the results of this test and others in this chapters, it is assumed that a minor factor in this discrepancy is the response of the cameras (this issue is specific to non-calibrated cameras) and that a more significant factor is the physical layout of the lab and the walls surrounding the CSE rig. A source of error in this test is that the pixel lines that are plotted are at in consistent heights up the CSE in relation to each other, due to the requirement that they correspond to the hot cable-sealing body and not the cooler shed fins.

Further testing should determine whether this is the case or not, by adjusting the camera positions around the CSEs to determine the effect it has on the output.

5.7 Difference in Thermal Profile Between CSEs

Throughout this chapter there is evidence of a difference between the two CSEs. In January 2020 the stress-cone within CSE 1 underwent an expedited aging process in separate laboratory experiments. Following this process, the CSE was re-constructed and the experiment rig was re-assembled.

This section aims to demonstrate the extent of the difference between the two CSEs and briefly discuss the potential causes. Figures 5.16 and 5.17 show clearly the significant difference in thermal output from February to August 2020. In February, both CSEs display a similar heating pattern, with CSE2 heating more than CSE1, but both settling in similar amounts of time. The output in August is dramatically different. CSE2 heats a moderate amount over the test-period, with an overall increase in digital counts of nearly 500. While this represents less heating than the February test, the laboratory was significantly warmer,

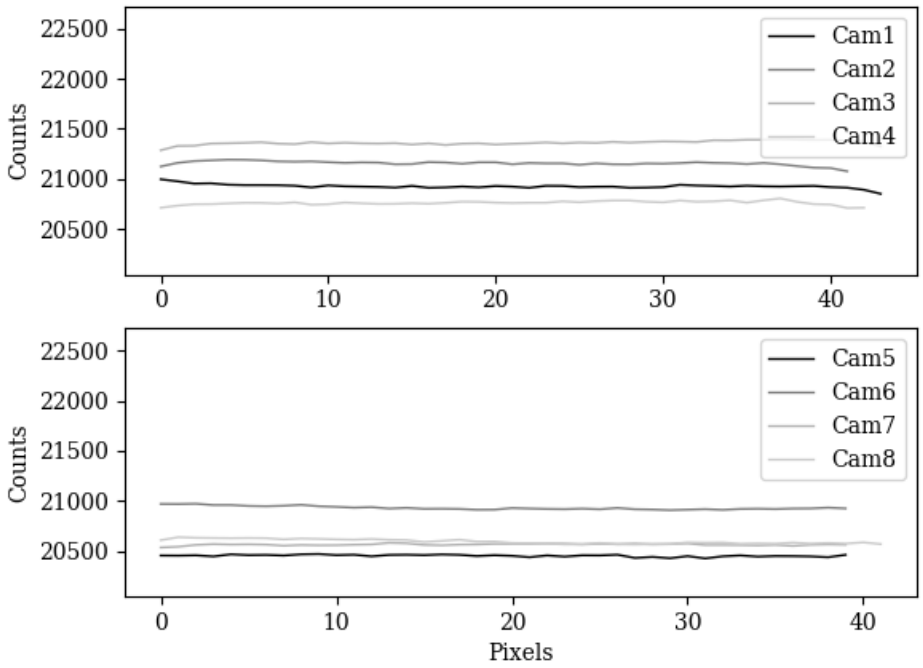

Figure 5.14: Plot of horizontal lines of pixels across the surface of the CSE, beginning of 36 hour step-test (11/08/2020 16:21).

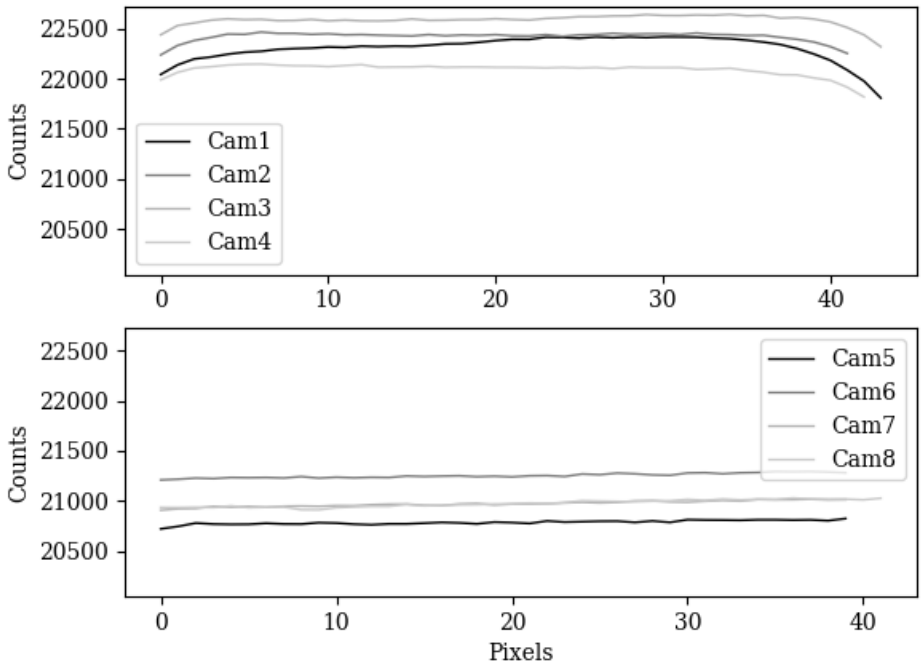

Figure 5.15: Plot of horizontal lines of pixels across the surface of the CSE, beginning of 36 hour step-test (13/08/2020 08:28).

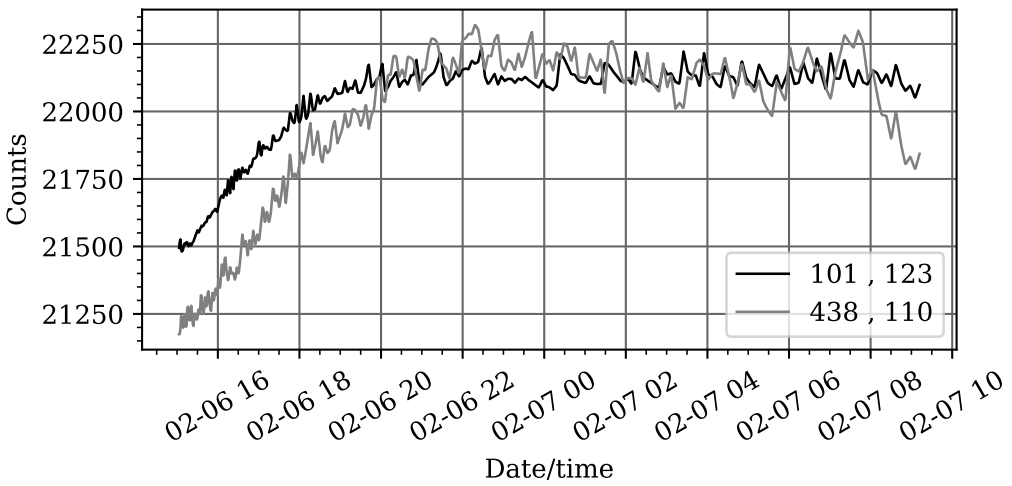

Figure 5.16: POI3 as viewed by cameras 1 and 5, February 2020

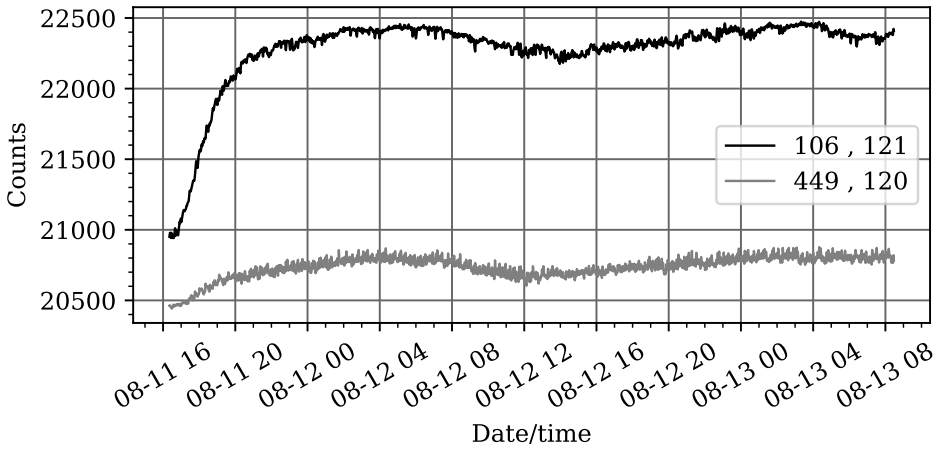

Figure 5.17: POI3 as viewed by cameras 1 and 5, August 2020

so this can be expected. Assuming CSE2 is still in a healthy state, it can be used as a reference to which the output of CSE1 can be compared. In August, CSE1 shows an increase of nearly 2500 digital counts. Given the assumption that a healthy CSE in the summer conditions would heat up less than in the winter conditions, that the aged CSE heated up significantly more than in winter suggests that the aging process may have stimulated measurable change in the system output. However, there are a number of other potential causes for the extra heating. Firstly, on occasion a small pool of oil was discovered underneath CSE1, suggesting a light leak. Due to the small size of the pool in comparison to the volume of oil the CSE holds it seems unlikely that this is the root cause of the issue, but measurements should be taken with the correct amount of oil for

a fair comparison.

The differences in heating observed for POI3 are reflected in POI1 (busbar clamp): POI1 corresponding to CSE1 shows significantly more heating in the summer test, while POI1 for CSE2 shows a smaller amount of heating. It is possible that the clamp has become loose, because of poor initial fitting or thermal cycling. While it seems unlikely to be the root cause of the extra heating observed in the CSE body, it would be recommended to check the connections are secure if investigating the potential presence of a fault thoroughly.

5.8 Discussion

This work has provided significant insight into the process of evaluating the output of non-calibrated thermal cameras. While there has been a consistent limitation in the inability to quantify what temperature a given digital count output corresponds to, insights have seen been gained. The expected output of 66 kV CSEs under load has been established, with some indication of how this output should be expected to vary over time and over the surface of the CSE. Furthermore, a CSE deviating from the expected response, following a stimulated aging process, has been recorded. That CSE1 heats considerably more than previously, in conditions under which it would be expected to heat less, may be a significant finding.

Non-calibrated thermal cameras have been found to be sensitive to inconsistencies in measurement. It has been found that, in order to achieve repeatable results, points of interest must be selected carefully and consistently. In the case where this has occurred, the output of these cameras can be consistent, and as described above, used to gain insight into thermal profiles both over space and time. A further nuance of these thermal cameras is that their high sensitivity can mean non-intuitive factors influence the output, for example, the hypothesised impact of the nearby walls and empty space of the laboratory, as described in Section 5.6.

The investigation reveals similarities in thermal output between the 66 kV CSE presented here and the 132 kV asset presented in [67]. Firstly, it is found that there is temperature gradient down the length of the CSE, and secondly, that minimal heating occurs on the outer edge of the shed surface. While this work has not quantified these findings into temperature values, it may be of value to know

that CSEs with different geometry have similar thermal output profiles under load in laboratory conditions. It is additionally notable, that when the internal conditions of the 66 kV CSE are changed, the shape of the heating profile does not appear to change, but the amplitude of the profile does. This is in contrast with the findings of the 132 kV CSE.

5.9 Summary

This chapter provides an opportunity to investigate the thermal output of 66 kV CSEs, providing valuable insight to the industrial sponsor, while providing additional contextual knowledge to assist in the modelling process described in Chapter 4. The manual investigation into the output of the cameras during various tests is useful when informing the modelling process and provides intuition with regards to the equipment, assisting in extracting information from it. It also is an opportunity to expand on existing work into the thermal monitoring of larger 132 kV CSEs, completed by Zachariades et al., evaluating whether the responses recorded during testing match what was found for the larger units. Lastly the creation of the experimental rig, with the ability to energise multiple assets in different conditions for days uninterrupted, with both high-voltage and high-current, while capturing a full coverage of thermal images, constitutes a contribution.

Furthermore, the experiment offers a significant contribution to existing work on thermal monitoring of oil-filled cable terminations such as CSEs. It expands on work completed by researchers within the university on 132 kV CSEs, supporting and supplementing a number of their findings, in the slightly different context of a 66 kV CSE.

Chapter 6

Conclusions and Future Work

6.1 Conclusions

This thesis presents two sets of experiments, associated data and analysis regarding electrical substation condition monitoring, comprising of four distinct contributions. Firstly, it is demonstrated that electrically loaded substation equipment can be modelled to a high degree of accuracy, using low-cost sensors, including non-calibrated FLIR Boson thermal cameras. The thesis demonstrates that LSTM recurrent neural networks provide the most accurate models of those that are utilised as part of the work but acknowledges it is not an exhaustive test of modelling techniques. The utilisation of LSTM recurrent neural networks in this context forms the second contribution. It shows that LSTM recurrent neural networks can provide accurate models for both simple electrical conductors, with a low thermal mass and fast response, and large cable terminations, comprised of solid and fluid insulators, with a high thermal mass and resulting slow thermal response, under the influence of significant environmental variation. The development of the large scale experimental rig facilitating the long-term monitoring of electrically loaded power equipment forms the third contribution. Finally, it presents a case study into the otherwise thermal output of a specific pair of 66 kV cable sealing ends, providing insight into the manner in which they heat over time and space, both in healthy condition and after undergoing an accelerated aging process, concluding that there may be differences in how different types of CSEs present indications of bad asset-health.

Broadly this thesis demonstrates that time-series modelling of environmental and electrical load data is a valid methodology for monitoring asset health of

electrically loaded substation equipment, provided sufficient variables and time are monitored.

In summary, accurate multivariate time series models of the thermal output of electrically loaded substation equipment under the influence of varying environmental conditions in a laboratory, have been generated. A broad understanding of the limitations of thermal imaging as a condition monitoring tool has been developed, and used to inform data collection and processing decisions. The work contributes an application of a modern, well-performing modelling technique, to a unique dataset, with results that encourage further investigation, providing value to the academic community and to the industrial sponsor.

6.2 Future Work

The primary recommendation based on the work presented in the thesis is for expanded and extended data collection. The significant limitation of the work presented is the ambiguity regarding it's genuine utility in significantly more volatile outdoor substation environments. A comprehensive data collection platform monitoring selected in-service substation assets would provide an opportunity to test the ability of the adopted methods to characterise the thermal output of equipment. Such a platform would be capable of capturing all significant environmental influences on thermal camera output, including precipitation, humidity, view distance and angle, solar radiation and shade, air temperature, wind speed and direction and load. Furthermore, in-service substation assets would be operating at high voltage in addition to high current, therefore effects not considered in this work such as localised heating due to partial discharge, could be modelled. Development of a suitable platform would not be trivial, involving integration of multiple discrete sensors, including a visible light camera and ideally a calibrated thermal camera. Further envisioned difficulties include extracting accurate and timely electrical load data, ensuring that the selected assets were at the desirable level of asset health, and dealing with any variation in emissivity over time, due to pollution, bird waste or similar. Ideally this monitoring platform would be installed for a period spanning over a year if situated in the UK, in order to ensure the domain of the time-series model included as much potential environmental variation as possible.

A secondary recommendation would be to explore the creation of a more

general LSTM based model. One of the significant benefits of neural network technologies is the scalability, therefore it is intuitive to extend the modelling architecture to model multiple point of interest time series simultaneously. The development of such a model may provide interesting challenges regarding the encoding of the pixel locations relative to the environmental effects, and provide a modelling challenge in capturing the difference in extent of environmental and loading effects at different points on an asset.

A number of potentially beneficial features of LSTM neural networks are not explored within the work. The cable-sealing end system is framed as consisting of multiple inputs and a single output, which corresponds to the thermal output as viewed through a thermal image. It is possible to manipulate the configuration of the neural network in order to receive more inputs and generate more outputs. The first potential utilisation of this would be to explore modelling multiple points of interest in a scene with a single model. This would improve the data pipeline, allow a model to attempt to characterise the different responses of different temperatures but also conceivably result in a longer training time and overall worse performance. The second potential utilisation of the technology would be to generate output sequences, corresponding to multiple-step forecasts, rather than single step forecasts. This could be useful in specific condition monitoring applications, for example, if a model were generated using data collection in a substation on a temporary basis, and on inspection, short-term data collection is completed and used to generate forecasts which are compared to the real output.

The fact that there are many variables contributing to the model, and that the influence of these variables is opaque to the user due to the form of the LSTM neural network, is a limitation. The network configuration selected for these experiments, with 30 lags of each input and 300 LSTM units, has 371,101 trainable parameters. It is possible to explore the effects of the inputs by generating artificial data, but this is non-trivial when compared to traditional regression analysis.

Given that the fault detection operation, that the models generated in this chapter are generated for, is not completed, it is difficult to interpret whether the models are suffiecient to be used in this manner. The models appear to perform well but this cannot be categorically confirmed without extending the work. Additionally, collecting data from faulty cable-sealing ends is inherently

difficult, as it is often not known whether a unit is faulty until failure has occurred. As discussed in Chapter 3, the range of thermal variation in the scenes imaged by the various cameras could be exploited to provide false 'faulty' data, in order to test the model-based fault detection method implementations, determining their sensitivity to the data and model performance. This is a strongly recommended body of future work and is a natural progression of what is presented in this chapter.

Further recommendations include extensive exploration of multiple time-series modelling techniques, with wide-ranging configuration testing, using the data acquired during this work. There are a wide range of available methods and there may be opportunities to explore modelling the data with greater ability to deterministically understand the influence of the input variables, for example, with models with significantly less parameters.

Other supplementary bodies of work could explore the issue of sensor coverage in large complex substations, analysing the cost and benefit of multiple static sensors as opposed to a mobile robotic platform. A potentially interesting area to explore would be how the directional influence of environmental effects should influence the optimal path planning of such a robotic inspection platform, for example: should an asset be observed from the north, in shade, for the best signal to noise ratio?

Finally, the work could be expanded into other condition monitoring sensing modalities. Partial discharge monitoring is a significant aspect of condition monitoring practise, and may generate useful insight if monitored in the long-term.

Bibliography

[1] S. Bagavathiappan, B. B. Lahiri, T. Saravanan, John Philip, and T. Jayakumar. Infrared thermography for condition monitoring – a review. *Infrared Physics & Technology*, 60:35–55, 2013.

[2] E. C. Bortoni, G. S. Bastos, L. dos Santos, and L. E. Souza. Wind-influence modeling for outdoor thermographic surveys. In *Thermosense XXXII*, volume 7661, page 43:48. International Society for Optics and Photonics, 2010.

[3] E. C. Bortoni, L. Santos, and G. S. Bastos. A model to extract wind influence from outdoor IR thermal inspections. *IEEE Transactions on Power Delivery*, 28(3):1969–1970, 2013.

[4] E. da Costa Bortoni, L. dos Santos, G. S. Bastos, L. E. de Souza, and M. A. C. Craveiro. Extracting load current influence from infrared thermal inspections. *IEEE Transactions on Power Delivery*, 26(2):501–506, 2011.

[5] Caledonian Cables. All aluminum alloy conductor (AAAC) cables - ABC cable – aerial bundled cables. `https://www.caledonian-cables.co.uk/products/abc-cables/al/aaac.shtml`, 2018. Accessed: 2018.

[6] K. Chatterjee and S. Tuli. Image enhancement in transient lock-in thermography through time series reconstruction and spatial slope correction. *IEEE Transactions on Instrumentation and Measurement*, 61(4):1079–1089, 2012. Conference Name: IEEE Transactions on Instrumentation and Measurement.

[7] Chongsheng Cheng and Zhigang Shen. Detecting concrete abnormality using time-series thermal imaging and supervised learning. *arXiv preprint arXiv:1804.05406*, page 8, 2018.

[8] Y. C. Chou and L. Yao. Automatic diagnostic system of electrical equipment using infrared thermography. In *2009 International Conference of Soft Computing and Pattern Recognition*, pages 155–160, 2009.

[9] David Dolezilek. Understanding, predicting, and enhancing the power system through equipment monitoring and analysis. In *Western Power Delivery Automation Conference*, 2000.

[10] A. Duarte, L. Carrão, M. Espanha, T. Viana, D. Freitas, P. Bártolo, P. Faria, and H. A. Almeida. Segmentation algorithms for thermal images. *Procedia Technology*, 16:1560–1569, 2014.

[11] R. A. Epperly, G. E. Heberlein, and L. G. Eads. A tool for reliability and safety: predict and prevent equipment failures with thermography. In *Record of Conference Papers. IEEE Industry Applications Society 44th Annual Petroleum and Chemical Industry Conference*, pages 59–68, 1997.

[12] Qiang Fang and Xavier Maldague. A method of defect depth estimation for simulated infrared thermography data with deep learning. *Applied Sciences*, 10(19):6819, 2020. Number: 19 Publisher: Multidisciplinary Digital Publishing Institute.

[13] Matthew Fox, David Coley, Steve Goodhew, and Pieter Wilde. *Comparing transient simulation with thermography time series*. IBPSA-England, 2012. Journal Abbreviation: Building Simulation and Optimization Conference Publication Title: Building Simulation and Optimization Conference.

[14] "National Grid". Case study: How national grid uses analytics to keeps the lights on. `https://www.computerweekly.com/news/2240241752/Case-Study-How-National-Grid-uses-analytics-to-keeps-the-lights-on`. Accessed: 2021.

[15] Y. Han and Y. H. Song. Condition monitoring techniques for electrical equipment-a literature survey. *IEEE Transactions on Power Delivery*, 18(1):4–13, 2003.

[16] Yangdong He and Jiabao Zhao. Temporal convolutional networks for anomaly detection in time series. *Journal of Physics: Conference Series*, 1213:042050, 2019. Publisher: IOP Publishing.

[17] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, 1997.

[18] Niancang Hou. The infrared thermography diagnostic technique of high-voltage electrical equipments with internal faults. In *1998 International Conference on Power System Technology, 1998. Proceedings. POWERCON '98*, volume 1, pages 110–115 vol.1, 1998.

[19] A. S. Nazmul Huda and Soib Taib. Application of infrared thermography for predictive/preventive maintenance of thermal defect in electrical equipment. *Applied Thermal Engineering*, 61(2):220–227, 2013.

[20] I. Hwang, S. Kim, Y. Kim, and C. E. Seah. A survey of fault detection, isolation, and reconfiguration methods. *IEEE Transactions on Control Systems Technology*, 18(3):636–653, 2010. Conference Name: IEEE Transactions on Control Systems Technology.

[21] Clemente Ibarra-Castanedo, Stefano Sfarra, Marc Genest, and Xavier Maldague. Infrared vision: Visual inspection beyond the visible spectrum. In Zheng Liu, Hiroyuki Ukida, Pradeep Ramuhalli, and Kurt Niel, editors, *Integrated Imaging and Vision Techniques for Industrial Inspection: Advances and Applications*, Advances in Computer Vision and Pattern Recognition, pages 41–58. Springer, 2015.

[22] IBM. IBM analytics powering national grid. `https://techmonitor.ai/techonology/data/ibm-analytics-powering-national-grid-4504139`. Accessed: 2021-03.

[23] Magna Power Electronics Inc. Magna power MS series IV DC power supplies operating and service manual. `https://magna-power.com/assets/files/manuals/user-manual_ms_4.0.pdf`. Accessed: 2017-08.

[24] M. S. Jadin, K. H. Ghazali, and S. Taib. Detecting ROIs in the thermal image of electrical installations. In *2014 IEEE International Conference on Control System, Computing and Engineering (ICCSCE 2014)*, pages 496–501, 2014.

[25] M. S. Jadin, K. H. Ghazali, S. Taib, and N. Huda. Finding ROIs in infrared image of electrical installation for qualitative thermal condition evaluation. In *2012 IEEE International Conference on Control System, Computing and Engineering*, pages 244–249, 2012.

[26] Mohd Shawal Jadin and Soib Taib. Recent progress in diagnosing the reliability of electrical equipment by using infrared thermography. *Infrared Physics & Technology*, 55(4):236–245, 2012.

[27] H. H. Jung and J. Lyou. Matching of thermal and color images with application to power distribution line fault detection. In *2015 15th International Conference on Control, Automation and Systems (ICCAS)*, pages 1389–1392, 2015.

[28] Julia Kelly, Natascha Kljun, Per-Ola Olsson, Laura Mihai, Bengt Liljeblad, Per Weslien, Leif Klemedtsson, and Lars Eklundh. Challenges and best practices for deriving temperature data from an uncalibrated UAV thermal infrared camera. *Remote Sensing*, 11(5):567, 2019. Number: 5 Publisher: Multidisciplinary Digital Publishing Institute.

[29] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv:1412.6980 [cs]*, 2017.

[30] A. Kingsmill, S. Jones, and J. Zhu. Application of new condition monitoring technologies in the electricity transmission industry. In *Sixth International Conference on Electrical Machines and Systems, 2003. ICEMS 2003.*, volume 2, pages 852–855 vol.2, 2003.

[31] R. H. Kingston. *Detection of Optical and Infrared Radiation*. Springer, 2013. Google-Books-ID: oB2hBwAAQBAJ.

[32] B. B. Lahiri, S. Bagavathiappan, T. Jayakumar, and John Philip. Medical applications of infrared thermography: A review. *Infrared Physics & Technology*, 55(4):221–235, 2012.

[33] Pedro Lara-Benítez, Manuel Carranza-García, José M. Luna-Romera, and José C. Riquelme. Temporal convolutional networks applied to energy-related time series forecasting. *Applied Sciences*, 10(7):2322, 2020. Number: 7 Publisher: Multidisciplinary Digital Publishing Institute.

[34] B. Li, X. Zhu, S. Zhao, and W. Niu. HV power equipment diagnosis based on infrared imaging analyzing. In *2006 International Conference on Power System Technology*, pages 1–4, 2006.

[35] K. B. Liland, K. Eidnes, K. Bjorneklett, and S. Hvidsten. Measurement of solubility and water content of insulating oils for HV XLPE cable terminations. In *Conference Record of the 2008 IEEE International Symposium on Electrical Insulation*, pages 7–10, 2008. ISSN: 1089-084X.

[36] K.B. Liland, Sverre Hvidsten, G. Birkenes, and Frank Mauseth. *Development of a simple method for conduction assessment of oil filled cable terminations.* Conference: Jicable 2011, 2011.

[37] Y. Lozanov, S. Tzvetkova, and A. Petleshkov. Use of machine learning techniques for classification of thermographic images. In *2020 12th Electrical Engineering Faculty Conference (BulEF)*, pages 1–4, 2020.

[38] R.E. Mackiewicz. *Overview of IEC 61850 and benefits.* IEEE, 2006. Journal Abbreviation: 2006 IEEE PES Power Systems Conference and Exposition, PSCE 2006 - Proceedings Pages: 630 Publication Title: 2006 IEEE PES Power Systems Conference and Exposition, PSCE 2006 - Proceedings.

[39] Sergio Marinetti and Pier Giorgio Cesaratto. Emissivity estimation for accurate quantitative thermography. *NDT & E International*, 51:127–134, 2012.

[40] Sergio Marinetti and Pier Giorgio Cesaratto. Emissivity estimation for accurate quantitative thermography. *NDT & E International*, 51:127–134, 2012.

[41] B. K. McAtee, A. J. Prata, and M. J. Lynch. The angular behavior of emitted thermal infrared radiation (8–12 um) at a semiarid site. *Journal of Applied Meteorology and Climatology*, 42(8):1060–1071, 2003. Publisher: American Meteorological Society Section: Journal of Applied Meteorology and Climatology.

[42] Nathan T. Moja and Andrew J. Willis. Classification of thermally condition-monitored components using statistical and neural network techniques. In Andrew G. Tescher, editor, *Proceedings Volume 3164, Applications of Digital Image Processing XX*, pages 573–581, 1997.

[43] Andrew Y. Ng. Feature selection, L_1 vs. L_2 regularization, and rotational invariance. In *Proceedings of the twenty-first international conference on Machine learning*, ICML '04, page 78. Association for Computing Machinery, 2004.

[44] Optotherm. Emissivity table. https://www.optotherm.com/emiss-table.htm. Accessed: 2021-03.

[45] N. Otsu. A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, 9(1):62–66, 1979.

[46] Aurabind Pal and Roma Dash. A paradigm shift in substation engineering: IEC 61850 approach. *Procedia Technology*, 21:8–14, 2015.

[47] National Grid UK ET plc. Safety rules and guidance - national grid UK electricity transmission plc. https://www.nationalgrid.com/uk/electricity-transmission/sites/et/files/documents/NGUKPLETSRGN%20Issue%207%20Final.pdf. Accessed: 2018-03.

[48] Ashok Prajapati, James Bechtel, and Subramaniam Ganesan. Condition based maintenance: a survey. *Journal of Quality in Maintenance Engineering*, 18(4):384–400, 2012. Publisher: Emerald Group Publishing Limited.

[49] A. Rahmani, J. Haddadnia, and O. Seryasat. Intelligent fault detection of electrical equipment in ground substations using thermo vision technique. In *2010 2nd International Conference on Mechanical and Electronics Engineering*, volume 2, pages V2–150–V2–154, 2010.

[50] K. S. Reddy, G. Veershetty, and T. Srihari Vikram. Effect of wind speed and direction on convective heat losses from solar parabolic dish modified cavity receiver. *Solar Energy*, 131:183–198, 2016.

[51] SMP Robotics. Electrical substation inspection robot. http://smprobotics.com/products_autonomous_ugv/electrical-substation-inspection-robot/. "Accessed: 2017-06".

[52] L. dos Santos, E. C. Bortoni, L. E. Souza, G. S. Bastos, and M. a. C. Craveiro. Infrared thermography applied for outdoor power substations. In *Thermosense XXX*, volume 6939, page 69390R. International Society for Optics and Photonics, 2008.

[53] J. Snell and J. Renowden. Improving results of thermographic inspections of electrical transmission and distribution lines. In *2000 IEEE ESMO - 2000*

IEEE 9th International Conference on Transmission and Distribution Construction, Operation and Live-Line Maintenance Proceedings. ESMO 2000 Proceedings. Global ESMO 2000. The Pow, pages 135–144, 2000.

[54] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014.

[55] Alastair Straker, Joaquin Carrasco, Frank Podd, Richard Gardner, and Ian Cotton. Improving thermal substation inspections utilising machine learning. In *Thermosense: Thermal Infrared Applications XLI*, volume 11004, page 1100406. International Society for Optics and Photonics, 2019.

[56] FLIR Commercial Systems. FLIR boson engineering datasheet. `https://flir.netx.net/file/asset/15754/original/attachment`. Accessed: 2017-07.

[57] FLIR Commercial Systems. FLIR lepton engineering datasheet. `https://www.flir.com/globalassets/imported-assets/document/flir-lepton-engineering-datasheet.pdf`. Accessed: 2017-07.

[58] Irfan Ullah, Rehan Ullah Khan, Fan Yang, and Lunchakorn Wuttisittikulkij. Deep learning image-based defect detection in high voltage electrical equipment. *Energies*, 13(2):392, 2020. Number: 2 Publisher: Multidisciplinary Digital Publishing Institute.

[59] Irfan Ullah, Fan Yang, Rehanullah Khan, Ling Liu, Haisheng Yang, Bing Gao, and Kai Sun. Predictive maintenance of power substation equipment by infrared thermography using a machine-learning approach. *Energies*, 10(12):1987, 2017. Number: 12 Publisher: Multidisciplinary Digital Publishing Institute.

[60] W. A. M. Ursine, J. L. Silvino, L. G. Fonseca, and R. M. de Andrade. Metal-oxide surge arrester's leakage current analysis and thermography. In *2013 International Symposium on Lightning Protection (XII SIPDA)*, pages 297–303, 2013.

[61] Rubén Usamentiaga, Venegas Pablo, Jon Gueredíaga, Laura Vega, Julio Molleda, and Francisco Bulnes. Infrared thermography for temperature measurement and non-destructive testing. *Sensors (Basel, Switzerland)*, 14:12305–12348, 2014.

[62] Vaisala. Vaisala WXT520 datasheet. `https://www.vaisala.com/sites/default/files/documents/M210906EN-C.pdf`.

[63] Renzhuo Wan, Shuping Mei, Jun Wang, Min Liu, and Fan Yang. Multivariate temporal convolutional network: A deep neural networks approach for multivariate time series forecasting. *Electronics*, 8(8):876, 2019. Number: 8 Publisher: Multidisciplinary Digital Publishing Institute.

[64] Binhai Wang, Rui Guo, Bingqiang Li, Lei Han, Yong Sun, and Mingrui Wang. SmartGuard: An autonomous robotic system for inspecting substation equipment. *Journal of Field Robotics*, 29(1):123–137, 2012.

[65] Q. Wang, Q. Liu, R. Xia, G. Li, J. Gao, H. Zhou, and B. Zhao. Defect depth determination in laser infrared thermography based on LSTM-RNN. *IEEE Access*, 8:153385–153393, 2020. Conference Name: IEEE Access.

[66] Hanyu Ye, Tobias Fechner, Xianzhang Lei, Yi Luo, Mingyu Zhou, Zhengyi Han, Haitian Wang, Qikai Zhuang, Ruoyu Xu, and Duo Li. Review on HVDC cable terminations. *High Voltage*, 3(2):79–89, 2018. _eprint: https://ietresearch.onlinelibrary.wiley.com/doi/pdf/10.1049/hve.2017.0144.

[67] C. Zachariades, V. Peesapati, R. Gardner, and O. Cwikowski. Electric field and thermal analysis of 132 kV ceramic oil-filled cable sealing ends. *IEEE Transactions on Power Delivery*, 36(1):311–319, 2021. Conference Name: IEEE Transactions on Power Delivery.

[68] Z. Zhao, G. Xu, and Y. Qi. Representation of binary feature pooling for detection of insulator strings in infrared images. *IEEE Transactions on Dielectrics and Electrical Insulation*, 23(5):2858–2866, 2016.

Appendix A

Software API

A.1 GitHub

A link to the API utilised in pre-processing and modelling the data during the work is provided: `https://github.com/astraekr/thermal_preprocess`