SIEMENS

Insights Hub

Insights Hub Energy Optimizer

System Manual

07/2025

Legal information

Warning notice system

This manual contains notices you have to observe in order to ensure your personal safety, as well as to prevent damage to property. The notices referring to your personal safety are highlighted in the manual by a safety alert symbol, notices referring only to property damage have no safety alert symbol. These notices shown below are graded according to the degree of danger.

⚠ DANGER

indicates that death or severe personal injury **will** result if proper precautions are not taken.

⚠ WARNING

indicates that death or severe personal injury **may** result if proper precautions are not taken.

⚠ CAUTION

indicates that minor personal injury can result if proper precautions are not taken.

NOTICE

indicates that property damage can result if proper precautions are not taken.

If more than one degree of danger is present, the warning notice representing the highest degree of danger will be used. A notice warning of injury to persons with a safety alert symbol may also include a warning relating to property damage.

Qualified Personnel

The product/system described in this documentation may be operated only by **personnel qualified** for the specific task in accordance with the relevant documentation, in particular its warning notices and safety instructions. Qualified personnel are those who, based on their training and experience, are capable of identifying risks and avoiding potential hazards when working with these products/systems.

Proper use of Siemens products

Note the following:

⚠ WARNING

Siemens products may only be used for the applications described in the catalog and in the relevant technical documentation. If products and components from other manufacturers are used, these must be recommended or approved by Siemens. Proper transport, storage, installation, assembly, commissioning, operation and maintenance are required to ensure that the products operate safely and without any problems. The permissible ambient conditions must be complied with. The information in the relevant documentation must be observed.

Trademarks

All names identified by ® are registered trademarks of Siemens AG. The remaining trademarks in this publication may be trademarks whose use by third parties for their own purposes could violate the rights of the owner.

Disclaimer of Liability

We have reviewed the contents of this publication to ensure consistency with the hardware and software described. Since variance cannot be precluded entirely, we cannot guarantee full consistency. However, the information in this publication is reviewed regularly and any necessary corrections are included in subsequent editions.

Siemens AG
Digital Industries
Postfach 48 48
90026 NÜRNBERG
GERMANY

7/2025 Subject to change

Table of contents

1 Introduction

Introduction

Insights Hub Energy Optimizer extends Insights Hub Energy Manager with data-driven insights that enhance the understanding of energy conservation across production processes. It calculates the optimal energy consumption, sends notifications if the actual energy consumption is higher and identifies where energy efficiency can be improved.

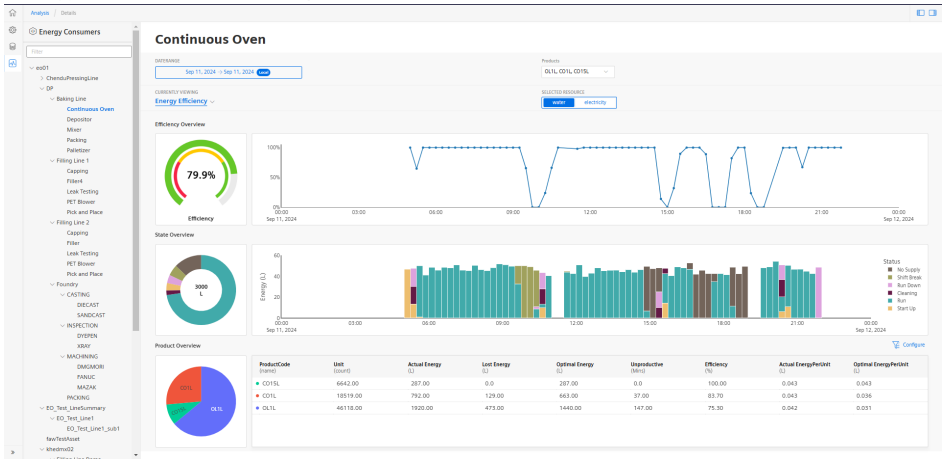

Insights Hub Energy Optimizer offers the following core capabilities:

- Calculate the Energy Efficiency metric as a KPI for multiple energy resources.
- Estimate the optimal energy consumption and compare to the actual for each energy resource consumed by each asset.
- Attribute energy resource consumption to different phases and states of assets to estimate productive and unproductive energy consumption.
- Estimate the energy consumed and the optimal energy per unit of production for every product.
- Aggregate the energy resources to calculate KPIs at the line and plant level.
- Identifiy a baseline energy standard for each product associated with each asset automatically.
- Send email notifications when actual energy consumption deviates from this standard.
- Ability to modify this standard to customize the frequency of email notifications.

Key concepts in Energy Optimizer include:

- **Energy Resource**: This represents a quantity such as water, gas, electricity or chemical compound, consumed during production, that the user wishes to monitor, control and optimize.
- **Energy Consumer**: This is an Insights Hub asset that has been configured with at least one Energy Resource variable. It would usually include other production context variables.
- **Energy Model**: This is a mathematical model relating Energy Resource to production context data on a specific Energy Consumer.

Prerequisites 2

Prerequisites

To fully enable the Insights Hub Energy Optimizer (EO) features, it is necessary to connect to energy resource data and other context variables from the asset of interest. Insights Hub EO does not require special devices to collect the data from the machines and lines. The only mandatory condition is that the data is stored in the timeseries database and structured properly. These energy variables can also be aligned with the configuration in Energy Manager.
The following table lists the variables that are used in various Energy Optimizer use cases. For each variable, there is a description on the use case that it enables. The usage of the variable is explained in detail in the configuration section.

Variable	Common Datatype	Comment
Resource	DOUBLE	This is the energy resource of interest. Multiple resources can be defined for each asset.
Status	BOOLEAN, INT, LONG, STRING	Status is the state or phase of the asset operation. It is necessary to fully enable energy efficiency calculations.
Product	INT, LONG, STRING	The product that is currently produced by the asset. Necessary to filter results by product and to setup product level energy notifications.
Batch	INT, LONG, STRING	The batch of products currently produced by the asset. Necessary to enable batch based aggregation of energy data for the notification configuration.
Reference	INT, LONG, DOUBLE	The quantity of parts produced on the asset. Necessary to calculate the KPIs on a per unit level.

User rights

3

User rights

Insights Hub Energy Optimizer adopts user rights from "Settings". The user rights depend on the following user roles:

- `mdsp:core:energyoptimizer.admin`
- `mdsp:core:energyoptimizer.user`
- `mdsp:core.energyoptimizer.creator`

The following table gives an overview of the permissions for the Admin, User and Creator roles:

Permission	Role: mdsp:core:energyoptimizer.admin	Role: mdsp:core:energyoptimizer.user	Role: Mdsp:core.energyoptimizer.creator
Configure Energy Resources	✓	-	-
Configure Energy Consumers	✓	-	-
Configure Status Mapping	✓	-	-
View Energy Efficiency	✓	✓	✓
View Energy Cost	✓	✓	✓
View Operator Alerts	✓	✓	✓
Create an Energy Model	-	-	✓
Edit an Energy Model	-	-	✓

Permission	Role: mdsp:core:energyoptimizer.admin	Role: mdsp:core:energyoptimizer.user	Role: Mdsp:core.energyoptimizer.creator
Deploy an Energy Model	-	-	✓

For more information on roles, refer to the Settings documentation.

User interface 4

Home page

The Insights Hub Energy Optimizer application can be accessed from the Launchpad using the following icon:

The Home page displays an overview of the number of configured Energy Consumers, Energy Resources and Energy Models in the application with quick links to each section. A "What's new" link informs the user about the new features.

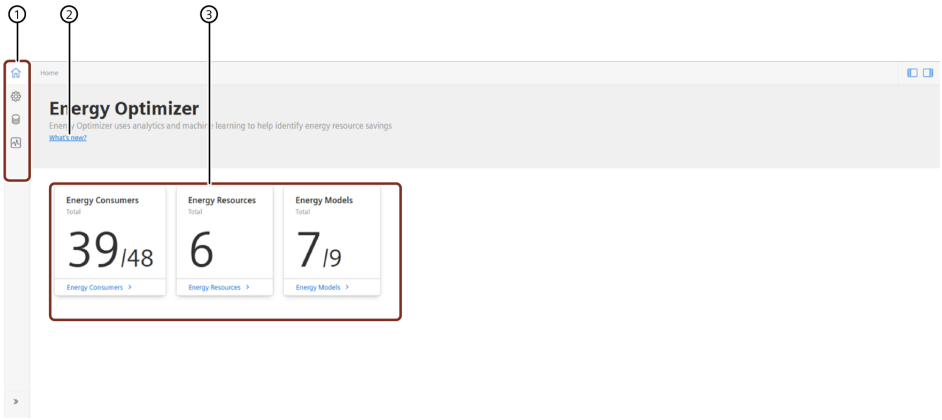

① Navigation tab with the following menu options:

- Configuration
- Models
- Analysis

② A link to the "What's new" page

③ Displays the number of configured Energy Consumers, Energy Resources and Energy Models

Configuration 5

5.1 Introduction to Configuration

Introduction to Configuration

To view the configuration screen, click the gear icon in the navigation tab. The main configuration screen is displayed as shown below:

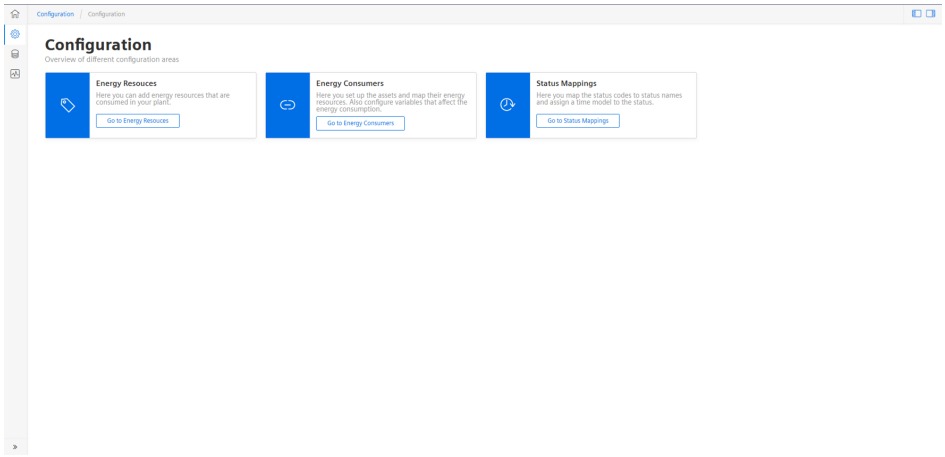

The main configuration screen has the following three panes:

- Energy Resources: setup the resources, such as electricity, gas or water.
- Energy Consumers: setup the asset data.
- Status Mappings: map the status value to a label and a production category.

For more information on the different panes, refer to the further sections.

5.2 Configuring Energy Resources

Configuring Energy Resources

In the "Manage Energy Resources" configuration, you can create multiple energy resources that are used within your plant. For each resource, you can set its name, units, description and a color. You can also set a contract to be used in the cost calculation. It is also possible to add more details to contracts in Insights Hub Energy Manager.

User interface

The user interface of "Manage Energy Resources" screen is as shown below:

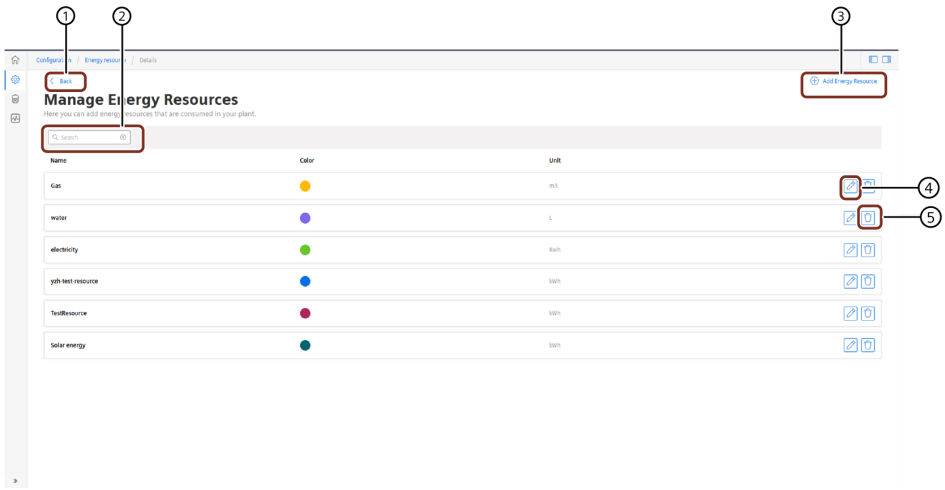

① Navigates back to the Configuration tab
② Searches the Energy Resources by name
③ Adds a new Energy Resource
④ Edits an existing Energy Resource
⑤ Deletes an existing Energy Resource

Creating a new energy resource

To create a new Energy Resource, proceed as follows:

1. From the navigation tab, click "Configuration" and select "Energy Resources".

2. In the "Manage Energy Resources" screen, click "Add Energy Resource".

3. Enter the "Name", "Description" and select a color to represent the resource.

4. Optionally, you can add one or more contracts.

5. For each contract, set a "Contract Name", "Cost", "Currency" and "CO2 equivalent". !!! note The contract data, if provided, is used in the energy cost calculation. More features for contracts are available in Energy Manager.

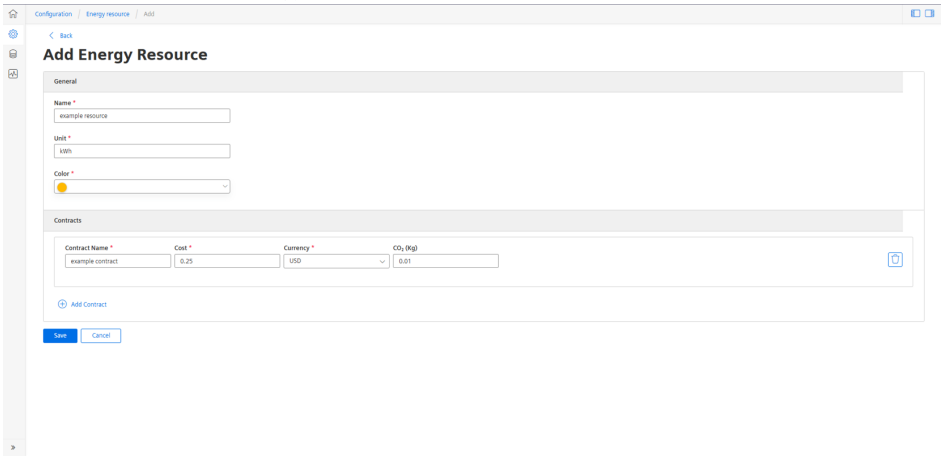

6. Click "Save".

The newly created energy resource will be displayed in the "Manage Energy Resources" page.

5.3 Configuring Status Mapping

Configuring Status Mappings

Status Mappings are used in the application to show how much energy is consumed in various machine states or production phases. Since a lot of energy resources are consumed during "unproductive" phases of a process, it is important to monitor this and set acceptable standards.
In the "Manage Status Mappings" configuration screen, you can create multiple status mappings that are used within your plant. Each status mapping relates the raw time series value for the status variable to a human readable value. It also relates the status to a time category. The time category by default is either "Productive" or "Unproductive". This is used to calculate the Energy Efficiency.

User interface

The user interface of "Manage Status Mappings" screen is as shown below:

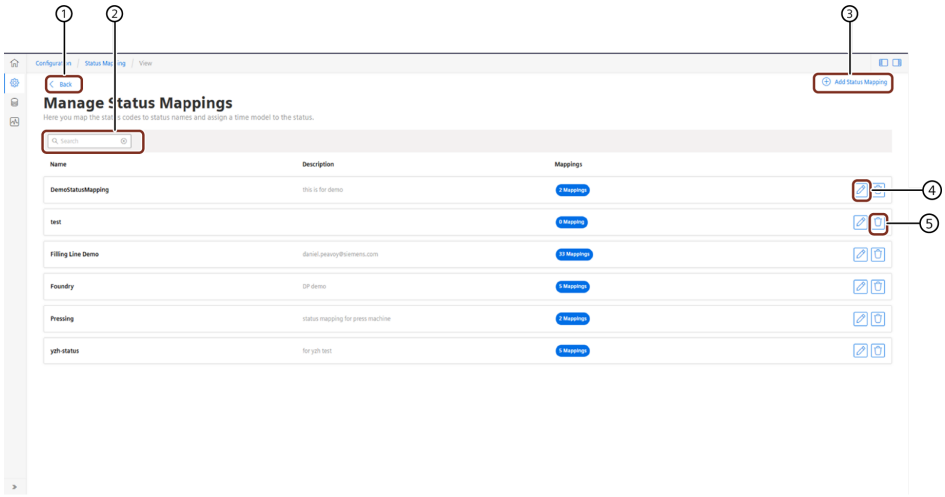

① Navigates back to the Configuration tab
② Searches the Status Mappings by name
③ Adds a new Status Mapping
④ Edits an existing Status Mapping
⑤ Deletes an existing Status Mapping

Creating a new Status Mapping

1. From the navigation tab, click "Configuration" and select "Status Mappings".
2. In "Manage Status Mappings", click "Add Status Mapping".
3. Enter a "Name" and "Description".
4. Click "Add" from the mapping section. Select a color from the drop-down or use the default option.

5. Set the value that is expected from time series and the "Label" and "Description" for this status.

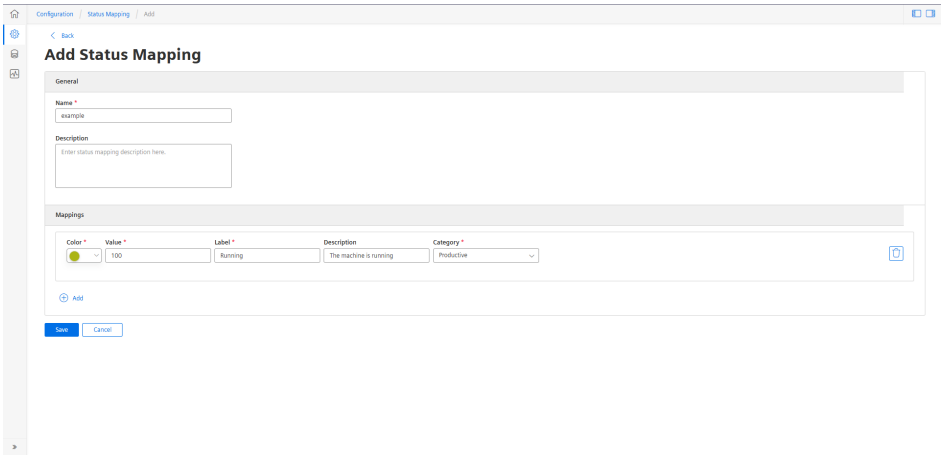

6. Select the category for this status. By default, this is either "Productive" or "Unproductive". You can add additional mappings by clicking the "Add" button.

7. Click "Save".

The newly added status mapping will be displayed in the "Manage Status Mappings" page.

5.4 Configuring Energy Consumers

Configuring Energy Consumers

In this section of the configuration, you can define the assets in your plant that consume energy resources. These assets should already be defined in Insights Hub Asset Manager. Therefore, you can inherit the hierarchy structure that is defined there. Typically, all the variables required to enable the Energy Optimizer use cases should be defined on aspects within the asset. For more information, refer to Asset Manager documentation.

In the "Manage Energy Consumers" page, the list of existing Energy Consumers in the asset hierarchy is displayed on the left. You can either select an existing consumer or to add a new Energy Consumer.

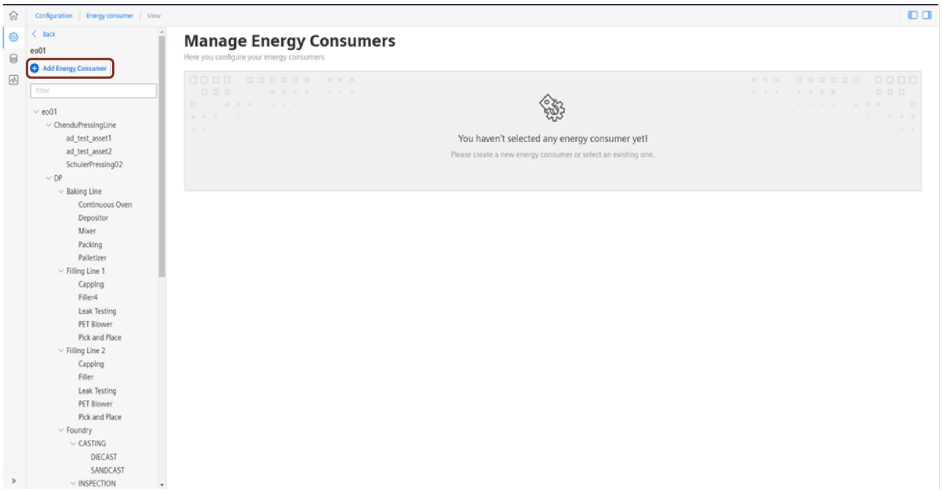

Adding a new Energy Consumer

1. In "Manage Energy Consumers", select "Add Energy Consumer".

2. Navigate to an existing asset defined in Insights Hub Asset Manager and click "Create".

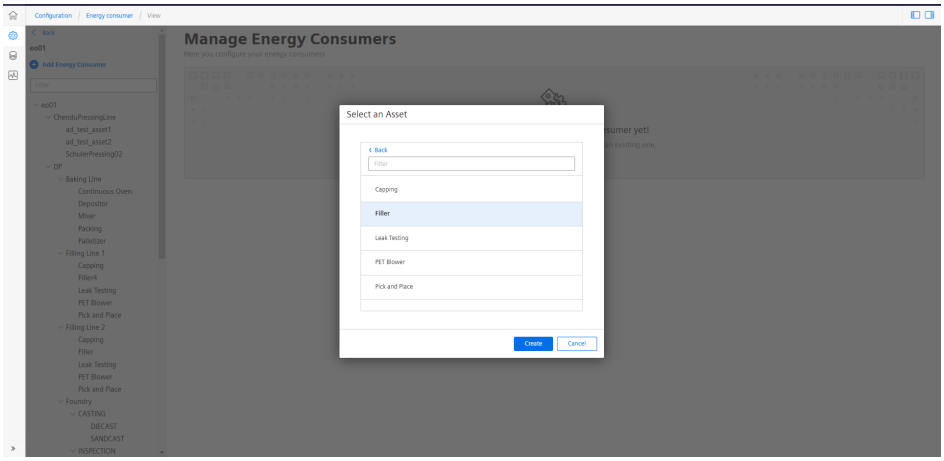

3. Set a "Name" and an optional "Description".

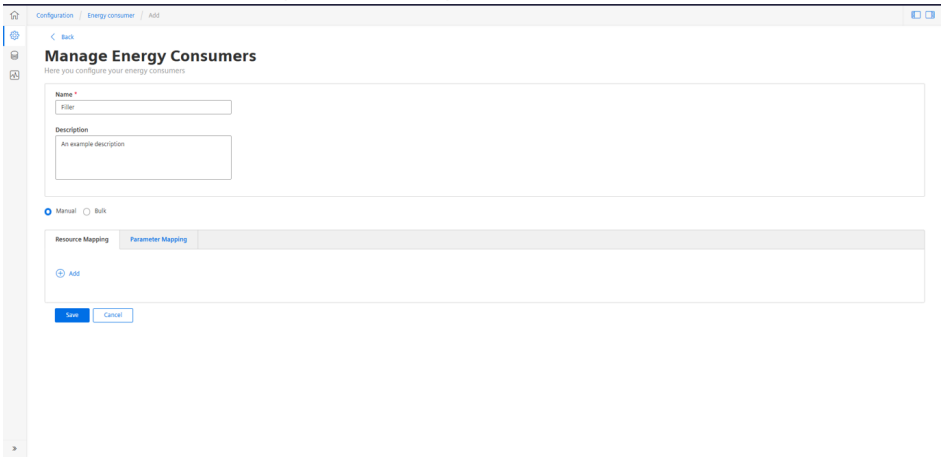

4. Select "Manual" option to map just a single asset.

5. Click "Save".

You can save this asset as a new Energy Consumer, but you need to add at least one Energy Resource mapping and set the context parameters to enable the use cases.

Adding a new Resource mapping to an Energy Consumer

1. In "Manage Energy Consumers", you can add a resource mapping. Select "Add" resource mapping to define at least one energy resource for this asset.

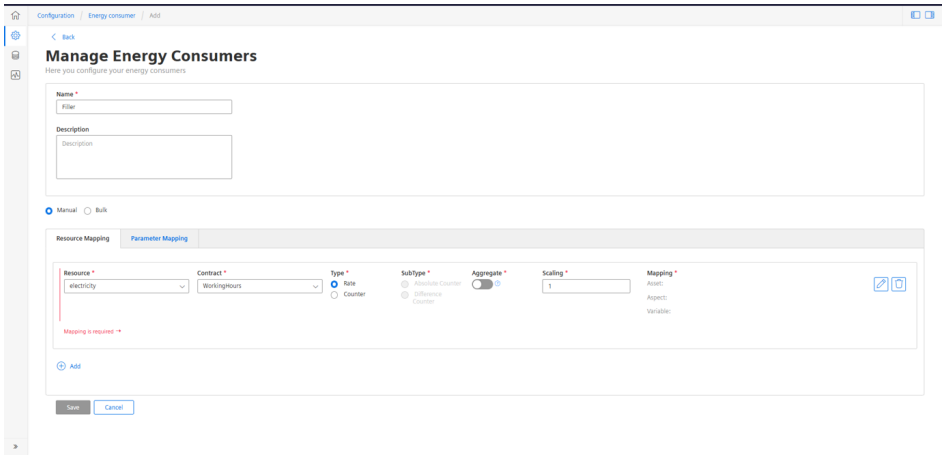

2. Select the resource from the drop-down menu and the associated contract. !!! note It is required to define the resource first in the Energy Resources section.

3. Specify whether the data for this resource is of type "Rate" or "Counter". A power or flow measurement, for example, would be of type "Rate", whereas energy or volume would be "Counter". !!! info In general, a "Rate" variable is a measure of change with respect to time and "Counter" accumulates with time.

4. If type "Counter" is selected, then select the "SubType". This specifies whether time series data are "Absolute" values or they are the "Difference" since the last reading.

5. If configuring an asset that has child assets, then you can choose whether the data is aggregated with the energy resource data defined on the child asset or whether it is to be treated as independent data. This could be important, for example, if there is an energy meter on a line that is measuring additional consumption to the other (child) assets on that line. In this case, you could choose to aggregate the data to give the total consumption for the line.

6. Select the time series variable to map to the energy resource.

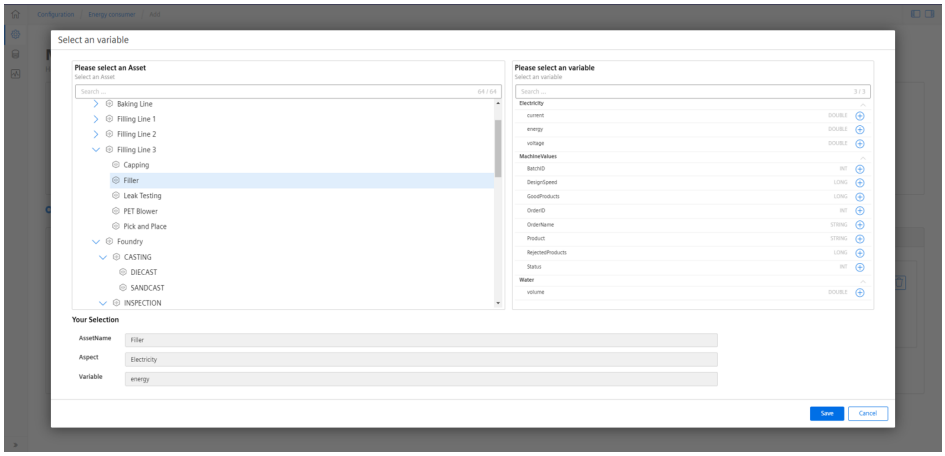

This is typically a variable defined on the asset, but it is possible to select any time series variable.

Additional resource mappings can be defined on the Energy Consumer. When all resource mappings are added, you can set the parameter mapping.

Adding a new Parameter mapping to an Energy Consumer

1. In "Manage Energy Consumers", select the "Parameter Mapping" tab.

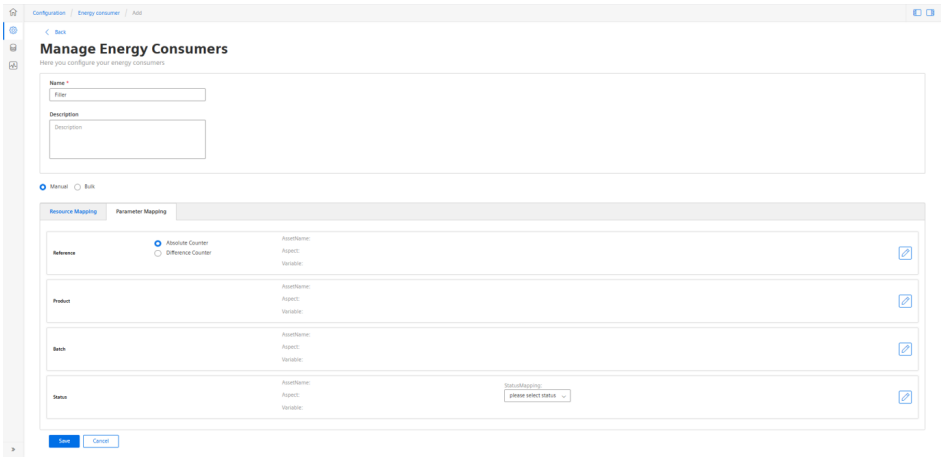

2. Define a "Reference" value.

This is the number of product units that have been produced on the asset. It is needed to calculate the "per unit" values of the KPIs in analysis pages. It can be defined as either an "Absolute Counter", where the value accumulates, or a "Difference Counter", where the value represents the increase since the last time series value. Select the time series variable using the button.

3. Map the "Product".

This is the time series for the product code that was being produced on the asset. It is needed to enable the product filter on the analysis pages. Navigate through the assets and select the aspect/variable.

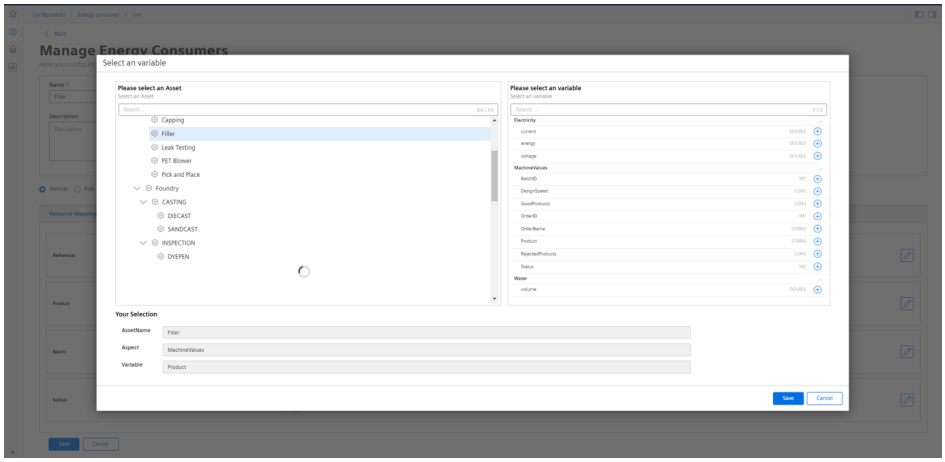

4. Map the "Batch".

This is the variable used to aggregate the energy data into "batches". This aggregation is used in defining the energy consumption standard in the "Operator Alerts"/Notifications use case.

5. Define a time series variable corresponding to the asset status. Select a status mapping from the dropdown. !!! note The status mapping needs to be defined already in the Status Mapping configuration.

When all variables are configured, click "Save" to see your new Energy Consumer in the asset hierarchy.

If some parameters are unmapped then the energy resource will still be displayed in the analysis pages but with limited functionality.

Bulk Configuration

You can configure multiple Energy Consumers using the Bulk configuration option. Select the radio button for "Bulk" and proceed similarly to the "Manual" procedure. Instead of selecting specific variables from aspects, you select variables from aspect types. This is then used as a template to build other Energy Consumers from assets in the hierarchy with the same aspect types. For example, the Energy Resource mapping asks the user to map the variable using an aspect type:

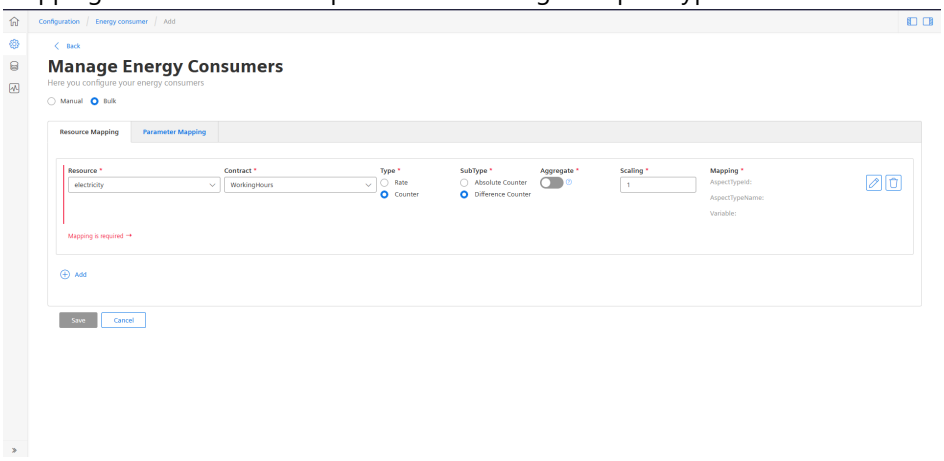

Analysis 6

6.1 Introduction to Analysis

Data Analysis

Introduction to Data Analysis

The Data Analysis area can be accessed by clicking the  icon from the left navigation. The "Data Analysis" screen provides several views to address energy related use cases. Each view displays specific charts and KPIs for the Energy Resource on the selected Energy Consumer. These views are used together to identify energy resource savings. They are explained in detail in the following sections:

- Energy Efficiency
- Energy Cost
- Operator Alerts
- Process Optimization

The user interface of "Analysis" screen is as shown below:

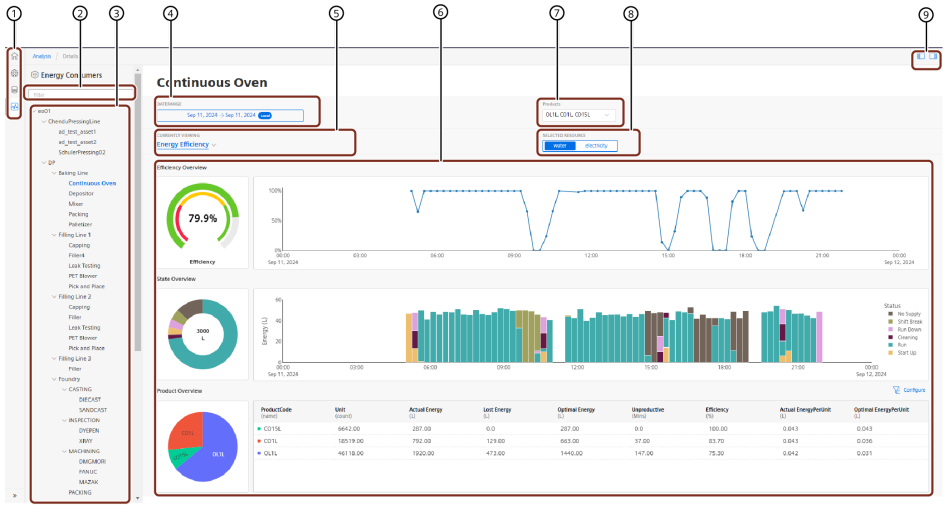

① Navigation tab
② Search for an Energy Consumer
③ Energy Consumer hierarchy which displays all assets that have been configured as Energy Consumers
④ Analysis view selection
⑤ Date range filter
⑥ Analysis view results pane
⑦ Product filter

⑧ Energy Resource selection
⑨ Hide/Show Asset Hierarchy

Analysis views

The Insights Hub Energy Optimizer addresses different use cases by providing specially designed analysis views. The "Energy Efficiency" page is displayed by default. To switch between analysis views, click on "Currently Viewing" dropdown.

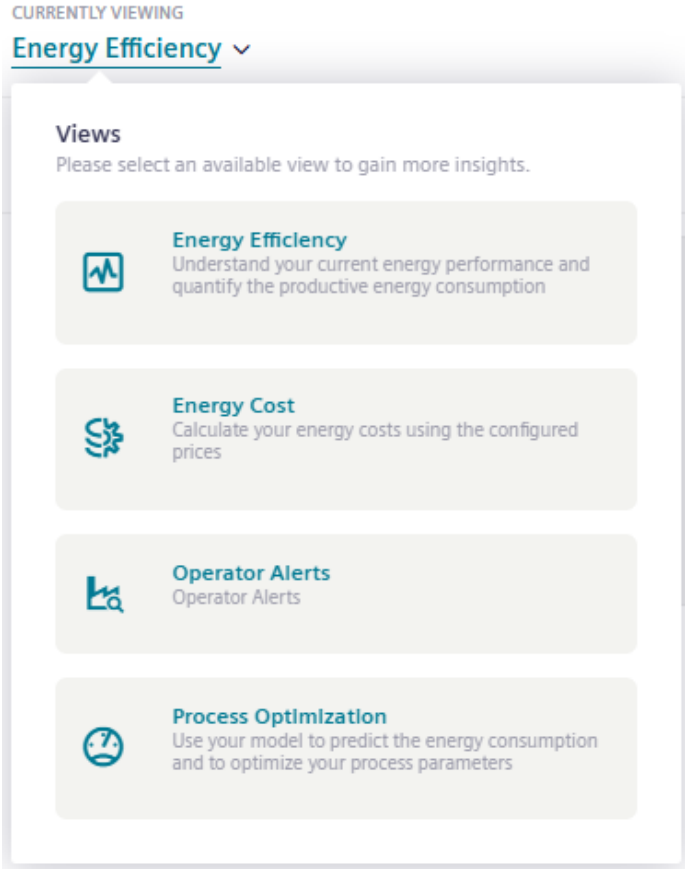

Energy Resource selection

These buttons enable switching between the multiple Energy Resources that can be defined on an Energy Consumer. In some views, such as the Energy Efficiency, only one Energy Resource can be selected. The other views, such as the Energy Cost can show results from multiple Energy Resources, so this filter becomes a multi-select.

Data filtering

The user can apply filters that persist when switching between views. On clicking the "Daterange" field, a window appears with the following options:

- **Absolute:** To select a specific date and/or time range.

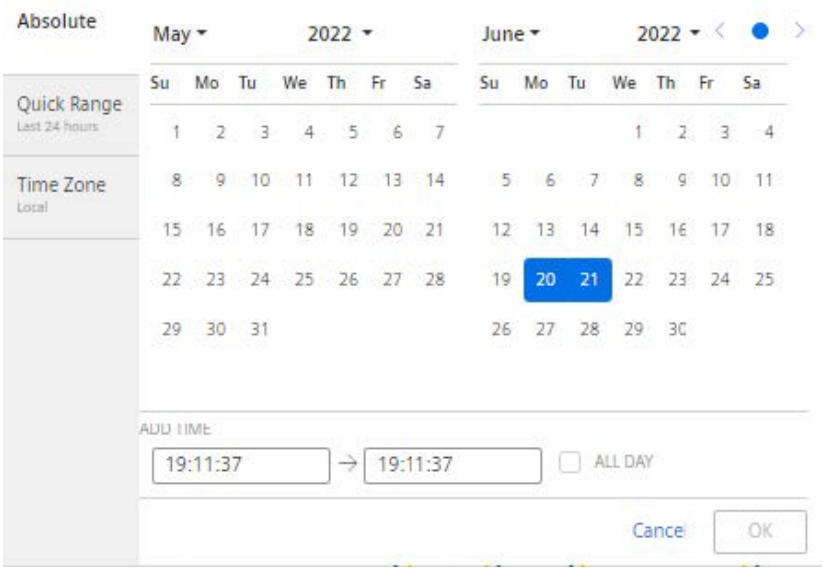

- **Quick Range:** To select the desired time frame.

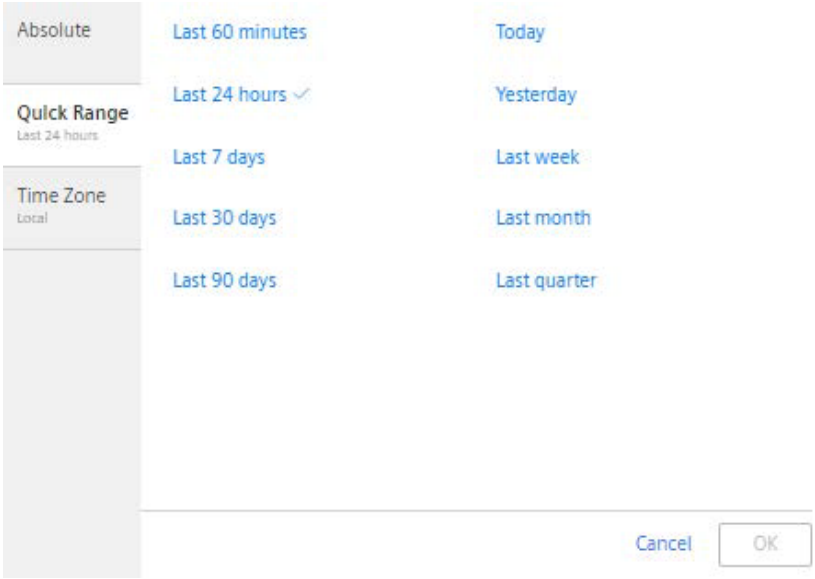

- **Time Zone:** To select a specific time zone or use the local time zone from your browser.

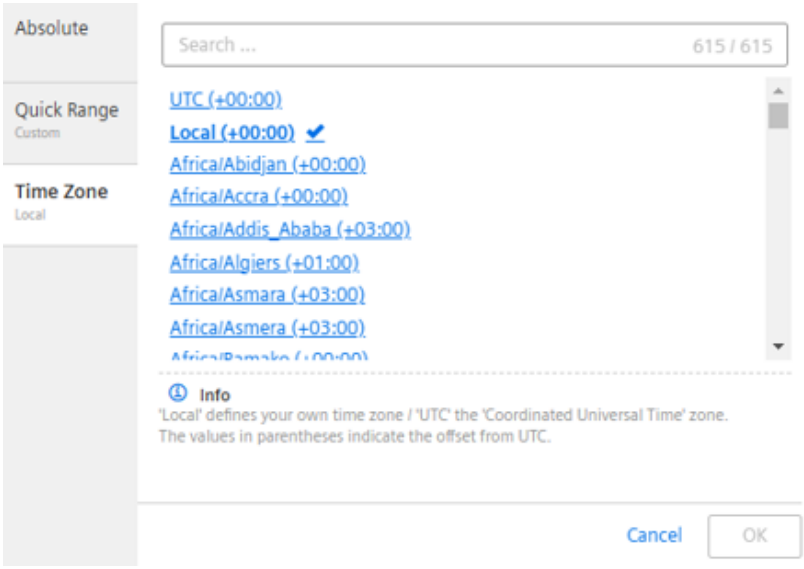

Additionally, the user can customize the view by selecting specific products using multi-select from the corresponding dropdown list:

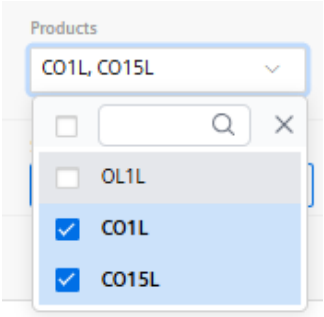

Visualization pane

Based on the selected view, different analysis results are displayed in the visualization pane.

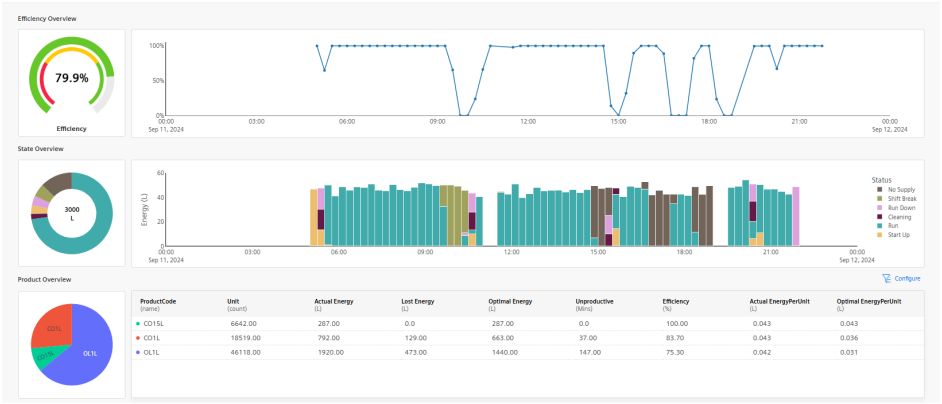

In all the views, the following functionalities are available for the user to explore the charts:

- Hover the mouse cursor over any chart points to get pop-up information about that particular point.

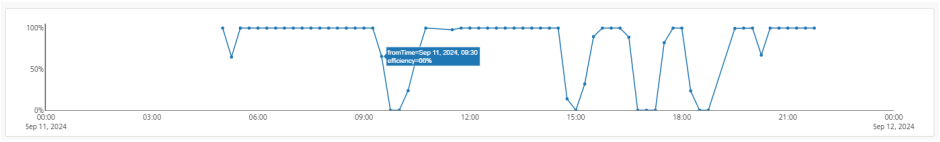

- Zoom into a date range by dragging a box over the chart. Double click to reset to the original date range.

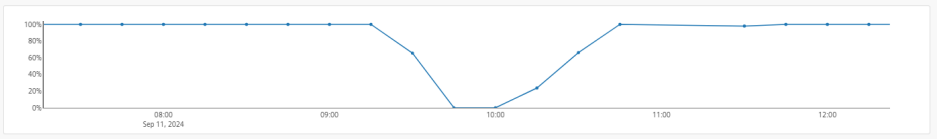

6.2 Energy Efficiency

Energy Efficiency

The Energy Efficiency is the central KPI to indicate how close we are to the 'optimal' or minimum energy resource consumption. It is the starting point to assess current and past energy performance and to monitor the impact of any process improvements. There are several ways to define the efficiency. Here, we are comparing the total productive with the unproductive energy consumption by using the status of the Energy Consumer and the associated time category from the status mapping. The efficiency for Energy Resource `j`, between dates t1 and t2, is calculated as:

$$\varepsilon_j = \frac{\sum_{t_1<t<t_2} \sum_{s \in prod} e_{sj}(t)}{E_j}$$

where E_j is the total consumption for resource `j` and where $\sum_{s \in prod}$ represents a sum of the consumption over productive states that occured during the date range. For example, productive states could be "running", "working", or "producing"; unproductive states could be "idle", "waiting" or "maintenance".

The Energy Efficiency view includes several sections and charts related to this KPI:

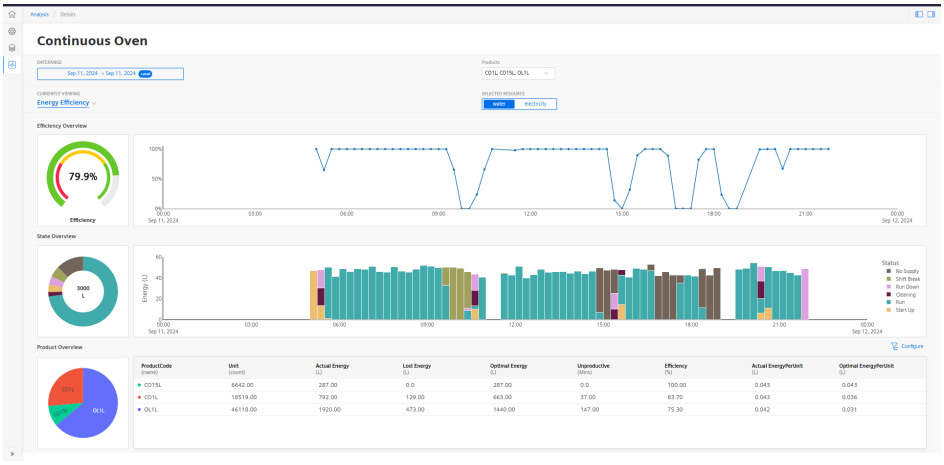

Efficiency Overview

- The gauge chart shows the total Energy Efficiency for the date range, with all applied filters.

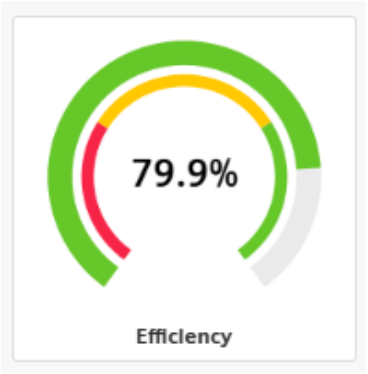

The red "error" region is set by default to less than 30%; the yellow "warning" region is less than 70%.

- The time chart shows the Energy Efficiency calculated in fixed time intervals. For date ranges less than or equal to 24 hours then the interval is 15 minutes; longer than 24 hours, but less than 7 days, then the interval is one hour; longer date ranges use daily intervals.

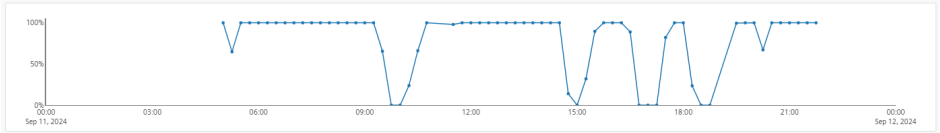

The time chart is useful to detect date ranges where the efficiency is low and to start investigating the root cause.

State Overview

- The donut chart shows the split in consumption according to each state that was active during the date range. 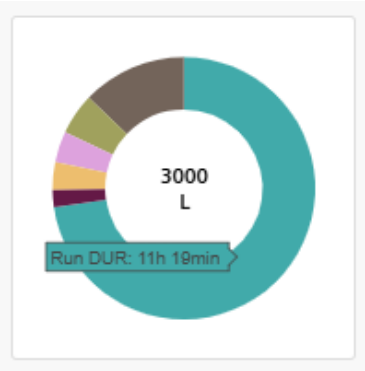

The total consumption is shown in the center. The tooltip shows the duration of each state.
- This grouped bar chart shows the energy consumed in each state for each time interval. Each state is color coded according to the "Status" legend.
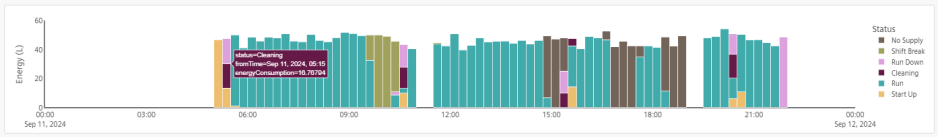

The tooltip shows the consumption in that state during the specific time interval. This chart is useful to show when Energy Consumers were operating in unproductive states and to quantify the associated energy "losses".

Product Overview

- The pie chart shows the split in consumption according to each product that was active in the date range and filters. 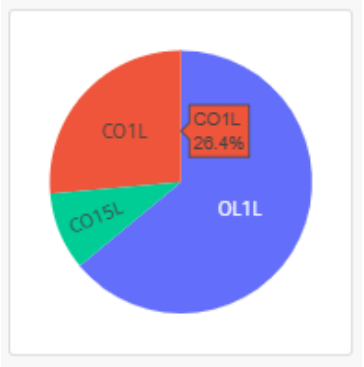

The tooltip shows the percentage consumption for each product.
- The product table shows multiple KPIs for each product. The KPIs shown in the table can be selected using the "Configure" filter at the top right. The KPIs include
    - Unit: the count of number of units for each product
    - Actual energy
    - Optimal energy
    - Lost energy: the difference between the optimal and actual energy
    - Unproductive time

- Efficiency for each product
- Actual energy per unit
- Optimal energy per unit

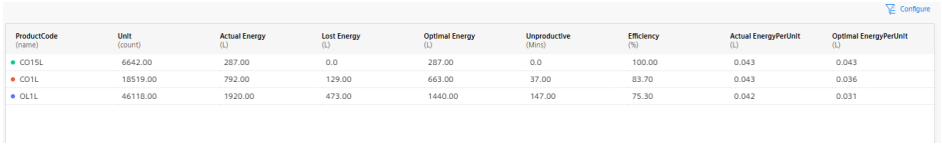

ProductCode (name)	Unit (count)	Actual Energy (L)	Lost Energy (L)	Optimal Energy (L)	Unproductive (Mins)	Efficiency (%)	Actual EnergyPerUnit (L)	Optimal EnergyPerUnit (L)
CO15L	6642.00	287.00	0.0	287.00	0.0	100.00	0.043	0.043
CO1L	18519.00	792.00	129.00	663.00	37.00	83.70	0.043	0.036
OL1L	46118.00	1920.00	473.00	1440.00	147.00	75.30	0.042	0.031

The table is useful to review numerical values for each KPI and identify any differences between product performance.

Energy Consumer aggregation

The Energy Efficiency is also calculated automatically at higher levels in the Energy Consumer hierarchy by aggregating the Energy Resource consumption across multiple Energy Consumers. The visualization is similar to the single asset view accept now the state variable is grouped according to its time category. This is so that the Energy Efficiency KPI can be calculated across machines that have different states and status mappings. In this example, the "Status" is shown as either Productive or Unproductive.

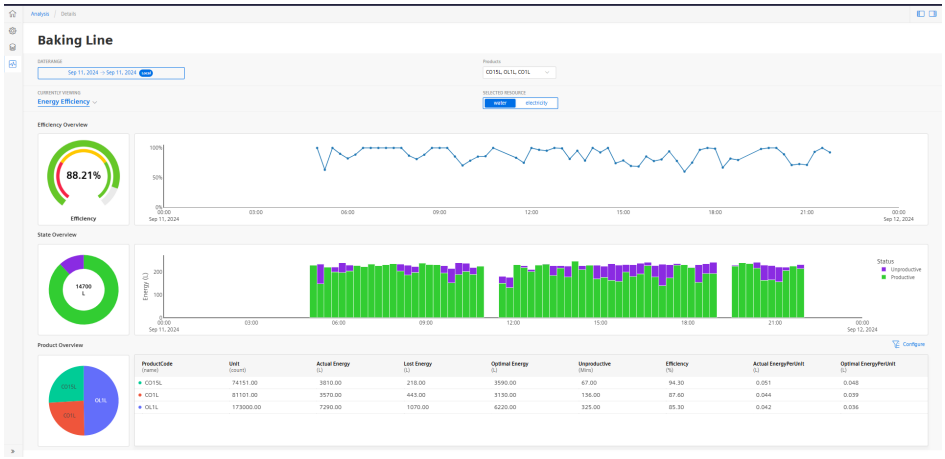

In this way, the user could monitor the overall Energy Efficiency for a line or an entire plant.

6.3 Energy Cost

Energy Cost

This view provides an overview of costs for each resource calculated from the contracts. Further analysis of costs can be achieved with Insights Hub Energy Manager. For more information, refer to documentation.

The user can select one or multiple resources to analyse by using the built-in charts and tables. In the following example, the costs associated with both electricity and water can be simultaneously viewed.

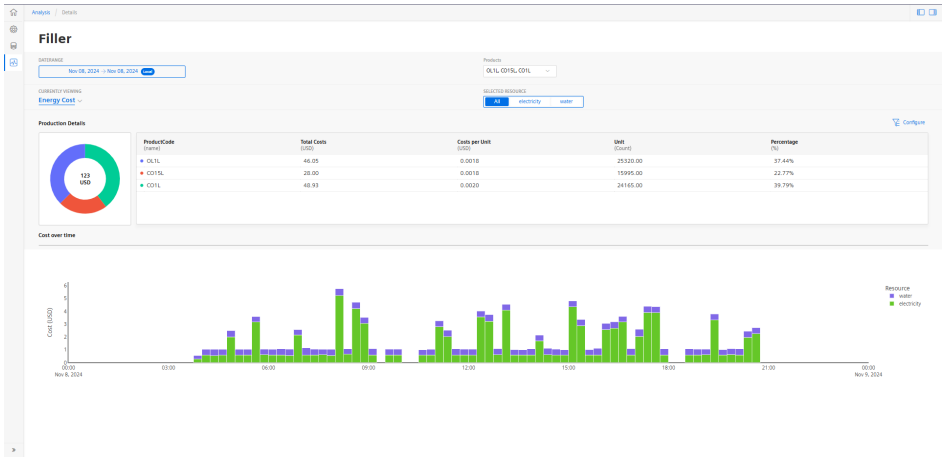

Production Details

- This section includes a donut chart that shows the split-in costs between each product.

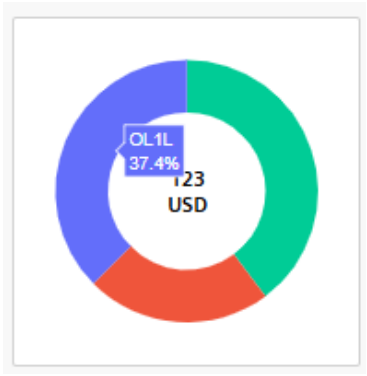

The tooltip shows the percentage of cost for that product.

- It includes a table with configurable KPIs for each product, that includes:
    - Total Costs
    - Cost per unit
    - Unit count
    - Percentage of cost

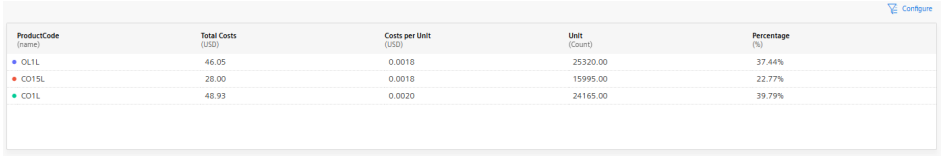

Cost over time

This time chart shows the costs associated with each Energy Resource for each time interval. It is useful for understanding the relative costs of each resource and variation over time.

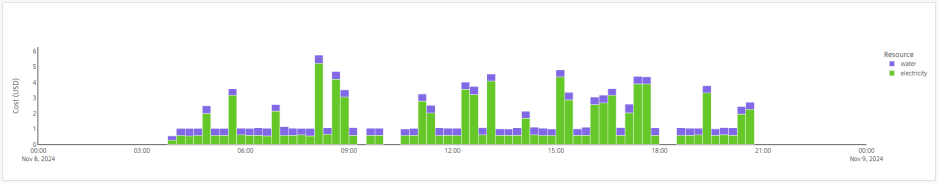

Cost Aggregation

Costs can be aggregated up the Energy Consumer hierarchy. If a parent level consumer is selected, then an "AGGREGATE IN HIERARCHY" slider is available. Activate this and the costs will be added from all the child assets for each Energy Resource.

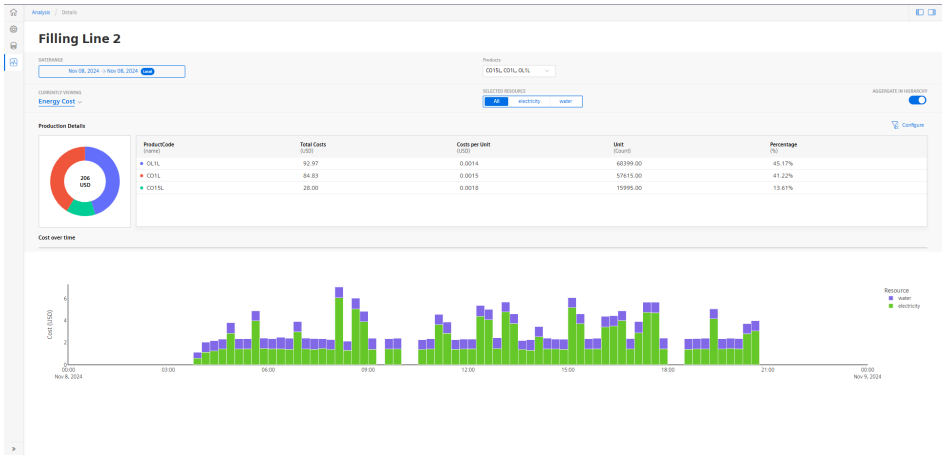

This makes it easy to calculate the costs all the way up the hierarchy to the line and plant level.

6.4 Operator Alerts

Operator Alerts

The alerting feature of Insights Hub Energy Optimizer sends notifications if the energy resource consumption per reference value is higher than it should be. The alert is sent via email from which the operator can directly open the "Operator Alert" view. This directs the operator to the relevant asset and enables an operator or any other staff to react quickly to consumption issues.

Overview

The "Operator Alert" view displays the Alert Data for the selected date range and filters. It shows the Energy Resource data aggregated at the batch level and indicates which batch consumption was recorded as "anomalous". In this case, anomalous means higher than expected.

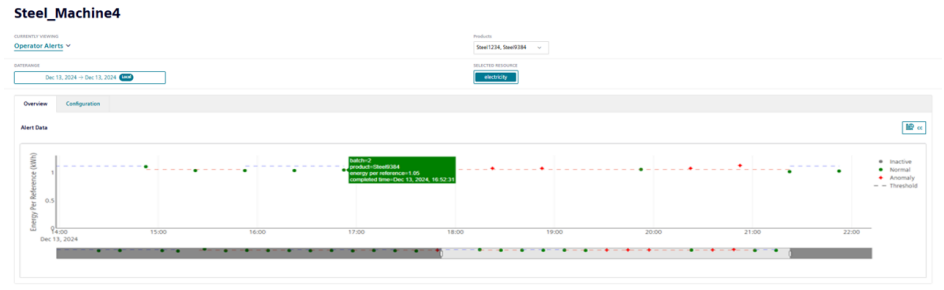

The threshold value is used to determine which batches are considered as anomalous. Values over the threshold are marked as red and those below are in green. The red values would have generated an email alert similar to the following:

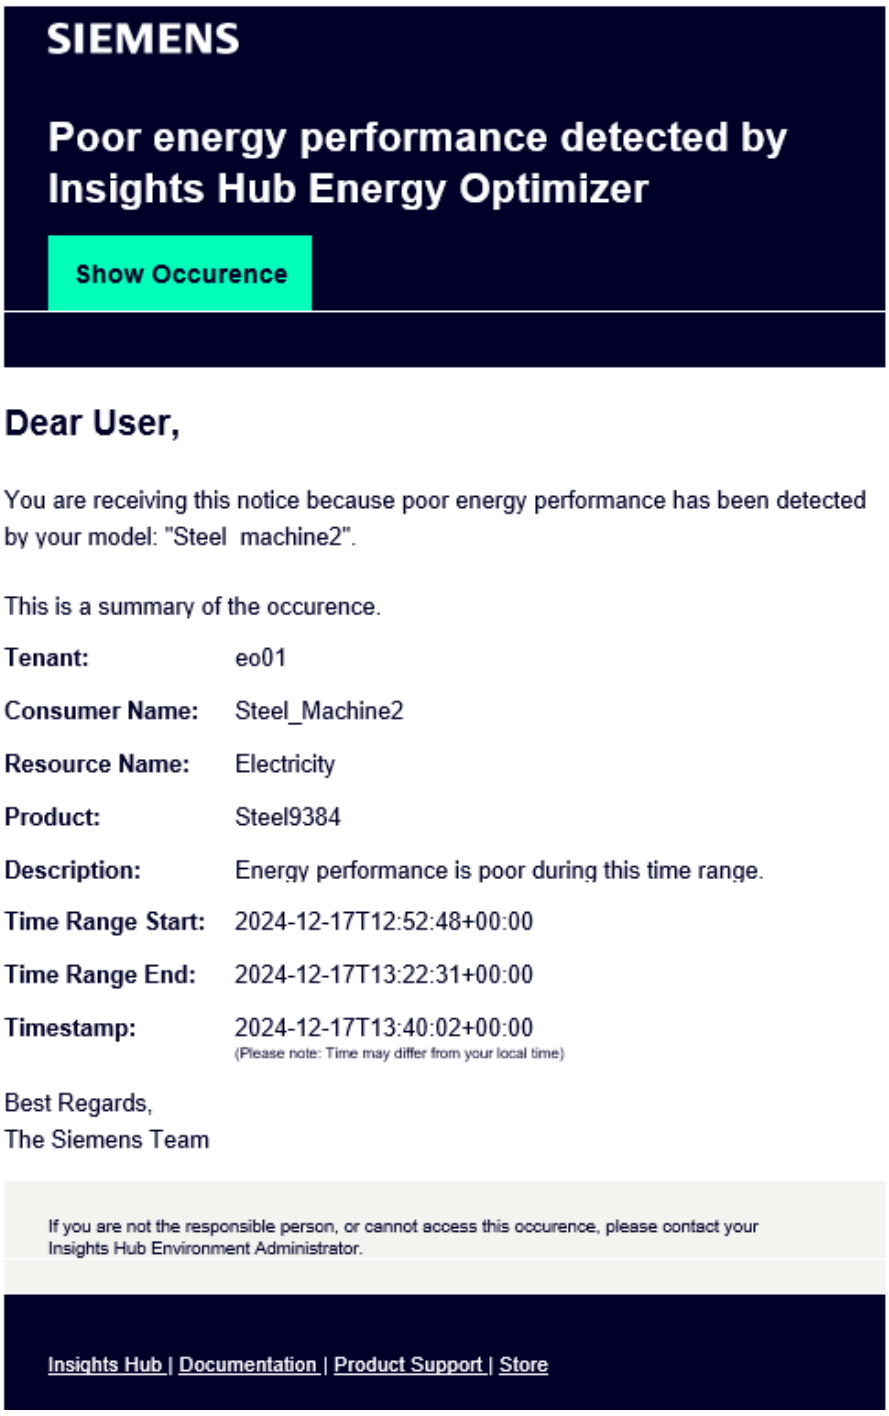

In this email, the details of the associated Energy Consumer and Resource, the current product and the exact time of the anomaly are shared. The "Show Occurence" includes a link to directly access the "Operator Alert view" for that user.

The anomaly threshold value varies depending on the product that is being produced. The thresholds are learned based on the distribution of data in the Active model. The model that is currently running for this energy resource on this energy consumer is shown the above the chart. For more information, refer to Model Config.

The threshold values can be modified by switching to the "Configuration" tab.

Configuration

Histograms are displayed for each product. A vertical line indicates the threshold value for the Energy per Reference for that product. By default, this is set at the 95th percentile of the distribution. These values can be adjusted manually. When the performance standard limit is entered manually, the probability of exceeding the performance limit changes from its default of 5%. Also, an estimate of the number of anomalies that would be detected each day using this value is displayed. This is a useful indication of the number of alerts that would be received each day.

You can setup an email address to send notifications whenever there is an alert.

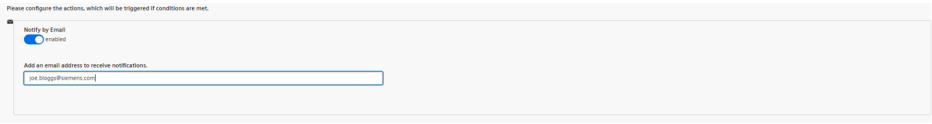

6.5 Process Optimization

Process Optimization

The "Process Optimization" analysis view assists the user to select the optimal variable settings for their process. These are the settings that minimize energy consumption while maintaining the process reference throughput value.

On the first page, a scatter plot of Energy versus Reference is displayed. The values in blue are the training points used in the model building. These are the real observations of energy consumption for each batch of the process. The red points are the best observations in the training set, i.e those batches that minimize energy for a given value of the reference variable. The line connecting these points is called the Pareto front. The objective is to use optimization to find variable settings that do better than the observed Pareto front. This optimization uses a model to predict the energy resource and reference based on the value of the process variables. The active model button is shown above the chart. This leads to the Training Board where the user can review the model accuracy and configuration.

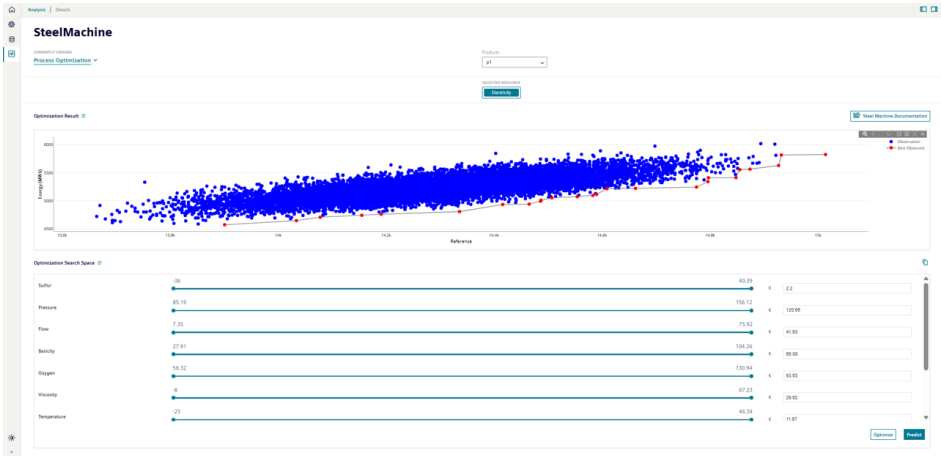

The application runs an optimization algorithm to search the space of variables. The limits on the variables are specified in the Optimization search space. For each variable, there is a slider showing its permissible range. The slider is used to set the upper and lower limits of the variable. There is also an input box where the user can enter a specific value for the variable. If any of these values change, then the "Predict" button is activated. This makes a single prediction of the energy and reference using the model. The result is shown in the chart as a yellow star. The latest prediction value is shown as a bold star, whereas the previous are transparent:

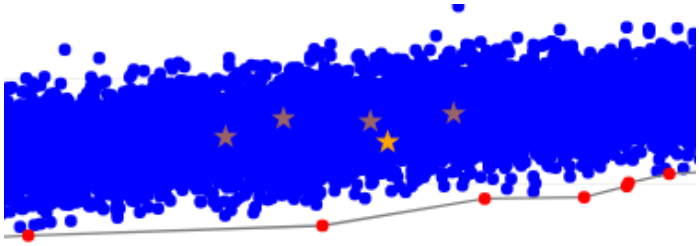

You can make multiple predictions to understand the relation between the process variables and the energy and find the best performing variable settings. However, with the optimization feature, this can be automatically done by clicking the optimize button. The built-in algorithm uses the model to search the variable space. The "Optimization in progress" symbol is displayed while waiting for the job to complete. After completion, the chart shows the optimization results in green.

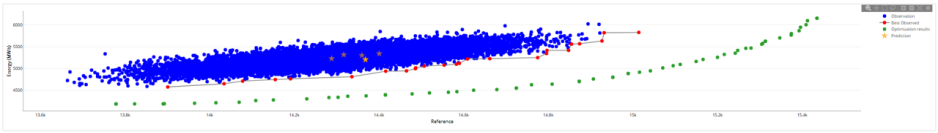

Each point is part of the new optimal set of solutions found by the algorithm. These solutions cover a range of reference values. So, you can select the one that corresponds to the current requirements of your production. For example, for some batches it may be sufficient to produce less than others. You can hover your mouse over the point and the tooltip will show the corresponding variable values. You can also select the point and this will enter the values into the variable input boxes. Click the button and the values will be copied to the your clipboard so that they can be pasted into other applications to help setup their process for the next batch.

By default, the Optimization search space includes the full range of variables used in the model. Instead, you can change the search space using the sliders next to each variable. Drag the bars to set the upper and lower limits of each. You can also use the to fix the space to the value entered in the input box. Click "Optimize" and a new set of results are displayed:

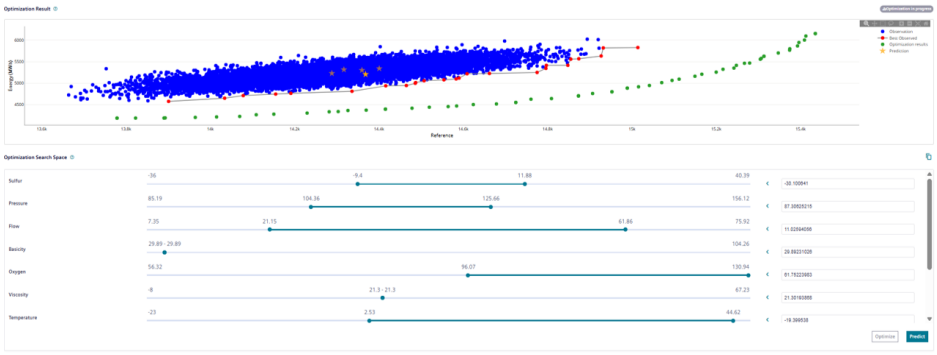

This process can be repeated. However, each new optimization result will overwrite any previous results. The latest optimization result persist when the application is closed and re-opened, but the prediction results are not preserved between sessions.

Models 7

7.1 Introduction to Models

Introduction to Models

Models in Insights Hub Energy Optimizer connects the Energy Resource consumption to production context data. Each model is applied to one Energy Consumer/Energy Resource combination and uses other data on the asset as predictors. In general, models can be used for root cause analysis, process optimization, forecasting and anomaly detection, which helps the user to control and optimize their Energy Resource consumption. The first modelling application available is for anomaly detection of batch or cyclic processes. This feature sends alerts to the user when the Energy per reference value exceeds a pre-defined limit. The model building technology in the application helps the user to define these limits based on learning from past data.

The users with `creator` role can access the modelling area in the application. For more information on assigning this user role, refer to Settings documentation.

The modelling landing page provides an overview of all the models configured on the tenant, as described in the Models Overview section. To build a new model, refer to Building a Model. To review the accuracy and features of an existing model you can use the Model Training Board.

7.2 Models Overview

Models Overview

The models overview page allows you to view and manage all your models on Insights Hub Energy Optimizer. This page can be accessed by clicking the  icon.

On the right pane, the information on the total number of models, plus those in "Trained" and "Active" states, is displayed.

If the model is in the "Trained" state, then this model is completely configured, the data set is loaded and the model has been built. However, the model is not sufficiently accurate to become "Active". Currently, the models need to achieve an R^2 score of 70% to be considered "Active". A model will automatically become "Active" when it meets the accuracy criteria. An "Active" model uses one unit of your model quota.

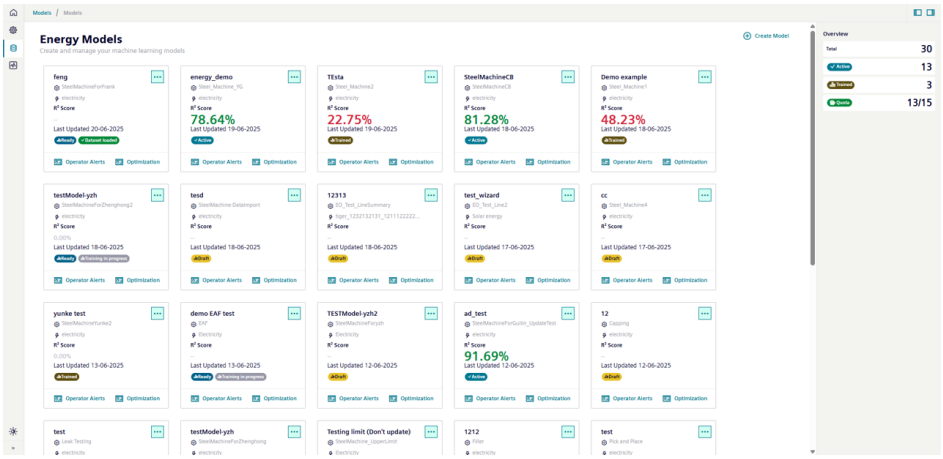

Each model's details is displayed as a model card where, the model name, the Energy Consumer and the Energy Resource details are displayed. In addition, the last updated date and its current status is also displayed. If the model is "Trained" or "Active", then the card shows the R^2 score:

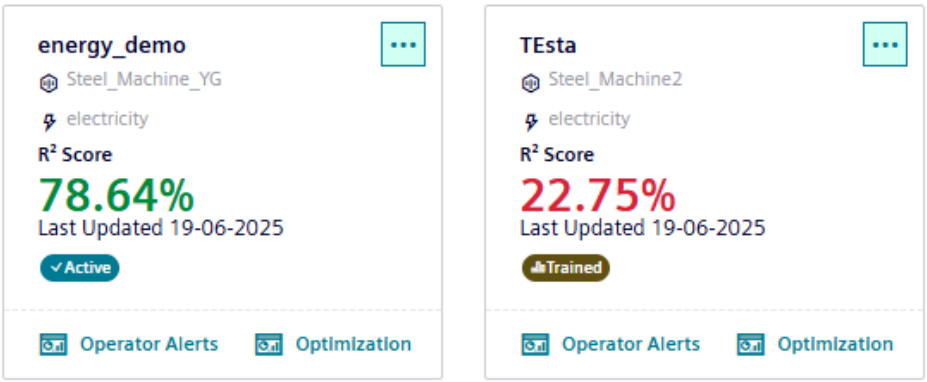

If the card shows "Ready" state, this means that the model is fully configured but has not yet been trained and there will be no R^2 score. In addition, a "Ready" model can be shown with "Dataset loaded" or "Training in progress":

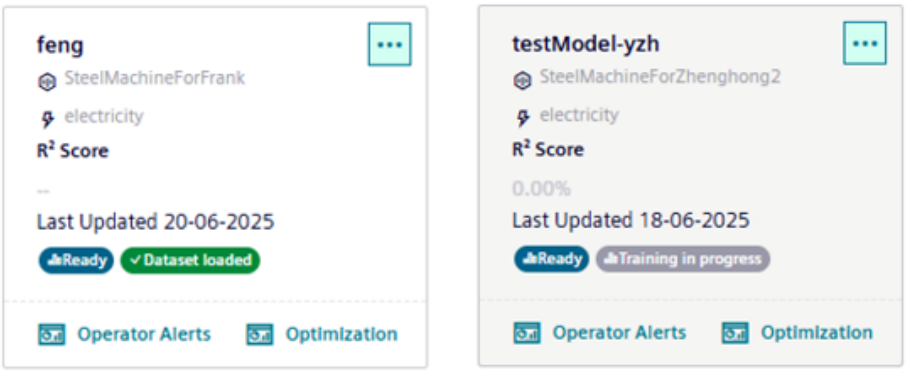

A model can also be in "Draft" state which means that it is not yet fully configured:

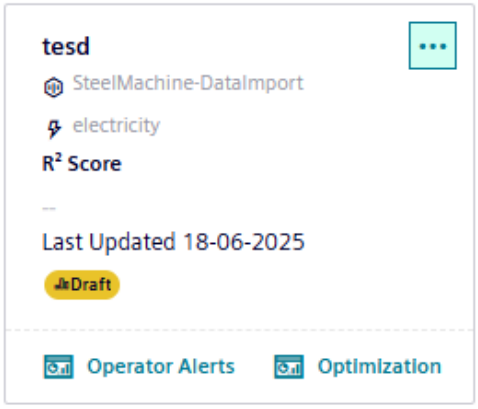

At the bottom of the cards, there are shortcuts to the analysis pages that use the model. Click the Operator Alerts button to navigate to the "Operator Alerts" analysis view for that Energy Consumer or the Optimization button to navigate to the "Process Optimization" view.

Click on the ••• icon to view the possible actions for that model.

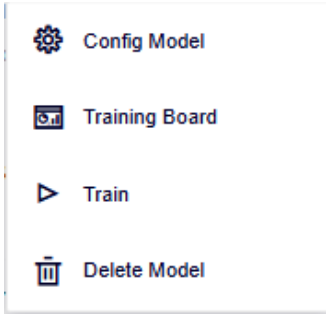

This includes multiple options:

- **Config Model**: This will open the model configuration wizard. This option is enabled for models in the "Draft", "Ready" and "Trained" states.
- **Training Board**: This shows the model build results, including the accuracy validation and feature importance analysis.
- **Train**: This builds the model from the current configuration. This option is available for models in the "Ready" state.
- **Delete Model**: All models can be deleted. If you delete an "Active" model, then this will free one unit of your model quota.

Note that "Active" models cannot be re-configured and so they cannot move to any other state. To remove a model from your "Active" quota it must be deleted.

To review the accuracy, features and configuration of a model, open the model's Training Board.

To create a new model, click the ⊕ Create Model button which takes you to the model building wizard, described in the Model Configuration section.

7.3 Building a Model

Model Setup

Model Setup can be accessed in either of the following ways:

- On the models overview page, click the "Create Model" button, or
- From an existing model card, click the "Config Model" button

"Models" in "Energy Optimizer" can use data both from Insights Hub time series and Integrated Data Lake (IDL). Data from time series is referred to as "Dynamic features", whereas data from IDL are "Batch features". The Model Building Wizard provides a guided workflow to the user to configure the model, select their Dynamic features Batch features. It also includes a step to remove anomalies from the training data. The wizard is split into the following steps:

- Model Information
- Dynamic Features
- Batch Features
- Select Training Data
- Review Training Data
- Configure Model

Model Information

In this section, configure the general details of the model such as the name, description and select the mandatory variables. A model requires the following four variables:

- Energy Consumer
- Energy Resource
- Batch variable
- Reference variable

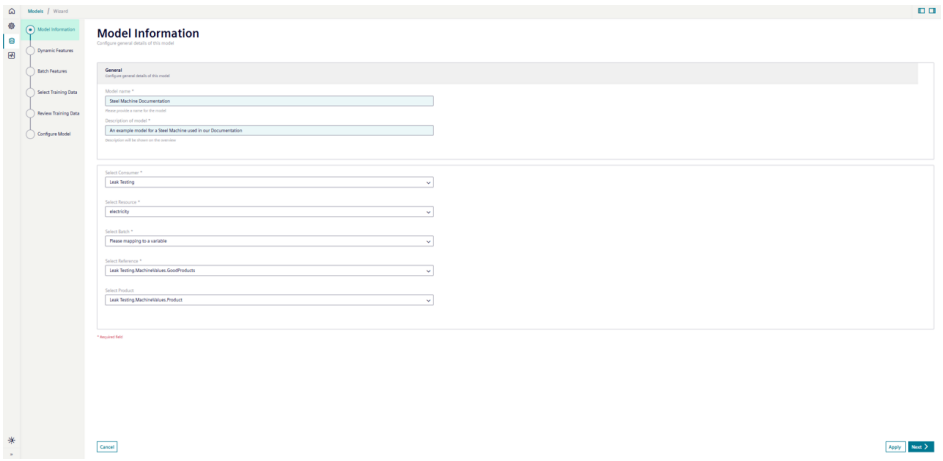

Most of these steps should be auto-filled based on the consumer configuration. After selecting a configured consumer from the first dropdown box, the variables for Batch and Reference will be pre-populated. Then, select the target Resource for the model. The Product variable is optional. For more information on configuring these variables for an energy consumer, see the Configuration.

Only one model can be defined on any Consumer/Resource combination.

When all variables are selected, click "Next" to proceed.

Dynamic Features

Dynamic features come directly from time series. First, select an asset and aspect that includes the Dynamic features for the model. These can come from any asset configured on the tenant, but currently only a single aspect is supported.

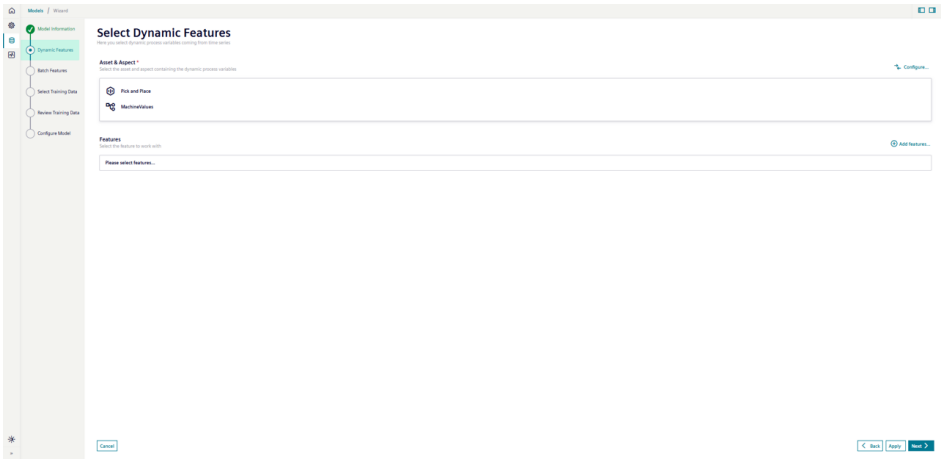

When an aspect is selected, select the individual variables from the pop-up box. Below you will see the list of selected features. In general, Dynamic features could have a timestamp that differs from the batch or energy data, so they must be first aggregated before they can be used in a model to predict the energy resource. A different aggregation method can be used for each feature.

Select the aggregation method such as sum, max, mean from the drop-down list.

Batch Features

Batch features come from a CSV file on IDL. Each row of the file corresponds to process data for a single batch. By definition, they are already aggregated at the batch level. However, each row of the file must also contain a time identifier so that it can be joined with the correct energy resource and dynamic feature data. The time identifier should be written in ISO 8601 format. This time should correspond to the start time for the batch.

Use the data source window to locate the file with the batch data. From here, you can browse all your files on IDL. When you select a file, the file schema configuration is displayed below. Select the checkbox next to each feature to include it in the model. The application will infer the data type for each feature, but there is also the option to change the feature data type. One of the features must be set as the time identifier.

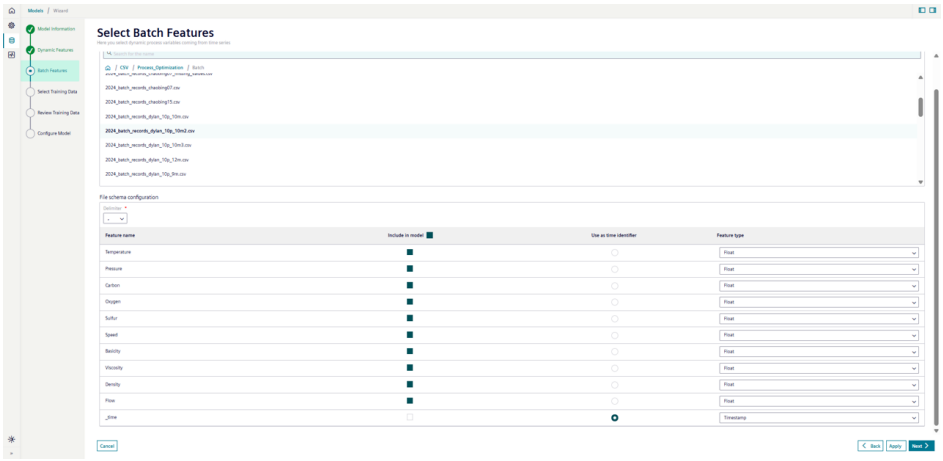

Once all Batch features are selected, click "Next".

Select Training Data

In this step, you can reduce the size of the training data by filtering on the date range. The filtered data will show in the time chart and scatterplot below.

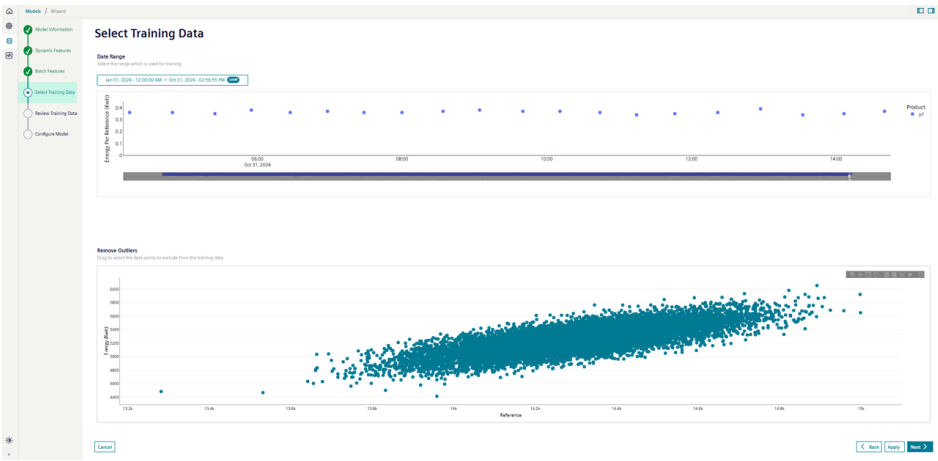

The remove outliers plot can be used to interactively remove anomalies from the data set. After selecting a point on the chart, it will be highlighted in red and will be removed from the training data. Either single point or box-select can be used to flag anomalies. Click on the points again to remove the anomaly flag.

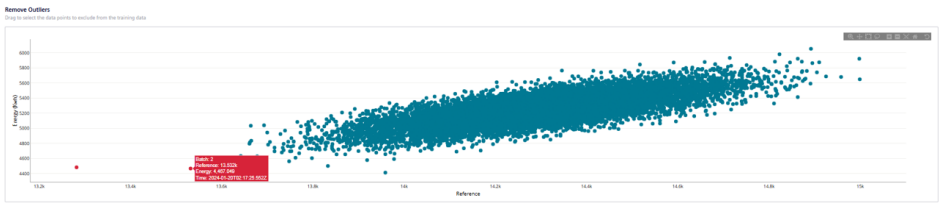

When all anomalies have been identified, move to the next step.

Review Training Data

On this page, the data from both Dynamic and Batch features are assembled into a single table for review. The total number of records and some summary statistics for each feature are displayed at the top.

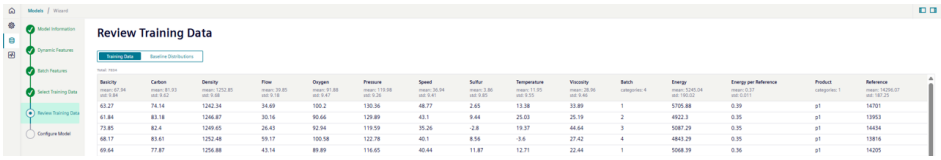

You can also review the distribution of data by switching to the "Baseline Distributions" tab.

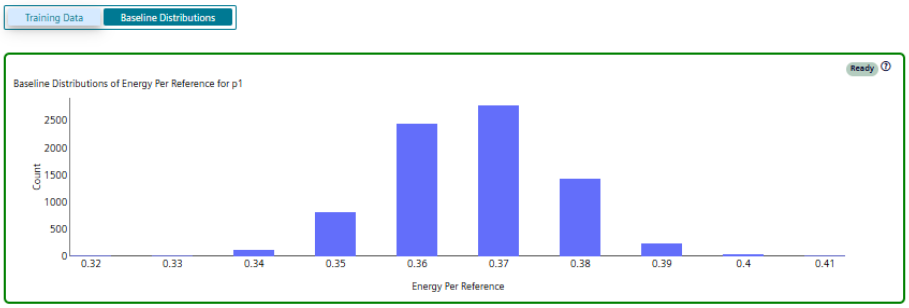

You can review the data here. Use the "Back" button to navigate back to any previous step. When returning to the "Review Training Data" page, the data set will be re-built. When everything is completed, select "Next".

Model Configuration

The final step is to specify whether the model should use the Energy variable or the Energy Per Reference as the target.
Click "Save" to return to the model card overview page and the model you have just created will be in the "Ready" state with "Dataset loaded". Click "Train" from the card menu and the model moves to "Training in progress". The app refreshes every 10 seconds to check whether training is complete.

7.4 Training Board

Training Board

The Training Board is where you can review the accuracy, validation tests, feature importance and configuration. You can also interact with the training data points to understand how they relate to the energy consumption and/or reference value. A Training Board is available for every model, but if the model has not been trained, then it will be displayed as "No Data".
In the Training Board header, you can view the the model name, description, state, the owner and the date range used for the training data. The Training Board contains the following tabs:

- Model Accuracy
- Prediction
- Feature Importance
- Explainability
- Input Data
- Data Analysis
- Configuration

Model Accuracy

This section shows the standard metrics for evaluating model accuracy.

- R^2, sometimes called coefficient of determination, reports the proportion of the variation in the target variable that is predictable.
- Root Mean Square Error measures the magnitude of the typical error that a model makes.
- Mean Absolute Error is similar to Root Mean Square Error.

The page shows these metrics for both the training and test data. In all cases, a 20% proportion of the data is used to create a test data set. The test set is data that the training algorithm has not used to build the model, so the scores on this are a much more reliable test of model performance.
Each model actually consists of two sub-models, one for energy and the other for the reference. Scores are displayed for each model on the separate rows.

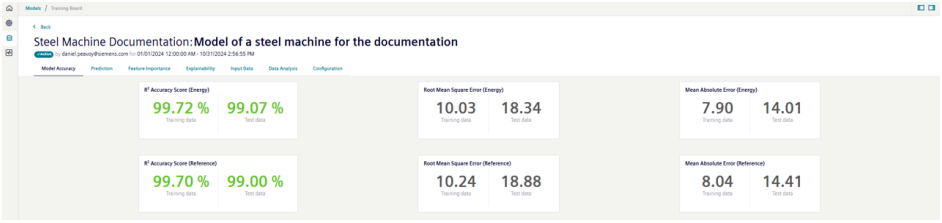

The R^2 on the test data are used to determine if a "Trained" model becomes an "Active" model. By default, they both need to be over 70% to be activated.

Prediction

This section shows the validation plots for both Energy and Reference models. The predicted vs actual chart shows the correlation between the Predicted value and the Actual value. If the model is accurate, then these points should be close to straight or diagonal line. This correlation between the predicted and actual is used to calculate R^2. The plot though gives more information than R^2 as it can be used to determine if there are any ranges of the target variable where the model is performing poorly.

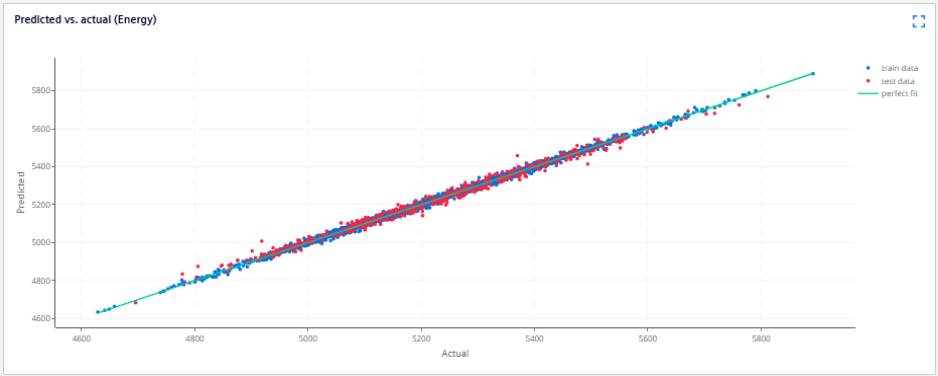

The Delta vs actual chart shows the distribution of residuals between the model prediction and the actual values for both the training and test data sets. A good model should have small residuals, should not be skewed and should be centered around 0. Ideally, the variance in the test data residuals is not a lot larger than for the training data.

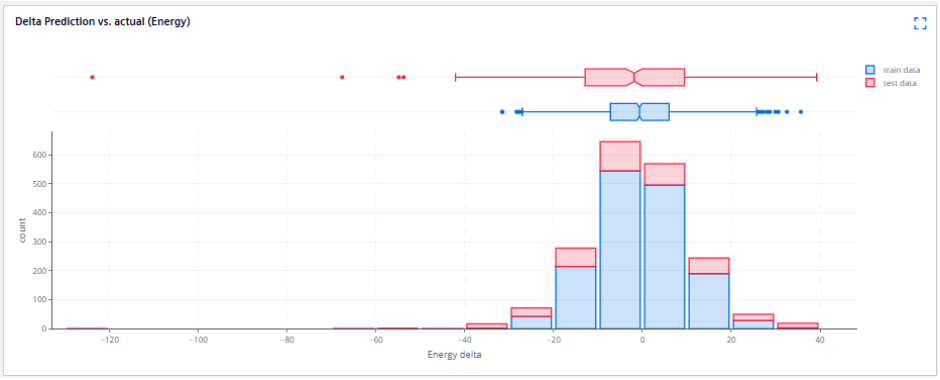

Feature Importance

This section assists the user to select the most important features of the model. They could focus only on these features during the optimization procedure, for example. There are three charts:

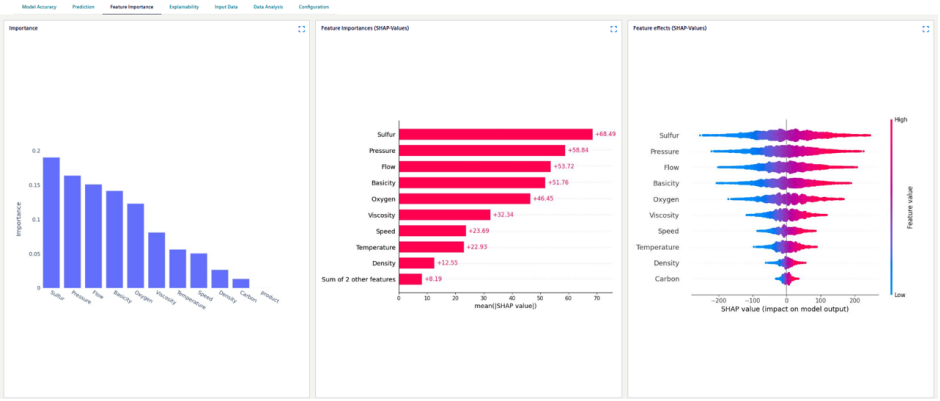

The first two charts show the relative importance of the features using standard measures such as SHAP values. Ideally, these two charts should agree on the order of importance. The third chart shows the feature effects. These indicate in which direction the features influence the model target. For example, in the chart above, the value of 'Sulfur' is positively correlated with the model output.

Explainability

The Explainability tab provides an interactive parallel coordinate plot for the user to see how the features are related to the models targets. The training data is displayed both in the chart and the table. The user can select a row in the table to see all the feature values and corresponding energy and reference values. Alternatively, the user can select a point directly in the plot.

The user can also box select a range for any of the variables on the plot. This could be useful if, for example, they would like to see the variable values corresponding to a range of energy and reference values.

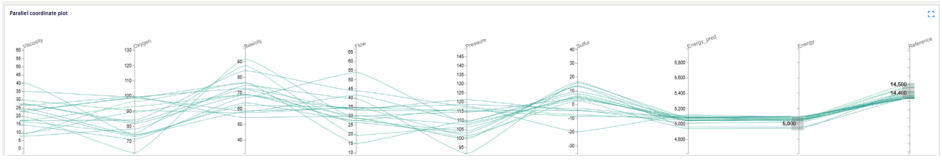

These plots assist the user find the best results within their training data. To identify new process settings that minimize energy, then the model can be used in the Process Optimization.

Input Data

These charts show all of the training data plotted as time series. The first chart shows the features, and the second both energy and reference targets. The user can zoom into particular time ranges to view the data in detail. This could be useful to see any particular time dependency or patterns in the training dat that could be affecting the model performance.

Data Analysis

This section shows the training data analysis results plotted as histograms, scatterplots and a correlation heatmap. The user can use these plots to detect patterns and correlations between the features and model targets. Each plot can be downloaded as a jpg by clicking on the training-board-analysis-download.

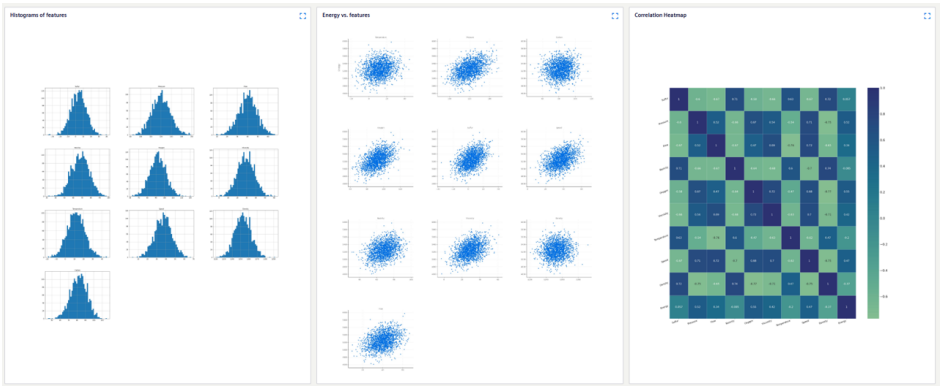

Configuration

The configuration includes all information on the features, target and training data that are needed to reproduce this model.

System Behavior 8

System Behavior

The application aggregates energy data from Insights Hub time series into 15 minute intervals after a waiting time of 5 minutes to allow for data delay. For example, energy data for the period 01:00 - 01:15 will not be processed in the app until at least 01:20. There could also be an additional delay of 5 minutes due to compute load balancing, delaying the data to 01:25.
If data is delayed so that there is no data in an aggregation window, then the application will re-try for these intervals only after new data arrives in subsequent intervals.
The aggregation will resolve changes in value for context data, for example the machine status. Also, when building models using energy data aggregated into batches, then the 15 minute interval is further sub-divided to record the batch level data and there is no loss of information.
When first building an Energy Model, up to 6 months of historical data can be imported and aggregated from time series.

Changelog

9

Insights Hub Energy Optimizer Release Changelog

In this section, you can find release related changelogs for the Insights Hub Energy Optimizer application.

EU1 Releases

June 2025

With this update we are releasing the Process Optimization feature. This includes,

- Changes to the Model Configuration Wizard. You can now select process variables from both time series and IDL, to build predictive models.
- A Model Training Board used to review the model validation results and model feature importances.
- A Process Optimization page with inputs for model parameter ranges and values.
- Automated model building, prediction and optimization algorithms.

January 2025

This is the first General Availability release including,

- Energy Resource, Consumer and Status Mapping configuration pages
- Energy Efficiency, Energy Cost and Operator Alerts analysis pages
- Model building wizard
- Baseline data development and model building for anomaly detection
- Notifications of energy performance below pre-defined standards