BMC Musculoskeletal Disorders

RESEARCH

Open Access

Automated detection and segmentation of chondromalacia patella on axial knee MRI using YOLOv11 and a custom CNN: a deep learning-based diagnostic model

Erdal Güngör[1*], Husam Vehbi[2], Mehmet Batu Ertan[3] and Ahmetcan Cansın[4]

Abstract

Objective To evaluate a deep learning pipeline using YOLOv11 for segmentation and a custom CNN for classification to automatically detect and assess chondromalacia patella on axial knee MRI, aligning with expert clinical evaluation.

Materials and methods A dataset of 650 axial knee MRIs was analyzed. YOLOv11 segmented the patellofemoral region, and a custom CNN classified chondromalacia. Performance was assessed using segmentation accuracy, classification accuracy, confidence scoring, and Grad-CAM-based visual explainability.

Results The CNN achieved a test accuracy of 82.30% on 113 images, with an AUC of 0.87, indicating promising but preliminary discriminative ability. Grad-CAM maps showed reasonable agreement with expert interpretation.

Conclusion The proposed YOLOv11-CNN pipeline demonstrated promising accuracy and may provide a potentially useful and interpretable solution for the detection and segmentation of chondromalacia patella on MRI, with the possibility of enhancing efficiency and consistency in orthopedic radiology workflows after further validation.

Keywords Chondromalacia patella, Knee MRI, Deep learning, YOLOv11, Convolutional neural network (CNN)

Introduction

Chondromalacia patella is a common and often underdiagnosed source of anterior knee pain, especially among young adults and physically active individuals. If not recognized early, it may progress to patellofemoral osteoarthritis, impacting quality of life and physical function [1]. Magnetic resonance imaging (MRI) is considered the gold standard for assessing articular cartilage and detecting chondral lesions [2]. However, manual interpretation of MRI images is often time-consuming and subject to significant interobserver variability, especially in axial plane evaluations [3].

Recent advances in artificial intelligence have introduced deep learning (DL) algorithms capable of automating image interpretation in clinical radiology [4–7]. These algorithms have shown promising results in classifying, segmenting, and detecting abnormalities in musculoskeletal MRI. Notably, Bien et al. [7] developed MRNet, a deep learning model for diagnosing knee pathologies using sagittal, coronal, and axial sequences. Xie et al. [8] used a deep CNN architecture for segmenting

*Correspondence:
Erdal Güngör
egungor_49@hotmail.com
[1]Department of Orthopaedic and Traumatology, İstanbul Medipol University Bahçelievler Hospital, İstanbul, Turkey
[2]Department of Radiology, Çam and Sakura City Hospital, İstanbul, Turkey
[3]Department of Orthopaedic and Traumatology, Medicana International Ankara Hospital, Ankara, Turkey
[4]İstanbul Medipol University International School of Medicine, İstanbul, Turkey

© The Author(s) 2025. **Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by-nc-nd/4.0/.

patellofemoral structures with high agreement compared to expert annotation.

The YOLO (You Only Look Once) family of models, including YOLOv11, has gained attention for real-time object detection due to its balance of speed and accuracy. YOLOv11 introduces an anchor-free segmentation head, streamlined convolutional backbone, and instance segmentation support, making it suitable for both classification and segmentation tasks in medical imaging [8].

This study aims to evaluate a pipeline using YOLOv11 for segmentation and a custom CNN for classification to detect and assess chondromalacia patella on axial knee MRI. In addition to assessing its diagnostic accuracy, we explore model interpretability using Grad-CAM visualization and compare the results with expert clinical evaluation.

Methods

This study was approved by the Istanbul Medipol University Ethics Committee (Approval No: E-10840098-202.3.3.3.02–3813, Date: 20/06/2025). Written informed consent was obtained from all individual participants included in the study. For retrospective analyses using anonymized patient data, the requirement for informed consent was waived by the Ethics Committee in accordance with national regulations. All procedures were performed in accordance with the Declaration of Helsinki.

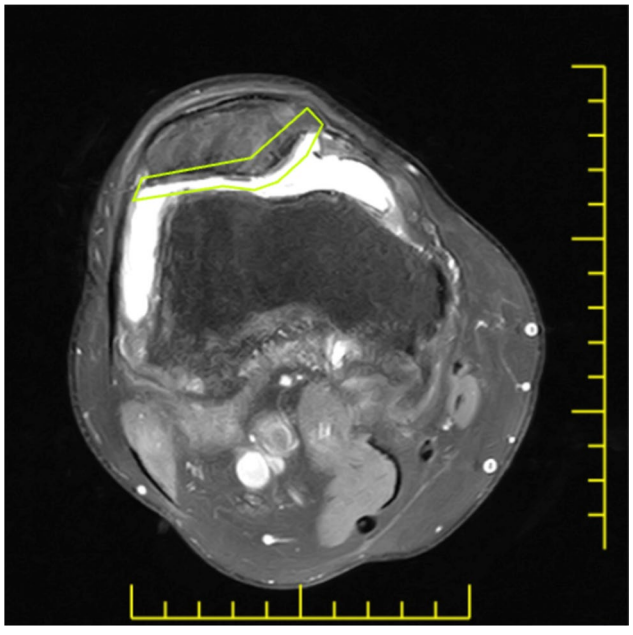


Fig. 1 Example of manual annotation for chondromalacia patella using Roboflow. The axial T2-weighted fat-saturated knee MRI image shows the region of interest along the patellar cartilage surface, delineated by a yellow line indicating cartilage irregularity

Study design and dataset

This retrospective study included 650 anonymized axial knee MRI scans from patients aged between 18 and 65 years. The images were reviewed and labeled by one board-certified radiologist and two orthopedic surgeons with over 10 years of clinical experience. Each case was labeled as either having chondromalacia patella or not, based on cartilage surface irregularity, thinning, and signal intensity alterations. For classification, the dataset included 923 train and 113 test images. For segmentation, 77 validation instances were used. Below is an example for annotation using Roboflow Fig 1.

For segmentation tasks, a total of 200 axial MRI slices were manually annotated with pixel-level labels. Of these, 123 images were used for training and 77 for validation and evaluation of the YOLOv11 segmentation model. The classification dataset was larger, as classification did not require pixel-level annotation. Some of the annotated images overlapped with the classification dataset, but no image was used simultaneously for segmentation validation and classification testing, ensuring that the two evaluations remained independent.

Input type (2D vs. 3D)

The model was trained exclusively on 2D axial slices rather than full 3D volumes. No interpolation across slices or volumetric normalization was applied.

Slice selection policy

Only axial slices that fully included the patellofemoral joint were selected for analysis. Slices not covering this region were excluded.

Image preprocessing

All images were resampled to an in-plane resolution of 0.5×0.5 mm, cropped to the knee joint region, and zero-padded when necessary to maintain input size. Intensity values were normalized to the range [0–1] before training. MRI images were independently reviewed by two orthopedic surgeons and one musculoskeletal radiologist. Inter-rater agreement was substantial (Cohen's $\kappa = 0.65$). In cases of disagreement, the final decision was made by consensus among the three evaluators.

Model architecture

The pipeline used YOLOv11 for segmenting the patellofemoral region, employing a CSPDarknet backbone and an anchor-free segmentation head. Segmented regions were resized to 224×224 pixels and input into a custom CNN for classification. YOLOv11 processes images at 640×640 pixels, performing localization and boundary delineation in a single forward pass. Both detection and segmentation were applied in the pipeline: detection was used to localize the patellofemoral region,

while segmentation provided more precise boundary delineation to remove background noise and anatomical structures not of interest. The CNN subsequently received cropped original MRIs derived from bounding boxes (resized to 224 × 224 pixels) rather than masked-out MRIs. This ensured that the classifier focused on the relevant anatomical region while retaining sufficient contextual information. A schematic representation of the pipeline is provided in Fig. 2. The custom CNN consisted of four convolutional layers with ReLU activation, each followed by max-pooling, a flatten layer, two fully connected layers with dropout regularization, and a final sigmoid output layer for binary classification.

Preprocessing and augmentation

MRI images were normalized to reduce inter-scan variability. Data augmentation techniques including horizontal flipping, brightness contrast adjustment, rotation, and zoom were applied to improve generalization and prevent overfitting. All images were resized to 640 × 640 for YOLOv11 and 224 × 224 for the CNN.

The dataset consisted of 923 training, 77 validation, and 113 test images. Data augmentation (horizontal flipping, brightness/contrast adjustment, rotation, and zoom) was applied only to the training set to improve model generalization. Class balance was maintained across all subsets, with approximately 50% chondromalacia-positive and 50% normal cases.

Detailed dataset distribution is provided in Supplementary Table 1.

Training and evaluation

The YOLOv11 model was trained on 80% of the labeled dataset, with 10% for validation and 10% for testing. The CNN was trained with a batch size of 16, a learning rate of 0.001, and the Adam optimizer over 15 epochs. Performance metrics included segmentation accuracy (IoU), classification accuracy, precision, recall, and F1-score. Validation and test performances were emphasized to evaluate the robustness and generalizability of the model. Grad-CAM visualizations were generated for interpretability. To avoid data leakage, all splits (train/validation/test) were performed at the patient level, ensuring that MRI scans from the same patient were included in only one subset. This guarantees fully disjoint partitions across training, validation, and testing stages. In addition, weight decay tuning was explored during preliminary experiments, but it did not yield further improvements. Cross-validation was considered; however, due to the computational burden and dataset size, we opted for a patient-level train/validation/test split to avoid data leakage (Fig. 3).

Explainability using Grad-CAM

Gradient-weighted Class Activation Mapping (Grad-CAM) was used to visualize where the CNN focused its attention during classification. These heatmaps were overlaid on the original MR images and compared to clinical labels for anatomical relevance (Fig 4).

Model output evaluation and figures

The outputs were evaluated by orthopedic surgeons and a musculoskeletal radiologist. Grad-CAM heatmaps confirmed model's focus on relevant regions (Fig 5).

Analysis

All model outputs were reviewed independently by one radiology specialist and two orthopedic surgeons to verify the presence, localization, and segmentation accuracy of chondromalacia lesions.

Statistical analysis

Model outputs were compared to expert annotations. Segmentation accuracy was quantified using IoU, with YOLOv11 achieving a Box mAP50 of 0.978 and Mask mAP50 of 0.978. Classification performance was assessed using confusion matrix metrics, with the CNN achieving a test accuracy of 82.30%.

Results

The YOLOv11 model analyzed 77 validation instances, achieving a Box mAP50 of 0.978, Mask mAP50 of 0.978, and mAP50-95 of 0.616 and 0.731 respectively. The CNN analyzed 113 test images, detecting chondromalacia with a test accuracy of 82.30%, train accuracy of 94.58%, and train loss of 0.1390. Grad-CAM maps aligned with expert understanding. Beyond qualitative agreement with expert assessment, Grad-CAM analysis revealed that the CNN primarily focused on the patellar articular cartilage, the subchondral bone interface, and the trochlear groove as the most influential anatomical landmarks. In chondromalacia-positive cases, activation maps concentrated on regions showing cartilage thinning, surface irregularity, and signal intensity changes, supporting the model's discriminative capacity (Fig. 3).

Additionally, the performance of the classification model was further evaluated using a receiver operating characteristic (ROC) curve. The model achieved an area under the curve (AUC) of 0.87, indicating strong



Fig. 2 Pipeline of the proposed YOLOv11–CNN framework

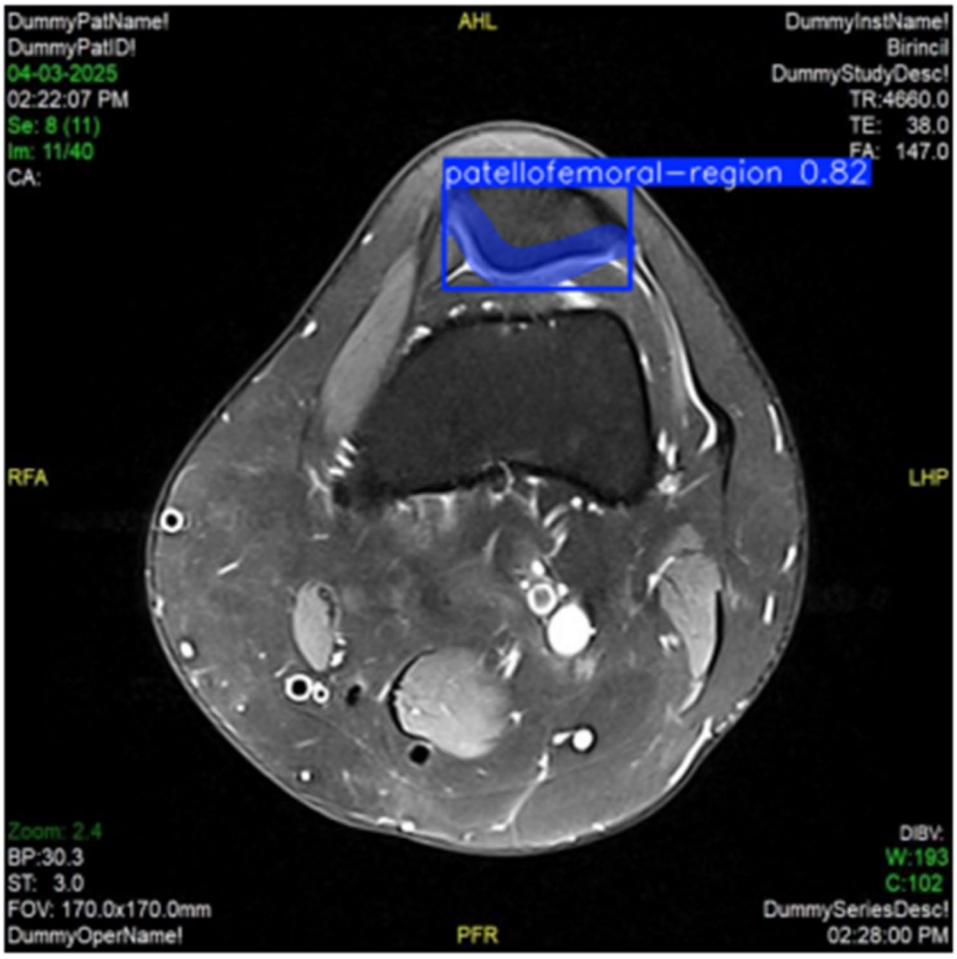


Fig. 3 Segmentation result of an axial knee MRI using YOLOv11, showing the patellofemoral region with a confidence score of 0.82, labeled using Roboflow

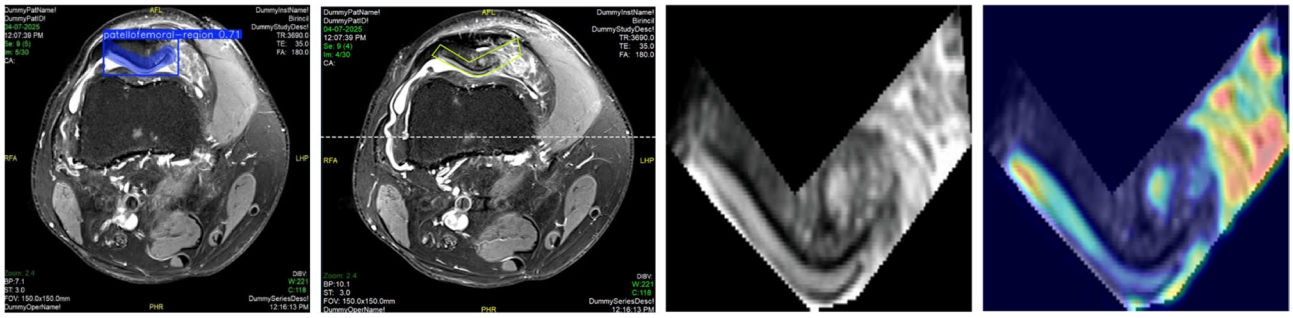


Fig. 4 Visualization of segmentation and classification outputs. **a** YOLOv11 detection with blue bounding box. **b** Segmentation boundaries overlaid on raw MRI (yellow). **c** Cropped cartilage region after YOLOv11 segmentation. **d** Grad-CAM heatmap overlay highlighting CNN focus regions

discriminatory power in distinguishing between chondromalacia patella and normal cartilage cases (Fig. 6).

A detailed comparison of performance metrics for both models, including accuracy, AUC, precision, recall, and F1-score, is summarized in Table 1.

The CNN was trained for 15 epochs based on validation loss monitoring, as extending training beyond this point led to overfitting without improvement in validation accuracy. Early stopping with a patience of three epochs was applied, and dropout layers were included in the fully connected stages to reduce overfitting and enhance generalization. In addition, data augmentation was employed to mitigate the relatively small dataset size Table 2.

In addition to IoU and mAP, the Dice Similarity Coefficient (DSC) was calculated, with YOLOv11 achieving a mean DSC of 0.89 on the validation set. The Dice

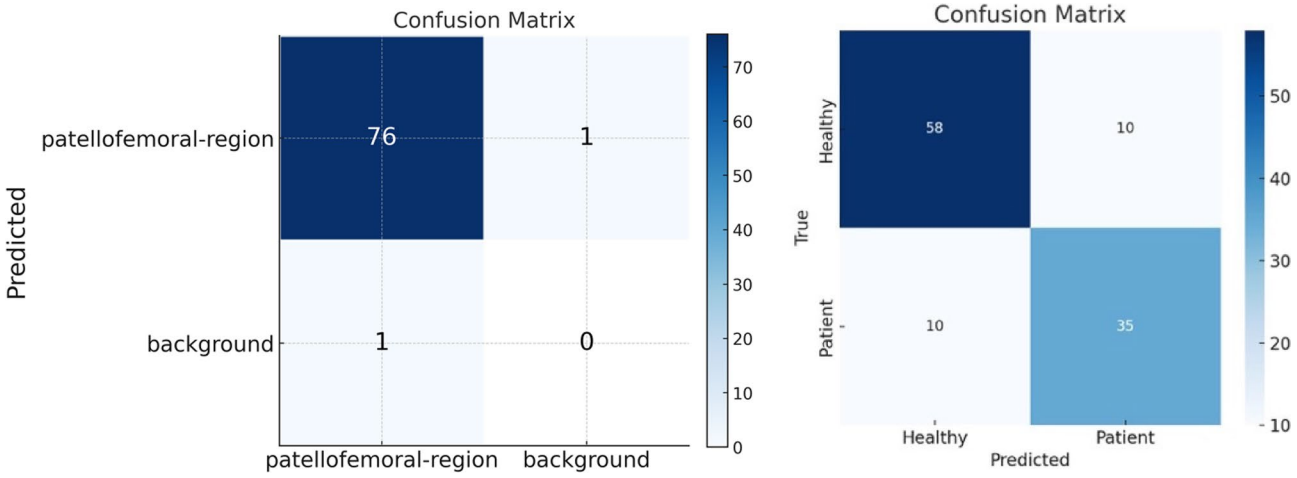


Fig. 5 Confusion matrices from the validation set. **a** Segmentation results for patellofemoral region vs. background. **b** Classification results for Healthy vs. Patient groups

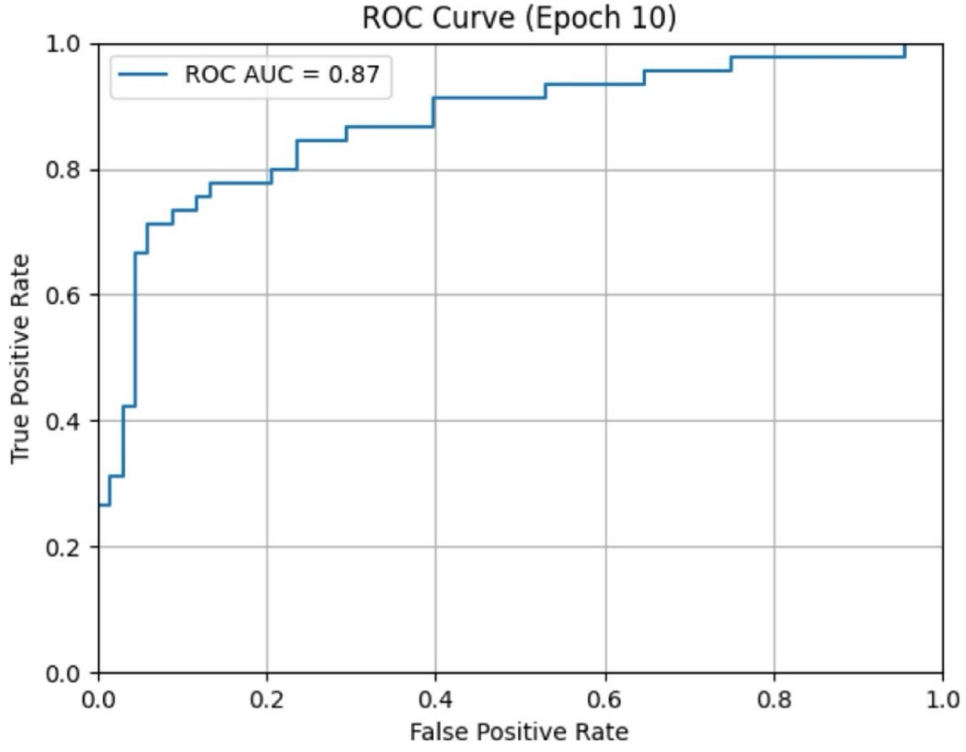


Fig. 6 Receiver Operating Characteristic (ROC) curve of the custom CNN model on the test set after 10 epochs, showing an AUC of 0.87, indicating strong discriminative ability between chondromalacia and normalcartilage

Table 1 Performance metrics of the custom CNN and YOLOv11 models for chondromalacia patella detection. Metrics include accuracy, AUC, precision, recall, and F1-score for the CNN, and mAP values for YOLOv11

Model	Accuracy (%)	AUC	Precision	Recall	F1-score
Custom CNN	82.30	0.87	0.83	0.82	0.825

Similarity Coefficient (DSC) was reported to allow easier comparison with existing medical imaging literature.

Table 2 The model achieved high segmentation accuracy across multiple metrics. In addition to Intersection over Union (IoU) and mean Average Precision (mAP), the Dice

Metric	Value
IoU	0.90
mAP50	0.978
mAP50-95	0.616
DSC	**0.89**

These segmentation performance metrics are summarized in Table 2.

In addition to qualitative agreement with expert assessment, Grad-CAM analysis demonstrated that the CNN model primarily focused on the patellar articular cartilage, the subchondral bone interface, and the trochlear groove when differentiating between normal and chondromalacia cases. In positive cases, activation maps concentrated on areas of cartilage thinning, surface irregularity, and signal intensity changes, supporting the model's discriminative capacity.

Discussion

Several other studies have contributed to the growing body of literature on deep learning in musculoskeletal and cartilage imaging. Panfilov et al. [9] demonstrated the utility of fully automated subregional cartilage assessment, aligning with our use of model-driven precision. Tajbakhsh et al. [10] highlighted how robust models can still emerge from imperfect datasets, supporting the effectiveness of our YOLOv11 model even in real-world imaging variability. Lundervold and Lundervold [11] and Erickson et al. [12] both provided comprehensive overviews of deep learning applications in MRI and medical imaging workflows, forming the theoretical and clinical basis for models like ours.

This study demonstrates the feasibility and effectiveness of a YOLOv11-based deep learning model in detecting and segmenting chondromalacia patella using axial knee MRI images. The model showed a segmentation accuracy of 0.98 and a classification accuracy of 0.82, indicating reasonably good performance in localizing and identifying cartilage degeneration. Our model utilized a combined detection and segmentation architecture, allowing for simultaneous localization and delineation of cartilage lesions in a single forward pass. Compared to traditional multi-step pipelines, this approach improves efficiency and reduces potential sources of error. Previous studies have reported substantial inter-rater variability among radiologists in the evaluation of chondromalacia on MRI, with kappa values ranging between 0.4 and 0.7 [13, 14]. Thus, automated models achieving accuracies above 80% may offer comparable or complementary diagnostic support.

Our findings resonate with Bien et al. [7], who introduced MRNet for knee pathology classification. While MRNet focused solely on classification, our YOLOv11 model also performs pixel-level segmentation, making it more comprehensive for assessing cartilage lesions. Xie et al. [8] also utilized segmentation in patellofemoral MRI and reported strong alignment with expert assessments, which our results similarly demonstrate. Furthermore, Panfilov et al. [9] applied deep learning for subregional cartilage analysis and highlighted the importance of spatial precision—something our model achieves through its anchor-free segmentation mechanism.

In terms of interpretability, Selvaraju et al. [6] introduced Grad-CAM to highlight model attention areas, which we used to ensure our network focused on clinically relevant regions. This visual validation may improve transparency in the decision-making process and could help support trust among medical professionals. Holzinger et al. [15] emphasized the need for explainability in AI, especially in high-stakes clinical environments.

YOLOv11 showed comparable performance to other CNN-based models, with the advantage of faster inference speed. Ravi et al. [4] and Esteva et al. [5] demonstrated early success in medical classification tasks using CNNs; however, these models lacked segmentation capabilities. YOLOv11 combines real-time detection with high-resolution segmentation, making it more versatile for radiological applications. Compared to Mask R-CNN [16], YOLOv11 offers improved speed and comparable accuracy, with simpler architecture. TransUNet [17], which integrates transformers into segmentation pipelines, shows excellent accuracy, but requires significantly higher computational resources, making YOLOv11 more deployable in clinical environments.

Similarly, nnU-Net [18] is widely recognized as a strong benchmark in biomedical image segmentation tasks. In this study, our primary aim was to develop a unified pipeline that combined detection, segmentation, and classification in a single framework. For this reason, we employed YOLOv11, which offers real-time detection and segmentation capabilities with relatively lower computational demands. We acknowledge that nnU-Net, particularly its 3D implementation, remains a widely used baseline for MRI segmentation tasks. Our choice of YOLOv11 should therefore be understood not as a claim of superiority over nnU-Net, but rather as a pragmatic decision aligned with the study design and feasibility. Our selection was primarily motivated by YOLOv11's faster inference speed and lower computational demands, which make it practical for clinical workflows with limited resources. A direct comparison with nnU-Net is planned in future work.

In contrast to traditional radiographic analysis methods, Antony et al. [19] developed a deep learning system to assess osteoarthritis from plain X-rays. Although their approach was based on 2D radiographs, it aligns with our findings in highlighting the value of AI in musculoskeletal diagnostics. Shen et al. [20] further emphasize the broad applicability of deep learning in medical imaging, supporting our adoption of a CNN-based architecture for MRI analysis.

Xie et al. [8] presented a holistically nested network for patellofemoral segmentation, achieving robust results across MRI sequences. Our model similarly targets

patellofemoral cartilage but offers a unified detection-segmentation framework that simplifies deployment. Additionally, Park et al. [21] implemented a deep learning system for automatic patellar height measurement. While their focus differed, their work reinforces the practicality and accuracy of AI models in analyzing patellofemoral structures.

Despite these strengths, our study has several limitations. First, it was conducted retrospectively in a single center, which may limit the generalizability of the findings. Differences in MRI acquisition parameters, including scanner type, field strength, and protocol variations, could also influence model performance and reduce external validity. In addition, the ground truth labels were derived from expert clinical interpretation rather than histological confirmation. While expert consensus increases diagnostic reliability, it does not fully replace histological validation as the gold standard.

Although we implemented measures such as data augmentation, dropout, and early stopping, a gap between training and test performance suggests a potential risk of overfitting. Furthermore, Grad-CAM visualizations in this study were assessed only qualitatively, since pixel-level lesion annotations were not available; quantitative validation of explainability methods will be necessary in future research.

External validation with larger and more diverse multi-center datasets is essential to confirm the robustness and applicability of the model. Future work should also investigate performance across different MRI sequences and scanner types, incorporate quantitative lesion annotations, and compare deep learning models against histological gold standards. Moreover, integrating advanced approaches such as self-supervised learning or transformer–CNN hybrid architectures may further improve adaptability and clinical utility.

Conclusion

Chondromalacia patella is a frequently overlooked cause of anterior knee pain, and its diagnosis remains challenging in routine clinical practice. In this study, we demonstrated the utility of a YOLOv11-based segmentation and custom CNN classification pipeline that integrates lesion detection, segmentation, and explainability tools such as Grad-CAM to assist in the evaluation of axial knee MRIs. The model provided reproducible assessments that were generally aligned with expert interpretation, although further validation with larger, multi-center datasets will be required before considering clinical application. While these findings are encouraging, additional external validation is necessary to confirm robustness and generalizability.

Abbreviations

MRI	Magnetic Resonance Imaging
DL	Deep Learning
CNN	Convolutional Neural Network
YOLO	You Only Look Once
Grad-CAM	Gradient-weighted Class Activation Mapping
AUC	Area Under the Curve
IoU	Intersection over Union
ROC	Receiver Operating Characteristic
mAP	Mean Average Precision
CSP	Cross Stage Partial
T2	T2-weighted Imaging
AI	Artificial Intelligence
nnU-Net	No-new-Net
TransUNet	Transformer-based U-Net

Supplementary Information

The online version contains supplementary material available at https://doi.org/10.1186/s12891-025-09275-7.

Supplementary Material 1.

Authors' contributions

Erdal Güngör: Conceptualization, Software, Supervision, Writing – review & editingHusam Vehbi: Data curation, Resources, Formal analysis, VisualizationMehmet Batu Ertan: Investigation, Resources, Writing – original draftAhmetcan Cansın: Software, Validation, Project administration, All authors read and approved the final manuscript.

Funding

The authors did not receive support from any organization for the submitted work.

Data availability

The datasets used and/or analysed during the current study available from the corresponding author on reasonable request.

Declarations

Consent for publication

Not applicable.

Competing interests

The authors declare no competing interests.

Received: 9 July 2025 / Accepted: 8 October 2025
Published online: 13 November 2025

References

1. Pihlajamäki HK, Kuikka PI, Leppänen VV, Jaatinen P, Mäkelä A, Visuri T, et al. Reliability of clinical findings and magnetic resonance imaging for the diagnosis of patellofemoral pain syndrome. Clin J Sport Med. 2010;20(6):456–61.
2. Stefanik JJ, Roemer FW, Zumwalt AC, Zhu Y, Jarraya M, Lynch JA, et al. Association between patella Alta and structural features of patellofemoral joint osteoarthritis on MRI: the multicenter osteoarthritis study. J Orthop Res. 2018;36(5):1435–42.
3. Litjens G, Kooi T, Bejnordi BE, Setio AAA, Ciompi F, Ghafoorian M, et al. A survey on deep learning in medical image analysis. Med Image Anal. 2017;42:60–88.
4. Ravi D, Wong C, Deligianni F, Berthelot M, Andreu-Perez J, Lo B, et al. Deep learning for health informatics. IEEE J Biomed Health Inf. 2017;21(1):4–17.
5. Esteva A, Kuprel B, Novoa RA, Ko J, Swetter SM, Blau HM, et al. Dermatologist-level classification of skin cancer with deep neural networks. Nature. 2017;542(7639):115–118.
6. Selvaraju RR, Cogswell M, Das A, Vedantam R, Parikh D, Batra D, et al. Grad-CAM: Visual explanations from deep networks via gradient-based localization. Proc IEEE Int Conf Comput Vis. 2017:618–626.

7. Bien N, Rajpurkar P, Ball RL, Irvin J, Park A, Jones E, et al. Deep-learning-assisted diagnosis for knee magnetic resonance imaging: development and retrospective validation of MRNet. PLoS Med. 2018;15(11):e1002699.
8. Xie Y, Zhang J, Xia Y, Shen C, Li W, Xie X, et al. Fully automated patellofemoral MRI segmentation using holistically nested networks. Magn Reson Med. 2020;83(2):731–42.
9. Panfilov E, Tiulpin A, Nieminen MT, Saarakkala S, Casula V. Deep learning-based segmentation of knee MRI for fully automatic subregional morphological assessment of cartilage tissues: data from the osteoarthritis initiative. J Orthop Res. 2022;40(5):1113–24.
10. Tajbakhsh N, Jeyaseelan L, Li Q, Chiang JN, Wu Z, Ding X. Embracing imperfect datasets: a review of deep learning solutions for medical image segmentation. Med Image Anal. 2020;63:101693.
11. Lundervold AS, Lundervold A. An overview of deep learning in medical imaging focusing on MRI. Z Med Phys. 2019;29(2):102–27.
12. Erickson BJ, Korfiatis P, Akkus Z, Kline TL. Machine learning for medical imaging. Radiographics. 2017;37(2):505–15.
13. Kornaat PR, Ceulemans RY, Kroon HM, Riyazi N, Kloppenburg M, Carter WO, et al. MRI assessment of knee osteoarthritis: knee osteoarthritis scoring system (KOSS)—inter-observer and intra-observer reproducibility of a compartment-based scoring system. Skeletal Radiol. 2005;34(2):95–102.
14. Jungmann PM, Welsch GH, Brittberg M, Trattnig S, Braun S, Imhoff AB, et al. Magnetic resonance imaging score and classification system (AMADEUS) for assessment of preoperative cartilage defect severity. Cartilage. 2017;8(3):272–82.
15. Holzinger A, Langs G, Denk H, Zatloukal K, Müller H. Causability and explainability of artificial intelligence in medicine. WIREs Data Min Knowl Discov. 2019;9(4):e1312.
16. He H, Xu H, Zhang Y, Gao K, Li H, Ma L, et al. Mask R-CNN based automated identification and extraction of oil well sites. Int J Appl Earth Obs Geoinf. 2022;112:102875.
17. Chen J, Lu Y, Yu Q, Luo X, Adeli E, Wang Y, et al. TransUNet: Transformers make strong encoders for medical image segmentation. ArXiv Preprint. 2021;arXiv:210204306.
18. Isensee F, Jaeger PF, Kohl SAA, Petersen J, Maier-Hein KH. Nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. Nat Methods. 2021;18(2):203–11.
19. Antony J, McGuinness K, O'Connor NE, Moran K, O'Neill C, Meehan L, et al. Automatic knee osteoarthritis diagnosis from plain radiographs using deep learning. Sci Rep. 2017;7(1):4662.
20. Shen D, Wu G, Suk HI. Deep learning in medical image analysis. Annu Rev Biomed Eng. 2017;19:221–48.
21. Park HJ, Lee SM, Kang BJ, Choi SH, Kim MJ, Yoo HJ, et al. Deep learning-based automatic measurement system for patellar height: A multicenter retrospective study. J Orthop Surg Res. 2024;19(1):45.

Publisher's note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.