



## FPGA-Based Acceleration of Convolutional Neural Networks for Image Recognition

John A Doe <sup>1\*</sup>, Dr. Femi Adeyemi <sup>2</sup>

<sup>1</sup> Department of Electrical Engineering, Massachusetts Institute of Technology (MIT), USA

<sup>2</sup> School of Engineering, University of Lagos, Nigeria

\* Corresponding Author: John A Doe

---

### Article Info

**Volume:** 01

**Issue:** 02

**March-April 2025**

**Received:** 28-02-2025

**Accepted:** 20-03-2025

**Page No:** 01-03

### Abstract

The rapid advancement of deep learning, particularly Convolutional Neural Networks (CNNs), has revolutionized image recognition. However, the computational intensity and power demands of CNNs present significant challenges for real-time, embedded, and edge applications. Field Programmable Gate Arrays (FPGAs) have emerged as a promising hardware platform to accelerate CNN inference, offering a balance between performance, flexibility, and energy efficiency. This paper presents a comprehensive overview of FPGA-based CNN acceleration for image recognition, covering architectural design, implementation strategies, performance analysis, and future directions. Recent research and practical implementations are discussed, highlighting the benefits and challenges of deploying CNNs on FPGAs.

**Keywords:** FPGA Acceleration, Convolutional Neural Networks (CNNs), Edge and Embedded AI, Energy-Efficient Computing, Quantization and Compression

---

### 1. Introduction

Image recognition is a cornerstone of modern artificial intelligence, powering applications from autonomous vehicles and medical diagnostics to industrial automation and smart devices. Convolutional Neural Networks (CNNs) have set new benchmarks in image recognition accuracy, but their deep architectures require substantial computational and memory resources.

Traditional hardware platforms—CPUs and GPUs—provide the necessary computational power but are often unsuitable for power-constrained, real-time, or embedded scenarios due to high energy consumption or lack of flexibility. FPGAs, with their reconfigurable logic and parallel processing capabilities, offer an attractive alternative for accelerating CNN inference, especially in edge and embedded systems<sup>53</sup>.

### 2. Motivation for FPGA Acceleration of CNNs

#### 2.1 Computational Demands of CNNs

CNNs consist of multiple layers—convolutional, pooling, activation, and fully connected layers—each involving intensive matrix and vector operations. Large-scale models such as VGG, ResNet, and YOLO require billions of operations per inference, making real-time processing challenging on general-purpose processors<sup>5</sup>.

#### 2.2 Limitations of CPUs and GPUs

- **CPUs:** Offer flexibility but limited parallelism, leading to high latency for deep CNNs.
- **GPUs:** Provide massive parallelism and high throughput but consume significant power and are less suitable for custom optimizations in embedded contexts<sup>[5]</sup>.
- **ASICs:** Deliver high performance and low power but lack post-fabrication flexibility and have high development costs.

#### 2.3 Advantages of FPGAs

- **Parallelism:** FPGAs can exploit the inherent parallelism of CNNs by mapping multiple operations to hardware.
  - **Reconfigurability:** Hardware can be tailored to specific CNN architectures and updated post-deployment.
-

- **Energy Efficiency:** FPGAs achieve high throughput at a fraction of the power consumption of GPUs, making them ideal for edge and embedded AI applications<sup>56</sup>.
- **Custom Precision:** Support for fixed-point and custom-precision arithmetic enables further optimization of speed and resource usage<sup>6</sup>.

### 3. FPGA Architecture for CNN Acceleration

#### 3.1 Basic Architecture

An FPGA-based CNN accelerator typically consists of:

- **Processing Elements (PEs):** Parallel units for convolution, pooling, and activation functions.
- **On-chip Memory:** Buffers for input data, weights, and intermediate feature maps.
- **DMA Controllers:** Manage data transfer between external memory and on-chip buffers.
- **Control Logic:** Orchestrates data flow and layer execution.

#### 3.2 Design Strategies

- **Pipelining:** Overlaps computation stages to maximize throughput.
- **Parallelism:** Distributes computations across multiple PEs for concurrent execution.
- **Resource Sharing:** Efficiently reuses hardware blocks for different CNN layers<sup>35</sup>.
- **Data Quantization:** Reduces bit-width of operations (e.g., 16-bit or 8-bit fixed-point) to save resources and increase speed<sup>6</sup>.

#### 3.3 Implementation Example

A recent implementation on a Xilinx Virtex-7 FPGA used Verilog to design an accelerator that balanced computational efficiency and resource usage, utilizing 588 Look-Up Tables (LUTs) and 353 Flip Flops<sup>3</sup>. Another design on a Z-7020 FPGA used 16-bit fixed-point operations and kernel binarization to optimize performance and minimize hardware cost<sup>6</sup>.

### 4. CNN Mapping and Optimization on FPGAs

#### 4.1 Layer-wise Mapping

- **Convolutional Layers:** Mapped to parallel PEs for simultaneous processing of multiple kernels and input channels.
- **Pooling Layers:** Implemented as simple max or average operations, often fused with convolution for efficiency.
- **Activation Functions:** Realized using lookup tables or piecewise linear approximations to minimize resource usage.

#### 4.2 Dataflow Optimization

- **Systolic Arrays:** Enable efficient matrix-vector multiplication by streaming data through a grid of PEs<sup>5</sup>.
- **Resource Multiplexing:** Allows the same hardware to be used for different operations (e.g., convolution and matrix multiplication), reducing area and power<sup>5</sup>.
- **Memory Hierarchy:** On-chip buffers reduce latency and external memory bandwidth requirements.

#### 4.3 Quantization and Compression

- **Fixed-Point Arithmetic:** Replacing floating-point with fixed-point reduces hardware complexity and power

consumption, with minimal impact on accuracy<sup>6</sup>.

- **Model Pruning and Binarization:** Removing redundant weights and binarizing kernels further reduces resource demands and accelerates computation<sup>6</sup>.

### 5. Performance Evaluation

#### 5.1 Benchmarking Metrics

- **Throughput (GOPS):** Giga Operations Per Second, measuring raw computational performance.
- **Latency:** Time taken for a single inference.
- **Power Consumption (W):** Total energy usage during operation.
- **Energy Efficiency (GOPS/W):** Throughput per watt, a key metric for embedded and edge applications<sup>5</sup>.

#### 5.2 Comparative Analysis

An FPGA-based resource reuse accelerator for 1D-CNN-LSTM achieved 7.34 GOPS at 5.022 W, with an energy efficiency of 1.46 GOPS/W, outperforming CPUs by 73 times and GPUs by over 9 times in energy efficiency for radar emitter signal recognition<sup>5</sup>. Another design achieved a recognition rate of 97.53% while maintaining low power consumption, demonstrating the suitability of FPGAs for real-time, power-constrained applications<sup>5</sup>.

#### 5.3 Application to Image Recognition

FPGA-accelerated CNNs have been successfully deployed for image recognition tasks such as object detection (YOLOv4-Tiny), achieving real-time performance on edge devices<sup>1</sup>. The flexibility of FPGAs allows adaptation to various CNN models and image recognition benchmarks.

### 6. Case Studies and Real-World Applications

#### 6.1 YOLOv4-Tiny Object Detection

A method for accelerating YOLOv4-Tiny on FPGA demonstrated the feasibility of deploying complex object detection networks in resource-constrained environments, balancing speed and accuracy<sup>1</sup>.

#### 6.2 Radar Signal Recognition

FPGA-based 1D-CNN-LSTM models have been used for radar emitter signal recognition, achieving high accuracy and energy efficiency, making them suitable for communication security and electronic support systems<sup>5</sup>.

#### 6.3 Edge AI and IoT

FPGAs are increasingly used in edge AI scenarios, where real-time image recognition is required under strict power and latency constraints, such as in smart cameras, drones, and autonomous vehicles<sup>46</sup>.

### 7. Challenges and Limitations

#### 7.1 Resource Constraints

FPGAs have limited on-chip memory and logic resources compared to GPUs, requiring careful design and optimization to fit large CNN models<sup>7</sup>.

#### 7.2 Design Complexity

Developing efficient FPGA accelerators requires expertise in hardware design, parallel programming, and deep learning, making the development process more complex than for software-based platforms<sup>7</sup>.

### 7.3 Scalability

Scaling FPGA designs to support very large models or batch processing can be challenging due to hardware limitations and memory bandwidth bottlenecks<sup>7</sup>.

### 7.4 Portability

FPGA designs are often tailored to specific devices or models, limiting portability and reusability across different platforms or CNN architectures<sup>7</sup>.

## 8. Future Directions

### 8.1 High-Level Synthesis (HLS) and Frameworks

The adoption of HLS tools (e.g., Intel OpenCL, Xilinx Vitis) and deep learning frameworks is streamlining FPGA development, enabling faster prototyping and deployment of CNN accelerators<sup>2</sup>.

### 8.2 Dynamic Reconfiguration

Future FPGAs may support dynamic reconfiguration, allowing hardware to adapt to different CNN layers or models at runtime for optimal performance and resource utilization<sup>7</sup>.

### 8.3 Support for Advanced Models

Research is ongoing to efficiently map more complex models (e.g., Transformers, 3D CNNs) and support mixed-precision or adaptive computation on FPGAs.

### 8.4 Integration with Edge and Cloud AI

FPGAs are expected to play a key role in heterogeneous computing platforms, complementing CPUs and GPUs in edge-cloud AI pipelines for scalable, energy-efficient image recognition<sup>27</sup>.

## 9. Conclusion

FPGAs offer a compelling platform for accelerating CNN inference in image recognition, combining high performance, energy efficiency, and adaptability. Through parallelism, custom precision, and resource sharing, FPGA-based accelerators can meet the stringent requirements of real-time, embedded, and edge AI applications. While challenges remain in design complexity and scalability, ongoing advances in tools, architectures, and model optimization continue to expand the potential of FPGAs in deep learning. As AI-powered image recognition becomes ubiquitous, FPGA acceleration will be central to enabling intelligent, responsive, and energy-efficient systems.

## 10. References

1. Zhang Y, *et al.* Acceleration and implementation of convolutional neural networks for YOLOv4-Tiny object detection algorithm accelerator based on FPGA. ScienceDirect. 2023. Available from: <https://www.sciencedirect.com/science/article/pii/S105120042300283X>
2. BittWare. FPGA Acceleration of Convolutional Neural Networks (CNNs). White Paper. 2023. Available from: <https://www.bittware.com/resources/cnn/>
3. Patel P, *et al.* Implementation of FPGA-based Accelerator for Convolutional Neural Networks. Int J Res Sci Eng. 2024;4(03):10–16. Available from: <https://journal.hmjournals.com/index.php/IJRSE/article/view/3838>
4. ITM Conferences. Implementing Convolutional Neural Networks on FPGA. 2023. Available

from: [https://www.itm-](https://www.itm-conferences.org/articles/itmconf/pdf/2023/02/itmconf_ocia2023_02004.pdf)

[conferences.org/articles/itmconf/pdf/2023/02/itmconf\\_ocia2023\\_02004.pdf](https://www.itm-conferences.org/articles/itmconf/pdf/2023/02/itmconf_ocia2023_02004.pdf)

5. Wang S, *et al.* Efficient FPGA Implementation of Convolutional Neural Networks and LSTM for Radar Emitter Signal Recognition. Frontiers in Neuroscience. 2024;18:10857097. Available from: <https://pmc.ncbi.nlm.nih.gov/articles/PMC10857097/>
6. Li X, *et al.* FPGA Implementation of a Convolutional Neural Network and Its Application in Agriculture. Agriculture. 2022;12(11):1849. Available from: <https://www.mdpi.com/2077-0472/12/11/1849>
7. Suda N, *et al.* A survey of FPGA-based accelerators for convolutional neural networks. Neural Comput Appl. 2019;32:1109–1139. Available from: <https://dl.acm.org/doi/10.1007/s00521-018-3761-1>