A novel method for automatic localization of joint area on knee plain radiographs

Aleksei Tiulpin¹, Jerome Thevenot¹, Esa Rahtu², and Simo Saarakkala^{1,3}

¹ Research Unit of Medical Imaging, Physics and Technology,
University of Oulu, Finland

² Center for Machine Vision and Signal Analysis,
University of Oulu, Finland

³ Department of Diagnostic Radiology, Oulu University Hospital, Finland

Abstract. Osteoarthritis (OA) is a common musculoskeletal condition typically diagnosed from radiographic assessment after clinical examination. However, a visual evaluation made by a practitioner suffers from subjectivity and is highly dependent on the experience. Computer-aided diagnostics (CAD) could improve the objectivity of knee radiographic examination. The first essential step of knee OA CAD is to automatically localize the joint area. However, according to the literature this task itself remains challenging. The aim of this study was to develop novel and computationally efficient method to tackle the issue. Here, three different datasets of knee radiographs were used ($n = 473/93/77$) to validate the overall performance of the method. Our pipeline consists of two parts: anatomically-based joint area proposal and their evaluation using Histogram of Oriented Gradients and the pre-trained Support Vector Machine classifier scores. The obtained results for the used datasets show the mean intersection over the union equals to: 0.84, 0.79 and 0.78. Using a high-end computer, the method allows to automatically annotate conventional knee radiographs within 14-16ms and high resolution ones within 170ms. Our results demonstrate that the developed method is suitable for large-scale analyses.

Keywords: Knee Radiographs, Medical Image Analysis, Object localization, Proposal generation

1 Introduction

Hip and knee osteoarthritis (OA) are globally ranked as the 11th highest contributor to disability [4]. Typically, OA is diagnosed at a late stage when there is no cure available anymore and when joint replacement surgery becomes the only remaining option. However, if the disease could be diagnosed at an early stage, its progression could be slowed down or even stopped. Such diagnostics is currently possible; however, it involves the use of expensive imaging modalities, which is clinically unfeasible in the primary health care.

X-ray imaging is the most popular and cheap method for knee OA diagnostics [6,3]. However, the diagnostics at an early stage based on this modality has still some limitations due to several factors. A visual evaluation made by a practitioner suffers from subjectivity and is highly dependent on the experience. Eventually it has been reported that radiologists misdiagnosed OA in 30% of the cases, and that in 20% of the cases a specialist disagrees with his/her own previous decision after a period of time [2,17]. Therefore, to make the diagnostic process more systematic and reliable, computer-aided diagnostics (CAD) can be used to reduce the impact of the subjective factors which alter the diagnostics.

In the literature, multiple attempts have been made to approach knee OA CAD from X-ray images [1,24,23,20,22]. These studies indicate existence of two parts in the diagnostic pipeline: the region of interest (ROI) localization and, subsequently, a classification – OA diagnostics from localized ROI(s). It has been reported (see, section below), that this problem remains challenging and requires a better solution.

The aim of our study was to propose a novel and efficient knee joint area localization algorithm, which is applicable for large-scale knee X-ray studies. Our study has the following novelties:

- We propose a new approach to generate and score knee joint area proposals.

- We report a cross-dataset validation, showing that the our method performs similarly for three different datasets and drastically outperforms the baseline.
- Finally, we show, that the developed method can be used to accurately annotate from hundreds of thousands to millions of X-ray images per day.

2 Related work

In the literature, multiple approaches have been used to localize ROIs within radiographs: manual [12,26], semi-automatic [8] and in a fully automatic manner [1,19,22]. To the best of our knowledge, only the studies by Anthony *et al.* [1] and Shamir *et al.* [22] focused on knee joint area localization. While the problem of knee joint area localization can be implicitly solved by annotating the anatomical bone landmarks using deformable models [21,15], it would be unfeasible to perform large-scale studies since the use of these algorithms is computationally costly. Currently, despite the presence of multiple approaches, no studies have reported so far their applicability among different datasets. This cross-dataset validation is crucial for the development of clinically applicable solutions.

In the recently published large scale study [1], two methods were analyzed: a template matching adapted from a previously published work [22] and a sliding-window approach. Both methods showed limited localization ability, however, the sliding-window approach demonstrated a better performance. In particular, this approach was designed to find knee joint centers in radiographic images using Sobel [7] gradients and a linear Support Vector Machine (SVM) classifier. For each sliding window, an SVM score was computed and eventually the patch having the best score was selected. Subsequently, a 300×300 pixels region was drawn around the selected patch center. As the localization metric, the authors used intersection over the union (IoU), which is also called the Jaccard index, between the drawn region and the manual annotation:

$$IoU = \frac{A \cap B}{A \cup B}, \quad (1)$$

where A is the manual annotation and B is the detected bounding box. While the sliding window approach was better than the template matching (mean IoU over the dataset was 0.36 vs. 0.1), the performance was still insufficient to perform further OA CAD, as it was indicated by the authors themselves. Consequently, there is a need for more effective methods for ROI localization since nowadays large scale studies are possible due to the availability of multiple image cohorts like Osteoarthritis Initiative (OAI) [9] and Multicenter Osteoarthritis study (MOST) [10]. These cohorts contain follow-up data of thousands of normal subjects and subjects with knee, hip and wrist OA.

Finally, it should be mentioned, that the problem of joint area localization is not limited to only knee radiographs. The attention of the research community is also focused on hand radiographs, where OA occurs as well. It has been recently shown how the anatomic structure of the image can be utilized to annotate hand radiographic images [13]. However, to the best of authors' knowledge, there are no studies in the knee domain where such information is used to segment or annotate the image. In this study, we show how such information can be used to annotate knee radiographs.

3 Method

In this section, we describe our joint area localization approach. The method consists of two main blocks: proposal generation and SVM-based scoring. First, we describe a limb anatomy-based proposal generation algorithm. As shown in previous object detection studies [25,27], object proposal approaches can significantly reduce the amount of candidate regions: from hundreds of thousands to a few thousands. Subsequently, we show the proposal scoring step, based on Histogram of Oriented Gradients (HoG) [5] and SVM classifier. Schematically, our approach is presented in Figure 1.

3.1 Data pre-processing

To approach the joint area localization problem on knee radiographs, it would be possible to utilize an exhaustive sliding window approach, however, it is time consuming and inadequate for big data

Fig. 1. Schematic representation of the proposed knee joint area localization method. For more details of the method, see the sections 3.2 and 3.3.

analysis. Thus, we propose an anatomically-based proposal generation approach to significantly reduce the amount of checked joint locations. Our method is trained for only one leg, despite the fact that two legs are usually imaged (see the leftmost block in Figure 1). We obtain the method working for both legs by mirroring the leg which is not used for training the method. Therefore, below we describe the problem for the image containing only one leg and denote it as $\mathbf{I}$ having size $C \times H$.

3.2 Region proposal

The core idea and novelty behind our proposal generation are to utilize the anatomic structure of the limb. As the prior information of the knee joint anatomy is known, it can be efficiently utilized. Considering marginal distribution of intensities of a leg, it can be noticed, that around the joint area, there is an intensity increase due to the presence of the patella, and then a sharp intensity drop due to the space between the joints. In our approach, we detect these changes and use their locations as Y-coordinates of region proposals.

At first, we take the middle of an input image, and then sum all the values horizontally to obtain the sequence I_y , which corresponds to a vertical marginal distribution of pixel intensities:

$$I_y[i - \alpha] = \sum_{j=\frac{1}{3}C}^{\frac{2}{3}C-1} \mathbf{I}[i, j], \forall i \in [\alpha, H - \alpha]. \quad (2)$$

Here, we do not sum all the values – instead, we use a margin α to ignore the outliers which are usually present in the top and the bottom of the image. This also reduces the computational complexity of the method. The next step is to identify the intensity peaks located near the patella. We apply the derivation to the obtained sequence in order to detect the anatomical changes and a moving average as a convolution with a window function $w[\cdot]$ of size s_w to reduce the number of peaks. Eventually, we use a sequence obtained by taking the absolute of each of the I_y values:

$$I_y[i] = |(I_y' * w)[i]|, \forall i \in [0, H - 2\alpha]. \quad (3)$$

At the last step, we obtain each k -th index out of the top $\tau\%$ of I_y values. It should be mentioned, that since the margin α was used, it has to be added to each selected index of I_y . A visualization of the procedure described above is given in Figure 2.

Finally, we take the proposed Y-coordinates and as X-coordinates derive as $x = \frac{1}{2}W + j, \forall j \in [d_1, d_1 + p, d_1 + 2p, \dots, d_2]$, where p is a displacement step, which can be estimated using a validation set (see, section 4.2). The procedure described above is repeated for each leg.

To generate the proposals at multiple scales, we use a data-driven approach, which requires a training set having manual annotations. Here, we consider a joint area to be within a square of size S_n , where S_n is proportional to an image height H with some factor $\frac{1}{Z_n}$. Using manual annotations of training data we estimate a set $\mathbf{Z}$ having scales Z_n for each training image. Eventually, having an image $\mathbf{I}$ of size $C \times H$, we use a set $\mathbf{S}$ of proposal sizes based on the following estimations: $S_n = \frac{1}{Z_n}H, \forall Z_n \in [\dots, \bar{Z} - \sigma(\mathbf{Z}), \bar{Z}, \bar{Z} + \sigma(\mathbf{Z}), \dots]$, where $\bar{Z}$ is the mean and $\sigma(\mathbf{Z})$ is the standard deviation of the scales set estimated from the training data.

Fig. 2. Visualization of the joint center Y-coordinate proposal procedure for the left and right leg images – left and right columns, correspondingly. a) and b) show the sub-areas of the leg images, the red lines indicate the zones which are used for the analysis in equation 2. c) and d) show the obtained sequences I_y from equation 2. e) and f) show the result of applying equation 3, red dots indicate the values used to obtain Y-coordinates.

To conclude, the exact amount of generated proposals for each image is calculated as

$$|S| [R + (R \bmod k \neq 0)] \cdot [d_2 - d_1 + ((d_2 - d_1) \bmod p \neq 0)], \quad (4)$$

where $R = 0.01\tau \cdot (H - 2\alpha)$. Here, we use a product rule and consider that all proposals belong to the image area. To estimate the number of proposals for the whole image with two leg image, their amount has to be multiplied by 2.

3.3 Proposal scoring

The next step in our pipeline is to train a classification algorithm in order to select the best candidate region among generated proposals. For that, we use HoG feature descriptor (to compactly describe the knee joint shape) and linear SVM classifier. This combination has been successfully used for human detection and has demonstrated excellent detection performance[5]. Additionally, the extraction of HoG features and SVM-based scoring are relatively fast pipeline blocks, which significantly supported our choice of the proposal scoring approach.

The HoG-SVM pipeline is implemented as follows: for each half of the image we generate the proposals and downscale each of them into a patch of 64×64 pixels. Then, we compute HoG features using the following parameters: 16 blocks with 50% overlapping, 4 cells in a block and 9 bins [5]. Finally, we classify vectors of 1,764 values (7 blocks vertically times 7 blocks horizontally times 4 cells times 9 bins). Basic maximum search is used to find the best scored proposal.

4 Experiments

4.1 Data

In this study we used several data sources: 1,574 high resolution knee radiographs from MOST cohort[10], 93 radiographs from [16] (Dataset obtained from Central Finland Central Hospital, Jyväskylä. The dataset is called "Jyväskylä dataset" in this paper), and 77 from OKOA dataset [18]. The images in all datasets contain knees with different severity of OA as well as implants. The images from MOST were used to create training, validation and test sets, while the remaining were only used to assess the generalization ability of the developed algorithm. More detailed information about the data and their usage is presented in Table 1.

We converted the original X-ray data from 16 bit DICOM files to a 8 bit format for standardization: we truncated their histograms between the 5th and the 99th percentiles to remove corrupted

Table 1. Description of datasets used in this study. The training data were used to train the algorithm, validation data to find hyperparameters for the algorithm, and the testing set to assess the localization and detection performance of the proposed method.

Dataset	Training set	Validation set	Test set	Average image size
MOST	991	110	473	3588×4279
Jyväskylä	-	-	93	2494×2048
OKOA	-	-	77	2671×2928

pixels. Then we normalized the images by scaling their dynamic range between 0 and 255. After that, we manually annotated all used images using ad-hoc MATLAB-based tool. Main criterion for the annotation was, that the ROI should include the joint itself and the fibula bone.

To create a dataset for training a HoG-SVM proposal scoring, we processed the original training data as follows: for each knee on the training images, we generated the proposals and marked them as positive if the IoU of the manual annotation and the generated proposal was greater or equal than 0.8. Such a strict threshold was selected to make positive proposals be as close to manual annotations as possible. Examples of positive and negative proposals are given in Figure 3.

Fig. 3. Examples of positive and negative train patches. These examples were generated automatically using the proposal generation algorithm and manual annotations.

To augment the training set for more robust training, we performed the following transformations: the amount of positive was increased by a factor of 5 using rotations in the range $[-2, 2]$ degrees with a step 0.8. Here, we trained the classifier for only one leg, since the legs on the image are symmetrical. At the classification step, we used flipped proposals of the left leg image before extracting HoG features.

4.2 Implementation details

To extract HoG, we used OpenCV [14] implementation and for SVM training we used a dual-form implementation from LIBLINEAR [11] package. To find the appropriate regularization constant C_s for SVM scorer, we used a validation set described earlier. We tried to scale the data before SVM to improve the classification results as well as hard-negative mining. However, neither of these approaches did not provided any improvement. Eventually, we found that $C_s = 0.01$ without data scaling and hard-negative mining gives the best precision and recall on the validation set.

In our pipeline, we fixed the smoothing window width $s_w = 11$ pixels, the displacement range to be $[d_1 = -\frac{1}{4}C, d_2 = \frac{1}{4}C]$ pixels, $k = 10$ and $\tau = 10$. Based on the manual annotations of the training data we estimated the set of scales $\mathbf{Z} = [3.0, 3.2, 3.4, 3.6, 3.8, 4.0, 5.0]$. After that, using the same validation set as for SVM parameters tuning, we adjusted the step p . Our main criterion was to find the best IoU having fast computational time. We found $p = 95$ pixels (IoU=0.843, time per image – 162ms).

4.3 Localization performance evaluation

Before performing the evaluation of our pipeline on the testing datasets, we estimated the quality of the generated proposals for each of them. The evaluations are presented in Figure 4.

Fig. 4. Proposals quality evaluation for each analyzed dataset – recall depending on the step value p and different IoU thresholds. Analysis shows, that it is feasible to reach a recall above 80% in every analyzed dataset having IoU threshold 0.8.

We varied the step of displacement p from 5 to 1,000 and evaluated each generated proposal by calculating IoU with the manual annotations. The best IoU score was used as a measure of quality. Eventually, we used different IoU thresholds to evaluate the best possible recall. It can be seen from Figure 4, that on all testing datasets for each given IoU threshold our proposal algorithm reaches at least 80% recall. Using the pre-trained SVM classifier described above, we also reached high mean IoU for all analyzed datasets (see, Table 2). Examples of localization are given in Figure 5.

Table 2. Localization performance evaluation. The computations were parallelized on Intel i7 5820k CPU. Time benchmarks were averaged over 3 runs.

Dataset	# of images	Mean IoU [%]	Average time [s]	Average time/image [ms]
MOST	473	0.8399	79.9922	169
Jyväskylä	93	0.7878	6.4195	14
OKOA	77	0.7761	7.5264	16

Apart from our own implementation, we also adopted the method described in [1] as a baseline (see, section 2). The only difference in our approach was that we used a larger window as a joint center patch – 40x40 pixels instead of 20x20, since with the latter, the baseline method did not perform well with our images. It can be seen from Figure 6, that our method clearly outperforms the baseline on each analyzed dataset. The mean IoU values for the baseline were 0.1, 0.33 and 0.26 for MOST, Jyväskylä and OKOA datasets, respectively.

5 Discussion and Conclusions

In this study, we presented a novel automatic method for knee joint localization on plain radiographs. We demonstrated our proposal generation approach, which allows avoiding using exhaustive sliding window approaches. Here, we also showed a new way to use the information about the anatomical knee joint structure for proposal generation. Moreover, we showed that the generated proposals in average are highly reliable, since a recall of above 80% can be reached for IoU

Fig. 5. Examples of bounding boxes produced by our method (OKOA dataset). Here, the IoU values for right and left joint are presented: a) 0.8953 and 0.77, b) 0.9373 and 0.7867, c) 0.7955 and 0.984.

Fig. 6. Evaluation of recall depending on different IoU thresholds. The curves indicate a significant advantage of our method a) in comparison to the baseline b) for each of the testing datasets.

thresholds 0.5, 0.7 and 0.8. We showed that our method significantly outperforms the baseline. In the presented results, we showed that the baseline method performs comparatively similar on Jyväskylä dataset as on OAI data in [1] (reported mean IoU was 0.36), however, the detector fails on MOST dataset. This can, most probably, be explained by the presence of artifacts – parts of knee positioning frame, which are detected as joint centers.

We demonstrated, that our method generalizes well: the trained model can be used for other datasets than the one used during the training stage. Moreover, the proposed method neither requires a complex training procedure nor much computational power and memory. The developed method is mainly designed to be used for large scale knee X-ray analysis – especially for a CAD of OA from plain knee radiographs. However, the applications are not limited to this domain only. Our approach can also be easily adapted, for example, to the hand radiographic images.

Nevertheless, some limitations of this study remain to be addressed. First of all, our results are biased due to the manual annotations – only one person annotated the images. Secondly, the used data augmentation might include some of the false positive regions in the positive set for training the scoring block of the pipeline, which can have a negative effect on the detection. Finally, our

method can be computationally optimized. For example, bias in the HoG-SVM block can explain a slight performance drop on Jyväskylä and OKOA datasets.

The method can be improved by applying the following optimizations. At first, the detection can be done on the downsampled images and then the detected joint area coordinates just need to be upsampled. We believe this optimization will significantly speed up the computations, since there is more than 10 times difference in performance between Jyväskylä dataset and MOST. However, this might require to find new hyperparameters. The second optimization would be to reuse Sobel gradients for the overlapping ROI proposals before computing HoG features, since in our current implementation, we recomputed them for each of the proposals. Furthermore, the following post-processing step can be applied for a possible improvement of the localization performance: localization regions could be centered at the joint center – this can be done by using a classifier from the baseline method [1]. However, the effect of these optimizations on the localization performance needs to be further investigated.

To conclude, despite the limitations, our method scales with a number of cores and can also be even efficiently parallelized on GPU to achieve a high-speed detection, due to the presence of loops. For example, it can be parallelized over X and Y coordinates of joints locations. It can be calculated using the values given in Table 2, that our data-parallel CPU implementation written in Python 2.7 already will allow to annotate more than 6,000,000 images of the size 2500×2000 per day (here, Jyväskylä dataset is taken as a reference), which makes the large-scale knee X-ray studies analysis possible. Eventually, our novel approach may enable the reliable and objective knee OA CAD, which will significantly benefit OA research as well as clinical OA diagnostics. The implementation of our method will be soon released on GitHub.

Acknowledgements

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.

The authors would also like to acknowledge the strategic funding of University of Oulu.

References

1. Antony, J., McGuinness, K., O'Connor, N.E., Moran, K.: Quantifying radiographic knee osteoarthritis severity using deep convolutional neural networks. In: Proceedings of 23rd International Conference on Pattern Recognition. ICPR (2016)
2. Berlin, L.: Malpractice issues in radiology. perceptual errors. *AJR. American journal of roentgenology* 167(3), 587–590 (1996)
3. Cibere, J.: Do we need radiographs to diagnose osteoarthritis? *Best Practice & Research Clinical Rheumatology* 20(1), 27–38 (2006)
4. Cross, M., Smith, E., Hoy, D., Nolte, S., Ackerman, I., Fransen, M., Bridgett, L., Williams, S., Guillemin, F., Hill, C.L., et al.: The global burden of hip and knee osteoarthritis: estimates from the global burden of disease 2010 study. *Annals of the rheumatic diseases* pp. annrheumdis–2013 (2014)
5. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05). vol. 1, pp. 886–893. IEEE (2005)
6. Demehri, S., Hafezi-Nejad, N., Carrino, J.A.: Conventional and novel imaging modalities in osteoarthritis: current state of the evidence. *Current opinion in rheumatology* 27(3), 295–303 (2015)
7. Duda, R.O., Hart, P.E., Stork, D.G., et al.: Pattern classification, vol. 2. Wiley New York (1973)
8. 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), 580–591 (2000)
9. Eckstein, F., Hudelmaier, M., Wirth, W., Kiefer, B., Jackson, R., Yu, J., Eaton, C., Schneider, E.: Double echo steady state magnetic resonance imaging of knee articular cartilage at 3 tesla: a pilot study for the osteoarthritis initiative. *Annals of the rheumatic diseases* 65(4), 433–441 (2006)

10. Englund, M., Guermazi, A., Roemer, F.W., Aliabadi, P., Yang, M., Lewis, C.E., Torner, J., Nevitt, M.C., Sack, B., Felson, D.T.: Meniscal tear in knees without surgery and the development of radiographic osteoarthritis among middle-aged and elderly persons: The multicenter osteoarthritis study. *Arthritis & Rheumatism* 60(3), 831–839 (2009)
11. 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, 1871–1874 (2008)
12. 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), 1724–1731 (2014)
13. Huo, Y., Vincken, K.L., Viergever, M.A., Lafeber, F.P.: Automatic joint detection in rheumatoid arthritis hand radiographs. In: 2013 IEEE 10th International Symposium on Biomedical Imaging. pp. 125–128. IEEE (2013)
14. Itseez: Open source computer vision library. <https://github.com/itseez/opencv> (2016)
15. Lindner, C., Thiagarajah, S., Wilkinson, J., Wallis, G., Cootes, T., arcOGEN Consortium, et al.: Development of a fully automatic shape model matching (fasmm) system to derive statistical shape models from radiographs: application to the accurate capture and global representation of proximal femur shape. *Osteoarthritis and Cartilage* 21(10), 1537–1544 (2013)
16. Multanen, J., Heinonen, A., Häkkinen, A., Kautiainen, H., Kujala, U., Lammintausta, E., Jämsä, T., Kiviranta, I., Nieminen, M.: Bone and cartilage characteristics in postmenopausal women with mild knee radiographic osteoarthritis and those without radiographic osteoarthritis. *J Musculoskeletal Neuronal Interact* 15(1), 69–77 (2015)
17. Pitman, A.: Perceptual error and the culture of open disclosure in australian radiology. *Australasian radiology* 50(3), 206–211 (2006)
18. Podlipská, J., Guermazi, A., Lehenkari, P., Niinimäki, J., Roemer, F.W., Arokoski, J.P., Kaukinen, P., Liukkonen, E., Lammintausta, E., Nieminen, M.T., et al.: Comparison of diagnostic performance of semi-quantitative knee ultrasound and knee radiography with mri: Oulu knee osteoarthritis study. *Scientific reports* 6 (2016)
19. Podsiadlo, P., Wolski, M., Stachowiak, G.: Automated selection of trabecular bone regions in knee radiographs. *Medical physics* 35(5), 1870–1883 (2008)
20. Podsiadlo, P., Cicuttini, F., Wolski, M., Stachowiak, G., Wluka, A.: Trabecular bone texture detected by plain radiography is associated with an increased risk of knee replacement in patients with osteoarthritis: a 6 year prospective follow up study. *Osteoarthritis and Cartilage* 22(1), 71–75 (2014)
21. Seise, M., McKenna, S.J., Ricketts, I.W., Wigderowitz, C.A.: Double contour active shape models. In: BMVC. Citeseer (2005)
22. 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), 1307–1312 (2009)
23. Stachowiak, G.W., Wolski, M., Woloszynski, T., Podsiadlo, P.: Detection and prediction of osteoarthritis in knee and hand joints based on the x-ray image analysis. *Biosurface and Biotribology* (2016)
24. 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. pp. 127–134. Springer (2015)
25. Uijlings, J.R., van de Sande, K.E., Gevers, T., Smeulders, A.W.: Selective search for object recognition. *International journal of computer vision* 104(2), 154–171 (2013)
26. Woloszynski, T., Podsiadlo, P., Stachowiak, G., Kurzynski, M.: A signature dissimilarity measure for trabecular bone texture in knee radiographs. *Medical physics* 37(5), 2030–2042 (2010)
27. Zitnick, C.L., Dollár, P.: Edge boxes: Locating object proposals from edges. In: European Conference on Computer Vision. pp. 391–405. Springer (2014)