

Convolutional Neural Networks Lab – Lab 1

A modern LeNet-5, built and verified from the inside

CSSE 313 Artificial Intelligence · Lab Manual · Lab 1 (35 points)

Name: _____ Date: _____

How This Lab Works

In this lab you implement a modern version of Yann LeCun's LeNet-5 (LeCun, Bottou, Bengio & Haffner, *Gradient-Based Learning Applied to Document Recognition*, 1998, section II.B) and train it on MNIST. As in the feedforward lab, the premise is that **the learning is not in the typing**. You may use AI to write, debug and refactor your code. What this manual grades is the reasoning around the code.

There are no points for code. Working code is still required, though, because every graded item asks you to predict, run, and explain something your own network does. Without a working network you cannot produce that evidence.

AI use policy

- **Permitted:** using AI to write and debug the model, and to explain concepts to you.
- **Not permitted for credit:** AI-written predictions, derivations, hand traces, mechanism explanations or reconciliations.

The Network You Will Build

The input is a 32×32 image: an MNIST digit (28×28, pixel values divided by 255) padded with two rows or columns of zeros on every side. All convolutions use stride 1 and no padding. All pooling is 2×2 with stride 2.

Layer	What it does in this lab
C1	6 kernels of size 5×5 over the input, plus one bias per map, then ReLU
S2	Max pooling, 2×2, stride 2. No trainable parameters, no bias, no activation
C3	16 kernels of size 5×5, each connected to all 6 S2 maps (full connectivity), plus one bias per map, then ReLU
S4	Max pooling, 2×2, stride 2

Layer	What it does in this lab
C5	120 kernels of size 5×5, each connected to all 16 S4 maps, plus one bias per map, then ReLU
F6	Fully connected, 84 units, then ReLU
Output	Fully connected, 10 units, then softmax. Loss: cross-entropy, $-\ln p(\text{correct label})$

Training: one example at a time (stochastic gradient descent), learning rate 0.01, no momentum, no regularization. All weights and biases start uniformly in $[-0.1, 0.1)$, drawn from `java.util.Random(seed)` so that runs are reproducible.

What changes compared with LeCun (1998)

	LeCun et al. 1998	This lab
Activation	Scaled tanh, $f(a) = 1.7159 \tanh(2a/3)$	ReLU
Subsampling (S2, S4)	Average of 2×2, times a trainable coefficient, plus a trainable bias, then sigmoid	Max of 2×2; no parameters
C3 connectivity	Sparse: each C3 map reads 3, 4 or 6 of the S2 maps (Table I)	Full: each C3 map reads all 6 S2 maps
Output layer	84 → 10 Euclidean RBF units compared with fixed 7×12 bitmaps	84 → 10 fully connected + softmax
Loss	MAP criterion on RBF outputs	Cross-entropy

Lab 1 — Architecture and Forward Pass by Hand (35 points)

Part 1A — The architecture on paper

20 points

- (15) Complete the table. *Connections* counts the multiply-adds of one forward pass, counting each bias as one connection. *Receptive field* is the size of the input region (in padded-input pixels) that can influence one unit of the layer.

Layer	Output shape	Trainable parameters	Connections	Receptive field
Input	1 @ 32×32	—	—	1×1
C1				
S2				
C3 (full)				
S4				
C5				
F6				—
Output				—
Total				

- (3) In C1, how many times is each individual weight used during one forward pass?
- (2) Which is the first layer whose units can see an entire MNIST digit? (Digits sit in a centered 20×20 box of the 28×28 image.)

Part 1B — A forward pass by hand

15 points

One input channel, one 3×3 kernel, bias $b = -2$. Compute a valid cross-correlation $Z = K \circledast X + b$ (4×4), $A = \text{ReLU}(Z)$, then $P = 2 \times 2$ max-pool of A (2×2).

$X =$	2 0 1 0 1 0	$K =$	1 0 -1	$b =$	-2
	3 1 2 1 2 3		2 1 0		
	1 3 2 1 2 3		0 -1 1		
	1 1 0 0 1 3				
	3 3 1 0 1 1				
	1 0 1 2 2 3				

- (8) Compute Z , A and P by hand. Show the full arithmetic for $Z[0][0]$ and $Z[3][1]$. For every pooling window, mark which entry of A wins.

2. (3) One pooling window contains a tie and one window is entirely zero after ReLU. Identify both. For the tie, which entry wins under the lab's convention, and does the convention change P?