

Interreg EUROPEAN UNION

North-West Europe

REAMIT

European Regional Development Fund

WP T2 - Deliverable 8.1

Deployment of the integrated IoT, Big Data, analytics, Decision support technology

Improving Resources Efficiency of Agribusiness supply chains by Minimizing waste using Internet of Things sensors (REAMIT)



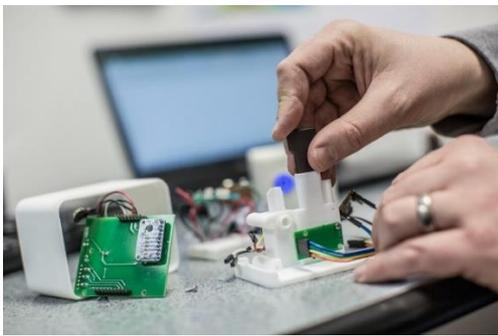
Table of Contents

1. Introduction.....	3
2. HMF.....	3
Preliminaries.....	3
Technology stack selected.....	5
Analytics	5
Pre-processing	5
Analysis	5
Interpretation	7
Deployment	7
3. Yumchop	9
Preliminaries.....	9
Analytics	9
Pre-processing	9
Analysis	9
Interpretation	11
4. Biogros	14
Preliminaries.....	14
Technology stack selected.....	15
Analytics	15
Pre-processing	15
Analysis	15
Interpretation	17
Deployment	18
5. Musgrave	20
Preliminaries.....	20
Technology stack selected.....	21
Analytics	21
Pre-processing	21
Analysis	21
Interpretation	22
Deployment	23
6. WD Meats & Burns Farm Meats	25
Preliminaries.....	25
Technology stack selected.....	26

Analytics	26
Pre-processing	26
Analysis	27
Interpretation	30
Deployment	30
Additional analysis	31
Burns Farm Meats.....	32
Alerting logic.....	33
7. Picnic	37
Preliminaries	37
Technology stack selected	38
Analytics	38
Pre-processing	38
Analysis	38
Interpretation	40
Deployment	41
8. Raman	42
Preliminaries	42
Technology stack selected	42
Analytics	42
Pre-processing	42
Fig. The necessary steps to pre-process Raman spectra	43
Analysis	43
Interpretation	45
Deployment	45
9. 3D Fluorescence (FreshDetect)	48
Preliminaries	48
Technology stack selected	48
Analytics	48
Pre-processing	48
Analysis	49
Interpretation	50
Deployment	50

1. Introduction

One of the most important activities in REAMIT project is to integrate and test the REAMIT technology in the pilot companies. The integrated technology has the capability to bring together the outputs of sensor technology and analytics to provide decision support to owners of food at risk, truck drivers and warehouse managers via the Smartphone App. This report presents REAMIT deliverable 8.1 on “Deployment of the integrated IoT / Big Data / analytics / Decision support technology”. The deployment of REAMIT technologies involves analysing the problems and challenges faced when dealing with the food at risk in the business process, food safety requirements, selection and deployment of suitable sensors and technologies, data collection, pre-processing and analysis. Eight pilot cases are presented in this report based on the deployment reporting framework.



2. HMF

Preliminaries

The Human Milk Foundation (HMF) plays a critical role in providing rapid delivery of breast milk to neonatal wards in South England using motorcycles and cars. However, due to a lack of means of temperature monitoring, there is presently a risk of spoilage if their cold storage fails without their knowledge until the milk is delivered to the hospital. To help reduce the chances of undetected cold storage failures during transportation, the REAMIT team proposed deploying temperature and humidity IoT sensors with HMF allowing them to monitor the parameters of the storage conditions and receive alerts if problems arose during milk delivery. With the temperature tracking sensors planned for use at HMF, however, another opportunity was discovered using the collected data.

Through analysis of the external parameters and their impact on ambient temperature changes during milk transportation, it was identified that the collected data from the temperature tracking sensors could be combined with other relevant business data collected on each journey and utilised to develop a regression model. With the regression model, HMF would have a quantitative tool to estimate, based on the specific planned journey parameters, the maximum duration and distance for delivery journeys without compromising the quality or safety of the milk due to temperature threshold abuse. This predictive capability would allow Human Milk Bank (HMF) to optimise their delivery planning and potentially extend their reach to sites located further away. Overall, the model could help HMF avoid breaches in temperature thresholds that could result in milk wastage by offering a tool to safely plan logistics so that this doesn't occur. Ultimately, the proposed model could enhance the operational efficiency of the milk delivery process, reduce the risk of milk wastage, and contribute to the overall success of the HMF. The development of a User Interface (UI) application was proposed, enabling end users to interact with the trained regression model and generate predictions for maximum journey

length. Detailed below is a list of the equipment installed, the accompanying labels provided by HMF for each journey, the outline of the proposed analysis, and the data dictionary of all the variables in the complete dataset.

Sensors (x 10)

- Digital Matter eagle data logger
 - Internally records acceleration and GPS and reports a “trip” indication if these both update within the space of 5 minutes.
- T9603 T/RH probe
- Binary sensor providing bag lid status (1/0)

Business provided labels

- Date, time, and length of journey
- Weight of milk on journey
- Size of bag used for journey
- External weather temperature
- Type of vehicle used (car or motorbike)

Proposed Analysis

Produce a regression model to estimate the maximum journey length based on user provided variables about the details of the journey. The model should be accessible and interfaceable through a user-friendly UI by staff at HMF so they can make use of the predictions provided by the model.

Data Dictionary

Source	Column	Data Type	Description
Sensor	device_id	Int	Numeric device ID, unique to the sensor
Sensor	datetime_measure	ISO8601 date-time format, "yyyy-MM-dd'T'HH:mm:ss.SSS'Z'"	Datetime stamp of sensor recording
Sensor	battery	Float	Battery level during sensor recording
Sensor	temperature1	Float	Temperature recording inside bag
Sensor	humidity	Float	Humidity recording inside bag
Sensor	bag_status	char(1)	Indication if the bag is open (0) or closed (1)
Sensor	device_name	nvarchar(50)	Descriptive name of device. Should help identify where device is installed, but often left as device_id.
Sensor	Trip	char(1)	Indication if the bike is being driven (1) or is stationary (0) when the recording was made.
HMF spreadsheet	device_id		Numeric device ID, unique to the sensor
HMF spreadsheet	Datetime		Date and time of journey
HMF spreadsheet	journey_length		Approximate journey length in minutes

HMF spreadsheet	milk_quantity		Milk quantity on delivery (litres or KG)
HMF spreadsheet	bag_size		Small, medium, or large
HMF spreadsheet	vehicle_type		Bike or car
HMF spreadsheet	external_temperature		External weather temperature in Harpenden at time of milk departure

Technology stack selected

The analysis and application were developed using R, a popular open-source programming language known for its extensive capabilities in statistical analysis and machine learning. R provides a wide range of packages and libraries specifically designed for statistical modelling, making it a suitable choice for building predictive models. In addition to its statistical capabilities, R also offers the Shiny package, which allows for the development of interactive web applications. Shiny enables the creation of user-friendly front-end interfaces that can be deployed as web-based applications, providing a convenient way to showcase and interact with the models and their results. By utilising R and its Shiny package, the analysis and application development process can benefit from the rich statistical functionality of R while also providing an intuitive and interactive user interface through a web application, allowing for a comprehensive solution that integrates statistical modelling, data visualization, and user interaction.

Analytics

Pre-processing

The first step in performing regression modelling was to prepare a training dataset for building the model. To prepare the training dataset, the data gathered from 117 trips was first compiled and aggregated. Then the data was pre-processed, cleaned, and transformed to a form that was useful for model building. Each trip was identified in the sensor dataset using the 'trip' variable logged by the Eagle. A new variable 'temperature rise' was created by calculating the difference between each temperature sensor reading during the trip. Then, the timestamp of the identified trip was cross-referenced with the business spreadsheet provided by HMF. Once the journey information from the spreadsheet had been identified for the trip, the journey data (journey length, milk quantity, bag size, vehicle type, external temperature) was appended to the corresponding sensor data. The pre-processed dataset was divided into a training dataset and a test dataset. The ratio chosen for this split was 1:4, meaning that 20% of the data was allocated for testing the model's performance.

Analysis

In our analysis, we aimed to estimate the total viable journey length through a regression analysis on temperature rise. This estimation process allowed us to derive the rate of temperature change within the bag, subsequently enabling an estimation of the total journey duration based on a maximum permissible temperature threshold, set at -10°C for the case of HMF. In this context, the dependent variable was the temperature rise, while the independent variables encompassed time, milk quantity, external temperature, vehicle type, and bag size. Following data preparation, we adopted various modelling techniques to construct predictive models for approximating the total journey length. Specifically, these techniques encompassed linear regression, decision trees, random forest, and

neural networks. Each model underwent training using the prepared dataset. Evaluative measures for the model's predictive performance were applied utilizing an independent test dataset. The assessment was quantified through multiple statistical indicators, notably including R2 and RMSE. In the case of the linear model, we further conducted checks for Variance Inflation Factor (VIF), Durbin-Watson Test, F score, and P value.

- R-squared (R2) is a statistical measure that represents the proportion of the variance in the dependent variable (outcome) that can be explained by the independent variables (predictors) in the model. It provides an indication of how well the model fits the data. Higher values of R2 indicate a better fit, with 1.0 indicating a perfect fit.
- Root Mean Squared Error (RMSE) measures the average magnitude of the residuals (prediction errors) of the model. It provides a measure of the model's predictive accuracy, with lower values indicating better predictive performance. RMSE is calculated by taking the square root of the average of the squared differences between the predicted and actual values.
- To assess multicollinearity, the Variance Inflation Factor (VIF) was calculated. VIF measures the extent to which the variance of the estimated regression coefficients is increased due to multicollinearity in the predictors. A high VIF value (normally >5) suggests high multicollinearity, which can affect the reliability of the model.
- The Durbin Watson Test was performed to check for the presence of autocorrelation in the model's residuals. This test examines whether there is a systematic pattern or correlation among the residuals, which can impact the model's accuracy.
- The F Score and P value characteristics were also evaluated. The F Score is a measure of overall significance of the model, indicating whether the model as a whole is statistically significant. The P value helps determine the significance of individual predictor variables, indicating whether they have a significant impact on the outcome.

By analysing these statistical measures for each model, it was possible to assess their performance and select the most suitable model for the prediction task. Presented in table 1 are the results from each of these models.

Table 1: Results from regression modelling to predict temperature rise on HMF dataset.

	R2	RMSE (°C)	VIF	Durbin Watson	F-score	p-value
Linear regression	0.7396	0.7242	Time: 1.376079 Temp. Outside 1.549767 Quantity of milk 1.753195 Car: 1.377830 Size of bag (big): 1.493121 Size of bag (small): 1.446634	550.5	< 2.2e-16	
Decision trees	0.8721	0.5332	---	---	---	---
Random forest	0.9673	0.3585	---	---	---	---
Neural networks	0.9514	0.3099	---	---	---	---

The results show that very impressive performance is displayed by both the random forest and neural network models. While random forest achieves a slightly better R-squared score (0.967 vs 0.951), the neural network achieves a slightly more impressive RMSE score (0.309 vs 0.358). Deciding which to

give more importance to depends on the application. In this instance, we chose to maximise the R-squared score and chose the random forest model for our prediction application.

Interpretation

The Increase in Node Impurity plot is one way to understand the contribution of individual variables (or features) to the decision-making process of the Random Forest model.

In a decision tree (the building block of Random Forest), the impurity of a node is a measure of how mixed the target values are at that node. It's used to determine how well a split separates the data. When building a Random Forest, the algorithm considers different variables to make splits, and the IncNodePurity plot visualizes the increase in impurity caused by splits involving a specific variable. The IncNodePurity plot of the random forest model for the HMF regression is plotted in figure 1.

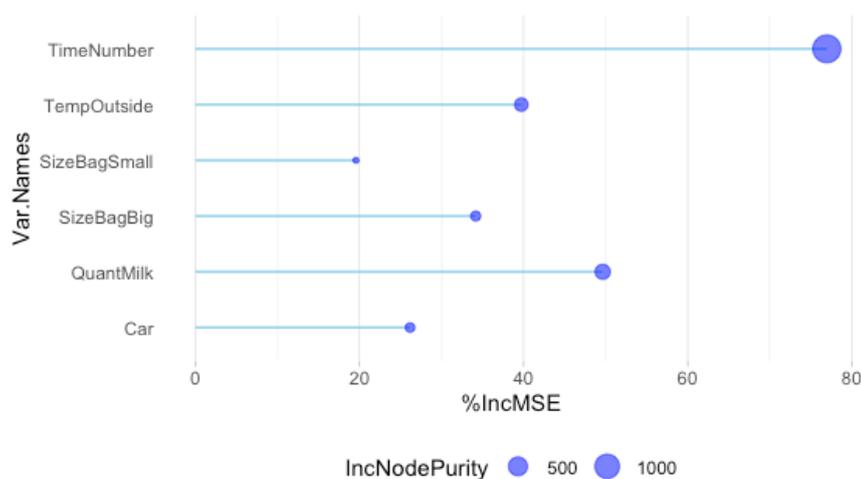


Figure 1: Increase in node impurity plot visualising the contribution of individual variables to the final model.

In the plot, the x-axis represents the increase in node impurity (measured by Mean Squared Error) caused by splits involving each variable and the y-axis represents the variables (features) used for splitting. Variables with higher values on the x-axis contribute more to the decision-making process and have a larger impact on the model's performance. Conversely, variables with lower values might not be contributing much and could potentially be candidates for removal or further investigation. From our plot, it is observed that the most important variable is time, followed by quantity of milk and outside temperature. If a large bag has been used for storage is the next most important, while vehicle type and small bag size has the least impact on the prediction.

Deployment

To simplify the use of the random forest model and empower end users to make valuable predictions, a web application was created (see figure 2). This application presents an intuitive user interface (UI) tailored for various end users, including delivery personnel and stakeholders. The UI application seamlessly interacts with a concealed machine learning model on the backend. Users input external journey parameters and activate the process by clicking the 'max trip duration & distance' estimation button. Upon clicking the button, the UI sends the entered data to the backend random forest model

for processing. The model uses the provided inputs for prediction and displays the returned estimation on the user screen. The UI's elegant simplicity shields users from the backend intricacies, promoting a user-friendly experience. Additionally, the application prioritises security by concealing the model and sensitive trip data.

The User Interface has been deployed using R Shiny, a web application framework for R programming language. The codebase includes an app.R script, which is responsible for creating the user interface and establishing the connection with the trained prediction model in the backend.

The application is accessible from the following URL:

<https://gautamsamriya.shinyapps.io/HMFAppPrediction/>

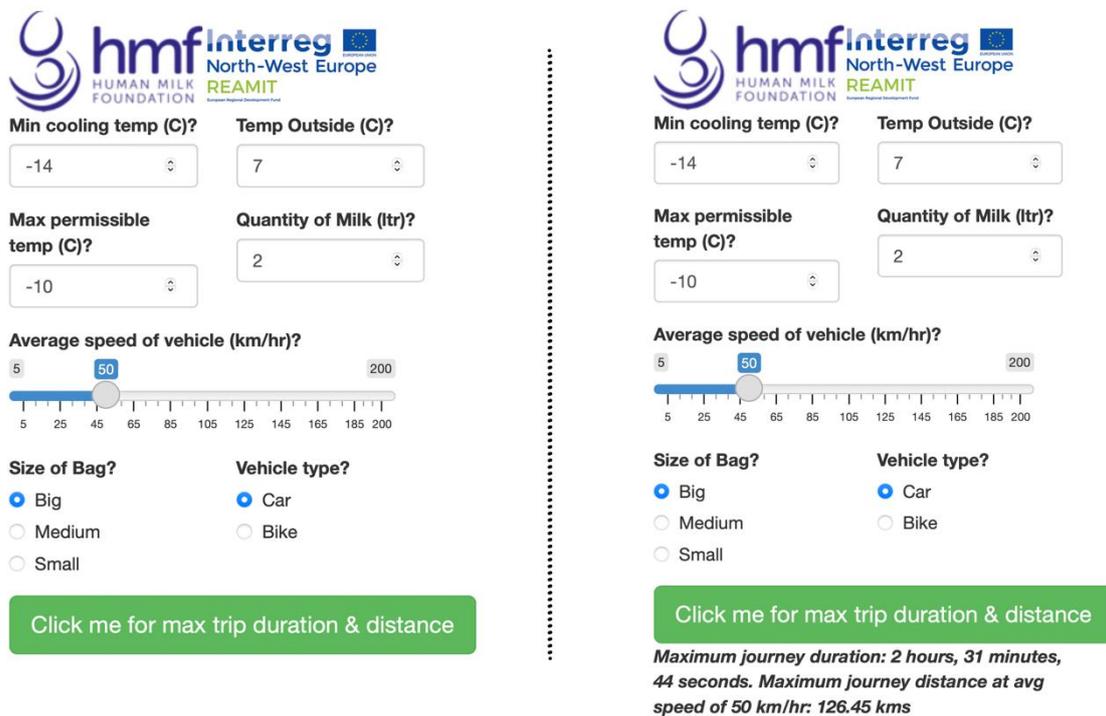


Figure 2: HMF maximum journey prediction app, deployed as a shiny.io application.

3. Yumchop

Preliminaries

Yumchop specialises in producing African flavoured frozen ready meals. They provide flavoursome and authentic food from around the world with an African twist that is frozen to retain its goodness and freshness and minimise waste. Yumchop is using locally sourced raw materials to prepare their ready-meal products. Most of the ingredients are supplied by local vendors, located locally from the production plant of Yumchop in Towcester, Northamptonshire, UK. Their tasty meals are distributed at institutions such as universities or hospitals through self-service automated vending machines. These unattended retailing kiosks have been fitted with an integrated microwave oven which enables them to warm the food upon purchase. However, Yumchop also delivers food to customers' homes through direct purchase at their website, enabling one-off purchases and monthly subscriptions that customers can customise to receive food at their preferred intervals. Moreover, they also supply directly to retailers and large organisations.

Equipment

- Sensors (x 10)
 - ELT Internal Antenna sensors (Elsys, Sweden)
 - DS18B20 temperature probes added to each ELT sensor.
 - 1 sensor was deployed per fridges/freezer
 - 10 sensors in total were used.

- Tektelic Kona Micro IoT Gateway (Tektelic, Canada)

Business provided labels

- None

Data Dictionary

Column	Data Type	Description
device_id	int	Numeric device ID, unique to the sensor
datetime_measure	ISO8601 date-time format, "yyyy-MM-dd'T'HH:mm:ss.SSS'Z'"	Datetime stamp of sensor recording
battery	float	Battery level during sensor recording
temperature	float	Temperature recording inside chamber
device_name	nvarchar(50)	Descriptive name of device. Should help identify where device is installed, but often left as device_id.

Analytics

Pre-processing

- Data cleaning and transformation using the Warp10 platform.

Analysis

In this pilot test, no specific information was provided regarding the specific analytical applications that would be beneficial for the data. However, despite identifying anomalies, none of them yielded significant interpretability. Therefore, our focus shifted towards designing tools to effectively describe the data.

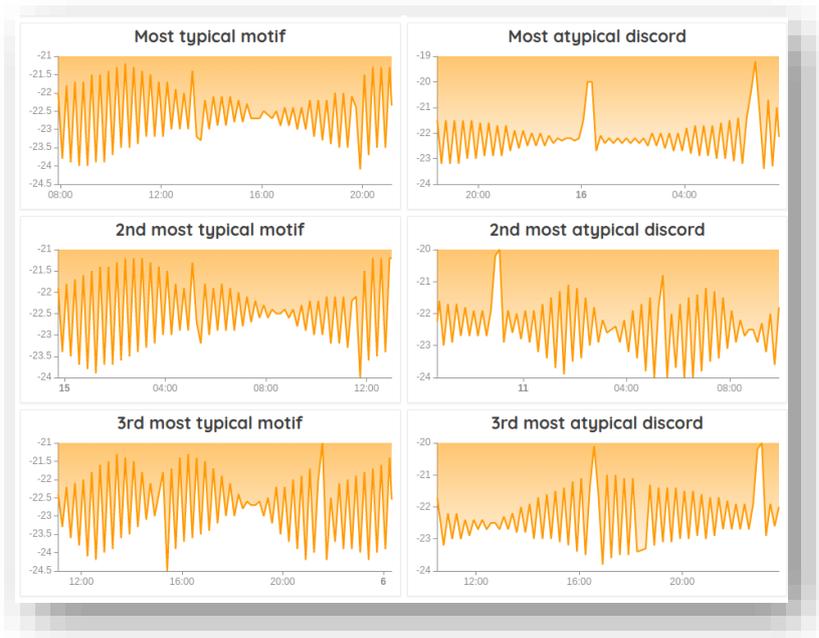
In the following figure, we have computed various features to provide a comprehensive description of the data. These features include the mean, standard deviation, minimum, maximum, median time between each record, potential seasonal cycles and associated weights.

```
In [4]: %%warpscript -v
        $data @remit/analysis/synthesis

In [5]: d = stack.pop()
        df = pd.DataFrame.from_dict(d, orient='index').transpose()
        df
```

Out[5]:

	classname	device_name	device_id	size	modes	mean	std	min	max	step_(s)	seasonality_1	val_1	seasonality_2	val_2
0	YUMCHOP:temperature	Vending Machine (066E)	A81758FFFE06066E	2003	48	-22.5	0.93	-24.6	-19.2	600.02	PT20M	0.69	PT19H10M	0.47
1	YUMCHOP:temperature	Zone D Fridge (066D)	A81758FFFE06066D	2009	79	-4.1	1.38	-5.5	11.0	599.99	PT1H30M	0.55	PT23H50M	0.32
2	YUMCHOP:temperature	Zone E fridge (0670)	A81758FFFE060670	2005	49	3.0	0.91	1.3	7.1	599.98	PT6H10M	0.4	PT6H10M	0.4
3	YUMCHOP:temperature	Zone B freezer 2 (0659)	A81758FFFE060659	2008	68	-30.0	1.51	-32.2	-25.5	600.03	PT146H5M	0.41	None	None
4	YUMCHOP:temperature	Zone E Freezer (065C)	A81758FFFE06065C	1996	60	-31.6	1.07	-32.8	-25.5	600.04	PT165H45M	0.36	PT120H25M	-0.01
5	YUMCHOP:temperature	Container cold room (0658)	A81758FFFE060658	1989	100	6.7	2.01	2.3	12.2	599.98	PT144H55M	0.25	PT94H50M	0.17
6	YUMCHOP:temperature	Zone D Coldroom Fridge (065B)	A81758FFFE06065B	2008	81	3.0	1.45	1.3	12.2	599.98	PT21H35M	0.46	PT21H35M	0.46
7	YUMCHOP:temperature	Zone D cold room freezer (05F7)	A81758FFFE0605F7	1976	70	-19.0	1.24	-22.0	-4.9	600.01	PT4H	0.45	PT4H	0.45
8	YUMCHOP:temperature	Zone B Freezer 1 (066F)	A81758FFFE06066F	2012	31	-23.6	0.66	-25.3	-22.0	600.02	PT30M	0.97	PT14H	0.46
9	YUMCHOP:temperature	A81758FFFE06065A	A81758FFFE06065A	2003	151	16.2	4.14	12.0	27.2	599.97	PT24H10M	0.2	PT126H30M	0.13
10	YUMCHOP:temperature	Zone B Standing freezer (066C)	A81758FFFE06066C	2011	105	-22.5	1.73	-26.1	-8.2	600.02	PT8H15M	0.57	PT8H15M	0.57

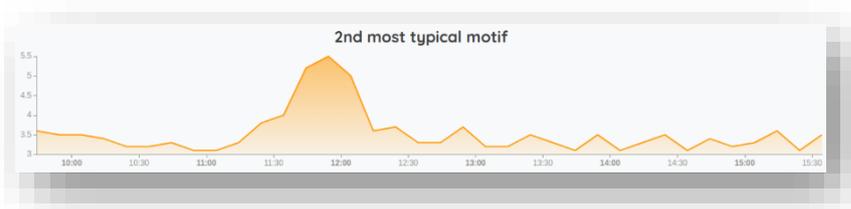


Next, we proceeded to explore patterns within the data using the matrix profile algorithm. This approach allowed us to identify recurring patterns as well as atypical occurrences. A resulting output of this analysis is depicted by the side screenshot, providing a visual representation of frequently occurring patterns and anomalies within the data.

These representations effectively capture the

data patterns for each sensor and highlight any deviations from those patterns, hence being useful for exploratory data analysis.

Interpretation



For sensors placed in refrigerators, the motif that typically appears as the second pattern in the matrix profile is likely

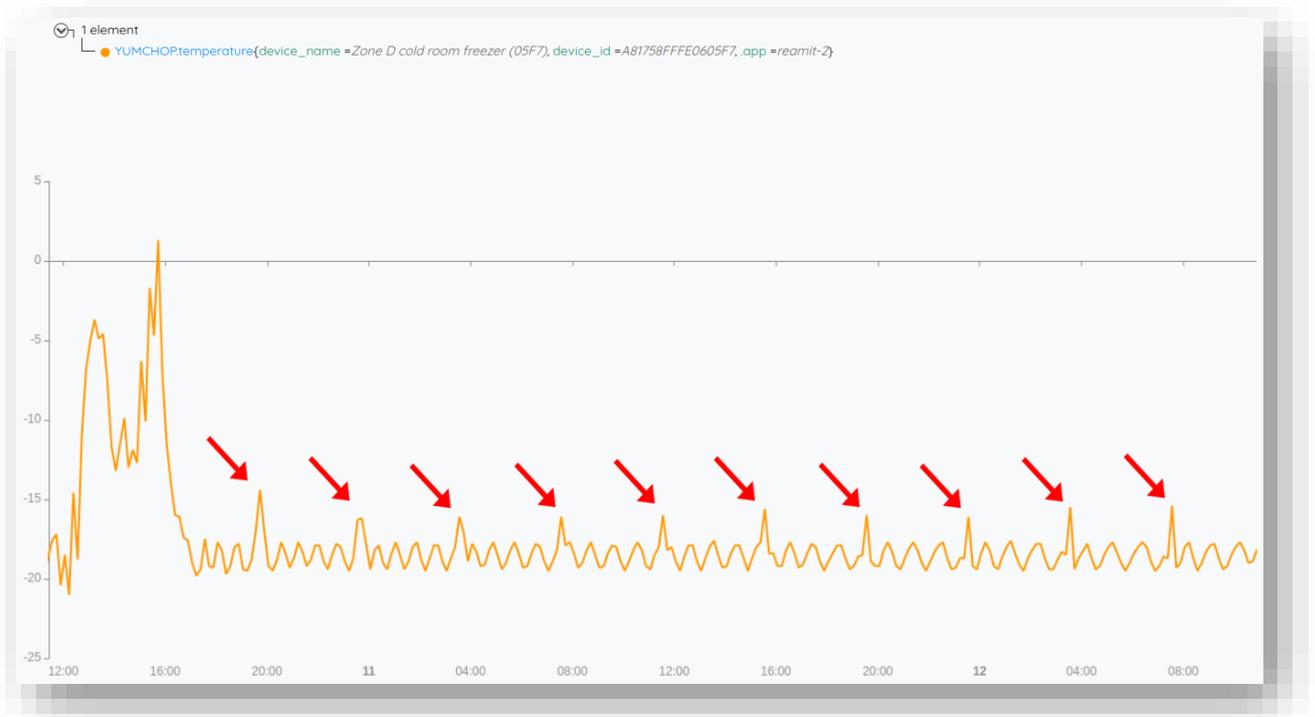
indicative of a defrost operation.

In the following figure, we can identify instances of this motif:



We can notice a consistent occurrence pattern of this motif, indicating that the defrost operation typically takes place every 6 hours under normal circumstances. In essence, this enables us to predict the timing of the next defrost operation unless any external event disrupts this pattern.

Likewise, when considering sensors placed in cold room freezers, we can observe consistent motifs that likely reflect the effects of defrost operations on temperature, as shown in the following figure.



In this particular context, these motifs tend to occur approximately every 4 hours.

Additionally, it is worth noting that occasionally, atypical discordant patterns like the one observed at the beginning can be found in the data. These patterns are likely a result of significant movements within the cold room and instances where the door remains open for a long time. The following figure shows an example of the alert dashboard.



4. Biogros

Preliminaries

Biogros is a well-established company that has been supplying high-quality organic food to customers in Luxembourg for over 25 years. They are known for their commitment to customer satisfaction, regularly adapting their product offerings to meet the evolving needs of their clientele. Biogros stocks a wide range of organic products, including well-known brands like Naturata, Rapunzel, and Lebensbaum, as well as products from smaller, lesser-known producers. Through their close collaboration with the Bio-Bauere-Genossenschaft Lëtzebuerg (BIOG), a cooperative of organic farmers, they also offer a diverse selection of regional organic products under the BIOG brand. With an extensive inventory of over 3,500 items in categories such as fruits and vegetables, dry goods, and dairy produce, customers can find a wide variety of organic options.



Biogros has established a commendable track record in sustainable transportation practices and minimizing waste packaging. To further their commitment to reducing food waste, they have partnered with REAMIT to explore potential solutions. As part of this collaboration, Biogros has installed humidity and temperature sensors in key areas such as their warehouse, cold stores, and refrigerated trucks. These sensors enable continuous monitoring and generate valuable data that the REAMIT team can analyze and leverage in their efforts to find innovative applications aimed at reducing food waste.

Following an in-depth data analysis, the REAMIT analytics team has discovered two promising avenues. One of the contributors, SenX, has developed an innovative anomaly scoring system using temperature sensor data. This system employs logical ranges and multiple temperature thresholds, drawing inspiration from the basic artificial neuron model by McCulloch & Pitt. With the inclusion of a multi-scaling layer that segregates information within each range, the score proves useful for anomaly detection, individual sensor forecasting and quality assessment. For example, by leveraging this approach, we identified stationary temperature sensors in cold storage facilities that exhibit seasonal forecast patterns. Biogros has provided a physical explanation for these patterns. We also

detected transportation periods where the quality of the refrigeration was not in the up to the desired standards.

Furthermore, Biogros is actively exploring opportunities to incorporate human feedback into their transportation operations. This feedback, specifically related to quality complaints regarding food delivery, can be instrumented to collectively refine the AI model based on each individual sensor's anomaly score neuron model. This would prove useful for them for root cause analysis across transportation routes and storage facilities.

Data Dictionary

Column	Data Type	Description
device_id	Int	Numeric device ID, unique to the sensor
datetime_measure	ISO8601 date-time format, "yyyy-MM-dd'T'HH:mm:ss.SSS'Z'"	Datetime stamp of sensor recording
Battery	Float	Battery level during sensor recording
temperature	Float	Temperature recording inside truck or cold store
Humidity	Float	Humidity recording inside truck or cold store
device_name	nvarchar(50)	Descriptive name of device. Should help identify where device is installed, but often left as device_id.

Technology stack selected

- Data is pulled from Whysor API and ingested into a Warp 10 platform
- Anomaly score is implemented using warpscript
- Temperature thresholds are stored in time series metadata and can be updated with an HTTP call
- Forecasting and anomaly detection is implemented in warpscript
- Visualisation is provided by Discovery and WarpStudio

Analytics

Pre-processing

- Date is parsed and formatted as Unix timestamp in microseconds.
- Time series model is applied. One time series per sensor
- A mask is applied on time series related to sensor in moving vehicle to only keep data points when the sensor is moving
- The series are split into sub series by detecting quiet periods of 4 hours with no data
- Synchronization is applied by affecting the last recorded value to every tick every 5 minute
- Linear interpolation fills missing values

Analysis

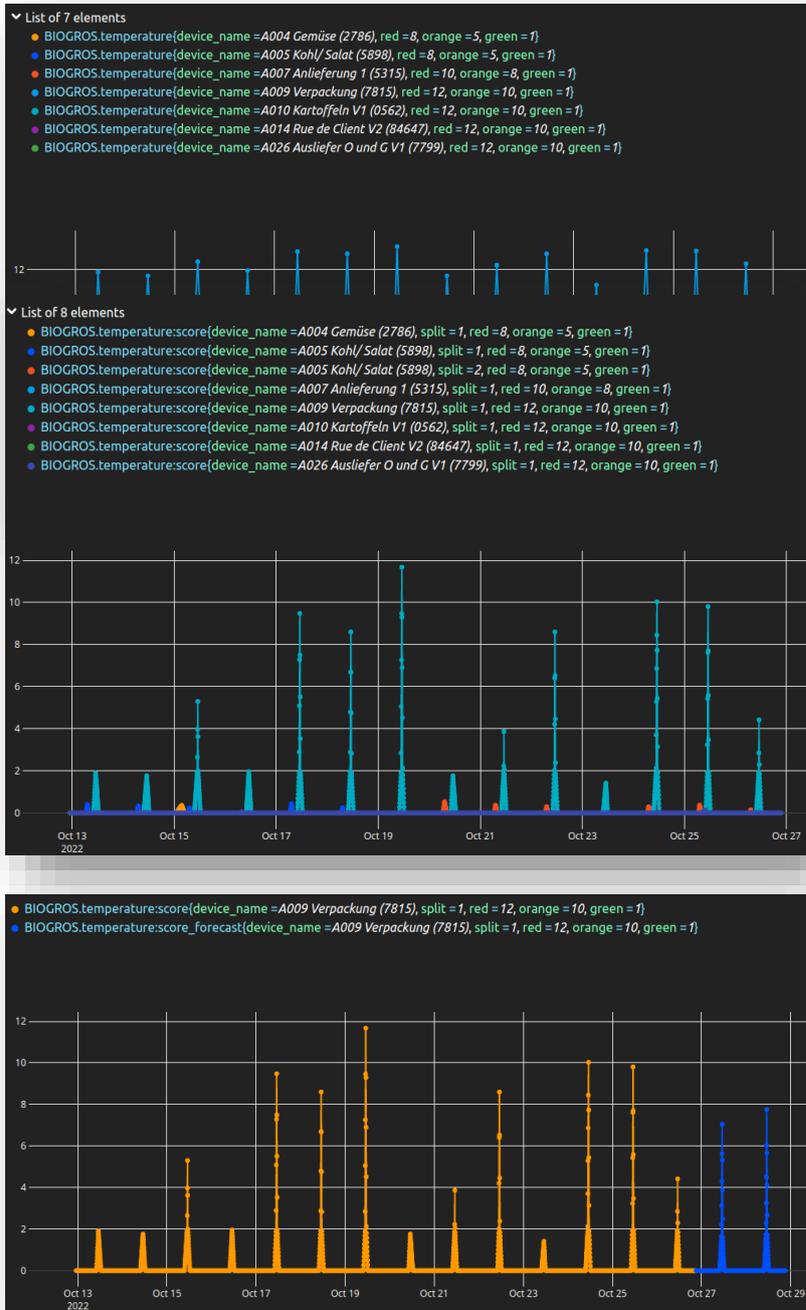
For each sensor, a set of two thresholds is provided. If the first one is crossed (orange threshold), then the cooling condition is not ideal but okay. If the second one is crossed (red threshold), then the conditions are bad. The formula of the anomaly score weighs in the temperature surplus when crossing these thresholds:

$$S(t) = \max(c_o \cdot \text{reLU}(t - t_o), c_r \cdot \text{reLU}(t - t_r))$$

where S is the anomaly score, t is the sensor's temperature reading, t_o is the orange threshold, t_r is the red threshold, and c_o , c_r are the associated penalty weights.

As a matter of fact, this formula can be interpreted as a basic two-neurons neural network, preceded by a multi-scaling operation and using max-pooling.

c_o and c_r are initialized at 1 and 10. Using consumer feedback, these values can be improved.



The first screenshot illustrates the plotting of data captured by temperature sensors in selected warehouses over a specific time period.

Afterward, the anomaly score is computed for each temperature time series by utilizing the logical ranges provided in the metadata of the time series. These scores effectively curate the time series data, enhancing the applicability of machine learning models.

An example of this is depicted in the provided screenshot, showcasing the plotting of the results generated by a forecasting model. Notably, this model demonstrates a clear seasonal forecasting pattern.



In this screenshot, an example of data from refrigerated trucks is plotted, specifically after applying a mask to retain only the data points corresponding to trips.

Presented here is an example of the anomaly score computed for a trip. A positive score indicates that the temperature exceeded the recommended threshold for transportation.

Interpretation

The analysis has revealed seasonal forecast patterns in the anomaly score for certain static sensors, indicating cyclical instances where the temperature exceeded the recommended threshold. These results were communicated to Biogros.

For moving sensors, a positive score suggests an inefficiency in freezing or cooling capabilities. As an illustration in the provided screenshot, the mean anomaly scores depict the quality of chilling conditions during product transportation. A score close to 0 indicates optimal chilling, meaning that the products were transported under suitable cooling conditions. If the score exceeds 1, it suggests that the chilling was acceptable but not ideal, indicating a need to slightly increase the cooling power. However, if the score surpasses 10, it usually indicates a lack of proper cooling throughout the transportation process.

Mean anomaly scores																	
device_name	trip 1	trip 2	trip 3	trip 4	trip 5	trip 6	trip 7	trip 8	trip 9	trip 10	trip 11	trip 12	trip 13	trip 14	trip 15	trip 16	trip 17
MAN KN 8715 (2424)	0.4	0.77															
MAN ZE 0001 (6084)	0.05	1.54	0.74	4.78	2.46	13.86	12.9	75.07	2.68	3.2	0.47	2.36	2.7	2.07	4.0	1.9	0.29
MB Antos SK 4508 (6407)	55.46	0.39	0.18	1.7	5.03	0.56	15.13	0.59	0.13	0.59	0.03	0.98					
MB Antos SK 4509 (8094)	0.01	0.51	0.52	0.15	18.02	45.08	55.73	0.6	0.01	0.12	0.52						
MB Ateco SK 4463 (6480)	0.01	1.16	0.45	1.92	13.84	0.73	1.58	1.88	0.14	2.9	2.07						
MB YH 5823 (5995)	2.17	1.84	1.92	0.72	0.0	0.7	0.02	0.14	0.49								
MB YH 5868 (3274)	0.0	0.1	1.52	0.87	0.53	0.15	12.57	0.12	0.01								
Sprinter YH 5827 (8681)	0.55	0.84	4.1	1.5	2.5	0.76	1.85	3.77	0.39	11.19							

Deployment

Using Warp10 discovery, we have successfully deployed 126 reporting dashboards. These dashboards consist of two types of reports. The first type is a comprehensive monthly quality report that provides detailed information about the transportation of a specific truck. The second type is a monthly overview report that gathers anomaly scores for all trips that took place within the month.

Below are examples of these interactive reports, which allow users to interact with the graphs using the mouse pointer.



5. Musgrave

Preliminaries

At Musgrave, last mile refrigerated delivery vehicles perform deliveries within the greater Belfast area. Presently, Musgrave have no means of temperature monitoring, meaning if a refrigeration unit has malfunctioned, they are unaware until the van makes its next stop. This is often too late and could result in hundreds of kilograms of spoiled stock. Instead, they wish to have an automated instant alerting system so that they know if the temperature in the rear of the van reaches above a defined threshold. Detailed below is a list of the equipment installed, the outline of the proposed analysis, and the data dictionary of all the variables in the complete dataset.

Sensors (x 5)

- Digital Matter eagle data logger
 - Internally records acceleration and GPS and reports a “trip” indication if these both update within the space of 5 minutes.
- DS18B20 Temperature probe
- T9603 T/RH probe

Business provided variables

- None

Proposed analysis

- Identify correct parameters for alerting logic based on observed temperature profile during monitored journeys. Identifying correct parameters will significantly reduce the number of false alerts indicating spoilage.

Data Dictionary

Column	Data Type	Description
device_id	Int	Numeric device ID, unique to the sensor
datetime_measure	ISO8601 date-time format, "yyyy-MM-dd'T'HH:mm:ss.SSS'Z'"	Datetime stamp of sensor recording
battery	Float	Battery level during sensor recording
temperature1	Float	Temperature 1 recording, in freezer of van
temperature2	Float	Temperature 2 recording, in fridge of van
humidity	Float	Humidity recording, in freezer of van
device_name	nvarchar(50)	Descriptive name of device. Should help identify where device is installed, but often left as device_id.
trip	char(1)	Indication if the van is being driven (1) or is stationary (0) when the recording was made.

Technology stack selected

Matlab 2021b was chosen as the software for analysing the last mile delivery data due to its popularity and strong reputation in the field of data analysis. Matlab is widely recognized for its advanced statistical analysis toolkits, which provide a comprehensive set of functions and algorithms for conducting various statistical analyses. Furthermore, Matlab is renowned for its advanced data visualization capabilities. It offers a wide range of plotting and graphing functions that allow researchers to visually explore and present the cooling profile data in a clear and meaningful way. The interactive nature of Matlab's data visualization tools makes it easy to customize plots, add annotations, and interact with the data, facilitating a deeper understanding of the underlying patterns and trends within the cooling profile data.

Analytics

Pre-processing

The sensors are configured to record temperature data every 5 minutes while in the delivery vehicle is in motion and every 12 hours when stationary. A data import and processing pipeline was therefore built in Matlab to identify deliveries from the Musgrave vans. Any recordings made while the van was stationary was discarded as the goal is to analyse the data which is recorded while the vans are performing deliveries. This prevents data recorded while the refrigeration unit is powered down being included in the analysis, which if left could result in an inaccurate cooldown period being identified. To ensure a real delivery is taking place, the trip status logged by the sensor needed to be true twice in a row. Although sensors were configured for 5-minute intervals, it was observed that sometimes two datapoints were uploaded at once, likely due to packet loss experienced using the cellular connection employed by the sensor for data upload. To ensure this did not affect the analysis, data windowing was performed using a spline interpolation, resulting in equally spaced 5-minute readings for each journey. The ISO8601 timestamp was converted to POSIX time (i.e. Unix / epoch time) and a time elapsed feature was computed for each journey.

Analysis

Analysis was performed on two instances of the Musgrave dataset. First, an analysis was carried out on data collected by 3 vans with sensors the sensors installed, recorded from April 2022 – November 2022 (180 days). Between the vans, 49839 datapoints were logged, representing 3154 deliveries. From this, we can calculate statistical properties of each journey. The average journey length during this period was 16.4 minutes per delivery, where the mean freezer temperature was -3.9°C and the mean fridge temperature was 9.05°C . Notice that these values are substantially larger than Musgraves desired threshold values of -18 and 5°C respectively. Even considering a cooldown period of 35 minutes, if alerting had been enabling during this period $\sim 81\%$ of journeys would have generated an alert.

The sensor location was moved, placing the probes at the 'air-on' location of the vans. This modification was thought to improve the accuracy of the readings.

The second instance of the Musgrave dataset was on sensor data collected from the repositioned sensor. Over a one-month period, from 21st Feb 2023 – 22nd March 2023 (29 days), 1 van had cumulatively recorded 5959 datapoints, representing 314 deliveries (10.8 deliveries per day). The average journey length in this second dataset was 25.75 minutes, while the mean freezer temperature was -9.2982 (previously this was -3.9°C) and the mean fridge temperature was 5.9037 (previously this was 9.05°C).

Interpretation

Analysis from the original dataset, presented in figure 3, showed that the best performance noted over a 6-month period was a 35-minute cooldown period to the desired -18°C . However, from speaking to Musgrave, there were no known problems. It was therefore deduced that the sensors must have been installed in the wrong location as 35 minutes was longer than the average delivery time of 16.4 minutes. Enabling the alerting platform would therefore have been undesirable until this problem was rectified, otherwise Musgrave would have received many false alerts suggesting that food was at risk of being spoiled when the freezer was in fact correctly functioning.

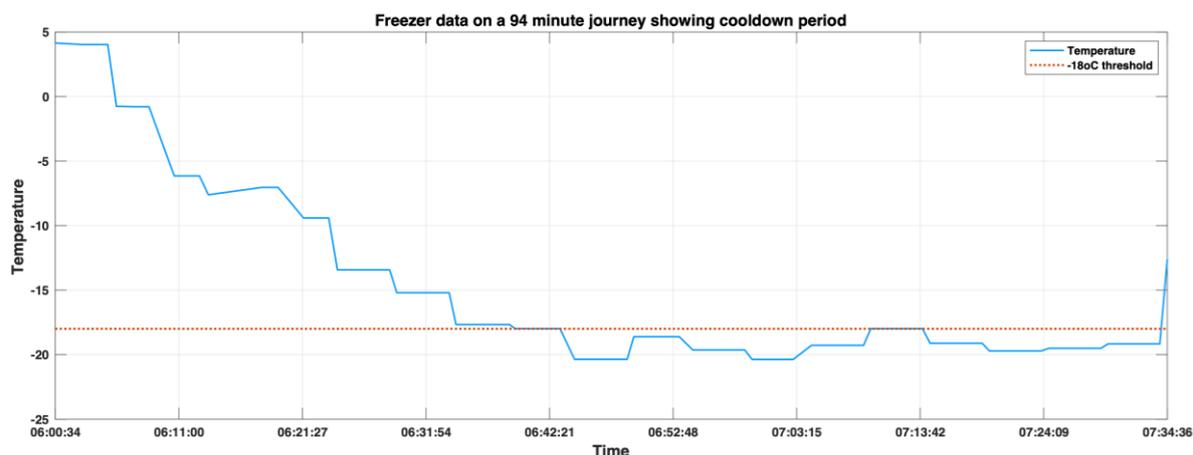


Figure 3: Cooldown period using the original dataset and sensor placement. The sensor takes 34 minutes to reach -18°C .

Analysis from the repositioned sensor, an example of which is presented in figure 4, shows that the new average cooldown period is 12.5 minutes. This is much more in line with what Musgrave thought the cooldown period should have been. Notice in the figure that on this particular 40-minute journey, the freezer had reached the correct temperature after only 10 minutes.

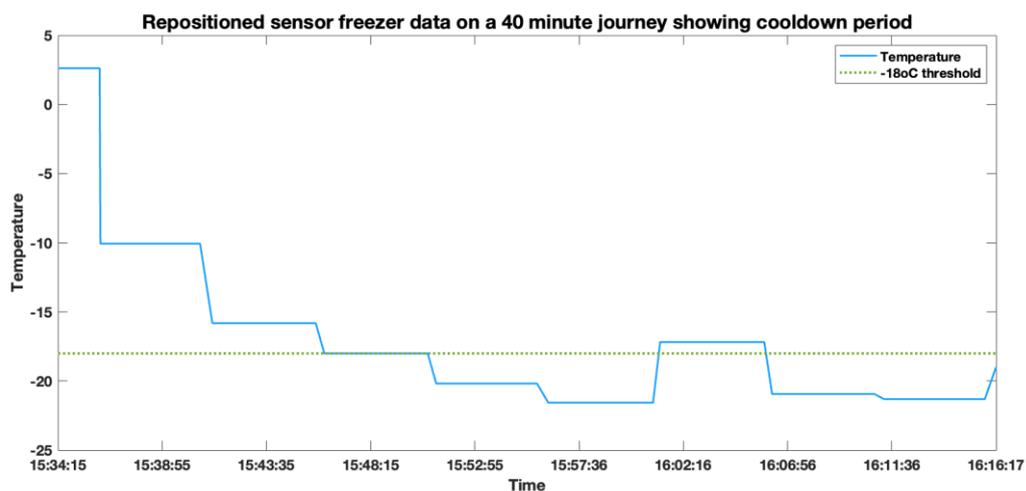


Figure 4: Cooldown period using the new dataset and sensor placement. The sensor takes 10 minutes to reach -18°C .

Deployment

One of the main REAMIT approaches to enhance efficiency and reduce food loss is the design and implementation of a suitable alerting system that can notify company representatives or staff members of potential anomalies in a timely manner. Following sensor installation, by continuously collecting data, the deployed sensors can identify potential issues or deviations that could lead to food spoilage by comparing each newly acquired environmental reading to a predefined safety threshold value. This way, real-time monitoring enables swift corrective actions, ensuring the quality and safety of food products.

In the context of this pilot test, and to best adapt to Musgrave's operations and procedures, the alerting system had to be such that it would not send alerts while the van was stationary or had not been running long enough to sufficiently cool down.

For the alerting algorithm, trip detection status is used as the first condition to ensure the vehicle is in motion. The analysis has shown that it takes approximately 12.5 minutes for a van to appropriately cool down. Therefore, the alerting was configured to ensure this time had elapsed. Since recordings are uploaded every 5 minutes, the trip condition should be true 3 times before progressing to the next stage of the alerting. Next, temperature threshold abuse is checked. To ensure detection anomalies are avoided, the temperature threshold should be passed twice before sending an alert. Figure 5 shows the alerting criteria configured on the Whysor dashboard. On the left, the configuration is shown for the sensor installed in the fridge, and on the right, the alerting configuration for the sensor installed in the freezer.

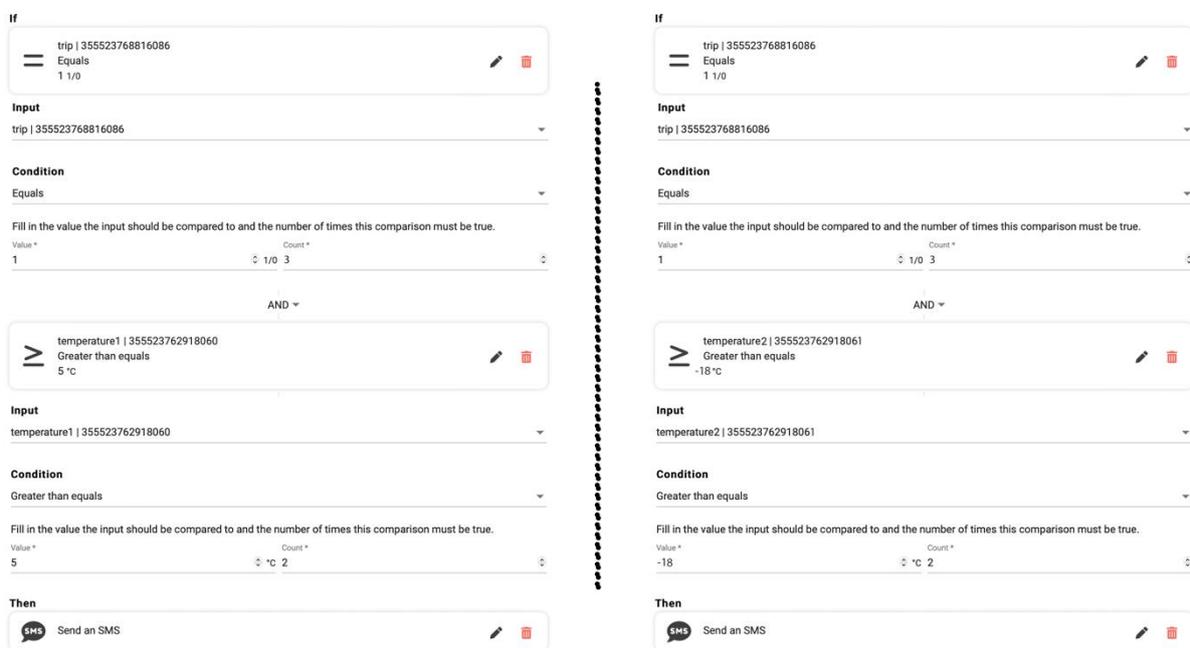


Figure 5: Configuration of the decision support technology on the Whysor platform for Musgrave.

An example of an SMS alert received based on exceeding the freezer threshold is shown in figure 6.

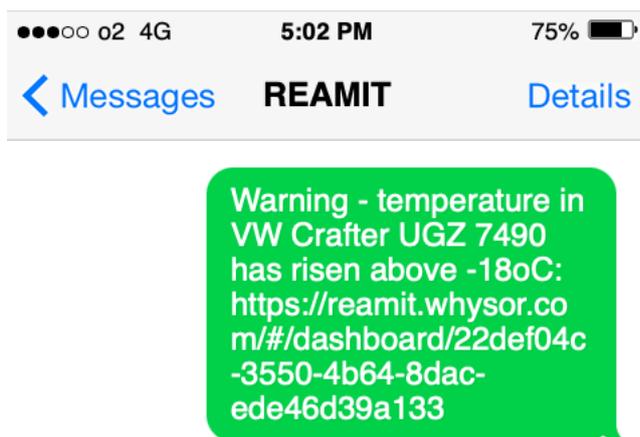


Figure 6: An example of the text alert received for end-users of the Musgrave decision support technology.

6. WD Meats & Burns Farm Meats

Preliminaries

Located in Coleraine, Northern Ireland, WD Meats have been supplying quality assured beef and innovative beef products to retail customers throughout UK, Europe, Africa and Asia for over 35 years. WD Meats select only the best local beef that Northern Ireland has to offer. Then they process and deliver it with the utmost care and attention to the animals, and to the highest standards that their customers demand.

At WD Meats, there was an opportunity for optimising the dry-ageing process. This is a 28-day cycle used for premium cuts of beef to both improve the tenderness of the beef and enhance the flavour, achieved by maintaining temperature and humidity in a sealed, refrigerated room. The ideal parameters for this room are still being explored, and so the REAMIT team installed sensors to the room to help map the conditions during the current ageing process and thus help identify more “ideal” parameters to reduce the weight percentage lost during the process, while avoiding the dark-face phenomenon. “Dark facing” meat forms when too much moisture is drawn from the hindquarter, which must be cut off (or trimmed) prior to sale. This meat is classified as food waste. Detailed below is a list of the equipment installed, the accompanying labels provided by WD Meats and Burns Meats for trim loss, the outline of the proposed analysis, and the data dictionary of all the variables in the complete dataset.

In total, 244 hindquarters were monitored during the dry-age process over the duration of 3 runs. Two runs were for 14 days, and one run was for 21 days. Each hindquarter has a start and end weight associated with it. The start weight is measured before the dry-ageing process begins, and the end weight is measured after the hindquarter is removed from the chamber. Therefore, this parameter represents the water weight lost during the dry age process. Additionally, the total trim loss average over the whole batch is provided. Therefore, there are 3 trim loss labels for the full dataset of 244 hindquarters.

Equipment

- Sensors (x 4)
 - Uرسالink UC-11 Internal Antenna sensors (Uرسالink, China)
 - The UC-11 consists of 2 built-in sensors: temperature, humidity
 - 4 sensors were deployed in the refrigeration chamber



- Multitech conduit mLinux IoT Gateway (Multitech, USA)

Business provided labels

- Hindquarter weights before and after dry ageing process, per hindquarter
- Trim loss weight % over the whole batch

Proposed Analysis

- Carcass/hindquarter weight loss during the dry ageing process (loss in water content).
- Trim loss (crust/surface of the meat that needs to be trimmed off) after the dry ageing process.

Data Dictionary

Column	Data Type	Description
device_id	Int	Numeric device ID, unique to the sensor
datetime_measure	ISO8601 date-time format, "yyyy-MM-dd'T'HH:mm:ss.SSS'Z'"	Datetime stamp of sensor recording
Battery	Float	Battery level during sensor recording
temperature	Float	Temperature recording inside chamber
Humidity	Float	Humidity recording inside chamber
device_name	nvarchar(50)	Descriptive name of device. Should help identify where device is installed, but often left as device_id.

Technology stack selected

Matlab 2021b was chosen as the software for analysing the trim loss data due to its popularity and strong reputation in the field of data analysis. Matlab is widely recognized for its advanced statistical analysis toolkits, which provide a comprehensive set of functions and algorithms for conducting various statistical analyses. Furthermore, Matlab is renowned for its advanced data visualization capabilities. It offers a wide range of plotting and graphing functions that allow researchers to visually explore and present the trim loss data in a clear and meaningful way. The interactive nature of Matlab's data visualization tools makes it easy to customize plots, add annotations, and interact with the data, facilitating a deeper understanding of the underlying patterns and trends within the trim loss data.

Analytics

Pre-processing

A data import and processing pipeline was built in Matlab to synchronise the sensor data recorded in the 4 zones of the dry-age chamber. This script processes input from the 4 sensor datafiles and converts the ISO8601 timestamp to POSIX time (i.e. Unix / epoch time), appending it to the dataframe. The popular spline interpolation technique is then used to synchronise sensor data within 1 minute.

Flexibility is built into the script through adjustable start and end date parameters, enabling users to extract the relevant sensor data corresponding to the dry-ageing period. Additionally, the script features a parameter for upsampling or downsampling the sensor data, catering to desired recording frequencies for model development.

Two distinct models were assessed during the experimentation phase:

Statistical Description Model:

In this approach, key statistical parameters – mean, maximum, minimum, and standard deviation – are computed for both temperature and humidity across the entirety of each of the dry-age periods for each of the four sensors. Each of these features are then appended to the start and end weight of the hindquarter dataframe for the corresponding dry-age experiment and zone in the chamber where the hindquarter was stored. 4 zones are defined, corresponding to the location where the sensor was installed. Zone 1 is back left, 2 is front left, 3 is back right, and 4 is front right. Note that the dry age chamber had one refrigeration unit, located at the back of the chamber. Hence, the back left and right zones are those closest to the fridge unit.

Full Sensor Dataset Model:

This model downsamples the temperature and humidity data to a one-hour recording frequency and uses this more detailed description of the dry-age process as the predictor feature set. A new feature representing hindquarter weight loss per hour is introduced. This weight loss feature serves as a proxy for time elapsed. It is generated using linear spacing techniques between the initial and final hindquarter weights. The downsampled sensor dataset size informs the spacing parameter for generating this feature.

Analysis

Water-weight loss on hindquarters

The initial question to examine was the relationship between temperature in the dry-age chamber and the weight-loss during the dry-age process after the first run of 14 days was completed. Before and after weights were collected on this first load of hindquarters, and the temperature during the dry-ageing process extracted. Figure 7 shows the temperature mapping over the 14-day dry-age process, which illustrates how the temperature is clearly colder at the front of the dry-age chamber closer to the refrigeration unit compared to at the rear of the chamber, furthest away from the dry-age chamber. Closest to the refrigerator, the mean temperature was -0.7954 , std 0.3274 over the 14-day dry age period. In comparison, furthest from the refrigerator, the mean temperature was 0.0164 , std 0.3459 over the same time period.

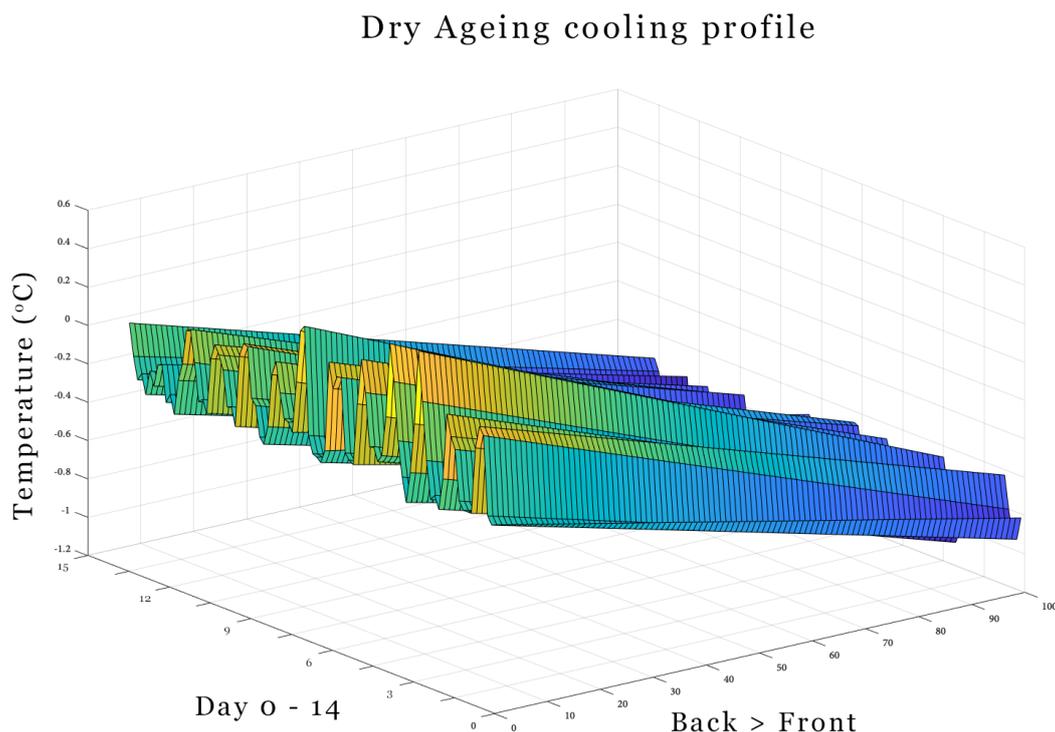


Figure 7: Surface plot showing the difference in temperatures between the front and back of the dry-ageing chamber.

The next step was to compare the weights of the hindquarters located in the front half of the chamber with those in the rear half. After the 14-day dry ageing was complete, the hindquarters in the front, cooler half of the chamber (closest to the fridge unit) had lost 3.87% total weight, while the hindquarters in the rear, warmer half had lost 4.33% total weight. An ANOVA (Analysis of Variance) test, a statistical test used to determine whether there are any significant differences between the means of two or more groups, was performed to check for statistical significance in weights between the hindquarter’s storage areas. The result of the ANOVA test was $P < 0.05$, telling us that the differences in weight loss between the front and rear were statistically significant.

From this initial analysis, we learned that the cooler the temperature in the dry-age chamber, the less weight loss occurs on the hindquarters.

After establishing that there was a statistical difference between the weights and the temperature parameters during the dry age process, an additional 2 runs of the dry-age process were undertaken to expand the size of the dataset and provide more statistical power to the findings. Ideally, edge case data would have been collected during this phase to allow a machine learning regression model to learn further insights and provide better estimates on the dry-age process, but due to the monetary costs this would have had to WD Meats, they did not want to adjust the parameters of the fridge unit significantly.

Regression modelling was undertaken on the full dataset testing the two modelling approaches outlined in the pre-processing step. Namely, these were a statistical description model, where the temperature and humidity data were summarised into a single feature vector, resulting in 244 observations. These observations consisted of the start and end weight of the hindquarter, the zone it was stored, and the statistical description of the temperature and humidity data over the dry age period for that zone. Secondly, the full sensor data model was tested, which used the sensor data downsampled to 1h recordings with a new weight loss per recording feature engineered, resulting in 99412 observations.

10-fold cross validation was used to validate the models. This rigorous validation method facilitates the assessment of model performance across distinct subsets of the data, fostering a comprehensive understanding of the model's predictive capabilities.

The results from the statistical feature regression are presented in table 2.

Table 2: Results from the regression models using the statistical description dataset. The best performing model is a linear regression and uses temperature, sensor, and zone information to make its prediction on weight loss, achieving an R-squared of 0.99 and RMSE of 1.0534 (kg).

Statistical features	Model	Dependent Variable	Independent variables	R-Squared	RMSE (kg)
Temperature only	Linear regression	End weight	Start weight, mean, min, max, std temperature	0.98	1.483
	Fine Tree			0.97	1.5267
Temperature and zone	Linear regression	End weight	Start weight, mean, min, max, std temperature, zone	0.98	1.444
	Fine Tree			0.97	1.5274
Humidity only	Linear regression	End weight	Start weight, mean, min, max, std humidity	0.96	1.9348
	Fine Tree			0.96	1.8545
Humidity and Zone	Linear regression	End weight	Start weight, mean, min, max, std humidity, zone	0.96	1.9668
	Fine Tree			0.96	1.8551

Temperature and humidity	Linear regression	End weight	Start weight, mean, min, max, std temperature and humidity	0.98	1.2548
	Fine Tree			0.98	1.4252
Temperature, humidity, and zone	Linear regression	End weight	Start weight, mean, min, max, std temperature and humidity, zone	0.99	1.0534
	Fine Tree			0.98	1.426

From the table, it is observed that the best performing models are the ones which uses all the features (independent variables) in the dataset, i.e. the statistical description of the temperature and humidity parameters and the zonal information. Of the linear regression and fine tree models, the linear regression performs best achieving an R-squared of 0.99 and an RMSE of 1.0534. The Predicted vs Actual plot of the linear regression is show in figure 8.

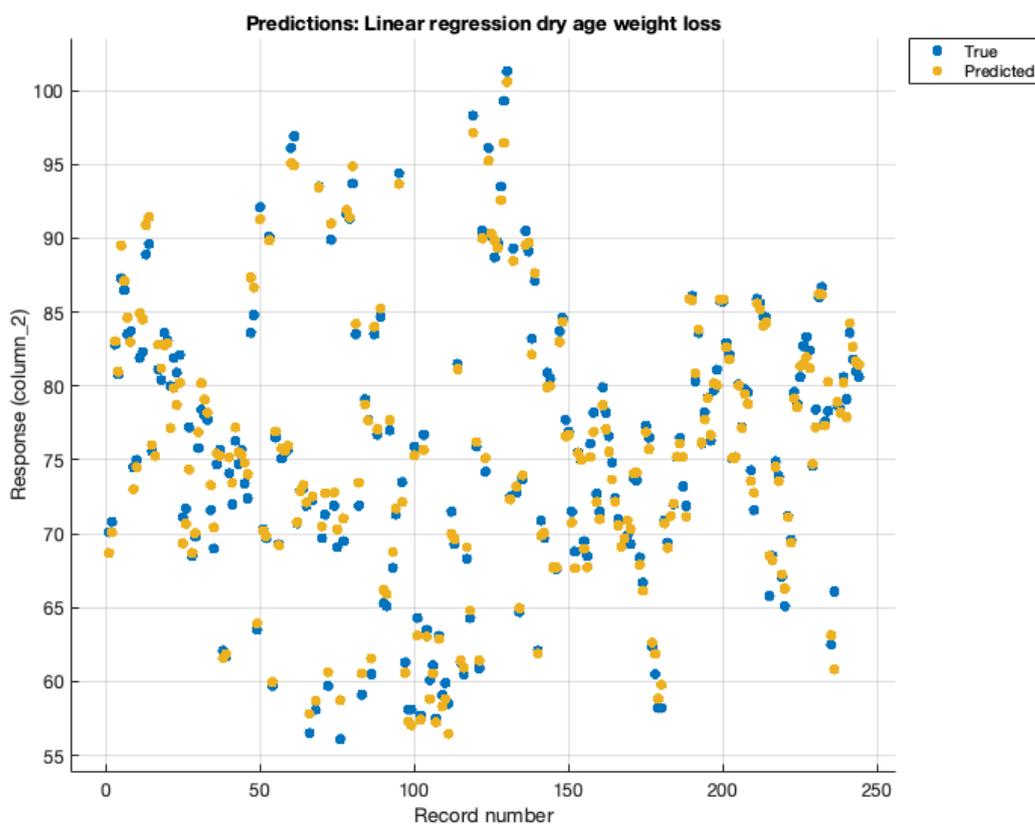


Figure 8: Predicted vs Actual plot of the dry age weight loss.

Based on the findings from the statistical prediction model, the best performing models combined both the temperature and humidity data to make the prediction on end weight. To this end, for the testing of the second model based on the full dataset downsampled to 1h observations, models were tested using the combined temperature and humidity dataset. The results from the tested models are presented in table 3. From the table, it is observed that the best performing model is the fine tree on the temperature, humidity, and zone information, achieving an R-squared of 0.99 and an RMSE of 1.0726 kg. These results are slightly worse than the best performing statistical model dataset, which had achieved an R-squared of 0.99 and an RMSE of 1.0534 kg using linear regression. As well as this,

the training time for the full feature models had increased since the number of observations had increased from 244 to 99412. Therefore, it was decided to take the statistical feature linear regression model forward and deploy that for the end weight predictor application.

Table 3: Results from the regression models using the full feature description dataset. The best performing model is a fine tree and uses temperature, sensor, and zone information to make its prediction on weight loss, achieving an R-squared of 0.99 and RMSE of 1.0726 (kg).

Statistical features	Model	Dependent Variable	Independent variables	R-Squared	RMSE (kg)
Temperature and humidity	Linear regression	Current weight	Start weight, temperature, humidity, time	0.97	1.6956
	Fine Tree			0.99	1.1399
Temperature, humidity, and zone	Linear regression	Current weight	Start weight, temperature, humidity, zone, time	0.97	1.6899
	Fine Tree			0.99	1.0726

Interpretation

Using the statistical feature set and linear regression, a robust model for estimating the water weight loss during the dry age process has been produced. Using this model, WD Meats can estimate how much their total yield will be on a full dry-age load, and by testing parameters, they have the potential to optimise the temperature and humidity settings and the distribution of hindquarters within the chamber. However, for the model to gain further statistical significance and generalise well to new data, a considerably larger dataset will be necessary. Adjustments to the dry-age process parameters are needed to observe the impact of varying temperature and humidity on hindquarter weight loss.

Deployment

A Matlab GUI application was developed which employs the temperature, humidity, and zonal statistical linear regression model for end weight predictions. Rather than relying on the user to provide the statistical depiction of temperature and humidity parameters, a practical approach was taken to allow the user to simply provide the temperature and humidity setting of the fridge to the model. In order for this approach to work, six regression models were daisy-chained to obtain estimated of the statistical features prior to making the ultimate end weight prediction. In total, a sequence of seven regression models were utilised.

The GUI application has text boxes for the user to provide the hindquarter's starting weight, the corresponding zone within the dry-age chamber, and the temperature/humidity settings configured for the chiller. Upon pressing the estimate end weight button, a linear regression model predicts the maximum temperature based on these parameters. Subsequently, this estimated maximum temperature, combined with the original parameters, serves as input for another regression model that predicts the minimum temperature, and so forth.

Upon completion of the six regression models' predictions for the statistical features (maximum, minimum, and standard deviation of temperature and humidity), derived from the provided temperature and humidity parameters, the final step involves utilising the temperature, humidity, and zonal statistical linear regression model (R-squared = 0.99, RMSE = 1.05kg). This model employs both the user-inputted data and the estimated features to yield the ultimate weight prediction. The result from the prediction appears in the estimated final weight text box. An example of the GUI being used to estimate the end-weight of a hindquarter weighing 80kg and stored in zone 3 (back right, closest to the chiller) is show in figure 9.



Figure 9: Hindquarter end-weight prediction application produced for WD Meats.

Additional analysis

Dark-facing trim loss

While optimising water-weight loss during the dry-ageing process was one goal of WD Meats, another question was how dark-facing trim loss could be minimised. As mentioned, “dark facing” meat forms when too much moisture is drawn from the hindquarter, which must be cut off (or trimmed) prior to sale. This meat is classified as food waste. Reducing the amount of dark facing on the beef was therefore of key interest to the REAMIT team as this was considered tangible food waste. Unfortunately, however, over the three trials conducted at WD Meats, only average dark-facing loss was calculated per batch of dry aged meat. This means that although we had 244 hindquarters monitored, we only had 3 trim loss values.

However, an exploratory analysis was carried out to see if, with the parameters collected, a meaningful model could be produced in the future to help predict and minimise the trim-loss of beef. The average batch trim loss value was appended to the statistical feature dataset, since this was the one which had performed best in the water-weight loss regression model. Therefore, the independent variables were starting weight, end weight, zone, mean, max, min, std temperature and humidity and the dependent variable was % dark facing trim loss. Linear regression modelling was applied to the dataset. Figure 10 shows the Predicted vs Actual plot of the dark facing trim loss. The model achieved an R-squared of 0.94 and an RMSE of 0.23%.

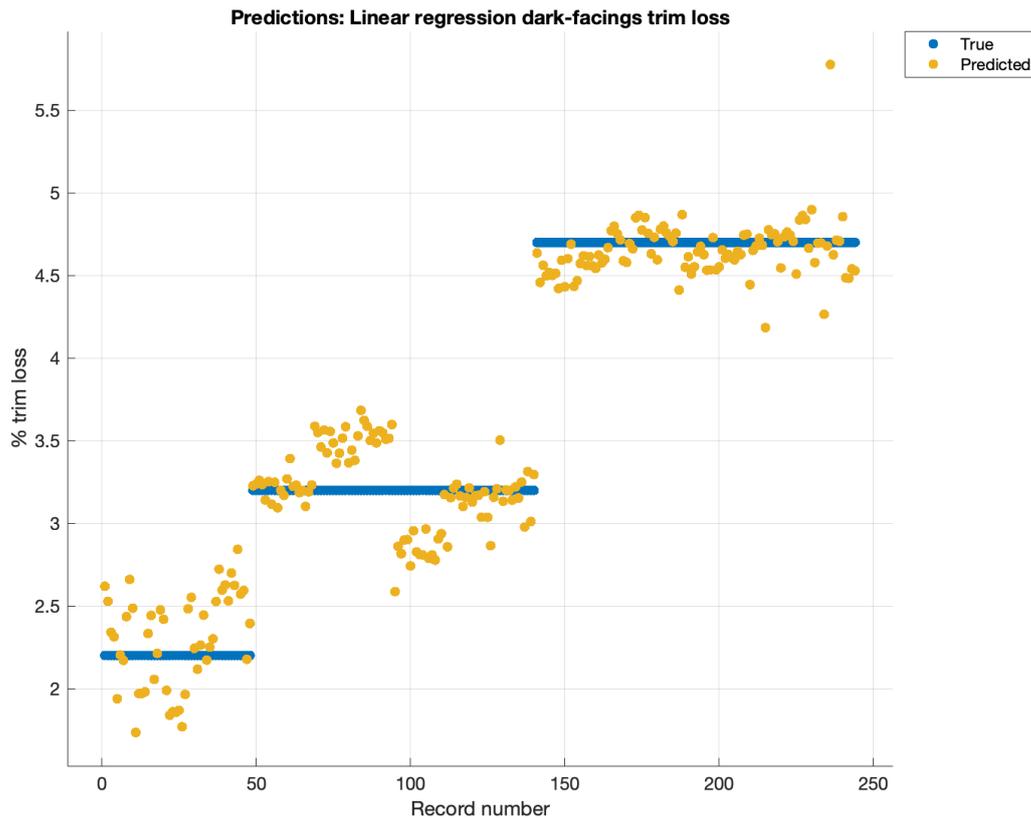


Figure 10: Predicted vs Actual plot of the dark facing trim loss.

This model shows promising potential going forward as a useful tool to estimate the dark-facing trim loss and serves as a starting point to optimise the other parameters so that both water weight loss and dark-facing trim loss are minimised. However, for the model to gain statistical significance and generalise well to new data, a considerably larger dataset is necessary. In addition to collecting dark facing data for each hindquarter rather than per batch, adjustments to other dry-age process parameters are needed to observe the impact of varying temperature and humidity on trim loss yield. This avenue remains a future endeavour for WD Meats to explore.

Burns Farm Meats

Pilot testing of a similar nature was also conducted with the pilot test company Burns Farm Meats. Located in north Sligo, Ireland, Burns Farm Meat Ltd. is a small-sized, family-owned company with main activities including farming, operation of an abattoir, and processing of organic meats. As part of these activities, being firmly committed to animal welfare and providing meat of the highest quality, they run a dry-ageing process to deliver tender cut meat of their own locally raised, fed and cared for animals.

This pilot test mainly consisted in the monitoring of environmental parameters - temperature and humidity - in the two refrigeration chambers that they use to carry out dry ageing of beef. To this company, it was important to obtain a system capable of monitoring environmental parameters in their refrigeration chambers that would ultimately provide them with a higher understanding and the means to enhance the quality of meat and reduce its loss.

Equipment

- Sensors (x 10)
 - ELT-2 Internal Antenna sensors (Elsys, Sweden)
 - The ELT-2 consists of 4 built-in sensors including the two necessary parameters for the pilot test: temperature and humidity
 - 6 sensors were deployed in the company's larger refrigeration chamber
 - 4 sensors were deployed in the company's smaller refrigeration chamber



Figure 11: showing the location of 6 sensors in the company's larger refrigeration chamber.

Alerting logic

As mentioned in the presentation of Musgrave's decision support technology, one of the main REAMIT approaches to enhance efficiency and reduce food loss is the design and implementation of a suitable alerting system that can notify company representatives or staff members of potential anomalies in a timely manner. Following sensor installation, by continuously collecting data, the deployed sensors can identify potential issues or deviations that could lead to food spoilage by comparing each newly acquired environmental reading to a predefined safety threshold value. This way, real-time monitoring enables swift corrective actions, ensuring the quality and safety of food products.

In the context of this pilot test, and to best adapt to Burns Farm Meats operations and procedures, the alerting system had to be such that it would not send alerts repeatedly at times when they were aware that the chambers were not at their adequate temperature, e.g., they were carrying out cleaning, loading or unloading procedures and the door of the chamber would remain open, leading to an increase of temperature inside. In some instances, it could also be that the refrigerator had been turned off as the chamber was not in use – this mainly concerned the smaller chamber.

For this reason, the REAMIT team decided to install door contact switch sensors on the doors with the purpose of refining the logic of the alerting system.

The first of the two alerts that was pilot tested was designed to alert of anomalies while the door of the chamber was closed (figure 12).

☰ Edit rule - 1 - Large chill: temperature alert while the door is closed
✕

General

Name *

1 - Large chill: temperature alert while the door is closed

Description

This alert is aimed at signalling any potential temperature anomaly in the chamber while the door is closed.

Active

If

Temperature | A81758FFFE05B7A0

> Greater than

5 °C

AND ▾

Temperature | A81758FFFE05B79F

> Greater than

5 °C

AND ▾

Temperature | A81758FFFE05B77E

> Greater than

5 °C

AND ▾

digitalInput | A81758FFFE05B7A0

= Equals

1 1/0

+

Then

Send an SMS

Figure 12: Alerting logic for the large chill at Burns Farm Meats when the door is closed.

In this case, the temperature threshold selected was 5°C and it had to be recorded consecutively 6 times by 3 different sensors. By implementing an “AND” logic, the team would avoid alerting if the selected sensor was malfunctioning. While that could also lead to not alerting

A second alert was also designed to signal potential issues while the door remained open (figure 13).

☰ Edit rule - 2 - Large chill: temperature alert while the door is open ✕

General

Name *

2 - Large chill: temperature alert while the door is open

Description

This alert is aimed at signalling any potential temperature anomaly in the chamber while the door is open. The ten

Active

If

Temperature | A81758FFFE05B7A0
> Greater than
7 °C

AND ▾

Temperature | A81758FFFE05B79F
> Greater than
7 °C

AND ▾

Temperature | A81758FFFE05B77E
> Greater than
7 °C

AND ▾

digitalInput | A81758FFFE05B7A0
= Equals
0 1/0

+

Then

Send an SMS

Cancel Save

Figure 13: Alerting logic for the large chill at Burns Farm Meats when the door is open.

An example of an SMS alert received based on exceeding a user-defined temperature threshold is shown in figure 14.



Figure 14: An example of the text alerting received at Burns Farm Meats.

Notice that the text message, as well as providing specific details about the sensor location and the rule that has been breached, links to the dashboard so the user can easily click the link and view further details about the status of the sensor. This is visible in figure 15.

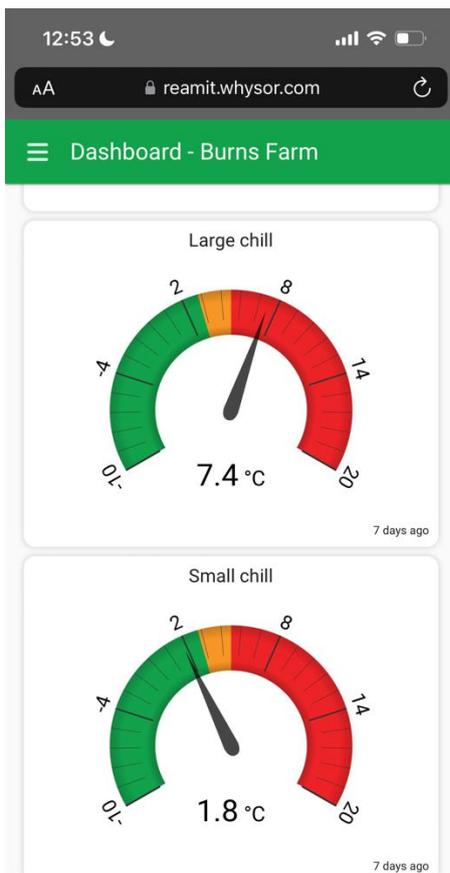


Figure 15: Upon clicking the link embedded in the text message, the Burns Farm Meats dashboard opens showing an overview of the current sensor status.

7. Picnic

Preliminaries

Picnic is a pure online last-mile grocery delivery company. It offers its users to find and order grocery items through their mobile phone. The products that it markets are vegetables, fruits, meat, fish, sweets, snacks, drinks, dairy, bread, but also non-food items. The company has presence in the Netherlands, Germany and France. They strive to provide affordable and sustainable service to customers using cutting-edge technology, efficient planning, and a fleet of electric vehicles. A typical vehicle on grocery delivery would have many food boxes and each food box would have grocery for delivery to customers based on their orders. Chilled goods are stored in some boxes, goods that do not require cooling are stored in ambient temperature, while some other boxes store frozen food below prespecified temperature thresholds. To maintain the food quality inside delivery vans, it is necessary that temperature inside the cooling boxes is always maintained below a certain threshold level throughout the journey. The company uses ice packs to maintain the temperature inside the cooling boxes.

For efficient operation, the company needed an ability to continuously monitor the temperature in each cooling box, make sure that the temperature inside is within an accepted pre-defined threshold, and check whether the capability of the process of cooling is acceptable. The company also need optimization of ice-packs usage in the cooling process of its delivery vehicles.

Equipment

- Sensors (x 10)
EMS Internal sensors (Elsys, Sweden) (Humidity, Temperature, Acceleration)
 - 1 sensor was deployed per crates, each box had 1 sensor installed.
 - LoRa coverage is provided by the netherlands KPN country-wide network.
 - 10 sensors in total were used.

Business provided labels

Optimizing usage of cooling ice-packs
Ensuring high quality food supply(Reducing food wastage/Reducing customer complaints)

Business defined variables

- acceptable temperature range / temperature threshold

Proposed alerting logic

- If temperature exceeds threshold send alert (SMS/email)
- If there are too many shocks or disorientation of crates
- If cooling process is not in statistical control or lacking capability

Proposed Analysis

Statistical process control, and process capability analysis.
Alerting in case of temperature crossing acceptable threshold.
Acceleration along x,y,z axes based orientation analysis of boxes/crates.
Optimizing usage of cooling ice-packs.

Data Dictionary

Column	Data Type	Description
--------	-----------	-------------

device_id	int	Numeric device ID, unique to the sensor
datetime_measure	ISO8601 date-time format, "yyyy-MM-dd'T'HH:mm:ss.SSS'Z'"	Datetime stamp of sensor recording
battery	float	Battery level during sensor recording
temperature	float	Temperature recording inside food crate/box
device_name	nvarchar(50)	Descriptive name of device. Should help identify where device is installed, but often left as device_id.
Motion_acceleration	int	The number of movement events during the measurement in a scale of 0 to 255
acceleration_x	Int	The forces on the X axis on this measurement in a scale of -127 to 127 (-2G to 2G, so 63 == 1G). This can be used to determine how the sensor has been placed, or if it's been handled roughly or moved
acceleration_y	Int	The forces on the Y axis on this measurement in a scale of -127 to 127 (-2G to 2G, so 63 == 1G).
acceleration_z	Int	The forces on the Z axis on this measurement in a scale of -127 to 127 (-2G to 2G, so 63 == 1G).
Humidity	Int	The humidity in %, in food crates/box.

Technology stack selected

The SPC analysis and PCA were developed using R, a popular open-source programming language known for its extensive capabilities in statistical analysis and machine learning. R provides a wide range of packages and libraries specifically designed for statistical modelling, making it a suitable choice for SPC and PCA. R Studio, Microsoft Excel and Python were used for the overall analysis.

Analytics

Pre-processing

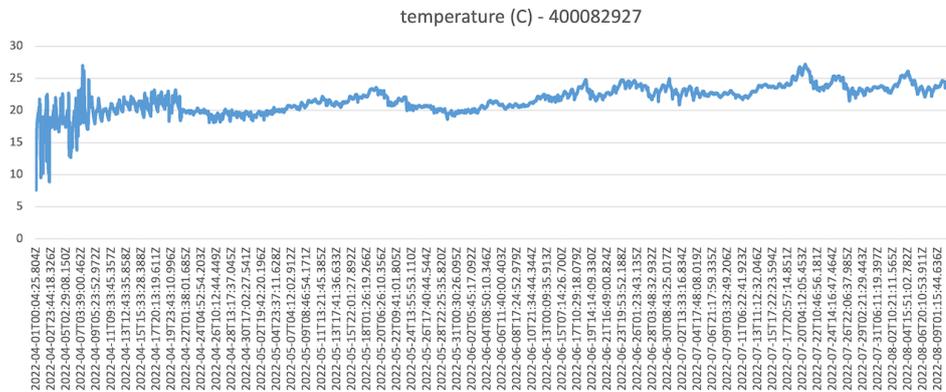
Data from temperature sensors was recorded at 10-minute intervals. Upon initial examination of the temperature data spanning from April 2022 to August 2022, noticeable temperature fluctuations were evident in April, which gradually diminished as the timeline progressed towards August 2022. To comprehensively explore this trend, Statistical Process Control analysis was employed on the sensor data during this period.

Given that the data pertained to measurable values, X-bar charts, Range Charts, and Process Capability Charts were constructed for the sensors. To conduct the analysis, a random selection of 5 temperature sensor values per hour was made for the initial consecutive days of both April 2022 and August 2022. Subsequently, the mean and range were calculated for the hourly temperature data across 24 hours, followed by the average mean and range over this 24-hour span. These charts were generated for the first 9 days of both April 2022 and August 2022.

The collected data underwent pre-processing to convert it from 10-minute intervals to hourly sampled data. Essential statistical metrics such as mean and standard deviation were calculated as part of this process.

Analysis

The variation of temperature from a single cooling box for the period from 1st April 2022 until 9th August 2022 was observed. It was noted that there were much larger fluctuations in temperature in the month of April 2022 compared to August 2022. This finding indicates the variability in temperature has reduced significantly in August compared to April. To look at the observations in more detail we decided to perform Statistical Process Control and Process Capability analysis based on the temperature data inside the vehicle.



Statistical Process Control

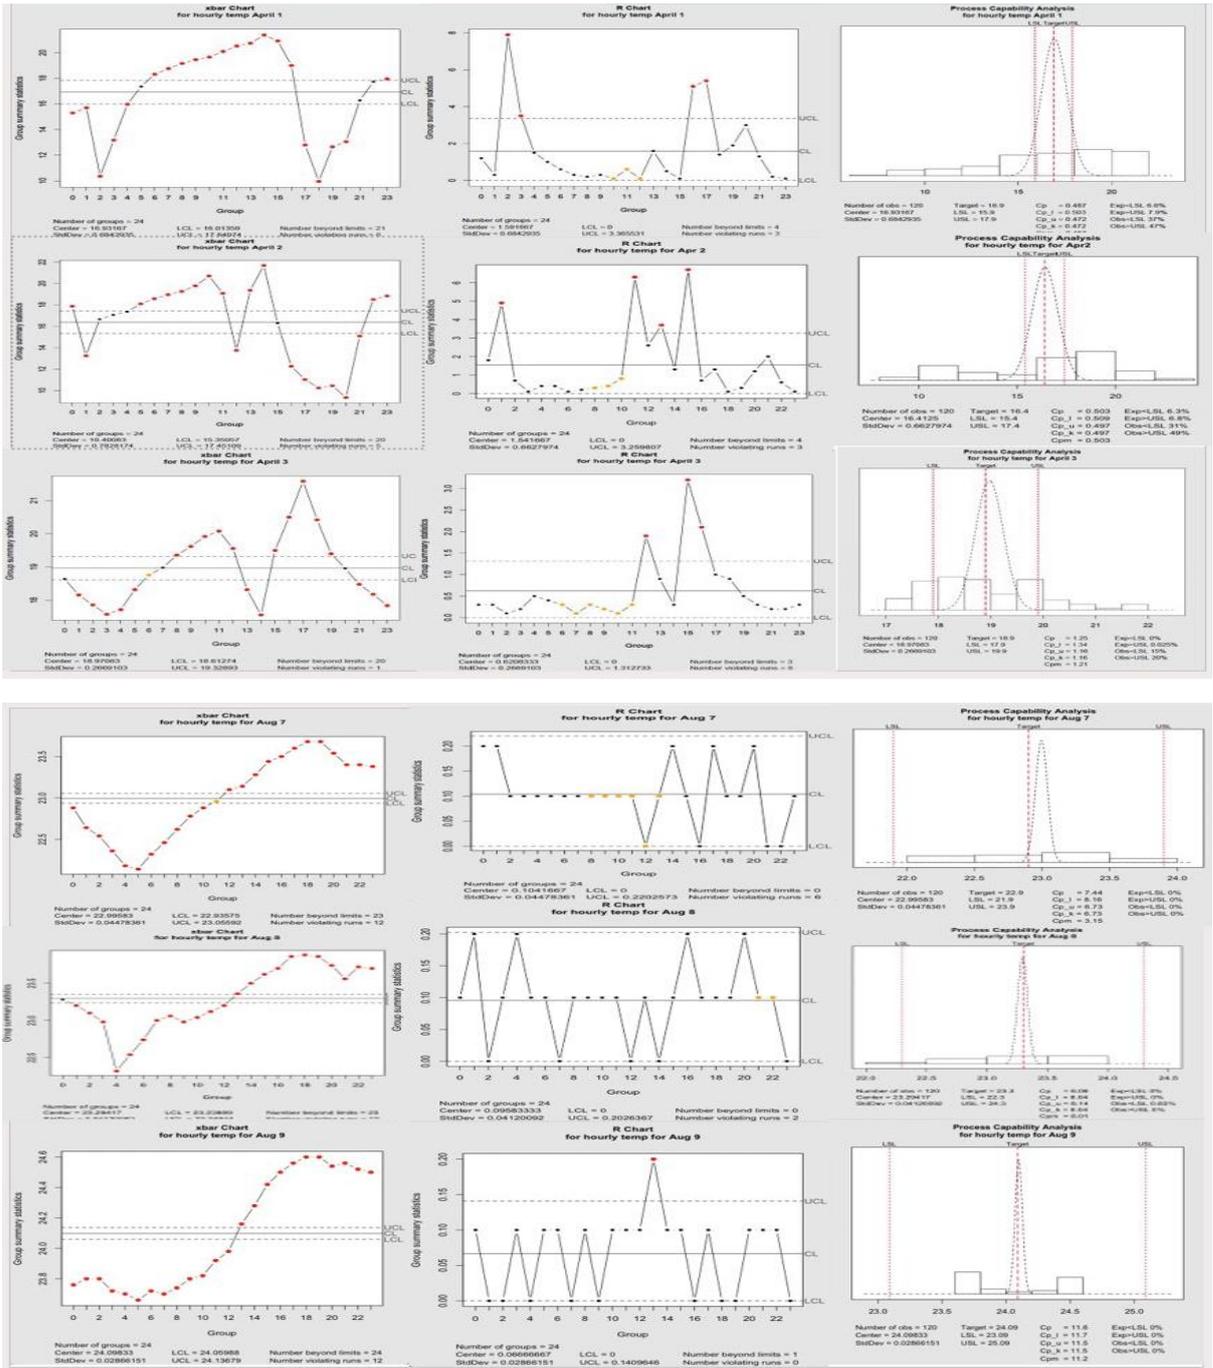
- Statistical methods-based technique to monitor, control, and improve the quality of a process or product
- Early identification of underlying causes of disturbance in processes, facilitating timely corrective action.
- SPC is based on averages and standard deviations of data over a duration, meaning the results are more reliable.

Process capability analysis (PCA)

- It is used to determine how well a process meets the desired specification limits.
- Measures how well the natural variation of the process meets the design specifications.

SPC & PCA are both widely used in many industries for decades, especially within the manufacturing supply chain for process improvement. The techniques are known to reduce overall costs by saving on wastage and process reworks, while enhancing operational efficiency & product recalls.

As mentioned, data from the temperature sensors is received every 10 minutes. Since we deal with measurable data, we have used X-bar and R-charts. Using the sample mean & range data, we draw the Process Capability Analysis charts for 3-sigmas.



Interpretation

Table 1: Process capability indices for April 01-09, 2022

Date	April 01	April 02	April 03	April 04	April 05	April 06	April 07	April 08	April 09
Cp	0.487,	0.503,	1.29	1.76	0.65	0.75	0.596	1.48	3.66
Avg Cp	1.24								

Table 2: Process capability indices for August 01-09, 2022

Date	Aug 01	Aug 02	Aug 03	Aug 04	Aug 05	Aug 06	Aug 07	Aug 08	Aug 09
Cp	7.75	8.28	8.88	6.86	7.75	6.42	7.44	6.06	11.6
Avg. Cp	7.8939								

The raw data shows larger fluctuations in April 2022, but the fluctuations have reduced over time as we moved to August 2022. The process capability has improved significantly between April 2022 (Avg. Cp = 1.24) and August 2022 (Avg. Cp = 7.9), indicating that the temperature control process being used by Picnic is becoming better at meeting the requirements. SPC charts based on the averages rather than individual data points, meaning the results presented are more reliable. This analysis substantiates the finding that for some underlying cause, there were larger fluctuations in April 2022 vs Aug 2022. Thus, SPC analysis can be a good quality control tool for temperature control in such food delivery processes.

Deployment

Incorporating SPC and PCA analysis into the operational framework of an agri-food business like Picnic offers a valuable means to assess and uphold the dependability of the cooling process. This analysis, combined with alerting mechanisms, constitute a robust toolkit for Picnic's operations. This deployment entails several significant benefits for the company, including enhanced monitoring capabilities, heightened responsiveness, augmentation of customer satisfaction and trust, reduction in costs, improved planning, adherence to legal and regulatory requirements, and the creation of new business opportunities. In this work, the analysis was limited to local (offline) calculations, however frameworks exist to integrate and automate SPC analysis with real-time sensor readings. Future work could examine developing online SPC analysis as a feature of the Whysor platform.

8. Raman

Preliminaries

The pilot program at UoN focused on upholding food quality during storage and transportation. It placed particular emphasis on utilizing an efficient and rapid technique called Raman spectroscopy. To accomplish this, a comprehensive plan was formulated. The initial stage involved laboratory testing of the Raman sensor on food samples, analyzing the acquired data, and subsequently uploading it to the server. This phase, known as "lab development," was followed by evaluating the technology under conditions resembling those encountered in refrigerated trucks, referred to as the "transitioning phase." Lastly, the technology will undergo testing in real-world conditions, representing the third and final stage of the pilot program.

The final system developed as a result of passing through all above phases is composed of:

- Raman spectrometer (QE Pro, Ocean Optics) - used to acquire the raman spectra of a sample
- Fiber-optic probe (InPhotonics, RPB78)
- Laser 785 nm (Oxxius, LBX-785-HPE) - laser source
- Laser controller (Oxxius, LaserBoxx series, ControlBoxx for LBX models) - control the power of the laser
- 3-axis motorized platform (Standa, 8-0026) - allow to move the sample in the desired location (X, Y, Z)
- 3-axis controller (Standa, 8SMC4-USB-B8-B9)
- A lockable box (QUIPO Locker - dimensions: height 450 mm x width 450 mm x depth 450 mm) covered on the inside with black opaque adhesive film (Adhesive vinyl, ME310-61). The box is equipped with four adjustable legs. - To prevent light from entering and interfering with the Raman signal
- A Raspberry Pi Kit – Allow for the Raman system to be an IOT-complaint device.

The aforementioned system was tested on a variety of food products, including chicken, resulting in the collection of over 6000 Raman spectra. These spectra were pre-processed (baseline correction, smoothing and normalization) and then analyzed using multivariate data analysis techniques such as Principal Component Analysis (PCA) and Kruskal-Wallis (KW) test. To facilitate further analysis, an automated script was developed to extract the data and transmit it through a server. Additionally, this script is capable of pre-processing the Raman spectra, conducting analysis, and providing the aforementioned results within a timeframe of 2-5 minutes, depending on the data size.

Technology stack selected

- Training data was provided offline by University of Nantes and ingested into a Warp 10 platform
- Python was used to design a pipeline for data ingestion. Python was later used at the model deployment stage.

Analytics

Pre-processing

Raman spectra were analyzed using Opus software (Bruker optics GmbH, V 7.2, Germany) and MATLAB software (version R2019b, MathWorks Inc, Natick, MA, USA). Initially, the raw data spanning a spectral range of 200–4000 cm^{-1} was restricted to the range of 500–3000 cm^{-1} . Subsequently, any interference caused by cosmic spikes within the defined spectral range (500–3000 cm^{-1}) was manually

removed prior to further data processing. To enhance the accuracy of the analysis, all spectra were subjected to baseline correction using an elastic concave method (with parameters of 64° and ten iterations), followed by smoothing using the Savitzky-Golay algorithm. Finally, the spectra were normalized using min-max normalization.

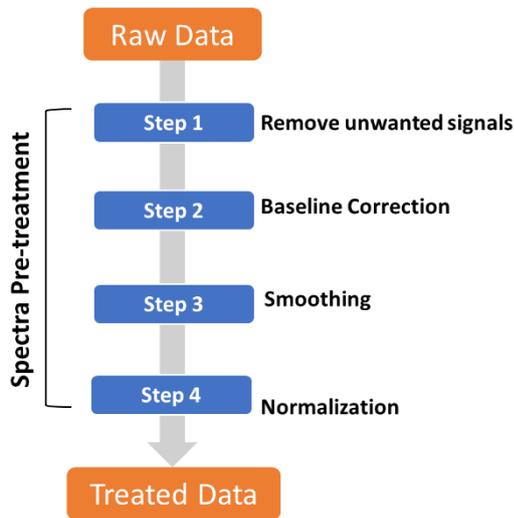
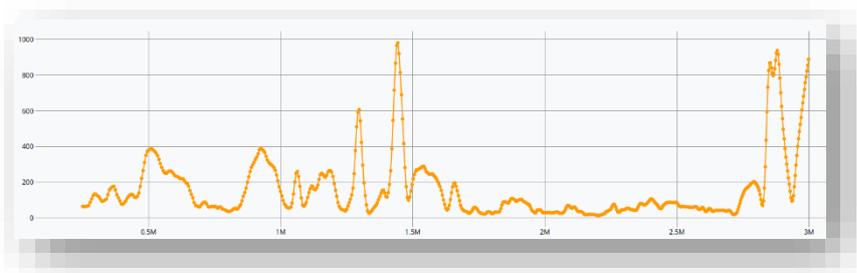
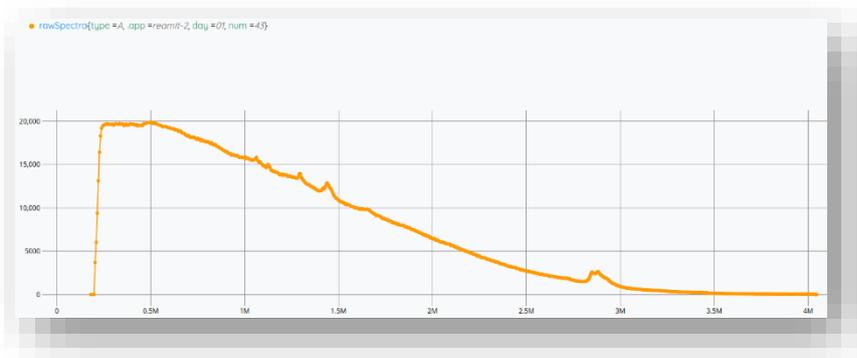


Fig. The necessary steps to pre-process Raman spectra

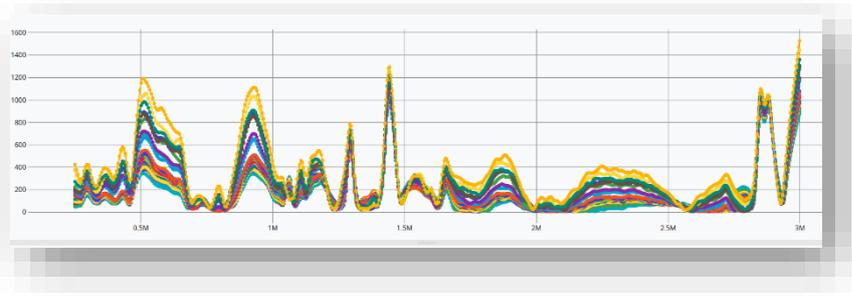
Analysis



An example of a raw sample is depicted here.

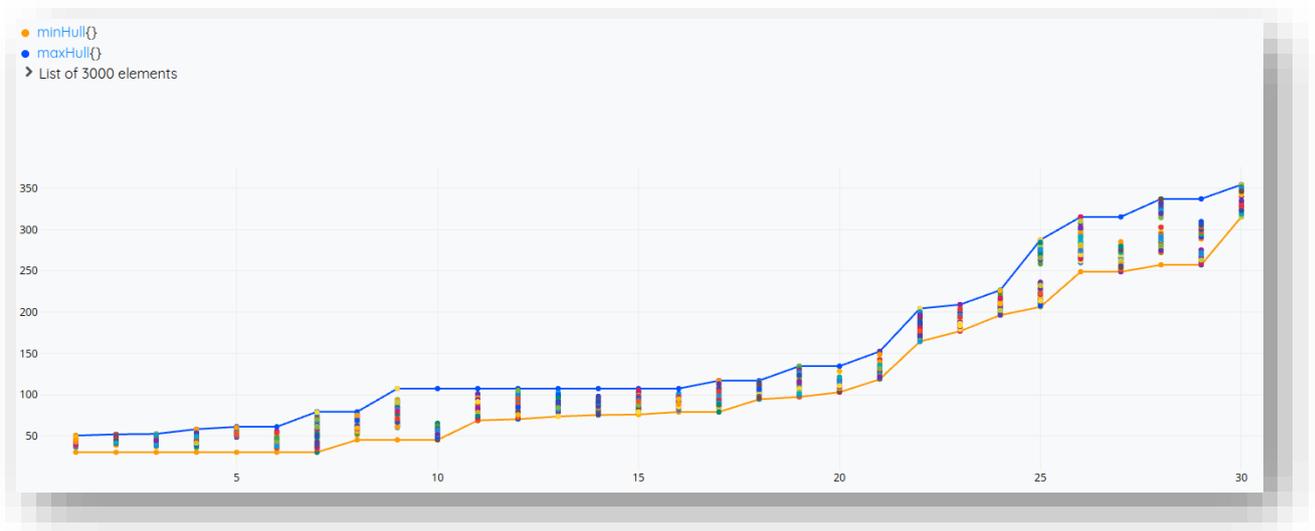


This is the same signal after it is has been treated.

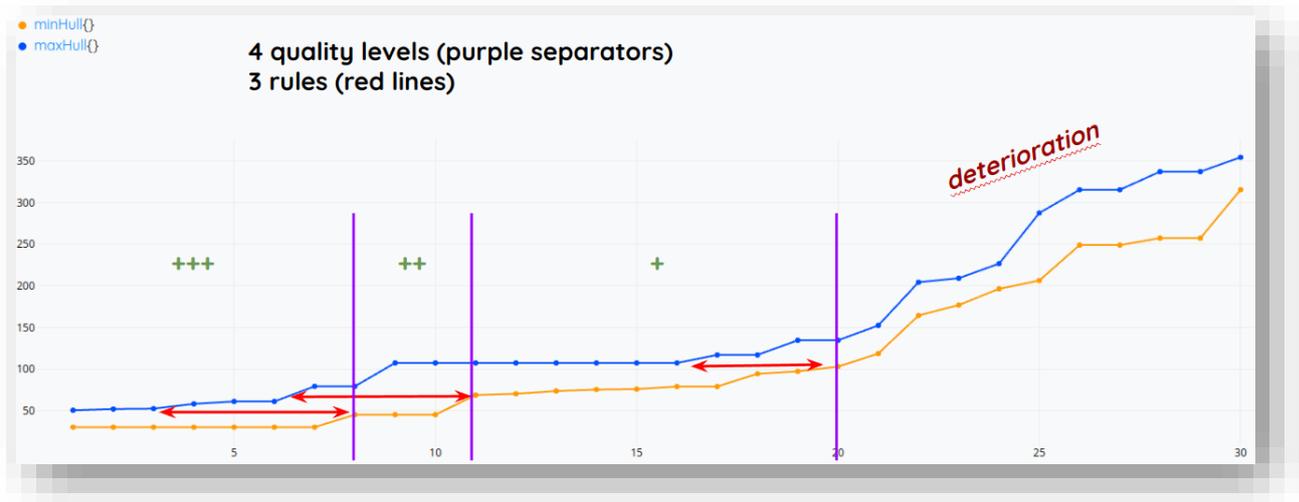


Here, we plotted the 30 day-mean series over the course of a month (each series aggregates 100 samples).

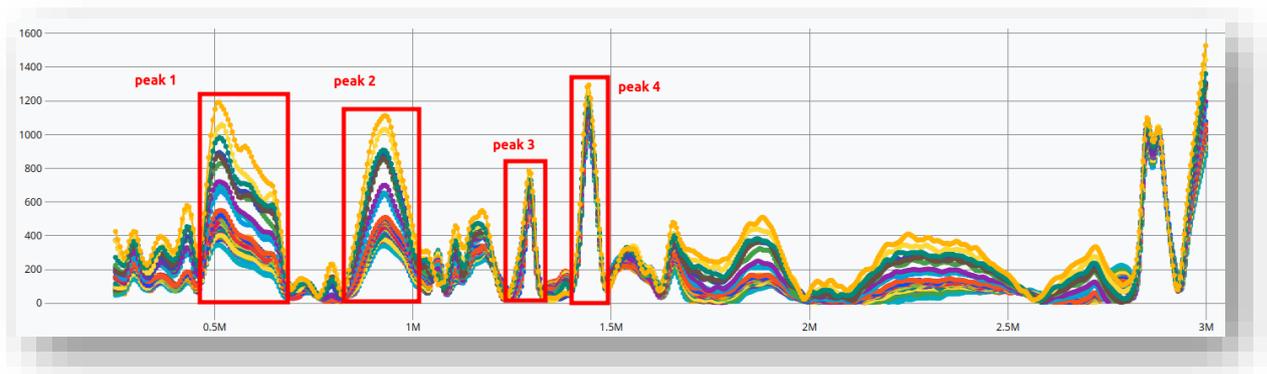
Next, we calculated the Euclidean distances of each sample from the mean series of the first day. In the following illustration, each point represents a result, with the day indicated on the x-axis. The blue and orange lines represent the upper and lower monotonic convex hulls, respectively.



As depicted in the figure below, the hull and the crossing points between its lower and upper lines can be utilized as boundaries to establish rules for an expert system. These rules can help define the quality level.



Even though the gist of the system given by those three rules, we have further enhanced the expert system by incorporating additional rules. By focusing on specific peaks and constraining the data, we have obtained a refined version of our model.



Interpretation

# samples	extra fresh	fresh	okay	expired
first model	384	513	774	1329
refined model	359	496	695	1450

Deployment

The 4-class chicken freshness model was successfully implemented and deployed on a local instance of the Warp10 platform running on a custom Raspberry Pi IoT module equipped with an LTE-M shield

for remote connectivity during transportation. By utilising the Warp10 platform on the Raspberry Pi module, real-time analysis and classification of chicken freshness could be achieved. The classifier is designed to categorise chicken into four distinct freshness classes, "extra-fresh", "fresh", "good", and "expired", allowing for efficient monitoring and management of poultry inventory. The system architecture diagram in Figure 16 depicts the overview of the online chicken analysis system.

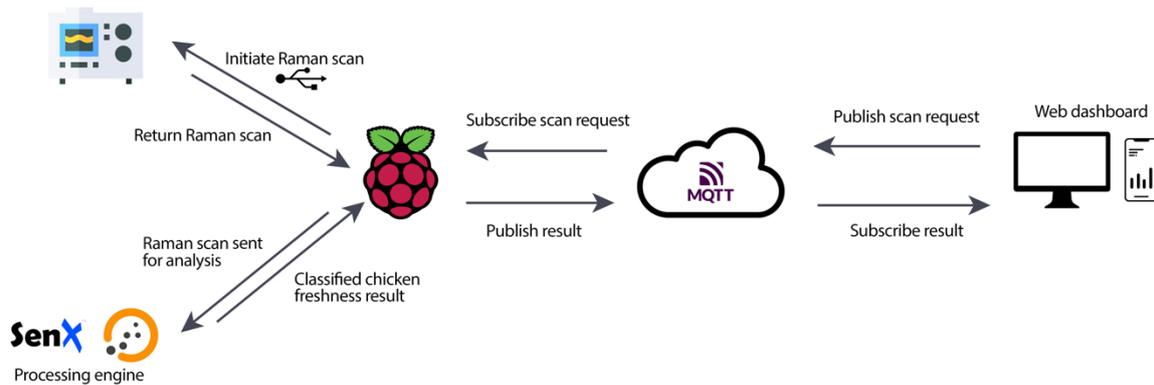


Figure 16: System architecture of the IoT Raman system for chicken classification.

To request a real-time scan and analysis, facilitate user interaction and provide a seamless experience, partners at Ulster University developed a frontend application. This application allows users to request Raman scans and, upon completion of the classification process, presents the results through a user-friendly interface. The system leverages MQTT (MQ Telemetry Transport), a lightweight, publish-subscribe, machine to machine network protocol for message queue/message queuing service. It is designed for connections with remote locations that have devices with resource constraints or limited network bandwidth, such as in the Internet of Things (IoT).

The frontend application was developed using PHP programming language and incorporates Bootstrap CSS, Chart.js, PHP-MQTT, and jQuery. To ensure security, the website requires a username and password for access, implemented through htaccess authentication. This authentication mechanism guarantees that only authorised users can initiate a scan and interact with the application, maintaining data integrity and confidentiality.

This interface enables end users to easily interpret and understand the freshness classification of the food samples. The user interface is accessible from <https://finch.ulster.ac.uk/reamit>

Towards real-time Raman analysis of chicken during cold-chain transportation

Food Quality Status Checker

Check chicken status



Figure 17: UI screen to initiate Raman scan.



Figure 18: Once the scan has been remotely triggered, scan data obtained, and classified using the Warp10 platform, the results are visualised on the front end using an easy to interpret freshness gauge for the end user. The status is also displayed as text below the visualisation.

9. 3D Fluorescence (FreshDetect)

Preliminaries

The pilot program at UU focused on testing a handheld non-invasive fluorescence spectrometer device called FreshDetect, which operates at an emission wavelength of 405nm. The REAMIT team had been tasked with expanding the application of FreshDetect beyond its original intended use with meat products, where it was able to estimate the total viable count of bacteria present. The REAMIT team at UU proposed examining its effectiveness in determining the freshness of whole milk (2%), fuelled by the removal of use-by dates on milk by some UK supermarkets in 2023, promoting alternative methods like the "sniff test" for determining milk spoilage. This pilot sought to explore the potential of utilising the portable handheld spectroscopy device as a quantitative tool for measuring milk quality. By doing so, it aimed to reduce the reliance on subjective olfaction techniques and potentially pave the way for the introduction of handheld spectrometers as a commonplace tool in households, in turn offering consumers a more reliable and convenient method for assessing the freshness of their milk. The objective was to examine the relationship between the fluorescence signals emitted by milk and attempt to correlate it to its freshness status.

The dataset compiled for the FreshDetect trial consisted of 12 recording sessions of fresh milk. The milk data was systematically recorded with spectra measurements acquired at 15-minute intervals, spanning a period of 48 hours per session. Notably, this recording process was conducted across two identical FreshDetect devices in parallel, effectively yielding a total of 3756 recorded spectra for model development. Each recording was executed on fresh whole milk (2%), sourced from a local convenience store. Since measurements happened in parallel between 2 devices, the 12 recording sessions resulted in the dataset being formed from 6 different milk samples. Milk pH was recorded on one batch of milk using an Orion Star A215 bench Ph meter sampled at 15-minute intervals.

Technology stack selected

Matlab 2021b was chosen as the software for analysing and model development of the milk spectra data due to its popularity and strong reputation in the field of data analysis. Matlab is widely recognised for its advanced statistical analysis and machine learning toolkits, which provide a comprehensive set of functions and algorithms for conducting various statistical analyses and producing ML models. Both classification models and regression models were produced in this study.

Analytics

Pre-processing

Fluorescence spectra were directly imported to Matlab 2021b using a custom data loading pipeline built for the use case. After the data is loaded, the ISO8601 timestamp in the spectra datafile is automatically converted to POSIX time (i.e. Unix / epoch time). Since the devices were turned off between each batch changeover, there were periods where the next recording was not exactly 15 minutes since the last. Using this information, batch labels were automatically generated by calculating the difference of the POSIX time between 2 scans. If the difference wasn't equal to 15 minutes, the script automatically saved this point as the changeover. A new time vector was then generated counting from 0 to 48 hours and appended to each identified milk batch spectra recording.

For the classification task, a simple binary classifier was selected, distinguishing between "fresh" and "not fresh" statuses. Spoilage in milk is typically marked by a pH below 6.4. From the pH dataset of a

single batch, it was observed that freshly purchased milk took 10.5 hours to reach this pH threshold. Given the absence of pH data for all milk batches, time was employed as a proxy for pH. Hence, a fresh/not fresh distinction was established at 10.5 hours. This resulted in two labels—true (spoiled) and false (not spoiled)—generated based on time, designating everything prior to 10.5 hours as false (not spoiled) and everything thereafter as true (spoiled).

Taking a data driven approach, the spectra was kept in its raw format and no data cleaning or processing was performed on the dataset beyond that which the device already applies internally.

Analysis

In classification analysis, two key metrics shine light on a model's performance: Sensitivity and Specificity. Sensitivity, or the True Positive Rate, gauges how well the model detects positive instances. It is calculated as:

$$\text{Sensitivity} = \text{TP} / (\text{TP} + \text{FN})$$

Meanwhile, Specificity, the True Negative Rate, assesses the model's precision in identifying negative instances, and is defined as:

$$\text{Specificity} = \text{TN} / (\text{TN} + \text{FP})$$

Balancing these metrics is crucial for refining a classification model's real-world utility.

In our model, positive labels are assigned to milk which is spoiled and negative labels assigned to the milk samples still considered fresh. In total, 21 models were tested for the binary classification task of classifying the spoilage status of the milk. These encompassed a range of techniques including Trees, Support Vector Machines (SVMs), Naïve bayes, K-NNs (K-nearest neighbour), ensembles, and neural networks. The best performing model was a neural network with 1 fully connected layer of 25 neurons, achieving an overall accuracy of 92.8% in classifying spoiled versus unspoiled milk samples. The confusion matrix for this model is presented in figure 19.

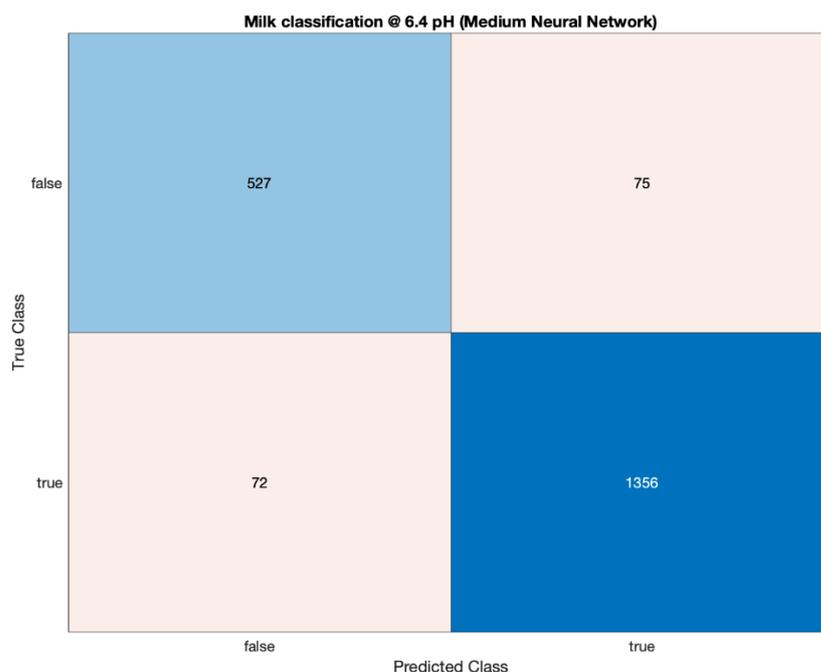


Figure 19: Confusion matrix for the fresh versus spoiled milk model.

Given the metrics from the confusion matrix:

TP = 1356

FN = 72

Sensitivity = $1356 / (1356 + 72) \approx 0.9492$

The True Positive Rate (Sensitivity) is approximately 94.92%, indicating that the model correctly identifies 94.92% of the truly spoiled instances out of all the actual spoiled instances in the dataset.

TN = 527

FP = 75

Specificity = $527 / (527 + 75) \approx 0.875$

The Specificity (True Negative Rate) is therefore approximately 0.875 or 87.5%, indicating that the model correctly identifies 87.5% of the truly fresh instances out of all the actual fresh instances in the dataset.

Interpretation

The calculation of the True Positive Rate reveals that the model possesses a Sensitivity of approximately 94.92%. This value encapsulates the model's efficacy in correctly identifying instances of spoiled milk from the dataset. In other words, out of all the genuinely spoiled instances present, the model successfully identifies around 94.92% of them. This high Sensitivity implies that the model is proficient at minimizing the number of instances where spoiled milk goes undetected. Such a characteristic is particularly relevant in scenarios where missing a case of spoiled milk carries significant consequences. For instance, in the food industry, the repercussions of distributing spoiled milk could extend to health concerns, regulatory issues, and economic losses.

On the flip side, the computation of Specificity, which stands at approximately 87.5%, underscores the model's aptitude in accurately classifying fresh milk instances. This metric signifies the model's ability to distinguish between truly fresh milk samples and identify them as such. Specificity holds importance in contexts where mistakenly flagging a fresh product as spoiled could lead to unnecessary waste or consumer dissatisfaction.

Given the nature of the problem at hand – classifying milk spoilage – a more cautious approach is preferred. Thus, the fact that the classifier places more significance on obtaining a higher Sensitivity is desirable. In this particular application, the ramifications of failing to identify spoiled milk outweigh the cost of occasional false alarms for fresh milk. Therefore, the model's adeptness in capturing a high proportion of truly spoiled instances lends confidence to its utility in ensuring consumer safety and product quality.

Deployment

The FreshDetect device, in its factory configuration, is equipped to run on-the-fly machine learning regression models to estimate the total viable count of bacteria in various meat products. These results are displayed to the end user on the device's screen. Consequently, the device is already primed for online evaluation of machine learning models. However, deploying the milk spoilage model onto the device will necessitate collaboration with the developers at FreshDetect. While this process

is not anticipated to be an extensive undertaking, it will require coordination to integrate the model seamlessly.

Once the deployment is accomplished, consumers who own a FreshDetect device will have the capability to conduct real-time measurements. This transition will empower users to assess the spoilage status of milk instantaneously. A mock-up of the FreshDetect user interface, showcasing milk classification, is presented in Figure 20. This depiction offers a glimpse into the user experience, illustrating how individuals will be able to leverage the device for milk spoilage assessment.



Figure 20: Mockup of the FreshDetect UI running the milk classification model.

Disclaimer: This document was prepared by James Gillespie (j.gillespie1@ulster.ac.uk), Jean-Charles Vialatte, Xavier Cama, Omar Dib, and Gautam Samriya to the best of their knowledge, in association with REAMIT project team. All information provided in this document are verified and found correct at the time of publication – June 2023.