CNN: Find a Line with One Shared Ruler

Plan 35-45 minutes: predict on paper, then check on one projected computer. The local-model lesson needs a computer to run its download. Split advanced calculations across two sessions.

Learning goal

Calculate one 1 × 3 window and distinguish feature-map locations from a global maximum.

Before starting

Complete layered networks and understand dot products and thresholds.

Materials

The new worksheet, a 5 × 5 grid and a paper strip with three windows.

Class sequence (5 / 10 / 10 / 10 min, plus 5-10 min discussion)

1. Predict: place filter [1,1,1] over 111 and then 010. What are the dot products?

2. Try: slide one column at a time across a 5 × 5 image. Three valid windows per row give 15 positions, stored in a 5 × 3 feature map.

3. Observe: take the global maximum and test >2.5 for three consecutive black pixels. Move the line and compare the location map with its maximum.

4. Challenge: rotate to vertical, break the line or move it off the edge. Shared weights save parameters without guaranteeing all orientations or deformations.

CNN: Find a Line with One Shared Ruler

Worked example

A row 01110 produces window scores 2,3,2. The filter has three shared weights rather than 45 separate weights across 15 positions.

Check understanding: ask before showing an answer

If you know only that the maximum is three, can you locate the row? Does an all-black image also pass?

Expected answers and worksheet check

No: global max pooling discards location. An all-black image also passes. Worksheet rows: 01110→2,3,2; 00100→1,1,1; 01010→1,2,1. Maxima are 3,1,2; only the first meets this detector's criterion.

Keep the claim within its limits

This hand-set sliding dot product is cross-correlation as commonly used in deep learning, not a trained CNN. It uses no padding and stride one; boundaries, rotation, stride and later layers affect translation behavior.

Adapt the lesson

For support, work through just the example, read prompts aloud and use counters or paper grids. For extension, change one condition at a time, write a prediction and test it. When answers differ, inspect rules, encoding and data before judging a child's answer.

At home

Ask a parent to predict one row's window scores, then verify with the paper filter.

Continue learning

Neural Networks with Layers / After Full Marks: Did the Model Really Learn? / Transformer: Calculate One Attention Step

Teacher reference (links also on the lesson page)

Dive into Deep Learning: Convolutions for Images: https://d2l.ai/chapter_convolutional-neural-networks/conv-layer.html