



Cat vs Dog Classifier

An AI Image Classification System

Muhammad Usama · Candidate No: 2345

LRN International GCSE Artificial Intelligence (7923) — Paper 2



The Problem

- The system decides whether a photo shows a cat or a dog.
- Humans do this easily, but writing fixed rules for it is very hard.
- So the model learns from example images instead of following rules.

Tool: Edge Impulse


- A beginner-friendly platform for building image models, and it lets me upload my own dataset of tens of thousands of images.
- It runs in the browser and needs no coding to train the model.
- It exports a model I can host on my own web page.



The Data

- A dataset of cat and dog images (source and licence noted in the report).
- Two labels: cat and dog.
- Data was cleaned and split 80/20 into training data and separate test data.

How It Works



The image is shrunk to 100×100, the model reads every pixel, and returns the more likely label.



Training and Testing

- The model was trained on the training images.
- It was then tested on separate images it had never seen.
- Test accuracy: 80%
- Testing on unseen data proves it learned patterns, not just memorised.

Improving the Model



BEFORE

- Failed or low confidence on hard images
- e.g. unusual breeds, dark or blurry photos

AFTER

- Added more varied training examples
- Retrained the model
- Higher confidence on the same hard images

Test → find failures → add examples → retrain → compare.



Ethics and Users

- Privacy: it runs locally in the browser, so no images are sent to a server.
- Limitation: with only two labels, it forces every image into cat or dog.
- Bias: results can be unfair if the dataset lacks variety of breeds.



Reflection

- Hardest part: finding diverse, good-quality datasets.
- What I learned: the variety of the training data matters more than the amount - it only became reliable once I added different breeds, backgrounds and lighting.
- Next improvement: add a 'neither' label and more varied data.