

Next-generation Digital Health:

Automating Machine Learning
(ML) Workflows in Digital Health
with Edge Impulse



Table of Contents

Introduction	1
Edge Impulse - 1	
Overview	2
About The Dataset	3
1. Collecting and Organizing Data	4
Data - 5	
Upload Portal - 5	
Option 1: Upload your data to Edge Impulse hosted bucket (a.k.a Upload Portal) - 5	
Step 1: Create a new upload portal - 5	
Step 2: Drag and drop your files -6	
Step 3: Follow the link to view your upload portal - 6	
Option 2: Upload your data to S3/GCP or S3 compatible bucket - 6	
Go to "Data" → Buckets and add a new bucket - 6	
2. Data Cleaning and Transformation	7
Step 1: Pre-process unclean raw data by pushing the pre-processing code to Edge Impulse in form of a Transformation Block - 7	
Step 2: Navigate to Custom blocks tab to view the pushed block in your Edge Impulse organization - 8	
Step 3: Under your organization, Navigate to "Data" and create new dataset - 9	
Step 4: Mount the S3 bucket or mount the Upload portal wherever your raw data resides - 9	
Step 5: Under your organization, navigate to "Data Transformation" tab and create a job - 10	
Step 6: Receive email with status of the transformation job - 11	
Step 7: Under your organization, navigate to "Projects" and create new project - 11	
Step 8: Set up Data pipeline #1 (from S3 into project) - 12	
Step 9: Navigate to "Projects" tab → "Data acquisition" - 13	

3. Feature Engineering and Model Training	14
Create an impulse - 14	
Input block - 14	
Processing block - 14	
EON Tuner for model optimization - 15	
Learning block - 16	
Step 1: Navigate to “Impulse design” → “Create impulse” - 17	
Step 2: Navigate to “Raw data” and “Generate features” - 20	
Step 3: Navigate to “Classifier” and start Training - 20	
Step 4: Navigate to the “Model testing” tab to test the trained model on unseen data - 22	
4. Model Deployment	23
Deployment block - 23	
EON Compiler (optional) - 24	
5. Data Pipelines	25
Retrain model & versioning - 25	
Step 1: Navigate to Project → Versioning tab → store your current project version - 25	
Data Pipelines - 26	
Step 2: Creating data pipeline #2 (model re-training to deployment) - 27	
Data Campaigns (optional) - 29	
Step 1: Setting up your dashboard - 29	
Step 2: Setting up your campaign - 30	
Step 3: Fill the form to create your campaign - 31	
6. Data Visualization and Exploration Capabilities	33
Step 1: Navigate to Project → Data acquisition → Data explorer to visualize the dataset - 33	
7. Collaboration Capabilities	34
Step 1: Navigate to Project → Dashboard → “Add Readme” and update the project information - 34	
Step 2: From the dashboard, add any collaborators you wish to add to the project - 34	
8. Security and Compliance	35
Summary	36
About The Author	37

Introduction

[Edge Impulse](#) is a cloud based machine learning platform for developing edge AI models for any hardware, from embedded MCUs to powerful CPUs and GPUs. It's low- to no-code interface provides seamless integration of the machine learning workflow from data collection to model optimization to hardware profiling and deployment. Edge Impulse (noted as EI throughout this document) offers a comprehensive solution that caters to various stages of a typical machine learning (ML) pipeline. It resembles the toolkit of a skilled architect designed for building skyscrapers. Just as an architect starts with a foundation, builds upwards, adjusts to challenges, and integrates new designs and technologies over time, EI provides the tools to start, adapt, and refine ML solutions, ensuring they're robust, current, and optimized for your data.


With a focus on data collection, cleaning, feature extraction, model training, testing, and deployment, Edge Impulse ensures that users have access to the necessary tools at each step. By seamlessly connecting these components through a unified software platform, users can efficiently manage the entire ML workflow from data ingestion to model deployment. Beyond its integrated approach, Edge Impulse provides APIs and SDKs, enabling users to customize their workflows and integrate with external tools when needed. Moreover, Edge Impulse acknowledges the importance of collaboration, incorporating collaboration tools while maintaining a strong focus on security and compliance, particularly relevant for teams working in the digital health industry. With these features, Edge Impulse empowers both beginners and experts to build and deploy ML models for the edge effectively.



In the upcoming sections of this guide, we will delve deeper into the features and capabilities that Edge Impulse offers, to help you understand and gain insights on how to use it effectively for your digital health ML applications.

Overview

This tutorial serves as a comprehensive guide for transforming your digital health Machine Learning (ML) applications, taking advantage of the enterprise-grade capabilities of the Edge Impulse (EI) platform.

We'll guide you step-by-step through the processes of collecting, integrating, transforming, visualizing, and engineering a time-series dataset. To round things off, we'll show you how to train, optimize, and deploy your trained model on the hardware of your choice!

Each segment provides an overview of the required EI components, which are highlighted with this  icon. This will help you comprehend what these elements are, why they are important, and how they fit in the Edge AI platform pipeline that we are building.

To aid understanding, each part of this guide is supplemented with a diagram , demonstrating the flow of inputs and outputs through these elements. We also provide screenshots  demonstrating the configuration steps along with concise explanations within each section.

Plus, for an in-depth understanding, we've provided links to our comprehensive documentation .



About the Dataset

Collecting information through brainwave analysis is a prevalent and crucial application in the digital healthcare sector. In this guide, the dataset includes EEG signal data from 10 college students, obtained while they watched a series of 20 MOOC video clips, categorized into non-confusing (basic algebra or geometry) and potentially confusing topics (Quantum Mechanics, Stem Cell Research). Each participant wore a single-channel wireless MindSet, measuring activity over the frontal lobe, and after viewing, they rated their confusion on a 1-7 scale. These subjective ratings were normalized to create binary confusion labels - 0 (not confused) and 1 (confused).

This transformed dataset is then stored into a new dataset called “HealWell_Customer_Dataset” and later fed into a project called “HealWell_customer_Project” to train the model and make predictions on unseen data.

Let’s begin by exploring ways to collect and organize your data using Edge Impulse.

🔄 Here is a typical high level ML workflow in EI.

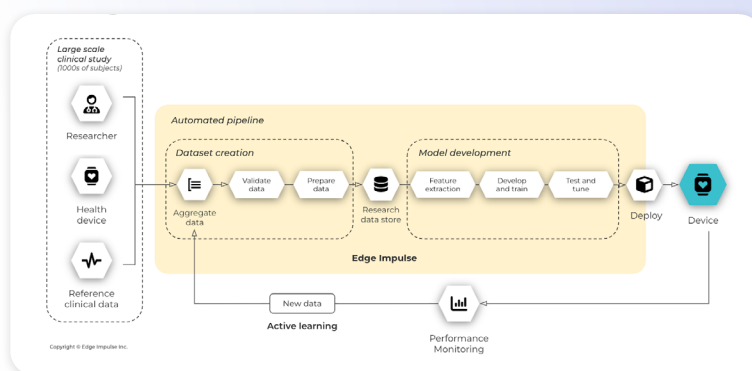


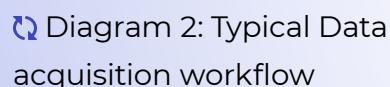
Diagram 1: High level ML workflow in EI

1. Collecting and Organizing Data

The process of creating a machine learning model begins with data. This data spans across wearable devices, electronic health records, medical imaging systems, patient surveys etc. Each data source presents its own unique data format, structure, and challenges.

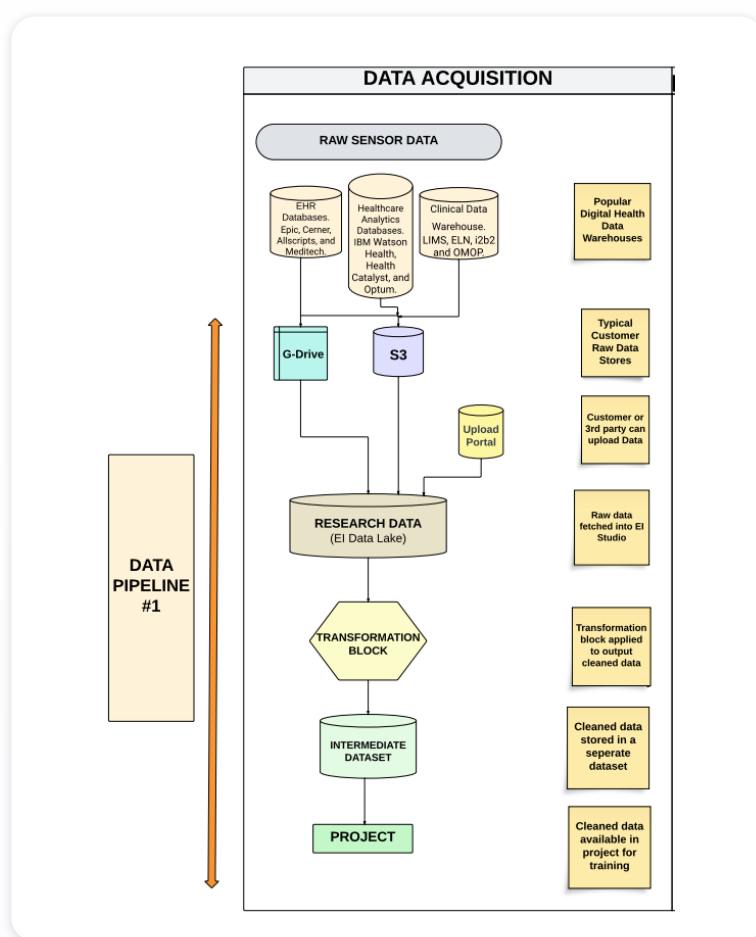
Edge Impulse's **Data** component (part of the data acquisition pipeline) is designed to ingest data from a variety of sources. This facilitates the collection of diverse, real-world data needed to train robust models.

Additionally, users can collect and store training and test data alongside their model and deployment code. Rather than relying on prebuilt datasets or requiring users to construct their own data gathering technology, Edge Impulse offers a variety of methods to gather data in real-world environments.

 Diagram 2: Typical Data acquisition workflow

[Data acquisition](#)

Note: “Research Data” in the screenshot below refers to the “Data” tab within EI UI.



Data

Often, data scientists and ML engineers like to work with a slice of the unprocessed data for preliminary tasks such as data quality checks and experimental trials. The “Data” component within EI organization stores the raw incoming data from your s3 bucket or a 3rd party research partner in a single location within EI. This data lake gives data scientists and ML engineers the freedom to segment the raw data as they please, using a query similar to SQL.

Data from your s3 bucket and/or data from a 3rd party research partner uploaded to an EI hosted bucket (a.k.a Upload Portal) feeds into a “dataset” within “Data” in your EI organization.

Upon preprocessing, semi-clean or clean data is outputted. The semi-clean data can be stored in a new or existing dataset for further cleaning or fed into a project for data exploration and/or conducting preliminary experiments. Alternatively, clean data can be fed into a project for model training.

*“Data” is located under the “Organization” tab within your EI account.
Follow the steps below to configure the data sources.*

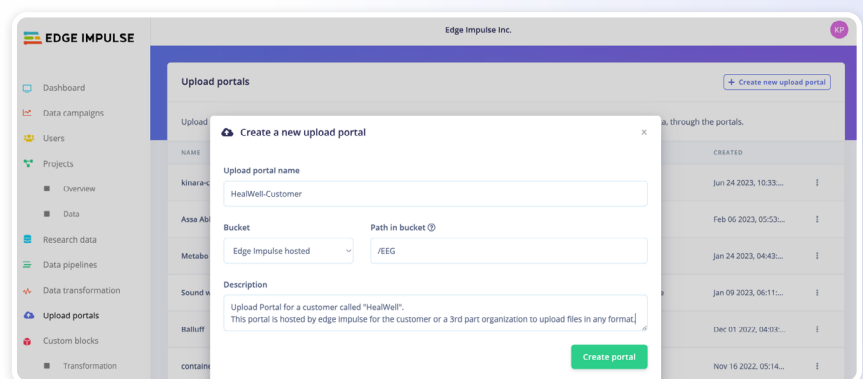
Upload Portal

Upload portals provide a secure way to let external parties upload data to your s3 bucket. Through an upload portal they get an easy user interface to add data, but they have no access to the content of the dataset, nor can they delete any files. Upload Portals offer a friendly user interface and does not require a technical person to use it, thus providing an easy way to use the data directly in Edge Impulse.

[Upload Portals](#)

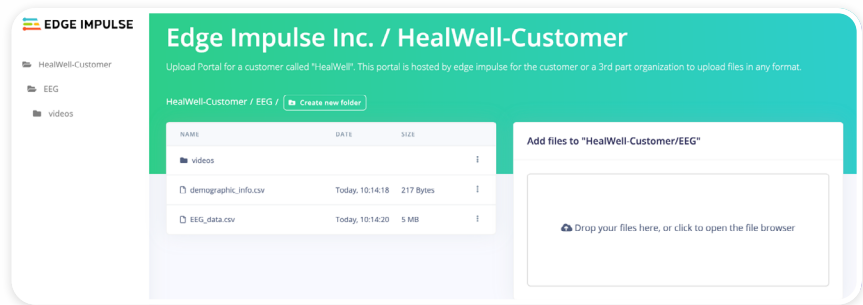
Option 1: Upload your data to Edge Impulse hosted bucket (a.k.a Upload Portal)

Step 1: Create a new upload portal




Step 2: Drag and drop your files

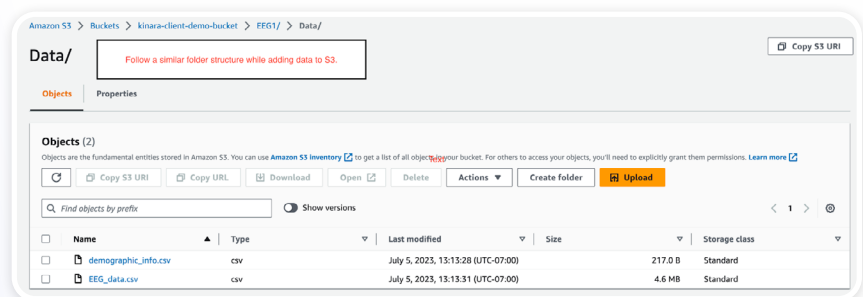
Sample dataset with details on the use case of the dataset are available [here](#).



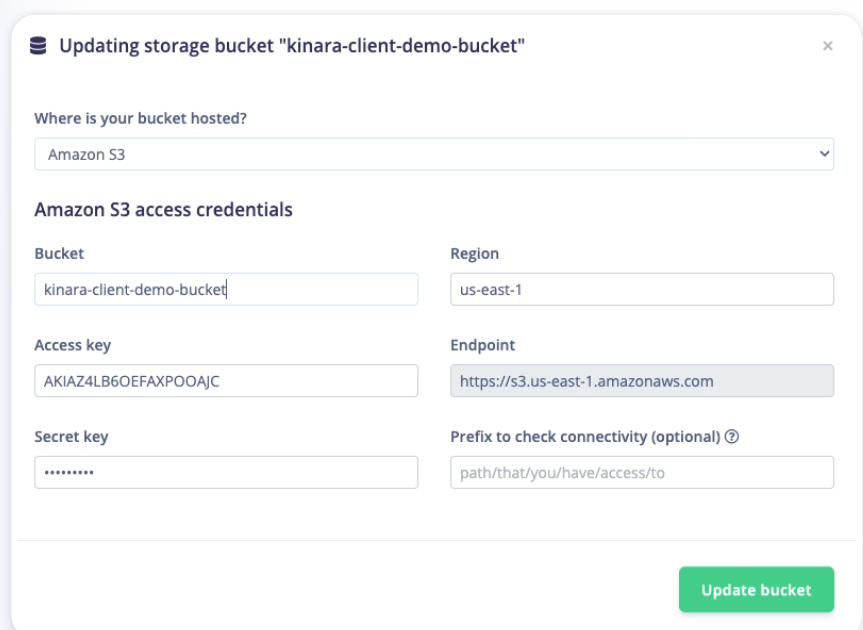
Step 3: Follow the link to view your upload portal



Option 2: Upload your data to S3/GCP or S3 compatible bucket



Go to “Data” → Buckets and add a new bucket



2. Data Cleaning and Transformation

Transformation Blocks

Raw data often arrives laden with errors, inconsistencies, and missing values, making it a daunting task to ensure data integrity. These errors are often detected at a later stage, wasting precious time and resources invested during data collection and data transfer.

To address these issues, we'll use the “*Transformation blocks*” feature of EI to implement comprehensive custom *data cleaning processes, validate data accuracy* and much more. By utilizing the power of parallelization, to ensure rapid discovery of potential errors, right at the source, these custom blocks save valuable time and resources.

Transformation blocks take raw data from the previous section and transform the data into a different dataset or files that can be loaded in an Edge Impulse project. In this tutorial, we use Transformation blocks to *load data from the csv files, merge dataframes, standardize the features, save the resulting data frame in EI compatible format etc.*

Transformation blocks are located under the “Organization” tab in your EI account. Follow the steps below to configure one as you progress through this tutorial.

[Additional documentation](#)

Step 1: Pre-process unclean raw data by pushing the pre-processing code to Edge Impulse in form of a Transformation Block

Prerequisite: Install [Docker](#) depending on your OS.

Open a new terminal and follow the steps in the screenshot below. You may replace the sample code within “transform.py” with your own custom code or clone [this](#) repository to use the same dataset and code within “transform.py”.

Next, push your new block to EI using command “*edge-impulse-blocks push*”

Example custom transformation
block code - [Github](#)

```
kinarapandya@kinarag2e79d4j transforms % edge-impulse-blocks init
Edge Impulse Blocks v1.19.2
WARN: You're running an outdated version of the Edge Impulse CLI tools
Upgrade via 'npm update -g edge-impulse-cli'
? In which organization do you want to create this block? Edge Impulse Inc.
Attaching block to organization 'Edge Impulse Inc.'
? Choose a type of block Transformation block
? Choose an option Create a new block
? Enter the name of your block eeg_data_cleaning
? Enter the description of your block merges eeg_data and demographic_info csv files, performs data cleaning on it and outputs a single cleaned csv file.
? What type of data does this block operate on? Data item (--in-directory passed into the block)
? Which buckets do you want to mount into this block (will be mounted under /mnt/s3fs/BUCKET_NAME, you can change these mount points in the Studio)?
? Would you like to download and load the example repository? yes
Template repository fetched!
Your new block 'eeg_data_cleaning' has been created in '/Users/kinarapandya/PycharmProjects/health_solution_guide/transforms'.
When you have finished building your transform block, run 'edge-impulse-blocks push' to update the block in Edge Impulse.
```

Step 2: Navigate to
Custom blocks tab
to view the pushed
block in your Edge
Impulse organization

Transformation blocks

+ Add new transformation block

When importing data into an Edge Impulse project you can use a transformation block to slice the data into windows, or to extract long running features.

NAME	DOCKER CONTAINER	DESCRIPTION	META...	CLI ARGUMENTS	CREATED
eeg_data_cleaning	eeg_data_cleaning	merges eeg_data and demographic...			Today, 12:57:27

Description

merges eeg_data and demographic_info csv files, performs data cleaning on it and outputs cleaned csv files that feed into the project.

Docker container

Managed through the Edge Impulse CLI, click to edit

CLI Arguments ?

E.g. --hyper-parameter 1

Allow passing in extra CLI arguments ?

☒

Operates on

☐ File (--in-file passed into the block)

☒ Data item (--in-directory passed into the block)

☐ Standalone (runs the container, but no files / data items passed in)

Pass in metadata ?

☒

Compute requests (optional)

Number of CPUs1

Memory (MiB)4096

Compute limits (optional)

Number of CPUs1

Memory (MiB)8192

Additional mount points

Bucket

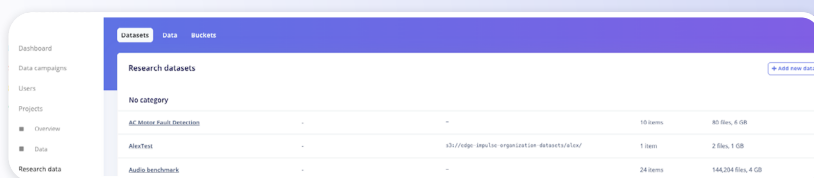
kinara-client-demo-

/EEG1/

Step 3: Under your organization, Navigate to “Data” and create new dataset

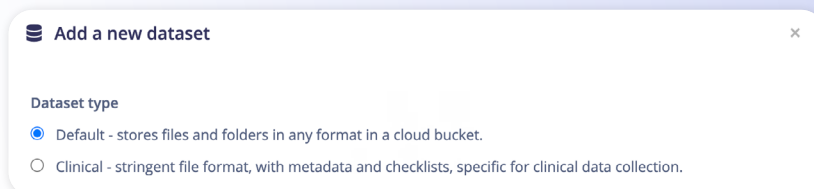
Select dataset type as “default”.
In future revisions, we will go deeper into “clinical” datasets.

Step 4: Mount the S3 bucket or mount the Upload portal wherever your raw data resides



The screenshot shows a table titled 'Research datasets' with columns for dataset name, category, item count, and storage size. The table lists three datasets: 'AC Motor Fault Detection', 'AlexText', and 'Audio benchmark'.

Dataset name	Category	Items	Size
AC Motor Fault Detection	-	10 items	40 MB, 4 GB
AlexText	ch/ridge-impulse-organization-datasets/alex/	1 item	2 files, 1 GB
Audio benchmark	-	24 items	144,204 files, 4 GB

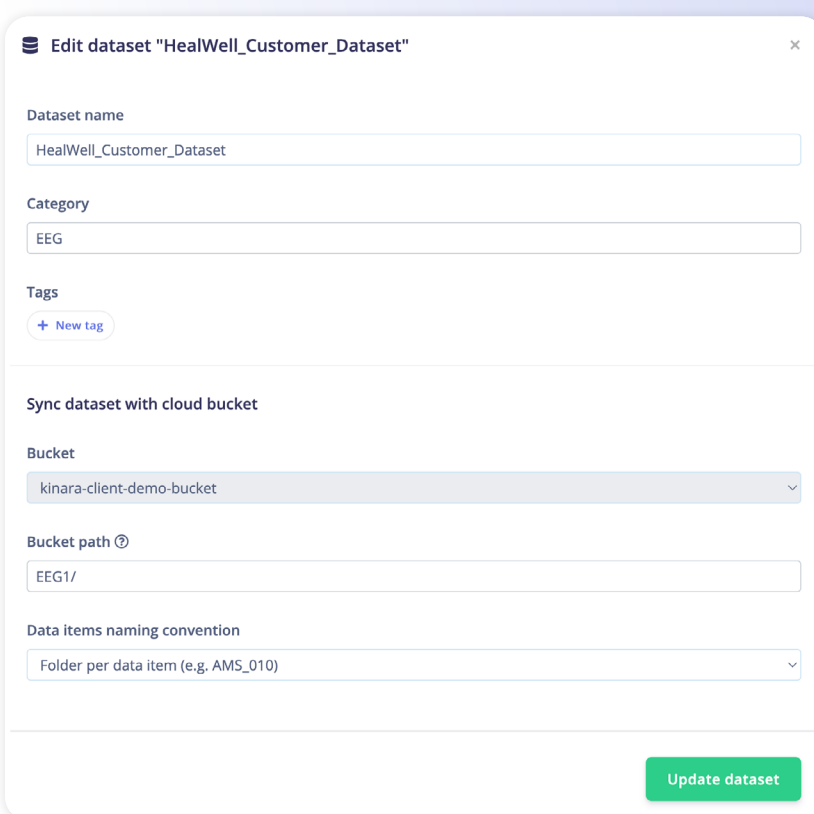


The dialog box titled 'Add a new dataset' shows the 'Dataset type' section with two options: 'Default' (selected) and 'Clinical'.

Add a new dataset

Dataset type

- ☒ Default - stores files and folders in any format in a cloud bucket.
- ☐ Clinical - stringent file format, with metadata and checklists, specific for clinical data collection.



The dialog box titled 'Edit dataset "HealWell_Customer_Dataset"' shows fields for dataset name, category, tags, sync settings, bucket, bucket path, and data items naming convention.

Edit dataset "HealWell_Customer_Dataset"

Dataset name
HealWell_Customer_Dataset

Category
EEG

Tags
[+ New tag](#)

Sync dataset with cloud bucket


Bucket
kinara-client-demo-bucket

Bucket path
EEG1/

Data items naming convention
Folder per data item (e.g. AMS_010)

Update dataset

You should be able to see a similar output as shown below. Here, it is fetching data from the s3 bucket.

 Add a new dataset

Creating job... OK (ID: 12914513)

Scheduling job in cluster...

Container image pulled!

Job started

Fetching folders for s3://kinara-client-demo-bucket/EEG1/

Fetching folders for s3://kinara-client-demo-bucket/EEG1/ OK

Fetching paths for 0 items...

Fetching paths for 0 items OK

Syncing 1 data items with cloud items...

Syncing 1 data items with cloud items OK

Updating data from cloud buckets...

[1/1] Updating data...

Updating data from cloud buckets OK

Refreshing cached datasets...

Refreshing cached datasets OK

Job completed

Cancel

Add dataset

Step 5: Under your organization, navigate to “Data Transformation” tab and create a job

This step creates a transformation job to save the transformed cleaned data into a new dataset called “HealWell customer dataset” created in the previous step.

Create transformation job

Name

Transform_Block_HealWell

Transformation block

eeg_data_cleaning

Extra CLI arguments for the transformation block (optional)

--dataset test

Dataset filter

name LIKE "Data"

Import data into

Project

☒ Dataset

Output dataset

HealWell_Customer_Dataset

Number of parallel jobs

3

Users to notify

Kinara Pandya

Copy as pipeline step

Start transformation job

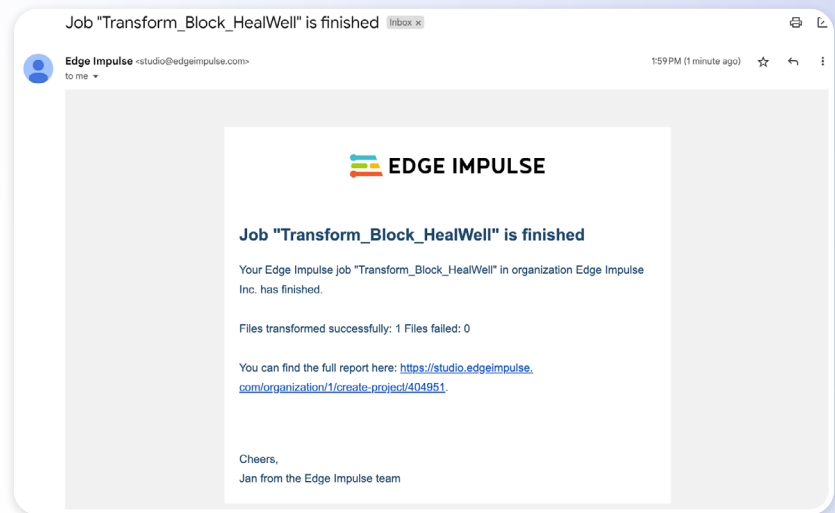
Next, select “start transformation job”. This will make the transformed data available in the Output dataset.

Make sure to click on “Copy as pipeline” and save the code. We will use it later while creating data pipelines in step 10.

Step 6: Receive email with status of the transformation job

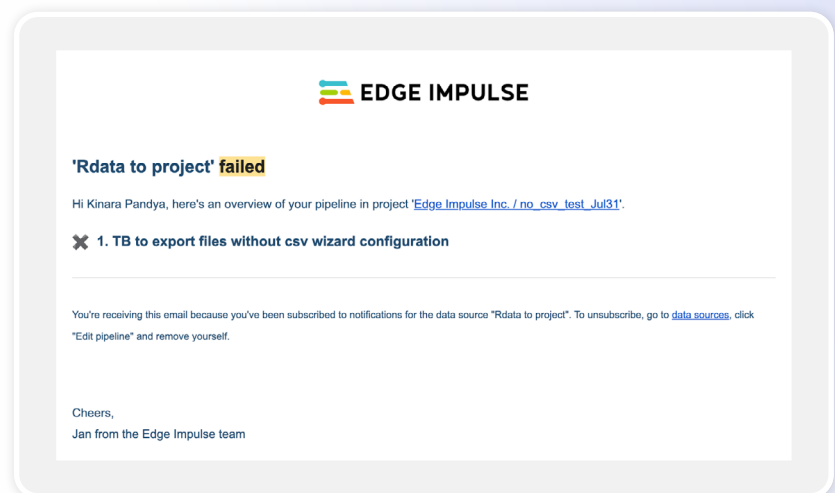
If the transformation was successful, you will receive an email as shown below.

If the transformation was not successful, you will receive a similar email as shown below.



Step 7: Under your organization, navigate to “Projects” and create new project

Quite often, data scientists and ML engineers want to experiment with a slice of data only. Now that we have clean data available in the dataset from the step above, we can feed this data into a project for feature extraction, visualization, model training and generating a C++ library for deployment on any target hardware.



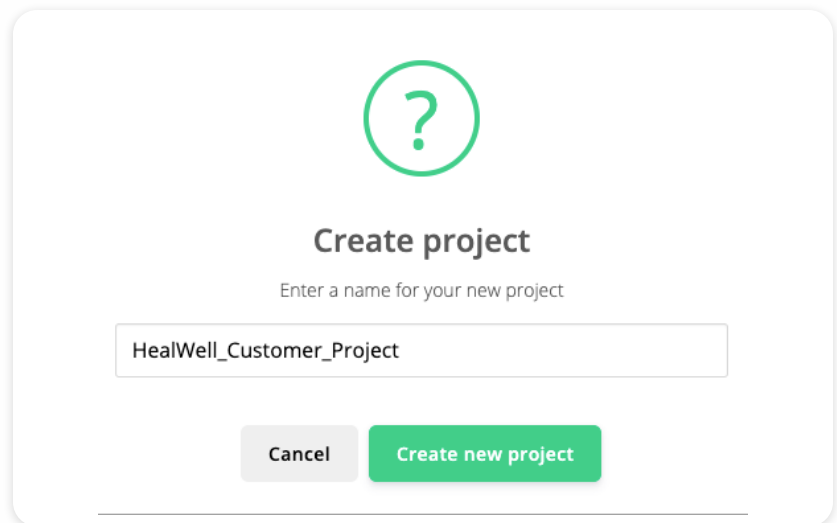
Step 8: Set up Data pipeline #1 (from S3 into project)

Let's automate all the steps until now (i.e fetching data from S3 bucket, running a transformation block every 30 days to clean any new data and feeding it into a dataset and project) using the “Data pipelines” feature within EI.

To configure data pipelines, follow the steps below:

In your organization, under “data pipelines” tab, “add new pipeline”

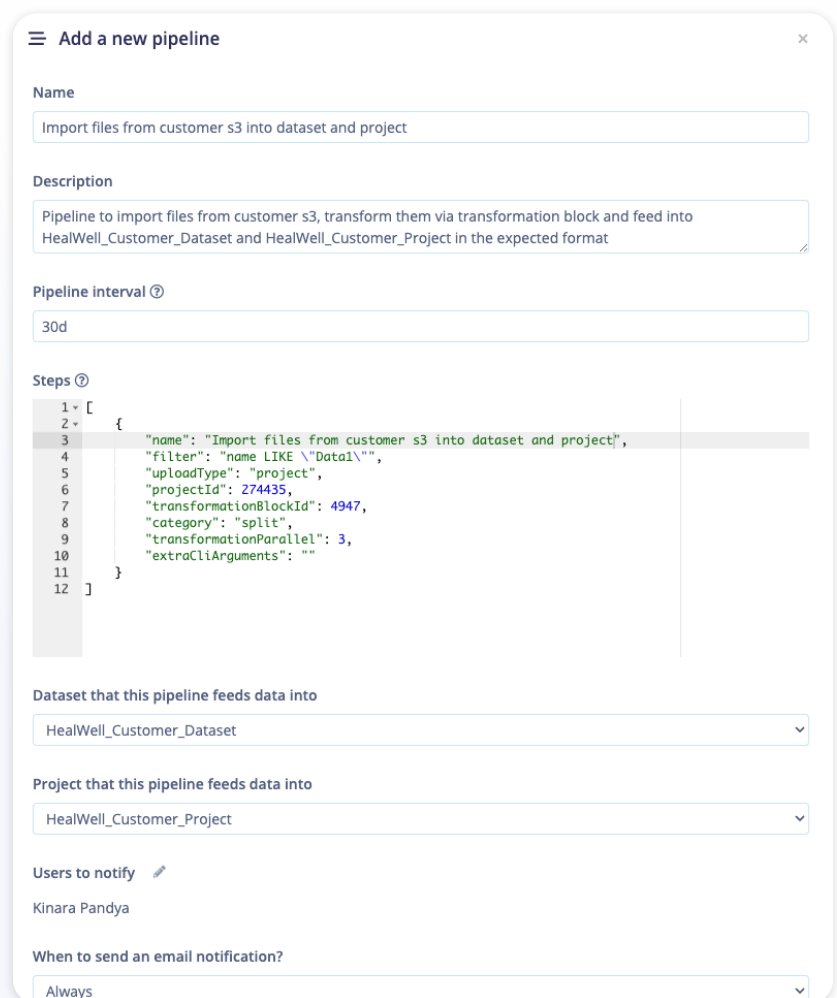
Paste the code saved in step 6 from “Copy as pipeline” step while configuring the new pipeline.



A dialog box titled "Create project" with a green question mark icon. It prompts the user to "Enter a name for your new project". A text input field contains "HealWell_Customer_Project". Below the input are two buttons: "Cancel" and "Create new project".



A header bar for the "Active data pipelines" section. It includes tabs for "Active pipelines" and "Archived pipelines", a "+ Add new pipeline" button, and a descriptive sentence: "Data pipelines periodically run multiple transformation blocks in series. Use them to automatically fetch, process or transform new data."



A form titled "Add a new pipeline" with a close button (X). It contains the following fields:

- Name:** "Import files from customer s3 into dataset and project"
- Description:** "Pipeline to import files from customer s3, transform them via transformation block and feed into HealWell_Customer_Dataset and HealWell_Customer_Project in the expected format"
- Pipeline interval:** "30d"
- Steps:** A code editor showing a JSON configuration:

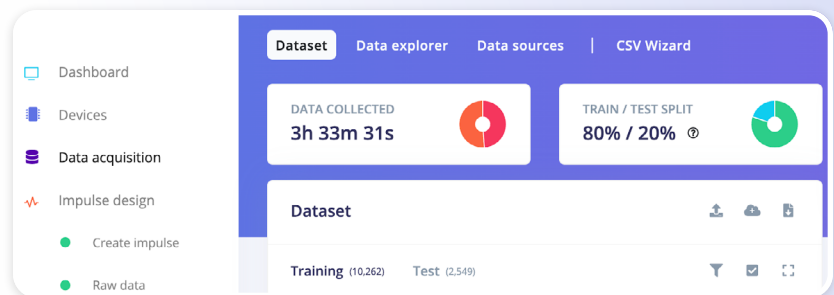
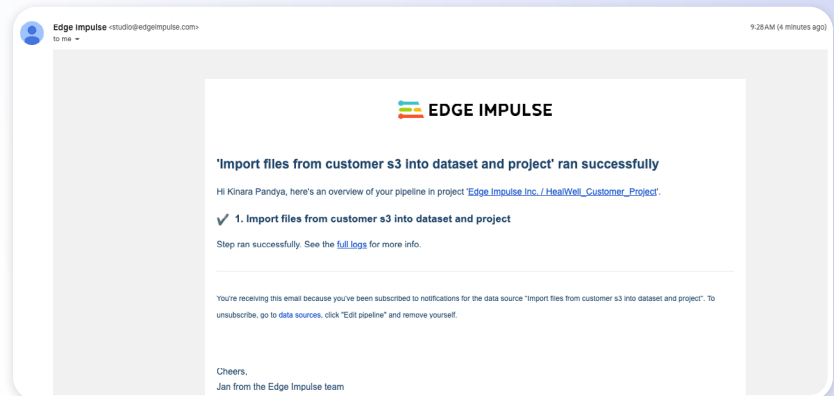

```
1 [
2   {
3     "name": "Import files from customer s3 into dataset and project",
4     "filter": "name LIKE '\\Data1\\'",
5     "uploadType": "project",
6     "projectId": "274435",
7     "transformationBlockId": "4947",
8     "category": "split",
9     "transformationParallel": 3,
10    "extraCliArguments": ""
11  }
12 ]
```
- Dataset that this pipeline feeds data into:** "HealWell_Customer_Dataset"
- Project that this pipeline feeds data into:** "HealWell_Customer_Project"
- Users to notify:** "Kinara Pandya" (with an edit icon)
- When to send an email notification?:** "Always"

Once the pipeline runs successfully, you should see an active pipeline under the “Active pipelines” tab.

Immediately, following completion of the pipeline, you should receive an update on the status of your pipeline via email.

Step 9: Navigate to “Projects” tab → “Data acquisition”

Your cleaned data is now populated into “HealWell_Customer_Project” as seen under the “Data acquisition” tab.



3. Feature Engineering and Model Training

Create an impulse

After the data is cleaned, meaningful features need to be extracted from it for the model to learn the most important patterns within the data. In order to achieve this, we need to first set up/create an impulse.

So what is an “Impulse” ?

A complete “Impulse” consists of 3 main building blocks: *input block*, processing block and a *learning block* mimicking the stages of a standard ML workflow of data input, feature extraction and training a model respectively.

The above capabilities of Edge Impulse help users derive insights from their data, which serve as the input for machine learning models.

Input block

The input block indicates the type of input data you are training your model with. For the purpose of this tutorial, we have chosen time series data.

Processing block

Extracting meaningful features from your data is crucial to building small and reliable machine learning models. In many cases, raw data can be extremely high dimensional, with many variables that may not be particularly informative for the task at hand. Feature extraction allows us to transform this raw data into a lower-dimensional space that captures the most important aspects of the data. It helps remove the noise and highlights only the important signals, making the model more reliable. With smaller, less complex models, computation and storage requirements are reduced which is the key for edge devices, which often have limited computational power and memory.

In Edge Impulse this is done through processing blocks.

A [processing block](#) is basically a feature extractor. It consists of DSP (Digital Signal Processing) operations such as normalization, scaling, flattening etc. that are used to extract features that our model learns on.



For example, to monitor vital signs such as temperature, heart rate, blood glucose, oxygen etc. in real time, a model trained on EI or outside [\(BYOM\)](#) could be deployed to a wearable device such as Polar H10 heart rate sensor, to continuously monitor heart rate. These signals are extracted to reduce noise and feed in the correct features for model training. EI provides a robust built-in **“PPG (Photoplethysmography) to Heart rate Variance (HRV)”** processing block that can be used right away for your digital signal processing needs. If the model detects unusual trends that might indicate a medical emergency, it could automatically alert emergency contacts or healthcare providers.

Another example would be for sleep stage analysis, where data from wearable devices such as smartwatches or [Oura ring](#) could be used to analyze users' sleep stages, helping users understand their sleep patterns. By detecting different sleep stages (e.g., REM, deep sleep, light sleep) using heart rate data from the *“PPG to HRV”* processing block, your application could provide personalized recommendations to improve sleep quality, such as adjusting bedtime routines or sleep environments.

To recap, previously, we used transformation blocks to clean the messy raw data stored in the S3 bucket. For the purpose of this tutorial, we have not used a processing block but rather just used the clean data as features which will feed into the model for training in the next step.

EON Tuner for model optimization

In machine learning, hyperparameters are the variables that govern the training process itself, such as the learning rate or the number of layers in a neural network. These are not learned from the data during training and must be set beforehand. Finding the optimal set of hyperparameters for a model can be a challenging and time-consuming task. It usually involves a lot of trial and error, and even with the same dataset, different sets of hyperparameters can produce very different results.

The EON Tuner (a proprietary algorithm developed by EI) automates this process by evaluating many candidate model architectures (selected based on your target device and latency requirements) concurrently to help you find the best performing architecture for your application.

It helps you find and select the best embedded machine learning model for your application within the constraints of your target device. It analyzes your input data, potential signal

processing blocks, and neural network architectures - and gives you an overview of possible model architectures that will fit your chosen device's latency and memory requirements.

For example, EON Tuner enables sleep monitoring wearables to comprehensively analyze sleep patterns and stages using data from various sensors, such as accelerometers and heart rate monitors. By tailoring the machine learning model to the unique constraints of the wearable sleep monitor, the EON Tuner ensures that users receive personalized insights into their sleep quality, helping them make informed decisions to improve their overall well-being.

Another example of a target device is a small wearable device that continuously records a patient's heart rate data to analyze their cardiac health. However, due to the limited computational resources and power constraints of the wearable, finding the right machine learning model architecture and hyperparameters becomes challenging. The EON Tuner steps in by evaluating various combinations of model architectures and hyperparameters that are suitable for the wearable's hardware limitations.

For time series data (in this tutorial), the tuner would test different sample window sizes and increment amounts.

Depending on the selected dataset category, the EON Tuner considers a variety of [Processing blocks](#) when evaluating model architectures. It then tests the different parameters and configurations of these processing blocks.

In summary, different model architectures, hyper-parameters, and even data augmentation techniques are evaluated by the EON Tuner. The tuner combines these different neural networks with the processing and input options described above, and then compares the end-to-end performance that balances accuracy, memory usage, and processing speed.

Learning block

Model training is the phase where your machine learning model learns to recognize patterns, correlations, and relationships in the input data. In a digital health application, a model might learn to predict health outcomes based on patterns in patient data. For example, consider a model trained to predict blood glucose levels for patients with diabetes. The model would learn from historical data, such as glucose levels, insulin doses, age, and dietary information. By recognizing patterns and correlations in this data, the model can predict blood glucose

levels in real time. This enables patients to proactively manage their diabetes by making informed decisions about insulin doses and dietary choices.

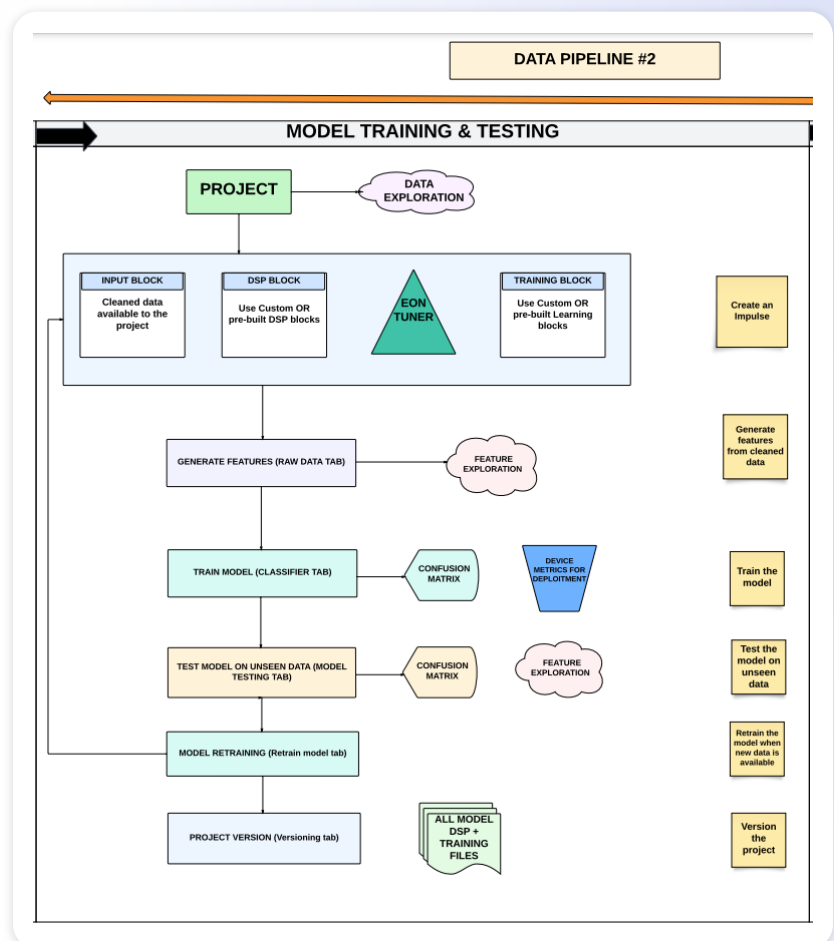
Another example would be, training a model to classify sleep stages based on data from wearable sensors like accelerometers and heart rate monitors. The model learns from a labeled sleep dataset and corresponding sensor data. As the model processes this data, it learns to identify patterns indicative of different sleep stages such as deep sleep, REM sleep, and light sleep. Once trained, the model can automatically classify sleep stages, providing valuable insights into sleep quality for users.

Essentially, without model training, machine learning wouldn't exist. It's the critical process that turns raw data into a tool for prediction and decision-making.

We achieve this within EI by using the Learning blocks. A [learning block](#) is simply a ML model, for e.g a neural network, that is trained to learn on your data. The basic idea is that a neural network classifier will take some input data, and output a probability score that indicates how likely it is that the input data belongs to a particular class.

After adding your [processing block](#) from the previous step, follow the steps below and add a [learning block](#) (mimics the train your model stage in a typical ML workflow) to make your impulse complete.

**Step 1: Navigate to
"Impulse design" →
"Create impulse"**



In this example, the “Input” block comes pre-populated from data acquisition. Configure the DSP (digital signal processing) and “Learning blocks”. They usually display a star to indicate the most recommended blocks based on your input data. You may use the build-in options or add custom blocks if the sensor type or model you wish to use is not supported.

Select “Raw Data” as a “processing block” for the purpose of this tutorial.

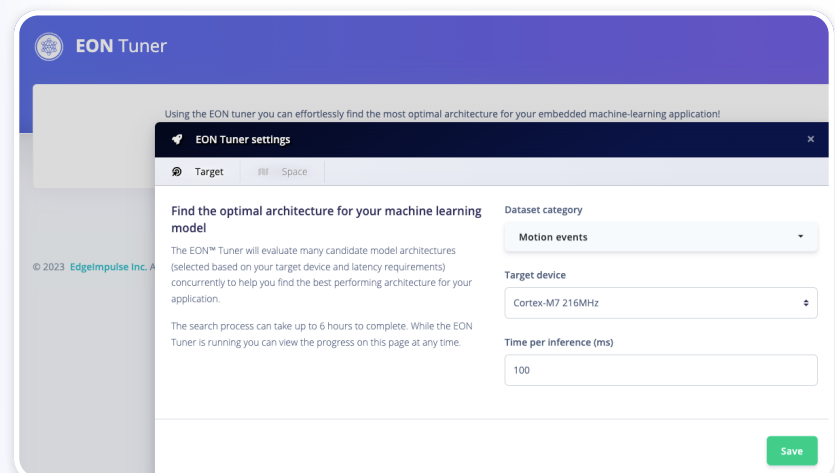
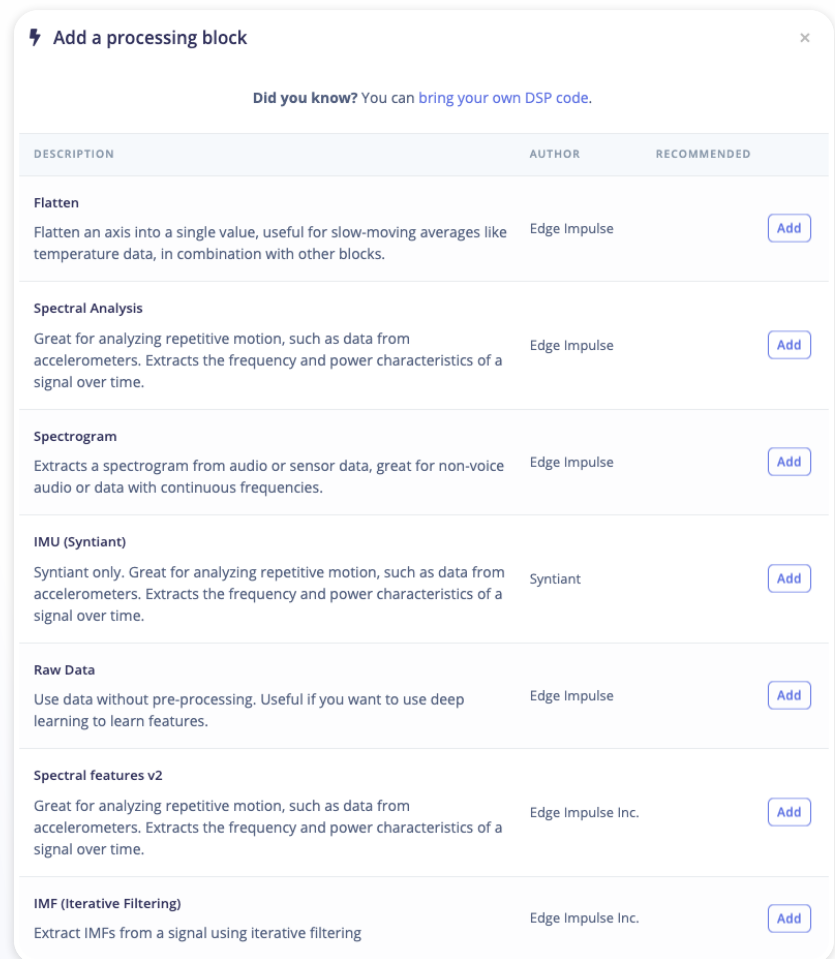
(Optional) EON Tuner:
Finds the optimal architecture for your model

The EON™ Tuner will evaluate many candidate model architectures (selected based on your target device and latency requirements) concurrently to help you find the best performing architecture for your resource constrained device.

The search process can take up to 6 hours to complete. While the EON Tuner is running you can view the progress on this page at any time.

(optional) Navigate to Project
→ EON Tuner tab and select the target hardware

Next, add a learning block.
Select the “Classification” block for this tutorial to predict if the student was confused or not while watching MOOC (Massive Open Online Courses) videos.



Add a learning block

Did you know? You can [bring your own model](#) in PyTorch, Keras or scikit-learn.

DESCRIPTION	AUTHOR	RECOMMENDED	
Regression Learns patterns from data, and can apply these to new data. Great for predicting numeric continuous values.	Edge Impulse	★	Add
Custom regression MLP Example of a regression block	Edge Impulse Inc.	★	Add
EfficientNetRegression EfficientNetRegression	Edge Impulse Inc.	★	Add
Classification Learns patterns from data, and can apply these to new data. Great for categorizing movement or recognizing audio.	Edge Impulse		Add
Anomaly Detection (K-means) Find outliers in new data. Good for recognizing unknown states, and to complement classifiers. Works best with low dimensionality features like the output of the spectral features block.	Edge Impulse		Add
Classification - BrainChip Akida™ Learns patterns from data, and can apply these to new data. Great for categorizing movement or recognizing audio. ONLY FOR: BrainChip AKD1000 MINI PCIe board	BrainChip		Add
PyTorch MLP example (20x10 hidden layers) Custom ML block with PyTorch	Edge Impulse Inc.		Add
MCSA A learning block designed specifically for the MCSA project.	Edge Impulse Inc.		Add

Once the entire flow of the impulse is configured. Save the impulse.

EDGE IMPULSE

Dashboard

Devices

Data acquisition

Impulse design

• Create impulse

• Raw data

• Classifier

• SKNN tuner

• Retrain model

• Live classification

• Model testing

• Versioning

• Deployment

GETTING STARTED

• Documentation

• Forums

Time series data

Input axes (15)

VideoID, Alpha2, Delta, Gamma1, Theta, Beta1, Alpha1, Anomalous, Gamma2, Beta2, SubjectID, MotionID, Raw, streaming, user-defined-label

Window size

1000 ms

Window increase

100 ms

Frequency (Hz)

0.5

Zero-pad data

7

Raw Data

Name

Raw data

Input axes (14)

VideoID

Alpha2

Delta

Gamma1

Theta

Beta1

Alpha1

Anomalous

Gamma2

Beta2

SubjectID

MotionID

Raw

streaming

user-defined-label

Classification

Name

Classifier

Input Features

Raw data

Output Features

2 (6, 1)

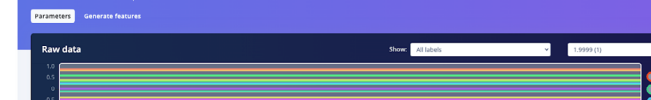

Add a learning block

Output features


1 (6, 1)

Save impulse

The screen also displays approx. on-device performance on a Cortex-M4 80MHz. This may be changed by navigating to the “Devices” tab.



The screenshot shows the Google Cloud AI Platform console for a Hyperparameter Tuning job. The 'Parameters' tab is selected, displaying various metrics and settings. The 'Raw data' section shows a line chart of 'loss' over 100 epochs, with a legend indicating different loss metrics. The 'Raw features' section lists features and their values. The 'DSP result' section shows the 'Processed features' and the 'On-device performance'.



Upon completion of the training, a confusion matrix along with a visual “data explorer” and “on-device performance” metrics are generated as seen in the screenshot below.

#1 [Click to set a description for this version](#)

Neural Network settings

Training settings

Number of training cycles ?

30

Learning rate ?

0.0005

Switch to Keras (expert) mode

Edit block locally

Last training performance (validation set)

%

ACCURACY

98.0%

📈

LOSS

0.07

Confusion matrix (validation set)

	0	1
0	97.9%	2.1%
1	2.0%	98.0%
F1 SCORE	0.98	0.98

Actual label
0

Feature explorer (full training set) ?

0 - correct

1 - correct

0 - incorrect

1 - incorrect

On-device performance ?

🕒

INFERENCEING ...

74 ms.

🏠

PEAK RAM USA...

2.9K

📱

FLASH USAGE

273.1K

 **EDGE IMPULSE**

Solutions Guide

21

Step 4: Navigate to the “Model testing” tab to test the trained model on unseen data.

Click “classify all” to test the trained model on unseen data. You may also “set confidence thresholds” depending on how critical your use case is. For e.g detecting a heart attack is critical and thus setting a low confidence threshold in this case would make sense. or “delete” the layer/s as shown below.

This lists all test data. You can manage this data through [Data acquisition](#).

Test data

[Classify all](#) ⋮

Set confidence thresholds

Set the 'expected outcome' for each sample to the desired outcome to automatically score the impulse.

SAMPLE NAME	EXPECTED OUTCOME	LENGTH	ACCURACY	RESULT	
1.9997	1	1s	100%	1.1	⋮
1.9993	1	1s	100%	1.1	⋮
1.9988	1	1s	100%	1.1	⋮
1.9983	1	1s	100%	1.1	⋮
1.9976	1	1s	100%	1.1	⋮

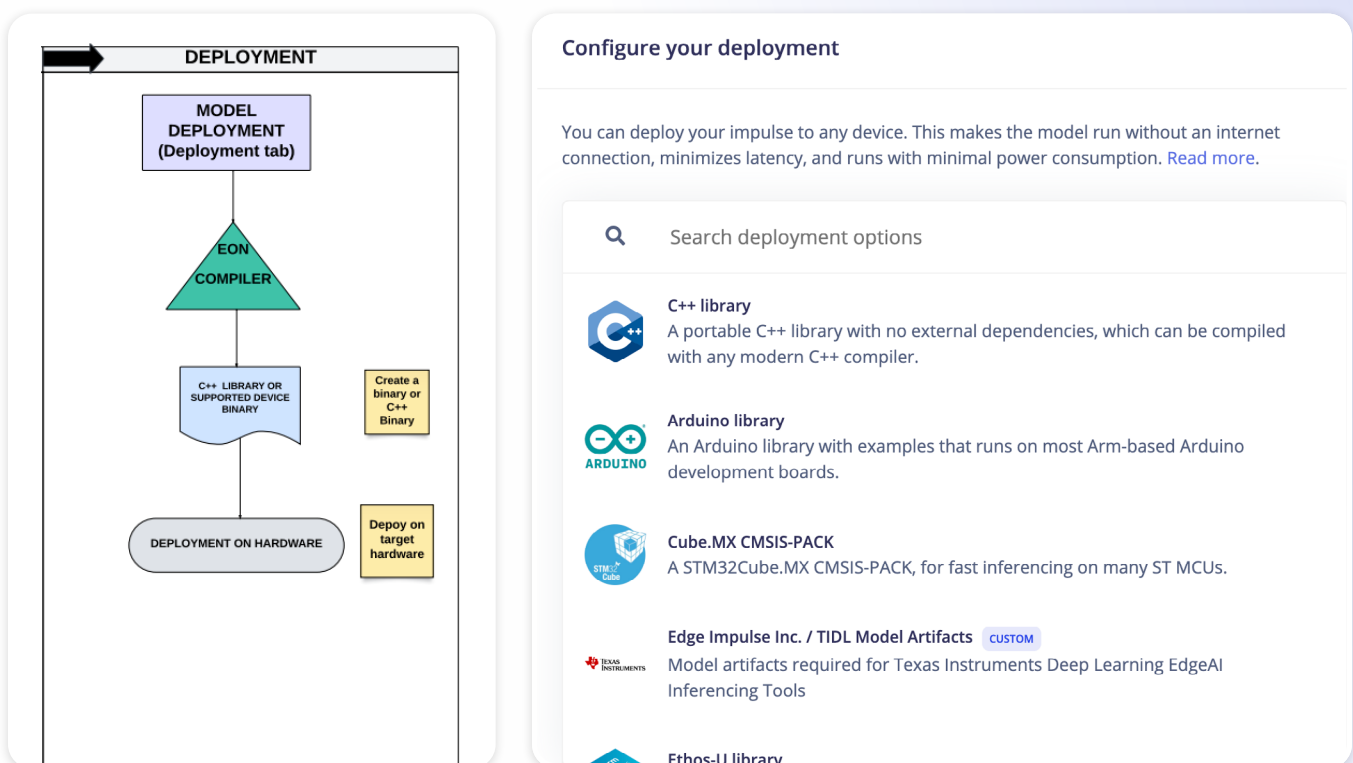
4. Model Deployment

Deployment block

Digital health use-cases often involve a variety of devices like wearables, mobile devices, or specialized medical equipment. Integrating with these devices and deploying the trained models on them in an efficient manner can be challenging.

One of the most powerful features in Edge Impulse are the built-in deployment targets (under **Deployment** in the Studio), which let you create ready-to-go binaries for development boards, or custom libraries for a wide variety of targets that incorporate your trained impulse.

Regardless of whether you are using a [fully supported development board](#) or not, you can deploy your impulse to any device. The C++ library and Edge Impulse SDK enable the model to run without an internet connection on the device, minimize latency, and with minimal power consumption.



EON Compiler (optional)

Many digital health applications need to operate in real-time. For instance, a wearable that monitors heart rate variability for early detection of cardiac issues, or a system that alerts for falls in the elderly, requires immediate response. A slow or lagging model could lead to delayed alerts, potentially risking lives.

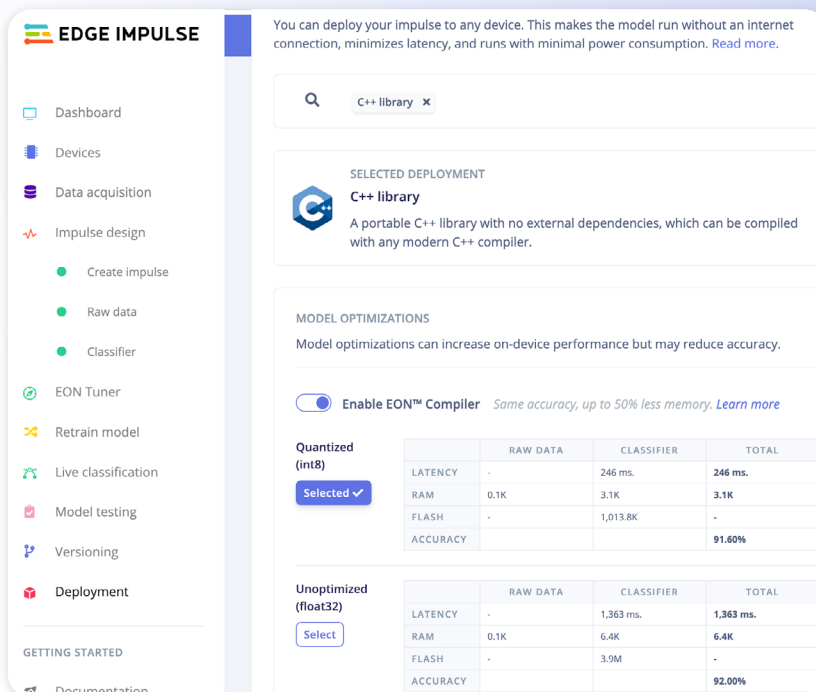
Additionally, the user experience is another important factor to consider. A device that is slow or unresponsive can frustrate users, leading to poor adoption rates. On the other hand, a device with a smooth, responsive user interface, facilitated by well-optimized models, can encourage user engagement. As a result, providing real-time monitoring and diagnosis, while maintaining a high standard of user experience and data privacy is crucial for digital health applications.

To support this goal, EI provides EON Compiler capability which helps you get the best performance from your device. By optimizing machine learning models to better utilize the device's specific hardware characteristics, the EON Compiler ensures your application runs as efficiently as possible.

5. Data Pipelines

Retrain model & versioning

One of the most common pain points in the digital health device space is the challenge of managing the constant influx of new data, drifted data distribution, model overfitting issues and ensuring seamless updates. Data pipelines enable automated data updates, effortlessly incorporating the latest information into your analysis. Moreover, pipeline versioning allows you to track and reproduce your findings, ensuring transparency and traceability.



The screenshot shows the Edge Impulse Studio interface. On the left is a sidebar with navigation options: Dashboard, Devices, Data acquisition, Impulse design, EON Tuner, Retrain model, Live classification, Model testing, Versioning, and Deployment. The main area displays the 'Retrain model' and 'Versioning' sections. The 'Retrain model' section includes a search bar, a 'C++ library' selection, and a 'MODEL OPTIMIZATIONS' section with a toggle for 'Enable EON™ Compiler'. Below this are two tables comparing 'Quantized (int8)' and 'Unoptimized (float32)' models across metrics like Latency, RAM, Flash, and Accuracy.

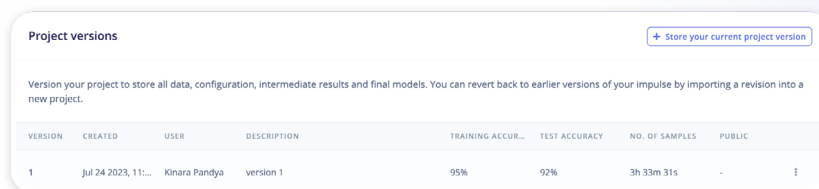
	RAW DATA	CLASSIFIER	TOTAL
LATENCY	-	246 ms.	246 ms.
RAM	0.1K	3.1K	3.1K
FLASH	-	1,013.8K	-
ACCURACY			91.60%

	RAW DATA	CLASSIFIER	TOTAL
LATENCY	-	1,363 ms.	1,363 ms.
RAM	0.1K	6.4K	6.4K
FLASH	-	3.9M	-
ACCURACY			92.00%

The Retrain model feature (optional in this tutorial) in Edge Impulse Studio comes handy during such scenarios. It uses already known parameters from your selected DSP and ML blocks then uses them to automatically regenerate new features and retrain the Neural Network model in one single step.

To version your current project, follow the steps below.

Step 1: Navigate to Project → Versioning tab → store your current project version



The screenshot shows the 'Project versions' table in Edge Impulse Studio. It includes a '+ Store your current project version' button and a table with columns: VERSION, CREATED, USER, DESCRIPTION, TRAINING ACCUR..., TEST ACCURACY, NO. OF SAMPLES, and PUBLIC. The table contains one row for version 1.

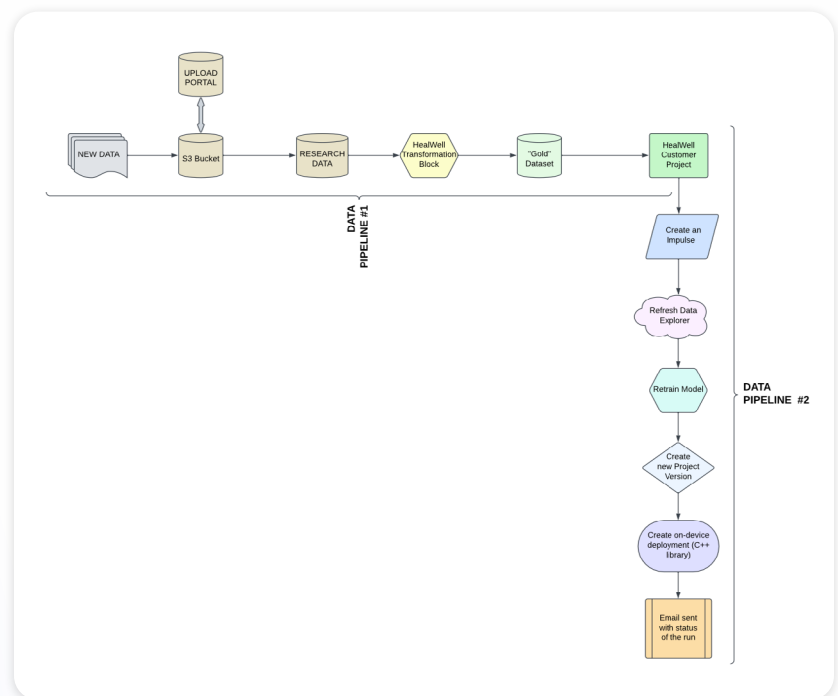
VERSION	CREATED	USER	DESCRIPTION	TRAINING ACCUR...	TEST ACCURACY	NO. OF SAMPLES	PUBLIC
1	Jul 24 2023, 11:...	Kinara Pandya	version 1	95%	92%	3h 33m 31s	-

Data Pipelines

Automated data pipelines can fuel machine learning models used for a variety of use cases such as predictive analytics by forecasting patient health outcomes, detecting potential health risks earlier, and guiding preventative care strategies.

In digital health, especially with the rise of wearable technologies and IoT devices, there's a constant stream of health data being produced. Data pipelines can help process this real-time data, allowing for continuous health monitoring and immediate response to critical changes.

To support the above requirements, EI provides the capability to build data pipelines where you can stack several transformation blocks which then feed transformed data into a dataset or a project.



We already created a data pipeline to fetch data from s3 into the project [here](#). In this section, we will extend the pipeline to retrain the model when new data becomes available.

As new data gets added, the pipeline smoothly manages the entire Edge AI platform process from start to finish: it recreates the data explorer, allowing insights into the newly acquired data; it generates and versions the retrained model, keeping track of the model's evolution over time; and also constructs a new C++ library, ensuring that the updated model is efficiently integrated into the target device of your choice. This approach not only ensures the model's adaptability to changing conditions but also optimizes its deployment, ensuring that the target device on which the model is deployed benefits from the most current and relevant insights provided by the retrained model.

Follow the steps below to configure this automated data pipeline.

Step 2: Creating data pipeline #2 (model re-training to deployment)

Navigate to your Project
→ Data acquisition → Data source → Add a new data source

Select “Don’t import data” →
“Next, setup actions”

Select all actions and “create pipeline”

All pipelines for the project are displayed under the “data sources” tab. Active data pipelines for a sample project are shown below.

Automatically fetch new data

You can automatically add data to your project from your cloud provider. This sync runs periodically, or you can trigger a sync from your own pipeline or application.

Where does your data live?

☐ Amazon S3

☐ Google Cloud Storage

☐ Somewhere else (S3 compatible)
For cloud providers with an Amazon S3 compatible storage solution.

☐ Upload portal
Upload portals are a secure way to let external parties upload data to your datasets.
Select portal:

☐ Transformation block
Transformation blocks let you combine, transform and import data in any form.
Select transformation block:

☒ Don't import data
Skip importing data automatically, just set up actions for this project.

Next, set up actions

Automatically fetch new data

What actions do you want to run in this pipeline?

☒ Recreate data explorer
The data explorer gives you a one-look view of your dataset, letting you quickly label unknown data. If you enable this you'll also get an email with a screenshot of the data explorer whenever there's new data.

☒ Retrain model
If needed, will retrain your model with the same impulse. If you enable this you'll also get an email with the new validation and test set accuracy.

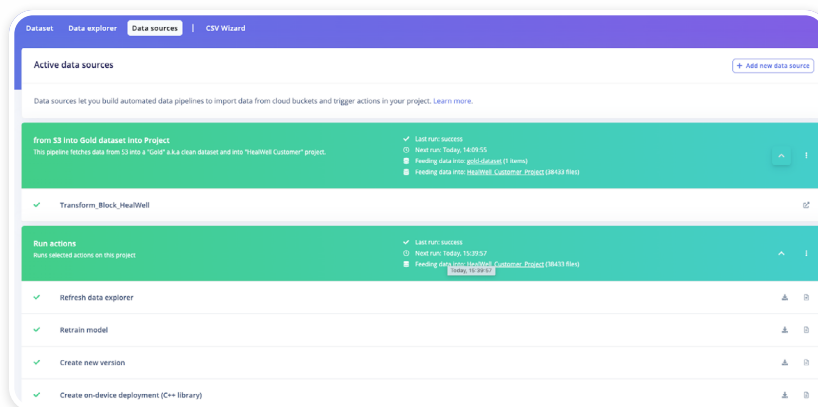
☒ Create new version
Store all data, configuration, intermediate results and final models.

☒ Create new deployment
Builds a new library or binary with your updated model. Requires 'Retrain model' to also be enabled.

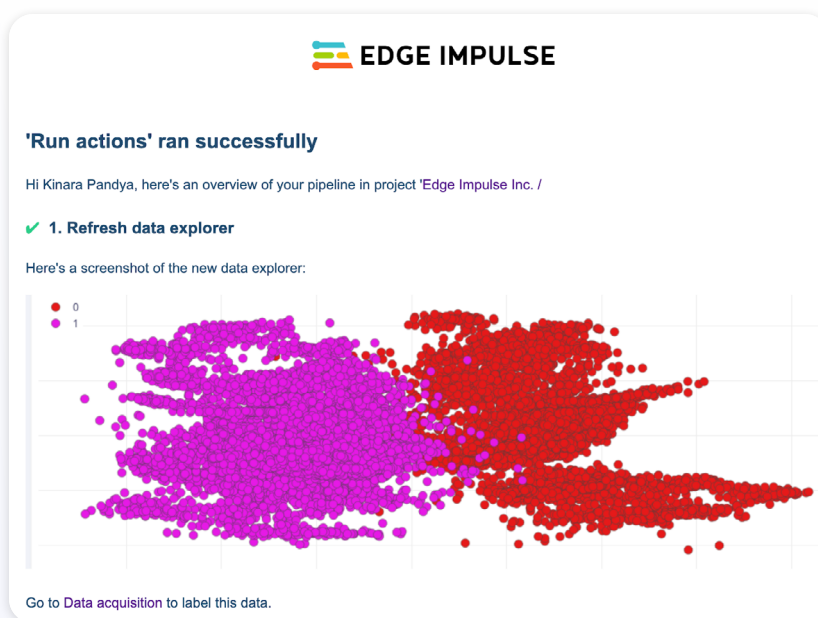
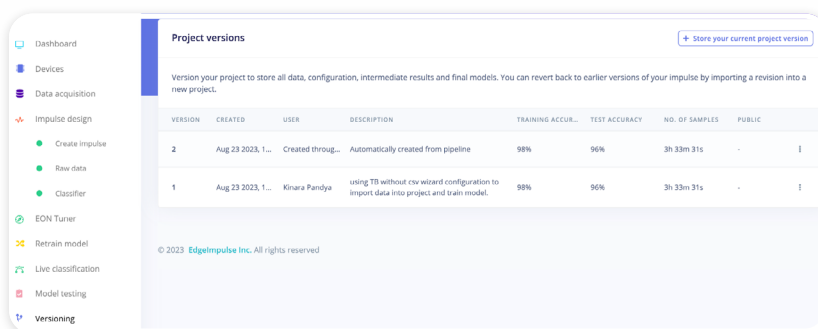
Optional: configure a webhook
You can enter a webhook URL here which will receive a notification whenever a step in this pipeline runs.

Create pipeline

You can also check the “Versioning” tab for the newly created version of the model.



Immediately following completion of the pipeline run, you will receive an email with the status of the run as shown below.



Data Campaigns (optional)

To keep the tutorial beginner friendly, this section is optional and for information only. “Data Campaigns” feature is an integral part of monitoring the health of your pipelines and model performance over time.

It allows you to quickly track your experiments as the model’s development progresses. It is an overview of your pipelines where you can easily extract useful information from your datasets and correlate those metrics with your model performances.

Step 1: Setting up your dashboard

(Data campaigns overview)

To get started, navigate to the Data campaigns tab in your organization:

(Data campaigns)

✓ 2. Retrain model

Step ran successfully. See the [full logs](#) for more info.

✓ 3. Create new version

Step ran successfully. See the [full logs](#) for more info.

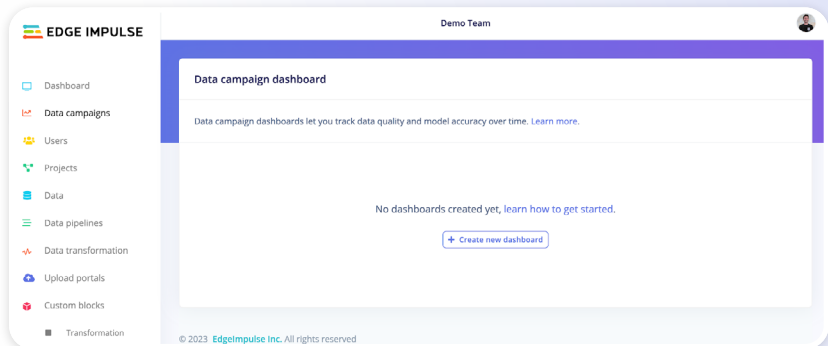
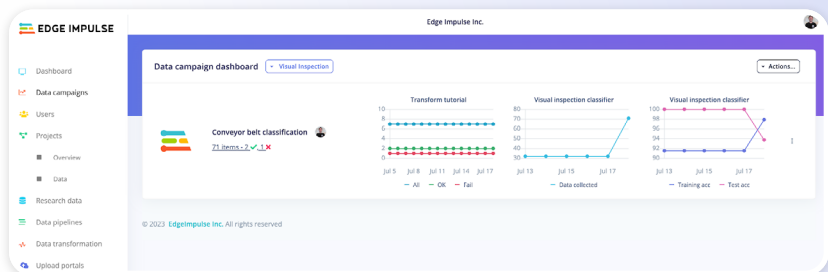
✓ 4. Create on-device deployment (C++ library)

Download the build artefacts from:

<https://studio.edgeimpulse.com/v1/api/274435/deployment/download?type=zip&modelType=int8>

You’re receiving this email because you’ve been subscribed to notifications for the data source “Run actions”. To unsubscribe, go to [data sources](#), click “Edit pipeline” and remove yourself.

Cheers,
Jan from the Edge Impulse team

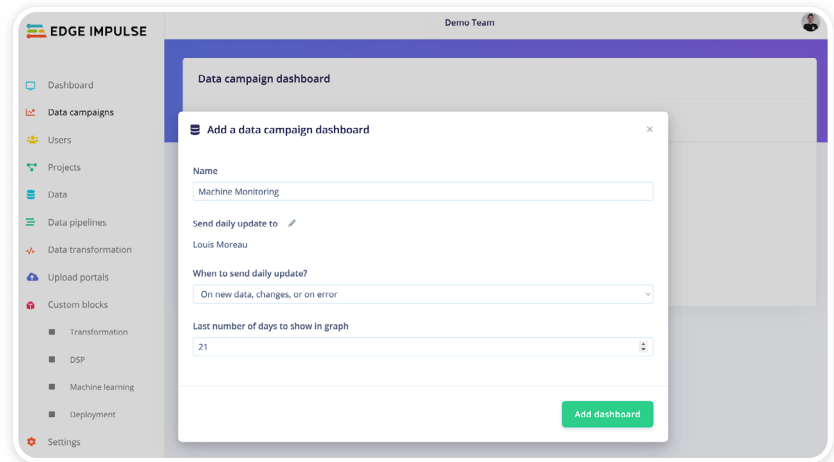


Click on **+ Create new dashboard**.

Give your dashboard a name, and select one or more collaborators to receive the daily updates by email. If you don't want to be spammed, you can select when you want to receive these updates, either *Always*, *On new data*, *changes or on error*, or *Never*. Finally, set the last number of days shown in the graphs.

Add data campaign dashboard

You can create as many dashboards as needed, simply click on **+ Create a new dashboard** from the dropdown available under your current dashboard:

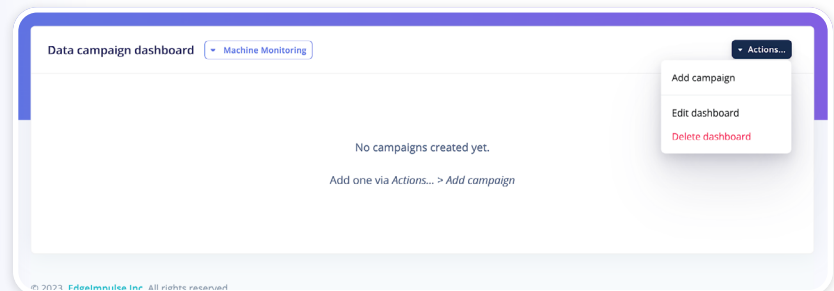
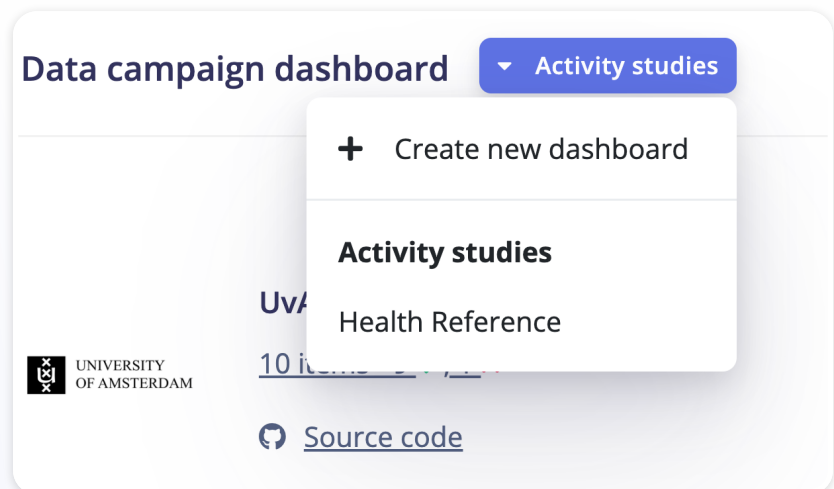


If you want to delete a dashboard, Click on **Actions... → Delete dashboard**

Step 2: Setting up your campaign

Once your dashboard is created, you can add your custom campaigns. It's where you will specify which metrics are important to you and your use case. Click on **Actions... → Add campaign**

Add data campaign



Step 3: Fill the form to create your campaign

Name: Name of your data campaign.

Description: Description of your data campaign.

Campaign coordinators: Add the collaborators that are engaged with this campaign

Datasets: Select the datasets you want to visualize in your campaign.

You can add several datasets.

Projects: Select the projects you want to visualize in your campaign.

You can add several projects.

Pipelines: Select the pipeline that is associated with your campaign. *Note that this is for reference only, it is currently not displayed in your campaign*

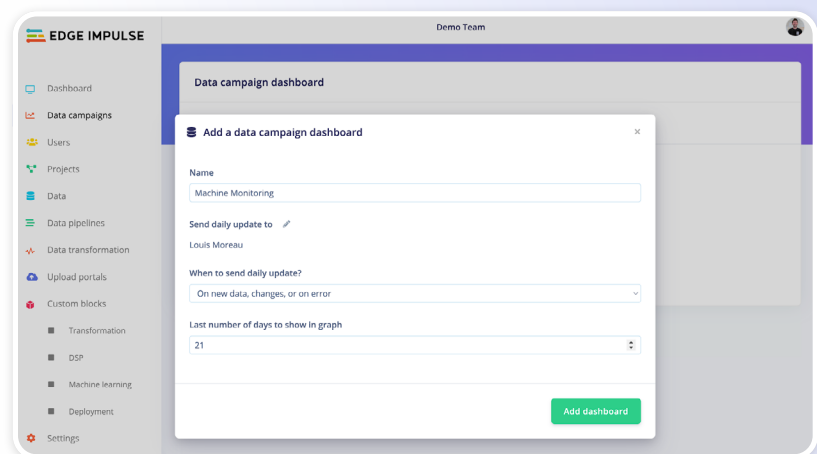
Links: Select between the link type you need. Options are Github, Spreadsheet, Text Document, Code repository, List or Folder. Add a name and the link. This place is useful for other collaborators to have all the needed information about your project, gathered in one place under your campaign.

Additional queries to track:

These queries are data filters that need to be written in the SQL WHERE format. For example `metadata->age >= 18`` will return the data samples from adult patients.

Save your data campaign and it will be added to your dashboard.

Create or edit data campaign

The screenshot shows the 'EDGE IMPULSE' interface with a sidebar on the left containing navigation links: Dashboard, Data campaigns, Users, Projects, Data, Data pipelines, Data transformation, Upload portals, Custom blocks, Transformation, DSP, Machine learning, Deployment, and Settings. The main content area is titled 'Data campaign dashboard' and features a modal window titled 'Add a data campaign dashboard'. This modal contains the following fields: 'Name' (with the value 'Machine Monitoring'), 'Send daily update to' (with the value 'Louis Moreau'), 'When to send daily update?' (a dropdown menu set to 'On new data, changes, or on error'), and 'Last number of days to show in graph' (a dropdown menu set to '21'). A green 'Add dashboard' button is located at the bottom right of the modal.


This dashboard shows the metrics' progress from the [Health reference design data](#).



Create or edit data campaign

If you want to edit or delete your campaign, click on the “⋮” button on the right side of your campaign:

[Data Campaigns](#)



Edit campaign

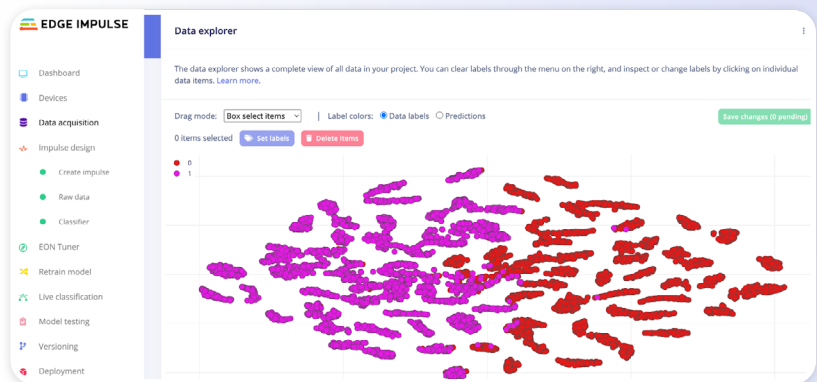

6. Data Visualization and Exploration Capabilities

As your organization expands, the need for efficiency and scalability becomes ever more critical. Traditional data processing methods struggle to keep pace with the growing volume and complexity of data. *Edge Impulse rises to the challenge by offering dataset and feature exploration graphs, scalable architectures, distributed computing, and parallel processing capabilities.* With these powerful tools, you can analyze vast datasets in a fraction of the time.

Step 1: Navigate to Project → Data acquisition → Data explorer to visualize the dataset

In the following screenshot you see the cleaned dataset with a pretty good separation between the labels 0 (not confused) and 1(confused).

In the screenshot below, you see the correct and incorrect predictions made by the model for the two labels.




7. Collaboration Capabilities

Collaboration and reproducibility are essential pillars of impactful product development and maintenance. *Edge Impulse* offers a unified environment where both embedded and ML teams can work together. It supports deployment of trained models directly onto resource-constrained embedded devices. This feature ensures that the ML models are integrated smoothly into the embedded system, eliminating potential integration challenges and fostering collaboration between the two teams.


Additionally, by documenting the origin of the dataset, data processing steps and transformations within your project's README, *Edge Impulse* fosters reproducible engineering practices, allowing your peers to replicate and validate your findings.

Step 1: Navigate to Project


→ Dashboard → “Add
Readme” and update the
project information



Step 2: From the
dashboard, add any
collaborators you wish
to add to the project



*From the Dashboard, you
may also make the project
public (optional).*



8. Security and Compliance

Data security and compliance are paramount in the digital health domain, where sensitive PHI (*Private Human information*) must be safeguarded. By allowing for on-device machine learning, Edge Impulse enables data to be processed without the need for transmission, improving data security and privacy. EI is designed to incorporate stringent security measures such as encryption, SOC2 regulatory compliance such as HIPAA, GDPR, access controls, and anonymization techniques.

Please note that our plan is to be HIPAA certified by Q2 2024.

[🔗 Security and Compliance](#)

Additional Resources

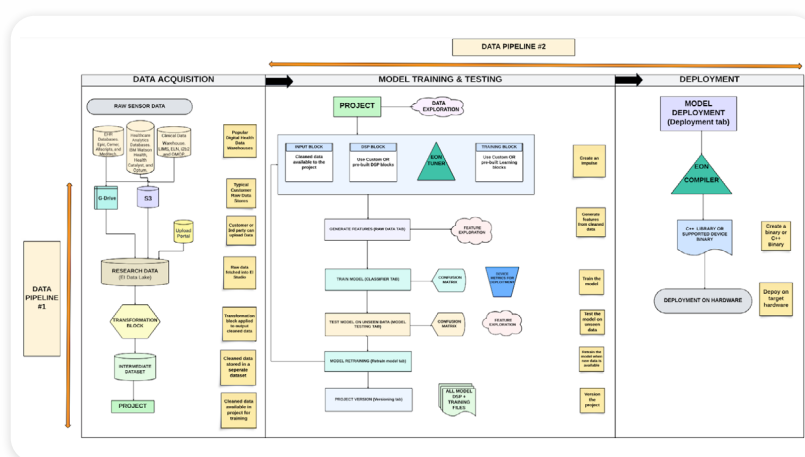
The Edge Impulse Python SDK is a library to help you develop machine learning (ML) applications for edge and Internet of Things (IoT) devices. While the Edge Impulse Studio is a great interface for guiding you through the process of collecting data and training a model, the [edgeimpulse](#) Python SDK allows you to programmatically Bring Your Own Model (BYOM), developed and trained on any platform.

[🔗 Edge Impulse Python SDK](#)



Summary

Edge Impulse's components are designed to address the different stages of a typical machine learning (ML) pipeline. The arrangement of these components reflects the common process of data collection, data cleaning/transformation, feature extraction, model training, testing, and deployment.



By mirroring the standard stages of the ML pipeline, Edge Impulse ensures users have all the tools they need at each stage. This integrated approach enables anyone who is a beginner or an expert to seamlessly build and deploy ML models.

The “glue” that connects all these components together in Edge Impulse is its unified software platform. Because of this centralized interface, users can manage each stage of the machine learning workflow right from data ingestion to model deployment. All the components are organized in a cohesive manner/ tightly integrated on the platform’s user interface. For example, data gathered and processed in one stage of the workflow is seamlessly available for the next, and so on. Additionally, Edge Impulse provides APIs and SDKs for integrating its platform with other tools and systems. This enables users to customize their workflows and use external tools when necessary.

The real power of Edge Impulse lies in its modular and flexible architecture. You can modify and iterate over your impulse design, data, and model as many times as you want until you achieve a satisfactory result.

Moreover, by incorporating collaboration tools and a focus on security and compliance, Edge Impulse addresses key considerations for teams working in the digital health industry.

About the Author

Kinara Pandya is a seasoned Machine Learning Solutions Engineer with a background in Computer Science, Data Science, and Biotechnology. With experience spanning digital health solutions and the life sciences industry, Kinara has shared insights on next-gen digital health technologies, including “Advancing Health with Edge AI”, at prominent conferences like HLTH in Las Vegas. Additionally, her expertise is showcased in her contributions to a report published on [MedRxiv](#).