# CS249r Fall 2023 Assignment 1: Setup Arduino Nicla Vision + Data Engineering for a Person Detection Transfer Learning Model

Part 1 Due: Monday Oct 2, 2023 at 11:59pm on Canvas Part 2 Due: Monday Oct 9, 2023 at 11:59pm on Canvas

## **Assignment Overview**

The purpose of this assignment is to get familiar with the Arduino Nicla Vision and the general workflow for designing and deploying a TinyML application. For this assignment, you can use open source data, web scraped data, synthetic data, or manually generated data. Our goal here is to get you to explore the effect of data quality on the performance of a model, as well as the quantity of data and other factors.

You need to collect 100 samples for the **"person"** class which is our target class and an additional 100 samples for the **"background/no-person**" class. This will constitute the negative class in a binary classification setting.

This assignment is split into two parts:

- In Part 1:
  - **What**: You will **set up the Arduino Nicla Vision** and submit an image of yourself taken with the Nicla. Your image will be included in the testing dataset for Part 2.
  - **Why:** This is to test your board and make sure it is working and get you familiar with some open source toolchains for working with embedded systems.
  - $\circ$   $\,$  How: You will use the Edge Impulse pipeline to bootstrap your system.
- In Part 2:
  - What:You will gather training data for a person detection model.
  - **Why:** The goal of Part 2 is to explore good data engineering techniques for training a TinyML model. We will provide you with information about the model but will hide the testing data which will be used to assess the quality of your training data. The testing data contains 500 images, including photos submitted by students taking this course.
  - **How:** You will use Edge Impulse or OpenMV tools to get your training data.

Please ask questions (and if you can, answer your fellow classmates' questions) in the #assignment1 channel of the course's Slack workspace. We also have office hours at different times in the week listed in the <u>Course Syllabus</u>.

## Deliverables

Part 1: Submit a single zip file containingi) an .mp4 video of length at most 2 mins showing your Nicla connected to Edge Impulse and capable of sampling images of your face

ii) a 96x96 image of yourself for the assignment testing data

Part 2: Submit a single zip file containing

i) a **training dataset** made up of a total of 200 images that belong to a **"person"** and **"non-person**" category. This dataset will be used to train MobileNetV2 96x96 0.1 transfer learning model in Edge Impulse.

ii) a PDF write up of your dataset design decisions.

For each submission, create a zip file containing the requested documents and upload it to canvas. You can collaborate with others on setting up your kit and using the Edge Impulse platform, **but you will need to collect your own dataset and complete each of the above deliverables individually**.

## **Instructions for Part 1**

1. Connect your Arduino Nicla Vision to your laptop using the micro-USB to USB-A cable that came with your kit.



2. Install all the <u>dependencies</u> required to set up your Arduino Nicla Vision device on your machine.

After installing the dependencies listed on the website linked above, you should be able to connect your Arduino Nicla to the Edge Impulse portal via command line.

- 3. Follow the steps below to take your picture:
  - a. Create a project in Edge Impulse.
  - b. In your terminal, run "edge-impulse-daemon". When prompted, provide your login details and select the project which you want to connect the device to from the dropdown menu (if you have multiple projects in Edge Impulse)

```
[MacBook-Pro-5:∼ jessicaquaye$ edge-impulse-daemon --clean
Edge Impulse serial daemon v1.16.0
  What is your user name or e-mail address (edgeimpulse.com)? jquaye@g.harvard.edu
What is your password? [hidden]
Endpoints:
    Websocket: wss://remote-mgmt.edgeimpulse.com
                 https://studio.edgeimpulse.com
     API:
    Ingestion: https://ingestion.edgeimpulse.com
 [SER] Connecting to /dev/tty.usbmodem2101
 SER] Serial is connected, trying to read config...
 [SER] Clearing configuration
 [SER] Clearing configuration OK
 [SER] Retrieved configuration
 [SER] Device is running AT command version 1.7.0
  To which project do you want to connect this device? (Use arrow keys)
  Jessica / tinyml-kws
Jessica / tinyml-kws-v2
  Jessica / mise
  Jessica / cs-249r
```

4. Name the device so that you can reference it in the Edge Impulse portal:



5. Once your device is named and connected to the project, you can start collecting data! Click on Data Acquisition from the left-hand navigation panel on the Edge Impulse website to view the real-time camera live stream in your web browser. In the "Collect Data" section, select your device (it will be listed as the name that you gave it in Step 4) and choose "Camera (96x96)" as the sensor:

Collect data	:
Device ③	
nicla	~
Label	Camera feed
Label name	6
Sensor	
Camera (96x96)	

6. Download your picture from the Edge Impulse platform for submission.

### Part 1 Deliverables

Submit a .zip file on Canvas containing

- 1. A video recording of your Nicla connected to Edge Impulse and streaming from the camera.
- 2. A 96x96 RGB image of yourself to be used in the testing dataset.

## **Instructions for Part 2**

If you are comfortable using the Edge Impulse interface, you do not have to complete this step. However, the data collection process can be simplified by using the OpenMV IDE:

- Download the <u>OpenMV IDE</u>.
   NOTE: Macbook Users might have to download <u>this Github release</u> to avoid the GUI from crashing when attempting to connect the Nicla to the OpenMV IDE.
- 2. Connect the Arduino Nicla to OpenMV IDE.
- 3. Once your Nicla is connected, navigate to File > Examples > HelloWorld > helloworld.py and click on the green button at the bottom left to Run the program. You should be able to see a real-time camera live stream in the Frame Buffer on the right.



4. Use the code on the <u>Nicla website</u> or the <u>Edge Impulse</u> website to take pictures which will be stored locally. You can then upload these pictures into the Edge Impulse platform for training.

## Using the Nicla Vision Camera

You can easily access the camera on the Nicla Vision through OpenMV IDE. Below is a short script that will set up the camera and take an image. The board will blink its LED to indicate when it will take the picture. The image can be seen in the frame buffer while the script is running.



#### Training and Assessing your Model Performance on Edge Impulse

In this section, you will create a training dataset for the MobileNetV2 96x96 0.1 transfer learning model. To help with the process of creating the Machine Learning model, we will be using EdgeImpulse Studio. This contains a suite of tools used to create the model and generate a C++ library that can be used for inference and image classification.

#### Step 1: Edge Impulse Project Setup

- 1. Once you have created an account, log in and create a new project. This will open the project setup dialogue where you can set the project name.
- 2. The project will open in the dashboard view with a menu in the left-hand margin for setting up the workflow. Scroll down and edit the Project Info on the bottom right as follows:
  - a. Labeling Method: One label per data item
  - b. Target Device: Arduino Nicla Vision

💳 EDGE IMPULSE	- <b>-</b>	E	սիի	Summary	
_					ES CONNECTED
Dashboard	Motion: Gesture recognition	Images: Object detection	Audio: Audio classification	0	
Devices				DATA	COLLECTED
Data acquisition	About this project		Add README		
✤ Impulse design					
Create impulse	Download block output			Project info	
Ø EON Tuner				Project ID	281614
🔀 Retrain model		No downloads available yet		Labeling method	✓ One label per data item
Live classification					Bounding boxes (object detection)
Model testing	Performance settings		(?)	Target device	Arduino Nicla Visic \vee
P Versioning					
📦 Deployment	Use GPU for training				
	Enterprise performance ⑦				
★ Try Enterprise Free	Job limit in minutes	20			
Get access to high job limits and training on GPUs.	Train job memory (MB) ③	8192			
Start free trial	DSP file size limit (MB)	4096			0

#### Step 2: Data Acquisition

- 1. Navigate to the project that you created for this assignment in EdgeImpulse.
- 2. Click the *Data Acquisition* activity from the left-hand navigation panel.
- 3. Click *Choose Files* and select the images from the relevant folder on your laptop.
- 4. Give the images a label ("person" or "non-person", depending on which class of images you're uploading) and begin the upload.
- 5. Repeat the process for each of your data classes.
- 6. You should now have all your data loaded into EdgeImpulse ready for processing. You can either upload your training dataset and your testing dataset separately or ask for EdgeImpulse to automatically split the data that you upload into 2 parts.

🔁 EDGE IMPULSE	Dataset Data expl	orer Data source:	s   CSV Wizard			
Dashboard	DATA COLLECTED 680 items	0	TRAIN / TEST SPLIT 41% / 59% ▲		Collect data	
<ul> <li>Devices</li> <li>Data acquisition</li> </ul>	Dataset		± a	Ð	Connect a device to start building your dataset.	
✓ Impulse design	Training (280) Tes	t (400)	T 🖸	0	raw data 202455	
<ul> <li>Create impulse</li> </ul>	SAMPLE NAME	LABEL	ADDED			
Image	202579	person	Sep 12 2023, 17:53:	:		
<ul> <li>Transfer learning</li> </ul>	202585	person	Sep 12 2023, 17:53:	÷		
🧭 EON Tuner	202586	person	Sep 12 2023, 17:53:			
Retrain model	202455	person	Sep 12 2023, 17:53:	:		
Model testing	202583	person	Sep 12 2023, 17:53:	1		
Versioning	202453	person	Sep 12 2023, 17:53:	:	Metadata	? +
*	202454	person	Sep 12 2023, 17:53:	÷	No metadata.	
Try Enterprise Free Get access to high job limits	202450	person	Sep 12 2023, 17:53:	:		
and training on GPUs.	202451	person	Sep 12 2023, 17:53:	1		
Start free trial	202582	person	Sep 12 2023, 17:53:	:		0

#### Step 3: Impulse Design

This step sets up all the parameters for building the ML model. EdgeImpulse guides you through this phase and indicates how accurate your model is likely to be after training.

The first step to designing an impulse is to **Create Impulse**.

- Click the Create Impulse activity in the menu bar.
- Add an image-processing block (*Image*) and set the image size to 96x96.
- Add a Transfer Learning block.
- There should be two *Output Features* (person, not\_person) automatically populated if your training data was labeled correctly.

💳 EDGE IMPULSE				
		and processing to extract features, and then us	as a learning block to classify new data	
	An impulse takes raw data, uses sig	gnal processing to extract features, and then us	es a learning block to classify new data.	
Dashboard				
Devices				
Data acquisition				
✤ Impulse design	Image data	Image 🦻	Transfer Learning	Output features
Create impulse		Name	Name	2 (non person person)
Create impuse	image	Image	Name	2 (non_person, person)
Image	Image width Image beight	inage	Transfer learning	
<ul> <li>Transfer learning</li> </ul>	96 96	Input axes (1)	Input features	Save Impulse
Ø EON Tuner		🖌 image	✓ Image	
	Resize mode			
🔀 Retrain model	Fit shortes		Output features	
Live classification	P 着		2 (non_person, person)	
Model testing	Eor optimal accuracy with			
<b>9</b>	transfer learning blocks, use a			
P versioning	96x96 or 160x160 image size.			
Try Enterprise Free		(	()	
Get access to high job limits		•	Д	
and training on GPUs.		,	-	
Start free trial		Add a processing block	Add a learning block	0

• Click Save Impulse.

The next stage, **Image**, sets the model parameters and generates a feature set from the training data images.

- Click the *Image* Item from the menu bar.
- Set the *Color Depth* to RGB.
- Click Save parameters.
- Then, select *Generate Features*. This will trigger a remote job in the EdgeImpulse Cloud to generate the model features.

At this stage, you can use the *Feature Explorer* to gauge whether the features are sufficiently distinct to give good results. The goal is to look for distinct clusters of features with as little overlap as possible.

🗮 EDGE IMPULSE	#1 ▼ Click to set a des Parameters Generate fe	scription for this version	
📮 Dashboard			
Devices	Training set		Feature explorer
Data acquisition	Data in training set	280 items	non_person
✤ Impulse design	Classes	2 (non_person, person)	• person
<ul> <li>Create impulse</li> </ul>			
Image		Generate featur	es
<ul> <li>Transfer learning</li> </ul>			
Ø EON Tuner			
🔀 Retrain model			
Live classification			
Model testing			
2 Versioning			
Try Enterprise Free			On during and annual @
Get access to high job limits			On-device performance @
and training on GPUs.			PROCESSING TIME PEAK RAM USAGE
Start free trial			1 ms. 4 KB

The final stage is **Transfer Learning**.

- Click *Transfer Learning* from the menu bar.
- Select MobileNetV2 96x96 0.35 model and set your parameters. The default parameters work well for our use case but you are welcome to experiment with them (number of training cycles, learning rate, number of neurons) and change them as you deem appropriate. Please note any changes to these parameters in your writeup.
- Click *Start Training*. The training can take several minutes depending on the amount of data and the complexity of the features.

When the training process is completed, Edge Impulse will provide statistics about how effective the model is and will show clusters of the training results in the Data Explorer view. You can hover over each point in the cluster and click on the point to see which image generated what result.



#### Step 4: Assessing model performance

The final step in the model creation process is to test the model with a labeled test data set to see how well it performs. In this assignment, we will assess each student's model with a hidden test dataset. The goal of this assignment is for you to creatively engineer a training dataset for the **MobileNetV2 96x96 0.35 model** to perform person detection in an image. That is, the model should be able to accurately classify whether an image contains a person or not. The test data will contain 96x96 images collected with an Arduino Nicla but with creative modifications (this assignment is supposed to encourage you to think about all the ways that an image can be slightly altered but still maintain the important information that a person is present or not). Additionally, we want you to think about the quality and diversity of the data that you collect. (Hint: It may not be a good idea for your training dataset to only contain images of your face  $\bigcirc$ )

You can create your own testing dataset to assess your model's performance, but we will use the training dataset that you submit to train the **MobileNetV2 96x96 0.35 model** and test it with our hidden dataset.

To assess the model's performance on your test dataset:

• If you asked for an automatic split of your dataset into training and testing, skip to the next step.

- Otherwise, upload your testing dataset by navigating back to the *Data Acquisition* using the left-hand navigation panel. There, you can upload data for your testing dataset.
- Select *Model Testing* from the left-hand navigation panel.
- Click Classify All.

Another remote job will start and run your test data against the trained model, showing you the results upon completion.

### Part 2 Deliverables

Submit a .zip folder on Canvas containing:

- 1. Two folders containing your training data separated into:
  - Person This folder will contain all images which should be classified as having a person present.
  - Not\_Person This folder will contain all images which should be classified as **not** having a person present.
- 2. A 2-page .pdf write up of your dataset design decisions. Record the accuracy of your validation set in your writeup, along with how many files were in your validation set. Also, in your write-up, summarize the key steps you took to build and evaluate your custom dataset. Discuss the iterations that you went through to improve the model's ability to detect the presence of a person. For example, you can touch upon some of the items below:
  - How did you curate the dataset? Did you supplement the images that you captured on the Nicla Vision with images from other sources?
  - How many images per class did you collect? Did increasing or decreasing the number of images per class help to increase the accuracy of your model?
  - What do you think about the quality of the dataset you collected? Did you use augmentation, any kind of filtering, or corruption of your training dataset to test robustness?
  - Did you find any corner cases in your dataset that cause an issue with accuracy? What did you do to mitigate it? Did fixing those help in your final accuracy?
  - Did you find any way to fool the model (false-positives)? What do you think was the cause? What improvements do you think can be made to fix this issue?
  - Was the live view tool in OpenMV or EdgeImpulse useful in framing yourself for inferencing? Did it help you combat lighting issues?

The training data submitted for Part 2 of this assignment will be used to train a MobileNetV2 96x96 0.35 model and will be assessed with our hidden test dataset.