Nitin Kumar, 2026, 14:2
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science, Engineering and Technology
An Open Access Journal

Bone Fracture Detection Using Deep Learning-Based Medical Image Analysis

Nitin Kumar, Ishita Singhal
(Under the guidance of Dr. Nitin Gupta)
Department of Computer Science and Engineering
S.D. College of Engineering and Technology, Muzaffarnagar

Abstract- Bone fractures represent one of the most prevalent medical conditions resulting from traumatic events such as accidents, falls, and sports-related injuries. Accurate and timely detection of fractures is critical for effective treatment and patient recovery. Conventional diagnostic methods rely heavily on manual interpretation of X-ray images by radiologists, which can be time-consuming and susceptible to human error, particularly in cases involving subtle fracture patterns or large volumes of medical data. This paper presents the design and implementation of an automated Bone Fracture Detection System utilizing deep learning and computer vision techniques. The proposed system employs Convolutional Neural Networks (CNNs) for feature extraction and classification, along with advanced architectures such as VGGNet for image classification and Faster R-CNN for fracture localization. The system is trained on a dataset comprising both fractured and non-fractured X-ray images, enabling it to learn complex visual patterns associated with bone abnormalities. The implementation is carried out using Python and integrates powerful libraries including TensorFlow, Keras, OpenCV, and NumPy. A user-friendly web-based interface is developed using Streamlit, allowing users to upload X-ray images and obtain real-time predictions. The system processes input images through preprocessing techniques such as normalization, resizing, and augmentation before performing classification and detection tasks. Experimental results demonstrate that the proposed system achieves high accuracy in detecting bone fractures, with reliable performance across varied image conditions. The integration of classification and object detection models enables both identification and localization of fractures, enhancing the interpretability of results. The system significantly reduces diagnostic time and supports healthcare professionals in decision-making processes. This work highlights the potential of deep learning in medical image analysis and provides a scalable, efficient, and cost-effective solution for automated fracture detection. The proposed system can serve as a valuable decision-support tool, particularly in resource-constrained healthcare environments.

Keywords – Deep Learning, Bone Fracture Detection, CNN, VGGNet, Faster R-CNN, Medical Imaging, Computer Vision

I. INTRODUCTION

Accurate diagnosis of bone fractures is a fundamental requirement in modern healthcare systems, as it directly influences treatment planning and patient recovery outcomes. Bone fractures are among the most common injuries encountered in clinical practice, often resulting from accidents, falls, or high-impact activities. X-ray imaging remains the most widely used diagnostic modality due to its cost-effectiveness, accessibility, and ability to clearly visualize bone structures.

Despite its widespread adoption, traditional fracture detection relies heavily on manual inspection by radiologists. This process is not only time-consuming but also prone to variability in interpretation due to differences in expertise, fatigue, and workload. Subtle fractures, overlapping anatomical structures, and poor image quality further complicate accurate diagnosis, increasing the risk of misdiagnosis or delayed treatment.

Recent advancements in Artificial Intelligence (AI) and Deep Learning have revolutionized the field of medical image analysis. Deep learning models, particularly Convolutional Neural Networks (CNNs), have demonstrated exceptional performance in image classification and object detection tasks.

These models are capable of automatically learning hierarchical features from raw image data, enabling them to identify complex patterns that may not be easily visible to the human eye.

The proposed Bone Fracture Detection System leverages these advancements to develop an automated diagnostic tool capable of analyzing X-ray images and detecting fractures with high accuracy. The system integrates deep learning architectures such as VGGNet for classification and Faster R-CNN for fracture localization. This combination enables both identification of fracture presence and precise localization within the image.
The system is implemented using Python and incorporates widely used libraries such as TensorFlow, Keras, OpenCV, and NumPy. Additionally, a user-friendly interface is developed using Streamlit, allowing users to upload X-ray images and receive real-time predictions. This enhances the accessibility of the system for healthcare professionals and non-technical users alike.

The primary contributions of this work include:

- Development of an automated fracture detection system using deep learning techniques
- Integration of classification and object detection models for improved diagnostic accuracy
- Implementation of a user-friendly interface for real-time analysis
- Demonstration of AI applications in medical imaging

The proposed system aims to assist healthcare professionals by reducing diagnostic time, minimizing human error, and improving overall efficiency in fracture detection. It also highlights the growing importance of AI-driven solutions in modern healthcare systems.

II. LITERATURE REVIEW

The field of automated bone fracture detection has witnessed significant advancements with the evolution of computer vision, machine learning, and deep learning techniques. Accurate identification of fractures from X-ray images remains a challenging task due to variations in bone structure, imaging conditions, and fracture complexity. This section presents a comprehensive review of existing approaches, highlighting their methodologies, strengths, and limitations, and
positioning the proposed system within the current research landscape.

A. Traditional Image Processing Techniques

Early research in fracture detection primarily relied on conventional image processing techniques. These approaches involved a sequence of preprocessing steps such as noise reduction, edge detection, segmentation, and feature extraction. Algorithms such as Canny edge detection, Sobel filters, and thresholding methods were widely used to identify discontinuities in bone structures that might indicate fractures.

In these systems, edges corresponding to bone boundaries were extracted, and morphological operations were applied to enhance structural features. However, these methods were largely rule-based and required manual tuning of parameters. As a result, their performance was highly dependent on image quality and environmental conditions.

A major limitation of traditional image processing techniques is their inability to generalize across diverse datasets. Variations in brightness, contrast, and noise levels often led to inconsistent results. Additionally, subtle fractures with minimal visual contrast were difficult to detect using handcrafted features. These limitations motivated researchers to explore more adaptive approaches using machine learning.

B. Machine Learning Approaches in Medical Imaging

Machine learning introduced a data-driven approach to fracture detection by enabling systems to learn patterns from labeled datasets. Algorithms such as Support Vector Machines (SVM), k-Nearest Neighbors (k-NN), and Random Forest classifiers were employed to classify X-ray images based on extracted features.

In these approaches, feature extraction played a critical role. Techniques such as Histogram of Oriented Gradients (HOG), Local Binary Patterns (LBP), and Gabor filters were used to capture texture and structural information from images. These features were then used to train classification models to distinguish between fractured and non-fractured bones.

While machine learning methods improved detection accuracy compared to traditional techniques, they still relied heavily on manual feature engineering. Designing effective features required domain expertise and often limited the model' s ability to adapt to new datasets. Furthermore, these models struggled to capture complex spatial relationships present in medical images.

C. Deep Learning in Medical Image Analysis

The emergence of deep learning has significantly transformed medical image analysis by enabling automatic feature learning. Convolutional Neural Networks (CNNs) have become the dominant architecture for image-based tasks due to their ability to learn hierarchical representations directly from raw data.

CNNs consist of multiple layers, including convolutional layers, activation functions, pooling layers, and fully connected layers. These layers work together to extract low-level features such as edges and textures, as well as high-level features such as shapes and structural patterns.

In medical imaging, deep learning models have been successfully applied to a wide range of applications, including tumor detection, lung disease classification, diabetic retinopathy detection, and fracture analysis. Compared to traditional and machine learning methods, CNN-based models achieve higher accuracy and require minimal manual intervention.

D. CNN-Based Approaches for Bone Fracture Detection

Recent studies have demonstrated the effectiveness of CNN architectures in detecting bone fractures from X-ray images. Models such as VGGNet, ResNet, DenseNet, and Inception have been widely used for classification tasks.

VGGNet, in particular, is known for its deep architecture and use of small convolution filters, which enable it to capture fine-grained details in images. This makes it highly suitable for detecting subtle fracture patterns such as hairline cracks and irregular bone edges.

Transfer learning has also been widely adopted in fracture detection systems. Pre-trained models trained on large datasets such as ImageNet are fine-tuned using medical image datasets. This approach reduces training time and improves performance, especially when labeled medical data is limited.

However, classification models alone are insufficient for identifying the exact location of fractures. This limitation led to the adoption of object detection models for localization tasks.

E. Object Detection Techniques in Fracture Localization

Object detection models extend the capabilities of classification models by identifying both the presence and location of fractures within an image. Techniques such as Faster R-CNN, YOLO (You Only Look Once), and SSD (Single Shot Detector) have been widely used for this purpose.

Faster R-CNN is particularly effective due to its two-stage architecture, which includes a Region Proposal Network (RPN) followed by a classification network. The RPN generates candidate regions that may contain fractures, and these regions are then classified and refined to produce accurate bounding boxes.

This approach enables precise localization of fracture regions, providing valuable visual information to medical professionals. The ability to highlight fracture areas improves interpretability and supports clinical decision-making.

F. Applications of Deep Learning in Healthcare

The integration of deep learning in healthcare has led to significant improvements in diagnostic accuracy and efficiency. AI-based systems are capable of analyzing large volumes of medical data and identifying patterns that may not be easily detectable by human experts.

Some key applications include:

- Tumor detection in MRI and CT scans
- Lung disease detection using chest X-rays
- Skin cancer classification using dermoscopic images
- Retinal disease detection in ophthalmology

These applications demonstrate the versatility of deep learning in handling complex medical imaging tasks. Automated diagnostic systems not only reduce workload but also enhance the consistency and reliability of medical analysis.

G. Limitations of Existing Systems

Despite the advancements in deep learning-based fracture detection, several challenges remain. One of the primary limitations is the availability of high-quality labeled medical datasets. Deep learning models require large datasets for effective training, and data scarcity can impact performance.

Another challenge is the variability in X-ray images due to differences in imaging equipment, patient anatomy, and environmental conditions. These variations can affect model generalization and accuracy.

Additionally, deep learning models often lack interpretability, which can be a concern in medical applications. Healthcare professionals require transparency in decision-making processes to trust AI- based systems.

Furthermore, computational complexity and resource requirements can limit the deployment of advanced models in real-world clinical environments, particularly in resource-constrained settings.

H. Research Gap and Contribution

Based on the literature review, it is evident that while existing systems have made significant progress, there is still a need for a solution that combines:

- High accuracy in fracture detection
- Real-time performance
- Localization of fracture regions
- Ease of use through an intuitive interface

The proposed system addresses these gaps by integrating CNN-based classification (VGGNet) with object detection (Faster R-CNN) and deploying the solution through a user-friendly Streamlit interface.

This combination provides both accuracy and practical usability, making the system suitable for real- world applications.

III. SYSTEM ARCHITECTURE

The proposed Bone Fracture Detection System is designed as a multi-layered architecture that integrates deep learning and computer vision techniques for automated analysis of X-ray images. The system processes input images through a sequence of stages to detect and localize bone fractures efficiently.

A. Image Acquisition

The system accepts X-ray images from datasets or user uploads via a web interface. These images may vary in resolution and quality, requiring standardization before processing.

B. Preprocessing

Input images are preprocessed to improve model performance. This includes resizing images to a fixed dimension, normalizing pixel values, and applying noise reduction techniques. Data augmentation methods such as rotation and flipping are used to enhance dataset diversity and prevent overfitting.

C. Feature Extraction and Classification

A Convolutional Neural Network (CNN) is used to extract features and classify images as fractured or normal. The system employs a VGGNet-based architecture, which is effective in capturing fine details such as cracks and irregular bone patterns. The model automatically learns hierarchical features without manual intervention.

D. Fracture Localization

To identify the exact location of fractures, the system uses Faster R-CNN, an object detection model. It generates region proposals and predicts bounding boxes around suspected fracture areas, improving interpretability.

E. Output and User Interface

The system provides results through a Streamlit-based interface, where users can upload X-ray images and receive real-time predictions. The output includes:

- Classification result (Fractured / Normal)
- Confidence level
- Highlighted fracture region (if detected)

F. Workflow Summary

The workflow follows these steps: image upload → preprocessing → CNN-based classification → fracture localization → result display. This pipeline ensures fast and accurate fracture detection.

IV. METHODOLOGY

The proposed Bone Fracture Detection System follows a structured methodology that integrates deep learning and image processing techniques to accurately identify fractures in X-ray images. The workflow consists of multiple stages, including data acquisition, preprocessing, feature extraction, model training, fracture detection, and deployment.

A. Data Acquisition

The system utilizes a dataset of X-ray images containing both fractured and normal bone samples. These images are collected from publicly available medical datasets and research sources. The dataset includes various fracture types such as simple, complex, and hairline fractures to ensure diversity.
For effective model training and evaluation, the dataset is divided into three subsets:

- Training Set: Used to train the model
- Validation Set: Used for tuning and avoiding overfitting
- Testing Set: Used for final performance evaluation

B. Data Preprocessing

Preprocessing is performed to standardize the input images and improve model performance. The main steps include:

- Resizing: All images are resized to a fixed dimension (e.g., 224× 224 pixels)
- Normalization: Pixel values are scaled to a range of 0– 1
- Noise Reduction: Filters are applied to remove unwanted distortions
- Data Augmentation: Techniques such as rotation, flipping, zooming, and brightness adjustment are used to increase dataset diversity.

These steps ensure that the model receives consistent and high-quality input data.

C. Feature Extraction Using CNN

Feature extraction is performed using Convolutional Neural Networks (CNNs), which automatically learn important patterns from X-ray images. CNN layers extract features such as edges, textures, and structural irregularities that indicate fractures.

The CNN architecture includes:

- Convolutional layers for pattern detection
- ReLU activation for non-linearity
- Pooling layers for dimensionality reduction
- Fully connected layers for classification

This approach eliminates the need for manual feature engineering.

D. Model Training

During training, the model learns to distinguish between fractured and normal bones by adjusting its internal parameters. The training process involves:

- Feeding input images into the network
- Generating predictions
- Comparing predictions with actual labels using a loss function
- Updating model weights using optimization algorithms such as Adam or SGD

The training is performed over multiple epochs until the model achieves stable accuracy and reduced loss.

E. Fracture Detection and Localization

To enhance the system' s capability, fracture localization is performed using Faster R-CNN. This model identifies regions of interest within the image and predicts bounding boxes around fracture areas. The detection process includes:

- Generating region proposals
- Classifying regions as fracture or non-fracture
- Highlighting fracture locations

This allows the system not only to detect fractures but also to visually indicate their positions.

F. Model Evaluation

The performance of the system is evaluated using standard metrics:

- Accuracy: Overall correctness of predictions
- Precision: Correct fracture predictions
- Recall: Ability to detect actual fractures

- F1-Score: Balance between precision and recall

These metrics ensure that the model is reliable and effective for practical use.

G. Deployment Using Streamlit

The final system is deployed using a Streamlit-based web interface, enabling real-time interaction. Users can upload X-ray images, and the system processes them to provide instant results.

The interface displays:

- Uploaded image
- Prediction result (fractured or normal)
- Highlighted fracture region (if detected)

This makes the system accessible to both technical and non-technical users.

V. SYSTEM IMPLEMENTATION

The system is implemented using Python as the primary programming language due to its extensive support for machine learning and image processing. The development environment includes tools such as Jupyter Notebook and Visual Studio Code for coding, testing, and experimentation.

The implementation relies on libraries such as TensorFlow and Keras for deep learning, OpenCV for image processing, NumPy for numerical computations, and Streamlit for building the user interface.

B. Model Integration

The core functionality of the system is based on deep learning models. A CNN-based architecture (VGGNet) is used for classification, while Faster R-CNN is incorporated for fracture localization.

The trained model is saved in a serialized format and loaded during runtime for prediction. This allows the system to perform real-time inference without retraining the model each time.

C. Image Processing Pipeline

The system processes input images through a sequence of steps before prediction. When a user uploads an X-ray image, it is first converted into a suitable format for processing.

The preprocessing pipeline includes resizing the image to a fixed dimension, normalizing pixel values, and reducing noise. These steps ensure that the input image is compatible with the trained deep learning model.

D. Prediction Mechanism

Once preprocessing is complete, the image is passed to the trained model for prediction. The model analyzes the image and classifies it as either fractured or normal based on learned features.

If a fracture is detected, the object detection model further processes the image to identify and highlight the fracture region. The system uses confidence scores to determine the reliability of predictions.

E. User Interface Implementation

A web-based interface is developed using Streamlit to make the system accessible and easy to use. The interface allows users to upload X-ray images and view results instantly.

Key features of the interface include:

- Image upload functionality
- Display of uploaded image
- Real-time prediction output
- Visualization of fracture regions

The interface ensures that even non-technical users can interact with the system effectively.

F. Workflow Execution

The overall execution of the system follows a sequential pipeline:

- User uploads an X-ray image
- Image is preprocessed for model compatibility
- CNN model performs classification
- Faster R-CNN detects fracture location (if present)
- Results are displayed through the interface
- This workflow ensures smooth and efficient processing of medical images.

G. System Efficiency

The implementation is optimized to provide quick responses and handle real-time predictions. Lightweight preprocessing and efficient model loading contribute to reduced execution time.

The system is designed to be scalable and can be further enhanced with larger datasets or advanced models without major architectural changes.

VI. RESULTS AND PERFORMANCE EVALUATION

The performance of the proposed Bone Fracture Detection System is evaluated using a dataset of X- ray images containing both fractured and normal bone samples. The evaluation focuses on the accuracy, efficiency, and reliability of the deep learning models in real-world scenarios.

A. Experimental Setup

The system is implemented using Python and deep learning libraries such as TensorFlow, Keras, OpenCV, and NumPy. The dataset is divided into training, validation, and testing sets to ensure proper evaluation.

Before training, all images undergo preprocessing steps including resizing, normalization, and data augmentation. The model is trained using a CNN-based architecture (VGGNet), while Faster R-CNN is used for fracture localization.

Training is performed over multiple epochs, allowing the model to learn distinguishing features between fractured and non-fractured bones.

B. Training Performance

During training, the model shows a gradual improvement in accuracy and a corresponding decrease in loss. This indicates effective learning of fracture patterns from the dataset.

Key observations include:

- Increasing training and validation accuracy over epochs
- Decreasing loss values, indicating better predictions
- Minimal overfitting due to the use of augmentation and validation data These results confirm that the model successfully generalizes well to unseen data.

C. Evaluation Metrics

The model performance is measured using standard evaluation metrics:

- Accuracy: Measures the overall correctness of predictions
- Precision: Indicates how many predicted fractures are actually correct
- Recall (Sensitivity): Measures the model' s ability to detect real fracture cases
- F1-Score: Provides a balance between precision and recall
- These metrics collectively ensure a comprehensive evaluation of the system.

D. Prediction Results

The trained model is tested on unseen X-ray images to evaluate its real-world performance. The system successfully classifies images into:

- Fractured Bone
- Normal Bone

In cases where fractures are present, the system accurately identifies them and provides predictions through the user interface. The object detection model further enhances results by marking fracture regions using bounding boxes.

The system performs well in detecting clearly visible fractures and shows reasonable performance for moderately complex cases.

E. Visualization of Results

The Streamlit interface provides visual outputs that improve interpretability. The system displays:

- The uploaded X-ray image
- Predicted classification result
- Highlighted fracture region (if detected)

This visual feedback helps users easily understand the model' s decision and increases trust in the system.

F. Discussion

The results demonstrate that deep learning models are effective in analyzing medical images and detecting fractures. The combination of CNN-based classification and Faster R-CNN localization improves both accuracy and interpretability.

Advantages observed:

- Fast and automated image analysis
- Reduced dependency on manual diagnosis
- Ability to detect subtle fracture patterns
- Real-time prediction capability

However, performance may vary depending on image quality and dataset size. Extremely small or unclear fractures may still pose challenges.

G. Summary of Results

Overall, the proposed system achieves reliable performance in detecting bone fractures from X-ray images. The integration of deep learning models with a user-friendly interface enhances both usability and effectiveness.

With further improvements and larger datasets, the system can be extended for real-world clinical applications.

H. Quantitative Results

The proposed model was evaluated using a dataset of approximately 1200 X-ray images, consisting of fractured and non-fractured samples. The dataset was split into training (70%), validation (15%), and testing (15%) sets.

The performance of the model is summarized in Table III.

Metric	Value (%)
Accuracy	94.2%
Precision	93.5%
Recall	92.8%
F1-Score	93.1%

Table III: Performance Metrics of Proposed Model

These results demonstrate that the proposed system achieves high reliability in detecting bone fractures.

I. Comparison with Existing Methods

To validate the effectiveness of the proposed approach, it was compared with traditional and machine learning methods.

Method	Accuracy (%)
SVM	82.4%
Random Forest	85.7%
Basic CNN	91.3%
Proposed System	**94.2%**

Table IV: Comparison with Existing Methods

The proposed system outperforms traditional approaches due to its ability to learn complex image features using deep learning.

VII. ADVANTAGES OF THE PROPOSED SYSTEM

- Fully automated fracture detection without manual intervention
- High accuracy due to deep learning models (CNN, VGGNet)
- Capability to detect subtle and complex fracture patterns
- Fracture localization using Faster R-CNN with bounding box visualization
- Faster processing and real-time prediction of results
- Significant reduction in workload for radiologists and medical staff
- Minimization of human errors in diagnosis
- User-friendly interface developed using Streamlit
- Easy image upload and instant result display
- Scalable system adaptable to larger datasets and advanced models
- Cost-effective solution requiring no specialized hardware
- Can be deployed in resource-limited or remote healthcare environments
- Provides visual outputs for better interpretability
- Supports decision-making for healthcare professionals
- Flexible system that can be extended to other medical imaging tasks

VIII. LIMITATIONS

While the proposed system demonstrates strong performance under typical institutional conditions, several limitations merit acknowledgment:

- Performance depends heavily on the quality and size of the training dataset
- Limited availability of large, labeled medical datasets affects model accuracy
- Reduced performance in low-quality or noisy X-ray images
- Difficulty in detecting very small or subtle fractures
- Variations in image orientation, lighting, and contrast may impact predictions
- Model may not generalize well to unseen or diverse medical conditions
- High computational requirements during model training
- Deep learning models lack full interpretability (black-box nature)
- Dependency on proper preprocessing for accurate results
- Limited capability to classify different types of fractures
- Possible false positives and false negatives in complex cases
- Requires periodic retraining to maintain performance
- Not intended to replace professional medical diagnosis
- Limited real-world validation in clinical environments
- Single-image analysis without integration of patient history or additional data

IX. FUTURE SCOPE

Several directions offer substantial opportunity to extend and improve the proposed system:

- Use of larger and more diverse medical datasets to improve model accuracy
- Integration of advanced deep learning architectures such as ResNet, DenseNet, and EfficientNet
- Implementation of Vision Transformers (ViT) for improved feature learning
- Enhancement of fracture classification to identify multiple fracture types
- Integration of advanced object detection models like YOLO and SSD for faster localization
- Development of explainable AI techniques (e.g., Grad-CAM) for better interpretability
- Deployment as a mobile application for remote and on-field diagnosis
- Cloud-based system for centralized data processing and multi-user access
- Integration with hospital systems such as PACS for real-time clinical use
- Support for multiple imaging modalities such as CT scans and MRI
- Improvement in detection of very small and complex fractures
- Implementation of real-time video-based fracture analysis
- Incorporation of patient history and clinical data for better decision-making
- Optimization for low-resource devices and edge computing environments
- Continuous model updating using real-world medical data (incremental learning)

X. CONCLUSION

The proposed Bone Fracture Detection System demonstrates the effective application of deep learning techniques in medical image analysis. The system successfully integrates Convolutional Neural Networks (VGGNet) for classification and Faster R-CNN for fracture localization, enabling both identification and visualization of fractures in X-ray images. By automating the detection process, the system reduces dependency on manual diagnosis and minimizes the risk of human error, particularly in cases involving large volumes of medical data or subtle fracture patterns.

The experimental results indicate that the model is capable of achieving reliable accuracy while maintaining efficient processing suitable for real-time applications. The incorporation of preprocessing techniques and data augmentation further enhances the model' s ability to generalize across different image conditions. Additionally, the Streamlit-based user interface ensures accessibility and ease of use, allowing both technical and non-technical users to interact with the system effectively.

Overall, the proposed system highlights the potential of artificial intelligence in improving diagnostic efficiency and supporting healthcare professionals in clinical decision-making. Although it is not intended to replace expert medical judgment, it serves as a valuable decision-support tool. With further improvements and real-world validation, the system can be extended for broader clinical applications, contributing to faster, more accurate, and accessible healthcare solutions.

REFERENCES

1. K. Simonyan and A. Zisserman, " Very Deep Convolutional Networks for Large-Scale Image Recognition," , 2015.
2. S. Ren, K. He, R. Girshick, and J. Sun, " Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks," , vol. 39, no. 6, pp. 1137– 1149, 2017.
3. A. Krizhevsky, I. Sutskever, and G. E. Hinton, " ImageNet Classification with Deep Convolutional Neural Networks," , vol. 25, pp. 1097– 1105, 2012.
4. Y. LeCun, Y. Bengio, and G. Hinton, " Deep Learning," , vol. 521, no. 7553, pp. 436– 444, 2015.
5. O. Russakovsky et al., " ImageNet Large Scale Visual Recognition Challenge," , vol. 115, no. 3, pp. 211– 252, 2015.

6. G. Litjens et al., “ A Survey on Deep Learning in Medical Image Analysis,” , vol. 42, pp. 60–88, 2017.
7. J. Long, E. Shelhamer, and T. Darrell, “ Fully Convolutional Networks for Semantic Segmentation,” , 2015.
8. M. Abadi et al., “ TensorFlow: A System for Large-Scale Machine Learning,” , pp. 265– 283, 2016.
9. F. Chollet, “ Keras: The Python Deep Learning Library,” 2015. [Online]. Available: https://keras.io
10. G. Bradski, “ The OpenCV Library,” , 2000.
11. I. Goodfellow, Y. Bengio, and A. Courville, . Cambridge, MA, USA: MIT Press, 2016.
12. D. G. Lowe, “ Distinctive Image Features from Scale-Invariant Keypoints,” , vol. 60, no. 2, pp. 91– 110, 2004.