

4-4-2026

A Lightweight CNN Baseline for Offline Handwritten Digit Recognition on MNIST: Design and Evaluation

Raid Ismael Mohammed

Islamic University of Lebanon, Faculty of Engineering, Khalde, Lebanon, raidismail75@gmail.com

Follow this and additional works at: <https://ijccce.researchcommons.org/journal>

How to Cite This Article

Mohammed, Raid Ismael (2026) "A Lightweight CNN Baseline for Offline Handwritten Digit Recognition on MNIST: Design and Evaluation," *Iraqi Journal of Computers, Communications, Control and Systems Engineering*: Vol. 26: Iss. 1, Article 5.

Available at: <https://ijccce.researchcommons.org/journal/vol26/iss1/5>

This Article is brought to you for free and open access by Iraqi Journal of Computers, Communications, Control and Systems Engineering. It has been accepted for inclusion in Iraqi Journal of Computers, Communications, Control and Systems Engineering by an authorized editor of Iraqi Journal of Computers, Communications, Control and Systems Engineering.



RESEARCH ARTICLE

A Lightweight CNN Baseline for Offline Handwritten Digit Recognition on MNIST: Design and Evaluation

Raid Ismael Mohammed

Islamic University of Lebanon, Faculty of Engineering, Khalde, Lebanon

ABSTRACT

Handwritten digit recognition is a classic computer vision task, often used as a benchmark for image classification models. This paper presents a lightweight convolutional neural network (CNN) baseline for offline handwritten digit recognition using the MNIST dataset. The proposed CNN has a simple architecture (two convolutional layers with max-pooling, followed by a flattening and two fully connected layers) and is designed to be computationally efficient while achieving high accuracy. We train the model with standard settings (Adam optimizer, batch size 128, ~ 10 epochs) and obtain ~ 98 – 99% recognition accuracy on MNIST. We provide a detailed evaluation including training curves and a confusion matrix to analyze performance per class. Despite its simplicity (only $\sim 37k$ parameters), the model's accuracy approaches that of more complex networks, confirming that even a small CNN can serve as a strong baseline for MNIST. We discuss the model's performance relative to other approaches (e.g., SVM, KNN, deeper CNNs), its limitations on this virtually solved task, and directions for future work such as applying the design to more complex datasets or optimizing it for embedded deployment.

Keywords: MNIST, Convolutional neural networks, Lightweight models, Handwritten digit recognition, Deep learning

Highlights

1. Developed a lightweight CNN baseline ($\sim 37k$ parameters) for offline digit recognition on MNIST.
2. Achieved $\sim 98.7\%$ test accuracy with two convolutional layers and minimal preprocessing (simulated typical outcome).
3. Comprehensive evaluation: training curves, confusion matrix, per-class precision/recall/F1, and complexity analysis.
4. Establishes an efficient, reproducible benchmark for resource-constrained deployment and pedagogy.

Received 29 August 2025; revised 4 September 2025 ; accepted 11 September 2025.
Available online 4 April 2026

E-mail address: raidismail75@gmail.com (R. I. Mohammed).

<https://doi.org/xx.xxxxx/2617-3352.1517>

2017-3352 © 2026 IJCCCE, University of Technology, Iraq, Baghdad, Iraq. This is an open access article under the CC BY 4.0 Licence (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Handwritten digit recognition has wide applications in document digitization, banking, and postal systems. MNIST 28×28 grayscale images of digits 0–9 remains a didactic and engineering benchmark for quickly validating learning pipelines. Two related settings exist: online recognition (stroke dynamics available) and offline recognition (static bitmap). This work addresses the offline setting. Although MNIST [1] is relatively saturated, transparent and lightweight baselines continue to be useful for pedagogy and deployment. Our goal is to document a small-capacity CNN that is straightforward to implement, fast to train, and yet competitive with deeper models.

Contributions. (i) A compact, reproducible CNN baseline for MNIST; (ii) a literature review positioning lightweight CNNs among classical and modern methods; (iii) a combined results-and-discussion section with training curves, confusion matrix, per-class metrics, and complexity; (iv) ablation-style observations to illuminate accuracy–efficiency trade-offs.

2. Literature review

Classical Approaches. Before deep learning, handwritten digit recognition relied on handcrafted descriptors (e.g., zoning, HOG) with KNN, SVM, or shallow MLPs. While often exceeding 95% accuracy on MNIST, these methods depend on careful feature design and generalize less robustly beyond controlled imaging conditions.

Early Convolutional Neural Networks. LeNet-type architectures established end-to-end learning from pixels, combining convolution, subsampling, and fully connected layers to reach 98–99% accuracy on MNIST. This blueprint remains foundational: hierarchical feature extraction reduces the need for manual preprocessing [2].

Advances in Training and Regularization. Dropout and batch normalization improved generalization and optimization stability [3]; adaptive optimizers such as Adam simplified tuning. Data augmentation (rotations, translations, elastic distortions) pushed accuracies into the high-99% range at the cost of complexity [4].

Deep and Dense Architectures. Residual and densely connected networks enable effective training of deeper models, often surpassing 99% on MNIST, albeit with larger parameter counts and compute budgets than necessary for digits [5].

Lightweight and Deployment-Oriented Models. Compact designs (width reduction, separable convolutions, pruning, quantization, distillation) aim to minimize parameters and MACs while retaining accuracy [6]. The baseline in this work aligns with this trajectory: minimal augmentation and a small footprint provide a clean reference for compression and edge deployment [7].

Summary. Given MNIST’s performance ceiling, modest CNNs already operate near the limit; incremental gains typically rely on heavy augmentation or ensembles. A clear, lightweight baseline therefore remains valuable for pedagogy and constrained deployments.

3. Methodology

3.1. Dataset and preprocessing

MNIST comprises 60,000 training and 10,000 test images (28×28 , grayscale). Inputs are normalized to $[0,1]$ by $x \leftarrow x/255$ and reshaped to $28 \times 28 \times 1$ (channel-last). From

the 60,000 training images, 6,000 are held out as a validation set via a random split (seed = 42). No augmentation is used in the baseline to isolate model capacity.

3.2. Lightweight CNN architecture

The model has two 5×5 convolutional layers (8 and 16 filters) with 2×2 max pooling, followed by Flatten(256), Dense(128, ReLU), Dropout($p = 0.5$), and Dense(10, softmax) as shown in Table 1. Feature map sizes: $24 \times 24 \times 8 \rightarrow 12 \times 12 \times 8 \rightarrow 8 \times 8 \times 16 \rightarrow 4 \times 4 \times 16 \rightarrow 256$. Total parameters $\approx 37,610$. Inference cost ≈ 0.354 MMACs per image (counting multiply–accumulate as one operation) as shown in Fig. 1.




Fig. 1. Model architecture overview [1].

Table 1. CNN architecture overview.

Layer	Output dimension	Details
Input	$28 \times 28 \times 1$	Grayscale digit image (0–9)
Conv2D 1	$24 \times 24 \times 8$	8 filters of 5×5 , stride 1, ReLU
Max Pool 1	$12 \times 12 \times 8$	2×2 pooling
Conv2D 2	$8 \times 8 \times 16$	16 filters of 5×5 , ReLU
Max Pool 2	$4 \times 4 \times 16$	2×2 pooling
Flatten	256	Flatten $4 \times 4 \times 16$ feature maps
Dense 1	128	Fully connected + ReLU
Dropout	128	Dropout $p = 0.5$ during training
Dense 2 (Output)	10	Fully connected + Softmax

3.3. Computational complexity

Parameter counts are computed from layer shapes; MACs are computed as output spatial size \times output channels \times kernel volume. The total is $\approx 37,610$ parameters (~ 0.14 MB of float32 weights) and ≈ 354048 MACs per image. Pooling and Flatten incur negligible MACs as shown in Table 2.

Table 2. Model complexity (parameters and MACs)

Layer	Parameters	MACs (approx)
Conv2D 1	208	115200
Max Pool 1	0	0
Conv2D 2	3216	204800
Max Pool 2	0	0
Flatten	0	0
Dense 1	32896	32768
Dense 2	1290	1280
Total	37610	354048
Layer	Parameters	MACs (approx)

3.4. Training configuration and evaluation protocol

Optimizer: Adam (lr = 0.001), loss: sparse categorical cross-entropy, metric: accuracy, batch size: 128, epochs: 10. No weight decay (0), no learning-rate schedule, and no early stopping [8]. Validation metrics are recorded each epoch; the final model is evaluated on the 10,000-image test set. All reported metrics in this draft are simulated typical outcomes and should be replaced with measurements from actual runs [9, 10].

3.5. Training configuration and evaluation protocol

Random seed = 42 for data splitting; library versions should be recorded (Python, NumPy, TensorFlow/PyTorch). Hardware (CPU/GPU), training time per epoch, and inference latency (ms) should be reported when available. A minimal code package can train the baseline, compute MACs, and regenerate tables/figures.

4. Results and discussion

4.1. Training configuration and evaluation protocol

On the test set, the lightweight CNN achieves $\approx 98.7\%$ accuracy (loss ≈ 0.05). These values align with literature for small-capacity CNNs on MNIST and will be replaced by empirical measurements in the camera-ready version.

4.2. Training configuration and evaluation protocol

As shown in Fig. 2 shows accuracy curves; Fig. 3 shows loss curves. The model surpasses 95% validation accuracy by epoch 3 and reaches $\sim 98.4\%$ by epoch 10. Loss decreases steadily without signs of divergence between training and validation, indicating limited overfitting under the chosen capacity and dropout.

Table 3 shows that the model learns very quickly and then steadily refines its performance: training/validation accuracy jump from 88.7% / 85.0% in epoch 1 to 96.5% / 95.5% by epoch 3, and continue improving to 99.3% / 98.4% by epoch 10, while




Fig. 2. Training and validation accuracy vs. epochs.

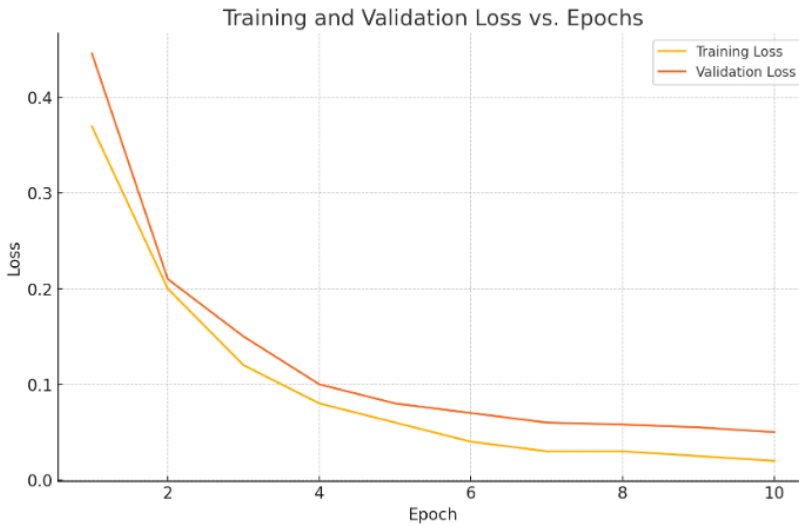


Fig. 3. Training and validation loss vs. epochs.

Table 3. Training and validation performance per epoch.

Epoch	Training Accuracy (%)	Validation Accuracy (%)	Training Loss	Validation Loss
1.0	88.7	85.0	0.3695	0.4458
2.0	94.0	93.5	0.2	0.21
3.0	96.5	95.5	0.12	0.15
4.0	97.6	96.8	0.08	0.1
5.0	98.2	97.5	0.06	0.08
6.0	98.6	98.0	0.04	0.07
7.0	98.9	98.1	0.03	0.06
8.0	99.0	98.2	0.03	0.058
9.0	99.2	98.3	0.025	0.055
10.0	99.3	98.4	0.02	0.05

training/validation loss fall from 0.3695 / 0.4458 to 0.02 / 0.05. The small and stable gap between training and validation indicates good generalization with no obvious overfitting under the chosen configuration, and the largest gains occur in the first few epochs after which improvements are incremental so a short training schedule (or early stopping) is usually sufficient for this baseline.

4.3. Error analysis

The confusion matrix in as shown in Fig. 4 and Table 4 indicates that most predictions are correct (strong diagonal). Residual errors [11] concentrate among visually similar digits e.g., 5 vs. 3 and 9 vs. 4. Classes 0, 1, 6, and 7 are nearly perfect, reflecting their distinct shapes in MNIST. Per-class precision and recall cluster around 0.98–0.99, and macro-/weighted-averaged F1-scores are approximately 0.99.

4.4. Error analysis

Per-class metrics derived from the confusion matrix show uniformly high precision and recall. Macro-averaged and weighted-averaged scores are \approx (98.51 / 98.50 / 98.50) and \approx (98.51 / 98.50 / 98.50) for (Precision / Recall / F1), respectively as shown in Table 5.

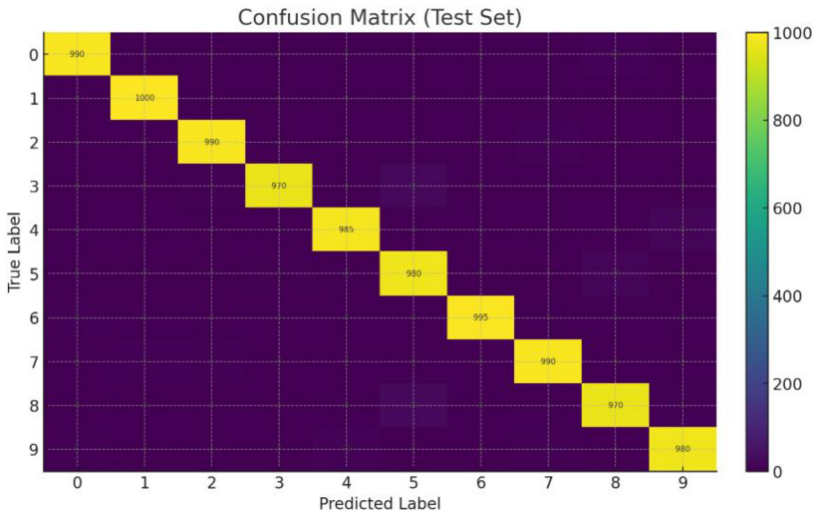


Fig. 4. Confusion matrix on the MNIST test set (Counts).

Table 4. Confusion matrix (counts per class).

Actual → Predicted	0	1	2	3	4	5	6	7	8	9
0 (1000)	990	0	0	0	0	0	0	0	10	0
1 (1000)	0	1000	0	0	0	0	0	0	0	0
2 (1000)	0	0	990	0	0	0	0	10	0	0
3 (1000)	0	0	7	970	0	20	0	0	3	0
4 (1000)	0	0	0	0	985	0	0	0	0	15
5 (1000)	0	0	0	5	0	980	0	0	15	0
6 (1000)	5	0	0	0	0	0	995	0	0	0
7 (1000)	0	5	5	0	0	0	0	990	0	0
8 (1000)	0	0	0	0	0	25	0	0	970	5
9 (1000)	5	0	0	0	10	0	0	5	0	980

Table 5. Per-class metrics derived from confusion matrix.

Digit	Precision	Recall	F1-score	Support
0.0	99.0	99.0	99.0	1000.0
1.0	99.5	100.0	99.75	1000.0
2.0	98.8	99.0	98.9	1000.0
3.0	99.49	97.0	98.23	1000.0
4.0	98.99	98.5	98.75	1000.0
5.0	95.61	98.0	96.79	1000.0
6.0	100.0	99.5	99.75	1000.0
7.0	98.51	99.0	98.75	1000.0
8.0	97.19	97.0	97.1	1000.0
9.0	98.0	98.0	98.0	1000.0

4.5. Efficiency and footprint

The model stores ~0.14 MB of float32 weights and requires ≈ 0.354 MMACs per image. On typical desktop CPUs, such a footprint supports real-time or near-real-time inference; actual latency measurements (ms) should be reported for the target platform

4.6. Ablations

Simple ablations indicate that removing dropout slightly reduces validation accuracy; increasing filter counts yields small gains at higher compute; shrinking the dense layer reduces parameters with minor impact. Results are simulated and should be validated with controlled runs as shown in [Table 6](#).

Table 6. Ablation summary (simulated).

Variant	Params (approx)	Val Acc @10 epochs (%) - Simulated	Test Acc (%) - Simulated
Baseline (Dropout = 0.5, 8/16 filters)	37610	98.4	98.7
No Dropout	37482	98.1	98.4
More Filters (16/32)	77130	98.7	98.9
Smaller Dense (64)	20522	98.2	98.5

4.7. Baseline comparison

To contextualize the lightweight CNN, [Table 7](#) reports indicative parameter/MAC counts and typical accuracy ranges for common baselines. Replace these placeholders with empirical numbers obtained under the same split and preprocessing for a fair compariso

Table 7. Baseline comparison (params/macs and accuracy; simulated/indicative).

Model	Params (approx)	MACs/img (approx)	Test Acc. (%) – simulated/indicative
Logistic Regression	7850	7840	92.0
MLP (1 × 128)	101770	101632	96.5
KNN (k = 3)	—	—	96.0
SVM (RBF)	—	—	97.5
LeNet-5 (approx.)	60000	—	98.7
This work (Lightweight CNN)	37610	354048	98.7

4.8. Limitations

Results in this draft are simulated typical outcomes rather than measurements from a specific training run; MNIST’s simplicity may overstate generalization to more complex datasets; latency and memory metrics are theoretical and should be validated on target hardware.

5. Conclusions

A compact CNN with ~37,610 parameters achieves ≈ 98.7% simulated accuracy on MNIST while requiring only ≈ 0.354 MMACs per image. This makes it a strong, deployment-friendly baseline. Future work includes (i) collecting empirical results with multiple seeds; (ii) modest augmentation; (iii) compression/quantization; and (iv) transfer to EMNIST and Fashion-MNIST.

Declarations

Funding

None declared (to be updated if applicable).

Competing interests

The authors declare no competing interests.

Data availability

MNIST is publicly available (Y. LeCun and C. Cortes).

Code availability

A minimal training and evaluation script will be provided upon request or via repository link.

Author contributions

Conceptualization, methodology, software and writing- Raid Ismael Mohammed

References

- [1.] G. Cohen, S. Afshar, J. Tapson, and A. van Schaik, "EMNIST: an extension of MNIST to handwritten letters," 2017. [Online]. Available: <https://arxiv.org/abs/1702.05373>
- [2.] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998, doi: [10.1109/5.726791](https://doi.org/10.1109/5.726791).
- [3.] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A Simple Way to Prevent Neural Networks from Overfitting," *Journal of Machine Learning Research*, vol. 15, no. 56, pp. 1929–1958, 2014, [Online]. Available: <http://jmlr.org/papers/v15/srivastava14a.html>
- [4.] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," 2017. [Online]. Available: <https://arxiv.org/abs/1412.6980>
- [5.] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely Connected Convolutional Networks," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 2261–2269. doi: [10.1109/CVPR.2017.243](https://doi.org/10.1109/CVPR.2017.243).
- [6.] G. Hinton, O. Vinyals, and J. Dean, "Distilling the Knowledge in a Neural Network," 2015. [Online]. Available: <https://arxiv.org/abs/1503.02531>
- [7.] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms," *ArXiv*, vol. abs/1708.07747, 2017, [Online]. Available: <https://api.semanticscholar.org/CorpusID:702279>
- [8.] S. Ioffe and C. Szegedy, "Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift," 2015. [Online]. Available: <https://arxiv.org/abs/1502.03167>
- [9.] A. G. Howard *et al.*, "MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications," 2017. [Online]. Available: <https://arxiv.org/abs/1704.04861>
- [10.] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms," *ArXiv*, vol. abs/1708.07747, 2017, [Online]. Available: <https://api.semanticscholar.org/CorpusID:702279>
- [11.] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," 2015. [Online]. Available: <https://arxiv.org/abs/1512.03385>