

Feature Learning to Automatically Assess Radiographic Knee Osteoarthritis Severity

Joseph Antony, Kevin McGuinness, Kieran Moran, Noel E O' Connor

Insight Centre for Data Analytics, Dublin City University, Ireland
`joseph.antony@insight-centre.org`

Abstract. Feature learning refers to techniques that learn to transform raw data input into an effective representation for further higher-level processing in many computer vision tasks. This chapter presents the investigations and the results of feature learning using convolutional neural networks to automatically assess knee osteoarthritis (OA) severity and the associated clinical and diagnostic features of knee OA from radiographs (X-ray images). Also, this chapter demonstrates that feature learning in a supervised manner is more effective than using conventional handcrafted features for automatic detection of knee joints and fine-grained knee OA image classification. In the general machine learning approach to automatically assess knee OA severity, the first step is to localize the region of interest that is to detect and extract the knee joint regions from the radiographs, and the next step is to classify the localized knee joints based on a radiographic classification scheme such as Kellgren and Lawrence grades. First, the existing approaches for detecting (or localizing) the knee joint regions based on handcrafted features are reviewed and outlined in this chapter. Next, three new approaches are introduced: 1) to automatically detect the knee joint region using a fully convolutional network, 2) to automatically assess the radiographic knee OA using CNNs trained from scratch for classification and regression of knee joint images to predict KL grades in ordinal and continuous scales, and 3) to quantify the knee OA severity optimizing a weighted ratio of two loss functions: categorical cross entropy and mean-squared error using multi-objective convolutional learning. The results from these methods show progressive improvement in the overall quantification of the knee OA severity. Two public datasets: the OAI and the MOST are used to evaluate the approaches with promising results that outperform existing approaches. In summary, this work primarily contributes to the field of automated methods for localization (automatic detection) and quantification (image classification) of radiographic knee OA.

Keywords: Feature learning, Handcrafted features, Convolutional neural networks, Kellgren and Lawrence grades, Automatic detection, Classification, Regression, Multi-objective convolutional learning

Table of Contents

Abstract	1
1 Introduction	3
1.1 Knee Osteoarthritis	5
1.1.1 Diagnostic Features of Knee OA	5
1.1.2 Radiographic Classification of Knee OA	6
1.2 Contributions	7
2 Related Work and Background	7
2.1 Detecting Knee Joints in Radiographs	8
2.1.1 Manual Methods	8
2.1.2 Semi-automatic Methods	8
2.1.3 Automatic Methods	9
2.2 Assessing Radiographic Knee OA Severity	10
2.2.1 Quantitative Analysis	11
2.2.2 Automatic Classification	11
2.3 Discussion	12
3 Public Knee OA Datasets	13
3.1 OAI Dataset	13
3.2 MOST Dataset	14
4 Automatic Detection of Knee Joints	14
4.1 Baseline Methods	14
4.1.1 Template Matching	15
4.1.2 SVM-Based Detection	18
4.2 Fully Convolutional Network Based Detection	20
4.2.1 Localisation with Reference to Knee Joint Centre	21
4.2.2 Localising the Region of Interest	32
4.3 Summary and Discussion	36
5 Automatic Assessment of Knee OA Severity	37
5.1 Baseline for Classifying Knee OA Radiographs	38
5.1.1 WNDCHRM Classification	38
5.1.2 Classification using Hand-crafted Features	40
5.2 Automatic Quantification using Convolutional Neural Networks	42
5.2.1 Off-the-shelf CNNs	42
5.2.2 Classification using CNN Feature Extraction	43
5.2.3 Transfer Learning	45
5.2.4 Training CNNs from Scratch	49
5.2.5 Multi-objective Convolutional Learning	62
5.2.6 Ordinal Regression	71
5.3 An Automatic Knee OA Diagnostic System	75
5.4 Summary and Discussion	77
6 Conclusion	78

1 Introduction

Traditionally, many handcrafted features have been successfully used in computer vision tasks, and they often simplify machine learning tasks. Nevertheless, they have a few limitations. These features are often low-level as prior knowledge is hand-encoded and features in one domain do not always generalize to other domains. In recent years, learning feature representations in a supervised manner also known as supervised feature learning is preferred over handcrafted features as they have outperformed the state-of-the-art in many computer vision tasks and have been highly successful. This chapter focuses on feature learning to automatically assess radiographic knee OA severity using convolutional neural networks (CNNs).

Clinically to assess knee OA severity, highly experienced clinicians or radiologists assess the knee joints in X-ray images [1, 2] and assign an ordinal grade based on a radiographic grading scheme. The most commonly used gradings, like the Kellgren and Lawrence (KL) grading scheme and Ahlback system, use distinctive grades (0 to 4). However, clinical features of knee OA are continuous in nature, and attributing distinctive grades is the subjective opinion of the graders. There are also uncertainties and variations in the subjective gradings. There is a need for automated methods to overcome the limitations arising from this subjectivity, and to improve the reliability in the measurements and classifications [2].

The automatic assessment of knee OA severity has been previously approached in the literature as an image classification problem [1, 3, 4], with the KL grading scale as the ground truth. WNDCHARM¹, a multi-purpose biomedical image classifier was used to classify knee OA images [4, 5]. High binary classification accuracies (80% to 91%) have been reported using the WNDCHARM classifier for classifying the extreme stages: grade 0 (normal) vs grade 4 (severe), grade 0 vs grade 3 (moderate). However, the classification accuracies of the images belonging to successive grades are low (55% to 65%) and the multi-class classification accuracy is low (35%). The overall classification accuracies of knee OA needs improvement for real-world computer aided diagnosis [1, 3, 6].

Radiographic features detected and learned through a computer-aided analysis can be useful to quantify knee OA severity and to predict the future development of knee OA [3]. Instead of manually designing features, the author proposes that learning feature representations using deep learning architectures can be a more effective approach for the classification of knee OA images. Traditionally, hand-crafted features based on pixel statistics, object and edge statistics, texture, histograms, and transforms, are typically used for multi-purpose medical image classification [4, 5, 7]. However, these features are not efficient for fine-grained classification such as classifying successive grades of knee OA images. Manually designed or hand-engineered features often

¹ Weighted Neighbor Distance using Compound Hierarchy of Algorithms Representing Morphology

simplify machine learning tasks. Nevertheless, they have a few disadvantages. The process of engineering features requires domain related expert knowledge and is often very time consuming [8]. These features are often low-level as prior knowledge is hand-encoded and features in one domain do not always generalize to other domains [9]. The next logical step is to automatically learn effective features for the desired task.

Over the last decade, learning feature representations or feature learning has been preferred to hand-crafted features in many computer vision tasks, particularly for fine-grained classification, because rich appearance and shape features are essential for describing subtle differences between categories [10]. Feature learning refers to techniques that learn to transform raw data input or pixels of an image to an effective representation for further higher-level processing such as object detection, automatic detection, segmentation, and classification. Feature learning approaches provide a natural way to capture cues by using a large number of code words (sparse coding) or neurons (deep networks), while traditional computer vision features, designed for basic-level category recognition, may eliminate many useful cues during feature extraction [10]. Deep learning architectures are multi-layered and they are used to learn feature representations in the hidden layer(s). These representations are subsequently used for classification or regression at the output layer. Feature learning is an integral part of deep learning [8, 11].

Even though many deep learning architectures have been proposed and have existed for decades, in recent times CNNs have become highly successful in the field of computer vision [12, 13]. AlexNet[14] won the ILSVRC² in 2012 by a large margin. CNNs have since then become more popular, widely-used and highly-successful in computer vision tasks such as object detection, image recognition, automatic detection and segmentation, content based image retrieval, and video classification [12]. Apart from computer vision tasks, CNNs are finding applications in natural language processing, hyper-spectral image processing, and medical image analysis [12, 15].

CNNs have also recently become successful in many medical applications such as knee cartilage segmentation [16] and brain tumour segmentation [17] in MRI scans, multi-modality iso-intense infant brain image segmentation [18], pancreas segmentation in CT images [19], and neuronal membrane segmentation in electron microscopy images [20]. Inspired by these success stories, the author proposes CNNs for classification of knee OA images and to improve the quantification of knee OA severity and knee OA diagnostic features. The author believes that this can lead to building a real-world knee OA diagnostic system that outperforms the existing approaches.

The remainder of this chapter is organized as follows. Section 1.1 introduces knee osteoarthritis (OA), the diagnostic features and the clinical evaluation of knee OA. Section 1.2 lists the contributions of this research. Section 2 provides an overview of the background, a comprehensive summary of the related work, and a critical analysis of the state-of-the-art in computer aided diagnosis of knee OA.

² ImageNet Large Scale Visual Recognition Challenge




Fig. 1: A healthy knee and a knee joint affected with OA.

Section 3 introduces the public datasets used in this study. Section 4 presents the baseline methods using hand-crafted features and the proposed approaches in this chapter for automatic detection of knee joints in the radiographs. Section 5 presents the baseline methods and the proposed approaches in this chapter to quantify knee OA severity using CNNs. Section 6 concludes this chapter by analyzing the current work and summarizing the research methodology, and providing future directions of research based on the proposed methods.

1.1 Knee Osteoarthritis

Knee Osteoarthritis (OA) is a debilitating joint disorder that mainly degrades the knee articular cartilage and in its severe stages it causes excruciating pain and often leads to total joint arthroplasty. In general, knee OA is characterized by joint pain, cartilage wear, and bony growths. Knee OA has a high-incidence among the elderly, obese, and those with a sedentary lifestyle. Early diagnosis is crucial for clinical treatments and pathology [3, 6].

1.1.1 Diagnostic Features of Knee OA

Clinically, the major pathological features for knee OA include joint space narrowing, osteophytes formation, and sclerosis [6, 21]. Figure 1 shows the anatomy of a healthy knee and a knee affected with osteoarthritis, and the characteristic features of knee OA. The causes for knee OA include mechanical abnormalities such as degradation of articular cartilage, menisci, ligaments, synovial tissue, and sub-chondral bone.

The major clinical features, joint space narrowing and osteophyte formation, are easily visualized using radiographs [6, 7, 22]. Despite the introduction of several imaging methods such as magnetic resonance imaging (MRI), computed tomography (CT), and ultrasound for augmented OA diagnosis, radiographs have traditionally been preferred [23, 22], and remain as the main accessible tool and “gold standard” for preliminary knee OA

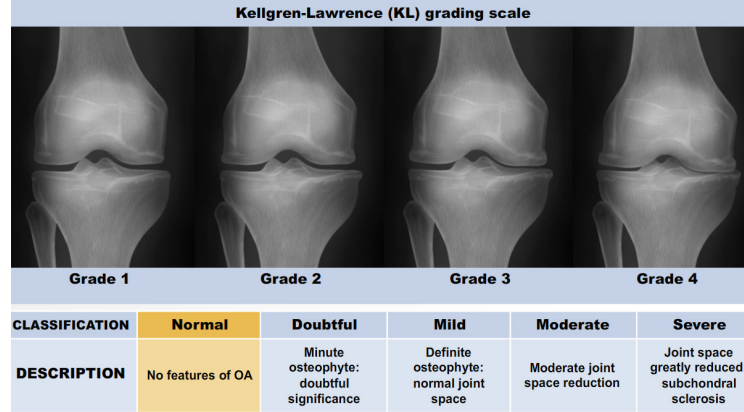


Fig. 2: The Kellgren and Lawrence grading system to assess the severity of knee OA.

diagnosis [1, 3]. Inspired by the previous successful approaches in the literature for early identification [3] and automatic assessment of knee OA severity [4, 6, 23], the focus is on radiographs in this work. More importantly, there are public datasets available that contain radiographs with associated ground truth. Public datasets for knee OA study, such as the OAI³ and the MOST⁴ datasets, provide radiographs with KL scores, and the OARSI⁵ readings for distinct knee OA features such as JSN, osteophytes, and sclerosis.

1.1.2 Radiographic Classification of Knee OA

Knee OA develops gradually over years and progresses in stages. In general, the severity of knee OA is divided into five stages. The first stage (stage 0) corresponds to normal healthy knee and the final stage (stage 4) corresponds to the most severe condition (see Figure 2). The most commonly used systems for grading knee OA are the International Knee Documentation Committee (IKDC) system, the Ahlback system, and the Kellgren & Lawrence (KL) grading system. The other widely used non-radiographic knee OA assessment system is WOMAC⁶, which measures pain, stiffness, and functional limitation. The public datasets, the OAI and the MOST used in this work, are provided with the KL grades and they are used as the ground truth to classify the knee OA X-ray images.

Kellgren and Lawrence Scores.

The KL grading scale was approved by the World Health Organisation as the reference standard for cross-sectional and longitudinal epidemiologic studies [7, 22, 24, 25]. The KL grading system is still considered the gold standard for

³ Osteoarthritis Initiative

⁴ Multicenter Osteoarthritis Study

⁵ Osteoarthritis Research Society International

⁶ Western Ontario and McMaster Universities Osteoarthritis Index

initial assessment of knee osteoarthritis severity in radiographs [1, 5–7]. Figure 2 shows the KL grading system. The KL grading system categorizes knee OA severity into five grades (grade 0 to 4). The KL grading scheme for quantifying knee OA severity from X-ray images is defined as follows [1, 5]:

- Grade 0 : absence of radiographic features (cartilage loss or osteophytes) of OA.
- Grade 1 : doubtful joint space narrowing (JSN), osteophytes sprouting, bone marrow oedema (BME), and sub-chondral cyst.
- Grade 2 : visible osteophytes formation and reduction in joint space width on the antero-posterior weight-bearing radiograph with BME and sub-chondral cyst.
- Grade 3 : multiple osteophytes, definite JSN, sclerosis, possible bone deformity.
- Grade 4 : large osteophytes, marked JSN, severe sclerosis, and definite bone deformity.

1.2 Contributions

The research contributions of this work are as follows.

- Proposing a novel and highly accurate technique to automatically detect and localise the knee joints from the X-ray images using a fully convolutional network (FCN).
- Developing a classifier based on a CNN to assess knee OA severity that is highly accurate in comparison to previous methods.
- Proposing a novel approach to train a CNN with a weighted ratio of two loss functions: categorical cross entropy and mean squared error with the natural benefit of predicting knee OA severity in ordinal (0, 1 to 4) and continuous (0–4) scales.
- Developing an ordinal regression approach using CNNs to automatically quantify knee OA severity in a continuous scale.
- Developing an automatic knee OA diagnostic system i.e. an end-to-end pipeline incorporating the FCN for automatically localising the knee joints and the CNN for automatically quantifying the localised knee joints.

2 Related Work and Background

The automatic assessment of knee OA severity from radiographs has been approached as an image classification problem [1, 3, 4]. According to the literature and in the machine learning approach to automatically assess knee OA severity, the first step is to localise the region of interest (ROI) that is to detect and extract the knee joint regions from the radiographs, and the next step is to classify the localised knee joints. First, the different approaches for detecting (or localising) the knee joint regions in the radiographs are outlined. Next, the approaches in the literature to assess knee OA severity are

investigated and the focus is on the automated methods. This section concludes with a discussion outlining the limitations in the state-of-the-art methods on automatic detection of knee joints and automatic assessment of knee OA severity, and how these limitations can be addressed.

2.1 Detecting Knee Joints in Radiographs

There are several approaches in the literature for detecting and segmenting knee joints and specific parts of the knee such as cartilage, menisci, and bones structures from 3D MRI and CT scan images [26, 27]. Nevertheless, the existing approaches are less accurate for automatically detecting the knee joints in radiographs [27, 28]. According to the literature, detecting knee joints remains a challenging task [27, 29]. In this chapter, automated methods for detecting knee joints in radiographs are investigated. The advantages of automatic methods are discussed and the need to investigate such methods are emphasized. Previous approaches in the literature that investigate the knee joints in radiographs can be categorized into manual, semi-automatic, and fully automatic, based on the level of manual intervention required [26, 27].

2.1.1 Manual Methods

Expert radiologists or trained physicians visually examine the knee joint regions and trace the structures using simple image processing and computer vision-based tools in radiographs, and may even use CAD-based measurements for assessing knee OA severity [27]. The expert knowledge-based manual segmentations are useful to build an atlas or template of anatomical structures, which are used to develop advanced interactive and automatic segmentation methods [27]. The knee joints labelled manually are reliable and are often used as ground truth for evaluating automatic methods [30, 31]. Nevertheless, such manual methods are subjective, highly experience-based, and they are laborious and time-consuming when a large number of subjects are to be examined.

There are previous studies in the literature that use manually-defined ROIs (knee joints) in radiographs for assessing knee OA severity. Hirvasniemi et al.[32] quantified the differences in bone density using texture analysis and local binary patterns (LBP) in plain radiographs to assess knee osteoarthritis. Woloszynski et al.[33] developed a signature dissimilarity measure for the classification of trabecular bone texture in knee radiographs. In both these methods, the ROIs are manually marked and extracted for texture analysis.

2.1.2 Semi-automatic Methods

Semi-automatic or interactive methods are developed to minimize manual interventions by automating essential steps in the detection and segmentation process [26, 34]. These methods often include manual initialization with low-level image processing, followed by manual evaluations and corrections of the results [35]. The main advantage of the semi-automatic methods are

flexibility in manual intervention that allow incorporating expert knowledge, plus the use of advanced computer vision-based tools to automate the essential steps. An expert may improve the detection and segmentation performance through tuning the essential parameters for instance seed region and threshold values in region growing, initial shape of active models, and delineating the required contour [27] to define the region of interest. However, these methods may not be reproducible due to inter-observer or inter-user variations and there is a possibility of oversight or human error in the manual evaluations.

There are some knee OA studies in the literature which use semi-automatic methods to detect the knee joints in radiographs. Knee OA computer aided diagnosis (KOACAD)[6] is an interactive method to measure the joint space narrowing, osteophytes formation and joint angulation in radiographs. In KOACAD, a Roberts filter is used to obtain the rough contour of tibia and femur bone structures and a vertical neighbourhood difference filter is used to identify points with high absolute values of difference of scales. The centre of all the points is calculated and a rectangular region around the centre, of size 480×200 pixels, is selected as the knee joint region. This system has purportedly provided accurate assessment of structural severity of knee OA after detecting the knee joint regions. However, human intervention is required for plotting various lines for the measurement, and automatic detection is not feasible with this system.

Knee images digital analysis (KIDA) is a tool to analyse knee radiographs interactively, proposed by Marijnissen et al. [21]. KIDA quantifies the individual radiographic features of knee OA like medial and lateral joint space width (JSW) measurements, subchondral bone densities and osteophytes. However, this interactive tool can only be used by experts for quantitative measurements and requires expert intervention for objective quantitative evaluation.

Duryea et al. [36] proposed a trainable rule-based algorithm (software) to measure the joint space width between the edges of the femoral condyle and the tibial plateau on knee radiographs. Contours marking the edges of the femur and tibia are automatically generated. This interactive method can be used to monitor joint space narrowing and the progression of knee osteoarthritis.

2.1.3 Automatic Methods

Automatic segmentation methods have become an essential part of computer aided diagnosis and clinical decision support systems [29]. These methods are fast and accurate, and they are highly beneficial in clinical trials and pathology [27]. According to the literature, there have been multiple attempts to automatically localise knee joints in radiographs. Nevertheless, this task still remains a challenge.

Podsiadlo et al. [37] proposed an automated system for the prediction and early diagnosis of knee OA. In this approach, active shape models and morphological operations are used to delineate the cortical bone plates and locate the ROIs in radiographs. This approach is developed for selection of

tibial trabecular bone regions in the knee joints as ROIs. Nevertheless, this approach can be extended to localise the entire knee joint. A set of 40 X-ray images are used for training and 132 X-ray images are used for testing in this method. The automatic detections from this method are compared to the gold standard, which contains manually annotated ROIs from the expert radiologists and the similarity indices (SI) are calculated. This method achieved SI of 0.83 for the medial and 0.81 for the lateral regions of the knee joints.

Shamir et al. [1] proposed template matching for automatic knee joint detection in radiographs. Template matching uses predefined joint centre images as templates and calculates Euclidean distances over every patch in an X-ray image using a sliding window. The image patch with the shortest distance is recorded as the detected knee joint centre. After detecting the centre, an image segment of 700×500 pixels around the centre is extracted as the knee joint region. The X-ray images from the BLSA dataset are used in this method. In total 55 X-ray images from each grade are used for the experiments, such that 20 images from each grade for training and 35 images from each grade for testing. Shamir et al. reported that template matching was successful in finding the knee joint centres in all the X-ray images in their dataset.

Anifah et al. [38] investigated template matching and contrast-limited adaptive histogram equalisation for detecting knee joints and quantifying joint space area. In total 98 X-ray images are used in this method. The detection accuracy achieved by this method varies from 83.3% to 100% for the left knees and 60.4% to 100% for the right knees. Template matching is a simple and relatively fast method. However, this method is ad hoc, entirely based on the set of templates used and is unlikely to generalise well for larger datasets.

Recently, Tuilpin et al. [29] investigated a SVM-based method to automatically localise knee joints in plain radiographs. This method uses knee anatomy-based region proposals, and the best candidate region from the proposals are selected using histogram of oriented Gradients (HOG) as feature descriptors and a SVM. This method generalises well in comparison to the previous methods and shows reasonable improvement in automatic detections with mean intersection over union (IOU) of 0.84, 0.79 and 0.78 on the public datasets MOST, Jyvaskyla, and OKOA.

2.2 Assessing Radiographic Knee OA Severity

The key pathological features of knee OA include joint space narrowing, osteophytes (bone spurs) formation, and sclerosis (bone hardening) [6, 21]. All these features are implicitly integrated in composite scoring systems, like Kellgren & Lawrence (KL) grading system, to quantify knee OA severity [6, 21], and the OARSI readings provide the gradings of distinct knee OA features. There are two common approaches for assessing knee OA severity in plain radiographs: 1) quantifying the distinct pathological features of knee OA, and 2) automatic classification based on composite scoring systems such as KL grades.

2.2.1 Quantitative Analysis

The most conventional system to assess radiographic knee OA severity has been KL gradings [3, 5, 6]. Nevertheless, some researchers [6, 21] argue that categorical systems like KL gradings are limited by incorrect assumptions that the progression of distinct OA features like JSN and osteophytes formation is linear and constant, their relationships are proportional, and such grading systems are less sensitive to small changes in distinct features. Therefore, quantification of individual features of knee OA is required to overcome the problems with KL gradings and to improve the overall radiographic assessment of knee OA [6, 21]. The OARSI has published a radiographic atlas of individual features to assess and to quantitatively evaluate the knee OA features [6].

Interactive methods (KOACAD [6] and KIDA [21]) measure individual knee OA radiographic features such as joint space width (JSW), osteophyte area, sub-chondral bone density, joint angle, and tibial eminence height as continuous variables. These measurements were compared to KL gradings and significant differences were found between healthy knees and knees with OA. In this context, a trainable rule-based algorithm has also been proposed [36] to measure the minimum joint space width (mJSW) between the edges of the femoral condyle and the tibial plateau, and thus to monitor the progression of knee OA. Podsiadlo et al. [37] have used a slightly different approach for quantitative knee OA analysis. In this method, the trabecular bone regions of the tibia are automatically located as the ROI after delineating the cortical bone plates using active shape models, followed by fractal analysis of bone textures for the diagnosis of knee OA. In a similar approach, Lee et al. [39] use active shape models to detect the tibia and femur joint boundaries, and calculate anatomical geometric parameters to diagnose knee OA.

Even though these methods are simple to implement, objective, and purportedly accurate in evaluating radiographic knee OA, a great deal of manual intervention is required. Hence, these methods become very time-consuming and laborious when large numbers of subjects are to be investigated. Furthermore, the measurements from these methods are prone to inter- and intra-observer variability.

2.2.2 Automatic Classification

After the introduction of radiography-based semi quantitative scoring systems like KL gradings, the assessment of radiographic knee OA severity has been approached as an image classification problem [2, 3, 40–42]. According to the literature, the most common approach to classify knee OA images includes two steps: 1) extracting image features from the knee joints, and 2) applying a classification algorithm on the extracted features. A brief review of such approaches follows.

Subramoniam et al. [40, 41] investigated two methods using: 1) the histograms of local binary pattern extracted from knee images and a k-Nearest neighbour classifier [40] and 2) Haralick features extracted from the ROI of knee images and a SVM [41]. Thomson et al. [2] proposed an automated method that uses features

derived from tibia and femur bone shapes, and image textures extracted from the tibia with a simple weighted sum of the outputs of two random forest classifiers. Deokar et al. [42] investigated an artificial neural network based approach for knee OA images classification using grey level co-occurrence matrix (GLCM) textures, shape, and statistical features. Even though these methods claim high accuracy, the datasets are not publicly available and these datasets contain only a few hundred radiographs. The classification accuracies of all these methods for public datasets like the OAI and the MOST need to be studied to derive conclusive results.

In this context, there are two approaches that use large public datasets like the OAI: 1) WNDCHRM, and 2) an artificial neural network-based scoring system. Shamir et al. proposed WNDCHRM, a multi purpose medical image classifier to automatically assess knee OA severity in radiographs[1,3]. A set of features based on polynomial decompositions, high contrast, pixel statistics, and textures are used in WNDCHRM. Besides extracting features from raw image pixels, features extracted from image transforms like Chebyshev, Chebyshev-Fourier, Radon, and Gabor wavelets are included to expand the feature space [1, 4, 5]. From the entire feature space, highly informative features are selected by assigning feature weights based on a Fisher discriminant score for all the extracted features [1,3,5]. WNDCHRM uses a variant of the k-Nearest Neighbour classifier.

Yoo et al. [43] have built a self-assessment scoring system and an artificial neural network (ANN) model for radiographic and symptomatic knee OA risk prediction. In a recent approach, Tiulpin et al. [44] presented a new computer-aided diagnostic approach based on deep Siamese CNNs, which were originally designed to learn a similarity metric between pairs of images. However, rather than comparing image pairs, the authors extend this idea to similarity in knee x-ray images (with 2 symmetric knee joints). Splitting the images at the central position and feeding both knee joints into a separate CNN branch allows the network to learn identical weights for both branches. They outperform the previous approaches by achieving an average multi-class testing accuracy score of 66.7 % on the entire OAI dataset.

2.3 Discussion

According to the literature, the automatic quantification of knee OA severity involves two steps: 1) automatically detecting the ROI, and 2) classifying the detected knee joints. Many previous studies investigated automatic methods for both localisation and classification of knee joint images, but still these tasks remain a challenge.

The common approaches in the literature for automatic detection of knee joints in radiographs include template matching [1, 38], active shape models and morphological operations [37], and a classifier-based sliding window method[29]. Template matching and active shape models based approaches do not generalise well and are slow for large datasets. Classifier-based methods that use hand-crafted features are subjective and the classification accuracy is influenced by

the choice of extracted features. Therefore, there is still a need for an automated method for detecting knee joints in radiographs which gives high accuracy and precision. A deep learning based method for this is investigated in this chapter.

There are several approaches in the literature for knee OA image classification that have extracted and tested many image features, such as Haralick textures [41], Gabor textures [2], GLCM textures [42], local binary patterns [40], shape, and statistical features of knee joints [42]. There is even an approach that uses a large set of features based on pixel statistics, object and edge statistics, texture, histograms, and transforms [7, 5, 4]. Different classifiers have been tested for knee OA images classification such as k-Nearest Neighbour [40, 1], SVM [41], and random forest classifiers [2]. However, all these approaches have achieved low multi-class classification accuracy, and in particular classifying successive grade knee OA images still remains a challenging task. There is a need for a highly accurate real world automated system that can be used as a support system by clinicians and medical practitioners for knee OA diagnosis.

In recent years, many methods using manually designed or hand-crafted features have been outperformed by approaches that learn feature representations using deep neural networks. In particular, convolutional neural networks (CNN) have become highly successful in many computer vision tasks like object detection, face recognition, content based image retrieval, pose estimation, and shape recognition, and even in medical applications such as knee cartilage segmentation in MRI scans[16], brain tumour segmentation in magnetic resonance imaging (MRI) scans[17], multi-modality iso-intense infant brain image segmentation[18], pancreas segmentation in CT images [19], and neuronal membrane segmentation in electron microscopy images [20]. CNNs for automatically quantifying knee OA severity is investigated in this work. The next section introduces the public knee OA datasets.

3 Public Knee OA Datasets

The data used for the experiments and analysis in this study are bilateral PA fixed flexion knee X-ray images. Figure 3 shows some samples of knee X-ray images from the dataset. Due to variations in X-ray imaging protocols, there are some visible artefacts in the X-ray images (Figure 3).

The datasets are from the Osteoarthritis Initiative (OAI) and Multicenter Osteoarthritis Study (MOST) in the University of California, San Francisco. These are standard public datasets used in knee osteoarthritis studies.

3.1 OAI Dataset

The baseline cohort of the OAI dataset contains MRI and X-ray images of 4,746 participants. In total 4,446 X-ray images are selected from the entire cohort based on the availability of KL grades for both knees as per the assessments by Boston University X-ray reading centre (BU). In total there are 8,892 knee images. Figure 4 shows the distribution as per the KL grades.




Fig. 3: Samples of bilateral PA fixed flexion knee OA radiographs.

3.2 MOST Dataset

The MOST dataset includes lateral knee radiograph assessments of 3,026 participants. In total 2,920 radiographs are selected in this study based on the availability of KL grades for both knees as per baseline to 84-month longitudinal knee radiograph assessments. There are 5,840 knee images in this dataset. Figure 5 shows the distribution as per KL grades.

4 Automatic Detection of Knee Joints

Classification of knee OA images and the assessment of severity conditions can be achieved by examining the characteristic features of knee OA: variations in the joint space width and the osteophytes (bone spurs) formations in the knee joints [6]. Radiologists and medical practitioners examine only the knee joint regions in the X-ray images to assess knee OA. Hence, the region of interest (ROI) for classifying knee OA images is only the knee joint regions (left and right knees). Figure 6 shows the ROI in a X-ray image. The author believes that it is better to focus on the ROI instead of the entire X-ray image for accurate classification and this is also computationally economical. For these reasons, automatically detecting and extracting the knee joint regions from the X-ray images becomes an essential pre-processing step, before classification.

4.1 Baseline Methods

First, template matching for automatic detection of the knee joints [3] is implemented as a baseline. Next, the author proposes an SVM-based




Fig. 4: The OAI baseline data set distribution based on KL grades.




Fig. 5: The MOST data set distribution based on KL grades.

automatic detector for this. The implementation details and outcomes of these methods are discussed in this section.

4.1.1 Template Matching

In digital image processing, template matching is a technique for finding portions of an image that are similar to a standard template image. Shamir et al. [3] proposed this approach for automatically detecting the centre of the knee joints. As a baseline, the template matching approach is adapted. The steps involved in this method are as follows:

- First, the radiographs are downscaled to 10% of the original size and subjected to histogram equalisation for intensity normalisation. This step is followed as proposed by Shamir et al. [3].
- An image patch (20×20 pixels) containing the centre of the knee joint is taken as a template. 5 image patches are taken from each grade, so that in total 25 patches are pre-selected as templates. Figure 7 shows the pre-




Fig. 6: A knee OA X-ray image with the region of interest: the knee joints.

selected knee joint centres of size 20×20 pixels extracted from the knee joint images as templates.

- Each image is scanned by an overlapping (20×20) sliding window. For each location at an interval of 10 pixels, distances (Euclidean) between an image patch (20×20 pixels) and 25 pre-selected templates (patches with knee joint centre) are computed using;

$$dist_{i,w} = \sqrt{\sum_{y=1}^{20} \sum_{x=1}^{20} (I_{x,y} - W_{x,y})^2},$$

where $I_{x,y}$ is the intensity of pixel (x, y) in the knee joint image I , $W_{x,y}$ is the intensity of pixel (x, y) in the sliding window, and $dist_{i,w}$ is the Euclidean distance between the knee joint image (I) and the sliding window W .

- In total, 25 different distances are calculated at each location of the sliding window for the 25 templates, and the shortest among the 25 distances is recorded.
- The window with the smallest Euclidean distance is selected as the centre of the knee joint after scanning the image with a sliding window, and a fixed size region (700×500 pixels) around this centre is extracted as the knee joint region from the X-ray image.
- The input X-ray images are horizontally split in half to isolate left and right knees separately and the sliding window is run on both halves.

Experiments and Results. For the experiments on template matching, the baseline data sample of 200 progression and incidence cohort subjects under the knee OA study is used. This dataset contains in total 191 X-ray images (382 knee joints) and it is a subset of the large OAI dataset.

In this implementation, five different sets of templates (each set with 25 templates) are used to show the influence of templates on knee joint detections.




Fig. 7: Pre-selected knee joint centres (20×20 pixels) extracted from knee joint images for template matching.

Table 1: Detection of knee joint centres using template matching method.

Templates	True Positives	False Positives	Precision
Set 1	87	295	22.8 %
Set 2	78	304	20.4 %
Set 3	99	283	25.9 %
Set 4	116	266	30.3 %
Set 5	55	327	14.4 %

The templates are selected from a separate training set. Visual inspection is used to evaluate the results of template matching by plotting a bounding box (20×20 pixels) on the image patch that recorded the shortest Euclidean distance after template matching. Table 1 shows the total number of true positives (the detected knee joint centres), the total number of false positives and the precision.

It is clearly evident from the results (Table 1) that template matching is not precise in detecting the knee joints and that the detections are heavily dependent on the choice of templates. The number of templates was increased to 50, but still there was no further improvement in the results. The reason for low-performance of the template matching is that the computations are mainly based on the intensity level difference between an image patch and a template. There are also possibilities for image patches not around the knee joint, having the shortest Euclidean distance to a template in the set and thus, being detected as matches. In the next section, a new SVM-based method is investigated to improve the detection of the knee joints.

4.1.2 SVM-Based Detection

Standard template matching is not scalable and produces poor detection accuracy on large datasets like the OAI. We proposed a classifier-based model to automatically detect the knee joints in the X-ray images [45]. The idea is to use well-known Sobel edge detection [46] for detecting the knee joints. The two major steps involved in this method are 1) training a classifier and 2) developing a sliding window detector.

Training a Classifier. First, image patches (20×20 pixels) are generated from the input X-ray images. The image patches containing the knee joint centre (20×20 pixels) are used as positive samples and randomly sampled patches excluding the knee joint centre are used as negative samples. In total, 200 positive and 600 negative samples are used. The image patches (samples) are split into training (70%) and test (30%) sets. Sobel horizontal image gradients are extracted as features from all these samples to train a classifier. The powerful and well-known SVM is used for classification. A linear SVM is fitted with default parameters ($C=1$, and linear kernel), using Sobel horizontal image gradients as the features.

Before settling on Sobel horizontal image gradients as features, the state-of-the-art features such as histogram of oriented gradients, Tamura and Haralick textures, and the Gabor features were tested. The HOG features are highly accurate and efficient in object detection and human detection [47]. The Tamura, Haralick, and Gabor features are highly influential and top-ranked among the features used in WNDCHRM for knee OA image classification [1, 3, 5, 6]. The Sobel operator or Sobel filter uses vertical and horizontal image gradients to emphasise the edges in images [46]. From these, the horizontal image gradients are used as the features for detecting the knee joints centres. Intuitively, the knee joint images primarily contain horizontal edges that are easy to detect.

Sliding Window Detector. To detect the knee joint centre from both left and right knees, input images are split in half to isolate left and right knees separately. A sliding window (20×20 pixels) is used on either half of the image, and the Sobel horizontal gradient features are extracted for every image patch. The image patch with the maximum score based on the SVM decision function is recorded as the detected knee joint centre, and the area (200×300 pixels) around the knee joint centre is extracted from the input images using the corresponding recorded coordinates. Figure 8 shows an instance of a detected knee joint and the extracted ROI in a X-ray image.

Results and Discussion. In total, 200 image patches with the knee joint centres as positive samples and 600 image patches that exclude the centre of knee joint as negative samples are used. These images are split into training (70%) and test (30%) sets. Fitting a linear SVM with the training data produced a 5-fold cross validation accuracy of **95.2%** and an accuracy of **94.2%** for the




Fig. 8: Detecting the knee joint centres and extracting the knee joints.

Table 2: Classification metrics of the SVM for detection.

Class	Precision	Recall	F_1 score
Positive	0.93	0.84	0.88
Negative	0.95	0.98	0.96
Mean	0.94	0.94	0.94

test data. Table 2 shows the precision, recall, and F_1 scores of this classification. To evaluate the automatic detection, the ground truth is generated by manually annotating the knee joint centres (20×20 pixels) in 4,446 radiographs using an annotation tool that we developed, which recorded the bounding box (20×20 pixels) coordinates of each annotation.

The well-known Jaccard index (JI) is used to give a matching score for each detected instance. The Jaccard index $JI(A, D)$ is given by,

$$JI(A, D) = \frac{A \cap D}{A \cup D} \quad (1)$$

where A , is the manually annotated and D is the automatically detected knee joint centre using the proposed method.

Table 3 shows the resulting average detection accuracies based on thresholding of Jaccard indices. The mean JI for the template matching and the classifier methods are **0.1** and **0.36**. From Table 3, it is evident that the proposed method is more accurate than template matching. This is due to the fact that template matching relies upon the intensity level difference across an input image. Thus, it is prone to matching a patch with small Euclidean distance that does not actually correspond to the knee joint centre. Also, the templates are varied in a set, and it is observed that the detection is highly dependent on the choice of templates. Template matching is similar to a k -nearest neighbour classifier with $k = 1$.

Table 3: Comparison of template matching and the proposed SVM-based method.

Method	$JI = 1$	$JI \geq 0.5$	$JI > 0$
Template Matching	0.3 %	8.3 %	54.4 %
Proposed Method	1.1 %	38.6 %	81.8 %

The reason for higher accuracy in the proposed method is the use of horizontal edge detection instead of intensity level differences. The knee joints primarily contain horizontal edges and thus are easily detected by the classifier using horizontal image gradients as features. The proposed method is approximately $80\times$ faster than template matching; for detecting all the knee joints in the dataset comprising 4,446 radiographs, the proposed method took ~ 9 minutes and the template matching method took ~ 798 minutes.

Despite sizeable improvements in accuracy and speed using the proposed approach, detection accuracy still falls short. Therefore, manual annotations for the incorrect detections from this method were substituted to investigate KL grade classification performance independently of knee joint detection. The next Section describes the proposed methods for automatically localising the knee joint region using fully convolutional neural networks.

4.2 Fully Convolutional Network Based Detection

A typical CNN architecture consists of three main types of layers: convolutional, pooling and fully-connected or dense layers. A fully convolutional network (FCN) is similar to a CNN, but the fully-connected layers are replaced by convolutional layers [48]. A FCN consists of mostly convolutional layers and if pooling layers are used, then suitable up-sampling layers are added before the last convolutional layer. The two major differences of FCNs over CNNs can be summarised as:

- FCNs are trained end-to-end to make pixel-wise predictions [48]. Even the decision-making layers at the last stage of the network use learned convolutional filters.
- The input image size need not be fixed as there are no fully-connected layers in the FCN. CNNs with fully connected layers can operate only on a fixed size input.

FCNs have achieved great success in semantic segmentations of general images [48]. Recent approaches using FCNs for medical image segmentation show promising results [49–51]. Motivated by this, the use of FCN is investigated in this chapter for automatically detecting the knee joints. Two approaches are developed for localising the knee joints: 1) training a FCN to detect the centre of knee joints and extract a fixed-size region around the detected centre, and 2) training a FCN to detect the ROI and thus extract the knee joints directly.




Fig. 9: Automatic localisation of knee joints with reference to the centre of the knee joints.

4.2.1 Localisation with Reference to Knee Joint Centre

In the initial approach to localise the knee joints in X-ray images using a FCN, a similar strategy to template matching and the SVM based methods is followed; that is to detect the centre of knee joints and to extract the ROI with reference to the detected centres. Figure 9 shows the steps involved in this method: training a FCN to detect the knee joint centres (20×20 pixels), computing the coordinates of the centres from the FCN output, and extracting a fixed size region as knee joints. In the next section, the experimental data and the ground truth used to train the FCNs are introduced.

Dataset and Ground Truth Generation. The data used for the experiments are taken from the baseline cohort of the OAI dataset. In total 4,446 X-ray images are selected from the entire dataset based on the availability of KL grades for both knee joints. The knee joint centres in all these X-ray images are manually annotated, after downscaling to 10% of the actual size. Binary masks of size 20×20 pixels are marked around the knee joint centres using the annotations. Figure 10 shows an instance of an input X-ray image and the binary mask annotations corresponding to the knee joint centres. The image patches from the masked region i.e. the knee joint centres, are taken as positive training samples and the patches from rest of the image are taken as the negative training samples to train an FCN. The dataset is split into training (3,333 images) and test (1,113 images) sets.

Training Fully Convolutional Neural Networks

To start, a FCN is configured with a lightweight architecture containing 4 convolutional layers followed by a fully convolutional layer, which is a convolutional layer with a kernel size $[1 \times 1]$ and that uses a *sigmoid activation*. FCNs use fully convolutional layers at the last stage to make pixel-wise predictions [48]. Table 4 shows the network configuration in detail. Each convolution layer is followed by a ReLU layer.

The network parameters are trained from scratch with training samples of knee OA radiographs from the OAI dataset. The dataset is split into training (3,333 images) and test (1,113 images) sets. The ground truth for training the network are binary images with masks specifying the ROI: the knee joints. The network is trained to minimise the total *binary cross entropy* between the




Fig. 10: (a) An input X-ray image and (b) The binary mask annotations for knee joint centres.

Table 4: Initial FCN Configuration for detecting the knee joint centres.

Layer	Kernel	Kernel Size
Conv1	32	3×3
Conv2	32	3×3
Conv3	64	3×3
Conv4	64	3×3
Conv5	1	1×1

predicted pixels and the ground truth. *Stochastic gradient descent* (SGD) with default parameters: learning rate = 0.01, decay = $1e^{-6}$, momentum = 0.9, and nesterov = True, is used. The network is trained for 40 epochs and the batch size is 10. Figure 11 shows an instance of the test input, the ground truth and the output (pixel-wise predictions) of the FCN. From the predictions of this FCN, it is observed that the network is able to slightly detect the edges of the knee joints and these are promising initial results. In an attempt to improve the detections, the FCN configurations are experimented and for this the hyper-parameters of the network are tuned.

Receptive Field. When dealing with high-dimensional inputs such as images, it is impractical to connect neurons in the current level to all the neurons in the previous volume. Instead, each neuron is only connected to a local region of the input volume. The spatial extent of this connectivity is a hyper-parameter called the receptive field of the neuron [52]. The receptive field size, otherwise termed the effective aperture size of a CNN, shows how much a convolutional node sees of the input pixels (patch) that affects a node’s output. The effective aperture size depends on kernel size and strides of the previous layers. For instance, a 3×3 kernel can see a 3×3 patch of the previous layer and a stride of 2 doubles what all succeeding layers can see.




Fig. 11: An instance of input, ground truth and output (predictions) of FCN.

Table 5: FCN for detecting the knee joint centres.

Layer	Kernel	Kernel Size
Conv1	32	7×7
Conv2	64	3×3
Conv3	96	3×3
Conv4 (fullyConv)	1	1×1

The receptive field size of neurons in the final layer of the FCNs is calculated and used to analyse the output of FCNs and the overall detection results. The receptive field size of a neuron in the final layer (Conv5) of the initial FCN configuration (Table 4.1) is 9, which is low and may be a reason for poor performance of this network. Larger convolutional kernel sizes to increase the receptive field of the network is investigated. The forthcoming Section will show that a network (Table 9) with larger receptive field gives the best results for detecting the knee joint centres.

Tuning the FCN Hyper-parameters. VGG-M-128 [53], the deep convolutional neural network developed by the Oxford visual geometry group (VGG) uses kernel size 7×7 in the first convolutional layer and 5×5 in the following convolutional layer. Inspired by this, kernel sizes of 5×5 , and 7×7 for the first convolutional layer are tested retaining the other settings. The kernel size 7×7 gives better results in this configuration. This is because of the larger receptive field size of the 7×7 kernel in comparison to the 3×3 kernel.

Next, the experiments are conducted by varying the number of convolutional layers and also the number of filters (kernel) in a convolutional layer, before obtaining the configuration that gave the best results based on visual observations. Table 5 shows the configuration of the network derived from the initial configuration and the receptive field size of a neuron in the final layer (Conv4) is 11. The networks are trained with 3,333 images and tested on 1,113 images from the OAI dataset.

There is an improvement in the detections using this network in comparison to the previously tested configurations. Figure 12 shows an instance of the output predictions of this network. To quantitatively evaluate the automatic detections, the well-known Jaccard Index is used.




Fig. 12: An input image, ground truth, and outcome of the final FCN.

Quantitative Evaluation. A simple contour detection is used and the Jaccard index i.e. the overlap statistics calculated by the Intersection over Union (IoU) to evaluate the automatic detections of the FCN. The steps involved are as follows:

- First, the objects are detected i.e. the knee joint regions from the output image of the FCN using simple contour detection [54]. Contours can be explained simply as a curve joining all the continuous points (along the boundary), having the same colour or intensity. The contours are a useful tool for shape analysis and simple object detection and recognition. In this method, first the images are converted to binary by applying Otsu’s threshold. Next, the contours of the objects or shapes in the binary image are automatically detected and recorded [54].
- Next, the detected objects in the image are sorted based on the area and from these the top two are selected. This is to eliminate noise or other faint edges picked up by the FCN.
- The centroids of the largest two detected regions are recorded as the knee joint centres.
- A binary mask of 20×20 pixels size is marked around each detected knee joint centre.
- The Jaccard index is computed for each image with the masks of predicted centres and the masks predefined using manual annotation i.e. the labels used for training FCN.

In total 1,113 X-ray images (2,226 knee joints) are included in the test set. The FCN with the final configuration detects 1,851 knee joints in the test set with Jaccard index ≥ 0.5 , the accuracy of detection is **83.2%** with a mean 0.66 and standard deviation 0.18. This is an improvement in comparison to previous approaches but still falls short of perfect detections. The pooling and up-sampling layers in the FCN are varied and experimented in an attempt to improve the detection accuracy. This may help to increase the receptive field size and in turn improve the overall detections.

FCN with Pooling and Up-sampling Layers. Two max pooling layers with stride 2 and up-sampling by a factor of 4 are included to the previous configuration (Table 5). Table 6 shows the FCN architecture in detail. Each convolutional layer is followed by a ReLU activation.

Table 6: FCN with pooling and up-sampling layers.

Layer	Kernel	Kernel Size	Strides
Conv1	32	7×7	1
MaxPool2	–	2×2	2
Conv3	64	3×3	1
MaxPool4	–	2×2	2
Conv5	96	3×3	1
UpSamp6	–	4×4	1
Conv7 (fullyConv)	1	1×1	1




Fig. 13: Prediction of the FCN with max pooling and up-sampling layers.

Figure 13 shows the output of this network for a test image. On visual observation, the output image contains less noise and the detections are improving compared to the previous approaches, even though the output image resolution is low. This is due to the inclusion of pooling and up-sampling stages to the network and this has increased the receptive field size of the final layer (Conv7) to 34. The number of convolutional-pooling stages is increased, to see if there is improvement in the detections. Table 7 shows the architecture of this network in detail.

From the output of this FCN, it can be observed that the detections become more precise in comparison to the previous networks even though the resolution is low in comparison to the previous networks. Figure 14 shows an instance of the input test image, ground truth and the FCN output.

The outcomes of this FCN are evaluated using Jaccard index and the detection accuracy is **96.7%**, that is in total 2,152 out of 2,226 knee joints are detected with a Jaccard index ≥ 0.5 . The Jaccard index mean is 0.74 and standard deviation is 0.13. The detection accuracy is high in comparison to the previous networks. Table 8 shows the detection accuracy of the FCN for the Jaccard index values at 0.25, 0.5 and 0.75.

This FCN (Table 7) has three convolutional-pooling stages. A configuration with 4 convolutional-pooling stages followed was tested by adding an up-sampling layer with kernel size (16×16). There was no improvement in the detection accuracy for this configuration.

Table 7: FCN with 3 Convolution-Pooling stages for detecting the knee joint centres.

Layer	Kernel	Kernel Size	Strides
Conv1	32	7×7	1
MaxPool2	–	2×2	2
Conv3	32	3×3	1
MaxPool4	–	2×2	2
Conv5	64	3×3	1
MaxPool6	–	2×2	2
Conv7	96	3×3	1
UpSamp8	–	8×8	1
Conv9 (fullyConv)	–	1×1	1




Fig. 14: Predictions of the FCN with 3 Convolution-Pooling stages.

Best Performing FCN for Detecting the Knee Joint Centres. Before settling on the final architecture, experiments were done by varying the number of convolution stages, the number of filters and kernel sizes in each convolution layer. The best performing FCN (Table 9) was selected based on a high detection accuracy on the test data. This network was trained with the OAI dataset containing 4,444 knee radiographs. The dataset was split into a training set containing 3,333 knee images and test set containing 1,113 knee images. The validation set (10%) was taken from the training set. The effective aperture size of this FCN (Table 9) for a node in the last convolutional layer (before up-sampling) is 66. The aperture size for the previous networks shown in Table 7 is 42 and Table 6 is 34. For the other tested configurations the effective aperture size is even lower (less than 30).

Table 9 shows the configuration of the best performing FCN for detecting the knee joint centres. This FCN is based on a lightweight architecture and the network parameters (in total 214,177) are trained from scratch. The network consists of 4 stages of convolutions with a max-pooling layer after each convolutional stage, and the final stage of convolutions is followed by an up-sampling and a fully-convolutional layer. The network uses a uniform $[3 \times 3]$ convolution and $[2 \times 2]$ max pooling. Each convolution layer is followed by a ReLU activation layer. After the final convolution layer, an $[8 \times 8]$ up-sampling is performed as the network uses 3 stages of $[2 \times 2]$ max pooling. The up-sampling is essential for an end-to-end learning by back propagation from

Table 8: Detection accuracy of FCN based on Jaccard Index.

Jaccard Index	JI \geq 0.25 JI \geq 0.5 JI \geq 0.75		
Detection Accuracy	98.5 %	96.7 %	39.6 %

Table 9: Best performing FCN for detecting the knee joint centres.

Layer	Kernel	Kernel Size	Strides
Conv1	32	3 \times 3	1
MaxPool1	–	2 \times 2	2
Conv2.1	32	3 \times 3	1
Conv2.2	32	3 \times 3	1
MaxPool2	–	2 \times 2	2
Conv3.1	64	3 \times 3	1
Conv3.2	64	3 \times 3	1
MaxPool3	–	2 \times 2	2
Conv4.1	96	3 \times 3	1
Conv4.1	96	3 \times 3	1
UpSamp5	–	8 \times 8	1
Conv5 (fullyConv)	–	1 \times 1	1

the pixel-wise loss and to obtain pixel-dense outputs [48]; when pooling layer(s) and strides more than one are used in the network. The final layer is a fully convolutional layer with a kernel size of $[1 \times 1]$ and uses a sigmoid activation for pixel-based classification. The input to the network is of size $[256 \times 256]$.

This network was trained to minimise the total binary cross entropy between the predicted pixels and the ground truth using *stochastic gradient descent* (SGD) with default parameters: learning rate = 0.01, decay = $1e^{-6}$, momentum = 0.9, and nesterov = True. This network was trained for 40 epochs with a batch size of 32. The validation (10%) data was taken from the training set. Figure 15 shows the learning curves when training this network and decrease in the validation and training losses.

Table 10 shows the results of the best performing FCN. This network achieved a detection accuracy of **97.1%**, in total 2,162 knee joints out of the 2,226 test samples detected with a Jaccard index 0.5. The Jaccard index mean is 0.76 and standard deviation is 0.12.

Error Analysis. The results of the best performing FCN (Table 10) show 99% detection accuracy for a Jaccard index ≥ 0.1 , in total 2,205 out of 2,226 knee joints are successfully detected. On observing the failed detections: 1% (in total 21 knee joints), there are two patterns.

1. The output of the FCN is very faint or no detections at all. Figure 16 shows two instances of input X-ray images, masks defining the knee joint centres as ground truth, and output of the best performing FCN with faint detections.




Fig. 15: Training and validation losses of the FCN.

Table 10: Detection accuracy of the best performing FCN.

Jaccard Index	JI ≥ 0.25	JI ≥ 0.5	JI ≥ 0.75
Detection Accuracy	98.9 %	97.1 %	43.3 %

The input images with variations in the local contrast and local luminance due to the imaging protocol variations appear to be the main cause for this error. Histogram equalisation is used as a pre-processing step to adjust the contrast of the input images. Even though this adjusts the contrast globally in an image, there are still contrast variations in portions of the image. Local contrast enhancement algorithms [55] or adaptive histogram equalisation [56] can be used to normalise the images for variations in the local contrast and local luminance.

2. The FCN output picks up noise along with the knee joints. Figure 17 shows two instances of input X-ray images, masks defining the knee joint centres as ground truth, and output of the best performing FCN with noise. The reason for this error appears to be due to the variations in the imaging protocol and resolution of the X-ray images, and presence of artefacts in the input X-ray images. Intuitively, the FCN uses horizontal edge detection along with other features to detect the knee joints. The artefacts with predominant horizontal edges are picked up by the FCN along with the centre of knee joints. When simple contour detection is applied on the FCN output, instead of the knee joints the artefacts are also detected.

Automatically Extracting the Knee Joints. After training FCNs to automatically detect the centre of the knee joints, the next step is to extract the ROI i.e. the knee joints with reference to the detected centres. The initial goal is to train an end-to-end network for localising the knee joints i.e. to




Fig.16: Error analysis: X-ray images, ground truth, FCN output - weak detections

directly predict the bounding box co-ordinates of the knee joints from the input X-ray images. A bounding box regression is investigated [57] that is a network trained on top of the FCN (Table 9) output, to achieve this. First, CNNs are trained with the masks (20×20) of knee joint centres as the input (256×256) and the bounding box coordinates of the left knee joint (x_1, y_1) and right knee joint (x_2, y_2) as the ground truth (labels). Next, CNNs are trained with the X-ray images as input and the targets (labels) are the bounding box coordinates instead of the binary masks. However in both the experiments, the networks trained to predict the bounding boxes give low accuracy. On considering the overall knee joint centres, there is no large variations in the centre coordinates. The reason for the low accuracy is that the networks are not learning discernible features to predict the bounding box coordinates. This affects the overall performance of the localisation. Therefore, a simple approach based on contour detection is used to calculate the centres and extract the knee joints. Figure 18 shows an X-ray image with the centres, the left and the right knee joints extracted from the X-ray image using the centroids. The steps involved in this method are as follows.

- First, the contour detection [54] is used on the FCN output to calculate the spatial coordinates of the knee joint centres. In the contour detection method, first the input images (FCN output) are converted to binary by applying Otsu's threshold. Next, the contours from the binary image are automatically detected and recorded. Finally, the centroids are calculated from the detected knee joint regions.
- The knee OA radiographs are resized to 2560×2560 , that is 10 times the size of the FCN output 256×256 .
- The detected knee joint centres are up-scaled to a factor of 10.




Fig. 17: Error analysis: X-ray images, ground truth, FCN output - detections with noise.




Fig. 18: A knee X-ray image with the detected centres and the extracted left and right knees.

- Fixed size regions (640×560) are extracted around the up-scaled centres as the knee joint regions. After testing and visualising different sizes for the knee joint crop, image patch with the size (640×560) is found to be mostly suitable and containing the required ROI for further quantification. Figure 18 shows an instance of the extracted left and right knee joints.

Localisation Results. The results of the FCN are compared to the previous methods: template matching and SVM-based method to automatically detect the centre of the knee joints. All these methods are evaluated based on the Jaccard index (JI). Table 11 shows the detection accuracy of the knee joint centres using FCN, SVM-based method, and template matching. The results show that the proposed method using FCN clearly outperforms the previous

Table 11: Comparison of methods used for localising the centre of the knee joints

Method	JI > 0	JI \geq 0.5	JI \geq 0.75	Mean	Std. Dev.
Template Matching	54.4%	8.3%	3.1%	0.1	0.2
SVM-based Method	81.8%	38.6%	10.2%	0.36	0.31
Fully ConvNet	98.9%	97.1%	43.3%	0.76	0.12

methods. This also demonstrates that feature learning using an FCN is a better approach for detecting the knee joints than using hand-crafted features such as Sobel gradients and the template matching method that is sensitive to intensity level variations. However, the extracted knee joints from this method have some limitations.

Limitations of this Method. In all three approaches; FCN-based, SVM-based and template matching, the centre of the knee joints are detected and these are used as reference for automatically localising the knee joints. There are some limitations in extracting a fixed size region as the ROI with reference to the detected centres due to the variations in the resolution of the X-ray images and the variations in the size of the knee joints.



Fig. 19: Anomalies in the automatic extraction of the ROI.




Fig. 20: The actual ROI for the knee joints in Figure 19.

All the images are resized to a fixed size $2,560 \times 2,560$ and extract a fixed size region 640×560 around the detected centres as the ROI. Due to this scaling

issue, portions of the knee joints are omitted in the automatic extraction of the ROI. Figure 19 shows such instances. Figure 20 shows the corresponding actual ROIs. Due to the varying sizes of the knee joints and a fixed size region being extracted as the ROI, there are differences in the aspect ratio of the extracted and the actual ROI. Figure 21 shows instances where the knee joints are small in comparison to the fixed size region extracted as the ROI. Figure 22 shows the actual ROIs.

The classification of the automatically extracted knee joints is compared to the manually extracted knee joints. There is a decrease in the accuracy by a margin of 3–4% when using the automatically extracted knee joints with reference to the detected centres. The discrepancies in the localisation of knee joints affects the overall classification of the knee OA images. To overcome these limitations, as the next approach FCNs are trained to detect the ROI itself, instead of detecting the knee joint centres.



Fig. 21: Variations in the aspect ratio of the extracted knee joints.




Fig. 22: The actual ROI for the extracted knee joints in Figure 21.

4.2.2 Localising the Region of Interest

The previous methods to localise the knee joints in the X-ray images with reference to the automatically detected centres have certain limitations. To overcome these limitations and to improve the localisation, FCNs are trained to detect the ROI directly [58]. Figure 23 shows the steps involved in this method.

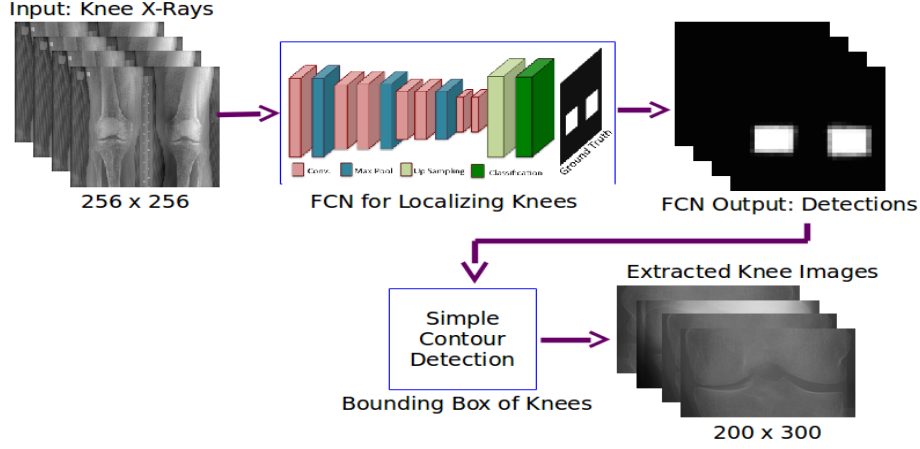


Fig. 23: Automatic localisation of the Region of Interest.

Dataset and Ground Truth. For the experiments in this approach, a new dataset from the MOST is used along with the data from the previous experiments, the baseline cohort of the OAI dataset. In total 4,446 X-ray images are selected from the OAI dataset and 2,920 X-ray images from the MOST dataset based on the availability of KL grades for both knee joints. The full ROI is manually annotated in all these X-ray images, after downscaling to 10% of the actual size. The down-sampling of the images is necessary to reduce the computational costs. Binary masks are generated based on the manual annotations. Figure 24 shows an instance of an input X-ray image and the binary mask annotations corresponding to the ROI. The image patches from the masked region (the knee joints) are taken as positive training samples and the patches from rest of the image are taken as the negative training samples to train a FCN. The datasets are split into a training/validation set (70%) and test set (30%). The training and test samples from the OAI dataset are 3,146 images and 1,300 images, and from the MOST dataset are 2,020 images and 900 images.

Training the FCN. First, a FCN is trained using the same architecture (Table 9) from the previous approach to detect the ROI. Initially, the network is trained with training samples from OAI dataset and test it with OAI and MOST datasets separately. Next, the training samples are increased by including the MOST training set where the test set is a combination of both OAI and MOST test sets. This network is trained to minimise the total binary cross entropy between the predicted pixels and the ground truth using the adaptive moment estimation (Adam) optimiser with default parameters: initial learning rate (α) = 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e^{-8}$. Adam optimiser gives faster convergence than standard SGD. Figure 25 shows the learning curves converging to small loss




Fig. 24: (a) An input X-ray image and (b) The binary mask annotations for the region of interest.




Fig. 25: Training and validation losses of the FCN.

when training this network. Figure 26 shows the output of this network for a test image.

A few other network configurations are tested by varying the number of convolutional-pooling stages, convolutional layers in each stage and the number of convolutional kernels in a convolutional layer. There was no further improvement in the detection accuracy on the validation set. Therefore, this configuration was settled as the final network for localising the knee joints.

Quantitative Evaluation. The Jaccard index, i.e. the intersection over Union (IoU) of the automatically detected and the annotated knee joint is used to quantitatively evaluate the automatic detections. For this evaluation, all the knee joints in both the OAI and MOST datasets are manually annotated using a fast annotation tool. Table 12 shows the number (percentage) of knee joint correctly detected based on the Jaccard index (JI) values greater than 0.25, 0.5 and 0.75 along with the mean and the standard deviation of JI. Table 12 also shows detection rates on the OAI and MOST test sets separately.




Fig. 26: An input X-ray image, ground truth and output prediction of the FCN.

Table 12: Comparison of automatic detection based on the Jaccard Index (JI).

Test Data	JI > 0	JI \geq 0.5	JI \geq 0.75	Mean	Std. Dev.
OAI	100%	100%	88%	0.82	0.06
MOST	99.7%	98.8%	80.6%	0.80	0.09
Combined OAI-MOST	100%	100%	92.2%	0.83	0.06

Considering the anatomical variations of the knee joints and the imaging protocol variations, the automatic detection with a FCN is highly accurate with 100% detection accuracy for $\text{JI} \geq 0.5$ and 92.2% (4,056 out of 4,400) of the knee joints for $\text{JI} \geq 0.75$ being correctly detected.

Qualitative Evaluation. Figures 27, 28, and 29 show a few instances of successful knee joint detections with the JI values for the left and right knee detections. Detecting the ROI directly gives high accuracy (100%) in comparison to the previous method (Section 4.2) to detect the knee joint centres and extracting a fixed size region as the ROI. The FCN in this method learns features from a relatively larger region (the actual ROI) in comparison to the previous method where the FCN is confined to learn features from a small region (20×20), the centre of the knee joints, and therefore, the detections are more accurate.

Error Analysis. This method is highly accurate with 100% detection accuracy for a $\text{JI} \geq 0.5$. Nevertheless, there are a few anomalies in the FCN detections due to variations in the imaging protocols, presence of artefacts and noise in the input images. Figures 30 and 31 show two instances where one knee has undergone joint-arthoplasty and the knee implants are visible in the X-ray images, and due to this the FCN detections are distorted. Figures 32, 33 and 34 show a few instances of X-ray images with noise and presence of artefacts due to imaging protocols. This adversely affects the FCN detections.

Extracting the Knee Joints. The bounding boxes of the knee joints are calculated using simple contour detection from the output predictions of the FCN. After converting the FCN output to binary image using Otsu's




Fig. 27: Qualitative Evaluation: An input X-ray image, ground truth, and FCN detections: left knee with $JI=0.98$, right knee with $JI=0.888$.




Fig. 28: Qualitative Evaluation: An input X-ray image, ground truth, and FCN detections: left knee with $JI=0.879$, right knee with $JI=0.969$.

threshold, the contours are detected using simple image analysis by calculating the zero order moments [54], which gives the perimeter of the detected object. The contours are recorded as bounding boxes. The knee joints are extracted from knee OA radiographs using the bounding boxes. The bounding boxes are up-scaled from the output of the FCN that is of size $[256 \times 256]$ to the original size of each knee OA radiograph, before extracting the knee joints so that the aspect ratio of the knee joints is preserved.

4.3 Summary and Discussion

Automatically localising the knee joints in X-ray images is an important and an essential step before quantifying knee OA severity. Previously, template matching was implemented as a baseline method to localise the knee joints, proposed by Shamir et al. [1, 3], and it was shown that the detection accuracy is low ($\sim 30\%$) in this method for large datasets like OAI. To improve the localisation, a SVM-based method with Sobel horizontal image gradients as features was proposed in this Section. This method showed a large improvement in detection accuracy (82%) but still falls short of perfect localisation. The anomalies in localised knee joints can affect the step involving classification of the localised knee joints to quantify knee OA severity.

Instead of using hand-crafted features, a deep learning-based solution was proposed in this Section to further improve localisation. FCNs were trained to automatically detect and extract the knee joints. All three methods: template




Fig. 29: Qualitative Evaluation: An input X-ray image, ground truth, and FCN detections: left knee with $JI=0.768$, right knee with $JI=0.984$.




Fig. 30: Error Analysis: An input X-ray image, ground truth, and FCN detections: left knee with $JI=0.83$, right knee with $JI=0.398$. The implants in the right knee is the reason for this localisation error.

matching, SVM-based and FCN-based were evaluated using a common metric: the Jaccard Index. This method achieved almost perfect detection with 100% accuracy for a Jaccard Index 0.5 and an accuracy of 92% for a Jaccard index greater than equal to 0.75. The author believes this performance is sufficient to localise and extract the knee images for classification. As such further improvements are left as future work. The localisation performance may be improved by including additional pre-processing steps to remove the artefacts and noise in the images, and to normalise the local contrast variations in the images. Using additional data for learning and data augmentation may improve the localisation performance.

5 Automatic Assessment of Knee OA Severity

Previous work on automated assessment of knee OA severity approached it as an image classification problem [2, 3, 40–42]. Previous methods have tested many hand-crafted features based on pixel statistics, textures, edge and object statistics, and transforms [3, 4, 6, 7, 40, 41]. Many classifiers such as the SVM [41], the k-nearest neighbour classifier [40], the weighted neighbour nearest classifier [3, 4], the random forest classifiers [2], and even artificial neural networks (ANN) [42, 43] have been tested for knee image classification. As a baseline (in Section 4.1), the state-of-the-art features successful in computer vision tasks, such as histogram of oriented gradients [47], local binary patterns




Fig. 31: Error Analysis: An input X-ray image, ground truth, and FCN detections: left knee with $JI=0.473$, right knee with $JI=0.837$. The implants in the left knee is the reason for this localisation error.




Fig. 32: Error Analysis: An input X-ray image, ground truth, and FCN detections: left knee with $JI=0.887$, right knee with $JI=0.356$. The noise in the right knee causes this localisation error.

[59], and Sobel Gradients [46] are tested. These features are not included in the previous studies to assess knee OA severity. All the previous approaches based on hand-crafted features give low multi-class classification accuracy when classifying knee images, and in particular classifying fine-grained successive knee OA grades remains a challenge. As a baseline, the state-of-the-art CNNs features (in Section 5.2.2) are also tested for knee images classification on a small baseline data set from OAI and this approach gave promising results. Motivated by this, the use of CNNs are investigated for quantifying knee OA severity.

5.1 Baseline for Classifying Knee OA Radiographs

5.1.1 WNDCHRM Classification

WNDCHRM is an open source utility for biological image analysis and medical image classification [3, 4, 60]. In WNDCHRM, a generic set of image features based on pixel statistics (multi-scale histograms, first four moments), textures (Haralick and Tamura features), factors from polynomial decomposition (Zernike polynomials), and transforms (Radon, Chebyshev statistics, Chebyshev-Fourier statistics) are extracted. For feature selection, every feature is assigned a *Fisher*




Fig. 33: Error Analysis: An input X-ray image, ground truth, and FCN detections: left knee with $JI=0.681$, right knee with $JI=0.488$. The localisation error in this image is due to the variation in the imaging protocol.




Fig. 34: Error Analysis: An input X-ray image, ground truth, and FCN detections: left knee with $JI=0.768$, right knee with $JI=0.507$. The variations in the local contrast and luminance affects the localisations.

$score^7$ and 85% of the features with lowest *Fisher scores* are rejected and the remaining 15% of the features are used for classification [3].

Experiments. The dataset used for the initial experiments to classify knee OA images using WNDCHRM are taken from the baseline data sample of 200 progression and incidence cohort. After histogram equalisation and mean normalisation of the X-ray images, the knee joints are extracted manually from the radiographs. The extracted knee joints are split into training (70%) and test (30%) sets. The WNDCHRM command line program is used to classify the extracted knee joint images. WNDCHRM uses a variant of k-nearest neighbour classifier.

Results and Discussion. The baseline dataset is not balanced and there are only 44 samples available in KL grade 4. Given the limited number of images in this class, only a small number of images are used for training and testing (35 images for training and 9 images for testing) for multi-class classification. For other classifications 100 images are used for training and 30 images for testing.

⁷ Fisher score is one of the widely used methods for determining the most relevant features for classification

Table 13: Results of WNDCHRM Classification.

Classification	Grades	Accuracy
Binary	G0 vs G1	66.7 %
	G1 vs G2	48.3 %
	G2 vs G3	60 %
	G3 vs G4	55 %
	G0 vs G2	48.3 %
	G0 vs G3	70 %
Multi-class	G0 to G4	28.3 %
	G0 to G3	35.8 %

It is evident from the results (Table 13) that the multi-class classification accuracy and successive grades classification accuracies are very low. The reason for low classification accuracy is that the features used for classification are not capable of capturing the minute structural and morphological variations in the knee joints between the successive grades. Next, the state-of-the-art hand-crafted features are investigated in an attempt to improve the classification accuracy.

5.1.2 Classification using Hand-crafted Features

Histogram of oriented gradients (HOG), local binary patterns, and Sobel Gradients are tested for classifying knee OA images [3, 1, 5]. These features are not used in the previous studies. HOG describes the local object shape and appearance within an image by the distribution of intensity gradients or edge directions and the HOG descriptor was successful in human detection [47]. LBP is powerful for image texture classification. LBP uses local spatial patterns and grey scale contrast as measures for texture classification [59].

Experiments with HOG, LBP and Sobel descriptors. Once again the images from the baseline data sample of 200 progression and incidence cohort is used. The HOG, LBP and Sobel descriptors are extracted from the knee joint images and a SVM is used for classification. Table 14 shows the classification results of successive grades of knee OA images using a SVM and the feature space included the HOG, LBP and Sobel gradients.

There is no large improvement in the classification accuracies using HOG, LBP and Sobel gradients features with SVM classification from the previous results with the WNDCHRM classification. To improve the classification, the features space is expanded by including highly effective and top-ranked features from the WNDCHRM classification.

Expanding the feature space. The features based on pixel statistics and textures such as Tamura, Haralick, Gabor and Zernike are used for

Table 14: Classification results of the proposed methods using hand-crafted features.

Grades	WNDCHRM	SVM classification with hand-crafted features			
		HOG	LBP	Sobel	Combining all
G0 vs G1	66.7 %	53.3 %	58.3 %	58.3 %	55 %
G1 vs G2	48.3 %	48.3 %	53.3 %	58.3 %	51.6 %
G2 vs G3	60 %	60 %	60 %	56.7 %	63.3 %
G3 vs G4	55 %	65 %	65 %	50 %	65 %

Table 15: Classification results of WNDCHRM and the proposed methods using hand-crafted features.

Grades	WNDCHRM	Proposed Methods		
		kNN	SVM	SVR
G0 vs G1	66.7%	55%	60%	60%
G1 vs G2	48.3%	61.7 %	46.7%	48.3%
G2 vs G3	60%	51.7%	55%	60%
G3 vs G4	55%	35%	50%	45%
G0 vs G2	48.3%	46.7%	55%	56.7%
G0 vs G3	70%	48.3%	58.3%	60%

classification. These features are used in the WNDCHRM classification. Tamura texture features represent contrast, coarseness and directionality of an image [61]. Haralick features are the statistics computed on the co-occurrence matrix of an image [62]. Gabor textures are based on Gabor wavelets and the image descriptors are computed using Gabor transform of an image [63]. Zernike features are obtained by the Zernike polynomial approximation of an image [64]. The feature space for classification is formed by simple concatenation of all the extracted features into a super vector following the early fusion approach.

Results and Discussion. First, a SVM is used with the extracted features for classifying knee OA images. Next, a k-nearest neighbour classifier and support vector regression (SVR) are tested for classification. In total, 100 knee joint images are taken for training and 30 for test set in each grade. Table 15 shows the classification accuracy of the WNDCHRM classifier and the classification using kNN, SVM, and SVR.

When comparing the classification results of the proposed methods (SVM, kNN, and SVR) to the WNDCHRM classification, for some cases the results are slightly better and promising. Nevertheless, there is a need for a more significant improvement in the classification results. In these experiments, a subset of features from WNDCHRM, such as Tamura & Haralick texture

features, Gabor wavelet features, and Zernike features, were extracted and used for classification. In addition to these features HOG, LBP, and Sobel Gradients were tested. It was found that by further expanding the feature space by including features from WNDCHRM based on transforms such as Radon, Chebyshev, FFT, and Wavelet, and compound image transforms such as Chebyshev-FFT, Chebyshev-Wavelet, and Wavelet-FFT classification can be improved. However, the author believes that learning feature representations can be more effective for fine-grained knee OA classification. In the following section, the state-of-the-art CNN features are investigated for classifying knee OA images.

5.2 Automatic Quantification using Convolutional Neural Networks

First, the use of off-the-shelf CNNs are investigated for quantifying knee OA severity through classification and regression. Two approaches are followed for this: 1) using a pre-trained CNN for fixed feature extraction, and 2) fine-tuning pre-trained CNN following a transfer learning approach. WNDCHRM, an open source utility for medical image classification[4, 3, 60] is used for benchmarking the classification results obtained from the proposed methods.

Next, three new methods are investigated to automatically quantify knee OA: 1) training a CNN from scratch for multi-class classification of knee OA images; 2) training a CNN to optimise a weighted ratio of two loss functions categorical cross-entropy for multi-class classification and mean-squared error for regression; and 3) training a CNN for ordinal regression of knee OA images. The results from these methods are compared to the previous methods. The classification results using both manual and automatic localisation of knee joints are also compared.

5.2.1 Off-the-shelf CNNs

The use of well-known off-the-shelf CNNs such as the VGG-16 network [53], and comparatively simpler networks like VGG-M-128 network [65], and BVLC reference CaffeNet [66, 67] (which is very similar to the widely-used *AlexNet* model [14]) are investigated to classify knee OA images. These networks are pre-trained for general image classification using a very large dataset: the ImageNet LSVRC dataset [68] which contains more than 1.2 million images in 1000 classes. Initially, features are extracted from the convolutional, pooling, and fully-connected layers of VGG-16, VGG-M-128, and BVLC CaffeNet, and used to train linear SVMs to classify knee OA images.

The pre-trained networks are fine-tuned for knee OA images classification motivated by the transfer learning approach [69]. Transfer learning is adopted as the OAI dataset is small, containing only a few thousand images. In transfer learning, a base network is first trained on external data, and then the weights of the initial n layers are transferred to a target network [69]. The new layers of the target network are randomly initialised following the Xavier weight initialisation procedure [70]. The random weights initialisations increase the likelihood of the training algorithms during the backpropagation to obtain a

global solution through the gradient descent instead of settling to a nearest local solution.

Intuitively, the lower layers of the networks contain more generic features such as edge or texture detectors useful for multiple tasks, whilst the upper layers progressively focus on more task specific cues [67, 69]. This approach is used for both classification and regression, adding new fully-connected layers, and backpropagation is used to fine-tune the weights for the complete network on the target loss.

5.2.2 Classification using CNN Feature Extraction

The VGG-16 network [53] is trained with the OAI dataset. Features are extracted from different layers of the VGG net such as fully-connected (fc7), pooling (pool5), and convolutional (conv5_2) layers to identify the most discriminating set of features. Linear SVMs (LIBLINEAR [71]) are trained with the extracted CNN features for classifying knee OA images, where the ground truth are images labelled with KL grades. Next, the use of simple pre-trained CNNs such as VGG-M-128 [65] and the BVLC CaffeNet model [66] are investigated for classifying the knee OA images. These networks have fewer layers and parameters in comparison to the VGG-16 network. The features are extracted from the fully-connected, pooling, and convolutional layers, using the VGG-M-128 net and the BVLC reference CaffeNet.

Experiments and Results. The knee joint images are split into training ($\sim 70\%$) and test ($\sim 30\%$) set based on the distribution of each KL grade. Features are extracted from fully-connected, pooling, and convolution layers of VGG-16, VGG-M-128, and BVLC CaffeNet. Linear SVMs are trained individually for binary and multi-class classifications on the extracted features. WNDCHRM is used for benchmarking the classification results from the proposed methods in this chapter [3, 4, 60]. WNDCHRM is trained with the same training data so that the classification results from WNDCHRM and CNN features can be compared. The knee OA images are classified in three ways as follows. Classifying healthy knee images (grade 0) with the progressive stages (grade 1, 2, 3, and 4), classifying the images belonging to the successive stages (grade 0 vs 1, grade 1 vs 2, ...) and multi-class classification to classify all the stages of knee OA images.

Table 16 shows the test set classification accuracies achieved by WNDCHRM and the CNN features. The CNN features consistently outperform WNDCHRM for classifying healthy knee samples against the progressive stages of knee OA. The features from conv4 layer with dimension $512 \times 13 \times 13$ and pool5 layer $256 \times 13 \times 13$ of VGG-M-128 net, and conv5 layer with dimension $512 \times 6 \times 6$ and pool5 layer with dimension $256 \times 6 \times 6$ of BVLC reference CaffeNet give higher classification accuracy in comparison to the fully-connected fc6 and fc7 layers of VGG nets and CaffeNet. Intuitively, the lower layers capture more discriminative low-level features such as edge or shape detectors, and the higher layers tend to contain high-level features

Table 16: Classification accuracy (%) achieved by the WNDCHRM and pre-trained CNN features.

Category	Classification	WNDCHRM	VGG-16 Net				VGG-M-128 Net				BVLG ref	CaffeNet
			fc7	pool5	conv5_2	fc6	pool5	conv4	fc7	pool5		
Progressive	Grade 0 vs Grade 1	51.5	56.3	61.3	63.5	56.5	63.2	64.7	62.0	64.3	63.3	
	Grade 0 vs Grade 2	62.6	68.6	74.3	76.7	67.8	75.5	77.6	69.6	73.6	73.9	
	Grade 0 vs Grade 3	70.6	86.4	91.4	92.4	88.5	90.2	92.9	87.9	92.5	91.5	
	Grade 0 vs Grade 4	82.8	98.1	98.6	99.3	98.8	99.3	99.2	98.5	99.4	99.1	
Successive	Grade 1 vs Grade 2	48.8	60.0	64.7	67.3	57.9	63.5	65.3	61.2	65.8	62.8	
	Grade 2 vs Grade 3	54.5	69.8	76.4	77.0	73.0	77.3	79.0	70.3	78.1	77.1	
	Grade 3 vs Grade 4	58.6	85.2	88.8	90.0	85.0	90.4	91.2	87.4	91.6	91.4	
Multi-class	Grade 0 to Grade 2	39.9	51.1	53.4	56.9	51.1	55.0	57.4	51.1	54.8	54.4	
	Grade 0 to Grade 3	32.0	44.6	48.7	53.9	45.4	50.2	53.3	46.9	51.6	50.2	
	Grade 0 to Grade 4	28.9	42.6	47.6	53.1	43.8	49.5	53.4	44.1	50.8	50.0	

specific to object classes as per the training data. Features are also extracted from lower layers such as pool4, conv4.2, pool3, pool2 and train classifiers on top of these features. As the dimension of the bottom layers are high, the training time is increased, however, no improvement in classification accuracy is observed.

In a fine-grained classification task such as knee OA image classification, the accuracy of classifying successive classes tends to be low, as the variations in the progressive stages of the disease are minimal, and only highly discriminant features can capture these variations. From the experimental results, as shown in Table 16, the features extracted from CNNs provide significantly higher classification accuracy in comparison to the WNDCHRM, and these features are effective and promising for classifying the consecutive stages of knee OA.

Multi-class classifications are performed using linear SVMs with the CNN features (Table 16, multi-class). Again, the CNN features outperform WNDCHRM. The classification accuracies obtained using convolutional (conv4, conv5) and pooling (pool5) layers are slightly higher in comparison to fully-connected layer features. There are minimal variations in classification accuracy obtained with the features extracted from VGG-M-128 net and BVLC reference CaffeNet in comparison to VGG-16.

5.2.3 Transfer Learning

As a next approach [45], the BVLC CaffeNet [66] and VGG-M-128 [65] networks are fine-tuned using transfer learning to classify knee images. These two smaller networks are chosen because they contain fewer layers and parameters ($\sim 62\text{M}$), over the much deeper VGG-16, which has $\sim 138\text{M}$ parameters. The top fully-connected layer of both networks is replaced and the model is retrained on the OAI dataset using backpropagation. The lower-level features in the bottom layers are also updated during fine-tuning. Standard softmax loss is used as the objective for classification, and accuracy layers are added to monitor the training progress. A Euclidean loss layer (mean squared error) is used for the regression experiments.

Experiments and Results. Table 17 shows the multi-class classification results for the fine-tuned BVLC CaffeNet. The VGG-16 network is omitted in these experiment since the variation in accuracy among the pre-trained CNNs is small, and fine-tuning VGG-16 is more computationally expensive.

The dataset is split into training (60%), validation (10%) and test (30%) sets for fine-tuning. The right-left flipped knee joint images are included in the training set to increase the number of training samples. The networks are fine-tuned for 20 epochs using a learning rate of 0.001 for the transferred layers, and 0.01 for the newly introduced layers. The performance of fine-tuned BVLC CaffeNet is slightly better than VGG-M-128. Hence, the results of fine-tuning BVLC CaffeNet is only shown here. Figure 35 shows the learning curves for training and validation loss, and validation accuracy. The decrease in loss and




Fig.35: Learning curves:training and validation losses (left), and validation accuracy (right) during fine-tuning.

Table 17: Classification accuracy (%) achieved with the features extracted from fine-tuned BVLC Net.

Classification	Before Fine-Tuning			After Fine-Tuning		
	fc7	pool5	conv5	fc7	pool5	conv5
grade 0 vs grade 1	62.0	64.3	63.3	63.3	64.3	61.9
grade 0 vs grade 2	69.6	73.6	73.9	76.3	77.2	74.1
grade 0 vs grade 3	87.9	92.5	91.5	96.7	96.0	96.3
grade 0 vs grade 4	98.5	99.4	99.1	99.8	99.7	99.7
grade 1 vs grade 2	61.2	65.8	62.8	63.3	66.7	62.7
grade 2 vs grade 3	70.3	78.1	77.1	85.8	83.9	83.3
grade 3 vs grade 4	87.4	91.6	91.4	94.4	93.6	92.6
grade 0 to grade 2	51.1	54.8	54.4	57.4	57.0	52.0
grade 0 to grade 3	46.9	51.6	50.2	57.2	56.5	51.8
grade 0 to grade 4	44.1	50.8	50.0	57.6	56.2	51.8

increase in accuracy shows that the fine-tuning is effective and makes the CNN features more discriminative, which improves classification accuracy (Table 16). The features extracted from the fully connected (fc7) layer provide slightly better classification in comparison to pooling (pool5) and convolution (conv5) layers.

Regression using Fine-tuned CNNs. Existing work on automatic assessment of knee OA severity treats it as an image classification problem, assigning each KL grade to a distinct category [1–3, 5]. To date, evaluation of automatic KL grading algorithms has been based on binary and multi-class classification accuracy with respect to these discrete KL grades [6, 3, 5]. Nevertheless, KL grades are not categorical, but rather represent an ordinal scale of increasing severity. Treating them as categorical during evaluation means that the penalty for incorrectly predicting that a subject with grade 0

OA has grade 4 is the same as the penalty for predicting that the same subject has grade 1 OA. Clearly the former represents a more serious error, yet this is not captured by evaluation measures that treat grades as categorical variables [45]. In this set up, permuting the ordering of the grades has no effect on classification performance. Moreover, the quantisation of the KL grades to discrete integer levels is essentially an artefact of convenience; the true progression of the disease in nature is continuous, not discrete.

The author proposes that it is more appropriate to measure the performance of an automatic knee OA severity assessment system using a continuous evaluation metric like mean squared error. Such a metric appropriately penalises errors in proportion to their distance from the ground truth, rather than treating all errors equally. Directly optimising mean squared error on a training set also naturally leads to the formulation of knee OA assessment as a standard regression problem. Treating it as such provides the model with more information on the structure and relationship between training examples with successive KL grades. It is demonstrated that the use of regression reduces both the mean squared error and improves the multi-class classification accuracy of the model [45].

The pre-trained BVLC CaffeNet model is fine-tuned using both classification loss (cross entropy on softmax outputs) and regression loss (mean squared error) to compare their performance in assessing knee OA severity. In both cases, the fully connected layer fc7 is replaced with a randomly initialised layer and fine-tuned for 20 epochs, selecting the model with the highest validation performance. The classification network uses a 5D fully connected layer and softmax following the fc7 layer, and the regression network uses a 1D fully connected node with a linear activation.

The models are compared using both mean squared error (MSE) and standard multi-class classification metrics. The mean squared error is calculated using the standard formula:

$$MSE = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2,$$

where n is the number of test samples, y_i is the true (integer) label and \hat{y}_i is the predicted label. For the classification network the predicted labels \hat{y}_i are integers and for the regression network they are real numbers. A configuration is tested, where the real outputs are rounded from the regression network to produce integer labels. Table 18 shows the MSE for classification using the WNDCHRM and the CNN trained with classification loss (CNN-Clsf), regression loss (CNN-Reg), and regression loss with rounding (CNN-Reg*). Regression loss clearly achieves significantly lower mean squared error than both the CNN classification network and the WNDCHRM features.

To demonstrate that the regression loss also produces better classification accuracy, the classification accuracy from the network trained with classification loss and the network trained with regression loss and rounded labels are compared. Rounding in this case is necessary to allow the use of

Table 18: MSE for classification and regression.

Classes	WNDCHRM	CNN-Clsf	CNN-Reg	CNN-Reg*
grade 0 to 4	2.459	0.836	0.504	0.576

Table 19: Comparison of classification performance using classification (left) and regression (right) losses.

Classification	Classification Loss			Regression Loss		
	Precision	Recall	F_1	Precision	Recall	F_1
0	0.53	0.64	0.58	0.57	0.92	0.71
1	0.25	0.19	0.22	0.32	0.14	0.20
2	0.44	0.32	0.37	0.71	0.46	0.56
3	0.37	0.47	0.41	0.78	0.73	0.76
4	0.56	0.54	0.55	0.89	0.73	0.80
Mean	0.43	0.44	0.43	0.61	0.62	0.59

standard classification metrics. Table 19 compares the resulting precision, recall, and F_1 scores. The multi-class (grade 0–4) classification accuracy of the network fine-tuned with regression loss is 59.6%. The network trained using regression loss clearly gives superior classification performance. The author suspects this is due to the fact that using regression loss gives the network more information about the ordinal relationship between the KL grades, allowing it to converge on parameters that better generalise to unseen data.

Discussion. The initial approach to quantify knee OA severity used features extracted from pre-trained CNNs. Three pre-trained networks are investigated and it is found that the BVLC reference CaffeNet and VGG-M-128 networks perform best. A linear SVM trained on features from these networks achieved significantly higher classification accuracy (53.4%) in comparison to the previous state-of-the-art (28.9%). The features from pooling and convolutional layers were found to be more accurate than the fully connected layers. Fine-tuning the networks by replacing the top fully connected layer gave further improvements in multi-class classification accuracy.

Previous studies have assessed their algorithms using binary and multi-class classification metrics. The author proposes that it is more suitable to treat KL grades as a continuous variable and assess accuracy using mean squared error. This approach allows the model to be trained using regression loss so that errors are penalised in proportion to their severity, producing more accurate predictions. This approach also has the nice property that the predictions can fall between grades, which aligns with continuous disease progression.

In summary, this section presented two approaches based on the existing pre-trained CNNs for quantifying knee OA severity: first, the CNNs were used for fixed feature extraction and next, the CNNs were fine-tuned using transfer

learning. Both the approaches outperformed the previous state-of-the-art, the WNDCHRM classifier, giving promising results. As a next logical step, CNNs are trained from scratch to investigate if this leads to further improvement in quantifying knee OA severity.

5.2.4 Training CNNs from Scratch

Training a CNN from scratch (or full training) is challenging and complicated, because it requires a large amount of annotated training data. The learning curves during training should ensure proper convergence to generalise well avoiding overfitting [15]. An alternative to full training is transfer learning, fine-tuning CNNs pre-trained in other domain (for instance ImageNet dataset with natural images) to a target domain, for instance medical domain. However, the knowledge transfer may be limited by the substantial differences between the source and the target domains, which may mitigate the performance of the fine tuned CNNs. Nevertheless, with sufficient labelled training data and carefully selected hyper-parameters, fully trained CNNs can outperform fine-tuned CNNs and hand-crafted alternatives [15, 12].

Fully trained CNNs have been found to be highly successful in many medical applications [15, 12]. Some of the applications that use fully trained CNNs for musculo-skeletal (including knee) image analysis are knee cartilage segmentation using multi-stream CNNs [16], total knee arthroplasty kinematics by real-time 2D/3D registration using CNN regressors [72], automated skeletal bone age assessment in X-ray images using deep learning [73], and posterior-element fractures detection on spine CT using deep convolutional networks [74]. Motivated by these approaches, CNNs are trained from scratch to quantify knee OA severity using both classification and regression.

Dataset and Preprocessing. The data used for the initial experiments are taken from the baseline OAI dataset. There are 4,446 X-ray images with the KL grade annotations in this dataset. The MOST dataset is included for later experiments and this dataset consists of 2,920 X-ray images with KL grade annotations. Two set of knee joint images are used separately for the experiments: 1) extracted after automatic localisation and 2) extracted after manual annotation of the ROI. This is to compare the quantification performance of the CNNs trained with knee joints from automatic localisation and manual annotation. As a preprocessing step, all the knee joint images are subjected to histogram equalisation for intensity level normalisation. The images were resized to 256×256 pixels for the initial experiments. Later, the input image size is changed to 200×300 . This size is chosen to approximately preserve the aspect ratio based on the mean aspect ratio (1.6) of all the extracted knee joints. Right-left flip of the knee joint images are used to generate more training data.

Initial Configuration. A CNN is configured with a lightweight architecture with 4 layers of learned weights: 3 convolutional layers and 1 fully connected

Table 20: Initial CNN configuration.

Layer	Kernels	Kernel Size	Strides	Output shape
conv1	32	11×11	2	$32 \times 128 \times 128$
maxPool1	–	3×3	3	$32 \times 42 \times 42$
conv2	96	7×7	1	$96 \times 42 \times 42$
maxPool2	–	3×3	3	$96 \times 14 \times 14$
conv3	128	3×3	1	$128 \times 14 \times 14$
maxPool3	–	3×3	2	$128 \times 4 \times 4$
fc4	–	–	–	2048
fc5	–	–	–	5

layer. As the training data set is relatively small, a lightweight architecture is considered with minimal (4.5 million) parameters in comparison to the existing CNNs. Table 20 shows the CNN configuration in detail. Each convolutional layer is followed by batch normalisation and a ReLU activation layer. A max pooling layer is included after each convolution stage. The final pooling layer is followed by a fully connected layer (fc4), and a softmax dense layer (fc5) with an output shape 5 for the multi-class classification of (0–4) ordinal KL grades. A drop out layer with a drop out ratio of 0.5 is included after the fully connected layer (fc5) to avoid overfitting. The input images are of size 256×256 pixels and fed to the network after sub-sampling by a factor of 2. So, the input size is 128×128 pixels.

Training Process and Initial Results. The network parameters are trained from scratch with the knee joint images as training samples and the KL grades (0, 1, 2, 3 or 4) as labels. To start, the knee joint images extracted manually from the radiographs of the OAI dataset are used. The dataset is split into training (70%) and test (30%) sets. The validation (10%) data is taken from the training set. The network is trained to minimise categorical cross entropy for multi-class classification. *Stochastic gradient descent* (SGD) is used with default parameters: decay = $1e^{-6}$, momentum = 0.9, and nesterov = True and the initial learning rate is set to 0.0001. The networks are trained with fixed learning rate in the initial experiments. The Adam optimiser with default parameters: initial learning rate (α) = 0.001, β_1 = 0.9, β_2 = 0.999, ϵ = $1e^{-8}$ is tested, instead of SGD for the later experiments. The benefits of the Adam optimiser are that it uses adaptive learning rates and provides faster convergence.

This network achieves a multi-class classification accuracy of 44.7% on the test data. The mean-squared error is 1.75. Table 21 shows the classification results: precision, recall, and F_1 score of the initial configuration. The results show that the classification performance is low and the mean-squared error is high. These are initial results and the hyper-parameters of this network are tuned to improve the classification performance. Further, the number of convolutional layers, convolutional-pooling stages, the number of convolutional kernels, kernel sizes and other parameters are experimented.

Table 21: Classification results of the initial CNN configuration.

grade	Precision	Recall	F_1 Score
0	0.45	0.92	0.60
1	0.24	0.07	0.11
2	0.49	0.18	0.26
3	0.50	0.39	0.44
4	1.00	0.01	0.02
Mean	0.45	0.45	0.37

The terms ‘parameters’ and ‘hyper-parameters’ in machine learning are often used interchangeably, but there is a difference between them. Parameters are learned by a classifier or a machine learning model from the training data, for instance weights or coefficients of the independent variables. Hyper-parameters are the settings used to optimise the performance of a classifier or a model and they are not fit based on the training data. The hyper-parameters for a CNN include the number and size of the hidden layers, learning rate and its decay, drop out regularisation, gradient clipping threshold and other settings.

Tuning Hyper-parameters. After the initial CNN configuration giving low classification accuracy (44.7%), as a first step the depth of the network is increased. A convolutional layer and a pooling layer are included. This increases the number of layers with learned weights to 5 layers: 4 convolutional layers and 1 fully connected layer. SGD with default parameters: decay = $1e^{-6}$, momentum = 0.9, and nesterov = True, is used for training this network. Learning rates from 0.0001 to 0.01 with an incremental increase by a factor 10 are tested, and the learning rate 0.001 is found to be the best. After experimenting with the convolutional kernel size, the number of kernels in the convolutional layer, the number of outputs of the fully connected layer and other parameters, the final architecture in this configuration is obtained. Table 22 shows the CNN architecture in detail.

After 20 epochs of training, this network gave a multi-class classification accuracy of 55.2% with a mean-squared error 0.803 on the validation data. After 35 epochs the network achieves the best results for this configuration with a classification accuracy of 60.4% and mean-squared error 0.838. Table 23 shows the classification results: precision, recall, and F_1 score of this network. There is an improvement in the overall classification results in comparison to the previous results (Table 21). Figure 36 shows the learning curves with increase in the training and validation accuracies, and decrease in the training and validation losses whilst training this network. It can be observed from the learning curves (Figure 36), after 32 epochs there is an increase in validation loss with decrease in training loss and also there is no further increase in validation accuracy whilst training accuracy increases: the network is starting to overfit. A drop out regularisation by a ratio of 0.5 is included after the fully

Table 22: CNN architecture (CNN-1) after tuning hyper-parameters.

Layer	Kernels	Kernel Size	Strides	Output shape
conv1	32	11×11	2	$32 \times 128 \times 128$
maxPool1	–	3×3	2	$32 \times 63 \times 63$
conv2	96	5×5	1	$64 \times 63 \times 63$
maxPool2	–	3×3	2	$64 \times 31 \times 31$
conv3	128	3×3	1	$128 \times 31 \times 31$
maxPool3	–	3×3	2	$128 \times 15 \times 15$
conv4	256	3×3	1	$256 \times 15 \times 15$
maxPool4	–	3×3	2	$256 \times 7 \times 7$
fc5	–	–	–	1024
fc6	–	–	–	5

Table 23: Classification results after tuning hyper-parameters.

	grade	Precision	Recall	F_1 Score
0	0.57	0.90	0.70	
1	0.31	0.11	0.16	
2	0.64	0.45	0.53	
3	0.74	0.77	0.76	
4	0.86	0.72	0.78	
Mean	0.58	0.60	0.57	

connected layer (fc5) to mitigate overfitting. Also, data augmentation is used to increase the training samples by including the right-left flip of the knee joints and this doubles the number of training samples. Drop out regularisation after convolutional layers and fully connected layers, and l2-norm weight regularisations are used to further mitigate overfitting in the next set of experiments.

Next, the depth of the network is further increased by increasing the number of layers with learned weights, continuing the experimentation with the other associated hyper-parameters. Up to 5 convolutional-pooling stages followed by two fully connected layers are tested. The classification accuracy with 4 convolutional-pooling stages is 60.8% and with 5 convolutional-pooling stages is 61%.

Previous networks use a single convolutional layer followed by a pooling layer. Next, cascaded convolutional layers are used in a convolution-pooling stage like VGG-16 model. Each convolutional layer is followed by a ReLU activation. Figure 24 shows the CNN architecture that gives the best results in this approach. This network gives a classification accuracy of 60.1% with a mean-squared error 0.838.

Inspired by the success of VGG networks [53], a network with cascaded convolutional layers of uniform (3×3) kernel size and (2×2) max pooling with stride 2 is trained, and the hyper-parameters are tuned. This network gives a




Fig. 36: Learning curves: training and validation losses, and accuracies of the fully trained CNN

Table 24: CNN architecture (CNN-2) after tuning hyper-parameters.

Layer	Kernels	Kernel Size	Strides	Output shape
conv1	32	11×11	2	$32 \times 128 \times 128$
maxPool1	—	3×3	2	$32 \times 63 \times 63$
conv2	64	5×5	1	$64 \times 63 \times 63$
maxPool2	—	3×3	2	$64 \times 31 \times 31$
conv3-1	64	3×3	1	$64 \times 31 \times 31$
conv3-2	64	3×3	1	$64 \times 31 \times 31$
maxPool3	—	3×3	2	$64 \times 15 \times 15$
conv4-1	96	3×3	1	$96 \times 15 \times 15$
conv4-2	96	3×3	1	$96 \times 15 \times 15$
maxPool4	—	3×3	2	$96 \times 7 \times 7$
fc5	—	—	—	1024
fc6	—	—	—	5

classification accuracy of 57.5% with a mean-squared error 0.961. There is no further improvement in the classification results in comparison to the previous results.

Training Off-the-shelf CNNs from Scratch. Earlier, the widely used off-the-shelf CNNs such as BVLC reference CaffeNet [66, 67] (which is very similar to the AlexNet model [14]), VGG-M-128 network [65], and VGG-16 network [53] were fine-tuned for knee images classification. The pre-trained VGG-16 network has ~ 138 million free parameters, and the other networks, Alexnet with ~ 62 million and the VGG-M-128 with ~ 26 million parameters, are relatively simple. Training these networks, in particular VGG-16, from scratch is computationally very expensive due to the depth and the number of free parameters. Previously

Table 25: AlexNet architecture.

Layer	Kernels	Kernel Size	Strides	Output shape
conv1	96	11×11	4	$96 \times 64 \times 64$
maxPool1	–	3×3	2	$96 \times 31 \times 31$
conv2	256	5×5	1	$256 \times 31 \times 31$
maxPool2	–	3×3	2	$256 \times 15 \times 15$
conv3	384	3×3	1	$384 \times 15 \times 15$
conv4	384	3×3	1	$384 \times 15 \times 15$
conv5	256	3×3	1	$256 \times 15 \times 15$
maxPool5	–	3×3	2	$256 \times 7 \times 7$
fc6	–	–	–	4096
fc7	–	–	–	4096
fc8	–	–	–	5

trained CNNs have relatively fewer parameters (~ 4 to 6 million) to suit the relatively small dataset with a few thousand of training examples.

Next, CNNs are fully trained using the AlexNet and the VGG-M-128 architectures. This is to compare the classification performance of these networks to the previously trained networks from scratch. Table 25 shows the AlexNet architecture in detail. The convolutional layers conv1 and conv2 in this network are followed by Relu and batch normalisation layers. The two fully connected layers (fc6) and (fc7) are followed by a drop out regularisation by a ratio 0.5. This network was pre-trained for 1,000 classes in the ImageNet [75] dataset. The output of the last fully connected layer (fc8) is replaced with a 5 output dense layer for multi-class knee OA image classification. This network is trained using SGD with default parameters. Learning rates from 0.00001 to 0.01 with an incremental increase by a factor 10 are tested. The learning rate set at 0.001 gives the best results.

The fully trained AlexNet gives a classification accuracy of 57.2% with a mean-squared error 0.741. Table 26 shows the classification results; precision, recall, and F_1 score of the fully trained AlexNet model. These results show that the classification accuracy achieved by the fully trained AlexNet is low (57.2%) in comparison to the accuracy (60.8%) achieved by previous networks. Moreover, this network is overfitting. This is evident from the learning curves (Figure 37) obtained whilst training this network. After 30 epochs, the learning curves show an increase in validation loss whilst the training loss is decreasing and there is no improvement in the validation accuracy whilst the training accuracy keeps increasing. The reason for overfitting is the number of training samples in the dataset ($\sim 10,000$) is very low in comparison to the number of free parameters (~ 62 million) in AlexNet. This model was originally developed and trained on datasets like ImageNet [75] that consists of more than ~ 1.2 million images. There are two fully connected layers with 4,096 outputs in the AlexNet and these layers contribute to more than 95% of the total free parameters in this network. Next,

Table 26: Classification results of the fully trained AlexNet.

grade	Precision	Recall	F_1 Score
0	0.65	0.61	0.63
1	0.29	0.36	0.32
2	0.59	0.55	0.57
3	0.75	0.73	0.74
4	0.77	0.79	0.78
Mean	0.59	0.57	0.58




Fig.37: Learning curves: training and validation losses, and accuracies of the fully trained AlexNet.

a relatively simple architecture (VGG-M-128) is investigated for the knee OA images classification.

The VGG-M-128 network is a simplified model of the AlexNet [53]. The last fully connected layer (fc7) of AlexNet has 4,096 outputs. The number of fc7 outputs is reduced to 128 in VGG-M-128. This reduces the number of free parameters and this network contains (~ 26 million) parameters in total. The AlexNet configuration is retained in the VGG-M-128 network with a few changes in the architecture. The kernel size of the first convolutional layer is reduced to (7×7) and the stride is reduced to 2. The number of filters is fixed to 512 in the conv3, conv4, and conv5 layers. Table 27 shows the architecture details. This network parameters are trained from scratch using SGD with default parameters: decay = $1e^{-6}$, momentum = 0.9, and nesterov = True. The learning rate is fixed to 0.001 after testing different rates like before.

This network gives a classification accuracy of 56.3% and the mean-squared error is 0.685. Table 28 shows the classification results of this network. The results show a slightly lower classification accuracy (56.3%) in comparison to the previous results. There is no significant difference in the precision, recall, and F_1 score of this network in comparison to the AlexNet classification results (Table 26). This network is also overfitting like the AlexNet. This is evident from the

Table 27: VGG-M-128 architecture.

Layer	Kernels	Kernel Size	Strides	Output shape
conv1	96	7×7	2	$96 \times 128 \times 128$
maxPool1	–	3×3	2	$96 \times 63 \times 63$
conv2	256	5×5	1	$256 \times 32 \times 32$
maxPool2	–	3×3	2	$256 \times 15 \times 15$
conv3	512	3×3	1	$512 \times 15 \times 15$
conv4	512	3×3	1	$512 \times 15 \times 15$
conv5	512	3×3	2	$512 \times 8 \times 8$
maxPool5	–	3×3	2	$512 \times 3 \times 3$
fc6	–	–	–	4096
fc7	–	–	–	128
fc8	–	–	–	5

Table 28: Classification results of the fully trained VGG-M-128.

	grade	Precision	Recall	F_1 Score
0	0.66	0.65	0.66	
1	0.27	0.42	0.33	
2	0.62	0.46	0.53	
3	0.77	0.69	0.72	
4	0.87	0.73	0.79	
Mean	0.60	0.56	0.58	

learning curves (Figure 38) of this network. The learning curves show increase in the validation loss after 30 epochs and the validation accuracy remains almost the same. The drop out regularisations after the fully connected layers fc6 and fc7 are not able to fully mitigate overfitting. The reason for overfitting remains the same as for AlexNet. The number of training samples is very low even for the number of free parameters in this network (~ 26 million).

Best Performing CNN for Classification. After experimenting with different configurations, the network in Table 29 is found to be the best for classifying knee images. This network is similar to the previous configuration (Table 24), but with slight variations. The network contains five layers of learned weights: four convolutional layers and a fully connected layer. The total number of free parameters in the network is ~ 5.4 million. Each convolutional layer in the network is followed by batch normalisation and a ReLU activation layer. After each convolutional stage there is a max pooling layer. The final pooling layer (maxPool4) is followed by a fully connected layer (fc5) and a softmax dense (fc6) layer. To avoid overfitting, a drop out layer with a drop out ratio of 0.25 is included after the last convolutional (conv4) layer and a drop out layer with a drop out ratio of 0.5 after the fully connected layer (fc5).




Fig.38: Learning curves: training and validation losses, and accuracies of the fully trained VGG-M-128 network.

Table 29: Best performing CNN for classifying the knee images.

Layer	Kernels	Kernel Size	Strides	Output shape
conv1	32	11×11	2	$32 \times 100 \times 150$
maxPool1	—	3×3	2	$32 \times 49 \times 74$
conv2	64	5×5	1	$64 \times 49 \times 74$
maxPool2	—	3×3	2	$64 \times 24 \times 36$
conv3	96	3×3	1	$96 \times 24 \times 36$
maxPool3	—	3×3	2	$96 \times 11 \times 17$
conv4	128	3×3	1	$128 \times 11 \times 17$
maxPool4	—	3×3	2	$128 \times 5 \times 8$
fc5	—	—	—	1024
fc6	—	—	—	5

Also, a L2-norm weight regularisation penalty of 0.01 is applied in the last two convolutional layers (conv3 and conv4) and the fully connected layer (fc5). Applying a regularisation penalty to other layers increases the training time whilst not introducing significant variation in the learning curves. The network is trained to minimise categorical cross-entropy loss using the Adam optimiser with default parameters: initial learning rate (α) = 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e^{-8}$. The inputs to the network are knee images of size 200×300 . This size is selected to approximately preserve the aspect ratio based on the mean aspect ratio (1.6) of all the extracted knee joints.

First, this network is trained using the OAI dataset like the previous network trainings. This network achieves a classification accuracy of 61% with a mean-squared error 0.861. Next, training samples are included from the MOST dataset. This network achieves a classification accuracy of 61.8% with a mean-squared error 0.735 for the combined OAI-MOST dataset. There is a slight increase in the classification accuracy (0.8%) and decrease in the mean-squared error (0.126).

Table 30: Classification results of the best performing fully trained CNN.

grade	Precision	Recall	F_1 Score
0	0.65	0.83	0.73
1	0.30	0.10	0.14
2	0.51	0.60	0.55
3	0.77	0.69	0.73
4	0.87	0.70	0.78
Mean	0.59	0.62	0.59




Fig.39: Learning curves: training and validation losses, and accuracies of the best performing fully trained CNN.

Table 30 shows the classification results: precision, recall, and F_1 score of this network for the combined OAI-MOST dataset. Figure 39 shows the learning curves whilst training this network. The learning curves show proper convergence of the training and validation losses with consistent increase in the training and validation accuracies till they reach constant values.

To sum up, a high classification accuracy (61%) is achieved with the CNN (Table 29) trained from scratch and outperform the VGG-M-128 and the AlexNet trained from scratch. The fully trained AlexNet gives a classification accuracy of 57.2% and VGG-M-128 gives an accuracy of 56.3%. The classification results of the methods proposed in this section and the previous state-of-the-art are compared in the next section.

Classification Results. The classification results of the fully trained network is compared to WNDCHARM, the multi purpose medical image classifier [5, 4, 60] that gave the previous best results for automatically classifying knee OA X-ray images, and to previous results (Table 17) on fine-tuning BVLC reference CaffeNet for this task (Section 5.2.3). WNDCHARM is trained with the data taken from the OAI and MOST datasets.

Table 31: Classification results of the proposed methods and the existing methods.

Method	Test Data	Accuracy	Mean-Squared Error
Wndchrm	OAI & MOST	34.8%	2.112
Fine-Tuned BVLC CaffeNet	OAI	57.6%	0.836
Fully trained CNN	OAI	61%	0.861
Fully trained CNN	OAI & MOST	61.8%	0.735

Table 31 shows the multi-class classification accuracy and mean-squared error of the fine-tuned BVLC CaffeNet, the network trained from scratch and WND-CHARM for the OAI and MOST datasets. The results show that the network trained from scratch for classifying knee OA images clearly outperforms WNDCHARM. This shows learning feature representations using CNNs for fine-grained knee OA images classification is highly effective and a better approach in comparison to using a combination of hand-crafted features in WNDCHARM. The other reason for low classification accuracy of WNDCHARM is that it uses only a balanced dataset for training. Both the OAI and MOST datasets are very unbalanced and in particular the number of knee images available in KL grade 4 is very small, $\sim 5\%$ in total.

Moreover, these results show an improvement over previous methods that used fine-tuned off-the-shelf networks such as VGG-M-128 and the BVLC Reference CaffeNet for classifying knee OA X-ray images through transfer learning. These improvements are due to the lightweight architecture of the network trained from scratch with less (~ 5.4 million) free parameters in comparison to 62 million free parameters of BVLC CaffeNet for the small amount of training data available. The off-the-shelf networks are trained using a large dataset like ImageNet containing millions of images, whereas the dataset used in this experiment contains much fewer ($\sim 10,000$) training samples. Furthermore, the results show an increase in classification accuracy from 61% to 61.8% when the MOST dataset is included in the training set. This result is promising and it shows that with more training data the CNN performance maybe further improved. Next, the use of regression by fully trained CNNs is investigated to improve the quantification performance.

Training CNNs for Regression. CNNs are trained from scratch to classify knee images in the previous approach. The outcomes are ordinal KL grades (0, 1, 2, 3 or 4) that quantify knee OA severity. CNNs are trained for regression in the next approach. This is to assess knee OA severity in a continuous scale (0–4). The author argued earlier (Section 5.2.2) that it is more appropriate to assess knee OA in a continuous scale as knee OA is progressive in nature, not discrete [45]. The existing CNNs are fine-tuned to quantify knee OA severity using regression.

Table 32: Results of the initial network trained for regression after rounding the predicted continuous grades.

Grade	Precision	Recall	F_1 Score
0	0.78	0.18	0.29
1	0.24	0.83	0.37
2	0.49	0.32	0.39
3	0.63	0.42	0.50
4	0.57	0.20	0.30
Mean	0.57	0.37	0.36

Initial Configuration. A CNN is trained for regression using almost the same architecture (Table 29) that gave the highest multi-class classification accuracy previously. The last fully connected layer (fc6) with softmax activation and an output shape 5 for multi-class classification is replaced with a linear activation with an output shape of 1 for regression. The CNN is trained to minimise mean-squared error using the Adam optimiser with default parameters: initial learning rate (α) = 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e^{-8}$. Like before, the inputs to the network are images of size 200×300 . The data for training is taken from both the OAI and the MOST datasets. Both these datasets contain discrete KL grade (0, 1, 2, 3 or 4) annotations for the knee joints. These labels are used in the previous approach to train classifiers. However, there is no ground truth for KL grades on a continuous scale for either of these datasets to train a network directly for regression output. Hence, the discrete KL grades are used as labels to train CNNs for regression.

Initial Results. This CNN gives a mean-squared error of 0.654 on the test data after training. In comparison to the mean-squared error achieved by the classifier (0.898) with almost the same architecture, there is definitely an improvement in the quantification using regression. The performance metrics: accuracy, precision, recall, and F_1 score are computed for the regression results by rounding the predicted continuous grade to the next integer value. Rounding, in this case, is necessary to allow the use of standard classification metrics and to compare the performances of classification and regression. Table 32 shows the precision, recall, and F_1 score for regression. In comparing these results to the previous classification results (Table 30), there is a decrease in precision, recall, and F_1 score. The classification accuracy achieved by regression is 36.9% with a mean-squared error 0.75. From these results it is evident that the regression performance is low in this initial configuration. Next, the hyper parameters of this network are tuned to improve the regression performance.

Tuning the Hyper-parameters. The experiment is continued by varying the number of layers with learned weights in the architecture, the number of

Table 33: Best performing CNN for regression of the knee images.

Layer	Kernels	Kernel Size	Strides	Output shape
conv1	32	11×11	2	$32 \times 100 \times 158$
maxPool1	–	3×3	2	$32 \times 49 \times 74$
conv2	64	5×5	1	$64 \times 49 \times 74$
maxPool2	–	3×3	2	$64 \times 24 \times 36$
conv3-1	64	3×3	1	$64 \times 24 \times 36$
conv3-2	64	3×3	1	$64 \times 24 \times 36$
maxPool3	–	3×3	2	$64 \times 11 \times 17$
conv4-1	128	3×3	1	$96 \times 11 \times 17$
conv4-2	128	3×3	1	$96 \times 11 \times 17$
maxPool4	–	3×3	2	$96 \times 5 \times 8$
fc5	–	–	–	1024
fc6	–	–	–	1

convolutional-pooling stages, the number of kernels and kernel sizes in the convolutional layers and the regularisations to avoid overfitting. The architecture in Table 33 is found to be the best for quantifying knee OA severity using regression. This network contains seven layers of learned weights: six convolutional layers and a fully connected layer. This network has ~ 5.6 million free parameters in total. Each convolutional layer is followed by batch normalisation and a ReLU layer. The last pooling layer (maxPool4) is followed by two dense layers: fc5 with ReLU and fc6 with linear activations. A drop out layer with a drop out ratio 0.5 is added after fc6. The network is trained to minimise the mean-squared error using the Adam optimiser with default parameters: initial learning rate (α) = 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e^{-8}$. The network is trained with knee images taken from the OAI and the MOST datasets. Figure 40 shows the learning curves: training and validation losses whilst training this network. The learning curves show convergence in the losses.

Comparison of Classification and Regression Results. The best performing CNN for regression gives a mean-squared error of 0.574. After rounding the continuous grade predictions, this network achieves a multi-class classification accuracy of 54.7% and the mean-squared error is 0.661. Table 34 shows the accuracy and mean-squared error for the fully trained CNN for classification and regression. The results show that the multi-class classification accuracy calculated after rounding the output is low for CNN-regression. The main reason for this likely is training the regression network with ordinal labels instead of continuous labels. There is also a decrease in accuracy due to the rounding of regression output and the rounding is necessary to compute standard classification metrics. On the other hand, the mean-squared error of the fully trained CNN for regression is low in both the cases before rounding (0.574) and after rounding (0.661) in comparison to the fully trained CNN for




Fig. 40: Learning curves: training and validation losses for the best performing CNN for regression.

Table 34: Comparison of classification and regression results.

Method	Accuracy	MSE (before rounding)	MSE (after rounding)
CNN-Classification	61.8%	0.735	–
CNN-Regression	54.7%	0.574	0.661

regression. Table 35 shows the precision, recall, and F_1 score of the rounded regression output and the classification output. These results show that the network trained with classification loss outperforms the regression loss. The reason for this is again likely the lack of continuous KL grade ground truth to train a CNN directly for regression output.

To sum up, training a CNN from scratch for regression output gives low mean-squared error. The lack of ground truth affects the performance of the regression. To overcome this drawback, in the next approach multi-objective convolutional learning is investigated to quantify knee OA severity.

5.2.5 Multi-objective Convolutional Learning

In general, assessing knee OA severity is based on the multi-class classification of knee images and assigning KL grade to each distinct category [3–6]. The author previously argued that assigning a continuous grade (0–4) to knee images through regression is a better approach for quantifying knee OA severity as the disease is progressive in nature. However, there is no ground truth i.e. KL grades on a continuous scale to train a network directly for regression output. Therefore, the networks are trained using multi-objective convolutional learning [76–80] to optimise a weighted-ratio of two loss functions: categorical cross-entropy and mean-squared error. Mean squared error gives the network information about the ordering of grades, and cross entropy gives information about the quantisation of grades. Intuitively, optimising a network with two loss functions provides

Table 35: Comparison of the regression and classification performances.

Grade	Regression			Classification		
	Precision	Recall	F_1	Precision	Recall	F_1
0	0.70	0.71	0.70	0.65	0.83	0.73
1	0.29	0.42	0.34	0.30	0.10	0.14
2	0.52	0.39	0.45	0.51	0.60	0.55
3	0.67	0.51	0.58	0.77	0.69	0.73
4	0.58	0.55	0.57	0.87	0.70	0.78
Mean	0.57	0.55	0.55	0.59	0.62	0.59

a stronger error signal and it is a step to improve the overall quantification, considering both classification and regression results.

Initial Configuration. The same architecture of the best performing CNN is used for classification (Table 29) as an initial configuration to jointly train a CNN for classification and regression outputs. Table 36 and Figure 41 shows the configuration details of the initial configuration. The network has five layers with learned weights: four convolutional layers and a fully connected layer. The total free parameters in the network are ~ 5.4 million. The last fully connected layer (fc5) is followed by two dense layers with softmax and linear activations for simultaneous multi-class classification and regression outputs. Drop out layers with a drop out ratio 0.25 are included after the conv4 layer and a drop out ratio 0.5 after the fc6 layer to avoid overfitting. In addition to this, a L2-norm weight regularisation penalty of 0.01 is applied in conv3, conv4 and fc6 layers to avoid overfitting. Applying a regularisation penalty to other layers did not introduce significant variations in the learning curves. Unlike the previous approaches, this network is trained to minimise a weighted ratio of two loss functions: categorical cross-entropy and mean-squared error. After testing different values from 0.2 to 0.6 for the weight of regression loss, a ratio of 0.5 is fixed, as this ratio gives optimal results.

The input to the network are knee images of size 200×300 . The knee images taken from the combined OAI-MOST dataset is used for training this network. The same train (70%) and test (30%) split are maintained from the previous experiments to make valid comparisons of the quantification results from the different methods. The right-left flip of the knee images is included to increase the training data. A validation split of 20% from the training data is used. This network is trained using the Adam optimiser with default parameters: initial learning rate (α) = 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e^{-8}$, as it gives faster convergence in comparison to the standard SGD.

Initial Results. Figure 42 shows the learning curves obtained whilst jointly training the CNN for classification and regression outputs. The learning curves show convergence in the validation and training losses with improvement in

Table 36: Initial configuration to jointly train a CNN for classification and regression outputs.

Layer	Kernels	Kernel Size	Strides	Output shape
conv1	32	11×11	2	$32 \times 100 \times 150$
maxPool1	—	3×3	2	$32 \times 49 \times 74$
conv2	64	5×5	1	$64 \times 49 \times 74$
maxPool2	—	3×3	2	$64 \times 24 \times 36$
conv3	96	3×3	1	$128 \times 24 \times 36$
maxPool3	—	3×3	2	$128 \times 11 \times 17$
conv4	128	3×3	1	$256 \times 11 \times 17$
maxPool4	—	3×3	2	$256 \times 5 \times 8$
fc5	—	—	—	1024
fc6-Clf	—	—	—	5
fc6-Reg	—	—	—	1

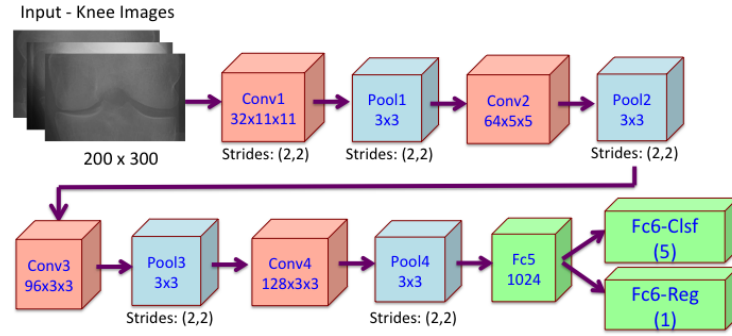




Fig.41: Initial configuration to jointly train a CNN for classification and regression outputs.

validation and classification accuracies. The jointly trained CNN with the initial configuration gives a classification accuracy of 60.8% with mean-squared error 0.795 for the classification outputs and 0.652 for the regression outputs. These results do not show improvement from the previous results. Previously, the network with the same configuration gave a classification accuracy of 61.8% and a mean-squared error 0.735 (Table 30) when trained to minimise only the classification loss. The same configuration after training to minimise only with the regression loss gave a mean squared error of 0.654 (Table 32). This configuration is optimal to minimise classification loss as it gave the highest classification accuracy (61.8%). However, this configuration is not optimal for regression as it gives a high mean-squared error (0.654). Next, the number of layers with learned weights and other hyper-parameters in this configuration are varied to find a good architecture that will give improved results for both classification and regression outputs.



(a) Classification



(b) Regression

Fig. 42: Learning curves for (a) classification and (b) regression in jointly trained CNN.

Tuning Hyper-parameters. Cascaded convolutional stages are included in the next configuration in an attempt to improve the regression outputs. The CNNs with cascaded convolutional stages gave best results for regression (Table 33) in the previous approach. Table 37 shows the network details. This network contains six layers of learned weights: five convolutional layers and a fully connected layer. The total free parameters in this network are ~ 7.8 million. The other settings remain the same from the previous network and the same training procedure is followed.

This network gives a multi-class classification accuracy of 62.9% and the mean-squared error is 0.754 for the classification output and 0.583 for the regression output. These results show improvement in the quantification performance in comparison to the previous results. Tuning the hyper parameters improves both the classification and the regression outcomes. Next, the depth of the architecture is increased and other related hyper parameters

Table 37: Jointly trained network for classification and regression outputs.

Layer	Kernels	Kernel Size	Strides	Output shape
conv1	32	11×11	2	$32 \times 100 \times 150$
maxPool1	–	3×3	3	$32 \times 33 \times 50$
conv2-1	64	3×3	1	$64 \times 33 \times 50$
conv2-2	64	3×3	1	$64 \times 33 \times 50$
maxPool2	–	3×3	2	$64 \times 16 \times 24$
conv3-1	96	3×3	1	$96 \times 16 \times 24$
conv3-2	96	3×3	1	$96 \times 16 \times 24$
maxPool3	–	3×3	2	$96 \times 7 \times 11$
fc4	–	–	–	1024
fc5-Clsf	–	–	–	5
fc5-Reg	–	–	–	1

are tuned to investigate further improvement in the classification and regression outputs.

Best Performing Jointly Trained CNN. The best configuration (Table 38) is obtained after experimenting with different settings for jointly training a CNN for classification and regression outputs. This network has eight layers with learned weights: seven convolutional layers and a fully connected layer. This network has ~ 2.9 million free parameters in total. This is a lightweight architecture with minimal parameters in comparison to the previous networks and the existing off-the-shelf CNNs. Each convolutional layer is followed by batch normalisation and a ReLU activation layer. The fc5 layer is followed by two dense layers with softmax and linear activations for multi-class classification and regression outputs. To avoid overfitting, drop out with ratio 0.3 is included after the last fully connected (fc5) layer. Also, a L2 weight regularisation penalty of 0.01 is applied to all the convolutional and fully connected layers except the first two convolutional layers. This network is trained to minimise a weighted ratio of two loss functions: categorical cross-entropy and mean-squared error. This network is trained using the Adam optimiser with default parameters: initial learning rate (α) = 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e^{-8}$.

Jointly Trained CNN Results. Figure 43 shows the learning curves obtained whilst jointly training the CNN for classification and regression outputs. The learning curves show convergence to the minimum of the validation and training losses with improvement in validation and classification accuracies. The combined OAI-MOST dataset is used to compute these results. The same train-test split is maintained from the previous experiments. This jointly trained network gives a multi-class classification accuracy of 64.6% with a mean-squared error 0.685 for the classification outputs and 0.507 for the regression outputs. Table 39 shows the precision, recall, and F_1 score of this network. There is an improvement in the results: the classification accuracy

Table 38: Jointly trained network for classification and regression outputs.


Layer	Kernels	Kernel Size	Strides	Output shape
conv1	32	11×11	2	$32 \times 100 \times 150$
maxPool1	–	3×3	2	$32 \times 49 \times 74$
conv2-1	64	3×3	1	$64 \times 49 \times 74$
conv2-2	64	3×3	1	$64 \times 49 \times 74$
maxPool2	–	3×3	2	$64 \times 24 \times 36$
conv3-1	96	3×3	1	$96 \times 24 \times 36$
conv3-2	96	3×3	1	$96 \times 24 \times 36$
maxPool3	–	3×3	2	$96 \times 11 \times 17$
conv4-1	128	3×3	1	$128 \times 11 \times 17$
conv4-2	128	3×3	1	$128 \times 11 \times 17$
maxPool4	–	3×3	2	$128 \times 5 \times 8$
fc5	–	–	–	512
fc6-Clsf	–	–	–	5
fc6-Reg	–	–	–	1

Table 39: Results of the best performing jointly trained CNN for classification and regression outputs.


	Grade	Precision	Recall	F_1 Score
	0	0.68	0.85	0.75
	1	0.34	0.07	0.12
	2	0.53	0.63	0.57
	3	0.74	0.77	0.75
	4	0.86	0.81	0.84
Mean	0.62	0.65	0.60	

increases to 64.6% from the initial configuration (60.8%), the mean-squared error for regression decreases to 0.507 from the initial configuration (0.652). Increasing the depth of the architecture by including more layers with learned weights to the initial configuration and tuning the other hyper-parameters improves both the classification and regression results. Intuitively, providing a stronger error signal using both the classification and regression loss to the network allow to fit more parameters.

Results Comparison. The results of the jointly trained CNN are compared to the previous CNNs trained separately for classification and regression outputs. Table 40 shows the multi-class classification accuracy and mean-squared error of the jointly trained CNN and the separately trained CNNs for classification and regression outputs. There is an improvement in the classification accuracy and also the mean-squared error decreases for the joint training. These results show that the network jointly trained for classification and regression learns a



(a) Classification



(b) Regression

Fig. 43: Learning curves for (a) classification and (b) regression in jointly trained CNN.

better representation in comparison to the previous network trained separately for classification and regression outputs.

In summary, CNNs are trained from scratch to quantify knee OA severity using three approaches: classification, regression and jointly training for simultaneous classification and regression. From the results it is evident that the joint training outperforms both the individual training for classification and regression outputs. This supports the hypothesis that training a CNN for optimising a weighted ratio of two loss functions can improve the overall quantification of knee OA severity.

Error analysis. A confusion matrix and the area under curve (AUC) after plotting the receiver operating characteristics are computed to perform an error analysis on the classification of the knee images by the jointly trained CNN. From the classification metrics (Table 39), the confusion matrix (Figure 44), and the

Table 40: Comparison of results from jointly trained CNN and individually trained CNNs for classification and regression results.

Method	Clsf-Accuracy	Clsf-MSE	Reg-MSE
CNN-Classification	61.8%	0.735	—
CNN-Regression	54.7%	—	0.574
Jointly trained CNN	64.6%	0.685	0.507




Fig.44: Confusion matrix for the multi-class classification using the jointly trained CNN.

receiver operating characteristic (ROC) curves (Figure 45), it is evident that classification of successive grades is challenging, and in particular classification metrics for grade 1 have low values in comparison to the other grades.

Figure 46 shows some examples of misclassification: grade 1 knee joints predicted as grade 0, 2, and 3. Figure 47 shows the misclassification of knee joints categorised as grade 0, 2 and 3 predicted as grade 1. These images show minimal variations in terms of joint space width and osteophytes formation, making them challenging to distinguish. Even the more serious misclassification in Figure 48, for instance grade 0 predicted as grade 3 and vice versa, do not show very distinguishable variations. Furthermore, when the knee X-ray images belonging to grade 0 and grade 1 severity are examined, it can be seen that there are very subtle variations in terms of the joint space width and osteophytes formation. Even better representations are needed to capture these fine-grained variations and to distinguish coarse grades: grade 0 and grade 1 images.

Discussion. Jointly training a CNN from scratch using the multi-objective convolutional approach improves the multi-class classification accuracy and minimises the mean-squared error. However, successive grade classification still




Fig. 45: ROC for the multi-class classification using the jointly trained CNN.




Fig. 46: Mis-classifications: grade 1 joints predicted as grade 0, 2, and 3.

remains a challenge. Even though the KL grades are widely used for assessing knee OA severity in clinical settings, there has been continued investigation and criticism over the use of KL grades as the individual categories are not equidistant from each other [22, 81, 25, 24, 82]. This could be a reason for the low multi-class classification accuracy in the automatic quantification. Using OARSIS readings instead of KL grades could possibly provide better results for automatic quantification as the knee OA features such as joint space narrowing, osteophytes formation, and sclerosis are separately graded. Moreover, when the knee X-ray images belonging to grade 0 and grade 1 severity are visually examined, it can be seen that there are very subtle variations in terms of the joint space width and osteophytes formation. To capture these variations and distinguish these coarse grades, for instance grade 0 versus grade 1, even better representations are required. Indeed, it should also be recognised that even medical experts do not always agree upon a particular KL grade e.g. either 0 or 1 attributed to the initial stage of knee OA [22, 81, 25].




Fig. 47: Misclassification: other grade knee joints predicted as grade 1.




Fig. 48: An instance of more severe misclassification: grade 0 and grade 3.

5.2.6 Ordinal Regression

Ordinal regression⁸ is an intermediate task between multi-class classification and regression, sharing the properties of both. The outcomes or predictions in multi-class classification are discrete values and there is a meaningful order in the classes in regression. Ordinal regression is useful to classify patterns using a categorical scale which shows a natural order between the labels [83,84]. The misclassification from a normal classifier are treated the same, that is no misclassification are worse than others [85]. However, some mis-classifications on the extreme grades: grade 0 to grade 4 is treated worse than the others. This implies that the distances between the classes need to be taken into account when training a classifier. When quantifying the stages of a physical disease, it is preferable to predict the stage as ‘mild’ or ‘doubtful’ than ‘absent’ when the true label is ‘severe’. Ordinal regression models formalise this notion of order by ensuring that predictions farther from the true label incur a greater penalty than those closer to the true label [84]. The author believes that the KL grades prediction based on ordinal regression can further improve classification performance by reducing the margin of error (mean-squared error), considering the progressive nature of knee OA and the ground truth or labels for training a CNN i.e. the KL grades in an ordinal scale (0–4).

⁸ <https://statistics.laerd.com/spss-tutorials/ordinal-regression-using-spss-statistics.php>




Fig. 49: The CNN configuration for ordinal regression.

CNN Configuration for Ordinal Regression. For ordinal regression output, the last stage of the CNN (Table 38) that gave best results on the joint training for multi-class classification and regression is modified. The previous approach on the joint training used two dense layers with softmax and linear activations in parallel (Figure 41) for simultaneous multi-class classification and regression outputs. To train the CNN for ordinal regression, fixed weights ($[w_0, w_1, w_2, w_3, w_4] = [0, 1, 2, 3, 4]$) are applied to the outputs (probabilities) from the dense layer (Clsf) with softmax activations and back-propagate through a dense layer (Reg) with linear activations, optimising the mean-squared error loss function. The dense layer with softmax activations is treated as a hidden layer in this configuration. This is similar to the approach proposed by Beckham et al. [85] for ordinal classification. Figure 49 shows the CNN configuration for ordinal regression.

CNN Training. The CNN for ordinal regression (Table 41) is based on a lightweight architecture with ~ 2.9 million free parameters in total and it contains eight layers with learned weights: seven convolutional layers and a fully connected layer. Each convolutional layer is followed by batch normalisation and a ReLU activation layer. To avoid overfitting, drop out with ratio 0.3 is applied after the last fully connected (fc5) layer. In addition to this, a L2 weight regularisation penalty of 0.01 is applied to all the convolutional and fully connected layers except the first two convolutional layers. The fc5 layer is followed by two dense layers with softmax (fc6-Clsf) and linear activations (fc7-Reg). The output of the softmax (fc6-Clsf) layer is multiplied (dot product) with fixed weights ($[0, 1, 2, 3, 4]$) and given as input to the last dense layer (fc7-Reg). This network is trained to minimise two loss functions: categorical cross-entropy and mean-squared error with equal weights. The dense layers (fc7-Reg) and (fc6-Clsf) provides the ordinal regression and multi-class classification outputs. The network is trained for 80 epochs with a batch size 32, using the Adam optimiser with default parameters: initial

Table 41: CNN architecture for ordinal regression


Layer	Kernels	Kernel Size	Strides	Output shape
input	—	—	—	$1 \times 200 \times 300$
conv1	32	11×11	2	$32 \times 100 \times 150$
maxPool1	—	3×3	2	$32 \times 49 \times 74$
conv2-1	64	3×3	1	$64 \times 49 \times 74$
conv2-2	64	3×3	1	$64 \times 49 \times 74$
maxPool2	—	3×3	2	$64 \times 24 \times 36$
conv3-1	96	3×3	1	$96 \times 24 \times 36$
conv3-2	96	3×3	1	$96 \times 24 \times 36$
maxPool3	—	3×3	2	$96 \times 11 \times 17$
conv4-1	128	3×3	1	$128 \times 11 \times 17$
conv4-2	128	3×3	1	$128 \times 11 \times 17$
maxPool4	—	3×3	2	$128 \times 5 \times 8$
fc5	—	—	—	512
fc6-Clf	—	—	—	5
input-weights	—	—	—	5
merge-product	—	—	—	1
fc7-Reg	—	—	—	1

learning rate (α) = 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e^{-8}$. The same training, validation, and test data are used from the joint training to make valid comparison of the results.


The CNN configuration in Table 41 gives the best results for ordinal regression and this configuration is similar to the network for joint training (Table 38) except the arrangement of the last two dense layers. Tuning the hyper-parameters of this CNN by increasing the number of layers with learned weights does not improve the quantification performance. Therefore, this CNN configuration is selected as the final network for ordinal regression.

Results. The learning curves (Figure 50) obtained whilst training the CNN for ordinal regression shows convergence of the validation and training losses with improvement in validation and classification accuracies. Figure 51 shows the classification accuracy of the trained CNN model after every epoch on the test data for the ordinal regression and classification output. After 40 epochs of training, there is no significant improvement in the classification accuracies. There is a slight decrease in the accuracy of the ordinal regression in comparison to the classification. This is likely due to the rounding of the output.

After training, the CNN gives a classification accuracy of 64.3% on the test data. In the previous method on jointly training a CNN for classification and regression (Section 5.4.5), a multi-class classification accuracy of 64.6% (Table 39) is achieved. As the same CNN configuration is used except the last stage (Figure 49) and other settings are retained, the classification performance remains almost the same for the CNN trained for ordinal regression in comparison to the jointly trained CNN.



(a) Classification



(b) Ordinal regression

Fig. 50: Learning curves for (a) classification and (b) ordinal regression .

The classification metrics for the ordinal regression output are computed by rounding the predictions to integer values (0, 1, 2, 3, or 4). After rounding, the classification accuracy for the ordinal regression output is 61.8% with mean-squared error 0.504 on the test data. The classification metrics for the regression output from the jointly trained CNN gives a classification accuracy 53.3% with mean squared error 0.595 on the test data. There is an improvement in the classification accuracy and mean-squared error for ordinal regression in comparison to the previous regression results (see Table 40). Table 42 shows the precision, recall, and F_1 score for regression and ordinal regression. From the results it is evident that ordinal regression is out performing regression for quantifying knee OA images in a continuous scale.

Results Comparison. Four approaches are investigated to automatically quantify knee OA severity. Table 43 shows the multi-class classification accuracy and mean-squared error for the four approaches: 1) fine-tuning




Fig. 51: The CNN configuration for ordinal regression.

Table 42: Comparison of classification metrics from regression and ordinal regression.

Grades	Regression			Ordinal Regression		
	Precision	Recall	F_1	Precision	Recall	F_1
0	0.78	0.58	0.66	0.71	0.79	0.75
1	0.29	0.55	0.38	0.31	0.33	0.32
2	0.51	0.50	0.50	0.59	0.44	0.50
3	0.65	0.53	0.58	0.74	0.73	0.73
4	0.63	0.33	0.43	0.81	0.76	0.78
Mean	0.60	0.53	0.55	0.62	0.62	0.61

off-the-shelf CNN (BVLC CaffeNet), 2) training CNNs from scratch individually for classification and regression, 3) jointly training a CNN based on multi-objective convolutional learning for classification and regression, and 4) training a CNN for ordinal regression. These results are compared to WNDCHRM, that gave the previous best results for automatically classifying knee OA radiographs. The results show that jointly trained CNN gives the best results for multi-class classification. The ordinal regression outperforms all the other methods for quantifying knee OA images in a continuous scale with low mean-squared error.

5.3 An Automatic Knee OA Diagnostic System

An automatic knee OA diagnostic system is developed combining the automatic localisation pipeline and the quantification pipeline developed in the previous sections. Figure 52 shows the proposed end-to-end diagnostic pipeline to automatically quantify knee OA severity based on KL grades. The input X-ray images are subjected to histogram equalisation, mean normalisation and resized to a fixed size 256×256 . A fully convolutional network (FCN) is used to

Table 43: Comparison of classification, regression and ordinal regression results.

Method	Classification Accuracy	Mean-Squared Error
WNDCHRM	34.8%	2.112
Fine-Tuned BVLC CaffeNet	57.6%	0.836
CNN-Classification	61.8%	0.735
CNN-Regression	54.7%	0.574
Jointly trained CNN	64.6%	0.507
Ordinal Regression	64.3%	0.480

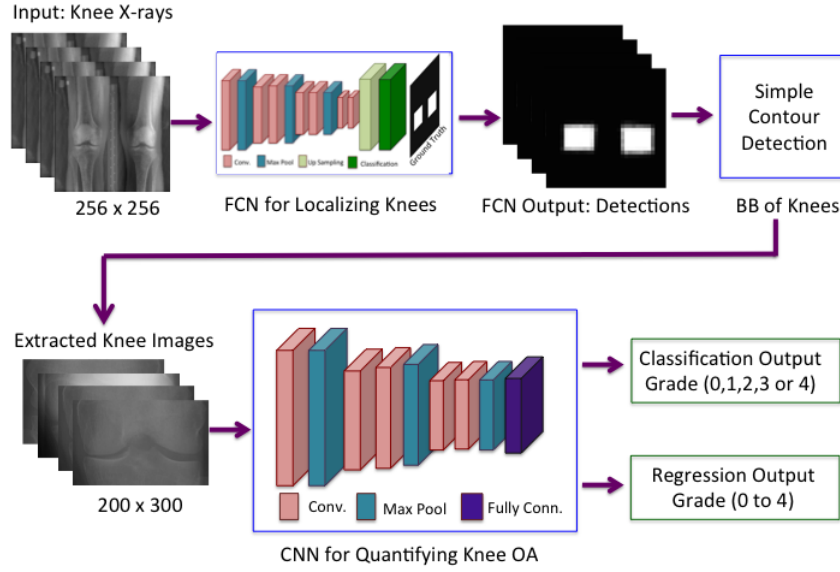


Fig. 52: The proposed pipeline for quantifying knee OA severity.

automatically detect the ROI, the knee joint regions. The bounding box coordinates of the ROI are calculated using simple contour detection. The knee joint regions are extracted from the knee radiographs using the bounding box coordinates. The localised and extracted knee images are resized to 200×300 to preserve the mean aspect ratio (~ 1.6) and fed to the jointly trained CNN. This system provides quantification of knee OA severity in both discrete grades (classification) and continuous grades (regression) simultaneously.

The major pathological features that indicate the onset of knee OA include: reduction in joint space width due to loss of knee cartilage, and the formation of bone spurs (osteophytes) or bony projections along the joint margins. The author believes that quantifying these features along with the KL grades can provide deeper insights to assess knee OA severity and to study the progression of knee OA. Therefore, a deep learning-based automatic knee OA diagnostic system

that can provide simultaneous predictions of KL grades, JSN, and osteophytes is developed.

5.4 Summary and Discussion

Four approaches are presented in this section to automatically assess knee OA severity using CNNs. First, the existing pre-trained CNNs are investigated for classifying knee images based on KL grades. Two methods are used in this approach: using the pre-trained CNNs for fixed feature extraction, and fine-tuning the pre-trained CNNs using the transfer learning approach. The predictions or outputs from these methods are ordinal KL grades (0,1,2,3 or 4). Furthermore, the author argued that quantifying knee OA severity in a continuous scale (0–4) is more appropriate as the OA degradation is progressive in nature, not discrete. Regression is used to quantify the knee OA severity on a continuous scale. The classification and regression results from the proposed methods in this chapter outperform the previous best results achieved by WNDCHRM, which uses many hand-crafted features with a variation of k-nearest neighbour classifier for classifying knee OA radiographs.

Second, CNNs are trained from scratch for classification and regression. The objective was to further improve the quantification results. As the training data is relatively scarce, a lightweight architectures with fewer (~ 4 to ~ 5 million) free parameters are used in the CNNs. The fully trained CNN for classification achieved high classification accuracy in comparison to the pre-trained CNNs. However, the fully trained CNN for regression did not achieve high-performing results as no ground truth of KL grades was available on a continuous scale. Therefore, the discrete KL grades are used to train the CNNs for regression.

Third, CNNs are fully trained using multi-objective learning for simultaneous classification and regression. The intuition behind this is optimising a CNN with two loss functions provide a stronger error signal and it is a step to improve the overall quantification, considering both classification and regression results. The jointly trained CNN achieved better quantification results with a high classification accuracy in comparison to the previous methods.

As a last approach, CNNs are fully trained for ordinal regression using a softmax dense layer as the hidden layer. This approach achieved low mean-squared error and outperformed other methods to quantify knee OA severity in a continuous scale. The added benefit of this method is to provide simultaneous multi-class classification output.

In summary, a progressive improvement is achieved in the quantification performance with an increase in classification accuracy and other performance metrics in the four approaches to automatically quantify knee OA severity. To conclude this Section, an error analysis was presented that discusses the possible reasons for the misclassification from the jointly trained CNN. The variations in the X-ray imaging protocols and discrepancies in the KL grades scoring need to be taken into account when analysing the mis-classification.

6 Conclusion

The main goal of this research is to advance the state-of-the-art in computer aided diagnostics of the severity of knee OA by developing deep learning based automatic methods. According to the literature, automatic assessment of knee OA severity has been previously approached as an image classification problem and existing approaches report low accuracy for multi-class and classification of successive grades. The state-of-the-art machine learning based methods are investigated for image classification, and we developed new methods using convolutional neural networks (CNNs) to automatically classify knee OA images. A significant outcome of this research is a new automatic knee OA diagnostic system that achieves high accuracy, on par with radiologic reliability readings, which are considered the gold standard for knee OA assessment.

In recent years, deep learning-based approaches, in particular CNNs, have become highly successful in many computer vision tasks and medical applications. Our research has mainly focused on developing a deep learning-based computer aided diagnostic system. The proposed approaches in this chapter are related to two main medical applications: localising or automatically detecting and extracting a region of interest (ROI) from a radiograph, and classifying the ROI to automatically assess disease severity. The FCN-based localisation approach could be extended to other medical applications such as localising a substructure or a ROI in MRI and CT scan images, object or lesion detection, locating anatomical landmarks or identifying imaging markers to study the disease progression. For instance, we followed a similar FCN-based approach to automatically detect and quantify ischemic injury (brain lesions) on diffusion-weighted MRI of infants, and we improved the state-of-the-art by achieving promising results.

In the author's opinion the most interesting research findings in this research are as follows. First, fine-tuning off-the-shelf CNNs pre-trained on very large datasets such as ImageNet (with $\sim 1\text{M}$ images) to classify knee images with relatively small datasets (with $\sim 10,000$ images) is promising for medical image classification. The main challenge in medical image classification is a lack of sufficient annotated data for training deep networks from scratch. Fine-tuning existing CNNs that have been trained using a large annotated dataset from a different application is possibly the best alternative to full training for medical applications. A second extremely interesting result is that training CNNs, optimising a weighted ratio of two loss functions for simultaneous classification and regression, provides a better error signal to the network and improves the overall classification performance. Many diseases are progressive by nature such as Alzheimer's disease, cancer, emphysema, tumours, lesions, and muscular dystrophy. Automatic quantification of such diseases using jointly trained CNNs may improve the quantification performance and provide insights to study the progression of the disease. Finally, it is very interesting that using multi-objective convolutional learning to jointly train CNNs based on different diagnostic features of a disease as the ground truth, can produce an overall improvement in the quantification

performance achieving, results on par with human accuracy. Multi-objective learning and joint prediction of multiple regression and classification variables may be useful to assess diseases involving multiple diagnostic features like Alzheimer's, multiple sclerosis, and multiple myeloma (cancer).

There are several potential directions for future work and further development of the research in this chapter. Some of the interesting extensions and prospects are outlined as follows.

Training an end-to-end deep learning model: The knee OA diagnostic pipeline consists of two steps: 1) localising the knee joints in radiographs and 2) assessing the knee OA severity from the localised knee joints. A FCN was trained for automatic localisation and a CNN was jointly trained for classification and regression of knee OA images. It would be interesting to train a single deep learning model integrating the FCN for localisation and the CNN for classification and/or regression, as this could further improve the automatic assessment of knee OA. In a recent work, Gorriz et al. [86] presented a CNN attention-based end-to-end architecture to automatically assess knee OA severity.

Using semantic segmentations to measure joint space width: Among the knee OA diagnostic features, joint space narrowing is highly sensitive to changes due to disease progression. The proposed approach to automatically localise the knee joints using fully convolutional network can be extended for semantic segmentation of the knee joints and can be used to automatically measure the joint space width (JSW) between the femur and tibia. However, pixel level knee joint annotations in radiographs are needed to measure the JSW.

Assessing the progression of knee OA severity: The automatic quantification methods developed in this study can be extended to assess the progression and early detection of knee OA severity. The baseline datasets are used from the OAI and the MOST dataset. Datasets are available for annual follow-up visits up to 9 years. These datasets could be used to detect the features predictive of radiographic knee OA progression. Shamir et al. reported a similar approach using WNDCHRM that predicted whether a knee would change from KL grade 0 to grade 3 with 72% accuracy using 20 years of data [3]. The findings of our current work has indicated the ability to improve upon the existing methods.

Relating the automatic quantification results to knee pain: The primary clinical features to assess knee OA are pain and radiographic evidence of deformity [87]. It would be interesting to study the relationship between the automatic assessments of the proposed methods (KL grades) to WOMAC scores for knee pain. WOMAC is one among the most widely used assessments in knee OA.

Relating the automatic quantification results to physiological variables: There are several pathological and physiological variables available in the OAI and the MOST datasets. These variables include potential predictors of knee pain status. It would be interesting to study the relationship

between the outcomes of the automatic methods and the predictions from the pathological and physiological variables. Abedin et al. [88] presented a comparative analysis to predict knee OA severity based on statistical models developed on physiological variables and a CNN model developed on features from X-ray images.

Investigating human level accuracy: The radiologic reliability readings from the OAI used 150 participants (300 knees) to evaluate the test-retest reliability of semi-quantitative readings. This is considered the current gold standard for knee OA assessment. Simple kappa coefficients and weighted kappa coefficients were used to evaluate the inter-rater agreement. Investigating the human level accuracy for a large sample and comparing it with the automatic quantification results would provide insight to help reduce the error involved in automatic assessments.

Acknowledgments.

This publication has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) under grant numbers SFI/12/RC/2289 and 15/SIRG/3283.

The OAI is a public-private partnership comprised of five contracts (N01-AR-2-2258; N01-AR-2-2259; N01-AR-2-2260; N01-AR-2-2261; N01-AR-2-2262) funded by the National Institutes of Health, a branch of the Department of Health and Human Services, and conducted by the OAI Study Investigators. Private funding partners include Merck Research Laboratories; Novartis Pharmaceuticals Corporation, GlaxoSmithKline; and Pfizer, Inc. Private sector funding for the OAI is managed by the Foundation for the National Institutes of Health.

MOST is comprised of four cooperative grants (Felson – AG18820; Torner – AG18832; Lewis – AG18947; and Nevitt – AG19069) funded by the National Institutes of Health, a branch of the Department of Health and Human Services, and conducted by MOST study investigators. This manuscript was prepared using MOST data and does not necessarily reflect the opinions or views of MOST investigators.

References

1. Shamir, L., Ling, S.M., Scott Jr, W.W., Bos, A., Orlov, N., Macura, T.J., Eckley, D.M., Ferrucci, L., Goldberg, I.G.: Knee X-ray image analysis method for automated detection of osteoarthritis. *IEEE Transactions on Biomedical Engineering* **56**(2) (2009)
2. Thomson, J., O'Neill, T., Felson, D., Cootes, T.: Automated shape and texture analysis for detection of osteoarthritis from radiographs of the knee. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer (2015) 127–134
3. Shamir, L., Ling, S.M., Scott, W., Hochberg, M., Ferrucci, L., Goldberg, I.G.: Early detection of radiographic knee osteoarthritis using computer-aided analysis. *Osteoarthritis and Cartilage* **17**(10) (2009) 1307–1312

4. Shamir, L., Orlov, N., Eckley, D.M., Macura, T., Johnston, J., Goldberg, I.G.: Wndchrm—an open source utility for biological image analysis. *Source code for biology and medicine* **3**(1) (2008) 13
5. Orlov, N., Shamir, L., Macura, T., Johnston, J., Eckley, D.M., Goldberg, I.G.: WND-CHARM: Multi-purpose image classification using compound image transforms. *Pattern recognition letters* **29**(11) (2008) 1684–1693
6. Oka, H., Muraki, S., Akune, T., Mabuchi, A., Suzuki, T., Yoshida, H., Yamamoto, S., Nakamura, K., Yoshimura, N., Kawaguchi, H.: Fully automatic quantification of knee osteoarthritis severity on plain radiographs. *Osteoarthritis and Cartilage* **16**(11) (2008) 1300–1306
7. Park, H.J., Kim, S.S., Lee, S.Y., Park, N.H., Park, J.Y., Choi, Y.J., Jeon, H.J.: A practical MRI grading system for osteoarthritis of the knee: association with Kellgren–Lawrence radiographic scores. *European journal of radiology* **82**(1) (2013) 112–117
8. Lee, H.: Unsupervised feature learning via sparse hierarchical representations. PhD thesis, Stanford University (2010)
9. Le, Q.V.: Scalable feature learning. PhD Thesis, Stanford University (2013)
10. Yang, S.: Feature Engineering in Fine-Grained Image Classification. PhD Thesis, University of Washington (2013)
11. Donoghue, C.R.: Analysis of MRI for Knee Osteoarthritis using Machine Learning. PhD Thesis, Imperial College London (2013)
12. Litjens, G., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciompi, F., Ghafoorian, M., van der Laak, J.A., van Ginneken, B., Sánchez, C.I.: A survey on deep learning in medical image analysis. *Osteoarthritis and Cartilage* **42**(1) (2017) 60–88
13. Shen, D., Wu, G., Suk, H.I.: Deep learning in medical image analysis. *Annual Review of Biomedical Engineering* (0) (2017)
14. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: *Advances in neural information processing systems*. (2012) 1097–1105
15. Tajbakhsh, N., Shin, J.Y., Gurudu, S.R., Hurst, R.T., Kendall, C.B., Gotway, M.B., Liang, J.: Convolutional neural networks for medical image analysis: full training or fine tuning? *IEEE transactions on medical imaging* **35**(5) (2016) 1299–1312
16. Prasoon, A., Petersen, K., Igel, C., Lauze, F., Dam, E., Nielsen, M.: Deep feature learning for knee cartilage segmentation using a triplanar convolutional neural network. In: *International conference on medical image computing and computer-assisted intervention*, Springer (2013) 246–253
17. Havaei, M., Davy, A., Warde-Farley, D., Biard, A., Courville, A., Bengio, Y., Pal, C., Jodoin, P.M., Larochelle, H.: Brain tumor segmentation with deep neural networks. *Medical image analysis* **35** (2017) 18–31
18. Zhang, W., Li, R., Deng, H., Wang, L., Lin, W., Ji, S., Shen, D.: Deep convolutional neural networks for multi-modality isointense infant brain image segmentation. *NeuroImage* **108** (2015) 214–224
19. Roth, H.R., Farag, A., Lu, L., Turkbey, E.B., Summers, R.M.: Deep convolutional networks for pancreas segmentation in ct imaging. In: *SPIE Medical Imaging, International Society for Optics and Photonics* (2015) 94131G–94131G
20. Ciresan, D., Giusti, A., Gambardella, L.M., Schmidhuber, J.: Deep neural networks segment neuronal membranes in electron microscopy images. In: *Advances in neural information processing systems*. (2012) 2843–2851

21. Marijnissen, A.C., Vincken, K.L., Vos, P.A., Saris, D., Viergever, M., Bijlsma, J., Bartels, L., Lafeber, F.: Knee images digital analysis (kida): a novel method to quantify individual radiographic features of knee osteoarthritis in detail. *Osteoarthritis and Cartilage* **16**(2) (2008) 234–243
22. Felson, D.T., McAlindon, T.E., Anderson, J.J., Weissman, B.W., Aliabadi, P., Evans, S., Levy, D., LaValley, M.P.: Defining radiographic osteoarthritis for the whole knee. *Osteoarthritis and Cartilage* **5**(4) (1997) 241–250
23. Braun, H.J., Gold, G.E.: Diagnosis of osteoarthritis: Imaging. *Bone* **51**(2) (2012) 278–288
24. Emrani, P.S., Katz, J.N., Kessler, C.L., Reichmann, W.M., Wright, E.A., McAlindon, T.E., Losina, E.: Joint space narrowing and kellygren–lawrence progression in knee osteoarthritis: an analytic literature synthesis. *Osteoarthritis and Cartilage* **16**(8) (2008) 873–882
25. Hart, D., Spector, T.: Kellgren & lawrence grade 1 osteophytes in the kneedoubtful or definite? *Osteoarthritis and cartilage* **11**(2) (2003) 149–150
26. Shaikh, H., Panbude, J., Joshi, A.: Image segmentation techniques and its applications for knee joints: A survey. *IOSR Journal of Electronics and Communication Engineering (IOSR-JECE)* **9**(5) (2014) 23–28
27. Sun, Y., Teo, E., Zhang, Q.: Discussions of knee joint segmentation. In: *Biomedical and Pharmaceutical Engineering, 2006. ICBPE 2006. International Conference on, IEEE* (2006)
28. Gornale, S.S., Patravali, P.U., Manza, R.R.: Detection of osteoarthritis using knee x-ray image analyses: A machine vision based approach. *International Journal of Computer Applications* **145**(1) (2016)
29. Tiulpin, A., Thevenot, J., Rahtu, E., Saarakkala, S.: A novel method for automatic localization of joint area on knee plain radiographs. In: *Scandinavian Conference on Image Analysis, Springer* (2017) 290–301
30. Stammberger, T., Eckstein, F., Michaelis, M., Englmeier, K.H., Reiser, M.: Interobserver reproducibility of quantitative cartilage measurements: comparison of b-spline snakes and manual segmentation. *Magnetic resonance imaging* **17**(7) (1999) 1033–1042
31. Cohen, Z.A., McCarthy, D.M., Kwak, S.D., Legrand, P., Fogarasi, F., Ciaccio, E.J., Ateshian, G.A.: Knee cartilage topography, thickness, and contact areas from mri: in-vitro calibration and in-vivo measurements. *Osteoarthritis and cartilage* **7**(1) (1999) 95–109
32. Hirvasniemi, J., Thevenot, J., Immonen, V., Liikavainio, T., Pulkkinen, P., Jämsä, T., Arokoski, J., Saarakkala, S.: Quantification of differences in bone texture from plain radiographs in knees with and without osteoarthritis. *Osteoarthritis and cartilage* **22**(10) (2014) 1724–1731
33. Woloszynski, T., Podsiadlo, P., Stachowiak, G., Kurzynski, M.: A signature dissimilarity measure for trabecular bone texture in knee radiographs. *Medical physics* **37**(5) (2010) 2030–2042
34. Zhao, F., Xie, X.: An overview of interactive medical image segmentation. *Annals of the BMVA* **2013**(7) (2013) 1–22
35. Pirnogi, C.D.: Articular cartilage segmentation and tracking in sequential MR images of the knee. PhD thesis, ETH Zurich (2005)
36. Duryea, J., Li, J., Peterfy, C., Gordon, C., Genant, H.: Trainable rule-based algorithm for the measurement of joint space width in digital radiographic images of the knee. *Medical physics* **27**(3) (2000) 580–591
37. Podsiadlo, P., Wolski, M., Stachowiak, G.: Automated selection of trabecular bone regions in knee radiographs. *Medical physics* **35**(5) (2008) 1870–1883

38. Anifah, L., Purnama, I.K.E., Hariadi, M., Purnomo, M.H.: Automatic segmentation of impaired joint space area for osteoarthritis knee on x-ray image using gabor filter based morphology process. *IPTEK The Journal for Technology and Science* **22**(3) (2011)
39. Lee, H.C., Lee, J.S., Lin, M.C.J., Wu, C.H., Sun, Y.N.: Automatic assessment of knee osteoarthritis parameters from two-dimensional x-ray image. In: *Innovative Computing, Information and Control, 2006. ICICIC'06. First International Conference on*. Volume 2., IEEE (2006) 673–676
40. Subramoniam, M., Rajini, V.: Local binary pattern approach to the classification of osteoarthritis in knee x-ray images. *Asian Journal of Scientific Research* **6**(4) (2013) 805
41. Subramoniam, B., et al.: A non-invasive computer aided diagnosis of osteoarthritis from digital x-ray images. *Biomedical Research* (2015)
42. Deokar, D.D., Patil, C.G.: Effective feature extraction based automatic knee osteoarthritis detection and classification using neural network. *International Journal of Engineering and Techniques* **1**(3) (2015)
43. Yoo, T.K., Kim, D.W., Choi, S.B., Park, J.S.: Simple scoring system and artificial neural network for knee osteoarthritis risk prediction: A cross-sectional study. *PloS one* **11**(2) (2016) e0148724
44. Tiulpin, A., Thevenot, J., Rahtu, E., Lehenkari, P., Saarakkala, S.: Automatic knee osteoarthritis diagnosis from plain radiographs: a deep learning-based approach. *Scientific reports* **8**(1) (2018) 1727
45. Antony, J., McGuinness, K., O'Connor, N.E., Moran, K.: Quantifying radiographic knee osteoarthritis severity using deep convolutional neural networks. In: *Pattern Recognition (ICPR), 2016 23rd International Conference on*, IEEE (2016) 1195–1200
46. Sobel, I.: An isotropic 3×3 image gradient operator. *Machine Vision for three-dimensional Sciences* (1990)
47. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*. Volume 1., IEEE (2005) 886–893
48. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2015) 3431–3440
49. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: *3D Vision (3DV), 2016 Fourth International Conference on*, IEEE (2016) 565–571
50. Kayalibay, B., Jensen, G., van der Smagt, P.: Cnn-based segmentation of medical imaging data. *arXiv preprint arXiv:1701.03056* (2017)
51. Christ, P.F., Ettlinger, F., Grün, F., Elshaera, M.E.A., Lipkova, J., Schlecht, S., Ahmaddy, F., Tatavarty, S., Bickel, M., Bilic, P., et al.: Automatic liver and tumor segmentation of ct and mri volumes using cascaded fully convolutional neural networks. *arXiv preprint arXiv:1702.05970* (2017)
52. Li, F.F., Karpathy, A., Johnson, J.: CS231n: Convolutional Neural Networks for Visual Recognition. (2016) Available at <http://cs231n.github.io/convolutional-networks/>.
53. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014)
54. Kilian, J.: Simple image analysis by moments. *OpenCV library documentation* (2001)

55. Bressan, M., Dance, C.R., Poirier, H., Arregui, D.: Local contrast enhancement. In: *Color Imaging: Processing, Hardcopy, and Applications*. (2007) 64930Y
56. Pizer, S.M., Amburn, E.P., Austin, J.D., Cromartie, R., Geselowitz, A., Greer, T., ter Haar Romeny, B., Zimmerman, J.B., Zuiderveld, K.: Adaptive histogram equalization and its variations. *Computer vision, graphics, and image processing* **39**(3) (1987) 355–368
57. Erhan, D., Szegedy, C., Toshev, A., Anguelov, D.: Scalable object detection using deep neural networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2014) 2147–2154
58. Antony, J., McGuinness, K., Moran, K., OConnor, N.E.: Automatic detection of knee joints and quantification of knee osteoarthritis severity using convolutional neural networks. In: *International Conference on Machine Learning and Data Mining in Pattern Recognition*, Springer (2017) 376–390
59. Ojala, T., Pietikainen, M., Maenpaa, T.: Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on pattern analysis and machine intelligence* **24**(7) (2002) 971–987
60. Shamir, L., Orlov, N., Eckley, D.M., Macura, T., Johnston, J., Goldberg, I.: Wndcharm: Multi-purpose image classifier. *Astrophysics Source Code Library* (2013)
61. Tamura, H., Mori, S., Yamawaki, T.: Textural features corresponding to visual perception. *IEEE Transactions on Systems, Man, and Cybernetics* **8**(6) (1978) 460–473
62. Haralick, R.M., Shanmugam, K., et al.: Textural features for image classification. *IEEE Transactions on systems, man, and cybernetics* **3**(6) (1973) 610–621
63. Grigorescu, S.E., Petkov, N., Kruizinga, P.: Comparison of texture features based on gabor filters. *IEEE Transactions on Image processing* **11**(10) (2002) 1160–1167
64. Teague, M.R.: Image analysis via the general theory of moments. *JOSA* **70**(8) (1980) 920–930
65. Chatfield, K., Simonyan, K., Vedaldi, A., Zisserman, A.: Return of the devil in the details: Delving deep into convolutional nets. In: *Proceedings of British Machine Vision Conference*. (2014)
66. Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., Darrell, T.: Caffe: Convolutional architecture for fast feature embedding. In: *Proceedings of the ACM International Conference on Multimedia*. (2014) 675–678
67. Karayev, S., Trentacoste, M., Han, H., Agarwala, A., Darrell, T., Hertzmann, A., Winnemoeller, H.: Recognizing image style. In: *Proceedings of the British Machine Vision Conference*, BMVA Press (2014)
68. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: ImageNet large scale visual recognition challenge. *International Journal of Computer Vision* **115**(3) (2015) 211–252
69. Yosinski, J., Clune, J., Bengio, Y., Lipson, H.: How transferable are features in deep neural networks? In: *Advances in Neural Information Processing Systems*. (2014) 3320–3328
70. Glorot, X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks. In: *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*. (2010) 249–256
71. Fan, R.E., Chang, K.W., Hsieh, C.J., Wang, X.R., Lin, C.J.: LIBLINEAR: A library for large linear classification. *Journal of machine learning research* **9** (2008) 1871–1874

72. Miao, S., Wang, Z.J., Zheng, Y., Liao, R.: Real-time 2d/3d registration via cnn regression. In: Biomedical Imaging (ISBI), 2016 IEEE 13th International Symposium on, IEEE (2016) 1430–1434
73. Spampinato, C., Palazzo, S., Giordano, D., Aldinucci, M., Leonardi, R.: Deep learning for automated skeletal bone age assessment in x-ray images. *Medical image analysis* **36** (2017) 41–51
74. Roth, H.R., Wang, Y., Yao, J., Lu, L., Burns, J.E., Summers, R.M.: Deep convolutional networks for automated detection of posterior-element fractures on spine ct. In: Proceedings Volume 9785, Medical Imaging 2016: Computer-Aided Diagnosis, SPIE Medical Imaging (2016)
75. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: A Large-Scale Hierarchical Image Database. In: CVPR09. (2009)
76. Liu, S., Yang, J., Huang, C., Yang, M.H.: Multi-objective convolutional learning for face labeling. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2015) 3451–3459
77. Tompson, J.J., Jain, A., LeCun, Y., Bregler, C.: Joint training of a convolutional network and a graphical model for human pose estimation. In: Advances in neural information processing systems. (2014) 1799–1807
78. Ranftl, R., Pock, T.: A deep variational model for image segmentation. In: German Conference on Pattern Recognition, Springer (2014) 107–118
79. Ning, F., Delhomme, D., LeCun, Y., Piano, F., Bottou, L., Barbano, P.E.: Toward automatic phenotyping of developing embryos from videos. *IEEE Transactions on Image Processing* **14**(9) (2005) 1360–1371
80. Rudd, E.M., Günther, M., Boulton, T.E.: Moon: A mixed objective optimization network for the recognition of facial attributes. In: European Conference on Computer Vision, Springer (2016) 19–35
81. Schiphof, D., Boers, M., Bierma-Zeinstra, S.M.: Differences in descriptions of kellgren and lawrence grades of knee osteoarthritis. *Annals of the rheumatic diseases* **67**(7) (2008) 1034–1036
82. Shamir, L., Felson, D.T., Ferrucci, L., Goldberg, I.G.: Assessment of osteoarthritis initiative–kellgren and lawrence scoring projects quality using computer analysis. *Journal of Musculoskeletal Research* **13**(04) (2010) 197–201
83. Gutiérrez, P.A., Perez-Ortiz, M., Sanchez-Monedero, J., Fernandez-Navarro, F., Hervas-Martinez, C.: Ordinal regression methods: survey and experimental study. *IEEE Transactions on Knowledge and Data Engineering* **28**(1) (2016) 127–146
84. Pedregosa, F., Bach, F., Gramfort, A.: On the consistency of ordinal regression methods. *Journal of Machine Learning Research* **18**(55) (2017) 1–35
85. Beckham, C., Pal, C.: A simple squared-error reformulation for ordinal classification. *arXiv preprint arXiv:1612.00775* (2016)
86. Górriz, M., Antony, J., McGuinness, K., Giró-i Nieto, X., OConnor, N.E.: Assessing knee oa severity with cnn attention-based end-to-end architectures. In: International Conference on Medical Imaging with Deep Learning. (2019) 197–214
87. Williams, D.A., Farrell, M.J., Cunningham, J., Gracely, R.H., Ambrose, K., Cupps, T., Mohan, N., Clauw, D.J.: Knee pain and radiographic osteoarthritis interact in the prediction of levels of self-reported disability. *Arthritis Care & Research* **51**(4) (2004) 558–561
88. Abedin, J., Antony, J., McGuinness, K., Moran, K., OConnor, N.E., Rebholz-Schuhmann, D., Newell, J.: Predicting knee osteoarthritis severity: comparative modeling based on patients data and plain x-ray images. *Scientific reports* **9**(1) (2019) 5761