Neural-network gradients and spatial filters — review sheet

Use this after attempting the worked examples and practice. Explain each claim in words before calculating.

Key ideas

- Affine maps combine features; nonlinearities make compositions richer.
- Loss specifies the fitting objective; gradients guide local parameter changes.
- Batch and spatial axes have different meanings.
- A shared image filter is a repeated local weighted sum.
- Fixed filters and a trained CNN are different claims.
- Choose fitting procedures before the final test evaluation.

Common errors and repairs

Any neural network trained on flattened images is a CNN. The sample trained network is fully connected. Its separate fixed-filter demonstration has convolutional structure but no learned filters.

A falling training loss establishes held-out performance. Training loss and held-out evaluation answer different questions.

A negative filter response is invalid. Signed directional filters encode which way intensity changes.

Perfect synthetic accuracy proves robust real-world recognition. It only demonstrates success on the particular held-out generator distribution.

Retrieve without your notes

- What is an epoch?
- Why keep training mean and scale for inference?
- What output shape results from an 8×8 image and a 3×3 valid stride 1 filter?

Readiness check

- Trace the shapes in a batched affine–nonlinearity–output calculation.
- Explain loss, gradient update, epoch and mini-batch in the actual runnable model.
- Interpret signed responses of two fixed 3×3 image filters and predict output shape.
- Compare a trained network with held-out baselines without using test data for fitting.
- State why synthetic test accuracy does not establish general image or clinical performance.