

# SHIFTING FOCUS: FROM GLOBAL SEMANTICS TO LOCAL PROMINENT FEATURES IN SWIN-TRANSFORMER FOR KNEE OSTEOARTHRITIS SEVERITY ASSESSMENT

Aymen Sekhri<sup>1\*</sup>, Marouane Tliba<sup>1\*</sup>, Mohamed Amine Kerkouri<sup>1\*</sup>, Yassine Nasser<sup>1</sup>,  
Aladine Chetouani<sup>1</sup>, Alessandro Bruno<sup>2</sup>, and Rachid Jennane<sup>3</sup>

<sup>1</sup>Laboratoire PRISME, Université d'Orléans, Orléans, France

<sup>2</sup>IULM Libera Università di Lingue e Comunicazione. 20143 Milan, Italy

<sup>3</sup>IDP - UMR CNRS 7013, Université d'Orléans. Orléans, France

## ABSTRACT

Conventional imaging diagnostics frequently encounter bottlenecks due to manual inspection, which can lead to delays and inconsistencies. Although deep learning offers a pathway to automation and enhanced accuracy, foundational models in computer vision often emphasize global context at the expense of local details, which are vital for medical imaging diagnostics. To address this, we harness the Swin Transformer's capacity to discern extended spatial dependencies within images through the hierarchical framework. Our novel contribution lies in refining local feature representations, orienting them specifically toward the final distribution of the classifier. This method ensures that local features are not only preserved but are also enriched with task-specific information, enhancing their relevance and detail at every hierarchical level. By implementing this strategy, our model demonstrates significant robustness and precision, as evidenced by extensive validation of two established benchmarks for Knee OsteoArthritis (KOA) grade classification. These results highlight our approach's effectiveness and its promising implications for the future of medical imaging diagnostics. Our implementation is available on Github.

**Index Terms**— Medical Imaging Diagnostics, Representation Learning, Knee OsteoArthritis (KOA) Grade Classification, Task-Relevant Feature Representation

## 1. INTRODUCTION

Knee OsteoArthritis (KOA) stands as a prevalent arthritis form, majorly afflicting the elderly, characterized by the degenerative wear and tear of knee joint cartilage. Occasionally, joint infections exacerbate this degeneration, manifesting as mobility constraints, pain, and swelling. Although various imaging modalities exist, radiography is preferred for initial KOA assessment due to its cost-effectiveness and widespread availability, revealing critical indicators like joint space narrowing and bone spur formation. Clinicians traditionally employ the Kellgren and Lawrence (KL) grading system [1] to evaluate KOA severity, spanning five stages from normal (KL-0) to severe (KL-4). Despite its utility, KOA's gradual progression introduces a degree

of subjectivity in these classifications, complicating the automation of diagnosis. The nuanced distinctions necessary for accurate KL grade determination from X-ray images underline the inherent challenges in automating KOA diagnosis, necessitating advanced solutions to enhance diagnostic precision and reliability.

Despite the advancements in deep learning models for image recognition, a persistent limitation exists in their ability to capture detailed local features and construct hierarchical representations concurrently, as the opted training strategies for this purpose push the network to be invariant to local feature transformation. In contrast, this limitation is particularly pronounced in medical imaging, where the nuanced differentiation of local textural features within a joint is crucial. For KOA diagnosis, the ability to discern subtle variations in joint texture while contextualizing the overall joint form is paramount [2]. Traditional deep-learning approaches and training strategies tend to excel in identifying broad patterns but often falter at isolating minute yet diagnostically critical details. This inadequacy is especially detrimental in the context of KOA, where the early detection of subtle textural changes can significantly influence treatment outcomes [2]. As such, there is a pressing need for a training strategy that can adeptly navigate the intricacies of medical images, capturing the fine-grained details necessary for accurate diagnosis without compromising on the broader contextual understanding.

Addressing this gap, our research makes two fundamental contributions:

- We leverage the Swin Transformer architecture's ability to capture localized features at each layer with a long range of feature space dependencies; combining this information with the global representation ensures a comprehensive representation that integrates both relevant local features and global information.
- We refine the final representation by employing a regularization term: we optimize the alignment between local features and the network's decision-making features distribution layer using a Negative Cosine Similarity Loss

\*Authors with equal contributions.

(NCSL). This technique works as a regularization term to the hierarchical representation abstraction learning process, ensuring critical local details are preserved and emphasized during the classification process.

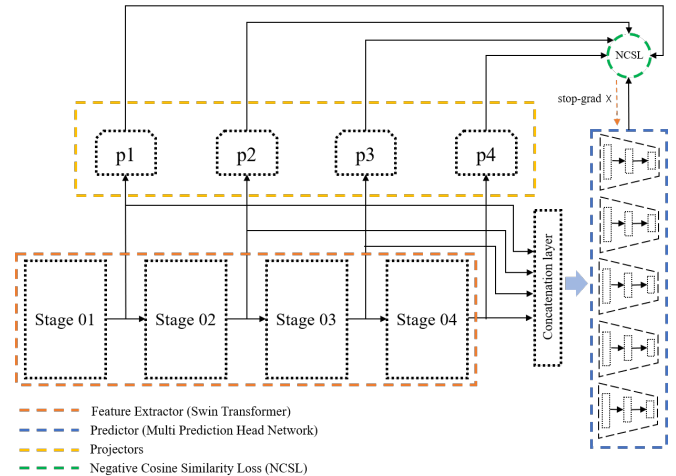
This approach ensures that while the network maintains its robustness to global context and transformations, it does not sacrifice the critical local information essential for precise medical diagnostics, thereby aligning the model’s learning process with the unique demands of medical imaging in KOA.

## 2. RELATED WORKS

In Knee Osteoarthritis (KOA) diagnosis, deep learning (DL) methodologies have significantly advanced, leveraging rich datasets from the Osteoarthritis Initiative (OAI) [3] and the Multicenter Osteoarthritis Study (MOST) [4]. Various approaches have sought to harness the potential of DL in discerning KOA severity and progression. Noteworthy endeavors include Antony et al.[5], who utilized Convolutional Neural Networks (CNNs) for a two-stage assessment involving knee joint localization followed by classification, innovatively combining classification and regression losses for enhanced accuracy. Similarly, Tiulpin et al. [6] employed a Siamese CNN structure to estimate KOA KL grades, amalgamating multiple model outputs for robust predictions. Chen et al. [7] introduced an ordinal loss-based strategy to optimize CNNs for KOA grading, leveraging the KL scale’s inherent ordinality. Nasser et al. [8] further contributed by developing a Discriminative Regularized Auto-Encoder (DRAE) targeting early KOA detection, emphasizing feature distinction via a discriminative penalty. Emerging research by Wang et al. [9] explored Vision Transformer architectures, complemented by novel data augmentation methods and a hybrid loss function to address early KOA identification and tackle data drift challenges in multi-center studies, a topic we have previously addressed to enhance cross-dataset model efficacy. Nasser et al. [10] devised a deep neural network that synergizes shape and texture analysis.

## 3. PROPOSED METHOD

Our proposed methodology integrates four key components: (1) A Swin Transformer is utilized as the primary feature extractor, capitalizing on its efficacy in hierarchical feature representation. (2) A multi-prediction head network is deployed for classification, ensuring robust and accurate categorization. (3) Skip connections are integrated within the architecture to facilitate enhanced information flow and feature retention. (4) The adoption of Negative Cosine Similarity Loss (NCSL) across the transformer stages serves as a regularization term, optimizing local feature refinement and aligning it with the global representation, thus enhancing model performance. For the architectural visualization, refer to Figure 1.



**Fig. 1:** Overview of the Proposed KOA Diagnostic Model. This figure outlines our advanced model, highlighting the Swin Transformer for hierarchical feature extraction, multiple prediction heads for detailed KL grade classification, skip connections for effective feature flow, and Negative Cosine Similarity Loss for feature optimization. Together, these elements illustrate our novel approach to balancing local detail recognition with global feature abstraction for improved KOA diagnostic accuracy.

### 3.1. Swin Transformer

The Swin Transformer, delineated in [11], is a pivotal advancement in adapting transformers for computer vision, featuring hierarchical feature maps and an optimized Multi-head Self Attention (MSA) framework to improve efficiency and effectiveness. Initially, the model partitions input images of size  $H \times W \times 3$  into  $4 \times 4$  patches, which are then embedded and processed through the model’s stages. The initial stage involves a linear embedding, followed by two Swin Transformer blocks using Window Multi-head Self-Attention (W-MSA) and Shifted Window Multi-head Self-Attention (SW-MSA). Subsequent stages merge and transform these patches, diminishing their quantity while increasing dimensionality through additional Swin Transformer blocks. By its final stage, the model refines the patches to a significantly reduced count with augmented feature representation, underscoring the Swin Transformer’s utility in computer vision tasks.

### 3.2. Multi-Prediction Head Network

The Multi-Prediction Head Network (MPHN), integral to our methodology and elaborated upon in prior work [12], deconstructs the multi-class task into simpler binary classifications. It features a series of Multilayer Perceptrons (MLPs), each with three layers, designed to identify individual KL grades, thereby encompassing five distinct MLPs tailored for specific grade classifications.

### 3.3. Skip Connection

Medical imaging diagnostics emphasize localized features, contrasting with the broader focus typical of natural image analysis. Utilizing the Swin Transformer’s hierarchical feature mapping, as expounded in Section 3.1, we derive and normalize stage-specific outputs. These, coupled with an average pooling step, culminate in refined local feature representations. By amalgamating these with the comprehensive output from the final stage, our model achieves a synergistic a balanced consideration of both detailed local attributes and global semantic context.

However, the amalgamation of the features representation is not valid without a distillation of the representation for the following considerations, first, the local features space may contain noises that shift the deep semantic representation, second, the two representations (local features, and global features) should be in the same features space dimension.

### 3.4. Optimization of Hierarchical Local Feature Representation

Our methodology focuses on enhancing the hierarchical representation of local features, we took insight from advances in Knowledge distillation [13] and self-supervised learning [14, 15] [16] from features similarities. We formulate our contribution as a regularization term that helps the network to keep the prominent local features within Swin transformer, denoted as a learnable function  $f_\theta$ , where  $\theta$  represents the model parameters. For each stage  $s$  of the transformer, we obtain an output  $O_s = f_\theta(I_s)$ , where  $I_s$  represents the input to that stage. To refine these outputs further and align their representation spaces, we introduce projection heads  $P_{\phi_s}$  for each stage  $s$ , parameterized by  $\phi_s$ . The projected representation for each stage is given by  $P_s = P_{\phi_s}(O_s)$ . The concatenated output  $C$  of these projected representations from all stages forms the enhanced feature representation and is given by:

$$C = \bigoplus_{s=1}^S P_s \quad (1)$$

where  $\bigoplus$  denotes the concatenation operation, and  $S$  is the total number of stages.

For classification, building upon the preliminary context discussed in Section 3.2, we instantiate separate classifier networks, denoted as  $C_{\psi_k}$  for each class  $k$ , where each network is parameterized by  $\psi_k$ . This architecture enables tailored processing for each specific class, ensuring independent and specialized distillation of the relevant KL grade representation during classification. Each classifier network, conceptualized to encapsulate the feature’ space distribution pertinent to its designated class, functions as a sophisticated decision maker in our framework.

Formally, the operation of the classifier for each class  $k$  can be delineated as follows: The decision maker network processes

the concatenated feature representation  $C$ , yielding an intermediate output  $D_k = C_{\psi_k}(C)$ . To transition from this multidimensional output to a scalar prediction value, we introduce a learnable parameter vector  $\omega_k$  for each class  $k$ , facilitating the necessary dimensionality reduction. Consequently, the predictive likelihood  $y_k$  for each class is computed via a sigmoid activation function  $\sigma$ , applied to the dot product of  $D_k$  and  $\omega_k$ :

$$y_k = \sigma(\omega_k^\top D_k) \quad (2)$$

where  $\sigma(\cdot)$  denotes the sigmoid function, ensuring that the final output  $y_k$  resides within the interval  $[0, 1]$ , representing the probability or likelihood of class  $k$  being the correct classification.

The empirical risk minimization for each class  $k$  is conducted using Binary Cross-Entropy (BCE), formulated as:

$$\mathcal{L}_{BCE_k} = - \sum_n [y_n^k \log(\hat{y}_n^k) + (1 - y_n^k) \log(1 - \hat{y}_n^k)] \quad (3)$$

where  $y_n^k$  and  $\hat{y}_n^k$  are the true and predicted labels for the  $n$ -th sample and class  $k$ , respectively.

In parallel, we aim to minimize the Negative Cosine Similarity Loss (NCSL) between the projected features  $P_s$  and the features distribution from the classifier, enhancing the feature space’s relevance. The NCSL inspired from [15] is defined as:

$$\mathcal{L}_{NCSL} = - \frac{1}{S} \sum_{s=1}^S \frac{P_s}{\|P_s\|_2} \cdot \frac{\text{sg}(D)}{\|\text{sg}(D)\|_2} \quad (4)$$

where  $D$  represents the aggregated feature distribution across  $k$  decision-makers. The function  $\text{sg}(\cdot)$  denotes the stop-gradient operation applied to  $D$ , essential for detaching the main branch during optimization. This selection is intentional, focusing on a singular optimized pathway to reinforce the regularization effect. The primary objective here is to meticulously align the local feature representations  $P_s$  with the global classifier-driven feature landscape  $D$  as ground truth target, fostering a synergistic enhancement in feature discernment and classification efficacy. Our implementation seeks to deeply integrate these local representations with the classifier’s insights, ensuring a robust and contextually aware feature synergy for improved model interpretability and performance.

The total loss  $\mathcal{L}$  combines the BCE loss for each class and the NCSL, providing a comprehensive learning signal:

$$\mathcal{L} = \sum_{k=1}^K \mathcal{L}_{BCE_k} + \lambda \mathcal{L}_{NCSL} \quad (5)$$

where  $K$  is the total number of classes and  $\lambda$  is a coefficient balancing the two loss components. This loss function orchestrates the overall learning process, emphasizing both global risk minimization and the fidelity of local feature representations.

## 4. RESULT ANALYSIS

### 4.1. Dataset

In this study, we employed two commonly used databases, the OsteoArthritis Initiative (OAI) [3] and the Multicenter Osteoarthritis Study (MOST) [4]. We adopted the identical data preprocessing approach as described in our earlier publication [12]. However, a noteworthy departure in our methodology was the use of the MOST dataset for training our feature extractor. Subsequently, we employed the pre-trained weights from this feature extractor as an initialization step for training on the OAI dataset.

Setups	Method				Metrics		
	A	B	C	D	ACC (%)	B-ACC (%)	F1-Score
1	✓	✗	✗	✗	70.89	69.07	0.698
2	✓	✗	✗	✓	71.68	69.13	0.687
3	✗	✓	✗	✗	71.44	67.46	0.690
4	✗	✓	✗	✓	<b>72.40</b>	<b>71.11</b>	<b>0.704</b>
5	✗	✗	✓	✗	66.55	54.05	0.528
6	✗	✗	✓	✓	67.63	54.29	0.534

**Table 1:** Results of the ablation study using the OAI test-set. A: SPHN, B: MPHN, C: MLPReg, D: NCSL.


### 4.2. Ablation Study

Our ablation study evaluates the efficacy of different decision frameworks within our model, contrasting three configurations: a Single Prediction Head Network (SPHN) that employs a universal MLP across all KL grades, a Multi Prediction Head Network (MPHN) as elaborated in Section 3.2, and an MLPRegressor that conceptualizes the task in a regression paradigm.

Key observations from our analysis, delineated in Table 1, are summarized as follows: **(1) SPHN Evaluation:** In Setup 1, we employed SPHN with integrated feature concatenation from the Swin Transformer, establishing our baseline. Enhancement was observed in Setup 2 where the NCSL application improved ACC from 70.89% to 71.68% and yielded a modest uplift in F1-Score. **(2) MPHN Evaluation:** Advancing to Setups 3 and 4 with the MPHN, marked improvements were noted. Significantly, Setup 4, which synergizes MPHN with NCSL, recorded the highest metrics: ACC (72.40%), B-ACC (71.11%), and F1-Score (0.704), attesting to the MPHN’s capacity to distinctively refine and exploit grade-specific features, as well as the impact of using NCSL compared to setup 3. **(3) MLPRegressor Evaluation:** The last frameworks, Setups 5 and 6, implementing an MLPRegressor, demonstrated marginal enhancements, peaking at an ACC of 67.63%, suggesting a constrained utility of the regression format for this classification context. To sum up, these findings substantiate the integration of MPHN and NCSL as the most efficacious approach.

### 4.3. Qualitative Analysis

We used GradCAM [17] as a tool to visualize the activations of the final layer of the feature extractor. Within our analysis,



**Fig. 2:** [17] demonstrate our model’s ability to discern progression in KOA severity from KL grade 0 (healthy) to KL grade 4 (most severe). These visualizations underscore the model’s focus on both local and global features, adjusting according to severity. They confirm our model’s effective use of various hierarchical feature levels without central region bias, relying on both joint’ local textures and global form

we carefully selected representative samples from each grade of knee OA (KL-0 to KL-4, ordered left to right). Visualizations are presented in Figure 2. As can be seen this Figure reveals that the model adeptly identifies crucial features such as osteophytes, joint space narrowing, and sclerosis. These elements hold significant clinical importance in assessing the severity of knee OA [18].

### 4.4. Comparison to State-Of-The-Art (SOTA)

Table 2 presents a comparison of the obtained results to SOTA methods.

Model	Accuracy (%) ↑	F1-Score ↑
Antony et al. 2016 [19]	53.40	0.43
Antony et al. 2017 [5]	63.60	0.59
Tiulpin et al. 2018 [20]	66.71	-
Chen et al. (Vgg19) 2019 [7]	69.60	-
Chen et al. (ResNet50) 2019 [7]	66.20	-
Chen et al. (ResNet101) 2019 [7]	65.50	-
Wang et al. 2021 [9]	69.18	-
Sekhri et al. 2023 [12]	70.17	0.67
<b>Ours - Setup 4</b>	<b>72.40</b>	<b>0.70</b>

**Table 2:** State of the art Comparison on OAI Test set

Antony et al. [19] and [5] reported their diagnostic accuracies as 53.40% and 63.60%, respectively, accompanied by F1-scores of 0.43 and 0.59. Meanwhile, Chen et al. [7] employed an ordinal loss strategy with various deep learning architectures. They achieved accuracies of 69.60%, 66.20%, and 65.50% using VGG19, ResNet50, and ResNet101, respectively. However, the corresponding F1 scores were not provided in their study. Tiulpin et al. [20] harnessed a Siamese network, achieving an ACC of 66.71%. Furthermore, Wang et al. [9] demonstrated an accuracy of 69.18%. Our proposed method surpasses all aforementioned previous methods, with an ACC of 70.17% and a F1-score of 0.67.

Our contributions, specifically setup 4 (Table 1), stands out as the best approach, surpassing all previous works. Notably, it achieves an enhanced accuracy of 72.40% and an improved F1-Score of 0.70. These findings reflect the ongoing refinement of our methodology by preserving the learned local features as well as aligning the local feature representations with the global classifier-driven feature.

## 5. CONCLUSION

In conclusion, our study presents a comprehensive evaluation of an advanced model tailored for the task of KOA grade classification. Throughout an extensive experiment, we have demonstrated the superior performance of our proposed approach, particularly when employing the MPHNL in conjunction with NCSL as regularization for feature refinement. The results indicate our model accuracy and reliability of KOA diagnostics. Furthermore, our proposed novel approach underscores the significance of adopting a specialized, hierarchical feature processing strategy. This research opens a significant path forward for the application of deep learning for medical imaging, paving the way for more precise and dependable automatic diagnostic tools in the field of computational radiology.

## 6. REFERENCES

- [1] Jonas H Kellgren and JS1006995 Lawrence, "Radiological assessment of osteo-arthritis," *Annals of the rheumatic diseases*, vol. 16, no. 4, pp. 494, 1957.
- [2] Ahmad Almhdi-Imjabbar, Hechmi Toumi, and Eric Lespessailles, "Radiographic biomarkers for knee osteoarthritis: A narrative review," *Life*, vol. 13, no. 1, 2023.
- [3] Osteoarthritis Initiative Investigators, "Osteoarthritis initiative (oai) data," 2006.
- [4] Nancy A Segal, Matthew C Nevitt, Keith D Gross, Jeffrey Hietpas, Natalie A Glass, Celeste E Lewis, and James C Torner, "The multicenter osteoarthritis study: opportunities for rehabilitation research," *PM R*, vol. 5, no. 8, pp. 647–654, 2013.
- [5] Joseph Antony, Kevin McGuinness, Kieran Moran, and Noel E O'Connor, "Automatic detection of knee joints and quantification of knee osteoarthritis severity using convolutional neural networks," in *Machine Learning and Data Mining in Pattern Recognition: 13th International Conference, MLDM 2017, New York, NY, USA, July 15-20, 2017, Proceedings 13*. Springer, 2017, pp. 376–390.
- [6] Aleksei Tiulpin and Simo Saarakkala, "Automatic grading of individual knee osteoarthritis features in plain radiographs using deep convolutional neural networks," *Diagnostics*, vol. 10, no. 11, pp. 932, 2020.
- [7] Pingjun Chen, Linlin Gao, Xiaoshuang Shi, Kyle Allen, and Lin Yang, "Fully automatic knee osteoarthritis severity grading using deep neural networks with a novel ordinal loss," *Computerized Medical Imaging and Graphics*, vol. 75, pp. 84–92, 2019.
- [8] Yassine Nasser, Rachid Jennane, Aladine Chetouani, Eric Lespessailles, and Mohammed El Hassouni, "Discriminative regularized auto-encoder for early detection of knee osteoarthritis: data from the osteoarthritis initiative," *IEEE transactions on medical imaging*, vol. 39, no. 9, pp. 2976–2984, 2020.
- [9] Yifan Wang, Xianan Wang, Tianning Gao, Le Du, and Wei Liu, "An automatic knee osteoarthritis diagnosis method based on deep learning: data from the osteoarthritis initiative," *Journal of Healthcare Engineering*, vol. 2021, pp. 1–10, 2021.
- [10] Yassine Nasser, Mohammed El Hassouni, and Rachid Jennane, "Discriminative deep neural network for predicting knee osteoarthritis in early stage," in *Predictive Intelligence in Medicine: 5th International Workshop, PRIME 2022, Held in Conjunction with MICCAI 2022, Singapore, September 22, 2022, Proceedings*. Springer, 2022, pp. 126–136.
- [11] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 10012–10022.
- [12] Aymen Sekhri, Mohamed A Kerkouri, Aladine Chetouani, Marouane Tliba, Yassine Nasser, Rachid Jennane, and Alessandro Bruno, "Automatic diagnosis of knee osteoarthritis severity using swin transformer," in *Proceedings of the 20th International Conference on Content-Based Multimedia Indexing*, New York, NY, USA, 2023, CBMI '23, p. 41–47, Association for Computing Machinery.
- [13] Xiatian Zhu, Shaogang Gong, et al., "Knowledge distillation by on-the-fly native ensemble," *Advances in neural information processing systems*, vol. 31, 2018.
- [14] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al., "Bootstrap your own latent—a new approach to self-supervised learning," *Advances in neural information processing systems*, vol. 33, pp. 21271–21284, 2020.
- [15] Xinlei Chen and Kaiming He, "Exploring simple siamese representation learning," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 15750–15758.
- [16] Marouane Tliba, Aladine Chetouani, Giuseppe Valenzise, and Frédéric Dufaux, "Balancing representation abstractions and local details preservation for 3d point cloud quality assessment," in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 1–5.
- [17] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra, "Grad-cam: Visual explanations from deep networks via gradient-based localization," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 618–626.
- [18] Michelle J Lespasio, Nicolas S Piuze, M Elaine Husni, George F Muschler, AJ Guarino, and Michael A Mont, "Knee osteoarthritis: a primer," *The Permanente Journal*, vol. 21, 2017.
- [19] Joseph Antony, Kevin McGuinness, Noel E O'Connor, and Kieran Moran, "Quantifying radiographic knee osteoarthritis severity using deep convolutional neural networks," in *2016 23rd international conference on pattern recognition (ICPR)*. IEEE, 2016, pp. 1195–1200.
- [20] Aleksei Tiulpin, Jérôme Thevenot, Esa Rahtu, Petri Lehenkari, and Simo Saarakkala, "Automatic knee osteoarthritis diagnosis from plain radiographs: a deep learning-based approach," *Scientific reports*, vol. 8, no. 1, pp. 1–10, 2018.