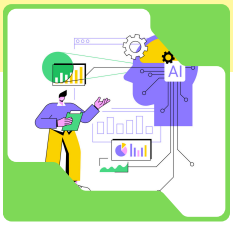

From Data to the Machine's Brain

Training, Interpreting and Imagining Social Uses of AI

AI exploration phase

P3 - Experimenting with Language Models

AI and Exploration Topics. What are we exploring?

Supervised classification: training a model to sort inputs into classes.
Training data and test data: what the model learns on versus what it is checked on.
Generalisation: behaving correctly on new, unseen cases.
Edge cases: tricky inputs that reveal the model's limits.
Model error: when and why the model gets it wrong.
Perceptron: a simple neuron combining weighted inputs against a threshold.

Materials Needed and to Create. What must be prepared before the activity?

Computer or laptop with internet (and webcam if training live), 1 per group
Training and test recording sheet (A4, double-sided), 1 per group
Post-its in different colours, 1 set per group
50 ml cups (4) and 250 ml graduated beakers (3), per group
Coloured indicator post-its, 1 pack per class
Water, 0.5 to 1 L per group
Teachable Machine (online)

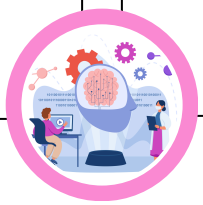

Exploration Steps and Description. What do participants actually do?

Young people train small classification models, analyse when they fail, build a physical perceptron to see how a neuron combines inputs, and imagine social uses of AI linked to the SDGs.

Activity 1 - Reactivate the machine (60 min)

Each group trains a **classifier** on a **visible feature** (glasses / no glasses, open hand / closed fist) and records training and test results. It identifies what the model gets right, what it confuses and where it fails, then introduces **edge cases** (poor light, different angle or distance, unseen accessory) to see when it breaks.

Activity 2 - Bias, error (15 min)

A short ethical discussion links **reliability to real cases of biased AI** (image recognition trained on too narrow a group, CV-screening systems), handled carefully so no young person is exposed.

Co-funded by the European Union

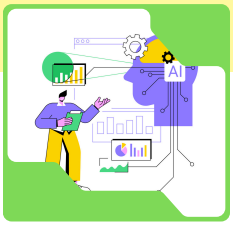

From Data to the Machine's Brain

Training, Interpreting and Imagining Social Uses of AI

Exploration Steps and Description. What do participants actually do?

Activity 3 - Explore a perceptron (75 min)

Groups build a **water-cup perceptron** where inputs are small cups (on or off), weights are the water poured, the bias is a threshold line on the big beaker, and the output is whether the level passes it. They tune the water to classify correctly, then face a diagonal challenge a single perceptron cannot solve, revealing **the need for a hidden layer and a neural network**.

Activity 4 - Brainstorm (20 min)

Each group picks a **social problem** (sorting waste, wasted energy, watering a garden, a full bin) and sketches what the model should recognise and what risk or limit it could have.

Expected Output. What do participants create, produce or gain from the exploration?

Young people train and stress-test a real model, identify what affects its performance, manipulate a physical perceptron, and propose a socially useful AI task linked to an SDG, with its data needs and one limit. A closing card collects one idea, one piece of evidence and one question.

The Hook and the Playfulness. What makes this fun/exciting for young people?

Teaching a machine and trying to fool it with edge cases, operating a neuron with water, and imagining useful local applications as a small team of AI designers with a social mission.

Success indicators. How do we know it worked? What shows participants learned?

- They train, test, record and analyse the model rather than just using it, talking about data, errors, edge cases and thresholds.
- They relate errors to light, angle, background or lack of variety in the data.
- They explain what each element of the perceptron represents and that a response depends on a combination of inputs.
- Each group proposes one SDG-linked application with at least one data need and one model limit.

Co-funded by the European Union

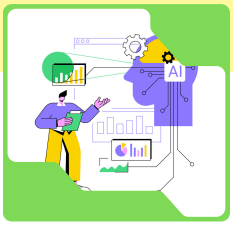

Appendix - From Data to the Machine's Brain

Appendix 1 | Perceptron solution: Training 1: the single-neuron water-cup diagram

Content of the appendix

Solution diagram for the first perceptron training (single neuron)

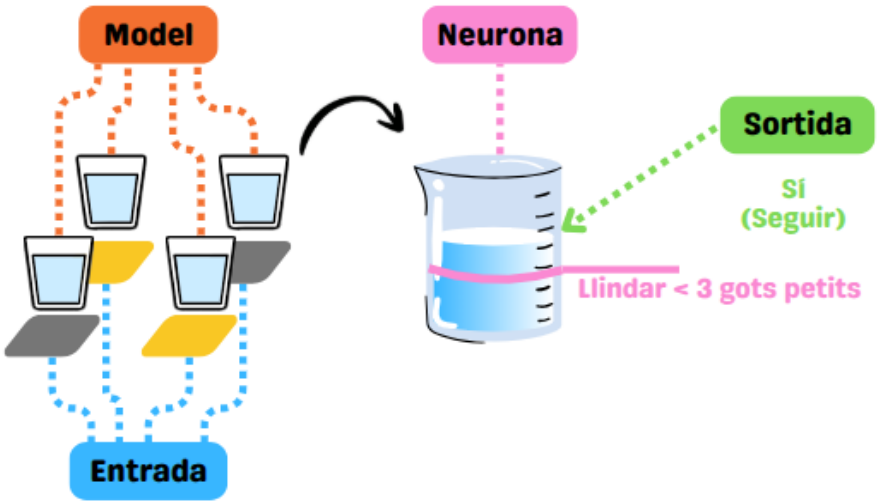

Labels on the diagram

- Model
- Input (the small cups)
- Neuron / Output: Yes (act)
- Threshold: less than 3 small cups

Explanation: For the big beaker to activate (water above the threshold), the small cups must be filled. Water is poured from as many small cups as there are lit squares. So with 3 or more lit squares, the big beaker holds the volume of 3 or more small cups, which is above the threshold set at the start (120 ml, less than 3 small cups), so the output neuron activates.

Co-funded by the European Union

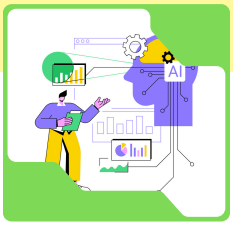

Appendix - From Data to the Machine's Brain

Appendix 2 | Perceptron solution: Final challenge: the two-hidden-neuron diagram

Content of the appendix

Solution diagram for the final challenge (a hidden layer with two neurons)

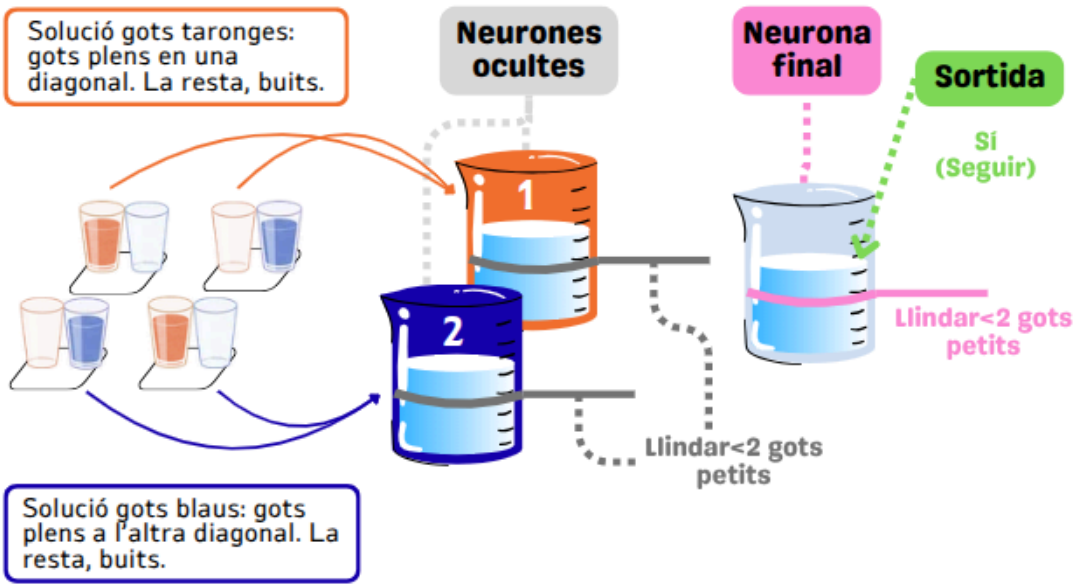

Labels on the diagram

- Orange small cups feed hidden neuron 1; blue small cups feed hidden neuron 2.
- Hidden neurons (1 and 2); final neuron; Output: Yes (act)
- Threshold (each neuron): less than 2 small cups

Explanation: Orange cups: decide how much water to put in the orange small cups, then pour it into hidden neuron 1. Blue cups: decide how much water to put in the blue small cups, then pour it into hidden neuron 2. Solution: orange cups full on one diagonal, the rest empty; blue cups full on the other diagonal, the rest empty. The cup colours can be marked with matching post-its.

Co-funded by the European Union