



Fork vs Spoon Classifier

An AI image classifier for returned canteen cutlery

Muhammad Usama · Candidate No: 2345

LRN International GCSE Artificial Intelligence (7923) — Paper 2



The Problem

- The system looks at a photo of a returned utensil and decides whether it is a fork, a spoon, or neither.
- The setting is a school canteen cutlery return point, where staff currently sort every item by hand.
- Fixed rules are hard because cutlery varies in size, finish and angle, so the model learns from example photos instead.



Tool: Edge Impulse

- A browser-based platform for building image models with no coding.
- I chose it because I can upload my own dataset and retrain quickly between versions.
- It exports to WebAssembly, which I deployed as a static web page.



The Data

- Dataset: my own photographs, taken by me - no third-party licence needed, and nobody is identifiable in them.
- Three labels: fork, spoon, and neither.
- Photographed on varied backgrounds and lighting, then split 80/20 into training and test data.

How It Works



The image is resized to 96x96, the model reads all 9,216 pixels, and returns the most likely of the three labels.

Training and Testing



- The model was trained on the training images.
- It was then tested on separate images it had never seen.
- Test accuracy: 82%
- Testing on unseen data proves it learned the shapes, not just memorised my photos.

Improving the Model



BEFORE

- Version 1 failed on unfamiliar backgrounds
- e.g. cutlery on wood, or held at an angle

AFTER

- Added more varied training photos
- Retrained as version 2
- Higher confidence on the same hard images

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



Ethics and Users

- Privacy: inference runs locally in the browser, so no photo of anyone's tray is uploaded.
- Limitation: three labels only - a knife falls into 'neither' rather than being identified.
- Bias: trained on the cutlery I own, so a canteen with different sets would score worse until its own photos are added.



Reflection

- Hardest part: building a 'neither' class broad enough to catch everything that is not a fork or a spoon.
- What I learned: keeping colour rather than greyscale mattered, because bright metal against a darker tray is most of the signal.
- Next improvement: add a knife label and photograph a real return point rather than a desk.