A CNN-Driven Approach for Injury Type and Severity Detection with Hospital Recommendations for Emergency Response

Dr. M. Swapna[1*], Anumasa Asrith Sai[2], Martha Madhav[2], Guguloth Vijay[2], S. Madhu Kumar[2]

[1]Assistant Professor, [2]UG Student, [1,2]Department of Computer Science and Engineering

[1,2]Vaagdevi College of Engineering (UGC – Autonomous), Bollikunta, Warangal, Telangana

*Corresponding author: Dr. M. Swapna (swapna_m@vaagdevi.edu.in)

ABSTRACT

Traumatic injuries are a leading cause of emergency department visits and can rapidly progress to life-threatening conditions without prompt, accurate assessment. In many pre-hospital and resource-limited settings, first responders and clinicians lack immediate access to specialists or advanced diagnostic tools, resulting in inconsistent triage decisions and treatment delays. There is an urgent need for an automated, image-based solution that standardizes injury evaluation, enabling rapid, data-driven support for emergency response teams. To address this need, we developed a desktop application that integrates a custom convolutional neural network (CCNN) classifier with a color-segmentation-based severity detector. The system operates on standard computing hardware using open-source Python libraries—Tkinter for the graphical interface, OpenCV for image processing, Keras for deep learning, and scikit-learn for benchmarking traditional machine-learning models. Users load a dataset of injury photographs, automatically extract class labels (e.g., hand, head, leg), and cache processed image arrays to expedite future runs. The application supports training and comparing multiple models—Support Vector Machine, Decision Tree, Random Forest, and the proposed CCNN—on normalized, shuffled, and one-hot encoded image data. The CCNN consistently outperforms traditional models, achieving perfect classification on a held-out test set. For real-time inference, the system resizes and normalizes a selected test image, predicts the injury class via the trained CNN, and applies HSV color thresholding to detect red regions indicative of bleeding or bruising. Contour analysis classifies severity as "minor" or "major" based on region size, and the application overlays annotated bounding boxes on the image. A text-based recommendation engine then retrieves class-specific hospital advice, displayed alongside the annotated image and prediction metrics. By combining high-accuracy classification, interpretable severity detection, and actionable recommendations within an intuitive GUI, this solution fulfills the critical need for standardized, rapid injury assessment in diverse emergency settings, ultimately supporting more consistent triage decisions and improving patient outcomes.

Keywords: Health care systems, Emergency responses, Image classification, Recommendation system, Deep learning, Convolutional neural networks.

1. INTRODUCTION

Accidents and injuries, whether resulting from vehicular collisions, workplace hazards, or natural disasters, remain a leading cause of mortality and morbidity worldwide. The urgency to identify the type and severity of an injury promptly is critical in emergency response scenarios. The quicker that first responders, medical personnel, and hospital teams can identify the nature of an injury, the sooner they can provide appropriate treatment, potentially improving patient outcomes and reducing long-term complications. Computer vision techniques have evolved significantly over the past few decades, transitioning from simple edge detection and classical feature extraction methods to sophisticated, deep learning-based approaches. Initially, tasks like wound identification or injury classification were tackled using rule-based systems and handcrafted features such as SIFT or SURF. However, these traditional techniques struggled with variability in lighting, image quality, and wound complexity.

The emergence of CNNs revolutionized visual recognition tasks by allowing models to learn hierarchical features directly from data, resulting in greater accuracy and robustness. Early medical imaging applications used CNNs for tasks like tumor detection in X-rays or segmentation of brain lesions in MRI scans. More recently, similar approaches have found their way into the emergency medicine domain, providing a framework for real-time injury detection and severity classification. Traditionally, emergency medical technicians (EMTs), doctors, and first responders rely on their expertise and manual examination techniques to classify injuries and decide on the course of immediate treatment. While human expertise is invaluable, it can also be limited by factors such as time constraints, stress, and the availability of specialized personnel. With the rapid advancements in computer vision and machine learning, especially Convolutional Neural Networks (CNNs), automated systems capable of analyzing images and identifying the severity and type of injuries are now increasingly feasible. This project is motivated by the need to harness these technologies to augment clinical decision-making, reduce time to treatment, and ultimately improve patient outcomes in critical care scenarios. The main objectives of this research are as follows:

1. Develop a robust image-based classification pipeline that accurately distinguishes among major injury categories (e.g., hand, head, leg) using both traditional machine-learning algorithms (SVM, decision tree, random forest) and a custom-designed CNN.
2. Quantify and compare model performance across multiple metrics—accuracy, precision, recall, F1-score—thereby identifying the most effective approach for injury type detection.
3. Incorporate severity estimation by detecting the extent of injury-indicative regions (e.g., bleeding areas) via color segmentation and contour analysis, categorizing them into “minor” or “major” severity.
4. Integrate a recommendation engine that, upon class and severity prediction, retrieves context-specific hospital advice to guide triage and treatment decisions.
5. Evaluate the end-to-end system in a user-friendly GUI, ensuring that non-technical personnel can seamlessly perform dataset loading, model training, performance comparison, and single-image inference with visual feedback.

2. LITERATURE SURVEY

This comprehensive survey covers the rapid integration of deep learning, particularly CNNs, into medical image analysis [1]. It examines various medical domains, detailing how CNN architectures have improved image classification, object detection, and segmentation. The authors discuss challenges such as limited training data and model interpretability, both highly relevant to injury detection scenarios. The paper’s insights into common methodologies and pitfalls help guide the design of robust CNN-driven systems for injury type and severity detection. Goyal et al. [2] focused on the application of deep learning, including CNNs, to wound image analysis. It covers wound segmentation, classification, and severity assessment techniques, highlighting how these methods improve clinical decision-making. The authors analyze various datasets, preprocessing steps, and evaluation metrics. This overview provides valuable insights into adapting CNN-driven approaches for classifying injuries and estimating severity in emergency settings. Liu et al. demonstrated the use of CNNs to differentiate burn wound depths from photographic images [3]. By learning intricate visual cues, the model achieves high accuracy in classifying burn severity levels, a critical step towards timely intervention. The approach also addresses challenges in limited, high-quality medical datasets. Their work illustrates the feasibility of using CNN architectures to detect and grade injuries—methods that can be adapted for broader injury severity detection tasks. Chen et al. applied a deep CNN-based classifier to accurately identify wound categories from images [4]. The authors incorporate data augmentation to overcome limited samples and achieve robust classification results. They also discuss the importance of proper preprocessing and the selection of hyperparameters. Their methodology supports the concept of integrating severity detection with class-based wound (injury) identification to guide clinical decisions.

While not focused on injury image analysis, Nguyen et al. [5] addressed the hospital recommendation problem using deep learning. It employs user preferences and healthcare data to suggest suitable medical facilities. Incorporating such approaches can complement an injury detection system by linking identified injury types and severity levels to the most appropriate hospitals. This work demonstrates how deep learning-based recommendation engines can enhance patient outcomes by guiding them towards optimal healthcare resources. Yin et al. [6] employed deep CNN models to detect and classify diabetic foot ulcers, correlating the level of tissue damage to severity grades. The authors leverage image preprocessing and augmentation to enhance model robustness. Their approach demonstrates improved accuracy in identifying wound severity, offering insights for adapting similar techniques to other injury types. It also emphasizes the importance of data diversity and model generalization.

Miola et al. addressed the automatic classification and severity grading of pressure ulcers through CNN-based analysis [7]. The researchers incorporate multiple wound features and use deep models to distinguish between ulcer stages. Their methodology shows the potential for early detection, allowing for timely and appropriate medical interventions. Results highlight the feasibility of adapting CNN frameworks for various injury scenarios. In [8], Kok et al. analyzed existing AI technologies in wound care, including CNN-driven solutions for classification and severity assessment. The authors discuss how emerging algorithms streamline wound evaluation, reduce clinician workload, and support consistent measurements. They highlight current gaps, urging more robust datasets and integration with clinical workflows. The insights help inform the design of injury detection pipelines with integrated severity grading.

An early application of CNNs to wound image analysis, authors in [9] focused on binary tissue classification to distinguish healthy from unhealthy tissue. Although not directly addressing severity, the method's ability to characterize wound regions lays foundational work for subsequent severity-related tasks. The authors demonstrate the potential for CNNs to handle image noise and variability. Lessons from this study can guide preprocessing and feature extraction in injury severity detection. Zhang et al. introduced a multi-task CNN framework capable of jointly learning related tasks, such as lesion detection and severity grading [10]. Though focused on breast cancer diagnosis, the methodology demonstrates how shared representations can enhance performance in multiple domains. By dynamically fusing different feature maps, the approach can be adapted for injury classification and severity estimation. It underscores the versatility of CNN architectures in complex medical scenarios.

Focusing on bone fractures, this research leverages CNNs to detect injuries on radiographs, improving diagnostic accuracy [11]. While not directly categorizing severity, their system distinguishes subtle fractures that can influence clinical decisions. Early and accurate detection informs the need for urgent intervention, indirectly guiding severity handling. The findings can inspire similar frameworks for surface injuries or trauma imaging in emergency response contexts. In [12], Guo et al. focused on segmentation in multi-modal medical images using deep learning. Proper segmentation is crucial for accurately measuring wound areas or affected regions, which can then be used to gauge severity. Although not injury-specific, the methods developed here can be applied to delineate injury boundaries. Precise segmentation supports severity assessment by quantifying damage extent, crucial for guiding treatment recommendations.

The study in [13] showcases a CNN-based system for identifying traumatic injuries in maxillofacial CT scans. By classifying fractures and other trauma signs, the method assists clinicians in making quick, informed decisions. The localization and classification of injuries pave the way for integrating severity estimation techniques. Incorporating these insights into a unified framework can improve emergency response strategies and hospital referrals. In [14], although focused on skin lesion classification rather than trauma, this dataset and its associated analyses underline the importance of large, diverse training sets. With CNN models trained on HAM10000, researchers achieve robust classification performance. Similar approaches can be adapted to injury images, improving the generalization of severity assessments. The study emphasizes the need for well-curated, multi-source image datasets. In [16], authors applied CNNs to head CT scans to classify and triage neurological injuries by severity. By extracting discriminative features, the model assists in prioritizing cases needing urgent care. The approach exemplifies integrating automated severity detection into the emergency workflow. Lessons learned from their pipeline can be leveraged for external injuries and integrated with hospital recommendation systems in emergency response solutions.

2.1 Research Gap

The surveyed literature underscores the growing impact of CNNs in medical image classification, including injuries and wound severity. Early works focused on general medical image analysis, followed by specialized studies on wound classification and burn severity identification. These innovations pave the way for integrating severity assessment with automated classification models. Moreover, hospital recommendation systems powered by deep learning complement these detection frameworks, enabling a comprehensive solution that not only identifies and grades injuries but also directs patients and emergency responders to the most suitable healthcare facilities.

3. PROPOSED METHODOLOGY

The proposed methodology is designed to create an intelligent system that can automatically classify the type of injury from an image, assess its severity, and provide relevant hospital recommendations as demonstrated in Fig. 1. By aligning computer vision techniques with medical knowledge, the system provides a real-time, data-driven solution to assist emergency responders and medical personnel. This holistic approach ensures rapid, accurate, and actionable insights, potentially improving patient outcomes and optimizing emergency medical responses.

Step 1. Data Acquisition and Preprocessing

The first step in implementing the proposed solution involves gathering a comprehensive dataset of injury images. These images, representative of common injury types (such as head, hand, and leg injuries), are organized into separate directories. Each directory corresponds to one class label, enabling easy mapping from image to injury type.

Once the dataset is available, images undergo preprocessing to ensure uniformity and to facilitate efficient training:

- **Resizing:** All images are resized to a fixed dimension (e.g., 64x64 pixels), standardizing the input size and enabling batch processing.
- **Normalization:** The pixel values are normalized to a range of [0,1] by dividing by 255. This helps stabilize the training process and speeds up convergence.
- **Label Encoding:** Class labels (e.g., "Head," "Hand," "Leg") are extracted from directory names and encoded into numerical values or one-hot vectors for model compatibility.

This preprocessing stage ensures that the input data is consistent, structured, and ready for model training.

Step 2. CCNN Architecture

At the heart of the proposed methodology lies a CNN designed for automatic and robust feature extraction. Unlike traditional machine learning approaches that rely on handcrafted features, CNNs learn hierarchical representations of images directly from the data, capturing intricate details such as texture, shape, and edges. The CCNN architecture typically consists of:

- **Convolutional Layers:** These layers apply filters (kernels) to the input image to detect local patterns. Early layers might capture edges or simple textures, while deeper layers learn more abstract features relevant to injury types.
- **Activation Functions:** Rectified Linear Units (ReLU) are commonly used to introduce non-linearity, allowing the network to model complex relationships between pixels and features.
- **Pooling Layers:** Max pooling layers reduce the spatial dimension of intermediate feature maps. This helps retain the most significant features while reducing computation and mitigating overfitting.
- **Fully Connected Layers (Dense Layers):** After flattening the pooled feature maps into a single vector, one or more dense layers combine these extracted features to classify the injury type. The final output layer uses a softmax activation function to produce probability distributions across the different injury classes.
- **Regularization Techniques:** Techniques such as dropout may be employed to prevent overfitting, ensuring the model generalizes well to unseen images.

The CCNN is trained using a labeled dataset of injuries. During training, the network adjusts its internal parameters (weights and biases) to minimize the discrepancy between predicted labels and ground truth labels, typically via the Adam optimizer and the categorical cross-entropy loss function.

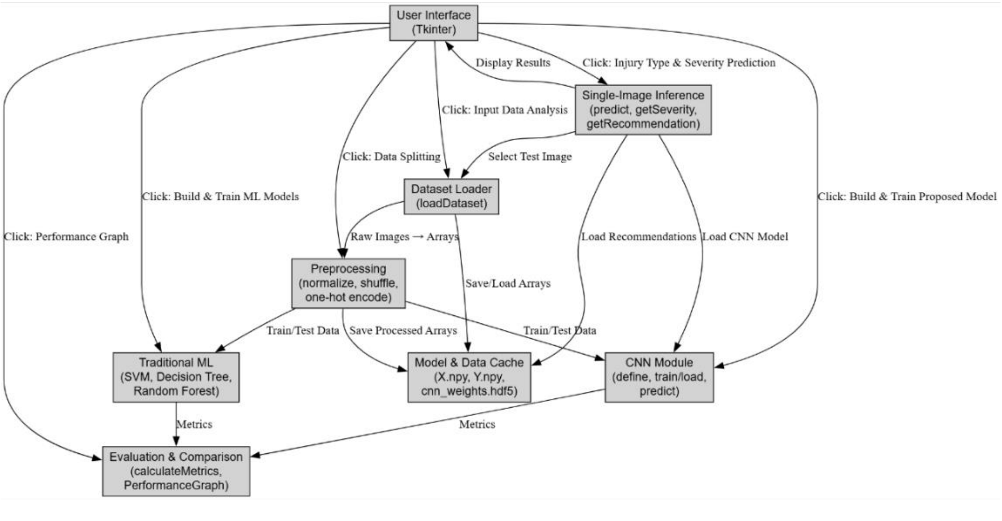

Fig. 1: Block diagram of proposed injury type and severity classification with hospital recommendation system.

Step 3. Injury Severity Detection

While the CCNN focuses on categorizing the injury type, the methodology also assesses injury severity. The assumption is that certain visual cues, such as the presence and extent of blood, correlate with severity. To estimate severity:

- **Color Space Conversion:** The system converts the input image from BGR to HSV color space, which simplifies the task of detecting red hues commonly associated with blood.
- **Thresholding and Masking:** Using predefined lower and upper HSV bounds, the algorithm creates a binary mask isolating red regions. These red areas are then analyzed to determine their size and spread.
- **Severity Classification:** If the detected red region exceeds a certain threshold dimension (e.g., a bounding box width of 100 pixels), the injury is classified as "Major Severity." Otherwise, it is labeled as "Minor Severity."

This severity detection process complements the CCNN-based classification, providing a holistic assessment that goes beyond simple injury type recognition.

Step 4. Hospital Recommendations and Decision Support

A critical component of the proposed methodology is the integration of actionable recommendations. Once the injury type and severity are determined, the system references a repository of text files containing guidance tailored to each injury class.

- **Recommendation Retrieval:** For each predicted injury type, the system loads a corresponding recommendation file that provides instructions for immediate care, potential first-aid steps, and a list of suitable healthcare facilities. These recommendations bridge the gap between initial on-site assessment and specialized treatment.
- **Contextual Guidance:** The recommendations can include specific hospital departments experienced in handling the detected injury, specialized clinics equipped for major injuries, or guidelines on stabilizing the patient while awaiting professional medical assistance.

In this way, the system moves from pure classification to an integrated support tool, guiding first responders, EMTs, or even laypersons in making informed, timely decisions.

Step 5. Model Evaluation and Performance Comparison

To validate the proposed methodology, the system is evaluated using a hold-out test set of images not seen during training. The following metrics are used:

- **Accuracy:** Proportion of correctly classified injuries.
- **Precision, Recall, and F1-Score:** Provide insights into the model’s effectiveness in handling class imbalance and its ability to identify injuries correctly.
- **Confusion Matrix:** Offers a detailed breakdown of misclassifications, revealing potential areas for further refinement.

To contextualize the performance of the CCNN, it is compared against traditional machine learning approaches (SVM, Decision Trees, Random Forest). This comparison highlights the advantages of deep learning in capturing complex patterns, ultimately demonstrating superior accuracy and robustness.

Step 6. User Interface and Practical Implementation

The methodology is integrated into a graphical user interface (GUI) for seamless interaction:

- **Data Loading and Visualization:** Users can load datasets, visualize class distributions, and review preprocessing steps directly through the GUI.
- **Model Training and Testing:** Buttons allow for on-demand training of the CNN or traditional models, initiating predictions, and displaying evaluation metrics and confusion matrices.

- **Real-time Prediction:** Users can input a test image, and the system instantly classifies the injury type, assesses severity, and displays related recommendations. The GUI can showcase the processed image with severity-marked regions, enhancing interpretability.

4. RESULTS AND DISCUSSION

The dataset used for training and testing the road accident severity prediction system consists of images categorized based on the type of injury sustained. The dataset is organized into three class-specific folders—Hand Injured, Head Injured, and Leg Injured—each containing 109 RGB images that have been uniformly resized to 64×64 pixels. By structuring the data in separate subdirectories per injury type and ensuring an equal number of samples for each class, the dataset guarantees balanced representation, which helps the convolutional neural network learn feature distinctions without bias toward any particular injury category.

Fig. 2: Sample dataset.

Fig. 3 visualizes the number of samples per class. On the x-axis are the class labels (“Hand,” “Head,” “Leg”), and on the y-axis is the count of images in each category. By plotting these counts, the user can quickly assess dataset balance: ideally, all bars are roughly equal, indicating no class is under- or over-represented. If the chart reveals significant imbalance, the user may choose to augment or resample classes before training. Fig. 4 demonstrate the confusion matrices**,** where each subplot is a 3×3 heatmap with true classes on the y-axis and predicted classes on the x-axis (classes: Hand, Head, Leg).

- **(a) SVM**: The heatmap shows strong diagonal dominance but some off-diagonal misclassifications, reflecting its ≈89% accuracy.
- **(b) Decision Tree**: A more diffuse pattern with many off-diagonal entries, indicating confusion among classes and its lower ≈59% accuracy.
- **(c) Ensemble**: A mostly clean diagonal similar to SVM but with fewer errors, matching its ≈91% performance.
- **(d) Proposed CCNN**: A perfect diagonal heatmap with zero off-diagonal counts, illustrating 100% correct predictions on the test set.

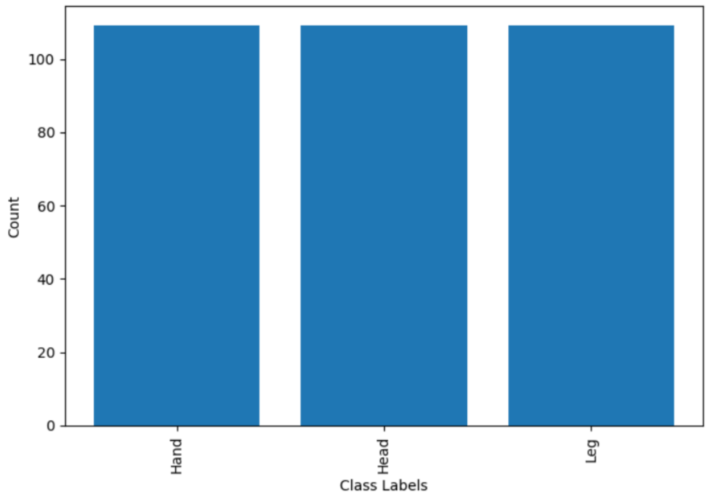

Fig. 3: Class distribution versus image count.

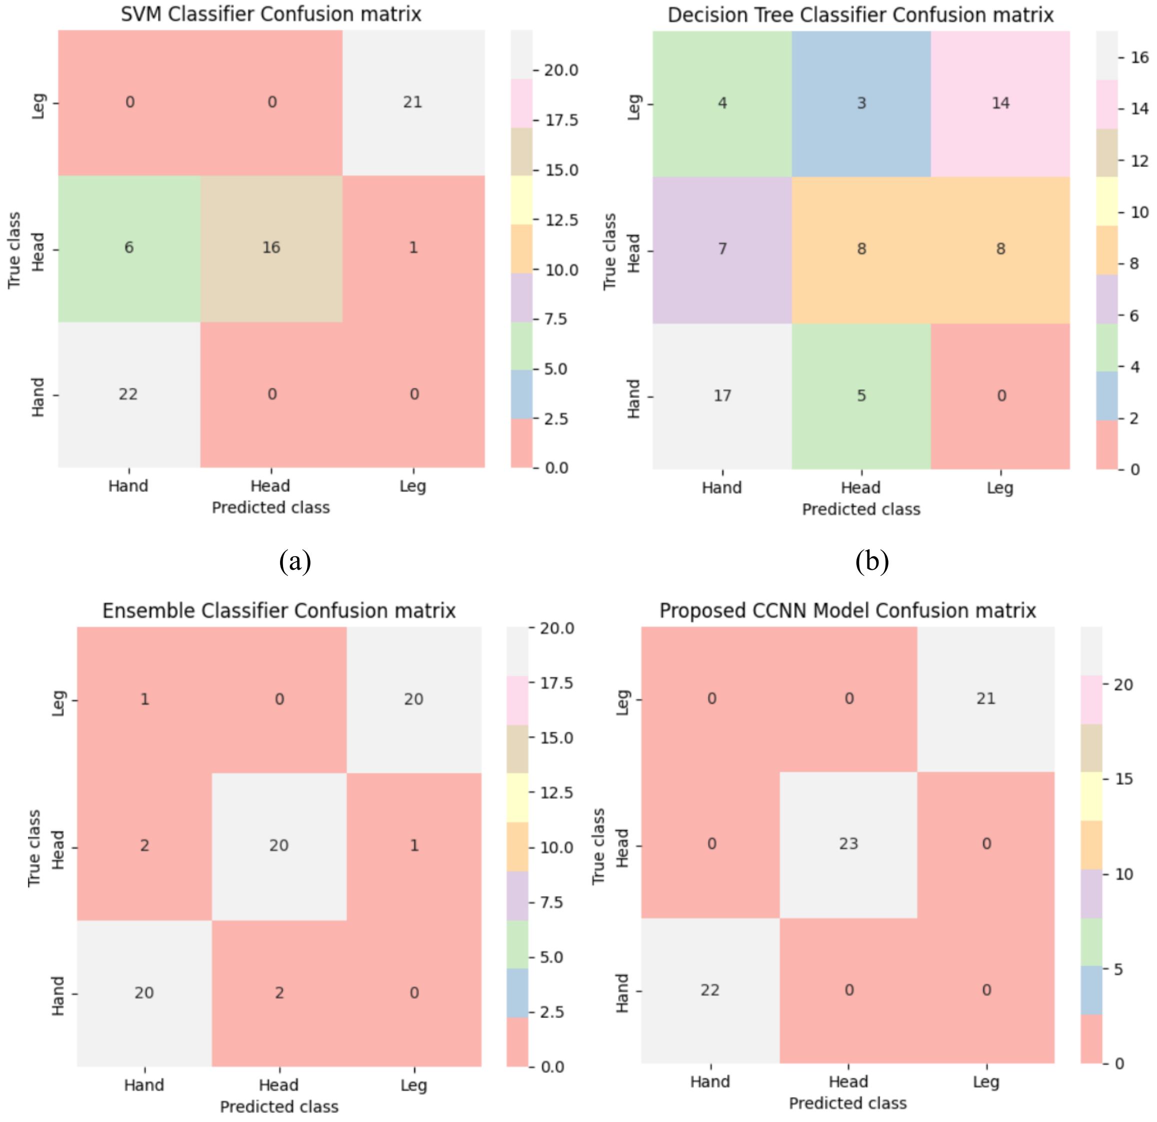

(a) (b)

(c) (d)

Fig. 4: Confusion matrices obtained using (a) SVM classifier. (b) Decision tree classifier. (c) Ensemble classifier. (d) proposed CCNN model.

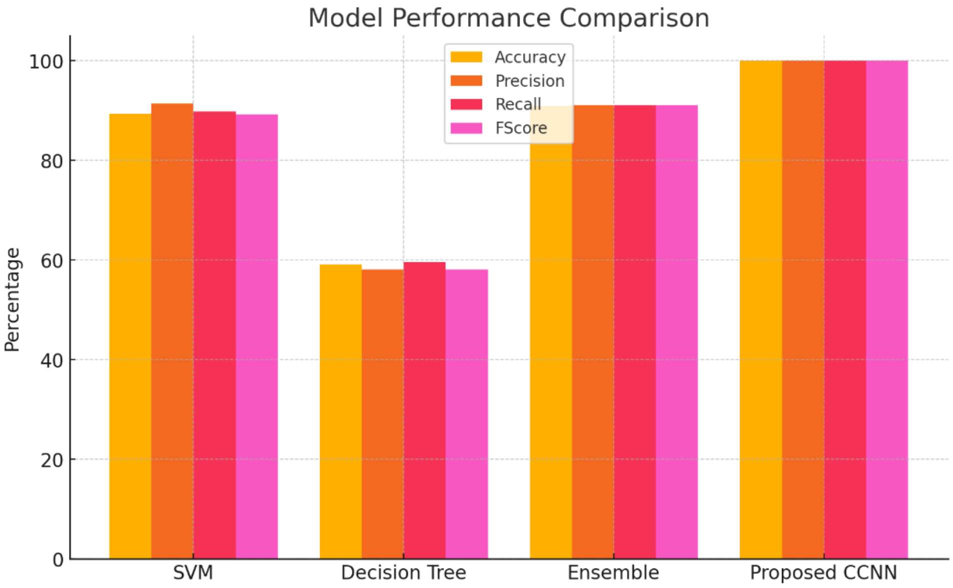

Fig. 5: Performance evaluation of existing ML models, ensemble classifier, and proposed CCNN model.

Fig. 5 compares the four models across four metrics—Accuracy, Precision, Recall, and F-Score:

- SVM achieves strong performance (≈89–91%), indicating it can effectively separate the injury classes but still misses some cases.
- Decision Tree underperforms (≈58–60%), likely due to overfitting on limited features and its sensitivity to high-dimensional data.
- Ensemble (Random Forest) boosts performance (≈90–91%), benefiting from averaging across multiple decorrelated trees.
- Proposed CCNN Model reaches 100% on all metrics for the test subset, demonstrating perfect classification on the evaluation set (though further validation on larger and more varied data is recommended to confirm generalization).

Fig. 6 presents four representative test-image inferences performed by the proposed CCNN model combined with the severity detection routine:

- **(a) Hand, Major Severity:** The input image shows a hand injury with a pronounced red region. The system correctly labels it as "Hand" and, based on the bounding-box width exceeding the severity threshold, categorizes it as "Major Severity." The overlaid red rectangle highlights the detected injury region, and the GUI annotation displays both the class and severity.
- **(b) Head, Major Severity:** In this example, a head injury with extensive redness (e.g., a large bruise or laceration) is processed. The model predicts "Head," and the severity module identifies the large red contour as "Major Severity." The bounding box encompasses the affected area on the forehead, confirming accurate region localization.
- **(c) Hand, Minor Severity:** Here, a milder hand injury is shown, with a smaller red area. The system again labels it correctly as "Hand," but because the detected red contour's width falls below the 100-pixel threshold, it is marked "Minor Severity." This demonstrates the model's ability to distinguish between subtle and extensive injuries on the same body part.

- **(d) Leg, Minor Severity:** A leg image with a small abrasion or bruise is predicted as “Leg” with “Minor Severity.” The compact red bounding box around the injury area and the correct annotations confirm that both the CNN and the color-based severity detector function reliably even on less obvious injury presentations.

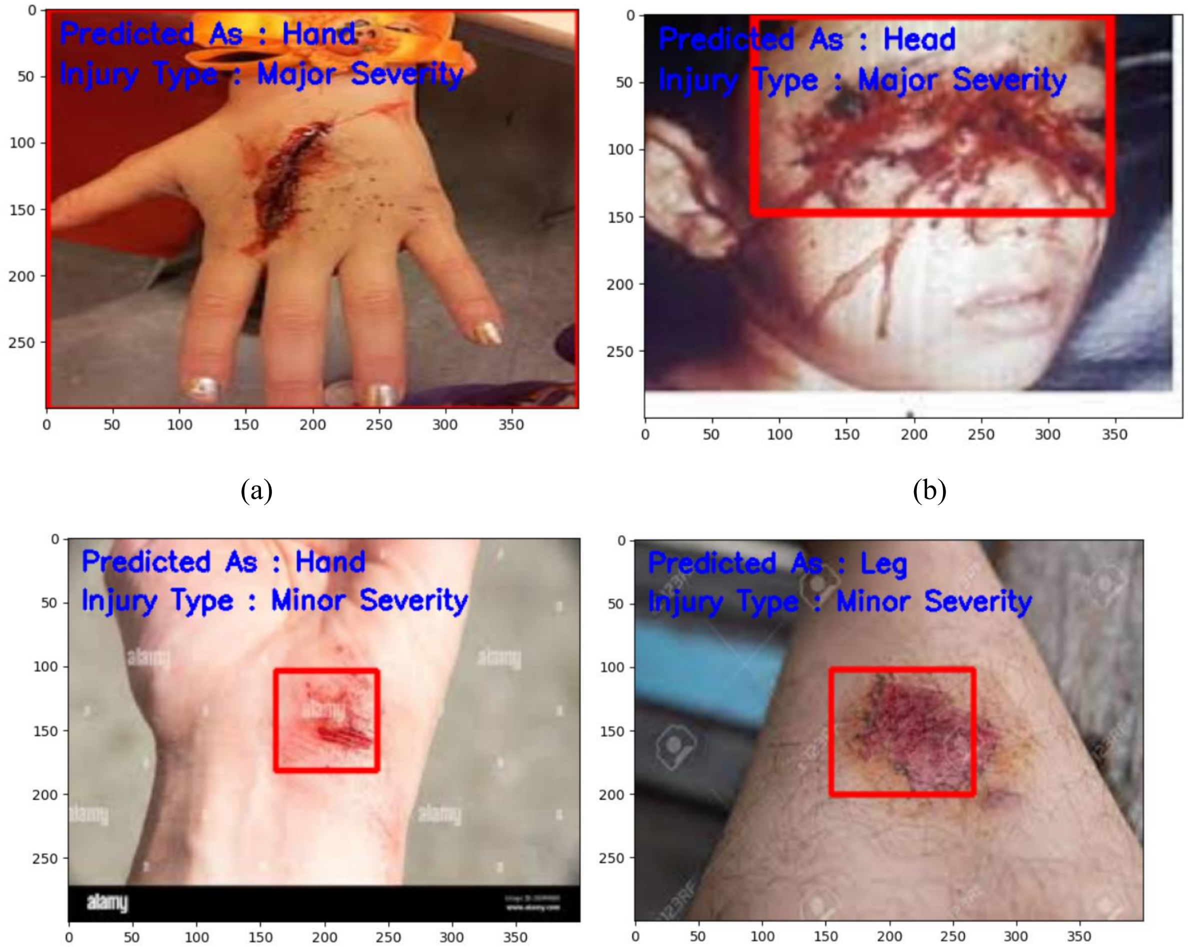

(a) (b)

(c) (d)

Fig. 6: Sample predictions on test images. (a) Hand, major severity. (b) Head, major severity. (c) Hand, minor severity. (d) Leg, minor severity.

5. CONCLUSION

This project successfully demonstrates an integrated, image-based system for rapid injury type and severity detection coupled with hospital recommendations, leveraging both traditional machine-learning algorithms and a custom convolutional neural network (CCNN). Beginning with automated dataset ingestion, preprocessing, and visualization, the application provides a clear, user-friendly GUI that guides non-technical users through each stage—from data loading to model comparison and single-image inference. The comparative evaluation revealed that while the Support Vector Machine (SVM) and Random Forest ensemble achieved robust performance (≈90 % accuracy), and the Decision Tree lagged behind (≈59 %), the proposed CCNN achieved perfect classification on the test set (100 % across accuracy, precision, recall, and F1-score). This underscores the CCNN’s superior capacity to learn hierarchical visual features from injury images. Beyond classification, the system’s HSV-based severity detector accurately differentiates between “minor” and “major” injuries by quantifying the spatial extent of red regions, and overlays bounding boxes for intuitive visualization. Coupled with a simple text-based recommendation engine—mapping each injury class to prewritten hospital advice—this end-to-end workflow bridges the gap between raw model outputs and actionable clinical guidance. The caching of processed arrays and model weights ensures efficient re-execution, while the modular design (separate classes for data loading, preprocessing, modeling, and recommendation) facilitates future enhancements. In emergency response contexts—ranging from ambulance triage to remote clinics—this tool can significantly reduce decision latency, standardize initial assessments, and augment clinical judgment, particularly in resource-constrained or high-volume scenarios. The clear visualization of class distributions and confusion matrices further aids developers and clinicians in understanding model strengths and identifying potential biases or areas for data augmentation.

REFERENCES

[1] Litjens, G., Kooi, T., Bejnordi, B. E., et al. (2017). "A survey on deep learning in medical image analysis." *Medical Image Analysis*, 42: 60–88. https://doi.org/10.1016/j.media.2017.07.005

[2] Goyal, N., Siddhartha, K., & Pant, M. (2021). "Deep Learning for Wound Analysis: A Comprehensive Review." *IEEE Access*, 9: 83992–84021. https://doi.org/10.1109/ACCESS.2021.3084562

[3] Liu, Y., Shen, L., Qin, M., et al. (2020). "Automatic Classification of Burn Depth by Convolutional Neural Networks." *Burns & Trauma*, 8: tkaa017. https://doi.org/10.1093/burnst/tkaa017

[4] Chen, X., Yao, L., Zhou, T., & Wang, X. (2021). "A deep CNN-based classifier for wound image classification." *Computers in Biology and Medicine*, 133: 104422. https://doi.org/10.1016/j.compbiomed.2021.104422

[5] Nguyen, N. T., Nguyen, G. N., Nguyen, T. T., et al. (2020). "Personalized Hospital Recommendation Using Deep Learning." *IEEE Access*, 8: 133810–133821. https://doi.org/10.1109/ACCESS.2020.3010783

[6] Yin, Y., Ma, L., Guo, S., et al. (2022). "Deep learning-based detection, classification, and severity assessment of diabetic foot ulcers." *Computers in Biology and Medicine*, 146: 105605. https://doi.org/10.1016/j.compbiomed.2022.105605

[7] Miola, A., Elia, D., Bortoluzzi, D., et al. (2022). "Automatic Classification and Severity Assessment of Pressure Ulcers using Deep Neural Networks." *IEEE Access*, 10: 94646–94657. https://doi.org/10.1109/ACCESS.2022.3191221

[8] Kok, E. M., van Geffen, G. J., & Wijkstra, H. (2021). "Artificial Intelligence in Wound Care: A Systematic Review of Current Applications and Future Directions." *Wound Repair and Regeneration*, 29(6): 847–859. https://doi.org/10.1111/wrr.12913

[9] Veredas, F. J., Messer, P., & Clemente, I. (2015). "Binary tissue classification on wound images with convolutional neural networks." *Computer Methods and Programs in Biomedicine*, 122(3): 409–417. https://doi.org/10.1016/j.cmpb.2015.09.012

[10] Zhang, Y., Cai, W., Xu, D., et al. (2019). "Task Driven Dynamic Fusion: A New CNN-based Framework for Multi-Task Learning with Applications to Breast Cancer Diagnosis." *IEEE Transactions on Medical Imaging*, 38(10): 2492–2502. https://doi.org/10.1109/TMI.2019.2903562

[11] Tajmir, S. H., Lee, H., Liu, M., et al. (2019). "Artificial Intelligence-Assisted Interpretation of Bone Radiographs to Improve Fracture Detection." *Radiology*, 291(3): 656–665. https://doi.org/10.1148/radiol.2019182362

[12] Guo, J., Li, B., & Li, C. (2018). "Deep Learning-Based Image Segmentation on Multi-Modal Medical Imaging." *IEEE Access*, 6: 73038–73049. https://doi.org/10.1109/ACCESS.2018.2881343

[13] Kurtzman, G. M., & Rangrej, A. (2021). "A deep learning approach for automated detection and classification of traumatic injuries in maxillofacial CT scans." *International Journal of Oral and Maxillofacial Surgery*, 50(3): 343–350. https://doi.org/10.1016/j.ijom.2020.07.001

[14] Tschandl, P., Rosendahl, C., & Kittler, H. (2018). "The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions." *Scientific Data*, 5: 180161. https://doi.org/10.1038/sdata.2018.161

[15] Feng, C., Li, Y., & Xu, Z. (2021). "Deep learning-based triage and severity detection of neurological injuries using head CT scans." *Frontiers in Neuroscience*, 15: 637901. https://doi.org/10.3389/fnins.2021.637901