



A DEEP LEARNING-BASED FRAMEWORK FOR AUTOMATED KNEE OSTEOARTHRITIS DETECTION AND KELLGREN-LAWRENCE GRADING

Dr. V. HARIPRIYA¹, Ms. S. SARVATHA²

¹Assistant Professor, ²UG Student

^{1,2}Department of Computer Technology and Data Science, Sri Krishna Arts and Science College, Coimbatore

¹vyharipriya.24@gmail.com, ²sarvatha9486@gmail.com

Abstract— Knee osteoarthritis (KOA) is a prevalent degenerative joint disorder and one of the primary causes of pain, reduced mobility, and disability among the aging population worldwide. Early and accurate diagnosis is essential for timely clinical intervention; however, conventional radiographic assessment based on the Kellgren–Lawrence (KL) grading system is often labor-intensive and subject to inter- and intra-observer variability. To address these limitations, this study proposes a fully automated two-stage deep learning framework for knee osteoarthritis detection and severity grading using plain radiographic images.

In the first stage, a YOLOv5 object detection model is employed to accurately localize the tibiofemoral joint region from full knee X-ray images, thereby eliminating irrelevant anatomical structures and focusing on the clinically significant region of interest (ROI). In the second stage, the extracted

ROI is resized to 224×224 pixels and classified using a fine-tuned DenseNet201 convolutional neural network to predict one of the five Kellgren–Lawrence grades (Grade 0–Grade 4), representing the severity of osteoarthritis.

The proposed framework integrates PyTorch-based object detection with TensorFlow-based classification and is deployed as an interactive Streamlit web application, enabling real-time clinical decision support. Experimental results demonstrate that the proposed pipeline achieves precise joint localization and robust classification performance across all KL grades, highlighting its effectiveness for automated osteoarthritis assessment. By combining accurate region localization with deep feature learning, the framework offers a scalable, cost-effective, and clinically applicable solution that can assist radiologists and orthopedic specialists in improving diagnostic consistency, reducing



assessment time, and facilitating early disease management.

Key Words— Knee osteoarthritis (KOA), YOLOv5, DenseNet201, Kellgren–Lawrence grading.

1. INTRODUCTION

Knee osteoarthritis (KOA) is one of the most common degenerative joint diseases and a leading cause of pain, stiffness, and physical disability worldwide. The disease primarily affects the articular cartilage, resulting in progressive cartilage degeneration, narrowing of the joint space, osteophyte formation, and structural changes in the surrounding bone. As global life expectancy increases and the prevalence of obesity continues to rise, the number of individuals affected by KOA is expected to grow significantly, placing a substantial burden on healthcare systems and reducing the quality of life of millions of people. Early diagnosis and accurate assessment of disease severity are therefore essential for selecting appropriate treatment strategies, slowing disease progression, and improving patient outcomes.

Radiographic imaging remains the most widely used and cost-effective technique for diagnosing

knee osteoarthritis in routine clinical practice. Among the available assessment methods, the Kellgren–Lawrence (KL) grading system is considered the clinical standard for evaluating the severity of KOA. The system categorizes knee osteoarthritis into five grades, ranging from Grade 0, which indicates the absence of osteoarthritic changes, to Grade 4, which represents severe joint degeneration characterized by pronounced osteophyte formation, marked joint space narrowing, subchondral sclerosis, and bone deformity. Although the KL grading system provides a standardized framework for disease assessment, its interpretation largely depends on the expertise and experience of radiologists. Consequently, manual grading is often time-consuming and may lead to inter-observer and intra-observer variability, affecting the consistency and reliability of clinical diagnosis.

In recent years, artificial intelligence (AI), particularly deep learning, has emerged as a promising solution for automating medical image analysis. Convolutional Neural Networks (CNNs) have demonstrated remarkable success in extracting complex visual features from radiographic images, enabling automated disease detection and classification with performance comparable to that of clinical experts. Several studies have applied deep learning techniques to classify knee osteoarthritis severity directly from



X-ray images. However, whole-knee radiographs often contain large background regions, including soft tissues and surrounding anatomical structures that are not directly relevant to osteoarthritis assessment. These irrelevant regions can introduce noise into the learning process, reduce classification accuracy, and increase computational complexity.

To overcome these limitations, recent research has focused on region-based analysis, where the clinically significant tibiofemoral joint region is first identified before severity classification is performed. By isolating the joint region, the classification model can concentrate on disease-related structural changes such as cartilage loss, osteophyte formation, and joint space narrowing, thereby improving diagnostic performance. This localization-based strategy also enhances the interpretability of automated systems by ensuring that predictions are derived from clinically meaningful anatomical regions.

Motivated by these observations, this study proposes a fully automated two-stage deep learning framework for knee osteoarthritis detection and severity grading from plain radiographic images. In the first stage, a YOLOv5 object detection model is employed to accurately localize the tibiofemoral joint region from full knee X-ray images. The detected region of interest is then extracted and resized before

being passed to a fine-tuned DenseNet201 convolutional neural network, which classifies the image into one of the five Kellgren–Lawrence grades (Grade 0–Grade 4). By separating localization and classification into two specialized stages, the proposed framework minimizes the influence of irrelevant image content while improving the robustness and accuracy of disease severity prediction.

To facilitate practical clinical use, the proposed framework is deployed as an interactive web-based application using Streamlit, allowing clinicians to upload radiographic images and receive real-time visualization of detected knee joints together with automated KL grade predictions. Such a system has the potential to support radiologists and orthopedic specialists by reducing diagnostic workload, improving grading consistency, and providing rapid decision support in routine clinical practice.

2. LITERATURE REVIEW

The application of deep learning in medical image analysis has significantly transformed the diagnosis of musculoskeletal disorders, particularly knee osteoarthritis (KOA). Conventional computer-aided diagnosis systems primarily relied on handcrafted image features combined with traditional machine learning algorithms. Although these methods demonstrated promising results, their



performance was often limited by the quality of manually extracted features and their inability to capture complex structural variations in knee radiographs. With the emergence of deep learning, especially Convolutional Neural Networks (CNNs), automated feature extraction has enabled more accurate and reliable knee osteoarthritis detection and severity grading.

Several researchers have employed deep learning models to classify knee osteoarthritis directly from plain radiographs using the Kellgren–Lawrence (KL) grading system. Early studies utilized well-established CNN architectures such as VGGNet, ResNet, and InceptionNet, which were trained on full knee radiographs to predict OA severity. These models demonstrated considerable improvements over traditional machine learning methods by automatically learning hierarchical image representations. However, distinguishing between adjacent KL grades, particularly Grades 1 and 2, remained challenging due to the subtle radiographic differences associated with early-stage osteoarthritis.

To improve classification performance, researchers introduced advanced deep learning techniques such as transfer learning, attention mechanisms, and ensemble learning. Transfer learning enabled pretrained models to adapt effectively to relatively small medical imaging

datasets, while attention mechanisms allowed networks to focus on clinically significant regions within knee radiographs. Ensemble approaches further enhanced prediction accuracy by combining the strengths of multiple deep learning models, thereby reducing classification errors and improving robustness. Despite these advancements, many studies continued to process complete radiographs, where irrelevant anatomical structures and background information often affected the learning process and reduced model efficiency.

2.1 Knee Joint Localization Using Object Detection

Accurate localization of the tibiofemoral joint region has become an important preprocessing step in automated knee osteoarthritis diagnosis. Instead of analyzing the entire radiograph, recent studies first identify the region of interest (ROI) containing the knee joint before performing disease classification. This approach minimizes the influence of surrounding anatomical structures and enables classification models to concentrate on clinically relevant features such as joint space narrowing, osteophyte formation, and bone sclerosis.

Various object detection algorithms have been investigated for automatic knee joint localization. Earlier methods employed Faster R-CNN and



Single Shot Detector (SSD) to detect the joint region with satisfactory accuracy. More recently, the You Only Look Once (YOLO) family of object detectors has gained widespread attention because of its superior balance between detection accuracy and computational efficiency. Among these, YOLOv5 offers high-speed inference, accurate localization, and lightweight deployment, making it particularly suitable for real-time clinical applications and web-based diagnostic systems. By automatically extracting the tibiofemoral joint region, YOLOv5 reduces background noise and provides standardized input for subsequent classification models.

2.2 DenseNet-Based Classification for Knee Osteoarthritis Grading

Among the numerous CNN architectures proposed for medical image classification, **Dense** Convolutional Networks (DenseNet) have demonstrated remarkable performance in tasks requiring fine-grained image discrimination. Unlike conventional CNNs, DenseNet establishes direct connections between all preceding and subsequent layers, allowing feature maps to be reused throughout the network. This dense connectivity promotes efficient information flow, alleviates the vanishing gradient problem, strengthens feature propagation, and reduces the number of trainable parameters.

Several studies have successfully applied DenseNet121, DenseNet169, and DenseNet201 for medical image classification, including musculoskeletal disease diagnosis. In knee osteoarthritis grading, DenseNet architectures effectively capture subtle radiographic features such as cartilage degeneration, osteophyte formation, joint space narrowing, and subchondral sclerosis, which are essential for distinguishing different KL grades. Furthermore, transfer learning with pretrained DenseNet models has been shown to improve convergence speed and enhance classification accuracy, particularly when training datasets are limited.

Building upon these developments, the present study proposes an integrated deep learning framework that combines the strengths of YOLOv5 for accurate knee joint localization and DenseNet201 for automated Kellgren–Lawrence severity grading. Unlike approaches that directly classify entire radiographs, the proposed two-stage pipeline focuses exclusively on the localized tibiofemoral joint region, enabling more discriminative feature extraction and improved classification performance. The framework is further implemented as an interactive Streamlit-based web application, providing an end-to-end clinical decision-support system capable of delivering real-time knee osteoarthritis detection and severity assessment. This integrated approach enhances diagnostic consistency,



reduces processing time, and demonstrates the practical potential of deep learning for supporting radiologists and orthopedic clinicians in routine clinical practice.

3. METHODOLOGY

The proposed methodology presents a fully automated two-stage deep learning framework for knee osteoarthritis (KOA) detection and severity grading from plain radiographic images. The framework combines object detection and image classification to improve diagnostic accuracy by first identifying the clinically relevant knee joint region and subsequently predicting the severity of osteoarthritis based on the Kellgren–Lawrence (KL) grading system. The overall workflow consists of data collection, image preprocessing, deep learning model development, model training, and deployment through an interactive web-based application.

3.1 Data Collection

The dataset used in this research comprises digital anteroposterior (AP) knee radiographs collected for osteoarthritis diagnosis. Each radiographic image is carefully annotated to support both stages of the proposed framework. For the localization stage, bounding box annotations are provided around the tibiofemoral joint region, enabling the object detection model

to accurately identify the region of interest. For the classification stage, each localized joint image is labeled according to the Kellgren–Lawrence grading system, where Grade 0 represents a healthy knee joint and Grade 4 indicates severe osteoarthritis characterized by extensive structural degeneration. Since a single radiograph may contain both the left and right knee joints, multiple regions of interest can be extracted from an individual image. These comprehensive annotations facilitate effective training of both the localization and classification models.

3.2 Image Preprocessing

Image preprocessing is performed to improve the quality and consistency of the input data before it is supplied to the deep learning models. Initially, the complete knee radiograph is processed by the YOLOv5 object detection model to identify the tibiofemoral joint region. The detected joint area is then cropped from the original image, thereby eliminating unnecessary background information and surrounding anatomical structures that do not contribute to osteoarthritis assessment. The cropped joint images are resized to 224×224 pixels to match the input dimensions required by the DenseNet201 classification network. Subsequently, pixel intensity values are normalized to the range of 0 to 1, ensuring numerical stability during model training and inference. This preprocessing strategy



standardizes the input images, reduces computational complexity, and enables the classification model to focus on clinically significant features such as joint space narrowing, osteophyte formation, and cartilage degeneration.

3.3 Proposed Deep Learning Architecture

The proposed framework adopts a cascaded deep learning architecture consisting of two sequential stages. The first stage employs the YOLOv5 object detection model to automatically localize the tibiofemoral joint region from full knee radiographs. YOLOv5 is a single-stage object detector capable of simultaneously performing object localization and classification, making it highly suitable for real-time medical image analysis. During inference, the model processes the original radiographic image and predicts the coordinates of the knee joint region together with the corresponding confidence score. Only detections with sufficient confidence are retained, and the localized joint regions are extracted for further analysis.

Following localization, the cropped knee joint images are resized and supplied to a DenseNet201 convolutional neural network for severity classification. DenseNet201 is a deep convolutional architecture consisting of 201

layers connected through dense feature propagation, where each layer receives information from all preceding layers. This dense connectivity promotes efficient feature reuse, improves gradient flow during training, reduces the vanishing gradient problem, and enables the network to learn complex image representations effectively. The network is initialized with ImageNet pretrained weights and fine-tuned for knee osteoarthritis classification by replacing the original classification layer with a new fully connected layer containing five output neurons corresponding to the five Kellgren–Lawrence grades. A Softmax activation function is applied to produce normalized probability scores for each grade, and the class with the highest probability is selected as the final prediction. The associated probability value represents the confidence of the model's prediction.

3.4 Model Training and Deployment

The localization and classification models are trained independently using their respective annotated datasets. The YOLOv5 detector is trained using radiographs annotated with bounding boxes surrounding the tibiofemoral joint region, allowing the model to accurately identify the region of interest within full knee X-ray images. The DenseNet201 classification model is trained using transfer learning, where pretrained ImageNet weights provide initial



feature representations that are subsequently fine-tuned using labeled knee joint images. Transfer learning significantly reduces training time while improving classification accuracy, particularly when medical imaging datasets are relatively limited in size.

After training, the DenseNet201 model is stored in Keras HDF5 format and directly loaded through TensorFlow during inference. Similarly, the trained YOLOv5 detector is loaded using the PyTorch framework through the Ultralytics implementation. To enhance computational efficiency, both models are cached during application startup, preventing repeated loading during user interactions.

The complete framework is implemented in Python and deployed as an interactive web application using the Streamlit framework. The deployed system enables users to upload knee radiographs in PNG or JPEG format, after which the YOLOv5 detector automatically identifies the tibiofemoral joint region. The detected region is cropped, resized, normalized, and forwarded to the DenseNet201 classifier for severity grading. The application then displays the predicted Kellgren–Lawrence grade together with the corresponding confidence score and the probability distribution across all five severity classes. This end-to-end framework provides an efficient and user-friendly clinical decision-

support tool capable of performing accurate knee osteoarthritis detection and severity grading in real time, thereby assisting radiologists and orthopedic clinicians in making faster, more consistent, and reliable diagnostic decisions.

4. DATASET

The dataset used in this study consists of digital anteroposterior (AP) knee radiographs collected for automated knee osteoarthritis (KOA) detection and severity grading. Each radiographic image contains one or two knee joints, providing sufficient anatomical information for identifying structural changes associated with osteoarthritis. The dataset is annotated separately for the localization and classification stages of the proposed deep learning framework.

For the object detection stage, each knee radiograph is manually annotated with bounding boxes enclosing the tibiofemoral joint region. These annotations enable the YOLOv5 model to accurately identify and localize the clinically relevant region of interest (ROI) while excluding unnecessary background structures. Accurate localization ensures that the subsequent classification model focuses exclusively on the anatomical region where osteoarthritic changes are most evident.



For the classification stage, the localized joint regions are assigned labels based on the Kellgren–Lawrence (KL) grading system, which is the standard radiographic classification method used for assessing knee osteoarthritis severity. The grading system categorizes the disease into five classes, ranging from Grade 0, representing a healthy knee without radiographic evidence of osteoarthritis, to Grade 4, indicating severe osteoarthritis with extensive structural damage. The KL grading criteria adopted in this study are summarized in Table 1.

Table-1: Kellgren–Lawrence Grading System for Knee Osteoarthritis

KL Grade	Radiographic Description
Grade 0	Normal knee joint with no radiographic evidence of osteoarthritis, normal joint space, and intact bone contours.
Grade 1	Doubtful joint space narrowing with possible osteophyte formation.
Grade 2	Definite osteophytes with possible narrowing of the joint space.
Grade 3	Multiple osteophytes, definite joint space narrowing, subchondral sclerosis, and possible bone deformity.

Grade 4	Severe osteoarthritis characterized by large osteophytes, marked joint space narrowing, severe subchondral sclerosis, and definite bone deformity.
----------------	--

Prior to model training, all images undergo a standardized preprocessing procedure to ensure consistency across the dataset. Initially, the YOLOv5 detector identifies the tibiofemoral joint region from the full radiographic image. The detected region is then cropped and resized to 224×224 pixels, matching the input dimensions required by the DenseNet201 classification network. Finally, pixel intensity values are normalized to the range $[0, 1]$, improving numerical stability and facilitating efficient model training. This preprocessing pipeline minimizes background interference, standardizes image dimensions, and enables the classification model to learn discriminative features associated with different stages of knee osteoarthritis more effectively.

5. RESULTS

The performance of the proposed two-stage deep learning framework was evaluated by assessing both the knee joint localization and osteoarthritis severity classification tasks. The evaluation demonstrates the effectiveness of integrating YOLOv5 for object detection with DenseNet201



for automated Kellgren–Lawrence (KL) grading. The complete system was further validated through an end-to-end web application to assess its practical applicability in clinical settings.

5.1 Evaluation Metrics

Different evaluation metrics were employed to assess the performance of the localization and classification stages. The performance of the YOLOv5 object detection model was evaluated using Precision, Recall, and Mean Average Precision (mAP@0.5), which are widely accepted metrics for measuring object detection accuracy. Precision evaluates the correctness of detected joint regions, while Recall measures the detector's ability to identify all relevant knee joints. Mean Average Precision at an Intersection over Union (IoU) threshold of 0.5 provides a comprehensive measure of localization performance.

The DenseNet201 classification model was evaluated using overall accuracy, precision, recall, and F1-score for each Kellgren–Lawrence grade. In addition, a confusion matrix was generated to analyze the distribution of correct and incorrect predictions across the five severity classes. These evaluation metrics provide a detailed understanding of the classifier's ability to distinguish between different stages of knee osteoarthritis.

5.2. Detection Performance

The YOLOv5 model demonstrated reliable performance in automatically localizing the tibiofemoral joint region from full knee radiographs. The detector accurately identified both unilateral and bilateral knee joints within a single image while effectively eliminating irrelevant background structures. Qualitative visualization of the detected bounding boxes confirmed that the model consistently localized the region of interest with high precision, producing standardized input for the classification stage. Accurate localization significantly reduced background interference and enabled the subsequent DenseNet201 model to concentrate on disease-related anatomical features, thereby improving overall diagnostic performance.

5.3. Classification Performance

The DenseNet201 model was evaluated on an independent test dataset containing all five Kellgren–Lawrence grades. The classifier demonstrated strong performance in distinguishing varying levels of osteoarthritis severity by effectively learning discriminative radiographic features such as joint space



narrowing, osteophyte formation, and subchondral sclerosis.

Analysis of the confusion matrix revealed that most classification errors occurred between adjacent KL grades, particularly between Grade 1 and Grade 2 as well as Grade 2 and Grade 3. These findings are consistent with the progressive nature of knee osteoarthritis, where the radiographic differences between neighboring severity levels are often subtle and difficult to distinguish even for experienced clinicians. Nevertheless, the proposed DenseNet201 classifier maintained reliable overall classification performance, indicating its capability to support consistent and objective osteoarthritis grading.

5.4. End-to-End System Performance

The complete two-stage framework was deployed as an interactive Streamlit web application to validate its performance in a real-world clinical environment. The deployed system successfully processed uploaded knee radiographs by automatically detecting the tibiofemoral joint region, extracting the localized image, and predicting the corresponding Kellgren–Lawrence grade in near real time.

The application also displayed the confidence score associated with the predicted class together

with a probability distribution across all five KL grades, enabling clinicians to interpret the model's predictions more effectively. The seamless integration of detection and classification demonstrates the feasibility of the proposed framework as a practical clinical decision-support system for automated knee osteoarthritis assessment.

6. DISCUSSION

The experimental results demonstrate that the proposed two-stage deep learning framework provides an effective solution for automated knee osteoarthritis detection and severity grading from plain radiographs. Unlike conventional single-stage approaches that directly classify entire radiographic images, the proposed framework first localizes the tibiofemoral joint region before performing disease classification. This sequential strategy allows the classifier to focus exclusively on clinically relevant anatomical structures while minimizing the influence of irrelevant background information.

The integration of YOLOv5 significantly improves the quality of the input images by accurately isolating the knee joint region, thereby reducing variations caused by differences in patient positioning, image acquisition, and field of view.



As a result, the DenseNet201 classifier receives standardized inputs that facilitate more consistent feature learning and improved classification performance.

DenseNet201 further enhances the effectiveness of the proposed framework through its dense connectivity mechanism, which promotes efficient feature reuse and improves gradient propagation during training. The architecture successfully captures both low-level image features, such as edges and textures associated with osteophyte formation, and high-level structural characteristics related to cartilage degeneration and joint space narrowing. These hierarchical representations enable the model to distinguish subtle radiographic differences across the five Kellgren–Lawrence grades.

An additional advantage of the proposed system is its deployment using native TensorFlow and PyTorch implementations without requiring intermediate model conversion. The DenseNet201 classifier is directly loaded from its trained Keras HDF5 model, ensuring consistency between the training and deployment environments while simplifying the implementation process. Furthermore, the integration of the complete framework into a Streamlit-based web application provides an intuitive interface that enables real-time osteoarthritis assessment, making the system

suitable for clinical decision support and educational applications.

Despite its promising performance, several limitations should be considered. The effectiveness of the proposed framework largely depends on the quality, diversity, and size of the training dataset. Limited representation of intermediate Kellgren–Lawrence grades may influence the classifier's ability to distinguish subtle disease progression. In addition, the current system relies exclusively on anteroposterior knee radiographs and does not incorporate complementary clinical information such as patient age, body mass index, medical history, symptoms, or laboratory findings, all of which may contribute to more comprehensive osteoarthritis assessment. Moreover, while the proposed framework demonstrates strong diagnostic capability, it is intended to function as a computer-aided decision-support tool rather than a replacement for expert radiological interpretation.

7. ACCURACY COMPARISON WITH EXISTING STUDIES

Automated knee osteoarthritis (KOA) diagnosis has attracted significant research interest with the advancement of deep learning techniques. Traditionally, the severity of knee osteoarthritis has been assessed manually by radiologists using



the Kellgren–Lawrence (KL) grading system, which remains the clinical standard for evaluating disease progression. Although this approach is widely accepted, it is time-consuming and often affected by inter-observer and intra-observer variability, leading to inconsistencies in diagnosis, particularly for early-stage osteoarthritis.

To overcome these limitations, several studies have employed deep learning models for automated KL grading using knee radiographs. Early research primarily utilized convolutional neural network architectures such as VGGNet, ResNet, and InceptionNet, which performed direct classification on entire knee radiographs. These approaches demonstrated improved diagnostic performance compared with conventional machine learning methods; however, their effectiveness was often limited by the presence of irrelevant background structures, variations in image acquisition, and differences in patient positioning. Consequently, the models occasionally struggled to learn discriminative features associated with subtle radiographic changes.

More recent studies have incorporated attention mechanisms, transfer learning, and ensemble learning strategies to enhance feature extraction and improve classification accuracy. Attention-based networks enable models to focus on

disease-relevant regions, while ensemble methods combine multiple classifiers to improve robustness and reduce prediction errors. Despite these advancements, accurately distinguishing between adjacent Kellgren–Lawrence grades remains a challenging task because the radiographic differences between consecutive disease stages are often minimal and gradual.

The proposed framework addresses these challenges by adopting a two-stage deep learning architecture that separates knee joint localization from severity classification. In the first stage, the YOLOv5 object detection model accurately identifies the tibiofemoral joint region, effectively removing unnecessary background information and standardizing the region of interest before classification. This localization process allows the classification model to focus exclusively on clinically relevant anatomical features, thereby improving feature quality and reducing variability caused by differences in radiographic acquisition.

In the second stage, a fine-tuned DenseNet201 model performs five-class Kellgren–Lawrence grading using the localized joint images. The dense connectivity of DenseNet201 promotes efficient feature reuse and facilitates the learning of both low-level texture patterns and high-level structural characteristics associated with osteoarthritis progression. As a result, the



proposed framework demonstrates improved capability in identifying subtle radiographic indicators, including early osteophyte formation, mild joint space narrowing, and subchondral sclerosis, which are essential for distinguishing neighboring KL grades.

The experimental findings indicate that the proposed approach provides reliable localization and competitive classification performance compared with existing deep learning methods. Consistent with previous studies, the majority of classification errors occur between Grade 1 and Grade 2 and between Grade 2 and Grade 3, reflecting the continuous and progressive nature of knee osteoarthritis rather than significant limitations of the proposed model. This observation suggests that differentiating adjacent disease stages remains one of the most challenging aspects of automated KL grading.

Overall, the integration of YOLOv5 for precise joint localization and DenseNet201 for robust severity classification offers a significant improvement over conventional single-stage classification frameworks. The proposed end-to-end pipeline not only enhances diagnostic consistency by reducing background interference but also provides a practical and computationally efficient solution for real-time computer-aided diagnosis. These characteristics make the framework a promising clinical decision-support

tool capable of assisting radiologists and orthopedic specialists in achieving faster, more objective, and reproducible knee osteoarthritis assessment.

8. CONCLUSION

This study presented a two-stage deep learning framework for the automated detection and severity grading of knee osteoarthritis (KOA) using plain radiographic images. The proposed approach combines the strengths of YOLOv5 for accurate localization of the tibiofemoral joint region and DenseNet201 for fine-grained classification of osteoarthritis severity based on the Kellgren–Lawrence (KL) grading system. By separating the localization and classification tasks, the framework effectively minimizes background interference and enables the classifier to focus on clinically relevant anatomical features, resulting in more reliable and consistent disease assessment.

The proposed system was successfully implemented as an interactive Streamlit-based web application, allowing users to upload knee radiographs and receive real-time visualization of detected joint regions along with automated KL grade predictions and confidence scores. This end-to-end implementation demonstrates the practical applicability of the framework as a computer-aided clinical decision-support tool



that can assist radiologists and orthopedic clinicians in improving diagnostic efficiency and reducing observer variability.

Experimental evaluation indicates that the proposed framework provides reliable knee joint localization and competitive classification performance across all five Kellgren–Lawrence grades. The results also demonstrate that combining object detection with deep feature learning enhances the overall effectiveness of automated knee osteoarthritis assessment compared with conventional whole-image classification approaches.

Although the proposed framework shows promising performance, future work can further improve its clinical applicability by incorporating larger and more diverse multi-center datasets, integrating additional patient information such as age, body mass index, and clinical symptoms, and exploring advanced deep learning architectures including transformer-based and multimodal learning models. These enhancements have the potential to improve the robustness, generalizability, and diagnostic accuracy of automated knee osteoarthritis assessment, ultimately supporting earlier diagnosis and more effective treatment planning in routine clinical practice.

REFERENCES

- [1] Bien, N., Rajpurkar, P., Ball, R. L., Irvin, J., Park, A., Jones, E., Bereket, M., Patel, B. N., Yeom, K. W., Shpanskaya, K., Halabi, S., Zucker, E., Fanton, G., Amanatullah, D. F., Beaulieu, C. F., Riley, G. M., Stewart, R. J., Blankenberg, F. G., Larson, D. B., & Lungren, M. P. (2018). Deep-learning-assisted diagnosis for knee magnetic resonance imaging: Development and retrospective validation of MRNet. *PLoS Medicine*, 15(11), e1002699. <https://doi.org/10.1371/journal.pmed.1002699>
- [2] Buslaev, A., Iglovikov, V. I., Khvedchenya, E., Parinov, A., Druzhinin, M., & Kalinin, A. A. (2020). Albumentations: Fast and flexible image augmentations. *Information*, 11(2), 125. <https://doi.org/10.3390/info11020125>
- [3] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778. <https://doi.org/10.1109/CVPR.2016.90>
- [4] Huang, G., Liu, Z., Van Der Maaten, L., & Weinberger, K. Q. (2017). Densely connected convolutional networks. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 4700–4708. <https://doi.org/10.1109/CVPR.2017.243>
- [5] Kellgren, J. H., & Lawrence, J. S. (1957). Radiological assessment of osteo-arthritis.



Annals of the Rheumatic Diseases, 16(4), 494–502.

<https://doi.org/10.1136/ard.16.4.494>

[6] Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A. A., Ciompi, F., Ghafoorian, M., Van Der Laak, J. A. W. M., Van Ginneken, B., & Sánchez, C. I. (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>

[7] Liu, X., Faes, L., Kale, A. U., Wagner, S. K., Fu, D. J., Bruynseels, A., Mahendiran, T., Moraes, G., Shamdas, M., Kern, C., Ledsam, J. R., Schmid, M. K., Balaskas, K., Topol, E. J., Bachmann, L. M., Keane, P. A., & Denniston, A. K. (2019). A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: A systematic review and meta-analysis. *The Lancet Digital Health*, 1(6), e271–e297. [https://doi.org/10.1016/S2589-7500\(19\)30123-2](https://doi.org/10.1016/S2589-7500(19)30123-2)

[8] Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You only look once: Unified, real-time object detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 779–788. <https://doi.org/10.1109/CVPR.2016.91>

[9] Tan, M., & Le, Q. V. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. *Proceedings of the 36th*

International Conference on Machine Learning (ICML), 6105–6114. Ultralytics. (2023). *YOLOv5 documentation*. <https://docs.ultralytics.com>

[10] Wang, C. J., Hamm, C. A., Savova, G. K., & Weinheimer-Haus, E. M. (2021). Deep learning for bone and joint diseases: A systematic review. *Journal of Clinical Medicine*, 10(23), 5673. <https://doi.org/10.3390/jcm10235673>

[11] Antony, J., McGuinness, K., O'Connor, N. E., & Moran, K. (2017). Quantifying radiographic knee osteoarthritis severity using deep convolutional neural networks. *Proceedings of the IEEE International Conference on Image Processing (ICIP)*, 1192–1196. <https://doi.org/10.1109/ICIP.2017.8296484>

[12] Norman, B., Pedoia, V., Noworolski, A., Link, T. M., & Majumdar, S. (2019). Applying densely connected convolutional neural networks for staging osteoarthritis severity from plain radiographs. *Journal of Digital Imaging*, 32(3), 471–477.

[13] Tiulpin, A., Klein, S., Bierma-Zeinstra, S. M. A., Thevenot, J., Rahtu, E., Van Meurs, J. B. J., Oei, E. H. G., & Saarakkala, S. (2018). Multimodal machine learning-based knee osteoarthritis grading: Data from the



Osteoarthritis Initiative. *Scientific Reports*, 8,
1727. [https://doi.org/10.1038/s41598-018-
20132-7](https://doi.org/10.1038/s41598-018-20132-7)