Artificial Intelligence in Agriculture

AI Insights for Dairy Cows

DS 413/513 & DS 413/513L

Maristela Rovai, DVM, Ph.D.

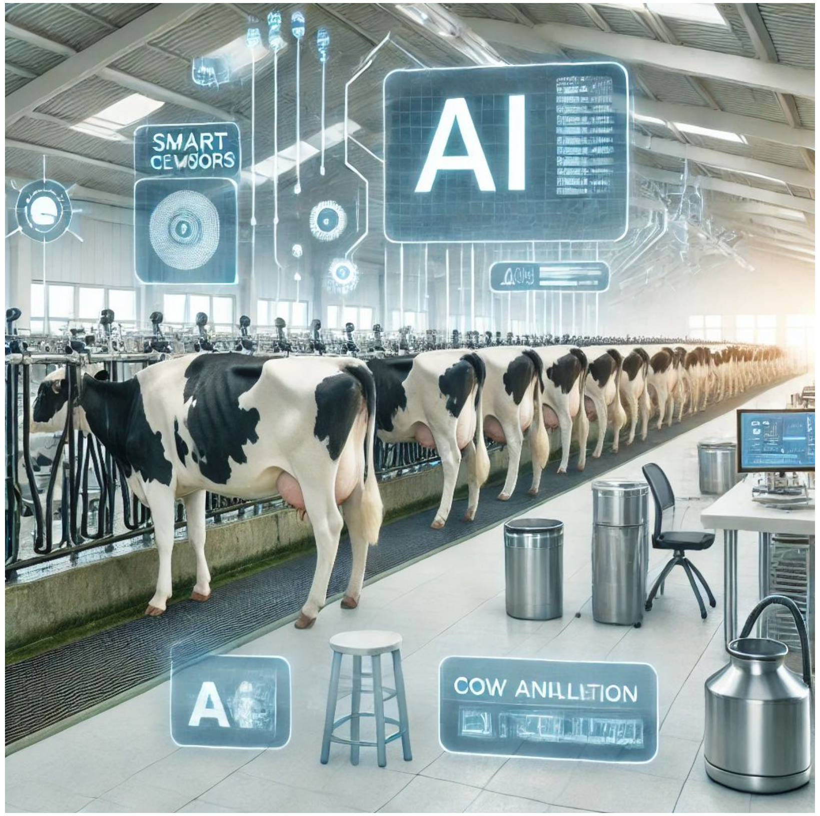

Prof. Marcelo J. Rovai
Federal University of Itajuba, Brazil

Marcelo Rovai is an educator and professional in the field of engineering and technology, holding the title of Professor Honoris Causa from the Federal University of Itajubá, Brazil. His educational background includes an Engineering degree from UNIFEI and an advanced specialization from the Polytechnic School of São Paulo University. Further enhancing his expertise, he earned an MBA from IBMEC (INSPER) and a Master's in Data Science from the Universidad del Desarrollo (UDD) in Chile.

With a career spanning several high-profile technology companies such as AVIBRAS Airspace, ATT, NCR, and IGT, where he served as Vice President for Latin America, he brings a wealth of industry experience to his academic endeavors. He is a prolific writer on electronics-related topics and shares his knowledge through open platforms like Hackster.io. In addition to his professional pursuits, he is dedicated to educational outreach, serving as a volunteer professor at UNIFEI and engaging with the TinyML4D group as a Co-Chair, promoting TinyML education in developing countries. His work underscores a commitment to leveraging technology for societal advancement.

Content

- AI Introduction and Applications
- Lab – Image Classification

Content

- AI Introduction and Applications
- Lab – Image Classification

What is AI?

Artificial Intelligence (AI) is the simulation of human intelligence in machines that can perform tasks like learning, reasoning, and problem-solving.

A Brief History of Artificial Intelligence

1950
Alan Turing proposes the Turing Test for machine intelligence

1956
The term '**Artificial Intelligence**' is coined at the Dartmouth Conference

1960s
Early AI programs: ELIZA (natural language processing - **NLP**)

1970s-80s
AI Winter: Funding and interest decline due to limitations of early AI

1998
Yann LeCun developed the **CNN** LeNet, trained with MNIST dataset.

2011
IBM Watson wins Jeopardy!, showcasing advanced natural language processing

2012
Deep learning **breakthroughs in image recognition (ImageNet competition)**

2016
Google's AlphaGo defeats world champion Go player Lee Sedol

2020s
Large language models (e.g., GPT-4) and AI in various industries,

Teaching computers how to learn a task directly from raw data

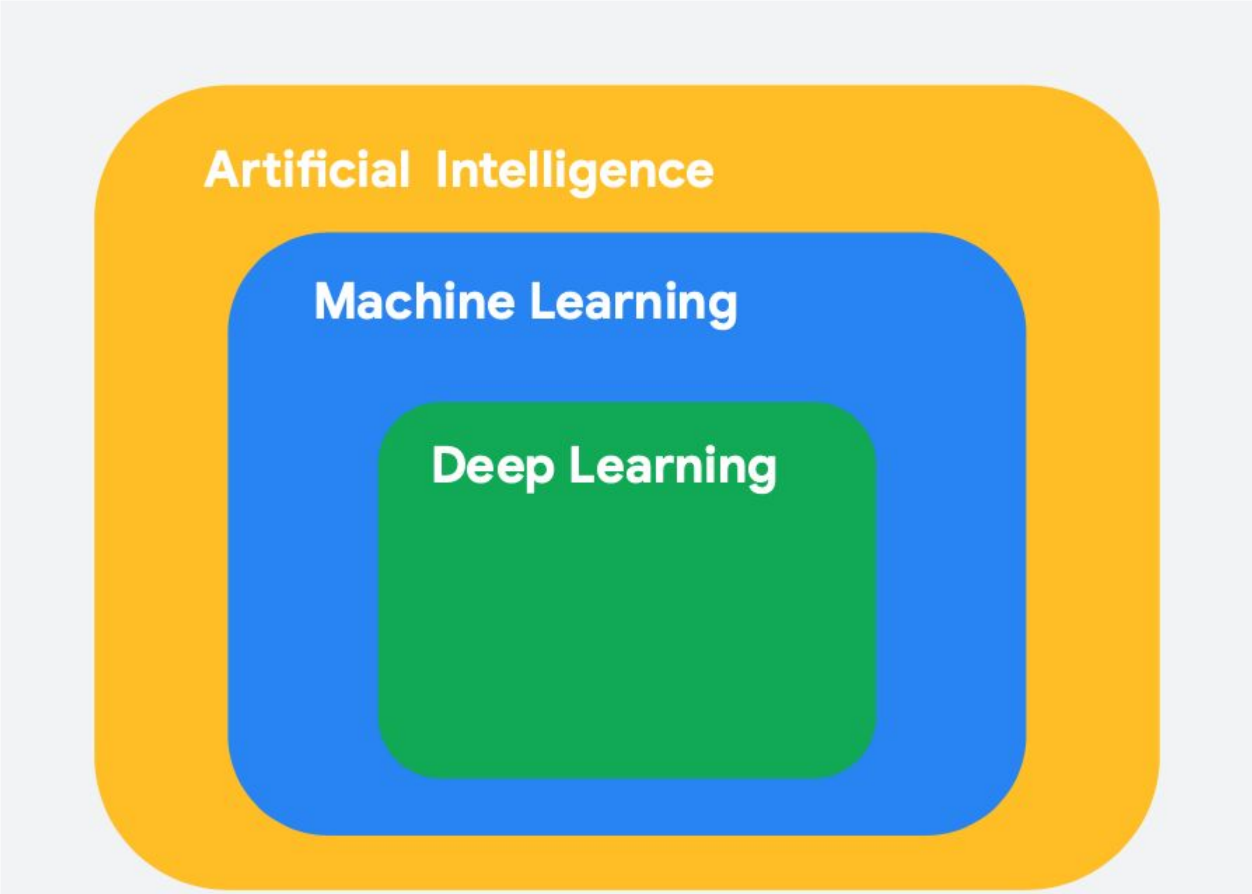

AI: Any technique that enables computers to mimic human behavior

ML: Ability to learn without explicitly being programed

DL: Extract patterns from data using neural networks

(Deep) Machine Learning

Deep Learning: Subset of Machine Learning in which multilayered neural networks learn from vast amounts of data

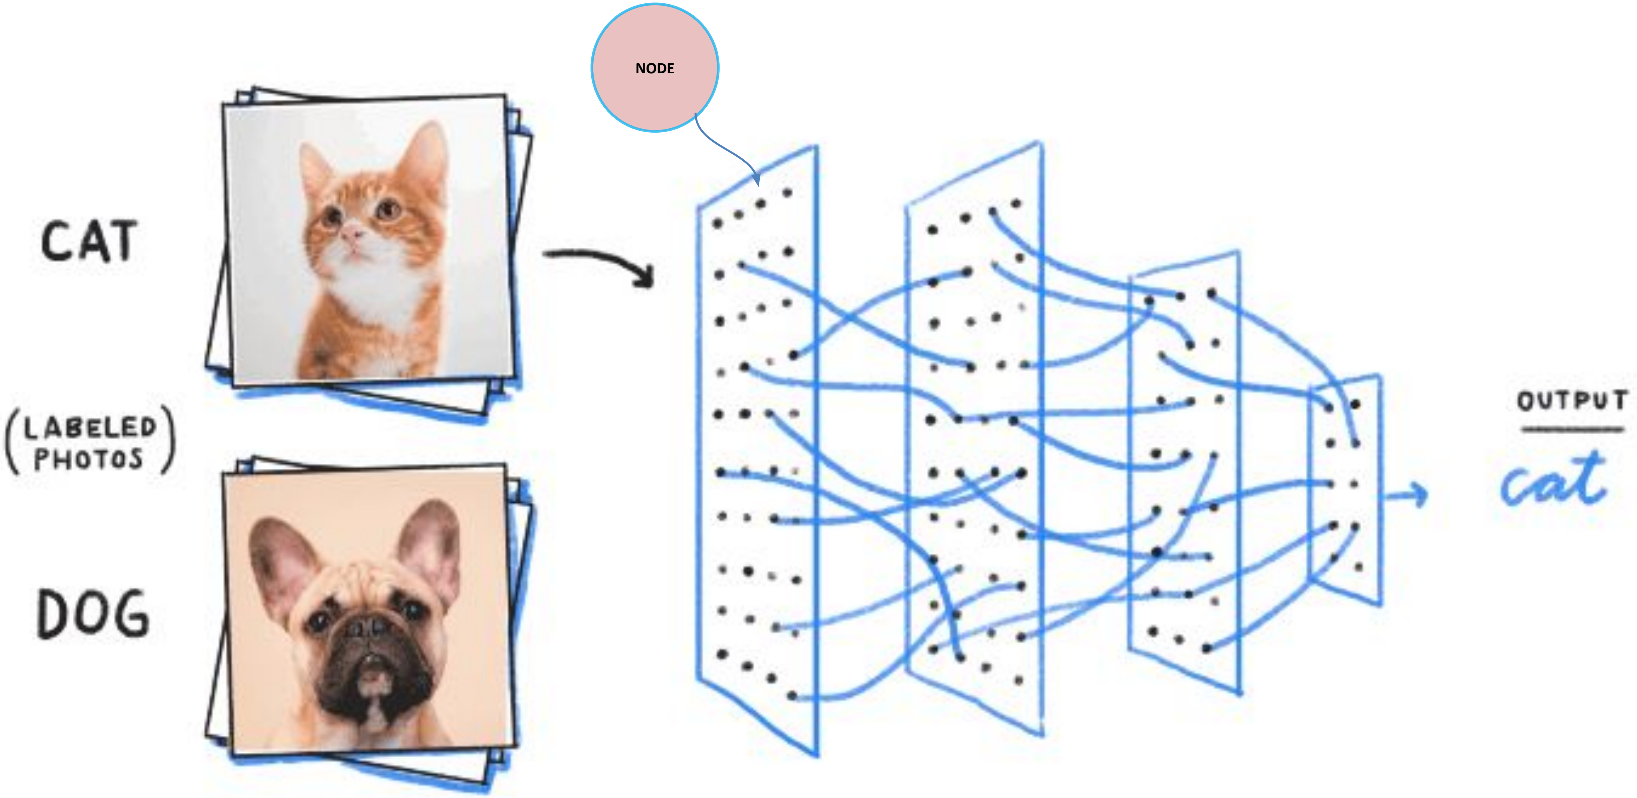

Neuron (Perceptron)

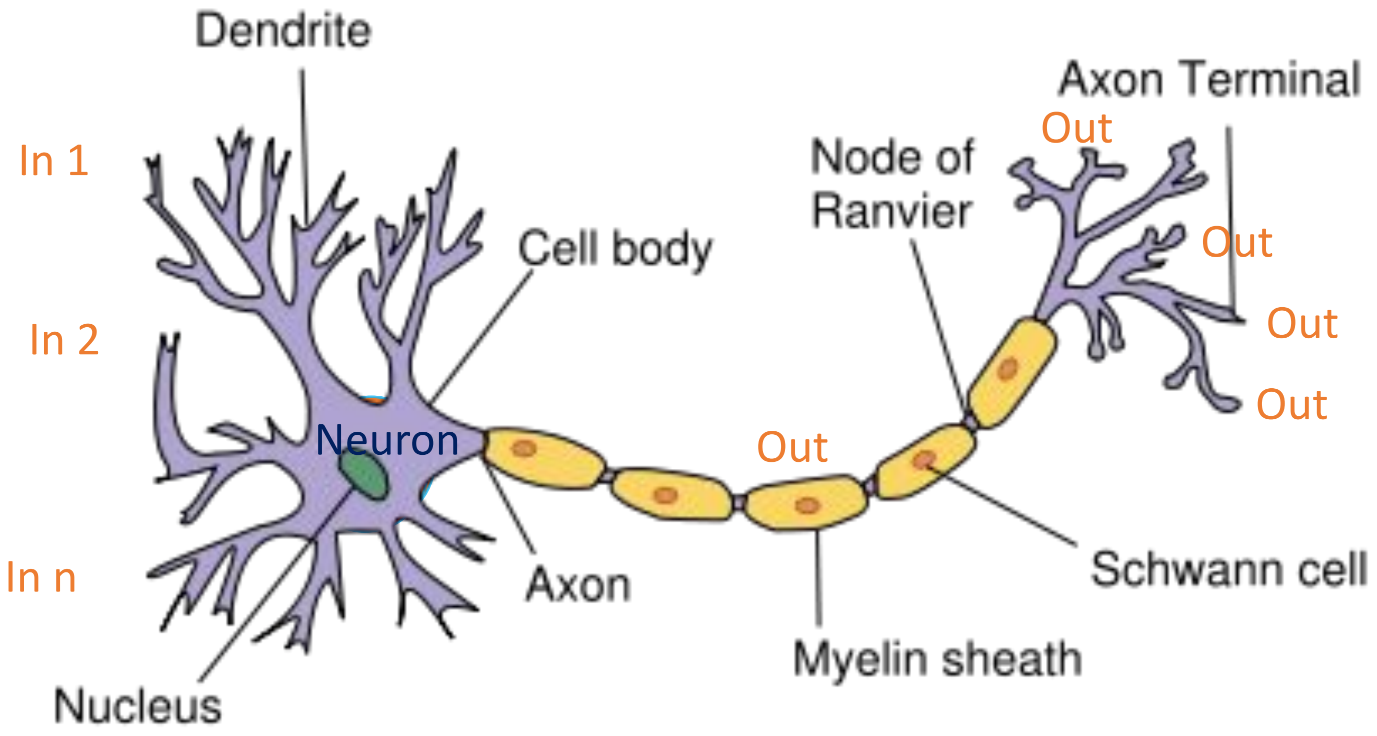

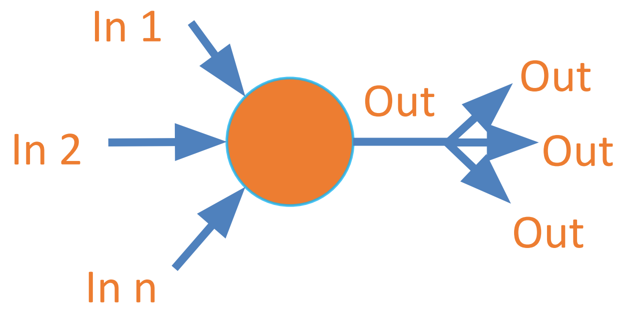

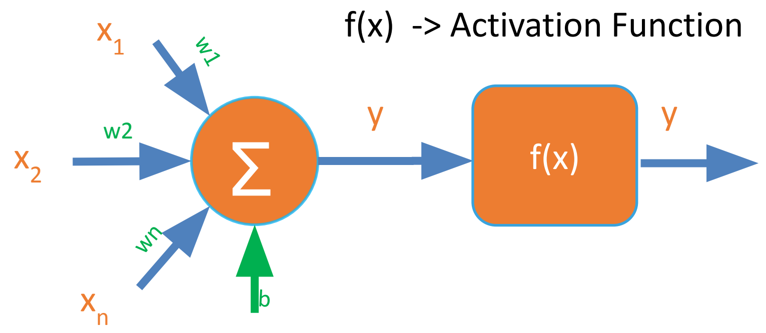

Parameters

$$y = f(\sum_{i=1}^{n} x_i w_i + b)$$

$$y = a x + b$$

$y = ax + b$

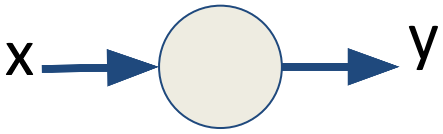

Traditional Computation

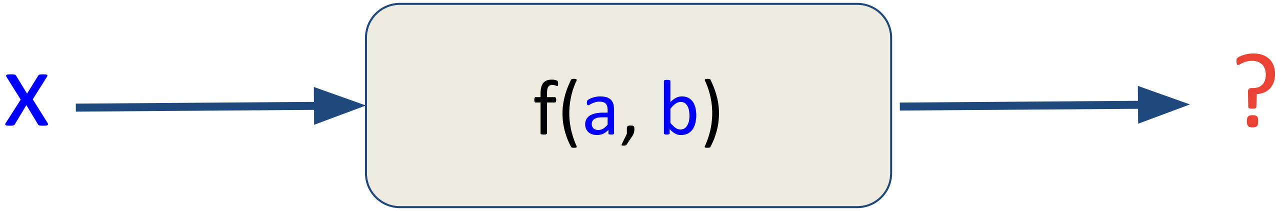

Machine Learning

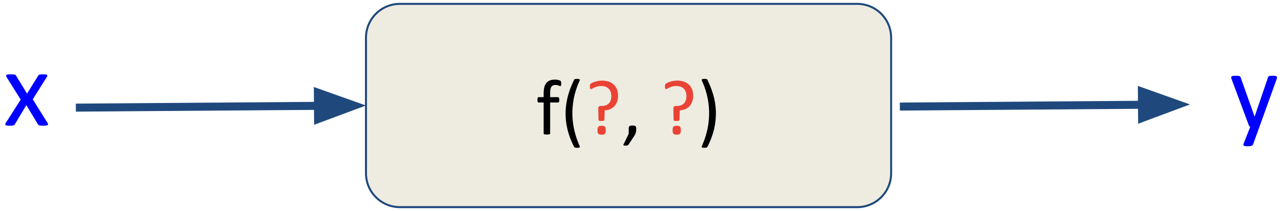

Neural Network Architectures

Vibration Analysis	Image Classification	Text Generation	Image Generation	Large Language Models- LLMs
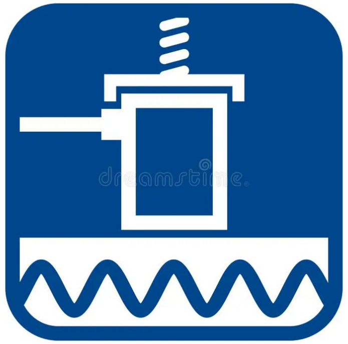				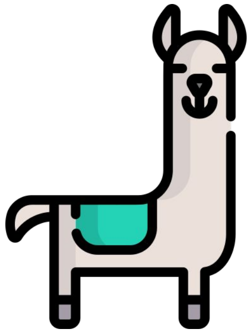
	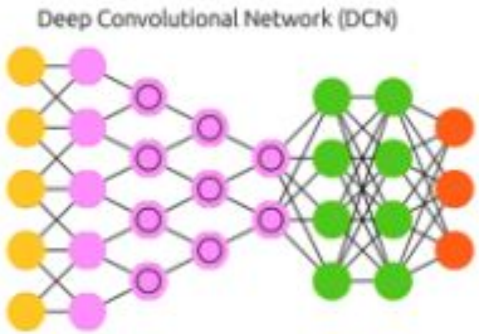	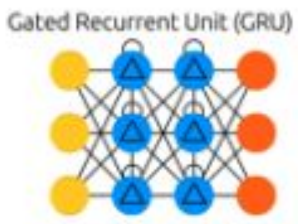	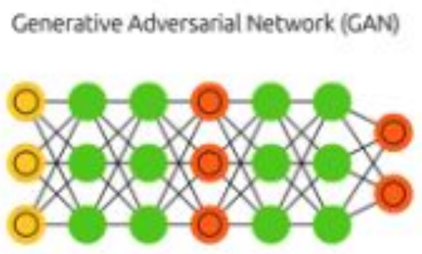	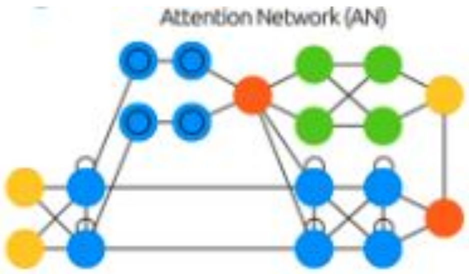

LLMs (Large Language Models)

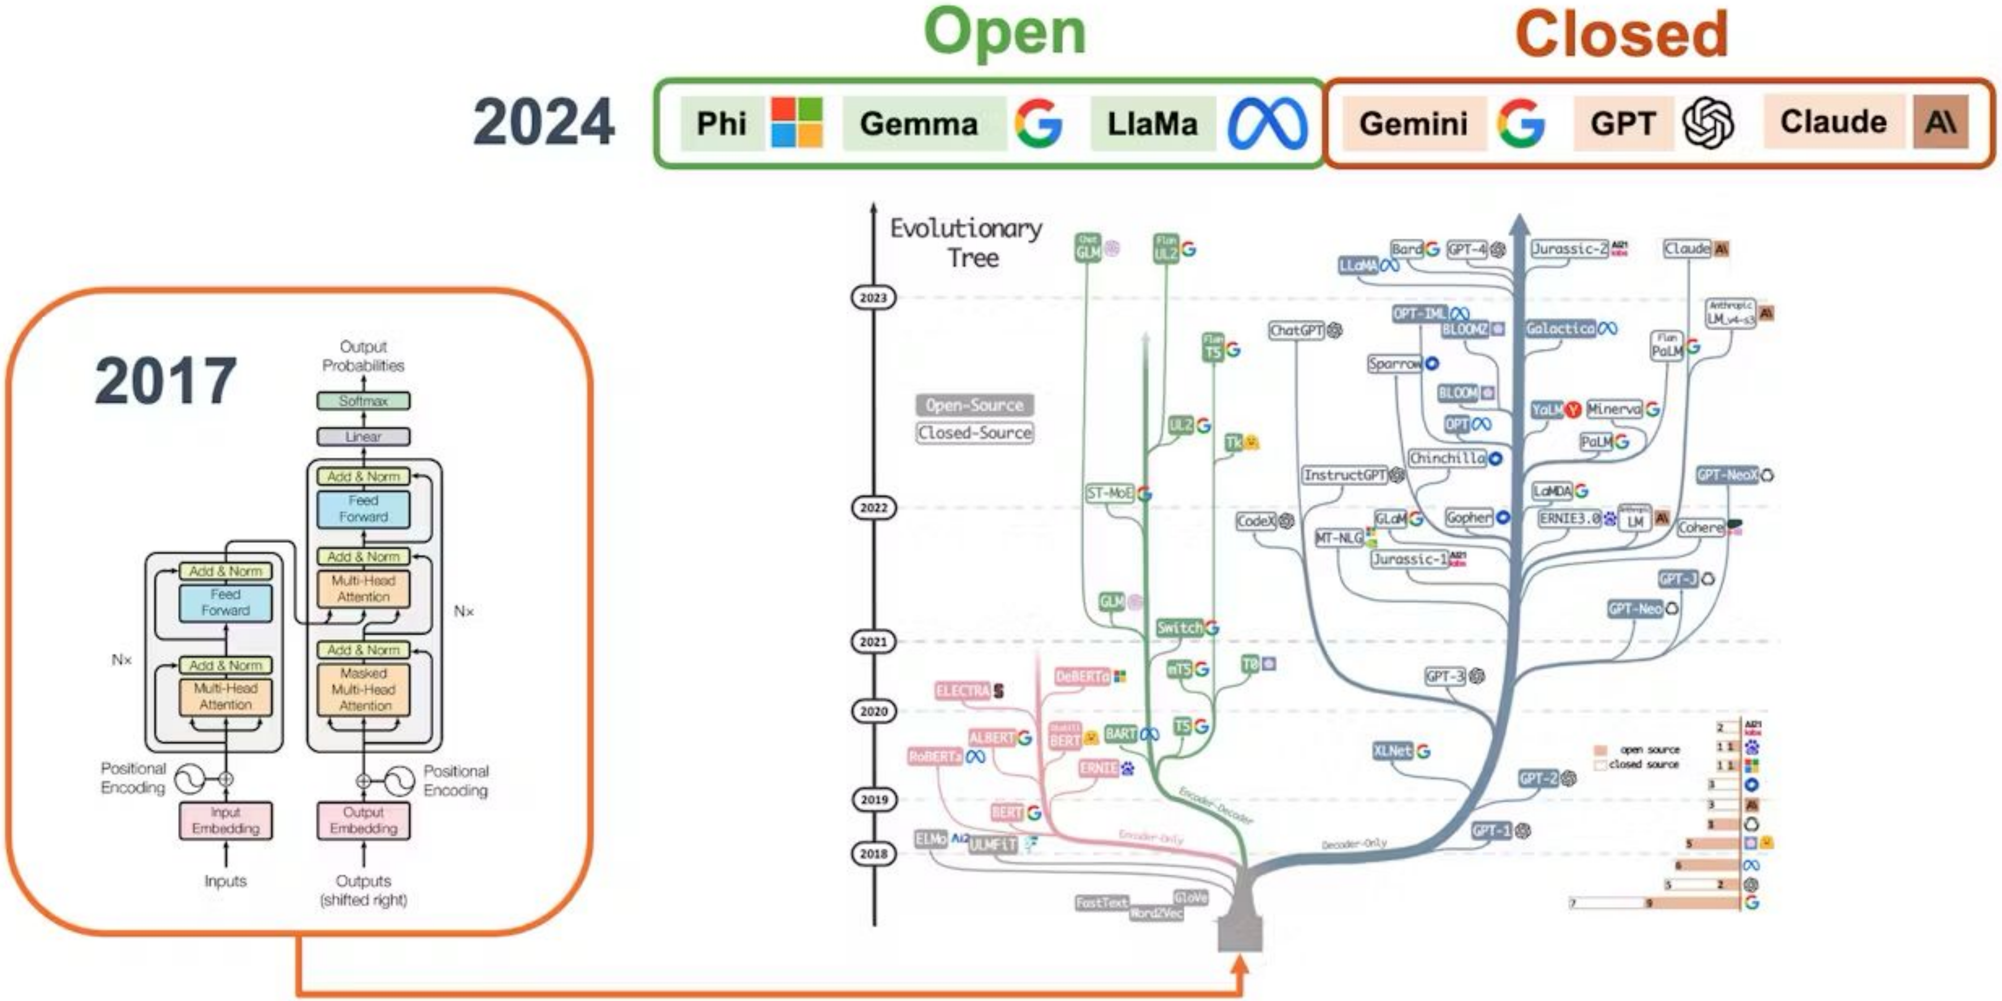

Deep Learning and AI: Why now?

Neural networks date back decades, so why do they dominate?

BIG
DATA

HARDWARE
(GPUs)

SOFTWARE
(NN Models)

AI (ML/DL) Applications

Examples

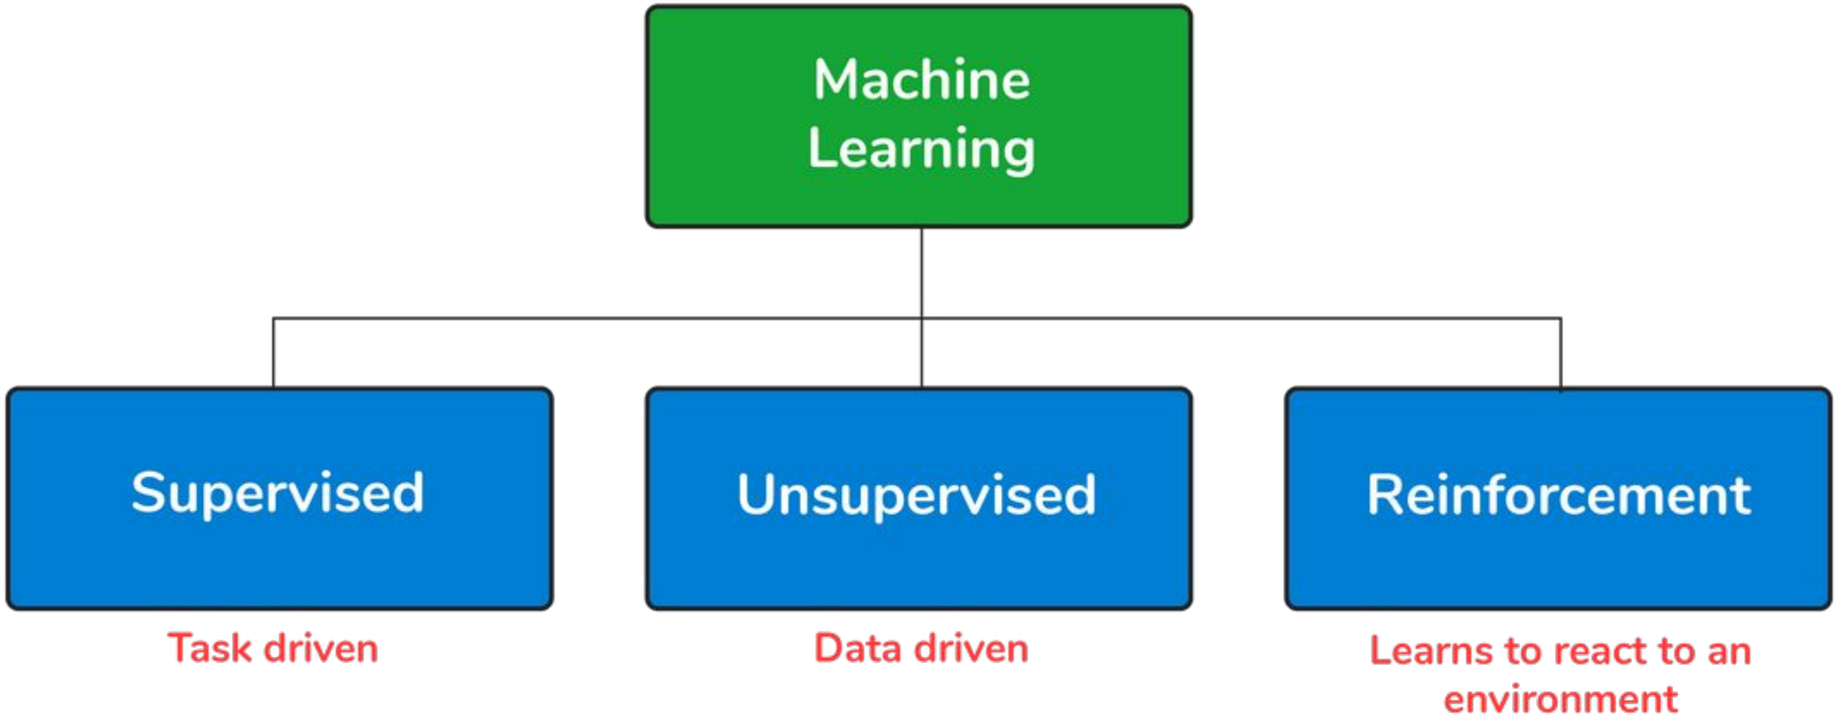

- Regression
- Classification
- Object Detection
- Time Series Forecasting

- Anomaly Detection
- Clustering
- Dimensionality Reduction

- Autonomous Navigation

Models

Sound

Vibration

Vision

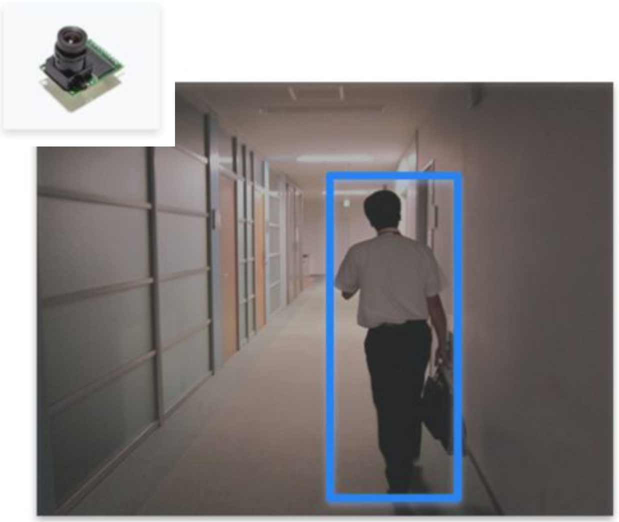

Sensors

Sound

Vibration

Vision

Personal Assistant

Classifying mosquito wingbeat sound using TinyML

Moez Altayeb
University of Khartoum, Sudan
ICTP, Trieste, Italy
mohedahmed@hotmail.com

Marcelo Rovai
Universidade Federal de Itajubá
Itajubá, Brazil
rovai@unifei.edu.br

Marco Zennaro
ICTP
Trieste, Italy
mzennaro@ictp.it

ABSTRACT

Every year more than one billion people are infected and more than one million people die from vector-borne diseases including malaria, dengue, zika and chikungunya. Mosquitoes are the best known disease vector and are geographically spread worldwide. It is important to raise awareness of mosquito proliferation by monitoring their incidence, especially in poor regions. Acoustic detection of mosquitoes has been studied for long and ML can be used to automatically identify mosquito species by their wingbeat. We present a prototype solution based on an openly available dataset, on the Edge Impulse platform and on three commercially-available TinyML devices. The proposed solution is low-power, low-cost and can run without human intervention in resource-constrained areas. This insect monitoring system can reach a global scale.

affected. People from poor communities with little access to health care and clean water sources are also at risk. Although anti-malarial drugs exist, there's currently no malaria vaccine.

Vector-borne diseases also exacerbate poverty. Illness prevent people from working and supporting themselves and their families, impeding economic development. Countries with intensive mal have much lower income levels than those that don't have ma

Countries affected by malaria turn to control rather than e tion. Vector control means decreasing contact between hum disease carriers on an area-by-area basis. It is therefore c be able to detect the presence of mosquitoes in a specific paper presents an approach based on TinyML and on embedded devices.

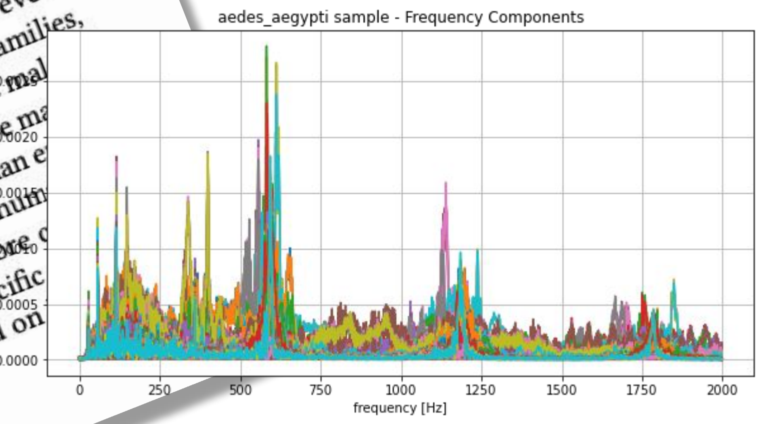

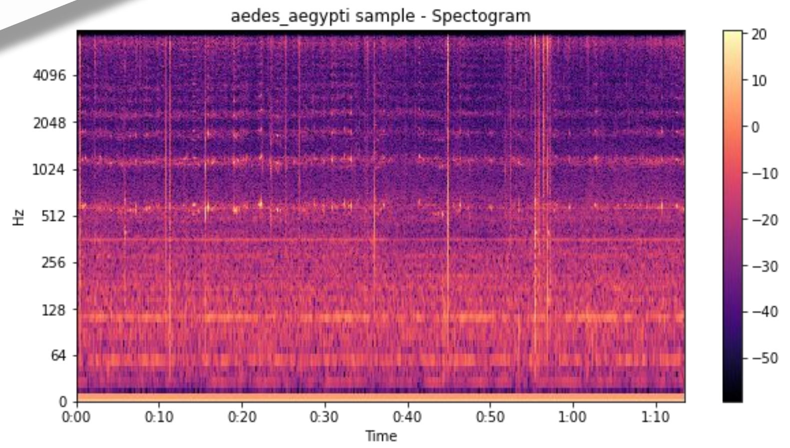

https://www.edgeimpulse.com/blog/the-world-is-all-abuzz-about-tinyml

Bionic Hand Voice Commands Module

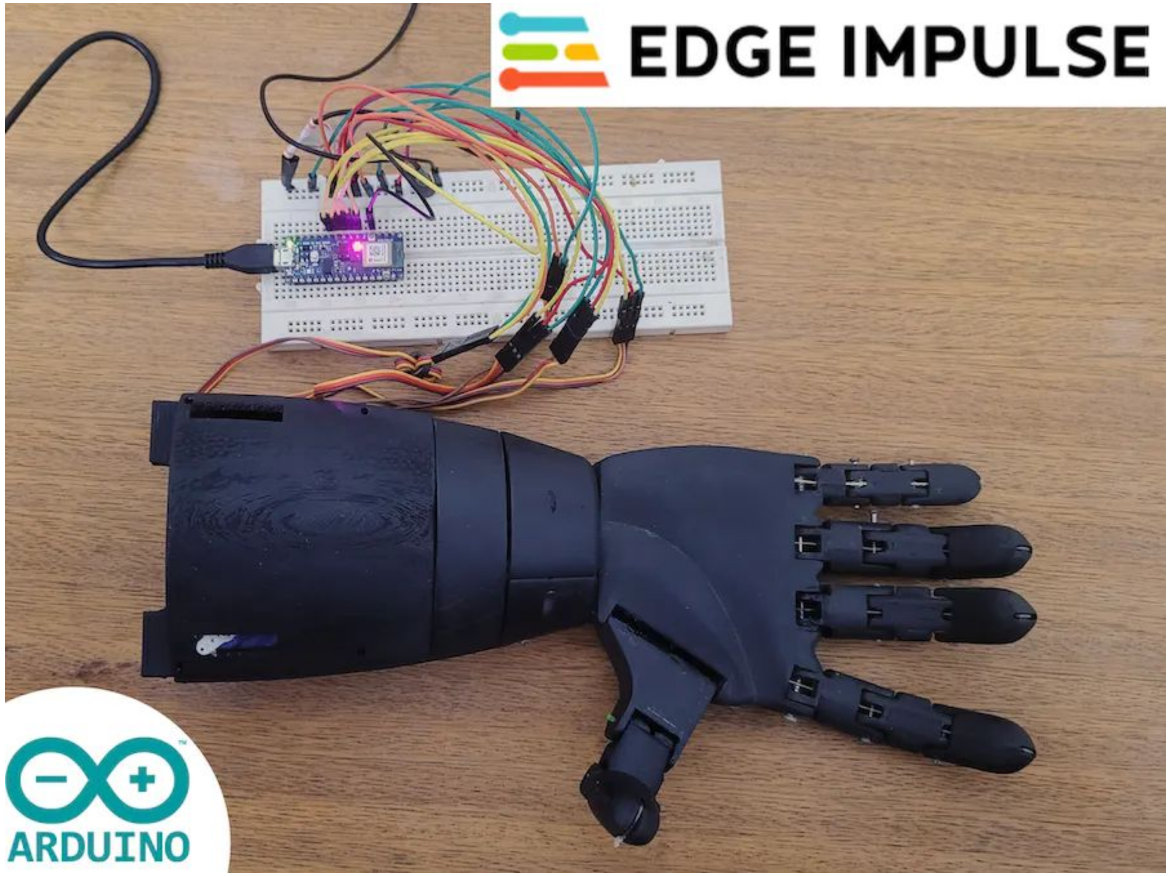

https://www.hackster.io/ex-machina/bionic-hand-voice-commands-module-w-edge-impulse-arduino-aa97e3

Automatic cough detection for BRD

(Bovine Respiratory Disease)

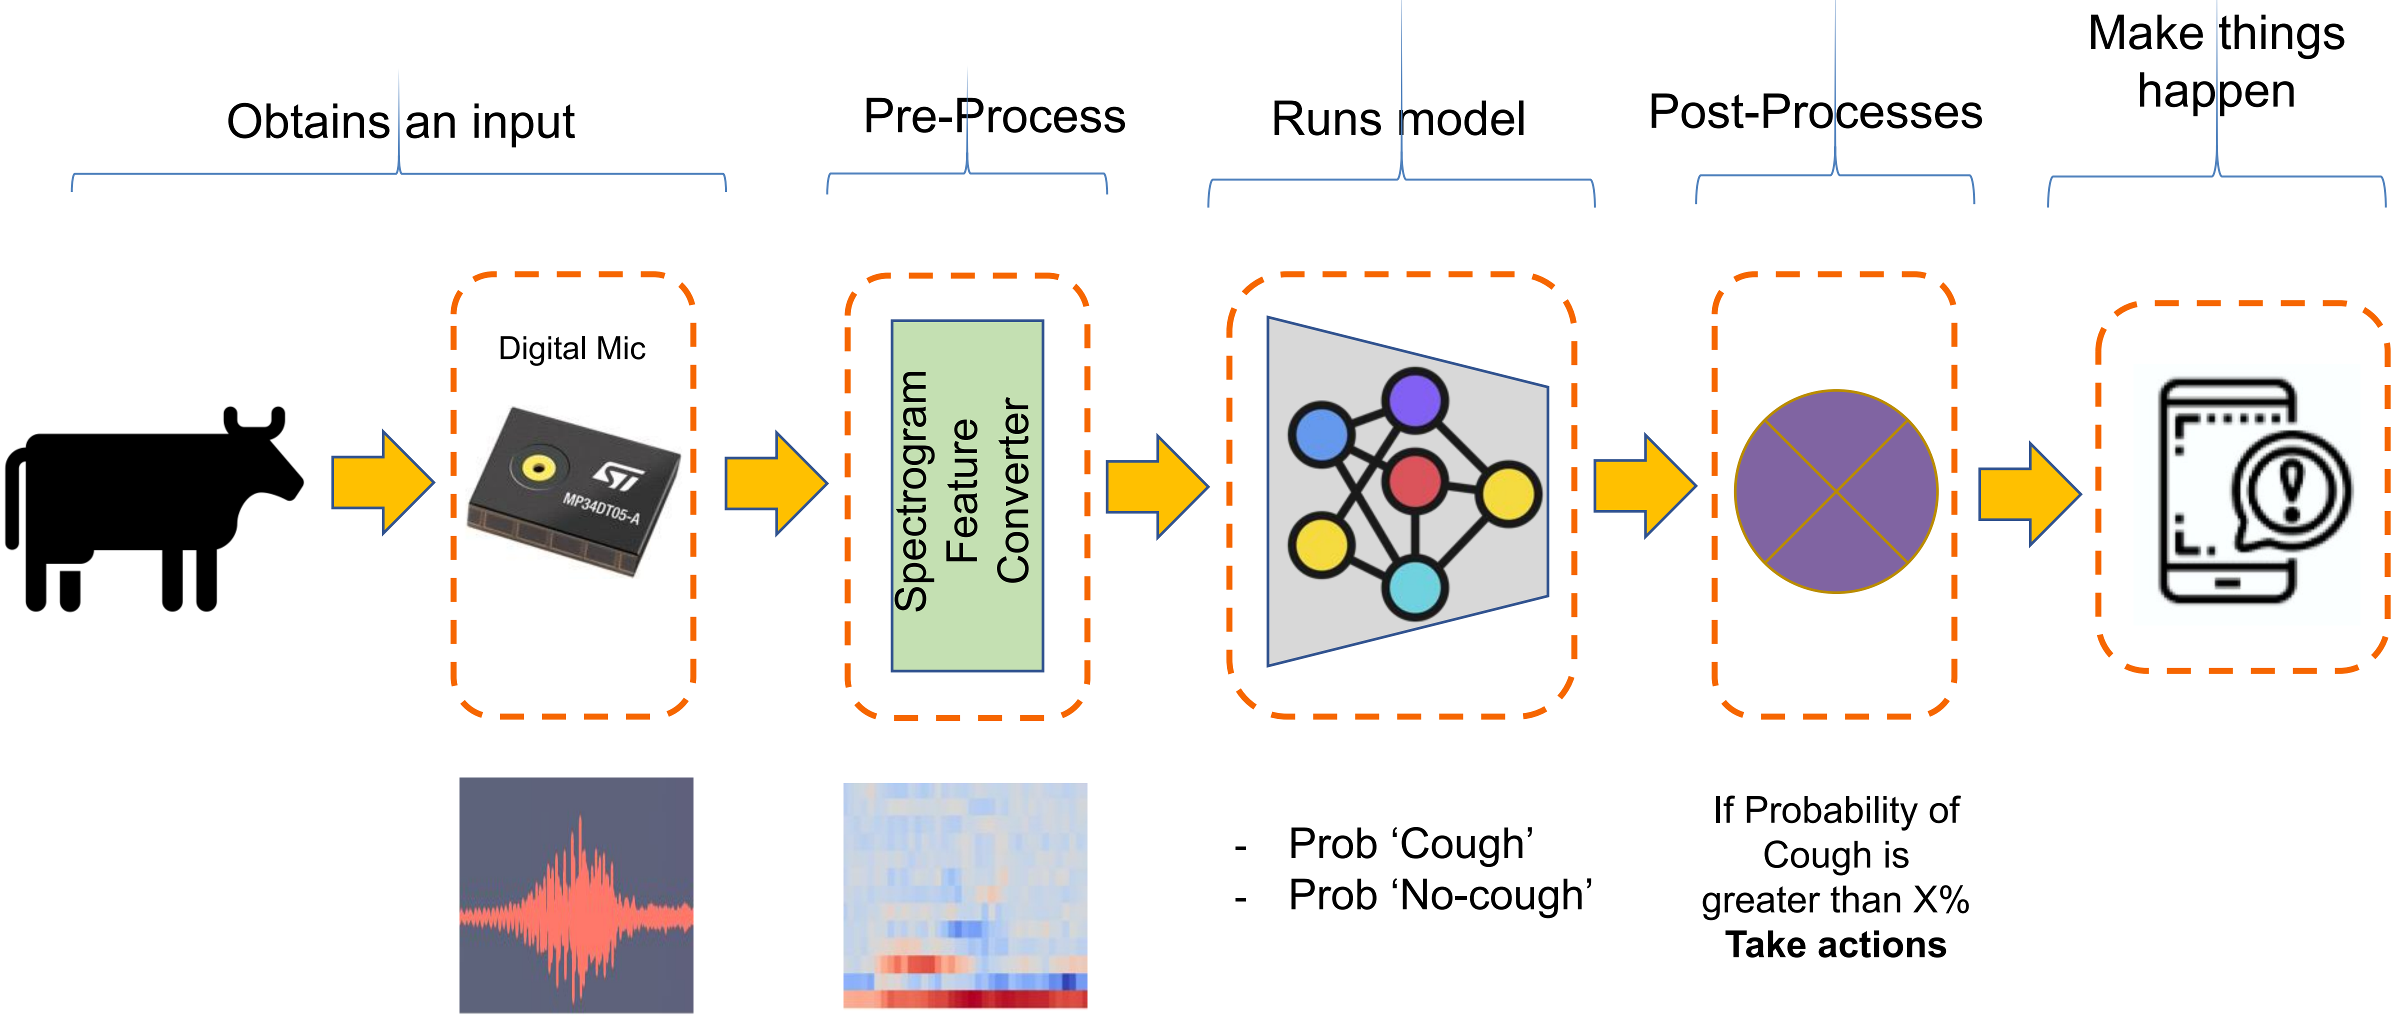

https://www.sciencedirect.com/science/article/abs/pii/S1537511017304762

Sound

Vibration

Vision

Movement Classification

CNPq UNIFEI

Predict and classify common Elephant behavior

Aggressive

Standing

Sleeping

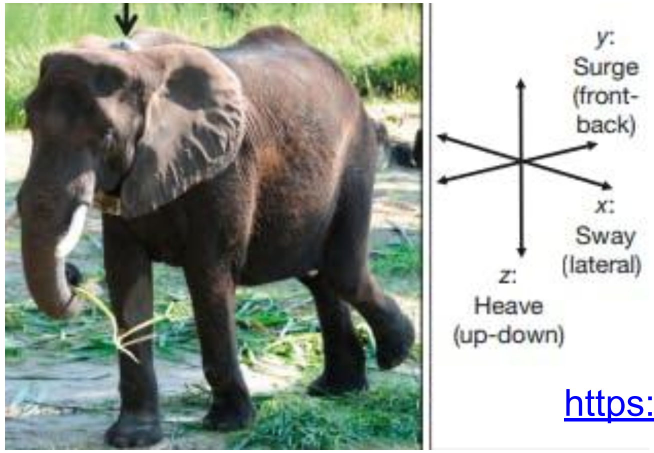

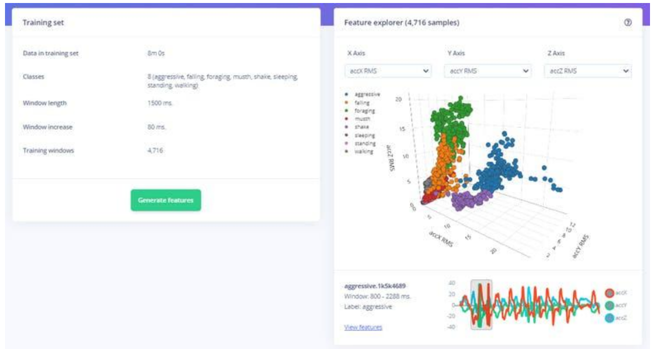

https://www.hackster.io/dhruvsheth_/eletect-tinyml-and-iot-based-smart-wildlife-tracker-c03e5a

Smart Cattle-Farm - Animal Behavior

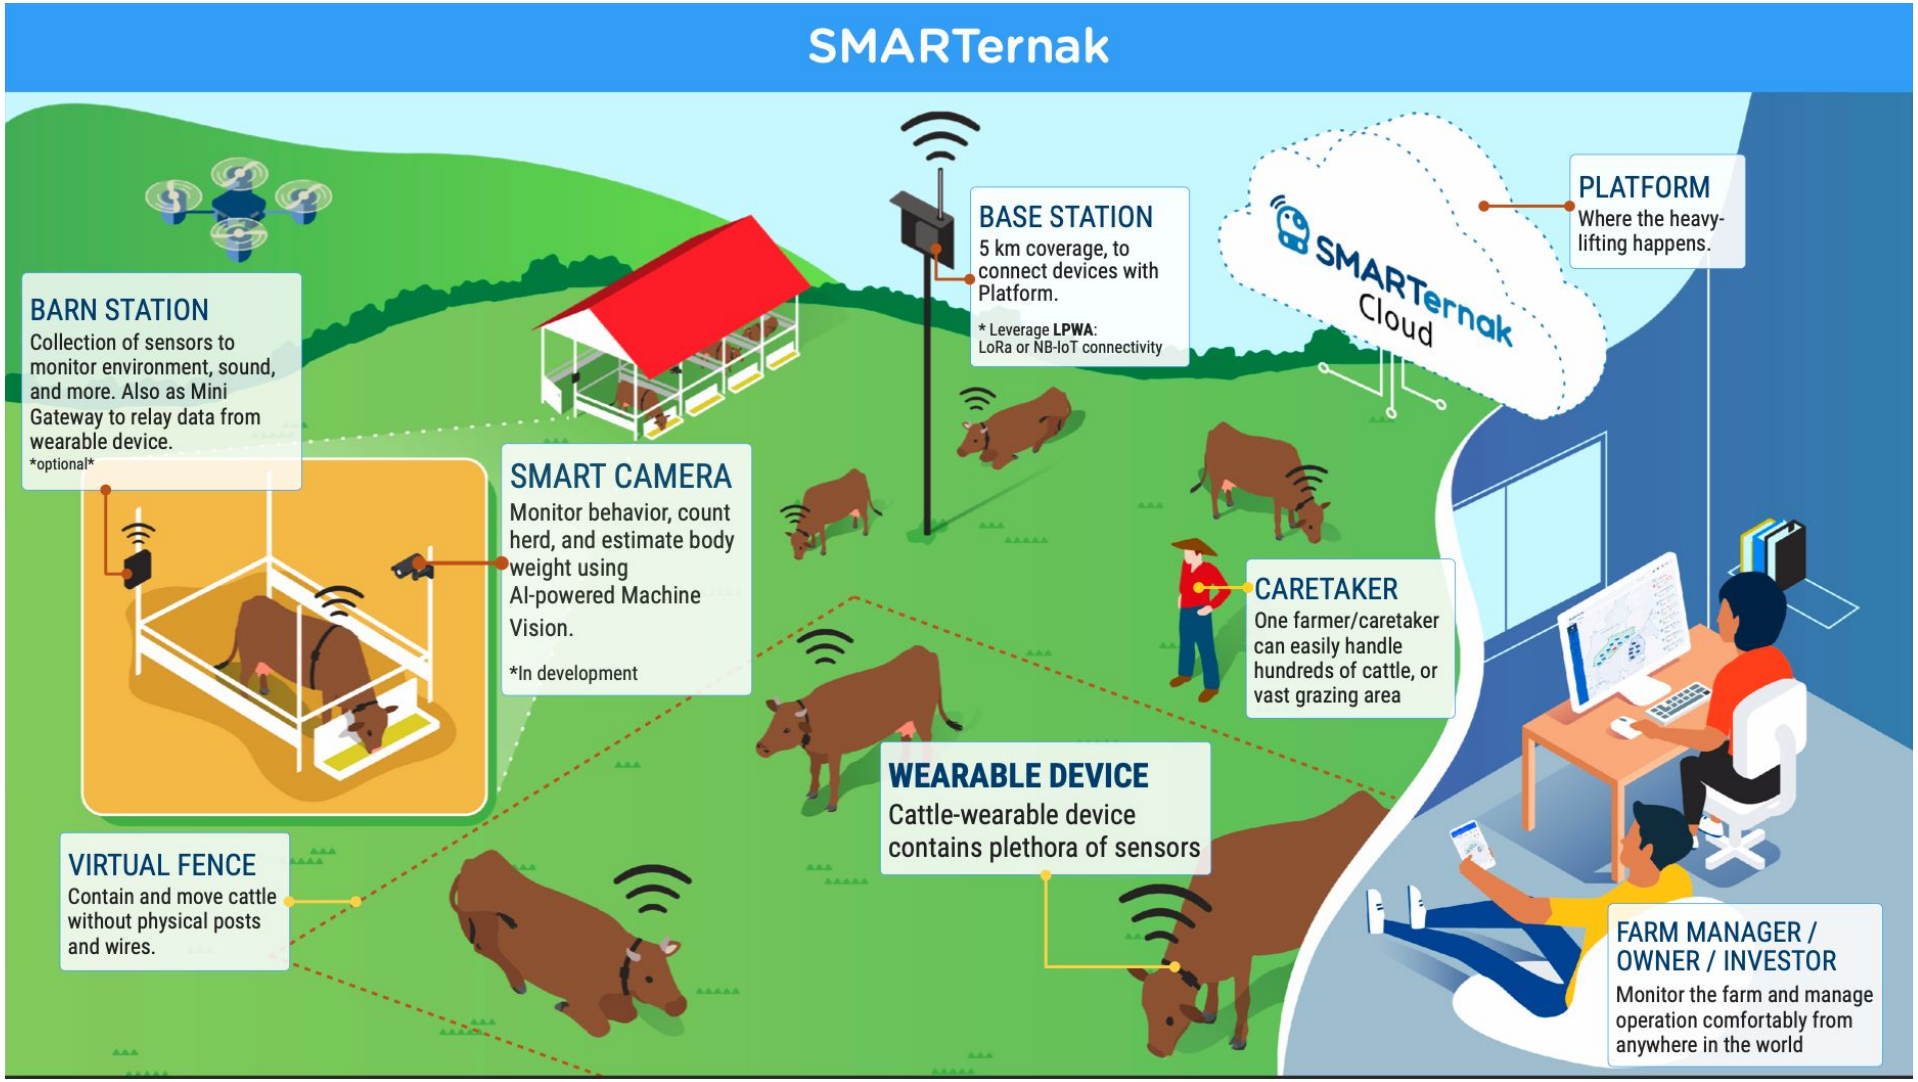

https://dycodex.com/smarternak/

On-device Activity Prediction

Video: http://bit.ly/st-feed-reg

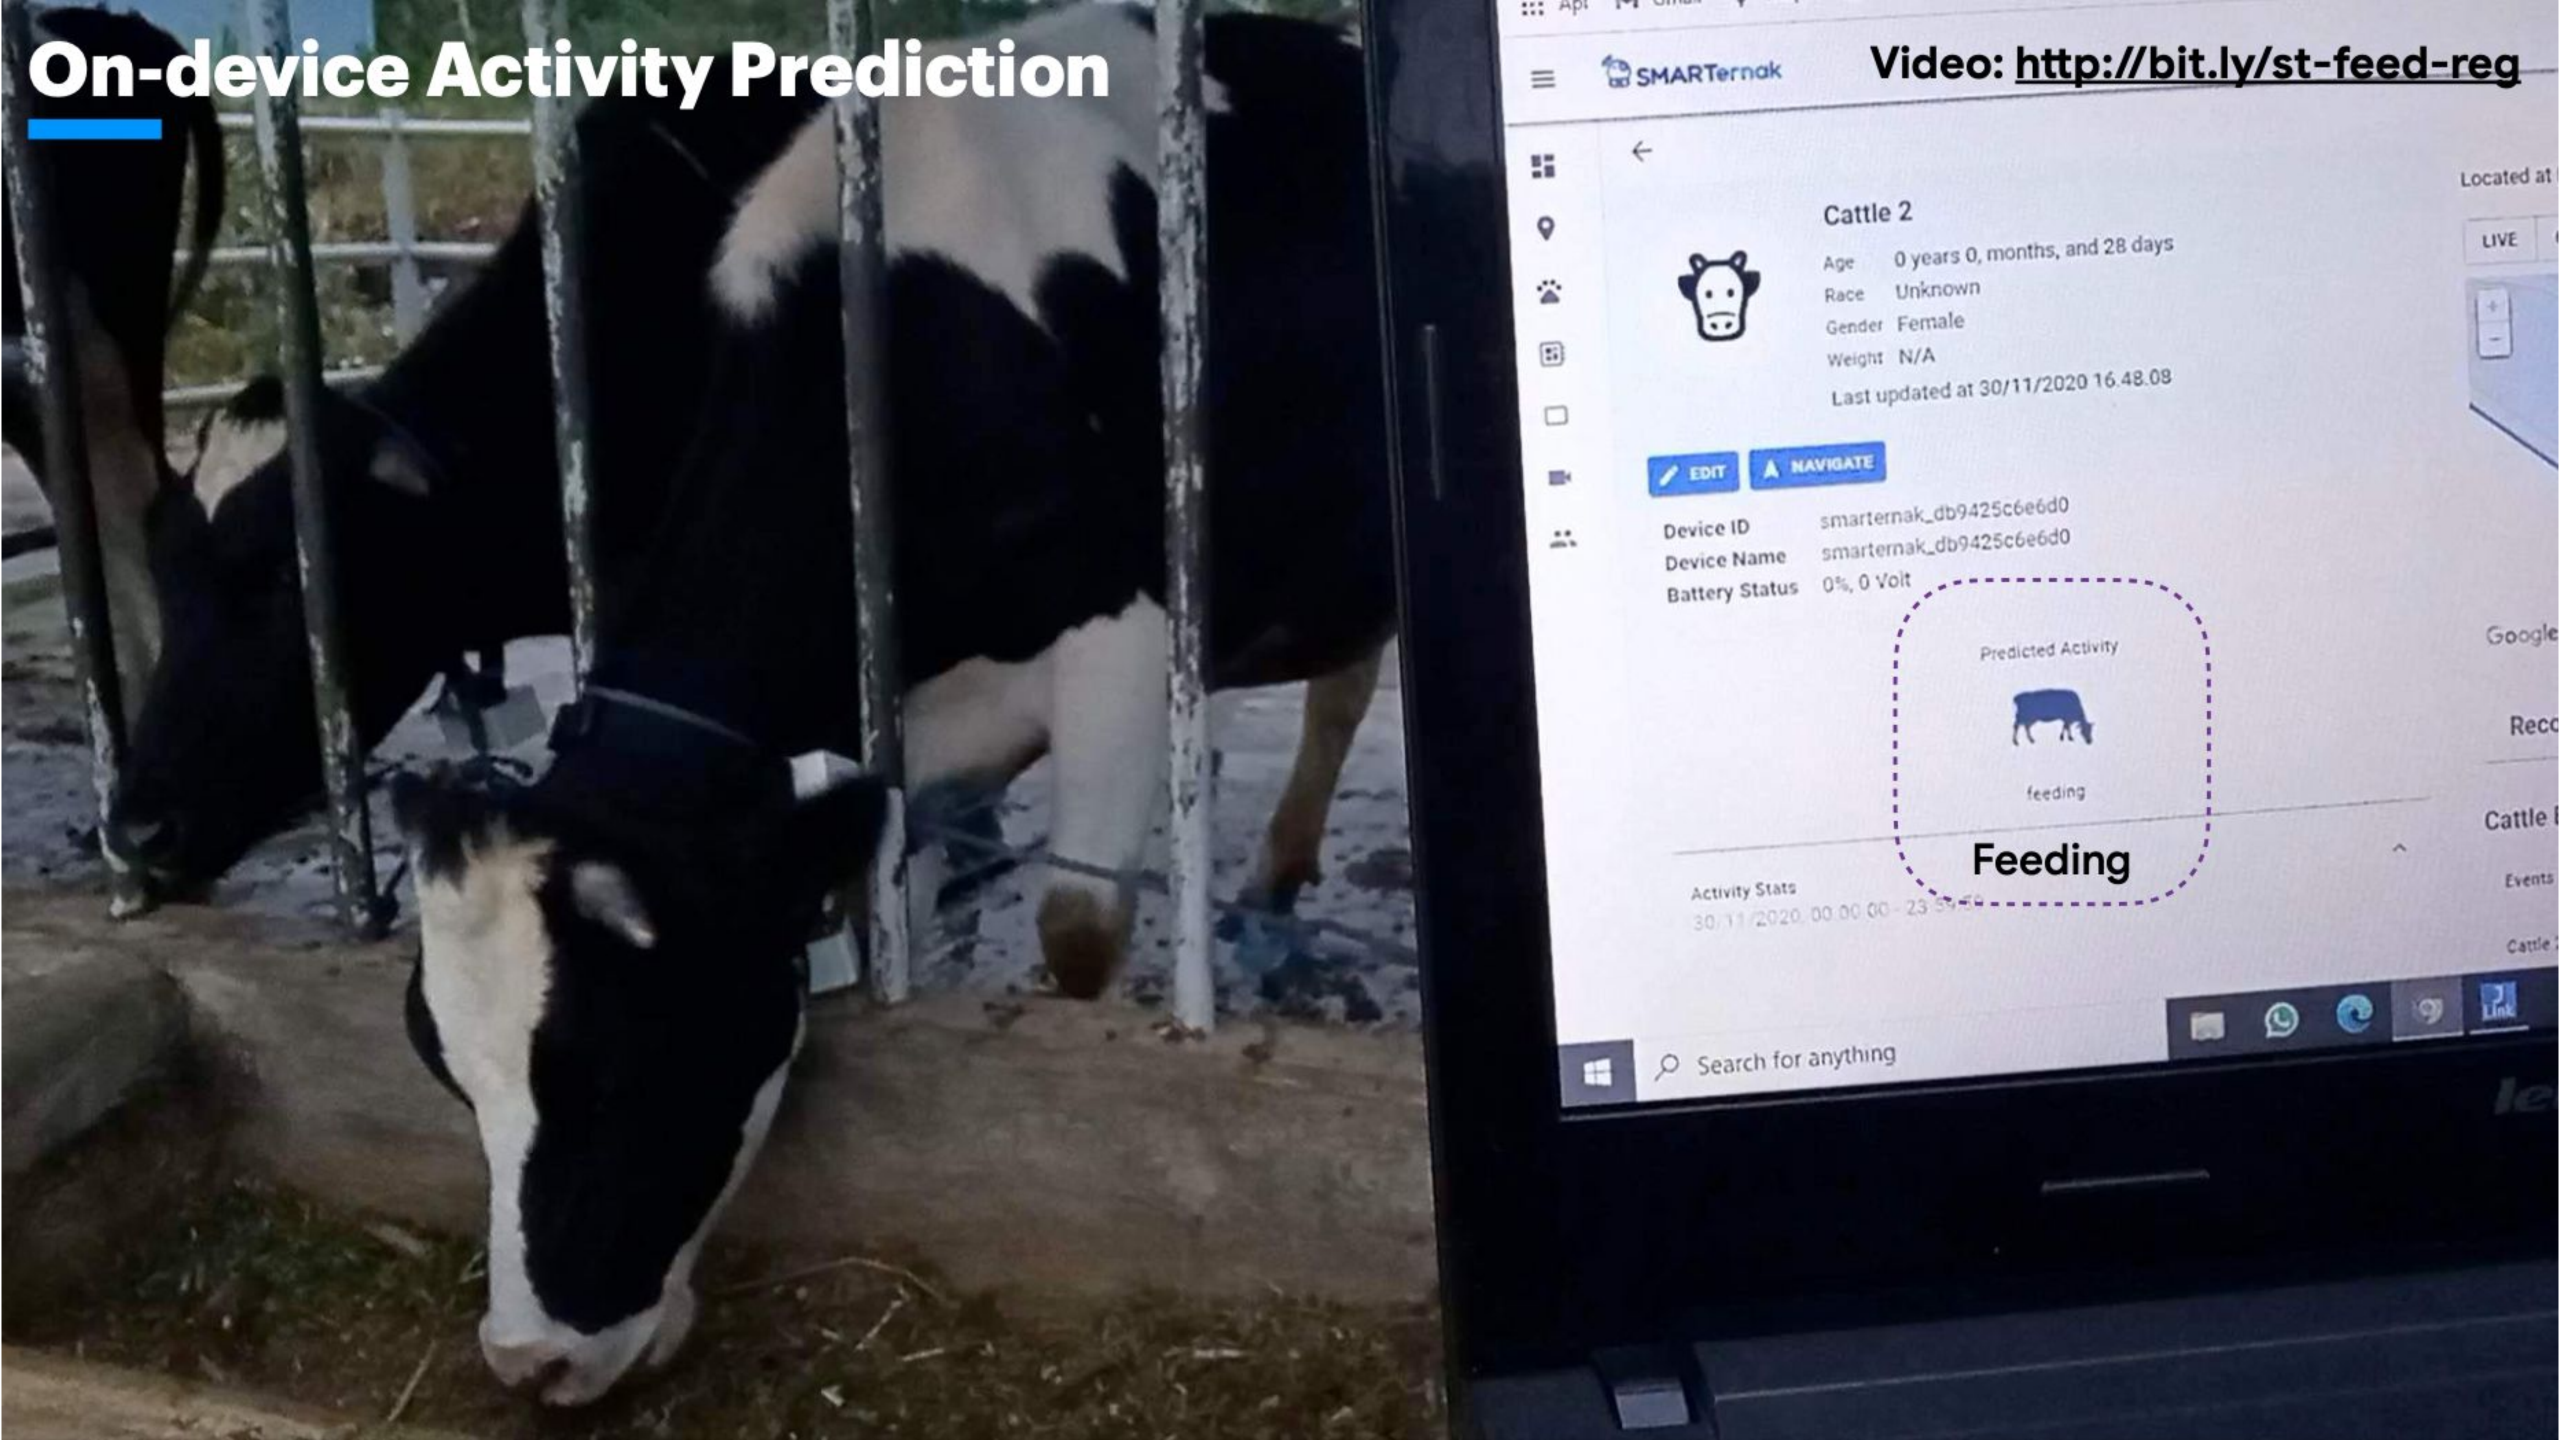

Sound

Vibration

Vision

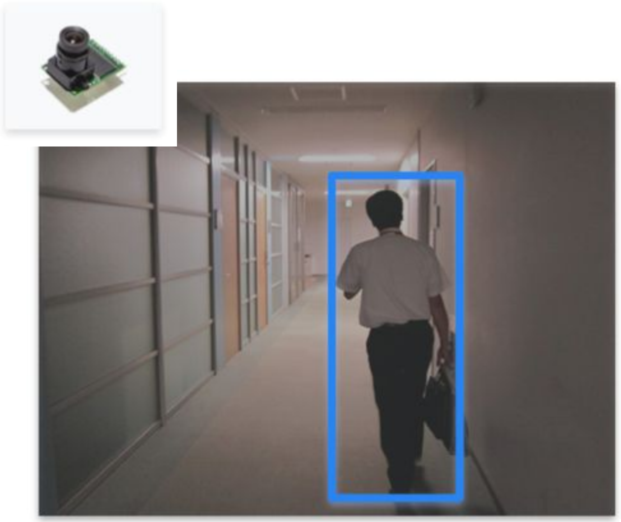

Computer Vision Recognition Tasks

Image Classification

(**Multi-Class** Classification)

Cat: 70%

Dog: 80%

Object Detection

Multi-Label Classification + Object **Localization**

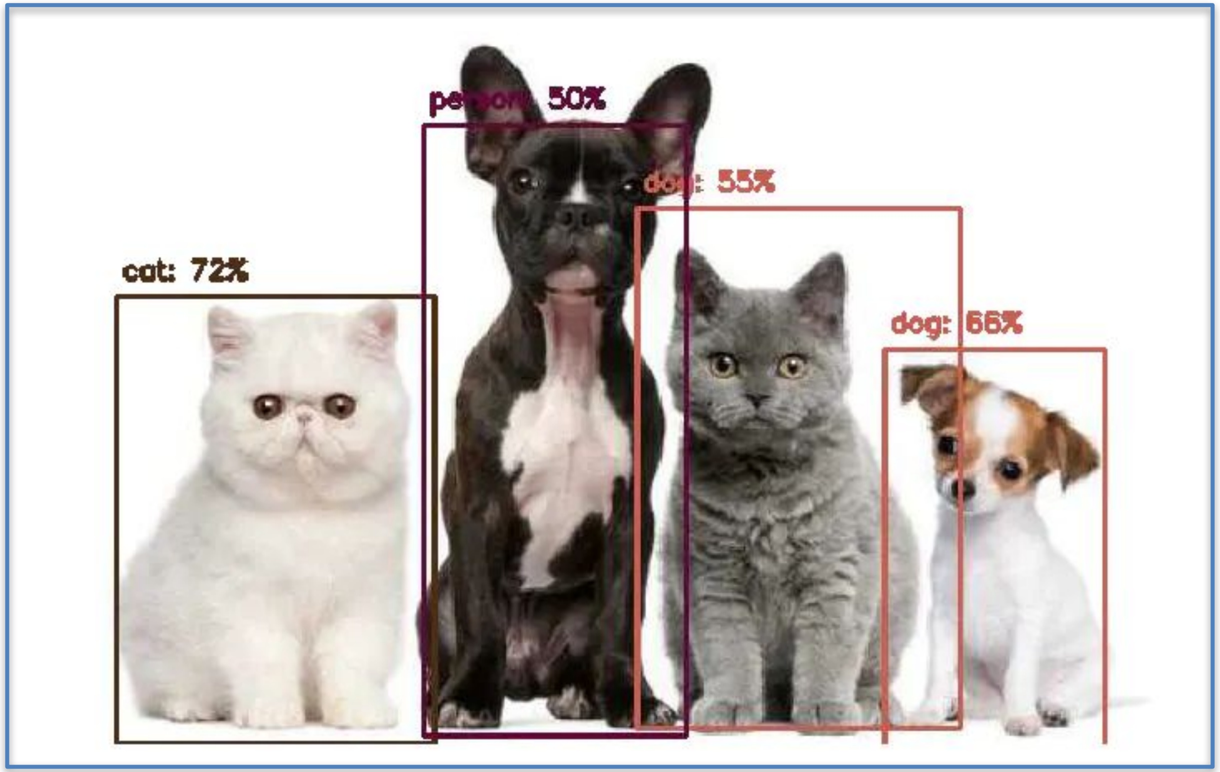

Computer Vision Recognition Tasks

Instance Segmentation

Each pixel in an image IS CLASSIFIED into a predefined category.

Pose Estimation

Key points (or landmarks) on the object, such as joints on a human or animal body are detected

Computer Vision Recognition Tasks

Image Classification

(**Multi-Class** Classification)

Cat: 70%

Dog: 80%

Object Detection

Multi-Label Classification + Object **Localization**

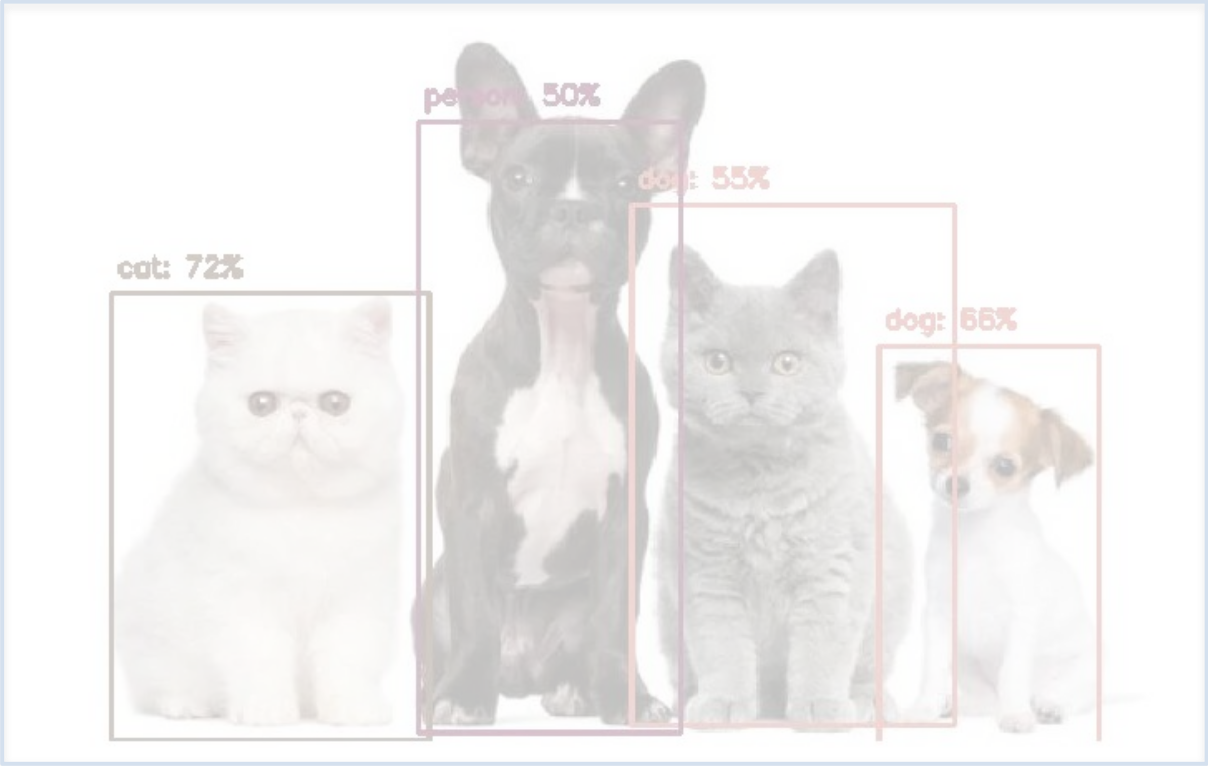

Detecting Diseases in the Bean plants

AIR Lab Makerere University

UGANDA

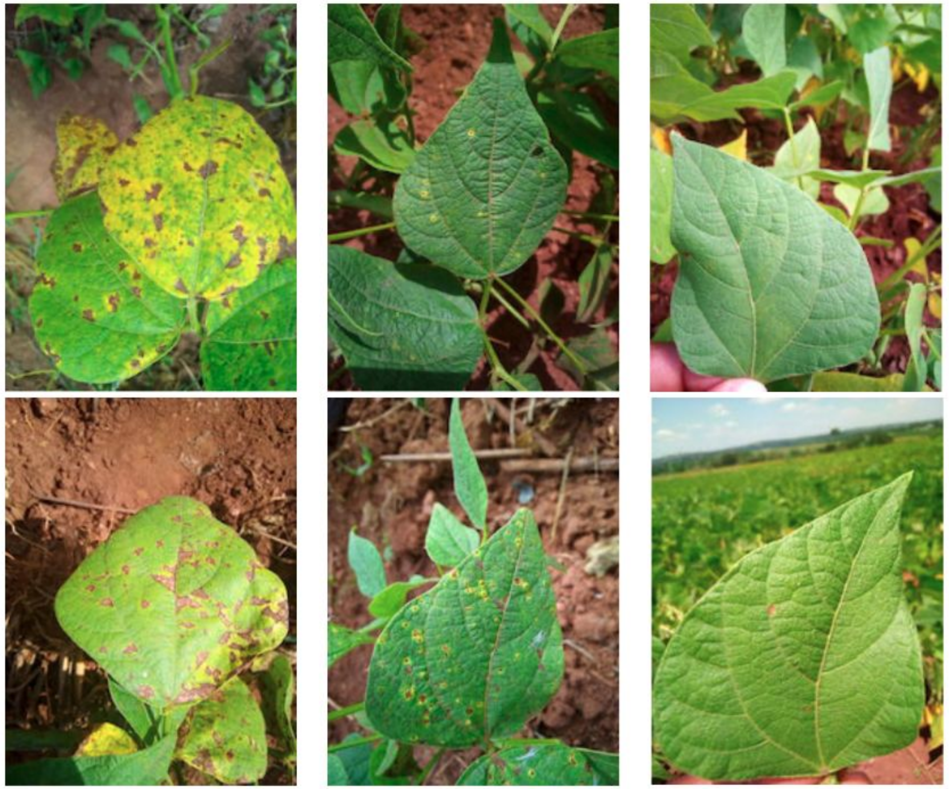

Angular Leaf Spot | Bean Rust | Healthy

Dataset: https://github.com/AI-Lab-Makerere/ibean/

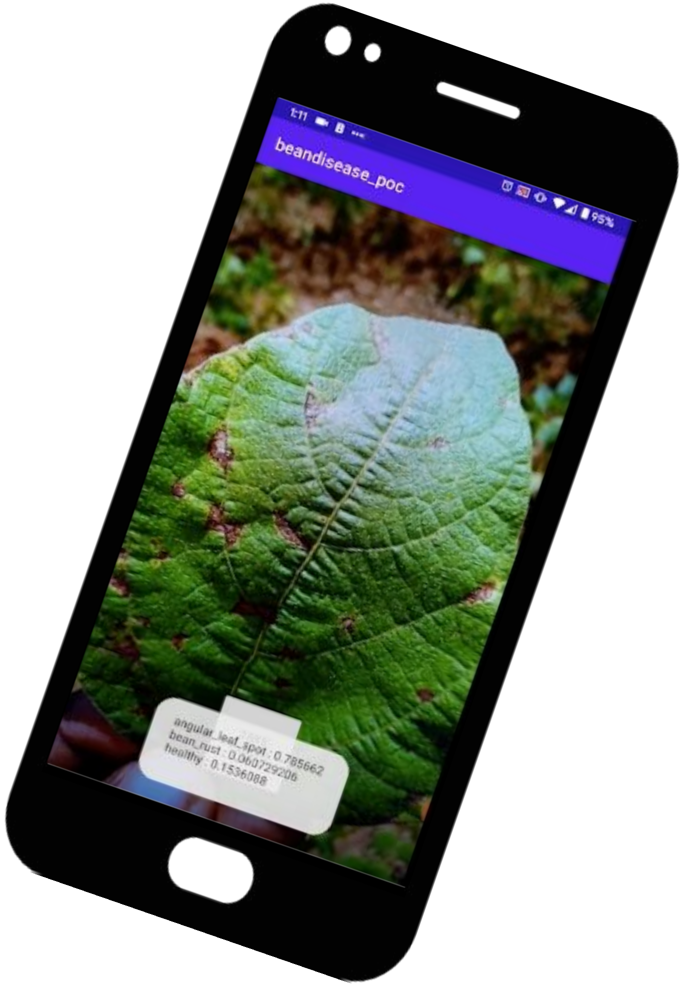

Learn the steps to build an app that detects crop diseases
(Android Studio)

Forest Fire Detection

[TinyML Aerial Forest Fire Detection](#)

[IESTI01 - Forest Fire Detection – Proof of Concept](#)

Coffee Disease Classification

João Vitor Yukio Bordin Yamashita
Graduando em Engenharia Eletrônica pela UNIFEI

https://www.hackster.io/Yukio/coffee-disease-classification-with-ml-b0a3fc

Automated Cow Body Condition Scoring (BCS)

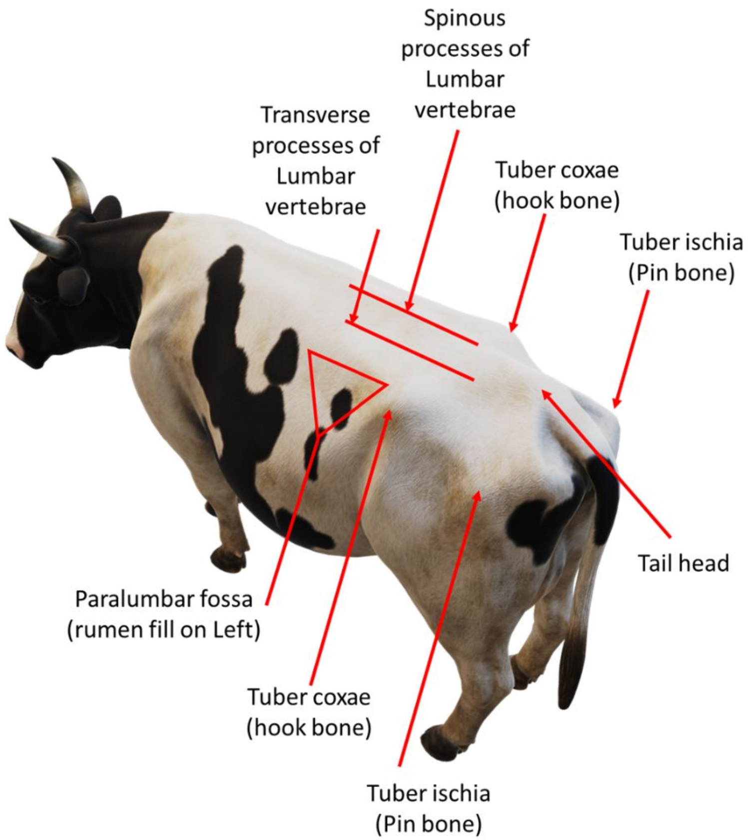

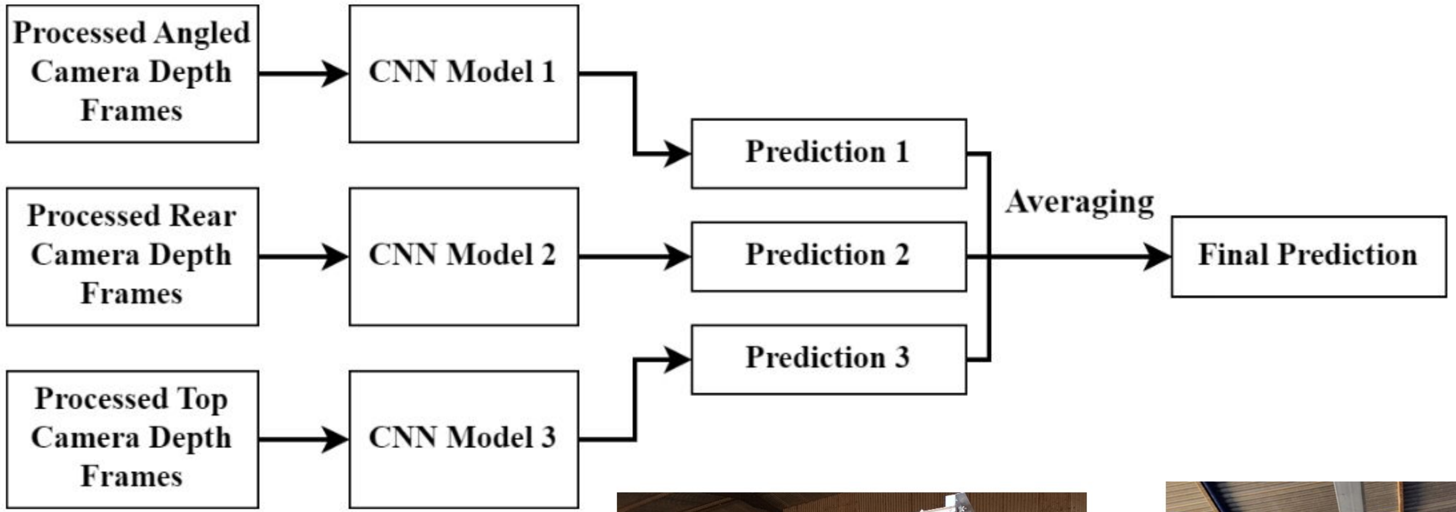

https://www.mdpi.com/1424-8220/23/22/9051

Computer Vision Recognition Tasks

Image Classification
(**Multi-Class** Classification)

Cat: 70%

Dog: 80%

Object Detection
Multi-Label Classification + Object **Localization**

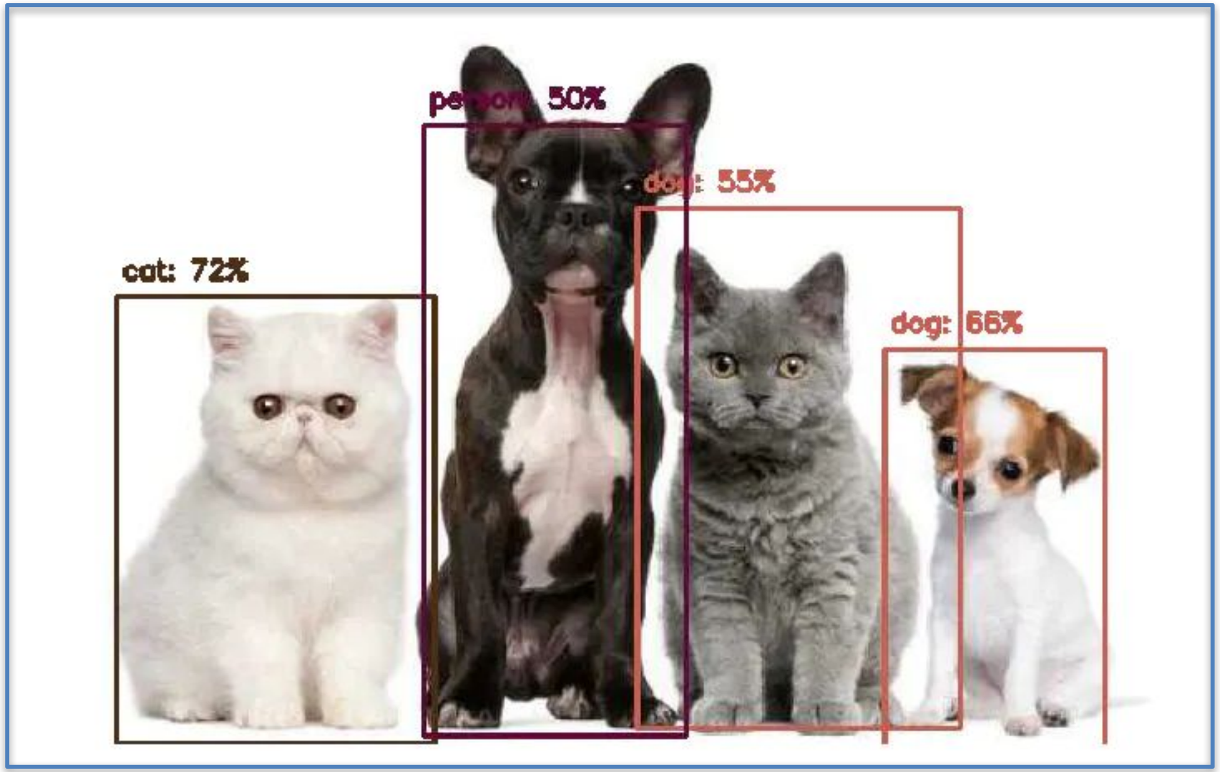

Bee Counting

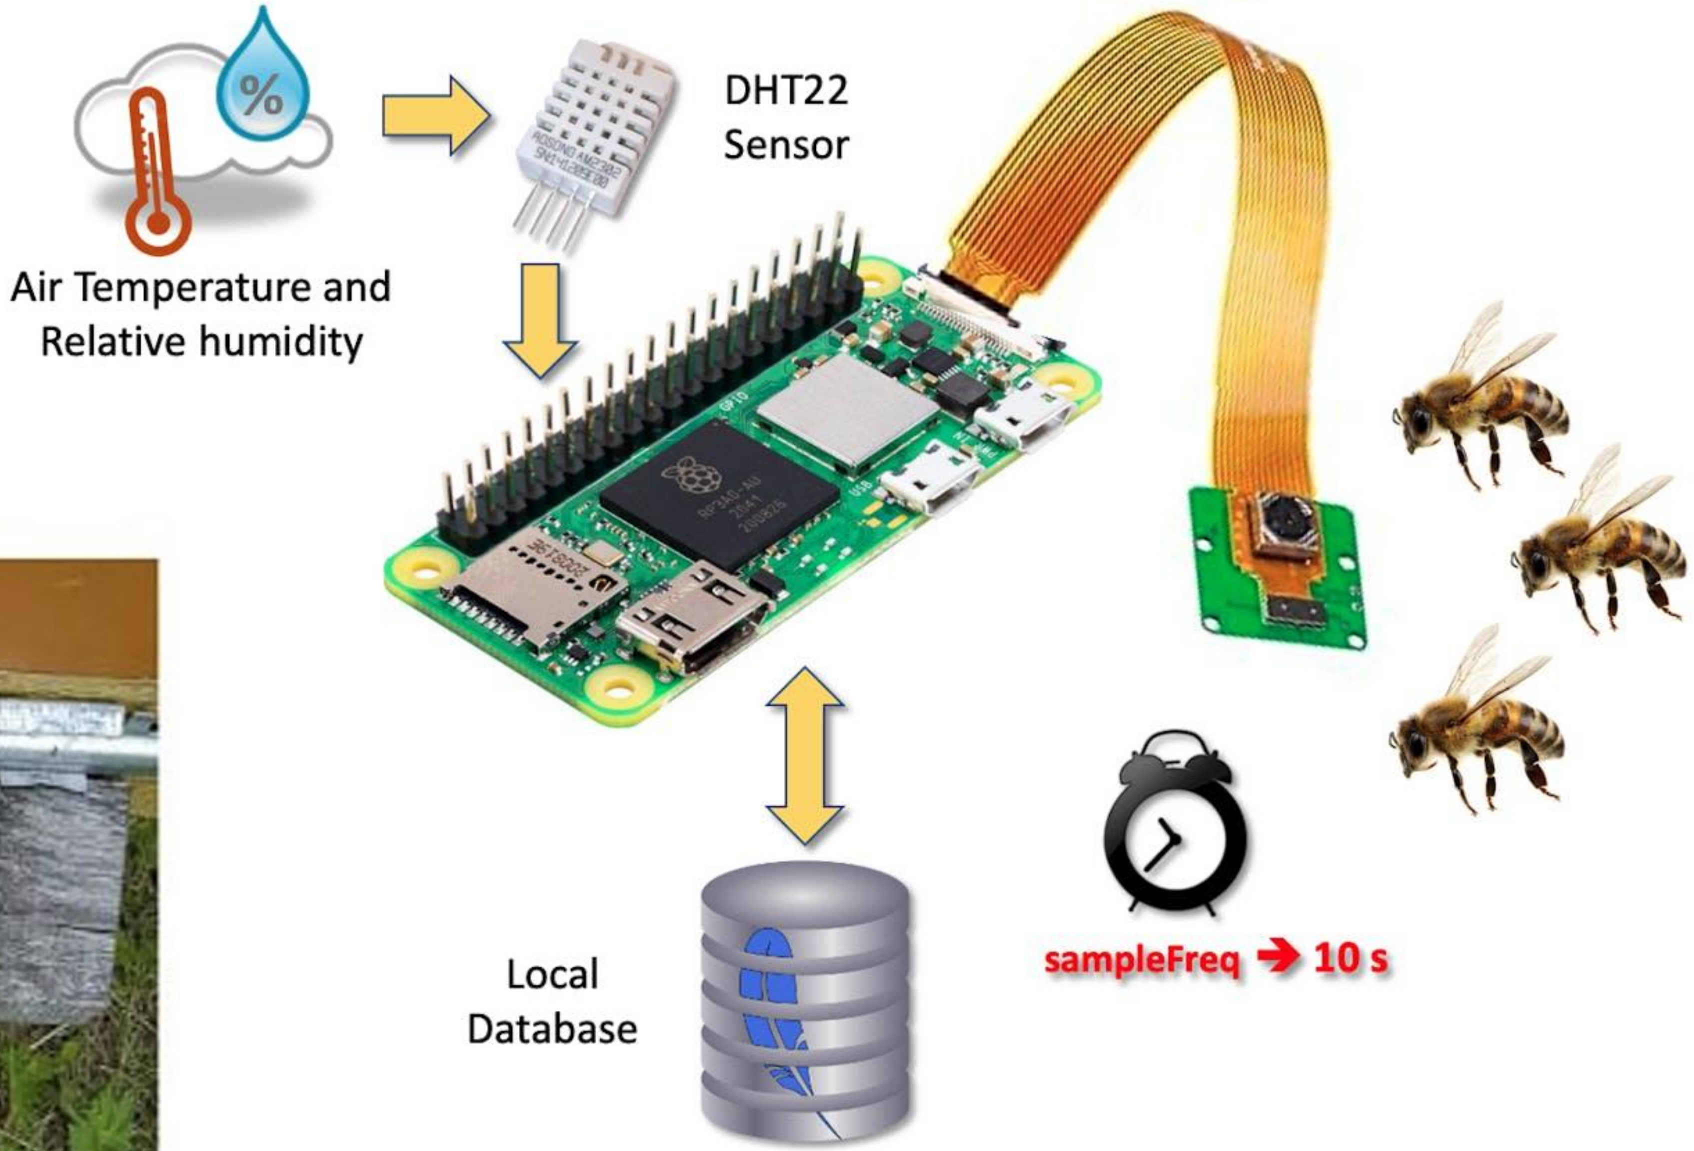

BuzzTech: Machine Learning at the Edge

Ant Detection

Ant Detection

Perimeter Fence Detection

Improper wiping https://www.cattle-care.com/individual

Breed Identification

https://www.folio3.ai/animal-detection/

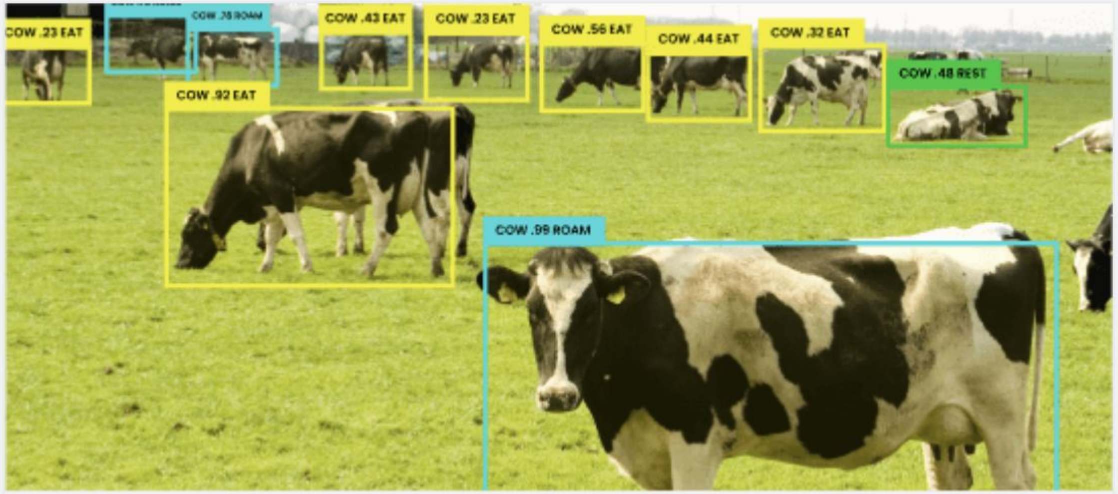

Cattle Gender/Pose Detection

Other Sensors / Models

Examples

Atrial Fibrillation Detection on ECG

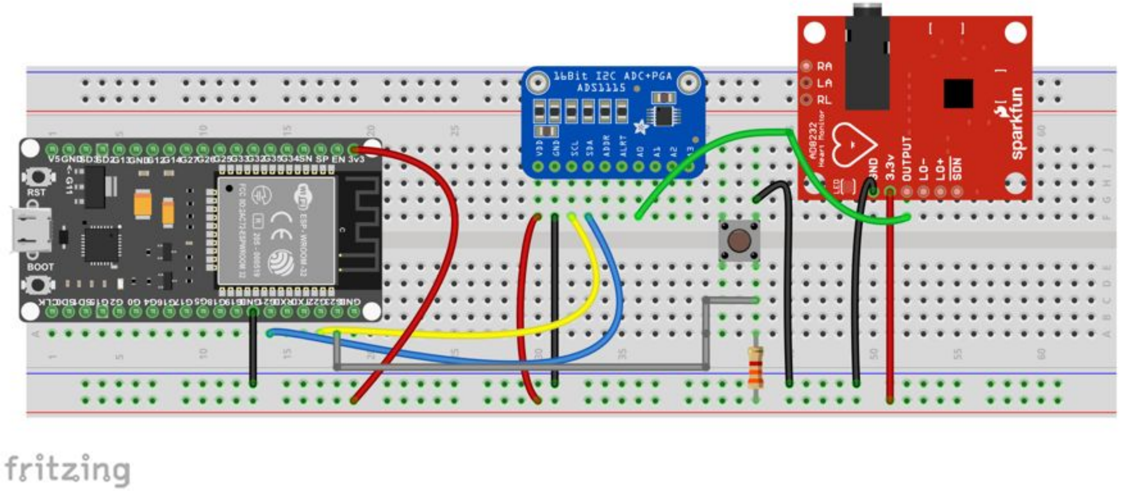

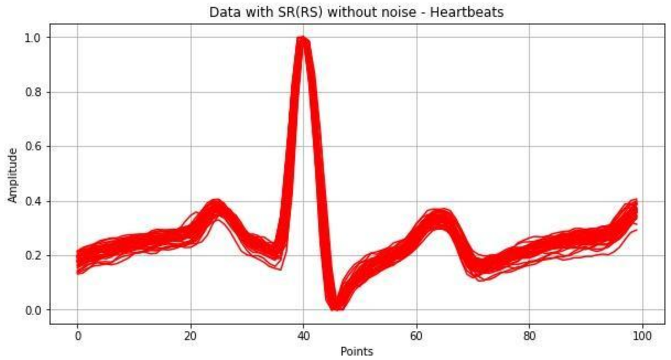

Guilherme Silva
Engenheiro - UNIFEI

[Atrial Fibrillation Detection on ECG using TinyML](#)
[Silva et al. UNIFEI 2021](#)

ECG: Detecting Cardiac Diseases in Cattle

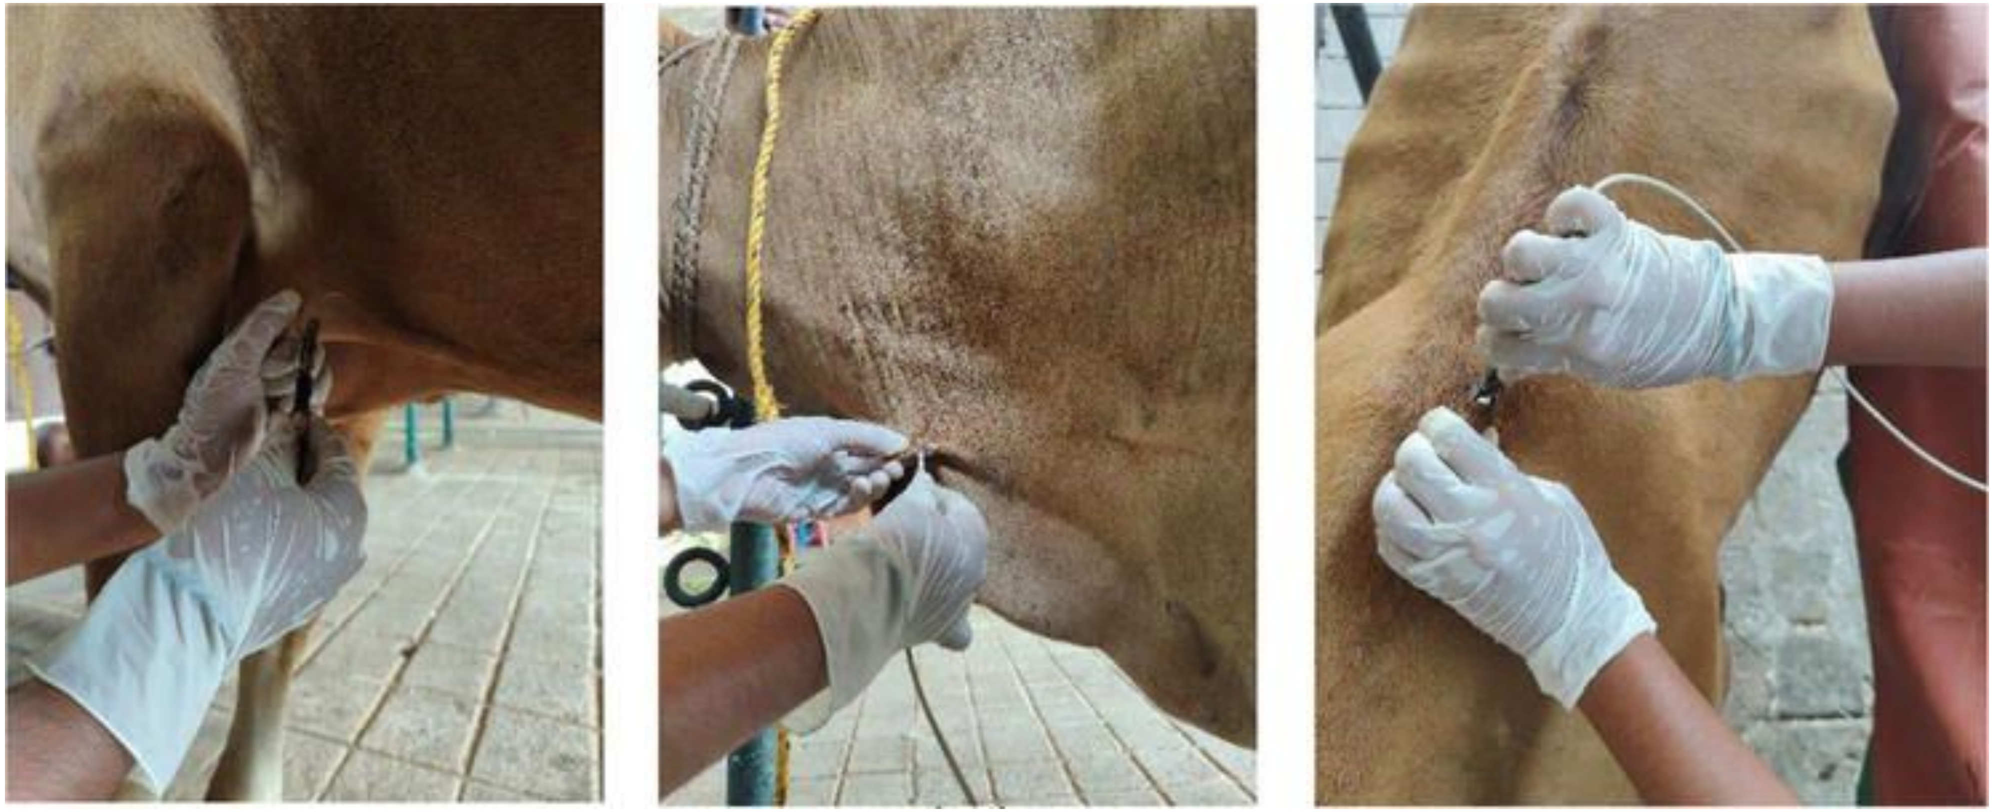

Interpretation of ECG in cattle is useful in many medical-veterinarian decisions because it permits the observation of many heart diseases in cattle (myocardial dystrophy, atrial and ventricular enlargements, pericardium effusion, etc.).

https://www.intechopen.com/chapters/82061

Humidity Forecasting (RNN-LSTM)

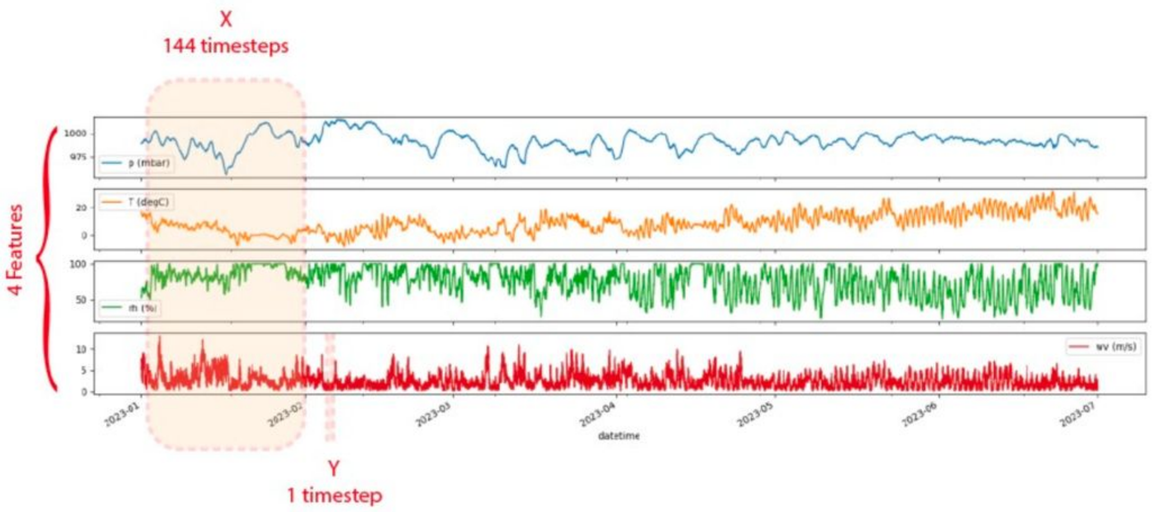

ESP32 LSTM Phenolic Sponge Moisture

Content

- AI Introduction and Applications
- Lab – Image Classification

Machine Learning Workflow

Machine Learning Workflow ("Where")

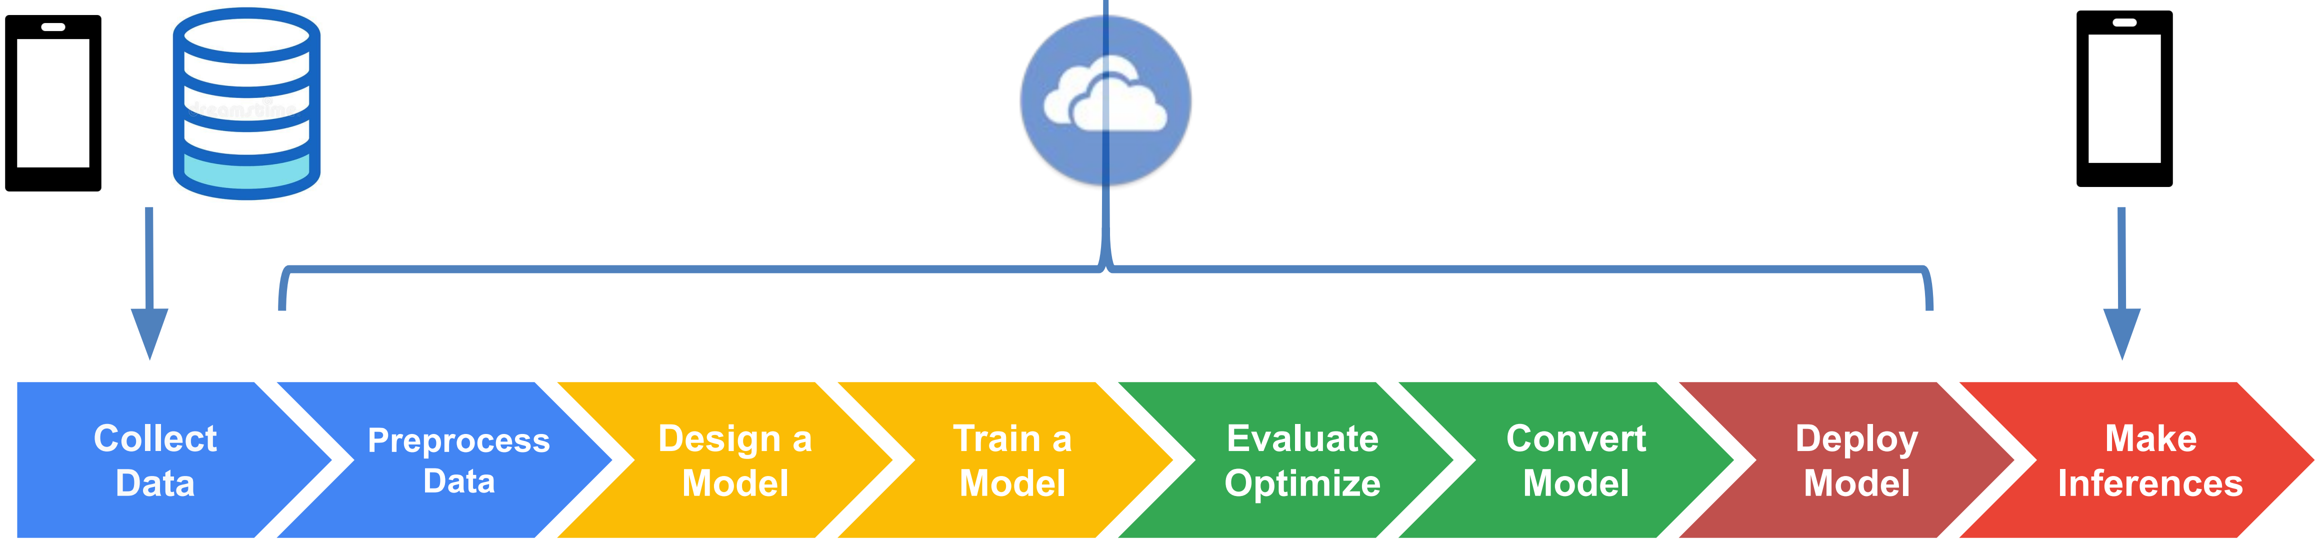

Machine Learning Workflow ("How")

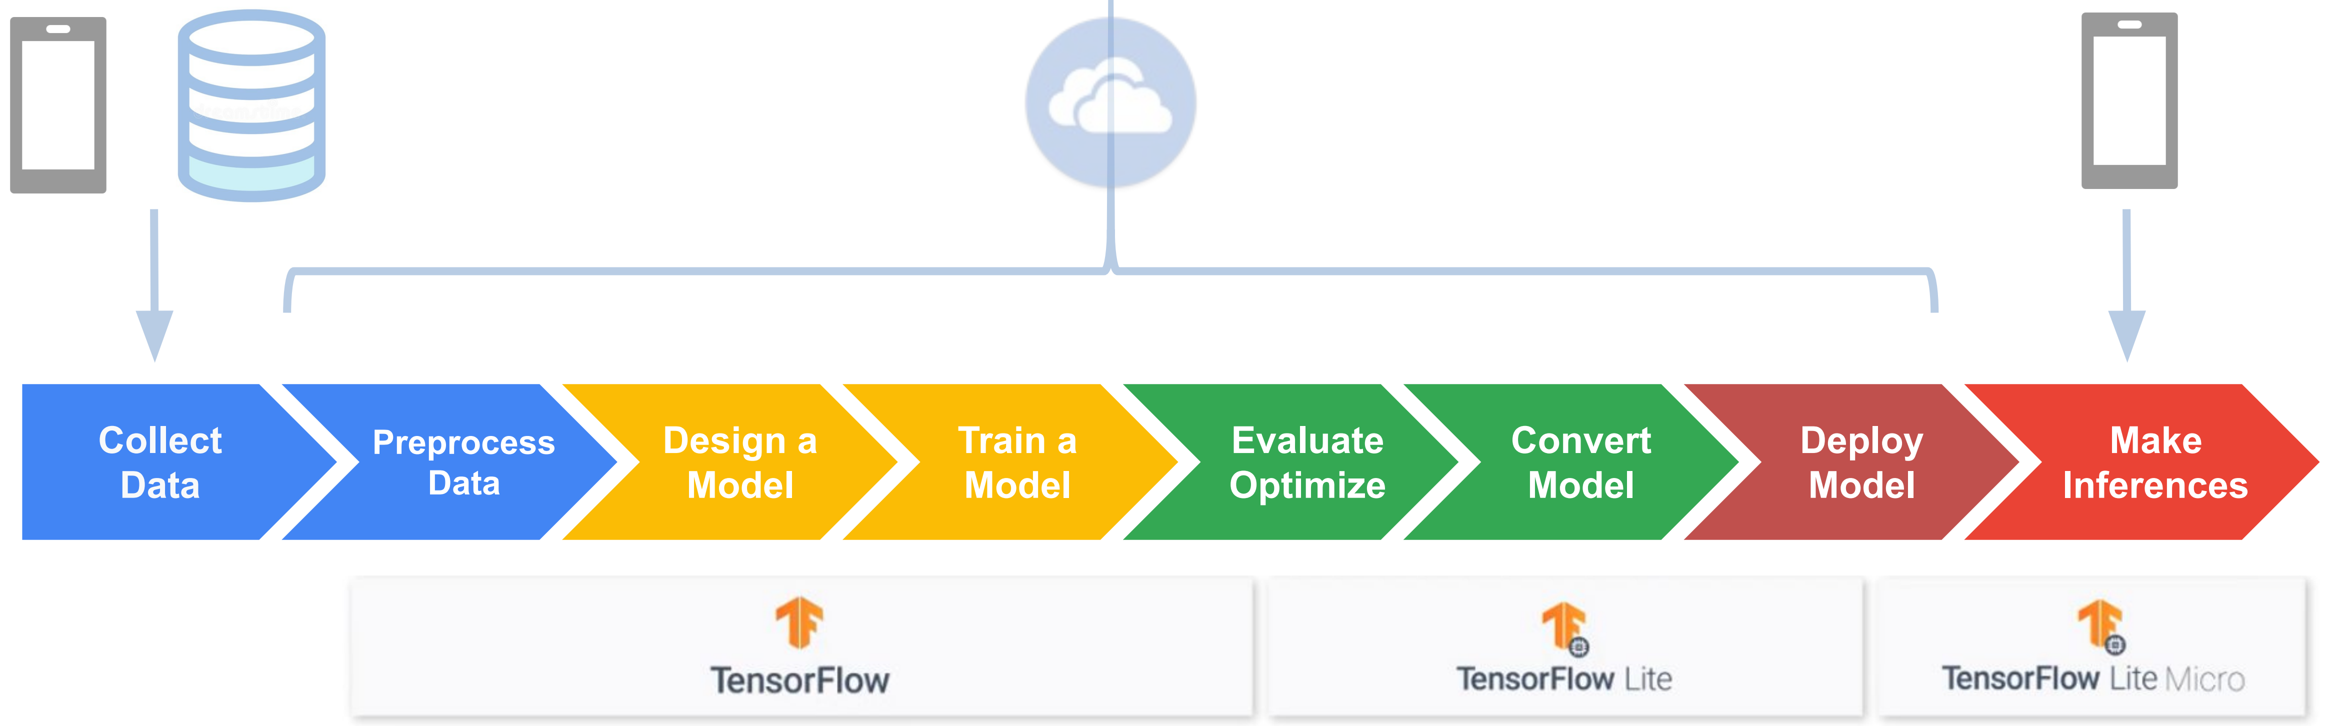

Machine Learning Workflow ("How")

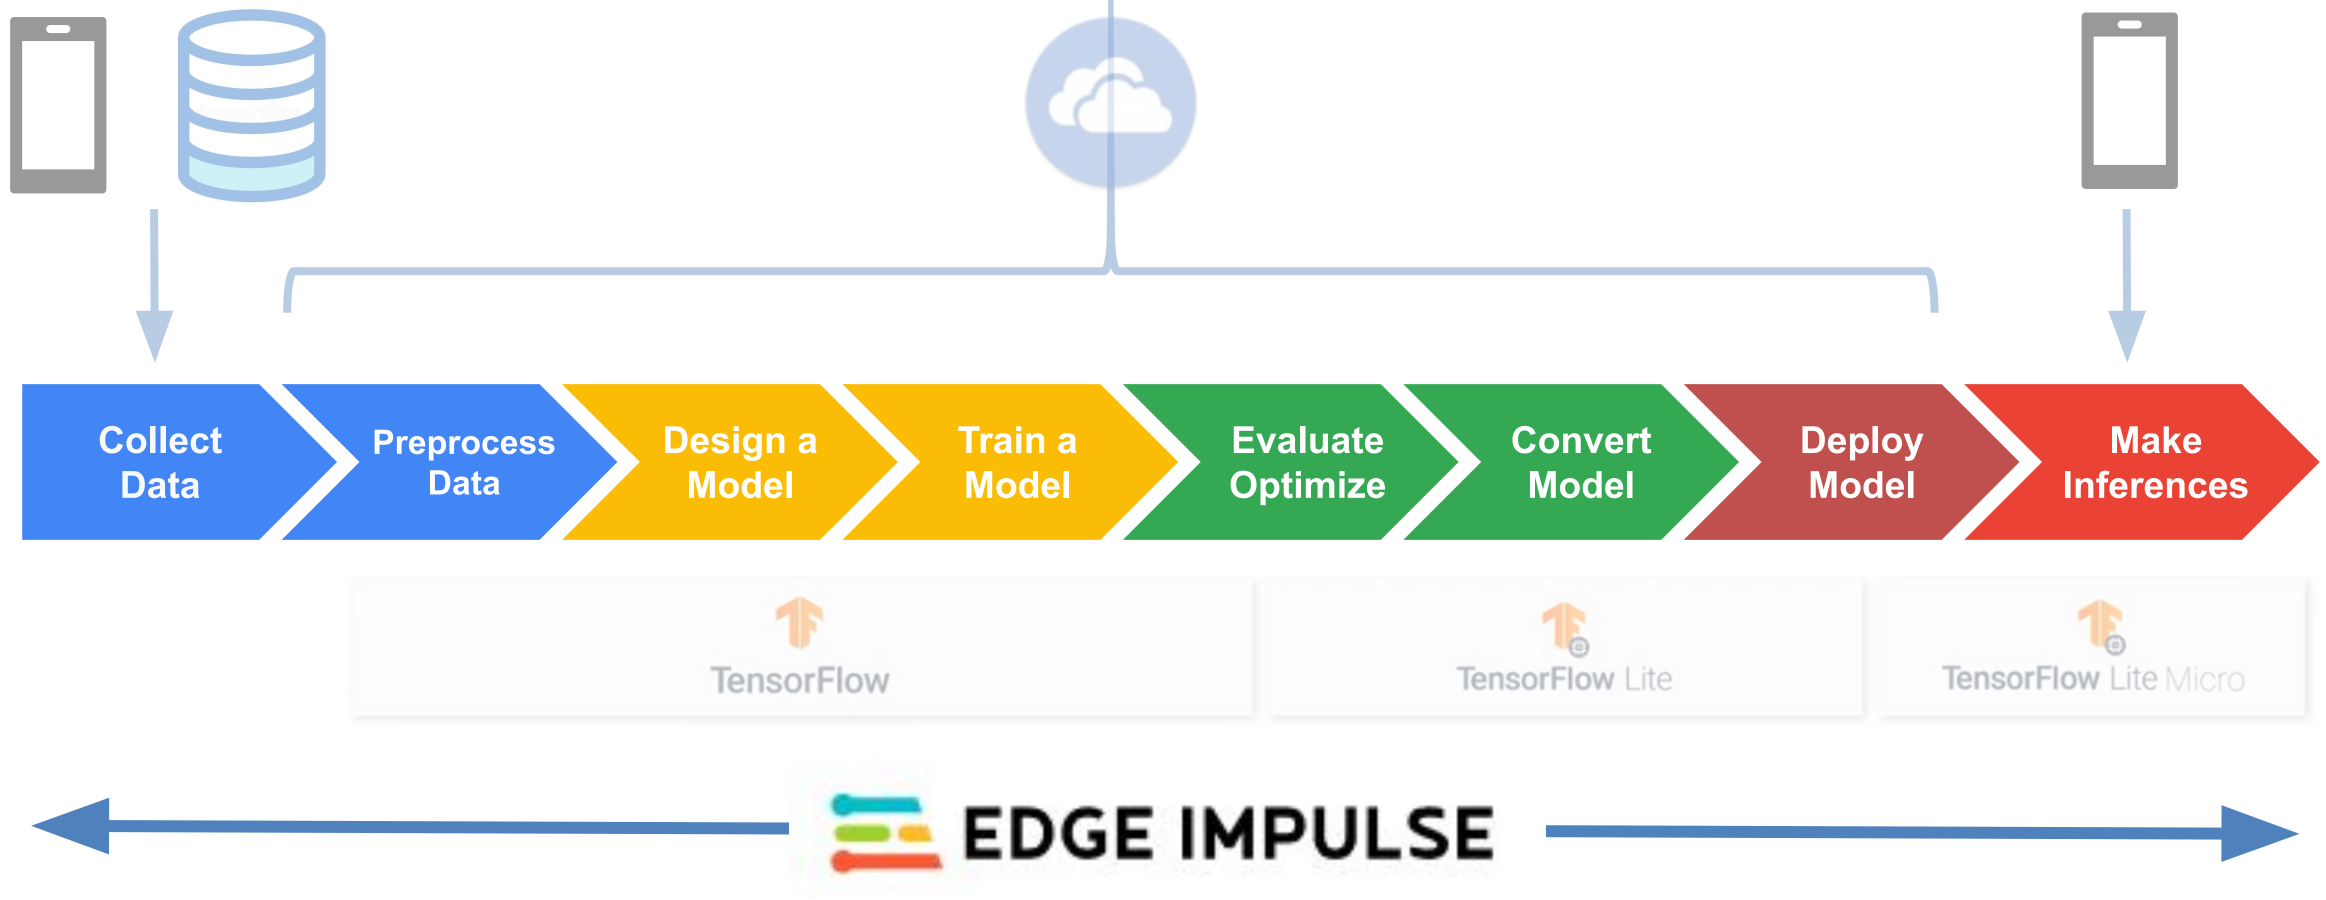

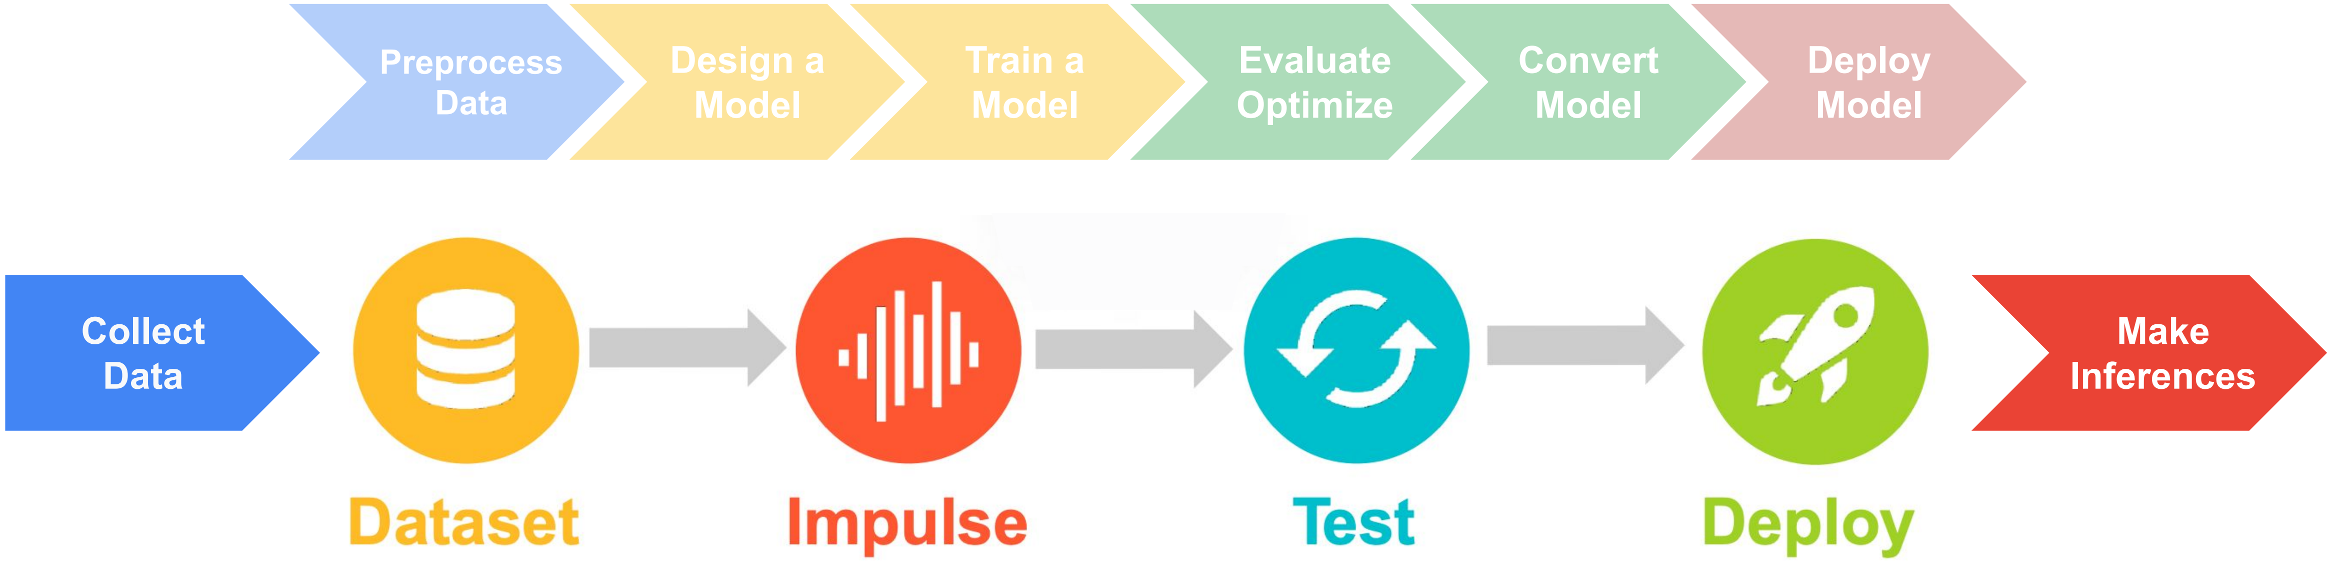
Preprocess Data
Design a Model
Train a Model
Evaluate Optimize
Convert Model
Deploy Model
Collect Data
Dataset
Impulse
Test
Deploy
Make Inferences

EI Studio - Embedded ML platform ("AutoML")

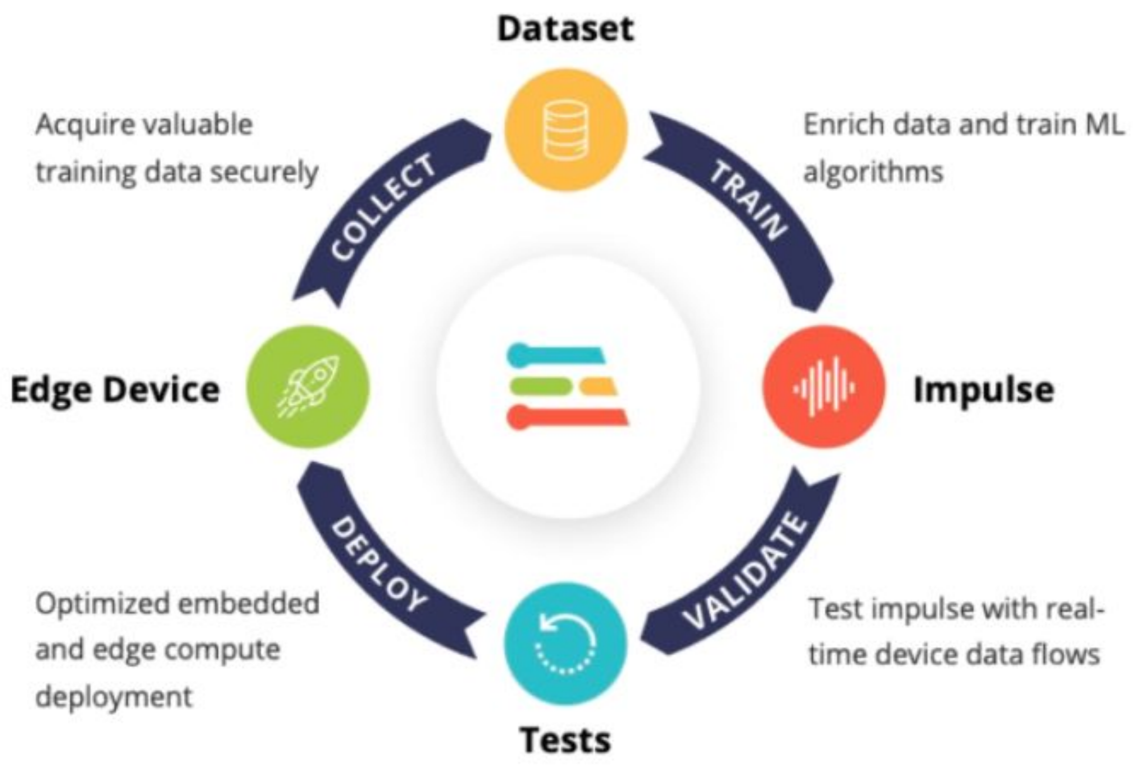

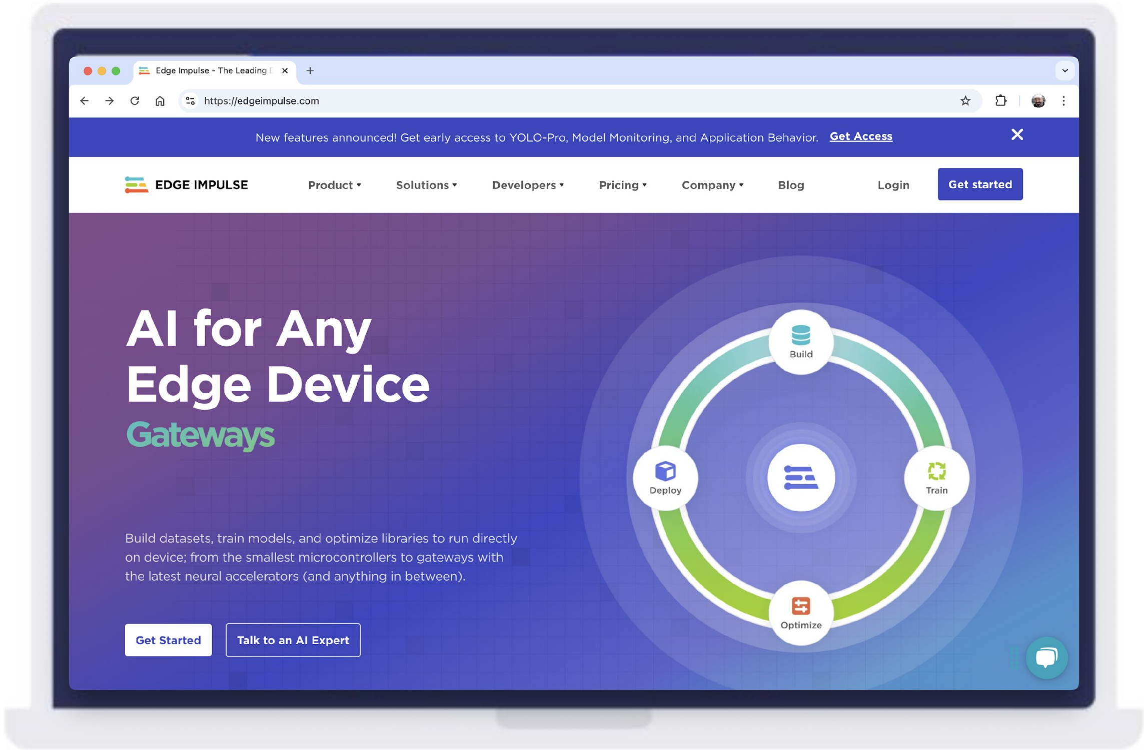

Learn more at http://edgeimpulse.com

[2023 SciTinyML Workshop - Edge Impulse Overview by Shawn Hymel]

Image Classification Project
Edge Impulse Studio

https://studio.edgeimpulse.com/public/540912/live

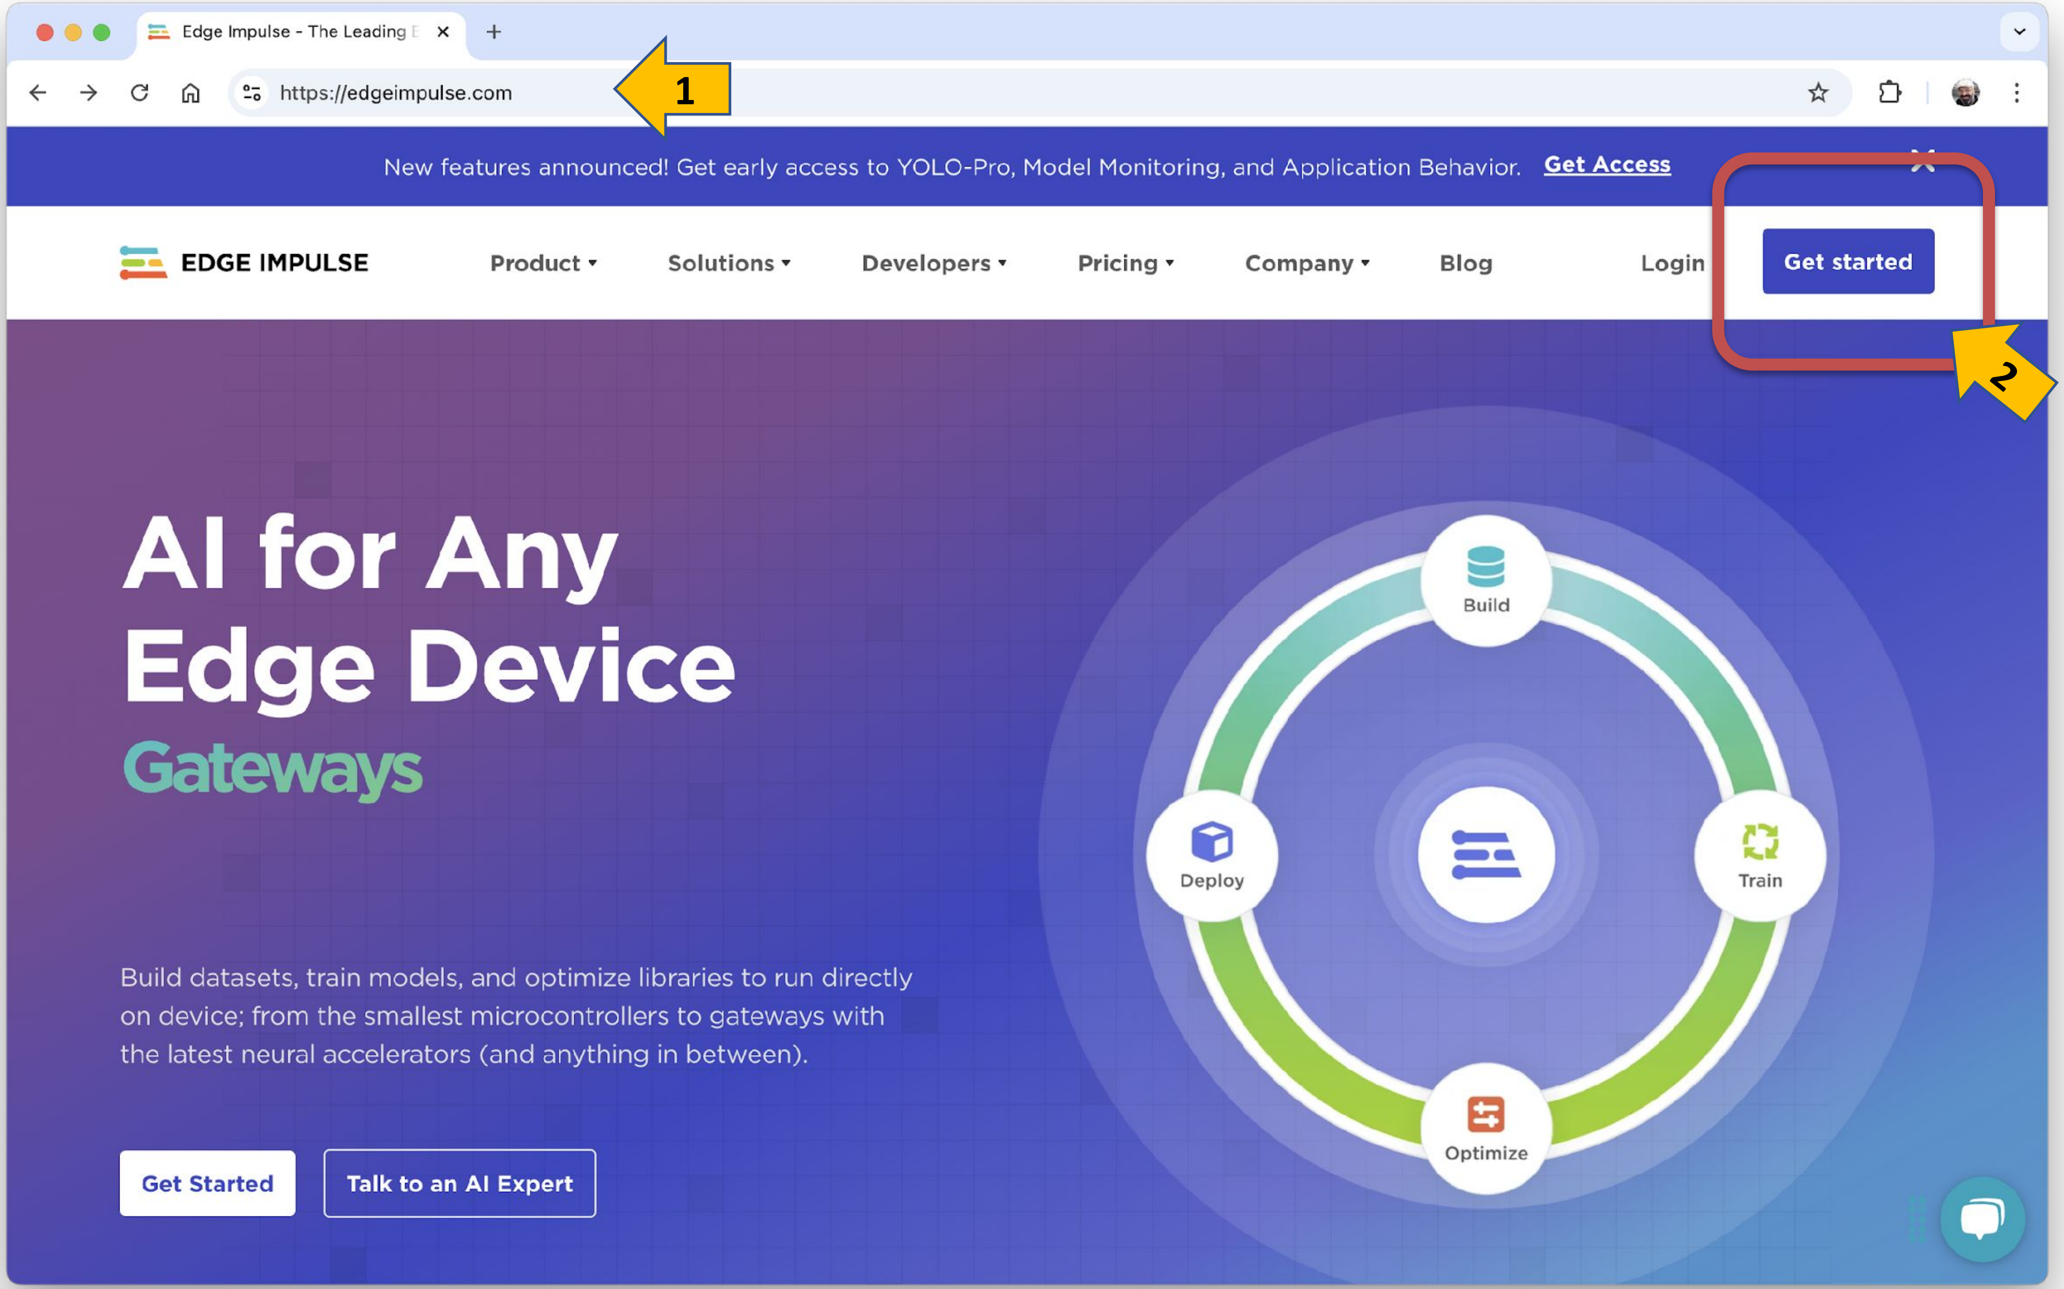
Edge Impulse - The Leading E
https://edgeimpulse.com
1
New features announced! Get early access to YOLO-Pro, Model Monitoring, and Application Behavior.
Get Access
EDGE IMPULSE
Product
Solutions
Developers
Pricing
Company
Blog
Login
Get started
2
AI for Any
Edge Device
Gateways
Build datasets, train models, and optimize libraries to run directly on device; from the smallest microcontrollers to gateways with the latest neural accelerators (and anything in between).
Get Started
Talk to an AI Expert
Build
Deploy
Train
Optimize

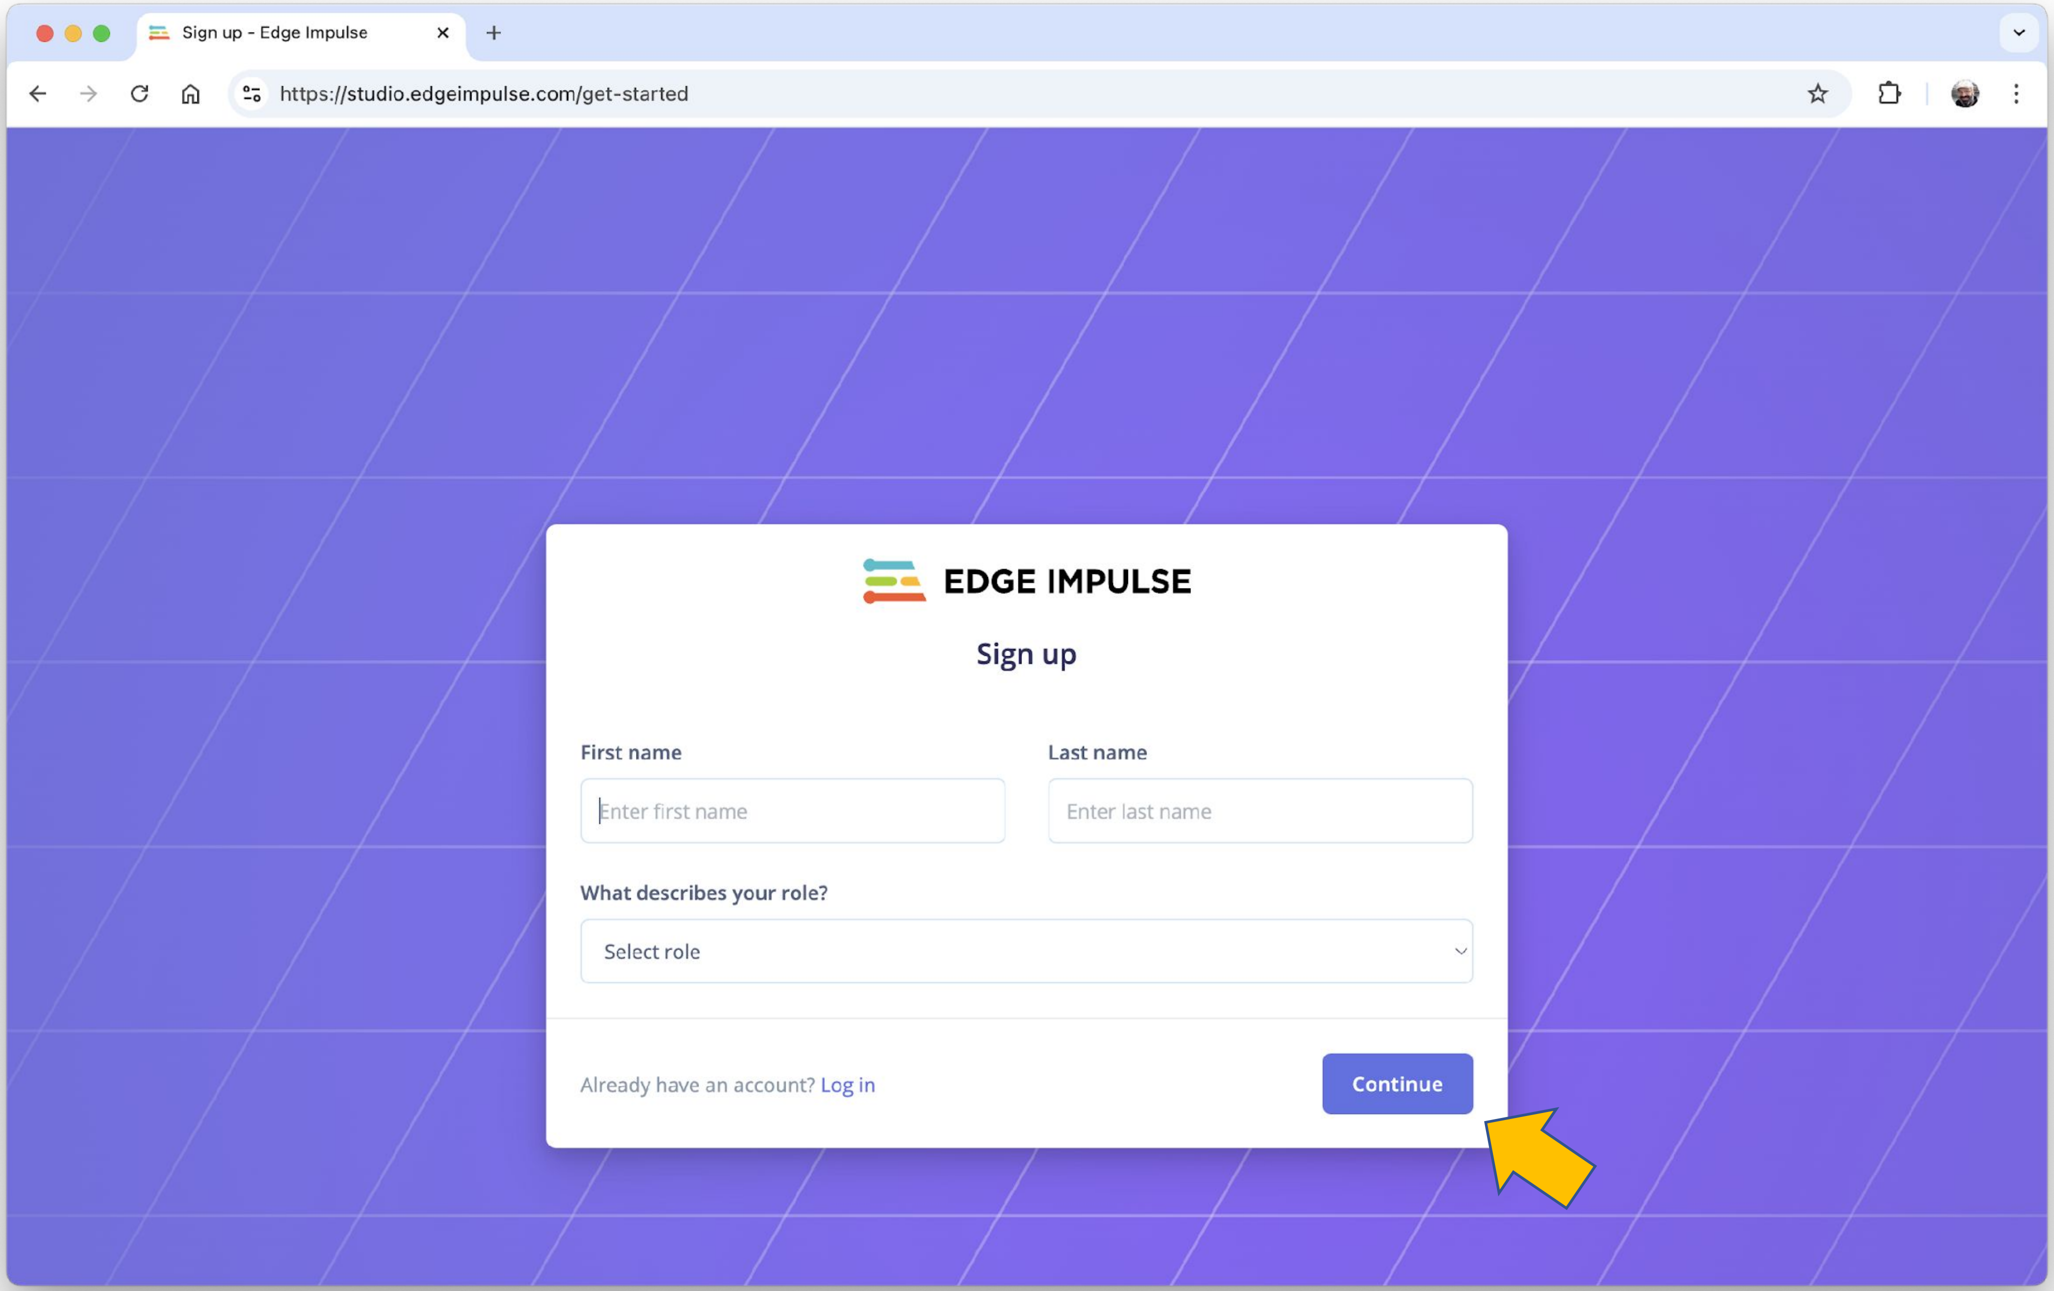
Sign up - Edge Impulse
https://studio.edgeimpulse.com/get-started
EDGE IMPULSE
Sign up
First name
Enter first name
Last name
Enter last name
What describes your role?
Select role
Already have an account? Log in
Continue

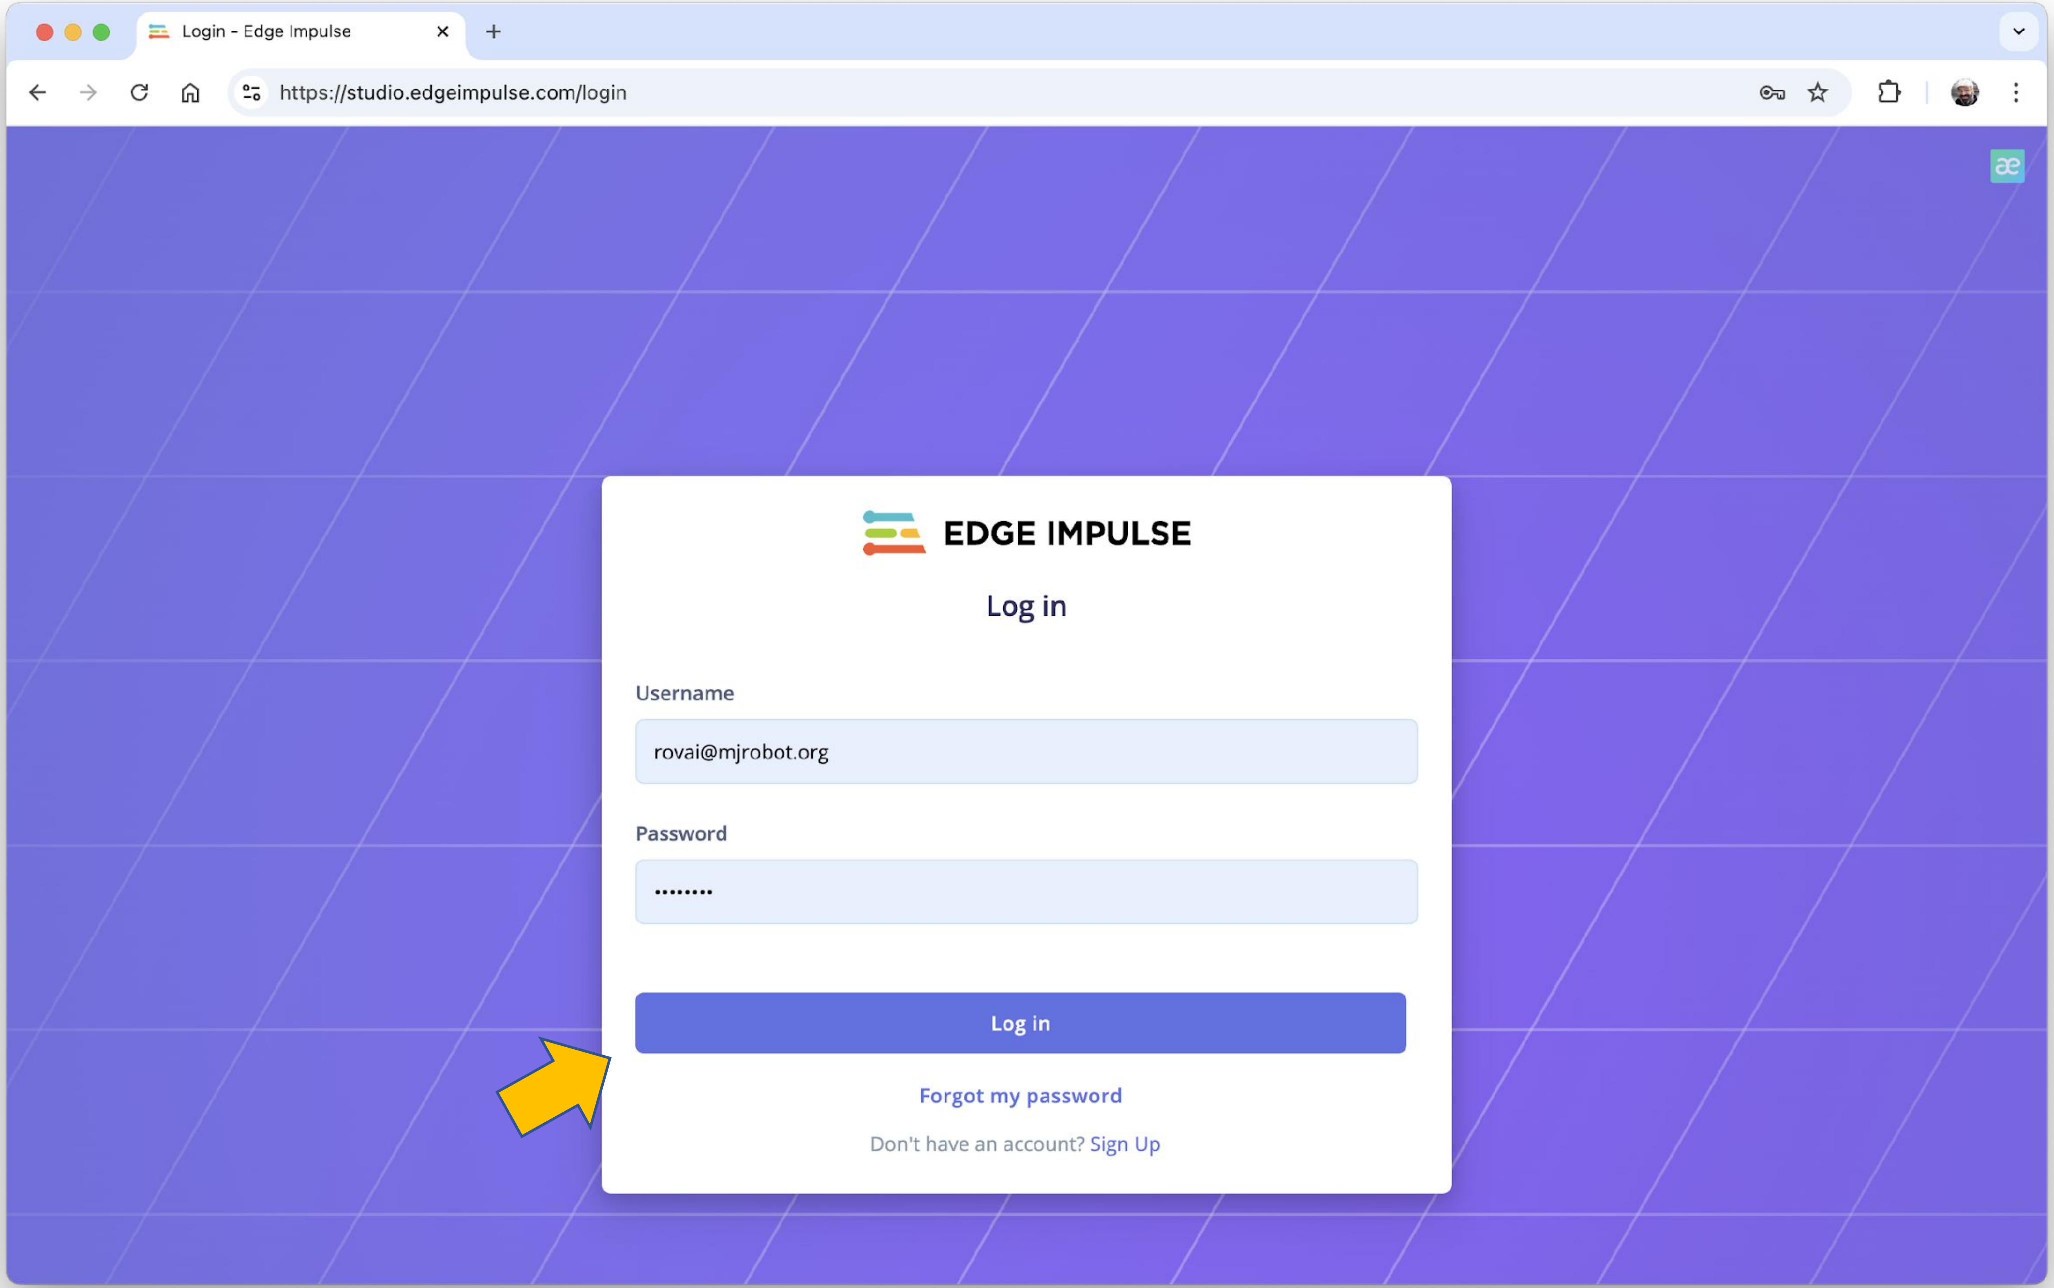
Login - Edge Impulse
https://studio.edgeimpulse.com/login
EDGE IMPULSE
Log in
Username
rovai@mjrobot.org
Password
Log in
Forgot my password
Don't have an account? Sign Up

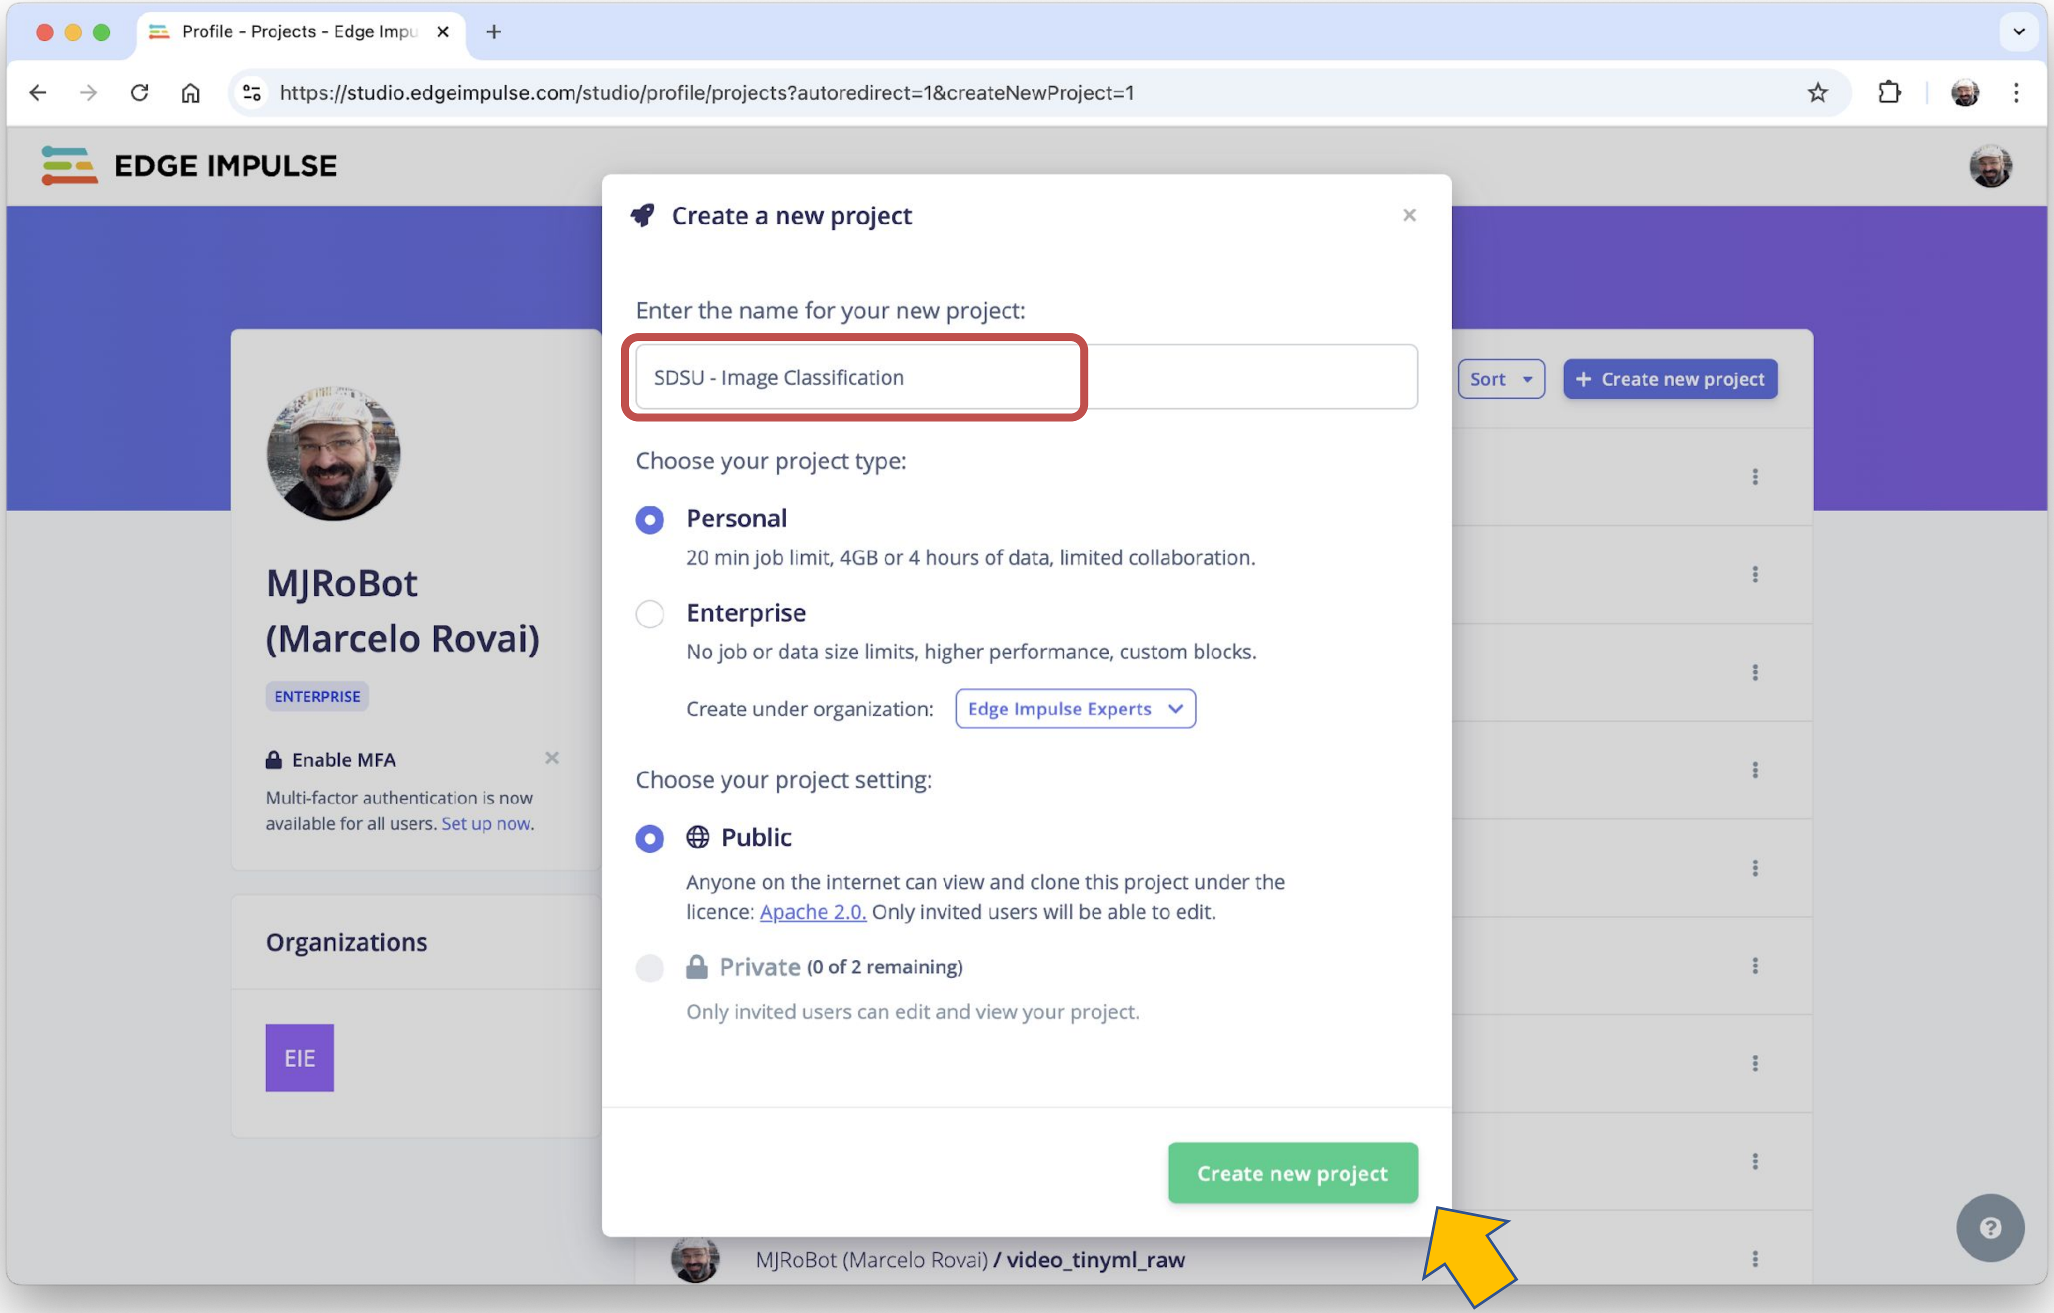
Profile - Projects - Edge Impu
https://studio.edgeimpulse.com/studio/profile/projects?autoredirect=1&createNewProject=1
EDGE IMPULSE
MJRoBot (Marcelo Rovai)
ENTERPRISE
Enable MFA
Multi-factor authentication is now available for all users. Set up now.
Organizations
EIE
Sort
+ Create new project
Create a new project
Enter the name for your new project:
SDSU - Image Classification
Choose your project type:
Personal
20 min job limit, 4GB or 4 hours of data, limited collaboration.
Enterprise
No job or data size limits, higher performance, custom blocks.
Create under organization:
Edge Impulse Experts
Choose your project setting:
Public
Anyone on the internet can view and clone this project under the licence: Apache 2.0. Only invited users will be able to edit.
Private (0 of 2 remaining)
Only invited users can edit and view your project.
Create new project
MJRoBot (Marcelo Rovai) / video_tinyml_raw

Medicine Classificator

- Classes:
  - prodR
  - prodD
  - prodT
  - backG

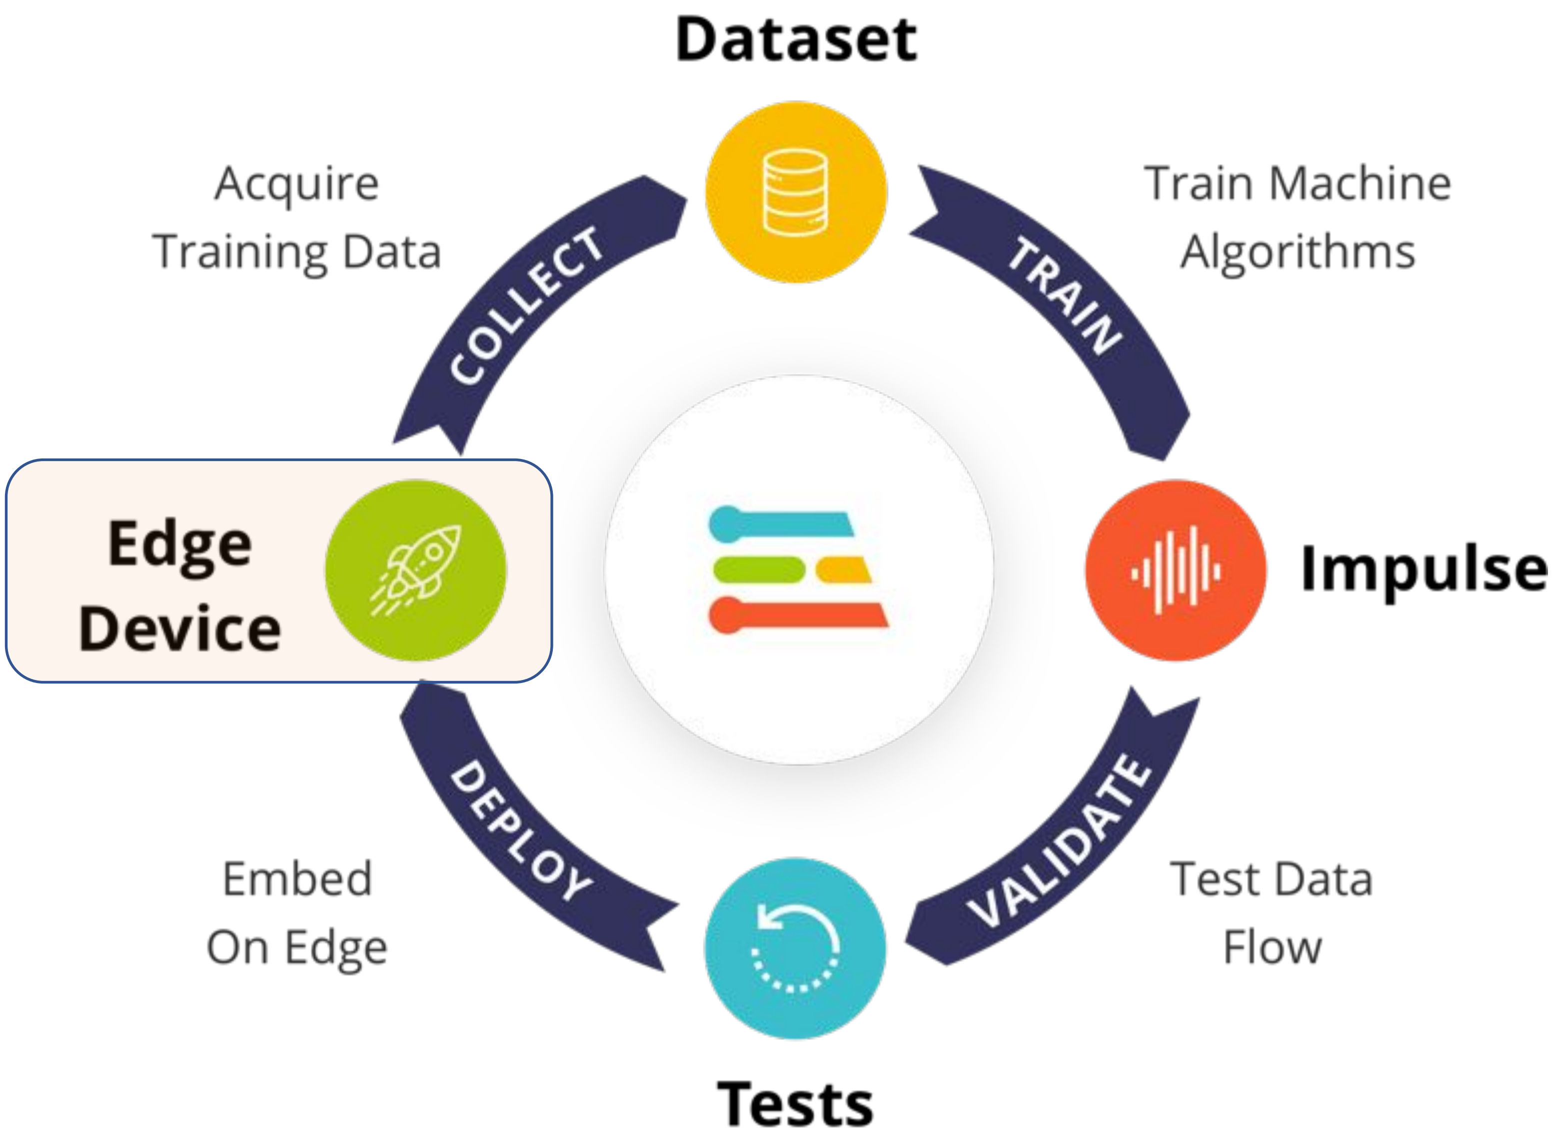

- Pre-Processing Data
- Design a Model
- Train a Model

Edge Device & Sensors

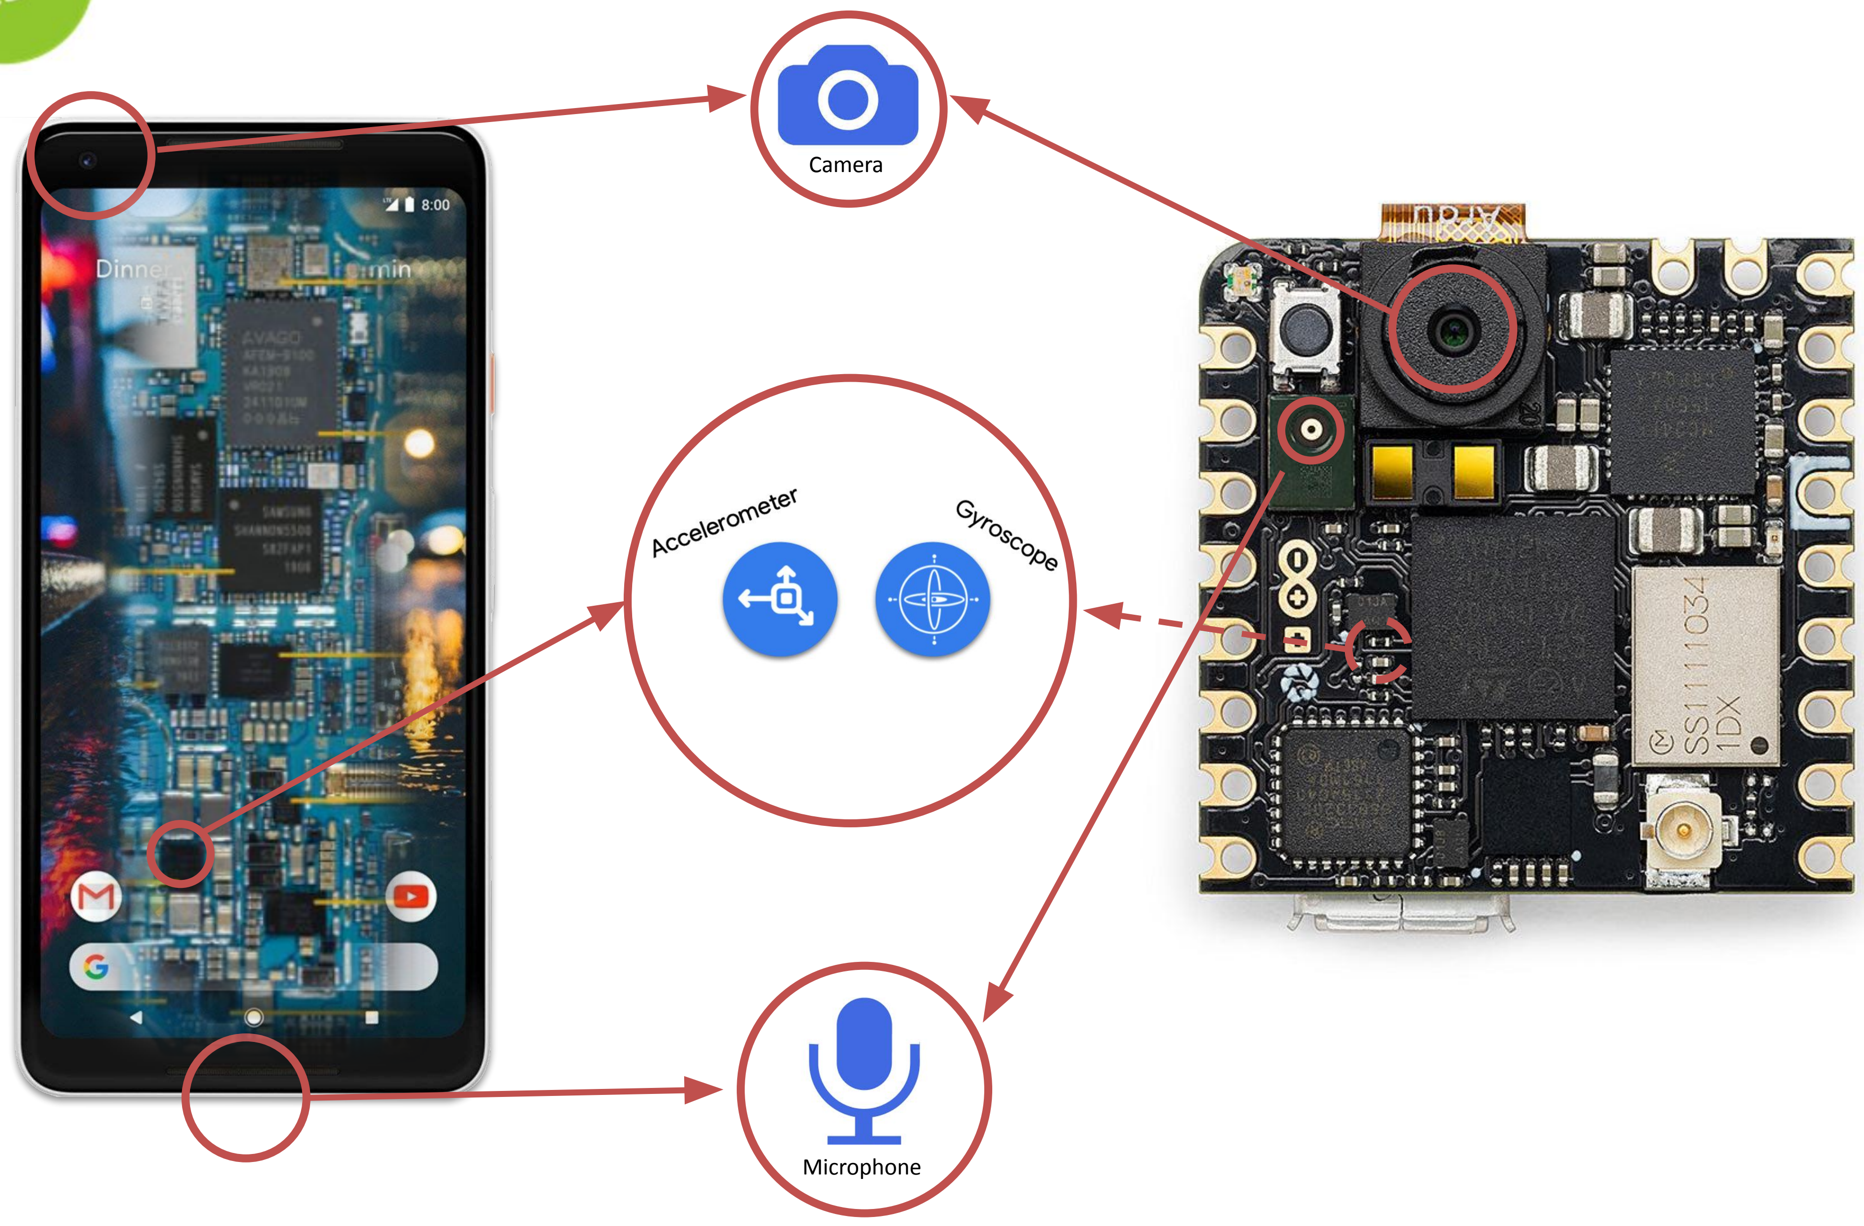

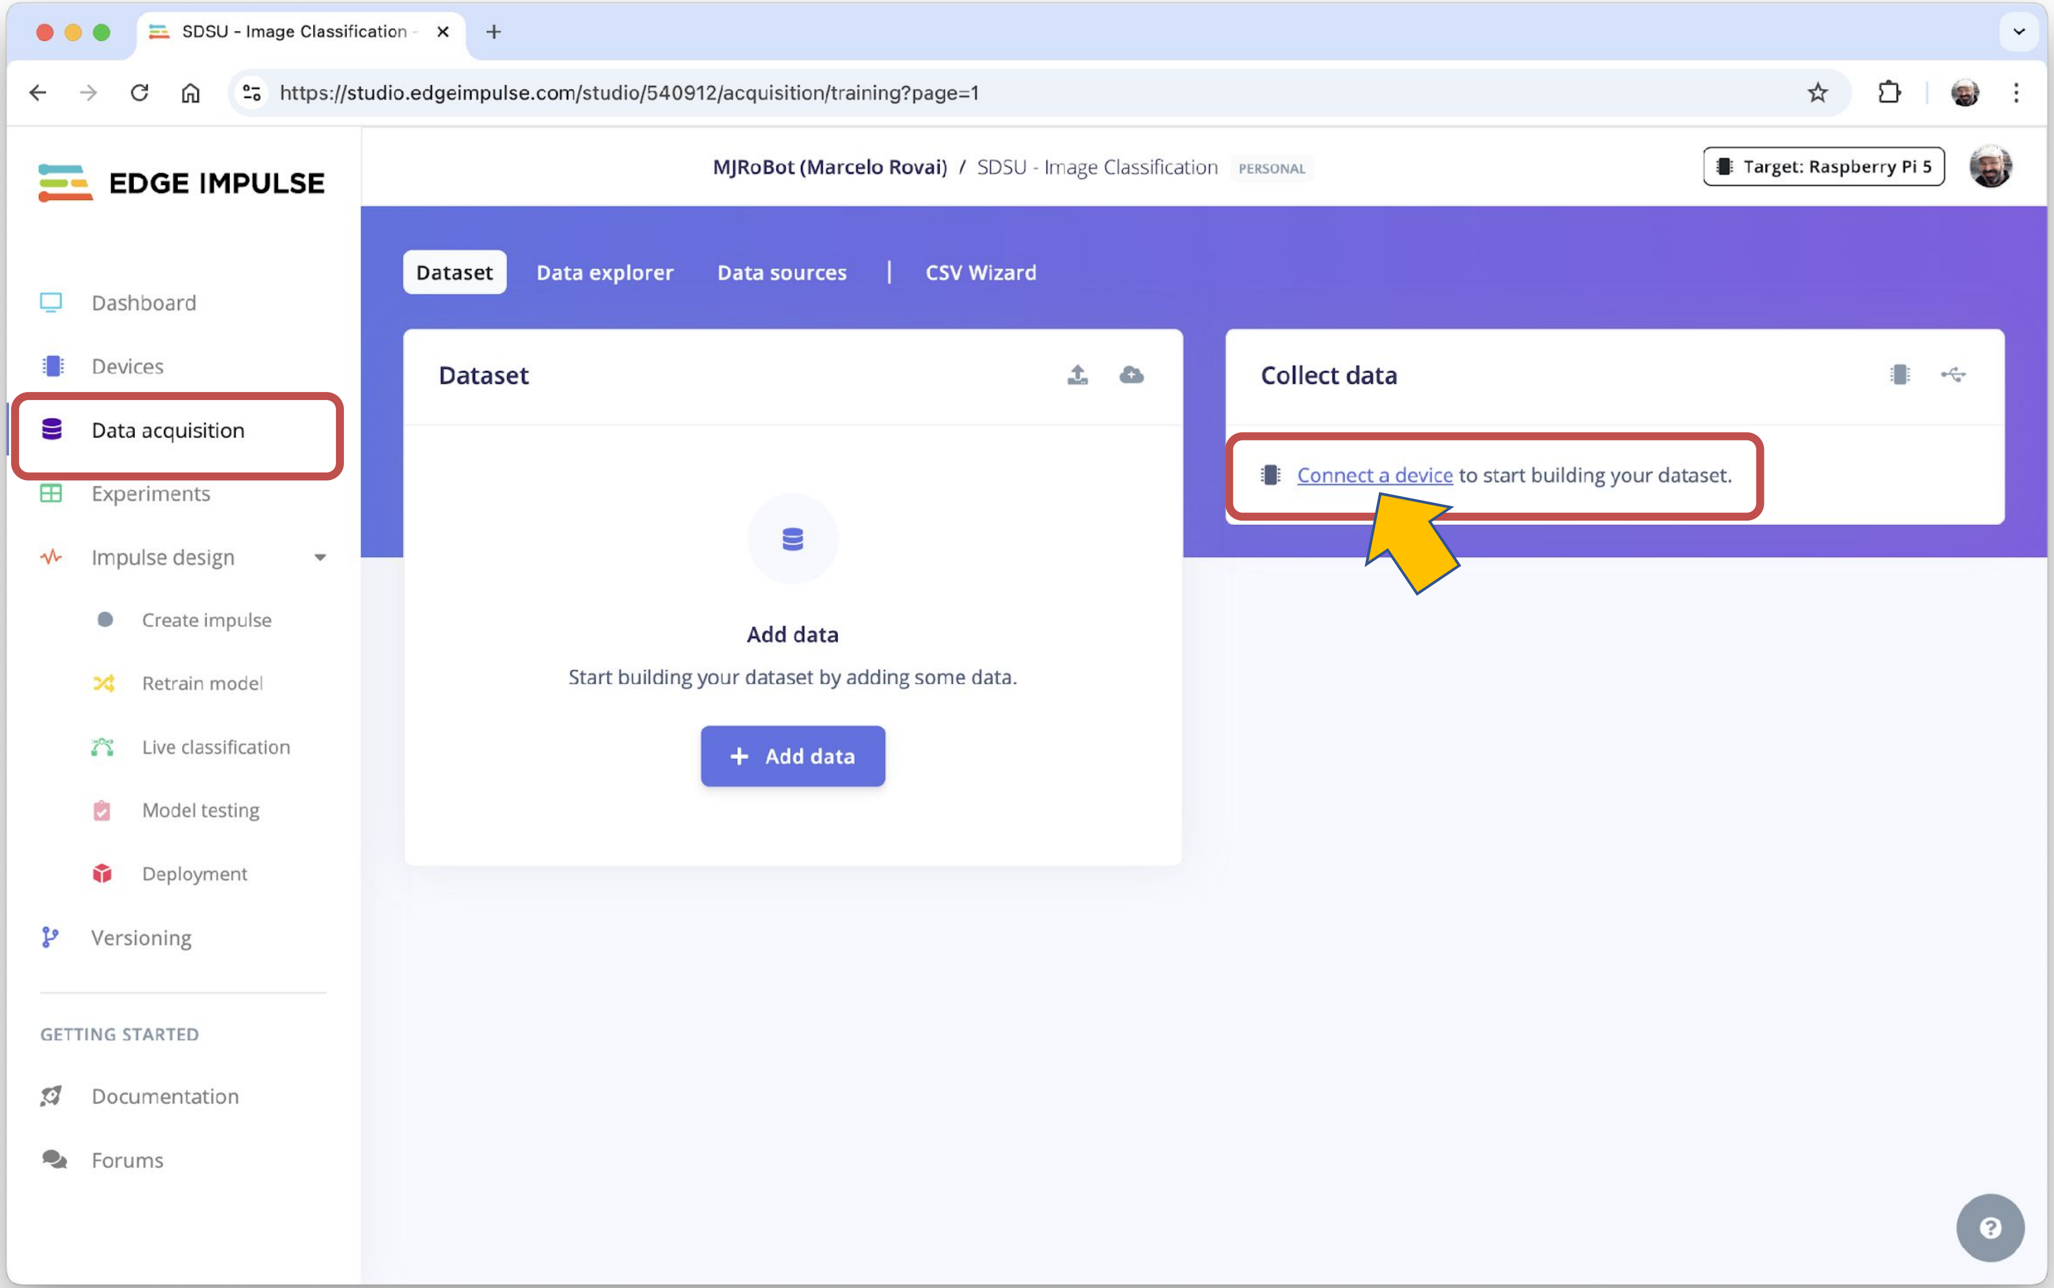
SDSU - Image Classification
https://studio.edgeimpulse.com/studio/540912/acquisition/training?page=1
EDGE IMPULSE
MJRoBot (Marcelo Rovai) / SDSU - Image Classification PERSONAL
Target: Raspberry Pi 5
Dashboard
Devices
Data acquisition
Experiments
Impulse design
Create impulse
Retrain model
Live classification
Model testing
Deployment
Versioning
GETTING STARTED
Documentation
Forums
Dataset
Data explorer
Data sources
CSV Wizard
Dataset
Add data
Start building your dataset by adding some data.
+ Add data
Collect data
Connect a device to start building your dataset.

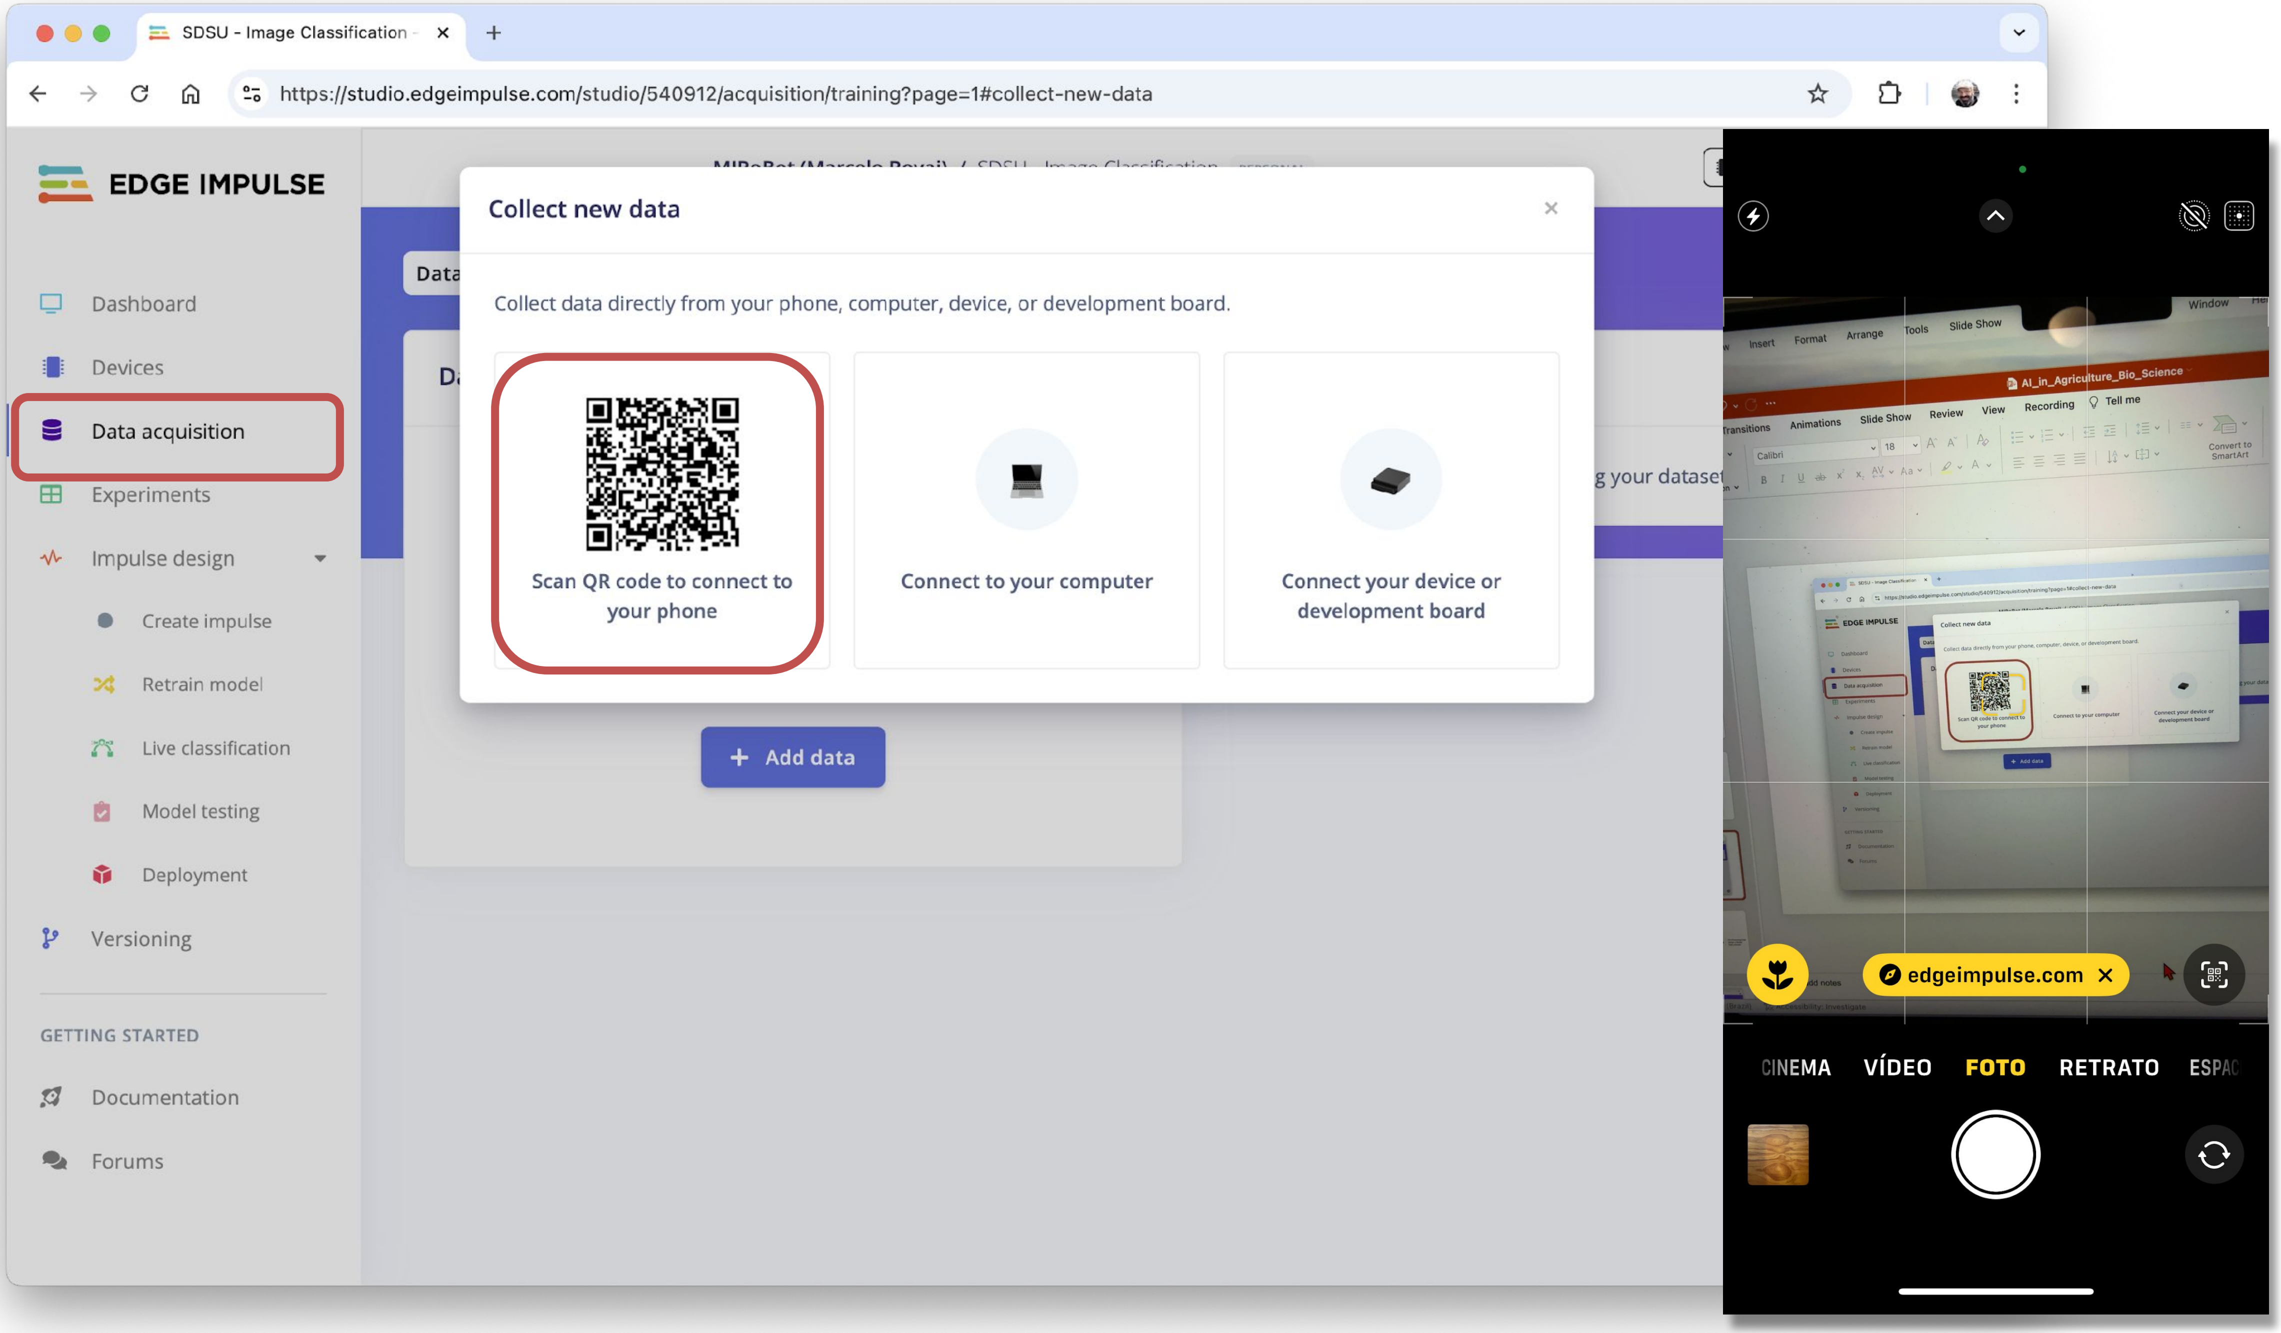

SDSU - Image Classification
https://studio.edgeimpulse.com/studio/540912/acquisition/training?page=1#collect-new-data
EDGE IMPULSE
Dashboard
Devices
Data acquisition
Experiments
Impulse design
Create impulse
Retrain model
Live classification
Model testing
Deployment
Versioning
GETTING STARTED
Documentation
Forums
Collect new data
Collect data directly from your phone, computer, device, or development board.
Scan QR code to connect to your phone
Connect to your computer
Connect your device or development board
Add data
edgeimpulse.com
CINEMA
VÍDEO
FOTO
RETRATO

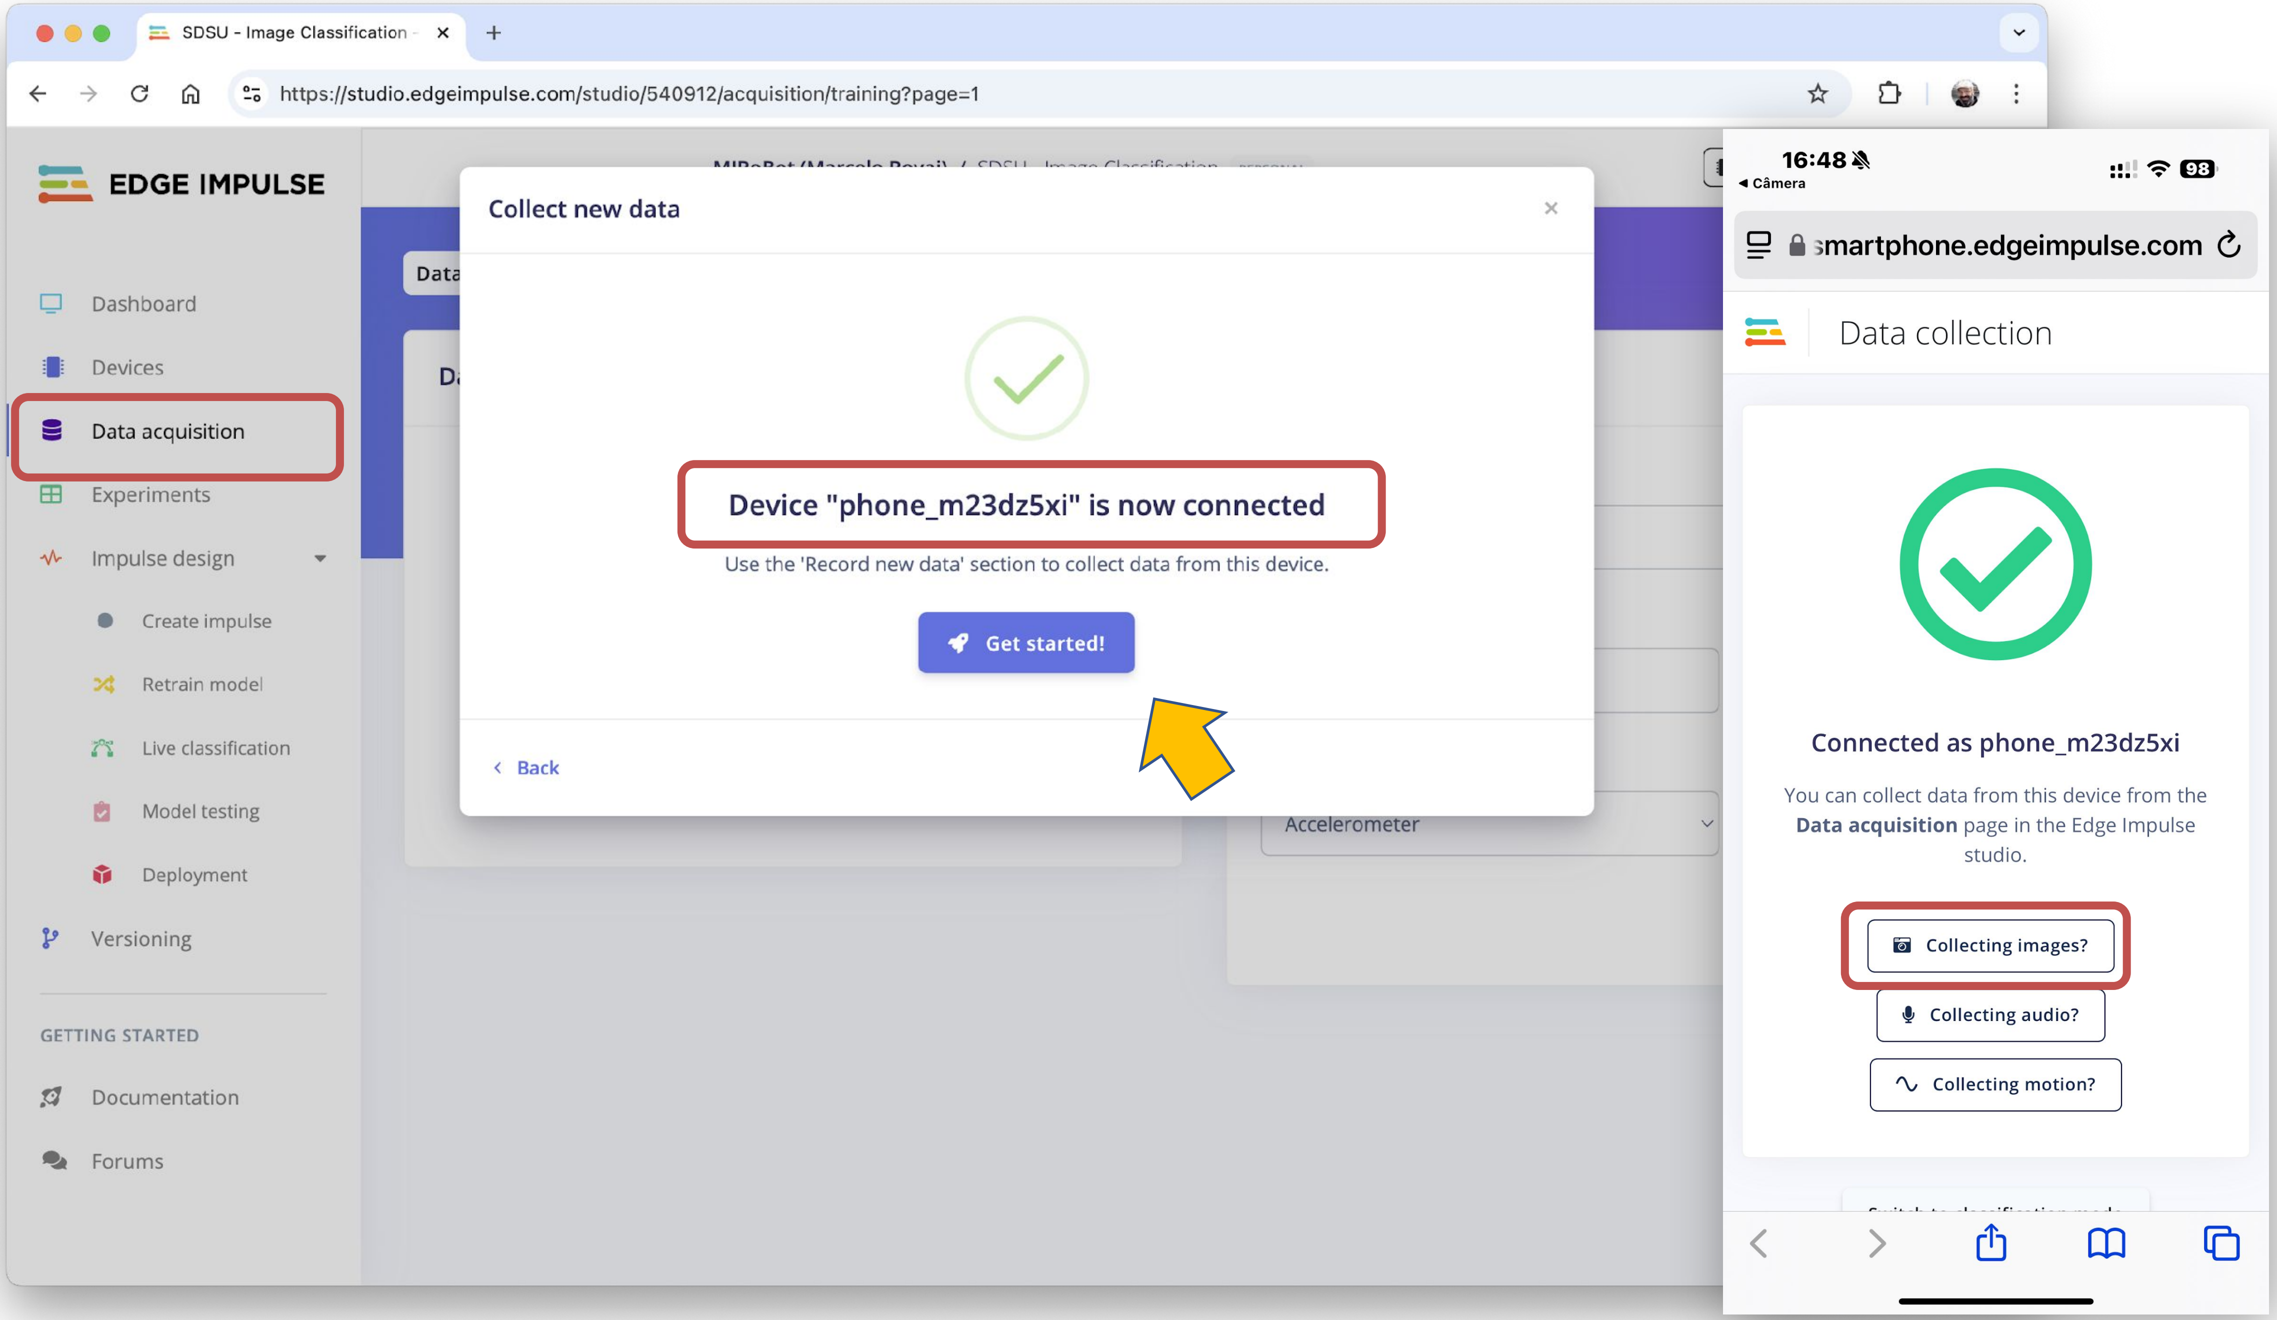
SDSU - Image Classification
https://studio.edgeimpulse.com/studio/540912/acquisition/training?page=1
EDGE IMPULSE
Dashboard
Devices
Data acquisition
Experiments
Impulse design
Create impulse
Retrain model
Live classification
Model testing
Deployment
Versioning
GETTING STARTED
Documentation
Forums
Collect new data
Device "phone_m23dz5xi" is now connected
Use the 'Record new data' section to collect data from this device.
Get started!
Back
Accelerometer
16:48
Câmera
smartphone.edgeimpulse.com
Data collection
Connected as phone_m23dz5xi
You can collect data from this device from the Data acquisition page in the Edge Impulse studio.
Collecting images?
Collecting audio?
Collecting motion?

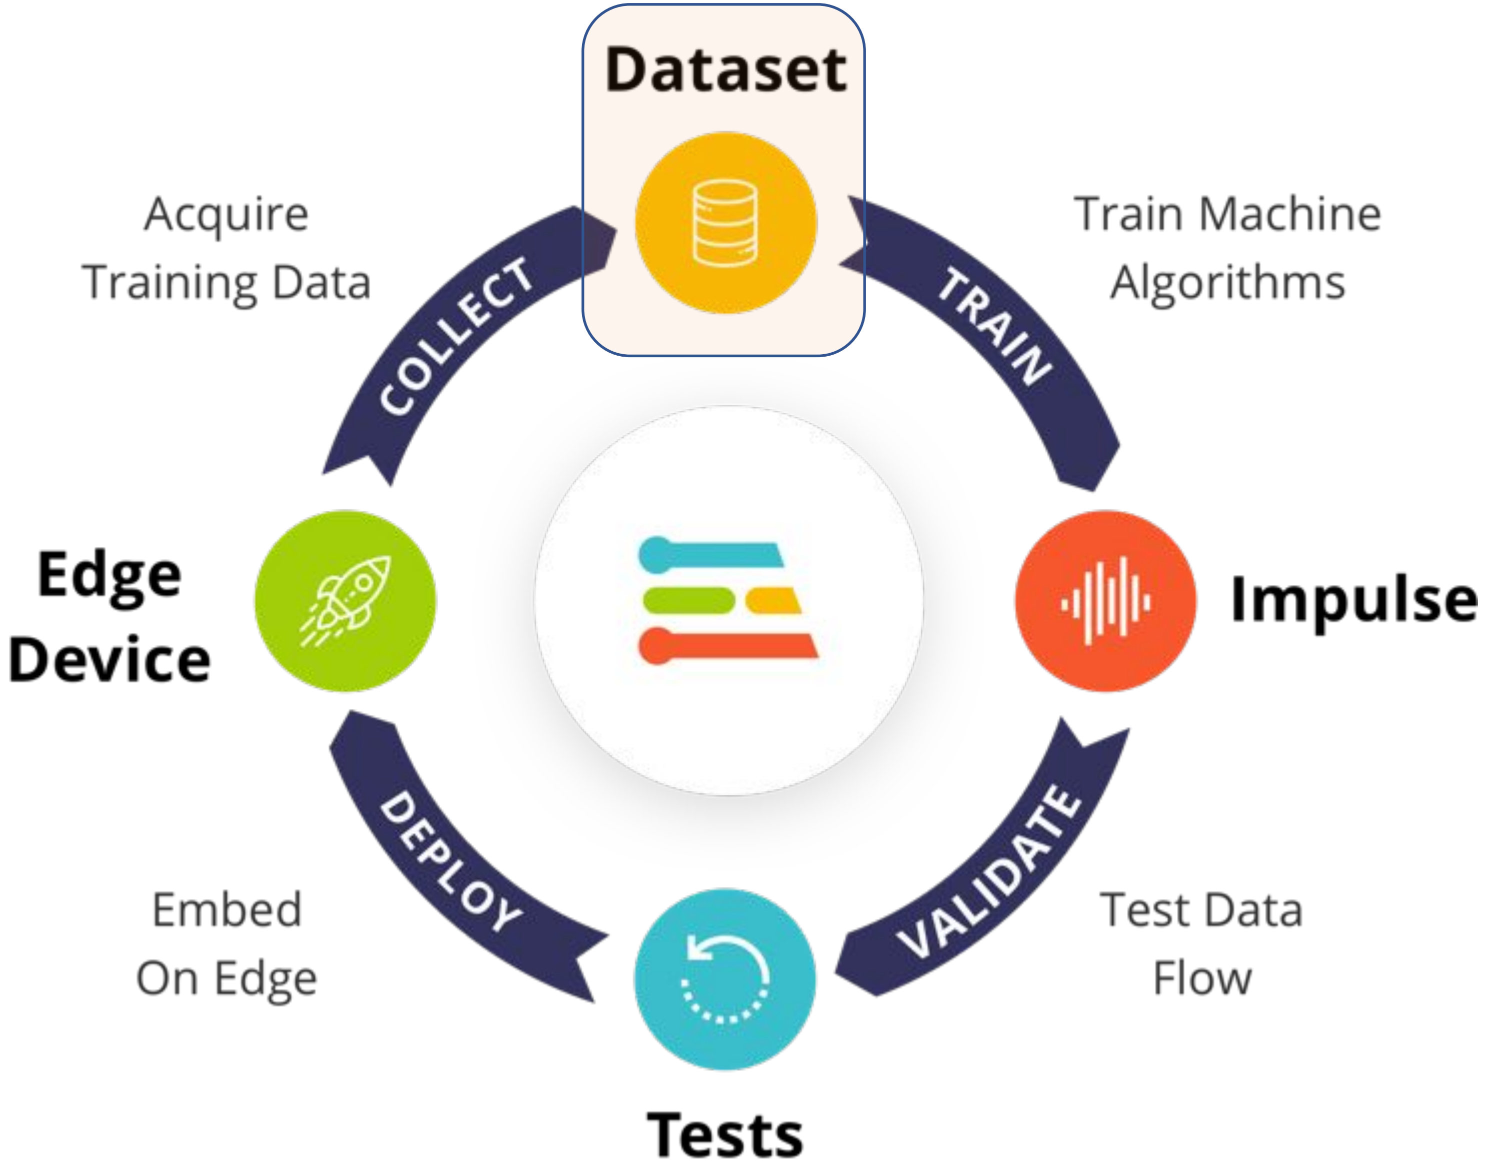

- Pre-Processing Data
- Design a Model
- Train a Model

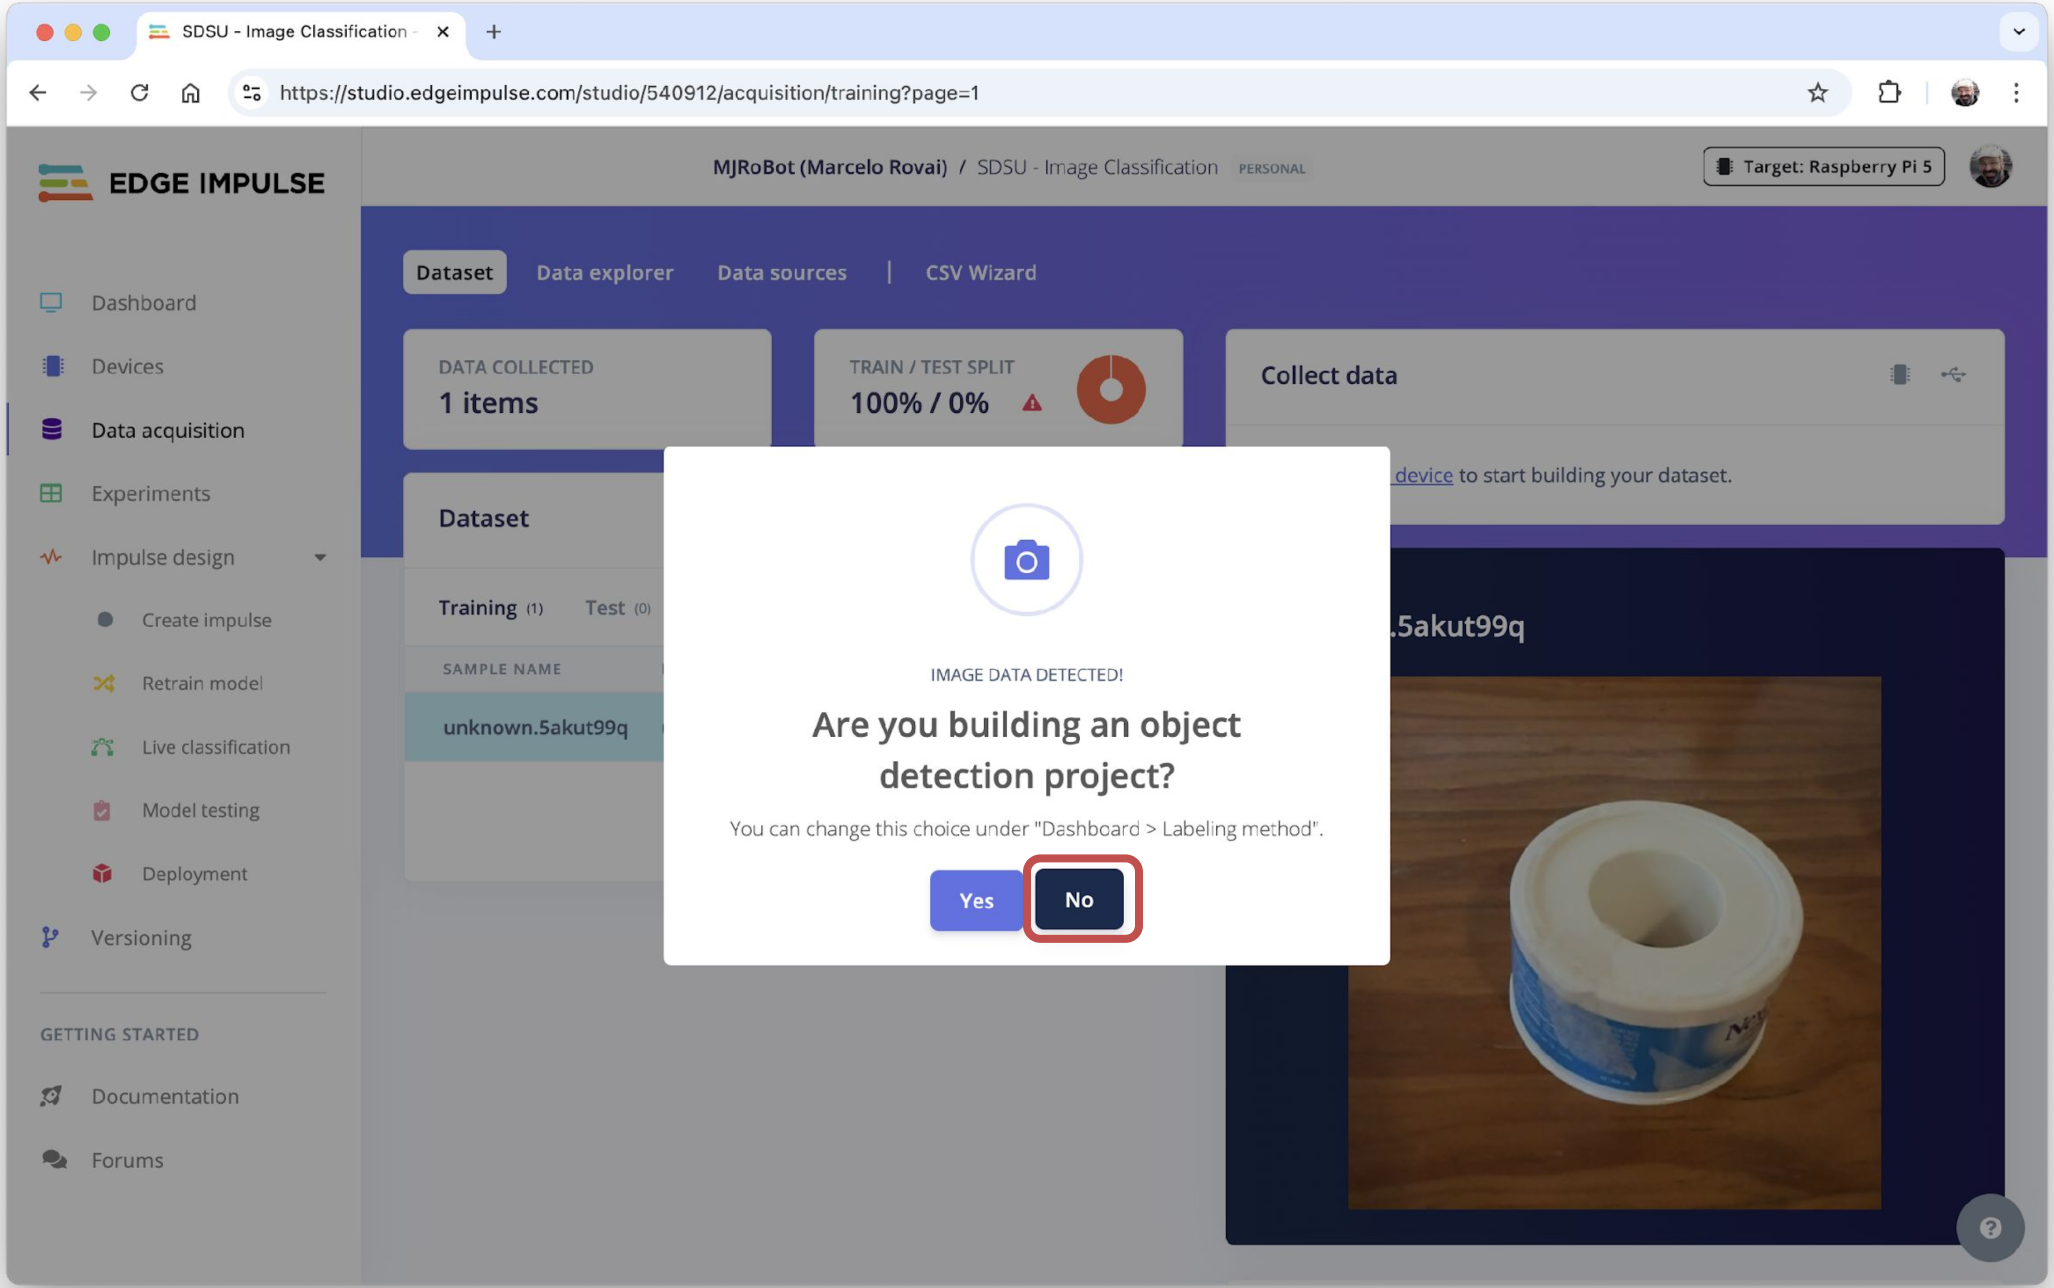
SDSU - Image Classification
https://studio.edgeimpulse.com/studio/540912/acquisition/training?page=1
EDGE IMPULSE
MJRoBot (Marcelo Rovai) / SDSU - Image Classification PERSONAL
Target: Raspberry Pi 5
Dashboard
Devices
Data acquisition
Experiments
Impulse design
Create impulse
Retrain model
Live classification
Model testing
Deployment
Versioning
GETTING STARTED
Documentation
Forums
Dataset
Data explorer
Data sources
CSV Wizard
DATA COLLECTED
1 items
TRAIN / TEST SPLIT
100% / 0%
Collect data
device to start building your dataset.
Dataset
Training (1)
Test (0)
SAMPLE NAME
unknown.5akut99q
.5akut99q
IMAGE DATA DETECTED!
Are you building an object detection project?
You can change this choice under "Dashboard > Labeling method".
Yes
No

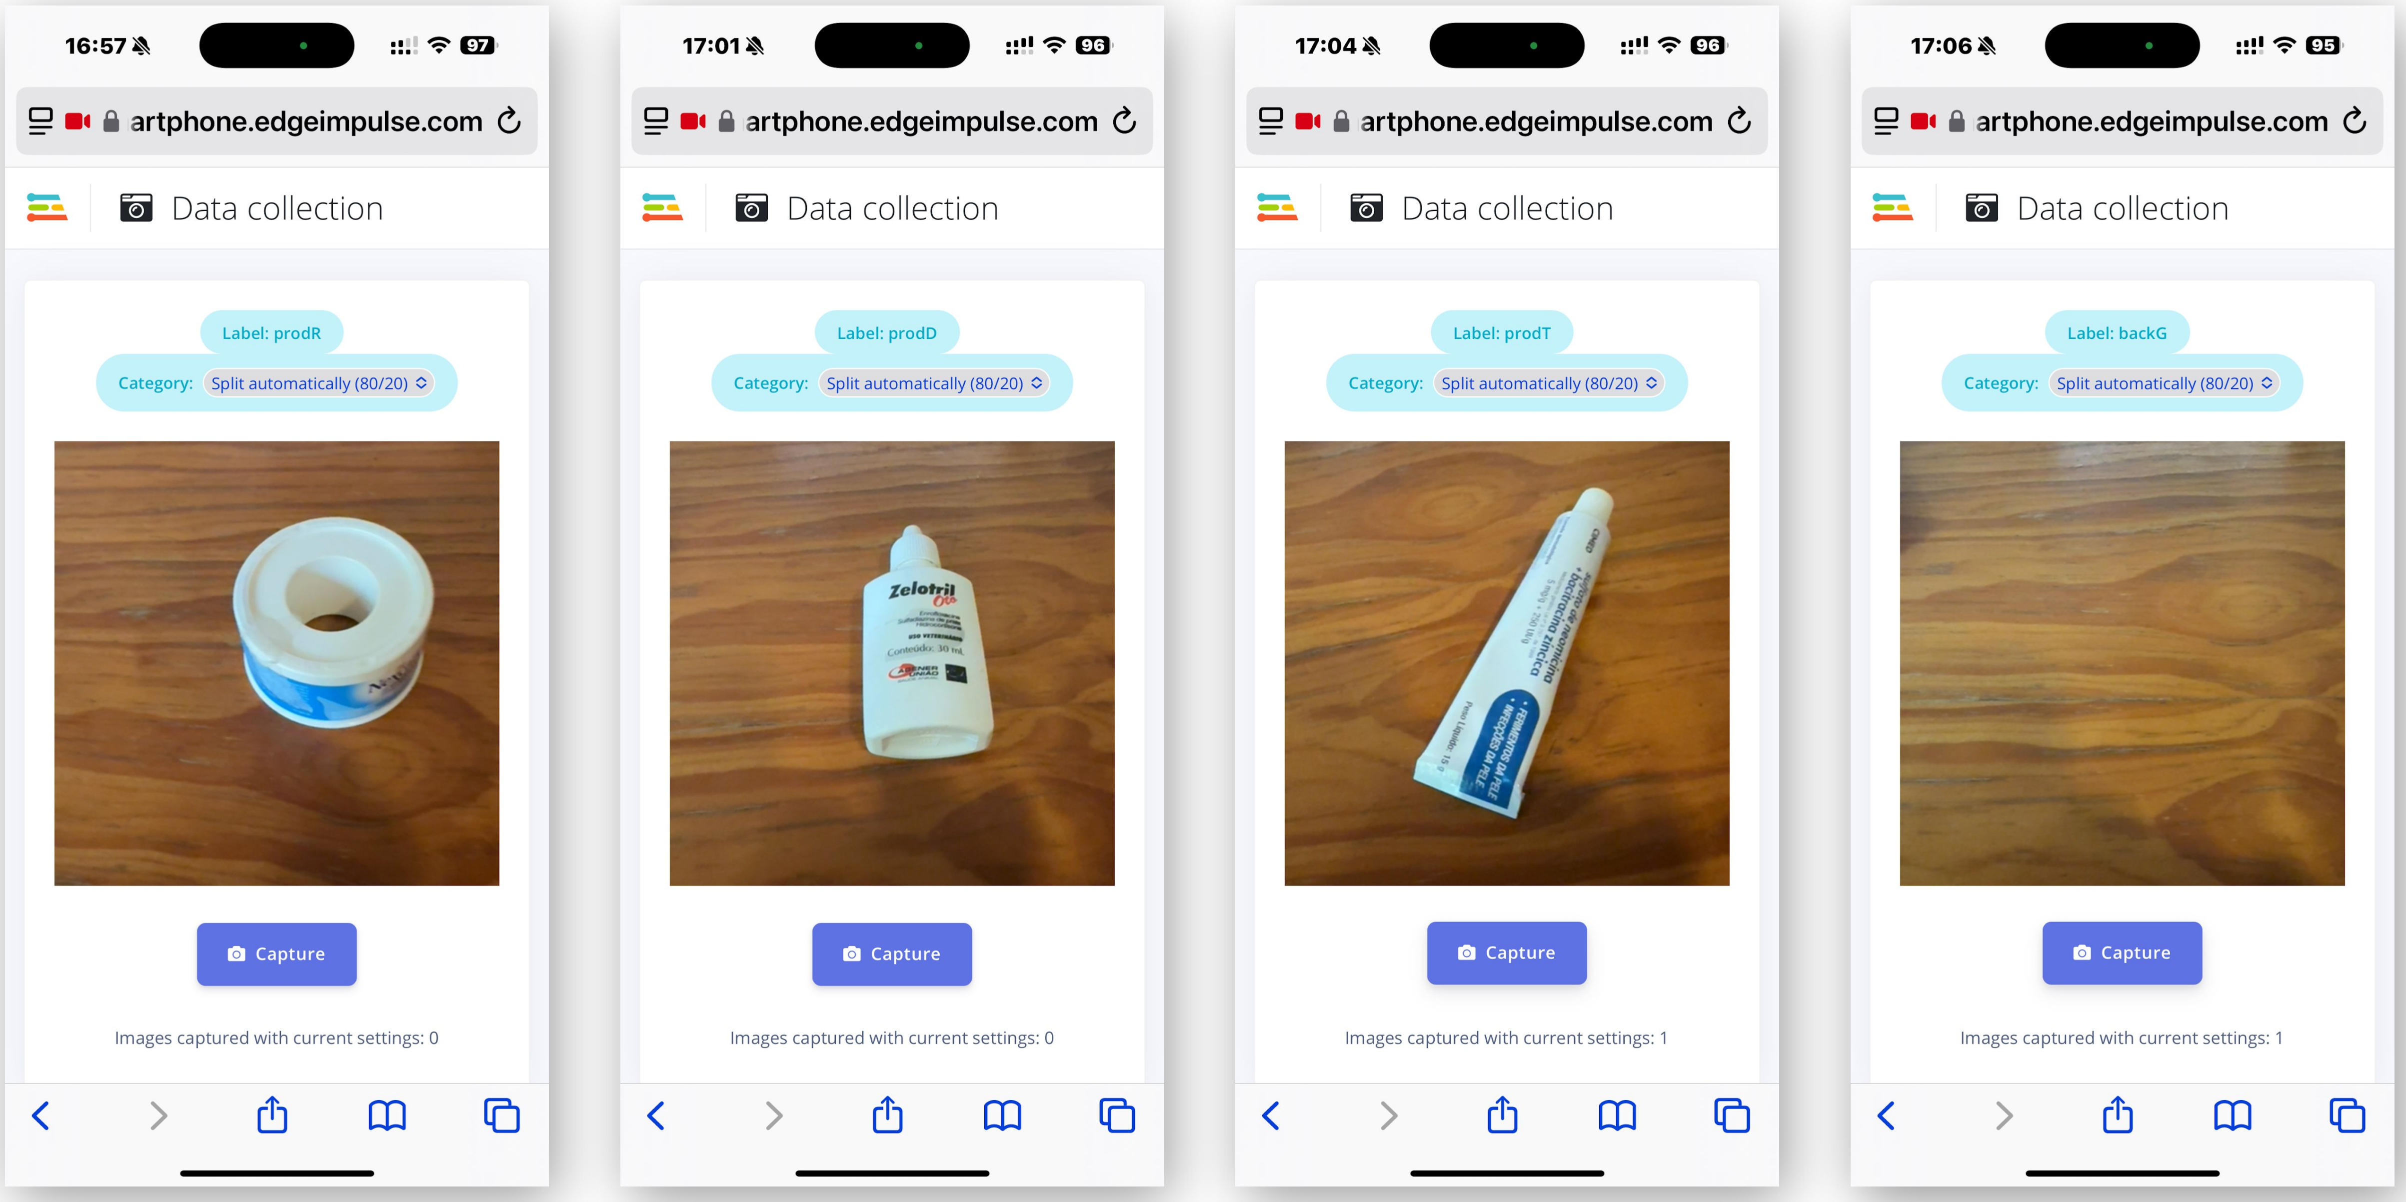
16:57
97
artphone.edgeimpulse.com
Data collection
Label: prodR
Category: Split automatically (80/20)
Capture
Images captured with current settings: 0
17:01
96
artphone.edgeimpulse.com
Data collection
Label: prodD
Category: Split automatically (80/20)
Zelotril
Capture
Images captured with current settings: 0
17:04
96
artphone.edgeimpulse.com
Data collection
Label: prodT
Category: Split automatically (80/20)
Capture
Images captured with current settings: 1
17:06
95
artphone.edgeimpulse.com
Data collection
Label: backG
Category: Split automatically (80/20)
Capture
Images captured with current settings: 1

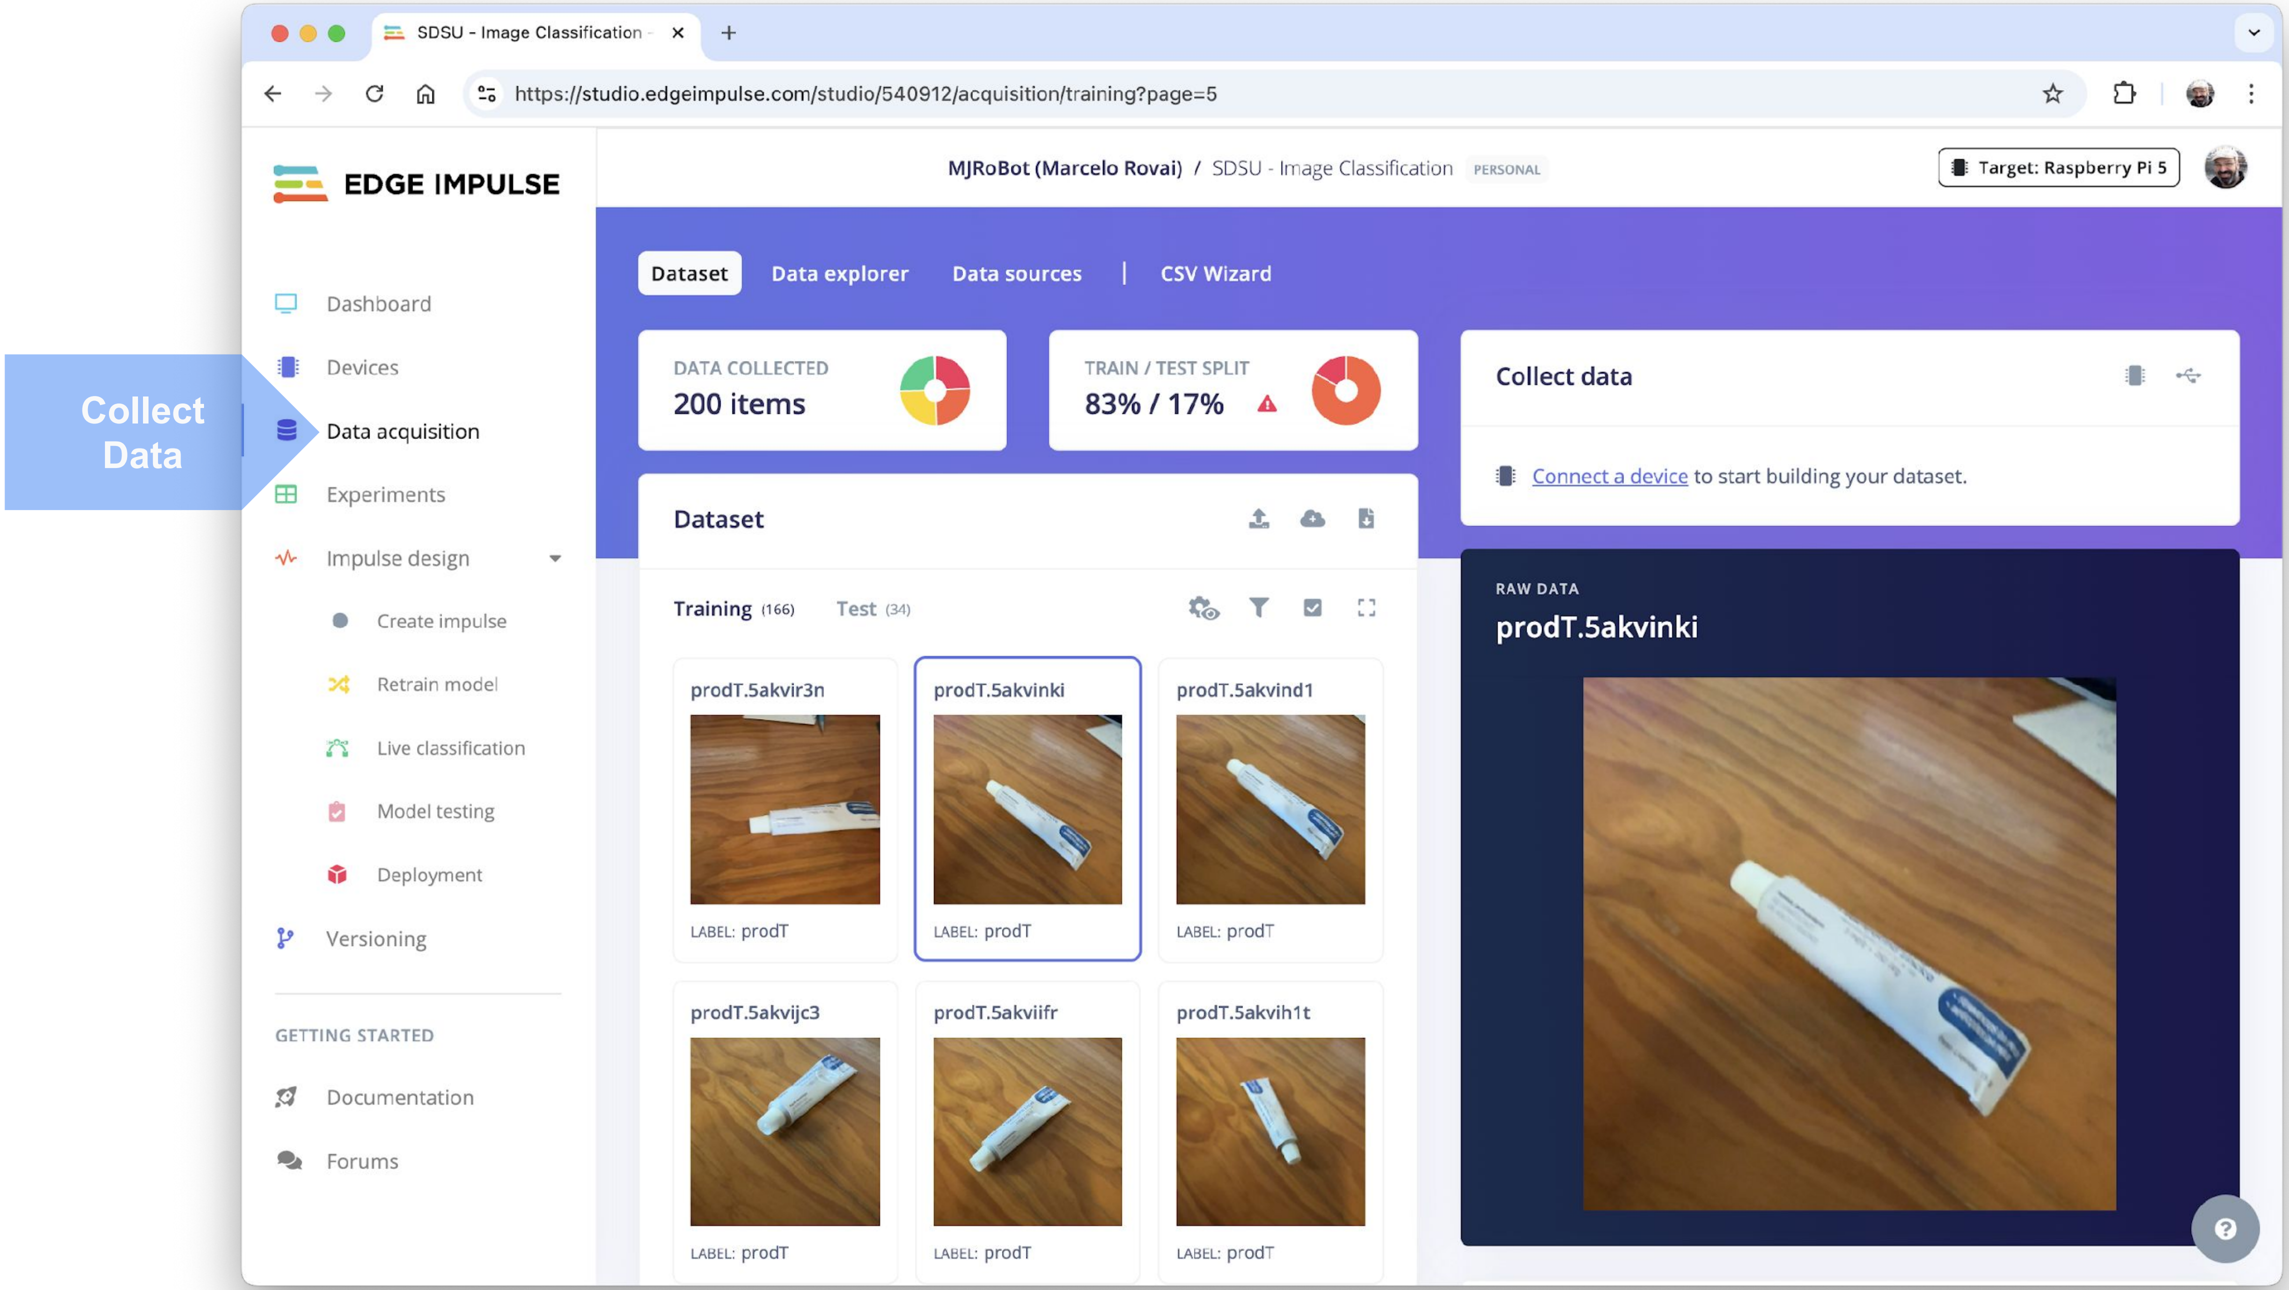
Collect Data
SDSU - Image Classification
https://studio.edgeimpulse.com/studio/540912/acquisition/training?page=5
EDGE IMPULSE
MJRoBot (Marcelo Rovai) / SDSU - Image Classification PERSONAL
Target: Raspberry Pi 5
Dataset
Data explorer
Data sources
CSV Wizard
Dashboard
Devices
Data acquisition
Experiments
Impulse design
Create impulse
Retrain model
Live classification
Model testing
Deployment
Versioning
GETTING STARTED
Documentation
Forums
DATA COLLECTED
200 items
TRAIN / TEST SPLIT
83% / 17%
Collect data
Connect a device to start building your dataset.
Dataset
Training (166)
Test (34)
prodT.5akvir3n
LABEL: prodT
prodT.5akvinki
LABEL: prodT
prodT.5akvind1
LABEL: prodT
prodT.5akvijc3
LABEL: prodT
prodT.5akviifr
LABEL: prodT
prodT.5akvih1t
LABEL: prodT
RAW DATA
prodT.5akvinki

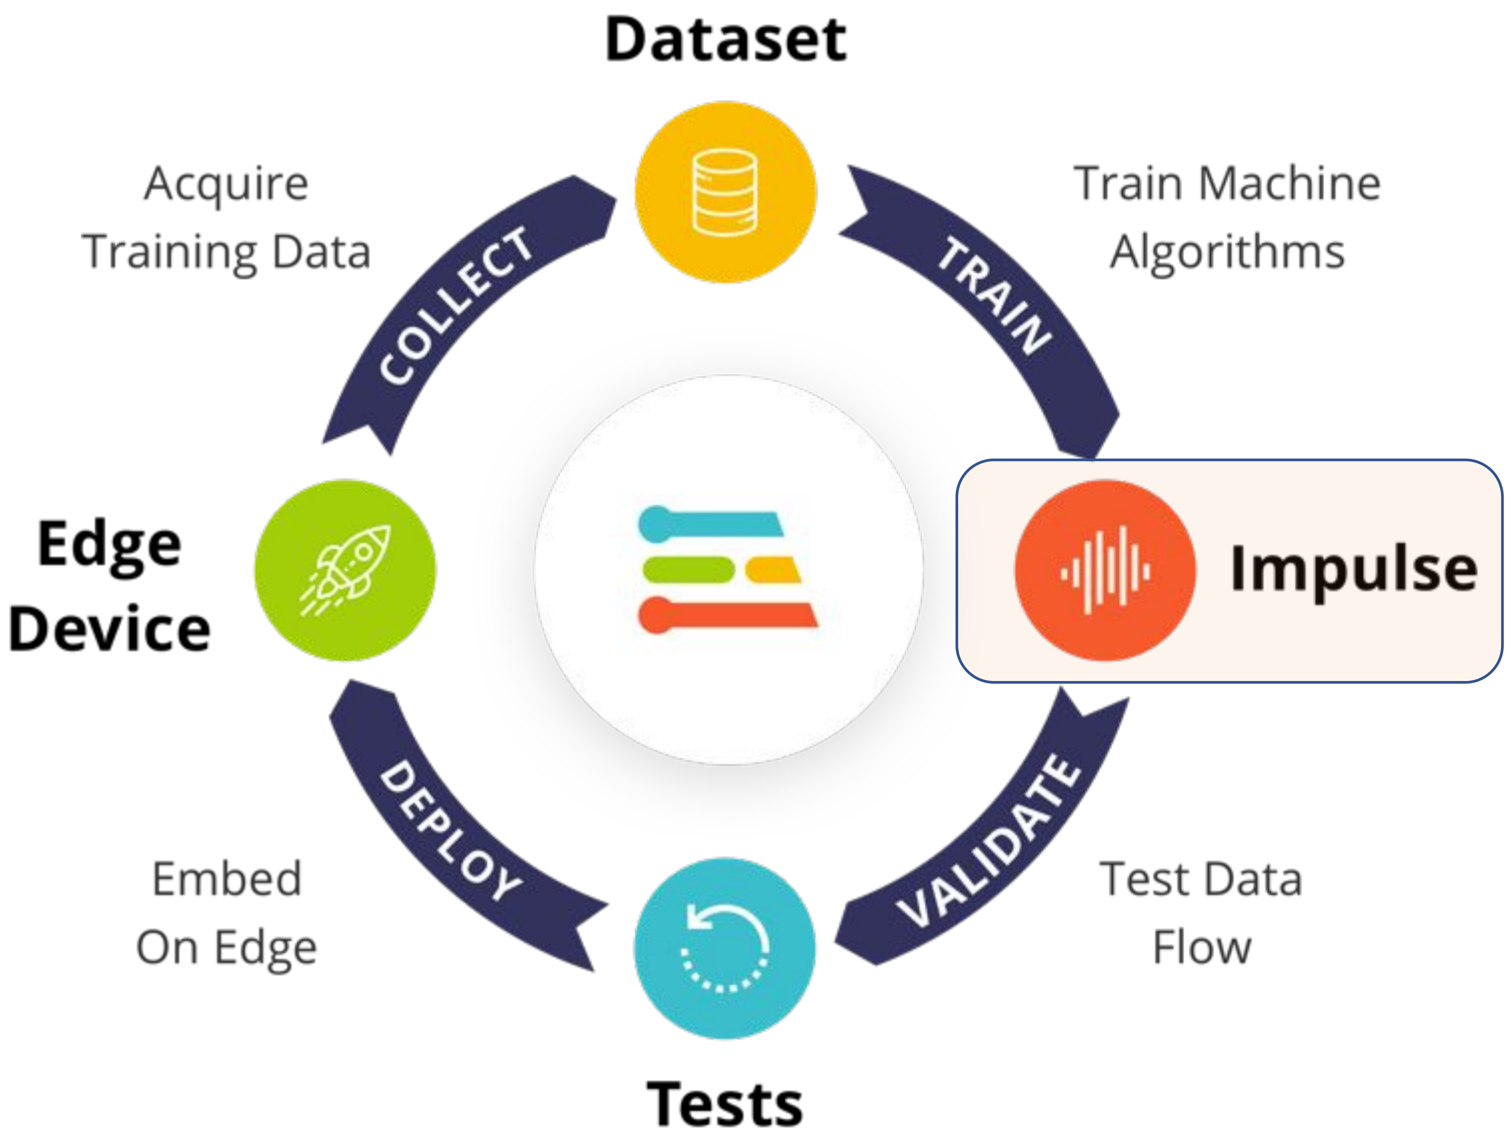

- Pre-Processing Data
- Design a Model
- Train a Model

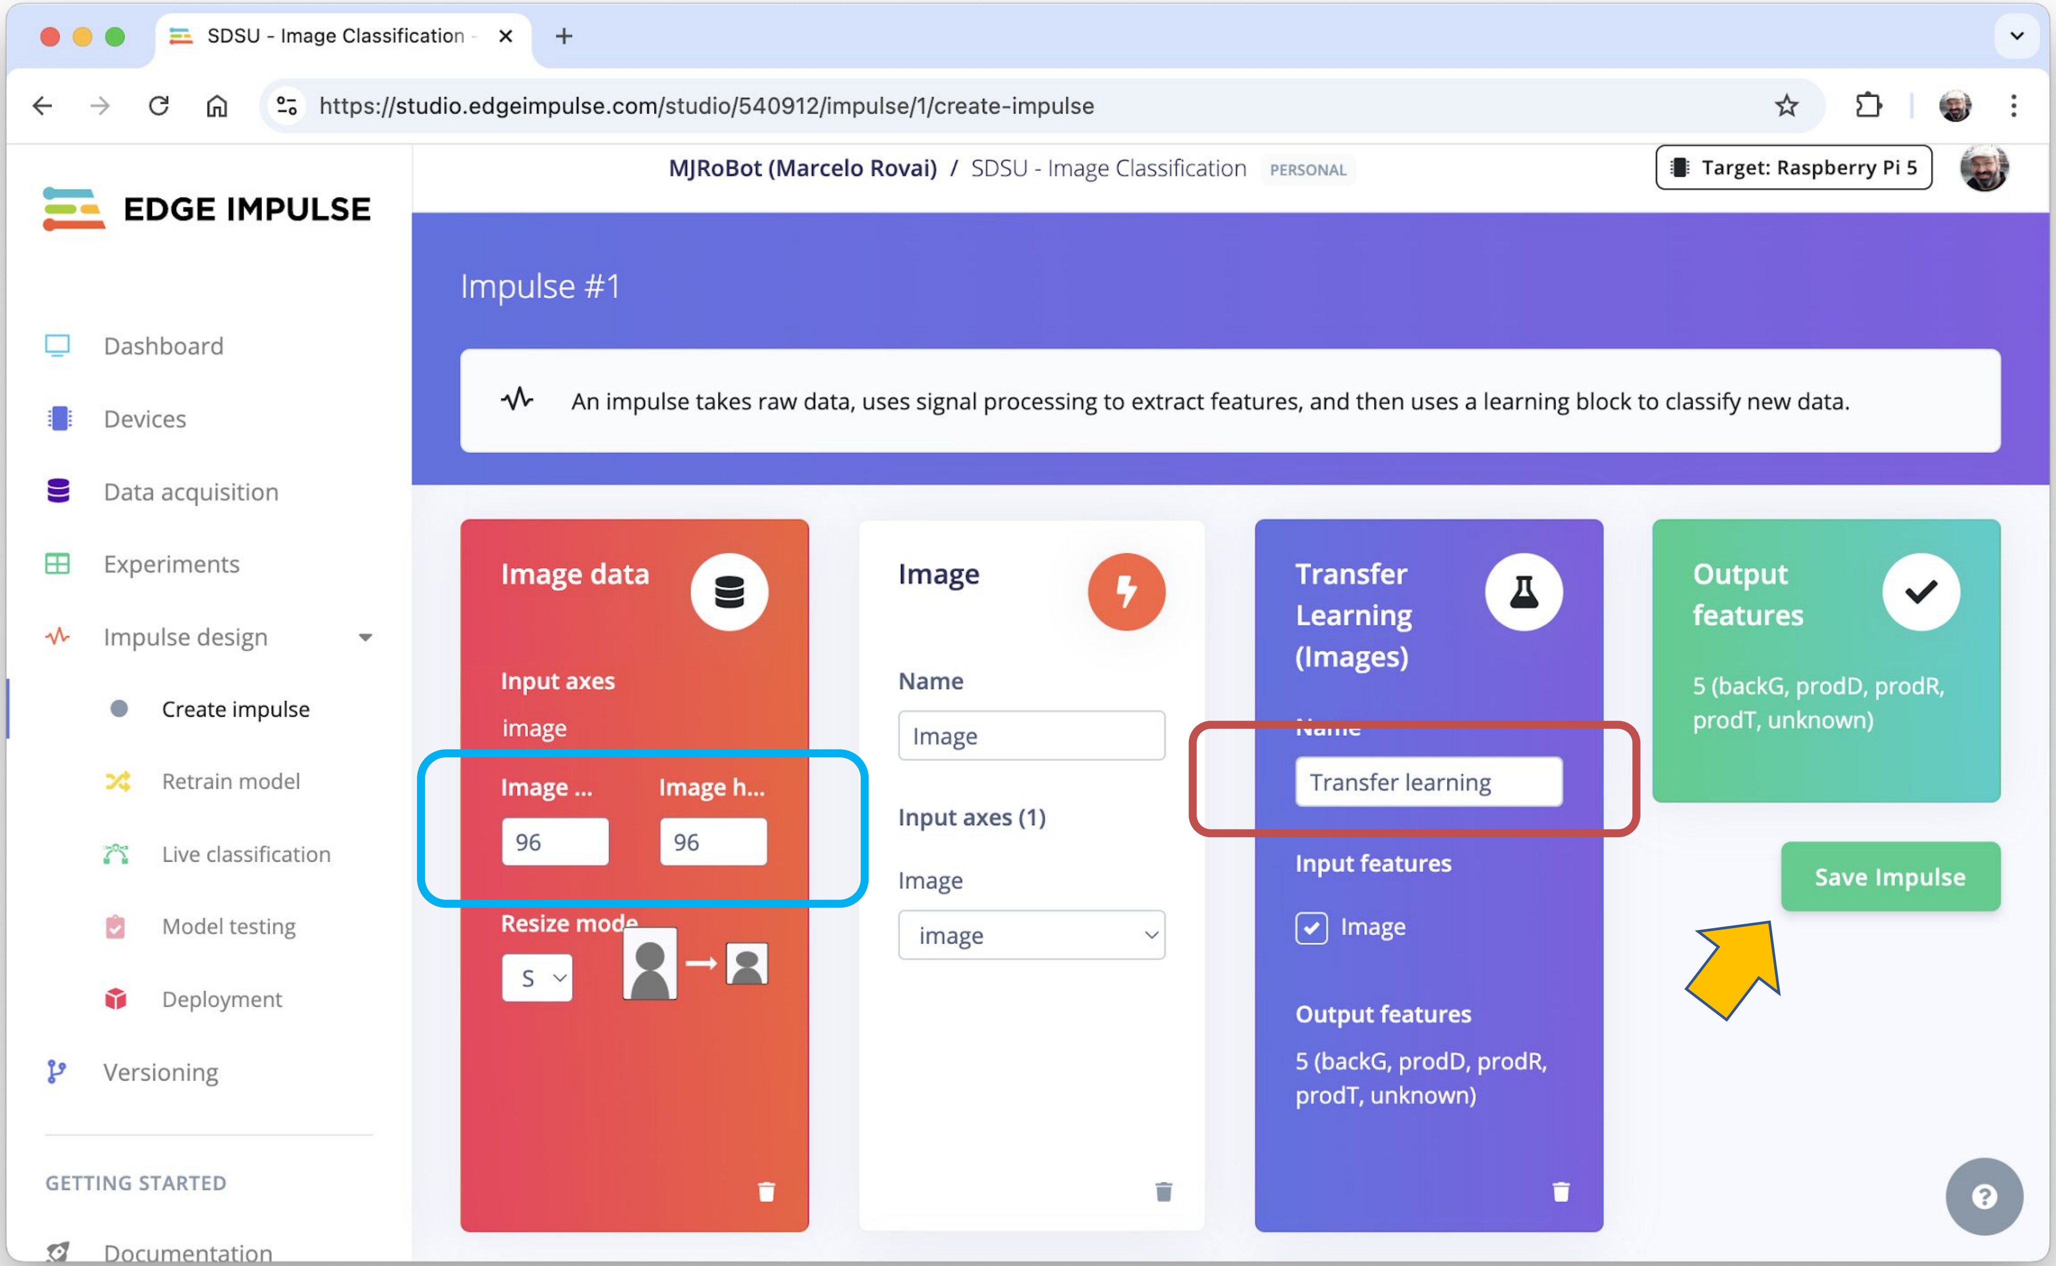
SDSU - Image Classification
https://studio.edgeimpulse.com/studio/540912/impulse/1/create-impulse
MJRoBot (Marcelo Rovai) / SDSU - Image Classification PERSONAL
Target: Raspberry Pi 5
EDGE IMPULSE
Dashboard
Devices
Data acquisition
Experiments
Impulse design
Create impulse
Retrain model
Live classification
Model testing
Deployment
Versioning
GETTING STARTED
Documentation
Impulse #1
An impulse takes raw data, uses signal processing to extract features, and then uses a learning block to classify new data.
Image data
Input axes
image
Image ...
Image h...
96
96
Resize mode
Image
Name
Image
Input axes (1)
Image
image
Transfer Learning (Images)
Name
Transfer learning
Input features
Image
Output features
5 (backG, prodD, prodR, prodT, unknown)
Output features
5 (backG, prodD, prodR, prodT, unknown)
Save Impulse

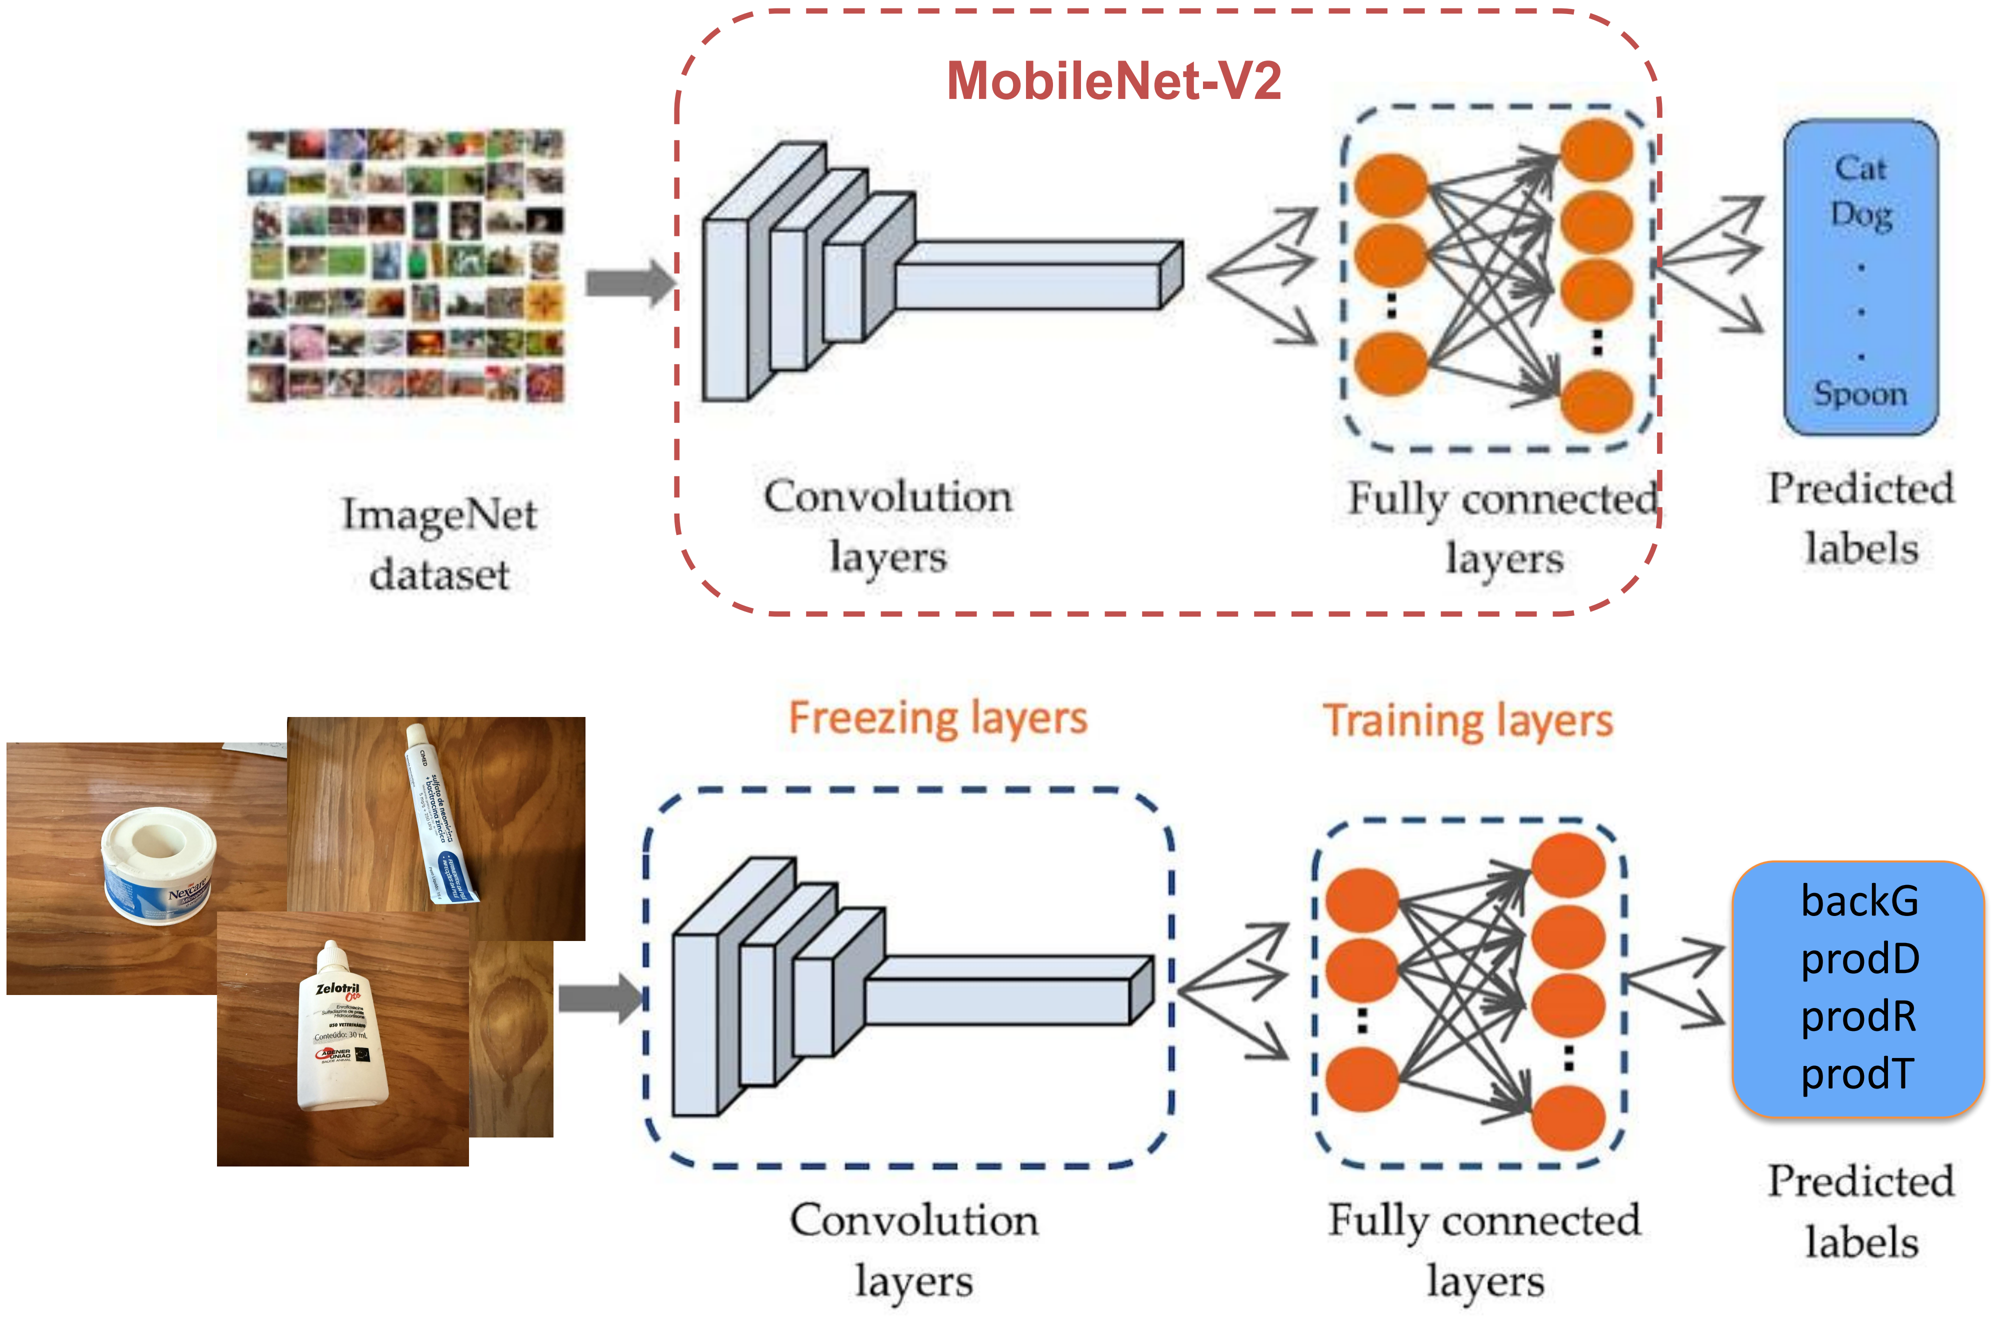
MobileNet-V2
ImageNet dataset
Convolution layers
Fully connected layers
Cat
Dog
Spoon
Predicted labels
Freezing layers
Training layers
Convolution layers
Fully connected layers
backG
prodD
prodR
prodT
Predicted labels

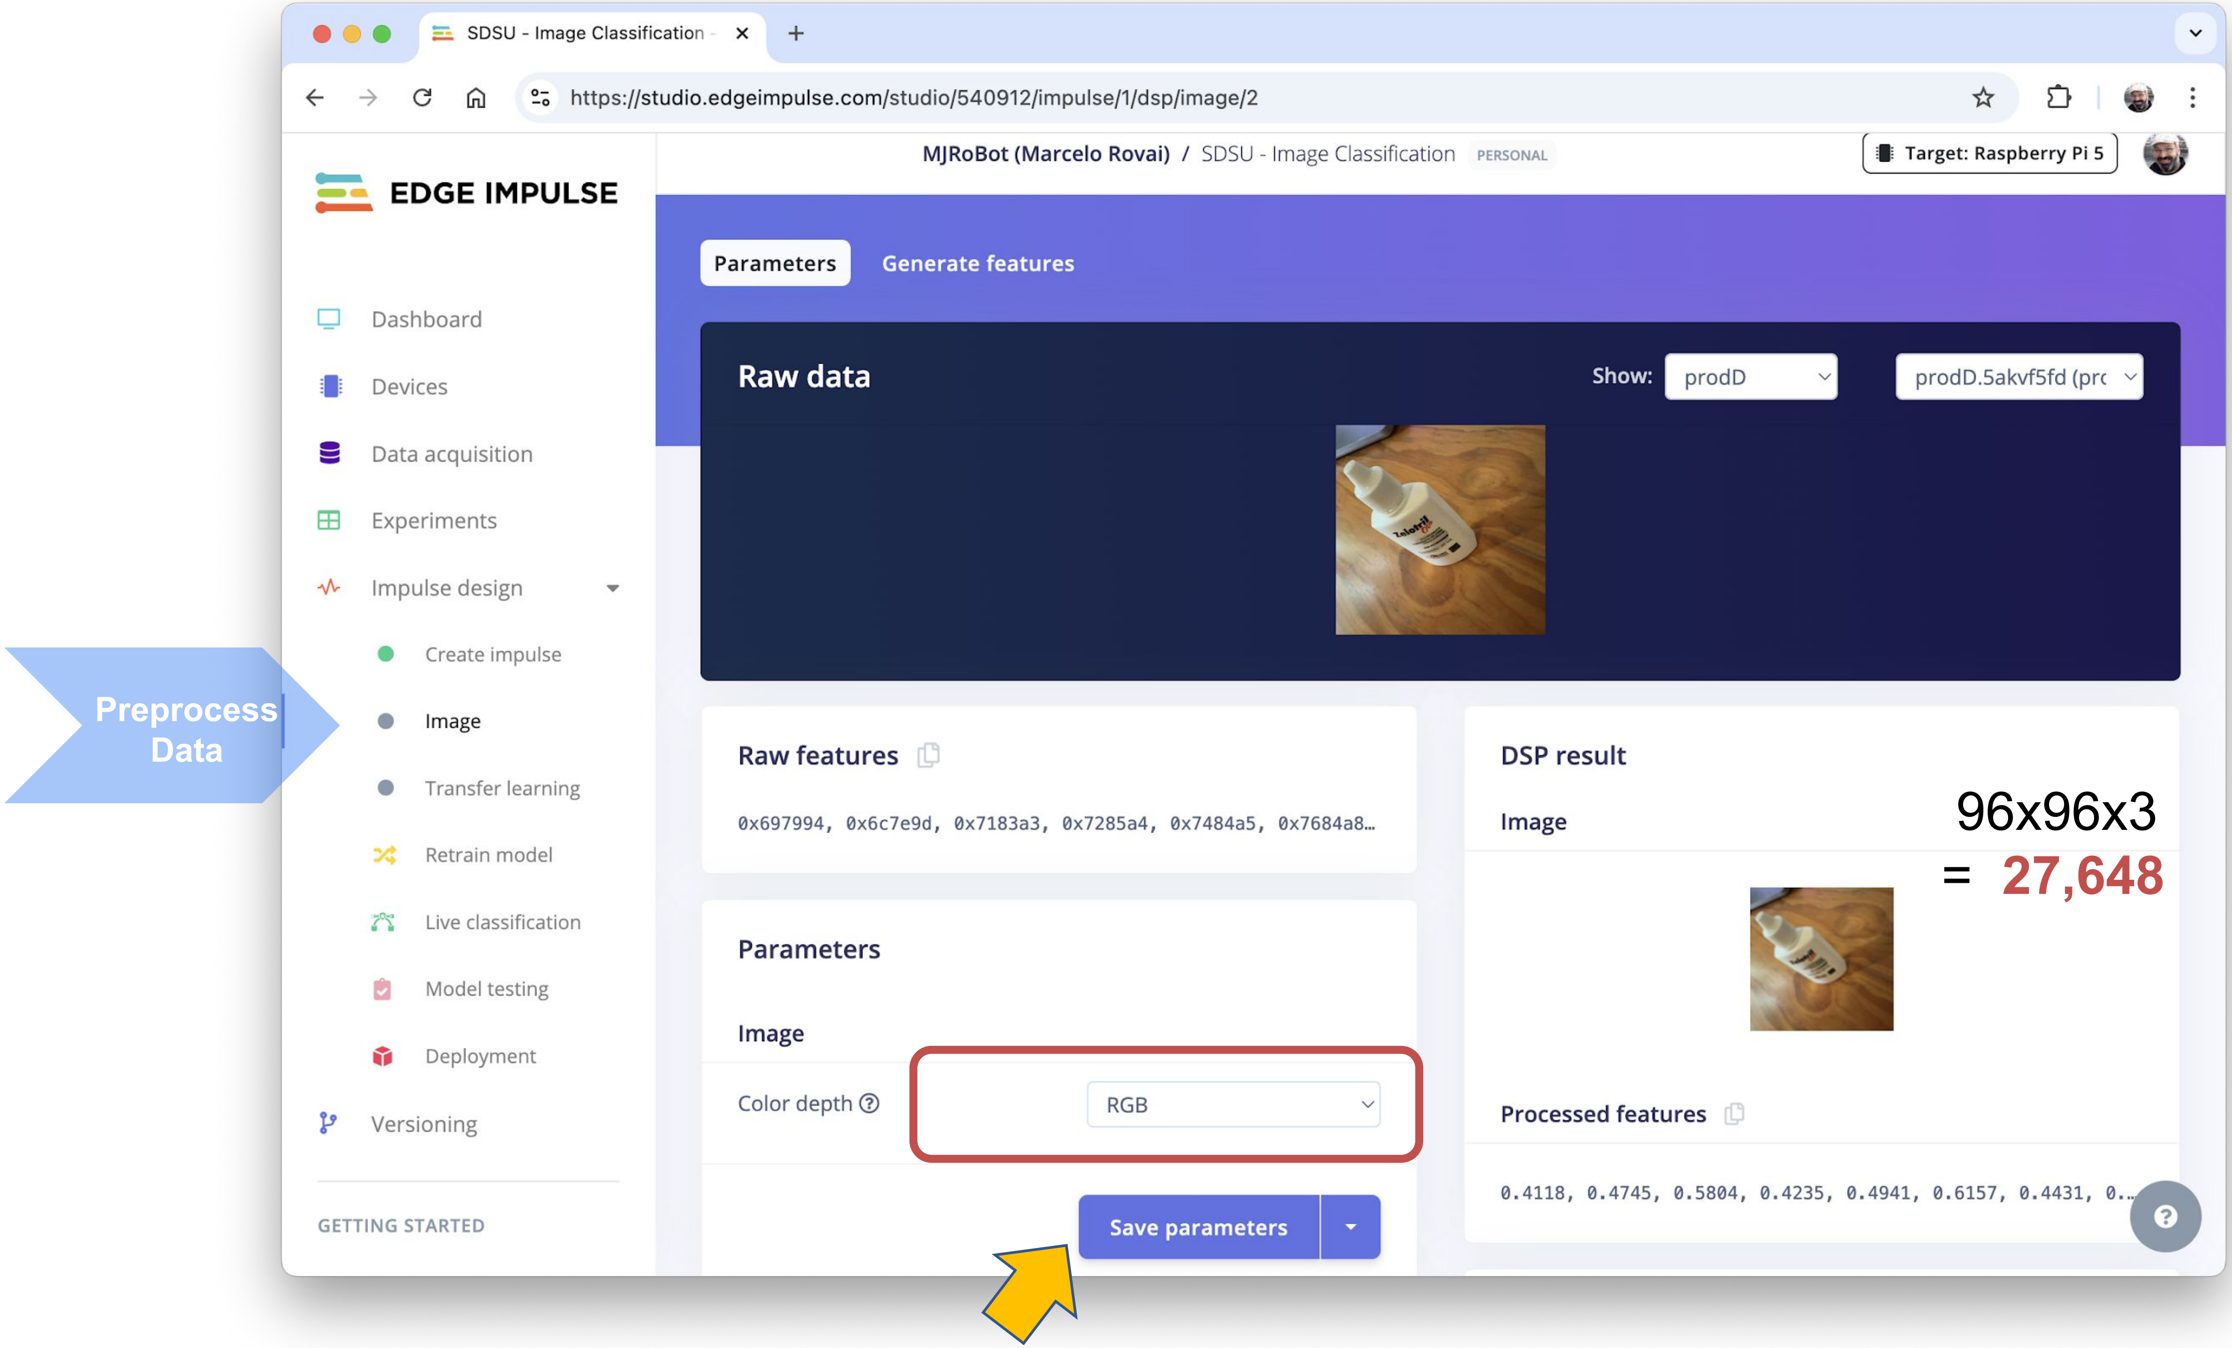

SDSU - Image Classification
https://studio.edgeimpulse.com/studio/540912/impulse/1/dsp/image/2
EDGE IMPULSE
MJRoBot (Marcelo Rovai) / SDSU - Image Classification
PERSONAL
Target: Raspberry Pi 5
Dashboard
Devices
Data acquisition
Experiments
Impulse design
Create impulse
Image
Transfer learning
Retrain model
Live classification
Model testing
Deployment
Versioning
GETTING STARTED
Preprocess Data
Parameters
Generate features
Raw data
Show:
prodD
prodD.5akvf5fd (prc
Raw features
0x697994, 0x6c7e9d, 0x7183a3, 0x7285a4, 0x7484a5, 0x7684a8…
Parameters
Image
Color depth
RGB
Save parameters
DSP result
Image
96x96x3
= 27,648
Processed features
0.4118, 0.4745, 0.5804, 0.4235, 0.4941, 0.6157, 0.4431, 0.…

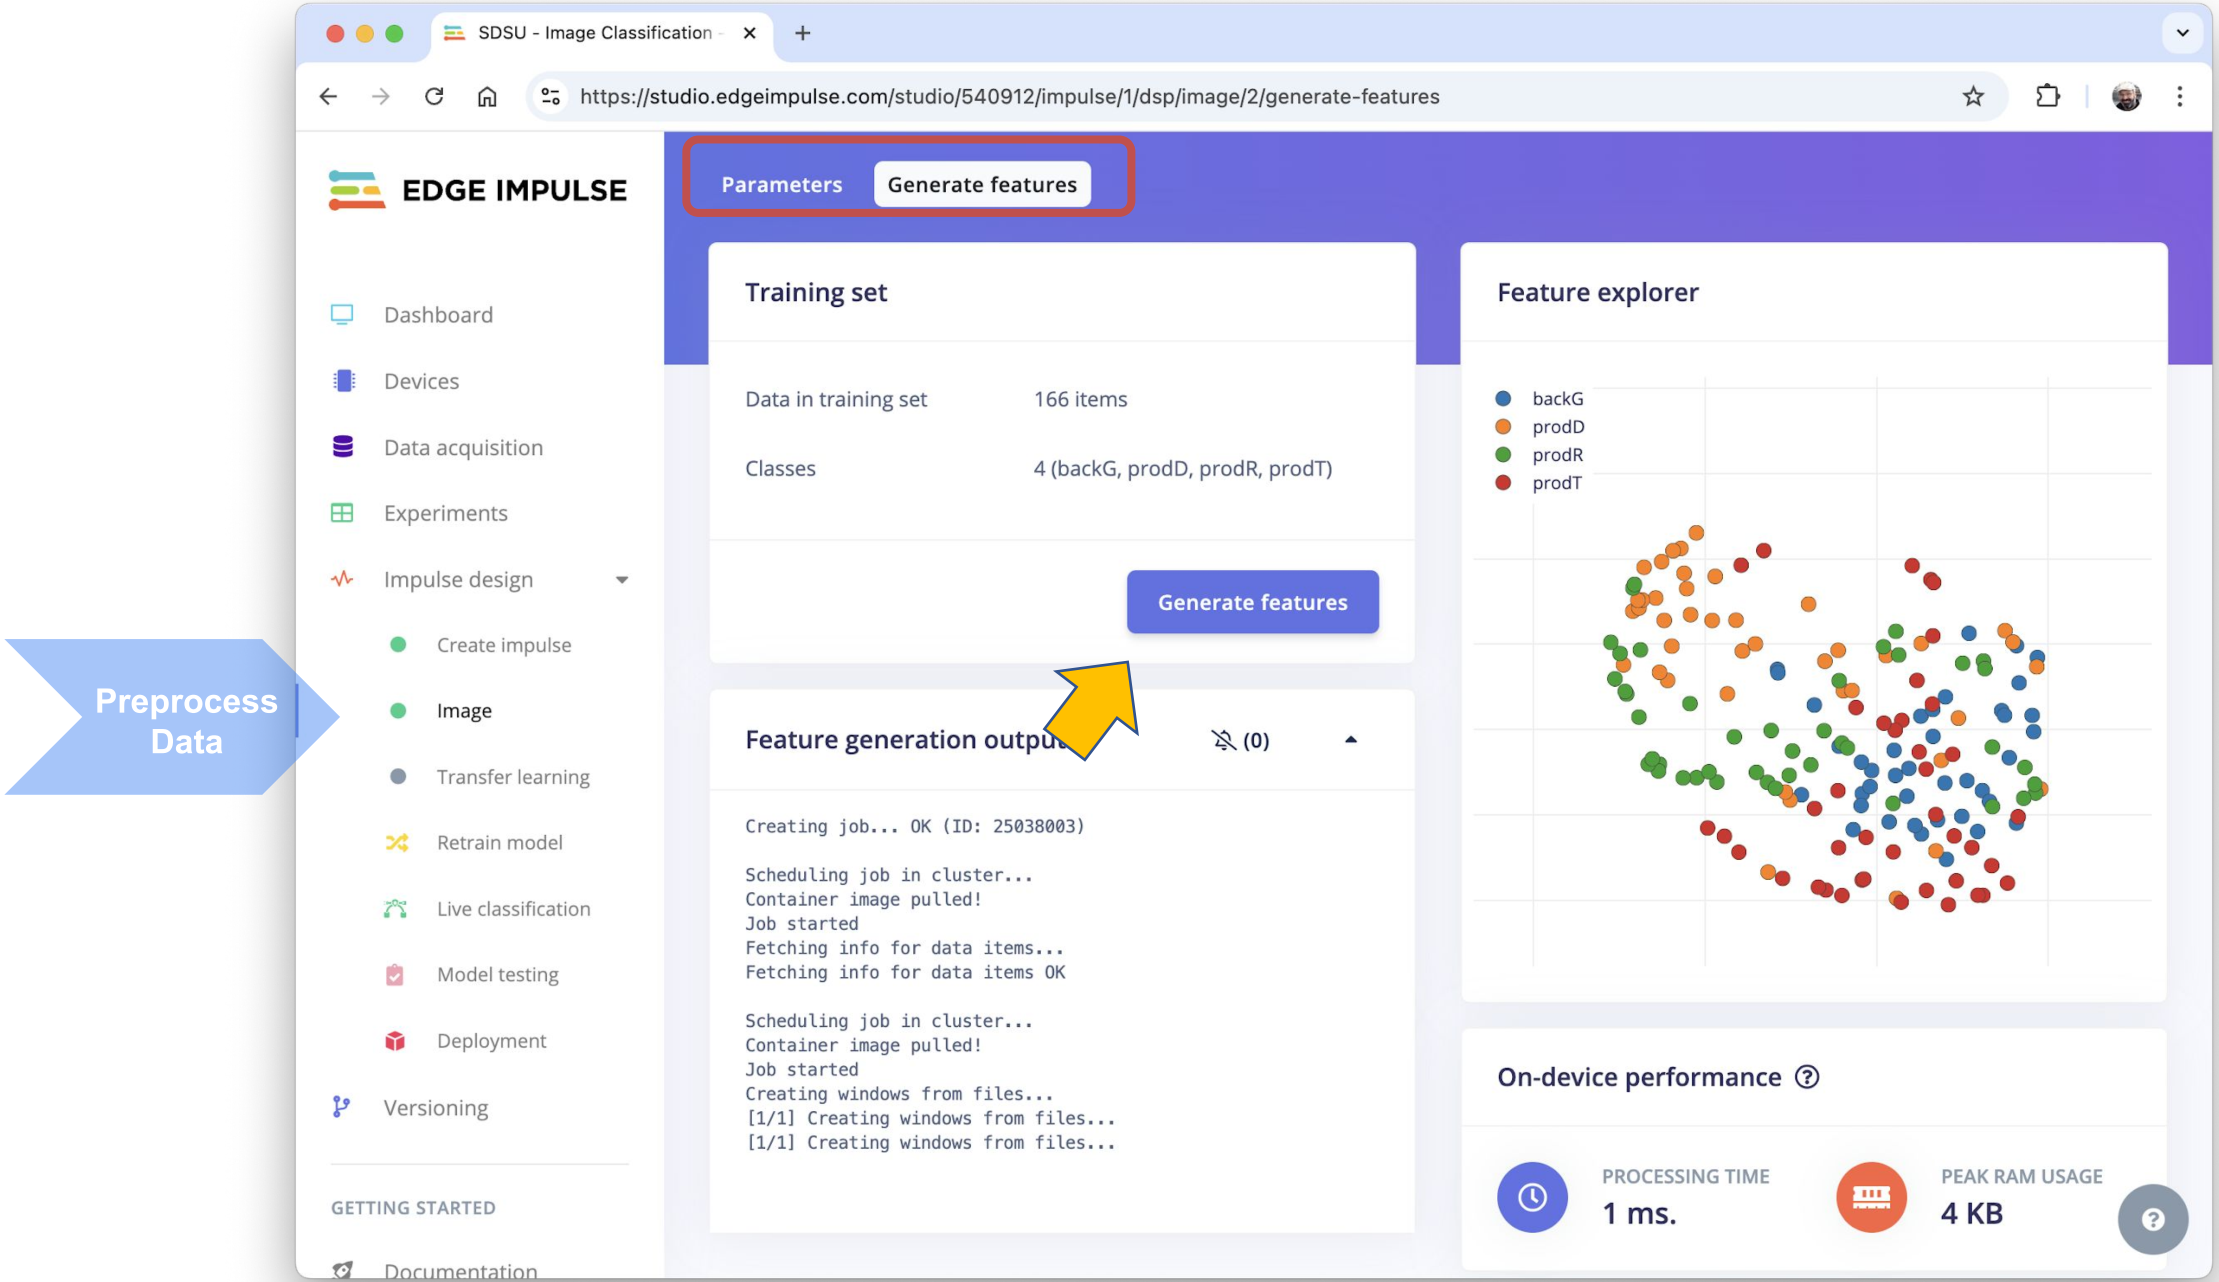
SDSU - Image Classification -
https://studio.edgeimpulse.com/studio/540912/impulse/1/dsp/image/2/generate-features
EDGE IMPULSE
Parameters
Generate features
Dashboard
Devices
Data acquisition
Experiments
Impulse design
Create impulse
Image
Transfer learning
Retrain model
Live classification
Model testing
Deployment
Versioning
GETTING STARTED
Documentation
Preprocess Data
Training set
Data in training set
166 items
Classes
4 (backG, prodD, prodR, prodT)
Generate features
Feature generation output
(0)
Creating job... OK (ID: 25038003)
Scheduling job in cluster...
Container image pulled!
Job started
Fetching info for data items...
Fetching info for data items OK
Scheduling job in cluster...
Container image pulled!
Job started
Creating windows from files...
[1/1] Creating windows from files...
[1/1] Creating windows from files...
Feature explorer
backG
prodD
prodR
prodT
On-device performance
PROCESSING TIME
1 ms.
PEAK RAM USAGE
4 KB

Neural Network settings

Hyperparameters

Training settings

Number of training cycles ⓘ	20
Use learned optimizer ⓘ	☐
Learning rate ⓘ	0.0005
Training processor ⓘ	CPU
Data augmentation ⓘ	☑

Advanced training settings

Validation set size ⓘ	20 %
Split train/validation set on metadata key ⓘ	
Batch size ⓘ	32
Auto-weight classes ⓘ	☐
Profile int8 model ⓘ	☑

Neural network architecture

Model Design

Input layer (27,648 features)

MobileNetV2 96x96 0.35 (final layer: 16 neurons, 0.1 dropout)

Choose a different model

Output layer (4 classes)

Save & train

Last training performance (validation set)

ACCURACY
85.3%

LOSS
0.43

Confusion matrix (validation set)

	BACKG	PRODD	PRODR	PRODT
BACKG	77.8%	11.1%	11.1%	0%
PRODD	0%	71.4%	0%	28.6%
PRODR	0%	0%	85.7%	14.3%
PRODT	0%	0%	0%	100%
F1 SCORE	0.88	0.77	0.86	0.88

Metrics (validation set)

METRIC	VALUE
Area under ROC Curve	0.98
Weighted average Precision	0.87
Weighted average Recall	0.85
Weighted average F1 score	0.85

Data explorer (full training set)

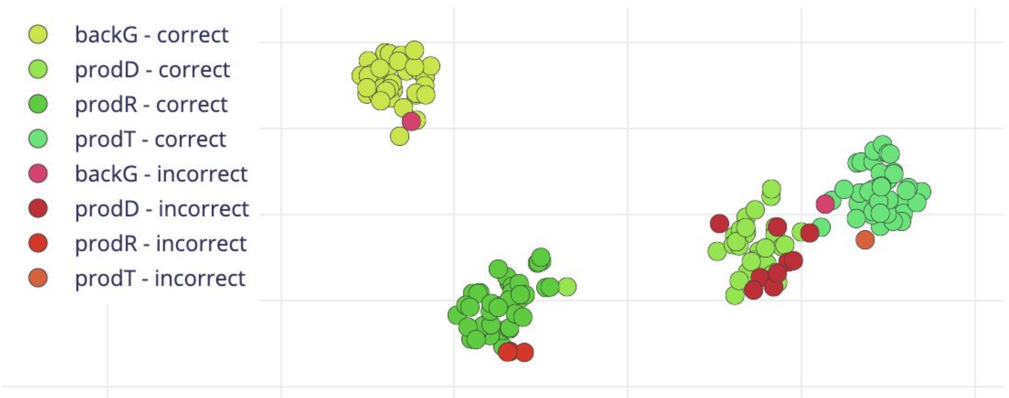

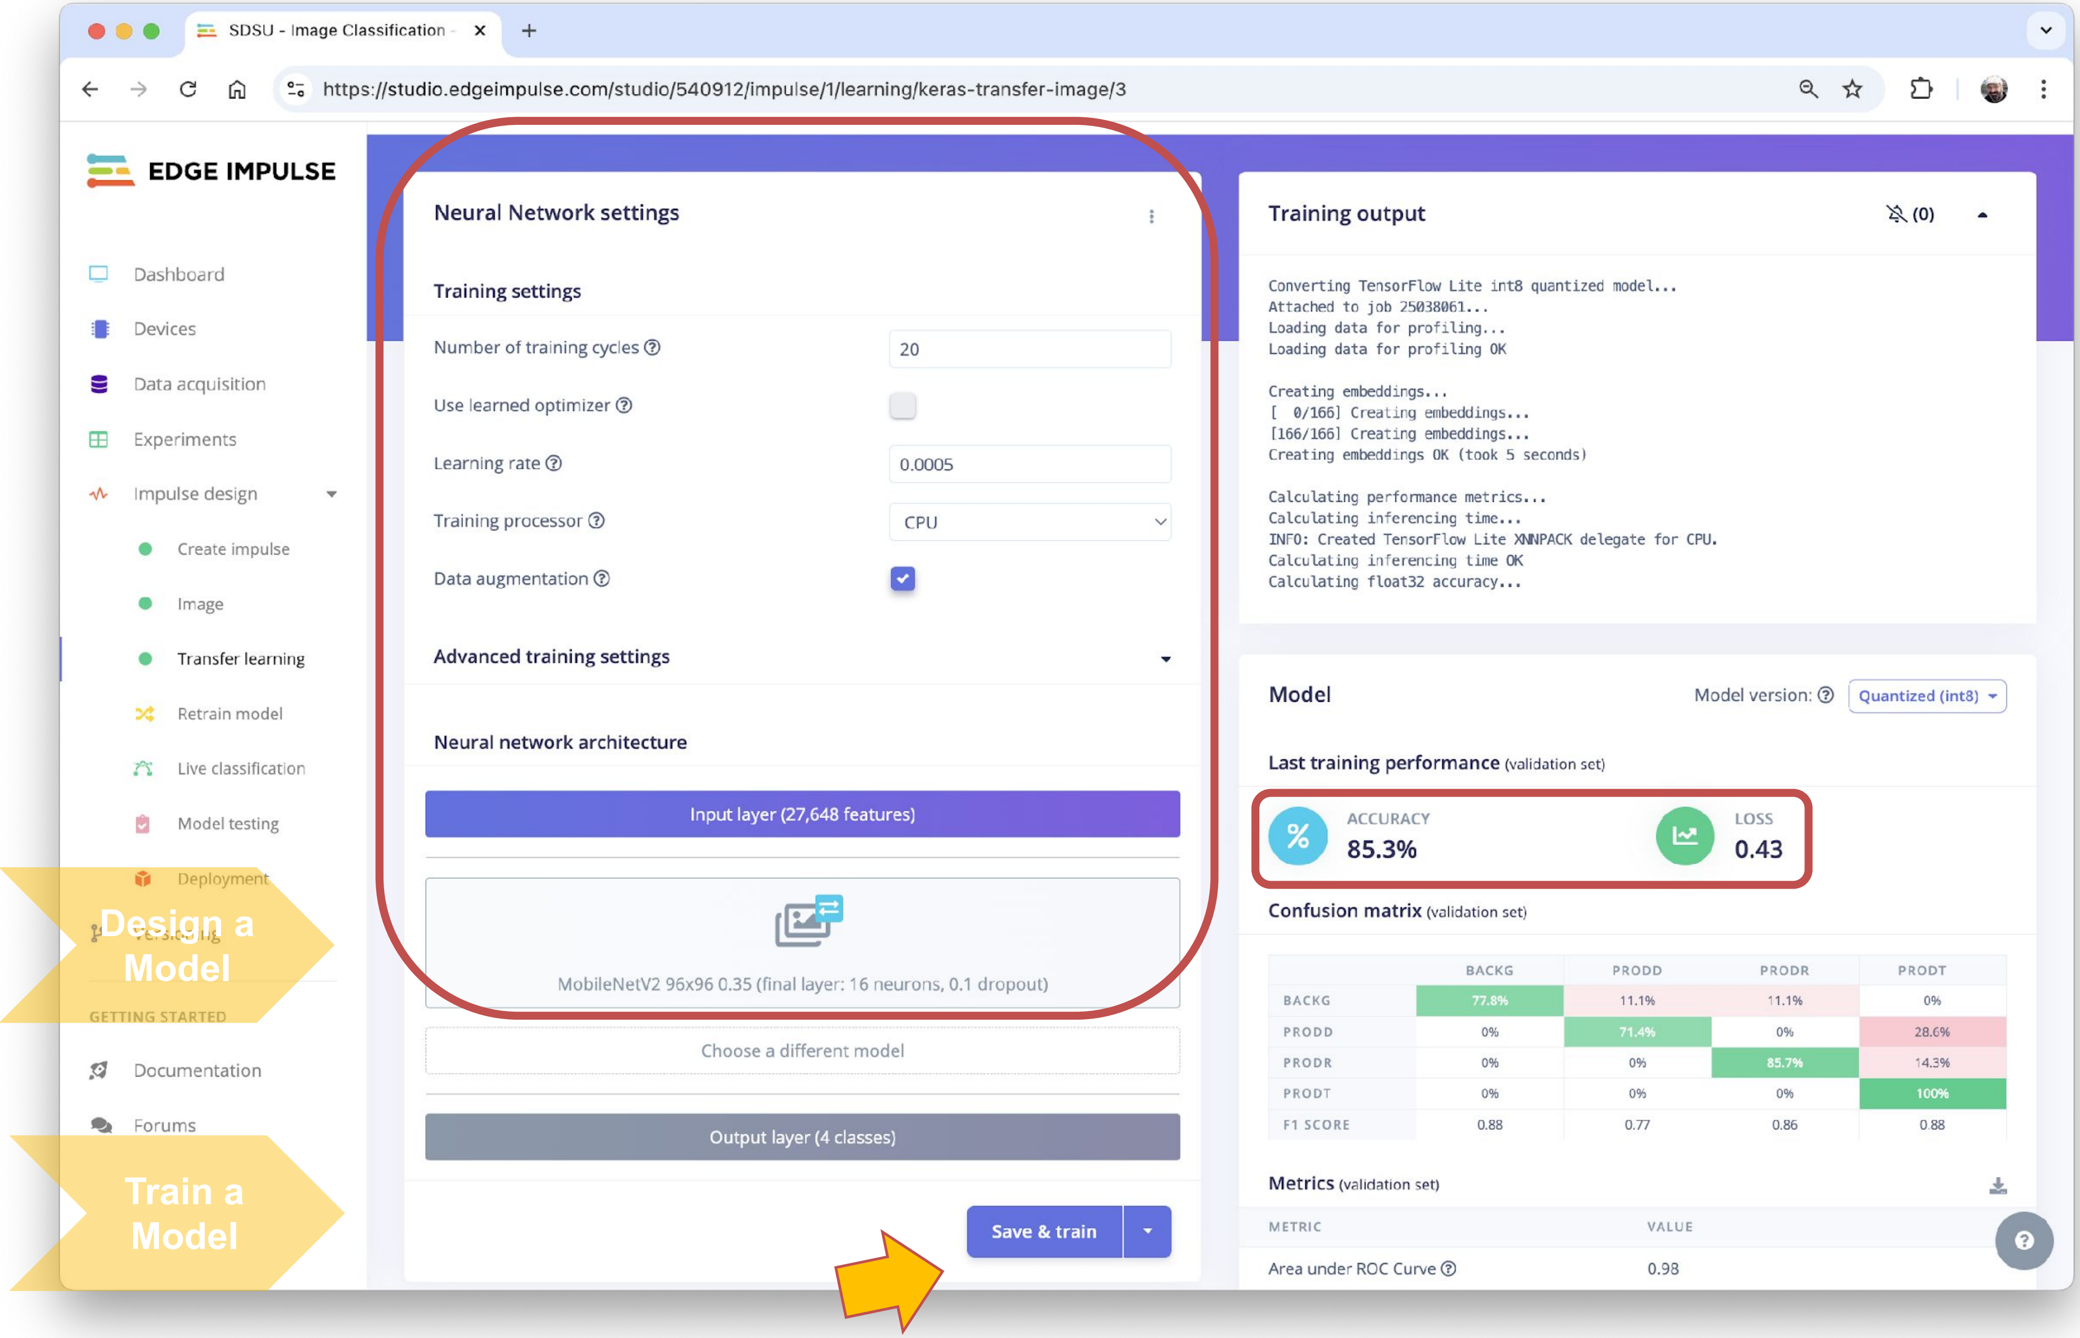
SDSU - Image Classification
https://studio.edgeimpulse.com/studio/540912/impulse/1/learning/keras-transfer-image/3
EDGE IMPULSE
Dashboard
Devices
Data acquisition
Experiments
Impulse design
Create impulse
Image
Transfer learning
Retrain model
Live classification
Model testing
Deployment
Design a Model
GETTING STARTED
Documentation
Forums
Train a Model
Neural Network settings
Training settings
Number of training cycles
20
Use learned optimizer
Learning rate
0.0005
Training processor
CPU
Data augmentation
Advanced training settings
Neural network architecture
Input layer (27,648 features)
MobileNetV2 96x96 0.35 (final layer: 16 neurons, 0.1 dropout)
Choose a different model
Output layer (4 classes)
Save & train
Training output
Converting TensorFlow Lite int8 quantized model...
Attached to job 25038061...
Loading data for profiling...
Loading data for profiling OK
Creating embeddings...
[0/166] Creating embeddings...
[166/166] Creating embeddings...
Creating embeddings OK (took 5 seconds)
Calculating performance metrics...
Calculating inferencing time...
INFO: Created TensorFlow Lite XNNPACK delegate for CPU.
Calculating inferencing time OK
Calculating float32 accuracy...
Model
Model version:
Quantized (int8)
Last training performance (validation set)
ACCURACY
85.3%
LOSS
0.43
Confusion matrix (validation set)
BACKG
PRODD
PRODR
PRODT
BACKG 77.8% 11.1% 11.1% 0%
PRODD 0% 71.4% 0% 28.6%
PRODR 0% 0% 85.7% 14.3%
PRODT 0% 0% 0% 100%
F1 SCORE 0.88 0.77 0.86 0.88
Metrics (validation set)
METRIC
VALUE
Area under ROC Curve
0.98

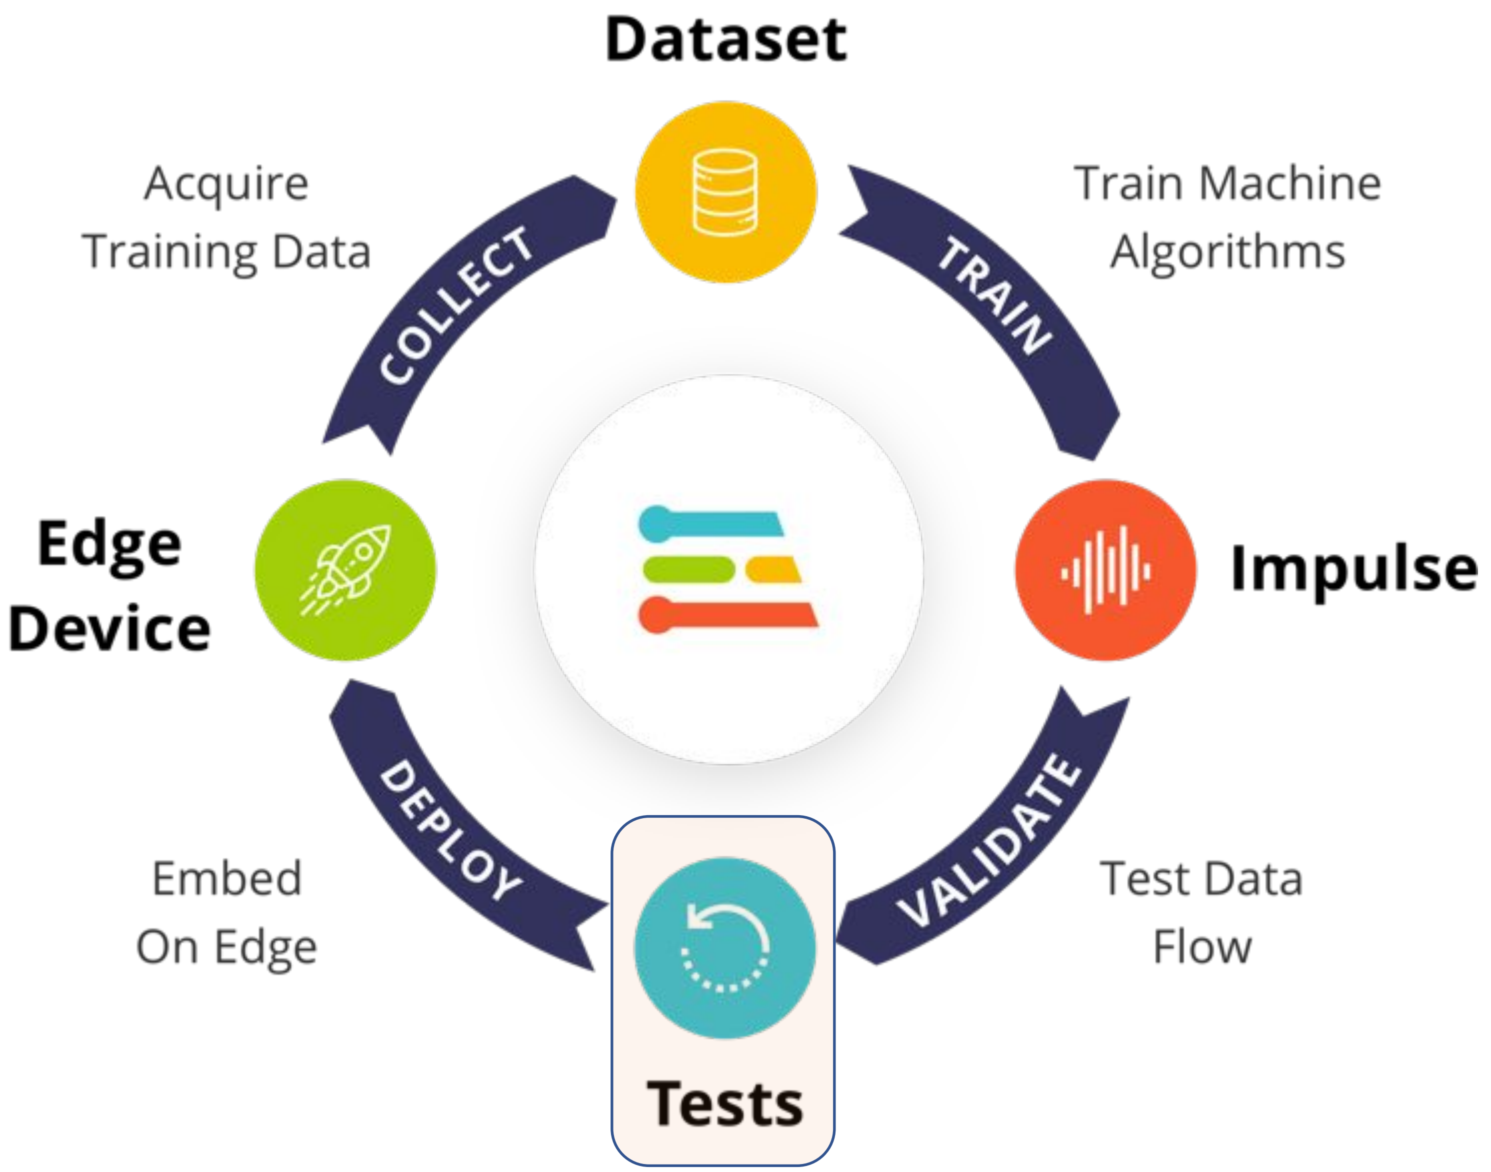

- Pre-Processing Data
- Design a Model
- Train a Model

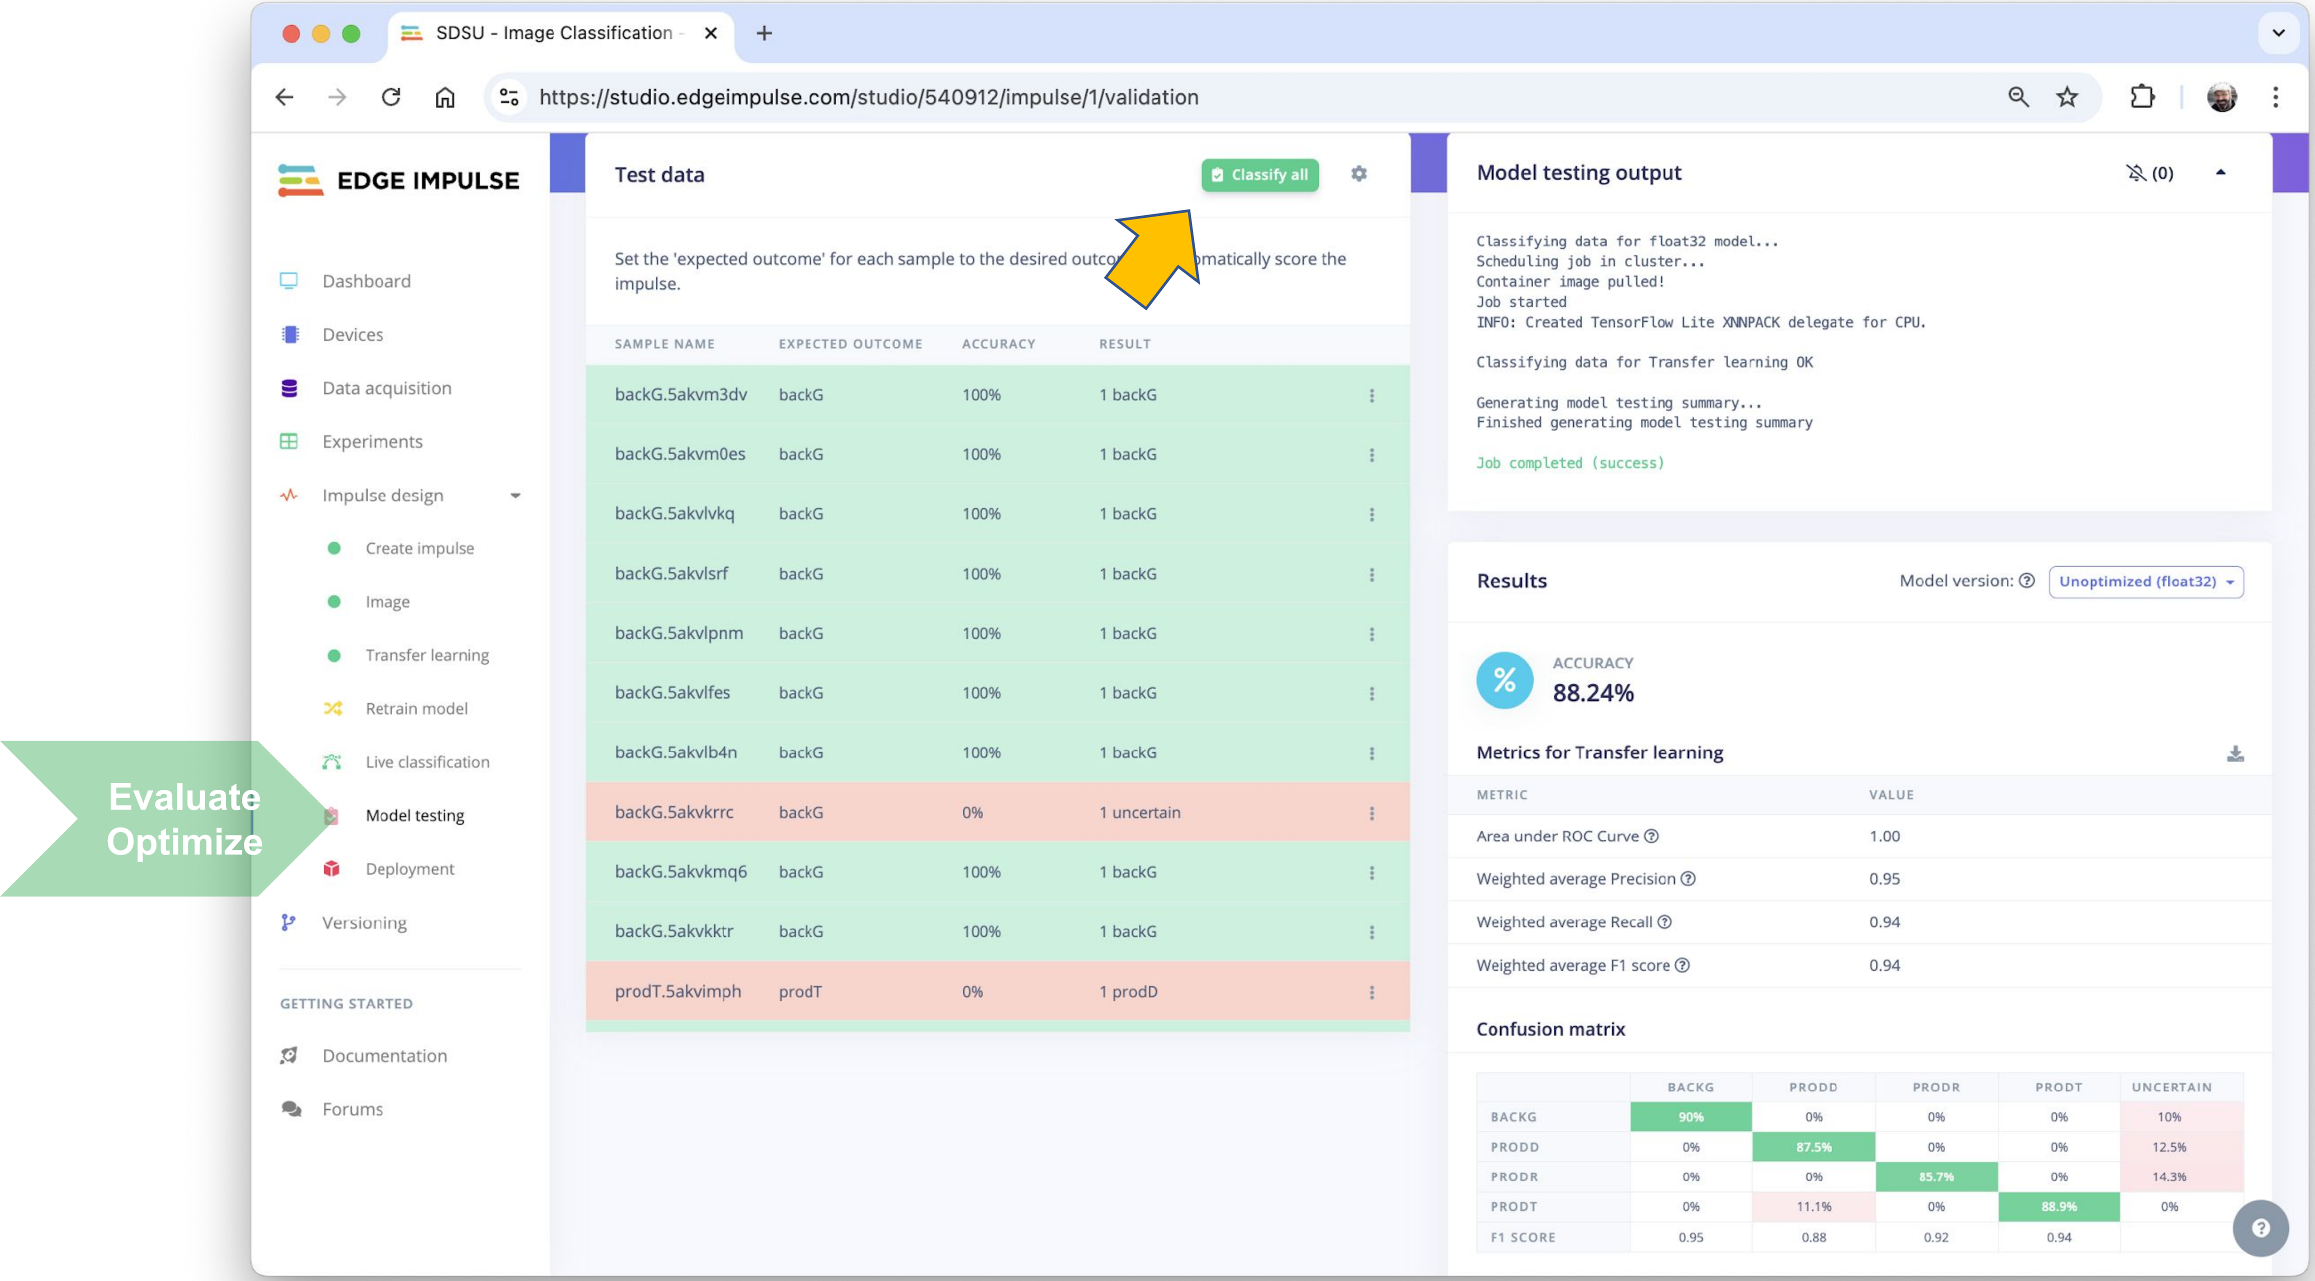
SDSU - Image Classification
https://studio.edgeimpulse.com/studio/540912/impulse/1/validation
EDGE IMPULSE
Dashboard
Devices
Data acquisition
Experiments
Impulse design
Create impulse
Image
Transfer learning
Retrain model
Live classification
Model testing
Deployment
Versioning
GETTING STARTED
Documentation
Forums
Evaluate
Optimize
Test data
Classify all
Set the 'expected outcome' for each sample to the desired outco
omatically score the
impulse.
SAMPLE NAME
EXPECTED OUTCOME
ACCURACY
RESULT
backG.5akvm3dv backG 100% 1 backG
backG.5akvm0es backG 100% 1 backG
backG.5akvlvkq backG 100% 1 backG
backG.5akvlsrf backG 100% 1 backG
backG.5akvlpnm backG 100% 1 backG
backG.5akvlfes backG 100% 1 backG
backG.5akvlb4n backG 100% 1 backG
backG.5akvkrrc backG 0% 1 uncertain
backG.5akvkmq6 backG 100% 1 backG
backG.5akvkktr backG 100% 1 backG
prodT.5akvimph prodT 0% 1 prodD
Model testing output
(0)
Classifying data for float32 model...
Scheduling job in cluster...
Container image pulled!
Job started
INFO: Created TensorFlow Lite XNNPACK delegate for CPU.
Classifying data for Transfer learning OK
Generating model testing summary...
Finished generating model testing summary
Job completed (success)
Results
Model version:
Unoptimized (float32)
ACCURACY
88.24%
Metrics for Transfer learning
METRIC
VALUE
Area under ROC Curve 1.00
Weighted average Precision 0.95
Weighted average Recall 0.94
Weighted average F1 score 0.94
Confusion matrix
BACKG PRODD PRODR PRODT UNCERTAIN
BACKG 90% 0% 0% 0% 10%
PRODD 0% 87.5% 0% 0% 12.5%
PRODR 0% 0% 85.7% 0% 14.3%
PRODT 0% 11.1% 0% 88.9% 0%
F1 SCORE 0.95 0.88 0.92 0.94

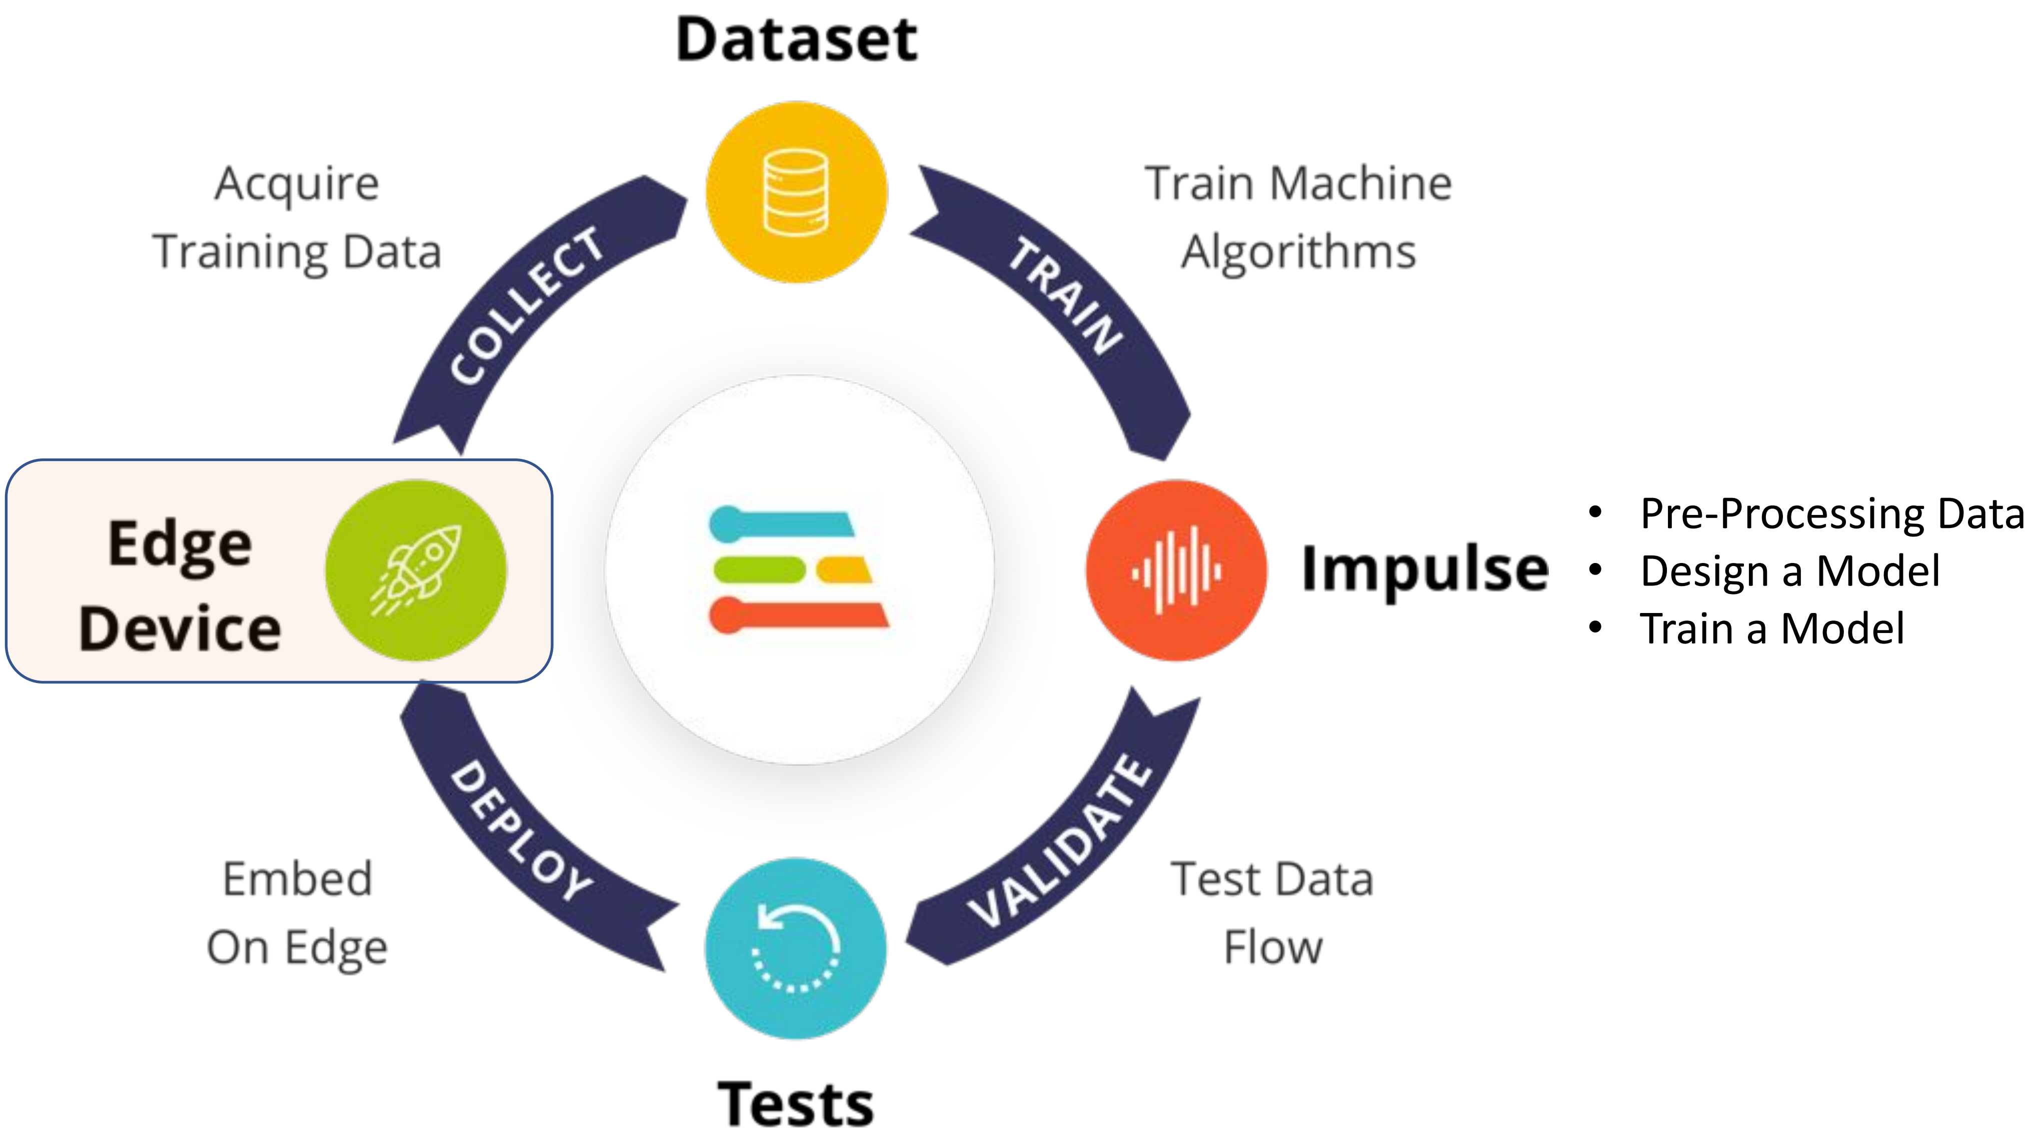
Dataset
Acquire
Training Data
COLLECT
Train Machine
Algorithms
TRAIN
Edge
Device
Impulse
Pre-Processing Data
Design a Model
Train a Model
DEPLOY
VALIDATE
Embed
On Edge
Test Data
Flow
Tests

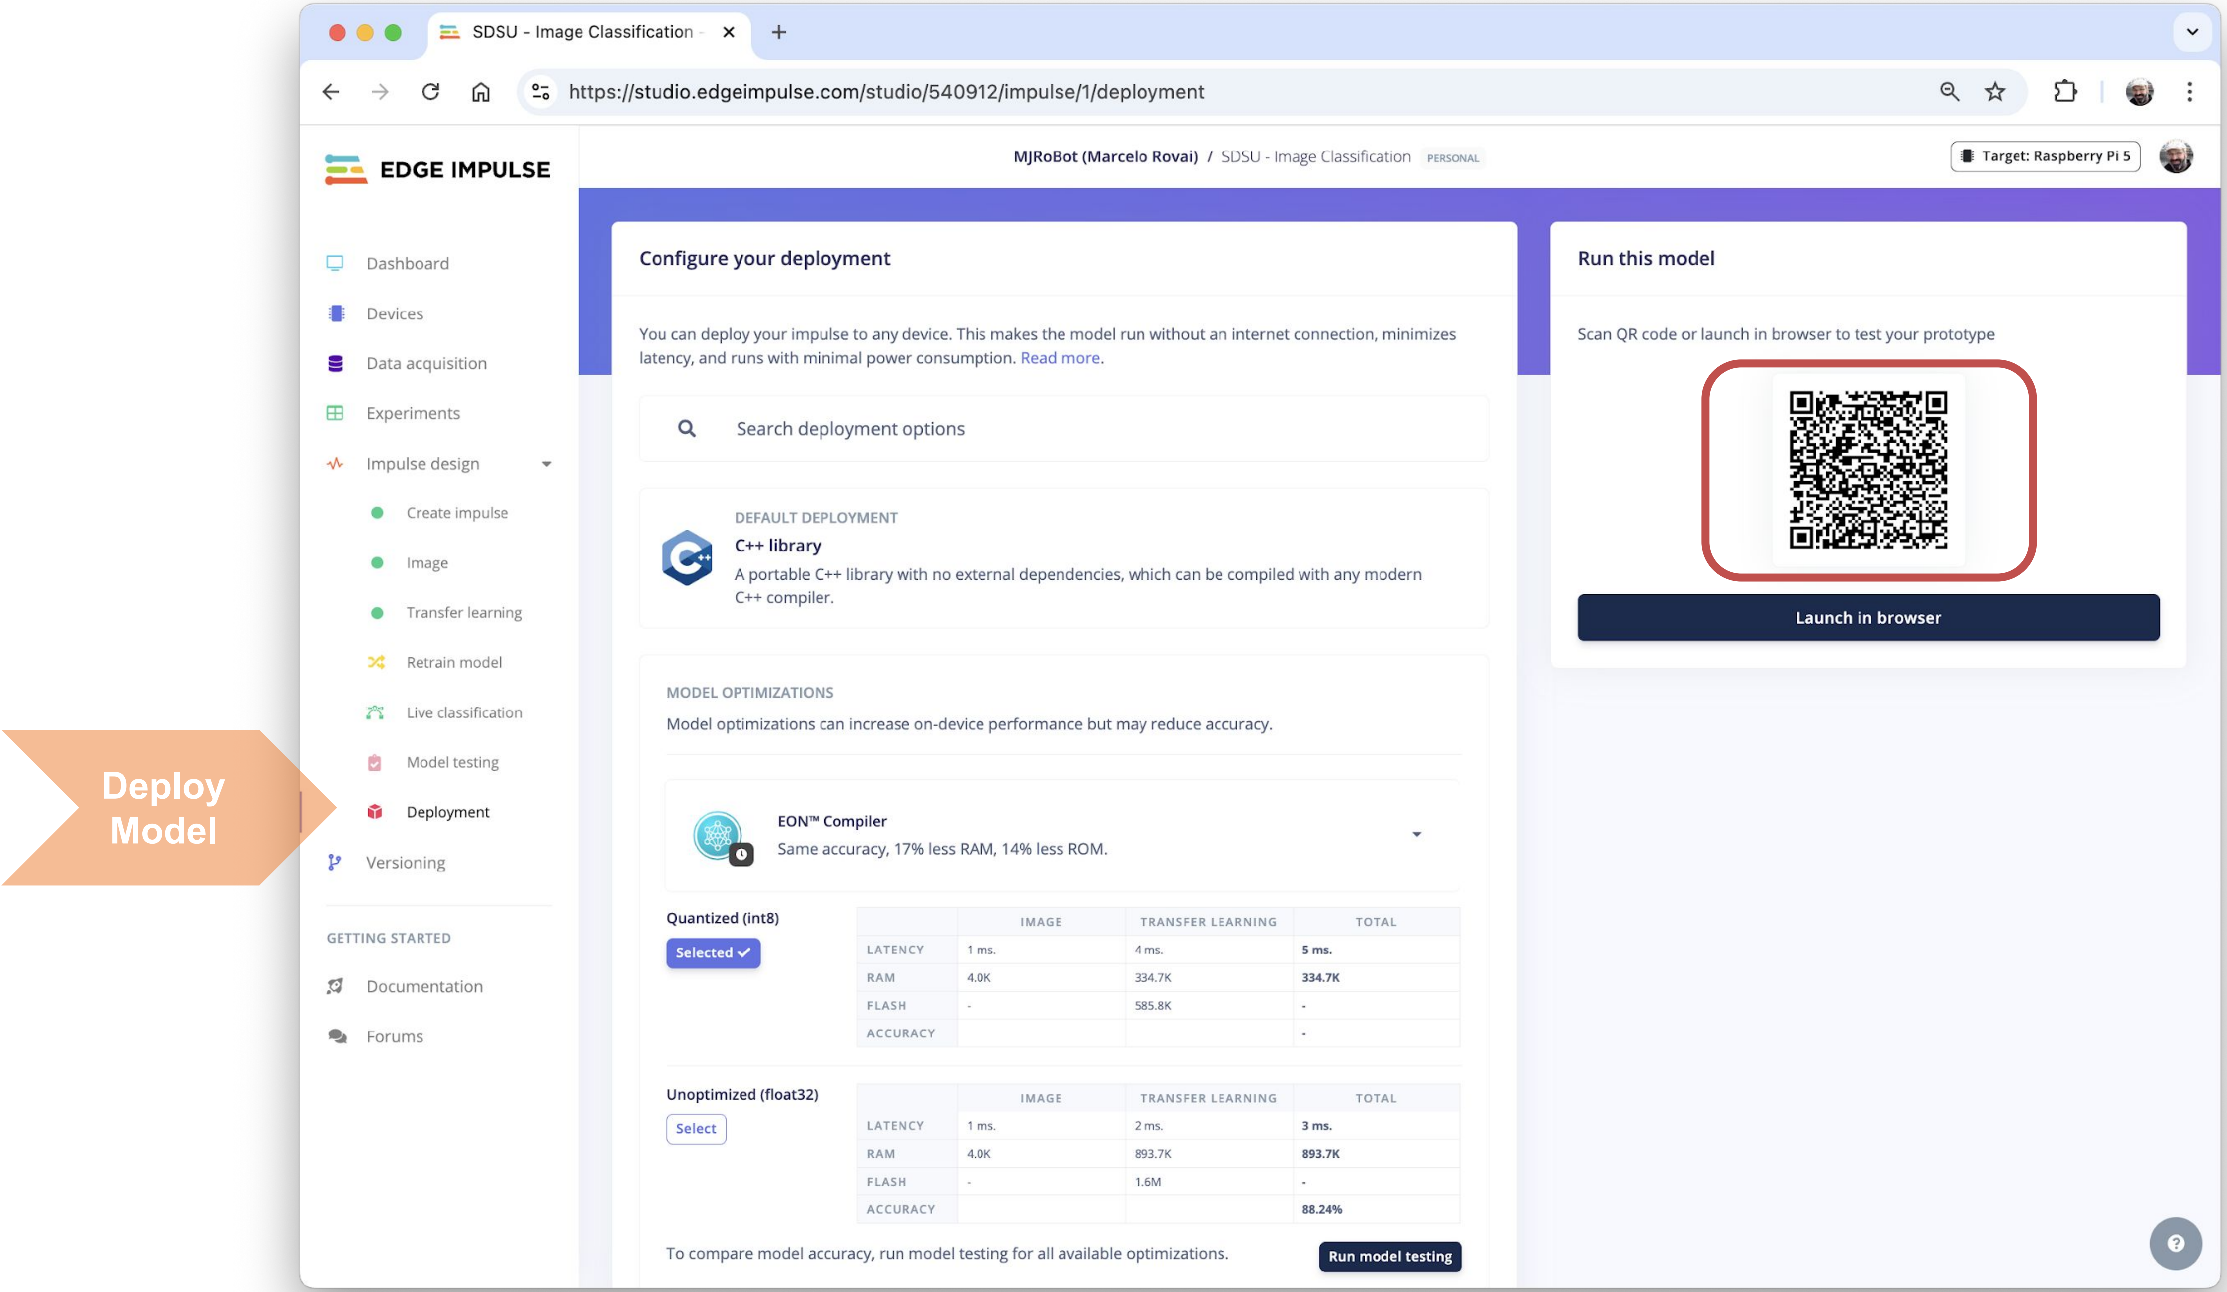
SDSU - Image Classification
https://studio.edgeimpulse.com/studio/540912/impulse/1/deployment
EDGE IMPULSE
MJRoBot (Marcelo Rovai) / SDSU - Image Classification PERSONAL
Target: Raspberry Pi 5
Dashboard
Devices
Data acquisition
Experiments
Impulse design
Create impulse
Image
Transfer learning
Retrain model
Live classification
Model testing
Deployment
Versioning
GETTING STARTED
Documentation
Forums
Deploy Model
Configure your deployment
You can deploy your impulse to any device. This makes the model run without an internet connection, minimizes latency, and runs with minimal power consumption. Read more.
Search deployment options
DEFAULT DEPLOYMENT
C++ library
A portable C++ library with no external dependencies, which can be compiled with any modern C++ compiler.
MODEL OPTIMIZATIONS
Model optimizations can increase on-device performance but may reduce accuracy.
EON™ Compiler
Same accuracy, 17% less RAM, 14% less ROM.
Quantized (int8)
Selected
IMAGE
TRANSFER LEARNING
TOTAL
LATENCY 1 ms. 4 ms. 5 ms.
RAM 4.0K 334.7K 334.7K
FLASH - 585.8K -
ACCURACY -
Unoptimized (float32)
Select
LATENCY 1 ms. 2 ms. 3 ms.
RAM 4.0K 893.7K 893.7K
FLASH - 1.6M -
ACCURACY 88.24%
To compare model accuracy, run model testing for all available optimizations.
Run model testing
Run this model
Scan QR code or launch in browser to test your prototype
Launch in browser

Make Inferences

Inference – Real World

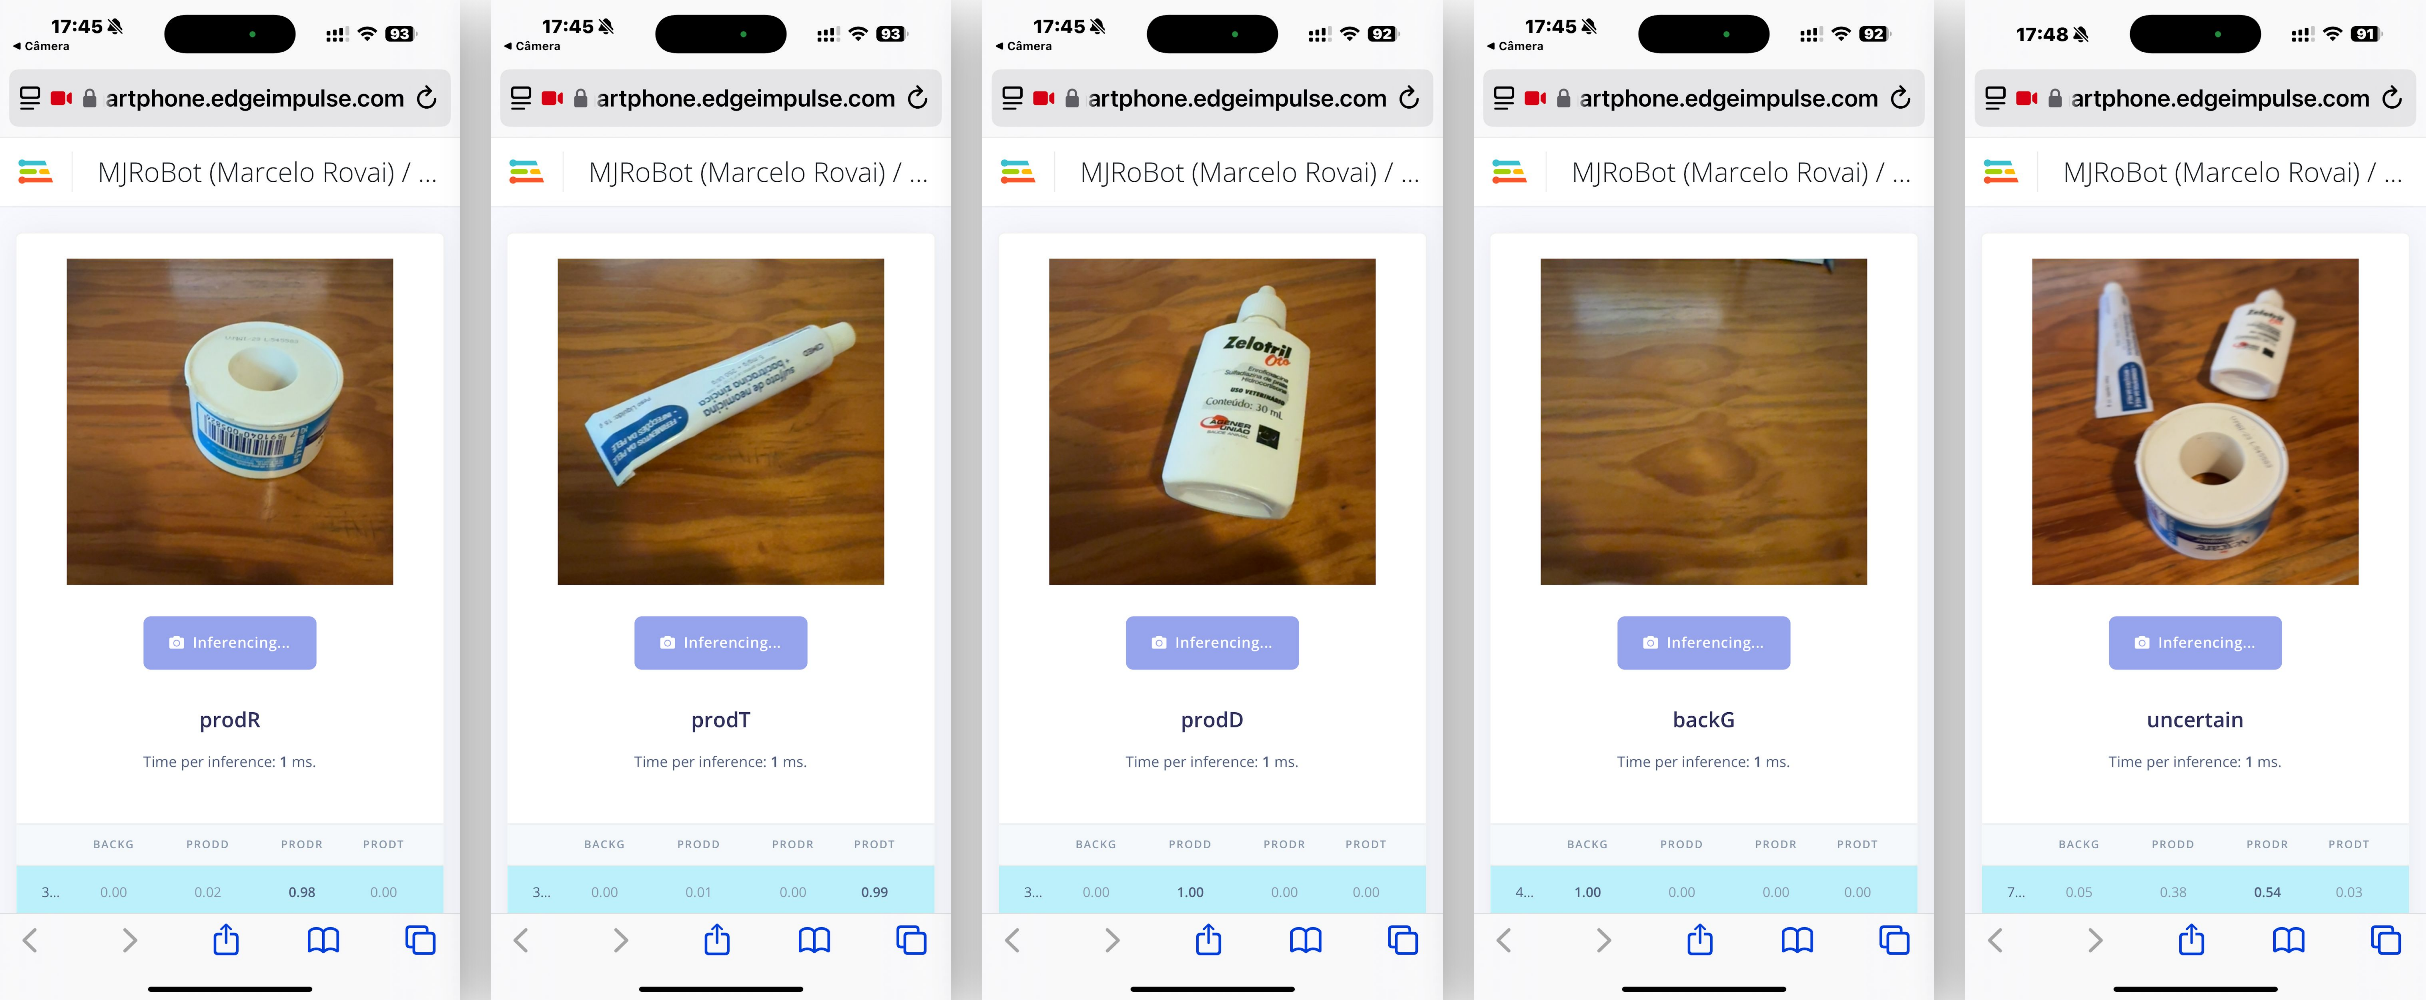

Conclusion

How AI is Shaping Agriculture (Animal Farms)

- AI applications in diagnostics, treatment, surgery assistance, behavior analysis, and livestock management are transforming veterinary care.
- AI tools in radiography, ultrasound, and MRI help detect animal diseases early. AI-based image recognition improves diagnostic accuracy.
- AI analyzes cow behavior and physiological data to detect early signs of disease like mastitis or lameness, allowing for timely intervention.
- Wearable devices with AI monitor animal health in real-time, tracking heart rate, activity, and vital signs, providing insights for early interventions.
- Machine learning models track and analyze animal behavior, identifying patterns in stress, anxiety, or abnormal behavior.
- In livestock farming, AI helps manage herd health, improve nutrition, and enhance overall productivity by monitoring animals' welfare in real-time.

Ethical Considerations and future of AI

- Ethical concerns in AI use for dairy farming include data privacy, algorithm biases, and ensuring that AI tools are used for the welfare of the animals.
- The future holds advanced AI tools like 3D-printed prosthetics, AI-enhanced diagnostics, and collaborations between AI researchers and veterinarians.
- AI advancements will lead to further integration of predictive analytics, automated health monitoring, and new tools for precision dairy farming.

To learn more …

Online Courses

Harvard School of Engineering and Applied Sciences - CS249r: Tiny Machine Learning
Professional Certificate in Tiny Machine Learning (TinyML) – edX/Harvard
Introduction to Embedded Machine Learning - Coursera/Edge Impulse
Computer Vision with Embedded Machine Learning - Coursera/Edge Impulse
UNIFEI-IESTI01 TinyML: “Machine Learning for Embedding Devices”

Books

“Python for Data Analysis” by Wes McKinney
“Deep Learning with Python” by François Chollet - GitHub Notebooks
“TinyML” by Pete Warden and Daniel Situnayake
“TinyML Cookbook 2nd Edition” by Gian Marco Iodice
“Technical Strategy for AI Engineers, In the Era of Deep Learning” by Andrew Ng
“AI at the Edge” book by Daniel Situnayake and Jenny Plunkett
“XIAO: Big Power, Small Board” by Lei Feng and Marcelo Rovai
“MACHINE LEARNING SYSTEMS for TinyML” by a collaborative effort

Projects Repository

Edge Impulse Expert Network

On the TinyML4D website, You can find lots of educational materials on TinyML. They are all free and open-source for educational uses – we ask that if you use the material, please cite them! TinyML4D is an initiative to make TinyML education available to everyone globally.

Questions?

rovai@unifei.edu.br