CHAPTER 03 / LOOK INSIDE THE MODEL

How a network learns.

Small calculations. Connected layers. Adjusted weights.

Signals forward. Learning adjusts the connections.

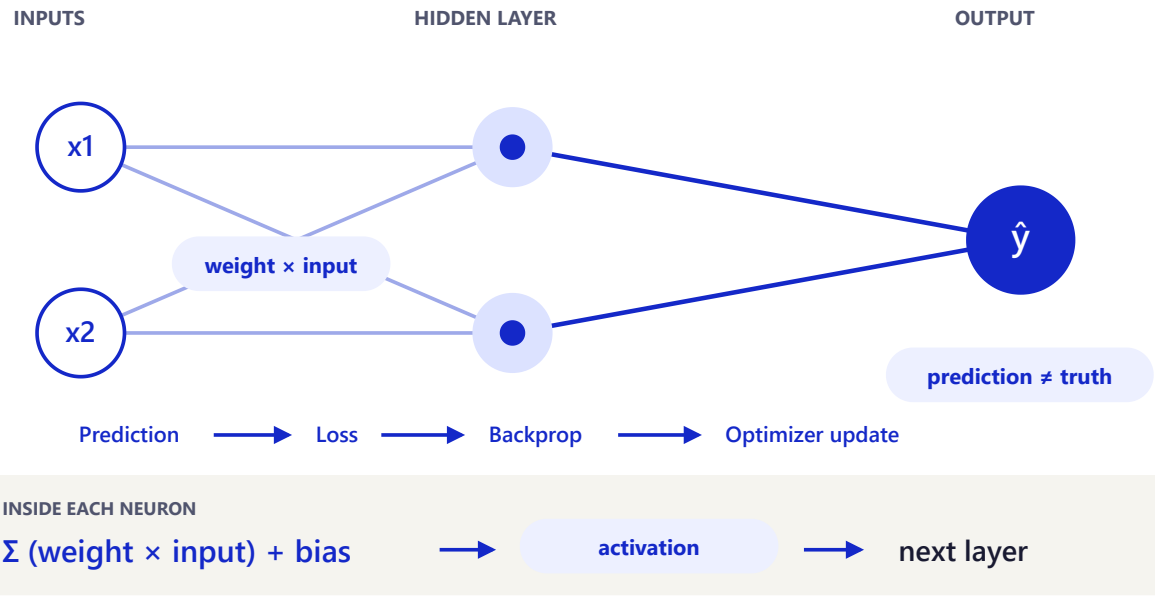

Handpicked 2-2-1 feed-forward example, matching the interactive topology. Adjusting a weight manually is not a training run.

A neuron

Combine weighted inputs and a bias, then apply an activation. It is a calculation, not a miniature brain.

A forward pass

Layers transform signals into an output. Learned representations depend on the architecture and training task.

Training

Backprop computes gradients; the optimizer updates parameters. Learning rate scales the update. Evaluate on held-out data.

↗ **Lower training loss does not guarantee better performance on new data.**

PUT IT TO WORK

For a defect detector, keep some labeled photos out of training. Test on those, then inspect real production mistakes and data drift.

CLOSE THE PAGE. EXPLAIN IT.

What changes during training, and what evidence would show that the network generalizes?