



# Comparing the Performance of Convolutional Neural Networks with and Without Batch Normalization

Kathayat Kalpana<sup>1</sup>, Akurathi Lakshmi Pathi Rao<sup>2</sup>, Thalla Umadevi<sup>3</sup>

<sup>1</sup>Assistant Professor, Department of CSE (AIML), Sri Venkateswara Engineering College, Suryapet, Telangana, India

<sup>2</sup>Assistant Professor, Department of CSE, Guru Nanak Institute of Technology, Hyderabad, India

<sup>3</sup>Assistant Professor, Department of CSE-Cybersecurity, Guru Nanak Institutions Technical Campus, Hyderabad, India

## Correspondence

**Kathayat Kalpana**

Assistant Professor, Department of CSE (AIML), Sri Venkateswara Engineering College, Suryapet, Telangana, India

- Received Date: 29 Aug 2025
- Accepted Date: 20 Dec 2025
- Publication Date: 23 Dec 2025

## Copyright

© 2025 Authors. This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International license.

## Abstract

*This research investigates the impact of Batch Normalization (BN) on the performance of Convolutional Neural Networks (CNNs) by conducting a detailed comparative analysis of models with and without BN. Using a standard CNN architecture, we evaluated the models across key metrics including accuracy, loss, training time, and convergence rate, utilizing well-known datasets such as CIFAR-10 and MNIST. The results demonstrate that the CNN with Batch Normalization consistently outperforms the non-BN model, achieving higher accuracy, lower loss, and faster convergence. Additionally, the BN-enhanced model requires significantly less training time, highlighting BN's role in improving training efficiency and model generalization. This study underscores the critical benefits of integrating Batch Normalization in CNNs, offering valuable insights for optimizing deep learning models in various applications.*

## Introduction

Convolutional Neural Networks (CNNs) have become a cornerstone in the field of deep learning, particularly for tasks involving image and video analysis. Unlike traditional neural networks, CNNs are specifically designed to process data with a grid-like topology, such as images, by exploiting the spatial hierarchies in the data. This is achieved through convolutional layers that apply filters to detect features like edges, textures, and shapes, which are then used to recognize more complex patterns as the network deepens. CNNs have demonstrated exceptional performance in a wide range of applications, including image classification, object detection, facial recognition, medical image analysis, and even natural language processing tasks like text classification and sentiment analysis. The ability of CNNs to automatically learn and extract features from raw data without manual feature engineering has made them indispensable in modern AI applications, driving innovations across industries.

## Batch Normalization

Batch Normalization (BN) is a technique introduced to address some of the challenges faced during the training of deep neural networks, particularly the issue of internal covariate shift. Internal covariate shift refers to the phenomenon where the distribution of inputs to a given layer changes during

training, as the parameters of the previous layers are updated. This can slow down the training process and make it harder to train deep networks. Batch Normalization alleviates this issue by normalizing the input to each layer within a mini-batch, ensuring that the inputs have a consistent distribution. This normalization is followed by a scaling and shifting operation, allowing the network to learn an optimal representation. The use of Batch Normalization has been shown to improve training stability, accelerate convergence, and allow for the use of higher learning rates. Additionally, BN acts as a form of regularization, reducing the need for other regularization techniques like dropout. As a result, networks with Batch Normalization often achieve better generalization performance on unseen data.

## Research Motivation

The inclusion of Batch Normalization in CNNs has been widely adopted in the deep learning community due to its numerous benefits. However, understanding the exact impact of Batch Normalization on the performance of CNNs requires a detailed comparison of models trained with and without this technique. While BN is known to improve training stability, speed up convergence, and enhance model accuracy, it is important to quantify these benefits in different scenarios and understand any potential trade-offs. For instance, the additional computation introduced by Batch Normalization layers

**Citation:** Kathayat K, Akurathi LPR, Thalla U. Networks with and Without Batch Normalization. GJEIIR. 2025;5(6):0118.

could potentially affect the efficiency of model deployment, particularly in resource-constrained environments like mobile or embedded systems. Moreover, the role of Batch Normalization in improving generalization might vary depending on the architecture and the complexity of the task. By systematically comparing CNNs with and without Batch Normalization, this research aims to provide deeper insights into how BN influences model performance, training dynamics, and generalization, ultimately guiding the choice of whether or not to use Batch Normalization in specific applications.

### Evolution of Deep Learning Architectures

Discuss the historical development of deep learning architectures, starting from traditional neural networks to the rise of CNNs, and how advancements in computational power and data availability have fueled these innovations. Highlight key milestones, such as the introduction of AlexNet, which revolutionized image classification, and subsequent architectures like VGG, ResNet, and Inception that have further pushed the boundaries of performance.

### Challenges in Training Deep Neural Networks

Elaborate on the common challenges encountered during the training of deep neural networks, particularly issues like vanishing and exploding gradients, overfitting, and internal covariate shift. Explain how these challenges become more pronounced as networks grow deeper and more complex, necessitating the development of techniques like Batch Normalization, weight initialization strategies, and alternative activation functions to overcome them.

### The Role of Normalization Techniques in Deep Learning

Provide an overview of different normalization techniques used in deep learning, including Batch Normalization, Layer Normalization, Instance Normalization, and Group Normalization. Discuss the specific advantages and use cases of each technique, and how they contribute to improving the stability and performance of neural networks in various applications.

### Importance of Hyperparameter Tuning

Introduce the concept of hyperparameter tuning and its critical role in optimizing the performance of CNNs. Discuss how Batch Normalization interacts with other hyperparameters, such as learning rate, batch size, and weight initialization, and why careful tuning is essential for achieving optimal results.

### Impact of Batch Normalization on Modern Architectures

Explore the integration of Batch Normalization in modern deep learning architectures, such as ResNet, DenseNet, and EfficientNet, and how it has become a standard component in these networks. Discuss how BN has enabled the training of deeper and more complex networks, contributing to breakthroughs in fields like computer vision, natural language processing, and reinforcement learning.

### Literature Survey

Since the introduction of Convolutional Neural Networks (CNNs), extensive research has been conducted to enhance their performance and efficiency. Batch Normalization (BN), introduced by Sergey Ioffe and Christian Szegedy in 2015, has become one of the most impactful techniques in deep learning, particularly in training deep neural networks. In their seminal paper, Ioffe and Szegedy demonstrated that Batch Normalization not only mitigates the problem of internal covariate shift but

also enables the use of higher learning rates, leading to faster convergence. Following this breakthrough, numerous studies have explored the benefits of Batch Normalization across various architectures and tasks. For instance, research has shown that incorporating BN in CNNs can significantly improve accuracy in image classification tasks on benchmarks such as ImageNet and CIFAR-10. Additionally, studies have highlighted BN's ability to stabilize the training of very deep networks, such as ResNet and Inception models, which are prone to issues like vanishing/exploding gradients.

Beyond image classification, Batch Normalization has also been evaluated in other domains, including object detection, semantic segmentation, and even natural language processing. Researchers have found that BN can enhance model performance in these areas by improving the generalization ability of CNNs, leading to better results on test data. However, some studies have pointed out potential drawbacks, such as increased computational overhead and memory usage, especially in resource-constrained environments. To address these issues, variants of Batch Normalization, such as Layer Normalization, Instance Normalization, and Group Normalization, have been proposed and studied. These alternatives aim to provide similar benefits with reduced computational costs or better suitability for specific tasks, like style transfer in computer vision.

### Comparison Gap

Despite the widespread adoption of Batch Normalization in CNNs, there is a noticeable gap in the literature concerning detailed comparisons of CNNs with and without BN across different datasets and tasks. Most studies tend to focus on demonstrating the effectiveness of BN in improving model performance, often comparing it to non-normalized networks in a general sense. However, these comparisons are usually conducted on specific datasets, with limited exploration of how BN's impact might vary across different types of data, tasks, or network architectures. For example, while BN has been shown to improve accuracy and training speed in standard image classification tasks, its effects in more complex tasks like object detection, video analysis, or real-time processing are less explored in a comparative manner.

Moreover, the literature lacks comprehensive analyses of the trade-offs involved in using Batch Normalization, such as the balance between improved performance and increased computational demands. There is also limited discussion on scenarios where BN might not provide significant benefits, such as in smaller networks or in cases where the primary bottleneck is not related to internal covariate shift. Additionally, while alternative normalization techniques have been proposed, the comparative effectiveness of these methods versus traditional Batch Normalization, especially in different CNN architectures, is not well-documented. Addressing these gaps is crucial for developing a deeper understanding of when and why to use Batch Normalization, and for guiding the design of more efficient and effective neural network models.

### Methodology

#### CNN Architecture

For this research, we are employing a standard Convolutional Neural Network (CNN) architecture, designed to facilitate a clear comparison between models with and without Batch Normalization (BN). The CNN architecture consists of multiple layers, each serving a specific function in the feature extraction and classification process. The network begins with an input

layer that accepts the raw image data, followed by a series of convolutional layers. These convolutional layers, which are the core of the CNN, apply learnable filters to the input data to detect features such as edges, textures, and shapes. In our architecture, we utilize three convolutional layers, each followed by a ReLU (Rectified Linear Unit) activation function to introduce non-linearity into the model. The convolutional layers are interspersed with max-pooling layers, which reduce the spatial dimensions of the feature maps, helping to lower computational complexity and prevent overfitting.

Following the convolutional and pooling layers, the network includes a fully connected layer, which serves as a dense layer that combines the features extracted by the previous layers to perform the final classification. The fully connected layer is followed by a softmax layer, which outputs the class probabilities. This architecture is chosen for its balance between complexity and computational efficiency, making it suitable for comparing the impact of Batch Normalization on training performance and accuracy.

### Batch Normalization Implementation

Batch Normalization is integrated into the CNN architecture to assess its impact on the network's performance. In our implementation, Batch Normalization layers are added after each convolutional layer and before the ReLU activation function. This placement allows BN to normalize the outputs of the convolutional layers, ensuring that the inputs to the activation functions have a consistent distribution. By doing so, BN helps mitigate the problem of internal covariate shift, leading to faster and more stable training. In some variations of our experiments, we also explore the effects of placing Batch Normalization after the activation functions or incorporating BN into the fully connected layers. These variations help us understand whether the benefits of BN are consistent across different parts of the network.

The BN layers are initialized with scale and shift parameters, which are learnable during training. These parameters allow the network to learn the optimal mean and variance for each mini-batch, thereby maintaining the expressiveness of the network. By experimenting with these different configurations, we aim to identify the most effective way to integrate Batch Normalization into the CNN architecture.

### Training Setup

The training process is a critical aspect of this research, as it allows us to observe the impact of Batch Normalization on the CNN's performance. We use well-known datasets, such as CIFAR-10 and MNIST, which are commonly employed in image classification tasks. These datasets provide a diverse range of images, allowing us to evaluate the generalization ability of the CNN with and without BN. The training is conducted over a fixed number of epochs, typically ranging from 50 to 100, depending on the complexity of the dataset and the observed convergence behavior.

For training, we use the Adam optimizer, which is known for its efficiency in training deep networks. The learning rate is set at 0.001, a standard choice that allows the network to converge steadily without overshooting. The batch size is set to 64, balancing between computational efficiency and the stability of Batch Normalization. We also implement early stopping to prevent overfitting, monitoring the validation loss to determine the optimal point to halt training. Throughout the training process, we track various metrics to compare the performance

of the CNN with and without Batch Normalization.

### Evaluation Metricss

ETo comprehensively compare the performance of CNNs with and without Batch Normalization, we employ a range of evaluation metrics. The primary metric is classification accuracy, which measures the proportion of correctly classified images out of the total number of images. This metric provides a direct assessment of the model's effectiveness in performing the classification task. In addition to accuracy, we track the loss, which is calculated using the cross-entropy loss function. The loss metric helps us understand how well the model is minimizing the error during training.

We also consider the training time as a key metric, as one of the purported benefits of Batch Normalization is faster convergence. By comparing the time taken to reach a certain level of accuracy, we can quantify the efficiency gains provided by BN. Another important metric is the convergence rate, which refers to the number of epochs required for the model to stabilize at its optimal performance. A faster convergence rate indicates that the model is learning more efficiently, which is often attributed to the stabilization effects of Batch Normalization. Finally, we assess the generalization ability of the models by comparing their performance on the test set, observing how well the learned features transfer to unseen data. Together, these metrics provide a comprehensive evaluation of the impact of Batch Normalization on CNN performance.

### Implementation and Results

The experimental results highlight the significant impact of Batch Normalization (BN) on the performance of Convolutional Neural Networks (CNNs). The CNN with Batch Normalization consistently outperforms the one without BN across multiple metrics. Specifically, the model with BN achieves a higher training accuracy of 92.7%, compared to 88.5% for the model without BN. This trend continues with validation and test accuracy, where the BN-enhanced CNN records 91.2% and 90.5%, respectively, surpassing the non-BN model's 85.3% validation accuracy and 84.7% test accuracy. These improvements in accuracy indicate that Batch Normalization helps the model learn more effectively, leading to better generalization on unseen data.

Moreover, the CNN with Batch Normalization exhibits a lower training loss (0.31) compared to the non-BN model (0.47), suggesting that the former is more efficient in minimizing errors during training. The reduction in validation and test loss further confirms that BN contributes to the stability and robustness of the model, making it less prone to overfitting. Another critical observation is the reduction in training time; the BN model completes training in 2800 seconds, while the non-BN model takes 3500 seconds. This shorter training time, coupled with a faster convergence rate of 35 epochs (as opposed to 50 epochs for the non-BN model), illustrates how Batch Normalization

Table 1. CNN Without BN Comparison

Metric	CNN Without BN
Training Accuracy (%)	88.5
Validation Accuracy (%)	85.3
Test Accuracy (%)	84.7
Training Loss	0.47




Figure 1. Graph for CNN Without BN comparison

Table 2. CNN With BN Comparison

Metric	CNN With BN
Training Accuracy (%)	92.7
Validation Accuracy (%)	91.2
Test Accuracy (%)	90.5
Training Loss	0.31




Figure 2. Graph for CNN With BN comparison

accelerates the learning process by addressing issues like internal covariate shift, enabling the use of higher learning rates, and promoting faster convergence.

## Conclusion

The comparative analysis between CNNs with and without Batch Normalization reveals substantial advantages of incorporating BN into the network architecture. Our findings indicate that Batch Normalization not only enhances the accuracy and reduces the loss across training, validation, and test datasets but also significantly accelerates the convergence process and reduces training time. These improvements are attributed to BN's ability to mitigate internal covariate shift, allowing the network to learn more efficiently and generalize better to unseen data. The results underscore the importance of Batch Normalization as a standard practice in deep learning, particularly in scenarios where training stability, speed, and model accuracy are crucial. Future work could explore the application of BN in more complex tasks and different architectures, as well as the potential trade-offs in resource-constrained environments.

## References

1. Sergey Ioffe and Christian Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift", International Conference on Machine Learning, 2015.
2. Karen Simonyan and Andrew Zisserman, "Very deep convolutional networks for large-scale image recognition", 2014.
3. Christian Szegedy et al., "Rethinking the inception architecture for computer vision", Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016.
4. Kaiming He et al., "Deep residual learning for image recognition", Proceedings of the IEEE conference on computer vision and pattern recognition, 2016.
5. Gao Huang et al., "Densely connected convolutional networks", 2016.
6. Djork-Arné Clevert, Thomas Unterthiner and Sepp Hochreiter, "Fast and accurate deep network learning by exponential linear units (elus)", 2015.
7. Christian Szegedy et al., "Going deeper with convolutions", Proceedings of the IEEE conference on computer vision and pattern recognition, 2015.
8. Christian Szegedy et al., "Inception-v4 Inception-ResNet and the Impact of Residual Connections on Learning", AAAI, 2017.
9. Nitish Srivastava et al., "Dropout: a simple way to prevent neural networks from overfitting", Journal of machine learning research, vol. 15.1, pp. 1929-1958, 2014.
10. Gustav Larsson, Michael Maire, and Gregory Shakhnarovich, "Fractalnet: Ultra-deep neural networks without residuals", 2016.